

RAW
Internet-Draft
Intended status: Informational
Expires: 5 May 2021

N. Maeurer, Ed.
T. Graeupl, Ed.
German Aerospace Center (DLR)
C. Schmitt, Ed.
Research Institute CODE, UniBwM
1 November 2020

L-band Digital Aeronautical Communications System (LDACS)
draft-ietf-raw-ldacs-05

Abstract

This document provides an overview of the architecture of the L-band Digital Aeronautical Communications System (LDACS), which provides a secure, scalable and spectrum efficient terrestrial data link for civil aviation. LDACS is a scheduled, reliable multi-application cellular broadband system with support for IPv6. LDACS SHALL provide a data link for IP network-based aircraft guidance. High reliability and availability for IP connectivity over LDACS are therefore essential.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 5 May 2021.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights

and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Requirements Language	4
2. Terminology	4
3. Motivation and Use Cases	5
3.1. Voice Communications Today	5
3.2. Data Communications Today	6
4. Provenance and Documents	7
5. Applicability	8
5.1. Advances Beyond the State-of-the-Art	8
5.1.1. Priorities	8
5.1.2. Security	8
5.1.3. High Data Rates	9
5.2. Application	9
5.2.1. Air-to-Ground Multilink	9
5.2.2. Air-to-Air Extension for LDACS	9
5.2.3. Flight Guidance	10
5.2.4. Business Communication of Airlines	11
5.2.5. LDACS Navigation	11
6. Requirements to LDACS	11
7. Characteristics of LDACS	13
7.1. LDACS Sub-Network	13
7.2. Topology	14
7.3. LDACS Physical Layer	14
7.4. LDACS Data Link Layer	15
7.5. LDACS Mobility	15
8. Reliability and Availability	15
8.1. Layer 2	15
8.2. Beyond Layer 2	18
9. Protocol Stack	18
9.1. MAC Entity Services	19
9.2. DLS Entity Services	21
9.3. VI Services	22
9.4. LME Services	22
9.5. SNP Services	22
10. Security Considerations	22
10.1. Reasons for Wireless Digital Aeronautical Communications	22
10.2. Requirements for LDACS	23
10.3. Security Objectives for LDACS	24
10.4. Security Functions for LDACS	24
10.5. Security Architectural Details for LDACS	24

10.5.1.	Entities in LDACS Security Model	25
10.5.2.	Matter of LDACS Entity Identification	25
10.5.3.	Matter of LDACS Entity Authentication and Key Negotiation	25
10.5.4.	Matter of LDACS Message-in-transit Confidentiality, Integrity and Authenticity	26
10.6.	Security Architecture for LDACS	26
11.	Privacy Considerations	27
12.	IANA Considerations	27
13.	Acknowledgements	27
14.	Normative References	27
15.	Informative References	27
Appendix A.	Selected Information from DO-350A	30
Authors'	Addresses	32

1. Introduction

One of the main pillars of the modern Air Traffic Management (ATM) system is the existence of a communication infrastructure that enables efficient aircraft control and safe separation in all phases of flight. Current systems are technically mature but suffering from the VHF band's increasing saturation in high-density areas and the limitations posed by analogue radio communications. Therefore, aviation globally and the European Union (EU) in particular, strives for a sustainable modernization of the aeronautical communication infrastructure.

In the long-term, ATM communication SHALL transition from analogue VHF voice and VDLM2 communication to more spectrum efficient digital data communication. The European ATM Master Plan foresees this transition to be realized for terrestrial communications by the development (and potential implementation) of the L-band Digital Aeronautical Communications System (LDACS). LDACS SHALL enable IPv6 based air- ground communication related to the aviation safety and regularity of flight. The particular challenge is that no additional spectrum can be made available for terrestrial aeronautical communication. It was thus necessary to develop co-existence mechanism/procedures to enable the interference free operation of LDACS in parallel with other aeronautical services/systems in the same frequency band.

Since LDACS SHALL be used for aircraft guidance, high reliability and availability for IP connectivity over LDACS are essential.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Terminology

The following terms are used in the context of RAW in this document:

A2A Air-to-Air
AeroMACS Aeronautical Mobile Airport Communication System
A2G Air-to-Ground
ACARS Aircraft Communications Addressing and Reporting System
ADS-C Automatic Dependent Surveillance - Contract
AM(R)S Aeronautical Mobile (Route) Service
ANSP Air Traffic Network Service Provider
AOC Aeronautical Operational Control
AS Aircraft Station
ATC Air-Traffic Control
ATM Air-Traffic Management
ATN Aeronautical Telecommunication Network
ATS Air Traffic Service
CCCH Common Control Channel
COTS IP Commercial Off-The-Shelf
CM Context Management
CNS Communication Navigation Surveillance
CPDLC Controller Pilot Data Link Communication
DCCH Dedicated Control Channel
DCH Data Channel
DLL Data Link Layer
DLS Data Link Service
DME Distance Measuring Equipment
DSB-AM Double Side-Band Amplitude Modulation
FCI Future Communication Infrastructure
FL Forward Link
GNSS Global Navigation Satellite System
GS Ground-Station
GSC Ground-Station Controller
G2A Ground-to-Air
HF High Frequency
ICAO International Civil Aviation Organization
IP Internet Protocol
kbit/s kilobit per second
LDACS L-band Digital Aeronautical Communications System
LLC Logical Link Control
LME LDACS Management Entity
MAC Medium Access Layer

MF Multi Frame
OFDM Orthogonal Frequency-Division Multiplexing
OFDMA Orthogonal Frequency-Division Multiplexing Access
OSI Open Systems Interconnection
PHY Physical Layer
RL Reverse Link
SF Super-Frame
SNP Sub-Network Protocol
TDMA Time-Division Multiplexing-Access
VDLM1 VHF Data Link mode 1
VDLM2 VHF Data Link mode 2
VHF Very High Frequency
VI Voice Interface

3. Motivation and Use Cases

Aircraft are currently connected to Air-Traffic Control (ATC) and Aeronautical Operational Control (AOC) via voice and data communications systems through all phases of a flight. Within the airport terminal, connectivity is focused on high bandwidth communications, while during en-route high reliability, robustness, and range is the main focus. Voice communications MAY use the same or different equipment as data communications systems. In the following the main differences between voice and data communications capabilities are summarized. The assumed use cases for LDACS completes the list of use cases stated in [RAW-USE-CASES] and the list of reliable and available wireless technologies presented in [RAW-TECHNOS].

3.1. Voice Communications Today

Voice links are used for Air-to-Ground (A2G) and Air-to-Air (A2A) communications. The communication equipment is either ground-based working in the High Frequency (HF) or Very High Frequency (VHF) frequency band or satellite-based. All VHF and HF voice communications is operated via open broadcast channels without authentication, encryption or other protective measures. The use of well-proven communication procedures via broadcast channels helps to enhance the safety of communications by taking into account that other users MAY encounter communication problems and MAY be supported, if REQUIRED. The main voice communications media is still the analogue VHF Double Side-Band Amplitude Modulation (DSB-AM) communications technique, supplemented by HF Single Side-Band Amplitude Modulation and satellite communications for remote and oceanic areas. DSB-AM has been in use since 1948, works reliably and safely, and uses low-cost communication equipment. These are the main reasons why VHF DSB-AM communications is still in use, and it is

likely that this technology will remain in service for many more years. This however results in current operational limitations and impediments in deploying new Air-Traffic Management (ATM) applications, such as flight-centric operation with Point-to-Point communications.

3.2. Data Communications Today

Like for voice, data communications into the cockpit is currently provided by ground-based equipment operating either on HF or VHF radio bands or by legacy satellite systems. All these communication systems are using narrowband radio channels with a data throughput capacity in order of kilobits per second. While the aircraft is on ground some additional communications systems are available, like the Aeronautical Mobile Airport Communication System (AeroMACS) or public cellular networks, operating in the Airport (APT) domain and able to deliver broadband communication capability.

The data communication networks used for the transmission of data relating to the safety and regularity of the flight MUST be strictly isolated from those providing entertainment services to passengers. This leads to a situation that the flight crews are supported by narrowband services during flight while passengers have access to inflight broadband services. The current HF and VHF data links cannot provide broadband services now or in the future, due to the lack of available spectrum. This technical shortcoming is becoming a limitation to enhanced ATM operations, such as Trajectory-Based Operations and 4D trajectory negotiations.

Satellite-based communications are currently under investigation and enhanced capabilities are under development which will be able to provide inflight broadband services and communications supporting the safety and regularity of flight. In parallel, the ground-based broadband data link technology LDACS is being standardized by ICAO and has recently shown its maturity during flight tests [SCH20191]. The LDACS technology is scalable, secure and spectrum efficient and provides significant advantages to the users and service providers. It is expected that both - satellite systems and LDACS - will be deployed to support the future aeronautical communication needs as envisaged by the ICAO Global Air Navigation Plan.

4. Provenance and Documents

The development of LDACS has already made substantial progress in the Single European Sky ATM Research framework, short SESAR, and is currently being continued in the follow-up program SESAR2020 [RIH2018]. A key objective of these activities is to develop, implement and validate a modern aeronautical data link able to evolve with aviation needs over long-term. To this end, an LDACS specification has been produced [GRA2019] and is continuously updated; transmitter demonstrators were developed to test the spectrum compatibility of LDACS with legacy systems operating in the L-band [SAJ2014]; and the overall system performance was analyzed by computer simulations, indicating that LDACS can fulfil the identified requirements [GRA2011].

LDACS standardization within the framework of the ICAO started in December 2016. The ICAO standardization group has produced an initial Standards and Recommended Practices document [ICA2018]. It defines the general characteristics of LDACS. The ICAO standardization group plans to produce an ICAO technical manual – the ICAO equivalent to a technical standard – within the next years. Generally, the group is open to input from all sources and develops LDACS in the open.

Up to now LDACS standardization has been focused on the development of the physical layer and the data link layer, only recently have higher layers come into the focus of the LDACS development activities. There is currently no "IPv6 over LDACS" specification publicly available; however, SESAR2020 has started the testing of IPv6-based LDACS testbeds.

The IPv6 architecture for the aeronautical telecommunication network is called the Future Communications Infrastructure (FCI). FCI SHALL support quality of service, diversity, and mobility under the umbrella of the "multi-link concept". This work is conducted by ICAO Communication Panel working group WG-I.

In addition to standardization activities several industrial LDACS prototypes have been built. One set of LDACS prototypes has been evaluated in flight trials confirming the theoretical results predicting the system performance [GRA2018] [SCH20191].

5. Applicability

LDACS is a multi-application cellular broadband system capable of simultaneously providing various kinds of Air Traffic Services (including ATS-B3) and AOC communications services from deployed Ground-Stations (GS). The LDACS A2G sub-system physical layer and data link layer are optimized for data link communications, but the system also supports digital air-ground voice communications.

LDACS supports communication in all airspaces (airport, terminal maneuvering area, and en-route), and on the airport surface. The physical LDACS cell coverage is effectively de-coupled from the operational coverage REQUIRED for a particular service. This is new in aeronautical communications. Services requiring wide-area coverage can be installed at several adjacent LDACS cells. The handover between the involved LDACS cells is seamless, automatic, and transparent to the user. Therefore, the LDACS A2G communications concept enables the aeronautical communication infrastructure to support future dynamic airspace management concepts.

5.1. Advances Beyond the State-of-the-Art

LDACS offers several capabilities that are not provided in contemporarily deployed aeronautical communication systems.

5.1.1. Priorities

LDACS is able to manage services priorities, an important feature not available in some of the current data link deployments. Thus, LDACS guarantees bandwidth, low latency, and high continuity of service for safety critical ATS applications while simultaneously accommodating less safety-critical AOC services.

5.1.2. Security

LDACS is a secure data link with built-in security mechanisms. It enables secure data communications for ATS and AOC services, including secured private communications for aircraft operators and ANSPs (Air Traffic Network Service Providers). This includes concepts for key and trust management, mutual authenticated key exchange protocols, key derivation measures, user and control message-in-transit confidentiality and authenticity protection, secure logging and availability and robustness measures [MAE20181], [MAE20191], [MAE20192].

5.1.3. High Data Rates

The user data rate of LDACS is 315 kbit/s to 1428 kbit/s on the forward link (FL) for the connection Ground-to-Air (G2A), and 294 kbit/s to 1390 kbit/s on the reverse link (RF) for the connection A2G, depending on coding and modulation. This is 50 times the amount terrestrial digital aeronautical communications systems such as VDL M2 provide [SCH20191].

5.2. Application

LDACS SHALL be used by several aeronautical applications ranging from enhanced communication protocol stacks (multi-homed mobile IPv6 networks in the aircraft and potentially ad-hoc networks between aircraft) to classical communication applications (sending GBAS correction data) and integration with other service domains (using the communication signal for navigation).

5.2.1. Air-to-Ground Multilink

It is expected that LDACS together with upgraded satellite-based communications systems will be deployed within the FCI and constitute one of the main components of the multilink concept within the FCI.

Both technologies, LDACS and satellite systems, have their specific benefits and technical capabilities which complement each other. Especially, satellite systems are well-suited for large coverage areas with less dense air traffic, e.g. oceanic regions. LDACS is well-suited for dense air traffic areas, e.g. continental areas or hot-spots around airports and terminal airspace. In addition, both technologies offer comparable data link capacity and, thus, are well-suited for redundancy, mutual back-up, or load balancing.

Technically the FCI multilink concept SHALL be realized by multi-homed mobile IPv6 networks in the aircraft. The related protocol stack is currently under development by ICAO and the Single European Sky ATM Research framework.

5.2.2. Air-to-Air Extension for LDACS

A potential extension of the multi-link concept is its extension to ad-hoc networks between aircraft.

Direct A2A communication between aircrafts in terms of ad-hoc data networks is currently considered a research topic since there is no immediate operational need for it, although several possible use cases are discussed (digital voice, wake vortex warnings, and trajectory negotiation) [BEL2019]. It SHOULD also be noted that

currently deployed analog VHF voice radios support direct voice communication between aircraft, making a similar use case for digital voice plausible.

LDACS direct A2A is currently not part of standardization.

5.2.3. Flight Guidance

The FCI (and therefore LDACS) SHALL be used to host flight guidance. This is realized using three applications:

1. Context Management (CM): The CM application SHALL manage the automatic logical connection to the ATC center currently responsible to guide the aircraft. Currently this is done by the air crew manually changing VHF voice frequencies according to the progress of the flight. The CM application automatically sets up equivalent sessions.
2. Controller Pilot Data Link Communication (CPDLC): The CPDLC application provides the air crew with the ability to exchange data messages similar to text messages with the currently responsible ATC center. The CPDLC application SHALL take over most of the communication currently performed over VHF voice and enable new services that do not lend themselves to voice communication (e.g., trajectory negotiation).
3. Automatic Dependent Surveillance - Contract (ADS-C): ADS-C reports the position of the aircraft to the currently active ATC center. Reporting is bound to "contracts", i.e. pre-defined events related to the progress of the flight (i.e. the trajectory). ADS-C and CPDLC are the primary applications used to implement in-flight trajectory management.

CM, CPDLC, and ADS-C are available on legacy datalinks, but not widely deployed and with limited functionality.

Further ATC applications MAY be ported to use the FCI or LDACS as well. A notable application is GBAS for secure, automated landings: The Global Navigation Satellite System (GNSS) based Ground Based Augmentation System (GBAS) is used to improve the accuracy of GNSS to allow GNSS based instrument landings. This is realized by sending GNSS correction data (e.g., compensating ionospheric errors in the GNSS signal) to the aircraft's GNSS receiver via a separate data link. Currently the VDB data link is used. VDB is a narrow-band single-purpose datalink without advanced security only used to transmit GBAS correction data. This makes VDB a natural candidate for replacement by LDACS.

5.2.4. Business Communication of Airlines

In addition to air traffic services AOC services SHALL be transmitted over LDACS. AOC is a generic term referring to the business communication of airlines. Regulatory this is considered related to the safety and regularity of flight and MAY therefore be transmitted over LDACS.

AOC communication is considered the main business case for LDACS communication service providers since modern aircraft generate significant amounts of data (e.g., engine maintenance data).

5.2.5. LDACS Navigation

Beyond communication radio signals can always also be used for navigation. LDACS takes this into account.

For future aeronautical navigation, ICAO recommends the further development of GNSS based technologies as primary means for navigation. However, the drawback of GNSS is its inherent single point of failure – the satellite. Due to the large separation between navigational satellites and aircraft, the received power of GNSS signals on the ground is very low. As a result, GNSS disruptions might occasionally occur due to unintentional interference, or intentional jamming. Yet the navigation services MUST be available with sufficient performance for all phases of flight. Therefore, during GNSS outages, or blockages, an alternative solution is needed. This is commonly referred to as Alternative Positioning, Navigation, and Timing (APNT).

One of such APNT solution consists of integrating the navigation functionality into LDACS. The ground infrastructure for APNT is deployed through the implementation of LDACS's GSs and the navigation capability comes "for free".

LDACS navigation has already been demonstrated in practice in a flight measurement campaign [SCH20191].

6. Requirements to LDACS

The requirements to LDACS are mostly defined by its application area: Communication related to safety and regularity of flight.

A particularity of the current aeronautical communication landscape is that it is heavily regulated. Aeronautical data links (for applications related to safety and regularity of flight) MAY only use spectrum licensed to aviation and data links endorsed by ICAO. Nation states can change this locally, however, due to the global scale of the air transportation system adherence to these practices is to be expected.

Aeronautical data links for the Aeronautical Telecommunication Network (ATN) are therefore expected to remain in service for decades. The VDLM2 data link currently used for digital terrestrial internetworking was developed in the 1990es (the use of the Open Systems Interconnection (OSI) stack indicates that as well). VDLM2 is expected to be used at least for several decades. In this respect aeronautical communication (for applications related to safety and regularity of flight) is more comparable to industrial applications than to the open Internet.

Internetwork technology is already installed in current aircraft. Current ATS applications use either the Aircraft Communications Addressing and Reporting System (ACARS) or the OSI stack. The objective of the development effort LDACS as part of the FCI is to replace legacy OSI stack and proprietary ACARS internetwork technologies with industry standard IP technology. It is anticipated that the use of Commercial Off-The-Shelf (COTS) IP technology mostly applies to the ground network. The avionics networks on the aircraft will likely be heavily modified or proprietary.

AOC applications currently mostly use the same stack (although some applications, like the graphical weather service MAY use the commercial passenger network). This creates capacity problems (resulting in excessive amounts of timeouts) since the underlying terrestrial data links (VDLM1/2) do not provide sufficient bandwidth. The use of non-aviation specific data links is considered a security problem. Ideally the aeronautical IP internetwork and the Internet SHOULD be completely separated.

The objective of LDACS is to provide a next generation terrestrial data link designed to support IP and provide much higher bandwidth to avoid the currently experienced operational problems.

The requirement for LDACS is therefore to provide a terrestrial high-throughput data link for IP internetworking in the aircraft.

In order to fulfil the above requirement LDACS needs to be interoperable with IP (and IP-based services like Voice-over-IP) at the gateway connecting the LDACS network to other aeronautical ground networks (the totality of them being the ATN). On the avionics side in the aircraft aviation specific solutions are to be expected.

In addition to the functional requirements LDACS and its IP stack need to fulfil the requirements defined in RTCA DO-350A/EUROCAE ED-228A [DO350A]. This document defines continuity, availability, and integrity requirements at different scopes for each air traffic management application (CPDLC, CM, and ADS-C). The scope most relevant to IP over LDACS is the CSP (Communication Service Provider) scope.

Continuity, availability, and integrity requirements are defined in [DO350A] volume 1 Table 5-14, and Table 6-13. Appendix A presents the REQUIRED information.

In a similar vein, requirements to fault management are defined in the same tables.

7. Characteristics of LDACS

LDACS will become one of several wireless access networks connecting aircraft to the ATN implemented by the FCI and possibly ACARS/FANS networks [FAN2019].

The current LDACS design is focused on the specification of layer 2.

Achieving stringent the continuity, availability, and integrity requirements defined in [DO350A] will require the specification of layer 3 and above mechanisms (e.g. reliable crossover at the IP layer). Fault management mechanisms are similarly undefined. Input from the working group will be appreciated here.

7.1. LDACS Sub-Network

An LDACS sub-network contains an Access Router (AR), a Ground-Station Controller (GSC), and several GS, each of them providing one LDACS radio cell.

User plane interconnection to the ATN is facilitated by the AR peering with an A2G Router connected to the ATN. It is up to implementer's choice to keep AR and A2G Router functions separated, or to merge them.

The internal control plane of an LDACS sub-network is managed by the GSC. An LDACS sub-network is illustrated in Figure 1.

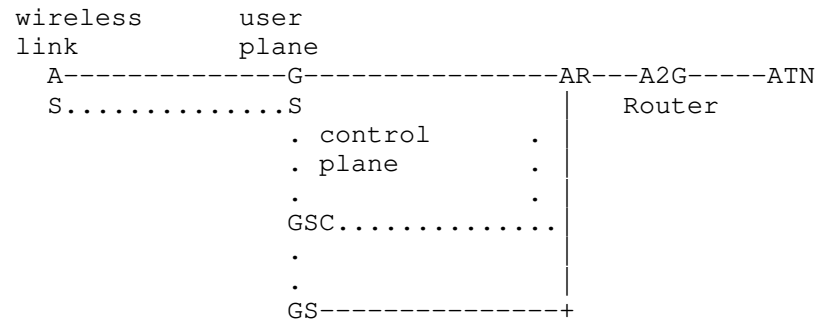


Figure 1: LDACS sub-network with two GSs and one AS

7.2. Topology

LDACS operating in A2G mode is a cellular point-to-multipoint system. The A2G mode assumes a star-topology in each cell where Aircraft Stations (AS) belonging to aircraft within a certain volume of space (the LDACS cell) is connected to the controlling GS. The LDACS GS is a centralized instance that controls LDACS A2G communications within its cell. The LDACS GS can simultaneously support multiple bi-directional communications to the ASs under its control. LDACS's GSs themselves are connected to a GSC controlling the LDACS sub-network.

Prior to utilizing the system an AS has to register with the controlling GS to establish dedicated logical channels for user and control data. Control channels have statically allocated resources, while user channels have dynamically assigned resources according to the current demand. Logical channels exist only between the GS and the AS.

The LDACS wireless link protocol stack defines two layers, the physical layer and the data link layer.

7.3. LDACS Physical Layer

The physical layer provides the means to transfer data over the radio channel. The LDACS GS supports bi-directional links to multiple aircraft under its control. The FL direction at the G2A connection and the RL direction at the A2G connection are separated by Frequency Division Duplex. FL and RL use a 500 kHz channel each. The GS transmits a continuous stream of Orthogonal Frequency-Division Multiplexing (OFDM) symbols on the FL. In the RL different aircraft are separated in time and frequency using a combination of Orthogonal

Frequency-Division Multiple-Access (OFDMA) and Time-Division Multiple-Access (TDMA). Aircraft thus transmit discontinuously on the RL with radio bursts sent in precisely defined transmission opportunities allocated by the GS.

7.4. LDACS Data Link Layer

The data-link layer provides the necessary protocols to facilitate concurrent and reliable data transfer for multiple users. The LDACS data link layer is organized in two sub-layers: The medium access sub-layer and the Logical Link Control (LLC) sub-layer. The medium access sub-layer manages the organization of transmission opportunities in slots of time and frequency. The LLC sub-layer provides acknowledged point-to-point logical channels between the aircraft and the GS using an automatic repeat request protocol. LDACS supports also unacknowledged point-to-point channels and G2A broadcast.

7.5. LDACS Mobility

LDACS supports layer 2 handovers to different LDACS channels. Handovers MAY be initiated by the aircraft (break-before-make) or by the GS (make-before-break). Make-before-break handovers are only supported for GSs connected to the same GSC.

External handovers between non-connected LDACS sub-networks or different aeronautical data links SHALL be handled by the FCI multi-link concept.

8. Reliability and Availability

8.1. Layer 2

LDACS has been designed with applications related to the safety and regularity of flight in mind. It has therefore been designed as a deterministic wireless data link (as far as this is possible).

Based on channel measurements of the L-band channel [SCHN2016] and respecting the specific nature of the area of application, LDACS was designed from the PHY layer up with robustness in mind.

In order to maximize the capacity per channel and to optimally use the available spectrum, LDACS was designed as an OFDM-based Frequency Division Duplex system, supporting simultaneous transmissions in FL at the G2A connection and RF at the A2G connection. The legacy systems already deployed in the L-band limit the bandwidth of both channels to approximately 500 kHz.

The LDACS physical layer design includes propagation guard times sufficient for the operation at a maximum distance of 200 nautical miles from the GS. In actual deployment, LDACS can be configured for any range up to this maximum range.

The LDACS FL physical layer is a continuous OFDM transmission. LDACS RL transmission is based on OFDMA-TDMA bursts, with silence between such bursts. The RL resources (i.e. bursts) are assigned to different ASs on demand by the GS.

The LDACS physical layer supports adaptive coding and modulation for user data. Control data is always encoded with the most robust coding and modulation (QPSK coding rate 1/2).

LDACS medium access on top of the physical layer uses a static frame structure to support deterministic timer management. As shown in Figure 3 and Figure 4, LDACS framing structure is based on Super-Frames (SF) of 240ms duration corresponding to 2000 OFDM symbols. FL and RL boundaries are aligned in time (from the GS perspective) allowing for deterministic sending windows for KEEP ALIVE messages and control and data channels in general.

LDACS medium access is always under the control of the GS of a radio cell. Any medium access for the transmission of user data has to be requested with a resource request message stating the requested amount of resources and class of service. The GS performs resource scheduling on the basis of these requests and grants resources with resource allocation messages. Resource request and allocation messages are exchanged over dedicated contention-free control channels.

The purpose of Quality-of-Service in LDACS medium access is to provide prioritized medium access at the bottleneck (the wireless link). The signaling of higher layer Quality-of-Service requirements to LDACS is yet to be defined. A DiffServ-based solution with a small number of priorities is to be expected.

LDACS has two mechanisms to request resources from the scheduler in the GS.

Resources can either be requested "on demand" with a given priority. On the FL, this is done locally in the GS, on the RL a dedicated contention-free control channel is used called Dedicated Control Channel (DCCH), which is roughly 83 bit every 60 ms. A resource allocation is always announced in the control channel of the FL, short Common Control Channel (CCCH) having variable size. Due to the spacing of the RL control channels every 60 ms, a medium access delay in the same order of magnitude is to be expected.

Resources can also be requested "permanently". The permanent resource request mechanism supports requesting recurring resources in given time intervals. A permanent resource request has to be canceled by the user (or by the GS, which is always in control).

User data transmissions over LDACS are therefore always scheduled by the GS, while control data uses statically (i.e. at cell entry) allocated recurring resources (DCCH and CCCH). The current specification specifies no scheduling algorithm. Scheduling of RL resources is done in physical Protocol Data Units of 112 bit (or larger if more aggressive coding and modulation is used). Scheduling on the FL is done Byte-wise since the FL is transmitted continuously by the GS.

In addition to having full control over resource scheduling, the GS can send forced Handover commands for off-loading or RF channel management, e.g. when the signal quality declines and a more suitable GS is in the AS reach. With robust resource management of the capacities of the radio channel, reliability and robustness measures are therefore also anchored in the LDACS management entity.

In addition, to radio resource management, the LDACS control channels are also used to send keep-alive messages, when they are not otherwise used. Since the framing of the control channels is deterministic, missing keep-alive messages can thus be immediately detected. This information is made available to the multi-link protocols for fault management.

The protocol used to communicate faults is not defined in the LDACS specification. It is assumed that vendors would use industry standard protocols like the Simple Network Management Protocol or the Network Configuration Protocol where security permits.

The LDACS data link layer protocol running on top of the medium access sub-layer uses ARQ to provide reliable data transmission on layer 2.

It employs selective repeat ARQ with transparent fragmentation and reassembly to the resource allocation size to achieve low latency and a low overhead without losing reliability. It ensures correct order of packet delivery without duplicates. In case of transmission errors it identifies lost fragments with deterministic timers synced to the medium access frame structure and initiates retransmission. Additionally, the priority mechanism of LDACS ensures the timely delivery of messages with high importance.

8.2. Beyond Layer 2

LDACS availability can be increased by appropriately deploying LDACS infrastructure: This means proliferating the number of terrestrial base stations. However, the scarcity of aeronautical spectrum for data link communication (in the case of LDACS: tens of MHz in the L-band) and the long range (in the case of LDACS: up to 400 km) make this quite hard. The deployment of a larger number of small cells is certainly possible, suffers, however, also from the scarcity of spectrum. An additional constraint to take into account, is that Distance Measuring Equipment (DME) is the primary user of the aeronautical L-band. That is, any LDACS deployment has to take DME frequency planning into account, too.

The aeronautical community has therefore decided not to rely on a single communication system or frequency band. It is envisioned to have multiple independent data link technologies in the aircraft (e.g., terrestrial and SatCom) in addition to legacy VHF voice.

However, as of now no reliability and availability mechanisms that could utilize the multi-link have been specified on Layer 3 and above.

Below Layer 2 aeronautics usually relies on hardware redundancy. To protect availability of the LDACS link, an aircraft equipped with LDACS will have access to two L-band antennae with triple redundant radio systems as REQUIRED for any safety relevant system by ICAO.

9. Protocol Stack

The protocol stack of LDACS is implemented in the AS, GS, and GSC: It consists of the Physical Layer (PHY) with five major functional blocks above it. Four are placed in the Data Link Layer (DLL) of the AS and GS: (1) Medium Access Layer (MAC), (2) Voice Interface (VI), (3) Data Link Service (DLS), and (4) LDACS Management Entity (LME). The last entity resides within the Sub-Network Layer: Sub-Network Protocol (SNP). The LDACS network is externally connected to voice units, radio control units, and the ATN Network Layer.

Figure 2 shows the protocol stack of LDACS as implemented in the AS and GS.

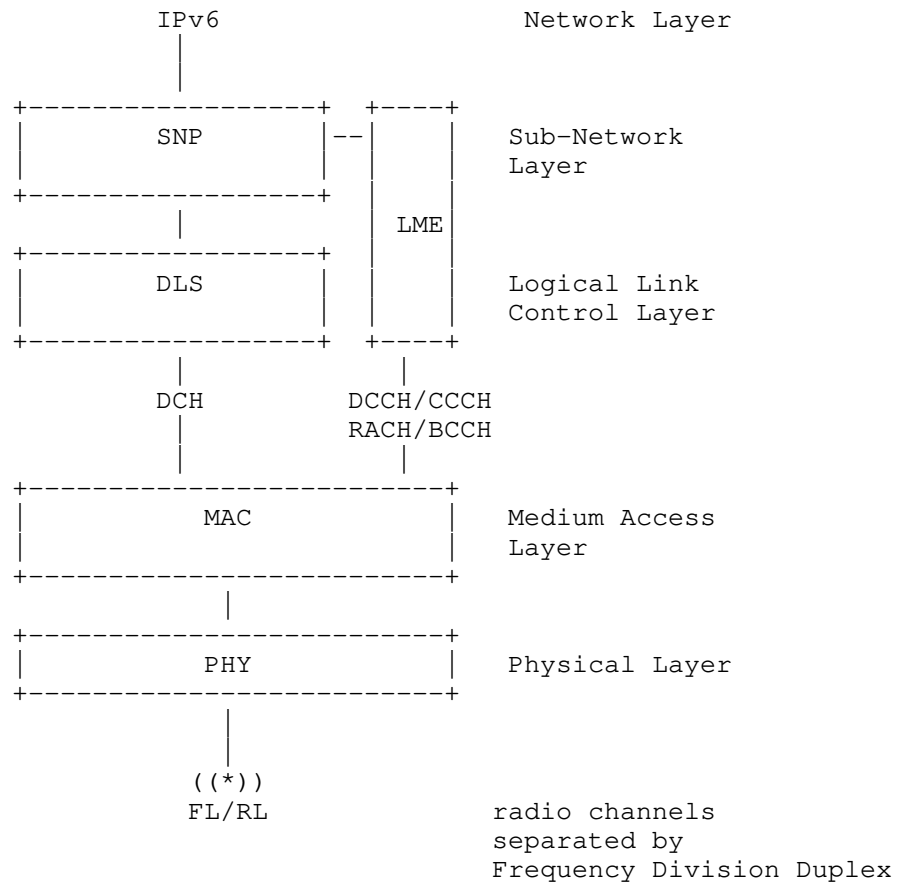


Figure 2: LDACS protocol stack in AS and GS

9.1. MAC Entity Services

The MAC time framing service provides the frame structure necessary to realize slot-based Time Division Multiplex access on the physical link. It provides the functions for the synchronization of the MAC framing structure and the PHY Layer framing. The MAC time framing provides a dedicated time slot for each logical channel.

The MAC Sub-Layer offers access to the physical channel to its service users. Channel access is provided through transparent logical channels. The MAC Sub-Layer maps logical channels onto the appropriate slots and manages the access to these channels. Logical channels are used as interface between the MAC and LLC Sub-Layers.

The LDACS framing structure for FL and RL is based on Super-Frames (SF) of 240 ms duration. Each SF corresponds to 2000 OFDM symbols. The FL and RL SF boundaries are aligned in time (from the view of the GS).

In the FL, an SF contains a Broadcast Frame of duration 6.72 ms (56 OFDM symbols) for the Broadcast Control Channel (BCCH), and four Multi-Frames (MF), each of duration 58.32 ms (486 OFDM symbols).

In the RL, each SF starts with a Random Access (RA) slot of length 6.72 ms with two opportunities for sending RL random access frames for the Random Access Channel (RACH), followed by four MFs. These MFs have the same fixed duration of 58.32 ms as in the FL, but a different internal structure

Figure 3 and Figure 4 illustrate the LDACS frame structure.

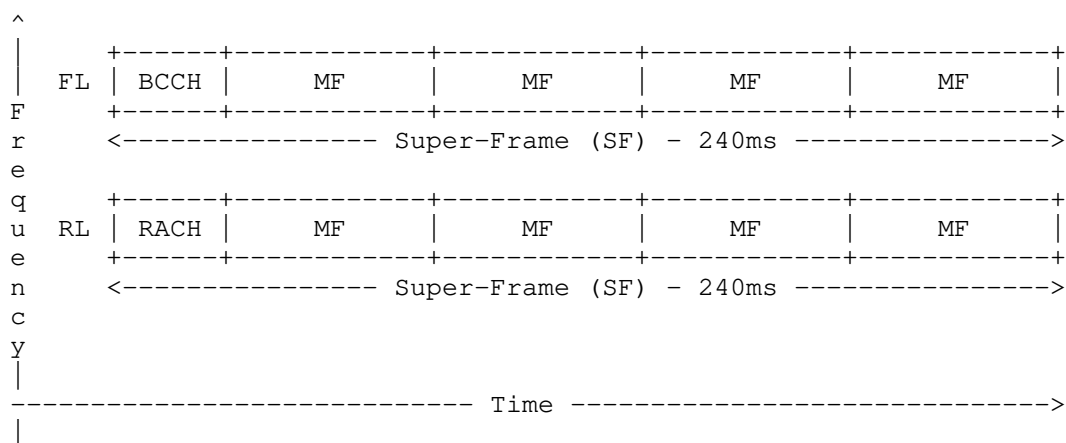


Figure 3: SF structure for LDACS




Figure 4: MF structure for LDACS

9.2. DLS Entity Services

The DLS provides acknowledged and unacknowledged (including broadcast and packet mode voice) bi-directional exchange of user data. If user data is transmitted using the acknowledged DLS, the sending DLS entity will wait for an acknowledgement from the receiver. If no acknowledgement is received within a specified time frame, the sender MAY automatically try to retransmit its data. However, after a certain number of failed retries, the sender will suspend further retransmission attempts and inform its client of the failure.

The DLS uses the logical channels provided by the MAC:

1. A GS announces its existence and access parameters in the Broadcast Channel (BC).
2. The RA channel enables AS to request access to an LDACS cell.
3. In the FL the CCCH is used by the GS to grant access to data channel resources.
4. The reverse direction is covered by the RL, where ASs need to request resources before sending. This happens via the DCCH.
5. User data itself is communicated in the Data Channel (DCH) on the FL and RL.

9.3. VI Services

The VI provides support for virtual voice circuits. Voice circuits MAY either be set-up permanently by the GS (e.g., to emulate voice party line) or MAY be created on demand. The creation and selection of voice circuits is performed in the LME. The VI provides only the transmission services.

9.4. LME Services

The mobility management service in the LME provides support for registration and de-registration (cell entry and cell exit), scanning RF channels of neighboring cells and handover between cells. In addition, it manages the addressing of aircraft/ ASs within cells. It is controlled by the network management service in the GSC.

The resource management service provides link maintenance (power, frequency and time adjustments), support for adaptive coding and modulation, and resource allocation.

9.5. SNP Services

The DLS provides functions REQUIRED for the transfer of user plane data and control plane data over the LDACS sub-network.

The security service provides functions for secure communication over the LDACS sub-network. Note that the SNP security service applies cryptographic measures as configured by the GSC.

10. Security Considerations

10.1. Reasons for Wireless Digital Aeronautical Communications

Aviation will require secure exchanges of data and voice messages for managing the air-traffic flow safely through the airspaces all over the world. Historically Communication Navigation Surveillance (CNS) wireless communications technology emerged from military and a threat landscape where inferior technological and financial capabilities of adversaries were assumed [STR2016]. The main communication method for ATC today is still an open analogue voice broadcast within the aeronautical VHF band. Currently, the information security is purely procedural based by using well-trained personnel and proven communications procedures. This communication method has been in service since 1948. However since the emergence of civil aeronautical CNS application and today, the world has changed. First of all civil applications have significant lower spectrum available than military applications. This means several military defense mechanisms such as frequency hopping or pilot symbol scrambling and

thus a defense-in-depth approach starting at the physical layer is impossible for civil systems. With the rise of cheap Software Defined Radios, the previously existing financial barrier is almost gone and open source projects such as GNU radio [GNU2012] allow the new type of unsophisticated listeners and possible attackers. Furthermore most CNS technology developed in ICAO relies on open standards, thus syntax and semantics of wireless digital aeronautical communications can be common knowledge for attackers. Finally with increased digitization and automation of civil aviation the human as control instance is being taken gradually out of the loop. Autonomous transport drones or single piloted aircraft demonstrate this trend. However without profound cybersecurity measures such as authenticity and integrity checks of messages in-transit on the wireless link or mutual entity authentication, this lack of a control instance can prove disastrous. Thus future digital communications waveforms will need additional embedded security features to fulfill modern information security requirements like authentication and integrity. However, these security features require sufficient bandwidth which is beyond the capabilities of a VHF narrowband communications system. For voice and data communications, sufficient data throughput capability is needed to support the security functions while not degrading performance. LDACS is a data link technology with sufficient bandwidth to incorporate security without losing too much user throughput.

As digitalization progresses even further with LDACS and automated procedures such as 4D-Trajectories allowing semi-automated en-route flying of aircraft, LDACS requires stronger cybersecurity measures.

10.2. Requirements for LDACS

Overall there are several business goals for cybersecurity to protect in FCI in civil aviation:

1. Safety: The system MUST sufficiently mitigate attacks, which contribute to safety hazards.
2. Flight regularity: The system MUST sufficiently mitigate attacks, which contribute to delays, diversions, or cancellations of flights.
3. Protection of business interests: The system MUST sufficiently mitigate attacks which result in financial loss, reputation damage, disclosure of sensitive proprietary information, or disclosure of personal information.

To further analyze assets and derive threats and thus protection scenarios several Threat-and Risk Analysis were performed for LDACS [MAE20181] , [MAE20191]. These results allowed deriving security scope and objectives from the requirements and the conducted Threat-and Risk Analysis.

10.3. Security Objectives for LDACS

Security considerations for LDACS are defined by the official Standards And Recommended Practices document by ICAO [ICA2018]:

1. LDACS SHALL provide a capability to protect the availability and continuity of the system.
2. LDACS SHALL provide a capability including cryptographic mechanisms to protect the integrity of messages in transit.
3. LDACS SHALL provide a capability to ensure the authenticity of messages in transit.
4. LDACS SHOULD provide a capability for nonrepudiation of origin for messages in transit.
5. LDACS SHOULD provide a capability to protect the confidentiality of messages in transit.
6. LDACS SHALL provide an authentication capability.
7. LDACS SHALL provide a capability to authorize the permitted actions of users of the system and to deny actions that are not explicitly authorized.
8. If LDACS provides interfaces to multiple domains, LDACS SHALL provide capability to prevent the propagation of intrusions within LDACS domains and towards external domains.

10.4. Security Functions for LDACS

These objectives were used to derive several security functions for LDACS REQUIRED to be integrated in the LDACS cybersecurity architecture: (1) Identification, (2) Authentication, (3) Authorization, (4) Confidentiality, (5) System Integrity, (6) Data Integrity, (7) Robustness, (8) Reliability, (9) Availability, and (10) Key and Trust Management. Several works investigated possible measures to implement these security functions [BIL2017], [MAE20181], [MAE20191]. Having identified security requirements, objectives and functions it MUST be ensured that they are applicable.

10.5. Security Architectural Details for LDACS

The requirements lead to a LDACS security model including different entities for identification, authentication and authorization purposes ensuring integrity, authenticity and confidentiality of data in-transit especially.

10.5.1. Entities in LDACS Security Model

A simplified LDACS architectural model requires the following entities: Network operators such as the Societe Internationale de Telecommunications Aeronautiques (SITA) [SIT2020] and ARINC [ARI2020] are providing access to the (1) Ground IPS network via an (2) A2G LDACS Router. This router is attached to a closed off LDACS Access Network (3) which connects via further (4) Access Routers to the different (5) LDACS Cell Ranges, each controlled by a (6) GSC and spanning a local LDACS Access Network connecting to the (7) GSs that serve one LDACS cell. Via the (8) A2G wireless LDACS data link (9) AS the aircraft is connected to the ground network and via the (10) aircraft's VI and (11) aircraft's network interface, aircraft's data can be sent via the AS back to the GS and the forwarded back via GSC, LDACS local access network, access routers, LDACS access network, A2G LDACS router to the ground IPS network.

10.5.2. Matter of LDACS Entity Identification

LDACS needs specific identities for (1) the AS, (2) the GS, (3) the GSC and (4) the Network Operator. The aircraft itself can be identified using the ICAO unique address of an aircraft, the call sign of that aircraft or the recently founded Privacy ICAO Address (PIA) program [FAA2020]. It is conceivable that the LDACS AS will use a combination of aircraft identification, radio component identification such as MAC addresses and even operator features identification to create a unique AS LDACS identification tag. Similar to a 4G's eNodeB Serving Network (SN) Identification tag, a GS could be identified using a similar field. And again similar to 4G's Mobility Management Entities (MME), a GSC could be identified using similar identification fields within the LDACS network. The identification of the network operator is again similar to 4G (e.g., E-Plus, AT&T, and TELUS), in the way that the aeronautical network operators are listed (e.g., ARINC [ARI2020] and SITA [SIT2020]).

10.5.3. Matter of LDACS Entity Authentication and Key Negotiation

In order to anchor Trust within the system all LDACS entities connected to the ground IPS network SHALL be rooted in an LDACS specific chain-of-trust and PKI solution, quite similar to AeroMACS approach [CRO2016]. These X.509 certificates [RFC5280] residing at the entities and incorporated in the LDACS PKI proof the ownership of their respective public key, include information about the identity of the owner and the digital signature of the entity that has verified the certificate's content. First all ground infrastructures MUST mutually authenticate to each other, negotiate and derive keys and, thus, secure all ground connections. How this process is handled in detail is still an ongoing discussion. However,

established methods to secure user plane by IPSec [RFC4301] and IKEv2 [RFC7296] or the application layer via TLS 1.3 [RFC8446] are conceivable. The LDACS PKI with their chain-of-trust approach, digital certificates and public entity keys lay the groundwork for this step. In a second step the AS with the LDACS radio approaches an LDACS cell and performs a cell entry with the corresponding GS. Similar to the LTE cell attachment process [TS33.401], where authentication happens after basic communication has been enabled between AS and GS (step 5a in the UE attachment process [TS33.401]), the next step is mutual authentication and key exchange. Hence, in step three using the identity based Station-to-Station (STS) protocol with Diffie-Hellman Key Exchange [MAE2020], AS and GS establish mutual trust by authenticating each other, exchanging key material and finally both ending up with derived key material. A key confirmation is mandatory before the communication channel between the AS and the GS can be opened for user-data communications.

10.5.4. Matter of LDACS Message-in-transit Confidentiality, Integrity and Authenticity

The subsequent key material from the previous step can then be used to protect LDACS Layer 2 communications via applying encryption and integrity protection measures on the SNP layer of the LDACS protocol stack. As LDACS transports AOC and ATS data, the integrity of that data is most important, while confidentiality only needs to be applied to AOC data to protect business interests [ICA2018]. This possibility of providing low layered confidentiality and integrity protection ensures a secure delivery of user data over the air gap. Furthermore it ensures integrity protection of LDACS control data.

10.6. Security Architecture for LDACS

A draft of the cybersecurity architecture of LDACS can be found in [ICA2018] and [MAE20182] and respective updates in [MAE20191], [MAE20192], and [MAE2020]. It proposes the use of an own LDACS PKI, identity management based on aircraft identities and network operator identities (e.g., SITA and ARINC), public key certificates incorporated in the PKI based chain-of-trust and stored in the entities allowing for mutual authentication and key exchange procedures, key derivation mechanisms for perfect forward secrecy and user/control plane message-in-transit integrity and confidentiality protection. This secures data traveling over the airgap between AS and GS and also between GS and ANSP regardless of the secure or unsecure nature of application data. Of course application data itself MUST be additionally secured to achieve end-to-end security (secure dialogue service), however the LDACS datalinks aims to provide an additional layer of protection just for this network segment.

11. Privacy Considerations

LDACS provides a Quality-of-Service, and the generic considerations for such mechanisms apply.

12. IANA Considerations

This memo includes no request to IANA.

13. Acknowledgements

Thanks to all contributors to the development of LDACS and ICAO PT-T.

Thanks to Klaus-Peter Hauf, Bart Van Den Einden, and Pierluigi Fantappie for further input to this draft.

Thanks to SBA Research Vienna for fruitful discussions on aeronautical communications concerning security incentives for industry and potential economic spillovers.

14. Normative References

- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<https://www.rfc-editor.org/info/rfc4301>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/info/rfc5280>>.
- [RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October 2014, <<https://www.rfc-editor.org/info/rfc7296>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

15. Informative References

- [SCHN2016] Schneckenburger, N., Jost, T., Shutin, D., Walter, M., Thiasiriphet, T., Schnell, M., and U.C. Fiebig, "Measurement of the L-band Air-to-Ground Channel for Positioning Applications", IEEE Transactions on Aerospace and Electronic Systems, 52(5), pp.2281-229 , 2016.
- [MAE20191] Maeurer, N., Graeupl, T., and C. Schmitt, "Evaluation of the LDACS Cybersecurity Implementation", IEEE 38th Digital Avionics Systems Conference (DACS), pp. 1-10, San Diego, CA, USA , 2019.
- [MAE20192] Maeurer, N. and C. Schmitt, "Towards Successful Realization of the LDACS Cybersecurity Architecture: An Updated Datalink Security Threat- and Risk Analysis", IEEE Integrated Communications, Navigation and Surveillance Conference (ICNS), pp. 1-13, Herndon, VA, USA , 2019.
- [GRA2019] Graeupl, T., Rihacek, C., and B. Haindl, "LDACS A/G Specification", SESAR2020 PJ14-02-01 D3.3.030 , 2019.
- [FAN2019] Pierattelli, S., Fantappie, P., Tamalet, S., van den Einden, B., Rihacek, C., and T. Graeupl, "LDACS Deployment Options and Recommendations", SESAR2020 PJ14-02-01 D3.4.020 , 2019.
- [MAE20182] Maeurer, N. and A. Bilzhause, "A Cybersecurity Architecture for the L-band Digital Aeronautical Communications System (LDACS)", IEEE 37th Digital Avionics Systems Conference (DASC), pp. 1-10, London, UK , 2017.
- [GRA2011] Graeupl, T. and M. Ehammer, "L-DACS1 Data Link Layer Evolution of ATN/IPS", 30th IEEE/AIAA Digital Avionics Systems Conference (DASC), pp. 1-28, Seattle, WA, USA , 2011.
- [GRA2018] Graeupl, T., Schneckenburger, N., Jost, T., Schnell, M., Filip, A., Bellido-Manganell, M.A., Mielke, D.M., Maeurer, N., Kumar, R., Osechas, O., and G. Battista, "L-band Digital Aeronautical Communications System (LDACS) flight trials in the national German project MICONAV", Integrated Communications, Navigation, Surveillance Conference (ICNS), pp. 1-7, Herndon, VA, USA , 2018.
- [SCH20191] Schnell, M., "DLR Tests Digital Communications Technologies Combined with Additional Navigation Functions for the First Time", 2019.

- [ICA2018] International Civil Aviation Organization (ICAO), "L-Band Digital Aeronautical Communication System (LDACS)", International Standards and Recommended Practices Annex 10 - Aeronautical Telecommunications, Vol. III - Communication Systems , 2018.
- [SAJ2014] Haindl, B., Meser, J., Sajatovic, M., Mueller, S., Arthaber, H., Faseth, T., and M. Zaisberger, "LDACS1 Conformance and Compatibility Assessment", IEEE/AIAA 33rd Digital Avionics Systems Conference (DASC), pp. 1-11, Colorado Springs, CO, USA , 2014.
- [RIH2018] Rihacek, C., Haindl, B., Fantappie, P., Pierattelli, S., Graeupl, T., Schnell, M., and N. Fistas, "L-band Digital Aeronautical Communications System (LDACS) Activities in SESAR2020", Integrated Communications Navigation and Surveillance Conference (ICNS), pp. 1-8, Herndon, VA, USA , 2018.
- [BEL2019] Bellido-Manganell, M. A. and M. Schnell, "Towards Modern Air-to-Air Communications: the LDACS A2A Mode", IEEE/AIAA 38th Digital Avionics Systems Conference (DASC), pp. 1-10, San Diego, CA, USA , 2019.
- [TS33.401] Zhang, D., "3GPP System Architecture Evolution (SAE); Security architecture", T33.401, 3GPP , 2012.
- [CRO2016] Crowe, B., "Proposed AeroMACS PKI Specification is a Model for Global and National Aeronautical PKI Deployments", WiMAX Forum at 16th Integrated Communications, Navigation and Surveillance Conference (ICNS), pp. 1-19, New York, NY, USA , 2016.
- [MAE2020] Maeurer, N., Graeupl, T., and C. Schmitt, "Comparing Different Diffie-Hellman Key Exchange Flavors for LDACS", IEEE/AIAA 39th Digital Avionics Systems Conference (DASC), pp. 1-10, San Antonio, TX, USA , 2020.
- [STR2016] Strohmeier, M., Schaefer, M., Pinheiro, R., Lenders, V., and I. Martinovic, "On Perception and Reality in Wireless Air Traffic Communication Security", IEEE Transactions on Intelligent Transportation Systems, 18(6), pp. 1338-1357, New York, NY, USA , 2016.

- [BIL2017] Bilzhause, A., Belgacem, B., Mostafa, M., and T. Graeupl, "Datalink Security in the L-band Digital Aeronautical Communications System (LDACS) for Air Traffic Management", IEEE Aerospace and Electronic Systems Magazine, 32(11), pp. 22-33, New York, NY, USA , 2017.
- [MAE20181] Maeurer, N. and A. Bilzhause, "Paving the Way for an IT Security Architecture for LDACS: A Datalink Security Threat and Risk Analysis", IEEE Integrated Communications, Navigation, Surveillance Conference (ICNS), pp. 1-11, New York, NY, USA , 2018.
- [FAA2020] FAA, "Federal Aviation Administration. ADS-B Privacy.", August 2020, <<https://www.faa.gov/nextgen/equipadsb/privacy/>>.
- [GNU2012] GNU Radio project, "GNU radio", August 2012, <<http://gnuradio.org>>.
- [SIT2020] SITA, "Societe Internationale de Telecommunications Aeronautiques", August 2020, <<https://www.sita.aero/>>.
- [ARI2020] ARINC, "Aeronautical Radio Incorporated", August 2020, <<https://www.aviation-ia.com/>>.
- [DO350A] RTCA SC-214, "Safety and Performance Standard for Baseline 2 ATS Data Communications (Baseline 2 SPR Standard)", May 2016, <<https://standards.globalspec.com/std/10003192/rtca-do-350-volume-1-2>>.
- [RAW-TECHNOS]
Thubert, P., Cavalcanti, D., Vilajosana, X., Schmitt, C., and J. Farkas, "Reliable and Available Wireless Technologies", Work in Progress, Internet-Draft, draft-ietf-raw-technologies-00, 20 October 2020, <<https://tools.ietf.org/html/draft-ietf-raw-technologies-00>>.
- [RAW-USE-CASES]
Papadopoulos, G., Thubert, P., Theoleyre, F., and C. Bernardos, "RAW use cases", Work in Progress, Internet-Draft, draft-ietf-raw-use-cases-00, 23 October 2020, <<https://tools.ietf.org/html/draft-ietf-raw-use-cases-00>>.

Appendix A. Selected Information from DO-350A

This appendix includes the continuity, availability, and integrity requirements interesting for LDACS defined in [DO350A].

The following terms are used here:

CPDLC Controller Pilot Data Link Communication
 DT Delivery Time (nominal) value for RSP
 ET Expiration Time value for RCP
 FH Flight Hour
 MA Monitoring and Alerting criteria
 OT Overdue Delivery Time value for RSP
 RCP Required Communication Performance
 RSP Required Surveillance Performance
 TT Transaction Time (nominal) value for RCP

	ECP 130	ECP 130
Parameter	ET	TT95%
Transaction Time (sec)	130	67
Continuity	0.999	0.95
Availability	0.989	0.989
Integrity	1E-5 per FH	1E-5 per FH

Table 1: CPDLC Requirements for ECP

	RCP 240	RCP 240	RCP 400	RCP 400
Parameter	ET	TT95%	ET	TT95%
Transaction Time (sec)	240	210	400	350
Continuity	0.999	0.95	0.999	0.95
Availability	0.989 (safety)	0.989 (efficiency)	0.989	0.989
Integrity	1E-5 per FH	1E-5 per FH	1E-5 per FH	1E-5 per FH

Table 2: CPDLC Requirements for RCP

RCP Monitoring and Alerting Criteria in case of CPDLC:

- MA-1: The system SHALL be capable of detecting failures and configuration changes that would cause the communication service no longer meet the RCP specification for the intended use.
- MA-2: When the communication service can no longer meet the RCP specification for the intended function, the flight crew and/or the controller SHALL take appropriate action.

	RSP 160	RSP 160	RSP 180	RSP 180	RSP 400	RSP 400
Parameter	OT	DT95%	OT	DT95%	OT	DT95%
Transaction Time (sec)	160	90	180	90	400	300
Continuity	0.999	0.95	0.999	0.95	0.999	0.95
Availability	0.989	0.989	0.989 (safety)	0.989 (efficiency)	0.989	0.989
Integrity	1E-5 per FH	1E-5 per FH	1E-5 per FH	1E-5 per FH	1E-5 per FH	1E-5 per FH

Table 3: ADS-C Requirements

RCP Monitoring and Alerting Criteria:

- MA-1: The system SHALL be capable of detecting failures and configuration changes that would cause the ADS-C service no longer meet the RSP specification for the intended function.
- MA-2: When the ADS-C service can no longer meet the RSP specification for the intended function, the flight crew and/or the controller SHALL take appropriate action.

Authors' Addresses

Nils Maeurer (editor)
German Aerospace Center (DLR)
Muenchner Strasse 20
82234 Wessling
Germany

Email: Nils.Maeurer@dlr.de

Thomas Graeupl (editor)
German Aerospace Center (DLR)
Muenchner Strasse 20
82234 Wessling
Germany

Email: Thomas.Graeupl@dlr.de

Corinna Schmitt (editor)
Research Institute CODE, UniBwM
Werner-Heisenberg-Weg 28
85577 Neubiberg
Germany

Email: corinna.schmitt@unibw.de

RAW
Internet-Draft
Intended status: Informational
Expires: 22 September 2022

N. Maeurer, Ed.
T. Graeupl, Ed.
German Aerospace Center (DLR)
C. Schmitt, Ed.
Research Institute CODE, UniBwM
21 March 2022

L-band Digital Aeronautical Communications System (LDACS)
draft-ietf-raw-ldacs-10

Abstract

This document gives an overview of the architecture of the L-band Digital Aeronautical Communications System (LDACS), which provides a secure, scalable and spectrum efficient terrestrial data link for civil aviation. LDACS is a scheduled, reliable multi-application cellular broadband system with support for IPv6. LDACS provides a data link for IPv6 network-based aircraft guidance. High reliability and availability for IP connectivity over LDACS, as well as security, are therefore essential.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 22 September 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights

and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
2. Terminology	5
3. Motivation and Use Cases	6
3.1. Voice Communications Today	7
3.2. Data Communications Today	7
4. Provenance and Documents	8
5. Applicability	9
5.1. Advances Beyond the State-of-the-Art	9
5.1.1. Priorities	9
5.1.2. Security	10
5.1.3. High Data Rates	10
5.2. Application	10
5.2.1. Air/Ground Multilink	10
5.2.2. Air/Air Extension for LDACS	11
5.2.3. Flight Guidance	11
5.2.4. Business Communications of Airlines	12
5.2.5. LDACS-based Navigation	12
6. Requirements	13
7. Characteristics	14
7.1. LDACS Sub-Network	14
7.2. Topology	15
7.3. LDACS Protocol Stack	16
7.3.1. LDACS Physical Layer	17
7.3.2. LDACS Data Link Layer	18
7.3.3. LDACS Sub-Network Layer and Protocol Services	19
7.4. LDACS Mobility	20
8. Reliability and Availability	20
8.1. Below Layer 1	20
8.2. Layer 1 and 2	20
8.3. Beyond Layer 2	23
9. Security	23
9.1. Security in Wireless Digital Aeronautical Communications	24
9.2. LDACS Requirements	25
9.3. LDACS Security Objectives	25
9.4. LDACS Security Functions	26
9.5. LDACS Security Architecture	26
9.5.1. Entities	26
9.5.2. Entity Identification	27
9.5.3. Entity Authentication and Key Establishment	27

9.5.4. Message-in-transit Confidentiality, Integrity and Authenticity	28
10. IANA Considerations	28
11. Acknowledgements	28
12. Normative References	28
13. Informative References	28
Appendix A. Selected Information from DO-350A	34
Authors' Addresses	36

1. Introduction

One of the main pillars of the modern Air Traffic Management (ATM) system is the existence of a communications infrastructure that enables efficient aircraft control and safe aircraft separation in all phases of flight. Current systems are technically mature but suffering from the Very High Frequency (VHF) band's increasing saturation in high- density areas and the limitations posed by analogue radio communications. Therefore, aviation globally, and the European Union (EU) in particular, strives for a sustainable modernization of the aeronautical communications infrastructure.

This modernization is realized in two steps: (1) the transition of communications datalinks from analogue to digital technologies and, (2) the introduction of IPv6 based networking protocols in aeronautical networks [RFC4291], [RFC7136], [ICAO2015].

Step (1) is realized via ATM communications transitioning from analogue VHF voice [KAMA2010] to more spectrum efficient digital data communication. For terrestrial communications the European ATM Master Plan foresees this transition to be realized by the development of the L-band Digital Aeronautical Communications System (LDACS). Since central Europe has been identified as the area of the world, that suffers the most from increased saturation of the VHF band, the initial roll-out of LDACS will likely start there, and continue to other increasingly saturated zones as the east- and west-cost of the US and parts of Asia [ICAO2018].

Technically LDACS enables IPv6 based air- ground communication related to aviation safety and regularity of flight [ICAO2015]. Passenger communication and similar services are not supported, since only communications related to "safety and regularity of flight" are permitted in protected aviation frequency bands. The particular challenge is that no additional frequencies can be made available for terrestrial aeronautical communication. It was thus necessary to develop co-existence mechanism/procedures to enable the interference free operation of LDACS in parallel with other aeronautical services/systems in the protected frequency band. Since LDACS will be used for aircraft guidance, high reliability and availability for IP connectivity over LDACS are essential.

Step (2) is a strategy for the worldwide roll-out of IPv6 capable digital aeronautical inter-networking. This is called the Aeronautical Telecommunications Network (ATN)/Internet Protocol Suite (IPS) (hence, ATN/IPS). It is specified in the International Civil Aviation Organization (ICAO) document Doc 9896 [ICAO2015], the Radio Technical Commission for Aeronautics (RTCA) document DO-379 [RTCA2019], the European Organization for Civil Aviation Equipment (EUROCAE) document ED-262 [EURO2019], and the Aeronautical Radio Incorporated (ARINC) document P858 [ARI2021]. LDACS is subject to these regulations since it provides access subnets to the ATN/IPS.

ICAO has chosen IPv6 as basis for the ATN/IPS mostly for historical reasons, since a previous architecture based on ISO/OSI protocols, the ATN/OSI, failed in the market place.

In the context of safety-related communications, LDACS will play a major role in future ATM. ATN/IPS datalinks will provide diversified terrestrial and space-based connectivity in a multi-link concept, called the Future Communications Infrastructure (FCI) [VIR2021]. From a technical point of view the FCI will realize airborne multi-homed IPv6 networks connected to a global ground network via at least two independent communication technologies. This is considered in more detail in related IETF work in progress [I-D.haindl-lisp-gb-atn] [I-D.ietf-rtwg-atn-bgp].

In the context of the Reliable and Available Wireless (RAW) working group, developing options, such as intelligent switching between datalinks, for reliably delivering content from and to endpoints, is foreseen. As LDACS is part of such a concept, the work of RAW is immediately applicable. In general, with the aeronautical communications system transitioning to ATN/IPS, and data being transported via IPv6, closer cooperation and collaboration between the aeronautical and IETF community is desirable.

LDACS standardization within the framework of ICAO started in December 2016. The ICAO standardization group has produced an initial Standards and Recommended Practices (SARPS) document [ICA2018]. It defines the general characteristics of LDACS. The ICAO standardization group plans to produce an ICAO technical manual - the ICAO equivalent to a technical standard - within the next years. Generally, the group is open to input from all sources and encourages cooperation between the aeronautical and the IETF community.

2. Terminology

The following terms are used in the context of RAW in this document:

A/A Air/Air
A/G Air/Ground
A2G Air-to-Ground
ACARS Aircraft Communications Addressing and Reporting System
ADS-B Automatic Dependent Surveillance - Broadcast
ADS-C Automatic Dependent Surveillance - Contract
AeroMACS Aeronautical Mobile Airport Communications System
ANSP Air Traffic Network Service Provider
AOC Aeronautical Operational Control
AR Access Router
ARINC Aeronautical Radio, Incorporated
ARQ Automatic Repeat reQuest
AS Aircraft Station
ATC Air Traffic Control
ATM Air Traffic Management
ATN Aeronautical Telecommunication Network
ATS Air Traffic Service
BCCH Broadcast Channel
CCCH Common Control Channel
CM Context Management
CNS Communication Navigation Surveillance
COTS Commercial Off-The-Shelf
CPDLC Controller Pilot Data Link Communications
CRL Certificate Revocation List
CSP Communications Service Provider
DCCH Dedicated Control Channel
DCH Data Channel
DiffServ Differentiated Services
DLL Data Link Layer
DLS Data Link Service
DME Distance Measuring Equipment
DSB-AM Double Side-Band Amplitude Modulation
DTLS Datagram Transport Layer Security
EUROCAE European Organization for Civil Aviation Equipment

FAA Federal Aviation Administration
FCI Future Communications Infrastructure
FDD Frequency Division Duplex
FL Forward Link
GANP Global Air Navigation Plan
GBAS Ground Based Augmentation System
GNSS Global Navigation Satellite System
GS Ground-Station
G2A Ground-to-Air
HF High Frequency
ICAO International Civil Aviation Organization
IP Internet Protocol
IPS Internet Protocol Suite
kbit/s kilobit per second
LDACS L-band Digital Aeronautical Communications System
LLC Logical Link Control
LME LDACS Management Entity
MAC Medium Access Control
MF Multi Frame
OFDM Orthogonal Frequency-Division Multiplexing
OFDMA Orthogonal Frequency-Division Multiplexing Access
OSI Open Systems Interconnection
PHY Physical Layer
QPSK Quadrature Phase-Shift Keying
RACH Random Access Channel
RL Reverse Link
RTCA Radio Technical Commission for Aeronautics
SARPS Standards and Recommended Practices
SDR Software Defined Radio
SESAR Single European Sky ATM Research
SF Super-Frame
SNP Sub-Network Protocol
VDLm2 VHF Data Link mode 2
VHF Very High Frequency
VI Voice Interface

3. Motivation and Use Cases

Aircraft are currently connected to Air Traffic Control (ATC) and Aeronautical Operational Control (AOC) services via voice and data communications systems through all phases of flight. ATC refers to communication for flight guidance. AOC is a generic term referring to the business communication of airlines. It refers to the mostly proprietary exchange of data between the aircraft of the airline, its operation centers, and its service partners. ARINC document 633 was developed and first released in 2007 [ARI2019] with the goal to standardize these messages for interoperability, e.g., messages

between the airline and fueling or de-icing companies. Within the airport terminal, connectivity is focused on high bandwidth communications, while during en-route, high reliability, robustness, and range is the main focus. Voice communications may use the same or different equipment as data communications systems. In the following, the main differences between voice and data communications capabilities are summarized. The assumed use cases for LDACS complements the list of use cases stated in [RAW-USE-CASES] and the list of reliable and available wireless technologies presented in [RAW-TECHNOS].

3.1. Voice Communications Today

Voice links are used for Air/Ground (A/G) and Air/Air (A/A) communications. The communications equipment is either ground-based working in the High Frequency (HF) or VHF frequency band or satellite-based. All VHF and HF voice communications are operated via open broadcast channels without authentication, encryption or other protective measures. The use of well-proven communications procedures via broadcast channels, such as phraseology or read-backs, requiring well-trained personnel, help to enhance the safety of communications, but does not replace necessary cryptographical security mechanisms. The main voice communications media is still the analogue VHF Double Side-Band Amplitude Modulation (DSB-AM) communications technique, supplemented by HF single side-band amplitude modulation and satellite communications for remote and oceanic regions. DSB-AM has been in use since 1948, works reliably and safely, and uses low-cost communication equipment. These are the main reasons why VHF DSB-AM communications are still in use, and it is likely that this technology will remain in service for many more years. This however, results in current operational limitations and impediments in deploying new ATM applications, such as flight-centric operation with point-to-point communications between pilots and air traffic control officers. [BOE2019]

3.2. Data Communications Today

Like for voice, data communications into the cockpit, are currently provided by ground-based equipment operating either on HF or VHF radio bands or by legacy satellite systems. All these communication systems are using narrowband radio channels with a data throughput capacity in the order of kilobits per second. While the aircraft is on ground, some additional communications systems are available, like the Aeronautical Mobile Airport Communications System (AeroMACS) or public cellular networks, operating in the Airport (APT) domain and able to deliver broadband communications capability. [BOE2019]

For regulatory reasons, the data communications networks, used for the transmission of data relating to the safety and regularity of flight, must be strictly isolated from those providing entertainment services to passengers. This leads to a situation that the flight crews are supported by narrowband services during flight while passengers have access to inflight broadband services. The current HF and VHF data links cannot provide broadband services now or in the future, due to the lack of available spectrum. This technical shortcoming is becoming a limitation to enhanced ATM operations, such as trajectory-based operations and 4D trajectory negotiations. [BOE2019]

Satellite-based communications are currently under investigation and enhanced capabilities are under development which will be able to provide inflight broadband services and communications supporting the safety and regularity of flight. In parallel the ground-based broadband data link technology LDACS is being standardized by ICAO and has recently shown its maturity during flight tests [MAE2021] [BEL2021]. The LDACS technology is scalable, secure and spectrum efficient and provides significant advantages to the users and service providers. It is expected that both - satellite systems and LDACS - will be deployed to support the future aeronautical communication needs as envisaged by the ICAO Global Air Navigation Plan (GNAP). [BOE2019]

4. Provenance and Documents

The development of LDACS has already made substantial progress in the Single European Sky ATM Research (SESAR) framework and is currently being continued in the follow-up program SESAR2020 [RIH2018]. A key objective of these activities is to develop, implement and validate a modern aeronautical data link able to evolve with aviation needs over long-term. To this end, an LDACS specification has been produced [GRA2020] and is continuously updated; transmitter demonstrators were developed to test the spectrum compatibility of LDACS with legacy systems operating in the L-band [SAJ2014]; and the overall system performance was analyzed by computer simulations, indicating that LDACS can fulfil the identified requirements [GRA2011].

Up to now LDACS standardization has been focused on the development of the physical layer and the data link layer. Only recently have higher layers have come into the focus of the LDACS development activities. There is currently no "IPv6 over LDACS" specification publicly available; however, SESAR2020 has started the testing of IPv6-based LDACS testbeds.

The IPv6 architecture for the aeronautical telecommunication network is called the FCI. The FCI will support quality of service, diversity, and mobility under the umbrella of the "multi-link concept". This work is led by ICAO Communication Panel working group WG-I.

In addition to standardization activities several industrial LDACS prototypes have been built. One set of LDACS prototypes has been evaluated in flight trials confirming the theoretical results predicting the system performance [GRA2018] [MAE2021] [BEL2021].

5. Applicability

LDACS is a multi-application cellular broadband system capable of simultaneously providing various kinds of Air Traffic Services (ATS) including ATS-B3, and AOC communications services from deployed Ground-Stations (GS). The physical layer and data link layer of LDACS are optimized for controller-pilot data link communications, but the system also supports digital air-ground voice communications.

LDACS supports communications in all airspaces (airport, terminal maneuvering area, and en-route), and on the airport surface. The physical LDACS cell coverage is effectively de-coupled from the operational coverage required for a particular service. This is new in aeronautical communications. Services requiring wide-area coverage can be installed at several adjacent LDACS cells. The handover between the involved LDACS cells is seamless, automatic, and transparent to the user. Therefore, the LDACS communications concept enables the aeronautical communication infrastructure to support future dynamic airspace management concepts.

5.1. Advances Beyond the State-of-the-Art

LDACS offers several capabilities, not yet provided in contemporarily deployed aeronautical communications systems.

5.1.1. Priorities

LDACS is able to manage service priorities, an important feature not available in some of the current data link deployments. Thus, LDACS guarantees bandwidth availability, low latency, and high continuity of service for safety critical ATS applications while simultaneously accommodating less safety-critical AOC services.

5.1.2. Security

LDACS is a secure data link with built-in security mechanisms. It enables secure data communications for ATS and AOC services, including secured private communications for aircraft operators and Air traffic Network Service Providers (ANSP). This includes concepts for key and trust management, mutual authentication and key establishment protocols, key derivation measures, user and control message-in-transit protection, secure logging and availability and robustness measures [MAE20182] [MAE2021].

5.1.3. High Data Rates

The user data rate of LDACS is 315 kbit/s to 1428 kbit/s on the Forward Link (FL) for the Ground-to-Air (G2A) connection, and 294 kbit/s to 1390 kbit/s on the Reverse Link (RL) for the Air-to-Ground (A2G) connection, depending on coding and modulation. This is up to two orders of magnitude greater than current terrestrial digital aeronautical communications systems, such as the VHF Data Link mode 2 (VDLm2), provide [ICAO2019] [GRA2020].

5.2. Application

LDACS will be used by several aeronautical applications ranging from enhanced communications protocol stacks (multi-homed mobile IPv6 networks in the aircraft and potentially ad-hoc networks between aircraft) to broadcast communication applications (sending Ground Based Augmentation System (GBAS) correction data) and integration with other service domains (using the communications signal for navigation) [MAE20211].

5.2.1. Air/Ground Multilink

It is expected that LDACS, together with upgraded satellite-based communications systems, will be deployed within the FCI and constitute one of the main components of the multilink concept within the FCI.

Both technologies, LDACS and satellite systems, have their specific benefits and technical capabilities which complement each other. Especially, satellite systems are well-suited for large coverage areas with less dense air traffic, e.g. oceanic regions. LDACS is well-suited for dense air traffic areas, e.g., continental areas or hot-spots around airports and terminal airspace. In addition, both technologies offer comparable data link capacity and, thus, are well-suited for redundancy, mutual back-up, or load balancing.

Technically the FCI multilink concept will be realized by multi-homed mobile IPv6 networks in the aircraft. The related protocol stack is currently under development by ICAO, within SESAR, and the IETF [I-D.haindl-lisp-gb-atn] [I-D.ietf-rtgwg-atn-bgp].

5.2.2. Air/Air Extension for LDACS

A potential extension of the multi-link concept is its extension to the integration of ad-hoc networks between aircraft.

Direct A/A communication between aircraft in terms of ad-hoc data networks are currently considered a research topic since there is no immediate operational need for it, although several possible use cases are discussed (Automatic Dependent Surveillance - Broadcast (ADS-B), digital voice, wake vortex warnings, and trajectory negotiation) [BEL2019]. It should also be noted, that currently deployed analog VHF voice radios support direct voice communication between aircraft, making a similar use case for digital voice plausible.

LDACS A/A is currently not part of the standardization process and will not be covered within this document.

5.2.3. Flight Guidance

The FCI (and therefore LDACS) is used to provide flight guidance. This is realized using three applications:

1. Context Management (CM): The CM application manages the automatic logical connection to the ATC center currently responsible to guide the aircraft. Currently this is done by the air crew manually changing VHF voice frequencies according to the progress of the flight. The CM application automatically sets up equivalent sessions.
2. Controller Pilot Data Link Communications (CPDLC): The CPDLC application provides the air crew with the ability to exchange data messages similar to text messages with the currently responsible ATC center. The CPDLC application takes over most of the communication currently performed over VHF voice and enables new services that do not lend themselves to voice communication (i.e., trajectory negotiation).
3. Automatic Dependent Surveillance - Contract (ADS-C): ADS-C reports the position of the aircraft to the currently active ATC center. Reporting is bound to "contracts", i.e., pre-defined events related to the progress of the flight (i.e., the trajectory). ADS-C and CPDLC are the primary applications used for implementing in-flight trajectory management.

CM, CPDLC, and ADS-C are available on legacy datalinks, but are not widely deployed and with limited functionality.

Further ATC applications may be ported to use the FCI or LDACS as well. A notable application is GBAS for secure, automated landings: The Global Navigation Satellite System (GNSS) based GBAS is used to improve the accuracy of GNSS to allow GNSS based instrument landings. This is realized by sending GNSS correction data (e.g., compensating ionospheric errors in the GNSS signal) to the aircraft's GNSS receiver via a separate data link. Currently the VDB data link is used. VDB is a narrow-band single-purpose datalink without advanced security only used to transmit GBAS correction data. This makes VDB a natural candidate for replacement by LDACS [MAE20211].

5.2.4. Business Communications of Airlines

In addition to air traffic services, AOC services are transmitted over LDACS. AOC is a generic term referring to the business communication of airlines, between the airlines and service partners on the ground and their own aircraft in the air. Regulatory-wise, this is considered related to safety and regularity of flight and may therefore be transmitted over LDACS. AOC communication is considered the main business case for LDACS communications service providers since modern aircraft generate significant amounts of data (e.g., engine maintenance data).

5.2.5. LDACS-based Navigation

Beyond communications, radio signals can always also be used for navigation. This fact is used for the LDACS navigation concept.

For future aeronautical navigation, ICAO recommends the further development of GNSS based technologies as primary means for navigation. Due to the large separation between navigational satellites and aircraft, the power of the GNSS signals received by the aircraft is, however, very low. As a result, GNSS disruptions might occasionally occur due to unintentional interference, or intentional jamming. Yet the navigation services must be available with sufficient performance for all phases of flight. Therefore, during GNSS outages, or blockages, an alternative solution is needed. This is commonly referred to as Alternative Positioning, Navigation, and Timing (APNT).

One of such APNT solutions consists of exploiting the built-in navigation capabilities of LDACS operation. That is, the normal operation of LDACS for ATC and AOC communications would also directly enable the aircraft to navigate and obtain a reliable timing reference from the LDACS GSs.

LDACS navigation has already been demonstrated in practice in two flight measurement campaigns [SHU2013] [BEL2021] [MAE2021]. .

6. Requirements

The requirements for LDACS are mostly defined by its application area: Communications related to safety and regularity of flight.

A particularity of the current aeronautical communication landscape is that it is heavily regulated. Aeronautical data links (for applications related to safety and regularity of flight) may only use spectrum licensed to aviation and data links endorsed by ICAO. Nation states can change this locally, however, due to the global scale of the air transportation system, adherence to these practices is to be expected.

Aeronautical data links for the ATN are therefore expected to remain in service for decades. The VDLm2 data link currently used for digital terrestrial internetworking was developed in the 1990ies (the use of the Open Systems Interconnection (OSI) stack indicates that as well). VDLm2 is expected to be used at least for several decades. In this respect aeronautical communications (for applications related to safety and regularity of flight) is more comparable to industrial applications than to the open Internet.

Internetwork technology is already installed in current aircraft. Current ATS applications use either Aircraft Communications Addressing and Reporting System (ACARS) or the OSI stack. The objective of the development effort of LDACS, as part of the FCI, is to replace legacy OSI stack and proprietary ACARS internetwork technologies with industry standard IP technology. It is anticipated that the use of Commercial Off-The-Shelf (COTS) IP technology mostly applies to the ground network. The avionics networks on the aircraft will likely be heavily modified versions of Ethernet or proprietary.

AOC applications currently mostly use the same stack (although some applications, like the graphical weather service may use the commercial passenger network). This creates capacity problems (resulting in excessive amounts of timeouts) since the underlying terrestrial data links do not provide sufficient bandwidth (i.e., with VDLm2 currently in the order of 10 kbit/s). The use of non-aviation specific data links is considered a security problem. Ideally the aeronautical IP internetwork, hence the ATN over which only communications related to safety and regularity of flight is handled, and the Internet should be completely separated at Layer 3.

The objective of LDACS is to provide a next generation terrestrial data link designed to support IP addressing and provide much higher bandwidth to avoid the currently experienced operational problems.

The requirement for LDACS is therefore to provide a terrestrial high-throughput data link for IP internetworking in the aircraft.

In order to fulfil the above requirement LDACS needs to be interoperable with IP (and IP-based services like Voice-over-IP) at the gateway connecting the LDACS network to other aeronautical ground networks (i.e., the ATN). On the avionics side, in the aircraft, aviation specific solutions are to be expected.

In addition to these functional requirements, LDACS and its IP stack need to fulfil the requirements defined in RTCA DO-350A/EUROCAE ED-228A [DO350A]. This document defines continuity, availability, and integrity requirements at different scopes for each air traffic management application (CPDLC, CM, and ADS-C). The scope most relevant to IP over LDACS is the Communications Service Provider (CSP) scope.

Continuity, availability, and integrity requirements are defined in [DO350A] volume 1 Table 5-14, and Table 6-13. Appendix A presents the required information.

In a similar vein, requirements to fault management are defined in the same tables.

7. Characteristics

LDACS will become one of several wireless access networks connecting aircraft to the ATN implemented by the FCI.

The current LDACS design is focused on the specification of layer one and two. However, for the purpose of this work, only layer two details are discussed here.

Achieving the stringent continuity, availability, and integrity requirements defined in [DO350A] will require the specification of layer 3 and above mechanisms (e.g. reliable crossover at the IP layer). Fault management mechanisms are similarly undefined.

7.1. LDACS Sub-Network

An LDACS sub-network contains an Access Router (AR) and several GS, each of them providing one LDACS radio cell.

User plane interconnection to the ATN is facilitated by the AR peering with an A/G Router connected to the ATN.

The internal control plane of an LDACS sub-network interconnects the GSs. An LDACS sub-network is illustrated in Figure 1.

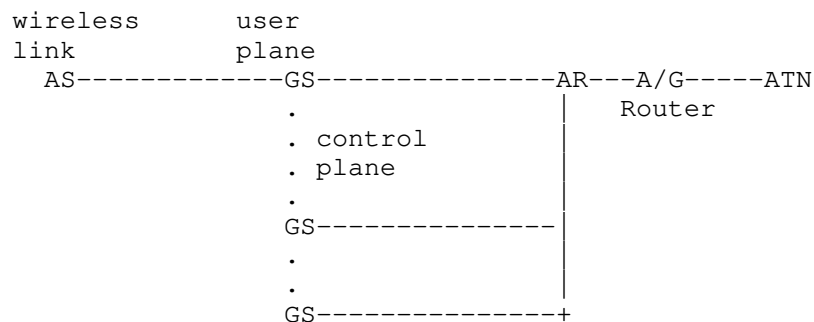


Figure 1: LDACS sub-network with three GSs and one AS

7.2. Topology

LDACS is a cellular point-to-multipoint system. It assumes a star-topology in each cell where Aircraft Stations (AS) belonging to aircraft within a certain volume of space (the LDACS cell) is connected to the controlling GS. The LDACS GS is a centralized instance that controls LDACS A/G communications within its cell. The LDACS GS can simultaneously support multiple bi-directional communications to the ASs under its control. LDACS's GSs themselves are connected to each other and the AR.

Prior to utilizing the system an aircraft has to register with the controlling GS to establish dedicated logical channels for user and control data. Control channels have statically allocated resources, while user channels have dynamically assigned resources according to the current demand. Logical channels exist only between the GS and the AS.

7.3. LDACS Protocol Stack

The protocol stack of LDACS is implemented in the AS and GS: It consists of the Physical Layer (PHY) with five major, functional blocks above it. Four are placed in the Data Link Layer (DLL) of the AS and GS: (1) Medium Access Control (MAC) Layer, (2) Voice Interface (VI), (3) Data Link Service (DLS), and (4) LDACS Management Entity (LME). The last entity resides within the sub-network layer: the Sub-Network Protocol (SNP). The LDACS network is externally connected to voice units, radio control units, and the ATN network layer.

LDACS is considered an ATN/IPS radio access technology, from the view of ICAO's regulatory framework. Hence, the interface between ATN and LDACS must be IPv6 based, as regulatory documents, such as ICAO Doc 9896 [ICAO2015] and DO-379 [RTCA2019] clearly foresee that. The translation between IPv6 layer and SNP layer is currently subject of ongoing standardization efforts and at the time of writing not finished yet.

Figure 2 shows the protocol stack of LDACS as implemented in the AS and GS. Acronyms used here are introduced throughout the upcoming sections.




Figure 2: LDACS protocol stack in AS and GS

7.3.1. LDACS Physical Layer

The physical layer provides the means to transfer data over the radio channel. The LDACS GS supports bi-directional links to multiple aircraft under its control. The FL direction at the G2A connection and the RL direction at the A2G connection are separated by Frequency Division Duplex (FDD). FL and RL use a 500 kHz channel each. The GS transmits a continuous stream of Orthogonal Frequency-Division Multiplexing Access (OFDM) symbols on the FL. In the RL different aircraft are separated in time and frequency using Orthogonal Frequency-Division Multiple Access (OFDMA). Aircraft thus transmit discontinuously on the RL via short radio bursts sent in precisely defined transmission opportunities allocated by the GS.

7.3.2. LDACS Data Link Layer

The data-link layer provides the necessary protocols to facilitate concurrent and reliable data transfer for multiple users. The LDACS data link layer is organized in two sub-layers: The medium access sub-layer and the Logical Link Control (LLC) sub-layer. The medium access sub-layer manages the organization of transmission opportunities in slots of time and frequency. The LLC sub-layer provides acknowledged point-to-point logical channels between the aircraft and the GS using an Automatic Repeat reQuest (ARQ) protocol. LDACS supports also unacknowledged point-to-point channels and G2A Broadcast transmission.

7.3.2.1. Medium Access Control (MAC) Services

The MAC time framing service provides the frame structure necessary to realize slot-based time-division multiplex-access on the physical link. It provides the functions for the synchronization of the MAC framing structure and the PHY Layer framing. The MAC time framing provides a dedicated time slot for each logical channel.

The MAC sub-layer offers access to the physical channel to its service users. Channel access is provided through transparent logical channels. The MAC sub-layer maps logical channels onto the appropriate slots and manages the access to these channels. Logical channels are used as interface between the MAC and LLC sub-layers.

7.3.2.2. Data Link Service (DLS) Services

The DLS provides acknowledged and unacknowledged (including broadcast and packet mode voice) bi-directional exchange of user data. If user data is transmitted using the acknowledged DLS, the sending DLS entity will wait for an acknowledgement from the receiver. If no acknowledgement is received within a specified time frame, the sender may automatically try to retransmit its data. However, after a certain number of failed retries, the sender will suspend further retransmission attempts and inform its client of the failure.

The DLS uses the logical channels provided by the MAC:

1. A GS announces its existence and access parameters in the Broadcast Channel (BCCH).
2. The Random Access Channel (RACH) enables AS to request access to an LDACS cell.
3. In the FL the Common Control Channel (CCCH) is used by the GS to grant access to data channel resources.
4. The reverse direction is covered by the RL, where ASs need to

request resources before sending. This happens via the Dedicated Control Channel (DCCH).

5. User data itself is communicated in the Data Channel (DCH) on the FL and RL.

Access to the FL and RL data channel is granted by the scheduling mechanism implemented in the LME discussed below.

7.3.2.3. Voice Interface (VI) Services

The VI provides support for virtual voice circuits. Voice circuits may either be set-up permanently by the GS (e.g., to emulate voice party line) or may be created on demand.

7.3.2.4. LDACS Management Entity (LME) Services

The mobility management service in the LME provides support for registration and de-registration (cell entry and cell exit), scanning RF channels of neighboring cells and handover between cells. In addition, it manages the addressing of aircraft within cells.

The resource management service provides link maintenance (power, frequency and time adjustments), support for adaptive coding and modulation, and resource allocation.

The resource management service accepts resource requests from/for different AS and issues resource allocations accordingly. While the scheduling algorithm is not specified and a point of possible vendor differentiation, it is subject to the following requirements:

1. Resource scheduling must provide channel access according to the priority of the request
2. Resource scheduling must support "one-time" requests.
3. Resource scheduling must support "permanent" requests that reserve a resource until the request is canceled e.g. for digital voice circuits.

7.3.3. LDACS Sub-Network Layer and Protocol Services

Lastly, the SNP handles the transition from IPv6 packets to LDACS internal packet structures. This work is ongoing and not part of this document. The DLS provides functions required for the transfer of user plane data and control plane data over the LDACS sub-network. The security service provides functions for secure user data communication over the LDACS sub-network. Note that the SNP security service applies cryptographic measures as configured by the GS.

7.4. LDACS Mobility

LDACS supports layer 2 handovers to different LDACS cells. Handovers may be initiated by the aircraft (break-before-make) or by the GS (make-before-break). Make-before-break handovers are only supported between GSs connected to each other.

External handovers between non-connected LDACS sub-networks or different aeronautical data links are handled by the FCI multi-link concept.

8. Reliability and Availability

8.1. Below Layer 1

Below Layer 2, aeronautics usually relies on hardware redundancy. To protect availability of the LDACS link, an aircraft equipped with LDACS will have access to two L-band antennae with triple redundant radio systems as required for any safety relevant aeronautical systems by ICAO.

8.2. Layer 1 and 2

LDACS has been designed with applications related to the safety and regularity of flight in mind. It has therefore been designed as a deterministic wireless data link (as far as this is possible).

Based on channel measurements of the L-band channel LDACS was designed from the PHY layer up with robustness in mind. Channel measurements of the L-band channel [SCH2016] confirmed LDACS to be well adapted to its channel.

In order to maximize the capacity per channel and to optimally use the available spectrum, LDACS was designed as an OFDM-based FDD system, supporting simultaneous transmissions in FL in the G2A connection and RL in the A2G connection. The legacy systems already deployed in the L-band limit the bandwidth of both channels to approximately 500 kHz.

The LDACS physical layer design includes propagation guard times sufficient for the operation at a maximum distance of 200 nautical miles from the GS. In actual deployment, LDACS can be configured for any range up to this maximum range.

The LDACS physical layer supports adaptive coding and modulation for user data. Control data is always encoded with the most robust coding and modulation (FL: Quadrature Phase-Shift Keying (QPSK), coding rate 1/2, RL: QPSK, coding rate 1/3).

LDACS medium access layer on top of the physical layer uses a static frame structure to support deterministic timer management. As shown in Figure 3 and Figure 4, LDACS framing structure is based on Super-Frames (SF) of 240ms duration corresponding to 2000 OFDM symbols. FL and RL boundaries are aligned in time (from the GS perspective) allowing for deterministic slots for control and data channels. This initial AS time synchronization and time synchronization maintenance is based on observing the synchronization symbol pairs that repetitively occur within the FL stream, being sent by the controlling GS [GRA2020].




Figure 3: SF structure for LDACS




Figure 4: MF structure for LDACS

LDACS cell entry is conducted with an initial control message exchange via the RACH and the BCCH.

After cell entry, LDACS medium access is always under the control of the GS of a radio cell. Any medium access for the transmission of user data on a DCH has to be requested with a resource request message stating the requested amount of resources and class of service. The GS performs resource scheduling on the basis of these requests and grants resources with resource allocation messages. Resource request and allocation messages are exchanged over dedicated contention-free control channels (DCCH and CCCH).

The purpose of quality-of-service in LDACS medium access is to provide prioritized medium access at the bottleneck (the wireless link). The signaling of higher layer quality-of-service requirements to LDACS is yet to be defined. A Differentiated Services- (DiffServ) based solution with a small number of priorities is to be expected.

In addition to having full control over resource scheduling, the GS can send forced handover commands for off-loading or channel management, e.g., when the signal quality declines and a more suitable GS is in the AS's reach. With robust resource management of the capacities of the radio channel, reliability and robustness measures are therefore also anchored in the LME.

In addition to radio resource management, the LDACS control channels are also used to send keep-alive messages, when they are not otherwise used. Since the framing of the control channels is deterministic, missing keep-alive messages can thus be immediately detected. This information is made available to the multi-link protocols for fault management.

The protocol used to communicate faults is not defined in the LDACS specification. It is assumed that vendors would use industry standard protocols like the Simple Network Management Protocol or the Network Configuration Protocol, where security permits.

The LDACS data link layer protocol, running on top of the medium access sub-layer, uses ARQ to provide reliable data transmission on the data channel.

It employs selective repeat ARQ with transparent fragmentation and reassembly to the resource allocation size to minimize latency and overhead without losing reliability. It ensures correct order of packet delivery without duplicates. In case of transmission errors, it identifies lost fragments with deterministic timers synced to the medium access frame structure and initiates retransmission.

8.3. Beyond Layer 2

LDACS availability can be increased by appropriately deploying LDACS infrastructure: This means proliferating the number of terrestrial ground stations. However, the scarcity of aeronautical spectrum for data link communication (in the case of LDACS: tens of MHz in the L-band) and the long range (in the case of LDACS: up to 200 nautical miles) make this quite hard. While the deployment of a larger number of small cells is one possible solution, this also suffers from the spectrum scarcity. An additional constraint to consider, is that Distance Measuring Equipment (DME) is the primary user of the aeronautical L-band. That is, any LDACS deployment has to take DME frequency planning into account.

The aeronautical community has therefore decided not to rely on a single communication system or frequency band. It is envisioned to have multiple independent data link technologies in the aircraft (e.g., terrestrial and satellite communications) in addition to legacy VHF voice.

However, as of now, no reliability and availability mechanisms that could utilize the multi-link architecture, have been specified on Layer 3 and above. Even if LDACS has been designed for reliability, the wireless medium presents significant challenges to achieve deterministic properties such as low packet error rate, bounded consecutive losses, and bounded latency. Support for high reliability and availability for IP connectivity over LDACS is certainly highly desirable but needs to be adapted to the specific use case.

9. Security

ICAO Doc 9896 foresees transport layer security [ICAO2015] for all aeronautical data as described in ARINC P858 [ARI2021], most likely realized via Datagram Transport Layer Security (DTLS) [RFC6012] [RFC6347].

LDACS also needs to comply with in-depth security requirements, stated in P858, for the radio access technologies transporting ATN/IPS data [ARI2021]. These requirements imply that LDACS must provide layer 2 security in addition to any higher layer mechanisms.

9.1. Security in Wireless Digital Aeronautical Communications

Aviation will require secure exchanges of data and voice messages for managing the air traffic flow safely through the airspaces all over the world. Historically Communication Navigation Surveillance (CNS) wireless communications technology emerged from military and a threat landscape where inferior technological and financial capabilities of adversaries were assumed [STR2016]. The main communications method for ATC today is still an open analogue voice broadcast within the aeronautical VHF band. Currently, information security is mainly procedural, based by using well-trained personnel and proven communications procedures. This communication method has been in service since 1948. However, since the emergence of civil aeronautical CNS applications in the 70s, and today, the world has changed.

Civil applications have significant lower spectrum available than military applications. This means several military defense mechanisms, such as frequency hopping or pilot symbol scrambling and, thus, a defense-in-depth approach starting at the physical layer, is infeasible for civil systems. With the rise of cheap Software Defined Radios (SDRs), the previously existing financial barrier is almost gone and open source projects such as GNU radio [GNU2021] allow a new type of unsophisticated listeners and possible attackers.

Most CNS technology developed in ICAO relies on open standards, thus syntax and semantics of wireless digital aeronautical communications should be expected to be common knowledge for attackers. With increased digitization and automation of civil aviation, the human as control instance, is being taken gradually out of the loop. Autonomous transport drones or single piloted aircraft demonstrate this trend. However, without profound cybersecurity measures such as authenticity and integrity checks of messages in-transit on the wireless link or mutual entity authentication, this lack of a control instance can prove disastrous. Thus, future digital communications will need additional embedded security features to fulfill modern information security requirements like authentication and integrity. These security features require sufficient bandwidth which is beyond the capabilities of currently deployed VHF narrowband communications systems. For voice and data communications, sufficient data throughput capability is needed to support the security functions while not degrading performance. LDACS is a data link technology with sufficient bandwidth to incorporate security without losing too much user data throughput.

9.2. LDACS Requirements

Overall, there are several business goals for cybersecurity to protect, within the FCI in civil aviation:

1. Safety: The system must sufficiently mitigate attacks, which contribute to safety hazards.
2. Flight regularity: The system must sufficiently mitigate attacks, which contribute to delays, diversions, or cancellations of flights.
3. Protection of business interests: The system must sufficiently mitigate attacks which result in financial loss, reputation damage, disclosure of sensitive proprietary information, or disclosure of personal information.

To further analyze assets and derive threats and thus protection scenarios several threat-and risk analyses were performed for LDACS [MAE20181] , [MAE20191]. These results allowed deriving security scope and objectives from the requirements and the conducted threat-and risk analysis.

9.3. LDACS Security Objectives

Security considerations for LDACS are defined by the official SARPS document by ICAO [ICA2018]:

1. LDACS shall provide a capability to protect the availability and continuity of the system.
2. LDACS shall provide a capability including cryptographic mechanisms to protect the integrity of messages in transit.
3. LDACS shall provide a capability to ensure the authenticity of messages in transit.
4. LDACS should provide a capability for nonrepudiation of origin for messages in transit.
5. LDACS should provide a capability to protect the confidentiality of messages in transit.
6. LDACS shall provide an authentication capability.
7. LDACS shall provide a capability to authorize the permitted actions of users of the system and to deny actions that are not explicitly authorized.
8. If LDACS provides interfaces to multiple domains, LDACS shall provide capability to prevent the propagation of intrusions within LDACS domains and towards external domains.

Currently, a change request for these SARPS aims to limit the "non-repudiation of origin of messages in transit" requirement only to the authentication and key establishment messages at the beginning of every session.

9.4. LDACS Security Functions

These objectives were used to derive several security functions for LDACS required to be integrated in the LDACS cybersecurity architecture: Identification, Authentication, Authorization, Confidentiality, System Integrity, Data Integrity, Robustness, Reliability, Availability, and Key and Trust Management. Several works investigated possible measures to implement these security functions [BIL2017], [MAE20181], [MAE20191].

9.5. LDACS Security Architecture

The requirements lead to a LDACS security model, including different entities for identification, authentication and authorization purposes, ensuring integrity, authenticity and confidentiality of data. A draft of the cybersecurity architecture of LDACS can be found in [ICA2018] and [MAE20182] and respective updates in [MAE20191], [MAE20192], [MAE2020], and most recently [MAE2021].

9.5.1. Entities

A simplified LDACS architectural model requires the following entities: Network operators such as the Societe Internationale de Telecommunications Aeronautiques (SITA) [SIT2020] and ARINC [ARI2020] are providing access to the ground IPS network via an A/G LDACS router. This router is attached to a closed off LDACS access network, which connects via further access routers to the different LDACS cell ranges, each controlled by a GS (serving one LDACS cell), with several interconnected GS spanning a local LDACS access network. Via the A/G wireless LDACS data link AS the aircraft is connected to the ground network and via the aircraft's VI and aircraft's network interface, aircraft's data can be sent via the AS back to the GS, then to the LDACS local access network, access routers, LDACS access network, A/G LDACS router and finally to the ground IPS network [ICAO2015].

9.5.2. Entity Identification

LDACS needs specific identities for the AS, the GS, and the network operator. The aircraft itself can be identified using the ICAO unique address of an aircraft, the call sign of that aircraft or the recently founded privacy ICAO address of the Federal Aviation Administration (FAA) program with the same name [FAA2020]. It is conceivable that the LDACS AS will use a combination of aircraft identification, radio component identification and even operator feature identification to create a unique AS LDACS identification tag. Similar to a 4G's eNodeB serving network identification tag, a GS could be identified using a similar field. The identification of the network operator is again similar to 4G (e.g., E-Plus, AT&T, and TELUS), in the way that the aeronautical network operators are listed (e.g., ARINC [ARI2020] and SITA [SIT2020]).

9.5.3. Entity Authentication and Key Establishment

In order to anchor trust within the system, all LDACS entities connected to the ground IPS network will be rooted in an LDACS specific chain-of-trust and PKI solution, quite similar to AeroMACS's approach [CRO2016]. These certificates, residing at the entities and incorporated in the LDACS PKI, providing proof the ownership of their respective public key, include information about the identity of the owner and the digital signature of the entity that has verified the certificate's content. First, all ground infrastructures must mutually authenticate to each other, negotiate and derive keys and, thus, secure all ground connections. How this process is handled in detail is still an ongoing discussion. However, established methods to secure user plane by IPsec [RFC4301] and IKEv2 [RFC7296] or the application layer via TLS 1.3 [RFC8446] are conceivable. The LDACS PKI with their chain-of-trust approach, digital certificates and public entity keys lay the groundwork for this step. In a second step, the AS with the LDACS radio aboard, approaches an LDACS cell and performs a cell-attachment procedure with the corresponding GS. This procedure consists of (1) the basic cell entry [GRA2020] and (2) a Mutual Authentication and Key Establishment (MAKE) procedure [MAE2021].

Note, that LDACS will foresee multiple security levels. To address the issue of the long service life of LDACS (i.e., possibly >30 years) and the security of current pre-quantum cryptography, these security levels include pre- and post-quantum cryptographic solutions. Limiting security data on the LDACS datalink as much as possible, to reserve as much space for actual user data transmission, is key in the LDACS security architecture, this is also reflected in the underlying cryptography: Pre-quantum solutions will rely on elliptic curves [KOB1987], while post-quantum solutions consider

Falcon [SON2021] [MAE2021] or similar lightweight PQC signature schemes, and SIKE or SABER as key establishment options [SIK2021] [ROY2020].

9.5.4. Message-in-transit Confidentiality, Integrity and Authenticity

The key material from the previous step can then be used to protect LDACS Layer 2 communications via applying encryption and integrity protection measures on the SNP layer of the LDACS protocol stack. As LDACS transports AOC and ATS data, the integrity of that data is most important, while confidentiality only needs to be applied to AOC data to protect business interests [ICA2018]. This possibility of providing low layered confidentiality and integrity protection ensures a secure delivery of user data over the wireless link. Furthermore, it ensures integrity protection of LDACS control data.

10. IANA Considerations

This memo includes no request to IANA.

11. Acknowledgements

Thanks to all contributors to the development of LDACS and ICAO PT-T.

Thanks to Klaus-Peter Hauf, Bart Van Den Einden, and Pierluigi Fantappie for further input to this draft.

Thanks to the Chair for Network Security and the research institute CODE for their comments and improvements.

Thanks to SBA Research Vienna for fruitful discussions on aeronautical communications concerning security incentives for industry and potential economic spillovers.

Thanks to the Aeronautical Communications group at the Institute of Communications and Navigation of the German Aerospace Center (DLR). With that, the authors would like to explicitly thank Miguel Angel Bellido-Manganell and Lukas Marcel Schalk for their thorough feedback.

12. Normative References

13. Informative References

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.

- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<https://www.rfc-editor.org/info/rfc4301>>.
- [RFC6012] Salowey, J., Petch, T., Gerhards, R., and H. Feng, "Datagram Transport Layer Security (DTLS) Transport Mapping for Syslog", RFC 6012, DOI 10.17487/RFC6012, October 2010, <<https://www.rfc-editor.org/info/rfc6012>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.
- [RFC7136] Carpenter, B. and S. Jiang, "Significance of IPv6 Interface Identifiers", RFC 7136, DOI 10.17487/RFC7136, February 2014, <<https://www.rfc-editor.org/info/rfc7136>>.
- [RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October 2014, <<https://www.rfc-editor.org/info/rfc7296>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [GRA2020] Graeupl, T., Rihacek, C., and B. Haindl, "LDACS A/G Specification", SESAR2020 PJ14-02-01 D3.3.030 , 2020, <https://www.ldacs.com/wp-content/uploads/2013/12/SESAR2020_PJ14-W2-60_D3_1_210_Initial_LDACS_AG_Specification_00_01_00-1_0_updated.pdf>.
- [ARI2021] ARINC, "Internet Protocol Suite (IPS) For Aeronautical Safety Services Part 1- Airborne IP System Technical Requirements, ARINC SPECIFICATION 858 P1", June 2021, <<https://standards.globalspec.com/std/14391274/858p1>>.
- [EURO2019] European Organization for Civil Aviation Equipment (EUROCAE), "Technical Standard of Aviation Profiles for ATN/IPS, ED-262", September 2019, <<https://eshop.eurocae.net/eurocae-documents-and-reports/ed-262/>>.
- [ICAO2015] International Civil Aviation Organization (ICAO), "Manual on the Aeronautical Telecommunication Network (ATN) using Internet Protocol Suite (IPS) Standards and Protocols, Doc 9896", January 2015, <<https://standards.globalspec.com/std/10026940/icao-9896>>.

- [RTCA2019] Radio Technical Commission for Aeronautics (RTCA), "Internet Protocol Suite Profiles, DO-379", September 2019, <<https://www.rtca.org/products/do-379/>>.
- [SCH2016] Schneckenburger, N., Jost, T., Shutin, D., Walter, M., Thiasiriphet, T., Schnell, M., and U.C. Fiebig, "Measurement of the L-band Air-to-Ground Channel for Positioning Applications", IEEE Transactions on Aerospace and Electronic Systems, 52(5), pp.2281-229 , 2016.
- [MAE20191] Maeurer, N., Graeupl, T., and C. Schmitt, "Evaluation of the LDACS Cybersecurity Implementation", IEEE 38th Digital Avionics Systems Conference (DACS), pp. 1-10, San Diego, CA, USA , 2019.
- [MAE20192] Maeurer, N. and C. Schmitt, "Towards Successful Realization of the LDACS Cybersecurity Architecture: An Updated Datalink Security Threat- and Risk Analysis", IEEE Integrated Communications, Navigation and Surveillance Conference (ICNS), pp. 1-13, Herndon, VA, USA , 2019.
- [MAE20182] Maeurer, N. and A. Bilzhause, "A Cybersecurity Architecture for the L-band Digital Aeronautical Communications System (LDACS)", IEEE 37th Digital Avionics Systems Conference (DASC), pp. 1-10, London, UK , 2017.
- [GRA2011] Graeupl, T. and M. Ehammer, "L-DACS1 Data Link Layer Evolution of ATN/IPS", 30th IEEE/AIAA Digital Avionics Systems Conference (DASC), pp. 1-28, Seattle, WA, USA , 2011.
- [GRA2018] Graeupl, T., Schneckenburger, N., Jost, T., Schnell, M., Filip, A., Bellido-Manganell, M.A., Mielke, D.M., Maeurer, N., Kumar, R., Osechas, O., and G. Battista, "L-band Digital Aeronautical Communications System (LDACS) flight trials in the national German project MICONAV", Integrated Communications, Navigation, Surveillance Conference (ICNS), pp. 1-7, Herndon, VA, USA , 2018.
- [ICA2018] International Civil Aviation Organization (ICAO), "L-Band Digital Aeronautical Communication System (LDACS)", International Standards and Recommended Practices Annex 10 - Aeronautical Telecommunications, Vol. III - Communication Systems , 2018.

- [SAJ2014] Haindl, B., Meser, J., Sajatovic, M., Mueller, S., Arthaber, H., Faseth, T., and M. Zaisberger, "LDACS1 Conformance and Compatibility Assessment", IEEE/AIAA 33rd Digital Avionics Systems Conference (DASC), pp. 1-11, Colorado Springs, CO, USA , 2014.
- [RIH2018] Rihacek, C., Haindl, B., Fantappie, P., Pierattelli, S., Graeupl, T., Schnell, M., and N. Fistas, "L-band Digital Aeronautical Communications System (LDACS) Activities in SESAR2020", Integrated Communications Navigation and Surveillance Conference (ICNS), pp. 1-8, Herndon, VA, USA , 2018.
- [BEL2019] Bellido-Manganell, M. A. and M. Schnell, "Towards Modern Air-to-Air Communications: the LDACS A2A Mode", IEEE/AIAA 38th Digital Avionics Systems Conference (DASC), pp. 1-10, San Diego, CA, USA , 2019.
- [CRO2016] Crowe, B., "Proposed AeroMACS PKI Specification is a Model for Global and National Aeronautical PKI Deployments", WiMAX Forum at 16th Integrated Communications, Navigation and Surveillance Conference (ICNS), pp. 1-19, New York, NY, USA , 2016.
- [MAE2020] Maeurer, N., Graeupl, T., and C. Schmitt, "Comparing Different Diffie-Hellman Key Exchange Flavors for LDACS", IEEE/AIAA 39th Digital Avionics Systems Conference (DASC), pp. 1-10, San Antonio, TX, USA , 2020.
- [STR2016] Strohmeier, M., Schaefer, M., Pinheiro, R., Lenders, V., and I. Martinovic, "On Perception and Reality in Wireless Air Traffic Communication Security", IEEE Transactions on Intelligent Transportation Systems, 18(6), pp. 1338-1357, New York, NY, USA , 2016.
- [BIL2017] Bilzhause, A., Belgacem, B., Mostafa, M., and T. Graeupl, "Datalink Security in the L-band Digital Aeronautical Communications System (LDACS) for Air Traffic Management", IEEE Aerospace and Electronic Systems Magazine, 32(11), pp. 22-33, New York, NY, USA , 2017.
- [MAE20181] Maeurer, N. and A. Bilzhause, "Paving the Way for an IT Security Architecture for LDACS: A Datalink Security Threat and Risk Analysis", IEEE Integrated Communications, Navigation, Surveillance Conference (ICNS), pp. 1-11, New York, NY, USA , 2018.

- [FAA2020] FAA, "Federal Aviation Administration. ADS-B Privacy.", August 2020, <<https://www.faa.gov/nextgen/equipadsb/privacy/>>.
- [GNU2021] GNU Radio project, "GNU radio", October 2021, <<http://gnuradio.org>>.
- [SIT2020] SITA, "Societe Internationale de Telecommunications Aeronautiques", August 2020, <<https://www.sita.aero/>>.
- [ARI2020] ARINC, "Aeronautical Radio Incorporated", August 2020, <<https://www.aviation-ia.com/>>.
- [DO350A] RTCA SC-214, "Safety and Performance Standard for Baseline 2 ATS Data Communications (Baseline 2 SPR Standard)", May 2016, <<https://standards.globalspec.com/std/10003192/rtca-do-350-volume-1-2>>.
- [ICAO2019] International Civil Aviation Organization (ICAO), "Manual on VHF Digital Link (VDL) Mode 2, Doc 9776", January 2019, <<https://store.icao.int/en/manual-on-vhf-digital-link-vdl-mode-2-doc-9776>>.
- [KAMA2010] Kamali, B., "An Overview of VHF Civil Radio Network and the Resolution of Spectrum Depletion", Integrated Communications, Navigation, and Surveillance Conference, pp. F4-1-F4-8 , May 2010.
- [SON2021] Soni, D., Basu, K., Nabeel, M., Aaraj, N., Manzano, M., and R. Karri, "FALCON", Hardware Architectures for Post-Quantum Digital Signature Schemes, pp. 31-41 , November 2021.
- [SIK2021] SIKE, "SIKE Supersingular Isogeny Key Encapsulation", October 2021, <<https://sike.org/>>.
- [ROY2020] Roy, S.S.. and A. Basso, "High-Speed Instruction-Set Coprocessor For Lattice-Based Key Encapsulation Mechanism: Saber In Hardware", IACR Transactions on Cryptographic Hardware and Embedded Systems, 443-466. , August 2020.
- [RAW-TECHNOS]
Thubert, P., Cavalcanti, D., Vilajosana, X., Schmitt, C., and J. Farkas, "Reliable and Available Wireless Technologies", Work in Progress, Internet-Draft, draft-ietf-raw-technologies-05, 2 February 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-raw-technologies-05>>.

[RAW-USE-CASES]

Bernardos, C. J., Papadopoulos, G. Z., Thubert, P., and F. Theoleyre, "RAW use-cases", Work in Progress, Internet-Draft, draft-ietf-raw-use-cases-05, 23 February 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-raw-use-cases-05>>.

[I-D.haindl-lisp-gb-atn]

Haindl, B., Lindner, M., Rahman, R., Comeras, M. P., Moreno, V., Maino, F., and B. Venkatachalapathy, "Ground-Based LISP for the Aeronautical Telecommunications Network", Work in Progress, Internet-Draft, draft-haindl-lisp-gb-atn-06, 6 March 2021, <<https://datatracker.ietf.org/doc/html/draft-haindl-lisp-gb-atn-06>>.

[I-D.ietf-rtgwg-atn-bgp]

Templin, F. L., Saccone, G., Dawra, G., Lindem, A., and V. Moreno, "A Simple BGP-based Mobile Routing System for the Aeronautical Telecommunications Network", Work in Progress, Internet-Draft, draft-ietf-rtgwg-atn-bgp-14, 14 February 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-rtgwg-atn-bgp-14>>.

[ICAO2018] International Civil Aviation Organization (ICAO), "Handbook on Radio Frequency Spectrum Requirements for Civil Aviation, Doc 9718, Volume 1, ICAO Spectrum Strategy, Policy Statements and Related Information", July 2018, <[https://www.icao.int/safety/FSMP/Documents/Doc9718/Doc9718_Vol_I_2nd_ed_\(2018\)corrl.pdf](https://www.icao.int/safety/FSMP/Documents/Doc9718/Doc9718_Vol_I_2nd_ed_(2018)corrl.pdf)>.

[ARI2019] ARINC, "AOC Air-Ground Data And Message Exchange Format, ARINC 633", January 2019, <<https://standards.globalspec.com/std/13152055/ARINC%20633>>.

[VIR2021] Viridia, A., Stea, G., and G. Dini, "SAPIENT: Enabling Real-Time Monitoring and Control in the Future Communication Infrastructure of Air Traffic Management", IEEE Transactions on Intelligent Transportation Systems, 22(8):4864-4875, August 2021.

[SHU2013] Shutin, D., Schneckenburger, N., Walter, M., and M. Schnell, "LDACS1 Ranging Performance - An Analysis Of Flight Measurement Results", IEEE 32th Digital Avionics Systems Conference (DASC), pp. 1-10, East Syracuse, NY, USA, October 2013.

- [BEL2021] Bellido-Manganell, M.A., Graeupl, T., Heirich, O., Maeurer, N., Filip-Dhaubhadel, A., Mielke, D.M., Schalk, L.M., Becker, D., Schneckenburger, N., and M. Schnell, "LDACS Flight Trials: Demonstration and Performance Analysis of the Future Aeronautical Communications System", IEEE Transactions on Aerospace and Electronic Systems, pp. 1-19 , September 2021.
- [MAE2021] Maeurer, N., Graeupl, T., Gentsch, C., Guggemos, T., Tiepelt, M., Schmitt, C., and G. Dreo Rodosek, "A Secure Cell-Attachment Procedure for LDACS", 1st Workshop on Secure and Reliable Communication and Navigation in the Aerospace Domain (SRCNAS), pp. 1-10, Vienna, Austria , September 2021.
- [MAE20211] Maeurer, N., Graeupl, T., Bellido-Manganell, M.A., Mielke, D.M., Filip-Dhaubhadel, A., Heirich, O., Gerberth, D., Flux, M., Schalk, L.M., Becker, D., Schneckenburger, N., and M. Schnell, "Flight Trial Demonstration of Secure GBAS via the L-band Digital Aeronautical Communications System (LDACS)", IEEE Aerospace and Electronic Systems Magazine, 36(4), pp. 8-17 , April 2021.
- [BOE2019] Boegl, T., Rautenberg, M., Haindl, R., Rihacek, C., Meser, J., Fantappie, P., Pringvanich, N., Micallef, J., Klauspeter, H., MacBride, J., Sacre, P., v.d. Eiden, B., Graeupl, T., and M. Schnell, "LDACS White Paper - A Roll-out Scenario", International Civil Aviation Organization, Communications Panel - Data Communications Infrastructure Working Group - Third Meeting, pp. 1-8, Montreal, Canada , October 2019.

Appendix A. Selected Information from DO-350A

This appendix includes the continuity, availability, and integrity requirements applicable for LDACS defined in [DO350A].

The following terms are used here:

CPDLC Controller Pilot Data Link Communication
DT Delivery Time (nominal) value for RSP
ET Expiration Time value for RCP
FH Flight Hour
MA Monitoring and Alerting criteria
OT Overdue Delivery Time value for RSP
RCP Required Communication Performance
RSP Required Surveillance Performance
TT Transaction Time (nominal) value for RCP

	RCP 130	RCP 130
Parameter	ET	TT95%
Transaction Time (sec)	130	67
Continuity	0.999	0.95
Availability	0.989	0.989
Integrity	1E-5 per FH	1E-5 per FH

Table 1: CPDLC Requirements for RCP 130

	RCP 240	RCP 240	RCP 400	RCP 400
Parameter	ET	TT95%	ET	TT95%
Transaction Time (sec)	240	210	400	350
Continuity	0.999	0.95	0.999	0.95
Availability	0.989	0.989	0.989	0.989
Integrity	1E-5 per FH	1E-5 per FH	1E-5 per FH	1E-5 per FH

Table 2: CPDLC Requirements for RCP 240/400

RCP Monitoring and Alerting Criteria in case of CPDLC:

- MA-1: The system shall be capable of detecting failures and configuration changes that would cause the communication service no longer meet the RCP specification for the intended use.
- MA-2: When the communication service can no longer meet the RCP specification for the intended function, the flight crew and/or the controller shall take appropriate action.

	RSP 160	RSP 160	RSP 180	RSP 180	RSP 400	RSP 400
Parameter	OT	DT95%	OT	DT95%	OT	DT95%
Transaction Time (sec)	160	90	180	90	400	300
Continuity	0.999	0.95	0.999	0.95	0.999	0.95
Availability	0.989	0.989	0.989	0.989	0.989	0.989
Integrity	1E-5 per FH	1E-5 per FH	1E-5 per FH	1E-5 per FH	1E-5 per FH	1E-5 per FH

Table 3: ADS-C Requirements

RCP Monitoring and Alerting Criteria:

- MA-1: The system shall be capable of detecting failures and configuration changes that would cause the ADS-C service no longer meet the RSP specification for the intended function.
- MA-2: When the ADS-C service can no longer meet the RSP specification for the intended function, the flight crew and/or the controller shall take appropriate action.

Authors' Addresses

Nils Maeurer (editor)
 German Aerospace Center (DLR)
 Muenchner Strasse 20
 82234 Wessling
 Germany
 Email: Nils.Maeurer@dlr.de

Thomas Graeupl (editor)
 German Aerospace Center (DLR)
 Muenchner Strasse 20
 82234 Wessling
 Germany
 Email: Thomas.Graeupl@dlr.de

Corinna Schmitt (editor)
Research Institute CODE, UniBwM
Werner-Heisenberg-Weg 28
85577 Neubiberg
Germany
Email: corinna.schmitt@unibw.de

RAW
Internet-Draft
Intended status: Informational
Expires: 23 April 2021

P. Thubert, Ed.
Cisco Systems
D. Cavalcanti
Intel
X. Vilajosana
Universitat Oberta de Catalunya
C. Schmitt
Research Institute CODE, UniBwM
J. Farkas
Ericsson
20 October 2020

Reliable and Available Wireless Technologies
draft-ietf-raw-technologies-00

Abstract

This document presents a series of recent technologies that are capable of time synchronization and scheduling of transmission, making them suitable to carry time-sensitive flows with high reliability and availability.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 23 April 2021.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document.

Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Terminology	3
3. On Scheduling	4
3.1. Benefits of Scheduling on Wires	4
3.2. Benefits of Scheduling on Wireless	5
4. IEEE 802.11	6
4.1. Provenance and Documents	6
4.2. 802.11ax High Efficiency (HE)	8
4.2.1. General Characteristics	8
4.2.2. Applicability to deterministic flows	9
4.3. 802.11be Extreme High Throughput (EHT)	10
4.3.1. General Characteristics	10
4.3.2. Applicability to deterministic flows	11
4.4. 802.11ad and 802.11ay (mmWave operation)	12
4.4.1. General Characteristics	12
4.4.2. Applicability to deterministic flows	13
5. IEEE 802.15.4	13
5.1. Provenance and Documents	13
5.2. TimeSlotted Channel Hopping	15
5.2.1. General Characteristics	15
5.2.2. Applicability to Deterministic Flows	16
6. 5G	30
6.1. Provenance and Documents	30
6.2. General Characteristics	32
6.3. Deployment and Spectrum	33
6.4. Applicability to Deterministic Flows	34
6.4.1. System Architecture	34
6.4.2. Overview of The Radio Protocol Stack	36
6.4.3. Radio (PHY)	37
6.4.4. Scheduling and QoS (MAC)	39
6.4.5. Time-Sensitive Networking (TSN) Integration	41
6.5. Summary	44
7. L-band Digital Aeronautical Communications System	45
7.1. Provenance and Documents	45
7.2. General Characteristics	46
7.3. Applicability to Deterministic Flows	47
8. IANA Considerations	48
9. Security Considerations	48
10. Contributors	48
11. Acknowledgments	49

12. Normative References	49
13. Informative References	49
Authors' Addresses	57

1. Introduction

When used in math or philosophy, the term "deterministic" generally refers to a perfection where all aspect are understood and predictable. A perfectly Deterministic Network would ensure that every packet reach its destination following a predetermined path along a predefined schedule to be delivered at the exact due time. In a real and imperfect world, a Deterministic Network must highly predictable, which is a combination of reliability and availability. On the one hand the network must be reliable, meaning that it will perform as expected for all packets and in particular that it will always deliver the packet at the destination in due time. On the other hand, the network must be available, meaning that it is resilient to any single outage, whether the cause is a software, a hardware or a transmission issue.

RAW (Reliable and Available Wireless) is an effort to provide Deterministic Networking on across a path that include a wireless physical layer. Making Wireless Reliable and Available is even more challenging than it is with wires, due to the numerous causes of loss in transmission that add up to the congestion losses and the delays caused by overbooked shared resources. In order to maintain a similar quality of service along a multihop path that is composed of wired and wireless hops, additional methods that are specific to wireless must be leveraged to combat the sources of loss that are also specific to wireless.

Such wireless-specific methods include per-hop retransmissions (HARQ) and P2MP overhearing whereby multiple receivers are scheduled to receive the same transmission, which balances the adverse effects of the transmission losses that are experienced when a radio is used as pure P2P. Those methods are collectively referred to as PAREO functions in the "Reliable and Available Wireless Architecture/Framework" [I-D.pthubert-raw-architecture].

2. Terminology

This specification uses several terms that are uncommon on protocols that ensure best effort transmissions for stochastic flows, such as found in the traditional Internet and other statistically multiplexed packet networks.

ARQ: Automatic Repeat Request, enabling an acknowledged transmission

and retries. ARQ is a typical model at Layer-2 on a wireless medium. It is typically avoided end-to-end on deterministic flows because it introduces excessive indetermination in latency, but a limited number of retries within a bounded time may be used over a wireless link and yet respect end-to-end constraints.

Available: That is exempt of unscheduled outage, the expectation for a network being that the flow is maintained in the face of any single breakage.

FEC: Forward error correction, sending redundant coded data to help the receiver recover transmission errors without the delays incurred with ARQ.

HARQ: Hybrid ARQ, a combination of FEC and ARQ.

PCE: Path Computation Element.

PAREO (functions): the wireless extension of DetNet PREOF. PAREO functions include scheduled ARQ at selected hops, and expect the use of new operations like overhearing where available.

Reliable: That consistently performs as expected, the expectation for a network being to always deliver a packet in due time.

Track: A DODAG oriented to a destination, and that enables Packet ARQ, Replication, Elimination, and Ordering Functions.

3. On Scheduling

The operations of a Deterministic Network often rely on precisely applying a tight schedule, in order to avoid collision loss and guarantee the worst-case time of delivery. To achieve this, there must be a shared sense of time throughout the network. The sense of time is usually provided by the lower layer and is not in scope for RAW.

3.1. Benefits of Scheduling on Wires

A network is reliable when the statistical effects that affect the packet transmission are eliminated. This involves maintaining at all time the amount of critical packets within the physical capabilities of the hardware and that of the radio medium. This is achieved by controlling the use of time-shared resources such as CPUs and buffers, by shaping the flows and by scheduling the time of transmission of the packets that compose the flow at every hop.

Equipment failure, such as an access point rebooting, a broken radio adapter, or a permanent obstacle to the transmission, is a secondary source of packet loss. When a breakage occurs, multiple packets are lost in a row before the flows are rerouted or the system may recover. This is not acceptable for critical applications such as related to safety. A typical process control loop will tolerate an occasional packet loss, but a loss of several packets in a row will cause an emergency stop (e.g., after 4 packets lost, within a period of 1 second).

Network Availability is obtained by making the transmission resilient against hardware failures and radio transmission losses due to uncontrolled events such as co-channel interferers, multipath fading or moving obstacles. The best results are typically achieved by pseudo randomly cumulating all forms of diversity, in the spatial domain with replication and elimination, in the time domain with ARQ and diverse scheduled transmissions, and in the frequency domain with frequency hopping or channel hopping between frames.

3.2. Benefits of Scheduling on Wireless

In addition to the benefits listed in Section 3.1, scheduling transmissions provides specific value to the wireless medium.

On the one hand, scheduling avoids collisions between scheduled transmissions and can ensure both time and frequency diversity between retries in order to defeat co-channel interference from uncontrolled transmitters as well as multipath fading. Transmissions can be scheduled on multiple channels in parallel, which enables to use the full available spectrum while avoiding the hidden terminal problem, e.g., when the next packet in a same flow interferes on a same channel with the previous one that progressed a few hops farther.

On the other hand, scheduling optimizes the bandwidth usage: compared to classical Collision Avoidance techniques, there is no blank time related to inter-frame space (IFS) and exponential back-off in scheduled operations. A minimal Clear Channel Assessment may be needed to comply with the local regulations such as ETSI 300-328, but that will not detect a collision when the senders are synchronized. And because scheduling allows a time-sharing operation, there is no limit to the ratio of isolated critical traffic.

Finally, scheduling plays a critical role to save energy. In IOT, energy is the foremost concern, and synchronizing sender and listener enables to always maintain them in deep sleep when there is no scheduled transmission. This avoids idle listening and long preambles and enables long sleep periods between traffic and resynchronization, allowing battery-operated nodes to operate in a mesh topology for multiple years.

4. IEEE 802.11

4.1. Provenance and Documents

With an active portfolio of nearly 1,300 standards and projects under development, IEEE is a leading developer of industry standards in a broad range of technologies that drive the functionality, capabilities, and interoperability of products and services, transforming how people live, work, and communicate.

The IEEE 802 LAN/MAN Standards Committee (SC) develops and maintains networking standards and recommended practices for local, metropolitan, and other area networks, using an open and accredited process, and advocates them on a global basis. The most widely used standards are for Ethernet, Bridging and Virtual Bridged LANs, Wireless LAN, Wireless PAN, Wireless MAN, Wireless Coexistence, Media Independent Handover Services, and Wireless RAN. An individual Working Group provides the focus for each area. Standards produced by the IEEE 802 SC are freely available from the IEEE GET Program after they have been published in PDF for six months.

The IEEE 802.11 LAN standards define the underlying MAC and PHY layers for the Wi-Fi technology. Wi-Fi/802.11 is one of the most successful wireless technologies, supporting many application domains. While previous 802.11 generations, such as 802.11n and 802.11ac, have focused mainly on improving peak throughput, more recent generations are also considering other performance vectors, such as efficiency enhancements for dense environments in 802.11ax, and latency and support for Time-Sensitive Networking (TSN) capabilities in 802.11be.

IEEE 802.11 already supports some 802.1 TSN standards and it is undergoing efforts to support for other 802.1 TSN capabilities required to address the use cases that require time synchronization and timeliness (bounded latency) guarantees with high reliability and availability. The IEEE 802.11 working group has been working in collaboration with the IEEE 802.1 group for several years extending 802.1 features over 802.11. As with any wireless media, 802.11 imposes new constraints and restrictions to TSN-grade QoS, and tradeoffs between latency and reliability guarantees must be

considered as well as managed deployment requirements. An overview of 802.1 TSN capabilities and their extensions to 802.11 are discussed in [Cavalcanti_2019].

Wi-Fi Alliance (WFA) is the worldwide network of companies that drives global Wi-Fi adoption and evolution through thought leadership, spectrum advocacy, and industry-wide collaboration. The WFA work helps ensure that Wi-Fi devices and networks provide users the interoperability, security, and reliability they have come to expect.

The following [IEEE Std. 802.11] specifications/certifications are relevant in the context of reliable and available wireless services and support for time-sensitive networking capabilities:

Time Synchronization: IEEE802.11-2016 with IEEE802.1AS; WFA TimeSync Certification.

Congestion Control: IEEE802.11-2016 Admission Control; WFA Admission Control.

Security: WFA Wi-Fi Protected Access, WPA2 and WPA3.

Interoperating with IEEE802.1Q bridges: [IEEE Std. 802.11ak].

Stream Reservation Protocol (part of [IEEE Std. 802.1Qat]): AIEEE802.11-2016

Scheduled channel access: IEEE802.11ad Enhancements for very high throughput in the 60 GHz band [IEEE Std. 802.11ad].

802.11 Real-Time Applications: Topic Interest Group (TIG) ReportDoc [IEEE_doc_11-18-2009-06].

In addition, major amendments being developed by the IEEE802.11 Working Group include capabilities that can be used as the basis for providing more reliable and predictable wireless connectivity and support time-sensitive applications:

IEEE 802.11ax D4.0: Enhancements for High Efficiency (HE). [IEEE Std. 802.11ax]

IEEE 802.11be Extreme High Throughput (EHT). [IEEE 802.11be WIP]

IEEE 802.11ay Enhanced throughput for operation in license-exempt bands above 45 GHz. [IEEE Std. 802.11ay]

The main 802.11ax and 802.11be capabilities and their relevance to RAW are discussed in the remainder of this document.

4.2. 802.11ax High Efficiency (HE)

4.2.1. General Characteristics

The next generation Wi-Fi (Wi-Fi 6) is based on the IEEE802.11ax amendment [IEEE Std. 802.11ax], which includes new capabilities to increase efficiency, control and reduce latency. Some of the new features include higher order 1024-QAM modulation, support for uplink multi-user MIMO, OFDMA, trigger-based access and Target Wake time (TWT) for enhanced power savings. The OFDMA mode and trigger-based access enable scheduled operation, which is a key capability required to support deterministic latency and reliability for time-sensitive flows. 802.11ax can operate in up to 160 MHz channels and it includes support for operation in the new 6 GHz band, which is expected to be open to unlicensed use by the FCC and other regulatory agencies worldwide.

4.2.1.1. Multi-User OFDMA and Trigger-based Scheduled Access

802.11ax introduced a new orthogonal frequency-division multiple access (OFDMA) mode in which multiple users can be scheduled across the frequency domain. In this mode, the Access Point (AP) can initiate multi-user (MU) Uplink (UL) transmissions in the same PHY Protocol Data Unit (PPDU) by sending a trigger frame. This centralized scheduling capability gives the AP much more control of the channel, and it can remove contention between devices for uplink transmissions, therefore reducing the randomness caused by CSMA-based access between stations. The AP can also transmit simultaneously to multiple users in the downlink direction by using a Downlink (DL) MU OFDMA PPDU. In order to initiate a contention free Transmission Opportunity (TXOP) using the OFDMA mode, the AP still follows the typical listen before talk procedure to acquire the medium, which ensures interoperability and compliance with unlicensed band access rules. However, 802.11ax also includes a multi-user Enhanced Distributed Channel Access (MU-EDCA) capability, which allows the AP to get higher channel access priority.

4.2.1.2. Improved PHY Robustness

The 802.11ax PHY can operate with 0.8, 1.6 or 3.2 microsecond guard interval (GI). The larger GI options provide better protection against multipath, which is expected to be a challenge in industrial environments. The possibility to operate with smaller resource units (e.g. 2 MHz) enabled by OFDMA also helps reduce noise power and improve SNR, leading to better packet error rate (PER) performance.

802.11ax supports beamforming as in 802.11ac, but introduces UL MU MIMO, which helps improve reliability. The UL MU MIMO capability is also enabled by the trigger based access operation in 802.11ax.

4.2.1.3. Support for 6GHz band

The 802.11ax specification [IEEE Std. 802.11ax] includes support for operation in the new 6 GHz band. Given the amount of new spectrum available as well as the fact that no legacy 802.11 device (prior 802.11ax) will be able to operate in this new band, 802.11ax operation in this new band can be even more efficient.

4.2.2. Applicability to deterministic flows

TSN capabilities, as defined by the IEEE 802.1 TSN standards, provide the underlying mechanism for supporting deterministic flows in a Local Area Network (LAN). The 802.11 working group has already incorporated support for several TSN capabilities, so that time-sensitive flow can experience precise time synchronization and timeliness when operating over 802.11 links. TSN capabilities supported over 802.11 (which also extends to 802.11ax), include:

1. 802.1AS based Time Synchronization (other time synchronization techniques may also be used)
2. Interoperating with IEEE802.1Q bridges
3. Time-sensitive Traffic Stream identification

The existing 802.11 TSN capabilities listed above, and the 802.11ax OFDMA and scheduled access provide a new set of tools to better server time-sensitive flows. However, it is important to understand the tradeoffs and constraints associated with such capabilities, as well as redundancy and diversity mechanisms that can be used to provide more predictable and reliable performance.

4.2.2.1. 802.11 Managed network operation and admission control

Time-sensitive applications and TSN standards are expected to operate under a managed network (e.g. industrial/enterprise network). Thus, the Wi-Fi operation must also be carefully managed and integrated with the overall TSN management framework, as defined in the [IEEE8021Qcc] specification.

Some of the random-access latency and interference from legacy/unmanaged devices can be minimized under a centralized management mode as defined in [IEEE8021Qcc], in which admission control procedures are enforced.

Existing traffic stream identification, configuration and admission control procedures defined in [IEEE Std. 802.11] QoS mechanism can be re-used. However, given the high degree of determinism required by many time-sensitive applications, additional capabilities to manage interference and legacy devices within tight time-constraints need to be explored.

4.2.2.2. Scheduling for bounded latency and diversity

As discussed earlier, the [IEEE Std. 802.11ax] OFDMA mode introduces the possibility of assigning different RUs (frequency resources) to users within a PPDU. Several RU sizes are defined in the specification (26, 52, 106, 242, 484, 996 subcarriers). In addition, the AP can also decide on MCS and grouping of users within a given OFDMA PPDU. Such flexibility can be leveraged to support time-sensitive applications with bounded latency, especially in a managed network where stations can be configured to operate under the control of the AP.

As shown in [Cavalcanti_2019], it is possible to achieve latencies in the order of 1msec with high reliability in an interference free environment. Obviously, there are latency, reliability and capacity tradeoffs to be considered. For instance, smaller Resource Units (RU)s result in longer transmission durations, which may impact the minimal latency that can be achieved, but the contention latency and randomness elimination due to multi-user transmission is a major benefit of the OFDMA mode.

The flexibility to dynamically assign RUs to each transmission also enables the AP to provide frequency diversity, which can help increase reliability.

4.3. 802.11be Extreme High Throughput (EHT)

4.3.1. General Characteristics

The [IEEE 802.11be WIP] is the next major 802.11 amendment (after [IEEE Std. 802.11ax]) for operation in the 2.4, 5 and 6 GHz bands. 802.11be is expected to include new PHY and MAC features and it is targeting extremely high throughput (at least 30 Gbps), as well as enhancements to worst case latency and jitter. It is also expected to improve the integration with 802.1 TSN to support time-sensitive applications over Ethernet and Wireless LANs.

The 802.11be Task Group started its operation in May 2019, therefore, detailed information about specific features is not yet available. Only high level candidate features have been discussed so far, including:

1. 320MHz bandwidth and more efficient utilization of non-contiguous spectrum.
2. Multi-band/multi-channel aggregation and operation.
3. 16 spatial streams and related MIMO enhancements.
4. Multi-Access Point (AP) Coordination.
5. Enhanced link adaptation and retransmission protocol, e.g. Hybrid Automatic Repeat Request (HARQ).
6. Any required adaptations to regulatory rules for the 6 GHz spectrum.

4.3.2. Applicability to deterministic flows

The 802.11 Real-Time Applications (RTA) Topic Interest Group (TIG) provided detailed information on use cases, issues and potential solution directions to improve support for time-sensitive applications in 802.11. The RTA TIG report [IEEE_doc_11-18-2009-06] was used as input to the 802.11be project scope.

Improvements for worst-case latency, jitter and reliability were the main topics identified in the RTA report, which were motivated by applications in gaming, industrial automation, robotics, etc. The RTA report also highlighted the need to support additional TSN capabilities, such as time-aware (802.1Qbv) shaping and packet replication and elimination as defined in 802.1CB.

802.11be is expected to build on and enhance 802.11ax capabilities to improve worst case latency and jitter. Some of the enhancement areas are discussed next.

4.3.2.1. Enhanced scheduled operation for bounded latency

In addition to the throughput enhancements, 802.11be will leverage the trigger-based scheduled operation enabled by 802.11ax to provide efficient and more predictable medium access. 802.11be is expected to include enhancements to reduce overhead and enable more efficient operation in managed network deployments [IEEE_doc_11-19-0373-00].

4.3.2.2. Multi-AP coordination

Multi-AP coordination is one of the main new candidate features in 802.11be. It can provide benefits in throughput and capacity and has the potential to address some of the issues that impact worst case latency and reliability. Multi-AP coordination is expected to address the contention due to overlapping Basic Service Sets (OBSS), which is one of the main sources of random latency variations. 802.11be can define methods to enable better coordination between APs, for instance, in a managed network scenario, in order to reduce latency due to unmanaged contention.

Several multi-AP coordination approaches have been discussed with different levels of complexities and benefits, but specific coordination methods have not yet been defined.

4.3.2.3. Multi-band operation

802.11be will introduce new features to improve operation over multiple bands and channels. By leveraging multiple bands/channels, 802.11be can isolate time-sensitive traffic from network congestion, one of the main causes of large latency variations. In a managed 802.11be network, it should be possible to steer traffic to certain bands/channels to isolate time-sensitive traffic from other traffic and help achieve bounded latency.

4.4. 802.11ad and 802.11ay (mmWave operation)

4.4.1. General Characteristics

The IEEE 802.11ad amendment defines PHY and MAC capabilities to enable multi-Gbps throughput in the 60 GHz millimeter wave (mmWave) band. The standard addresses the adverse mmWave signal propagation characteristics and provides directional communication capabilities that take advantage of beamforming to cope with increased attenuation. An overview of the 802.11ad standard can be found in [Nitsche_2015] .

The IEEE 802.11ay is currently developing enhancements to the 802.11ad standard to enable the next generation mmWave operation targeting 100 Gbps throughput. Some of the main enhancements in 802.11ay include MIMO, channel bonding, improved channel access and beamforming training. An overview of the 802.11ay capabilities can be found in [Ghasempour_2017]

4.4.2. Applicability to deterministic flows

The high data rates achievable with 802.11ad and 802.11ay can significantly reduce latency down to microsecond levels. Limited interference from legacy and other unlicensed devices in 60 GHz is also a benefit. However, directionality and short range typical in mmWave operation impose new challenges such as the overhead required for beam training and blockage issues, which impact both latency and reliability. Therefore, it is important to understand the use case and deployment conditions in order to properly apply and configure 802.11ad/ay networks for time sensitive applications.

The 802.11ad standard include a scheduled access mode in which stations can be allocated contention-free service periods by a central controller. This scheduling capability is also available in 802.11ay, and it is one of the mechanisms that can be used to provide bounded latency to time-sensitive data flows. An analysis of the theoretical latency bounds that can be achieved with 802.11ad service periods is provided in [Cavalcanti_2019].

5. IEEE 802.15.4

5.1. Provenance and Documents

The IEEE802.15.4 Task Group has been driving the development of low-power low-cost radio technology. The IEEE802.15.4 physical layer has been designed to support demanding low-power scenarios targeting the use of unlicensed bands, both the 2.4 GHz and sub GHz Industrial, Scientific and Medical (ISM) bands. This has imposed requirements in terms of frame size, data rate and bandwidth to achieve reduced collision probability, reduced packet error rate, and acceptable range with limited transmission power. The PHY layer supports frames of up to 127 bytes. The Medium Access Control (MAC) sublayer overhead is in the order of 10-20 bytes, leaving about 100 bytes to the upper layers. IEEE802.15.4 uses spread spectrum modulation such as the Direct Sequence Spread Spectrum (DSSS).

The Timeslotted Channel Hopping (TSCH) mode was added to the 2015 revision of the IEEE802.15.4 standard [IEEE Std. 802.15.4]. TSCH is targeted at the embedded and industrial world, where reliability, energy consumption and cost drive the application space.

At the IETF, the 6TiSCH Working Group (WG) [TiSCH] deals with best effort operation of IPv6 [RFC8200] over TSCH. 6TiSCH has enabled distributed scheduling to exploit the deterministic access capabilities provided by TSCH. The group designed the essential mechanisms to enable the management plane operation while ensuring IPv6 is supported. Yet the charter did not focus to providing a

solution to establish end to end Tracks while meeting quality of service requirements. 6TiSCH, through the RFC8480 [RFC8480] defines the 6P protocol which provides a pairwise negotiation mechanism to the control plane operation. The protocol supports agreement on a schedule between neighbors, enabling distributed scheduling. 6P goes hand-in-hand with a Scheduling Function (SF), the policy that decides how to maintain cells and trigger 6P transactions. The Minimal Scheduling Function (MSF) [I-D.ietf-6tisch-msf] is the default SF defined by the 6TiSCH WG; other standardized SFs can be defined in the future. MSF extends the minimal schedule configuration, and is used to add child-parent links according to the traffic load.

Time sensitive networking on low power constrained wireless networks have been partially addressed by ISA100.11a [ISA100.11a] and WirelessHART [WirelessHART]. Both technologies involve a central controller that computes redundant paths for industrial process control traffic over a TSCH mesh. Moreover, ISA100.11a introduces IPv6 capabilities with a Link-Local Address for the join process and a global unicast address for later exchanges, but the IPv6 traffic typically ends at a local application gateway and the full power of IPv6 for end-to-end communication is not enabled. Compared to that state of the art, work at the IETF and in particular at RAW could provide additional techniques such as optimized P2P routing, PAREO functions, and end-to-end secured IPv6/CoAP connectivity.

The 6TiSCH architecture [I-D.ietf-6tisch-architecture] identifies different models to schedule resources along so-called Tracks (see Section 5.2.2.2) exploiting the TSCH schedule structure however the focus at 6TiSCH is on best effort traffic and the group was never chartered to produce standard work related to Tracks.

Useful References include:

1. IEEE Std 802.15.4: "IEEE Std. 802.15.4, Part. 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks" [IEEE Std. 802.15.4]. The latest version at the time of this writing is dated year 2015.
2. Morell, A. , Vilajosana, X. , Vicario, J. L. and Watteyne, T. (2013), Label switching over IEEE802.15.4e networks. Trans. Emerging Tel. Tech., 24: 458-475. doi:10.1002/ett.2650" [morell13].

3. De Armas, J., Tuset, P., Chang, T., Adelantado, F., Watteyne, T., Vilajosana, X. (2016, September). Determinism through path diversity: Why packet replication makes sense. In 2016 International Conference on Intelligent Networking and Collaborative Systems (INCoS) (pp. 150-154). IEEE. [dearmas16].
4. X. Vilajosana, T. Watteyne, M. Vucinic, T. Chang and K. S. J. Pister, "6TiSCH: Industrial Performance for IPv6 Internet-of-Things Networks," in Proceedings of the IEEE, vol. 107, no. 6, pp. 1153-1165, June 2019. [vilajosana19].

5.2. TimeSlotted Channel Hopping

5.2.1. General Characteristics

As a core technique in IEEE802.15.4, TSCH splits time in multiple time slots that repeat over time. A set of timeslots constructs a Slotframe (see Section 5.2.2.1.4). For each timeslot, a set of available frequencies can be used, resulting in a matrix-like schedule (see Figure 1).

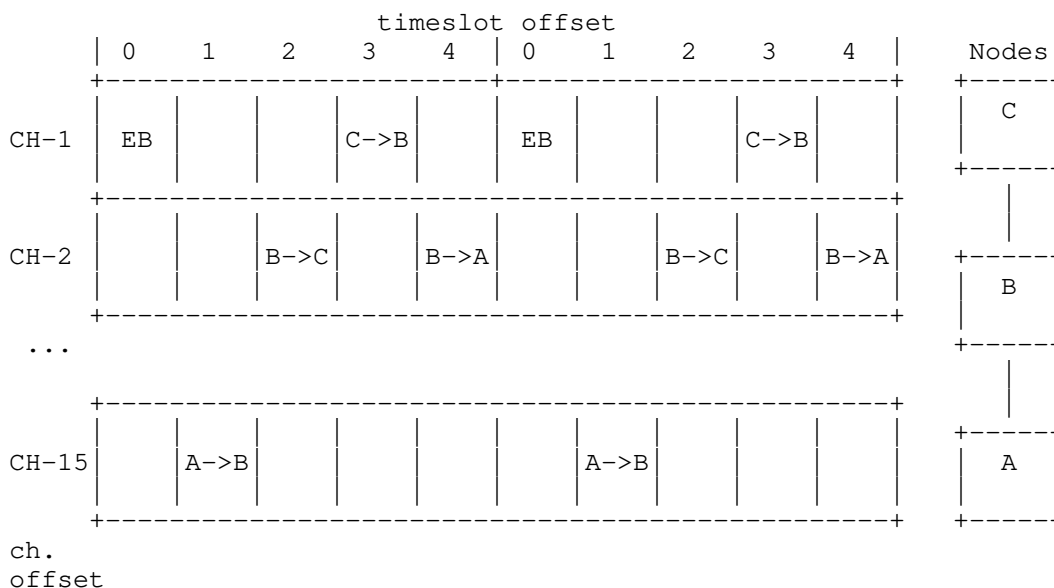


Figure 1: Slotframe example with scheduled cells between nodes A, B and C

This schedule represents the possible communications of a node with its neighbors, and is managed by a Scheduling Function such as the Minimal Scheduling Function (MSF) [I-D.ietf-6tisch-msf]. Each cell

in the schedule is identified by its slotoffset and channeloffset coordinates. A cell's timeslot offset indicates its position in time, relative to the beginning of the slotframe. A cell's channel offset is an index which maps to a frequency at each iteration of the slotframe. Each packet exchanged between neighbors happens within one cell. The size of a cell is a timeslot duration, between 10 to 15 milliseconds. An Absolute Slot Number (ASN) indicates the number of slots elapsed since the network started. It increments at every slot. This is a 5 byte counter that can support networks running for more than 300 years without wrapping (assuming a 10 ms timeslot). Channel hopping provides increased reliability to multi-path fading and external interference. It is handled by TSCH through a channel hopping sequence referred as macHopSeq in the IEEE802.15.4 specification.

The Time-Frequency Division Multiple Access provided by TSCH enables the orchestration of traffic flows, spreading them in time and frequency, and hence enabling an efficient management of the bandwidth utilization. Such efficient bandwidth utilization can be combined to OFDM modulations also supported by the IEEE802.15.4 standard [IEEE Std. 802.15.4] since the 2015 version.

In the RAW context, low power reliable networks should address non-critical control scenarios such as Class 2 and monitoring scenarios such as Class 4 defined by the RFC5673 [RFC5673]. As a low power technology targeting industrial scenarios radio transducers provide low data rates (typically between 50kbps to 250kbps) and robust modulations to trade-off performance to reliability. TSCH networks are organized in mesh topologies and connected to a backbone. Latency in the mesh network is mainly influenced by propagation aspects such as interference. ARQ methods and redundancy techniques such as replication and elimination should be studied to provide the needed performance to address deterministic scenarios.

5.2.2. Applicability to Deterministic Flows

Nodes in a TSCH network are tightly synchronized. This enables to build the slotted structure and ensure efficient utilization of resources thanks to proper scheduling policies. Scheduling is a key to orchestrate the resources that different nodes in a Track or a path are using. Slotframes can be split in resource blocks reserving the needed capacity to certain flows. Periodic and bursty traffic can be handled independently in the schedule, using active and reactive policies and taking advantage of overprovisionned cells to measure reth excursion. Along a Track, resource blocks can be chained so nodes in previous hops transmit their data before the next packet comes. This provides a tight control to latency along a Track. Collision loss is avoided for best effort traffic by

overprovisionning resources, giving time to the management plane of the network to dedicate more resources if needed.

5.2.2.1. Centralized Path Computation

In a controlled environment, a 6TiSCH device usually does not place a request for bandwidth between itself and another device in the network. Rather, an Operation Control System (OCS) invoked through an Human/Machine Interface (HMI) iprovides the Traffic Specification, in particular in terms of latency and reliability, and the end nodes, to a Path Computation element (PCE). With this, the PCE computes a Track between the end nodes and provisions every hop in the Track with per-flow state that describes the per-hop operation for a given packet, the corresponding timeSlots, and the flow identification to recognize which packet is placed in which Track, sort out duplicates, etc. In Figure 2, an example of Operational Control System and HMI is depicted.

For a static configuration that serves a certain purpose for a long period of time, it is expected that a node will be provisioned in one shot with a full schedule, which incorporates the aggregation of its behavior for multiple Tracks. The 6TiSCH Architecture expects that the programing of the schedule is done over CoAP as discussed in "6TiSCH Resource Management and Interaction using CoAP" [I-D.ietf-6tisch-coap].

But an Hybrid mode may be required as well whereby a single Track is added, modified, or removed, for instance if it appears that a Track does not perform as expected for, say, Packet Delivery Ratio (PDR). For that case, the expectation is that a protocol that flows along a Track (to be), in a fashion similar to classical Traffic Engineering (TE) [CCAMP], may be used to update the state in the devices. 6TiSCH provides means for a device to negotiate a timeSlot with a neighbor, but in general that flow was not designed and no protocol was selected and it is expected that DetNet will determine the appropriate end-to-end protocols to be used in that case.

Stream Management Entity

Operational Control System and HMI




Figure 2

5.2.2.1.1. Packet Marking and Handling

Section "Packet Marking and Handling" of [I-D.ietf-6tisch-architecture] describes the packet tagging and marking that is expected in 6TiSCH networks.

5.2.2.1.1.1. Tagging Packets for Flow Identification

For packets that are routed by a PCE along a Track, the tuple formed by the IPv6 source address and a local RPLInstanceID is tagged in the packets to identify uniquely the Track and associated transmit bundle of timeSlots.

It results that the tagging that is used for a DetNet flow outside the 6TiSCH LLN MUST be swapped into 6TiSCH formats and back as the packet enters and then leaves the 6TiSCH network.

Note: The method and format used for encoding the RPLInstanceID at 6lo is generalized to all 6TiSCH topological Instances, which includes Tracks.

5.2.2.1.1.2. Replication, Retries and Elimination

PRE establishes several paths in a network to provide redundancy and parallel transmissions to bound the end-to-end delay. Considering the scenario shown in Figure 3, many different paths are possible for S to reach R. A simple way to benefit from this topology could be to use the two independent paths via nodes A, C, E and via B, D, F. But more complex paths are possible as well.




Figure 3: A Typical Ladder Shape with Two Parallel Paths Toward the Destination

By employing a Packet Replication function, each node forwards a copy of each data packet over two different branches. For instance, in Figure 4, the source node S transmits the data packet to nodes A and B, in two different timeslots within the same TSCH slotframe.




Figure 4: Packet Replication: S transmits twice the same data packet, to its DP (A) and to its AP (B).

By employing Packet Elimination function once a node receives the first copy of a data packet, it discards the subsequent copies. Because the first copy that reaches a node is the one that matters, it is the only copy that will be forwarded upward.

Considering that the wireless medium is broadcast by nature, any neighbor of a transmitter may overhear a transmission. By employing the Promiscuous Overhearing function, nodes will have multiple opportunities to receive a given data packet. For instance, in Figure 4, when the source node S transmits the data packet to node A, node B may overhear this transmission.

6TiSCH expects elimination and replication of packets along a complex Track, but has no position about how the sequence numbers would be tagged in the packet.

As it goes, 6TiSCH expects that timeSlots corresponding to copies of a same packet along a Track are correlated by configuration, and does not need to process the sequence numbers.

The semantics of the configuration MUST enable correlated timeSlots to be grouped for transmit (and respectively receive) with a 'OR' relations, and then a 'AND' relation MUST be configurable between

groups. The semantics is that if the transmit (and respectively receive) operation succeeded in one timeSlot in a 'OR' group, then all the other timeSlots in the group are ignored. Now, if there are at least two groups, the 'AND' relation between the groups indicates that one operation must succeed in each of the groups.

On the transmit side, timeSlots provisioned for retries along a same branch of a Track are placed a same 'OR' group. The 'OR' relation indicates that if a transmission is acknowledged, then further transmissions SHOULD NOT be attempted for timeSlots in that group. There are as many 'OR' groups as there are branches of the Track departing from this node. Different 'OR' groups are programmed for the purpose of replication, each group corresponding to one branch of the Track. The 'AND' relation between the groups indicates that transmission over any of branches MUST be attempted regardless of whether a transmission succeeded in another branch. It is also possible to place cells to different next-hop routers in a same 'OR' group. This allows to route along multi-path Tracks, trying one next-hop and then another only if sending to the first fails.

On the receive side, all timeSlots are programmed in a same 'OR' group. Retries of a same copy as well as converging branches for elimination are converged, meaning that the first successful reception is enough and that all the other timeSlots can be ignored.

5.2.2.1.1.3. Differentiated Services Per-Hop-Behavior

Additionally, an IP packet that is sent along a Track uses the Differentiated Services Per-Hop-Behavior Group called Deterministic Forwarding, as described in [I-D.svshah-tsvwg-deterministic-forwarding].

5.2.2.1.2. Topology and capabilities

6TiSCH nodes are usually IoT devices, characterized by very limited amount of memory, just enough buffers to store one or a few IPv6 packets, and limited bandwidth between peers. It results that a node will maintain only a small number of peering information, and will not be able to store many packets waiting to be forwarded. Peers can be identified through MAC or IPv6 addresses.

Neighbors can be discovered over the radio using mechanism such as Enhanced Beacons, but, though the neighbor information is available in the 6TiSCH interface data model, 6TiSCH does not describe a protocol to pro-actively push the neighborhood information to a PCE. This protocol should be described and should operate over CoAP. The protocol should be able to carry multiple metrics, in particular the same metrics as used for RPL operations [RFC6551].

The energy that the device consumes in sleep, transmit and receive modes can be evaluated and reported. So can the amount of energy that is stored in the device and the power that it can be scavenged from the environment. The PCE SHOULD be able to compute Tracks that will implement policies on how the energy is consumed, for instance balance between nodes, ensure that the spent energy does not exceeded the scavenged energy over a period of time, etc...

5.2.2.1.3. Schedule Management by a PCE

6TiSCH supports a mixed model of centralized routes and distributed routes. Centralized routes can for example be computed by a entity such as a PCE [PCE]. Distributed routes are computed by RPL [RFC6550].

Both methods may inject routes in the Routing Tables of the 6TiSCH routers. In either case, each route is associated with a 6TiSCH topology that can be a RPL Instance topology or a Track. The 6TiSCH topology is indexed by a Instance ID, in a format that reuses the RPLInstanceID as defined in RPL.

Both RPL and PCE rely on shared sources such as policies to define Global and Local RPLInstanceIDs that can be used by either method. It is possible for centralized and distributed routing to share a same topology. Generally they will operate in different slotFrames, and centralized routes will be used for scheduled traffic and will have precedence over distributed routes in case of conflict between the slotFrames.

5.2.2.1.4. SlotFrames and Priorities

A slotFrame is the base object that a PCE needs to manipulate to program a schedule into an LLN node. Elaboration on that concept can be found in section "SlotFrames and Priorities" of [I-D.ietf-6tisch-architecture]

IEEE802.15.4 TSCH avoids contention on the medium by formatting time and frequencies in cells of transmission of equal duration. In order to describe that formatting of time and frequencies, the 6TiSCH architecture defines a global concept that is called a Channel Distribution and Usage (CDU) matrix; a CDU matrix is a matrix of cells with an height equal to the number of available channels (indexed by ChannelOffsets) and a width (in timeSlots) that is the period of the network scheduling operation (indexed by slotOffsets) for that CDU matrix. The size of a cell is a timeSlot duration, and values of 10 to 15 milliseconds are typical in 802.15.4 TSCH to accommodate for the transmission of a frame and an acknowledgement, including the security validation on the receive side which may take up to a few milliseconds on some device architecture.

The frequency used by a cell in the matrix rotates in a pseudo-random fashion, from an initial position at an epoch time, as the matrix iterates over and over.

A CDU matrix is computed by the PCE, but unallocated timeSlots may be used opportunistically by the nodes for classical best effort IP traffic. The PCE has precedence in the allocation in case of a conflict.

In a given network, there might be multiple CDU matrices that operate with different width, so they have different durations and represent different periodic operations. It is recommended that all CDU matrices in a 6TiSCH domain operate with the same cell duration and are aligned, so as to reduce the chances of interferences from slotted-aloha operations. The PCE MUST compute the CDU matrices and shared that knowledge with all the nodes. The matrices are used in particular to define slotFrames.

A slotFrame is a MAC-level abstraction that is common to all nodes and contains a series of timeSlots of equal length and precedence. It is characterized by a slotFrame_ID, and a slotFrame_size. A slotFrame aligns to a CDU matrix for its parameters, such as number and duration of timeSlots.

Multiple slotFrames can coexist in a node schedule, i.e., a node can have multiple activities scheduled in different slotFrames, based on the precedence of the 6TiSCH topologies. The slotFrames may be aligned to different CDU matrices and thus have different width. There is typically one slotFrame for scheduled traffic that has the highest precedence and one or more slotFrame(s) for RPL traffic. The timeSlots in the slotFrame are indexed by the SlotOffset; the first cell is at SlotOffset 0.

The 6TiSCH architecture introduces the concept of chunks ([I-D.ietf-6tisch-architecture]) to operate such spectrum distribution for a whole group of cells at a time. The CDU matrix is formatted into a set of chunks, each of them identified uniquely by a chunk-ID, see Figure 5. The PCE MUST compute the partitioning of CDU matrices into chunks and shared that knowledge with all the nodes in a 6TiSCH network.

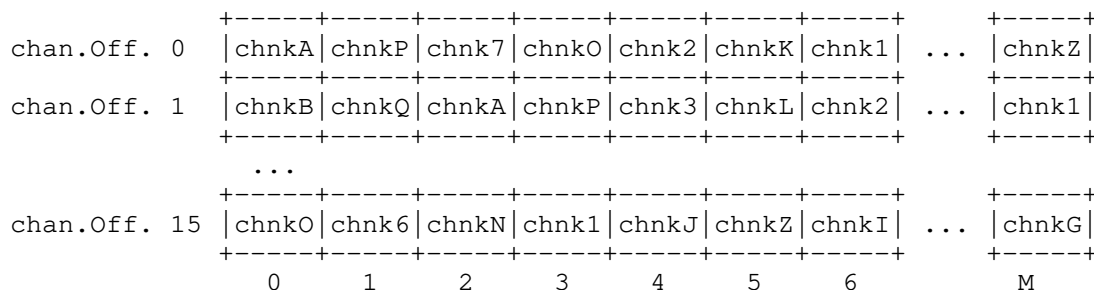


Figure 5: CDU matrix Partitioning in Chunks

The appropriation of a chunk can be requested explicitly by the PCE to any node. After a successful appropriation, the PCE owns the cells in that chunk, and may use them as hard cells to set up Tracks. Then again, 6TiSCH did not propose a method for chunk definition and a protocol for appropriation. This is to be done at RAW.

5.2.2.2. 6TiSCH Tracks

A Track at 6TiSCH is the application to wireless of the concept of a path in the Detnet architecture [I-D.ietf-detnet-architecture]. A Track can follow a simple sequence of relay nodes or can be structured as a more complex Destination Oriented Directed Acyclic Graph (DODAG) to a unicast destination. Along a Track, 6TiSCH nodes reserve the resources to enable the efficient transmission of packets while aiming to optimize certain properties such as reliability and ensure small jitter or bounded latency. The Track structure enables Layer-2 forwarding schemes, reducing the overhead of taking routing decisions at the Layer-3.

Serial Tracks can be understood as the concatenation of cells or bundles along a routing path from a source towards a destination. The serial Track concept is analogous to the circuit concept where resources are chained through the multi-hop topology. For example, A bundle of Tx Cells in a particular node is paired to a bundle of Rx Cells in the next hop node following a routing path.

Whereas scheduling ensures reliable delivery in bounded time along any Track, high availability requires the application of PAREO functions along a more complex DODAG Track structure. A DODAG has forking and joining nodes where the concepts such as Replication and Elimination can be exploited. Spatial redundancy increases the overall energy consumption in the network but improves significantly the availability of the network as well as the packet delivery ratio. A Track may also branch off and rejoin, for the purpose of the so-called Packet Replication and Elimination (PRE), over non congruent branches. PRE may be used to complement layer-2 Automatic Repeat reQuest (ARQ) and receiver-end Ordering to form the PAREO functions. PAREO functions enable to meet industrial expectations in PDR within bounded delivery time over a Track that includes wireless links, even when the Track extends beyond the 6TiSCH network.

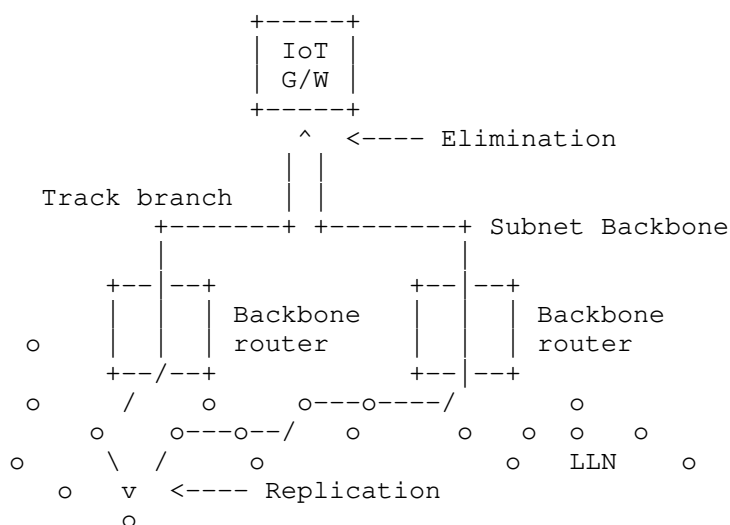


Figure 6: End-to-End deterministic Track

In the example above (see Figure 6), a Track is laid out from a field device in a 6TiSCH network to an IoT gateway that is located on a IEEE802.1 TSN backbone.

The Replication function in the field device sends a copy of each packet over two different branches, and a PCE schedules each hop of both branches so that the two copies arrive in due time at the gateway. In case of a loss on one branch, hopefully the other copy of the packet still makes it in due time. If two copies make it to the IoT gateway, the Elimination function in the gateway ignores the extra packet and presents only one copy to upper layers.

At each 6TiSCH hop along the Track, the PCE may schedule more than one timeSlot for a packet, so as to support Layer-2 retries (ARQ). It is also possible that the field device only uses the second branch if sending over the first branch fails.

In current deployments, a TSCH Track does not necessarily support PRE but is systematically multi-path. This means that a Track is scheduled so as to ensure that each hop has at least two forwarding solutions, and the forwarding decision is to try the preferred one and use the other in case of Layer-2 transmission failure as detected by ARQ.

Methods to implement complex Tracks are described in [I-D.papadopoulos-paw-pre-reqs] and complemented by extensions to the RPL routing protocol in [I-D.ietf-roll-nsa-extension] for best effort traffic, but a centralized routing technique such as promoted in DetNet is still missing.

5.2.2.2.1. Track Scheduling Protocol

Section "Schedule Management Mechanisms" of the 6TiSCH architecture describes 4 paradigms to manage the TSCH schedule of the LLN nodes: Static Scheduling, neighbor-to-neighbor Scheduling, remote monitoring and scheduling management, and Hop-by-hop scheduling. The Track operation for DetNet corresponds to a remote monitoring and scheduling management by a PCE.

Early work at 6TiSCH on a data model and a protocol to program the schedule in the 6TiSCH device was never concluded as the group focussed on best effort traffic. This work would be revived by RAW:

The 6top interface document [RFC8480] (to be reopened at RAW) was intended to specify the generic data model that can be used to monitor and manage resources of the 6top sublayer. Abstract methods were suggested for use by a management entity in the device. The data model also enables remote control operations on the 6top sublayer.

[I-D.ietf-6tisch-coap] (to be reopened at RAW) was intended to define a mapping of the 6top set of commands, which is described in RFC 8480, to CoAP resources. This allows an entity to interact with the 6top layer of a node that is multiple hops away in a RESTful fashion.

[I-D.ietf-6tisch-coap] also defined a basic set CoAP resources and associated RESTful access methods (GET/PUT/POST/DELETE). The payload (body) of the CoAP messages is encoded using the CBOR format. The PCE commands are expected to be issued directly as

CoAP requests or to be mapped back and forth into CoAP by a gateway function at the edge of the 6TiSCH network. For instance, it is possible that a mapping entity on the backbone transforms a non-CoAP protocol such as PCEP into the RESTful interfaces that the 6TiSCH devices support.

5.2.2.2.2. Track Forwarding

By forwarding, this specification means the per-packet operation that allows to deliver a packet to a next hop or an upper layer in this node. Forwarding is based on pre-existing state that was installed as a result of the routing computation of a Track by a PCE. The 6TiSCH architecture supports three different forwarding model, G-MPLS Track Forwarding (TF), 6LoWPAN Fragment Forwarding (FF) and IPv6 Forwarding (6F) which is the classical IP operation [I-D.ietf-6tisch-architecture]. The DetNet case relates to the Track Forwarding operation under the control of a PCE.

A Track is a unidirectional path between a source and a destination. In a Track cell, the normal operation of IEEE802.15.4 Automatic Repeat-reQuest (ARQ) usually happens, though the acknowledgment may be omitted in some cases, for instance if there is no scheduled cell for a retry.

Track Forwarding is the simplest and fastest. A bundle of cells set to receive (RX-cells) is uniquely paired to a bundle of cells that are set to transmit (TX-cells), representing a layer-2 forwarding state that can be used regardless of the network layer protocol. This model can effectively be seen as a Generalized Multi-protocol Label Switching (G-MPLS) operation in that the information used to switch a frame is not an explicit label, but rather related to other properties of the way the packet was received, a particular cell in the case of 6TiSCH. As a result, as long as the TSCH MAC (and Layer-2 security) accepts a frame, that frame can be switched regardless of the protocol, whether this is an IPv6 packet, a 6LoWPAN fragment, or a frame from an alternate protocol such as WirelessHART or ISA100.11a.

A data frame that is forwarded along a Track normally has a destination MAC address that is set to broadcast - or a multicast address depending on MAC support. This way, the MAC layer in the intermediate nodes accepts the incoming frame and 6top switches it without incurring a change in the MAC header. In the case of IEEE802.15.4, this means effectively broadcast, so that along the Track the short address for the destination of the frame is set to 0xFFFF.

A Track is thus formed end-to-end as a succession of paired bundles, a receive bundle from the previous hop and a transmit bundle to the next hop along the Track, and a cell in such a bundle belongs to at most one Track. For a given iteration of the device schedule, the effective channel of the cell is obtained by adding a pseudo-random number to the channelOffset of the cell, which results in a rotation of the frequency that used for transmission. The bundles may be computed so as to accommodate both variable rates and retransmissions, so they might not be fully used at a given iteration of the schedule. The 6TiSCH architecture provides additional means to avoid waste of cells as well as overflows in the transmit bundle, as follows:

In one hand, a TX-cell that is not needed for the current iteration may be reused opportunistically on a per-hop basis for routed packets. When all of the frame that were received for a given Track are effectively transmitted, any available TX-cell for that Track can be reused for upper layer traffic for which the next-hop router matches the next hop along the Track. In that case, the cell that is being used is effectively a TX-cell from the Track, but the short address for the destination is that of the next-hop router. It results that a frame that is received in a RX-cell of a Track with a destination MAC address set to this node as opposed to broadcast must be extracted from the Track and delivered to the upper layer (a frame with an unrecognized MAC address is dropped at the lower MAC layer and thus is not received at the 6top sublayer).

On the other hand, it might happen that there are not enough TX-cells in the transmit bundle to accommodate the Track traffic, for instance if more retransmissions are needed than provisioned. In that case, the frame can be placed for transmission in the bundle that is used for layer-3 traffic towards the next hop along the Track as long as it can be routed by the upper layer, that is, typically, if the frame transports an IPv6 packet. The MAC address should be set to the next-hop MAC address to avoid confusion. It results that a frame that is received over a layer-3 bundle may be in fact associated to a Track. In a classical IP link such as an Ethernet, off-Track traffic is typically in excess over reservation to be routed along the non-reserved path based on its QoS setting. But with 6TiSCH, since the use of the layer-3 bundle may be due to transmission failures, it makes sense for the receiver to recognize a frame that should be re-Tracked, and to place it back on the appropriate bundle if possible. A frame should be re-Tracked if the Per-Hop-Behavior group indicated in the Differentiated Services Field in the IPv6 header is set to Deterministic Forwarding, as discussed in Section 5.2.2.1.1. A frame is re-Tracked by scheduling it for transmission over the transmit bundle associated to the Track, with the destination MAC address set to broadcast.

There are 2 modes for a Track, transport mode and tunnel mode.

5.2.2.2.2.1. Transport Mode

In transport mode, the Protocol Data Unit (PDU) is associated with flow-dependant meta-data that refers uniquely to the Track, so the 6top sublayer can place the frame in the appropriate cell without ambiguity. In the case of IPv6 traffic, this flow identification is transported in the Flow Label of the IPv6 header. Associated with the source IPv6 address, the Flow Label forms a globally unique identifier for that particular Track that is validated at egress before restoring the destination MAC address (DMAC) and punting to the upper layer.




Figure 7: Track Forwarding, Transport Mode

5.2.2.2.2.2. Tunnel Mode

In tunnel mode, the frames originate from an arbitrary protocol over a compatible MAC that may or may not be synchronized with the 6TiSCH network. An example of this would be a router with a dual radio that is capable of receiving and sending WirelessHART or ISA100.11a frames with the second radio, by presenting itself as an Access Point or a Backbone Router, respectively.

In that mode, some entity (e.g. PCE) can coordinate with a WirelessHART Network Manager or an ISA100.11a System Manager to specify the flows that are to be transported transparently over the Track.




Figure 8: Track Forwarding, Tunnel Mode

In that case, the flow information that identifies the Track at the ingress 6TiSCH router is derived from the RX-cell. The dmac is set to this node but the flow information indicates that the frame must be tunneled over a particular Track so the frame is not passed to the upper layer. Instead, the dmac is forced to broadcast and the frame is passed to the 6top sublayer for switching.

At the egress 6TiSCH router, the reverse operation occurs. Based on metadata associated to the Track, the frame is passed to the appropriate link layer with the destination MAC restored.

5.2.2.2.2.3. Tunnel Metadata

Metadata coming with the Track configuration is expected to provide the destination MAC address of the egress endpoint as well as the tunnel mode and specific data depending on the mode, for instance a service access point for frame delivery at egress. If the tunnel egress point does not have a MAC address that matches the configuration, the Track installation fails.

In transport mode, if the final layer-3 destination is the tunnel termination, then it is possible that the IPv6 address of the destination is compressed at the 6LoWPAN sublayer based on the MAC address. It is thus mandatory at the ingress point to validate that the MAC address that was used at the 6LoWPAN sublayer for compression matches that of the tunnel egress point. For that reason, the node

that injects a packet on a Track checks that the destination is effectively that of the tunnel egress point before it overwrites it to broadcast. The 6top sublayer at the tunnel egress point reverts that operation to the MAC address obtained from the tunnel metadata.

5.2.2.2.2.4. OAM

An Overview of Operations, Administration, and Maintenance (OAM) Tools [RFC7276] provides an overview of the existing tooling for OAM [RFC6291]. Tracks are complex paths and new tooling is necessary to manage them, with respect to load control, timing, and the Packet Replication and Elimination Functions (PREF).

An example of such tooling can be found in the context of BIER [RFC8279] and more specifically BIER Traffic Engineering [I-D.ietf-bier-te-arch] (BIER-TE): [I-D.thubert-bier-replication-elimination] leverages BIER-TE to control the process of PREF, and to provide traceability of these operations, in the deterministic dataplane, along a complex Track. For the 6TiSCH type of constrained environment, [I-D.thubert-6lo-bier-dispatch] enables an efficient encoding of the BIER bitmap within the 6LoRH framework.

6. 5G

6.1. Provenance and Documents

The 3rd Generation Partnership Project (3GPP) incorporates many companies whose business is related to cellular network operation as well as network equipment and device manufacturing. All generations of 3GPP technologies provide scheduled wireless segments, primarily in licensed spectrum which is beneficial for reliability and availability.

In 2016, the 3GPP started to design New Radio (NR) technology belonging to the fifth generation (5G) of cellular networks. NR has been designed from the beginning to not only address enhanced Mobile Broadband (eMBB) services for consumer devices such as smart phones or tablets but is also tailored for future Internet of Things (IoT) communication and connected cyber-physical systems. In addition to eMBB, requirement categories have been defined on Massive Machine-Type Communication (M-MTC) for a large number of connected devices/sensors, and Ultra-Reliable Low-Latency Communication (URLLC) for connected control systems and critical communication as illustrated in Figure 9. It is the URLLC capabilities that make 5G a great candidate for reliable low-latency communication. With these three corner stones, NR is a complete solution supporting the connectivity needs of consumers, enterprises, and public sector for both wide area and local area, e.g. indoor deployments. A general overview of NR can be found in [TS38300].




Figure 9: 5G Application Areas

As a result of releasing the first NR specification in 2018 (Release 15), it has been proven by many companies that NR is a URLLC-capable technology and can deliver data packets at 10^{-5} packet error rate within 1ms latency budget [TR37910]. Those evaluations were consolidated and forwarded to ITU to be included in the [IMT2020] work.

In order to understand communication requirements for automation in vertical domains, 3GPP studied different use cases [TR22804] and released technical specification with reliability, availability and latency demands for a variety of applications [TS22104].

As an evolution of NR, multiple studies have been conducted in scope of 3GPP Release 16 including the following two, focusing on radio aspects:

1. Study on physical layer enhancements for NR ultra-reliable and low latency communication (URLLC) [TR38824].
2. Study on NR industrial Internet of Things (I-IoT) [TR38825].

In addition, several enhancements have been done on system architecture level which are reflected in System architecture for the 5G System (5GS) [TS23501].

6.2. General Characteristics

The 5G Radio Access Network (5G RAN) with its NR interface includes several features to achieve Quality of Service (QoS), such as a guaranteeably low latency or tolerable packet error rates for selected data flows. Determinism is achieved by centralized admission control and scheduling of the wireless frequency resources, which are typically licensed frequency bands assigned to a network operator.

NR enables short transmission slots in a radio subframe, which benefits low-latency applications. NR also introduces mini-slots, where prioritized transmissions can be started without waiting for slot boundaries, further reducing latency. As part of giving priority and faster radio access to URLLC traffic, NR introduces preemption where URLLC data transmission can preempt ongoing non-URLLC transmissions. Additionally, NR applies very fast processing, enabling retransmissions even within short latency bounds.

NR defines extra-robust transmission modes for increased reliability both for data and control radio channels. Reliability is further improved by various techniques, such as multi-antenna transmission, the use of multiple frequency carriers in parallel and packet duplication over independent radio links. NR also provides full mobility support, which is an important reliability aspect not only for devices that are moving, but also for devices located in a changing environment.

Network slicing is seen as one of the key features for 5G, allowing vertical industries to take advantage of 5G networks and services. Network slicing is about transforming a Public Land Mobile Network (PLMN) from a single network to a network where logical partitions are created, with appropriate network isolation, resources, optimized topology and specific configuration to serve various service

requirements. An operator can configure and manage the mobile network to support various types of services enabled by 5G, for example eMBB and URLLC, depending on the different customers' needs.

Exposure of capabilities of 5G Systems to the network or applications outside the 3GPP domain have been added to Release 16 [TS23501]. Via exposure interfaces, applications can access 5G capabilities, e.g., communication service monitoring and network maintenance.

For several generations of mobile networks, 3GPP has considered how the communication system should work on a global scale with billions of users, taking into account resilience aspects, privacy regulation, protection of data, encryption, access and core network security, as well as interconnect. Security requirements evolve as demands on trustworthiness increase. For example, this has led to the introduction of enhanced privacy protection features in 5G. 5G also employs strong security algorithms, encryption of traffic, protection of signaling and protection of interfaces.

One particular strength of mobile networks is the authentication, based on well-proven algorithms and tightly coupled with a global identity management infrastructure. Since 3G, there is also mutual authentication, allowing the network to authenticate the device and the device to authenticate the network. Another strength is secure solutions for storage and distribution of keys fulfilling regulatory requirements and allowing international roaming. When connecting to 5G, the user meets the entire communication system, where security is the result of standardization, product security, deployment, operations and management as well as incident handling capabilities. The mobile networks approach the entirety in a rather coordinated fashion which is beneficial for security.

6.3. Deployment and Spectrum

The 5G system allows deployment in a vast spectrum range, addressing use-cases in both wide-area as well as local networks. Furthermore, 5G can be configured for public and non-public access.

When it comes to spectrum, NR allows combining the merits of many frequency bands, such as the high bandwidths in millimeter Waves (mmW) for extreme capacity locally, as well as the broad coverage when using mid- and low frequency bands to address wide-area scenarios. URLLC is achievable in all these bands. Spectrum can be either licensed, which means that the license holder is the only authorized user of that spectrum range, or unlicensed, which means that anyone who wants to use the spectrum can do so.

A prerequisite for critical communication is performance predictability, which can be achieved by the full control of the access to the spectrum, which 5G provides. Licensed spectrum guarantees control over spectrum usage by the system, making it a preferable option for critical communication. However, unlicensed spectrum can provide an additional resource for scaling non-critical communications. While NR is initially developed for usage of licensed spectrum, the functionality to access also unlicensed spectrum was introduced in 3GPP Release 16.

Licensed spectrum dedicated to mobile communications has been allocated to mobile service providers, i.e. issued as longer-term licenses by national administrations around the world. These licenses have often been associated with coverage requirements and issued across whole countries, or in large regions. Besides this, configured as a non-public network (NPN) deployment, 5G can provide network services also to a non-operator defined organization and its premises such as a factory deployment. By this isolation, quality of service requirements, as well as security requirements can be achieved. An integration with a public network, if required, is also possible. The non-public (local) network can thus be interconnected with a public network, allowing devices to roam between the networks.

In an alternative model, some countries are now in the process of allocating parts of the 5G spectrum for local use to industries. These non-service providers then have a choice of applying for a local license themselves and operating their own network or cooperating with a public network operator or service provider.

6.4. Applicability to Deterministic Flows

6.4.1. System Architecture

The 5G system [TS23501] consists of the User Equipment (UE) at the terminal side, and the Radio Access Network (RAN) with the gNB as radio base station node, as well as the Core Network (CN). The core network is based on a service-based architecture with the central functions: Access and Mobility Management Function (AMF), Session Management Function (SMF) and User Plane Function (UPF) as illustrated in Figure 10.

The gNB's main responsibility is the radio resource management, including admission control and scheduling, mobility control and radio measurement handling. The AMF handles the UE's connection status and security, while the SMF controls the UE's data sessions. The UPF handles the user plane traffic.

The SMF can instantiate various Packet Data Unit (PDU) sessions for the UE, each associated with a set of QoS flows, i.e., with different QoS profiles. Segregation of those sessions is also possible, e.g., resource isolation in the RAN and in the CN can be defined (slicing).

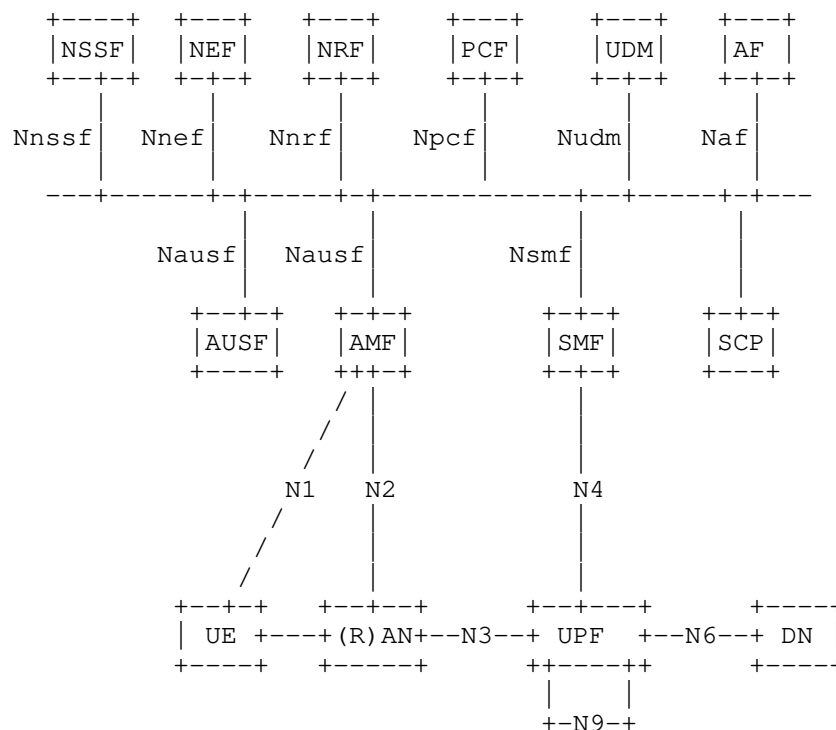


Figure 10: 5G System Architecture

To allow UE mobility across cells/gNBs, handover mechanisms are supported in NR. For an established connection, i.e., connected mode mobility, a gNB can configure a UE to report measurements of received signal strength and quality of its own and neighbouring cells, periodically or event-based. Based on these measurement reports, the gNB decides to handover a UE to another target cell/gNB. Before triggering the handover, it is hand-shaked with the target gNB based on network signalling. A handover command is then sent to the UE and the UE switches its connection to the target cell/gNB. The Packet Data Convergence Protocol (PDCP) of the UE can be configured to avoid data loss in this procedure, i.e., handle retransmissions if needed. Data forwarding is possible between source and target gNB as well. To improve the mobility performance further, i.e., to avoid connection failures, e.g., due to too-late handovers, the mechanism of conditional handover is introduced in Release 16 specifications.

Therein a conditional handover command, defining a triggering point, can be sent to the UE before UE enters a handover situation. A further improvement introduced in Release 16 is the Dual Active Protocol Stack (DAPS), where the UE maintains the connection to the source cell while connecting to the target cell. This way, potential interruptions in packet delivery can be avoided entirely.

6.4.2. Overview of The Radio Protocol Stack

The protocol architecture for NR consists of the L1 Physical layer (PHY) and as part of the L2, the sublayers of Medium Access Control (MAC), Radio Link Control (RLC), Packet Data Convergence Protocol (PDCP), as well as the Service Data Adaption Protocol (SDAP).

The PHY layer handles signal processing related actions, such as encoding/decoding of data and control bits, modulation, antenna precoding and mapping.

The MAC sub-layer handles multiplexing and priority handling of logical channels (associated with QoS flows) to transport blocks for PHY transmission, as well as scheduling information reporting and error correction through Hybrid Automated Repeat Request (HARQ).

The RLC sublayer handles sequence numbering of higher layer packets, retransmissions through Automated Repeat Request (ARQ), if configured, as well as segmentation and reassembly and duplicate detection.

The PDCP sublayer consists of functionalities for ciphering/deciphering, integrity protection/verification, re-ordering and in-order delivery, duplication and duplicate handling for higher layer packets, and acts as the anchor protocol to support handovers.

The SDAP sublayer provides services to map QoS flows, as established by the 5G core network, to data radio bearers (associated with logical channels), as used in the 5G RAN.

Additionally, in RAN, the Radio Resource Control (RRC) protocol, handles the access control and configuration signalling for the aforementioned protocol layers. RRC messages are considered L3 and thus transmitted also via those radio protocol layers.

To provide low latency and high reliability for one transmission link, i.e., to transport data (or control signaling) of one radio bearer via one carrier, several features have been introduced on the user plane protocols for PHY and L2, as explained in the following.

6.4.3. Radio (PHY)

NR is designed with native support of antenna arrays utilizing benefits from beamforming, transmissions over multiple MIMO layers and advanced receiver algorithms allowing effective interference cancellation. Those antenna techniques are the basis for high signal quality and effectiveness of spectral usage. Spatial diversity with up to 4 MIMO layers in UL and up to 8 MIMO layers in DL is supported. Together with spatial-domain multiplexing, antenna arrays can focus power in desired direction to form beams. NR supports beam management mechanisms to find the best suitable beam for UE initially and when it is moving. In addition, gNBs can coordinate their respective DL and UL transmissions over the backhaul network keeping interference reasonably low, and even make transmissions or receptions from multiple points (multi-TRP). Multi-TRP can be used for repetition of data packet in time, in frequency or over multiple MIMO layers which can improve reliability even further.

Any downlink transmission to a UE starts from resource allocation signaling over the Physical Downlink Control Channel (PDCCH). If it is successfully received, the UE will know about the scheduled transmission and may receive data over the Physical Downlink Shared Channel (PDSCH). If retransmission is required according to the HARQ scheme, a signaling of negative acknowledgement (NACK) on the Physical Uplink Control Channel (PUCCH) is involved and PDCCH together with PDSCH transmissions (possibly with additional redundancy bits) are transmitted and soft-combined with previously received bits. Otherwise, if no valid control signaling for scheduling data is received, nothing is transmitted on PUCCH (discontinuous transmission - DTX), and the base station upon detecting DTX will retransmit the initial data.

An uplink transmission normally starts from a Scheduling Request (SR) - a signaling message from the UE to the base station sent via PUCCH. Once the scheduler is informed about buffer data in UE, e.g., by SR, the UE transmits a data packet on the Physical Uplink Shared Channel (PUSCH). Pre-scheduling not relying on SR is also possible (see following section).

Since transmission of data packets require usage of control and data channels, there are several methods to maintain the needed reliability. NR uses Low Density Parity Check (LDPC) codes for data channels, Polar codes for PDCCH, as well as orthogonal sequences and Polar codes for PUCCH. For ultra-reliability of data channels, very robust (low spectral efficiency) Modulation and Coding Scheme (MCS) tables are introduced containing very low (down to 1/20) LDPC code rates using BPSK or QPSK. Also, PDCCH and PUCCH channels support multiple code rates including very low ones for the channel robustness.

A connected UE reports downlink (DL) quality to gNB by sending Channel State Information (CSI) reports via PUCCH while uplink (UL) quality is measured directly at gNB. For both uplink and downlink, gNB selects the desired MCS number and signals it to the UE by Downlink Control Information (DCI) via PDCCH channel. For URLLC services, the UE can assist the gNB by advising that MCS targeting 10^{-5} Block Error Rate (BLER) are used. Robust link adaptation algorithms can maintain the needed level of reliability considering a given latency bound.

Low latency on the physical layer is provided by short transmission duration which is possible by using high Subcarrier Spacing (SCS) and the allocation of only one or a few Orthogonal Frequency Division Multiplexing (OFDM) symbols. For example, the shortest latency for the worst case in DL can be 0.23ms and in UL can be 0.24ms according to (section 5.7.1 in [TR37910]). Moreover, if the initial transmission has failed, HARQ feedback can quickly be provided and an HARQ retransmission is scheduled.

Dynamic multiplexing of data associated with different services is highly desirable for efficient use of system resources and to maximize system capacity. Assignment of resources for eMBB is usually done with regular (longer) transmission slots, which can lead to blocking of low latency services. To overcome the blocking, eMBB resources can be pre-empted and re-assigned to URLLC services. In this way, spectrally efficient assignments for eMBB can be ensured while providing flexibility required to ensure a bounded latency for URLLC services. In downlink, the gNB can notify the eMBB UE about pre-emption after it has happened, while in uplink there are two pre-emption mechanisms: special signaling to cancel eMBB transmission and URLLC dynamic power boost to suppress eMBB transmission.

6.4.4. Scheduling and QoS (MAC)

One integral part of the 5G system is the Quality of Service (QoS) framework [TS23501]. QoS flows are setup by the 5G system for certain IP or Ethernet packet flows, so that packets of each flow receive the same forwarding treatment, i.e., in scheduling and admission control. QoS flows can for example be associated with different priority level, packet delay budgets and tolerable packet error rates. Since radio resources are centrally scheduled in NR, the admission control function can ensure that only those QoS flows are admitted for which QoS targets can be reached.

NR transmissions in both UL and DL are scheduled by the gNB [TS38300]. This ensures radio resource efficiency, fairness in resource usage of the users and enables differentiated treatment of the data flows of the users according to the QoS targets of the flows. Those QoS flows are handled as data radio bearers or logical channels in NR RAN scheduling.

The gNB can dynamically assign DL and UL radio resources to users, indicating the resources as DL assignments or UL grants via control channel to the UE. Radio resources are defined as blocks of OFDM symbols in spectral domain and time domain. Different lengths are supported in time domain, i.e., (multiple) slot or mini-slot lengths. Resources of multiple frequency carriers can be aggregated and jointly scheduled to the UE.

Scheduling decisions are based, e.g., on channel quality measured on reference signals and reported by the UE (cf. periodical CSI reports for DL channel quality). The transmission reliability can be chosen in the scheduling algorithm, i.e., by link adaptation where an appropriate transmission format (e.g., robustness of modulation and coding scheme, controlled UL power) is selected for the radio channel condition of the UE. Retransmissions, based on HARQ feedback, are also controlled by the scheduler. If needed to avoid HARQ round-trip time delays, repeated transmissions can be also scheduled beforehand, to the cost of reduced spectral efficiency.

In dynamic DL scheduling, transmission can be initiated immediately when DL data becomes available in the gNB. However, for dynamic UL scheduling, when data becomes available but no UL resources are available yet, the UE indicates the need for UL resources to the gNB via a (single bit) scheduling request message in the UL control channel. When thereupon UL resources are scheduled to the UE, the UE can transmit its data and may include a buffer status report, indicating the exact amount of data per logical channel still left to be sent. More UL resources may be scheduled accordingly. To avoid the latency introduced in the scheduling request loop, UL radio resources can also be pre-scheduled.

In particular for periodical traffic patterns, the pre-scheduling can rely on the scheduling features DL Semi-Persistent Scheduling (SPS) and UL Configured Grant (CG). With these features, periodically recurring resources can be assigned in DL and UL. Multiple parallels of those configurations are supported, in order to serve multiple parallel traffic flows of the same UE.

To support QoS enforcement in the case of mixed traffic with different QoS requirements, several features have recently been introduced. This way, e.g., different periodical critical QoS flows can be served together with best effort transmissions, by the same UE. Among others, these features (partly Release 16) are: 1) UL logical channel transmission restrictions allowing to map logical channels of certain QoS only to intended UL resources of a certain frequency carrier, slot-length, or CG configuration, and 2) intra-UE pre-emption, allowing critical UL transmissions to pre-empt non-critical transmissions.

When multiple frequency carriers are aggregated, duplicate parallel transmissions can be employed (beside repeated transmissions on one carrier). This is possible in the Carrier Aggregation (CA) architecture where those carriers originate from the same gNB, or in the Dual Connectivity (DC) architecture where the carriers originate from different gNBs, i.e., the UE is connected to two gNBs in this case. In both cases, transmission reliability is improved by this means of providing frequency diversity.

In addition to licensed spectrum, a 5G system can also utilize unlicensed spectrum to offload non-critical traffic. This version of NR is called NR-U, part of 3GPP Release 16. The central scheduling approach applies also for unlicensed radio resources, but in addition also the mandatory channel access mechanisms for unlicensed spectrum, e.g., Listen Before Talk (LBT) are supported in NR-U. This way, by using NR, operators have and can control access to both licensed and unlicensed frequency resources.

6.4.5. Time-Sensitive Networking (TSN) Integration

The main objective of Time-Sensitive Networking (TSN) is to provide guaranteed data delivery within a guaranteed time window, i.e., bounded low latency. IEEE 802.1 TSN [IEEE802.1TSN] is a set of open standards that provide features to enable deterministic communication on standard IEEE 802.3 Ethernet [IEEE802.3]. TSN standards can be seen as a toolbox for traffic shaping, resource management, time synchronization, and reliability.

A TSN stream is a data flow between one end station (Talker) to another end station (Listener). In the centralized configuration model, TSN bridges are configured by the Central Network Controller (CNC) [IEEE802.1Qcc] to provide deterministic connectivity for the TSN stream through the network. Time-based traffic shaping provided by Scheduled Traffic [IEEE802.1Qbv] may be used to achieve bounded low latency. The TSN tool for time synchronization is the generalized Precision Time Protocol (gPTP) [IEEE802.1AS]), which provides reliable time synchronization that can be used by end stations and by other TSN tools, e.g., Scheduled Traffic [IEEE802.1Qbv]. High availability, as a result of ultra-reliability, is provided for data flows by the Frame Replication and Elimination for Reliability (FRER) [IEEE802.1CB] mechanism.

3GPP Release 16 includes integration of 5G with TSN, i.e., specifies functions for the 5G System (5GS) to deliver TSN streams such that they meet their QoS requirements. A key aspect of the integration is the 5GS appears from the rest of the network as a set of TSN bridges, in particular, one virtual bridge per User Plane Function (UPF) on the user plane. The 5GS includes TSN Translator (TT) functionality for the adaptation of the 5GS to the TSN bridged network and for hiding the 5GS internal procedures. The 5GS provides the following components:

1. interface to TSN controller, as per [IEEE802.1Qcc] for the fully centralized configuration model
2. time synchronization via reception and transmission of gPTP PDUs [IEEE802.1AS]
3. low latency, hence, can be integrated with Scheduled Traffic [IEEE802.1Qbv]
4. reliability, hence, can be integrated with FRER [IEEE802.1CB]

Figure 10 shows an illustration of 5G-TSN integration where an industrial controller (Ind Ctrlr) is connected to industrial Input/Output devices (I/O dev) via 5G. The 5GS can directly transport Ethernet frames since Release 15, thus, end-to-end Ethernet connectivity is provided. The 5GS implements the required interfaces towards the TSN controller functions such as the CNC, thus adapts to the settings of the TSN network. A 5G user plane virtual bridge interconnects TSN bridges or connect end stations, e.g., I/O devices to the network. Note that the introduction of 5G brings flexibility in various aspects, e.g., more flexible network topology because a wireless hop can replace several wireline hops thus significantly reduce the number of hops end-to-end. [ETR5GTSN] dives more into the integration of 5G with TSN.

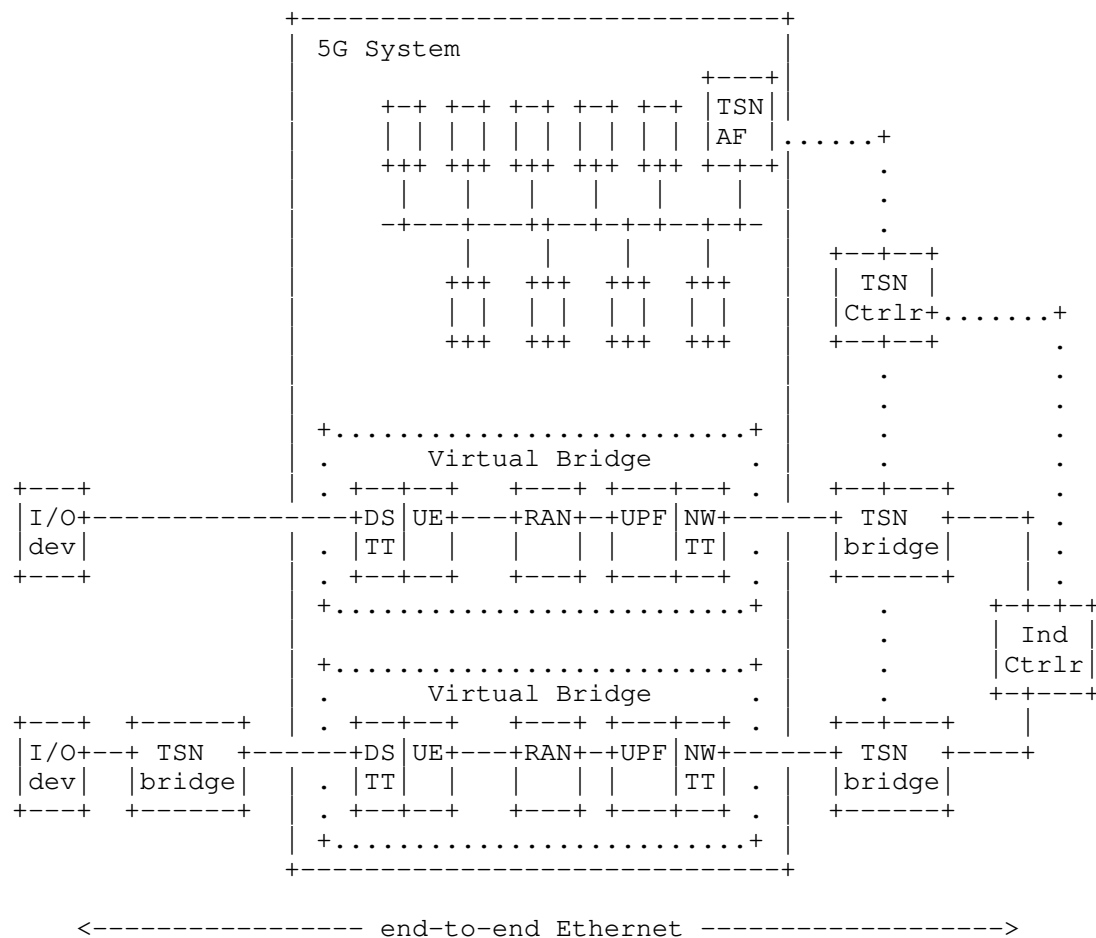


Figure 11: 5G - TSN Integration

NR supports accurate reference time synchronization in μ s accuracy level. Since NR is a scheduled system, an NR UE and a gNB are tightly synchronized to their OFDM symbol structures. A 5G internal reference time can be provided to the UE via broadcast or unicast signaling, associating a known OFDM symbol to this reference clock. The 5G internal reference time can be shared within the 5G network, i.e., radio and core network components. For the interworking with gPTP for multiple time domains, the 5GS acts as a virtual gPTP time-aware system and supports the forwarding of gPTP time synchronization information between end stations and bridges through the 5G user plane TTs. These account for the residence time of the 5GS in the time synchronization procedure. One special option is when the 5GS internal reference time is not only used within the 5GS, but also to the rest of the devices in the deployment, including connected TSN bridges and end stations.

Redundancy architectures were specified in order to provide reliability against any kind of failure on the radio link or nodes in the RAN and the core network, Redundant user plane paths can be provided based on the dual connectivity architecture, where the UE sets up two PDU sessions towards the same data network, and the 5G system makes the paths of the two PDU sessions independent as illustrated in Figure 13. There are two PDU sessions involved in the solution: the first spans from the UE via gNB1 to UPF1, acting as the first PDU session anchor, while the second spans from the UE via gNB2 to UPF2, acting as second the PDU session anchor. The independent paths may continue beyond the 3GPP network. Redundancy Handling Functions (RHF)s are deployed outside of the 5GS, i.e., in Host A (the device) and in Host B (the network). RHF can implement replication and elimination functions as per [IEEE802.1CB] or the Packet Replication, Elimination, and Ordering Functions (PREOF) of IETF Deterministic Networking (DetNet) [RFC8655].




Figure 12: Reliability with Single UE

An alternative solution is that multiple UEs per device are used for user plane redundancy as illustrated in Figure 13. Each UE sets up a PDU session. The 5GS ensures that those PDU sessions of the different UEs are handled independently internal to the 5GS. There is no single point of failure in this solution, which also includes RHF outside of the 5G system, e.g., as per FRER or as PREOF specifications.




Figure 13: Reliability with Dual UE

Note that the abstraction provided by the RHF and the location of the RHF being outside of the 5G system make 5G equally supporting integration for reliability both with FRER of TSN and PREOF of DetNet as they both rely on the same concept.

Note also that TSN is the primary subnetwork technology for DetNet. Thus, the DetNet over TSN work, e.g., [I-D.ietf-detnet-ip-over-tsn], can be leveraged via the TSN support built in 5G.

6.5. Summary

5G technology enables deterministic communication. Based on the centralized admission control and the scheduling of the wireless resources, licensed or unlicensed, quality of service such as latency and reliability can be guaranteed. 5G contains several features to achieve ultra-reliable and low latency performance, e.g., support for different OFDM numerologies and slot-durations, as well as fast processing capabilities and redundancy techniques that lead to achievable latency numbers of below 1ms with reliability guarantees up to 99.999%.

5G also includes features to support Industrial IoT use cases, e.g., via the integration of 5G with TSN. This includes 5G capabilities for each TSN component, latency, resource management, time synchronization, and reliability. Furthermore, 5G support for TSN

can be leveraged when 5G is used as subnet technology for DetNet, in combination with or instead of TSN, which is the primary subnet for DetNet. In addition, the support for integration with TSN reliability was added to 5G by making DetNet reliability also applicable, thus making 5G DetNet ready. Moreover, providing IP service is native to 5G.

Overall, 5G provides scheduled wireless segments with high reliability and availability. In addition, 5G includes capabilities for integration to IP networks.

7. L-band Digital Aeronautical Communications System

One of the main pillars of the modern Air Traffic Management (ATM) system is the existence of a communication infrastructure that enables efficient aircraft guidance and safe separation in all phases of flight. Although current systems are technically mature, they are suffering from the VHF band's increasing saturation in high-density areas and the limitations posed by analogue radio. Therefore, aviation globally and the European Union (EU) in particular, strives for a sustainable modernization of the aeronautical communication infrastructure.

In the long-term, ATM communication shall transition from analogue VHF voice and VDL2 communication to more spectrum efficient digital data communication. The European ATM Master Plan foresees this transition to be realized for terrestrial communications by the development and implementation of the L-band Digital Aeronautical Communications System (LDACS). LDACS shall enable IPv6 based air-ground communication related to the safety and regularity of the flight. The particular challenge is that no new frequencies can be made available for terrestrial aeronautical communication. It was thus necessary to develop procedures to enable the operation of LDACS in parallel with other services in the same frequency band.

7.1. Provenance and Documents

The development of LDACS has already made substantial progress in the Single European Sky ATM Research (SESAR) framework, and is currently being continued in the follow-up program, SESAR2020 [RIH18]. A key objective of the SESAR activities is to develop, implement and validate a modern aeronautical data link able to evolve with aviation needs over long-term. To this end, an LDACS specification has been produced [GRA19] and is continuously updated; transmitter demonstrators were developed to test the spectrum compatibility of LDACS with legacy systems operating in the L-band [SAJ14]; and the overall system performance was analyzed by computer simulations, indicating that LDACS can fulfil the identified requirements [GRA11].

LDACS standardization within the framework of the International Civil Aviation Organization (ICAO) started in December 2016. The ICAO standardization group has produced an initial Standards and Recommended Practices (SARPs) document [ICAO18]. The SARPs document defines the general characteristics of LDACS. The ICAO standardization group plans to produce an ICAO technical manual - the ICAO equivalent to a technical standard - within the next years. Generally, the group is open to input from all sources and develops LDACS in the open.

Up to now the LDACS standardization has been focused on the development of the physical layer and the data link layer, only recently have higher layers come into the focus of the LDACS development activities. There is currently no "IPv6 over LDACS" specification; however, SESAR2020 has started the testing of IPv6-based LDACS testbeds. The IPv6 architecture for the aeronautical telecommunication network is called the Future Communications Infrastructure (FCI). FCI shall support quality of service, diversity, and mobility under the umbrella of the "multi-link concept". This work is conducted by ICAO working group WG-I.

In addition to standardization activities several industrial LDACS prototypes have been built. One set of LDACS prototypes has been evaluated in flight trials confirming the theoretical results predicting the system performance [GRA18][SCH19].

7.2. General Characteristics

LDACS will become one of several wireless access networks connecting aircraft to the Aeronautical Telecommunications Network (ATN). The LDACS access network contains several ground stations, each of them providing one LDACS radio cell. The LDACS air interface is a cellular data link with a star-topology connecting aircraft to ground-stations with a full duplex radio link. Each ground-station is the centralized instance controlling all air-ground communications within its radio cell. A ground-station supports up to 512 aircraft.

The LDACS air interface protocol stack defines two layers, the physical layer and the data link layer.

The physical layer provides the means to transfer data over the radio channel. The LDACS ground-station supports bi-directional links to multiple aircraft under its control. The forward link direction (FL; ground-to-air) and the reverse link direction (RL; air-to-ground) are separated by frequency division duplex. Forward link and reverse link use a 500 kHz channel each. The ground-station transmits a continuous stream of OFDM symbols on the forward link. In the reverse link different aircraft are separated in time and frequency

using a combination of Orthogonal Frequency-Division Multiple-Access (OFDMA) and Time-Division Multiple-Access (TDMA). Aircraft thus transmit discontinuously on the reverse link with radio bursts sent in precisely defined transmission opportunities allocated by the ground-station. LDACS does not support beam-forming or Multiple Input Multiple Output (MIMO).

The data-link layer provides the necessary protocols to facilitate concurrent and reliable data transfer for multiple users. The LDACS data link layer is organized in two sub-layers: The medium access sub-layer and the logical link control sub-layer. The medium access sub-layer manages the organization of transmission opportunities in slots of time and frequency. The logical link control sub-layer provides acknowledged point-to-point logical channels between the aircraft and the ground-station using an automatic repeat request protocol. LDACS supports also unacknowledged point-to-point channels and ground-to-air broadcast.

The user data rate of LDACS is 315 kbit/s to 1428 kbit/s on the forward link, and 294 kbit/s to 1390 kbit/s on the reverse link, depending on coding and modulation. Due to strong interference from legacy systems in the L-band, the most robust coding and modulation should be expected for initial deployment i.e. 315/294 kbit/s on the forward/reverse link, respectively.

Since LDACS has been mainly designed for air traffic management communication it supports mutual entity authentication, integrity and confidentiality capabilities of user data messages and some control channel protection capabilities [MAE19].

7.3. Applicability to Deterministic Flows

LDACS has been designed with applications related to the safety and regularity of the flight in mind. It has therefore been designed as a deterministic wireless data link (as far as possible).

LDACS medium access is always under the control of the ground-station of a radio cell. Any medium access for the transmission of user data has to be requested with a resource request message stating the requested amount of resources and class of service. The ground-station performs resource scheduling on the basis of these requests and grants resources with resource allocation messages. Resource request and allocation messages are exchanged over dedicated contention-free control channels.

LDACS has two mechanisms to request resources from the scheduler in the ground-station. Resources can either be requested "on demand" with a given class of service. On the forward link, this is done

locally in the ground-station, on the reverse link a dedicated contention-free control channel is used (Dedicated Control Channel (DCCH); roughly 83 bit every 60 ms). A resource allocation is always announced in the control channel of the forward link (Common Control Channel (CCCH); variable sized). Due to the spacing of the reverse link control channels of every 60 ms, a medium access delay in the same order of magnitude is to be expected.

Resources can also be requested "permanently". The permanent resource request mechanism supports requesting recurring resources in given time intervals. A permanent resource request has to be canceled by the user (or by the ground-station, which is always in control). User data transmissions over LDACS are therefore always scheduled by the ground-station, while control data uses statically (i.e. at net entry) allocated recurring resources (DCCH and CCCH). The current specification documents specify no scheduling algorithm. However performance evaluations so far have used strict priority scheduling and round robin for equal priorities for simplicity. In the current prototype implementations LDACS classes of service are thus realized as priorities of medium access and not as flows. Note that this can starve out low priority flows. However, this is not seen as a big problem since safety related message always go first in any case. Scheduling of reverse link resources is done in physical Protocol Data Units (PDU) of 112 bit (or larger if more aggressive coding and modulation is used). Scheduling on the forward link is done Byte-wise since the forward link is transmitted continuously by the ground-station.

In order to support diversity, LDACS supports handovers to other ground-stations on different channels. Handovers may be initiated by the aircraft (break-before-make) or by the ground-station (make-before-break) if it is connected to an alternative ground-station via the same ground-station controller. Beyond this, FCI diversity shall be implemented by the multi-link concept.

8. IANA Considerations

This specification does not require IANA action.

9. Security Considerations

Most RAW technologies integrate some authentication or encryption mechanisms that were defined outside the IETF.

10. Contributors

Georgios Z. Papadopoulos: Contributed to the TSCH section.

Nils Mäurer: Contributed to the LDACS section.

Thomas Gräupl: Contributed to the LDACS section.

Janos Farkas, Torsten Dudda, Alexey Shapin, and Sara Sandberg: Contributed to the 5G section.

11. Acknowledgments

Many thanks to the participants of the RAW WG where a lot of the work discussed here happened.

12. Normative References

- [RFC8480] Wang, Q., Ed., Vilajosana, X., and T. Watteyne, "6TiSCH Operation Sublayer (6top) Protocol (6P)", RFC 8480, DOI 10.17487/RFC8480, November 2018, <<https://www.rfc-editor.org/info/rfc8480>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [RFC5673] Pister, K., Ed., Thubert, P., Ed., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low-Power and Lossy Networks", RFC 5673, DOI 10.17487/RFC5673, October 2009, <<https://www.rfc-editor.org/info/rfc5673>>.
- [I-D.ietf-detnet-architecture] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", Work in Progress, Internet-Draft, draft-ietf-detnet-architecture-13, 6 May 2019, <<https://tools.ietf.org/html/draft-ietf-detnet-architecture-13>>.
- [I-D.ietf-6tisch-architecture] Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", Work in Progress, Internet-Draft, draft-ietf-6tisch-architecture-29, 27 August 2020, <<https://tools.ietf.org/html/draft-ietf-6tisch-architecture-29>>.

13. Informative References

- [RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <<https://www.rfc-editor.org/info/rfc6550>>.
- [RFC6551] Vasseur, JP., Ed., Kim, M., Ed., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics Used for Path Calculation in Low-Power and Lossy Networks", RFC 6551, DOI 10.17487/RFC6551, March 2012, <<https://www.rfc-editor.org/info/rfc6551>>.
- [RFC6291] Andersson, L., van Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "Guidelines for the Use of the "OAM" Acronym in the IETF", BCP 161, RFC 6291, DOI 10.17487/RFC6291, June 2011, <<https://www.rfc-editor.org/info/rfc6291>>.
- [RFC7276] Mizrahi, T., Sprecher, N., Bellagamba, E., and Y. Weingarten, "An Overview of Operations, Administration, and Maintenance (OAM) Tools", RFC 7276, DOI 10.17487/RFC7276, June 2014, <<https://www.rfc-editor.org/info/rfc7276>>.
- [RFC8279] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", RFC 8279, DOI 10.17487/RFC8279, November 2017, <<https://www.rfc-editor.org/info/rfc8279>>.
- [I-D.ietf-6tisch-msf] Chang, T., Vucinic, M., Vilajosana, X., Duquennoy, S., and D. Dujovne, "6TiSCH Minimal Scheduling Function (MSF)", Work in Progress, Internet-Draft, draft-ietf-6tisch-msf-18, 12 September 2020, <<https://tools.ietf.org/html/draft-ietf-6tisch-msf-18>>.
- [I-D.pthubert-raw-architecture] Thubert, P., Papadopoulos, G., and R. Buddenberg, "Reliable and Available Wireless Architecture/Framework", Work in Progress, Internet-Draft, draft-pthubert-raw-architecture-04, 6 July 2020, <<https://tools.ietf.org/html/draft-pthubert-raw-architecture-04>>.

[I-D.ietf-roll-nsa-extension]

Koutsiamanis, R., Papadopoulos, G., Montavont, N., and P. Thubert, "Common Ancestor Objective Function and Parent Set DAG Metric Container Extension", Work in Progress, Internet-Draft, draft-ietf-roll-nsa-extension-09, 26 September 2020, <<https://tools.ietf.org/html/draft-ietf-roll-nsa-extension-09>>.

[I-D.papadopoulos-paw-pre-reqs]

Papadopoulos, G., Koutsiamanis, R., Montavont, N., and P. Thubert, "Exploiting Packet Replication and Elimination in Complex Tracks in LLNs", Work in Progress, Internet-Draft, draft-papadopoulos-paw-pre-reqs-01, 25 March 2019, <<https://tools.ietf.org/html/draft-papadopoulos-paw-pre-reqs-01>>.

[I-D.thubert-bier-replication-elimination]

Thubert, P., Eckert, T., Brodard, Z., and H. Jiang, "BIER-TE extensions for Packet Replication and Elimination Function (PREF) and OAM", Work in Progress, Internet-Draft, draft-thubert-bier-replication-elimination-03, 3 March 2018, <<https://tools.ietf.org/html/draft-thubert-bier-replication-elimination-03>>.

[I-D.thubert-6lo-bier-dispatch]

Thubert, P., Brodard, Z., Jiang, H., and G. Texier, "A 6loRH for BitStrings", Work in Progress, Internet-Draft, draft-thubert-6lo-bier-dispatch-06, 28 January 2019, <<https://tools.ietf.org/html/draft-thubert-6lo-bier-dispatch-06>>.

[I-D.ietf-bier-te-arch]

Eckert, T., Cauchie, G., and M. Menth, "Tree Engineering for Bit Index Explicit Replication (BIER-TE)", Work in Progress, Internet-Draft, draft-ietf-bier-te-arch-08, 13 July 2020, <<https://tools.ietf.org/html/draft-ietf-bier-te-arch-08>>.

[I-D.ietf-6tisch-coap]

Sudhaakar, R. and P. Zand, "6TiSCH Resource Management and Interaction using CoAP", Work in Progress, Internet-Draft, draft-ietf-6tisch-coap-03, 9 March 2015, <<https://tools.ietf.org/html/draft-ietf-6tisch-coap-03>>.

- [I-D.svshah-tsvwg-deterministic-forwarding]
Shah, S. and P. Thubert, "Deterministic Forwarding PHB",
Work in Progress, Internet-Draft, draft-svshah-tsvwg-
deterministic-forwarding-04, 30 August 2015,
<[https://tools.ietf.org/html/draft-svshah-tsvwg-
deterministic-forwarding-04](https://tools.ietf.org/html/draft-svshah-tsvwg-deterministic-forwarding-04)>.
- [IEEE Std. 802.15.4]
IEEE standard for Information Technology, "IEEE Std.
802.15.4, Part. 15.4: Wireless Medium Access Control (MAC)
and Physical Layer (PHY) Specifications for Low-Rate
Wireless Personal Area Networks".
- [IEEE Std. 802.11]
"IEEE Standard 802.11 - IEEE Standard for Information
Technology - Telecommunications and information exchange
between systems Local and metropolitan area networks -
Specific requirements - Part 11: Wireless LAN Medium
Access Control (MAC) and Physical Layer (PHY)
Specifications.".
- [IEEE Std. 802.11ak]
"802.11ak: Enhancements for Transit Links Within Bridged
Networks", 2017.
- [IEEE Std. 802.11ax]
"802.11ax D4.0: Enhancements for High Efficiency WLAN".
- [IEEE Std. 802.11ay]
"802.11ay: Enhanced throughput for operation in license-
exempt bands above 45 GHz".
- [IEEE Std. 802.11ad]
"802.11ad: Enhancements for very high throughput in the 60
GHz band".
- [IEEE 802.11be WIP]
"802.11be: Extreme High Throughput".
- [IEEE Std. 802.1Qat]
"802.1Qat: Stream Reservation Protocol".
- [IEEE8021Qcc]
"802.1Qcc: IEEE Standard for Local and Metropolitan Area
Networks--Bridges and Bridged Networks -- Amendment 31:
Stream Reservation Protocol (SRP) Enhancements and
Performance Improvements".

- [Cavalcanti_2019]
Dave Cavalcanti et al., "Extending Time Distribution and Timeliness Capabilities over the Air to Enable Future Wireless Industrial Automation Systems, the Proceedings of IEEE", June 2019.
- [Nitsche_2015]
Thomas Nitsche et al., "IEEE 802.11ad: directional 60 GHz communication for multi-Gigabit-per-second Wi-Fi", December 2014.
- [Ghasempour_2017]
Yasaman Ghasempour et al., "802.11ay: Next-Generation 60 GHz Communications for 100 Gb/s Wi-Fi", December 2017.
- [IEEE_doc_11-18-2009-06]
"802.11 Real-Time Applications (RTA) Topic Interest Group (TIG) Report", November 2018.
- [IEEE_doc_11-19-0373-00]
Kevin Stanton et al., "Time-Sensitive Applications Support in EHT", March 2019.
- [morell13] Antoni Morell et al., "Label switching over IEEE802.15.4e networks", April 2013.
- [dearmas16]
Jesica de Armas et al., "Determinism through path diversity: Why packet replication makes sense", September 2016.
- [vilajosana19]
Xavier Vilajosana et al., "6TiSCH: Industrial Performance for IPv6 Internet-of-Things Networks", June 2019.
- [ISA100.11a]
ISA/IEC, "ISA100.11a, Wireless Systems for Automation, also IEC 62734", 2011, <<http://www.isa100wci.org/en-US/Documents/PDF/3405-ISA100-WirelessSystems-Future-broch-WEB-ETSI.aspx>>.
- [WirelessHART]
www.hartcomm.org, "Industrial Communication Networks - Wireless Communication Network and Communication Profiles - WirelessHART - IEC 62591", 2010.
- [PCE]
IETF, "Path Computation Element", <<https://dataTracker.ietf.org/doc/charter-ietf-pce/>>.

- [CCAMP] IETF, "Common Control and Measurement Plane",
<<https://dataTracker.ietf.org/doc/charter-ietf-ccamp/>>.
- [TiSCH] IETF, "IPv6 over the TSCH mode over 802.15.4",
<<https://dataTracker.ietf.org/doc/charter-ietf-6tisch/>>.
- [RIH18] Rihacek, C., Haindl, B., Fantappie, P., Pierattelli, S.,
Gräupl, T., Schnell, M., and N. Fistas, "L-band Digital
Aeronautical Communications System (LDACS) Activities in
SESAR2020", Proceedings of the Integrated Communications
Navigation and Surveillance Conference (ICNS) Herndon, VA,
USA, April 2018.
- [GRA19] Gräupl, T., Rihacek, C., and B. Haindl, "LDACS A/G
Specification", SESAR2020 PJ14-02-01 D3.3.010, February
2019.
- [SAJ14] Sajatovic, M., Günzel, H., and S. Müller, "WA04 D22 Test
Report for Assessing LDACS1 Transmitter Impact upon DME/
TACAN Receivers", April 2014.
- [GRA11] Gräupl, T. and M. Ehammer, "L-DACS1 Data Link Layer
Evolution of ATN/IPS", Proceedings of the 30th IEEE/AIAA
Digital Avionics Systems Conference (DASC) Seattle, WA,
USA, October 2011.
- [ICAO18] International Civil Aviation Organization (ICAO), "L-Band
Digital Aeronautical Communication System (LDACS)",
International Standards and Recommended Practices Annex 10
- Aeronautical Telecommunications, Vol. III -
Communication Systems, July 2018.
- [GRA18] al., T. G. E., "L-band Digital Aeronautical Communications
System (LDACS) flight trials in the national German
project MICONAV", Proceedings of the Integrated
Communications, Navigation, Surveillance Conference
(ICNS) Herndon, VA, USA, April 2018.
- [SCH19] Schnell, M., "DLR tests digital communications
technologies combined with additional navigation functions
for the first time", 3 March 2019,
<https://www.dlr.de/dlr/en/desktopdefault.aspx/tabid-10081/151_read-32951/#/gallery/33877>.

- [MAE19] Mäurer, N. and C. Schmitt, "DLR tests digital communications technologies combined with additional navigation functions for the first time", Proceedings of the Integrated Communications, Navigation, Surveillance Conference (ICNS) Washington D.C., USA, April 2019.
- [TR37910] "3GPP TR 37.910, Study on self evaluation towards IMT-2020 submission",
<<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3190>>.
- [TR38824] "3GPP TR 38.824, Study on physical layer enhancements for NR ultra-reliable and low latency case (URLLC)",
<<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3498>>.
- [TR38825] "3GPP TR 38.825, Study on NR industrial Internet of Things (IoT)",
<<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3492>>.
- [TS22104] "3GPP TS 22.104, Service requirements for cyber-physical control applications in vertical domains",
<<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3528>>.
- [TR22804] "3GPP TR 22.804, Study on Communication for Automation in Vertical domains (CAV)",
<<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3187>>.
- [TS23501] "3GPP TS 23.501, System architecture for the 5G System (5GS)",
<<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144>>.
- [TS38300] "3GPP TS 38.300, NR Overall description",
<<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3191>>.
- [IMT2020] "ITU towards IMT for 2020 and beyond",
<<https://www.itu.int/en/ITU-R/study-groups/rsg5/rwp5d/imt-2020/Pages/default.aspx>>.
- [RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019,
<<https://www.rfc-editor.org/info/rfc8655>>.

[I-D.ietf-detnet-ip-over-tsn]

Varga, B., Farkas, J., Malis, A., and S. Bryant, "DetNet Data Plane: IP over IEEE 802.1 Time Sensitive Networking (TSN)", Work in Progress, Internet-Draft, draft-ietf-detnet-ip-over-tsn-03, 8 June 2020, <<https://tools.ietf.org/html/draft-ietf-detnet-ip-over-tsn-03>>.

[IEEE802.1TSN]

IEEE 802.1, "Time-Sensitive Networking (TSN) Task Group", <<http://www.ieee802.org/1/pages/tsn.html>>.

[IEEE802.1AS]

IEEE, "IEEE Standard for Local and metropolitan area networks -- Timing and Synchronization for Time-Sensitive Applications", IEEE 802.1AS-2020, <https://standards.ieee.org/content/ieee-standards/en/standard/802_1AS-2020.html>.

[IEEE802.1CB]

IEEE, "IEEE Standard for Local and metropolitan area networks -- Frame Replication and Elimination for Reliability", DOI 10.1109/IEEESTD.2017.8091139, IEEE 802.1CB-2017, <<https://ieeexplore.ieee.org/document/8091139>>.

[IEEE802.1Qbv]

IEEE, "IEEE Standard for Local and metropolitan area networks -- Bridges and Bridged Networks -- Amendment 25: Enhancements for Scheduled Traffic", IEEE 802.1Qbv-2015, <<https://ieeexplore.ieee.org/document/7440741>>.

[IEEE802.1Qcc]

IEEE, "IEEE Standard for Local and metropolitan area networks -- Bridges and Bridged Networks -- Amendment 31: Stream Reservation Protocol (SRP) Enhancements and Performance Improvements", IEEE 802.1Qcc-2018, <<https://ieeexplore.ieee.org/document/8514112>>.

[IEEE802.3]

IEEE, "IEEE Standard for Ethernet", IEEE 802.3-2018, <<https://ieeexplore.ieee.org/document/8457469>>.

[ETR5GTSN] Farkas, J., Varga, B., Miklos, G., and J. Sachs, "5G-TSN integration meets networking requirements for industrial automation", Ericsson Technology Review, Volume 9, No 7, August 2019, <<https://www.ericsson.com/en/reports-and-papers/ericsson-technology-review/articles/5g-tsn-integration-for-industrial-automation>>.

Authors' Addresses

Pascal Thubert (editor)
Cisco Systems, Inc
Building D
45 Allee des Ormes - BP1200
06254 MOUGINS - Sophia Antipolis
France

Phone: +33 497 23 26 34
Email: pthubert@cisco.com

Dave Cavalcanti
Intel Corporation
2111 NE 25th Ave
Hillsboro, OR, 97124
United States of America

Phone: 503 712 5566
Email: dave.cavalcanti@intel.com

Xavier Vilajosana
Universitat Oberta de Catalunya
156 Rambla Poblenou
08018 Barcelona Catalonia
Spain

Email: xvilajosana@uoc.edu

Corinna Schmitt
Research Institute CODE, UniBwM
Werner-Heisenberg-Weg 28
85577 Neubiberg
Germany

Email: corinna.schmitt@unibw.de

Janos Farkas
Ericsson
Budapest
Magyar tudosok korutja 11
1117
Hungary

Email: janos.farkas@ericsson.com

RAW
Internet-Draft
Intended status: Informational
Expires: 6 August 2022

P. Thubert, Ed.
Cisco Systems
D. Cavalcanti
Intel
X. Vilajosana
Universitat Oberta de Catalunya
C. Schmitt
Research Institute CODE, UniBw M
J. Farkas
Ericsson
2 February 2022

Reliable and Available Wireless Technologies
draft-ietf-raw-technologies-05

Abstract

This document presents a series of recent technologies that are capable of time synchronization and scheduling of transmission, making them suitable to carry time-sensitive flows with high reliability and availability.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 6 August 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document.

Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
2. Terminology	4
3. Towards Reliable and Available Networks	6
3.1. Scheduling for Reliability	6
3.2. Diversity for Availability	6
3.3. Benefits of Scheduling	7
4. IEEE 802.11	8
4.1. Provenance and Documents	8
4.2. 802.11ax High Efficiency (HE)	10
4.2.1. General Characteristics	10
4.2.2. Applicability to deterministic flows	12
4.3. 802.11be Extreme High Throughput (EHT)	13
4.3.1. General Characteristics	14
4.3.2. Applicability to deterministic flows	14
4.4. 802.11ad and 802.11ay (mmWave operation)	16
4.4.1. General Characteristics	16
4.4.2. Applicability to deterministic flows	16
5. IEEE 802.15.4	17
5.1. Provenance and Documents	17
5.2. TimeSlotted Channel Hopping	18
5.2.1. General Characteristics	18
5.2.2. Applicability to Deterministic Flows	20
6. 5G	34
6.1. Provenance and Documents	34
6.2. General Characteristics	36
6.3. Deployment and Spectrum	38
6.4. Applicability to Deterministic Flows	39
6.4.1. System Architecture	39
6.4.2. Overview of The Radio Protocol Stack	41
6.4.3. Radio (PHY)	42
6.4.4. Scheduling and QoS (MAC)	44
6.4.5. Time-Sensitive Communications (TSC)	46
6.5. Summary	51
7. L-band Digital Aeronautical Communications System	52
7.1. Provenance and Documents	52
7.2. General Characteristics	53
7.3. Deployment and Spectrum	54
7.4. Applicability to Deterministic Flows	55
7.4.1. System Architecture	55
7.4.2. Overview of The Radio Protocol Stack	55

7.4.3. Radio (PHY)	56
7.4.4. Scheduling, Frame Structure and QoS (MAC)	57
7.5. Summary	59
8. IANA Considerations	60
9. Security Considerations	60
10. Contributors	60
11. Acknowledgments	60
12. Normative References	60
13. Informative References	61
Authors' Addresses	70

1. Introduction

When used in math or philosophy, the term "deterministic" generally refers to a perfection where all aspect are understood and predictable. A perfectly deterministic network would ensure that every packet reach its destination following a predetermined path along a predefined schedule to be delivered at the exact due time. In a real and imperfect world, a deterministic network must highly predictable, which is a combination of reliability and availability.

- * On the one hand the network must be reliable, meaning that it will perform as expected for all but very few packets and in particular that it will deliver the packets at the destination within a pre-defined time interval. .
- * On the other hand, the network must be available, meaning that it has to be resilient to any single outage, independently of the cause of the failure, be it software, hardware, or a external, e.g., a physical event impacting the transmission channel.

RAW (Reliable and Available Wireless) is an effort to provide Deterministic Networking [RFC8557] in networking environments where at least some segments of a path are wireless. Making Wireless Reliable and Available is even more challenging than it is with wires, due to the numerous causes of loss in transmission that add up to the congestion losses and the delays caused by overbooked shared resources. In order to maintain a similar quality of service along a multihop path that is composed of wired and wireless hops, additional methods that are specific to wireless must be leveraged to combat the sources of loss that are also specific to wireless.

Such wireless-specific methods include per-hop retransmissions (HARQ) and P2MP overhearing whereby multiple receivers are scheduled to receive the same transmission, which balances the adverse effects of the transmission losses that are experienced when a radio is used as pure P2P. Those methods are collectively referred to as Packet (hybrid) ARQ, Replication, Elimination and Ordering (PAREO) functions in the "Reliable and Available Wireless Architecture/Framework" [I-D.ietf-raw-architecture].

2. Terminology

The terms Reliable and Available in the context of RAW require some discussion; that discussion takes place in the RAW architecture [I-D.ietf-raw-architecture]. Summarized definitions are provided below.

This specification uses several terms that are uncommon on protocols that ensure best effort transmissions for stochastic flows, such as found in the traditional Internet and other statistically multiplexed packet networks.

ARQ: Automatic Repeat Request, enabling an acknowledged transmission and retries. ARQ is a typical model at Layer-2 on a wireless medium. It is typically avoided end-to-end on deterministic flows because it introduces excessive indetermination in latency, but a limited number of retries within a bounded time may be used over a wireless link and yet respect end-to-end constraints.

Availability: Availability is a measure of the relative amount of time where a path operates in stated condition, in other words $(\text{uptime}) / (\text{uptime} + \text{downtime})$.

Available: That is exempt of unscheduled outage, the expectation for a network being that the flow is maintained in the face of any single breakage.

Deterministic Networking See section 2 of [RFC8557].

Deterministic Network See section 4.1.2 of [RFC8655].

Deterministic Flow Identifier (L2) : A tuple identified by a `stream_handle`, and provided by a bridge, in accordance with IEEE 802.1CB. The tuple comprises at least `src MAC`, `dst MAC`, `VLAN ID`, and `L2 priority`. Continuous streams are characterized by bandwidth and max packet size; scheduled streams are characterized by a repeating pattern of timed transmissions.

Deterministic Flow Identifier (L3): See section 3.3 of [RFC8938].

The classical IP 5-tuple that identifies a flow comprises the src IP, dst IP, src port, dest port, and the upper layer protocol (ULP). DetNet uses a 6-tuple where the extra field is the DSCP field in the packet. The IPv6 flow label is not used. for that purpose.

Uplink: Connection from end-devices to a data communication equipment. In the context of wireless, uplink refers to the connection between a station (STA) and a controller (AP) or a User Equipment (UE) to a Base Station (BS).

Downlink: The reverse direction.

Traffic type profile (IEEE): Corresponds to the traffic classification identifier provided to a deterministic flow. Traffic classes receive an identifier numbered from 0 to 8 (N-1), where N is the number of outbound queues associated with a given Bridge port. Each traffic class as a one-to-one correspondence with a specific egress/outbound queue for a port.

FEC: Forward error correction, sending redundant coded data to help the receiver recover transmission errors without the delays incurred with ARQ.

HARQ: Hybrid ARQ, a combination of FEC and ARQ.

PCE: Path Computation Element.

PREOF : DetNet Packet Replication, Elimination and Ordering Functions.

PAREO: Packet (hybrid) ARQ, Replication, Elimination and Ordering. PAREO is a superset Of DetNet's PREOF that includes radio-specific techniques such as short range broadcast, MUMIMO, constructive interference and overhearing, which can be leveraged separately or combined to increase the reliability.

Reliability: Reliability is a measure of the probability that an item will perform its intended function for a specified interval under stated conditions. For RAW, the service that is expected is delivery within a bounded latency and a failure is when the packet is either lost or delivered too late.

Track: A networking graph that can be used as a "path" to transport RAW packets with equivalent treatment; a Track may fork and rejoin to enable the PAREO operations.

3. Towards Reliable and Available Networks

3.1. Scheduling for Reliability

A packet network is reliable for critical (e.g., time-sensitive) packets when the undesirable statistical effects that affect the transmission of those packets, e.g., delay or loss, are eliminated.

The reliability of a Deterministic Network [RFC8655] often relies on precisely applying a tight schedule that controls the use of time-shared resources such as CPUs and buffers, and maintains at all time the amount of the critical packets within the physical capabilities of the hardware and that of the transmission medium. The schedule can also be used to shape the flows by controlling the time of transmission of the packets that compose the flow at every hop.

To achieve this, there must be a shared sense of time throughout the network. The sense of time is usually provided by the lower layer and is not in scope for RAW. As an example, the Precision Time Protocol, standardized as IEEE 1588 and IEC 61588, has mapping to Ethernet but also to a number of industrial and Smartrid protocols through profiles. Wi-Fi relies on IEEE Std 802.1AS, which incorporates a PTP profile, for Fine Timing Measurement.

3.2. Diversity for Availability

Equipment failure, such as a switch or an access point rebooting, a broken wire or radio adapter, or a fixed obstacle to the transmission, can be the cause of multiple packets lost in a row before the flows are rerouted or the system may recover.

This is not acceptable for critical applications such as related to safety. A typical process control loop will tolerate an occasional packet loss, but a loss of several packets in a row will cause an emergency stop (e.g., after 4 packets lost, within a period of 1 second). In an amusement park, a continuous loss of packet for a few 100ms may trigger an automatic interruption of the ride and cause the evacuation and the reboot of the game.

Network Availability is obtained by making the transmission resilient against hardware failures and radio transmission losses due to uncontrolled events such as co-channel interferers, multipath fading or moving obstacles. The best results are typically achieved by pseudo-randomly cumulating all forms of diversity, in the spatial domain with replication and elimination, in the time domain with ARQ and diverse scheduled transmissions, and in the frequency domain with frequency hopping or channel hopping between frames.

3.3. Benefits of Scheduling

Scheduling redundant transmissions of the critical packets of diverse paths improves the resiliency against breakages and statistical transmission loss, such as due to cosmic particles on wires, and interferences on wireless.

When required, the worst case time of delivery can be guaranteed as part of the end-to-end schedule, and the sense of time that must be shared throughout the network can be exposed to and leveraged by other applications.

In addition, scheduling provides specific value over the wireless medium:

- * Scheduling allows a time-sharing operation, where every transmission is assigned its own time/frequency resource. IOW, sender and receiver are synchronized and scheduled to talk on a given frequency resource at a given time and for a given duration. This way, scheduling can avoid collisions between scheduled transmissions and enable a high ratio of critical traffic compared to QoS-based priority.
- * Scheduling can be used as a technoque for both time and frequency diversity (e.g., between retries), allowing the next transmission to happen on a different frequency as programmed in both sender and receiver. This is useful to defeat co-channel interference from un-controlled transmitters as well as multipath fading.
- * Transmissions can be also scheduled on multiple channels in parallel, which enables to use the full available spectrum while avoiding the hidden terminal problem, e.g., when the next packet in a same flow interferes on a same channel with the previous one that progressed a few hops farther.
- * On the other hand, scheduling optimizes the bandwidth usage: compared to classical Collision Avoidance techniques, there is no blank time related to inter-frame space (IFS) and exponential back-off in scheduled operations. A minimal Clear Channel Assessment may be needed to comply with the local regulations such as ETSI 300-328, but that will not detect a collision when the senders are synchronized.

- * Finally, scheduling plays a critical role to save energy. In IoT, energy is the foremost concern, and synchronizing sender and listener enables to always maintain them in deep sleep when there is no scheduled transmission. This avoids idle listening and long preambles and enables long sleep periods between traffic and resynchronization, allowing battery-operated nodes to operate in a mesh topology for multiple years.

4. IEEE 802.11

4.1. Provenance and Documents

With an active portfolio of nearly 1,300 standards and projects under development, IEEE is a leading developer of industry standards in a broad range of technologies that drive the functionality, capabilities, and interoperability of products and services, transforming how people live, work, and communicate.

The IEEE 802 LAN/MAN Standards Committee (SC) develops and maintains networking standards and recommended practices for local, metropolitan, and other area networks, using an open and accredited process, and advocates them on a global basis. The most widely used standards are for Ethernet, Bridging and Virtual Bridged LANs, Wireless LAN, Wireless PAN, Wireless MAN, Wireless Coexistence, Media Independent Handover Services, and Wireless RAN. An individual Working Group provides the focus for each area. Standards produced by the IEEE 802 SC are freely available from the IEEE GET Program after they have been published in PDF for six months.

The IEEE 802.11 Wireless LAN (WLAN) standards define the underlying MAC and PHY layers for the Wi-Fi technology. Wi-Fi/802.11 is one of the most successful wireless technologies, supporting many application domains. While previous 802.11 generations, such as 802.11n and 802.11ac, have focused mainly on improving peak throughput, more recent generations are also considering other performance vectors, such as efficiency enhancements for dense environments in 802.11ax, latency, reliability and enhancements supporting Time-Sensitive Networking (TSN) [IEEE802.1TSN] capabilities in P802.11be.

IEEE Std 802.11-2012 introduced support for TSN time synchronization based on IEEE 802.1AS over 802.11 Timing Measurement protocol. IEEE Std 802.11-2016 extended the 802.1AS operation over 802.11 Fine Timing Measurement (FTM), as well as the Stream Reservation Protocol (IEEE 802.1Qat). 802.11 WLANs can also be part of a 802.1Q bridged networks with enhancements enabled by the 802.11ak amendment now retrofitted in IEEE Std 802.11-2020. Traffic classification based on 802.1Q VLAN tags is also supported in 802.11. Other 802.1 TSN

capabilities such as 802.1Qbv and 802.1CB, which are media agnostic, can already operate over 802.11. The IEEE Std 802.11ax-2021 adds new scheduling capabilities that can enhance the timeliness performance in the 802.11 MAC and achieve lower bounded latency. The IEEE 802.11be is undergoing efforts to enhance the support for 802.1 TSN capabilities especially related to worst-case latency, reliability and availability. The IEEE 802.11 working group has been working in collaboration with the IEEE 802.1 working group for several years extending some 802.1 features over 802.11. As with any wireless media, 802.11 imposes new constraints and restrictions to TSN-grade QoS, and tradeoffs between latency and reliability guarantees must be considered as well as managed deployment requirements. An overview of 802.1 TSN capabilities and challenges for their extensions to 802.11 are discussed in [Cavalcanti_2019].

Wi-Fi Alliance (WFA) is the worldwide network of companies that drives global Wi-Fi adoption and evolution through thought leadership, spectrum advocacy, and industry-wide collaboration. The WFA work helps ensure that Wi-Fi devices and networks provide users the interoperability, security, and reliability they have come to expect.

Avnu Alliance is also a global industry forum developing interoperability testing for TSN capable devices across multiple media including Ethernet, Wi-Fi, and 5G.

The following [IEEE Std 802.11] specifications/certifications are relevant in the context of reliable and available wireless services and support for time-sensitive networking capabilities:

Time Synchronization: IEEE802.11-2016 with IEEE802.1AS; WFA TimeSync Certification.

Congestion Control: IEEE Std 802.11-2016 Admission Control; WFA Admission Control.

Security: WFA Wi-Fi Protected Access, WPA2 and WPA3.

Interoperating with IEEE802.1Q bridges: IEEE Std 802.11-2020 incorporating 802.11ak.

Stream Reservation Protocol (part of [IEEE Std 802.1Qat]): AIEEE802.11-2016

Scheduled channel access: IEEE802.11ad Enhancements for very high throughput in the 60 GHz band [IEEE Std 802.11ad].

802.11 Real-Time Applications: Topic Interest Group (TIG) ReportDoc

[IEEE_doc_11-18-2009-06].

In addition, major amendments being developed by the IEEE802.11 Working Group include capabilities that can be used as the basis for providing more reliable and predictable wireless connectivity and support time-sensitive applications:

IEEE 802.11ax D4.0: Enhancements for High Efficiency (HE). [IEEE Std 802.11ax]

IEEE 802.11be Extreme High Throughput (EHT). [IEEE 802.11be WIP]

IEEE 802.11ay Enhanced throughput for operation in license-exempt bands above 45 GHz. [IEEE Std 802.11ay]

The main 802.11ax and 802.11be capabilities and their relevance to RAW are discussed in the remainder of this document.

4.2. 802.11ax High Efficiency (HE)

4.2.1. General Characteristics

The next generation Wi-Fi (Wi-Fi 6) is based on the IEEE802.11ax amendment [IEEE Std 802.11ax], which includes new capabilities to increase efficiency, control and reduce latency. Some of the new features include higher order 1024-QAM modulation, support for uplink multiple user (MU) multiple input multiple output (MIMO), orthogonal frequency-division multiple access (OFDMA), trigger-based access and Target Wake time (TWT) for enhanced power savings. The OFDMA mode and trigger-based access enable the AP, after reserving the channel using the clear channel assessment procedure for a given duration, to schedule multi-user transmissions, which is a key capability required to increase latency predictability and reliability for time-sensitive flows. 802.11ax can operate in up to 160 MHz channels and it includes support for operation in the new 6 GHz band, which is expected to be open to unlicensed use by the FCC and other regulatory agencies worldwide.

4.2.1.1. Multi-User OFDMA and Trigger-based Scheduled Access

802.11ax introduced a new OFDMA mode in which multiple users can be scheduled across the frequency domain. In this mode, the Access Point (AP) can initiate multi-user (MU) Uplink (UL) transmissions in the same PHY Protocol Data Unit (PPDU) by sending a trigger frame. This centralized scheduling capability gives the AP much more control of the channel in its Basic Service Set (BSS) and it can remove

contention between associated stations for uplink transmissions, therefore reducing the randomness caused by CSMA-based access between stations within the same BSS. The AP can also transmit simultaneously to multiple users in the downlink direction by using a Downlink (DL) MU OFDMA PPDU. In order to initiate a contention free Transmission Opportunity (TXOP) using the OFDMA mode, the AP still follows the typical listen before talk procedure to acquire the medium, which ensures interoperability and compliance with unlicensed band access rules. However, 802.11ax also includes a multi-user Enhanced Distributed Channel Access (MU-EDCA) capability, which allows the AP to get higher channel access priority than other devices in its BSS.

4.2.1.2. Traffic Isolation via OFDMA Resource Management and Resource Unit Allocation

802.11ax relies on the notion of OFDMA Resource Unit (RU) to allocate frequency chunks to different STAs over time. RUs provide a way to allow for multiple stations to transmit simultaneously, starting and ending at the same time. The way this is achieved is via padding, where extra bits are transmitted with the same power level. The current RU allocation algorithms provide a way to achieve traffic isolation per station which while per se does not support time-aware scheduling, is a key aspect to assist reliability, as it provides traffic isolation in a shared medium. IEEE 802.11be (see Section 4.3) is currently considering further and more flexible approaches concerning RU allocation.

4.2.1.3. Improved PHY Robustness

The 802.11ax PHY can operate with 0.8, 1.6 or 3.2 microsecond guard interval (GI). The larger GI options provide better protection against multipath, which is expected to be a challenge in industrial environments. The possibility to operate with smaller resource units (e.g. 2 MHz) enabled by OFDMA also helps reduce noise power and improve SNR, leading to better packet error rate (PER) performance.

802.11ax supports beamforming as in 802.11ac, but introduces UL MU MIMO, which helps improve reliability. The UL MU MIMO capability is also enabled by the trigger based access operation in 802.11ax.

4.2.1.4. Support for 6GHz band

The 802.11ax specification [IEEE Std 802.11ax] includes support for operation in the new 6 GHz band. Given the amount of new spectrum available as well as the fact that no legacy 802.11 device (prior 802.11ax) will be able to operate in this new band, 802.11ax operation in this new band can be even more efficient.

4.2.2. Applicability to deterministic flows

TSN capabilities, as defined by the IEEE 802.1 TSN standards, provide the underlying mechanism for supporting deterministic flows in a Local Area Network (LAN). The 802.11 working group has incorporated support for absolute time synchronization to extend the TSN 802.1AS protocol so that time-sensitive flow can experience precise time synchronization when operating over 802.11 links. As IEEE 802.11 and IEEE 802.1 TSN are both based on the IEEE 802 architecture, 802.11 devices can directly implement some TSN capabilities without the need for a gateway/translation protocol. Basic features required for operation in a 802.1Q LAN are already enabled for 802.11. Some TSN capabilities, such as 802.1Qbv, can already operate over the existing 802.11 MAC SAP [Sudhakaran2021]. Implementation and experimental results of TSN capabilities (802.1AS, 802.1Qbv, and 802.1CB) extended over standard Ethernet and Wi-Fi devices have also been described in [Fang_2021]. Nevertheless, the IEEE 802.11 MAC/PHY could be extended to improve the operation of IEEE 802.1 TSN features and achieve better performance metrics [Cavalcanti1287].

TSN capabilities supported over 802.11 (which also extends to 802.11ax), include:

1. 802.1AS based Time Synchronization (other time synchronization techniques may also be used)
2. Interoperating with IEEE802.1Q bridges
3. Time-sensitive Traffic Stream Classification

The existing 802.11 TSN capabilities listed above, and the 802.11ax OFDMA and AP-controlled access within a BSS provide a new set of tools to better serve time-sensitive flows. However, it is important to understand the tradeoffs and constraints associated with such capabilities, as well as redundancy and diversity mechanisms that can be used to provide more predictable and reliable performance.

4.2.2.1. 802.11 Managed network operation and admission control

Time-sensitive applications and TSN standards are expected to operate under a managed network (e.g. industrial/enterprise network). Thus, the Wi-Fi operation must also be carefully managed and integrated with the overall TSN management framework, as defined in the [IEEE802.1Qcc] specification.

Some of the random-access latency and interference from legacy/unmanaged devices can be reduced under a centralized management mode as defined in [IEEE802.1Qcc].

Existing traffic stream identification, configuration and admission control procedures defined in [IEEE Std 802.11] QoS mechanism can be re-used. However, given the high degree of determinism required by many time-sensitive applications, additional capabilities to manage interference and legacy devices within tight time-constraints need to be explored.

4.2.2.2. Scheduling for bounded latency and diversity

As discussed earlier, the [IEEE Std 802.11ax] OFDMA mode introduces the possibility of assigning different RUs (time/frequency resources) to users within a PPDU. Several RU sizes are defined in the specification (26, 52, 106, 242, 484, 996 subcarriers). In addition, the AP can also decide on MCS and grouping of users within a given OFDMA PPDU. Such flexibility can be leveraged to support time-sensitive applications with bounded latency, especially in a managed network where stations can be configured to operate under the control of the AP, in a controlled environment (which contains only devices operating on the unlicensed band installed by the facility owner and where unexpected interference from other systems and/or radio access technologies only sporadically happens), or in a deployment where channel/link redundancy is used to reduce the impact of unmanaged devices/interference.

When the network is lightly loaded, it is possible to achieve latencies under 1 msec when Wi-Fi is operated in contention-based (i.e., without OFDMA) mode. It has also been shown that it is possible to achieve 1 msec latencies in a controlled environment with higher efficiency when multi-user transmissions are used (enabled by OFDMA operation) [Cavalcanti_2019]. Obviously, there are latency, reliability and capacity tradeoffs to be considered. For instance, smaller RUs result in longer transmission durations, which may impact the minimal latency that can be achieved, but the contention latency and randomness elimination in an interference-free environment due to multi-user transmission is a major benefit of the OFDMA mode.

The flexibility to dynamically assign RUs to each transmission also enables the AP to provide frequency diversity, which can help increase reliability.

4.3. 802.11be Extreme High Throughput (EHT)

4.3.1. General Characteristics

The ongoing [IEEE 802.11be WIP] project is the next major 802.11 amendment (after IEEE Std 802.11ax-2021) for operation in the 2.4, 5 and 6 GHz bands. 802.11be is expected to include new PHY and MAC features and it is targeting extremely high throughput (at least 30 Gbps), as well as enhancements to worst case latency and jitter. It is also expected to improve the integration with 802.1 TSN to support time-sensitive applications over Ethernet and Wireless LANs.

The 802.11be Task Group started its operation in May 2019, therefore, detailed information about specific features is not yet available. Only high level candidate features have been discussed so far, including:

1. 320MHz bandwidth and more efficient utilization of non-contiguous spectrum.
2. Multi-link operation.
3. 16 spatial streams and related MIMO enhancements.
4. Multi-Access Point (AP) Coordination.
5. Enhanced link adaptation and retransmission protocol, e.g. Hybrid Automatic Repeat Request (HARQ).
6. Any required adaptations to regulatory rules for the 6 GHz spectrum.

4.3.2. Applicability to deterministic flows

The 802.11 Real-Time Applications (RTA) Topic Interest Group (TIG) provided detailed information on use cases, issues and potential solution directions to improve support for time-sensitive applications in 802.11. The RTA TIG report [IEEE_doc_11-18-2009-06] was used as input to the 802.11be project scope.

Improvements for worst-case latency, jitter and reliability were the main topics identified in the RTA report, which were motivated by applications in gaming, industrial automation, robotics, etc. The RTA report also highlighted the need to support additional TSN capabilities, such as time-aware (802.1Qbv) shaping and packet replication and elimination as defined in 802.1CB.

802.11be is expected to build on and enhance 802.11ax capabilities to improve worst case latency and jitter. Some of the enhancement areas are discussed next.

4.3.2.1. Enhanced scheduled operation for bounded latency

In addition to the throughput enhancements, 802.11be will leverage the trigger-based scheduled operation enabled by 802.11ax to provide efficient and more predictable medium access. 802.11be is expected to include enhancements to reduce overhead and enable more efficient operation in managed network deployments [IEEE_doc_11-19-0373-00].

4.3.2.2. Multi-AP coordination

Multi-AP coordination is one of the main new candidate features in 802.11be. It can provide benefits in throughput and capacity and has the potential to address some of the issues that impact worst case latency and reliability. Multi-AP coordination is expected to address the contention due to overlapping Basic Service Sets (OBSS), which is one of the main sources of random latency variations. 802.11be can define methods to enable better coordination between APs, for instance, in a managed network scenario, in order to reduce latency due to unmanaged contention.

Overall, multi-AP coordination algorithms consider three different phases: setup (where APs handling overlapping BSSs are assigned roles in a manual or automated way, e.g., coordinator and coordinated APs); coordination (where APs establish links among themselves, e.g., from a coordinating AP to coordinated APs; and then assign resources to served stations); transmission (where the coordinating APs optimize the distribution of the transmission opportunities).

Several multi-AP coordination approaches have been discussed with different levels of complexities and benefits, but specific coordination methods have not yet been defined. Out of the different categories, MAC-driven examples include: coordinated OFDMA (Co-OFDMA); Coordinated TDMA (Co-TDMA); HARQ; whereas PHY-driven examples include: Coordinated Spatial Reuse (Co-SR) and Coordinated Beamforming (Co-BF).

4.3.2.3. Multi-link operation

802.11be will introduce new features to improve operation over multiple links and channels. By leveraging multiple links/channels, 802.11be can isolate time-sensitive traffic from network congestion, one of the main causes of large latency variations. In a managed 802.11be network, it should be possible to steer traffic to certain links/channels to isolate time-sensitive traffic from other traffic

and help achieve bounded latency. The multi-link operation (MLO) has been already introduced in the 802.be Draft and it can also enhance latency and reliability by enabling data frames to be duplicated across links.

4.4. 802.11ad and 802.11ay (mmWave operation)

4.4.1. General Characteristics

The IEEE 802.11ad amendment defines PHY and MAC capabilities to enable multi-Gbps throughput in the 60 GHz millimeter wave (mmWave) band. The standard addresses the adverse mmWave signal propagation characteristics and provides directional communication capabilities that take advantage of beamforming to cope with increased attenuation. An overview of the 802.11ad standard can be found in [Nitsche_2015] .

The IEEE 802.11ay is currently developing enhancements to the 802.11ad standard to enable the next generation mmWave operation targeting 100 Gbps throughput. Some of the main enhancements in 802.11ay include MIMO, channel bonding, improved channel access and beamforming training. An overview of the 802.11ay capabilities can be found in [Ghasempour_2017]

4.4.2. Applicability to deterministic flows

The high data rates achievable with 802.11ad and 802.11ay can significantly reduce latency down to microsecond levels. Limited interference from legacy and other unlicensed devices in 60 GHz is also a benefit. However, directionality and short range typical in mmWave operation impose new challenges such as the overhead required for beam training and blockage issues, which impact both latency and reliability. Therefore, it is important to understand the use case and deployment conditions in order to properly apply and configure 802.11ad/ay networks for time sensitive applications.

The 802.11ad standard includes a scheduled access mode in which the central controller, after contending and reserving the channel for a dedicated period, can allocate to stations contention-free service periods. This scheduling capability is also available in 802.11ay, and it is one of the mechanisms that can be used to provide bounded latency to time-sensitive data flows in interference-free scenarios. An analysis of the theoretical latency bounds that can be achieved with 802.11ad service periods is provided in [Cavalcanti_2019].

5. IEEE 802.15.4

5.1. Provenance and Documents

The IEEE802.15.4 Task Group has been driving the development of low-power low-cost radio technology. The IEEE802.15.4 physical layer has been designed to support demanding low-power scenarios targeting the use of unlicensed bands, both the 2.4 GHz and sub GHz Industrial, Scientific and Medical (ISM) bands. This has imposed requirements in terms of frame size, data rate and bandwidth to achieve reduced collision probability, reduced packet error rate, and acceptable range with limited transmission power. The PHY layer supports frames of up to 127 bytes. The Medium Access Control (MAC) sublayer overhead is in the order of 10-20 bytes, leaving about 100 bytes to the upper layers. IEEE802.15.4 uses spread spectrum modulation such as the Direct Sequence Spread Spectrum (DSSS).

The Timeslotted Channel Hopping (TSCH) mode was added to the 2015 revision of the IEEE802.15.4 standard [IEEE Std 802.15.4]. TSCH is targeted at the embedded and industrial world, where reliability, energy consumption and cost drive the application space.

At the IETF, the 6TiSCH Working Group (WG) [TiSCH] deals with best effort operation of IPv6 [RFC8200] over TSCH. 6TiSCH has enabled distributed scheduling to exploit the deterministic access capabilities provided by TSCH. The group designed the essential mechanisms to enable the management plane operation while ensuring IPv6 is supported. Yet the charter did not focus to providing a solution to establish end to end Tracks while meeting quality of service requirements. 6TiSCH, through the RFC8480 [RFC8480] defines the 6P protocol which provides a pairwise negotiation mechanism to the control plane operation. The protocol supports agreement on a schedule between neighbors, enabling distributed scheduling. 6P goes hand-in-hand with a Scheduling Function (SF), the policy that decides how to maintain cells and trigger 6P transactions. The Minimal Scheduling Function (MSF) [RFC9033] is the default SF defined by the 6TiSCH WG; other standardized SFs can be defined in the future. MSF extends the minimal schedule configuration, and is used to add child-parent links according to the traffic load.

Time sensitive networking on low power constrained wireless networks have been partially addressed by ISA100.11a [ISA100.11a] and WirelessHART [WirelessHART]. Both technologies involve a central controller that computes redundant paths for industrial process control traffic over a TSCH mesh. Moreover, ISA100.11a introduces IPv6 capabilities with a Link-Local Address for the join process and a global unicast address for later exchanges, but the IPv6 traffic typically ends at a local application gateway and the full power of

IPv6 for end-to-end communication is not enabled. Compared to that state of the art, work at the IETF and in particular at RAW could provide additional techniques such as optimized P2P routing, PAREO functions, and end-to-end secured IPv6/CoAP connectivity.

The 6TiSCH architecture [RFC9030] identifies different models to schedule resources along so-called Tracks (see Section 5.2.2.2) exploiting the TSCH schedule structure however the focus at 6TiSCH is on best effort traffic and the group was never chartered to produce standard work related to Tracks.

Useful References include:

1. IEEE Std 802.15.4: "IEEE Std 802.15.4, Part. 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks" [IEEE Std 802.15.4]. The latest version at the time of this writing is dated year 2015.
2. Morell, A. , Vilajosana, X. , Vicario, J. L. and Watteyne, T. (2013), Label switching over IEEE802.15.4e networks. Trans. Emerging Tel. Tech., 24: 458-475. doi:10.1002/ett.2650 [morell13].
3. De Armas, J., Tuset, P., Chang, T., Adelantado, F., Watteyne, T., Vilajosana, X. (2016, September). Determinism through path diversity: Why packet replication makes sense. In 2016 International Conference on Intelligent Networking and Collaborative Systems (INCoS) (pp. 150-154). IEEE. [dearmas16].
4. X. Vilajosana, T. Watteyne, M. Vucinic, T. Chang and K. S. J. Pister, "6TiSCH: Industrial Performance for IPv6 Internet-of-Things Networks," in Proceedings of the IEEE, vol. 107, no. 6, pp. 1153-1165, June 2019. [vilajosana19].

5.2. TimeSlotted Channel Hopping

5.2.1. General Characteristics

As a core technique in IEEE802.15.4, TSCH splits time in multiple time slots that repeat over time. A set of timeslots constructs a Slotframe (see Section 5.2.2.1.4). For each timeslot, a set of available frequencies can be used, resulting in a matrix-like schedule (see Figure 1).




Figure 1: Slotframe example with scheduled cells between nodes A, B and C

This schedule represents the possible communications of a node with its neighbors, and is managed by a Scheduling Function such as the Minimal Scheduling Function (MSF) [RFC9033]. Each cell in the schedule is identified by its slotoffset and channeloffset coordinates. A cell's timeslot offset indicates its position in time, relative to the beginning of the slotframe. A cell's channel offset is an index which maps to a frequency at each iteration of the slotframe. Each packet exchanged between neighbors happens within one cell. The size of a cell is a timeslot duration, between 10 to 15 milliseconds. An Absolute Slot Number (ASN) indicates the number of slots elapsed since the network started. It increments at every slot. This is a 5 byte counter that can support networks running for more than 300 years without wrapping (assuming a 10 ms timeslot). Channel hopping provides increased reliability to multi-path fading and external interference. It is handled by TSCH through a channel hopping sequence referred as macHopSeq in the IEEE802.15.4 specification.

The Time-Frequency Division Multiple Access provided by TSCH enables the orchestration of traffic flows, spreading them in time and frequency, and hence enabling an efficient management of the bandwidth utilization. Such efficient bandwidth utilization can be combined to OFDM modulations also supported by the IEEE802.15.4 standard [IEEE Std 802.15.4] since the 2015 version.

TSCH networks operate in ISM bands in which the spectrum is shared by different coexisting technologies. Regulations such as FCC, ETSI and ARIB impose duty cycle regulations to limit the use of the bands but yet interference may constraint the probability to deliver a packet. Part of these reliability challenges are addressed at the MAC introducing redundancy and diversity, thanks to channel hopping, scheduling and ARQ policies. Yet, the MAC layer operates with a 1-hop vision, being limited to local actions to mitigate underperforming links.

In the RAW context, low power reliable networks should address non-critical control scenarios such as Class 2 and monitoring scenarios such as Class 4 defined by the RFC5673 [RFC5673]. As a low power technology targeting industrial scenarios radio transducers provide low data rates (typically between 50kbps to 250kbps) and robust modulations to trade-off performance to reliability. TSCH networks are organized in mesh topologies and connected to a backbone. Latency in the mesh network is mainly influenced by propagation aspects such as interference. ARQ methods and redundancy techniques such as replication and elimination should be studied to provide the needed performance to address deterministic scenarios.

5.2.2. Applicability to Deterministic Flows

Nodes in a TSCH network are tightly synchronized. This enables to build the slotted structure and ensure efficient utilization of resources thanks to proper scheduling policies. Scheduling is a key to orchestrate the resources that different nodes in a Track or a path are using. Slotframes can be split in resource blocks reserving the needed capacity to certain flows. Periodic and bursty traffic can be handled independently in the schedule, using active and reactive policies and taking advantage of overprovisionned cells to measure retransmission. Along a Track, resource blocks can be chained so nodes in previous hops transmit their data before the next packet comes. This provides a tight control to latency along a Track. Collision loss is avoided for best effort traffic by overprovisionning resources, giving time to the management plane of the network to dedicate more resources if needed.

5.2.2.1. Centralized Path Computation

In a controlled environment, a 6TiSCH device usually does not place a request for bandwidth between itself and another device in the network. Rather, an Operation Control System (OCS) invoked through an Human/Machine Interface (HMI) provides the Traffic Specification, in particular in terms of latency and reliability, and the end nodes, to a Path Computation element (PCE). With this, the PCE computes a Track between the end nodes and provisions every hop in the Track

with per-flow state that describes the per-hop operation for a given packet, the corresponding timeSlots, and the flow identification to recognize which packet is placed in which Track, sort out duplicates, etc. In Figure 2, an example of Operational Control System and HMI is depicted.

For a static configuration that serves a certain purpose for a long period of time, it is expected that a node will be provisioned in one shot with a full schedule, which incorporates the aggregation of its behavior for multiple Tracks. The 6TiSCH Architecture expects that the programming of the schedule is done over CoAP as discussed in "6TiSCH Resource Management and Interaction using CoAP" [I-D.ietf-6tisch-coap].

But an Hybrid mode may be required as well whereby a single Track is added, modified, or removed, for instance if it appears that a Track does not perform as expected for, say, Packet Delivery Ratio (PDR). For that case, the expectation is that a protocol that flows along a Track (to be), in a fashion similar to classical Traffic Engineering (TE) [CCAMP], may be used to update the state in the devices. 6TiSCH provides means for a device to negotiate a timeSlot with a neighbor, but in general that flow was not designed and no protocol was selected and it is expected that DetNet will determine the appropriate end-to-end protocols to be used in that case.

Stream Management Entity

Operational Control System and HMI

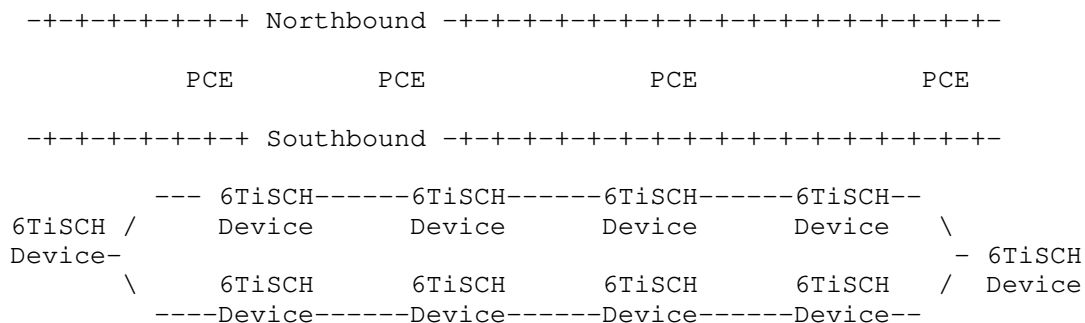


Figure 2

5.2.2.1.1. Packet Marking and Handling

Section "Packet Marking and Handling" of [RFC9030] describes the packet tagging and marking that is expected in 6TiSCH networks.

5.2.2.1.1.1. Tagging Packets for Flow Identification

For packets that are routed by a PCE along a Track, the tuple formed by the IPv6 source address and a local RPLInstanceID is tagged in the packets to identify uniquely the Track and associated transmit bundle of timeSlots.

It results that the tagging that is used for a DetNet flow outside the 6TiSCH LLN MUST be swapped into 6TiSCH formats and back as the packet enters and then leaves the 6TiSCH network.

Note: The method and format used for encoding the RPLInstanceID at 6lo is generalized to all 6TiSCH topological Instances, which includes Tracks.

5.2.2.1.1.2. Replication, Retries and Elimination

The 6TiSCH Architecture [RFC9030] leverages the Packet Replication, Retries and Elimination (PRE) functions (PREF), the precursor to what the RAW Architecture [I-D.ietf-raw-architecture] calls PAREO functions. PREF establishes several paths in a network to provide redundancy and parallel transmissions to bound the end-to-end delay. Considering the scenario shown in Figure 3, many different paths are possible for S to reach R. A simple way to benefit from this topology could be to use the two independent paths via nodes A, C, E and via B, D, F. But more complex paths are possible as well.




Figure 3: A Typical Ladder Shape with Two Parallel Paths Toward the Destination

By employing a Packet Replication function, each node forwards a copy of each data packet over two different branches. For instance, in Figure 4, the source node S transmits the data packet to nodes A and B, in two different timeslots within the same TSCH slotframe.




Figure 4: Packet Replication: S transmits twice the same data packet, to its DP (A) and to its AP (B).

By employing Packet Elimination function once a node receives the first copy of a data packet, it discards the subsequent copies. Because the first copy that reaches a node is the one that matters, it is the only copy that will be forwarded upward.

Considering that the wireless medium is broadcast by nature, any neighbor of a transmitter may overhear a transmission. By employing the Promiscuous Overhearing function, nodes will have multiple opportunities to receive a given data packet. For instance, in Figure 4, when the source node S transmits the data packet to node A, node B may overhear this transmission.

6TiSCH expects elimination and replication of packets along a complex Track, but has no position about how the sequence numbers would be tagged in the packet.

As it goes, 6TiSCH expects that timeSlots corresponding to copies of a same packet along a Track are correlated by configuration, and does not need to process the sequence numbers.

The semantics of the configuration MUST enable correlated timeSlots to be grouped for transmit (and respectively receive) with a 'OR' relations, and then a 'AND' relation MUST be configurable between groups. The semantics is that if the transmit (and respectively receive) operation succeeded in one timeSlot in a 'OR' group, then all the other timeSlots in the group are ignored. Now, if there are at least two groups, the 'AND' relation between the groups indicates that one operation must succeed in each of the groups.

On the transmit side, timeSlots provisioned for retries along a same branch of a Track are placed a same 'OR' group. The 'OR' relation indicates that if a transmission is acknowledged, then further transmissions SHOULD NOT be attempted for timeSlots in that group. There are as many 'OR' groups as there are branches of the Track departing from this node. Different 'OR' groups are programmed for the purpose of replication, each group corresponding to one branch of the Track. The 'AND' relation between the groups indicates that transmission over any of branches MUST be attempted regardless of whether a transmission succeeded in another branch. It is also possible to place cells to different next-hop routers in a same 'OR' group. This allows to route along multi-path Tracks, trying one next-hop and then another only if sending to the first fails.

On the receive side, all timeSlots are programmed in a same 'OR' group. Retries of a same copy as well as converging branches for elimination are converged, meaning that the first successful reception is enough and that all the other timeSlots can be ignored.

5.2.2.1.2. Topology and capabilities

6TiSCH nodes are usually IoT devices, characterized by very limited amount of memory, just enough buffers to store one or a few IPv6 packets, and limited bandwidth between peers. It results that a node will maintain only a small number of peering information, and will not be able to store many packets waiting to be forwarded. Peers can be identified through MAC or IPv6 addresses.

Neighbors can be discovered over the radio using mechanism such as Enhanced Beacons, but, though the neighbor information is available in the 6TiSCH interface data model, 6TiSCH does not describe a protocol to pro-actively push the neighborhood information to a PCE. This protocol should be described and should operate over CoAP. The protocol should be able to carry multiple metrics, in particular the same metrics as used for RPL operations [RFC6551].

The energy that the device consumes in sleep, transmit and receive modes can be evaluated and reported. So can the amount of energy that is stored in the device and the power that it can be scavenged from the environment. The PCE SHOULD be able to compute Tracks that will implement policies on how the energy is consumed, for instance balance between nodes, ensure that the spent energy does not exceeded the scavenged energy over a period of time, etc...

5.2.2.1.3. Schedule Management by a PCE

6TiSCH supports a mixed model of centralized routes and distributed routes. Centralized routes can for example be computed by a entity such as a PCE [PCE]. Distributed routes are computed by RPL [RFC6550].

Both methods may inject routes in the Routing Tables of the 6TiSCH routers. In either case, each route is associated with a 6TiSCH topology that can be a RPL Instance topology or a Track. The 6TiSCH topology is indexed by a Instance ID, in a format that reuses the RPLInstanceID as defined in RPL.

Both RPL and PCE rely on shared sources such as policies to define Global and Local RPLInstanceIDs that can be used by either method. It is possible for centralized and distributed routing to share a same topology. Generally they will operate in different slotFrames, and centralized routes will be used for scheduled traffic and will have precedence over distributed routes in case of conflict between the slotFrames.

5.2.2.1.4. SlotFrames and Priorities

A slotFrame is the base object that a PCE needs to manipulate to program a schedule into an LLN node. Elaboration on that concept can be found in section "SlotFrames and Priorities" of [RFC9030]

IEEE802.15.4 TSCH avoids contention on the medium by formatting time and frequencies in cells of transmission of equal duration. In order to describe that formatting of time and frequencies, the 6TiSCH architecture defines a global concept that is called a Channel Distribution and Usage (CDU) matrix; a CDU matrix is a matrix of cells with an height equal to the number of available channels (indexed by ChannelOffsets) and a width (in timeSlots) that is the period of the network scheduling operation (indexed by slotOffsets) for that CDU matrix. The size of a cell is a timeSlot duration, and values of 10 to 15 milliseconds are typical in 802.15.4 TSCH to accommodate for the transmission of a frame and an acknowledgement, including the security validation on the receive side which may take up to a few milliseconds on some device architecture.

The frequency used by a cell in the matrix rotates in a pseudo-random fashion, from an initial position at an epoch time, as the matrix iterates over and over.

A CDU matrix is computed by the PCE, but unallocated timeSlots may be used opportunistically by the nodes for classical best effort IP traffic. The PCE has precedence in the allocation in case of a conflict.

In a given network, there might be multiple CDU matrices that operate with different width, so they have different durations and represent different periodic operations. It is recommended that all CDU matrices in a 6TiSCH domain operate with the same cell duration and are aligned, so as to reduce the chances of interferences from slotted-aloha operations. The PCE MUST compute the CDU matrices and shared that knowledge with all the nodes. The matrices are used in particular to define slotFrames.

A slotFrame is a MAC-level abstraction that is common to all nodes and contains a series of timeSlots of equal length and precedence. It is characterized by a slotFrame_ID, and a slotFrame_size. A slotFrame aligns to a CDU matrix for its parameters, such as number and duration of timeSlots.

Multiple slotFrames can coexist in a node schedule, i.e., a node can have multiple activities scheduled in different slotFrames, based on the precedence of the 6TiSCH topologies. The slotFrames may be aligned to different CDU matrices and thus have different width. There is typically one slotFrame for scheduled traffic that has the highest precedence and one or more slotFrame(s) for RPL traffic. The timeSlots in the slotFrame are indexed by the SlotOffset; the first cell is at SlotOffset 0.

The 6TiSCH architecture introduces the concept of chunks ([RFC9030]) to operate such spectrum distribution for a whole group of cells at a time. The CDU matrix is formatted into a set of chunks, each of them identified uniquely by a chunk-ID, see Figure 5. The PCE MUST compute the partitioning of CDU matrices into chunks and shared that knowledge with all the nodes in a 6TiSCH network.




Figure 5: CDU matrix Partitioning in Chunks

The appropriation of a chunk can be requested explicitly by the PCE to any node. After a successful appropriation, the PCE owns the cells in that chunk, and may use them as hard cells to set up Tracks. Then again, 6TiSCH did not propose a method for chunk definition and a protocol for appropriation. This is to be done at RAW.

5.2.2.2. 6TiSCH Tracks

A Track at 6TiSCH is the application to wireless of the concept of a path in the "Detnet architecture" [RFC8655]. A Track can follow a simple sequence of relay nodes or can be structured as a more complex Destination Oriented Directed Acyclic Graph (DODAG) to a unicast destination. Along a Track, 6TiSCH nodes reserve the resources to enable the efficient transmission of packets while aiming to optimize certain properties such as reliability and ensure small jitter or bounded latency. The Track structure enables Layer-2 forwarding schemes, reducing the overhead of taking routing decisions at the Layer-3.

Serial Tracks can be understood as the concatenation of cells or bundles along a routing path from a source towards a destination. The serial Track concept is analogous to the circuit concept where resources are chained through the multi-hop topology. For example, A bundle of Tx Cells in a particular node is paired to a bundle of Rx Cells in the next hop node following a routing path.

Whereas scheduling ensures reliable delivery in bounded time along any Track, high availability requires the application of PAREO functions along a more complex DODAG Track structure. A DODAG has forking and joining nodes where the concepts such as Replication and Elimination can be exploited. Spatial redundancy increases the overall energy consumption in the network but improves significantly the availability of the network as well as the packet delivery ratio. A Track may also branch off and rejoin, for the purpose of the so-called Packet Replication and Elimination (PRE), over non congruent branches. PRE may be used to complement layer-2 Automatic Repeat reQuest (ARQ) and receiver-end Ordering to form the PAREO functions. PAREO functions enable to meet industrial expectations in PDR within bounded delivery time over a Track that includes wireless links, even when the Track extends beyond the 6TiSCH network.

The RAW Track described in the RAW Architecture [I-D.ietf-raw-architecture] inherits directly from that model. RAW extends the graph beyond a DODAG as long as a given packet cannot loop within the Track.




Figure 6: End-to-End deterministic Track

In the example above (see Figure 6), a Track is laid out from a field device in a 6TiSCH network to an IoT gateway that is located on a IEEE802.1 TSN backbone.

The Replication function in the field device sends a copy of each packet over two different branches, and a PCE schedules each hop of both branches so that the two copies arrive in due time at the gateway. In case of a loss on one branch, hopefully the other copy of the packet still makes it in due time. If two copies make it to the IoT gateway, the Elimination function in the gateway ignores the extra packet and presents only one copy to upper layers.

At each 6TiSCH hop along the Track, the PCE may schedule more than one timeSlot for a packet, so as to support Layer-2 retries (ARQ). It is also possible that the field device only uses the second branch if sending over the first branch fails.

In current deployments, a TSN Track does not necessarily support PRE but is systematically multi-path. This means that a Track is scheduled so as to ensure that each hop has at least two forwarding solutions, and the forwarding decision is to try the preferred one and use the other in case of Layer-2 transmission failure as detected by ARQ.

Methods to implement complex Tracks are described in [I-D.papadopoulos-paw-pre-reqs] and complemented by extensions to the RPL routing protocol in [I-D.ietf-roll-nsa-extension] for best effort traffic, but a centralized routing technique such as promoted in DetNet is still missing.

5.2.2.2.1. Track Scheduling Protocol

Section "Schedule Management Mechanisms" of the 6TiSCH architecture describes 4 paradigms to manage the TSCH schedule of the LLN nodes: Static Scheduling, neighbor-to-neighbor Scheduling, remote monitoring and scheduling management, and Hop-by-hop scheduling. The Track operation for DetNet corresponds to a remote monitoring and scheduling management by a PCE.

Early work at 6TiSCH on a data model and a protocol to program the schedule in the 6TiSCH device was never concluded as the group focussed on best effort traffic. This work would be revived by RAW:

The 6top interface document [RFC8480] (to be reopened at RAW) was intended to specify the generic data model that can be used to monitor and manage resources of the 6top sublayer. Abstract methods were suggested for use by a management entity in the device. The data model also enables remote control operations on the 6top sublayer.

[I-D.ietf-6tisch-coap] (to be reopened at RAW) was intended to define a mapping of the 6top set of commands, which is described in RFC 8480, to CoAP resources. This allows an entity to interact with the 6top layer of a node that is multiple hops away in a RESTful fashion.

[I-D.ietf-6tisch-coap] also defined a basic set CoAP resources and associated RESTful access methods (GET/PUT/POST/DELETE). The payload (body) of the CoAP messages is encoded using the CBOR format. The PCE commands are expected to be issued directly as CoAP requests or to be mapped back and forth into CoAP by a gateway function at the edge of the 6TiSCH network. For instance, it is possible that a mapping entity on the backbone transforms a non-CoAP protocol such as PCEP into the RESTful interfaces that the 6TiSCH devices support.

5.2.2.2.2. Track Forwarding

By forwarding, this specification means the per-packet operation that allows to deliver a packet to a next hop or an upper layer in this node. Forwarding is based on pre-existing state that was installed as a result of the routing computation of a Track by a PCE. The 6TiSCH architecture supports three different forwarding model, G-MPLS Track Forwarding (TF), 6LoWPAN Fragment Forwarding (FF) and IPv6 Forwarding (6F) which is the classical IP operation [RFC9030]. The DetNet case relates to the Track Forwarding operation under the control of a PCE.

A Track is a unidirectional path between a source and a destination. In a Track cell, the normal operation of IEEE802.15.4 Automatic Repeat-reQuest (ARQ) usually happens, though the acknowledgment may be omitted in some cases, for instance if there is no scheduled cell for a retry.

Track Forwarding is the simplest and fastest. A bundle of cells set to receive (RX-cells) is uniquely paired to a bundle of cells that are set to transmit (TX-cells), representing a layer-2 forwarding state that can be used regardless of the network layer protocol. This model can effectively be seen as a Generalized Multi-protocol Label Switching (G-MPLS) operation in that the information used to switch a frame is not an explicit label, but rather related to other properties of the way the packet was received, a particular cell in the case of 6TiSCH. As a result, as long as the TSCH MAC (and Layer-2 security) accepts a frame, that frame can be switched regardless of the protocol, whether this is an IPv6 packet, a 6LoWPAN fragment, or a frame from an alternate protocol such as WirelessHART or ISA100.11a.

A data frame that is forwarded along a Track normally has a destination MAC address that is set to broadcast - or a multicast address depending on MAC support. This way, the MAC layer in the intermediate nodes accepts the incoming frame and 6top switches it without incurring a change in the MAC header. In the case of IEEE802.15.4, this means effectively broadcast, so that along the Track the short address for the destination of the frame is set to 0xFFFF.

A Track is thus formed end-to-end as a succession of paired bundles, a receive bundle from the previous hop and a transmit bundle to the next hop along the Track, and a cell in such a bundle belongs to at most one Track. For a given iteration of the device schedule, the effective channel of the cell is obtained by adding a pseudo-random number to the channelOffset of the cell, which results in a rotation of the frequency that used for transmission. The bundles may be

computed so as to accommodate both variable rates and retransmissions, so they might not be fully used at a given iteration of the schedule. The 6TiSCH architecture provides additional means to avoid waste of cells as well as overflows in the transmit bundle, as follows:

In one hand, a TX-cell that is not needed for the current iteration may be reused opportunistically on a per-hop basis for routed packets. When all of the frame that were received for a given Track are effectively transmitted, any available TX-cell for that Track can be reused for upper layer traffic for which the next-hop router matches the next hop along the Track. In that case, the cell that is being used is effectively a TX-cell from the Track, but the short address for the destination is that of the next-hop router. It results that a frame that is received in a RX-cell of a Track with a destination MAC address set to this node as opposed to broadcast must be extracted from the Track and delivered to the upper layer (a frame with an unrecognized MAC address is dropped at the lower MAC layer and thus is not received at the 6top sublayer).

On the other hand, it might happen that there are not enough TX-cells in the transmit bundle to accommodate the Track traffic, for instance if more retransmissions are needed than provisioned. In that case, the frame can be placed for transmission in the bundle that is used for layer-3 traffic towards the next hop along the Track as long as it can be routed by the upper layer, that is, typically, if the frame transports an IPv6 packet. The MAC address should be set to the next-hop MAC address to avoid confusion. It results that a frame that is received over a layer-3 bundle may be in fact associated to a Track. In a classical IP link such as an Ethernet, off-Track traffic is typically in excess over reservation to be routed along the non-reserved path based on its QoS setting. But with 6TiSCH, since the use of the layer-3 bundle may be due to transmission failures, it makes sense for the receiver to recognize a frame that should be re-Track, and to place it back on the appropriate bundle if possible. A frame should be re-Track if the Per-Hop-Behavior group indicated in the Differentiated Services Field in the IPv6 header is set to Deterministic Forwarding, as discussed in Section 5.2.2.1.1. A frame is re-Track by scheduling it for transmission over the transmit bundle associated to the Track, with the destination MAC address set to broadcast.

There are 2 modes for a Track, transport mode and tunnel mode.

5.2.2.2.2.1. Transport Mode

In transport mode, the Protocol Data Unit (PDU) is associated with flow-dependant meta-data that refers uniquely to the Track, so the 6top sublayer can place the frame in the appropriate cell without ambiguity. In the case of IPv6 traffic, this flow identification is transported in the Flow Label of the IPv6 header. Associated with the source IPv6 address, the Flow Label forms a globally unique identifier for that particular Track that is validated at egress before restoring the destination MAC address (DMAC) and punting to the upper layer.




Figure 7: Track Forwarding, Transport Mode

5.2.2.2.2.2. Tunnel Mode

In tunnel mode, the frames originate from an arbitrary protocol over a compatible MAC that may or may not be synchronized with the 6TiSCH network. An example of this would be a router with a dual radio that is capable of receiving and sending WirelessHART or ISA100.11a frames with the second radio, by presenting itself as an Access Point or a Backbone Router, respectively.

In that mode, some entity (e.g. PCE) can coordinate with a WirelessHART Network Manager or an ISA100.11a System Manager to specify the flows that are to be transported transparently over the Track.

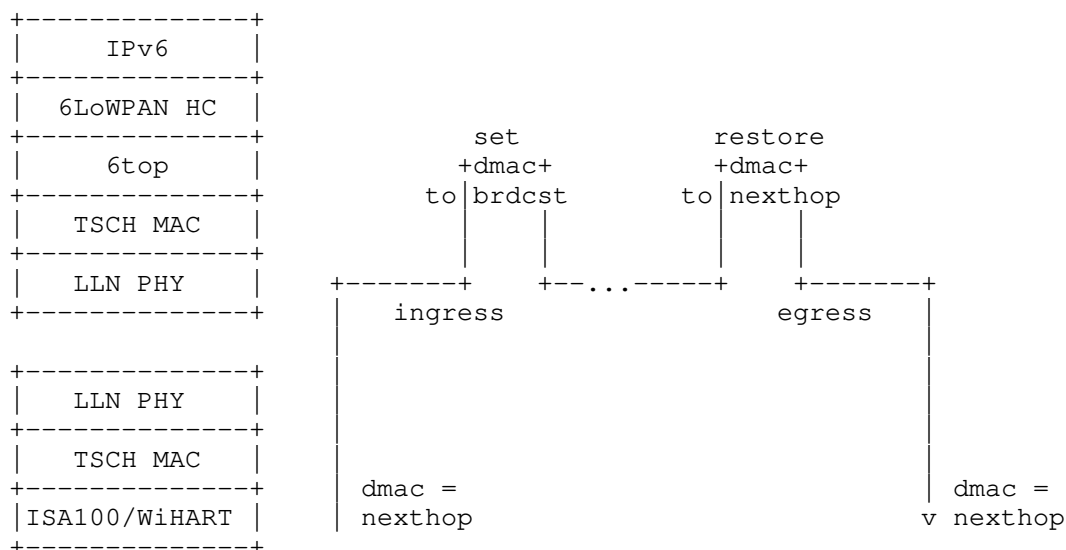


Figure 8: Track Forwarding, Tunnel Mode

In that case, the flow information that identifies the Track at the ingress 6TiSCH router is derived from the RX-cell. The dmac is set to this node but the flow information indicates that the frame must be tunneled over a particular Track so the frame is not passed to the upper layer. Instead, the dmac is forced to broadcast and the frame is passed to the 6top sublayer for switching.

At the egress 6TiSCH router, the reverse operation occurs. Based on metadata associated to the Track, the frame is passed to the appropriate link layer with the destination MAC restored.

5.2.2.2.2.3. Tunnel Metadata

Metadata coming with the Track configuration is expected to provide the destination MAC address of the egress endpoint as well as the tunnel mode and specific data depending on the mode, for instance a service access point for frame delivery at egress. If the tunnel egress point does not have a MAC address that matches the configuration, the Track installation fails.

In transport mode, if the final layer-3 destination is the tunnel termination, then it is possible that the IPv6 address of the destination is compressed at the 6LoWPAN sublayer based on the MAC address. It is thus mandatory at the ingress point to validate that the MAC address that was used at the 6LoWPAN sublayer for compression matches that of the tunnel egress point. For that reason, the node

that injects a packet on a Track checks that the destination is effectively that of the tunnel egress point before it overwrites it to broadcast. The 6top sublayer at the tunnel egress point reverts that operation to the MAC address obtained from the tunnel metadata.

5.2.2.2.2.4. OAM

An Overview of Operations, Administration, and Maintenance (OAM) Tools [RFC7276] provides an overview of the existing tooling for OAM [RFC6291]. Tracks are complex paths and new tooling is necessary to manage them, with respect to load control, timing, and the Packet Replication and Elimination Functions (PREF).

An example of such tooling can be found in the context of BIER [RFC8279] and more specifically BIER Traffic Engineering [I-D.ietf-bier-te-arch] (BIER-TE): [I-D.thubert-bier-replication-elimination] leverages BIER-TE to control the process of PREF, and to provide traceability of these operations, in the deterministic dataplane, along a complex Track. For the 6TiSCH type of constrained environment, [I-D.thubert-6lo-bier-dispatch] enables an efficient encoding of the BIER bitmap within the 6LoRH framework.

6. 5G

6.1. Provenance and Documents

The 3rd Generation Partnership Project (3GPP) incorporates many companies whose business is related to cellular network operation as well as network equipment and device manufacturing. All generations of 3GPP technologies provide scheduled wireless segments, primarily in licensed spectrum which is beneficial for reliability and availability.

In 2016, the 3GPP started to design New Radio (NR) technology belonging to the fifth generation (5G) of cellular networks. NR has been designed from the beginning to not only address enhanced Mobile Broadband (eMBB) services for consumer devices such as smart phones or tablets but is also tailored for future Internet of Things (IoT) communication and connected cyber-physical systems. In addition to eMBB, requirement categories have been defined on Massive Machine-Type Communication (M-MTC) for a large number of connected devices/sensors, and Ultra-Reliable Low-Latency Communication (URLLC) for connected control systems and critical communication as illustrated in Figure 9. It is the URLLC capabilities that make 5G a great candidate for reliable low-latency communication. With these three corner stones, NR is a complete solution supporting the connectivity needs of consumers, enterprises, and public sector for both wide area and local area, e.g. indoor deployments. A general overview of NR can be found in [TS38300].




Figure 9: 5G Application Areas

As a result of releasing the first NR specification in 2018 (Release 15), it has been proven by many companies that NR is a URLLC-capable technology and can deliver data packets at 10^{-5} packet error rate within 1ms latency budget [TR37910]. Those evaluations were consolidated and forwarded to ITU to be included in the [IMT2020] work.

In order to understand communication requirements for automation in vertical domains, 3GPP studied different use cases [TR22804] and released technical specification with reliability, availability and latency demands for a variety of applications [TS22104].

As an evolution of NR, multiple studies have been conducted in scope of 3GPP Release 16 including the following two, focusing on radio aspects:

1. Study on physical layer enhancements for NR ultra-reliable and low latency communication (URLLC) [TR38824].
2. Study on NR industrial Internet of Things (I-IoT) [TR38825].

Resulting of these studies, further enhancements to NR have been standardized in 3GPP Release 16, hence, available in [TS38300], and continued in 3GPP Release 17 standardization (according to [RP210854]).

In addition, several enhancements have been done on system architecture level which are reflected in System architecture for the 5G System (5GS) [TS23501]. These enhancements include multiple features in support of Time-Sensitive Communications (TSC) by Release 16 and Release 17. Further improvements will be considered in the 3GPP Release 18 time frame, e.g., support for DetNet [SP211633].

The adoption and the use of 5G is facilitated by multiple organizations. For instance, the 5G Alliance for Connected Industries and Automation (5G-ACIA) brings together widely varying 5G stakeholders including Information and Communication Technology (ICT) players and Operational Technology (OT) companies, e.g.: industrial automation enterprises, machine builders, and end users. Another example is the 5G Automotive Association (5GAA), which bridges ICT and automotive technology companies to develop end-to-end solutions for future mobility and transportation services.

6.2. General Characteristics

The 5G Radio Access Network (5G RAN) with its NR interface includes several features to achieve Quality of Service (QoS), such as a guaranteeably low latency or tolerable packet error rates for selected data flows. Determinism is achieved by centralized admission control and scheduling of the wireless frequency resources, which are typically licensed frequency bands assigned to a network operator.

NR enables short transmission slots in a radio subframe, which benefits low-latency applications. NR also introduces mini-slots, where prioritized transmissions can be started without waiting for slot boundaries, further reducing latency. As part of giving priority and faster radio access to URLLC traffic, NR introduces preemption where URLLC data transmission can preempt ongoing non-URLLC transmissions. Additionally, NR applies very fast processing, enabling retransmissions even within short latency bounds.

NR defines extra-robust transmission modes for increased reliability both for data and control radio channels. Reliability is further improved by various techniques, such as multi-antenna transmission, the use of multiple frequency carriers in parallel and packet duplication over independent radio links. NR also provides full mobility support, which is an important reliability aspect not only for devices that are moving, but also for devices located in a changing environment.

Network slicing is seen as one of the key features for 5G, allowing vertical industries to take advantage of 5G networks and services. Network slicing is about transforming a Public Land Mobile Network (PLMN) from a single network to a network where logical partitions are created, with appropriate network isolation, resources, optimized topology and specific configuration to serve various service requirements. An operator can configure and manage the mobile network to support various types of services enabled by 5G, for example eMBB and URLLC, depending on the different customers' needs.

Exposure of capabilities of 5G Systems to the network or applications outside the 3GPP domain have been added to Release 16 [TS23501]. Via exposure interfaces, applications can access 5G capabilities, e.g., communication service monitoring and network maintenance.

For several generations of mobile networks, 3GPP has considered how the communication system should work on a global scale with billions of users, taking into account resilience aspects, privacy regulation, protection of data, encryption, access and core network security, as well as interconnect. Security requirements evolve as demands on trustworthiness increase. For example, this has led to the introduction of enhanced privacy protection features in 5G. 5G also employs strong security algorithms, encryption of traffic, protection of signaling and protection of interfaces.

One particular strength of mobile networks is the authentication, based on well-proven algorithms and tightly coupled with a global identity management infrastructure. Since 3G, there is also mutual authentication, allowing the network to authenticate the device and the device to authenticate the network. Another strength is secure

solutions for storage and distribution of keys fulfilling regulatory requirements and allowing international roaming. When connecting to 5G, the user meets the entire communication system, where security is the result of standardization, product security, deployment, operations and management as well as incident handling capabilities. The mobile networks approach the entirety in a rather coordinated fashion which is beneficial for security.

6.3. Deployment and Spectrum

The 5G system allows deployment in a vast spectrum range, addressing use-cases in both wide-area as well as local networks. Furthermore, 5G can be configured for public and non-public access.

When it comes to spectrum, NR allows combining the merits of many frequency bands, such as the high bandwidths in millimeter Waves (mmW) for extreme capacity locally, as well as the broad coverage when using mid- and low frequency bands to address wide-area scenarios. URLLC is achievable in all these bands. Spectrum can be either licensed, which means that the license holder is the only authorized user of that spectrum range, or unlicensed, which means that anyone who wants to use the spectrum can do so.

A prerequisite for critical communication is performance predictability, which can be achieved by the full control of the access to the spectrum, which 5G provides. Licensed spectrum guarantees control over spectrum usage by the system, making it a preferable option for critical communication. However, unlicensed spectrum can provide an additional resource for scaling non-critical communications. While NR is initially developed for usage of licensed spectrum, the functionality to access also unlicensed spectrum was introduced in 3GPP Release 16. Moreover, URLLC features are enhanced in Release 17 [RP210854] to be better applicable to unlicensed spectrum.

Licensed spectrum dedicated to mobile communications has been allocated to mobile service providers, i.e. issued as longer-term licenses by national administrations around the world. These licenses have often been associated with coverage requirements and issued across whole countries, or in large regions. Besides this, configured as a non-public network (NPN) deployment, 5G can provide network services also to a non-operator defined organization and its premises such as a factory deployment. By this isolation, quality of service requirements, as well as security requirements can be achieved. An integration with a public network, if required, is also possible. The non-public (local) network can thus be interconnected with a public network, allowing devices to roam between the networks.

In an alternative model, some countries are now in the process of allocating parts of the 5G spectrum for local use to industries. These non-service providers then have a choice of applying for a local license themselves and operating their own network or cooperating with a public network operator or service provider.

6.4. Applicability to Deterministic Flows

6.4.1. System Architecture

The 5G system [TS23501] consists of the User Equipment (UE) at the terminal side, and the Radio Access Network (RAN) with the gNB as radio base station node, as well as the Core Network (CN). The core network is based on a service-based architecture with the central functions: Access and Mobility Management Function (AMF), Session Management Function (SMF) and User Plane Function (UPF) as illustrated in Figure 10.

The gNB's main responsibility is the radio resource management, including admission control and scheduling, mobility control and radio measurement handling. The AMF handles the UE's connection status and security, while the SMF controls the UE's data sessions. The UPF handles the user plane traffic.

The SMF can instantiate various Packet Data Unit (PDU) sessions for the UE, each associated with a set of QoS flows, i.e., with different QoS profiles. Segregation of those sessions is also possible, e.g., resource isolation in the RAN and in the CN can be defined (slicing).

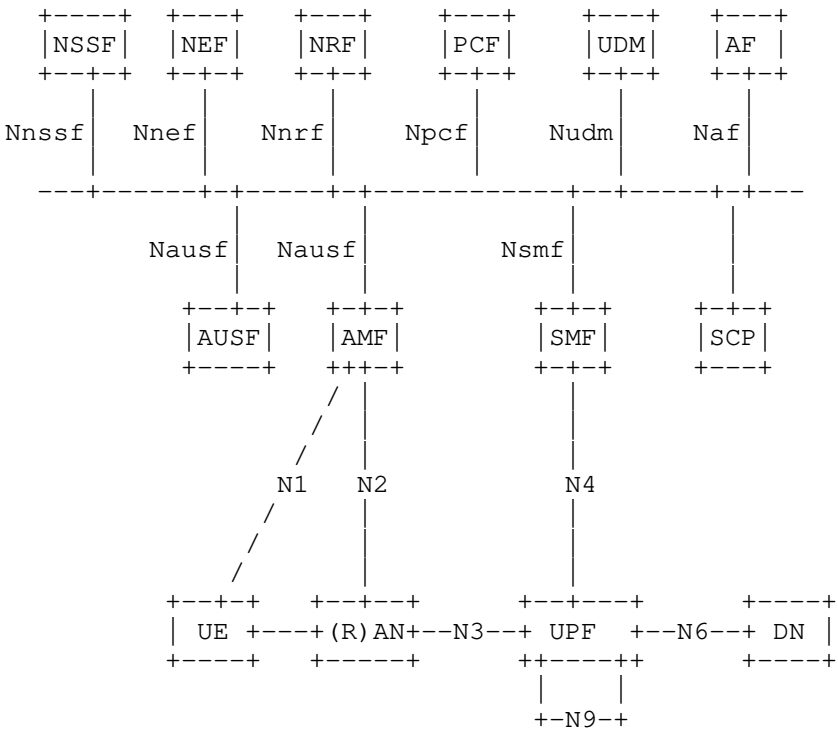


Figure 10: 5G System Architecture

To allow UE mobility across cells/gNBs, handover mechanisms are supported in NR. For an established connection, i.e., connected mode mobility, a gNB can configure a UE to report measurements of received signal strength and quality of its own and neighbouring cells, periodically or event-based. Based on these measurement reports, the gNB decides to handover a UE to another target cell/gNB. Before triggering the handover, it is hand-shaked with the target gNB based on network signalling. A handover command is then sent to the UE and the UE switches its connection to the target cell/gNB. The Packet Data Convergence Protocol (PDCP) of the UE can be configured to avoid data loss in this procedure, i.e., handle retransmissions if needed. Data forwarding is possible between source and target gNB as well. To improve the mobility performance further, i.e., to avoid connection failures, e.g., due to too-late handovers, the mechanism of conditional handover is introduced in Release 16 specifications. Therein a conditional handover command, defining a triggering point, can be sent to the UE before UE enters a handover situation. A further improvement that has been introduced in Release 16 is the Dual Active Protocol Stack (DAPS), where the UE maintains the connection to the source cell while connecting to the target cell. This way, potential interruptions in packet delivery can be avoided entirely.

6.4.2. Overview of The Radio Protocol Stack

The protocol architecture for NR consists of the L1 Physical layer (PHY) and as part of the L2, the sublayers of Medium Access Control (MAC), Radio Link Control (RLC), Packet Data Convergence Protocol (PDCP), as well as the Service Data Adaption Protocol (SDAP).

The PHY layer handles signal processing related actions, such as encoding/decoding of data and control bits, modulation, antenna precoding and mapping.

The MAC sub-layer handles multiplexing and priority handling of logical channels (associated with QoS flows) to transport blocks for PHY transmission, as well as scheduling information reporting and error correction through Hybrid Automated Repeat Request (HARQ).

The RLC sublayer handles sequence numbering of higher layer packets, retransmissions through Automated Repeat Request (ARQ), if configured, as well as segmentation and reassembly and duplicate detection.

The PDCP sublayer consists of functionalities for ciphering/deciphering, integrity protection/verification, re-ordering and in-order delivery, duplication and duplicate handling for higher layer packets, and acts as the anchor protocol to support handovers.

The SDAP sublayer provides services to map QoS flows, as established by the 5G core network, to data radio bearers (associated with logical channels), as used in the 5G RAN.

Additionally, in RAN, the Radio Resource Control (RRC) protocol, handles the access control and configuration signalling for the aforementioned protocol layers. RRC messages are considered L3 and thus transmitted also via those radio protocol layers.

To provide low latency and high reliability for one transmission link, i.e., to transport data (or control signaling) of one radio bearer via one carrier, several features have been introduced on the user plane protocols for PHY and L2, as explained in the following.

6.4.3. Radio (PHY)

NR is designed with native support of antenna arrays utilizing benefits from beamforming, transmissions over multiple MIMO layers and advanced receiver algorithms allowing effective interference cancellation. Those antenna techniques are the basis for high signal quality and effectiveness of spectral usage. Spatial diversity with up to 4 MIMO layers in UL and up to 8 MIMO layers in DL is supported. Together with spatial-domain multiplexing, antenna arrays can focus power in desired direction to form beams. NR supports beam management mechanisms to find the best suitable beam for UE initially and when it is moving. In addition, gNBs can coordinate their respective DL and UL transmissions over the backhaul network keeping interference reasonably low, and even make transmissions or receptions from multiple points (multi-TRP). Multi-TRP can be used for repetition of data packet in time, in frequency or over multiple MIMO layers which can improve reliability even further.

Any downlink transmission to a UE starts from resource allocation signaling over the Physical Downlink Control Channel (PDCCH). If it is successfully received, the UE will know about the scheduled transmission and may receive data over the Physical Downlink Shared Channel (PDSCH). If retransmission is required according to the HARQ scheme, a signaling of negative acknowledgement (NACK) on the Physical Uplink Control Channel (PUCCH) is involved and PDCCH together with PDSCH transmissions (possibly with additional redundancy bits) are transmitted and soft-combined with previously received bits. Otherwise, if no valid control signaling for scheduling data is received, nothing is transmitted on PUCCH (discontinuous transmission - DTX), and the base station upon detecting DTX will retransmit the initial data.

An uplink transmission normally starts from a Scheduling Request (SR) – a signaling message from the UE to the base station sent via PUCCH. Once the scheduler is informed about buffer data in UE, e.g., by SR, the UE transmits a data packet on the Physical Uplink Shared Channel (PUSCH). Pre-scheduling not relying on SR is also possible (see following section).

Since transmission of data packets require usage of control and data channels, there are several methods to maintain the needed reliability. NR uses Low Density Parity Check (LDPC) codes for data channels, Polar codes for PDCCH, as well as orthogonal sequences and Polar codes for PUCCH. For ultra-reliability of data channels, very robust (low spectral efficiency) Modulation and Coding Scheme (MCS) tables are introduced containing very low (down to 1/20) LDPC code rates using BPSK or QPSK. Also, PDCCH and PUCCH channels support multiple code rates including very low ones for the channel robustness.

A connected UE reports downlink (DL) quality to gNB by sending Channel State Information (CSI) reports via PUCCH while uplink (UL) quality is measured directly at gNB. For both uplink and downlink, gNB selects the desired MCS number and signals it to the UE by Downlink Control Information (DCI) via PDCCH channel. For URLLC services, the UE can assist the gNB by advising that MCS targeting 10^{-5} Block Error Rate (BLER) are used. Robust link adaptation algorithms can maintain the needed level of reliability considering a given latency bound.

Low latency on the physical layer is provided by short transmission duration which is possible by using high Subcarrier Spacing (SCS) and the allocation of only one or a few Orthogonal Frequency Division Multiplexing (OFDM) symbols. For example, the shortest latency for the worst case in DL can be 0.23ms and in UL can be 0.24ms according to (section 5.7.1 in [TR37910]). Moreover, if the initial transmission has failed, HARQ feedback can quickly be provided and an HARQ retransmission is scheduled.

Dynamic multiplexing of data associated with different services is highly desirable for efficient use of system resources and to maximize system capacity. Assignment of resources for eMBB is usually done with regular (longer) transmission slots, which can lead to blocking of low latency services. To overcome the blocking, eMBB resources can be pre-empted and re-assigned to URLLC services. In this way, spectrally efficient assignments for eMBB can be ensured while providing flexibility required to ensure a bounded latency for URLLC services. In downlink, the gNB can notify the eMBB UE about pre-emption after it has happened, while in uplink there are two pre-emption mechanisms: special signaling to cancel eMBB transmission and URLLC dynamic power boost to suppress eMBB transmission.

6.4.4. Scheduling and QoS (MAC)

One integral part of the 5G system is the Quality of Service (QoS) framework [TS23501]. QoS flows are setup by the 5G system for certain IP or Ethernet packet flows, so that packets of each flow receive the same forwarding treatment, i.e., in scheduling and admission control. QoS flows can for example be associated with different priority level, packet delay budgets and tolerable packet error rates. Since radio resources are centrally scheduled in NR, the admission control function can ensure that only those QoS flows are admitted for which QoS targets can be reached.

NR transmissions in both UL and DL are scheduled by the gNB [TS38300]. This ensures radio resource efficiency, fairness in resource usage of the users and enables differentiated treatment of the data flows of the users according to the QoS targets of the flows. Those QoS flows are handled as data radio bearers or logical channels in NR RAN scheduling.

The gNB can dynamically assign DL and UL radio resources to users, indicating the resources as DL assignments or UL grants via control channel to the UE. Radio resources are defined as blocks of OFDM symbols in spectral domain and time domain. Different lengths are supported in time domain, i.e., (multiple) slot or mini-slot lengths. Resources of multiple frequency carriers can be aggregated and jointly scheduled to the UE.

Scheduling decisions are based, e.g., on channel quality measured on reference signals and reported by the UE (cf. periodical CSI reports for DL channel quality). The transmission reliability can be chosen in the scheduling algorithm, i.e., by link adaptation where an appropriate transmission format (e.g., robustness of modulation and coding scheme, controlled UL power) is selected for the radio channel condition of the UE. Retransmissions, based on HARQ feedback, are also controlled by the scheduler. The feedback transmission in HARQ

loop introduces delays, but there are methods to minimize it by using short transmission formats, sub-slot feedback reporting and PUCCH carrier switching. If needed to avoid HARQ round-trip time delays, repeated transmissions can be also scheduled beforehand, to the cost of reduced spectral efficiency.

In dynamic DL scheduling, transmission can be initiated immediately when DL data becomes available in the gNB. However, for dynamic UL scheduling, when data becomes available but no UL resources are available yet, the UE indicates the need for UL resources to the gNB via a (single bit) scheduling request message in the UL control channel. When thereupon UL resources are scheduled to the UE, the UE can transmit its data and may include a buffer status report, indicating the exact amount of data per logical channel still left to be sent. More UL resources may be scheduled accordingly. To avoid the latency introduced in the scheduling request loop, UL radio resources can also be pre-scheduled.

In particular for periodical traffic patterns, the pre-scheduling can rely on the scheduling features DL Semi-Persistent Scheduling (SPS) and UL Configured Grant (CG). With these features, periodically recurring resources can be assigned in DL and UL. Multiple parallels of those configurations are supported, in order to serve multiple parallel traffic flows of the same UE.

To support QoS enforcement in the case of mixed traffic with different QoS requirements, several features have recently been introduced. This way, e.g., different periodical critical QoS flows can be served together with best effort transmissions, by the same UE. Among others, these features (partly Release 16) are: 1) UL logical channel transmission restrictions allowing to map logical channels of certain QoS only to intended UL resources of a certain frequency carrier, slot-length, or CG configuration, and 2) intra-UE pre-emption and multiplexing, allowing critical UL transmissions to either pre-empt non-critical transmissions or being multiplexed with non-critical transmissions keeping different reliability targets.

When multiple frequency carriers are aggregated, duplicate parallel transmissions can be employed (beside repeated transmissions on one carrier). This is possible in the Carrier Aggregation (CA) architecture where those carriers originate from the same gNB, or in the Dual Connectivity (DC) architecture where the carriers originate from different gNBs, i.e., the UE is connected to two gNBs in this case. In both cases, transmission reliability is improved by this means of providing frequency diversity.

In addition to licensed spectrum, a 5G system can also utilize unlicensed spectrum to offload non-critical traffic. This version of NR is called NR-U, part of 3GPP Release 16. The central scheduling approach applies also for unlicensed radio resources, but in addition also the mandatory channel access mechanisms for unlicensed spectrum, e.g., Listen Before Talk (LBT) are supported in NR-U. This way, by using NR, operators have and can control access to both licensed and unlicensed frequency resources.

6.4.5. Time-Sensitive Communications (TSC)

Recent 3GPP releases have introduced various features to support multiple aspects of Time-Sensitive Communication ((TSC), which includes Time-Sensitive Networking (TSN) and beyond as described in this section.

The main objective of Time-Sensitive Networking (TSN) is to provide guaranteed data delivery within a guaranteed time window, i.e., bounded low latency. IEEE 802.1 TSN [IEEE802.1TSN] is a set of open standards that provide features to enable deterministic communication on standard IEEE 802.3 Ethernet [IEEE802.3]. TSN standards can be seen as a toolbox for traffic shaping, resource management, time synchronization, and reliability.

A TSN stream is a data flow between one end station (Talker) to another end station (Listener). In the centralized configuration model, TSN bridges are configured by the Central Network Controller (CNC) [IEEE802.1Qcc] to provide deterministic connectivity for the TSN stream through the network. Time-based traffic shaping provided by Scheduled Traffic [IEEE802.1Qbv] may be used to achieve bounded low latency. The TSN tool for time synchronization is the generalized Precision Time Protocol (gPTP) [IEEE802.1AS]), which provides reliable time synchronization that can be used by end stations and by other TSN tools, e.g., Scheduled Traffic [IEEE802.1Qbv]. High availability, as a result of ultra-reliability, is provided for data flows by the Frame Replication and Elimination for Reliability (FRER) [IEEE802.1CB] mechanism.

3GPP Release 16 includes integration of 5G with TSN, i.e., specifies functions for the 5G System (5GS) to deliver TSN streams such that they meet their QoS requirements. A key aspect of the integration is the 5GS appears from the rest of the network as a set of TSN bridges, in particular, one virtual bridge per User Plane Function (UPF) on the user plane. The 5GS includes TSN Translator (TT) functionality for the adaptation of the 5GS to the TSN bridged network and for hiding the 5GS internal procedures. The 5GS provides the following components:

1. interface to TSN controller, as per [IEEE802.1Qcc] for the fully centralized configuration model
2. time synchronization via reception and transmission of gPTP PDUs [IEEE802.1AS]
3. low latency, hence, can be integrated with Scheduled Traffic [IEEE802.1Qbv]
4. reliability, hence, can be integrated with FRER [IEEE802.1CB]

3GPP Release 17 [TS23501] introduced enhancements to generalize support for Time-Sensitive Communications (TSC) beyond TSN. This includes IP communications to provide time-sensitive service to, e.g., Video, Imaging and Audio for Professional Applications (VIAPA). The system model of 5G acting as a "TSN bridge" in Release 16 has been reused to enable the 5GS acting as a "TSC node" in a more generic sense (which includes TSN bridge and IP node). In the case of TSC that does not involve TSN, requirements are given via exposure interface and the control plane provides the service based on QoS and time synchronization requests from an Application Function (AF).

Figure 10 shows an illustration of 5G-TSN integration where an industrial controller (Ind Ctrlr) is connected to industrial Input/Output devices (I/O dev) via 5G. The 5GS can directly transport Ethernet frames since Release 15, thus, end-to-end Ethernet connectivity is provided. The 5GS implements the required interfaces towards the TSN controller functions such as the CNC, thus adapts to the settings of the TSN network. A 5G user plane virtual bridge interconnects TSN bridges or connect end stations, e.g., I/O devices to the network. Note that the introduction of 5G brings flexibility in various aspects, e.g., more flexible network topology because a wireless hop can replace several wireline hops thus significantly reduce the number of hops end-to-end. [TSN5G] dives more into the integration of 5G with TSN.




Figure 11: 5G - TSN Integration

NR supports accurate reference time synchronization in 1us accuracy level. Since NR is a scheduled system, an NR UE and a gNB are tightly synchronized to their OFDM symbol structures. A 5G internal reference time can be provided to the UE via broadcast or unicast signaling, associating a known OFDM symbol to this reference clock. The 5G internal reference time can be shared within the 5G network, i.e., radio and core network components. Release 16 has introduced interworking with gPTP for multiple time domains, where the 5GS acts as a virtual gPTP time-aware system and supports the forwarding of gPTP time synchronization information between end stations and bridges through the 5G user plane TTs. These account for the residence time of the 5GS in the time synchronization procedure. One special option is when the 5GS internal reference time is not only

used within the 5GS, but also to the rest of the devices in the deployment, including connected TSN bridges and end stations. Release 17 includes further improvements, i.e., methods for propagation delay compensation in RAN, further improving the accuracy for time synchronization over-the-air, as well as the possibility for the TSN grandmaster clock to reside on the UE side. More extensions and flexibility were added to the time synchronization service making it general for TSC with additional support of other types of clocks and time distribution such as boundary clock, transparent clock peer-to-peer, transparent clock end-to-end, aside from the time-aware system used for TSN. Additionally, it is possible to use internal access stratum signaling to distribute timing (and not the usual (g)PTP messages), for which the required accuracy can be provided by the AF [TS23501]. The same time synchronization service is expected to be further extended and enhanced in Release 18 to support Timing Resiliency (according to study item [SP211634]), where the 5G system can provide a back-up or alternative timing source for the failure of the local GNSS source (or other primary timing source) used by the vertical.

IETF Deterministic Networking (DetNet) is the technology to support time-sensitive communications at the IP layer. 3GPP Release 18 includes a study [SP211633] on whether and how to enable 3GPP support for DetNet such that a mapping is provided between DetNet and 5G. The support for DetNet is considered to be added via the TSC framework introduced for Release 17. The study includes what information needs to be exposed by the 5G System and the translation of DetNet flow specification to 5G QoS parameters. Note that TSN is the primary subnetwork technology for DetNet. Thus, the DetNet over TSN work, e.g., [RFC9023], can be leveraged via the TSN support built in 5G. As the standards are ready for such an approach, it is out of scope for the 3GPP Release 18 study item [SP211633].

Redundancy architectures were specified in order to provide reliability against any kind of failure on the radio link or nodes in the RAN and the core network, Redundant user plane paths can be provided based on the dual connectivity architecture, where the UE sets up two PDU sessions towards the same data network, and the 5G system makes the paths of the two PDU sessions independent as illustrated in Figure 13. There are two PDU sessions involved in the solution: the first spans from the UE via gNB1 to UPF1, acting as the first PDU session anchor, while the second spans from the UE via gNB2 to UPF2, acting as second the PDU session anchor. The independent paths may continue beyond the 3GPP network. Redundancy Handling Functions (RHF) are deployed outside of the 5GS, i.e., in Host A (the device) and in Host B (the network). RHF can implement replication and elimination functions as per [IEEE802.1CB] or the Packet Replication, Elimination, and Ordering Functions (PREOF) of IETF Deterministic Networking (DetNet) [RFC8655].




Figure 12: Reliability with Single UE

An alternative solution is that multiple UEs per device are used for user plane redundancy as illustrated in Figure 13. Each UE sets up a PDU session. The 5GS ensures that those PDU sessions of the different UEs are handled independently internal to the 5GS. There is no single point of failure in this solution, which also includes RHF outside of the 5G system, e.g., as per FRER or as PREOF specifications.




Figure 13: Reliability with Dual UE

Note that the abstraction provided by the RHF and the location of the RHF being outside of the 5G system make 5G equally supporting integration for reliability both with FRER of TSN and PREOF of DetNet as they both rely on the same concept.

6.5. Summary

5G technology enables deterministic communication. Based on the centralized admission control and the scheduling of the wireless resources, licensed or unlicensed, quality of service such as latency and reliability can be guaranteed. 5G contains several features to achieve ultra-reliable and low latency performance, e.g., support for different OFDM numerologies and slot-durations, as well as fast processing capabilities and redundancy techniques that lead to achievable latency numbers of below 1ms with 99.999% or higher confidence.

5G also includes features to support Industrial IoT use cases, e.g., via the integration of 5G with TSN. This includes 5G capabilities for each TSN component, latency, resource management, time synchronization, and reliability. Furthermore, 5G support for TSN can be leveraged when 5G is used as subnet technology for DetNet, in combination with or instead of TSN, which is the primary subnet for DetNet. In addition, the support for integration with TSN reliability was added to 5G by making DetNet reliability also applicable, thus making 5G DetNet ready. Moreover, providing IP service is native to 5G and adding direct support for DetNet is in scope of the upcoming 3GPP Release 18.

Overall, 5G provides scheduled wireless segments with high reliability and availability. In addition, 5G includes capabilities for integration to IP networks.

7. L-band Digital Aeronautical Communications System

One of the main pillars of the modern Air Traffic Management (ATM) system is the existence of a communication infrastructure that enables efficient aircraft guidance and safe separation in all phases of flight. Although current systems are technically mature, they are suffering from the VHF band's increasing saturation in high-density areas and the limitations posed by analogue radio. Therefore, aviation globally and the European Union (EU) in particular, strives for a sustainable modernization of the aeronautical communication infrastructure.

In the long-term, ATM communication shall transition from analogue VHF voice and VDL2 communication to more spectrum efficient digital data communication. The European ATM Master Plan foresees this transition to be realized for terrestrial communications by the development and implementation of the L-band Digital Aeronautical Communications System (LDACS). LDACS shall enable IPv6 based air-ground communication related to the safety and regularity of the flight. The particular challenge is that no new frequencies can be made available for terrestrial aeronautical communication. It was thus necessary to develop procedures to enable the operation of LDACS in parallel with other services in the same frequency band.

7.1. Provenance and Documents

The development of LDACS has already made substantial progress in the Single European Sky ATM Research (SESAR) framework, and is currently being continued in the follow-up program, SESAR2020 [RIH18]. A key objective of the SESAR activities is to develop, implement and validate a modern aeronautical data link able to evolve with aviation needs over long-term. To this end, an LDACS specification has been produced [GRA19] and is continuously updated; transmitter demonstrators were developed to test the spectrum compatibility of LDACS with legacy systems operating in the L-band [SAJ14]; and the overall system performance was analyzed by computer simulations, indicating that LDACS can fulfill the identified requirements [GRA11].

LDACS standardization within the framework of the International Civil Aviation Organization (ICAO) started in December 2016. The ICAO standardization group has produced an initial Standards and Recommended Practices (SARPs) document [ICAO18]. The SARPs document defines the general characteristics of LDACS. The ICAO standardization group plans to produce an ICAO technical manual - the ICAO equivalent to a technical standard - within the next years. Generally, the group is open to input from all sources and develops LDACS in the open.

Up to now the LDACS standardization has been focused on the development of the physical layer and the data link layer, only recently have higher layers come into the focus of the LDACS development activities. There is currently no "IPv6 over LDACS" specification; however, SESAR2020 has started the testing of IPv6-based LDACS testbeds. The IPv6 architecture for the aeronautical telecommunication network is called the Future Communications Infrastructure (FCI). FCI shall support quality of service, diversity, and mobility under the umbrella of the "multi-link concept". This work is conducted by ICAO working group WG-I.

In addition to standardization activities several industrial LDACS prototypes have been built. One set of LDACS prototypes has been evaluated in flight trials confirming the theoretical results predicting the system performance [GRA18][SCH19].

7.2. General Characteristics

LDACS will become one of several wireless access networks connecting aircraft to the Aeronautical Telecommunications Network (ATN). The LDACS access network contains several ground stations, each of them providing one LDACS radio cell. The LDACS air interface is a cellular data link with a star-topology connecting aircraft to ground-stations with a full duplex radio link. Each ground-station is the centralized instance controlling all air-ground communications within its radio cell.

The user data rate of LDACS is 315 kbit/s to 1428 kbit/s on the forward link, and 294 kbit/s to 1390 kbit/s on the reverse link, depending on coding and modulation. Due to strong interference from legacy systems in the L-band, the most robust coding and modulation SHOULD be expected for initial deployment i.e. 315/294 kbit/s on the forward/reverse link, respectively.

In addition to the communications capability, LDACS also offers a navigation capability. Ranging data, similar to DME (Distance Measuring Equipment), is extracted from the LDACS communication links between aircraft and LDACS ground stations. This results in LDACS providing an APNT (Alternative Position, Navigation and Timing) capability to supplement the existing on-board GNSS (Global Navigation Satellite System) without the need for additional bandwidth. Operationally, there will be no difference for pilots whether the navigation data are provided by LDACS or DME. This capability was flight tested and proven during the MICONAV flight trials in 2019 [BAT19].

In previous works and during the MICONAV flight campaign in 2019, it was also shown, that LDACS can be used for surveillance capability. Filip et al. [FIL19] shown passive radar capabilities of LDACS and Automatic Dependence Surveillance - Contract (ADS-C) was demonstrated via LDACS during the flight campaign 2019 [SCH19].

Since LDACS has been mainly designed for air traffic management communication it supports mutual entity authentication, integrity and confidentiality capabilities of user data messages and some control channel protection capabilities [MAE18], [MAE19], [MAE192], [MAE20].

Overall this makes LDACS the world's first truly integrated CNS system and is the worldwide most mature, secure, terrestrial long-range CNS technology for civil aviation.

7.3. Deployment and Spectrum

LDACS has its origin in merging parts of the B-VHF [BRA06], B-AMC [SCH08], TIA-902 (P34) [HAI09], and WiMAX IEEE 802.16e technologies [EHA11]. In 2007 the spectrum for LDACS was allocated at the World Radio Conference (WRC).

It was decided to allocate the spectrum next to Distance Measuring Equipment (DME), resulting in an in-lay approach between the DME channels for LDAC [SCH14].

LDACS is currently being standardized by ICAO and several roll-out strategies are discussed:

The LDACS data link provides enhanced capabilities to existing Aeronautical communications infrastructure enabling them to better support user needs and new applications. The deployment scalability of LDACS allows its implementation to start in areas where most needed to Improve immediately the performance of already fielded infrastructure. Later the deployment is extended based on operational demand. An attractive scenario for upgrading the existing VHF communication systems by adding an additional LDACS data link is described below.

When considering the current VDL Mode 2 infrastructure and user base, a very attractive win-win situation comes about, when the technological advantages of LDACS are combined with the existing VDL mode 2 infrastructure. LDACS provides at least 50 time more capacity than VDL Mode 2 and is a natural enhancement to the existing VDL Mode 2 business model. The advantage of this approach is that the VDL Mode 2 infrastructure can be fully reused. Beyond that, it opens the way for further enhancements which can increase business efficiency and minimize investment risk. [ICAO19]

7.4. Applicability to Deterministic Flows

As LDACS is a ground-based digital communications system for flight guidance and communications related to safety and regularity of flight, time-bounded deterministic arrival times for safety critical messages are a key feature for its successful deployment and roll-out.

7.4.1. System Architecture

Up to 512 Aircraft Station (AS) communicate to an LDACS Ground Station (GS) in the Reverse Link (RL). GS communicate to AS in the Forward Link (FL). Via an Access-Router (AC-R) GSs connect the LDACS sub-network to the global Aeronautical Telecommunications Network (ATN) to which the corresponding Air Traffic Services (ATS) and Aeronautical Operational Control (AOC) end systems are attached.

7.4.2. Overview of The Radio Protocol Stack

The protocol stack of LDACS is implemented in the AS and GS: It consists of the Physical Layer (PHY) with five major functional blocks above it. Four are placed in the Data Link Layer (DLL) of the AS and GS: (1) Medium Access Layer (MAC), (2) Voice Interface (VI), (3) Data Link Service (DLS), and (4) LDACS Management Entity (LME). The last entity resides within the Sub-Network Layer: Sub-Network Protocol (SNP). The LDACS network is externally connected to voice units, radio control units, and the ATN Network Layer.

Figure 14 shows the protocol stack of LDACS as implemented in the AS and GS.




Figure 14: LDACS protocol stack in AS and GS

7.4.3. Radio (PHY)

The physical layer provides the means to transfer data over the radio channel. The LDACS ground-station supports bi-directional links to multiple aircraft under its control. The forward link direction (FL; ground-to-air) and the reverse link direction (RL; air-to-ground) are separated by frequency division duplex. Forward link and reverse link use a 500 kHz channel each. The ground-station transmits a continuous stream of OFDM symbols on the forward link. In the reverse link different aircraft are separated in time and frequency using a combination of Orthogonal Frequency-Division Multiple-Access (OFDMA) and Time-Division Multiple-Access (TDMA). Aircraft thus transmit discontinuously on the reverse link with radio bursts sent

in precisely defined transmission opportunities allocated by the ground-station. The most important service on the PHY layer of LDACS is the PHY time framing service, which indicates that the PHY layer is ready to transmit in a given slot and to indicate PHY layer framing and timing to the MAC time framing service. LDACS does not support beam-forming or Multiple Input Multiple Output (MIMO).

7.4.4. Scheduling, Frame Structure and QoS (MAC)

The data-link layer provides the necessary protocols to facilitate concurrent and reliable data transfer for multiple users. The LDACS data link layer is organized in two sub-layers: The medium access sub-layer and the logical link control sub-layer. The medium access sub-layer manages the organization of transmission opportunities in slots of time and frequency. The logical link control sub-layer provides acknowledged point-to-point logical channels between the aircraft and the ground-station using an automatic repeat request protocol. LDACS supports also unacknowledged point-to-point channels and ground-to-air broadcast. Before going more into depth about the LDACS medium access, the frame structure of LDACS is introduced:

The LDACS framing structure for FL and RL is based on Super-Frames (SF) of 240 ms duration. Each SF corresponds to 2000 OFDM symbols. The FL and RL SF boundaries are aligned in time (from the view of the GS).

In the FL, an SF contains a Broadcast Frame of duration 6.72 ms (56 OFDM symbols) for the Broadcast Control Channel (BCCH), and four Multi-Frames (MF), each of duration 58.32 ms (486 OFDM symbols).

In the RL, each SF starts with a Random Access (RA) slot of length 6.72 ms with two opportunities for sending RL random access frames for the Random Access Channel (RACH), followed by four MFs. These MFs have the same fixed duration of 58.32 ms as in the FL, but a different internal structure

Figure 15 and Figure 16 illustrate the LDACS frame structure.

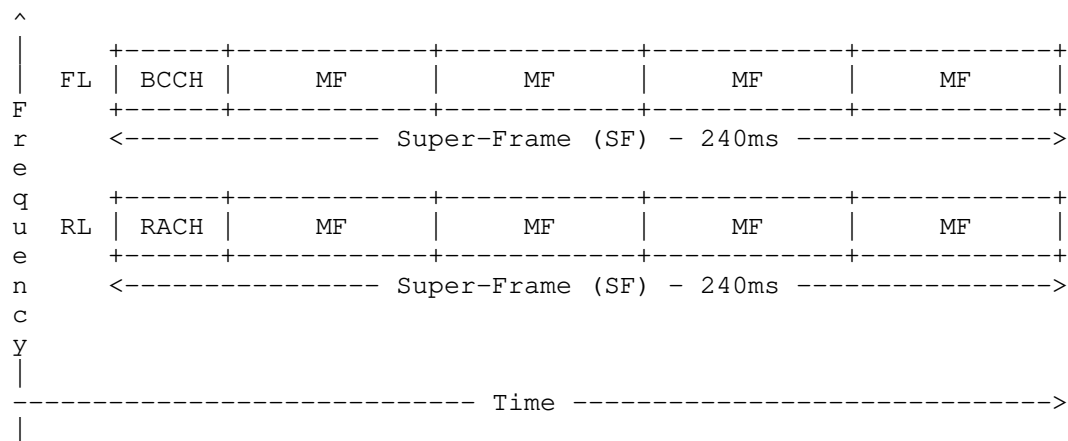


Figure 15: SF structure for LDACS




Figure 16: MF structure for LDACS

This fixed frame structure allows for a reliable and dependable transmission of data. Next, the LDACS medium access layer is introduced:

LDACS medium access is always under the control of the ground-station of a radio cell. Any medium access for the transmission of user data has to be requested with a resource request message stating the requested amount of resources and class of service. The ground-station performs resource scheduling on the basis of these requests

and grants resources with resource allocation messages. Resource request and allocation messages are exchanged over dedicated contention-free control channels.

LDACS has two mechanisms to request resources from the scheduler in the ground-station. Resources can either be requested "on demand" with a given class of service. On the forward link, this is done locally in the ground-station, on the reverse link a dedicated contention-free control channel is used (Dedicated Control Channel (DCCH); roughly 83 bit every 60 ms). A resource allocation is always announced in the control channel of the forward link (Common Control Channel (CCCH); variable sized). Due to the spacing of the reverse link control channels of every 60 ms, a medium access delay in the same order of magnitude is to be expected.

Resources can also be requested "permanently". The permanent resource request mechanism supports requesting recurring resources in given time intervals. A permanent resource request has to be canceled by the user (or by the ground-station, which is always in control). User data transmissions over LDACS are therefore always scheduled by the ground-station, while control data uses statically (i.e. at net entry) allocated recurring resources (DCCH and CCCH). The current specification documents specify no scheduling algorithm. However performance evaluations so far have used strict priority scheduling and round robin for equal priorities for simplicity. In the current prototype implementations LDACS classes of service are thus realized as priorities of medium access and not as flows. Note that this can starve out low priority flows. However, this is not seen as a big problem since safety related message always go first in any case. Scheduling of reverse link resources is done in physical Protocol Data Units (PDU) of 112 bit (or larger if more aggressive coding and modulation is used). Scheduling on the forward link is done Byte-wise since the forward link is transmitted continuously by the ground-station.

In order to support diversity, LDACS supports handovers to other ground-stations on different channels. Handovers may be initiated by the aircraft (break-before-make) or by the ground-station (make-before-break). Beyond this, FCI diversity shall be implemented by the multi-link concept.

7.5. Summary

LDACS has been designed with applications related to the safety and regularity of the flight in mind. It has therefore been designed as a deterministic wireless data link (as far as possible).

It is a secure, scalable and spectrum efficient data link with embedded navigation capability and thus, is the first truly integrated CNS system recognized by ICAO. During flight tests the LDACS capabilities have been successfully demonstrated. A viable roll-out scenario has been developed which allows gradual introduction of LDACS with immediate use and revenues. Finally, ICAO is developing LDACS standards to pave the way for a successful roll-out in the near future.

8. IANA Considerations

This specification does not require IANA action.

9. Security Considerations

Most RAW technologies integrate some authentication or encryption mechanisms that were defined outside the IETF.

10. Contributors

This document aggregates articles from authors specialized in each technologies. Beyond the main authors listed in the front page, the following contributors proposed additional text and refinement that improved the document greatly!

Georgios Z. Papadopoulos: Contributed to the TSCH section.

Nils Maeurer: Contributed to the LDACS section.

Thomas Graeupl: Contributed to the LDACS section.

Torsten Dudda, Alexey Shapin, and Sara Sandberg: Contributed to the 5G section.

Rocco Di Taranto: Contributed to the Wi-Fi section

Rute Sofia: Contributed to the Introduction and Terminology sections

11. Acknowledgments

Many thanks to the participants of the RAW WG where a lot of the work discussed here happened, and Malcolm Smith for his review of the 802.11 section.

12. Normative References

- [RFC5673] Pister, K., Ed., Thubert, P., Ed., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low-Power and Lossy Networks", RFC 5673, DOI 10.17487/RFC5673, October 2009, <<https://www.rfc-editor.org/info/rfc5673>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [RFC8557] Finn, N. and P. Thubert, "Deterministic Networking Problem Statement", RFC 8557, DOI 10.17487/RFC8557, May 2019, <<https://www.rfc-editor.org/info/rfc8557>>.
- [RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <<https://www.rfc-editor.org/info/rfc8655>>.
- [RFC9030] Thubert, P., Ed., "An Architecture for IPv6 over the Time-Slotted Channel Hopping Mode of IEEE 802.15.4 (6TiSCH)", RFC 9030, DOI 10.17487/RFC9030, May 2021, <<https://www.rfc-editor.org/info/rfc9030>>.

13. Informative References

- [RFC8480] Wang, Q., Ed., Vilajosana, X., and T. Watteyne, "6TiSCH Operation Sublayer (6top) Protocol (6P)", RFC 8480, DOI 10.17487/RFC8480, November 2018, <<https://www.rfc-editor.org/info/rfc8480>>.
- [RFC8938] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., and S. Bryant, "Deterministic Networking (DetNet) Data Plane Framework", RFC 8938, DOI 10.17487/RFC8938, November 2020, <<https://www.rfc-editor.org/info/rfc8938>>.
- [RFC9033] Chang, T., Ed., Vuini, M., Vilajosana, X., Duquennoy, S., and D. Dujovne, "6TiSCH Minimal Scheduling Function (MSF)", RFC 9033, DOI 10.17487/RFC9033, May 2021, <<https://www.rfc-editor.org/info/rfc9033>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <<https://www.rfc-editor.org/info/rfc6550>>.

- [RFC6551] Vasseur, JP., Ed., Kim, M., Ed., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics Used for Path Calculation in Low-Power and Lossy Networks", RFC 6551, DOI 10.17487/RFC6551, March 2012, <<https://www.rfc-editor.org/info/rfc6551>>.
- [RFC6291] Andersson, L., van Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "Guidelines for the Use of the "OAM" Acronym in the IETF", BCP 161, RFC 6291, DOI 10.17487/RFC6291, June 2011, <<https://www.rfc-editor.org/info/rfc6291>>.
- [RFC7276] Mizrahi, T., Sprecher, N., Bellagamba, E., and Y. Weingarten, "An Overview of Operations, Administration, and Maintenance (OAM) Tools", RFC 7276, DOI 10.17487/RFC7276, June 2014, <<https://www.rfc-editor.org/info/rfc7276>>.
- [RFC8279] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", RFC 8279, DOI 10.17487/RFC8279, November 2017, <<https://www.rfc-editor.org/info/rfc8279>>.
- [RFC9023] Varga, B., Ed., Farkas, J., Malis, A., and S. Bryant, "Deterministic Networking (DetNet) Data Plane: IP over IEEE 802.1 Time-Sensitive Networking (TSN)", RFC 9023, DOI 10.17487/RFC9023, June 2021, <<https://www.rfc-editor.org/info/rfc9023>>.
- [I-D.ietf-raw-architecture]
Thubert, P. and G. Z. Papadopoulos, "Reliable and Available Wireless Architecture", Work in Progress, Internet-Draft, draft-ietf-raw-architecture-03, 14 January 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-raw-architecture-03>>.
- [I-D.ietf-roll-nsa-extension]
Koutsiamanis, R., Papadopoulos, G., Montavont, N., and P. Thubert, "Common Ancestor Objective Function and Parent Set DAG Metric Container Extension", Work in Progress, Internet-Draft, draft-ietf-roll-nsa-extension-10, 29 October 2020, <<https://datatracker.ietf.org/doc/html/draft-ietf-roll-nsa-extension-10>>.
- [I-D.papadopoulos-paw-pre-reqs]
Papadopoulos, G., Koutsiamanis, R., Montavont, N., and P. Thubert, "Exploiting Packet Replication and Elimination in

Complex Tracks in LLNs", Work in Progress, Internet-Draft, draft-papadopoulos-paw-pre-reqs-01, 25 March 2019, <<https://datatracker.ietf.org/doc/html/draft-papadopoulos-paw-pre-reqs-01>>.

[I-D.thubert-bier-replication-elimination]

Thubert, P., Eckert, T., Brodard, Z., and H. Jiang, "BIER-TE extensions for Packet Replication and Elimination Function (PREF) and OAM", Work in Progress, Internet-Draft, draft-thubert-bier-replication-elimination-03, 3 March 2018, <<https://datatracker.ietf.org/doc/html/draft-thubert-bier-replication-elimination-03>>.

[I-D.thubert-6lo-bier-dispatch]

Thubert, P., Brodard, Z., Jiang, H., and G. Texier, "A 6loRH for BitStrings", Work in Progress, Internet-Draft, draft-thubert-6lo-bier-dispatch-06, 28 January 2019, <<https://datatracker.ietf.org/doc/html/draft-thubert-6lo-bier-dispatch-06>>.

[I-D.ietf-bier-te-arch]

Eckert, T., Cauchie, G., and M. Menth, "Tree Engineering for Bit Index Explicit Replication (BIER-TE)", Work in Progress, Internet-Draft, draft-ietf-bier-te-arch-11, 15 November 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-bier-te-arch-11>>.

[I-D.ietf-6tisch-coap]

Sudhaakar, R. S. and P. Zand, "6TiSCH Resource Management and Interaction using CoAP", Work in Progress, Internet-Draft, draft-ietf-6tisch-coap-03, 9 March 2015, <<https://datatracker.ietf.org/doc/html/draft-ietf-6tisch-coap-03>>.

[IEEE Std 802.15.4]

IEEE standard for Information Technology, "IEEE Std 802.15.4, Part. 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks".

[IEEE Std 802.11]

IEEE standard for Information Technology, "IEEE Standard 802.11 - IEEE Standard for Information Technology - Telecommunications and information exchange between systems Local and metropolitan area networks - Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications.", <<https://ieeexplore.ieee.org/document/9363693>>.

- [IEEE Std 802.11ax]
IEEE standard for Information Technology, "802.11ax:
Enhancements for High Efficiency WLAN", 2021,
<<https://ieeexplore.ieee.org/document/9442429>>.
- [IEEE Std 802.11ay]
IEEE standard for Information Technology, "802.11ay:
Enhanced throughput for operation in license-exempt bands
above 45 GHz", 2021,
<<https://ieeexplore.ieee.org/document/9502046/>>.
- [IEEE Std 802.11ad]
"802.11ad: Enhancements for very high throughput in the 60
GHz band", 2012,
<<https://ieeexplore.ieee.org/document/6392842/>>.
- [IEEE 802.11be WIP]
IEEE standard for Information Technology, "802.11be:
Extreme High Throughput PAR",
<[https://mentor.ieee.org/802.11/dcn/18/11-18-1231-04-0eht-
eht-draft-proposed-par.docx](https://mentor.ieee.org/802.11/dcn/18/11-18-1231-04-0eht-eht-draft-proposed-par.docx)>.
- [IEEE Std 802.1Qat]
"802.1Qat: Stream Reservation Protocol".
- [Cavalcanti_2019]
Dave Cavalcanti et al., "Extending Time Distribution and
Timeliness Capabilities over the Air to Enable Future
Wireless Industrial Automation Systems, the Proceedings of
IEEE", June 2019.
- [Nitsche_2015]
Thomas Nitsche et al., "IEEE 802.11ad: directional 60 GHz
communication for multi-Gigabit-per-second Wi-Fi",
December 2014.
- [Ghasempour_2017]
Yasaman Ghasempour et al., "802.11ay: Next-Generation 60
GHz Communications for 100 Gb/s Wi-Fi", December 2017.
- [IEEE_doc_11-18-2009-06]
IEEE standard for Information Technology, "802.11 Real-
Time Applications (RTA) Topic Interest Group (TIG)
Report", November 2018.
- [IEEE_doc_11-19-0373-00]
Kevin Stanton et Al., "Time-Sensitive Applications Support
in EHT", March 2019.

- [morell13] Antoni Morell et al., "Label switching over IEEE802.15.4e networks", April 2013.
- [dearmas16] Jesica de Armas et al., "Determinism through path diversity: Why packet replication makes sense", September 2016.
- [vilajosana19] Xavier Vilajosana et al., "6TiSCH: Industrial Performance for IPv6 Internet-of-Things Networks", June 2019.
- [ISA100.11a] ISA/IEC, "ISA100.11a, Wireless Systems for Automation, also IEC 62734", 2011, <<http://www.isa100wci.org/en-US/Documents/PDF/3405-ISA100-WirelessSystems-Future-broch-WEB-ETSI.aspx>>.
- [WirelessHART] www.hartcomm.org, "Industrial Communication Networks - Wireless Communication Network and Communication Profiles - WirelessHART - IEC 62591", 2010.
- [PCE] IETF, "Path Computation Element", <<https://dataTracker.ietf.org/doc/charter-ietf-pce/>>.
- [CCAMP] IETF, "Common Control and Measurement Plane", <<https://dataTracker.ietf.org/doc/charter-ietf-ccamp/>>.
- [TiSCH] IETF, "IPv6 over the TSCH mode over 802.15.4", <<https://dataTracker.ietf.org/doc/charter-ietf-6tisch/>>.
- [RIH18] Rihacek, C., Haindl, B., Fantappie, P., Pierattelli, S., Gräupl, T., Schnell, M., and N. Fistas, "L-band Digital Aeronautical Communications System (LDACS) Activities in SESAR2020", Proceedings of the Integrated Communications Navigation and Surveillance Conference (ICNS) Herndon, VA, USA, April 2018.
- [GRA19] Gräupl, T., Rihacek, C., and B. Haindl, "LDACS A/G Specification", SESAR2020 PJ14-02-01 D3.3.010, February 2019.
- [SAJ14] Sajatovic, M., Günzel, H., and S. Müller, "WA04 D22 Test Report for Assessing LDACS1 Transmitter Impact upon DME/TACAN Receivers", April 2014.

- [GRA11] Gräupl, T. and M. Ehammer, "L-DACS1 Data Link Layer Evolution of ATN/IPS", Proceedings of the 30th IEEE/AIAA Digital Avionics Systems Conference (DASC) Seattle, WA, USA, October 2011.
- [ICAO18] International Civil Aviation Organization (ICAO), "L-Band Digital Aeronautical Communication System (LDACS)", International Standards and Recommended Practices Annex 10 - Aeronautical Telecommunications, Vol. III - Communication Systems, July 2018.
- [GRA18] al., T. G. E., "L-band Digital Aeronautical Communications System (LDACS) flight trials in the national German project MICONAV", Proceedings of the Integrated Communications, Navigation, Surveillance Conference (ICNS) Herndon, VA, USA, April 2018.
- [SCH19] Schnell, M., "DLR tests digital communications technologies combined with additional navigation functions for the first time", 3 March 2019, <https://www.dlr.de/dlr/en/desktopdefault.aspx/tabid-10081/151_read-32951/#/gallery/33877>.
- [TR37910] 3GPP, "Study on self evaluation towards IMT-2020 submission", 3GPP TR 37.910, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3190>>.
- [TR38824] 3GPP, "Study on physical layer enhancements for NR ultra-reliable and low latency case (URLLC)", 3GPP TR 38.824, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3498>>.
- [TR38825] 3GPP, "Study on NR industrial Internet of Things (IoT)", 3GPP TR 38.825, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3492>>.
- [TS22104] 3GPP, "Service requirements for cyber-physical control applications in vertical domains", 3GPP TS 22.104, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3528>>.
- [TR22804] 3GPP, "Study on Communication for Automation in Vertical domains (CAV)", 3GPP TS 22.804, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3187>>.

- [TS23501] 3GPP, "System architecture for the 5G System (5GS)", 3GPP TS 23.501, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144>>.
- [TS38300] 3GPP, "NR Overall description", 3GPP TS 38.300, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3191>>.
- [RP210854] 3GPP, "Revised WID: Enhanced Industrial Internet of Things (IoT) and ultra-reliable and low latency communication (URLLC) support for NR", 3GPP RP-210854, March 2021, <https://www.3gpp.org/ftp/tsg_ran/TSG_RAN/TSGR_91e/Docs/RP-210854.zip>.
- [SP211633] 3GPP, "New SID: Extensions to the TSC Framework to support DetNet", 3GPP SP-211633, December 2021, <https://www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_94E_Electronic_2021_12/Docs/SP-211633.zip>.
- [SP211634] 3GPP, "Study on 5G Timing Resiliency, TSC, and URLLC enhancements", 3GPP SP-211634, December 2021, <https://www.3gpp.org/ftp/tsg_sa/TSG_SA/TSGS_94E_Electronic_2021_12/Docs/SP-211634.zip>.
- [IMT2020] "ITU towards IMT for 2020 and beyond", <<https://www.itu.int/en/ITU-R/study-groups/rsg5/rwp5d/imt-2020/Pages/default.aspx>>.
- [IEEE802.1TSN]
IEEE 802.1, "Time-Sensitive Networking (TSN) Task Group", <<http://www.ieee802.org/1/pages/tsn.html>>.
- [IEEE802.1AS]
IEEE, "IEEE Standard for Local and metropolitan area networks -- Timing and Synchronization for Time-Sensitive Applications", IEEE 802.1AS-2020, <https://standards.ieee.org/content/ieee-standards/en/standard/802_1AS-2020.html>.
- [IEEE802.1CB]
IEEE, "IEEE Standard for Local and metropolitan area networks -- Frame Replication and Elimination for Reliability", DOI 10.1109/IEEEStd2017.8091139, IEEE 802.1CB-2017, <<https://ieeexplore.ieee.org/document/8091139>>.

- [IEEE802.1Qbv]
IEEE, "IEEE Standard for Local and metropolitan area networks -- Bridges and Bridged Networks -- Amendment 25: Enhancements for Scheduled Traffic", IEEE 802.1Qbv-2015, <<https://ieeexplore.ieee.org/document/7440741>>.
- [IEEE802.1Qcc]
IEEE, "IEEE Standard for Local and metropolitan area networks -- Bridges and Bridged Networks -- Amendment 31: Stream Reservation Protocol (SRP) Enhancements and Performance Improvements", IEEE 802.1Qcc-2018, <<https://ieeexplore.ieee.org/document/8514112>>.
- [IEEE802.3]
IEEE, "IEEE Standard for Ethernet", IEEE 802.3-2018, <<https://ieeexplore.ieee.org/document/8457469>>.
- [TSN5G] 5G-ACIA, "Integration of 5G with Time-Sensitive Networking for Industrial Communications", 5G-ACIA whitepaper, <<https://5g-acia.org/whitepapers/integration-of-5g-with-time-sensitive-networking-for-industrial-communications>>.
- [MAE18] Maeurer, N. and A. Bilzhause, "A Cybersecurity Architecture for the L-band Digital Aeronautical Communications System (LDACS)", IEEE 37th Digital Avionics Systems Conference (DASC), pp. 1-10, London, UK , 2017.
- [MAE191] Maeurer, N. and C. Schmitt, "Towards Successful Realization of the LDACS Cybersecurity Architecture: An Updated Datalink Security Threat- and Risk Analysis", IEEE Integrated Communications, Navigation and Surveillance Conference (ICNS), pp. 1-13, Herndon, VA, USA , 2019.
- [ICAO19] International Civil Aviation Organization (ICAO), "TLDACS White PaperA Roll-out Scenario", Working Paper COMMUNICATIONS PANEL-DATA COMMUNICATIONS INFRASTRUCTURE WORKING GROUP, Montreal, Canada , October 2019.
- [MAE192] Maeurer, N., Graeupl, T., and C. Schmitt, "Evaluation of the LDACS Cybersecurity Implementation", IEEE 38th Digital Avionics Systems Conference (DACS), pp. 1-10, San Diego, CA, USA , September 2019.
- [MAE20] Maeurer, N., Graeupl, T., and C. Schmitt, "Comparing Different Diffie-Hellman Key Exchange Flavors for LDACS", IEEE 39th Digital Avionics Systems Conference (DACS), pp. 1-10, San Diego, CA, USA , October 2019.

- [FIL19] Filip-Dhaubhadel, A. and D. Shutin, "LDACS- Based Non-Cooperative Surveillance Multistatic Radar Design and Detection Coverage Assessment", IEEE 38th Digital Avionics Systems Conference (DACS), pp. 1-10, San Diego, CA, USA , September 2019.
- [BAT19] Battista, G., Osechas, O., Narayanan, S., Crespillo, O.G., Gerbeth, D., Maeurer, N., Mielke, D., and T. Graeupl, "Real-Time Demonstration of Integrated Communication and Navigation Services Using LDACS", IEEE Integrated Communications, Navigation and Surveillance Conference (ICNS), pp. 1-12, Herndon, VA, USA , 2019.
- [BRA06] Brandes, S., Schnell, M., Rokitansky, C.H., Ehammer, M., Graeupl, T., Steendam, H., Guenach, M., Rihacek, C., and B. Haindl, "B-VHF -Selected Simulation Results and Final Assessment", IEEE 25th Digital Avionics Systems Conference (DACS), pp. 1-12, New York, NY, USA , September 2019.
- [SCH08] Schnell, M., Brandes, S., Gligorevic, S., Rokitansky, C.H., Ehammer, M., Graeupl, T., Rihacek, C., and M. Sajatovic, "B-AMC - Broadband Aeronautical Multi-carrier Communications", IEEE 8th Integrated Communications, Navigation and Surveillance Conference (ICNS), pp. 1-13, New York, NY, USA , April 2008.
- [HAI09] Haindl, B., Rihacek, C., Sajatovic, M., Phillips, B., Budinger, J., Schnell, M., Kamiano, D., and W. Wilson, "Improvement of L-DACS1 Design by Combining B-AMC with P34 and WiMAX Technologies", IEEE 9th Integrated Communications, Navigation and Surveillance Conference (ICNS), pp. 1-8, New York, NY, USA , May 2009.
- [EHA11] Ehammer, M. and T. Graeupl, "AeroMACS - An Airport Communications System", IEEE 30th Digital Avionics Systems Conference (DACS), pp. 1-16, New York, NY, USA , September 2011.
- [SCH14] Schnell, M., Epple, U., Shutin, D., and N. Schneckenburger, "LDACS: Future Aeronautical Communications for Air- Traffic Management", IEEE Communications Magazine, 52(5), 104-110 , 2017.
- [Cavalcanti1287] Cavalcanti, D., Venkatesan, G., Cariou, L., and C. Cordeiro, "TSN support in 802.11 and potential extensions for TGbe", 2019,
<<https://mentor.ieee.org/802.11/dcn/19/11-19-1287>>.

[Sudhakaran2021]

Sudhakaran, S., Montgomery, K., Kashef, M., Cavalcanti, D., and R. Candell, "Wireless Time Sensitive Networking for Industrial Collaborative Robotic Workcells", 17th IEEE International Conference on Factory Communication Systems (WFCS) , 2021, <<https://ieeexplore.ieee.org/abstract/document/9483447>>.

[Fang_2021]

Fang, J., Cavalcanti, D., Cordeiro, C., and C. Cheng, "Wireless TSN with Multi-Radio Wi-Fi", IEEE International Conference on Standards for Communications and Networking, December 2021. , 2021.

Authors' Addresses

Pascal Thubert (editor)
Cisco Systems, Inc
Building D
45 Allee des Ormes - BP1200
06254 MOUGINS - Sophia Antipolis
France

Phone: +33 497 23 26 34
Email: pthubert@cisco.com

Dave Cavalcanti
Intel Corporation
2111 NE 25th Ave
Hillsboro, OR, 97124
United States of America

Phone: 503 712 5566
Email: dave.cavalcanti@intel.com

Xavier Vilajosana
Universitat Oberta de Catalunya
156 Rambla Poblenou
08018 Barcelona Catalonia
Spain

Email: xvilajosana@uoc.edu

Corinna Schmitt
Research Institute CODE, UniBw M
Werner-Heisenberg-Weg 39
85577 Neubiberg
Germany

Email: corinna.schmitt@unibw.de

Janos Farkas
Ericsson
Budapest
Magyar tudosok korutja 11
1117
Hungary

Email: janos.farkas@ericsson.com

RAW
Internet-Draft
Intended status: Standards Track
Expires: April 26, 2021

G. Papadopoulos
IMT Atlantique
P. Thubert
Cisco
F. Theoleyre
CNRS
CJ. Bernardos
UC3M
October 23, 2020

RAW use cases
draft-ietf-raw-use-cases-00

Abstract

The wireless medium presents significant specific challenges to achieve properties similar to those of wired deterministic networks. At the same time, a number of use cases cannot be solved with wires and justify the extra effort of going wireless. This document presents wireless use cases demanding reliable and available behavior.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 26, 2021.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Aeronautical Communications	5
2.1. Problem Statement	5
2.2. Specifics	5
2.3. Challenges	6
2.4. The Need for Wireless	7
2.5. Requirements for RAW	7
3. Amusement Parks	7
3.1. Use Case Description	7
3.2. Specifics	8
3.3. The Need for Wireless	8
3.4. Requirements for RAW	9
4. Wireless for Industrial Applications	9
4.1. Use Case Description	9
4.2. Specifics	10
4.2.1. Control Loops	10
4.2.2. Unmeasured Data	10
4.3. The Need for Wireless	11
4.4. Requirements for RAW	11
5. Pro Audio and Video	12
5.1. Use Case Description	12
5.2. Specifics	12
5.2.1. Uninterrupted Stream Playback	12
5.2.2. Synchronized Stream Playback	12
5.3. The Need for Wireless	12
5.4. Requirements for RAW	13
6. Wireless Gaming	13
6.1. Use Case Description	13
6.2. Specifics	14
6.3. The Need for Wireless	14
6.4. Requirements for RAW	14
7. UAV platooning and control	15
7.1. Use Case Description	15
7.2. Specifics	15
7.3. The Need for Wireless	15
7.4. Requirements for RAW	16
8. Edge Robotics control	16
8.1. Use Case Description	16
8.2. Specifics	17

8.3. The Need for Wireless	17
8.4. Requirements for RAW	17
9. Emergencies: Instrumented emergency vehicle	17
9.1. Use Case Description	17
9.2. Specifics	18
9.3. The Need for Wireless	18
9.4. Requirements for RAW	18
10. IANA Considerations	19
11. Security Considerations	19
12. Acknowledgments	19
13. Informative References	19
Authors' Addresses	22

1. Introduction

Based on time, resource reservation, and policy enforcement by distributed shapers, Deterministic Networking provides the capability to carry specified unicast or multicast data streams for real-time applications with extremely low data loss rates and bounded latency, so as to support time-sensitive and mission-critical applications on a converged enterprise infrastructure.

Deterministic Networking in the IP world is an attempt to eliminate packet loss for a committed bandwidth while ensuring a worst case end-to-end latency, regardless of the network conditions and across technologies. It can be seen as a set of new Quality of Service (QoS) guarantees of worst-case delivery. IP networks become more deterministic when the effects of statistical multiplexing (jitter and collision loss) are mostly eliminated. This requires a tight control of the physical resources to maintain the amount of traffic within the physical capabilities of the underlying technology, e.g., by the use of time-shared resources (bandwidth and buffers) per circuit, and/or by shaping and/or scheduling the packets at every hop.

Key attributes of Deterministic Networking include:

- o time synchronization on all the nodes,
- o centralized computation of network-wide deterministic paths,
- o multi-technology path with co-channel interference minimization,
- o frame preemption and guard time mechanisms to ensure a worst-case delay, and
- o new traffic shapers within and at the edge to protect the network.

Wireless operates on a shared medium, and transmissions cannot be fully deterministic due to uncontrolled interferences, including self-induced multipath fading. RAW (Reliable and Available Wireless) is an effort to provide Deterministic Networking Mechanisms on across a path that include a wireless physical layer. Making Wireless Reliable and Available is even more challenging than it is with wires, due to the numerous causes of loss in transmission that add up to the congestion losses and the delays caused by overbooked shared resources.

The wireless and wired media are fundamentally different at the physical level, and while the generic Problem Statement [RFC8557] for DetNet applies to the wired as well as the wireless medium, the methods to achieve RAW necessarily differ from those used to support Time-Sensitive Networking over wires.

So far, Open Standards for Deterministic Networking have prevalently been focused on wired media, with Audio/Video Bridging (AVB) and Time Sensitive Networking (TSN) at the IEEE and DetNet [RFC8655] at the IETF. But wires cannot be used in a number of cases, including mobile or rotating devices, rehabilitated industrial buildings, wearable or in-body sensory devices, vehicle automation and multiplayer gaming.

Purpose-built wireless technologies such as [ISA100], which incorporates IPv6, were developed and deployed to cope for the lack of open standards, but they yield a high cost in OPEX and CAPEX and are limited to very few industries, e.g., process control, concert instruments or racing.

This is now changing [I-D.thubert-raw-technologies]:

- o IMT-2020 has recognized Ultra-Reliable Low-Latency Communication (URLLC) as a key functionality for the upcoming 5G.
- o IEEE 802.11 has identified a set of real-applications [ieee80211-rt-tig] which may use the IEEE802.11 standards. They typically emphasize strict end-to-end delay requirements.
- o The IETF has produced an IPv6 stack for IEEE Std. 802.15.4 TimeSlotted Channel Hopping (TSCH) and an architecture [I-D.ietf-6tisch-architecture] that enables Reliable and Available Wireless (RAW) on a shared MAC.

This draft extends the "Deterministic Networking Use Cases" document [RFC8578] and describes a number of additional use cases which require "reliable/predictable and available" flows over wireless links and possibly complex multi-hop paths called Tracks. This is

covered mainly by the "Wireless for Industrial Applications" use case, as the "Cellular Radio" is mostly dedicated to the (wired) transport part of a Radio Access Network (RAN). Whereas the "Wireless for Industrial Applications" use case certainly covers an area of interest for RAW, it is limited to 6TiSCH, and thus its scope is narrower than the use cases described next in this document.

2. Aeronautical Communications

Aircraft are currently connected to ATC (Air-Traffic Control) and AOC (Airline Operational Control) via voice and data communications systems through all phases of a flight. Within the airport terminal, connectivity is focused on high bandwidth communications while during en-route high reliability, robustness and range is the main focus.

2.1. Problem Statement

Worldwide civil air traffic is expected to grow by 84% until 2040 compared to 2017 [EURO20]. Thus, legacy systems in air traffic management (ATM) are likely to reach their capacity limits and the need for new aeronautical communication technologies becomes apparent. Especially problematic is the saturation of VHF band in high density areas in Europe, the US, and Asia [KEAV20] [FAA20] calling for suitable new digital approaches such as AeroMACS for airport communications, SatCOM for remote domains, and LDACS as long-range terrestrial aeronautical communications system. Making the frequency spectrum's usage more efficient a transition from analogue voice to digital data communication [PLA14] is necessary to cope with the expected growth of civil aviation and its supporting infrastructure. A promising candidate for long range terrestrial communications, already in the process of being standardized in the International Civil Aviation Organization (ICAO), is the L-band Digital Aeronautical Communications System (LDACS) [ICAO18] [I-D.maeurer-raw-ldacs].

2.2. Specifics

During the creation process of new communications system, analogue voice is replaced by digital data communication. This sets a paradigm shift from analogue to digital wireless communications and supports the related trend towards increased autonomous data processing that the Future Communications Infrastructure (FCI) in civil aviation must provide. The FCI is depicted in Figure 1:




Figure 1: The Future Communication Infrastructure (FCI): AeroMACS for APT/TMA domain, LDACS A/G for TMA/ENR domain, LDACS A/G for ENR/ORP domain, SatCOM for ORP domain communications

2.3. Challenges

This paradigm change brings a lot of new challenges:

- o **Efficiency:** It is necessary to keep latency, time and data overhead (routing, security) of new aeronautical datalinks at a minimum.

- o Modularity: Systems in avionics usually operate up to 30 years, thus solutions must be modular, easily adaptable and updatable.
- o Interoperability: All 192 members of the international Civil Aviation Organization (ICAO) must be able to use these solutions.

2.4. The Need for Wireless

In a high mobility environment such as aviation, the envisioned solutions to provide worldwide coverage of data connections with in-flight aircraft require a multi-system, multi-link, multi-hop approach. Thus air, ground and space based datalink providing technologies will have to operate seamlessly together to cope with the increasing needs of data exchange between aircraft, air traffic controller, airport infrastructure, airlines, air network service providers (ANSPs) and so forth. Thus making use of wireless technologies is a must in tackling this enormous need for a worldwide digital aeronautical datalink infrastructure.

2.5. Requirements for RAW

Different safety levels need to be supported, from extremely safety critical ones requiring low latency, such as a WAKE warning - a warning that two aircraft come dangerously close to each other - and high resiliency, to less safety critical ones requiring low-medium latency for services such as WXGRAPH - graphical weather data.

Overhead needs to be kept at a minimum since aeronautical data links provide comparatively small data rates in the order of kbit/s.

Policy needs to be supported when selecting data links. The focus of RAW here should be on the selectors, responsible for the routing path a packet takes to reach its end destination. This would minimize the amount of routing information that has to travel inside the network because of precomputed routing tables with the selector being responsible for choosing the most appropriate option according to policy and safety.

3. Amusement Parks

3.1. Use Case Description

The digitalization of Amusement Parks is expected to decrease significantly the cost for maintaining the attractions. Such deployment is a mix between industrial automation (aka. Smart Factories) and multimedia entertainment applications.

Attractions may rely on a large set of sensors and actuators, which react in real time. Typical applications comprise:

- o Emergency: safety has to be preserved, and must stop the attraction when a failure is detected.
- o Video: augmented and virtual realities are integrated in the attraction. Wearable mobile devices (e.g., glasses, virtual reality headset) need to offload one part of the processing tasks.
- o Real-time interactions: visitors may interact with an attraction, like in a real-time video game. The visitors may virtually interact with their environment, triggering actions in the real world (through actuators) [robots].
- o Geolocation: visitors are tracked with a personal wireless tag so that their user experience is improved.
- o Predictive maintenance: statistics are collected to predict the future failures, or to compute later more complex statistics about the attraction's usage, the downtime, its popularity, etc.

3.2. Specifics

Amusement parks comprise a variable number of attractions, mostly outdoor, over a large geographical area. The IT infrastructure is typically multi-scale:

- o Local area: the sensors and actuators controlling the attractions are co-located. Control loops trigger only local traffic, with a small end-to-end delay, typically inferior than 10 milliseconds, like classical industrial systems [ieee80211-rt-tig].
- o Wearable mobile devices are free to move in the park. They exchange traffic locally (identification, personalization, multimedia) or globally (billing, child tracking).
- o Computationally intensive applications offload some tasks. Edge computing seems an efficient way to implement real-time applications with offloading. Some non time-critical tasks may rather use the cloud (predictive maintenance, marketing).

3.3. The Need for Wireless

Amusement parks cover large areas and a global interconnection would require a huge length of cables. Wireless also increases the reconfigurability, enabling to update cheaply the attractions. The frequent renewal helps to increase customer loyalty.

Some parts of the attraction are mobile, e.g., trucks of a roller-coaster, robots. Since cables are prone to frequent failures in this situation, wireless transmissions are recommended.

Wearable devices are extensively used for a user experience personalization. They typically need to support wireless transmissions. Personal tags may help to reduce the operating costs [disney-VIP] and to increase the number of charged services provided to the audience (VIP tickets, interactivity, etc.) Some applications rely on more sophisticated wearable devices such as digital glasses or Virtual Reality (VR) headsets for an immersive experience.

3.4. Requirements for RAW

The network infrastructure has to support heterogeneous traffic, with very different critical requirements. Thus, flow isolation has to be provided.

We have to schedule appropriately the transmissions, even in presence of mobile devices. While the [I-D.ietf-6tisch-architecture] already proposes an architecture for synchronized, IEEE Std. 802.15.4 Time-Slotted Channel Hopping (TSCH) networks, we still need multi-technology solutions, able to guarantee end-to-end requirements across heterogeneous technologies, with strict SLA requirements.

Nowadays, long-range wireless transmissions are used mostly for best-effort traffic. On the contrary, [IEEE802.1TSN] is used for critical flows using Ethernet devices. However, we need an IP enabled technology to interconnect large areas, independent of the PHY and MAC layers.

We expect to deploy several different technologies (long vs. short range) which have to cohabit in the same area. Thus, we need to provide layer-3 mechanisms able to exploit multiple co-interfering technologies.

4. Wireless for Industrial Applications

4.1. Use Case Description

A major use case for networking in Industrial environments is the control networks where periodic control loops operate between a sensor that measures a physical property such as the temperature of a fluid, a Programmable Logic Controller (PLC) that decides an action such as warm up the mix, and an actuator that performs the required action, e.g., inject power in a resistor.

4.2. Specifics

4.2.1. Control Loops

Process Control designates continuous processing operations, e.g., heating Oil in a refinery or mixing drinking soda. Control loops in the Process Control industry operate at a very low rate, typically 4 times per second. Factory Automation, on the other hand, deal with discrete goods such as individual automobile parts, and requires faster loops, in the order of 10ms. Motion control that monitors dynamic activities may require even faster rates in the order of a few ms. Finally, some industries exhibit hybrid behaviors, like canned soup that will start as a process industry while mixing the food and then operate as a discrete manufacturing when putting the final product in cans and shipping them.

In all those cases, a packet must flow reliably between the sensor and the PLC, be processed by the PLC, and sent to the actuator within the control loop period. In some particular use cases that inherit from analog operations, jitter might also alter the operation of the control loop. A rare packet loss is usually admissible, but typically 4 losses in a row will cause an emergency halt of the production and incur a high cost for the manufacturer.

4.2.2. Unmeasured Data

A secondary use case deals with monitoring and diagnostics. This so-called unmeasured data is essential to improve the performances of a production line, e.g., by optimizing real-time processing or maintenance windows using Machine Learning predictions. For the lack of wireless technologies, some specific industries such as Oil and Gas have been using serial cables, literally by the millions, to perform their process optimization over the previous decades. But few industries would afford the associated cost and the Holy Grail of the Industrial Internet of Things is to provide the same benefits to all industries, including SmartGrid, Transportation, Building, Commercial and Medical. This requires a cheap, available and scalable IP-based access technology.

Inside the factory, wires may already be available to operate the Control Network. But unmeasured data are not welcome in that network for a number of reasons. On the one hand it is rich and asynchronous, meaning that using they may influence the deterministic nature of the control operations and impact the production. On the other hand, this information must be reported to the carpeted floor over IP, which means the potential for a security breach via the interconnection of the Operational Technology (OT) network with the Internet technology (IT) network and possibly enable a rogue access.

4.3. The Need for Wireless

Ethernet cables used on a robot arm are prone to breakage after a few thousands flexions, a lot faster than a power cable that is wider in diameter, and more resilient. In general, wired networking and mobile parts are not a good match, mostly in the case of fast and recurrent activities, as well as rotation.

When refurbishing older premises that were built before the Internet age, power is usually available everywhere, but data is not. It is often impractical, time consuming and expensive to deploy an Ethernet fabric across walls and between buildings. Deploying a wire may take months and cost tens of thousands of US Dollars.

Even when wiring exists, e.g., in an existing control network, asynchronous IP packets such as diagnostics may not be welcome for operational and security reasons (see Section 4.2.1). An alternate network that can scale with the many sensors and actuators that equip every robot, every valve and fan that are deployed on the factory floor and may help detect and prevent a failure that could impact the production. IEEE Std. 802.15.4 Time-Slotted Channel Hopping (TSCH) [RFC7554] is a promising technology for that purpose, mostly if the scheduled operations enable to use the same network by asynchronous and deterministic flows in parallel.

4.4. Requirements for RAW

As stated by the "Deterministic Networking Problem Statement" [RFC8557], a Deterministic Network is backwards compatible with (capable of transporting) statistically multiplexed traffic while preserving the properties of the accepted deterministic flows. While the [I-D.ietf-6tisch-architecture] serves that requirement, the work at 6TiSCH was focused on best-effort IPv6 packet flows. RAW should be able to lock so-called hard cells for use by a centralized scheduler, and program so-called end-to-end Tracks over those cells.

Over the course of the recent years, major Industrial Protocols, e.g., [ODVA] with EtherNet/IP [EIP] and [Profinet], have been migrating towards Ethernet and IP. In order to unleash the full power of the IP hourglass model, it should be possible to deploy any application over any network that has the physical capacity to transport the industrial flow, regardless of the MAC/PHY technology, wired or wireless, and across technologies. RAW mechanisms should be able to setup a Track over a wireless access segment such as TSCH and a backbone segment such as Ethernet or WI-Fi, to report a sensor data or a critical monitoring within a bounded latency. It is also important to ensure that RAW solutions are interoperable with existing wireless solutions in place, and with legacy equipment which

capabilities can be extended using retrofitting. Maintainability, as a broader concept than reliability is also important in industrial scenarios [square-peg].

5. Pro Audio and Video

5.1. Use Case Description

Many devices support audio and video streaming by employing 802.11 wireless LAN. Some of these applications require low latency capability. For instance, when the application provides interactive play, or when the audio takes plays in real time (i.e. live) for public addresses in train stations or in theme parks.

The professional audio and video industry ("ProAV") includes:

- o Virtual Reality / Augmented Reality (VR/AR)
- o Public address, media and emergency systems at large venues (airports, train stations, stadiums, theme parks).

5.2. Specifics

5.2.1. Uninterrupted Stream Playback

Considering the uninterrupted audio or video stream, a potential packet losses during the transmission of audio or video flows cannot be tackled by re-trying the transmission, as it is done with file transfer, because by the time the packet lost has been identified it is too late to proceed with packet re-transmission. Buffering might be employed to provide a certain delay which will allow for one or more re-transmissions, however such approach is not efficient in application where delays are not acceptable.

5.2.2. Synchronized Stream Playback

In the context of ProAV, latency is the time between the transmitted signal over a stream and its reception. Thus, for sound to remain synchronized to the movement in the video, the latency of both the audio and video streams must be bounded and consistent.

5.3. The Need for Wireless

The devices need the wireless communication to support video streaming via 802.11 wireless LAN for instance.

During the public address, the deployed announcement speakers, for instance along the platforms of the train stations, need the wireless communication to forward the audio traffic in real time.

5.4. Requirements for RAW

The network infrastructure needs to support heterogeneous types of traffic (including QoS).

Content delivery with bounded (lowest possible) latency.

The deployed network topology should allow for multipath. This will enable for multiple streams to have different (and multiple) paths through the network to support redundancy.

6. Wireless Gaming

6.1. Use Case Description

The gaming industry includes [IEEE80211RTA] real-time mobile gaming, wireless console gaming and cloud gaming. For RAW, wireless console gaming is the most relevant one. We next summarize the three:

- o Real-time Mobile Gaming: Different from traditional games, real time mobile gaming is very sensitive to network latency and stability. The mobile game can connect multiple players together in a single game session and exchange data messages between game server and connected players. Real-time means the feedback should present on screen as users operate in game. For good game experience, the end to end latency plus game servers processing time should not be noticed by users as they play the game.
- o Wireless Console Gaming: Playing online on a console has 2 types of internet connectivity, which is either wired or Wi-Fi. Most of the gaming consoles today support Wi-Fi 5. But Wi-Fi has an especially bad reputation among the gaming community. The main reasons are high latency, lag spikes and jitter.
- o Cloud Gaming: The cloud gaming requires low latency capability as the user commands in a game session need to be sent back to the cloud server, the cloud server would update game context depending on the received commands, and the cloud server would render the picture/video to be displayed at user devices and stream the picture/video content to the user devices. User devices might very likely be connected wirelessly.

6.2. Specifics

While a lot of details can be found on [IEEE80211RTA], we next summarize the main requirements in terms of latency, jitter and packet loss:

- o Intra BSS latency: less than 5 ms.
- o Jitter variance: less than 2 ms.
- o Packet loss: less than 0.1 percent.

6.3. The Need for Wireless

It is clear that gaming is evolving towards wireless, as players demand being able to play anywhere. Besides, the industry is changing towards playing from mobile phones, which are inherently connected via wireless technologies.

6.4. Requirements for RAW

- o Time sensitive networking extensions. Extensions, such as time-aware shaping and redundancy (FRE) can be explored to address congestion and reliability problems present in wireless networks.
- o Priority tagging (Stream identification). One basic requirement to provide better QoS for time-sensitive traffic is the capability to identify and differentiate time-sensitive packets from other (e.g. best-effort) traffic.
- o Time-aware shaping. This capability (defined in IEEE 802.1Qbv) consists of gates to control the opening/closing of queues that share a common egress port within an Ethernet switch. A scheduler defines the times when each queue opens or close, therefore eliminating congestion and ensuring that frames are delivered within the expected latency bounds.
- o Dual/multiple link. Due to the competitions and interference are common and hardly in control under wireless network, in order to improve the latency stability, dual/multiple link proposal is brought up to address this issue. Two modes are defined: duplicate and joint.
- o Admission Control. Congestion is a major cause of high/variable latency and it is well known that if the traffic load exceeds the capability of the link, QoS will be degraded. QoS degradation maybe acceptable for many applications today, however emerging time-sensitive applications are highly susceptible to increased

latency and jitter. In order to better control QoS, it is important to control access to the network resources.

7. UAV platooning and control

7.1. Use Case Description

Unmanned Aerial Vehicles (UAVs) are becoming very popular for many different applications, including military and civil use cases. The term drone is commonly used to refer to a UAV.

UAVs can be used to perform aerial surveillance activities, traffic monitoring (e.g., Spanish traffic control has recently introduced a fleet of drones for quicker reactions upon traffic congestion related events), support of emergency situations, and even transportation of small goods.

UAVs typically have various forms of wireless connectivity:

- o cellular: for communication with the control center, for remote maneuvering as well as monitoring of the drone;
- o IEEE 802.11: for inter-drone communications (e.g., platooning) and providing connectivity to other devices (e.g., acting as Access Point).

7.2. Specifics

Some of the use cases/tasks involving drones require coordination among drones. Others involve complex compute tasks that might not be performed using the limited computing resources that a drone typically has. These two aspects require continuous connectivity with the control center and among drones.

Remote maneuvering of a drone might be performed over a cellular network in some cases, however, there are situations that need very low latencies and deterministic behavior of the connectivity. Examples involve platooning of drones or share of computing resources among drones (e.g., a drone offload some function to a neighboring drone).

7.3. The Need for Wireless

UAVs cannot be connected through any type of wired media, so it is obvious that wireless is needed.

7.4. Requirements for RAW

The network infrastructure is actually composed by the UAVs themselves, requiring self-configuration capabilities.

Heterogeneous types of traffic need to be supported, from extremely critical ones requiring ultra low latency and high resiliency, to traffic requiring low-medium latency.

When a given service is decomposed into functions -- hosted at different drones -- chained, each link connecting two given functions would have a well-defined set of requirements (latency, bandwidth and jitter) that have to be met.

8. Edge Robotics control

8.1. Use Case Description

The Edge Robotics scenario consists of several robots, deployed in a given area (for example a shopping mall), inter-connected via an access network to a network's edge device or a data center. The robots are connected to the edge so complex computational activities are not executed locally at the robots, but offloaded to the edge. This brings additional flexibility in the type of tasks that the robots do, as well as reducing the costs of robot manufacturing (due to their lower complexity), and enabling complex tasks involving coordination among robots (that can be more easily performed if robots are centrally controlled).

A simple example of the use of multiples robots is cleaning, delivering of goods from warehouses to shops or video surveillance. Multiple robots are simultaneously instructed to perform individual tasks by moving the robotic intelligence from the robots to the network's edge (e.g., data center). That enables easy synchronization, scalable solution and on-demand option to create flexible fleet of robots.

Robots would have various forms of wireless connectivity:

- o IEEE 802.11: for connection to the edge and also inter-robot communications (e.g., for coordinated actions).
- o Cellular: as an additional communication link to the edge, though primarily as backup, since ultra low latencies are needed.

8.2. Specifics

Some of the use cases/tasks involving robots might benefit from decomposition of a service in small functions that are distributed and chained among robots and the edge. These require continuous connectivity with the control center and among drones.

Robot control is an activity requiring very low latencies between the robot and the location where the control intelligence resides (which might be the edge or another robot).

8.3. The Need for Wireless

Deploying robots in scenarios such as shopping malls for the aforementioned applications cannot be done via wired connectivity.

8.4. Requirements for RAW

The network infrastructure needs to support heterogeneous types of traffic, from robot control to video streaming.

When a given service is decomposed into functions -- hosted at different robots -- chained, each link connecting two given functions would have a well-defined set of requirements (latency, bandwidth and jitter) that have to be met.

9. Emergencies: Instrumented emergency vehicle

9.1. Use Case Description

An instrumented ambulance would be one that has a LAN to which are connected these end systems:

- o vital signs sensors attached to the casualty in the ambulance. Relay medical data to hospital emergency room,
- o radionavigation sensor to relay position data to various destinations including dispatcher,
- o voice communication for ambulance attendant (e.g. consult with ER doctor),
- o voice communication between driver and dispatcher,
- o etc.

The LAN needs to be routed through radio-WANs to complete the internetwork linkage.

9.2. Specifics

What we have today is multiple communications systems to reach the vehicle:

- o A dispatching system,
- o a cellphone for the attendant,
- o a special purpose telemetering system for medical data,
- o etc.

This redundancy of systems, because of its stovepiping, does not contribute to availability as a whole.

Most of the scenarios involving the use of an instrumented ambulance are composed of many different flows, each of them with slightly different requirements in terms of reliability and latency. Destinations might be either at the ambulance itself (local traffic), at a near edge cloud or at the general Internet/cloud.

9.3. The Need for Wireless

Local traffic between the first responders/ambulance staff and the ambulance equipment cannot be done via wired connectivity as the responders perform initial treatment outside of the ambulance. The communications from the ambulance to external services has to be wireless as well.

9.4. Requirements for RAW

We can derive some pertinent requirements from this scenario:

- o High availability of the internetwork is required.
- o The internetwork needs to operate in damaged state (e.g. during an earthquake aftermath, heavy weather, wildfire, etc.). In addition to continuity of operations, rapid restoral is a needed characteristic.
- o End-to-end security, both authenticity and confidentiality, is required of traffic. All data needs to be authenticated; some (such as medical) needs to be confidential.
- o The radio-WAN has characteristics similar to cellphone -- the vehicle will travel from one radio footprint to another.

10. IANA Considerations

This document has no IANA actions.

11. Security Considerations

This document covers a number of representative applications and network scenarios that are expected to make use of RAW technologies. Each of the potential RAW use cases will have security considerations from both the use-specific perspective and the RAW technology perspective. [I-D.ietf-detnet-security] provides a comprehensive discussion of security considerations in the context of Deterministic Networking, which are generally applicable also to RAW.

12. Acknowledgments

Nils Maeurer, Thomas Graeupl and Corinna Schmitt have contributed significantly to this document, providing input for the Aeronautical communications section. Rex Buddenberg has also contributed to the document, providing input to the Emergency: instrumented emergency vehicle section.

The authors would like to thank Toerless Eckert, Xavi Vilajosana Guillen and Rute Sofia for their valuable comments on previous versions of this document.

The work of Carlos J. Bernardos in this draft has been partially supported by the H2020 5Growth (Grant 856709) and 5G-DIVE projects (Grant 859881).

13. Informative References

[disney-VIP]

Wired, "Disney's \$1 Billion Bet on a Magical Wristband",
March 2015,
<<https://www.wired.com/2015/03/disney-magicband/>>.

[EIP]

<http://www.odva.org/>, "EtherNet/IP provides users with the network tools to deploy standard Ethernet technology (IEEE 802.3 combined with the TCP/IP Suite) for industrial automation applications while enabling Internet and enterprise connectivity data anytime, anywhere.",
<http://www.odva.org/Portals/0/Library/Publications_Numbered/PUB00138R3_CIP_Adv_Tech_Series_EtherNetIP.pdf>.

- [EURO20] EUROCONTROL, "EUROCONTROL's Challenges of Growth 2018 Study Report", 2018, <<https://www.eurocontrol.int/sites/default/files/content/documents/officialdocuments/reports/challenges-of-growth-2018.pdf>>.
- [FAA20] U.S. Department of Transportation Federal Aviation Administration (FAA), "Next Generation Air Transportation System", 2019, <<https://www.faa.gov/nextgen/>>.
- [I-D.ietf-6tisch-architecture]
Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", draft-ietf-6tisch-architecture-29 (work in progress), August 2020.
- [I-D.ietf-detnet-security]
Grossman, E., Mizrahi, T., and A. Hacker, "Deterministic Networking (DetNet) Security Considerations", draft-ietf-detnet-security-12 (work in progress), October 2020.
- [I-D.maeurer-raw-ldacs]
Maeurer, N., Graeupl, T., and C. Schmitt, "L-band Digital Aeronautical Communications System (LDACS)", draft-maeurer-raw-ldacs-06 (work in progress), October 2020.
- [I-D.thubert-raw-technologies]
Thubert, P., Cavalcanti, D., Vilajosana, X., Schmitt, C., and J. Farkas, "Reliable and Available Wireless Technologies", draft-thubert-raw-technologies-05 (work in progress), May 2020.
- [ICAO18] International Civil Aviation Organization (ICAO), "L-Band Digital Aeronautical Communication System (LDACS)", International Standards and Recommended Practices Annex 10 - Aeronautical Telecommunications, Vol. III - Communication Systems , 2018.
- [IEEE802.1TSN]
IEEE standard for Information Technology, "IEEE 802.1AS-2011 - IEEE Standard for Local and Metropolitan Area Networks - Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks".
- [ieee80211-rt-tig]
IEEE, "IEEE 802.11 Real Time Applications TIG Report", Nov. 2018,
<http://www.ieee802.org/11/Reports/rtatig_update.htm>.

- [IEEE80211RTA] IEEE standard for Information Technology, "IEEE 802.11 Real Time Applications TIG Report", Nov 2018.
- [ISA100] ISA/ANSI, "ISA100, Wireless Systems for Automation", <<https://www.isa.org/isa100/>>.
- [KEAV20] T. Keaveney and C. Stewart, "Single European Sky ATM Research Joint Undertaking", 2019, <<https://www.sesarju.eu/>>.
- [ODVA] <http://www.odva.org/>, "The organization that supports network technologies built on the Common Industrial Protocol (CIP) including EtherNet/IP."
- [PLA14] Plass, S., Hermenier, R., Luecke, O., Gomez Depoorter, D., Tordjman, T., Chatterton, M., Amirfeiz, M., Scotti, S., Cheng, Y., Pillai, P., Graeupl, T., Durand, F., Murphy, K., Marriott, A., and A. Zaytsev, "Flight Trial Demonstration of Seamless Aeronautical Networking", IEEE Communications Magazine, vol. 52, no. 5 , May 2014.
- [Profinet] <http://us.profinet.com/technology/profinet/>, "PROFINET is a standard for industrial networking in automation.", <<http://us.profinet.com/technology/profinet/>>.
- [RFC7554] Watteyne, T., Ed., Palattella, M., and L. Grieco, "Using IEEE 802.15.4e Time-Slotted Channel Hopping (TSCH) in the Internet of Things (IoT): Problem Statement", RFC 7554, DOI 10.17487/RFC7554, May 2015, <<https://www.rfc-editor.org/info/rfc7554>>.
- [RFC8557] Finn, N. and P. Thubert, "Deterministic Networking Problem Statement", RFC 8557, DOI 10.17487/RFC8557, May 2019, <<https://www.rfc-editor.org/info/rfc8557>>.
- [RFC8578] Grossman, E., Ed., "Deterministic Networking Use Cases", RFC 8578, DOI 10.17487/RFC8578, May 2019, <<https://www.rfc-editor.org/info/rfc8578>>.
- [RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <<https://www.rfc-editor.org/info/rfc8655>>.

[robots] Kober, J., Glisson, M., and M. Mistry, "Playing catch and juggling with a humanoid robot.", 2012, <<https://doi.org/10.1109/HUMANOIDS.2012.6651623>>.

[square-peg] Martinez, B., Cano, C., and X. Vilajosana, "A Square Peg in a Round Hole: The Complex Path for Wireless in the Manufacturing Industry", 2019, <<https://ieeexplore.ieee.org/document/8703476>>.

Authors' Addresses

Georgios Z. Papadopoulos
IMT Atlantique
Office B00 - 114A
2 Rue de la Chataigneraie
Cesson-Sevigne - Rennes 35510
FRANCE

Phone: +33 299 12 70 04
Email: georgios.papadopoulos@imt-atlantique.fr

Pascal Thubert
Cisco Systems, Inc
Building D
45 Allee des Ormes - BP1200
MOUGINS - Sophia Antipolis 06254
FRANCE

Phone: +33 497 23 26 34
Email: pthubert@cisco.com

Fabrice Theoleyre
CNRS
ICube Lab, Pole API
300 boulevard Sebastien Brant - CS 10413
Illkirch 67400
FRANCE

Phone: +33 368 85 45 33
Email: theoleyre@unistra.fr
URI: <http://www.theoleyre.eu>

Carlos J. Bernardos
Universidad Carlos III de Madrid
Av. Universidad, 30
Leganes, Madrid 28911
Spain

Phone: +34 91624 6236
Email: cjbc@it.uc3m.es
URI: <http://www.it.uc3m.es/cjbc/>

RAW
Internet-Draft
Intended status: Informational
Expires: 27 August 2022

CJ. Bernardos, Ed.
UC3M
G.Z. Papadopoulos
IMT Atlantique
P. Thubert
Cisco
F. Theoleyre
CNRS
23 February 2022

RAW use-cases
draft-ietf-raw-use-cases-05

Abstract

The wireless medium presents significant specific challenges to achieve properties similar to those of wired deterministic networks. At the same time, a number of use-cases cannot be solved with wires and justify the extra effort of going wireless. This document presents wireless use-cases (such as aeronautical communications, amusement parks, industrial applications, pro audio and video, gaming, UAV and V2V control, edge robotics and emergency vehicles) demanding reliable and available behavior.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 27 August 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
2. Aeronautical Communications	5
2.1. Problem Statement	5
2.2. Specifics	6
2.3. Challenges	7
2.4. The Need for Wireless	8
2.5. Requirements for RAW	8
2.5.1. Non-latency critical considerations	9
3. Amusement Parks	9
3.1. use-case Description	9
3.2. Specifics	9
3.3. The Need for Wireless	10
3.4. Requirements for RAW	10
3.4.1. Non-latency critical considerations	11
4. Wireless for Industrial Applications	11
4.1. use-case Description	11
4.2. Specifics	11
4.2.1. Control Loops	11
4.2.2. Unmeasured Data	12
4.3. The Need for Wireless	12
4.4. Requirements for RAW	13
4.4.1. Non-latency critical considerations	14
5. Pro Audio and Video	14
5.1. use-case Description	14
5.2. Specifics	14
5.2.1. Uninterrupted Stream Playback	14
5.2.2. Synchronized Stream Playback	14
5.3. The Need for Wireless	15
5.4. Requirements for RAW	15
5.4.1. Non-latency critical considerations	15
6. Wireless Gaming	15
6.1. use-case Description	15
6.2. Specifics	16
6.3. The Need for Wireless	16
6.4. Requirements for RAW	16
6.4.1. Non-latency critical considerations	17

7. Unmanned Aerial Vehicles and Vehicle-to-Vehicle platooning and control	17
7.1. use-case Description	17
7.2. Specifics	18
7.3. The Need for Wireless	18
7.4. Requirements for RAW	18
7.4.1. Non-latency critical considerations	19
8. Edge Robotics control	19
8.1. use-case Description	19
8.2. Specifics	20
8.3. The Need for Wireless	20
8.4. Requirements for RAW	20
8.4.1. Non-latency critical considerations	20
9. Emergencies: Instrumented emergency vehicle	20
9.1. use-case Description	20
9.2. Specifics	21
9.3. The Need for Wireless	21
9.4. Requirements for RAW	21
9.4.1. Non-latency critical considerations	22
10. Summary	22
11. IANA Considerations	22
12. Security Considerations	22
13. Acknowledgments	22
14. References	23
14.1. Normative References	23
14.2. Informative References	23
Authors' Addresses	26

1. Introduction

Based on time, resource reservation, and policy enforcement by distributed shapers, Deterministic Networking provides the capability to carry specified unicast or multicast data streams for real-time applications with extremely low data loss rates and bounded latency, to support time-sensitive and mission-critical applications on a converged enterprise infrastructure.

Deterministic Networking in the IP world is an attempt to eliminate packet loss for a committed bandwidth while ensuring a worst case end-to-end latency, regardless of the network conditions and across technologies. By leveraging on lower (L2 and below) capabilities, L3 can exploit the use of a service layer, steering over multiple technologies, and using media independent signaling to provide high reliability, precise time delivery, and rate enforcement. Deterministic networking can be seen as a set of new Quality of Service (QoS) guarantees of worst-case delivery. IP networks become more deterministic when the effects of statistical multiplexing (jitter and collision loss) are mostly eliminated. This requires a

tight control of the physical resources to maintain the amount of traffic within the physical capabilities of the underlying technology, e.g., using time-shared resources (bandwidth and buffers) per circuit, and/or by shaping and/or scheduling the packets at every hop.

Key attributes of Deterministic Networking include:

- * time synchronization on all the nodes,
- * centralized computation of network-wide deterministic paths,
- * multi-technology path with co-channel interference minimization,
- * frame preemption and guard time mechanisms to ensure a worst-case delay, and
- * new traffic shapers within and at the edge to protect the network.

Wireless operates on a shared medium, and transmissions cannot be guaranteed to be fully deterministic due to uncontrolled interferences, including self-induced multipath fading. Reliable and Available Wireless (RAW) is an effort to provide Deterministic Networking Mechanisms on across a multi-hop path that includes a wireless physical layer. Making Wireless Reliable and Available is even more challenging than it is with wires, due to the numerous causes of loss in transmission that add up to the congestion losses and the delays caused by overbooked shared resources.

The wireless and wired media are fundamentally different at the physical level, and while the generic Problem Statement [RFC8557] for DetNet applies to the wired as well as the wireless medium, the methods to achieve RAW necessarily differ from those used to support Time-Sensitive Networking over wires, e.g., due to the wireless radio channel specifics.

So far, Open Standards for Deterministic Networking have prevalently been focused on wired media, with Audio/Video Bridging (AVB) and Time Sensitive Networking (TSN) at the IEEE and DetNet [RFC8655] at the IETF. But wires cannot be used in several cases, including mobile or rotating devices, rehabilitated industrial buildings, wearable or in-body sensory devices, vehicle automation and multiplayer gaming.

Purpose-built wireless technologies such as [ISA100], which incorporates IPv6, were developed and deployed to cope for the lack of open standards, but they yield a high cost in OPEX and CAPEX and are limited to very few industries, e.g., process control, concert instruments or racing.

This is now changing [I-D.ietf-raw-technologies]:

- * IMT-2020 has recognized Ultra-Reliable Low-Latency Communication (URLLC) as a key functionality for the upcoming 5G.
- * IEEE 802.11 has identified a set of real-applications [IEEE80211-RT-TIG] which may use the IEEE802.11 standards. They typically emphasize strict end-to-end delay requirements.
- * The IETF has produced an IPv6 stack for IEEE Std. 802.15.4 TimeSlotted Channel Hopping (TSCH) and an architecture [RFC9030] that enables RAW on a shared MAC.

Experiments have already been conducted with IEEE802.1 TSN over IEEE802.11be [IEEE80211BE]. This mode enables time synchronization, and time-aware scheduling (trigger based access mode) to support TSN flows.

This draft extends the "Deterministic Networking use-cases" document [RFC8578] and describes several additional use-cases which require "reliable/predictable and available" flows over wireless links and possibly complex multi-hop paths called Tracks. This is covered mainly by the "Wireless for Industrial Applications" use-case, as the "Cellular Radio" is mostly dedicated to the (wired) transport part of a Radio Access Network (RAN). Whereas the "Wireless for Industrial Applications" use-case certainly covers an area of interest for RAW, it is limited to 6TiSCH, and thus its scope is narrower than the use-cases described next in this document.

2. Aeronautical Communications

Aircraft are currently connected to ATC (Air-Traffic Control) and AOC (Airline Operational Control) via voice and data communications systems through all phases of a flight. Within the airport terminal, connectivity is focused on high bandwidth communications while during en-route high reliability, robustness and range are the focus.

2.1. Problem Statement

Up to 2020, civil air traffic has been growing constantly at a compound rate of 5.8% per year [ACI19] and despite the severe impact of the COVID-19 pandemic, air traffic growth is expected to resume very quickly in post-pandemic times [IAT20] [IAC20]. Thus, legacy systems in air traffic management (ATM) are likely to reach their capacity limits and the need for new aeronautical communication technologies becomes apparent. Especially problematic is the saturation of VHF band in high density areas in Europe, the US, and Asia [KEAV20] [FAA20] calling for suitable new digital approaches

such as AeroMACS for airport communications, SatCOM for remote domains, and LDACS as long-range terrestrial aeronautical communications system. Making the frequency spectrum's usage more efficient a transition from analog voice to digital data communication [PLA14] is necessary to cope with the expected growth of civil aviation and its supporting infrastructure. A promising candidate for long range terrestrial communications, already in the process of being standardized in the International Civil Aviation Organization (ICAO), is the L-band Digital Aeronautical Communications System (LDACS) [ICAO18] [I-D.ietf-raw-ldacs].

2.2. Specifics

During the creation process of new communications system, analog voice is replaced by digital data communication. This sets a paradigm shift from analog to digital wireless communications and supports the related trend towards increased autonomous data processing that the Future Communications Infrastructure (FCI) in civil aviation must provide. The FCI is depicted in Figure 1:




Figure 1: The Future Communication Infrastructure (FCI): AeroMACS for APT/TMA domain, LDACS A/G for TMA/ENR domain, LDACS A/G for ENR/GRP domain, SatCOM for GRP domain communications

2.3. Challenges

This paradigm change brings a lot of new challenges:

- * **Efficiency:** It is necessary to keep latency, time and data overhead (routing, security) of new aeronautical datalinks at a minimum.

- * **Modularity:** Systems in avionics usually operate up to 30 years, thus solutions must be modular, easily adaptable and updatable.
- * **Interoperability:** All 192 members of the international Civil Aviation Organization (ICAO) must be able to use these solutions.
- * **Dynamicity:** the communication infrastructure needs to accommodate mobile devices (airplanes) that move extremely fast.

2.4. The Need for Wireless

In a high mobility environment such as aviation, the envisioned solutions to provide worldwide coverage of data connections with in-flight aircraft require a multi-system, multi-link, multi-hop approach. Thus air, ground and space-based datalink providing technologies will have to operate seamlessly together to cope with the increasing needs of data exchange between aircraft, air traffic controller, airport infrastructure, airlines, air network service providers (ANSPs) and so forth. Thus, making use of wireless technologies is a must in tackling this enormous need for a worldwide digital aeronautical datalink infrastructure.

2.5. Requirements for RAW

Different safety levels need to be supported, from extremely safety critical ones requiring low latency, such as a WAKE warning - a warning that two aircraft come dangerously close to each other - and high resiliency, to less safety critical ones requiring low-medium latency for services such as WXGRAPH - graphical weather data.

Overhead needs to be kept at a minimum since aeronautical data links provide comparatively small data rates in the order of kbit/s.

Policy needs to be supported when selecting data links. The focus of RAW here should be on the selectors, responsible for the track a packet takes to reach its end destination. This would minimize the amount of routing information that must travel inside the network because of precomputed routing tables with the selector being responsible for choosing the most appropriate option according to policy and safety.

2.5.1. Non-latency critical considerations

Achieving low latency is a requirement for aeronautics communications, though the expected latency is not extremely low and what it is important is to keep the overall latency bounded under a certain threshold. This use-case is not latency-critical from that view point. On the other hand, given the controlled environment, end-to-end mechanisms can be applied to guarantee bounded latency where needed.

3. Amusement Parks

3.1. use-case Description

The digitalization of Amusement Parks is expected to decrease significantly the cost for maintaining the attractions. Such deployment is a mix between industrial automation (i.e., Smart Factories) and multimedia entertainment applications.

Attractions may rely on a large set of sensors and actuators, which react in real time. Typical applications comprise:

- * Emergency: safety has to be preserved, and must stop the attraction when a failure is detected.
- * Video: augmented and virtual realities are integrated in the attraction. Wearable mobile devices (e.g., glasses, virtual reality headset) need to offload one part of the processing tasks.
- * Real-time interactions: visitors may interact with an attraction, like in a real-time video game. The visitors may virtually interact with their environment, triggering actions in the real world (through actuators) [KOB12].
- * Geolocation: visitors are tracked with a personal wireless tag so that their user experience is improved.
- * Predictive maintenance: statistics are collected to predict the future failures, or to compute later more complex statistics about the attraction's usage, the downtime or its popularity for example.

3.2. Specifics

Amusement parks comprise a variable number of attractions, mostly outdoor, over a large geographical area. The IT infrastructure is typically multi-scale:

- * Local area: the sensors and actuators controlling the attractions are co-located. Control loops trigger only local traffic, with a small end-to-end delay, typically inferior to 10 ms, like classical industrial systems [IEEE80211-RT-TIG].
- * Wearable mobile devices are free to move in the park. They exchange traffic locally (identification, personalization, multimedia) or globally (billing, child tracking).
- * Computationally intensive applications offload some tasks. Edge computing seems an efficient way to implement real-time applications with offloading. Some non-time-critical tasks may rather use the cloud (predictive maintenance, marketing).

3.3. The Need for Wireless

Amusement parks cover large areas, and a global interconnection would require a huge length of cables. Wireless also increases the reconfigurability, enabling to update an attraction at a lower cost. The frequent renewal helps to increase the customer loyalty.

Some parts of the attraction are mobile, like trucks of a roller-coaster or robots. Since cables are prone to frequent failures in this situation, wireless transmissions are recommended.

Wearable devices are extensively used for a user experience personalization. They typically need to support wireless transmissions. Personal tags may help to reduce the operating costs [DISNEY15] and to increase the number of charged services provided to the audience (e.g., VIP tickets or interactivity). Some applications rely on more sophisticated wearable devices such as digital glasses or Virtual Reality (VR) headsets for an immersive experience.

3.4. Requirements for RAW

The network infrastructure must support heterogeneous traffic, with very different critical requirements. Thus, flow isolation must be provided.

The transmissions must be scheduled appropriately even in presence of mobile devices. While the [RFC9030] already proposes an architecture for synchronized, IEEE Std. 802.15.4 Time-Slotted Channel Hopping (TSCH) networks, the industry requires a multi-technology solution, able to guarantee end-to-end requirements across heterogeneous technologies, with strict SLA requirements.

Nowadays, long-range wireless transmissions are used mostly for best-effort traffic. On the contrary, [IEEE802.1TSN] is used for critical flows using Ethernet devices. However, IP enabled technology is required to interconnect large areas, independent of the PHY and MAC layers.

It is expected that several different technologies (long vs. short range) are deployed, which have to cohabit in the same area. Thus, we need to provide layer-3 mechanisms able to exploit multiple co-interfering technologies.

3.4.1. Non-latency critical considerations

While some of the applications in this use-case involve control loops (e.g., sensors and actuators) that require bounded latencies below 10 ms, that can therefore be considered latency critical, there are other applications as well that mostly demand reliability (e.g., safety related, or maintenance).

4. Wireless for Industrial Applications

4.1. use-case Description

A major use-case for networking in Industrial environments is the control networks where periodic control loops operate between a collection of sensors that measure a physical property such as the temperature of a fluid, a Programmable Logic Controller (PLC) that decides an action such as warm up the mix, and actuators that perform the required action, such as the injection of power in a resistor.

4.2. Specifics

4.2.1. Control Loops

Process Control designates continuous processing operations, like heating Oil in a refinery or mixing drinking soda. Control loops in the Process Control industry operate at a very low rate, typically four times per second. Factory Automation, on the other hand, deals with discrete goods such as individual automobile parts, and requires faster loops, in the order of milliseconds. Motion control that monitors dynamic activities may require even faster rates in the order of and below the millisecond. Finally, some industries exhibit hybrid behaviors, like canned soup that will start as a process industry while mixing the food and then operate as a discrete manufacturing when putting the final product in cans and shipping them.

In all those cases, a packet must flow reliably between the sensor and the PLC, be processed by the PLC, and sent to the actuator within the control loop period. In some particular use-cases that inherit from analog operations, jitter might also alter the operation of the control loop. A rare packet loss is usually admissible, but typically 4 losses in a row will cause an emergency halt of the production and incur a high cost for the manufacturer.

Additional details and use-cases related to Industrial applications and their RAW requirements can be found in [I-D.ietf-raw-industrial-requirements].

4.2.2. Unmeasured Data

A secondary use-case deals with monitoring and diagnostics. This so-called unmeasured data is essential to improve the performances of a production line, e.g., by optimizing real-time processing or maintenance windows using Machine Learning predictions. For the lack of wireless technologies, some specific industries such as Oil and Gas have been using serial cables, literally by the millions, to perform their process optimization over the previous decades. But few industries would afford the associated cost and the Holy Grail of the Industrial Internet of Things is to provide the same benefits to all industries, including SmartGrid, Transportation, Building, Commercial and Medical. This requires a cheap, available and scalable IP-based access technology.

Inside the factory, wires may already be available to operate the Control Network. But unmeasured data are not welcome in that network for several reasons. On the one hand it is rich and asynchronous, meaning that it may influence the deterministic nature of the control operations and impact the production. On the other hand, this information must be reported to the carpeted floor over IP, which means the potential for a security breach via the interconnection of the Operational Technology (OT) network with the Internet technology (IT) network and possibly enable a rogue access.

4.3. The Need for Wireless

Ethernet cables used on a robot arm are prone to breakage after a few thousands of flexions, a lot faster than a power cable that is wider in diameter, and more resilient. In general, wired networking and mobile parts are not a good match, mostly in the case of fast and recurrent activities, as well as rotation.

When refurbishing older premises that were built before the Internet age, power is usually available everywhere, but data is not. It is often impractical, time consuming and expensive to deploy an Ethernet fabric across walls and between buildings. Deploying a wire may take months and cost tens of thousands of US Dollars.

Even when wiring exists, like in the case of an existing control network, asynchronous IP packets such as diagnostics may not be welcome for operational and security reasons. For those packets, the option to create a parallel wireless network offers a credible solution that can scale with the many sensors and actuators that equip every robot, every valve and fan that are deployed on the factory floor. It may also help detect and prevent a failure that could impact the production, like the degradation (vibration) of a cooling fan on the ceiling. IEEE Std. 802.15.4 Time-Slotted Channel Hopping (TSCH) [RFC7554] is a promising technology for that purpose, mostly if the scheduled operations enable to use the same network by asynchronous and deterministic flows in parallel.

4.4. Requirements for RAW

As stated by the "Deterministic Networking Problem Statement" [RFC8557], a Deterministic Network is backwards compatible with (capable of transporting) statistically multiplexed traffic while preserving the properties of the accepted deterministic flows. While the 6TiSCH Architecture [RFC9030] serves that requirement, the work at 6TiSCH was focused on best-effort IPv6 packet flows. RAW should be able to lock so-called hard cells for use by a centralized scheduler, and leverage time and spatial diversity over a graph of end-to-end paths called a Track that is based on those cells.

Over the course of the recent years, major Industrial Protocols (e.g., [ODVA] with EtherNet/IP [EIP] and [PROFINET]) have been migrating towards Ethernet and IP. In order to unleash the full power of the IP hourglass model, it should be possible to deploy any application over any network that has the physical capacity to transport the industrial flow, regardless of the MAC/PHY technology, wired or wireless, and across technologies. RAW mechanisms should be able to setup a Track over a wireless access segment and a wired or wireless backbone to report both sensor data and critical monitoring within a bounded latency and maintain the high reliability of the flows over time. It is also important to ensure that RAW solutions are interoperable with existing wireless solutions in place, and with legacy equipment which capabilities can be extended using retrofitting. Maintainability, as a broader concept than reliability is also important in industrial scenarios [MAR19].

4.4.1. Non-latency critical considerations

Monitoring and diagnostics applications do not require latency critical communications, but demand reliable and scalable communications. On the other hand, process control applications involve control loops that require a bounded latency, thus are latency critical, but can be managed end-to-end, and therefore DetNet mechanisms can be applied in conjunction with RAW mechanisms.

5. Pro Audio and Video

5.1. use-case Description

Many devices support audio and video streaming by employing 802.11 wireless LAN. Some of these applications require low latency capability. For instance, when the application provides interactive play, or when the audio plays in real time - meaning live for public addresses in train stations or in theme parks.

The professional audio and video industry ("ProAV") includes:

- * Virtual Reality / Augmented Reality (VR/AR)
- * Production and post-production systems such as CD and Blue-Ray disk mastering.
- * Public address, media and emergency systems at large venues (e.g., airports, train stations, stadiums, and theme parks).

5.2. Specifics

5.2.1. Uninterrupted Stream Playback

Considering the uninterrupted audio or video stream, a potential packet loss during the transmission of audio or video flows cannot be tackled by re-trying the transmission, as it is done with file transfer, because by the time the packet lost has been identified it is too late to proceed with packet re-transmission. Buffering might be employed to provide a certain delay which will allow for one or more re-transmissions, however such approach is not efficient in application where delays are not acceptable.

5.2.2. Synchronized Stream Playback

In the context of ProAV, latency is the time between the transmitted signal over a stream and its reception. Thus, for sound to remain synchronized to the movement in the video, the latency of both the audio and video streams must be bounded and consistent.

5.3. The Need for Wireless

The devices need the wireless communication to support video streaming via IEEE 802.11 wireless LAN for instance. Wireless communications provide huge advantages in terms of simpler deployments in many scenarios, where the use of a wired alternative would not be feasible. Similarly, in live events, mobility support makes wireless communications the only viable approach.

Deployed announcement speakers, for instance along the platforms of the train stations, need the wireless communication to forward the audio traffic in real time.

5.4. Requirements for RAW

The network infrastructure needs to support heterogeneous types of traffic (including QoS).

Content delivery with bounded (lowest possible) latency.

The deployed network topology should allow for multipath. This will enable for multiple streams to have different (and multiple) paths (tracks) through the network to support redundancy.

5.4.1. Non-latency critical considerations

For synchronized streaming, latency must be bounded, and therefore, depending on the actual requirements, this can be considered as latency critical. However, the most critical requirement of this use-case is reliability, by the network providing redundancy. Note that in many cases, wireless is only present in the access, where RAW mechanisms could be applied, but other wired segments are also involved (like the Internet), and therefore latency cannot be guaranteed.

6. Wireless Gaming

6.1. use-case Description

The gaming industry includes [IEEE80211RTA] real-time mobile gaming, wireless console gaming and cloud gaming. For RAW, wireless console gaming is the most relevant one. We next summarize the three:

- * Real-time Mobile Gaming: Different from traditional games, real time mobile gaming is very sensitive to network latency and stability. The mobile game can connect multiple players together in a single game session and exchange data messages between game server and connected players. Real-time means the feedback should

present on screen as users operate in game. For good game experience, the end-to-end (E2E) latency plus game servers processing time must be the same for all players and should not be noticeable as the game is played.

- * **Wireless Console Gaming:** Playing online on a console has 2 types of internet connectivity, which is either wired or Wi-Fi. Most of the gaming consoles today support Wi-Fi 5. But Wi-Fi has an especially bad reputation among the gaming community. The main reasons are high latency, lag spikes, and jitter.
- * **Cloud Gaming:** The cloud gaming requires low latency capability as the user commands in a game session need to be sent back to the cloud server, the cloud server would update game context depending on the received commands, and the cloud server would render the picture/video to be displayed at user devices and stream the picture/video content to the user devices. User devices might very likely be connected wirelessly.

6.2. Specifics

While a lot of details can be found on [IEEE80211RTA], we next summarize the main requirements in terms of latency, jitter and packet loss:

- * Intra BSS latency is less than 5 ms.
- * Jitter variance is less than 2 ms.
- * Packet loss is less than 0.1 percent.

6.3. The Need for Wireless

Gaming is evolving towards wireless, as players demand being able to play anywhere, and the game requires a more immersive experience including body movements. Besides, the industry is changing towards playing from mobile phones, which are inherently connected via wireless technologies. Wireless controllers are the rule in modern gaming, with increasingly sophisticated interactions (e.g., haptic feedback, augmented reality).

6.4. Requirements for RAW

- * **Time sensitive networking extensions:** extensions, such as time-aware shaping and redundancy (FRE) can be explored to address congestion and reliability problems present in wireless networks.

- * Priority tagging (Stream identification): one basic requirement to provide better QoS for time-sensitive traffic is the capability to identify and differentiate time-sensitive packets from other (like best-effort) traffic.
- * Time-aware shaping: this capability (defined in IEEE 802.1Qbv) consists of gates to control the opening/closing of queues that share a common egress port within an Ethernet switch. A scheduler defines the times when each queue opens or close, therefore eliminating congestion and ensuring that frames are delivered within the expected latency bounds. Note thought, that while this requirement needs to be signalled by RAW mechanisms, it would be actually served by the lower layer.
- * Dual/multiple link: due to the competitions and interference are common and hardly in control under wireless network, to improve the latency stability, dual/multiple link proposal is brought up to address this issue.
- * Admission Control: congestion is a major cause of high/variable latency and it is well known that if the traffic load exceeds the capability of the link, QoS will be degraded. QoS degradation maybe acceptable for many applications today, however emerging time-sensitive applications are highly susceptible to increased latency and jitter. To better control QoS, it is important to control access to the network resources.

6.4.1. Non-latency critical considerations

Depending on the actual scenario, and on use of Internet to interconnect different users, the communication's requirements of this use-case might be considered as latency critical due to the need of bounded latency. But note that in most of these scenarios, part of the communication path is not wireless and DetNet mechanisms cannot be applied easily (e.g., when the public Internet is involved), and therefore in these cases, reliability is the critical requirement.

7. Unmanned Aerial Vehicles and Vehicle-to-Vehicle platooning and control

7.1. use-case Description

Unmanned Aerial Vehicles (UAVs) are becoming very popular for many different applications, including military and civil use-cases. The term drone is commonly used to refer to a UAV.

UAVs can be used to perform aerial surveillance activities, traffic monitoring (i.e., the Spanish traffic control has recently introduced a fleet of drones for quicker reactions upon traffic congestion related events), support of emergency situations, and even transportation of small goods (e.g., medicine in rural areas).

Similarly to UAVs, other time of vehicles (such as cars) can also travel in platoons. Most of the considerations made for UAVs in this section apply to vehicle-to-vehicle (V2V) scenarios.

UAVs/vehicles typically have various forms of wireless connectivity:

- * Cellular: for communication with the control center, for remote maneuvering as well as monitoring of the drone;
- * IEEE 802.11: for inter-drone communications (i.e., platooning) and providing connectivity to other devices (i.e., acting as Access Point).

Note that autonomous cars share many of the characteristics of the aforementioned UAV case, and therefore it is of interest for RAW.

7.2. Specifics

Some of the use-cases/tasks involving UAVs require coordination among UAVs. Others involve complex compute tasks that might not be performed using the limited computing resources that a drone typically has. These two aspects require continuous connectivity with the control center and among UAVs.

Remote maneuvering of a drone might be performed over a cellular network in some cases, however, there are situations that need very low latency and deterministic behavior of the connectivity. Examples involve platooning of drones or share of computing resources among drones (like, a drone offload some function to a neighboring drone).

7.3. The Need for Wireless

UAVs cannot be connected through any type of wired media, so it is obvious that wireless is needed.

7.4. Requirements for RAW

The network infrastructure is composed by the UAVs themselves, requiring self-configuration capabilities.

Heterogeneous types of traffic need to be supported, from extremely critical ones requiring ultra-low latency and high resiliency, to traffic requiring low-medium latency.

When a given service is decomposed into functions -- hosted at different UAVs -- chained, each link connecting two given functions would have a well-defined set of requirements (e.g., latency, bandwidth and jitter) that must be met.

7.4.1. Non-latency critical considerations

Today's solutions keep local the processing operations that are critical and would demand an ultra-low latency communication to be offloaded. Therefore, in this use-case, the critical requirement is reliability, and only for some platooning and inter-drone communications latency is critical.

8. Edge Robotics control

8.1. use-case Description

The Edge Robotics scenario consists of several robots, deployed in a given area (like a shopping mall), inter-connected via an access network to a network's edge device or a data center. The robots are connected to the edge so complex computational activities are not executed locally at the robots but offloaded to the edge. This brings additional flexibility in the type of tasks that the robots do, as well as reducing the costs of robot manufacturing (due to their lower complexity), and enabling complex tasks involving coordination among robots (that can be more easily performed if robots are centrally controlled).

Simple examples of the use of multiples robots are cleaning, video surveillance, search and rescue operations, and delivering of goods from warehouses to shops. Multiple robots are simultaneously instructed to perform individual tasks by moving the robotic intelligence from the robots to the network's edge (like a data center). That enables easy synchronization, scalable solution, and on-demand option to create flexible fleet of robots.

Robots would have various forms of wireless connectivity:

- * IEEE 802.11: for connection to the edge and also inter-robot communications (i.e., for coordinated actions).
- * Cellular: as an additional communication link to the edge, though primarily as backup, since ultra-low latency is needed.

8.2. Specifics

Some of the use-cases/tasks involving robots might benefit from decomposition of a service in small functions that are distributed and chained among robots and the edge. These require continuous connectivity with the control center and among drones.

Robot control is an activity requiring very low latency between the robot and the location where the control intelligence resides (which might be the edge or another robot).

8.3. The Need for Wireless

Deploying robots in scenarios such as shopping malls for the applications mentioned cannot be done via wired connectivity.

8.4. Requirements for RAW

The network infrastructure needs to support heterogeneous types of traffic, from robot control to video streaming.

When a given service is decomposed into functions -- hosted at different robots -- chained, each link connecting two given functions would have a well-defined set of requirements (latency, bandwidth and jitter) that must be met.

8.4.1. Non-latency critical considerations

This use-case might combine multiple communication flows, with some of them being latency critical (like those related to robot control tasks). Note that there are still many communication flows (like some offloading tasks) that only demand reliability and availability.

9. Emergencies: Instrumented emergency vehicle

9.1. use-case Description

An instrumented ambulance would be one that has a LAN to which are connected these end systems such as:

- * vital signs sensors attached to the casualty in the ambulance. Relay medical data to hospital emergency room,
- * radio-navigation sensor to relay position data to various destinations including dispatcher,
- * voice communication for ambulance attendant (like to consult with ER doctor), and

- * voice communication between driver and dispatcher.

The LAN needs to be routed through radio-WANs to complete the inter-network linkage.

9.2. Specifics

What we have today is multiple communications systems to reach the vehicle via:

- * A dispatching system,
- * a cellphone for the attendant,
- * a special purpose telemetering system for medical data,
- * etc.

This redundancy of systems, because of its stove-piping, does not contribute to availability.

Most of the scenarios involving the use of an instrumented ambulance are composed of many different flows, each of them with slightly different requirements in terms of reliability and latency. Destinations might be either at the ambulance itself (local traffic), at a near edge cloud or at the general Internet/cloud.

9.3. The Need for Wireless

Local traffic between the first responders/ambulance staff and the ambulance equipment cannot be done via wired connectivity as the responders perform initial treatment outside of the ambulance. The communications from the ambulance to external services must be wireless as well.

9.4. Requirements for RAW

We can derive some pertinent requirements from this scenario:

- * High availability of the inter-network is required.
- * The inter-network needs to operate in damaged state (e.g. during an earthquake aftermath, heavy weather, wildfire, etc.). In addition to continuity of operations, rapid restore is a needed characteristic.

- * E2E security, both authenticity and confidentiality, is required of traffic. All data needs to be authenticated; some like medical needs to be confidential.
- * The radio-WAN has characteristics similar to cellphone -- the vehicle will travel from one radio footprint to another.

9.4.1. Non-latency critical considerations

In this case, all applications identified do not require latency critical communication, but do need of high reliability and availability.

10. Summary

This document enumerates several use-cases and applications that need RAW technologies, focusing on the requirements from reliability, availability and latency. Whereas some use-cases are latency-critical, there are also several applications that are non-latency critical, but that do pose strict reliability and availability requirements. Future revisions of this document will include specific text devoted to highlight this non-latency critical requirements.

11. IANA Considerations

This document has no IANA actions.

12. Security Considerations

This document covers several representative applications and network scenarios that are expected to make use of RAW technologies. Each of the potential RAW use-cases will have security considerations from both the use-specific perspective and the RAW technology perspective. [RFC9055] provides a comprehensive discussion of security considerations in the context of Deterministic Networking, which are generally applicable also to RAW.

13. Acknowledgments

Nils Maeurer, Thomas Graeupl and Corinna Schmitt have contributed significantly to this document, providing input for the Aeronautical communications section. Rex Buddenberg has also contributed to the document, providing input to the Emergency: instrumented emergency vehicle section.

The authors would like to thank Toerless Eckert, Xavi Vilajosana Guillen, Rute Sofia and Corinna Schmitt for their valuable comments on previous versions of this document.

The work of Carlos J. Bernardos in this draft has been partially supported by the H2020 5Growth (Grant 856709) and 5G-DIVE projects (Grant 859881), and UNICO 5G I+D 6G-DATADRIVEN-04 project.

14. References

14.1. Normative References

[I-D.ietf-raw-technologies]
Thubert, P., Cavalcanti, D., Vilajosana, X., Schmitt, C., and J. Farkas, "Reliable and Available Wireless Technologies", Work in Progress, Internet-Draft, draft-ietf-raw-technologies-05, 2 February 2022, <<https://www.ietf.org/archive/id/draft-ietf-raw-technologies-05.txt>>.

14.2. Informative References

- [ACI19] Airports Council International (ACI), "Annual World Airport Traffic Report 2019", November 2019, <<https://store.aci.aero/product/annual-world-airport-traffic-report-2019/>>.
- [DISNEY15] Wired, "Disney's \$1 Billion Bet on a Magical Wristband", March 2015, <<https://www.wired.com/2015/03/disney-magicband/>>.
- [EIP] <http://www.odva.org/>, "EtherNet/IP provides users with the network tools to deploy standard Ethernet technology (IEEE 802.3 combined with the TCP/IP Suite) for industrial automation applications while enabling Internet and enterprise connectivity data anytime, anywhere.", <http://www.odva.org/Portals/0/Library/Publications_Numbered/PUB00138R3_CIP_Adv_Tech_Series_EtherNetIP.pdf>.
- [FAA20] U.S. Department of Transportation Federal Aviation Administration (FAA), "Next Generation Air Transportation System", 2019, <<https://www.faa.gov/nextgen/>>.
- [I-D.ietf-raw-industrial-requirements]
Sofia, R. C., Kovatsch, M., and P. M. Mendes, "Requirements for Reliable Wireless Industrial Services", Work in Progress, Internet-Draft, draft-ietf-raw-

industrial-requirements-00, 10 December 2021,
<<https://www.ietf.org/archive/id/draft-ietf-raw-industrial-requirements-00.txt>>.

[I-D.ietf-raw-ldacs]

Maeurer, N., Graeupl, T., and C. Schmitt, "L-band Digital Aeronautical Communications System (LDACS)", Work in Progress, Internet-Draft, draft-ietf-raw-ldacs-09, 22 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-raw-ldacs-09.txt>>.

[IAC20] Iacus, S.M., Natale, F., Santamaria, C., Spyrtos, S., and V. Michele, "Estimating and projecting air passenger traffic during the COVID-19 coronavirus outbreak and its socio-economic impact", Safety Science 129 (2020) 104791 , 2020.

[IAT20] International Air Transport Association (IATA), "Economic Performance of the Airline Industry", November 2020, <<https://www.iata.org/en/iata-repository/publications/economic-reports/airline-industry-economic-performance---november-2020---report/>>.

[ICAO18] International Civil Aviation Organization (ICAO), "L-Band Digital Aeronautical Communication System (LDACS)", International Standards and Recommended Practices Annex 10 - Aeronautical Telecommunications, Vol. III - Communication Systems , 2018.

[IEEE802.1TSN]

IEEE standard for Information Technology, "IEEE 802.1AS-2011 - IEEE Standard for Local and Metropolitan Area Networks - Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks".

[IEEE80211-RT-TIG]

IEEE, "IEEE 802.11 Real Time Applications TIG Report", November 2018, <http://www.ieee802.org/11/Reports/rtatig_update.htm>.

[IEEE80211BE]

Cavalcanti, D. and G. Venkatesan, "802.1 TSN over 802.11 with updates from developments in 802.11be", IEEE plenary meeting , November 2020, <<https://www.ieee802.org/1/files/public/docs2020/new-Cavalcanti-802-1TSN-over-802-11-1120-v02.pdf>>.

- [IEEE80211RTA] IEEE standard for Information Technology, "IEEE 802.11 Real Time Applications TIG Report", November 2018.
- [ISA100] ISA/ANSI, "ISA100, Wireless Systems for Automation", <<https://www.isa.org/isa100/>>.
- [KEAV20] T. Keaveney and C. Stewart, "Single European Sky ATM Research Joint Undertaking", 2019, <<https://www.sesarju.eu/>>.
- [KOB12] Kober, J., Glisson, M., and M. Mistry, "Playing catch and juggling with a humanoid robot.", 2012, <<https://doi.org/10.1109/HUMANOIDS.2012.6651623>>.
- [MAR19] Martinez, B., Cano, C., and X. Vilajosana, "A Square Peg in a Round Hole: The Complex Path for Wireless in the Manufacturing Industry", 2019, <<https://ieeexplore.ieee.org/document/8703476>>.
- [ODVA] <http://www.odva.org/>, "The organization that supports network technologies built on the Common Industrial Protocol (CIP) including EtherNet/IP.".
- [PLA14] Plass, S., Hermenier, R., Luecke, O., Gomez Depoorter, D., Tordjman, T., Chatterton, M., Amirfeiz, M., Scotti, S., Cheng, Y.J., Pillai, P., Graeupl, T., Durand, F., Murphy, K., Marriott, A., and A. Zaytsev, "Flight Trial Demonstration of Seamless Aeronautical Networking", IEEE Communications Magazine, vol. 52, no. 5 , May 2014.
- [PROFINET] <http://us.profinet.com/technology/profinet/>, "PROFINET is a standard for industrial networking in automation.", <<http://us.profinet.com/technology/profinet/>>.
- [RFC7554] Watteyne, T., Ed., Palattella, M., and L. Grieco, "Using IEEE 802.15.4e Time-Slotted Channel Hopping (TSCH) in the Internet of Things (IoT): Problem Statement", RFC 7554, DOI 10.17487/RFC7554, May 2015, <<https://www.rfc-editor.org/info/rfc7554>>.
- [RFC8557] Finn, N. and P. Thubert, "Deterministic Networking Problem Statement", RFC 8557, DOI 10.17487/RFC8557, May 2019, <<https://www.rfc-editor.org/info/rfc8557>>.
- [RFC8578] Grossman, E., Ed., "Deterministic Networking Use Cases", RFC 8578, DOI 10.17487/RFC8578, May 2019, <<https://www.rfc-editor.org/info/rfc8578>>.

- [RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas,
"Deterministic Networking Architecture", RFC 8655,
DOI 10.17487/RFC8655, October 2019,
<<https://www.rfc-editor.org/info/rfc8655>>.
- [RFC9030] Thubert, P., Ed., "An Architecture for IPv6 over the Time-
Slotted Channel Hopping Mode of IEEE 802.15.4 (6TiSCH)",
RFC 9030, DOI 10.17487/RFC9030, May 2021,
<<https://www.rfc-editor.org/info/rfc9030>>.
- [RFC9055] Grossman, E., Ed., Mizrahi, T., and A. Hacker,
"Deterministic Networking (DetNet) Security
Considerations", RFC 9055, DOI 10.17487/RFC9055, June
2021, <<https://www.rfc-editor.org/info/rfc9055>>.

Authors' Addresses

Carlos J. Bernardos (editor)
Universidad Carlos III de Madrid
Av. Universidad, 30
28911 Leganes, Madrid
Spain
Phone: +34 91624 6236
Email: cjbc@it.uc3m.es
URI: <http://www.it.uc3m.es/cjbc/>

Georgios Z. Papadopoulos
IMT Atlantique
Office B00 - 114A
2 Rue de la Chataigneraie
35510 Cesson-Sevigne - Rennes
France
Phone: +33 299 12 70 04
Email: georgios.papadopoulos@imt-atlantique.fr

Pascal Thubert
Cisco Systems, Inc
Building D
45 Allée des Ormes - BP1200
06254 MOUGINS - Sophia Antipolis
France
Phone: +33 497 23 26 34
Email: pthubert@cisco.com

Fabrice Theoleyre
CNRS
ICube Lab, Pole API
300 boulevard Sebastien Brant - CS 10413
67400 Illkirch
France
Phone: +33 368 85 45 33
Email: fabrice.theoleyre@cnrs.fr
URI: <https://fabrice.theoleyre.cnrs.fr/>

RAW
Internet-Draft
Intended status: Informational
Expires: 7 January 2021

P. Thubert, Ed.
Cisco Systems
G.Z. Papadopoulos
IMT Atlantique
R. Buddenberg
6 July 2020

Reliable and Available Wireless Architecture/Framework
draft-pthubert-raw-architecture-04

Abstract

Due to uncontrolled interferences, including the self-induced multipath fading, deterministic networking can only be approached on wireless links. The radio conditions may change -way- faster than a centralized routing can adapt and reprogram, in particular when the controller is distant and connectivity is slow and limited. RAW separates the routing time scale at which a complex path is recomputed from the forwarding time scale at which the forwarding decision is taken for an individual packet. RAW operates at the forwarding time scale. The RAW problem is to decide, within the redundant solutions that are proposed by the routing, which will be used for each individual packet to provide a DetNet service while minimizing the waste of resources.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 7 January 2021.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Terminology	5
3. Related Work at The IETF	6
4. Use Cases and Requirements Served	6
4.1. Radio Access Protection	7
4.2. End-to-End Protection in a Wireless Mesh	8
5. RAW Considerations	8
5.1. Reliability and Availability	8
5.1.1. High Availability Engineering Principles	8
5.1.2. Applying Reliability Concepts to Networking	11
5.1.3. Reliability in the Context of RAW	11
5.2. RAW Scope and Prerequisites	13
5.3. Routing Time Scale vs. Forwarding Time Scale	14
6. RAW Architecture Elements	15
6.1. Wireless Tracks	15
6.2. PAREO Functions	16
6.2.1. Packet Replication	17
6.2.2. Packet Elimination	18
6.2.3. Promiscuous Overhearing	18
6.2.4. Constructive Interference	18
7. RAW Architecture	19
7.1. PCE vs. PSE	20
7.2. RAW OAM	22
7.3. Flow Identification vs. Path Identification	23
7.4. Source-Routed vs. Distributed Forwarding Decision	25
7.5. Encapsulation and Decapsulation	26
8. Security Considerations	26
8.1. Forced Access	26
9. IANA Considerations	26
10. Contributors	26
11. Acknowledgments	27
12. References	27
12.1. Normative References	27
12.2. Informative References	28
Authors' Addresses	30

1. Introduction

Bringing determinism in a packet network means eliminating the statistical effects of multiplexing that result in probabilistic jitter and loss. This can be approached with a tight control of the physical resources to maintain the amount of traffic within a budgetted volume of data per unit of time that fits the physical capabilities of the underlying technology, and the use of time-shared resources (bandwidth and buffers) per circuit, and/or by shaping and/or scheduling the packets at every hop.

Wireless networks operate on a shared medium where uncontrolled interference, including the self-induced multipath fading, cause random transmission losses and add new dimensions to the statistical effects that affect the delivery. Scheduling can alleviate those effects by enabling diverse transmissions in the spatial, time, code, and frequency domains. Reliable and Available Wireless (RAW) leverages scheduling and all possible forms of diversity to defeat the possible causes of transmission loss while preserving energy and optimizing the use of the shared spectrum.

Deterministic Networking is an attempt to emulate the properties of a serial link over a switched fabric, by providing a bounded latency and eliminating congestion loss, even when co-existing with best-effort traffic. This innovation was introduced on wired networks with IEEE 802.1 TSN (for Ethernet LANs) and IETF DetNet. It is getting traction in various industries including professional A/V, manufacturing, online gaming, and smartgrid automation, enabling cost and performance optimizations (e.g., vs. loads of P2P cables).

The wireless and wired media are fundamentally different at the physical level, and while the generic "Deterministic Networking Problem Statement" [RFC8557] applies to both the wired and the wireless media, the methods to achieve RAW must extend those used to support time-sensitive networking over wires, as a RAW solution has to address less consistent transmissions, energy conservation and shared spectrum efficiency.

Uncontrolled interference and transmission obstacles may impede the wireless transmission, causing rapid variations of the throughput and packet delivery ratio (PDR) of the link. This uncertainty limits the volume and/or duration of traffic that can be safely transmitted on the same link while conforming to a RAW Service Level Agreement (SLA). Techniques such as beamforming with Multi-User MIMO can only alleviate some of those issues, and the term deterministic is usually not associated with a short range radio link, in particular one operated in the ISM band.

This increased complexity explains why the development of deterministic wireless technologies has been lagging behind the similar efforts for wired systems, both at the IEEE and the IETF. But recent progress on scheduled radios such as TSCH and OFDMA indicates that wireless is finally catching up at the lower layers. Sitting at the layer above, RAW takes up the challenge of providing highly available and reliable end-to-end performances in a network with scheduled wireless segments.

RAW provides DetNet elements that are specialized for short range radios. From this inheritance, RAW stays agnostic to the radio layer underneath though the capability to schedule transmissions is assumed. How the PHY is programmed to do so, and whether the radio is single-hop or meshed, are unknown at the IP layer and not part of the RAW abstraction.

The "Deterministic Networking Architecture" [RFC8655] is composed of three planes: the Application (User) Plane, the Controller Plane, and the Network Plane. The RAW Architecture extends the DetNet Network Plane, to accommodate one or multiple hops of homogeneous or heterogeneous wireless technologies, e.g. a Wi-Fi6 Mesh or parallel CBRS access links federated by a 5G backhaul.

The establishment of a path is not in-scope for RAW. It may be the product of a centralized Controller Plane as described for DetNet. As opposed to wired networks, the action of installing a path over a set of wireless links may be very slow relative to the speed at which the radio conditions vary, and it makes sense in the wireless case to provide redundant forwarding solutions along a complex path and to leave it to the Network Plane to select which of those forwarding solutions are to be used for a given packet based on the current conditions.

RAW distinguishes the longer time scale at which routes are computed from the the shorter forwarding time scale where per-packet decisions are made. RAW operates within the Network Plane at the forwarding time scale on one DetNet flow over a complex path called a Track. The Track is preestablished and installed by means outside of the scope of RAW; it may be strict or loose depending on whether each or just a subset of the hops are observed and controlled by RAW.

The RAW Architecture covers Network Plane protocol elements such as Operations, Administration and Maintenance (OAM) to observe some or all hops along a Track as well as the end-to-end packet delivery, and in-band control to optimize the use of redundancy to achieve the required SLA with minimal use of constrained resources.

2. Terminology

RAW reuses terminology defined for DetNet in the "Deterministic Networking Architecture" [RFC8655], e.g., PREOF for Packet Replication, Elimination and Ordering Functions.

RAW also reuses terminology defined for 6TiSCH in [6TiSCH-ARCHI] such as the term Track. A Track as a complex path with associated PAREO operations. The concept is abstract to the underlaying technology and applies to any fully or partially wireless mesh, including, e.g., a Wi-Fi mesh. RAW specifies strict and loose Tracks depending on whether the path is fully controlled by RAW or traverses an opaque network where RAW cannot observe and control the individual hops.

RAW uses the term OAM as defined in [RFC6291].

RAW defines the following terms:

PAREO: Packet (hybrid) ARQ, Replication, Elimination and Ordering. PAREO is a superset Of DetNet's PREOF that includes radio-specific techniques such as short range broadcast, MUMIMO, constructive interference and overhearing, which can be leveraged separately or combined to increase the reliability.

Flapping: In the context of RAW, a link flaps when the reliability of the wireless connectivity drops abruptly for a short period of time, typically of a subsecond to seconds duration.

In the context of the RAW work, Reliability and Availability are defined as follows:

Reliability: Reliability is a measure of the probability that an item will perform its intended function for a specified interval under stated conditions. For RAW, the service that is expected is delivery within a bounded latency and a failure is when the packet is either lost or delivered too late. RAW expresses reliability in terms of Mean Time Between Failure (MTBF) and Maximum Consecutive Failures (MCF). More in [NASA].

Availability: Availability is a measure of the relative amount of time where a path operates in stated condition, in other words (uptime)/(uptime+downtime). Because a serial wireless path may not be good enough to provide the required availability, and even 2 parallel paths may not be over a longer period of time, the RAW availability implies a path that is a lot more complex than what DetNet typically envisages (a Track).

3. Related Work at The IETF

RAW intersects with protocols or practices in development at the IETF as follows:

- * The Dynamic Link Exchange Protocol (DLEP) [RFC8175] from [MANET] can be leveraged at each hop to derive generic radio metrics (e.g., based on LQI, RSSI, queueing delays and ETX) on individual hops.
- * OAM work at [detnet] such as [DetNet-IP-OAM] for the case of the IP Data Plane observes the state of DetNet paths, typically MPLS and IPv6 pseudowires [DetNet-DP-FW], in the direction of the traffic. RAW needs feedback that flows on the reverse path and gathers instantaneous values from the radio receivers at each hop to inform back the source and replicating relays so they can make optimized forwarding decisions. The work named ICAN may be related as well.
- * [BFD] detect faults in the path between an Ingress and an Egress forwarding engines, but is unaware of the complexity of a path with replication, and expects bidirectionality. BFD considers delivery as success whereas with RAW the bounded latency can be as important as the delivery itself.
- * [SPRING] and [BIER] define in-band signaling that influences the routing when decided at the head-end on the path. There's already one RAW-related draft at BIER [BIER-PREF] more may follow. RAW will need new in-band signaling when the decision is distributed, e.g., required chances of reliable delivery to destination within latency. This signaling enables relays to tune retries and replication to meet the required SLA.
- * [CCAMP] defines protocol-independent metrics and parameters (measurement attributes) for describing links and paths that are required for routing and signaling in technology-specific networks. RAW would be a source of requirements for CCAMP to define metrics that are significant to the focus radios.


4. Use Cases and Requirements Served

In order to focus on real-worlds issues and assert the feasibility of the proposed capabilities, RAW focuses on selected technologies that can be scheduled at the lower layers: IEEE Std. 802.15.4 timeslotted channel hopping (TSCH), 3GPP 5G ultra-reliable low latency communications (URLLC), IEEE 802.11ax/be where 802.11be is extreme high throughput (EHT), and L-band Digital Aeronautical Communications System (LDACS). See [RAW-TECHNOS] for more.

"Deterministic Networking Use Cases" [RFC8578] presents a number of wireless use cases including Wireless, such as application to Industrial Applications, Pro-Audio, and SmartGrid Automation. [RAW-USE-CASES] adds a number of use cases that demonstrate the need for RAW capabilities for new applications such as Pro-Gaming and drones. The use cases can be abstracted in two families, Loose Protection, e.g., protecting the first hop in Radio Access Protection and Strict Protection, e.g., providing End-to-End Protection in a wireless mesh.

4.1. Radio Access Protection

To maintain the required SLA at all times, a wireless Host may use more than one Radio Access Network (RAN) in parallel.



*** = flapping at this time \$\$\$ expensive

Figure 1: Radio Access Protection


The RANs may be heterogeneous, e.g., 3GPP 5G [RAW-5G] and Wi-Fi [RAW-TECHNOS] for high-speed communication, in which case a Layer-3 abstraction becomes useful to select which of the RANs are used at a particular point of time, and the amount of traffic that is distributed over each RAN.

The idea is that the rest of the path to the destination(s) is protected separately (e.g., uses non-congruent paths, leverages DetNet / TSN, etc...) and is a lot more reliable, e.g., wired. In that case, RAW observes the reliability of the end-to-end operation through each of the RANs but only observes and controls the wireless operation the first hop.

A variation of that use case has a pair of wireless Hosts connected over a wired core / backbone network. In that case, RAW observes and controls the Ingress and Egress RANs, while neglecting the hops in the core. The resulting loose Track may be instantiated, e.g., using tunneling or loose source routing between the RANs.

4.2. End-to-End Protection in a Wireless Mesh

In radio technologies that support mesh networking (e.g., Wi-Fi and TSCH), a Track is a complex path with distributed PAREO capabilities. In that case, RAW operates through the multipath and makes decisions either at the Ingress or at every hop (more in Section 6.1).



zzz = flapping now

Figure 2: End-to-End Protection

The Protection may be imposed by the source based on end-to-end OAM, or performed hop-by-hop, in which case the OAM must enables the intermediate Nodes to estimate the quality of the rest of the feasible paths in the remainder of the Track to the destination.

5. RAW Considerations

5.1. Reliability and Availability

5.1.1. High Availability Engineering Principles

The reliability criteria of a critical system pervade through its elements, and if the system comprises a data network then the data network is also subject to the inherited reliability and availability criteria. It is only natural to consider the art of high availability engineering and apply it to wireless communicaitons in the context of RAW.

There are three principles [pillars] of high availability engineering:

1. elimination of single points of failure
2. reliable crossover
3. prompt detection of failures as they occur.

These principles are common to all high availability systems, not just ones with Internet technology at the center. Examples of both non-Internet and Internet are included.

5.1.1.1. Elimination of Single Points of Failure

Physical and logical components in a system happen to fail, either as the effect of wear and tear, when used beyond acceptable limits, or due to a software bug. It is necessary to decouple component failure from system failure to avoid the latter. This allows failed components to be restored while the rest of the system continues to function.

IP Routers leverage routing protocols to compute alternate routes in case of a failure. There is a rather open-ended issue over alternate routes -- for example, when links are cabled through the same conduit, they form a shared risk link group (SRLG), and will share the same fate if the bundle is cut. The same effect can happen with virtual links that end up in a same physical transport through the games of encapsulation. In a same fashion, an interferer or an obstacle may affect multiple wireless transmissions at the same time, even between different sets of peers.

Intermediate network Nodes such as routers, switches and APs, wire bundles and the air medium itself can become single points of failure. For High Availability, it is thus required to use physically link- and Node-disjoint paths; in the wireless space, it is also required to use the highest possible degree of diversity in the transmissions over the air to combat the additional causes of transmission loss.

From an economics standpoint, executing this principle properly generally increases capitalization expense because of the redundant equipment. In a constrained network where the waste of energy and bandwidth should be minimized, an excessive use of redundant links must be avoided; for RAW this means that the extra bandwidth must be used wisely and with parcimony.

5.1.1.2. Reliable Crossover

Having a backup equipment has a limited value unless it can be reliably switched into use within the down-time parameters. IP Routers execute reliable crossover continuously because the routers will use any alternate routes that are available [RFC0791]. This is due to the stateless nature of IP datagrams and the dissociation of the datagrams from the forwarding routes they take. The "IP Fast Reroute Framework" [FRR] analyzes mechanisms for fast failure detection and path repair for IP Fast-Reroute, and discusses the case of multiple failures and SRLG. Examples of FRR techniques include Remote Loop-Free Alternate [RLFA-FRR] and backup label-switched path (LSP) tunnels for the local repair of LSP tunnels using RSVP-TE [RFC4090].

Deterministic flows, on the contrary, are attached to specific paths where dedicated resources are reserved for each flow. This is why each DetNet path must inherently provide sufficient redundancy to provide the guaranteed SLA at all times. The DetNet PREOF typically leverages 1+1 redundancy whereby a packet is sent twice, over non-congruent paths. This avoids the gap during the fast reroute operation, but doubles the traffic in the network.

In the case of RAW, the expectation is that multiple transient faults may happen in overlapping time windows, in which case the 1+1 redundancy with delayed reestablishment of the second path will not provide the required guarantees. The Data Plane must be configured with a sufficient degree of redundancy to select an alternate redundant path immediately upon a fault, without the need for a slow intervention from the controller plane.

5.1.1.3. Prompt Notification of Failures

The execution of the two above principles is likely to render a system where the user will rarely see a failure. But someone needs to in order to direct maintenance.

There are many reasons for system monitoring (FCAPS for fault, configuration, accounting, performance, security is a handy mental checklist) but fault monitoring is sufficient reason.

"An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks" [STD 62] describes how to use SNMP to observe and correct long-term faults.

"Overview and Principles of Internet Traffic Engineering" [TE] discusses the importance of measurement for network protection, and provides abstract an method for network survivability with the analysis of a traffic matrix as observed by SNMP, probing techniques, FTP, IGP link state advertisements, and more.

Those measurements are needed in the context of RAW to inform the controller and make the long term reactive decision to rebuild a complex path. But RAW itself operates in the Network Plane at a faster time scale. To act on the Data Plane, RAW needs live information from the Operational Plane , e.g., using Bidirectional Forwarding Detection [BFD] and its variants (bidirectional and remote BFD) to protect a link, and OAM techniques to protect a path.

5.1.2. Applying Reliability Concepts to Networking

The terms Reliability and Availability are defined for use in RAW in Section 2 and the reader is invited to read [NASA] for more details on the general definition of Reliability. Practically speaking a number of nines is often used to indicate the reliability of a data link, e.g., 5 nines indicate a Packet Delivery Ratio (PDR) of 99.999%.

This number is typical in a wired environment where the loss is due to a random event such as a solar particle that affects the transmission of a particular frame, but does not affect the previous or next frame, nor frames transmitted on other links. Note that the QoS requirements in RAW may include a bounded latency, and a packet that arrives too late is a fault and not considered as delivered.

For a periodic networking pattern such as an automation control loop, this number is proportional to the Mean Time Between Failures (MTBF). When a single fault can have dramatic consequences, the MTBF expresses the chances that the unwanted fault event occurs. In data networks, this is rarely the case. Packet loss cannot never be fully avoided and the systems are built to resist to one loss, e.g., using redundancy with Retries (HARQ) or Packet Replication and Elimination (PRE), or, in a typical control loop, by linear interpolation from the previous measurements.

But the linear interpolation method cannot resist multiple consecutive losses, and a high MTBF is desired as a guarantee that this will not happen, IOW that the number of losses-in-a-row can be bounded. In that case, what is really desired is a Maximum Consecutive Failures (MCF). If the number of losses in a row passes the MCF, the control loop has to abort and the system, e.g., the production line, may need to enter an emergency stop condition.

Engineers that build automated processes may use the network reliability expressed in nines or as an MTBF as a proxy to indicate an MCF, e.g., as described in section 7.4 of the "Deterministic Networking Use Cases" [RFC8578].

5.1.3. Reliability in the Context of RAW

In contrast with wired networks, errors in transmission are the predominant source of packet loss in wireless networks.

The root cause for the loss may be of multiple origins, calling for the use of different forms of diversity:

Multipath Fading: A destructive interference by a reflection of the

original signal.

A radio signal may be received directly (line-of-sight) and/or as a reflection on a physical structure (echo). The reflections take a longer path and are delayed by the extra distance divided by the speed of light in the medium. Depending on the frequency, the echo lands with a different phase which may add up to (constructive interference) or cancel the direct signal (destructive interference).

The affected frequencies depend on the relative position of the sender, the receiver, and all the reflecting objects in the environment. A given hop will suffer from multipath fading for multiple packets in a row till the something moves that changes the reflection patterns.

Co-channel Interference: Energy in the spectrum used for the transmission confuses the receiver.

The wireless medium itself is a Shared Risk Link Group (SRLG) for nearby users of the same spectrum, as an interference may affect multiple co-channel transmissions between different peers within the interference domain of the interferer, possibly even when they use different technologies.

Obstacle in Fresnel Zone: The optimal transmission happens when the Fresnel Zone between the sender and the receiver is free of obstacles.

As long as a physical object (e.g., a metallic trolley between peers) that affects the transmission is not removed, the quality of the link is affected.

In an environment that is rich of metallic structures and mobile objects, a single radio link will provide a fuzzy service, meaning that it cannot be trusted to transport the traffic reliably over a long period of time.

Transmission losses are typically not independent, and their nature and duration are unpredictable; as long as a physical object (e.g., a metallic trolley between peers) that affects the transmission is not removed, or as long as the interferer (e.g., a radar) keeps transmitting, a continuous stream of packets will be affected.

The key technique to combat those unpredictable losses is diversity. Different forms of diversity are necessary to combat different causes of loss and the use of diversity must be maximised to optimize the PDR.

A single packet may be sent at different times (time diversity) over diverse paths (spatial diversity) that rely on diverse radio channels (frequency diversity) and diverse PHY technologies, e.g., narrowband vs. spread spectrum, or diverse codes. Using time diversity will defeat short-term interferences; spatial diversity combats very local causes such as multipath fading; narrowband and spread spectrum are relatively innocuous to one another and can be used for diversity in the presence of the other.

5.2. RAW Scope and Prerequisites

A prerequisite to the RAW work is that an end-to-end routing function computes a complex sub-topology along which forwarding can happen between a source and one or more destinations. The concept of Track is specified in the 6TiSCH Architecture [6TiSCH-ARCHI] to represent that complex sub-topology. Tracks provide a high degree of redundancy and diversity and enable the DetNet PREOF, network coding, and possibly RAW specific techniques such as PAREO, leveraging frequency diversity, time diversity, and possibly other forms of diversity as well.

How the routing operation (e.g., PCE) in the Controller Plane computes the Track is out of scope for RAW. The scope of the RAW operation is one Track, and the goal of the RAW operation is to optimize the use of the Track at the forwarding timescale to maintain the expected SLA while optimizing the usage of constrained resources such as energy and spectrum.

Another prerequisite is that an IP link can be established over the radio with some guarantees in terms of service reliability, e.g., it can be relied upon to transmit a packet within a bounded latency and provides a guaranteed BER/PDR outside rare but existing transient outage windows that can last from split seconds to minutes. The radio layer can be programmed with abstract parameters, and can return an abstract view of the state of the Link to help the Network Layer forwarding decision (think DLEP from MANET).

How the radio interface manages its lower layers is out of control and out of scope for RAW. In the same fashion, the non-RAW portion along a loose Track is by definition out of control and out of scope for RAW. Whether it is a single hop or a mesh is also unknown and out of scope.

5.3. Routing Time Scale vs. Forwarding Time Scale

With DetNet, the Controller Plane Function that handles the routing computation and maintenance (the PCE) can be centralized and can reside outside the network. In a wireless mesh, the path to the PCE can be expensive and slow, possibly going across the whole mesh and back. Reaching to the PCE can also be slow in regards to the speed of events that affect the forwarding operation at the radio layer.

Due to that cost and latency, the Controller Plane is not expected to be sensitive/reactive to transient changes. The abstraction of a link at the routing level is expected to use statistical metrics that aggregate the behavior of a link over long periods of time, and represent its properties as shades of gray as opposed to numerical values such as a link quality indicator, or a boolean value for either up or down.




Figure 3: Time Scales

In the case of wireless, the changes that affect the forwarding decision can happen frequently and often for short durations, e.g., a mobile object moves between a transmitter and a receiver, and will cancel the line of sight transmission for a few seconds, or a radar measures the depth of a pool and interferes on a particular channel for a split second.

There is thus a desire to separate the long term computation of the route and the short term forwarding decision. In that model, the routing operation computes a complex Track that enables multiple Non-Equal Cost Multi-Path (N-ECMP) forwarding solutions, and leaves it to the Data Plane to make the per-packet decision of which of these possibilities should be used.

In the wired world, and more specifically in the context of Traffic Engineering (TE), an alternate path can be used upon the detection of a failure in the main path, e.g., using OAM in MPLS-TP or BFD over a collection of SD-WAN tunnels. RAW formalizes a forwarding time scale that is an order(s) of magnitude shorter than the controller plane routing time scale, and separates the protocols and metrics that are used at both scales. Routing can operate on long term statistics such as delivery ratio over minutes to hours, but as a first approximation can ignore flapping. On the other hand, the RAW forwarding decision is made at the scale of the packet rate, and uses information that must be pertinent at the present time for the current transmission(s).

6. RAW Architecture Elements

A RAW Network Plane may be strict or loose, depending on whether RAW observes and takes actions on all hops or not. For instance, the packets between two wireless entities may be relayed over a wired infrastructure such as a Wi-Fi extended service set (ESS) or a 5G Core; in that case, RAW observes and control the transmission over the wireless first and last hops, as well as end-to-end metrics such as latency, jitter, and delivery ratio. This operation is loose since the structure and properties of the wired infrastructure are ignored, and may be either controlled by other means such as DetNet/TSN, or neglected in the face of the wireless hops.

6.1. Wireless Tracks

The "6TiSCH Architecture" [6TiSCH-ARCHI] introduces the concept of Track. RAW extends the concept to any wireless mesh technology, including, e.g., Wi-Fi. A simple Track is composed of a direct sequence of reserved hops to ensure the transmission of a single packet from a source Node to a destination Node across a multihop path.

A Complex Track provides multiple non-equal cost multipath (NECM) forwarding solutions. The Complex Track enables to support multi-path redundant forwarding by employing PRE functions [RFC8655] and the ingress and within the Track. For example, a Complex Track may branch off and rejoin over non-congruent segments.

In the context of RAW, some links or segments in the Track may be reversible, meaning that they can be used in either direction. In that case, an indication in the packet signals the direction of the reversible links or segments that the packet traverses and thus places a constraint that prevents loops from occurring. An individual packet follows a destination-oriented directed acyclic graph (DODAG) towards a destination Node inside the Complex Track.

6.2. PAREO Functions

RAW may control whether and how to use packet replication and elimination (PRE), Automatic Repeat reQuest (ARQ), Hybrid ARQ (HARQ) that includes Forward Error Correction (FEC) and coding, and other wireless-specific techniques such as overhearing and constructive interferences, in order to increase the reliability and availability of the end-to-end transmission.

Collectively, those function are called PAREO for Packet (hybrid) ARQ, Replication, Elimination and Ordering. By tuning dynamically the use of PAREO functions, RAW avoids the waste of critical resources such as spectrum and energy while providing that the guaranteed SLA, e.g., by adding redundancy only when a spike of loss is observed.

In a nutshell, PAREO establishes several paths in a network to provide redundancy and parallel transmissions to bound the end-to-end delay to traverse the network. Optionally, promiscuous listening between paths is possible, such that the Nodes on one path may overhear transmissions along the other path. Considering the scenario shown in Figure 4, many different paths are possible for S to reach R. A simple way to benefit from this topology could be to use the two independent paths via Nodes A, C, E and via B, D, F. But more complex paths are possible by interleaving transmissions from the lower level of the path to the upper level.




Figure 4: A Ladder Shape with Two Parallel Paths

PAREO may also take advantage of the shared properties of the wireless medium to compensate for the potential loss that is incurred with radio transmissions.

For instance, when the source sends to Node A, Node B may listen promiscuously and get a second chance to receive the frame without an additional transmission. Note that B would not have to listen if it already received that particular frame at an earlier timeslot in a dedicated transmission towards B.

The PAREO model can be implemented in both centralized and distributed scheduling approaches. In the centralized approach, a Path Computation Element (PCE) scheduler calculates a Track and schedules the communication. In the distributed approach, the Track is computed within the network, and signaled in the packets, e.g., using BIER-TE, Segment Routing, or a Source Routing Header.

6.2.1. Packet Replication

By employing a Packet Replication procedure, a Node forwards a copy of each data packet to more than one successor. To do so, each Node (i.e., Ingress and intermediate Node) sends the data packet multiple times as separate unicast transmissions. For instance, in Figure 5, the Ingress Node is transmitting the packet to both successors, nodes A and B, at two different times.




Figure 5: Packet Replication

An example schedule is shown in Table 1. This way, the transmission leverages with the time and spatial forms of diversity.

Channel	0	1	2	3	4	5	6
0	S->A	S->B	B->C	B->D	C->F	E->R	F->R
1		A->C	A->D	C->E	D->E	D->F	

Table 1: Packet Replication: Sample schedule

6.2.2. Packet Elimination

The replication operation increases the traffic load in the network, due to packet duplications. This may occur at several stages inside the Track, and to avoid an explosion of the number of copies, a Packet Elimination procedure must be applied as well. To this aim, once a Node receives the first copy of a data packet, it discards the subsequent copies.

The logical functions of Replication and Elimination may be collocated in an intermediate Node, the Node first eliminating the redundant copies and then sending the packet exactly once to each of the selected successors.

6.2.3. Promiscuous Overhearing

Considering that the wireless medium is broadcast by nature, any neighbor of a transmitter may overhear a transmission. By employing the Promiscuous Overhearing operation, the next hops have additional opportunities to capture the data packets. In Figure 6, when Node A is transmitting to its DP (Node C), the AP (Node D) and its sibling (Node B) may decode this data packet as well. As a result, by employing correlated paths, a Node may have multiple opportunities to receive a given data packet.




Figure 6: Unicast with Overhearing

6.2.4. Constructive Interference

Constructive Interference can be seen as the reverse of Promiscuous Overhearing, and refers to the case where two senders transmit the exact same signal in a fashion that the emitted symbols add up at the receiver and permit a reception that would not be possible with a single sender at the same PHY mode and the same power level.

Constructive Interference was proposed on 5G, Wi-Fi7 and even tested on IEEE Std 802.14.5. The hard piece is to synchronize the senders to the point that the signals are emitted at slightly different time to offset the difference of propagation delay that corresponds to the difference of distance of the transmitters to the receiver at the speed of light to the point that the symbols are superposed long enough to be recognizable.

7. RAW Architecture

RAW inherits the conceptual model described in section 4 of the DetNet Architecture [RFC8655]. RAW extends the DetNet service layer to provide additional agility against transmission loss.

A Controller Plane Function (CPF) called the Path Computation Element (PCE) [RFC4655] interacts with RAW Nodes over a Southbound API. The RAW Nodes are DetNet relays that are capable of additional diversity mechanisms and measurement functions related to the radio interface, in particular the PAREO diversity mechanisms.

The PCE defines a complex Track between an Ingress End System and an Egress End System, and indicates to the RAW Nodes where the PAREO operations may be actionned in the Network Plane. The Track may be expressed loosely to enable traversing a non-RAW subnetwork. In that case, the expectation is that the non-RAW subnetwork can be neglected in the RAW computation, that is, considered infinitely fast, reliable and/or available in comparison with the links between RAW nodes.

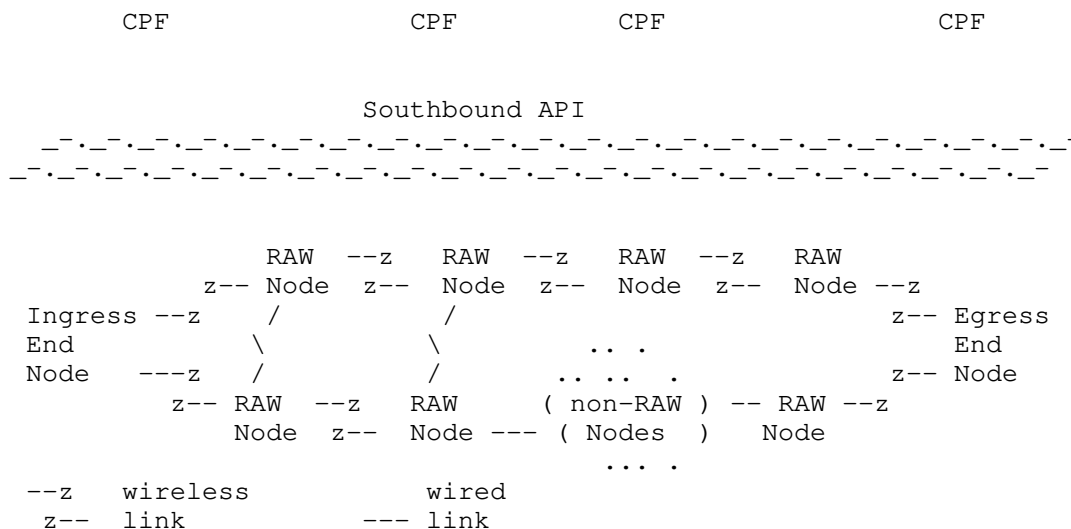


Figure 7: RAW Nodes

The Link-Layer metrics are reported to the PCE in a time-aggregated, e.g., statistical fashion. Example Link-Layer metrics include typical Link bandwidth (the medium speed depends dynamically on the PHY mode and the number of users sharing the spectrum) and average and mean squared deviation of availability and reliability figures such as Packet Delivery Ratio (PDR) over long periods of time.

Based on those metrics, the PCE installs the Track with enough redundant forwarding solutions to ensure that the Network Plane can reliably deliver the packets within a System Level Agreement (SLA) associated to the flows that it transports. The SLA defines end-to-end reliability and availability requirements, where reliability may be expressed as a successful delivery in order and within a bounded delay of at least one copy of a packet.

Depending on the use case and the SLA, the Track may comprise non-RAW segments, either interleaved inside the Track, or all the way to the Egress End Node (e.g., a server in the Internet). RAW observes the Lower-Layer Links between RAW nodes (typically, radio links) and the end-to-end Network Layer operation to decide at all times which of the PAREO diversity schemes is actioned by which RAW Nodes.

Once a Track is established, per-segment and end-to-end reliability and availability statistics are periodically reported to the PCE to assure that the SLA can be met or have it recompute the Track if not.

7.1. PCE vs. PSE

RAW separates the path computation time scale at which a complex path is recomputed from the path selection time scale at which the forwarding decision is taken for one or a few packets (more in Section 5.3). RAW operates at the path selection time scale. The RAW problem is to decide, within the redundant solutions that are proposed by the PCE, which will be used for each packet to provide a Reliable and Available service while minimizing the waste of constrained resources.

To that effect, RAW defines the Path Selection Engine (PSE) that is the counter-part of the PCE to perform rapid local adjustments of the forwarding tables within the diversity that the PCE has selected for the Track. The PSE enables to exploit the richer forwarding capabilities with PAREO and scheduled transmissions at a faster time scale over the smaller domain that is the Track, in either a loose or a strict fashion.

Compared to the PCE, the PSE operates on metrics that evolve faster, but that needs to be advertised at a fast rate but only locally, within the Track. The forwarding decision may also change rapidly, but with a scope that is also contained within the Track, with no visibility to the other Tracks and flows in the network. This is as opposed to the PCE that needs to observe the whole network, and optimize all the Tracks globally, which can only be done at a slow pace and using long-term statistical metrics, as presented in Table 2.

	PCE (Not in Scope)	PSE (In Scope)
Operation	Centralized	Source-Routed or Distributed
Communication	Slow, expensive	Fast, local
Time Scale	hours and above	seconds and below
Network Size	Large, many Tracks to optimize globally	Small, within one Track
Considered Metrics	Averaged, Statistical, Shade of grey	Instant values / boolean condition

Table 2: PCE vs. PSE

The PSE sits in the DetNet Service sub-Layer of Edge and Relay Nodes. On the one hand, it operates on the packet flow, learning the Track and path selection information from the packet, possibly making local decision and retagging the packet to indicate so. On the other hand, the PSE interacts with the lower layers and with its peers to obtain up-to-date information about its radio links and the quality of the overall Track, respectively, as illustrated in Figure 8.

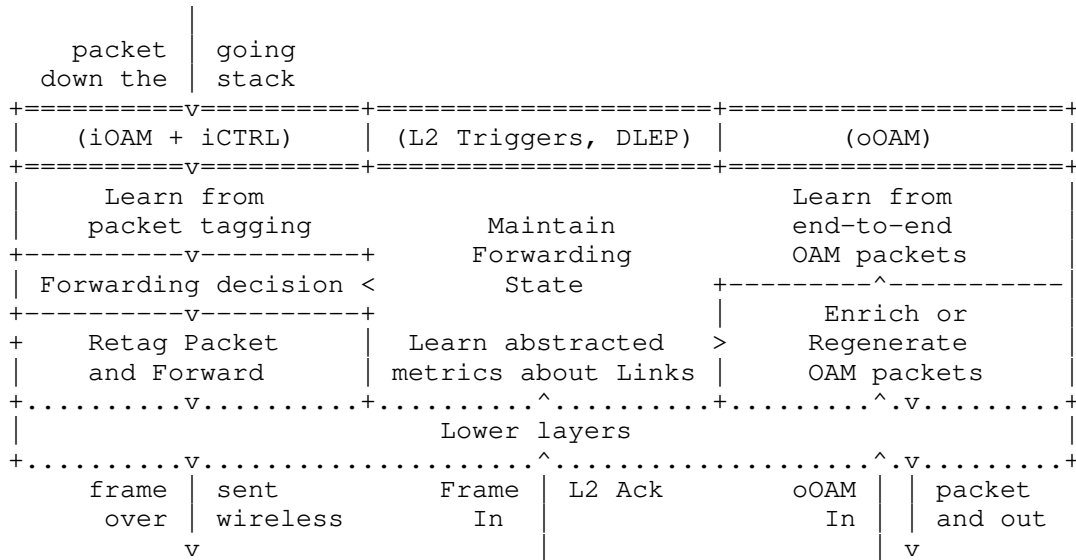


Figure 8: PSE

7.2. RAW OAM

The RAW OAM operation in the Network Plane observes either a full Track or subTracks that are being used at this time. This observation feeds the RAW PSE that makes the decision on which PAREO function is actioned at which RAW Node, for one a small continuous series of packets.

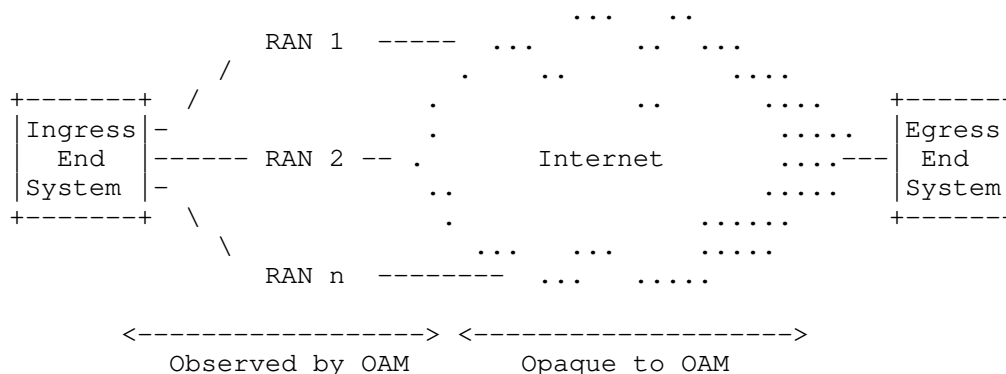


Figure 9: Observed Links in Radio Access Protection

In the case of a End-to-End Protection in a Wireless Mesh, the Track is strict and congruent with the path so all links are observed. Conversely, in the case of Radio Access Protection, the Track is Loose and in that case only the first hop is observed; the rest of the path is abstracted and considered infinitely reliable.

In the case of the Radio Access Protection, only the first hop is protected; the loss of a packet that was sent over one of the possible first hops is attributed to that first hop, even if a particular loss effectively happens farther down the path.

The Links that are not observed by OAM are opaque to it, meaning that the OAM information is carried across and possibly echoed as data, but there is no information capture in intermediate nodes. In the example above, the Internet is opaque and not controlled by RAW; still the RAW OAM measures the end-to-end latency and delivery ratio for packets sent via each of RAN 1, RAN 2 and RAN 3, and determines whether a packet should be sent over either or a collection of those access links.

7.3. Flow Identification vs. Path Identification

Section 4.7 of the DetNet Architecture [RFC8655] ties the app-flow identification which is an application layer concept with the network path identification that depends on the networking technology by "exporting of flow identification", e.g., to a MPLS label.

With RAW, this exporting operation is injective but not bijective. e.g., a flow is fully placed within one RAW Track, but not all packets along that Track are necessarily part of the same flow. For instance, out-of-band OAM packets must circulate in the exact same fashion as the flows that they observe. It results that the flow identification that maps to app-flow at the network layer must be separate from the path identification that is used to forward a packet.

Section 3.4 of the DetNet data-plane framework [DetNet-DP-FW] indicates that for a DetNet IP Data Plane, a flow is identified by an IPv6 6-tuple. With RAW, that 6-tuple is not what indicates the Track, in other words, the flow ID is not the Track ID.

For instance, the 6TiSCH Architecture [6TiSCH-ARCHI] uses a combination of the address of the Egress End System and an instance identifier in a Hop-by-hop option to indicate a Track. This way, if a packet "escapes" the Track, it will reach the Track Egress point through normal routing and be treated at the service layer through, say, elimination and reordering.

The RAW service includes forwarding over a subset of the Links that form the Track (a subTrack). Packets from the same or a different flow that are routed through the same Track will not necessarily traverse the same Links. The PSE selects a subTrack for a packet based on the links that are preferred and those that should be avoided at this time.

Each packet is forwarded within the subTrack that provides the best adequation with the SLA of the flow and the energy and bandwidth constraints of the network.

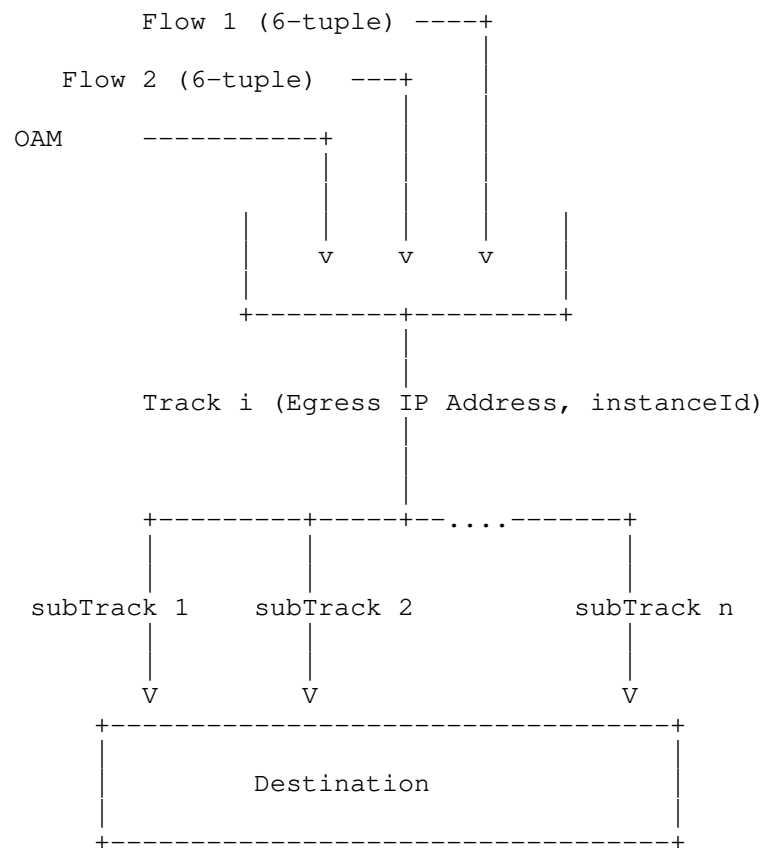


Figure 10: Flow Injection

With 6TiSCH, packets are tagged with the same (destination address, instance ID) will experience the same RAW service regardless of the IPv6 6-tuple that indicates the flow. The forwarding does not depend on whether the packets transport application flows or OAM. In the generic case, the Track or the subTrack can be signaled in the packet through other means, e.g., encoded in the suffix of the destination address as a Segment Routing Service Instruction [SR-ARCHI], or leveraging Bit Index Explicit Replication [BIER] Traffic Engineering [BIER-TE].

7.4. Source-Routed vs. Distributed Forwarding Decision

Within a large routed topology, the route-over mesh operation builds a particular complex Track with one source and one or more destinations; within the Track, packets may follow different paths and may be subject to RAW forwarding operations that include replication, elimination, retries, overhearing and reordering.

The RAW forwarding decisions include the selection of points of replication and elimination, how many retries can take place, and a limit of validity for the packet beyond which the packet should be destroyed rather than forwarded uselessly further down the Track.

The decision to apply the RAW techniques must be done quickly, and depends on a very recent and precise knowledge of the forwarding conditions within the complex Track. There is a need for an observation method to provide the RAW Data Plane with the specific knowledge of the state of the Track for the type of flow of interest (e.g., for a QoS level of interest). To observe the whole Track in quasi real time, RAW considers existing tools such as L2-triggers, DLEP, BFD and leverages in-band and out-of-band OAM to capture and report that information to the PSE.

One possible way of making the RAW forwarding decisions within a Track is to position a unique PSE at the Ingress and express its decision in-band in the packet, which requires the explicit signaling of the subTrack within the Track. In that case, the RAW forwarding operation along the Track is encoded by the source, e.g., by indicating the subTrack in the Segment Routing (SRv6) Service Instruction, or by leveraging BIER-TE such as done with [BIER-PREF].

The alternate way is to operate the PSE in each forwarding Node, which makes the RAW forwarding decisions for a packet on its own, based on its knowledge of the expectation (timeliness and reliability) for that packet and a recent observation of the rest of the way across the possible paths based on OAM. Information about the desired service should be placed in the packet and matched with the forwarding Node's capabilities and policies.

In either case, a per-track/subTrack state is installed in all the intermediate Nodes to recognize the packets that are following a Track and determine the forwarding operation to be applied.

7.5. Encapsulation and Decapsulation

In the generic case where the Track Ingress Node is not the source of the Packet, the Ingress Node needs to encapsulate IP-in-IP to ensure that the Destination IP Address is that of the Egress Node and that the necessary Headers (Router Header, Segment Routing Header and/or Hop-By-Hop Header) can be added to the packet to signal the Track or the subTrack, conforming [IPv6] that discourages the insertion of a Header on the fly.

In the specific case where the Ingress Node is the source of the packet, the encapsulation can be avoided, provided that the source adds the necessary headers and that the destination is set to the Egress Node. Forwarding to a final destination beyond the Egress Node is possible, e.g., with a Segment Routing Header that signals the rest of the way. In that case a Hop-by-Hop Header is not recommended since its validity is within the Track only.

8. Security Considerations

RAW uses all forms of diversity including radio technology and physical path to increase the reliability and availability in the face of unpredictable conditions. While this is not done specifically to defeat an attacker, the amount of diversity used in RAW makes an attack harder to achieve.

8.1. Forced Access

RAW will typically select the cheapest collection of links that matches the requested SLA, for instance, leverage free WI-Fi vs. paid 3GPP access. By defeating the cheap connectivity (e.g., PHY-layer interference) the attacker can force an End System to use the paid access and increase the cost of the transmission for the user.

9. IANA Considerations

This document has no IANA actions.

10. Contributors

Xavi Vilajosana: Wireless Networks Research Lab, Universitat Oberta de Catalunya

Remous-Aris Koutsiamanis: IMT Atlantique

Nicolas Montavont: IMT Atlantique

11. Acknowledgments

TBD

12. References

12.1. Normative References

[6TiSCH-ARCHI]

Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", Work in Progress, Internet-Draft, draft-ietf-6tisch-architecture-28, 29 October 2019, <<https://tools.ietf.org/html/draft-ietf-6tisch-architecture-28>>.

[RAW-TECHNOS]

Thubert, P., Cavalcanti, D., Vilajosana, X., Schmitt, C., and J. Farkas, "Reliable and Available Wireless Technologies", Work in Progress, Internet-Draft, draft-thubert-raw-technologies-05, 18 May 2020, <<https://tools.ietf.org/html/draft-thubert-raw-technologies-05>>.

[RAW-USE-CASES]

Papadopoulos, G., Thubert, P., Theoleyre, F., and C. Bernardos, "RAW use cases", Work in Progress, Internet-Draft, draft-bernardos-raw-use-cases-03, 8 March 2020, <<https://tools.ietf.org/html/draft-bernardos-raw-use-cases-03>>.

[RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <<https://www.rfc-editor.org/info/rfc4655>>.

[BFD] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <<https://www.rfc-editor.org/info/rfc5880>>.

[RFC6291] Andersson, L., van Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "Guidelines for the Use of the "OAM" Acronym in the IETF", BCP 161, RFC 6291, DOI 10.17487/RFC6291, June 2011, <<https://www.rfc-editor.org/info/rfc6291>>.

[RFC8578] Grossman, E., Ed., "Deterministic Networking Use Cases", RFC 8578, DOI 10.17487/RFC8578, May 2019, <<https://www.rfc-editor.org/info/rfc8578>>.

- [IPv6] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [SR-ARCHI] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [BIER] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", RFC 8279, DOI 10.17487/RFC8279, November 2017, <<https://www.rfc-editor.org/info/rfc8279>>.
- [RFC8175] Ratliff, S., Jury, S., Satterwhite, D., Taylor, R., and B. Berry, "Dynamic Link Exchange Protocol (DLEP)", RFC 8175, DOI 10.17487/RFC8175, June 2017, <<https://www.rfc-editor.org/info/rfc8175>>.
- [RFC8557] Finn, N. and P. Thubert, "Deterministic Networking Problem Statement", RFC 8557, DOI 10.17487/RFC8557, May 2019, <<https://www.rfc-editor.org/info/rfc8557>>.
- [RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <<https://www.rfc-editor.org/info/rfc8655>>.

12.2. Informative References

- [RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.
- [TE] Awduche, D., Chiu, A., Elwalid, A., Widjaja, I., and X. Xiao, "Overview and Principles of Internet Traffic Engineering", RFC 3272, DOI 10.17487/RFC3272, May 2002, <<https://www.rfc-editor.org/info/rfc3272>>.
- [STD 62] Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks", STD 62, RFC 3411, DOI 10.17487/RFC3411, December 2002, <<https://www.rfc-editor.org/info/rfc3411>>.

- [RFC4090] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, DOI 10.17487/RFC4090, May 2005, <<https://www.rfc-editor.org/info/rfc4090>>.
- [FRR] Shand, M. and S. Bryant, "IP Fast Reroute Framework", RFC 5714, DOI 10.17487/RFC5714, January 2010, <<https://www.rfc-editor.org/info/rfc5714>>.
- [RLFA-FRR] Bryant, S., Filsfils, C., Previdi, S., Shand, M., and N. So, "Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)", RFC 7490, DOI 10.17487/RFC7490, April 2015, <<https://www.rfc-editor.org/info/rfc7490>>.
- [BIER-PREF] Thubert, P., Eckert, T., Brodard, Z., and H. Jiang, "BIER-TE extensions for Packet Replication and Elimination Function (PREF) and OAM", Work in Progress, Internet-Draft, draft-thubert-bier-replication-elimination-03, 3 March 2018, <<https://tools.ietf.org/html/draft-thubert-bier-replication-elimination-03>>.
- [DetNet-IP-OAM] Mirsky, G., Chen, M., and D. Black, "Operations, Administration and Maintenance (OAM) for Deterministic Networks (DetNet) with IP Data Plane", Work in Progress, Internet-Draft, draft-mirsky-detnet-ip-oam-02, 23 March 2020, <<https://tools.ietf.org/html/draft-mirsky-detnet-ip-oam-02>>.
- [DetNet-DP-FW] Varga, B., Farkas, J., Berger, L., Malis, A., and S. Bryant, "DetNet Data Plane Framework", Work in Progress, Internet-Draft, draft-ietf-detnet-data-plane-framework-06, 6 May 2020, <<https://tools.ietf.org/html/draft-ietf-detnet-data-plane-framework-06>>.
- [RAW-5G] Farkas, J., Dudda, T., Shapin, A., and S. Sandberg, "5G - Ultra-Reliable Wireless Technology with Low Latency", Work in Progress, Internet-Draft, draft-farkas-raw-5g-00, 1 April 2020, <<https://tools.ietf.org/html/draft-farkas-raw-5g-00>>.
- [BIER-TE] Eckert, T., Cauchie, G., and M. Menth, "Tree Engineering for Bit Index Explicit Replication (BIER-TE)", Work in Progress, Internet-Draft, draft-ietf-bier-te-arch-07, 9 March 2020, <<https://tools.ietf.org/html/draft-ietf-bier-te-arch-07>>.

- [NASA] Adams, T., "RELIABILITY: Definition & Quantitative Illustration", <<https://kscddms.ksc.nasa.gov/Reliability/Documents/150814-3bWhatIsReliability.pdf>>.
- [MANET] IETF, "Mobile Ad hoc Networking", <<https://dataTracker.ietf.org/doc/charter-ietf-manet/>>.
- [detnet] IETF, "Deterministic Networking", <<https://dataTracker.ietf.org/doc/charter-ietf-detnet/>>.
- [SPRING] IETF, "Source Packet Routing in Networking", <<https://dataTracker.ietf.org/doc/charter-ietf-spring/>>.
- [BIER] IETF, "Bit Indexed Explicit Replication", <<https://dataTracker.ietf.org/doc/charter-ietf-bier/>>.
- [BFD] IETF, "Bidirectional Forwarding Detection", <<https://dataTracker.ietf.org/doc/charter-ietf-bfd/>>.
- [CCAMP] IETF, "Common Control and Measurement Plane", <<https://dataTracker.ietf.org/doc/charter-ietf-ccamp/>>.

Authors' Addresses

Pascal Thubert (editor)
Cisco Systems, Inc
Building D
45 Allée des Ormes - BP1200
06254 MOUGINS - Sophia Antipolis
France

Phone: +33 497 23 26 34
Email: pthubert@cisco.com

Georgios Z. Papadopoulos
IMT Atlantique
Office B00 - 114A
2 Rue de la Chataigneraie
35510 Cesson-Sevigne - Rennes
France

Phone: +33 299 12 70 04
Email: georgios.papadopoulos@imt-atlantique.fr

Rex Buddenberg
CA
United States of America

Email: buddenbergr@gmail.com

RAW
Internet-Draft
Intended status: Informational
Expires: 8 January 2022

P. Thubert, Ed.
Cisco Systems
G.Z. Papadopoulos
IMT Atlantique
L. Berger
LabN Consulting, L.L.C.
7 July 2021

Reliable and Available Wireless Architecture/Framework
draft-pthubert-raw-architecture-09

Abstract

Reliable and Available Wireless (RAW) provides for high reliability and availability for IP connectivity over a wireless medium. The wireless medium presents significant challenges to achieve deterministic properties such as low packet error rate, bounded consecutive losses, and bounded latency. This document defines the RAW Architecture. It builds on the DetNet Architecture and discusses specific challenges and technology considerations needed to deliver DetNet service utilizing scheduled wireless segments and other media, e.g., frequency/time-sharing physical media resources with stochastic traffic.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 8 January 2022.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. The RAW problem	5
2.1. Terminology	5
2.2. Reliability and Availability	7
2.2.1. High Availability Engineering Principles	8
2.2.2. Applying Reliability Concepts to Networking	10
2.2.3. Reliability in the Context of RAW	11
2.3. Use Cases and Requirements Served	12
2.3.1. Radio Access Protection	13
2.3.2. End-to-End Protection in a Wireless Mesh	13
2.4. Related Work at The IETF	14
3. The RAW Framework	15
3.1. Scope and Prerequisites	15
3.2. Routing Time Scale vs. Forwarding Time Scale	16
3.3. Wireless Tracks	17
3.4. PAREO Functions	18
3.4.1. Packet Replication	19
3.4.2. Packet Elimination	20
3.4.3. Promiscuous Overhearing	20
3.4.4. Constructive Interference	20
4. The RAW Architecture	21
4.1. The RAW Conceptual Model	21
4.2. The Path Selection Engine	23
4.3. RAW OAM	24
4.3.1. DetNet OAM	25
4.3.2. RAW Extensions	26
4.3.3. Observed Metrics	27
4.4. Flow Identification vs. Path Identification	27
4.5. Source-Routed vs. Distributed Forwarding Decision	30
4.6. Encapsulation and Decapsulation	31
5. Security Considerations	31
5.1. Forced Access	31
6. IANA Considerations	31
7. Contributors	31
8. Acknowledgments	32
9. References	32
9.1. Normative References	32

9.2. Informative References	34
Authors' Addresses	37

1. Introduction

Deterministic Networking is an attempt to emulate the properties of a serial link over a switched fabric, by providing a bounded latency and eliminating congestion loss, even when co-existing with best-effort traffic. It is getting traction in various industries including professional A/V, manufacturing, online gaming, and smartgrid automation, enabling cost and performance optimizations (e.g., vs. loads of P2P cables).

Bringing determinism in a packet network means eliminating the statistical effects of multiplexing that result in probabilistic jitter and loss. This can be approached with a tight control of the physical resources to maintain the amount of traffic within a budgetted volume of data per unit of time that fits the physical capabilities of the underlying network, and the use of time-shared resources (bandwidth and buffers) per circuit, and/or by shaping and/or scheduling the packets at every hop.

This innovation was initially introduced on wired networks, with IEEE 802.1 Time Sensitive networking (TSN) - for Ethernet LANs - and IETF DetNet. But the wired and the wireless media are fundamentally different at the physical level and in the possible abstractions that can be built for IP [IPoWIRELESS]. Wireless networks operate on a shared medium where uncontrolled interference, including the self-induced multipath fading, cause random transmission losses and add new dimensions to the statistical effects that affect reachability and packet delivery.

To defeat those additional causes of transmission delay and loss, Reliable and Available Wireless (RAW) leverages scheduled transmissions with redundancy and diversity in the spatial, time, code, and frequency domains. The challenge is to provide enough diversity and redundancy to ensure the timely packet delivery while preserving energy and optimizing the use of the shared spectrum.

While the generic "Deterministic Networking Problem Statement" [RFC8557] applies to both the wired and the wireless media, the methods to achieve RAW must extend those used to support time-sensitive networking over wires, as a RAW solution has to address less consistent transmissions, energy conservation and shared spectrum efficiency.

Uncontrolled interference and transmission obstacles may impede the wireless transmission, causing rapid variations of the throughput and packet delivery ratio (PDR) of the link. This uncertainty limits the volume and/or duration of traffic that can be safely transmitted on the same link while conforming to a RAW Service Level Agreement (SLA).

This increased complexity explains why the development of deterministic wireless technologies has been lagging behind the similar efforts for wired systems, both at the IEEE and the IETF. But recent progress on scheduled radios such as TSCH and OFDMA indicates that wireless is finally catching up at the lower layers. Sitting at the layer above, RAW takes up the challenge of providing highly available and reliable end-to-end performances in a network with scheduled wireless segments.

RAW provides DetNet elements that are specialized for short range radios. From this inheritance, RAW stays agnostic to the radio layer underneath though the capability to schedule transmissions is assumed. How the PHY is programmed to do so, and whether the radio is single-hop or meshed, are unknown at the IP layer and not part of the RAW abstraction.

The "Deterministic Networking Architecture" [RFC8655] is composed of three planes: the Application (User) Plane, the Controller Plane, and the Network Plane. The RAW Architecture extends the DetNet Network Plane, to accomodate one or multiple hops of homogeneous or heterogeneous wireless technologies, e.g. a Wi-Fi6 Mesh or parallel CBRS access links federated by a 5G backhaul.

The establishment of a path is not in-scope for RAW. It may be the product of a centralized Controller Plane as described for DetNet. As opposed to wired networks, the action of installing a path over a set of wireless links may be very slow relative to the speed at which the radio conditions vary, and it makes sense in the wireless case to provide redundant forwarding solutions along a complex path and to leave it to the Network Plane to select which of those forwarding solutions are to be used for a given packet based on the current conditions.

RAW distinguishes the longer time scale at which routes are computed from the the shorter forwarding time scale where per-packet decisions are made. RAW operates within the Network Plane at the forwarding time scale on one DetNet flow over a complex path called a Track. The Track is preestablished and installed by means outside of the scope of RAW; it may be strict or loose depending on whether each or just a subset of the hops are observed and controlled by RAW.

The RAW Architecture covers Network Plane protocol elements such as Operations, Administration and Maintenance (OAM) to observe some or all hops along a Track as well as the end-to-end packet delivery, and in-band control to optimize the use of redundancy to achieve the required SLA with minimal use of constrained resources.

2. The RAW problem

2.1. Terminology

RAW reuses terminology defined for DetNet in the "Deterministic Networking Architecture" [RFC8655], e.g., PREOF for Packet Replication, Elimination and Ordering Functions.

RAW also reuses terminology defined for 6TiSCH in [6TiSCH-ARCHI] such as the term Track. A Track as a complex path with associated PAREO operations. The concept is abstract to the underlaying technology and applies to any fully or partially wireless mesh, including, e.g., a Wi-Fi mesh. RAW specifies strict and loose Tracks depending on whether the path is fully controlled by RAW or traverses an opaque network where RAW cannot observe and control the individual hops.

RAW uses the following terminology:

PAREO: Packet (hybrid) ARQ, Replication, Elimination and Ordering.
PAREO is a superset Of DetNet's PREOF that includes radio-specific techniques such as short range broadcast, MUMIMO, constructive interference and overhearing, which can be leveraged separately or combined to increase the reliability.

Flow: A collection of consecutive packets that must be placed on the same Track to receive an equivalent treatment from Ingress to Egress within the Track. Multiple flows may be transported along the same Track. The subTrack that is selected for the flow may change over time under the control of the PSE.

Track: A networking graph that can be used as a "path" to transport RAW packets with equivalent treatment; as opposed to the usual understanding of a path (see for instance the definition of "path" in section 1.1 of [RFC9049]), a Track may fork and rejoin to enable the PAREO operations.

In DetNet [RFC8655] terms, a Track has the following properties:

- * A Track has one Ingress and one Egress nodes, which operate as DetNet Edge nodes.

- * A Track is reversible, meaning that packets can be routed against the flow of data packets, e.g., to carry OAM measurements or control messages back to the Ingress.
- * The vertices of the Track are DetNet Relay nodes that operate at the DetNet Service sublayer and provide the PAREO functions.
- * The topological edges of the graph are serial sequences of DetNet Transit nodes that operate at the DetNet Forwarding sublayer.

SubTrack: A Track within a Track. The RAW PSE selects a subTrack on a per-packet or a per-collection of packets basis to provide the desired reliability for the transported flows.

Segment: A serial path formed by a topological edge of a Track. East-West Segments are oriented from Ingress (East) to Egress (West). North/South Segments can be bidirectional; to avoid loops, measures must be taken to ensure that a given packet flows either Northwards or Southwards along a bidirectional Segment, but never bounces back.

Flapping: In the context of RAW, a link flaps when the reliability of the wireless connectivity drops abruptly for a short period of time, typically of a subsecond to seconds duration.

OAM: OAM stands for Operations, Administration, and Maintenance, and covers the processes, activities, tools, and standards involved with operating, administering, managing and maintaining any system. This document uses the terms Operations, Administration, and Maintenance, in conformance with the 'Guidelines for the Use of the "OAM" Acronym in the IETF' [RFC6291] and the system observed by the RAW OAM is the Track.

Active OAM: See [RFC7799]. In the context of RAXW, Active OAM is used to observe a particular Track, subTrack, or Segment of a Track regardless of whether it is used for traffic at that time.

In-Band OAM: An active OAM packet is considered in-band for the monitored Track when it traverses the same set of links and interfaces and if the OAM packet receives the same QoS and PAREO treatment as the packets of the data flows that are injected in the Track.

Out-of-Band OAM: Out-of-band OAM is an active OAM whose path is not topologically congruent to the Track, or its test packets receive a QoS and/or PAREO treatment that is different from that of the packets of the data flows that are injected in the Track, or both.

Limited OAM: An active OAM packet is a Limited OAM packet when it observes the RAW operation over a node, a segment, or a subTrack of the Track, though not from Ingress to Egress. It is injected in the datapath and extracted from the datapath around the particular function or subnetwork (e.g., around a relay providing a service layer replication point) that is being tested.

Reverse OAM: A Reverse OAM packet is an Out-of-Band OAM packet that traverses the Track from egress to ingress on the reverse direction, to capture and report OAM measurements upstream. The collection may capture all information along the whole Track, or it may only learn select data across all, or only a particular subTrack, or Segment of a Track.

[DetNet-OAM] provides additional terminology related to OAM in the context of DetNet and by extension of RAW, whereas [RFC7799] defines the Active, Passive, and Hybrid OAM methods.

In the context of the RAW work, Reliability and Availability are defined as follows:

Reliability: Reliability is a measure of the probability that an item will perform its intended function for a specified interval under stated conditions. For RAW, the service that is expected is delivery within a bounded latency and a failure is when the packet is either lost or delivered too late. RAW expresses reliability in terms of Mean Time Between Failure (MTBF) and Maximum Consecutive Failures (MCF). More in [NASA].

Availability: Availability is a measure of the relative amount of time where a path operates in stated condition, in other words $(\text{uptime})/(\text{uptime}+\text{downtime})$. Because a serial wireless path may not be good enough to provide the required reliability, and even 2 parallel paths may not be over a longer period of time, the RAW availability implies a path that is a lot more complex than what DetNet typically envisages (a Track).

Residence Time: A residence time (RT) is defined as the time period between the reception of a packet starts and the transmission of the packet begins. In the context of RAW, RT is useful for a transit node, not ingress or egress.

2.2. Reliability and Availability

2.2.1. High Availability Engineering Principles

The reliability criteria of a critical system pervade through its elements, and if the system comprises a data network then the data network is also subject to the inherited reliability and availability criteria. It is only natural to consider the art of high availability engineering and apply it to wireless communications in the context of RAW.

There are three principles [pillars] of high availability engineering:

1. elimination of single points of failure
2. reliable crossover
3. prompt detection of failures as they occur.

These principles are common to all high availability systems, not just ones with Internet technology at the center. Examples of both non-Internet and Internet are included.

2.2.1.1. Elimination of Single Points of Failure

Physical and logical components in a system happen to fail, either as the effect of wear and tear, when used beyond acceptable limits, or due to a software bug. It is necessary to decouple component failure from system failure to avoid the latter. This allows failed components to be restored while the rest of the system continues to function.

IP Routers leverage routing protocols to compute alternate routes in case of a failure. There is a rather open-ended issue over alternate routes -- for example, when links are cabled through the same conduit, they form a shared risk link group (SRLG), and will share the same fate if the bundle is cut. The same effect can happen with virtual links that end up in a same physical transport through the games of encapsulation. In a same fashion, an interferer or an obstacle may affect multiple wireless transmissions at the same time, even between different sets of peers.

Intermediate network Nodes such as routers, switches and APs, wire bundles and the air medium itself can become single points of failure. For High Availability, it is thus required to use physically link- and Node-disjoint paths; in the wireless space, it is also required to use the highest possible degree of diversity in the transmissions over the air to combat the additional causes of transmission loss.

From an economics standpoint, executing this principle properly generally increases capitalization expense because of the redundant equipment. In a constrained network where the waste of energy and bandwidth should be minimized, an excessive use of redundant links must be avoided; for RAW this means that the extra bandwidth must be used wisely and with parcimony.

2.2.1.2. Reliable Crossover

Having a backup equipment has a limited value unless it can be reliably switched into use within the down-time parameters. IP Routers execute reliable crossover continuously because the routers will use any alternate routes that are available [RFC0791]. This is due to the stateless nature of IP datagrams and the dissociation of the datagrams from the forwarding routes they take. The "IP Fast Reroute Framework" [FRR] analyzes mechanisms for fast failure detection and path repair for IP Fast-Reroute, and discusses the case of multiple failures and SRLG. Examples of FRR techniques include Remote Loop-Free Alternate [RLFA-FRR] and backup label-switched path (LSP) tunnels for the local repair of LSP tunnels using RSVP-TE [RFC4090].

Deterministic flows, on the contrary, are attached to specific paths where dedicated resources are reserved for each flow. This is why each DetNet path must inherently provide sufficient redundancy to provide the guaranteed SLA at all times. The DetNet PREOF typically leverages 1+1 redundancy whereby a packet is sent twice, over non-congruent paths. This avoids the gap during the fast reroute operation, but doubles the traffic in the network.

In the case of RAW, the expectation is that multiple transient faults may happen in overlapping time windows, in which case the 1+1 redundancy with delayed reestablishment of the second path will not provide the required guarantees. The Data Plane must be configured with a sufficient degree of redundancy to select an alternate redundant path immediately upon a fault, without the need for a slow intervention from the controller plane.

2.2.1.3. Prompt Notification of Failures

The execution of the two above principles is likely to render a system where the user will rarely see a failure. But someone needs to in order to direct maintenance.

There are many reasons for system monitoring (FCAPS for fault, configuration, accounting, performance, security is a handy mental checklist) but fault monitoring is sufficient reason.

"An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks" [STD 62] describes how to use SNMP to observe and correct long-term faults.

"Overview and Principles of Internet Traffic Engineering" [TE] discusses the importance of measurement for network protection, and provides abstract an method for network survivability with the analysis of a traffic matrix as observed by SNMP, probing techniques, FTP, IGP link state advertisements, and more.

Those measurements are needed in the context of RAW to inform the controller and make the long term reactive decision to rebuild a complex path. But RAW itself operates in the Network Plane at a faster time scale. To act on the Data Plane, RAW needs live information from the Operational Plane , e.g., using Bidirectional Forwarding Detection [BFD] and its variants (bidirectional and remote BFD) to protect a link, and OAM techniques to protect a path.

2.2.2. Applying Reliability Concepts to Networking

The terms Reliability and Availability are defined for use in RAW in Section 2.1 and the reader is invited to read [NASA] for more details on the general definition of Reliability. Practically speaking a number of nines is often used to indicate the reliability of a data link, e.g., 5 nines indicate a Packet Delivery Ratio (PDR) of 99.999%.

This number is typical in a wired environment where the loss is due to a random event such as a solar particle that affects the transmission of a particular frame, but does not affect the previous or next frame, nor frames transmitted on other links. Note that the QoS requirements in RAW may include a bounded latency, and a packet that arrives too late is a fault and not considered as delivered.

For a periodic networking pattern such as an automation control loop, this number is proportional to the Mean Time Between Failures (MTBF). When a single fault can have dramatic consequences, the MTBF expresses the chances that the unwanted fault event occurs. In data networks, this is rarely the case. Packet loss cannot never be fully avoided and the systems are built to resist to one loss, e.g., using redundancy with Retries (HARQ) or Packet Replication and Elimination (PRE), or, in a typical control loop, by linear interpolation from the previous measurements.

But the linear interpolation method cannot resist multiple consecutive losses, and a high MTBF is desired as a guarantee that this will not happen, IOW that the number of losses-in-a-row can be bounded. In that case, what is really desired is a Maximum

Consecutive Failures (MCF). If the number of losses in a row passes the MCF, the control loop has to abort and the system, e.g., the production line, may need to enter an emergency stop condition.

Engineers that build automated processes may use the network reliability expressed in nines or as an MTBF as a proxy to indicate an MCF, e.g., as described in section 7.4 of the "Deterministic Networking Use Cases" [RFC8578].

2.2.3. Reliability in the Context of RAW

In contrast with wired networks, errors in transmission are the predominant source of packet loss in wireless networks.

The root cause for the loss may be of multiple origins, calling for the use of different forms of diversity:

Multipath Fading: A destructive interference by a reflection of the original signal.

A radio signal may be received directly (line-of-sight) and/or as a reflection on a physical structure (echo). The reflections take a longer path and are delayed by the extra distance divided by the speed of light in the medium. Depending on the frequency, the echo lands with a different phase which may add up to (constructive interference) or cancel the direct signal (destructive interference).

The affected frequencies depend on the relative position of the sender, the receiver, and all the reflecting objects in the environment. A given hop will suffer from multipath fading for multiple packets in a row till the something moves that changes the reflection patterns.

Co-channel Interference: Energy in the spectrum used for the transmission confuses the receiver.

The wireless medium itself is a Shared Risk Link Group (SRLG) for nearby users of the same spectrum, as an interference may affect multiple co-channel transmissions between different peers within the interference domain of the interferer, possibly even when they use different technologies.

Obstacle in Fresnel Zone: The optimal transmission happens when the Fresnel Zone between the sender and the receiver is free of obstacles.

As long as a physical object (e.g., a metallic trolley between peers) that affects the transmission is not removed, the quality of the link is affected.

In an environment that is rich of metallic structures and mobile objects, a single radio link will provide a fuzzy service, meaning that it cannot be trusted to transport the traffic reliably over a long period of time.

Transmission losses are typically not independent, and their nature and duration are unpredictable; as long as a physical object (e.g., a metallic trolley between peers) that affects the transmission is not removed, or as long as the interferer (e.g., a radar) keeps transmitting, a continuous stream of packets will be affected.

The key technique to combat those unpredictable losses is diversity. Different forms of diversity are necessary to combat different causes of loss and the use of diversity must be maximised to optimize the PDR.

A single packet may be sent at different times (time diversity) over diverse paths (spatial diversity) that rely on diverse radio channels (frequency diversity) and diverse PHY technologies, e.g., narrowband vs. spread spectrum, or diverse codes. Using time diversity will defeat short-term interferences; spatial diversity combats very local causes such as multipath fading; narrowband and spread spectrum are relatively innocuous to one another and can be used for diversity in the presence of the other.


2.3. Use Cases and Requirements Served

In order to focus on real-worlds issues and assert the feasibility of the proposed capabilities, RAW focuses on selected technologies that can be scheduled at the lower layers: IEEE Std. 802.15.4 timeslotted channel hopping (TSCH), 3GPP 5G ultra-reliable low latency communications (URLLC), IEEE 802.11ax/be where 802.11be is extreme high throughput (EHT), and L-band Digital Aeronautical Communications System (LDACS). See [RAW-TECHNOS] for more.

"Deterministic Networking Use Cases" [RFC8578] presents a number of wireless use cases including Wireless, such as application to Industrial Applications, Pro-Audio, and SmartGrid Automation. [RAW-USE-CASES] adds a number of use cases that demonstrate the need for RAW capabilities for new applications such as Pro-Gaming and drones. The use cases can be abstracted in two families, Loose Protection, e.g., protecting the first hop in Radio Access Protection and Strict Protection, e.g., providing End-to-End Protection in a wireless mesh.

2.3.1. Radio Access Protection

To maintain the required SLA at all times, a wireless Host may use more than one Radio Access Network (RAN) in parallel.



*** = flapping at this time \$\$\$ expensive

Figure 1: Radio Access Protection


The RANs may be heterogeneous, e.g., 3GPP 5G [RAW-5G] and Wi-Fi [RAW-TECHNOS] for high-speed communication, in which case a Layer-3 abstraction becomes useful to select which of the RANs are used at a particular point of time, and the amount of traffic that is distributed over each RAN.

The idea is that the rest of the path to the destination(s) is protected separately (e.g., uses non-congruent paths, leverages DetNet / TSN, etc...) and is a lot more reliable, e.g., wired. In that case, RAW observes the reliability of the end-to-end operation through each of the RANs but only observes and controls the wireless operation the first hop.

A variation of that use case has a pair of wireless Hosts connected over a wired core / backbone network. In that case, RAW observes and controls the Ingress and Egress RANs, while neglecting the hops in the core. The resulting loose Track may be instantiated, e.g., using tunneling or loose source routing between the RANs.

2.3.2. End-to-End Protection in a Wireless Mesh

In radio technologies that support mesh networking (e.g., Wi-Fi and TSCH), a Track is a complex path with distributed PAREO capabilities. In that case, RAW operates through the multipath and makes decisions either at the Ingress or at every hop (more in Section 3.3).



zzz = flapping now

Figure 2: End-to-End Protection

The Protection may be imposed by the source based on end-to-end OAM, or performed hop-by-hop, in which case the OAM must enables the intermediate Nodes to estimate the quality of the rest of the feasible paths in the remainder of the Track to the destination.

2.4. Related Work at The IETF

RAW intersects with protocols or practices in development at the IETF as follows:

- * The Dynamic Link Exchange Protocol (DLEP) [RFC8175] from [MANET] can be leveraged at each hop to derive generic radio metrics (e.g., based on LQI, RSSI, queueing delays and ETX) on individual hops.
- * [detnet] provides an OAM framework with [DetNet-OAM] that applies within the DetNet dataplane described in [DetNet-DP], which is typically based on MPLS or IPv6 pseudowires.
- * [BFD] detect faults in the path between an Ingress and an Egress forwarding engines, but is unaware of the complexity of a path with replication, and expects bidirectionality. BFD asynchronous mode considers delivery as success whereas with DetNet and RAW, the bounded latency can be as important as the delivery itself, and delivering too late is actually a failure. Note that the BFD Demand mode with unsolicited notifications may be more suitable than the Asynchronous BFD mode. The use of the Demand mode in MPLS is analyzed in [I-D.mirsky-bfd-mpls-demand] and similar considerations could apply to IP as well.
- * [SPRING] and [BIER] define in-band signaling that influences the routing when decided at the head-end on the path. There's already one RAW-related draft at BIER [BIER-PREF] more may follow. RAW will need new in-band signaling when the decision is distributed, e.g., required chances of reliable delivery to destination within latency. This signaling enables relays to tune retries and replication to meet the required SLA.

- * [CCAMP] defines protocol-independent metrics and parameters (measurement attributes) for describing links and paths that are required for routing and signaling in technology-specific networks. RAW would be a source of requirements for CCAMP to define metrics that are significant to the focus radios.
- * [IPPM] develops and maintains standard metrics that can be applied to the quality, performance, and reliability of Internet data delivery services and applications running over transport layer protocols (e.g. TCP, UDP) over IP.

3. The RAW Framework

3.1. Scope and Prerequisites

A prerequisite to the RAW operation is that an end-to-end routing function computes a complex sub-topology along which forwarding can happen between a source and one or more destinations. The concept of Track is specified in the 6TiSCH Architecture [6TiSCH-ARCHI] to represent that complex sub-topology. Tracks provide a high degree of redundancy and diversity and enable the DetNet PREOF, network coding, and possibly RAW specific techniques such as PAREO, leveraging frequency diversity, time diversity, and possibly other forms of diversity as well.

How the routing operation (e.g., PCE) in the Controller Plane computes the Track is out of scope for RAW. The scope of the RAW operation is one Track, and the goal of the RAW operation is to optimize the use of the Track at the forwarding timescale to maintain the expected SLA while optimizing the usage of constrained resources such as energy and spectrum.

Another prerequisite is that an IP link can be established over the radio with some guarantees in terms of service reliability, e.g., it can be relied upon to transmit a packet within a bounded latency and provides a guaranteed BER/PDR outside rare but existing transient outage windows that can last from split seconds to minutes. The radio layer can be programmed with abstract parameters, and can return an abstract view of the state of the Link to help the Network Layer forwarding decision (think DLEP from MANET).

How the radio interface manages its lower layers is out of control and out of scope for RAW. In the same fashion, the non-RAW portion along a loose Track is by definition out of control and out of scope for RAW. Whether it is a single hop or a mesh is also unknown and out of scope.

3.2. Routing Time Scale vs. Forwarding Time Scale

With DetNet, the Controller Plane Function that handles the routing computation and maintenance (the PCE) can be centralized and can reside outside the network. In a wireless mesh, the path to the PCE can be expensive and slow, possibly going across the whole mesh and back. Reaching to the PCE can also be slow in regards to the speed of events that affect the forwarding operation at the radio layer.

Due to that cost and latency, the Controller Plane is not expected to be sensitive/reactive to transient changes. The abstraction of a link at the routing level is expected to use statistical metrics that aggregate the behavior of a link over long periods of time, and represent its properties as shades of gray as opposed to numerical values such as a link quality indicator, or a boolean value for either up or down.




Figure 3: Time Scales

In the case of wireless, the changes that affect the forwarding decision can happen frequently and often for short durations, e.g., a mobile object moves between a transmitter and a receiver, and will cancel the line of sight transmission for a few seconds, or a radar measures the depth of a pool and interferes on a particular channel for a split second.

There is thus a desire to separate the long term computation of the route and the short term forwarding decision. In that model, the routing operation computes a complex Track that enables multiple Non-Equal Cost Multi-Path (N-ECMP) forwarding solutions, and leaves it to the Data Plane to make the per-packet decision of which of these possibilities should be used.

In the wired world, and more specifically in the context of Traffic Engineering (TE), an alternate path can be used upon the detection of a failure in the main path, e.g., using OAM in MPLS-TP or BFD over a collection of SD-WAN tunnels. RAW formalizes a forwarding time scale that is an order(s) of magnitude shorter than the controller plane routing time scale, and separates the protocols and metrics that are used at both scales. Routing can operate on long term statistics such as delivery ratio over minutes to hours, but as a first approximation can ignore flapping. On the other hand, the RAW forwarding decision is made at the scale of the packet rate, and uses information that must be pertinent at the present time for the current transmission(s).

3.3. Wireless Tracks

The "6TiSCH Architecture" [6TiSCH-ARCHI] introduces the concept of Track. RAW extends the concept to any wireless mesh technology, including, e.g., Wi-Fi. A simple Track is composed of a direct sequence of reserved hops to ensure the transmission of a single packet from a source Node to a destination Node across a multihop path.

A Complex Track provides multiple N-ECMP forwarding solutions. The Complex Track enables to support multi-path redundant forwarding by employing PRE functions [RFC8655] and the ingress and within the Track. For example, a Complex Track may branch off and rejoin over non-congruent segments.

In the context of RAW, some links or segments in the Track may be reversible, meaning that they can be used in either direction. In that case, an indication in the packet signals the direction of the reversible links or segments that the packet traverses and thus places a constraint that prevents loops from occurring. An individual packet follows a destination-oriented directed acyclic graph (DODAG) towards a destination Node inside the Complex Track.

3.4. PAREO Functions

RAW may control whether and how to use packet replication and elimination (PRE), Automatic Repeat reQuest (ARQ), Hybrid ARQ (HARQ) that includes Forward Error Correction (FEC) and coding, and other wireless-specific techniques such as overhearing and constructive interferences, in order to increase the reliability and availability of the end-to-end transmission.

Collectively, those function are called PAREO for Packet (hybrid) ARQ, Replication, Elimination and Ordering. By tuning dynamically the use of PAREO functions, RAW avoids the waste of critical resources such as spectrum and energy while providing that the guaranteed SLA, e.g., by adding redundancy only when a spike of loss is observed.

In a nutshell, PAREO establishes several paths in a network to provide redundancy and parallel transmissions to bound the end-to-end delay to traverse the network. Optionally, promiscuous listening between paths is possible, such that the Nodes on one path may overhear transmissions along the other path. Considering the scenario shown in Figure 4, many different paths are possible for to traverse the network from ingress to egress. A simple way to benefit from this topology could be to use the two independent paths via Nodes A, C, E and via B, D, F. But more complex paths are possible by interleaving transmissions from the lower level of the path to the upper level.




Figure 4: A Ladder Shape with Two Parallel Paths

PAREO may also take advantage of the shared properties of the wireless medium to compensate for the potential loss that is incurred with radio transmissions.

For instance, when the source sends to Node A, Node B may listen promiscuously and get a second chance to receive the frame without an additional transmission. Note that B would not have to listen if it already received that particular frame at an earlier timeslot in a dedicated transmission towards B.

The PAREO model can be implemented in both centralized and distributed scheduling approaches. In the centralized approach, a Path Computation Element (PCE) scheduler calculates a Track and schedules the communication. In the distributed approach, the Track is computed within the network, and signaled in the packets, e.g., using BIER-TE, Segment Routing, or a Source Routing Header.

3.4.1. Packet Replication

By employing a Packet Replication procedure, a Node forwards a copy of each data packet to more than one successor. To do so, each Node (i.e., Ingress and intermediate Node) sends the data packet multiple times as separate unicast transmissions. For instance, in Figure 5, the Ingress Node is transmitting the packet to both successors, nodes A and B, at two different times.




Figure 5: Packet Replication

An example schedule is shown in Table 1. This way, the transmission leverages with the time and spatial forms of diversity.

Channel	0	1	2	3	4	5	6
0	S->A	S->B	B->C	B->D	C->F	E->R	F->R
1		A->C	A->D	C->E	D->E	D->F	

Table 1: Packet Replication: Sample schedule

3.4.2. Packet Elimination

The replication operation increases the traffic load in the network, due to packet duplications. This may occur at several stages inside the Track, and to avoid an explosion of the number of copies, a Packet Elimination procedure must be applied as well. To this aim, once a Node receives the first copy of a data packet, it discards the subsequent copies.

The logical functions of Replication and Elimination may be collocated in an intermediate Node, the Node first eliminating the redundant copies and then sending the packet exactly once to each of the selected successors.

3.4.3. Promiscuous Overhearing

Considering that the wireless medium is broadcast by nature, any neighbor of a transmitter may overhear a transmission. By employing the Promiscuous Overhearing operation, the next hops have additional opportunities to capture the data packets. In Figure 6, when Node A is transmitting to its DP (Node C), the AP (Node D) and its sibling (Node B) may decode this data packet as well. As a result, by employing correlated paths, a Node may have multiple opportunities to receive a given data packet.




Figure 6: Unicast with Overhearing

3.4.4. Constructive Interference

Constructive Interference can be seen as the reverse of Promiscuous Overhearing, and refers to the case where two senders transmit the exact same signal in a fashion that the emitted symbols add up at the receiver and permit a reception that would not be possible with a single sender at the same PHY mode and the same power level.

Constructive Interference was proposed on 5G, Wi-Fi7 and even tested on IEEE Std 802.14.5. The hard piece is to synchronize the senders to the point that the signals are emitted at slightly different time to offset the difference of propagation delay that corresponds to the difference of distance of the transmitters to the receiver at the speed of light to the point that the symbols are superposed long enough to be recognizable.

4. The RAW Architecture

4.1. The RAW Conceptual Model

RAW inherits the conceptual model described in section 4 of the DetNet Architecture [RFC8655]. RAW extends the DetNet service layer to provide additional agility against transmission loss.

A RAW Network Plane may be strict or loose, depending on whether RAW observes and takes actions on all hops or not. For instance, the packets between two wireless entities may be relayed over a wired infrastructure such as a Wi-Fi extended service set (ESS) or a 5G Core; in that case, RAW observes and control the transmission over the wireless first and last hops, as well as end-to-end metrics such as latency, jitter, and delivery ratio. This operation is loose since the structure and properties of the wired infrastructure are ignored, and may be either controlled by other means such as DetNet/TSN, or neglected in the face of the wireless hops.

A Controller Plane Function (CPF) called the Path Computation Element (PCE) [RFC4655] interacts with RAW Nodes over a Southbound API. The RAW Nodes are DetNet relays that are capable of additional diversity mechanisms and measurement functions related to the radio interface, in particular the PAREO diversity mechanisms.

The PCE defines a complex Track between an Ingress End System and an Egress End System, and indicates to the RAW Nodes where the PAREO operations may be actionned in the Network Plane. The Track may be expressed loosely to enable traversing a non-RAW subnetwork. In that case, the expectation is that the non-RAW subnetwork can be neglected in the RAW computation, that is, considered infinitely fast, reliable and/or available in comparison with the links between RAW nodes.

```

CPF                                CPF                                CPF                                CPF
Southbound API
-----
Ingress  --z  RAW  --z  RAW  --z  RAW  --z  RAW  --z
          z-- Node z-- Node z-- Node z-- Node --z
End      \      \      . . .      End
Node    ---z  /      /      . . .      z-- Node
          z-- RAW  --z  RAW      ( non-RAW ) -- RAW  --z
          Node  z-- Node --- ( Nodes )      Node
          wireless      wired
          z-- link      --- link

```

Figure 7: RAW Nodes

The Link-Layer metrics are reported to the PCE in a time-aggregated, e.g., statistical fashion. Example Link-Layer metrics include typical Link bandwidth (the medium speed depends dynamically on the PHY mode and the number of users sharing the spectrum) and average and mean squared deviation of availability and reliability figures such as Packet Delivery Ratio (PDR) over long periods of time.

Based on those metrics, the PCE installs the Track with enough redundant forwarding solutions to ensure that the Network Plane can reliably deliver the packets within a System Level Agreement (SLA) associated to the flows that it transports. The SLA defines end-to-end reliability and availability requirements, where reliability may be expressed as a successful delivery in order and within a bounded delay of at least one copy of a packet.

Depending on the use case and the SLA, the Track may comprise non-RAW segments, either interleaved inside the Track, or all the way to the Egress End Node (e.g., a server in the Internet). RAW observes the Lower-Layer Links between RAW nodes (typically, radio links) and the end-to-end Network Layer operation to decide at all times which of the PAREO diversity schemes is actioned by which RAW Nodes.

Once a Track is established, per-segment and end-to-end reliability and availability statistics are periodically reported to the PCE to assure that the SLA can be met or have it recompute the Track if not.

4.2. The Path Selection Engine

RAW separates the path computation time scale at which a complex path is recomputed from the path selection time scale at which the forwarding decision is taken for one or a few packets (more in Section 3.2). RAW operates at the path selection time scale. The RAW problem is to decide, within the redundant solutions that are proposed by the PCE, which will be used for each packet to provide a Reliable and Available service while minimizing the waste of constrained resources.

To that effect, RAW defines the Path Selection Engine (PSE) that is the counter-part of the PCE to perform rapid local adjustments of the forwarding tables within the diversity that the PCE has selected for the Track. The PSE enables to exploit the richer forwarding capabilities with PAREO and scheduled transmissions at a faster time scale over the smaller domain that is the Track, in either a loose or a strict fashion.

Compared to the PCE, the PSE operates on metrics that evolve faster, but that needs to be advertised at a fast rate but only locally, within the Track. The forwarding decision may also change rapidly, but with a scope that is also contained within the Track, with no visibility to the other Tracks and flows in the network. This is as opposed to the PCE that needs to observe the whole network, and optimize all the Tracks globally, which can only be done at a slow pace and using long-term statistical metrics, as presented in Table 2.

	PCE (Not in Scope)	PSE (In Scope)
Operation	Centralized	Source-Routed or Distributed
Communication	Slow, expensive	Fast, local
Time Scale	hours and above	seconds and below
Network Size	Large, many Tracks to optimize globally	Small, within one Track
Considered Metrics	Averaged, Statistical, Shade of grey	Instant values / boolean condition

Table 2: PCE vs. PSE

The PSE sits in the DetNet Service sub-Layer of Edge and Relay Nodes. On the one hand, it operates on the packet flow, learning the Track and path selection information from the packet, possibly making local decision and retagging the packet to indicate so. On the other hand, the PSE interacts with the lower layers and with its peers to obtain up-to-date information about its radio links and the quality of the overall Track, respectively, as illustrated in Figure 8.

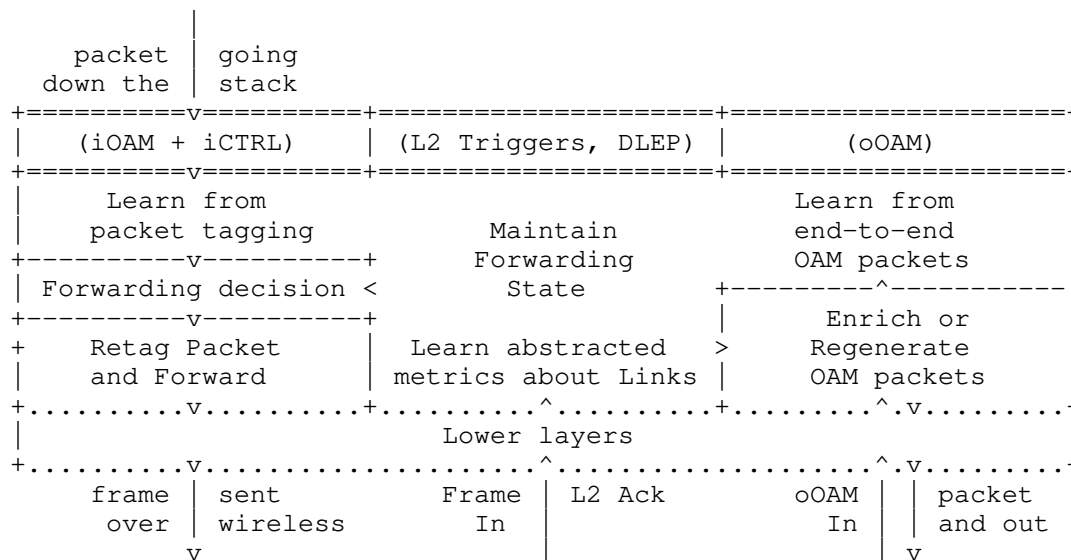


Figure 8: PSE

4.3. RAW OAM

RAW In-situ OAM operation in the Network Plane may observe either a full Track or subTracks that are being used at this time. Active RAW OAM may be needed to observe the unused segments and evaluate the desirability of a rerouting decision. Finally, the RAW Service Layer Assurance may observe the individual PAREO operation of a relay node to ensure that it is conforming; this might require injecting an OAM packet at an upstream point inside the Track and extracting that packet at another point downstream before it reaches the egress.

This observation feeds the RAW PSE that makes the decision on which PAREO function is actioned at which RAW Node, for one a small continuous series of packets.

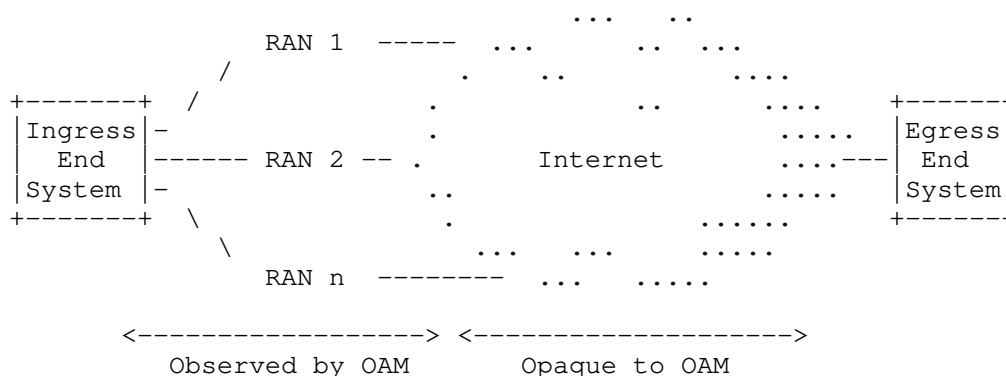


Figure 9: Observed Links in Radio Access Protection

In the case of a End-to-End Protection in a Wireless Mesh, the Track is strict and congruent with the path so all links are observed. Conversely, in the case of Radio Access Protection, the Track is Loose and in that case only the first hop is observed; the rest of the path is abstracted and considered infinitely reliable.

In the case of the Radio Access Protection, only the first hop is protected; the loss of a packet that was sent over one of the possible first hops is attributed to that first hop, even if a particular loss effectively happens farther down the path.

The Links that are not observed by OAM are opaque to it, meaning that the OAM information is carried across and possibly echoed as data, but there is no information capture in intermediate nodes. In the example above, the Internet is opaque and not controlled by RAW; still the RAW OAM measures the end-to-end latency and delivery ratio for packets sent via each if RAN 1, RAN 2 and RAN 3, and determines whether a packet should be sent over either or a collection of those access links.

4.3.1. DetNet OAM

[detnet] provides an OAM framework with [DetNet-OAM] that applies within the DetNet dataplane described in [DetNet-DP], which is typically based on MPLS or IPv6 pseudowires. How the framework applies to IPv6 is detailed in [DetNet-IP-OAM]. Within that framework, OAM messages follow the same forward path as the data packets and gather information about their individual treatment at each hop. When the destination receives an OAM message, it gets a view on the full path or at least of a segment of the path from the source of the flow.

In-situ OAM (IOAM) adds telemetry information about the experience of one packet within the packet itself [I-D.ietf-ippm-ioam-data], with the caveats that the measurement and the consecutive update of the packet interfere with the operation being observed, e.g., may increase the latency of the packet for which it is measured and into which it is stamped.

Note: IOAM and analogous on-path telemetry methods are capable of facilitating collection of useful telemetry information that characterizes the state of a system as experienced by the packet. But because of statistical character of a packet network, these methods may not be used to monitor the continuity of a path (Track) or proper connectivity of the Track (no leaking packets across Tracks).

This effect can be alleviated by measuring on the fly but reporting later, e.g., by exporting the data as a separate management packet [I-D.ietf-ippm-ioam-direct-export]. [I-D.mirsky-ippm-hybrid-two-step] proposes an hybrid two-steps method (HTS) where a trigger message starts the measurement and a follow up along the Track packet gathers the measured data.

"Error Performance Measurement" [I-D.mirsky-ippm-epm] uses Fault Management (FM) and Performance Management (PM) OAM mechanisms to determine availability/unavailability of a path according to predefined SLA.

4.3.2. RAW Extensions

Classical OAM typically measures information at the transmitter, e.g., residence time in the node or transmit queue size. With RAW, there is a need to combine information at the sender (number of retries) with that at the receiver (LQI, RSSI). This doubles the operating cost of an IOAM processing that would gather the experience of a single packet.

The RAW PSE may be centralized at the Track Ingress, or distributed long the Track. Either way, the PSE needs instant information about the rest of the way to the destination over the possible next-hop adjacencies along the Track in order to decide how to perform simple forwarding, load balancing, and/or replication, as well as determining how much latency credit is available for ARQ.

To provide that information timely, it makes sense that the OAM packets that gather instantaneous values from the radio senders and receivers at each hop flow on the reverse path and inform the PSE at the source and/or the PAREO relays about the state of the rest of the way. This is achieved using Reverse OAM packets that flow along the Reversed Track, West to East.

Because the quality of transmission over a wireless medium varies continuously, it is important that RAW OAM captures the state of the medium across an adjacency over multiple transmission and over a recent period of time, whether the transmitted packets belong to this flow or another. Some of the measured information relates to the medium itself. In other words, the captured information does not only relate to the experience of one packet as is the case for IOAM, but also to the medium itself. This makes an approach like HTS more suitable as it can trigger the capture of multiple measurements over a short period of time. On the other hand, the PSE needs a continuous measurement stream where a single trigger is followed by a periodic follow up capture.

In other words, the best suited OAM method to enable the PSE make accurate PAREO forwarding decisions is a periodic variation of the two-steps method flowing along the reverse Track, as a Reverse OAM technique. [RAW-OAM] provides more information on the RAW OAM problem and solution approaches.

4.3.3. Observed Metrics

The Dynamic Link Exchange Protocol (DLEP) [RFC8175] from [MANET] can be leveraged at each hop to derive generic radio metrics (e.g., based on LQI, RSSI, queueing delays and ETX) on individual hops.

Those lower-layer metrics are aggregated along a multihop segment into abstract layer 3 information that reflect the instant reliability and latency of the observed path.

4.4. Flow Identification vs. Path Identification

Section 4.7 of the DetNet Architecture [RFC8655] ties the app-flow identification which is an application layer concept with the network path identification that depends on the networking technology by "exporting of flow identification", e.g., to a MPLS label.

With RAW, this exporting operation is injective but not bijective. e.g., a flow is fully placed within one RAW Track, but not all packets along that Track are necessarily part of the same flow. For instance, out-of-band OAM packets must circulate in the exact same fashion as the flows that they observe. It results that the flow

identification that maps to to app-flow at the network layer must be separate from the path identification that is used to forward a packet.

Section 3.4 of the DetNet data-plane framework [DetNet-DP] indicates that for a DetNet IP Data Plane, a flow is identified by an IPv6 6-tuple. With RAW, that 6-tuple is not what indicates the Track, in other words, the flow ID is not the Track ID.

For instance, the 6TiSCH Architecture [6TiSCH-ARCHI] uses a combination of the address of the Egress End System and an instance identifier in a Hop-by-hop option to indicate a Track. This way, if a packet "escapes" the Track, it will reach the Track Egress point through normal routing and be treated at the service layer through, say, elimination and reordering.

The RAW service includes forwarding over a subset of the Links that form the Track (a subTrack). Packets from the same or a different flow that are routed through the same Track will not necessarily traverse the same Links. The PSE selects a subTrack for a packet based on the links that are preferred and those that should be avoided at this time.

Each packet is forwarded within the subTrack that provides the best adequation with the SLA of the flow and the energy and bandwidth constraints of the network.




Figure 10: Flow Injection

With 6TiSCH, packets are tagged with the same (destination address, instance ID) will experience the same RAW service regardless of the IPv6 6-tuple that indicates the flow. The forwarding does not depend on whether the packets transport application flows or OAM. In the generic case, the Track or the subTrack can be signaled in the packet through other means, e.g., encoded in the suffix of the destination address as a Segment Routing Service Instruction [SR-ARCHI], or leveraging Bit Index Explicit Replication [BIER] Traffic Engineering [BIER-TE].

4.5. Source-Routed vs. Distributed Forwarding Decision

Within a large routed topology, the route-over mesh operation builds a particular complex Track with one source and one or more destinations; within the Track, packets may follow different paths and may be subject to RAW forwarding operations that include replication, elimination, retries, overhearing and reordering.

The RAW forwarding decisions include the selection of points of replication and elimination, how many retries can take place, and a limit of validity for the packet beyond which the packet should be destroyed rather than forwarded uselessly further down the Track.

The decision to apply the RAW techniques must be done quickly, and depends on a very recent and precise knowledge of the forwarding conditions within the complex Track. There is a need for an observation method to provide the RAW Data Plane with the specific knowledge of the state of the Track for the type of flow of interest (e.g., for a QoS level of interest). To observe the whole Track in quasi real time, RAW considers existing tools such as L2-triggers, DLEP, BFD and leverages in-band and out-of-band OAM to capture and report that information to the PSE.

One possible way of making the RAW forwarding decisions within a Track is to position a unique PSE at the Ingress and express its decision in-band in the packet, which requires the explicit signaling of the subTrack within the Track. In that case, the RAW forwarding operation along the Track is encoded by the source, e.g., by indicating the subTrack in the Segment Routing (SRv6) Service Instruction, or by leveraging BIER-TE such as done with [BIER-PREF].

The alternate way is to operate the PSE in each forwarding Node, which makes the RAW forwarding decisions for a packet on its own, based on its knowledge of the expectation (timeliness and reliability) for that packet and a recent observation of the rest of the way across the possible paths based on OAM. Information about the desired service should be placed in the packet and matched with the forwarding Node's capabilities and policies.

In either case, a per-track/subTrack state is installed in all the intermediate Nodes to recognize the packets that are following a Track and determine the forwarding operation to be applied.

4.6. Encapsulation and Decapsulation

In the generic case where the Track Ingress Node is not the source of the Packet, the Ingress Node needs to encapsulate IP-in-IP to ensure that the Destination IP Address is that of the Egress Node and that the necessary Headers (Routing Header, Segment Routing Header and/or Hop-By-Hop Header) can be added to the packet to signal the Track or the subTrack, conforming [IPv6] that discourages the insertion of a Header on the fly.

In the specific case where the Ingress Node is the source of the packet, the encapsulation can be avoided, provided that the source adds the necessary headers and that the destination is set to the Egress Node. Forwarding to a final destination beyond the Egress Node is possible, e.g., with a Segment Routing Header that signals the rest of the way. In that case a Hop-by-Hop Header is not recommended since its validity is within the Track only.

5. Security Considerations

RAW uses all forms of diversity including radio technology and physical path to increase the reliability and availability in the face of unpredictable conditions. While this is not done specifically to defeat an attacker, the amount of diversity used in RAW makes an attack harder to achieve.

5.1. Forced Access

RAW will typically select the cheapest collection of links that matches the requested SLA, for instance, leverage free WI-Fi vs. paid 3GPP access. By defeating the cheap connectivity (e.g., PHY-layer interference) the attacker can force an End System to use the paid access and increase the cost of the transmission for the user.

6. IANA Considerations

This document has no IANA actions.

7. Contributors

The editor wishes to thank:

Xavi Vilajosana: Wireless Networks Research Lab, Universitat Oberta de Catalunya

Remous-Aris Koutsiamanis: IMT Atlantique

Nicolas Montavont: IMT Atlantique

Rex Buddenberg: Individual contributor

Greg Mirsky: ZTE

for their contributions to the text and ideas exposed in this document.

8. Acknowledgments

TBD

9. References

9.1. Normative References

[6TiSCH-ARCHI]

Thubert, P., Ed., "An Architecture for IPv6 over the Time-Slotted Channel Hopping Mode of IEEE 802.15.4 (6TiSCH)", RFC 9030, DOI 10.17487/RFC9030, May 2021, <<https://www.rfc-editor.org/info/rfc9030>>.

[RAW-TECHNOS]

Thubert, P., Cavalcanti, D., Vilajosana, X., Schmitt, C., and J. Farkas, "Reliable and Available Wireless Technologies", Work in Progress, Internet-Draft, draft-ietf-raw-technologies-01, 19 February 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-raw-technologies-01>>.

[RAW-USE-CASES]

Papadopoulos, G. Z., Thubert, P., Theoleyre, F., and C. J. Bernardos, "RAW use cases", Work in Progress, Internet-Draft, draft-ietf-raw-use-cases-01, 21 February 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-raw-use-cases-01>>.

[RFC4655]

Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <<https://www.rfc-editor.org/info/rfc4655>>.

[BFD]

Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <<https://www.rfc-editor.org/info/rfc5880>>.

- [RFC6291] Andersson, L., van Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "Guidelines for the Use of the "OAM" Acronym in the IETF", BCP 161, RFC 6291, DOI 10.17487/RFC6291, June 2011, <<https://www.rfc-editor.org/info/rfc6291>>.
- [RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <<https://www.rfc-editor.org/info/rfc7799>>.
- [RFC8578] Grossman, E., Ed., "Deterministic Networking Use Cases", RFC 8578, DOI 10.17487/RFC8578, May 2019, <<https://www.rfc-editor.org/info/rfc8578>>.
- [IPv6] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [SR-ARCHI] Filssils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [BIER] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", RFC 8279, DOI 10.17487/RFC8279, November 2017, <<https://www.rfc-editor.org/info/rfc8279>>.
- [RFC8175] Ratliff, S., Jury, S., Satterwhite, D., Taylor, R., and B. Berry, "Dynamic Link Exchange Protocol (DLEP)", RFC 8175, DOI 10.17487/RFC8175, June 2017, <<https://www.rfc-editor.org/info/rfc8175>>.
- [RFC8557] Finn, N. and P. Thubert, "Deterministic Networking Problem Statement", RFC 8557, DOI 10.17487/RFC8557, May 2019, <<https://www.rfc-editor.org/info/rfc8557>>.
- [RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <<https://www.rfc-editor.org/info/rfc8655>>.
- [RFC9049] Dawkins, S., Ed., "Path Aware Networking: Obstacles to Deployment (A Bestiary of Roads Not Taken)", RFC 9049, DOI 10.17487/RFC9049, June 2021, <<https://www.rfc-editor.org/info/rfc9049>>.

9.2. Informative References

- [RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.
- [TE] Awduche, D., Chiu, A., Elwalid, A., Widjaja, I., and X. Xiao, "Overview and Principles of Internet Traffic Engineering", RFC 3272, DOI 10.17487/RFC3272, May 2002, <<https://www.rfc-editor.org/info/rfc3272>>.
- [STD 62] Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks", STD 62, RFC 3411, DOI 10.17487/RFC3411, December 2002, <<https://www.rfc-editor.org/info/rfc3411>>.
- [RFC4090] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, DOI 10.17487/RFC4090, May 2005, <<https://www.rfc-editor.org/info/rfc4090>>.
- [FRR] Shand, M. and S. Bryant, "IP Fast Reroute Framework", RFC 5714, DOI 10.17487/RFC5714, January 2010, <<https://www.rfc-editor.org/info/rfc5714>>.
- [RLFA-FRR] Bryant, S., Filsfils, C., Previdi, S., Shand, M., and N. So, "Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)", RFC 7490, DOI 10.17487/RFC7490, April 2015, <<https://www.rfc-editor.org/info/rfc7490>>.
- [DetNet-DP] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., and S. Bryant, "Deterministic Networking (DetNet) Data Plane Framework", RFC 8938, DOI 10.17487/RFC8938, November 2020, <<https://www.rfc-editor.org/info/rfc8938>>.
- [BIER-PREF] Thubert, P., Eckert, T., Brodard, Z., and H. Jiang, "BIER-TE extensions for Packet Replication and Elimination Function (PREF) and OAM", Work in Progress, Internet-Draft, draft-thubert-bier-replication-elimination-03, 3 March 2018, <<https://datatracker.ietf.org/doc/html/draft-thubert-bier-replication-elimination-03>>.
- [DetNet-IP-OAM] Mirsky, G., Chen, M., and D. Black, "Operations, Administration and Maintenance (OAM) for Deterministic

Networks (DetNet) with IP Data Plane", Work in Progress, Internet-Draft, draft-ietf-detnet-ip-oam-02, 30 March 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-detnet-ip-oam-02>>.

[RAW-5G] Farkas, J., Dudda, T., Shapin, A., and S. Sandberg, "5G - Ultra-Reliable Wireless Technology with Low Latency", Work in Progress, Internet-Draft, draft-farkas-raw-5g-00, 1 April 2020, <<https://datatracker.ietf.org/doc/html/draft-farkas-raw-5g-00>>.

[BIER-TE] Eckert, T., Cauchie, G., and M. Menth, "Tree Engineering for Bit Index Explicit Replication (BIER-TE)", Work in Progress, Internet-Draft, draft-ietf-bier-te-arch-09, 30 October 2020, <<https://datatracker.ietf.org/doc/html/draft-ietf-bier-te-arch-09>>.

[IPoWIRELESS]

Thubert, P., "IPv6 Neighbor Discovery on Wireless Networks", Work in Progress, Internet-Draft, draft-thubert-6man-ipv6-over-wireless-09, 17 May 2021, <<https://datatracker.ietf.org/doc/html/draft-thubert-6man-ipv6-over-wireless-09>>.

[RAW-OAM] Theoleyre, F., Papadopoulos, G. Z., Mirsky, G., and C. J. Bernardos, "Operations, Administration and Maintenance (OAM) features for RAW", Work in Progress, Internet-Draft, draft-ietf-raw-oam-support-02, 3 June 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-raw-oam-support-02>>.

[I-D.ietf-ippm-ioam-direct-export]

Song, H., Gafni, B., Zhou, T., Li, Z., Brockners, F., Bhandari, S., Sivakolundu, R., and T. Mizrahi, "In-situ OAM Direct Exporting", Work in Progress, Internet-Draft, draft-ietf-ippm-ioam-direct-export-03, 17 February 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-ippm-ioam-direct-export-03>>.

[DetNet-OAM]

Mirsky, G., Theoleyre, F., Papadopoulos, G. Z., and C. J. Bernardos, "Framework of Operations, Administration and Maintenance (OAM) for Deterministic Networking (DetNet)", Work in Progress, Internet-Draft, draft-ietf-detnet-oam-framework-01, 19 May 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-detnet-oam-framework-01>>.

- [I-D.mirsky-ippm-hybrid-two-step]
Mirsky, G., Lingqiang, W., Zhui, G., and H. Song, "Hybrid Two-Step Performance Measurement Method", Work in Progress, Internet-Draft, draft-mirsky-ippm-hybrid-two-step-09, 30 March 2021, <<https://datatracker.ietf.org/doc/html/draft-mirsky-ippm-hybrid-two-step-09>>.
- [I-D.mirsky-ippm-epm]
Mirsky, G., Min, X., and L. Han, "Error Performance Measurement in Packet-switched Networks", Work in Progress, Internet-Draft, draft-mirsky-ippm-epm-03, 26 March 2021, <<https://datatracker.ietf.org/doc/html/draft-mirsky-ippm-epm-03>>.
- [I-D.mirsky-bfd-mpls-demand]
Mirsky, G., "BFD in Demand Mode over Point-to-Point MPLS LSP", Work in Progress, Internet-Draft, draft-mirsky-bfd-mpls-demand-09, 30 March 2021, <<https://datatracker.ietf.org/doc/html/draft-mirsky-bfd-mpls-demand-09>>.
- [I-D.ietf-ippm-ioam-data]
Brockners, F., Bhandari, S., and T. Mizrahi, "Data Fields for In-situ OAM", Work in Progress, Internet-Draft, draft-ietf-ippm-ioam-data-12, 21 February 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-ippm-ioam-data-12>>.
- [NASA] Adams, T., "RELIABILITY: Definition & Quantitative Illustration", <<https://ksdcdms.ksc.nasa.gov/Reliability/Documents/150814-3bWhatIsReliability.pdf>>.
- [MANET] IETF, "Mobile Ad hoc Networking", <<https://dataTracker.ietf.org/doc/charter-ietf-manet/>>.
- [detnet] IETF, "Deterministic Networking", <<https://dataTracker.ietf.org/doc/charter-ietf-detnet/>>.
- [SPRING] IETF, "Source Packet Routing in Networking", <<https://dataTracker.ietf.org/doc/charter-ietf-spring/>>.
- [BIER] IETF, "Bit Indexed Explicit Replication", <<https://dataTracker.ietf.org/doc/charter-ietf-bier/>>.
- [BFD] IETF, "Bidirectional Forwarding Detection", <<https://dataTracker.ietf.org/doc/charter-ietf-bfd/>>.

[CCAMP] IETF, "Common Control and Measurement Plane",
<<https://dataTracker.ietf.org/doc/charter-ietf-ccamp/>>.

[IPPM] IETF, "IP Performance Measurement",
<<https://dataTracker.ietf.org/doc/charter-ietf-ippm/>>.

Authors' Addresses

Pascal Thubert (editor)
Cisco Systems, Inc
Building D
45 Allee des Ormes - BP1200
06254 MOUGINS - Sophia Antipolis
France

Phone: +33 497 23 26 34
Email: pthubert@cisco.com

Georgios Z. Papadopoulos
IMT Atlantique
Office B00 - 114A
2 Rue de la Chataigneraie
35510 Cesson-Sevigne - Rennes
France

Phone: +33 299 12 70 04
Email: georgios.papadopoulos@imt-atlantique.fr

Lou Berger
LabN Consulting, L.L.C.
United States of America

Email: lberger@labn.net

RAW
Internet-Draft
Intended status: Standards Track
Expires: April 28, 2021

F. Theoleyre
CNRS
G. Papadopoulos
IMT Atlantique
G. Mirsky
ZTE Corp.
October 25, 2020

Operations, Administration and Maintenance (OAM) features for RAW
draft-theoleyre-raw-oam-support-04

Abstract

Some critical applications may use a wireless infrastructure. However, wireless networks exhibit a bandwidth of several orders of magnitude lower than wired networks. Besides, wireless transmissions are lossy by nature; the probability that a packet cannot be decoded correctly by the receiver may be quite high. In these conditions, guaranteeing the network infrastructure works properly is particularly challenging, since we need to address some issues specific to wireless networks. This document lists the requirements of the Operation, Administration, and Maintenance (OAM) features recommended to construct a predictable communication infrastructure on top of a collection of wireless segments. This document describes the benefits, problems, and trade-offs for using OAM in wireless networks to achieve Service Level Objectives (SLO).

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 28, 2021.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Terminology	4
1.2. Acronyms	5
1.3. Requirements Language	5
2. Role of OAM in RAW	5
2.1. Link concept and quality	6
2.2. Broadcast Transmissions	6
2.3. Complex Layer 2 Forwarding	7
3. Operation	7
3.1. Information Collection	7
3.2. Continuity Check	7
3.3. Connectivity Verification	7
3.4. Route Tracing	8
3.5. Fault Verification/detection	8
3.6. Fault Isolation/identification	8
4. Administration	9
4.1. Worst-case metrics	9
4.2. Efficient data retrieval	10
5. Maintenance	10
5.1. Dynamic Resource Reservation	11
5.2. Reliable Reconfiguration	11
6. IANA Considerations	11
7. Security Considerations	11
8. Acknowledgments	11
9. Informative References	11
Authors' Addresses	13

1. Introduction

Reliable and Available Wireless (RAW) is an effort that extends DetNet to approach end-to-end deterministic performances over a network that includes scheduled wireless segments. In wired networks, many approaches try to enable Quality of Service (QoS) by implementing traffic differentiation so that routers handle each type of packets differently. However, this differentiated treatment was expensive for most applications.

Deterministic Networking (DetNet) [RFC8655] has proposed to provide a bounded end-to-end latency on top of the network infrastructure, comprising both Layer 2 bridged and Layer 3 routed segments. Their work encompasses the data plane, OAM, time synchronization, management, control, and security aspects.

However, wireless networks create specific challenges. First of all, radio bandwidth is significantly lower than for wired networks. In these conditions, the volume of signaling messages has to be very limited. Even worse, wireless links are lossy: a layer 2 transmission may or may not be decoded correctly by the receiver, depending on a broad set of parameters. Thus, providing high reliability through wireless segments is particularly challenging.

Wired networks rely on the concept of links. All the devices attached to a link receive any transmission. The concept of a link in wireless networks is somewhat different from what many are used to in wireline networks. A receiver may or may not receive a transmission, depending on the presence of a colliding transmission, the radio channel's quality, and the external interference. Besides, a wireless transmission is broadcast by nature: any neighboring device may be able to decode it. The document includes detailed information on what the implications for the OAM features are.

Last but not least, radio links present volatile characteristics. If the wireless networks use an unlicensed band, packet losses are not anymore temporally and spatially independent. Typically, links may exhibit a very bursty characteristic, where several consecutive packets may be dropped. Thus, providing availability and reliability on top of the wireless infrastructure requires specific Layer 3 mechanisms to counteract these bursty losses.

Operations, Administration, and Maintenance (OAM) Tools are of primary importance for IP networks [RFC7276]. It defines a toolset for fault detection, isolation, and performance measurement.

The primary purpose of this document is to detail the specific requirements of the OAM features recommended to construct a

predictable communication infrastructure on top of a collection of wireless segments. This document describes the benefits, problems, and trade-offs for using OAM in wireless networks to provide availability and predictability.

In this document, the term OAM will be used according to its definition specified in [RFC6291]. We expect to implement an OAM framework in RAW networks to maintain a real-time view of the network infrastructure, and its ability to respect the Service Level Objectives (SLO), such as delay and reliability, assigned to each data flow.

1.1. Terminology

We re-use here the same terminology as [detnet-oam]:

- o OAM entity: a data flow to be controlled;
- o Maintenance End Point (MEP): OAM devices crossed when entering/exiting the network. In RAW, it corresponds mostly to the source or destination of a data flow. OAM message can be exchanges between two MEPs;
- o Maintenance Intermediate endPoint (MIP): OAM devices along the flow; OAM messages can be exchanged between a MEP and a MIP;
- o control/data plane: while the control plane expects to configure and control the network (long-term), the data plane takes the individual decision;
- o passive / active methods (as defined in [RFC7799]): active methods send additionnal control information (inserting novel fields, generating novel control packets). Passive methods infer information just by observing unmodified existing flows.
- o active methods may implement one of these two strategies:
 - * In-band: control information follows the same path as the data packets. In other words, a failure in the data plane may prevent the control information to reach the destination (e.g., end-device or controller).
 - * out-of-band: control information is sent separately from the data packets. Thus, the behavior of control vs. data packets may differ;

We also adopt the following terminology, which is particularly relevant for RAW segments.

- o piggybacking vs. dedicated control packets: control information may be encapsulated in specific (dedicated) control packets. Alternatively, it may be piggybacked in existing data packets, when the MTU is larger than the actual packet length. Piggybacking makes specifically sense in wireless networks: the cost (bandwidth and energy) is not linear with the packet size.
- o router-over vs. mesh under: a control packet is either forwarded directly to the layer-3 next hop (mesh under) or handled hop-by-hop by each router. While the latter option consumes more resource, it allows to collect additionnal intermediary information, particularly relevant in wireless networks.
- o Defect: a temporary change in the network (e.g., a radio link which is broken due to a mobile obstacle);
- o Fault: a definite change which may affect the network performance, e.g., a node runs out of energy.

1.2. Acronyms

OAM Operations, Administration, and Maintenance

DetNet Deterministic Networking

SLO Service Level Objective

QoS Quality of Service

SNMP Simple Network Management Protocol

SDN Software-Defined Network

1.3. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Role of OAM in RAW

RAW networks expect to make the communications reliable and predictable on top of a wireless network infrastructure. Most critical applications will define an SLO to be required for the data flows it generates. RAW considers network plane protocol elements

such as OAM to improve the RAW operation at the service and the forwarding sub-layers.

To respect strict guarantees, RAW relies on an orchestrator able to monitor and maintain the network. Typically, a Software-Defined Network (SDN) controller is in charge of scheduling the transmissions in the deployed network, based on the radio link characteristics, SLO of the flows, the number of packets to forward. Thus, resources have to be provisioned a priori to handle any defect. OAM represents the core of the pre-provisioning process and maintains the network operational by updating the schedule dynamically.

Fault-tolerance also assumes that multiple paths have to be provisioned so that an end-to-end circuit keeps on existing whatever the conditions. The Packet Replication and Elimination Function ([PREF-draft]) on a node is typically controlled by a central controller/orchestrator. OAM mechanisms can be used to monitor that PREOF is working correctly on a node and within the domain.

To be energy-efficient, reserving some dedicated out-of-band resources for OAM seems idealistic, and only in-band solutions are considered here.

RAW supports both proactive and on-demand troubleshooting.

The specific characteristics of RAW are discussed below.

2.1. Link concept and quality

In wireless networks, a `_link_` does not exist physically. A common convention is to define a wireless link as a pair of devices that have a non-null probability of exchanging a packet that the receiver can decode. Similarly, we designate as **neighbor** any device with a radio link with a specific transmitter.

Each wireless link is associated with a link quality, often measured as the Packet Delivery Ratio (PDR), i.e., the probability that the receiver can decode the packet correctly. It is worth noting that this link quality depends on many criteria, such as the level of external interference, the presence of concurrent transmissions, or the radio channel state. This link quality is even time-variant.

2.2. Broadcast Transmissions

In modern switching networks, the unicast transmission is delivered uniquely to the destination. Wireless networks are much closer to the ancient **shared access** networks. Practically, unicast and broadcast frames are handled similarly at the physical layer. The

link layer is just in charge of filtering the frames to discard irrelevant receptions (e.g., different unicast MAC address).

However, contrary to wired networks, we cannot be sure that a packet is received by **all** the devices attached to the layer-2 segment. It depends on the radio channel state between the transmitter(s) and the receiver(s). In particular, concurrent transmissions may be possible or not, depending on the radio conditions (e.g., do the different transmitters use a different radio channel or are they sufficiently spatially separated?)

2.3. Complex Layer 2 Forwarding

Multiple neighbors may receive a transmission. Thus, anycast layer-2 forwarding helps to maximize the reliability by assigning multiple receivers to a single transmission. That way, the packet is lost only if **none** of the receivers decode it. Practically, it has been proven that different neighbors may exhibit very different radio conditions, and that reception independency may hold for some of them [anycast-property].

3. Operation

OAM features will enable RAW with robust operation both for forwarding and routing purposes.

3.1. Information Collection

The model to exchange information should be the same as for detnet network, for the sake of inter-operability. YANG may typically fulfill this objective.

However, RAW networks imply specific constraints (e.g., low bandwidth, packet losses, cost of medium access) that may require to minimize the volume of information to collect. Thus, we discuss in Section 4.2 the different ways to collect information, i.e., transfer physically the OAM information from the emitter to the receiver.

3.2. Continuity Check

Similarly to detnet, we need to verify that the source and the destination are connected (at least one valid path exists)

3.3. Connectivity Verification

As in detnet, we have to verify the absence of misconnection. We will focus here on the RAW specificities.

Because of radio transmissions' broadcast nature, several receivers may be active at the same time to enable anycast Layer 2 forwarding. Thus, the connectivity verification must test any combination. We also consider priority-based mechanisms for anycast forwarding, i.e., all the receivers have different probabilities of forwarding a packet. To verify a delay SLO for a given flow, we must also consider all the possible combinations, leading to a probability distribution function for end-to-end transmissions. If this verification is implemented naively, the number of combinations to test may be exponential and too costly for wireless networks with low bandwidth.

3.4. Route Tracing

Wireless networks are meshed by nature: we have many redundant radio links. These meshed networks are both an asset and a drawback: while several paths exist between two endpoints, and we should choose the most efficient one(s), concerning specifically the reliability, and the delay.

Thus, multipath routing can be considered to make the network fault-tolerant. Even better, we can exploit the broadcast nature of wireless networks to exploit meshed multipath routing: we may have multiple Maintenance Intermediate Endpoints (MIE) for each hop in the path. In that way, each Maintenance Intermediate Endpoint has several possible next hops in the forwarding plane. Thus, all the possible paths between two maintenance endpoints should be retrieved, which may quickly become untractable if we apply a naive approach.

3.5. Fault Verification/detection

Wired networks tend to present stable performances. On the contrary, wireless networks are time-variant. We must consequently make a distinction between normal evolutions and malfunction.

3.6. Fault Isolation/identification

The network has isolated and identified the cause of the fault. While detnet already expects to identify malfunctions, some problems are specific to wireless networks. We must consequently collect metrics and implement algorithms tailored for wireless networking.

For instance, the decrease in the link quality may be caused by several factors: external interference, obstacles, multipath fading, mobility. It is fundamental to be able to discriminate the different causes to make the right decision.

4. Administration

The RAW network has to expose a collection of metrics to support an operator making proper decisions, including:

- o Packet losses: the time-window average and maximum values of the number of packet losses have to be measured. Many critical applications stop to work if a few consecutive packets are dropped;
- o Received Signal Strength Indicator (RSSI) is a very common metric in wireless to denote the link quality. The radio chipset is in charge of translating a received signal strength into a normalized quality indicator;
- o Delay: the time elapsed between a packet generation / enqueueing and its reception by the next hop;
- o Buffer occupancy: the number of packets present in the buffer, for each of the existing flows.

These metrics should be collected per device, virtual circuit, and path, as detnet already does. However, we have to face in RAW to a finer granularity:

- o per radio channel to measure, e.g., the level of external interference, and to be able to apply counter-measures (e.g., blacklisting).
- o per link to detect misbehaving link (assymetrical link, fluctuating quality).
- o per resource block: a collision in the schedule is particularly challenging to identify in radio networks with spectrum reuse. In particular, a collision may not be systematic (depending on the radio characteristics and the traffic profile)

4.1. Worst-case metrics

RAW inherits the same requirements as detnet: we need to know the distribution of a collection of metrics. However, wireless networks are known to be highly variable. Changes may be frequent, and may exhibit a periodical pattern. Collecting and analyzing this amount of measurements is challenging.

Wireless networks are known to be lossy, and RAW has to implement strategies to improve reliability on top of unreliable links. Hybrid Automatic Repeat reQuest (ARQ) has typically to enable

retransmissions based on the end-to-end reliability and latency requirements.

4.2. Efficient data retrieval

We have to minimize the number of statistics / measurements to exchange:

- o energy efficiency: low-power devices have to limit the volume of monitoring information since every bit consumes energy.
- o bandwidth: wireless networks exhibit a bandwidth significantly lower than wired, best-effort networks.
- o per-packet cost: it is often more expensive to send several packets instead of combining them in a single link-layer frame.

In conclusion, we have to take care of power and bandwidth consumption. The following techniques aim to reduce the cost of such maintenance:

on-path collection: some control information is inserted in the data packets if they do not fragment the packet (i.e., the MTU is not exceeded). Information Elements represent a standardized way to handle such information;

flags/fields: we have to set-up flags in the packets to monitor to be able to monitor the forwarding process accurately. A sequence number field may help to detect packet losses. Similarly, path inference tools such as [ipath] insert additional information in the headers to identify the path followed by a packet a posteriori.

hierarchical monitoring; localized and centralized mechanisms have to be combined together. Typically, a local mechanism should continuously monitor a set of metrics and trigger distant OAM exchanges only when a fault is detected (but possibly not identified). For instance, local temporary defects must not trigger expensive OAM transmissions.

5. Maintenance

RAW needs to implement a self-healing and self-optimization approach. The network must continuously retrieve the state of the network, to judge about the relevance of a reconfiguration, quantifying:

the cost of the sub-optimality: resources may not be used optimally (e.g., a better path exists);

the reconfiguration cost: the controller needs to trigger some reconfigurations. For this transient period, resources may be twice reserved, and control packets have to be transmitted.

Thus, reconfiguration may only be triggered if the gain is significant.

5.1. Dynamic Resource Reservation

Wireless networks exhibit time-variant characteristics. Thus, the network has to provide additional resources along the path to fit the worst-case performance. This time-variant characteristics make the resource reservation very challenging: over-reaction waste radio and energy resources. Inversely, under-reaction jeopardize the network operations, and some SLO may be violated.

5.2. Reliable Reconfiguration

Wireless networks are known to be lossy. Thus, commands may be received or not by the node to reconfigure. Unfortunately, inconsistent states may create critical misconfigurations, where packets may be lost along a path because it has not been properly configured.

We have to propose mechanisms to guarantee that the network state is always consistent, even if some control packets are lost. Timeouts and retransmissions are not sufficient since the reconfiguration duration would be, in that case, unbounded.

6. IANA Considerations

This document has no actionable requirements for IANA. This section can be removed before the publication.

7. Security Considerations

This section will be expanded in future versions of the draft.

8. Acknowledgments

TBD

9. Informative References

- [anycast-property] Teles Hermeto, R., Gallais, A., and F. Theoleyre, "Is Link-Layer Anycast Scheduling Relevant for IEEE 802.15.4-TSCH Networks?", 2019, <<https://doi.org/10.1109/LCNSymposium47956.2019.9000679>>.
- [detnet-oam] Theoleyre, F., Papadopoulos, G. Z., Mirsky, G., and C. J. Bernardos, "Operations, Administration and Maintenance (OAM) features for detnet", 2020, <<https://tools.ietf.org/html/draft-theoleyre-detnet-oam-support>>.
- [ipath] Gao, Y., Dong, W., Chen, C., Bu, J., Wu, W., and X. Liu, "iPath: path inference in wireless sensor networks.", 2016, <<https://doi.org/10.1109/TNET.2014.2371459>>.
- [PREF-draft] Thubert, P., Eckert, T., Brodard, Z., and H. Jiang, "BIER-TE extensions for Packet Replication and Elimination Function (PREF) and OAM", 2018, <<https://tools.ietf.org/html/draft-thubert-bier-replication-elimination>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC6291] Andersson, L., van Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "Guidelines for the Use of the "OAM" Acronym in the IETF", BCP 161, RFC 6291, DOI 10.17487/RFC6291, June 2011, <<https://www.rfc-editor.org/info/rfc6291>>.
- [RFC7276] Mizrahi, T., Sprecher, N., Bellagamba, E., and Y. Weingarten, "An Overview of Operations, Administration, and Maintenance (OAM) Tools", RFC 7276, DOI 10.17487/RFC7276, June 2014, <<https://www.rfc-editor.org/info/rfc7276>>.
- [RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <<https://www.rfc-editor.org/info/rfc7799>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas,
"Deterministic Networking Architecture", RFC 8655,
DOI 10.17487/RFC8655, October 2019,
<<https://www.rfc-editor.org/info/rfc8655>>.

Authors' Addresses

Fabrice Theoleyre
CNRS
Building B
300 boulevard Sebastien Brant - CS 10413
Illkirch - Strasbourg 67400
FRANCE

Phone: +33 368 85 45 33
Email: theoleyre@unistra.fr
URI: <http://www.theoleyre.eu>

Georgios Z. Papadopoulos
IMT Atlantique
Office B00 - 102A
2 Rue de la Chataigneraie
Cesson-Sevigne - Rennes 35510
FRANCE

Phone: +33 299 12 70 04
Email: georgios.papadopoulos@imt-atlantique.fr

Greg Mirsky
ZTE Corp.

Email: gregimirsky@gmail.com