

RAW
Internet-Draft
Intended status: Informational
Expires: 25 April 2020

P. Thubert, Ed.
Cisco Systems
G.Z. Papadopoulos
IMT Atlantique
23 October 2019

Reliable and Available Wireless Problem Statement
draft-pthubert-raw-problem-statement-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 forwarded 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 25 April 2020.

Copyright Notice

Copyright (c) 2019 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	2
2. Terminology	4
3. Use Cases and Requirements Served	5
4. Routing Time Scale vs. Forwarding Time Scale	5
5. Prerequisites	7
6. Related Work at The IETF	7
7. Problem Statement	8
8. Security Considerations	9
9. IANA Considerations	9
10. References	9
10.1. Normative References	9
10.2. Informative References	10
Authors' Addresses	11

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, adds another dimension to the statistical effects that affect the delivery. Scheduling transmissions can alleviate those effects by leveraging diversity in the spatial, time, code, and frequency domains, and provide a Reliable and Available service while preserving energy and optimizing the use of the shared spectrum.

Deterministic Networking is an attempt to mostly eliminate packet loss for a committed bandwidth with a guaranteed worst-case end-to-end latency, even when co-existing with best-effort traffic in a shared network. This innovation is enabled by recent developments in

technologies including IEEE 802.1 TSN (for Ethernet LANs) and IETF DetNet (for wired IP networks). It is getting traction in various industries including manufacturing, online gaming, professional A/V, cellular radio and others, making possible many cost and performance optimizations.

The DetNet architecture [DetNet-ARCH] is composed of three planes: a (User)Application Plane, a Controller Plane, and a Network Plane. Reliable and Available Wireless (RAW) extends DetNet to focus on issues that are mostly a concern on wireless links, and inherits the architecture and the planes. A RAW Network Plane is thus a Network Plane inherited by RAW from DetNet.

RAW networking aims at providing highly available and reliable end-to-end performances in a network with scheduled wireless segments. Uncontrolled interference and transmission obstacles may impede the transmission, and techniques such as beamforming with Multi-User MIMO can only alleviate some of those issues, so the term "deterministic" is usually not associated with short range radios, in particular in the ISM band. This uncertainty places limits to the amount of traffic that can be transmitted on a link while conforming to a RAW Service Level Agreement (SLA) that may vary rapidly.

The wireless and wired media are fundamentally different at the physical level, and while the generic Problem Statement for DetNet applies to the wired as well as the wireless medium, the methods to achieve RAW will differ from those used to support time-sensitive networking over wires, as a RAW solution will need to address less consistent transmissions, energy conservation and shared spectrum efficiency.

The development of RAW technologies has been lagging behind deterministic efforts for wired systems both at the IEEE and the IETF. But recent efforts at the IEEE and 3GPP indicate that wireless is finally catching up at the lower layer and that it is now possible for the IETF to extend DetNet for wireless segments that are capable of scheduled wireless transmissions.

The intent for RAW is to provide 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.

Still, in order to focus on real-worlds issues and assert the feasibility of the proposed capabilities, RAW will focus 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.

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 at the forwarding time scale on one DetNet flow over one path that is preestablished and installed by means outside of the scope of RAW. The scope of the RAW WG comprises Network plane protocol elements such as OAM and in-band control to improve the RAW operation at the Service and at the forwarding sub-layers, e.g., controlling whether to use packet replication, Hybrid ARQ and coding, with a constraint to limit the use of redundancy when it is really needed, e.g., when a spike of loss is observed. This is discussed in more details in Section 4 and the next sections.

2. Terminology

RAW reuses terminology defined for DetNet in [DetNet-ARCH], e.g., PEOF for Packet Replication, Elimination and Ordering Functions.

RAW also reuses terminology defined for 6TiSCH in [6TiSCH-ARCH] such as Track. 6TiSCH defined the term Track for that complex path with associated PAREO operations.

RAW defines the following terms:

PAREO: Packet (hybrid) ARQ, Replication, Elimination and Ordering.

PAREO is a superset Of DetNet's PEOF 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 wireless connectivity is interrupted for short transient times, typically of a subsecond duration.

This document reuses terms that are well-defined in the context of automation to networking and packet delivery, in particular for reliability and availability. In the context of the RAW work, they 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).

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 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. Use Cases and Requirements Served

[RFC8578] presents a number of wireless use cases including Wireless for Industrial Applications. [RAW-USE-CASES] adds a number of use cases that demonstrate the need for RAW capabilities in Pro-Audio, gaming and robotics.

4. Routing Time Scale vs. Forwarding Time Scale

With DetNet, the end-to-end routing can be centralized and can reside outside the network. In wireless, and in particular in a wireless mesh, the path to the controller that performs the route computation and maintenance expensive in terms of critical resources such as air time and energy.

Reaching to the routing computation can also be slow in regards to the speed of events that affect the forwarding operation at the radio layer. Due to the cost and latency to perform a route computation, 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 operational metrics that aggregate the behavior of a link over long periods of time, and represent its availability as shades of gray as opposed to either up or down.

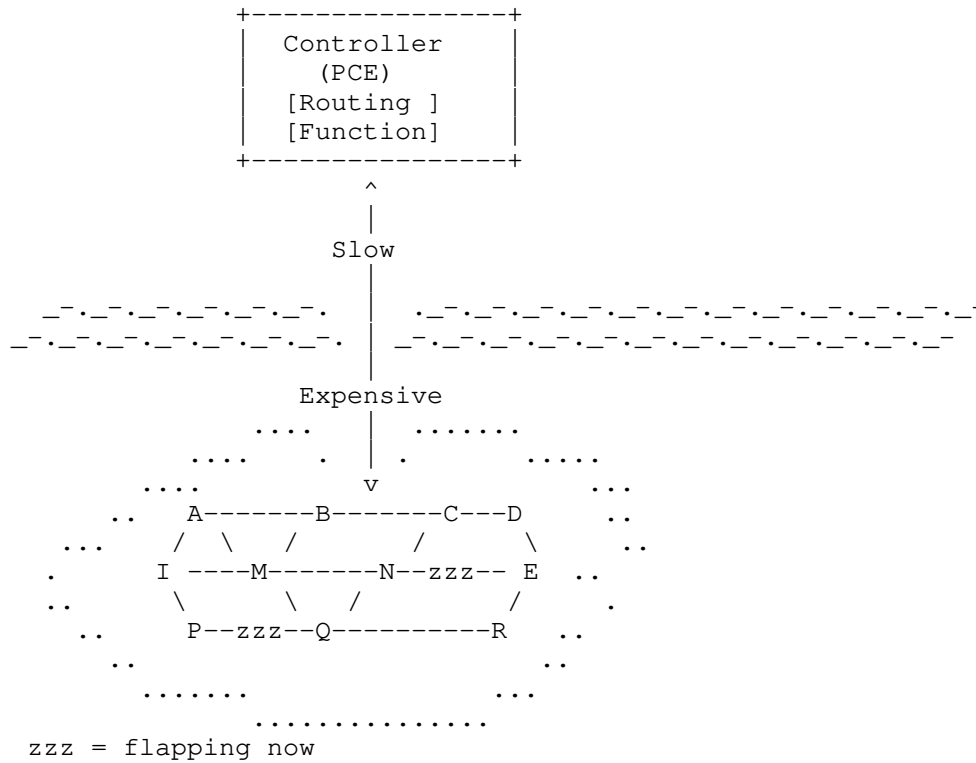


Figure 1: Time Scales

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 such a 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 forwarding plane to make the per-packet decision of which of these possibilities should be used.

In the case of wires, the concept is known in traffic engineering where 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 control 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 packet speed, and uses information that must be pertinent at the present time for the current transmission.

5. 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. For 6TiSCH, this is a Track. The concept of Track is specified in the 6TiSCH Architecture [6TiSCH-ARCH]. Tracks provide a high degree of redundancy and diversity and enable DetNet PREOF, end-to-end network coding, and possibly radio-specific abstracted techniques such as ARQ, overhearing, frequency diversity, time slotting, and possibly others.

How the routing operation 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 service 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 forwarding decision (think DLEP from MANET). In the layered approach, how the radio manages its PHY layer is out of control and out of scope. Whether it is single hop or meshed is also unknown and out of scope.

6. 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.
- * Operations, Administration and Maintenance (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.

7. Problem Statement

Within a large routed topology, the routing 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 forwarding 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 will consider existing tools such as L2-triggers, DLEP, BFD and in-band and out-of-band OAM.

One possible way of making the RAW forwarding decisions is to make them all at the ingress and express them in-band in the packet, which requires new loose or strict Hop-by-hop signaling. To control the RAW forwarding operation along a Track for the individual packets, RAW may leverage and extend known techniques such as DetNet tagging, Segment Routing (SRv6) or BIER-TE such as done with [BIER-PREF].

An alternate way is to enable each forwarding node to make 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 within the Track. Information about the service should be placed in the packet and matched with the forwarding node's capabilities and policies.

In either case, a per-flow state is installed in all intermediate nodes to recognize the flow and determine the forwarding policy to be applied.

8. Security Considerations

This document is a problem statement and does not propose a solution that could yield security issues.

9. IANA Considerations

This document has no IANA actions.

10. References

10.1. Normative References

[6TISCH-ARCH]

Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", Work in Progress, Internet-Draft, draft-ietf-6tisch-architecture-27, 18 October 2019, <<https://tools.ietf.org/html/draft-ietf-6tisch-architecture-27>>.

[DetNet-ARCH]

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>>.

[RAW-TECHNOS]

Thubert, P., Cavalcanti, D., Vilajosana, X., and C.

Schmitt, "Reliable and Available Wireless Technologies", Work in Progress, Internet-Draft, draft-thubert-raw-technologies-03, 1 July 2019, <<https://tools.ietf.org/html/draft-thubert-raw-technologies-03>>.

[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-00, 5 July 2019, <<https://tools.ietf.org/html/draft-bernardos-raw-use-cases-00>>.

[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>>.

[RFC8578] Grossman, E., Ed., "Deterministic Networking Use Cases", RFC 8578, DOI 10.17487/RFC8578, May 2019, <<https://www.rfc-editor.org/info/rfc8578>>.

10.2. Informative References

[BFD] IETF, "Bidirectional Forwarding Detection", October 2019, <<https://dataTracker.ietf.org/doc/charter-ietf-bfd/>>.

[BIER] IETF, "Bit Indexed Explicit Replication", October 2019, <<https://dataTracker.ietf.org/doc/charter-ietf-bier/>>.

[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>>.

[CCAMP] IETF, "Common Control and Measurement Plane", October 2019, <<https://dataTracker.ietf.org/doc/charter-ietf-ccamp/>>.

[DetNet] IETF, "Deterministic Networking", October 2019, <<https://dataTracker.ietf.org/doc/charter-ietf-detnet/>>.

[DetNet-DP-FW]

Varga, B., Farkas, J., Berger, L., Fedyk, D., Malis, A., Bryant, S., and J. Korhonen, "DetNet Data Plane

Framework", Work in Progress, Internet-Draft, draft-ietf-detnet-data-plane-framework-02, 13 September 2019, <<https://tools.ietf.org/html/draft-ietf-detnet-data-plane-framework-02>>.

[DetNet-IP-OAM]

Mirsky, G. and M. Chen, "Operations, Administration and Maintenance (OAM) for Deterministic Networks (DetNet) with IP Data Plane", Work in Progress, Internet-Draft, draft-mirsky-detnet-ip-oam-00, 8 July 2019, <<https://tools.ietf.org/html/draft-mirsky-detnet-ip-oam-00>>.

[MANET]

IETF, "Mobile Ad hoc Networking", October 2019, <<https://dataTracker.ietf.org/doc/charter-ietf-manet/>>.

[SPRING]

IETF, "Source Packet Routing in Networking", October 2019, <<https://dataTracker.ietf.org/doc/charter-ietf-spring/>>.

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

RAW
Internet-Draft
Intended status: Standards Track
Expires: May 6, 2020

F. Theoleyre
CNRS
G. Papadopoulos
IMT Atlantique
November 3, 2019

Operations, Administration and Maintenance (OAM) features for RAW
draft-theoleyre-raw-oam-support-01

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 some of these use-cases.

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 May 6, 2020.

Copyright Notice

Copyright (c) 2019 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	2
1.1. Terminology	3
2. Needs for OAM in RAW	3
3. Operation	4
3.1. Connectivity Verification	4
3.2. Route Tracing	4
3.3. Fault verification / detection	4
3.4. Fault isolation / identification	5
4. Administration	5
4.1. Worst-case metrics	6
4.2. Energy efficiency constraint	6
5. Maintenance	6
5.1. Multipath	7
5.2. Replication / Elimination	7
5.3. Resource Reservation	7
5.4. Soft transition after reconfiguration	7
6. Informative References	8
Authors' Addresses	8

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. The wireless and wired media are fundamentally different at the physical level. Enabling thus reliable and available wireless communications is even more challenging than it is in wired IP networks, due to the numerous causes of loss in transmission that add up to the congestion losses and the delays caused by overbooked shared resources. To provide quality of service along a multihop path that is composed of wired and wireless hops, additional methods needs to be considered to leverage the potential lossy wireless communication.

Traceability belongs to Operations, Administration, and Maintenance (OAM) which is the toolset for fault detection and isolation, and for performance measurement. More can be found on OAM Tools in [RFC7276].

The main purpose of this document is to detail the 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 Agreements (SLA), such as delay and reliability, assigned to each data flow.

1.1. Terminology

- o OAM entity: a data flow to be controlled;
- o OAM end-devices: the source or destination of a data flow;
- o defect: a temporary change in the network characteristics (e.g. link quality degradation because of temporary external interference, a mobile obstacle)
- o fault: a definite change which may affect the network performance, e.g. a node runs out of energy,

2. Needs for 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 a SLA to respect for the data flows it generates. RAW considers network plane protocol elements such as OAM to improve the RAW operation at the service and at the forwarding sub-layers.

To respect strict guarantees, RAW relies on a Path Computation Element (PCE) which will be responsible to schedule the transmissions in the deployed network. Thus, resources have to be provisioned a priori to handle any defect. OAM represents the core of the over provisioning process, and maintains the network operational by updating the schedule dynamically.

Fault-tolerance also assumes that multiple path have to be provisioned so that an end-to-end circuit keeps on existing whatever the conditions. OAM is in charge of controlling the replication/elimination processes.

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.

3. Operation

OAM features will enable RAW with robust operation both for forwarding and routing purposes.

3.1. Connectivity Verification

We need to verify that two endpoints are connected with each other. Since we reserve resources along the path independently for each flow, we must be able to verify that the path exists for a given flow label.

The control and data packets may not follow the same path, and the connectivity verification has to be triggered in-band without impacting the data traffic. In particular, the control plane may work while the data plane may be broken.

The ping packets must be labeled in the same way as the data packets of the flow to monitor.

3.2. Route Tracing

Ping and traceroute are two very common tools for diagnostic. They help to identify the list of routers in the route. However, to be predictable, resources are reserved per flow in RAW. Thus, we need to define route tracing tools able to track the route for a specific flow.

Because the network has to be fault-tolerant, multipath can be considered, with multiple Maintenance Intermediate Endpoints for each hop in the path. Thus, all the possible paths between two maintenance endpoints should be retrieved.

3.3. Fault verification / detection

RAW expects to operate fault-tolerant networks. Thus, we need mechanisms able to detect faults, before they impact the network performance.

The network has to detect when a fault occurred, i.e. the network has deviated from its expected behavior. While the network must report an alarm, the cause may not be identified precisely. For instance, the end-to-end reliability has decreased significantly, or a buffer overflow occurs.

We have to minimize the amount 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: is often more expensive to send several packets instead of combining them in a single link-layer frame.

Thus, localized and centralized mechanisms have to be combined together, and additional control packets have to be triggered only after a fault detection.

3.4. Fault isolation / identification

The network has isolated and identified the cause of the fault. For instance, the quality of a specific link has decreased, requiring more retransmissions, or the level of external interference has locally increased.

4. Administration

To take proper decisions, the network has to expose a collection of metrics, including:

- o Packet losses: the time-window average and maximum values of the number of packet losses has 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:

- o per virtual circuit to measure the end-to-end performance for a given flow. Each of the paths has to be isolated in multipath strategies;

- 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 device to detect misbehaving node, when it relays the packets of several flows.

4.1. Worst-case metrics

RAW aims to enable real-time communications on top of an heterogeneous architecture. Since wireless networks are known to be lossy, RAW has to implement strategies to improve the 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.

To take correct decisions, the controller needs to know the distribution of packet losses for each flow, and for each hop of the paths. In other words, average end-to-end statistics are not enough. They must allow the controller to predict the worst-case.

4.2. Energy efficiency constraint

RAW targets also low-power wireless networks, where energy represents a key constraint. Thus, we have to take care of the energy and bandwidth consumption. The following techniques aim to reduce the cost of such maintenance:

piggybacking: some control information are 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.

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. Multipath

To be fault-tolerant, several paths can be reserved between two maintenance endpoints. They must be node-disjoint, so that a path can be available at any time.

5.2. Replication / Elimination

When multiple paths are reserved between two maintenance endpoints, they may decide to replicate the packets to introduce redundancy, and thus to alleviate transmission errors and collisions. For instance, in Figure 1, the source node S is transmitting the packet to both parents, nodes A and B. Each maintenance endpoint will decide to trigger the replication / elimination process when a set of metrics passes through a threshold value.




Figure 1: Packet Replication: S transmits twice the same data packet, to its DP (A) and to its AP (B).

5.3. Resource Reservation

Because the QoS criteria associated to a path may degrade, the network has to provision additional resources along the path. We need to provide mechanisms to patch a schedule (changing the channel offset, allocating more timeslots, changing the path, etc.).

5.4. Soft transition after reconfiguration

Since RAW expects to support real-time flows, we have to support soft-reconfiguration, where the novel resources are reserved before the ancient ones are released. Some mechanisms have to be proposed so that packets are forwarded through the novel track only when the

resources are ready to be used, while maintaining the global state consistent (no packet re-ordering, duplication, etc.)

6. Informative References

- [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>>.
- [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>>.

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

