

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: November 20, 2015

Z. Cao
Leibniz University of Hannover
C. Gomez
Universitat Politecnica de Catalunya/i2CAT
M. Kovatsch
ETH Zurich
H. Tian
China Academy of Telecommunication Research
X. He
Hitachi China R&D Corporation
May 19, 2015

Energy Efficient Implementation of IETF Constrained Protocol Suite draft-ietf-lwig-energy-efficient-03

Abstract

This document summarizes the problems and current practices of energy efficient protocol implementation on constrained devices, mostly about how to make the protocols within IETF scope behave energy friendly. This document also summarizes the impact of link layer protocol power saving behaviors to the upper layer protocols, so that they can coordinately make the system energy efficient.

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 <http://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 November 20, 2015.

Copyright Notice

Copyright (c) 2015 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 (<http://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.	Conventions used in this document	3
1.2.	Terminology	3
2.	Overview	3
3.	MAC and Radio Duty Cycling	5
3.1.	Radio Duty Cycling techniques	6
3.2.	Latency and buffering	7
3.3.	Throughput	7
3.4.	Radio interface tuning	7
3.5.	Power save services available in example low-power radios	7
3.5.1.	Power Save Services Provided by IEEE 802.11	8
3.5.2.	Power Save Services Provided by Bluetooth LE	9
3.5.3.	Power Save Services in IEEE 802.15.4	10
3.5.4.	Power Save Services in DECT ULE	11
4.	IP Adaptation and Transport Layer	13
5.	Routing Protocols	14
6.	Application Layer	15
6.1.	Energy efficient features in CoAP	15
6.2.	Sleepy node support	15
6.3.	CoAP timers	16
7.	Summary	16
8.	Contributors	17
9.	Acknowledgments	17
10.	IANA Considerations	17
11.	Security Considerations	17
12.	References	17
12.1.	Normative References	17
12.2.	Informative References	18
	Authors' Addresses	20

[1. Introduction](#)

In many scenarios, the network systems comprise many battery-powered or energy-harvesting devices. For example, in an environmental monitoring system or a temperature and humidity monitoring system in a data center, there are no always-on and handy sustained power

supplies for the potentially large number of constrained devices. In such deployment environments, it is necessary to optimize the energy consumption of the entire system, including computing, application layer behavior, and lower layer communication.

Significant research efforts have been spent on this "energy efficiency" problem. Most of this research has focused on how to optimize the system's power consumption regarding a certain deployment scenario or how could an existing network function such as routing or security be more energy-efficient. Only few efforts were spent on energy-efficient designs for IETF protocols and standardized network stacks for such constrained devices [[I-D.kovatsch-lwig-class1-coap](#)].

The IETF has developed a suite of Internet protocols suitable for such constrained devices, including 6LoWPAN ([[RFC6282](#)],[[RFC6775](#)],[[RFC4944](#)]), RPL[[RFC6550](#)], and CoAP[[I-D.ietf-core-coap](#)]. This document tries to summarize the design considerations of making the IETF protocol suite as energy-efficient as possible. While this document does not provide detailed and systematic solutions to the energy efficiency problem, it summarizes the design efforts and analyzes the design space of this problem. In particular, it provides a comprehensive overview of the techniques used by the lower layers to save energy and how these may impact on the upper layers.

After reviewing the energy-efficient design of each layer, an overall conclusion is summarized. Though the lower layer communication optimization is the key part of energy efficient design, the protocol design at the upper layers is also important to make the device energy-efficient.

1.1. Conventions used in this document

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 [[RFC2119](#)]

1.2. Terminology

The terminologies used in this document can be referred to [[RFC7228](#)].

2. Overview

The IETF has developed multiple protocols to enable end-to-end IP communication between constrained nodes and fully capable nodes. This work has witnessed the evolution of the traditional Internet protocol stack to a light-weight Internet protocol stack. As shown

in Figure 1 below, the IETF has developed CoAP as the application layer and 6LoWPAN as the adaption layer to run IPV6 over IEEE 802.15.4 and Bluetooth Low-Energy, with the support of routing by RPL and efficient neighbor discovery by 6LoWPAN-ND. 6LoWPAN is currently being adapted by the 6lo working group to support IPV6 over various other technologies, such as ITU-T G.9959, DECT ULE, MS/TP-BACnet and NFC.

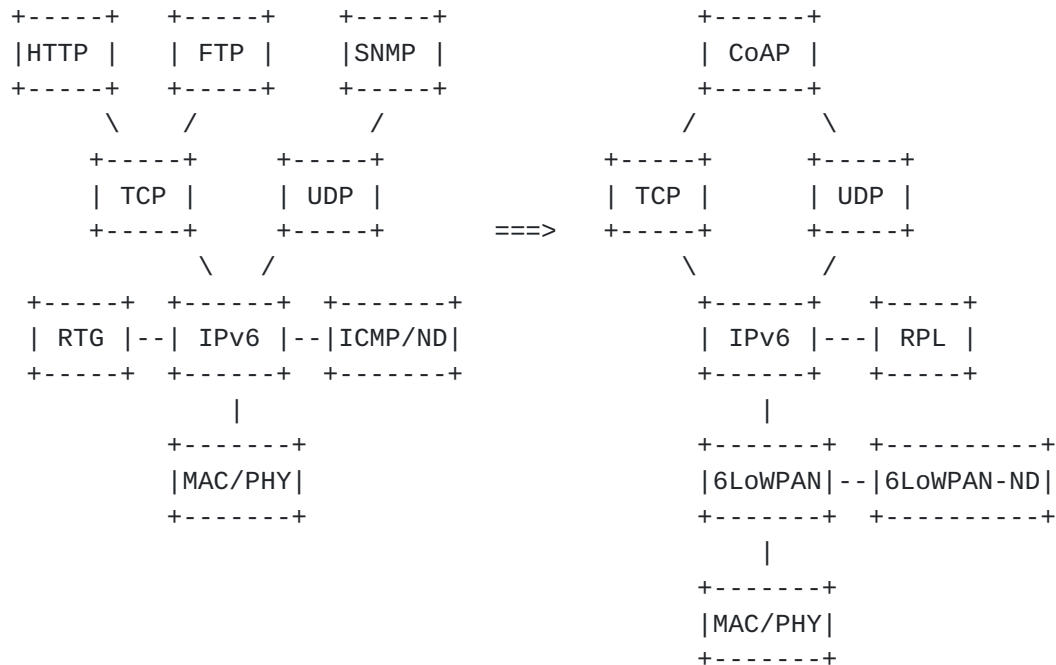


Figure 1: Traditional and Light-weight Internet Protocol Stack

There are numerous published studies reporting comprehensive measurements of wireless communication platforms [[Powertrace](#)]. As an example, below we list the energy consumption profile of the most common atom operations on a prevalent sensor node platform. The measurement was based on the Tmote Sky with ContikiMAC [[ContikiMAC](#)] as the radio duty cycling algorithm. From this and many other measurement reports (e.g. [[AN053](#)]), we can see that the energy consumption of optimized transmission and reception may be in the same order. For IEEE 802.15.4 and UWB radios, transmitting may actually be even cheaper than receiving. Only for broadcast and non-synchronized communication transmissions become costly in terms of energy because they need to flood the medium for a long time.

Activity	Energy (uJ)
Broadcast reception	178
Unicast reception	222
Broadcast transmission	1790
Non-synchronized unicast transmission	1090
Synchronized unicast transmission	120
Unicast TX to awake receiver	96

Figure 2: Power consumption of atom operations on the Tmote Sky with ContikiMAC

3. MAC and Radio Duty Cycling

In low-power wireless networks, communication and power consumption are intertwined. The communication device is typically the most power-consuming component, but merely refraining from transmissions is not enough to attain a low power consumption: the radio may consume as much power in listen mode as when actively transmitting. This augments the key problem known as idle listening, whereby the radio of a device may be in receive mode (ready to receive any message), even if no message is being transmitted to that device. Idle listening consumes a huge amount of energy unnecessarily. To reduce power consumption, the radio must be switched completely off -- duty-cycled -- as much as possible. By applying duty-cycling, the lifetime of a device operating on a common button battery may be in the order of years, whereas otherwise the battery may be exhausted in a few days or even hours. Duty-cycling is a technique generally exploited by devices that use the P1 strategy [RFC7228], which need to be able to communicate on a relatively frequent basis. Note that a more aggressive approach to save energy relies on the P0, Normally-off strategy, whereby devices sleep for very long periods and communicate infrequently, even though they spend energy in network reattachment procedures.

From the perspective of MAC&RDC, all upper layer protocols, such as routing, RESTful communication, adaptation, and management flows, are all applications. Since the duty cycling algorithm is the key to energy-efficiency of the wireless medium, it synchronizes the transmission and/or reception request from the higher layer.

The MAC&RDC are not in the scope of the IETF, yet lower layer designers and chipset manufactures take great care of the problem. For the IETF protocol designers, however, it is good to know the behaviors of lower layers so that the designed protocols can work perfectly with them.

Once again, the IETF protocols we are going to talk about in the following sections are the customers of the lower layers. If the different protocol layers want to get better service in a cooperative way, they should be considerate and understand each other.

3.1. Radio Duty Cycling techniques

This subsection describes the main three RDC techniques. Note that more than one of the presented techniques may be available or can even be combined in a specific radio technology:

a) Channel sampling. In this solution, the radio interface of a device periodically monitors the channel for very short time intervals (i.e. with a low duty cycle) with the aim of detecting incoming transmissions. In order to make sure that a receiver can correctly receive a transmitted data unit, the sender may prepend a preamble of a duration at least the sampling period to the data unit to be sent. Another option for the sender is to repeatedly transmit the data unit, instead of sending a preamble before the data unit. Once a transmission is detected by a receiver, the receiver may stay awake until the complete reception of the data unit. Examples of radio technologies that use preamble sampling include ContikiMAC, the Coordinated Sampled Listening (CSL) mode of IEEE 802.15.4e, and the Frequently Listening (FL) mode of ITU-T G.9959.

b) Scheduled transmissions. This approach allows a device to know the instants in which it should be awake (during some time interval) in order to receive data units. Otherwise, the device may remain in sleep mode. The decision on the instants that will be used for communication is reached by means of some form of negotiation between the involved devices. Such negotiation may be performed per transmission or per session/connection. Bluetooth Low Energy (Bluetooth LE) is an example of a radio technology based on this mechanism.

c) Listen after send. This technique allows a node to remain in sleep mode by default, wake up and poll a sender (which must be ready to receive a poll message) for pending transmissions. After sending the poll message, the node remains in receive mode, ready for a potential incoming transmission. After a certain time interval, the node may go back to sleep. For example, the Receiver Initiated Transmission (RIT) mode of 802.15.4e, and the transmission of data

between a coordinator and a device in IEEE 802.15.4-2003 use this technique.

3.2. Latency and buffering

The latency of a data unit transmission to a duty-cycled device is equal to or greater than the latency of transmitting to an always-on device. Therefore, duty-cycling leads to a trade-off between energy consumption and latency. Note that in addition to a latency increase, RDC may introduce latency variance, since the latency increase is a random variable (which is uniformly distributed if duty-cycling follows a periodical behavior).

On the other hand, due to the latency increase of duty-cycling, a sender waiting for a transmission opportunity may need to store subsequent outgoing packets in a buffer, increasing memory requirements and potentially incurring queuing waiting time that contributes to the packet overall delay and increases the probability of buffer overflow, leading to losses.

3.3. Throughput

Although throughput is not typically a key concern in constrained node network applications, it is indeed important in some services in this kind of networks, such as over-the-air software updates or when off-line sensors accumulate measurements that have to be quickly transferred when there is a connectivity opportunity.

Since RDC introduces inactive intervals in energy-constrained devices, it reduces the throughput that can be achieved when communicating with such devices. There exists a trade-off between the achievable throughput and energy consumption.

3.4. Radio interface tuning

The parameters controlling the radio duty cycle have to be carefully tuned to achieve the intended application and/or network requirements. On the other hand, upper layers should take into account the expected latency and/or throughput behavior due to RDC. The next subsection provides details on key parameters controlling RDC mechanisms, and thus fundamental trade-offs, for various examples of relevant low-power radio technologies.

3.5. Power save services available in example low-power radios

This subsection presents power save services and techniques used in a few relevant examples of wireless low-power radios: IEEE 802.11, Bluetooth LE and IEEE 802.15.4. For a more detailed overview of each

technology, the reader may refer to the literature or to the corresponding specifications.

3.5.1. Power Save Services Provided by IEEE 802.11

IEEE 802.11 defines the Power Save Mode (PSM) whereby a station may indicate to an Access Point (AP) that it will enter a sleep mode state. While the station is sleeping, the AP buffers any frames that should be sent to the sleeping station. The station wakes up every Listen Interval (which can be a multiple of the Beacon Interval) in order to receive beacons. The AP signals in the beacon whether there is data pending for the station or not. If there are not frames to be sent to the station, the latter may get back to sleep mode. Otherwise, the station may send a message requesting the transmission of the buffered data and stay awake in receive mode.

IEEE 802.11v [[IEEE80211v](#)] further defines mechanisms and services for power save of stations/nodes that include flexible multicast service (FMS), proxy ARP advertisement, extended sleep modes, traffic filtering. It would be useful if upper layer protocols knows such capabilities provided by the lower layer, so that they can coordinate with each other.

These services include:

Proxy ARP: The Proxy ARP capability enables an Access Point (AP) to indicate that the non-AP station (STA) will not receive ARP frames. The Proxy ARP capability enables the non-AP STA to remain in power-save for longer periods of time.

Basic Service Set (BSS) Max Idle Period management enables an AP to indicate a time period during which the AP does not disassociate a STA due to non-receipt of frames from the STA. This supports improved STA power saving and AP resource management.

FMS: A service in which a non-access point (non-AP) station (STA) can request a multicast delivery interval longer than the delivery traffic indication message (DTIM) interval for the purposes of lengthening the period of time a STA may be in a power save state.

Traffic Filtering Service (TFS): A service provided by an access point (AP) to a non-AP station (STA) that can reduce the number of frames sent to the non-AP STA by not forwarding individually addressed frames addressed to the non-AP STA that do not match traffic filters specified by the non-AP STA.

Using the above services provided by the lower layer, the constrained nodes can achieve either client initiated power save (via TFS) or network assisted power save (Proxy-ARP, BSS Max Idle Period and FMS).

Upper layer protocols would better synchronize with the parameters such as FMS interval and BSS MAX Idle Period, so that the wireless transmissions are not triggered periodically.

3.5.2. Power Save Services Provided by Bluetooth LE

Bluetooth LE is a wireless low-power communications technology that is the hallmark component of the Bluetooth 4.0, 4.1 and 4.2 specifications [[Bluetooth42](#)]. BT-LE has been designed for the goal of ultra-low-power consumption. Currently, it is possible to run IPv6 over Bluetooth LE networks by using a 6LoWPAN variant adapted to BT-LE [[I-D.ietf-6lowpan-btle](#)].

Bluetooth LE networks comprise a master and one or more slaves which are connected to the master. The Bluetooth LE master is assumed to be a relatively powerful device, whereas a slave is typically a constrained device (e.g. a class 1 device).

Medium access in Bluetooth LE is based on a TDMA scheme which is coordinated by the master. This device determines the start of connection events, in which communication between the master and a slave takes place. At the beginning of a connection event, the master sends a poll message, which may encapsulate data, to the slave. The latter must send a response, which may also contain data. The master and the slave may continue exchanging data until the end of the connection event. The next opportunity for communication between the master and the slave will be in the next connection event scheduled for the slave.

The time between consecutive connection events is defined by the `connInterval` parameter, which may range between 7.5 ms and 4 s. The slave may remain in sleep mode since the end of its last connection event until the beginning of its next connection event. Therefore, Bluetooth LE is duty-cycled by nature. Furthermore, after having replied to the master, a slave is not required to listen to the master (and thus may keep the radio in sleep mode) for `connSlaveLatency` consecutive connection events. `connSlaveLatency` is an integer parameter between 0 and 499 which should not cause link inactivity for more than `connSupervisionTimeout` time. The `connSupervisionTimeout` parameter is in the range between 100 ms and 32 s.

Upper layer protocols should take into account the medium access and duty-cycling behavior of Bluetooth LE. In particular, `connInterval`,

connSlaveLatency and connSupervisionTimeout determine the time between two consecutive connection events for a given slave. The upper layer packet generation pattern and rate should be consistent with the settings of the aforementioned parameters (and vice versa).

3.5.3. Power Save Services in IEEE 802.15.4

IEEE 802.15.4 is a family of standard radio interfaces for low-rate, low-power wireless networking [[fifteendotfour](#)]. Since the publication of its first version in 2003, IEEE 802.15.4 has become the de-facto choice for a wide range of constrained node network application domains and has been a primary target technology of various IETF working groups such as 6LoWPAN [[RFC6282](#)], [[RFC6775](#)], [[RFC4944](#)] and 6TiSCH [[I-D.ietf-6tisch-architecture](#)]. IEEE 802.15.4 specifies PHY and MAC layer functionality.

IEEE 802.15.4 defines three roles called device, coordinator and PAN coordinator. The device role is adequate for nodes that do not implement the complete IEEE 802.15.4 functionality, and is mainly targeted for constrained nodes with a limited energy source. The coordinator role includes synchronization capabilities and is suitable for nodes that do not suffer severe constraints (e.g. a mains-powered node). The PAN coordinator is a special type of coordinator that acts as a principal controller in an IEEE 802.15.4 network.

IEEE 802.15.4 has mainly defined two types of networks depending on their configuration: beacon-enabled and nonbeacon-enabled networks. In the first network type, coordinators periodically transmit beacons. The time between beacons is divided in three main parts: the Contention Access Period (CAP), the Contention Free Period (CFP) and an inactive period. In the first period, nodes use slotted CSMA/CA for data communication. In the second one, a TDMA scheme controls medium access. During the idle period, communication does not take place, thus the inactive period is a good opportunity for nodes to turn the radio off and save energy. The coordinator announces in each beacon the list of nodes for which data will be sent in the subsequent period. Therefore, devices may remain in sleep mode by default and wake up periodically to listen to the beacons sent by their coordinator. If a device wants to transmit data, or learns from a beacon that it is an intended destination, then it will exchange messages with the coordinator and will thus consume energy. An underlying assumption is that when a message is sent to a coordinator, the radio of the latter will be ready to receive the message.

The beacon interval and the duration of the beacon interval active portion (i.e. the CAP and the CFP), and thus the duty cycle, can be configured. The parameters that control these times are called `macBeaconOrder` and `macSuperframeOrder`, respectively. As an example, when IEEE 802.15.4 operates in the 2.4 GHz PHY, both times can be (independently) set to values in the range between 15.36 ms and 251.6 s.

In the beaconless mode, nodes use unslotted CSMA/CA for data transmission. The device may be in sleep mode by default and may activate its radio to either i) request to the coordinator whether there is pending data for the device, or ii) to transmit data to the coordinator. The wake-up pattern of the device, if any, is out of the scope of IEEE 802.15.4.

Communication between the two ends of an IEEE 802.15.4 link may also take place in a peer-to-peer configuration, whereby both link ends assume the same role. In this case, data transmission can happen at any moment. Nodes must have their radio in receive mode, and be ready to listen to the medium by default (which for battery-enabled nodes may lead to a quick battery depletion), or apply synchronization techniques. The latter are out of the scope of IEEE 802.15.4.

The main MAC layer IEEE 802.15.4 amendment to date is IEEE 802.15.4e. This amendment includes various new MAC layer modes, some of which include mechanisms for low energy consumption. Among these, the Time-Slotted Channel Hopping (TSCH) is an outstanding mode which offers robust features for industrial environments, among others. In order to provide the functionality needed to enable IPv6 over TSCH, the 6TiSCH working group has been recently created. TSCH is based on a TDMA schedule whereby a set of time slots are used for frame transmission and reception, and other time slots are unscheduled. The latter time slots may be used by a dynamic scheduling mechanism, otherwise nodes may keep the radio off during the unscheduled time slots, thus saving energy. The minimal schedule configuration specified in [[I-D.ietf-6tisch-minimal](#)] comprises 101 time slots, whereby 95 of these time slots are unscheduled and the time slot duration is 15 ms.

Other 802.15.4e modes, which are in fact designed for low energy, are the previously mentioned CSL and RIT.

3.5.4. Power Save Services in DECT ULE

DECT Ultra Low Energy (DECT ULE) is a wireless technology building on the key fundamentals of traditional DECT / CAT-iq [[EN300](#)] but with specific changes to significantly reduce the power consumption on the

expense of data throughput as specified in [\[TS102\]](#). DECT ULE devices typically operates on special power optimized silicon, but can connect to a DECT Gateway supporting traditional DECT / CAT-iq for cordless telephony and data as well as the DECT ULE extensions. It is possible to run IPv6 over DECT ULE by using a 6LoWPAN variant adapted for DECT ULE [\[I-D.ietf-6lo-dect-ule\]](#).

DECT terminology operates with two major role definitions: The Portable Part (PP) is the power constrained device, while the Fixed Part (FP) is the Gateway or base station in a star topology. DECT is operating in license free and reserved frequency bands based on TDMA/FDMA and TDD using dynamic channel allocation for interference avoidance. It provides good indoor (~50 m) and outdoor (~300 m) coverage. It is using a frame length of 10 ms, which is divided into 24 timeslots and it is supporting connection oriented, packet data and connection less services.

The FP usually transmits a so-called dummy bearer (beacon) that is used to broadcast synchronization, system and paging information. The slot/carrier position of this dummy bearer can automatically be reallocated in order to avoid mutual interference with other DECT signals.

At MAC level DECT ULE communications between FP and PP are initiated by the PP. A FP can initiate communication indirectly by sending paging signal to a PP. The PP determines the timeslot and frequency on which the communication between FP and PP takes place. The PP verifies the radio timeslot/frequency position is unoccupied before it initiates its transmitter. An access-request message, which usually carries data, is sent to the FP. The FP sends a confirm message, which also may carry data. More data can be sent in subsequent frames. A MAC level automatic retransmission scheme improves data transfer reliability significant. A segmentation and reassembly scheme supports transfer of larger higher layer SDUs and provides data integrity check. The DECT ULE packet data service ensures data integrity, proper sequencing, duplicate protection, but does not guaranteed delivery. Higher layers protocols have to take this into considerations.

The FP may send paging information to PPs to trigger connection setup and indicate required service type. The interval between paging information to a specific PP can be defined in range 10 ms to 327 seconds. The PP may enter sleep mode to save power. The listening interval is defined by the PP application. For short sleep intervals (below ~10 seconds) the PP may be able to retain synchronization to the FP dummy bearer and only turn on the receiver during the expected timeslot. For longer sleep intervals the PP can't keep synchronization and has to search for and resynchronize to the FP

dummybearer. Hence, longer sleep interval reduces the average energy consumption, but adds a energy consumption penalty for acquiring synchronization to the FP dummy bearer. The PP can obtain all information to determine paging and acquire synchronization information in a single reception of one full timeslot.

Packet data latency is normally 30 ms for short packets (below or equal to 32 octets), however if retry and back-off scenarios occur, the latency is increased. The latency can actually be reduced to about 10 ms by doing energy consuming RSSI scanning in advance. In the direction from FP to PP the latency is usually increased by the used paging interval and the sleep interval. The MAC layer can piggyback commands to improve efficiency (reduce latency) of higher layer protocols. Such commands can instruct the PP to initiate a new packet transfer in N frames without the need for resynchronization and listening to paging or instruct the PP to stay in a higher duty cycle paging detection mode.

The DECT ULE technology allows per PP configuration of paging interval, MTU size, reassembly window size and higher layer service negotiation and protocol.

4. IP Adaptation and Transport Layer

6LoWPAN is the adaption layer to run IPv6 over IEEE 802.15.4 MAC&PHY. It was born to fill the gap that the IPv6 layer does not support fragmentation and assembly of <1280-byte packets while IEEE 802.15.4 only supports a MTU of 127 bytes.

IPv6 is the basis for the higher layer protocols, including both TCP/UDP transport and applications. So they are quite ignorant of the lower layers, and are almost neutral to the energy-efficiency problem.

What the network stack can optimize is to save the computing power. For example the Contiki implementation has multiple cross layer optimizations for buffers and energy management, e.g., the computing and validation of UDP/TCP checksums without the need of reading IP headers from a different layer. These optimizations are software implementation techniques, and out of the scope of IETF and the LWIG working group.

6LoWPAN contributes to the energy-efficiency problem in two ways. First of all, it swaps computing with communication. 6LoWPAN applies compression of the IPv6 header. This means less amount of data will be handled by the lower layer, but both the sender and receiver should spend more computing power on the compression and decompression of the packets over the air. Secondly, the 6LoWPAN

working group developed the energy-efficient Neighbor Discovery called 6LoWPAN-ND, which is an energy efficient replacement of the IPv6 ND in constrained environments. IPv6 Neighbor Discovery was not designed for non-transitive wireless links, as its heavy use of multicast makes it inefficient and sometimes impractical in a low-power and lossy network. 6LoWPAN-ND describes simple optimizations to IPv6 Neighbor Discovery, its addressing mechanisms, and duplicate address detection for Low-power Wireless Personal Area Networks and similar networks. However, 6LoWPAN ND does not modify Neighbor Unreachability Detection (NUD) timeouts, which are very short (by default three transmissions spaced one second apart). NUD timeout settings should be tuned taking into account the latency that may be introduced by duty-cycled mechanisms at the link layer, or alternative, less impatient NUD algorithms should be considered [[I-D.ietf-6man-impatient-nud](#)].

5. Routing Protocols

The routing protocol designed by the IETF for constrained environments is called RPL [[RFC6550](#)]. As a routing protocol, RPL has to exchange messages periodically and keep routing states for each destination. RPL is optimized for the many-to-one communication pattern, where network nodes primarily send data towards the border router, but has provisions for any-to-any routing as well.

The authors of the Powertrace tool [[Powertrace](#)] studied the power profile of RPL. It divides the routing protocol into control and data traffic. The control channel uses ICMP messages to establish and maintain the routing states. The data channel is any application that uses RPL for routing packets. The study has shown that the power consumption of the control traffic goes down over time in a relatively stable network. The study also reflects that the routing protocol should keep the control traffic as low as possible to make it energy-friendly. The amount of RPL control traffic can be tuned by setting the Trickle algorithm parameters (i.e. I_{min} , I_{max} and k) to adequate values. However, there exists a trade-off between energy consumption and other performance parameters such as network convergence time and robustness.

[RFC 6551](#) [[RFC6551](#)] defines routing metrics and constraints to be used by RPL in route computation. Among others, [RFC 6551](#) specifies a Node Energy object that allows to provide information related to node energy, such as the energy source type or the estimated percentage of remaining energy. Appropriate use of energy-based routing metrics may help to balance energy consumption of network nodes, minimize network partitioning and increase network lifetime.

6. Application Layer

6.1. Energy efficient features in CoAP

CoAP [[RFC7252](#)] was designed as a RESTful application protocol, connecting the services of smart devices to the World Wide Web. CoAP is not a chatty protocol, it provides basic communication services such as service discovery and GET/POST/PUT/DELETE methods with a binary header.

The energy-efficient design is implicitly included in the CoAP protocol design. CoAP uses a fixed-length binary header of only four bytes that may be followed by binary options. To reduce regular and frequent queries of the resources, CoAP provides an observe mode, in which the requester registers its interest of a certain resource and the responder will report the value whenever it was updated. This reduces the request response round trips while keeping information exchange a ubiquitous service and, most importantly, it allows an energy-constrained server to remain in sleep mode during the period between observe notification transmissions.

Furthermore, [[RFC7252](#)] defines CoAP proxies which can cache resource representations previously provided by sleepy CoAP servers. The proxies themselves may respond to client requests if the corresponding server is sleeping and the resource representation is recent enough. Otherwise, a proxy may attempt to obtain the resource from the sleepy server.

6.2. Sleepy node support

Beyond these features of CoAP, there have been a number of proposals to further support sleepy nodes at the application layer by leveraging CoAP mechanisms. A good summary of such proposals can be found in [[I-D.rahman-core-sleepy-nodes-do-we-need](#)]. The different approaches include exploiting the use of proxies, leveraging the Resource Directory [[I-D.ietf-core-resource-directory](#)] or signaling when a node is awake to the interested nodes. A more recent work defines publish-subscribe and message queuing extensions to CoAP and the Resource Directory in order to support devices that spend most of their time in a sleeping state [[I-D.koster-core-coap-pubsub](#)]. As of the writing, none of these proposals has been adopted by the CoRE working group.

In addition to the work within the scope of CoAP to support sleepy nodes, other specifications define application layer functionality for the same purpose. The Lightweight Machine-to-Machine (LWM2M) specification from the Open Mobile Alliance (OMA) defines a Queue Mode whereby an LWM2M Server queues requests to an LWM2M Client until

the latter (which may often stay in sleep mode) is online. LWM2M functionality operates on top of CoAP.

On the other hand, oneM2M defines a CoAP binding with an application layer mechanism for sleepy nodes.

6.3. CoAP timers

CoAP offers mechanisms for reliable communication between two CoAP endpoints. A CoAP message may be signaled as a confirmable (CON) message, and an acknowledgment (ACK) is issued by the receiver if the CON message is correctly received. The sender starts a Retransmission TimeOut (RTO) for every CON message sent. The initial RTO value is chosen randomly between 2 and 3 s. If an RTO expires, the new RTO value is doubled (unless a limit on the number of retransmissions has been reached). Since duty-cycling at the link layer may lead to long latency (i.e. even greater than the initial RTO value), CoAP RTO parameters should be tuned accordingly in order to avoid spurious RTOs which would unnecessarily waste node energy and other resources.

7. Summary

We find a summary section necessary although most IETF documents do not contain it. The points we would like to summarize are as follows.

- a. All Internet protocols, which are in the scope of the IETF, are customers of the lower layers (PHY, MAC, and Duty-cycling). In order to get a better service, the designers of higher layers should know them better.
- b. The IETF has developed multiple protocols for constrained networked devices. A lot of implicit energy efficient design principles have been used in these protocols. The latter should be fine-tuned to exploit the collaboration with the lower layer protocols. Layers should offer interfaces that can be exploited by other layers in order to optimize global protocol stack performance.
- c. The power trace analysis of different protocol operations showed that for radio-duty-cycled networks broadcasts should be avoided. Saving unnecessary states maintenance is also an effective method to be energy-friendly.

8. Contributors

Jens T. Petersen, RTX, contributed the section on power save services in DECT ULE.

9. Acknowledgments

Carles Gomez has been supported by Ministerio de Economia y Competitividad and FEDER through project TEC2012-32531.

Authors would like to thank the review and feedback from a number of experts in this area: Carsten Bormann, Ari Keranen, Hannes Tschofenig, Dominique Barthel.

The text of this document was improved based on IESG Document Editing session during IETF87. Thank Ted Lemon, Joel Jaeggli, and efforts to initiate this facilities.

10. IANA Considerations

This document has no IANA requests.

11. Security Considerations

This document discusses the energy efficient protocol design, and does not incur any changes or challenges on security issues besides what the protocol specifications have analyzed.

12. References

12.1. Normative References

[Bluetooth42]

"Bluetooth Core Specification Version 4.2", 2014.

[EN300]

""Digital Enhanced Cordless Telecommunications (DECT); Common Interface (CI);"", 2013.

[IEEE80211v]

IEEE, , "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications, Amendment 8: IEEE 802.11 Wireless Network Management.", February 2012.

[RFC2119]

Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", [RFC 4944](#), September 2007.
- [RFC6282] Hui, J. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", [RFC 6282](#), September 2011.
- [RFC6550] Winter, T., Thubert, P., 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](#), March 2012.
- [RFC6551] Vasseur, JP., Kim, M., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics Used for Path Calculation in Low-Power and Lossy Networks", [RFC 6551](#), March 2012.
- [RFC6690] Shelby, Z., "Constrained RESTful Environments (CoRE) Link Format", [RFC 6690](#), August 2012.
- [RFC6775] Shelby, Z., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", [RFC 6775](#), November 2012.
- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", [RFC 7228](#), May 2014.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", [RFC 7252](#), June 2014.
- [TS102] ""Digital Enhanced Cordless Telecommunications (DECT); Ultra Low Energy (ULE); Machine to Machine Communications; Part 1: Home Automation Network (phase 1)"" , 2013.
- [fifteendotfour] "802.15.4-2011", 2011.

12.2. Informative References

- [AN053] Selvig, B., "Measuring power consumption with CC2430 and Z-Stack".
- [Announcementlayer] Dunkels, A., "The Announcement Layer: Beacon Coordination for the Sensornet Stack. In Proceedings of EWSN 2011".

[ContikiMAC]

Dunkels, A., "The ContikiMAC Radio Duty Cycling Protocol, SICS Technical Report T2011:13", December 2011.

[I-D.ietf-6lo-dect-ule]

Mariager, P., Petersen, J., Shelby, Z., Logt, M., and D. Barthel, "Transmission of IPv6 Packets over DECT Ultra Low Energy", [draft-ietf-6lo-dect-ule-01](#) (work in progress), January 2015.

[I-D.ietf-6lowpan-btle]

Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "Transmission of IPv6 Packets over BLUETOOTH Low Energy", [draft-ietf-6lowpan-btle-12](#) (work in progress), February 2013.

[I-D.ietf-6man-impatient-nud]

Nordmark, E. and I. Gashinsky, "Neighbor Unreachability Detection is too impatient", [draft-ietf-6man-impatient-nud-07](#) (work in progress), October 2013.

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

Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", [draft-ietf-6tisch-architecture-08](#) (work in progress), May 2015.

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

Vilajosana, X. and K. Pister, "Minimal 6TiSCH Configuration", [draft-ietf-6tisch-minimal-06](#) (work in progress), March 2015.

[I-D.ietf-core-coap]

Shelby, Z., Hartke, K., and C. Bormann, "Constrained Application Protocol (CoAP)", [draft-ietf-core-coap-18](#) (work in progress), June 2013.

[I-D.ietf-core-resource-directory]

Shelby, Z. and C. Bormann, "CoRE Resource Directory", [draft-ietf-core-resource-directory-02](#) (work in progress), November 2014.

[I-D.ietf-lwig-terminology]

Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained Node Networks", [draft-ietf-lwig-terminology-07](#) (work in progress), February 2014.

[I-D.koster-core-coap-pubsub]

Koster, M., Keranen, A., and J. Jimenez, "Publish-Subscribe Broker for the Constrained Application Protocol (CoAP)", [draft-koster-core-coap-pubsub-01](#) (work in progress), March 2015.

[I-D.kovatsch-lwig-class1-coap]

Kovatsch, M., "Implementing CoAP for Class 1 Devices", [draft-kovatsch-lwig-class1-coap-00](#) (work in progress), October 2012.

[I-D.rahman-core-sleepy-nodes-do-we-need]

Rahman, A., "Sleepy Devices: Do we need to Support them in CORE?", [draft-rahman-core-sleepy-nodes-do-we-need-01](#) (work in progress), February 2014.

[Powertrace]

Dunkels, , Eriksson, , Finne, , and Tsiftes, "Powertrace: Network-level Power Profiling for Low-power Wireless Networks", March 2011.

Authors' Addresses

Zhen Cao (Ed.)
Leibniz University of Hannover
P.R.China

Email: zhencao.ietf@gmail.com

Carles Gomez
Universitat Politecnica de Catalunya/i2CAT
C/Esteve Terradas, 7
Castelldefels 08860
Spain

Email: carlesgo@entel.upc.edu

Matthias Kovatsch
ETH Zurich
Universitaetstrasse 6
Zurich, CH-8092
Switzerland

Email: kovatsch@inf.ethz.ch

Hui Tian
China Academy of Telecommunication Research
Huayuanbeilu No.52
Beijing, Haidian District 100191
China

Email: tianhui@mail.ritt.com.cn

Xuan He
Hitachi China R&D Corporation
301, Tower C North, Raycom, 2 Kexuyuan Nanlu, Haidian District
Beijing 100190
P.R.China

Email: xhe@hitachi.cn

