Andreas Weigel (Autor)

Forwarding Strategies for 6LoWPAN-Fragmented IPv6 Datagrams

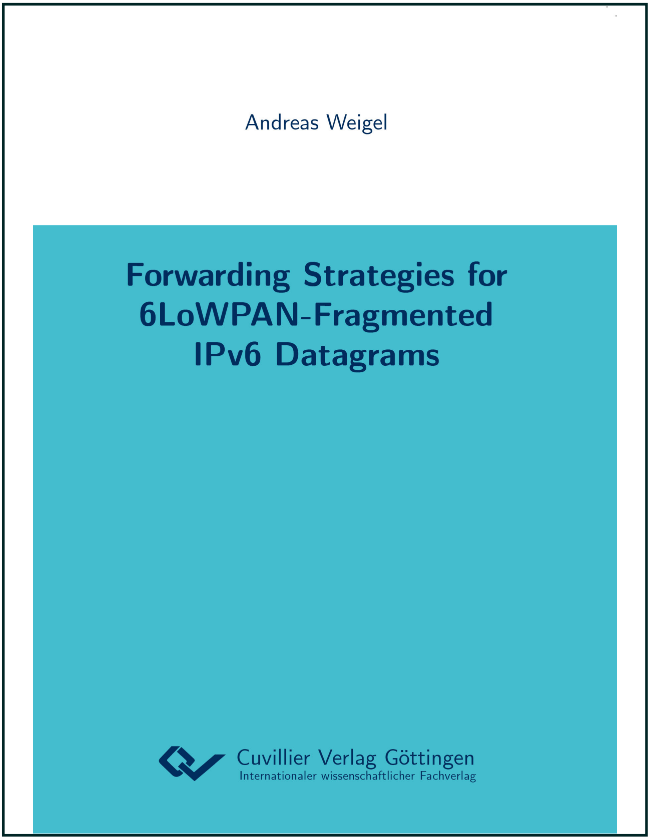

Andreas Weigel

Forwarding Strategies for 6LoWPAN-Fragmented IPv6 Datagrams

Cuvillier Verlag Göttingen
Internationaler wissenschaftlicher Fachverlag

https://cuvillier.de/de/shop/publications/7672

1 Introduction

Since the beginning of the 2000s, a class of networks termed wireless sensor networks (WSNs) have received the attention of a large research community. These usually feature a large number of small, resource- and energy-constrained devices that form a wireless mesh network to realize a sensing, monitoring or control task. At the time of writing, a typical node can be expected to possess from 4 to 32 KiB of RAM and 64 to 512 KiB of program memory, drawing current in the order of a few tens of mA with the transceiver active and a few µA when it is sleeping. Due to the nature of wireless communication channels, links between devices are often asymmetric and lossy, prone to interference by other wireless technologies and of transient nature due to changes in the environment. These properties also lead to the classification of low power, lossy networks (LLNs) for typical WSNs.

In recent years new names like cyber-physical systems, internet of things or industry 4.0 have emerged and show that the interest in ubiquitous autonomously communicating systems is unbroken.

Said attention brought forth a large number of protocols specifically tailored to cater the specialties of WSNs. Ranging from the “alphabet soup” of MAC protocols ([Ali+06]) over a plethora of routing protocols and corresponding link-quality metrics to transport protocols replacing the ubiquitous but for wireless lossy communication not terribly well-suited TCP, all layers of the communication stack have received due attention. When the dust settled, standardization efforts were launched to order the chaos.

Industry standards like ZigBee and WirelessHART based on the IEEE Standard for Local and metropolitan area networks–Part 15.4: Low-Rate Wireless Personal Area Networks (IEEE 802.15.4) were among the first of such efforts. Some years later, the Internet Engineering Task Force (IETF) instituted several working groups dealing with standardization of protocols for LLNs. Among them, the “IPv6 over Networks of Resource-Constrained Nodes” (6lo) defined mechanisms to enable the transmission of IPv6 datagrams over IEEE 802.15.4 networks, called 6LoWPAN. Its main responsibilities are compression of the comparatively large IPv6 and UDP/TCP headers to prevent the huge control overhead in combination with 127 B payload in standard IEEE 802.15.4 frames and fragmentation of large datagrams that do not fit a IEEE 802.15.4 frame even after compression. The Routing Protocol for Low Power and Lossy Networks (RPL) and the Constrained Application Protocol (CoAP) complete the fully standardized stack for that class of networks.

Considering the fragmentation of large datagrams it is intuitively clear that splitting up a datagram and transmitting the individual fragments one after the other does not improve the overall reliability of the reception of a datagram. Every single fragment has to arrive for the datagram to be successfully received and on paths that incorporate several wireless transmission hops, sending out a whole bunch of them can further degrade the reliability when frames belonging to the same datagram content with each other to acquire the wireless channel, that is, if they they can even “hear” each other – otherwise, senders along the same path are likely to cause hidden terminal collisions between consecutive fragment transmissions.

While arguably a large number of applications can be satisfied with small data payloads and low data rates and therefore are not overly concerned the issue of fragmentation, a number of applications with demand for large payloads and data rates exist. Examples for such applications are smart metering and structural health monitoring. Both produce comparatively large application data that periodically has to be collected and forwarded or processed. In the other traffic direction, over-the-air-programming (OTAP) of nodes usually is concerned with the transport of large data blobs to reprogram nodes within a wireless network.

With fragmentation being expected to have some impact on the performance of transmissions of large datagrams, it is desirable to have some quantitative information available on exactly how strong this impact can be. At the time of writing, several studies exist that examine the performance of 6LoWPAN fragmentation using either analytical models or in most cases very simple experimental setups with only a few number of wireless hops. All of them only cover the most basic forwarding strategies. While some problems are identified, at the moment no comprehensive evaluation of the 6LoWPAN fragmentation in more realistic multi-hop network environments and considering enhancements to the forwarding exist.

This dissertation aims at providing such a comprehensive overview over 6LoWPAN fragmentation and contains several contributions towards this aim. To be able to better assess the influence of fragmentation on reliability, an extension to an existing analytical model that better captures the realities of current 6LoWPAN implementations is presented. With the help of the model, it is possible to get an estimate of the impact of fragmentation in multi-hop networks.

Furthermore, a parameter study with regard to the IEEE 802.15.4 MAC and implementation parameters like the available data buffer size is carried out in simulation and a testbed of 13 nodes. It provides an overview about suitable configuration of the underlying IEEE 802.15.4 MAC in multi-hop traffic collection scenarios.

Because of initially inexplicable results in various testbed setups that deviated strongly from corresponding simulations, a detailed examination of the state of the MAC and PHY layers was carried out and revealed that the implementation of the so-called “extended operating mode” of the used transceiver hardware caused the reliability of transmissions to drop dramatically. While this effect is especially strong for the traffic pattern caused by 6LoWPAN fragmentation, it can be generalized to other scenarios as well and may cause bias to experiment results, whenever the extended operating mode of this transceiver or similar modes of operation on other transceivers is used.

To improve the overall reliability of fragmented transmissions, a novel forwarding strategy is proposed. The 6LoWPAN ordered forwarding (6LoOF) protocol is designed to reduce contention for the wireless channel between nodes especially in collection traffic scenarios, while being compliant to the 6LoWPAN standard. A thorough evaluation of 6LoOF is presented in simulation and two testbed scenarios, facilitating the open experiment platform of the IoTLab.

Some of the above mentioned contributions have also been published as a research paper or article. Some chapters of this dissertation are based on and reuse parts of these papers. The following list provides an overview on the publications reappearing in this dissertation and clarifies the part of work done by me and the other authors.

- Chapter 3 is based on [WT14], which was created by me.
- Chapter 5 is based on [Wei+14b]. Martin Ringwelski provided the majority of the 6LoWPAN implementation for CometOS, the idea to the progress-based retry control (PRC) forwarding mode and was involved in the evaluation. I developed the Direct-ARR mode, created most simulation scenarios and was responsible for a major part of the evaluation. Andreas Timm-Giel and Volker Turau gave feedback and made suggestions with regard to evaluation and editing.
- Chapter 6 is based on [WT15], which was created by me.
- Chapter 4 introduces CometOS ([UWT12]), which was initiated by Stefan Unterschütz and developed by Stefan Unterschütz, me and Florian Kauer.

This dissertation is structured as follows: Chapter 2 introduces the problem domain, points out the most important protocols and approaches and defines the research goals of the dissertation. The analytical model developed as part of this dissertation is introduced in Chapter 3. This chapter also discusses the output of the model for a certain sets of inputs, including different paths lengths, number of fragments, retransmissions and different forwarding strategies. Chapter 4 describes the used frameworks and tools for simulation environment and testbed deployments and the simulation model. Furthermore, the approach to derive simulation models from testbed deployments that serves as a validation mechanism of the simulation model is introduced. Chapter 5 contains a study on basic forwarding strategies for 6LoWPAN fragments of large IPv6 datagrams. Due to a combination of sub-optimal physical link layer model and the transceiver's operating mode chosen for the testbed, this chapter can be characterized as a "lessons learned" chapter. The issue of this operating mode is examined in detail in Chapter 6. An experimental methodology to assess the impact of the used transceivers "extended operating mode" is developed and applied. Chapter 7 contains a revised parameter study of 6LoWPAN forwarding strategies using simulations and a testbed environment. The 6LoOF protocol is introduced and described in detail in Chapter 8. Furthermore, an evaluation of the 6LoOF protocol in comparison to the basic forwarding modes is presented. The dissertation is concluded in chapter 9.

2 Problem Statement

This chapter introduces IEEE 802.15.4 and 6LoWPAN, discusses basic forwarding strategies for 6LoWPAN fragmentation, and introduces a typical protocol stack for the Internet of Things (IoT). Application scenarios are presented to support the significance of evaluating 6LoWPAN fragmentation performance. Furthermore, the operation in energy-constrained networks is discussed and goals of the experimental and simulative evaluation carried out in this dissertation are stated.

In this and the following chapters, the text refers to units of data that are transmitted by a node or a protocol layer, i.e., “packets”. To avoid confusion, in this dissertation the following nomenclature is used:

- **frame**: A data frame (header + PHY service data unit (PSDU)) in context of the IEEE 802.15.4 protocol, i.e., data packets used by the link layer.
- **fragment**: A frame carrying 6LoWPAN fragmentation information and part of a datagram as payload.
- **datagram**: An IPv6 datagram. Represented by multiple fragments, if 6LoWPAN fragmentation is applied.
- **packet**: A datagram carrying UDP header and payload of the TCP/IP application layer. In the context of this dissertation it translates into a single IPv6 datagram.

2.1 IEEE 802.15.4

The “IEEE Standard for Local and metropolitan area networks – Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs)” (in this thesis referred to as IEEE 802.15.4) was first published in 2003, with major revisions in 2006, 2011 and 2016 [06; 11a; 16]. It defines several physical (PHY) layers and a medium-access layer (with several extensions) for low-rate wireless networks. One of the most widely used PHYs for sensing application, for which also a large number of transceivers is available, is the one operating in the 2.4 GHz ISM band. IEEE 802.15.4 is used as PHY and MAC layer for the industry standard ZigBee [12b] and the IETF standard 6LoWPAN [Mon+07].

IEEE 802.15.4 defines two general operating modes: beacon-enabled and non-beacon-enabled. The former employs so-called beacons, which are regularly broadcasted by the coordinator of a personal area network (PAN coordinator). By means of those beacons, a superframe structure is established, which consists of a contention access period (CAP) and a contention-free period (CFP). In the former, nodes content for access of the channel using a slotted carrier sense multiple access with collision avoidance (CSMA/CA) protocol and may also try to allocate a guaranteed time slot

(GTS) from the CFP. In the latter, nodes can use a previously allocated GTS to communicate with the PAN coordinator. During the guaranteed time slots that are not allocated to a node, this node may turn off its transceiver to save energy, which is the only possibility for duty cycling explicitly defined by the standard. Other protocols may define low-power listening or low-power probing techniques but are out of the scope of the IEEE 802.15.4 standard.

Until recently, the beacon-enabled mode only supported single hop, i.e., star topologies. An extension to IEEE 802.15.4, the distributed synchronous multi-channel extension to IEEE 802.15.4 (DSME) [12a], describes a method to extend this TDMA scheme to multiple hops and multiple channels. A similar direction takes another extension named TSCH [WPG15], which also uses a TDMA scheme, albeit without making use of the beacon-enabled mode's superframe structure. TSCH is derived from the industry standard WirelessHART [10]. While TSCH defines the mechanisms for nodes to communicate according to an existing communication schedule, it does not provide any protocols to actually establish such a schedule. Recent research efforts in that direction are Orchestra [Duq+15] and 6top, which is a standardization effort by the IETF currently in draft state [WV16].

The non-beacon mode operates without any superframe structure or regular beacons. Nodes transmit frames using an unslotted CSMA/CA protocol, which includes a random backoff period, a clear channel assessment (CCA) and subsequent backoffs in case the channel is considered busy.

Independently of the mode used, IEEE 802.15.4 defines retransmissions and acknowledgment frames for unicast transmissions. Receivers of a frame transmit an acknowledgment after they receive a unicast frame with the `Ack Request` flag set in the frame control field. The ACK is sent after a short delay without executing the CSMA/CA mechanism.

2.2 6LoWPAN

Being the protocol this thesis examines, the 6LoWPAN protocol is introduced in this section, with the main focus on 6LoWPAN fragmentation.

2.2.1 Compression and Fragmentation

The reasons for the existence of the 6LoWPAN protocol are twofold: First, IPv6 specifies a minimal maximum transmission unit (MTU) of 1280 B for any link-layer protocol below it. To transport IPv6 datagrams over an IEEE 802.15.4 link-layer with a maximum PHY layer payload of 127 B, fragmentation at the 6LoWPAN layer is necessary to present an interface to the IPv6 layer that supports a sufficiently large MTU.

Secondly, the headers for IPv6 (40 B) and UDP (8 B) or TCP (20 B) are large compared to the typical maximum PHY frame payload of 127 B. The IEEE 802.15.4 MAC header occupies up to 25 B and AES-CCM-128 encryption may use up another 21 B. Hence, the available payload size for a UDP packet in one frame can be reduced to 33 B. Complete use of these 33 B yields an overhead ratio of 74.4 % (adding 2 B for PHY header and the start of frame delimiter (SFD)). To reduce this high overhead ratio, 6LoWPAN defines several compression algorithms (HC1 and HC2), which in

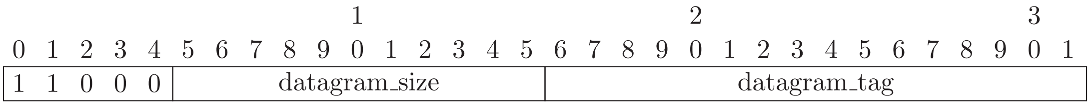

6LoWPAN FRAG_1 header

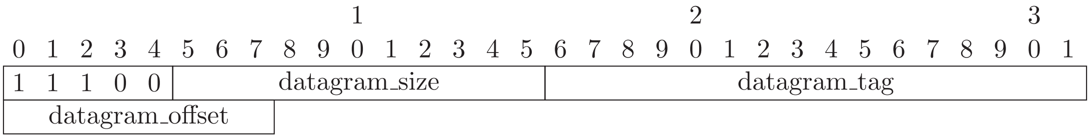

6LoWPAN FRAG_N header

Figure 2.1: 6LoWPAN fragmentation headers

turn are updated by 6LoWPAN IPHC header compression (IPHC) [HT11]. While compression is an important topic especially for small IPv6 datagrams with only a handful of bytes payload, it is seen as a problem orthogonal to the performance of the fragmentation mechanism.

To implement fragmentation, 6LoWPAN defines two different fragmentation headers, one for a first fragment (FRAG_1) and a different one for any subsequent fragment (FRAG_N; Fig. 2.1). Both include the uncompressed size of the IPv6 datagram and a tag to identify the datagram the fragment belongs to. The FRAG_N header additionally carries an offset field, which defines the position of the fragment within the whole datagram given in a unit of 8 octets.

Provided with experience concerning the fragmentation of large data blocks while working at the iEZMesh project, I expected fragmentation to amplify existing problems in multi-hop wireless mesh networks. iEZMesh was a project funded by the German government. One of the application requirements identified for the project was the collection of smart meter measurement tables sized 1 kB to 3 kB. With the link-layer supporting frame sizes of 128 B, the preconditions are similar to those found with large fragmented IPv6 datagrams over IEEE 802.15.4. To satisfy the requirements, we implemented a fragmentation mechanism at the transport layer, together with an actual retransmission scheme for individual fragments based on negative acknowledgments. The evaluation of the performance yielded significant reliability issues for the large data blocks [Wei+14a], even in the presence of the mechanism for retransmissions.

Considering that the loss of a single fragment leads to the loss of a whole datagram and the fact that typical wireless transmissions are inherently lossy due to the properties of wireless channels, collisions and interference, I expect that 6LoWPAN fragmentation is confronted with similar issues and that transmissions of large IPv6 datagrams via 6LoWPAN may exhibit low reliability. Further it is to be expected that the impact of fragmentation increases with the number of fragments and the length of a route. These considerations motivate the evaluation of the performance of fragmentation and several forwarding strategies and the development of the new forwarding protocol 6LoOF, which are presented in this thesis.

2.2.2 6LoWPAN Routing Schemes

With 6LoWPAN, routing in general can be performed at two different layers. First, a layer at the level 2.5 of the 6LoWPAN adaptation layer implements the routing. In that case, some mesh routing protocol has to emulate a full broadcast domain at the physical level for the IPv6 layer. This variant is called ımesh-under routing [HC08; Cho+09]. Thereby, link-local addresses and link-local multi-cast can easily be used from an IPv6 layer perspective and IPv6-based protocols can theoretically be left unchanged. An example for this is the neighbor discovery protocol [Nar+07]. Neighbor discovery makes extensive use of link-layer multicasts, which have to be translated to flooding the mesh network. This means that the mesh routing protocol has to provide potentially complex mechanisms to offer reliable operation over a multi-hop mesh network, which is far from trivial. Moreover, such mechanisms are already available at the IPv6 network layer and have to be recreated for the layer 2.5 mesh routing [HC08].

The other possibility is to delegate routing decisions to the IPv6 layer: Every hop in a meshed 6LoWPAN network becomes an IPv6 routing hop. This routing scheme is called ıroute-over [HC08]. Using route-over has several implications. First, global IPv6 addresses ([Nar+07]) have to be used, because IPv6 forbids routing of link-local addresses. That makes the original HC1 and HC2 compression algorithms impractical and is one reason for the introduction of the IPHC and 6LoWPAN next header compression (NHC) [HT11]. With regard to fragmentation, route-over also means that datagrams – sticking to a strict separation of layers – have to be reassembled at each intermediate hop of the 6LoWPAN mesh network, because fragmentation is then handled below the network layer.

With the creation of RPL [Win+12], a standardized routing protocol for low-power and lossy networks at the IPv6 layer is available to be used with a route-over routing scheme, which perfectly fits the usual demands on routing protocols for typical 6LoWPAN wireless mesh networks. Considering the arguments, I decided to focus on the evaluation of route-over, as it allows building a completely standardized protocol stack and avoids the awkward emulation of a single hop broadcast domain above a lossy multi-hop wireless mesh network.

2.2.3 Basic Route-Over Forwarding Techniques

As described in Sect. 2.2.2, using the route-over routing scheme implies reassembling a datagram at every intermediate node of the wireless mesh network. This is the first basic and most straightforward forwarding strategy and is called *Assembly* or *Assembly mode* throughout the thesis. During the whole process of reassembling the datagram, all fragments have to be stored in some buffer, even at intermediate nodes, which are not concerned with the content of the datagram. Hence, for each datagram in transit, buffer space for the whole datagram has to be available. Considering typical resource-constraint hardware for wireless sensor networks, this is a non-negligible issue. Furthermore, reassembling at every intermediate hop prevents pipelining of fragments on longer (> 3) paths. Therefore, an unnecessary large end-to-end latency can be expected (Fig. 2.2).

In contrast to this approach, which strictly preserves layer separation, a cross-layer approach can be employed, which is called *Direct* or *Direct mode* in the remainder of

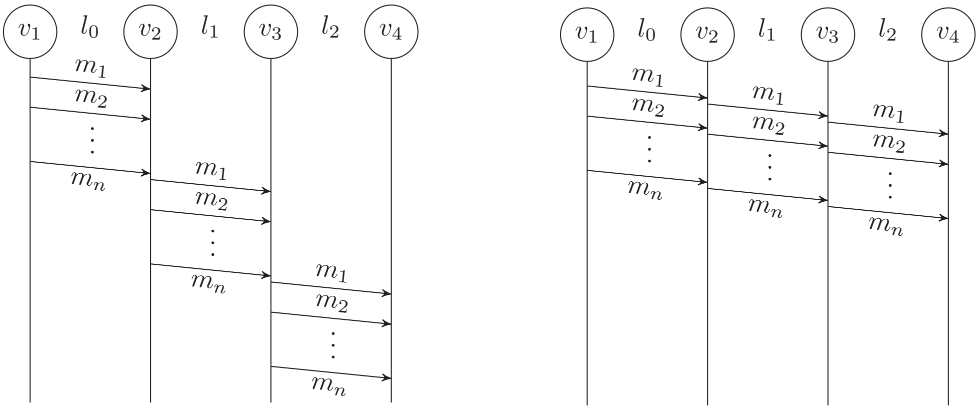

(a) Assembly (b) Direct

Figure 2.2: Message flow in Assembly and Direct modes. The Direct mode has potential for pipelining as well as an increased probability for collisions.

the thesis. For each incoming first fragment, the information necessary to identify and process subsequent fragments of the datagram is stored in a “virtual fragment buffer”. The buffer is called virtual, because it does not store the payload data the fragment carries, but only the metadata, i.e., information about progress and identity of the fragment. The fragment itself is immediately scheduled for transmission to the next hop, which is queried directly from the IPv6 layer. This is always possible, because the IPv6 header does always fit the first fragment. Subsequent fragments then are matched against the entries in the virtual fragment buffer and routed along the same path. Note that the Assembly mode also uses the same path, but the moment at which the routing decision is taken is different: with Assembly, it is the reception of the last fragment, with Direct, the reception of the first fragment.

Using the Direct mode, datagrams have only to be stored for reassembly at their IPv6 destination (or the 6LoWPAN border router). Hence, it provides good potential for saving buffer space. Furthermore, pipelining on long paths becomes possible and thereby the overall latency can potentially be reduced. On the other hand, immediate forwarding also gives rise to self-interference. Fragments are prone to interfere with their predecessors, which have already advanced on the routing path. This can be especially harmful at the node two hops farther down the path, because such a node usually will be a hidden-terminal. In this situation, the clear-channel assessment part of IEEE 802.15.4’s CSMA/CA algorithm is not able to prevent a collision. This increased potential for collisions is also implied in Fig. 2.2.

2.2.4 Adjacent Protocols

The IPv6-based standardized protocol stack for resource constrained networks and devices then could be the one shown in Fig. 2.3.

Above the described combination of a route-over 6LoWPAN adaption layer and the IPv6 network layer with RPL as routing protocol, the combination of UDP [Pos80] and CoAP [SHB14] is used at the transport and application layers.

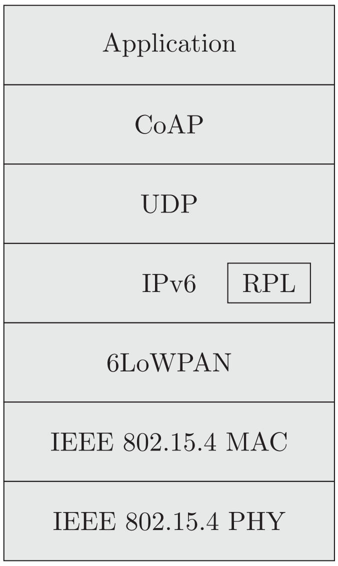

Figure 2.3: Standard protocol stack for low-power lossy networks

CoAP is similar in spirit to Hypertext Transfer Protocol (HTTP) in defining addressable resources as RESTful services. Additionally, it defines simple reliability features, a basic congestion control mechanism and a binary representation to increase efficiency. Hence, it takes on some responsibilities of a transport layer. Current standardization and research efforts for CoAP focus, among others, on congestion control [BGD15; Bet+15; JDK15] and blockwise CoAP transport [BS16]. The latter introduces a mechanism to CoAP to split up large payloads into smaller blocks with the target of avoiding to burden lower layers with "conversation state that is better managed in the application layer" [BS16]. This includes 6LoWPAN fragmentation and aims at reducing the need for it and thereby is fundamentally different from the approach presented in this thesis, which is to improve the performance of 6LoWPAN fragmentation.

The RPL protocol [Win+12] defines a tree-based routing for low-power lossy networks, such as wireless sensor networks or cyber-physical systems. Its main ideas are derived from the collection tree protocol [Gna+09]. It builds bi-directional routing trees by two core mechanisms:

- Beacons called DIOs are used to form routes towards a single destination: the root of the destination oriented directed acyclic graph (DODAG). The DIOs propagate from the DODAG root through the network governed by the Trickle algorithm [Lev+11].

- Communication in the opposite direction is enabled by letting each node in the tree periodically send so-called DAOs to the DODAG root. That way, reverse routes are installed either at each intermediate node (storing mode) or exclusively at the root, which uses source routing to transmit to arbitrary nodes.

RPL has been extensively evaluated [TOV10; HC11; YCI13; CMN14; KG14; Ise+15] and is emerging as the protocol of choice for static low power and lossy networks.