This item was submitted to [Loughborough's Research Repository](#) by the author.
Items in Figshare are protected by copyright, with all rights reserved, unless otherwise indicated.

IPv6 multicast forwarding in RPL-based wireless sensor networks

PLEASE CITE THE PUBLISHED VERSION

<http://dx.doi.org/10.1007/s11277-013-1250-5>

PUBLISHER

© Springer Science+Business Media

VERSION

AM (Accepted Manuscript)

LICENCE

CC BY-NC-ND 4.0

REPOSITORY RECORD

Oikonomou, George, Iain W. Phillips, and Theo Tryfonas. 2019. "Ipv6 Multicast Forwarding in Rpl-based Wireless Sensor Networks". figshare. <https://hdl.handle.net/2134/13147>.

This item was submitted to Loughborough's Institutional Repository (<https://dspace.lboro.ac.uk/>) by the author and is made available under the following Creative Commons Licence conditions.

For the full text of this licence, please go to:
<http://creativecommons.org/licenses/by-nc-nd/2.5/>

IPv6 Multicast Forwarding in RPL-Based Wireless Sensor Networks

George Oikonomou · Iain Phillips ·
Theo Tryfonas

Received: date / Accepted: date

Abstract In wireless sensor deployments, network layer multicast can be used to improve the bandwidth and energy efficiency for a variety of applications, such as service discovery or network management. However, despite efforts to adopt IPv6 in networks of constrained devices, multicast has been somewhat overlooked. The *Multicast Forwarding Using Trickle* (Trickle Multicast) internet draft is one of the most noteworthy efforts. The specification of the *IPv6 Routing Protocol for Low power and Lossy Networks* (RPL) also attempts to address the area but leaves many questions unanswered. In this paper we highlight our concerns about both these approaches. Subsequently, we present our alternative mechanism, called Stateless Multicast RPL Forwarding algorithm (SMRF), which addresses the aforementioned drawbacks. Having extended the TCP/IP engine of the Contiki embedded operating system to support both Trickle Multicast (TM) and SMRF, we present an in-depth comparison, backed by simulated evaluation as well as by experiments conducted on a multi-hop hardware testbed. Results demonstrate that SMRF achieves significant delay and energy efficiency improvements at the cost of a small increase in packet loss. The outcome of our hardware experiments show that simulation results were realistic. Lastly, we evaluate both algorithms in terms of code size and memory requirements, highlighting SMRF's low implementation complexity. Both implementations have been made available to the community for adoption.

George Oikonomou (✉)

University of Bristol, Faculty of Engineering, Merchant Venturers Building, Woodland Road, Clifton, BS8 1UB, UK

E-mail: g.oikonomou@bristol.ac.uk

Previous address: Loughborough University, Computer Science, Loughborough, LE11 3TU, UK

Iain Phillips

Loughborough University, Computer Science, Loughborough, LE11 3TU, UK

E-mail: i.w.phillips@lboro.ac.uk

Theo Tryfonas

University of Bristol, Faculty of Engineering, Queen's Building, University Walk, Clifton, BS8 1TR, UK

E-mail: theo.tryfonas@bristol.ac.uk

Keywords 6LoWPAN · Wireless Sensor Networks · IPv6 Multicast · Trickle

1 Introduction and the Need for Multicast in Wireless Sensor Networks

Over the past years, the research community has invested considerable efforts towards the seamless integration of wireless sensor networks (WSNs) with the Internet. Previous work has demonstrated that pure IPv6-based WSN architectures are not only viable but can also outperform application-centric designs [19]. Significant standardisation efforts have contributed to mature, interoperable implementations of embedded IPv6 stacks, such as uIPv6 which is distributed as part of the Contiki embedded Operating System. Among those standards are RFC 4944 [31] and RFC 6282 [20]. Published by IETF's 6LoWPAN work group, they discuss techniques for IPv6 datagram fragmentation and header compression, in order to achieve their efficient transmission within IEEE 802.15.4 low power radio frames. For those networks (6LoWPANs), the most widely adopted standard for routing is the "IPv6 Routing Protocol for Low-Power and Lossy Networks" (RPL), which is specified in RFC 6550 [44].

The importance of network layer multicast forwarding in 6LoWPANs stems from its ability to improve the efficiency of applications adopting a one-to-many communication paradigm. Examples of services which can benefit by multicast include service discovery [2, 21], network management and publish/subscribe schemes. For instance, Extended Multicast DNS (xmdns) builds on the mDNS specification [6] and expands it *"to site-local scope in order to support multi-hop LANs that forward multicast packets but do not provide a unicast DNS service"* [29]. uBonjour [21] is a service discovery scheme for resource-constrained devices in the wireless embedded internet. It is essentially Bonjour's lightweight variant and is based on multicast DNS (mDNS) and DNS service discovery (DNS-SD) [5], which have recently had their message sizes optimised for 6LoWPANs [22]. Additionally, the Constrained Application Protocol (COAP) is an emerging standard aiming at the integration of WSN devices with the web. It operates over UDP (mandatory) or TCP (optionally) and has been designed so that its messages can be easily translated to HTTP. It targets embedded devices with severe memory and power supply restrictions. The current version of its specification provides an extensive discussion on its operation over IPv6 multicast [40].

Previous research efforts in the area of multicast for WSNs focus on bespoke network stack designs and do not investigate IPv6-specific challenges. The majority of multicast forwarding algorithms encountered in current literature are based on geographic routing [39, 38, 3, 14, 23, 42]. However, most of those approaches have certain characteristics and make assumptions which render them unsuitable for IPv6-based WSN deployments. For instance, many of them assume that, for every multicast message, the sender is aware of the addresses or IDs of all intended destinations. Additionally, some efforts suffer from poor scalability while others rely on unrealistically large network packets. Lastly, they are only applicable in situations where the source as well as all destinations are within the WSN boundaries.

Despite the number of existing efforts, IPv6-based multicast has been somewhat overlooked by the 6LoWPAN research community, as we discuss further in section 2. The "Multicast Forwarding Using Trickle" internet draft (Trickle Multicast - TM) [17]

discusses an algorithm which poses among the most suitable candidates. The RPL RFC also briefly mentions multicast, but the discussion focuses on group management without providing a sufficient level of detail in terms of forwarding.

The open issues outlined above have motivated us to design and implement SMRF, a lightweight Stateless Multicast RPL Forwarding algorithm. In this paper, we disclose its design in depth and we highlight how it addresses current open issues while maintaining high speed, energy efficiency and low complexity. Compared to geographic multicast algorithms, SMRF does not require geolocation information (neither explicit nor via a location service) and does not suffer from the aforementioned scalability and datagram size issues. Moreover, by recommending a forwarding algorithm, SMRF fills the gaps left open by the RPL RFC. In this context, this paper's contributions are the following:

- We compare the performance of SMRF against TM. For the evaluation, we consider four metrics: i) Packet delivery ratio, ii) End-to-end delay, iii) Out-of-order datagram delivery ratio and iv) Energy consumption. Evaluation is performed through simulations as well as on a multi-hop hardware testbed. We also investigate the complexity of the two algorithms and we compare their code size and memory requirements. Simulation and testbed results demonstrate that SMRF is less complex and that it outperforms TM on three of the four metrics.
- Based on the outcome of the comparative evaluation, we present a host of criteria, which can be used by network engineers and designers in order to select the more suitable between SMRF and TM, depending on their deployment's specific needs.
- We have extended Contiki's TCP/IP stack to support both algorithms. Both implementations have been released¹ to the community for adoption and further scrutiny as a part of our port of the Contiki OS^{2,3} [32].

This paper extends our previous work [33], providing the following additional contributions: i) Extended design details for SMRF, ii) Evaluation of an additional metric: the ratio of datagrams delivered out of order by TM (sec 5.4), iii) Additional simulation experiments in a different topology for the evaluation of both algorithms on a hop-by-hop basis (sections 5.2.1, 5.3.1 and 5.5.1), iv) Results from the evaluation of both algorithms on a hardware testbed (sec. 6) and lastly v) Discussion on the code size and memory requirements for both algorithms (sec. 7).

2 Related Work

2.1 Multicast in Traditional WSNs

Previous research efforts in the area of multicast for WSNs have primarily been focusing on traditional, application-centric network designs and as such do not address IPv6-specific challenges. Multicast forwarding algorithms based on geographic routing are dominant in existing bibliography and can be broadly classified as either

¹ <https://github.com/g-oikonomou/contiki-sensinode/tree/mcast-forward>

² <https://github.com/g-oikonomou/contiki-sensinode/wiki>

³ <http://nets-www.lboro.ac.uk/george/contiki-sensinode/>

purely geographic [39,38,3,14] or *hybrid* [23,42], whereby the geographic component is complemented by features of other approaches, such as hierarchical routing.

The Geographic Multicast Routing (GMR) algorithm builds on existing unicast geographic routing approaches. By adapting them, it aims to achieve multicast message delivery to all intended destinations while maintaining minimum bandwidth consumption. Nodes exchange position information with their neighbours through periodic beacons. GMR is characterised by low computational complexity and a small memory footprint [39]. According to subsequent works, GMR scales better than some of its predecessors [38] but still suffers from scalability issues when dealing with large deployments [23].

Using periodic beacons can have negative side-effects such as collisions and increased energy consumption [38]. Based on this observation, BRUMA attempts geographic multicasting without beacons, whereby neighbour positions are discovered reactively. Next hop selection happens opportunistically through a mechanism which only requires a low number of control messages. The authors demonstrate that BRUMA is more efficient than GMR [38].

Carzaniga et al. propose a compact and completely decentralised multicast scheme with asymptotically optimal network congestion properties [3]. It operates by building a multicast forwarding tree over an underlying geographic unicast routing service, which allows nodes to send messages to a destination defined by a coordinate pair (x, y) .

Receiver-Based Multicast (RBMulticast) [14] is another purely geographic approach. Its principal novelty lies in the next-hop determination phase: potential next hops contend for the channel based on their contribution towards delivering a packet to its destination. Nodes offering the highest forward progress have higher probability of getting selected as next hop. By adopting this approach, RBMulticast can operate without routing tables and without maintaining a forwarding tree. To achieve this, RBMulticast embeds the geographic location of all destinations in the packet header. This raises questions regarding its scalability in large deployments.

The Hierarchical Geographic Multicast Routing (HGMR) [23] is a hybrid algorithm combining the key design concepts of GMR [39] and the Hierarchical Rendezvous Point Multicast (HRPM) protocol [8]. The resulting HGMR algorithm is further optimised to be more energy efficient and scalable. HGMR divides multicast groups into subgroups by using HRPM's geolocation hashing. It takes advantage of layer 2 reliability mechanisms by using HRPM's unicast forwarding approach for long, sparse paths and reverts to layer 2 broadcasting in areas of high density in order to reduce the number of transmissions. When unicast forwarding is in use, HRPM (and therefore HGMR) uses source routing along the branches of an overlay tree generated by the traffic source. In this work the authors conducted a performance evaluation of the three algorithms in a simulated IEEE 802.11 network. It is therefore difficult to understand how the algorithms would behave under the frame size and bandwidth limitations or the loss characteristics related to IEEE 802.15.4 networks.

The Multicast Routing with Branch Information Nodes (MR.BIN) [42] protocol is a hybrid approach, combining geographic unicast routing with state-based multicast. It maintains multicast states only on branch nodes of the forwarding tree. Communication between non-branching nodes takes place with geographic unicast. In order

to perform datagram forwarding, each node maintains a potentially long list of next hops. Messages are tagged with a 2-byte multicast group identifier but the management scheme for those IDs is not discussed; it is unclear how a node can choose which group ID to join and how two different groups are prevented from having the same ID.

Adaptive Geo-Source Multicast Routing (AGSMR) [43] is a geographic unicast and source multicast hybrid. It relies on generating a forwarding tree at the traffic source. Path information is embedded in packet headers in a compressed format which uses 2 bytes per hop, an additional 2 bytes per branch and reduced by $2(n - 2)$ bytes for each n -node long, non-branching path. The authors demonstrate that, with this compression scheme in place, AGSMR's packets are smaller than GMR's. However, reported packet sizes have an order of magnitude of KBytes (e.g. about 9KB for a tree with about 50 subscribers in a network of 1024 nodes). Packet size increases with the number of subscribers as well as the total number of network nodes.

Branch Aggregation Multicast (BAM) [34] differs from the aforementioned efforts in that it operates without knowledge of node geolocation. BAM's design has two components: S-BAM achieves single-hop aggregation at branching nodes and M-BAM aims to reduce the number of branches. The former tries to combine multiple layer 2 unicasts within a single radio frame, by extending headers to list the addresses of all intended destinations. Recipients reply with an acknowledgement frame, using a random back off to avoid ACK collisions. M-BAM relies on existing forwarding tables, which can be populated by any routing protocol. The authors make the assumption that the routing table contains multiple candidate next hops for the same destination and they propose an algorithm for path aggregation. However, routing table size has an impact on scalability with increasing network size [32] even when the table only lists a single next hop per destination. Storing multiple candidate next hops per destination would impose further memory overheads⁴ and should not be considered common practise. Furthermore, BAM does not manage multicast groups internally. Instead, it assumes that the network adopts a data centric routing model, whereby nodes broadcast requests for data and their neighbours remember querying node addresses in order to forward relevant data accordingly [34].

All aforementioned algorithms assume that the message source is aware of and maintains a list of all destinations, uniquely identified by an attribute such as a node ID, address or name. This is untrue in the case of IP and IPv6 multicast, whereby datagram destination is expressed as the multicast group's IP or IPv6 address and the traffic source is oblivious with respect to the unicast address of each individual group member.

The geographic multicast routing protocols discussed above also require knowledge of the geolocation of all destinations. When this information is not available, a geographic routing service can be used in conjunction with a location service, such as MLS [15] or GLS [27]. Querying the location service incurs time overheads which have an impact on delivery delay, compared to non-geographic approaches which do not rely on this information. However, this is not considered at all in relevant work.

⁴ In the version of the Contiki OS used for this research, each entry in the IPv6 routing table occupies approximately 48 bytes of RAM, the exact number depending on the hardware platform and toolchain.

Many of the above multicast forwarding algorithms embed a list of all destinations in the packet header [39,38,14,42,43,34]. This increases the byte overhead associated with each data transmission and has an adverse impact on the protocol's scalability [23].

Lastly, with all those approaches traffic is confined within the deployment's boundaries. In order to be able to communicate beyond the WSN's borders, one would need a dedicated gateway. This does not apply to 6LoWPANs and IPv6 multicast forwarding protocols, such as those discussed later on in this paper.

2.2 Multicast in 6LoWPANs

Previous research with focus on 6LoWPANs is very limited. Sá Silva et al. investigated the applicability and usefulness of traditional multicast paradigms in WSNs [37]. In this work, the authors evaluate Multicast Ad hoc On-Demand Distance Vector (MAODV) [36] and Source-Specific Multicast (SSM) [1] in a simulated sensor network. They demonstrate that, in their simulated scenarios, multicast can offer significant advantages over traditional network-wide broadcast flooding and unicast. Results also suggest that MAODV and SSM are comparable in terms of bandwidth consumption and energy efficiency.

Clausen and Herberg conducted relevant research with focus on RPL networks [7]. Despite the fact that this work only discusses network-wide broadcast, it offers some important insight on related issues and techniques. In their contribution, the authors conduct experiments by simulating an IEEE 802.11b network, which cannot capture the duty cycling aspects of modern WSNs, nor the low-power, lossy nature of IEEE 802.15.4-compliant radio hardware.

The RPL RFC [44] briefly discusses built-in multicast support. In RPL networks, nodes advertise unicast downward paths inside Destination Advertisement Object (DAO) messages. An RPL instance is administratively configured with one out of a possible four Modes of Operation (MOP). Nodes may only participate in the network as routers if they support the advertised MOP, otherwise they may only join as leaf nodes. In the *Storing with multicast support* MOP, DAO messages are also used for group management by advertising multicast prefixes. Unfortunately, the approach leaves many open issues, as we discuss in greater detail in sec. 4.

3 Multicast Forwarding with Trickle

In wired networks, multicast mechanisms rely on topology maintenance in order to forward packets to their intended destinations. Due to memory restrictions, this is a very challenging task in networks of constrained nodes. The *Multicast Forwarding with Trickle* algorithm (Trickle Multicast - TM) addresses this challenge by providing a means of supporting IPv6 multicast without having to rely on topology information. To control the frequency of datagram exchange, TM adopts the pre-existing Trickle algorithm.

Trickle [26,25] is a mechanism that governs the frequency of periodic information exchange among neighbouring nodes in a low power, lossy network. Strictly

speaking, trickle only specifies the dynamic behaviour of periodic timers. Its aim is to provide a method of propagating state information efficiently, without constantly flooding the network with control messages. In simple terms: when two single-hop neighbours share the same knowledge (*agree*), control message exchange rate slows down exponentially, achieving energy and bandwidth efficiency. Conversely, when an inconsistency is detected, the timer's interval (I) is reset to a minimum value (called I_{min}) and changes propagate within milliseconds. After a long period without changes, the trickle timer reaches its maximum interval called I_{max} , which is expressed as the maximum number of I_{min} doublings. Thus, for a trickle timer configured with (I_{min}, I_{max}) , the maximum trickle interval in time units will be $I_{min} * 2^{I_{max}}$ and can result in very infrequent message exchange, with nodes sending only a few packets per hour [25]. Trickle only dictates the behaviour of timers, in other words *when* nodes should exchange messages, not *how* nor their format. This trait makes it very attractive for any protocol involving periodic exchange of state information. Trickle was originally designed for data dissemination and network re-programming [24] and has been adopted by multiple works, such as Deluge [16] and DIP [28]. It also handles the frequency of RPL DIO (upward route advertisement) messages [44]. Lastly, trickle underpins the TM forwarding algorithm which is discussed in the remainder of this section.

With TM, each multicast datagram must carry a *Multicast Option* header in the shape of an IPv6 Hop-by-Hop Option (HBHO) extension header. The multicast option tags the packet with a sequence number, a single-bit M parameter and the unique identifier (Seed ID) of the sender, which may be different than the datagram's IPv6 source address. The contents of the multicast option remain unchanged en route. Each network node maintains a cache of recently seen multicast packets, uniquely identified by the information in the HBHO. Upon reception of a multicast datagram, a node inspects the multicast option and, if the packet is new, it gets added to the cache.

Neighbouring nodes use ICMPv6 datagrams to exchange information about their cache contents, at a frequency controlled by trickle timers. If a node's cache contents don't match the information in a received ICMPv6 datagram, the node resets its trickle timer to its minimum interval (I_{min}) in order to facilitate quick propagation of new packets. Inconsistency is also triggered upon reception of a new multicast datagram. At every trickle interval, nodes forward inconsistent datagrams to their single-hop neighbours inside link-layer broadcast frames.

3.1 Advantages

By design, TM has some very significant advantages. More specifically:

Generality: TM will work, without modifications, alongside any routing protocol.

Reliability: By caching datagrams and maintaining per-packet state information, TM increases its reliability (high packet delivery ratio / low loss). The exact reliability levels are heavily influenced by the choice of I_{min} and the underlying duty cycling algorithm.

Guaranteed no duplicates: With the help of the HBHO, TM guarantees that each node in the network will receive each individual datagram at most once.

3.2 Concerns

The design of the TM algorithm raises a number of concerns which are outlined below and further analysed in the evaluation sections (sections 5, 6 and 7).

Scalability: One of the arguments for specifying the TM algorithm in the first place was that maintaining topology information is hindered by memory constraints. The algorithm bypasses this requirement and as a result scales well with the number of nodes in a network. However, topology maintenance is replaced with bespoke, per-packet state maintenance. This raises concerns regarding scalability with traffic volume, cache size and number of multicast traffic sources.

Performance: In order to avoid duplication, nodes never forward multicast datagrams immediately. Instead, they cache them and wait for ICMPv6 control messages. When an inconsistency is detected, the packets causing it are scheduled for transmission during the next trickle interval. This forwarding delay has an impact on end-to-end delay and can be heavily influenced by trickle parameters. The trickle RFC [25] dictates that “A protocol specification that uses Trickle *MUST* specify: Default values for *Imin*, *Imax*, and *k*...”. Currently, this is not the case for TM; its internet draft only outlines examples with indicative values [17]. Furthermore, as we demonstrate in sections 5 and 6, the values used in these examples are sub-optimal. We outline alternative recommendations, supported by experimental results.

Complexity: Nodes maintain two trickle timers, a sliding window for each source of multicast traffic and a cache of recent multicast datagrams. They also need to be able to create and process a new type of ICMPv6 message and a new type of HBHO extension header. Especially in the case of incoming ICMPv6 messages, a node needs to compare all entries in the message against all cached messages. This raises concerns in terms of complexity, code size and memory requirements. This is further investigated in section 7.

Multicast vs Broadcast: Due to lack of topology maintenance and group registrations, TM forwards all multicast messages to all parts of the network, irrespective of whether they are needed or not. Any datagram with a routable multicast IPv6 destination address is in practice treated as a *network-wide broadcast*. In sparse multicast topologies (where only a small percentage of nodes is interested in a multicast flow), adopting TM leads to energy and bandwidth inefficiencies.

Arrival Order: Due to its store and forward nature and per-packet state maintenance, TM is susceptible to out-of-order datagram arrivals. Depending on the application relying on multicast, this trait may or may not be a problem. For instance, code dissemination applications generally distribute an image in a number of chunks. Changes only get committed after all chunks have been received successfully, thus arrival order is highly irrelevant. However, in other types of applications (e.g. network management) a packet arriving out of order could cause undesirable behaviour. In this case, the application would have to employ a technique to detect out of order datagram delivery and mitigate its impact.

Trickle Multicast in RPL Networks: Even though there is no technical reason preventing TM and RPL from happily coexisting in the same network, contradiction in the respective specifications makes the situation less clear. More specifically, the TM internet draft specifies that “The Trickle Multicast option is carried in an IPv6 Hop-

by-Hop Options header, immediately following the IPv6 header” [17]. The RPL RFC specifies that all data plane datagrams also carry an HBHO, used for loop detection and avoidance. The format and functionality of this header is further specified in a separate RFC, which states that “*The RPL Option is carried in an IPv6 Hop-by-Hop Options header, immediately following the IPv6 header*” [18]. With both HBHOs immediately following the IPv6 header, co-existence of the two protocols becomes less straightforward.

4 Stateless Multicast RPL Forwarding - SMRF

In this paper, we contribute a multicast forwarding algorithm called SMRF, as an alternative to TM for RPL networks. The principal rationale behind SMRF is that nodes participating in an RPL network exchange topology information in order to build the basic RPL construct (called a Destination-Oriented Directed Acyclic Graph - DODAG) and to populate their routing tables. The DODAG is a tree structure and is thus particularly attractive to form the basis of multicast forwarding. Since network nodes perceive the network as a tree, we can capitalise on it in order to perform multicast forwarding without defining new message formats.

RPL nodes advertise downward paths inside Destination Advertisement Object (DAO) messages (data messages in RPL networks can only flow up or down the DODAG). According to its specification, one of RPL’s modes of operation (MOP) is *Storing with multicast support*. In this MOP, unicast DAO messages are also used to relay multicast group registrations up the DODAG. Those DAOs are identical to the ones conveying unicast information except for the type of prefix being advertised, which is a routable multicast IPv6 address.

Nodes can join a multicast group by advertising the group’s multicast address in their outgoing DAO messages, which only travel upwards in the DODAG. Upon reception of such message from one of its children, a router makes an entry in its routing table for the advertised multicast address. Conceptually, this entry indicates that *a node under us in the DODAG is a member of this group*. This router will then i) advertise this prefix in its own DAOs and ii) relay multicast datagrams addressed to this destination.

This RPL built-in mechanism solves the problem of propagating group membership information towards the DODAG root. However, it suffers from two limitations: i) It lacks a method which would prevent a node from accepting the same datagram twice or more. ii) The RPL RFC specifies that each router should copy multicast datagrams to a subset of its link layer neighbours, for instance only its preferred parent or only those children that are registered group members. This destination filtering can only be achieved by using frames with a unicast destination at the link layer. Thus, a node would have to transmit each datagram multiple times, once per intended recipient. This would incur additional costs in terms of traffic, delay and processing time and would be largely inefficient in dense networks. It would also increase memory requirements, since each router would need to maintain associations between multicast groups and neighbour subsets.

Fig. 1 The SMRF algorithm.

SMRF is a multicast forwarding algorithm that uses information provided by RPL's group membership scheme and addresses both those drawbacks. Its operation, illustrated in the flowchart in Fig. 1, is the following:

- A node will accept an incoming multicast datagram if and only if the datagram's link layer source address is the link layer address of the node's preferred RPL parent (which can be looked up in the node's neighbour cache).
- If the message gets accepted, it will get delivered up the network stack locally if and only if the node is a member of the multicast group.
- If the message gets accepted, it will get forwarded if and only if there is an entry for the datagram's IPv6 destination address (multicast group) in the node's routing table (a node below us in the DODAG is a group member).

4.1 Cross-layer optimizations

With SMRF, multicast datagrams are always transmitted as layer 2 broadcast frames, with some cross-layer optimisations in place so as to improve performance. To better understand how these work, it is necessary to visit the concept of duty cycling

Fig. 2 Broadcast Packet Transmission with ContikiMAC.

and 802.15.4 frame transmission with ContikiMAC [9,10], which is one of the main duty cycling algorithms used by the Contiki OS. Despite the misleading ‘MAC’ suffix, ContikiMAC is actually a duty cycling mechanism which can operate in conjunction with a variety of MAC layers (e.g. CSMA). In very simple terms, each node wakes up every few milliseconds and checks the channel for traffic. This interval is called Channel Check Interval (CCI) or Channel Sampling Period. If no traffic is present, the radio transceiver is turned back off. If traffic is detected, the node stays on until complete reception. To send a frame, a node will transmit it repeatedly (strokes) for slightly longer than CCI, waiting for a brief time interval between two strokes for a potential acknowledgement frame (ACK). This repeated transmission, often called a *packet train*, lasts long enough for intended recipients to wake up, detect the packet and receive it, irrespective of exactly when they last went to sleep. This removes the complexity of maintaining synchronisation between neighbours. In the case of unicast packets, the receiver will send an ACK frame, causing the sender to terminate its chain of strokes and thus conserve energy. However, broadcast frames must be received by all neighbours and there are no ACKs; the sender always has to go through the entire packet train, as illustrated in Fig. 2. Thus, broadcast transmissions are fundamentally more costly. The actual implementation of the ContikiMAC algorithm is more sophisticated and optimised (e.g. phase locks, burst support), but the concept remains the same.

A side effect of the algorithm discussed above is that when a node receives a broadcast frame, it should not attempt to transmit before the sender has gone through its entire packet train. Immediate transmission would signal a collision and the outgoing packet would be dropped. Thus, SMRF introduces a short delay (D), defined as $D = \max(Fmin, CCI)$ where CCI is the Channel Sampling Interval as reported by the

Table 1 Examples of SMRF configuration parameters and resulting forwarding delay

Duty Cycling		Configuration		Outcome	
Algorithm	CC1 (ms)	F_{min} (ms)	$Spread$	D (ms)	Final Delay (ms)
ContikiMAC	125	≤ 125	1	125	125
ContikiMAC	125	≤ 125	4	125	[125 , 500]
NullRDC	0	0	Ignored	0	Immediate
NullRDC	0	31.25	8	31.25	[31.25 , 250]

underlying duty cycling algorithm and F_{min} is a configuration parameter. Configuring SMRF with a non-zero value for F_{min} is particularly useful in the case of duty cycling algorithms which keep RF hardware always on ($CCI = 0ms$). One such duty cycling algorithm is Contiki's NullRDC.

In order to mitigate the negative effect of hidden terminals, SMRF can also optionally further delay datagram forwarding by a random factor. This is parametrised on $Spread$, a positive integer. The final forwarding delay is a random number in $[D, Spread * D]$ with granularity equal to D . Table 1 outlines the resulting forwarding delays for various configuration values and duty cycling algorithms.

4.2 Benefits and Drawbacks

With SMRF, multicast traffic can only travel downwards in the DODAG. This makes the algorithm useful for applications such as service discovery or network management. Since each node will only consider datagrams received from its preferred parent and will forward each packet at most once, it guarantees that each datagram can be received at most once per node, without need for a method of uniquely identifying messages.

The gain in comparison to TM is multi-fold: SMRF uses multicast groups to differentiate between nodes that are interested in a flow and those that are not. Instead of blindly forwarding all datagrams to all nodes, multicast datagrams will only reach parts of the network that have expressed an interest in the flow by joining a multicast group.

SMRF does not define any control messages of its own. It operates based on RPL parent information and on multicast group membership information, carried inside RPL DAO messages, as defined in [44].

Nodes do not need to maintain per-packet state. A *drop or forward* decision is taken for each datagram individually and is based on information available at the moment of its arrival. A positive side-effect of this *on the spot* approach is that SMRF can't entangle datagram ordering. Thus, unless message order gets shuffled around by an underlying layer, SMRF will deliver them in the correct order.

SMRF is very lightweight in terms of complexity, code footprint and memory requirements, as demonstrated in section 7. This makes it a very attractive option for severely constrained hardware.

Compared to TM, SMRF achieves lower end-to-end delays and demonstrates better energy efficiency. This is further analysed in sections 5.3 and 5.5.

Table 2 Simulation Configuration

<i>Nodes</i>	21 Sky motes (1 traffic source, 20 sinks)
<i>Radio Medium</i>	Unit Disk Graph Medium (UDGM)
<i>Ranges</i>	TX: 50m, Interference: 60m
<i>PHY and MAC</i>	IEEE 802.15.4 with CSMA
<i>Duty Cycling</i>	ContikiMAC (CC1 125ms) & NullRDC
<i>Iterations</i>	10 for each parameter permutation
<i>RNG Seeds</i>	New seed each iteration
<i>Duration</i>	5 minutes of actual time each iteration
<i>Traffic Pattern</i>	CBR (exact rate discussed in text)
<i>Message Size</i>	4 app. layer bytes (variable size on link)
<i>TM Params</i>	$I_{min} \in \{125, 250, 375, 500, 625, 750\}$ $k = 3$, I_{max} variable (see text)
<i>SMRF Params</i>	$F_{min} = 0ms, Spread = 1$ $F_{min} = 31.25ms, Spread \in \{2, 4, 8\}$

The trade-off in order to achieve the aforementioned improvements, is a decrease in packet delivery ratio (increased packet loss), compared to TM which is by design more reliable. Packet delivery ratio is scrutinised for different traffic rates under multiple network topologies in section 5.2.

5 Simulation Results

In order to evaluate the algorithms, we performed a series of experiments in Contiki's Cooja simulator, with common parameters outlined in Table 2. Our discussion in the following paragraphs uses the term *network density*. In this context, network density ND is defined in the same way as the density of an undirected graph with edge set E and set of vertices V (eq. 1). ND can take values between 0 and 1 inclusive ($ND = 0$ for an *edgeless graph* and $ND = 1$ for a *complete graph*).

$$ND = \frac{2|E|}{|V|(|V| - 1)} \quad (1)$$

This is a link layer metric: an edge between nodes A and B exists if and only if the two nodes are single-hop neighbours (can directly hear each other). This definition only makes sense if radio links are symmetric, which is true for Cooja's UDGM environment (Table 2) but not always the case for real deployments. In case of non-symmetric links (e.g. when A can hear B, but B cannot hear A), the link layer topology would have to be modelled as a directed graph. Investigating the behaviour and performance of multicast algorithms in an environment with non-symmetric links is part of our future plans.

5.1 Simulation Topologies and Configuration

We ran our experiments in four different topologies, as illustrated in Fig. 3: i) A topology where all devices are placed in a line 40m apart, ii) Three different tree topologies

Fig. 3 Simulated topologies differentiated by network density. Very sparse line topology with $ND \approx 0.09$. Three tree topologies: *solid line* with $ND \approx 0.14$, *dashed* ($ND \approx 0.36$) and *dotted* ($ND \approx 0.71$). The *solid black* node acts as RPL root and multicast traffic source.

differentiated by network density. In all scenarios, all network nodes are configured to join the multicast group in order to provide us more detailed measurements.

In the case of the line topology, the chosen maximum transmission range of 50m (Table 2) meant that each node could directly exchange radio messages with a maximum of two nodes. This topology helps us examine the algorithms in an extremely sparse topology ($ND \approx 0.09$ for 21 nodes). Another useful feature is that the shape of the RPL DODAG is predictable: it always matches the physical topology. Previous works evaluating network protocols on a per-hop basis have been limiting their investigation to seven or eight hops [35]. Thus, the selected maximum distance of 20 hops is considered adequate.

For each of those topologies, we experimented with two radio duty cycling algorithms: NullRDC and ContikiMAC. Choosing these two mechanisms brings out potential performance differences between a duty cycling network (ContikiMAC) and a network where nodes keep their radios always on (NullRDC). ContikiMAC typically operates with a channel sampling rate of 16Hz [10] or 8Hz [9], resulting in wake-up intervals of approximately 62.5ms and 125ms respectively. In our simulations we used a sampling rate of 8Hz, which is the default value used by Contiki's port for Sky motes.

For TM, we used six different configurations of I_{min} , also listed in Table 2. For each value of I_{min} , we set I_{max} so that the longest possible trickle interval ($I_{min} * 2^{I_{max}}$) would not overflow the boundaries of the variable holding its value. For SMRF, we simulated four different (F_{min} , $Spread$) pairs (Table 2). We ran ten iterations (each one with a different random seed) per topology, per configuration, per traffic rate. For each permutation we evaluated three metrics: i) packet delivery ratio, ii) end to end delay and iii) energy consumption. For TM only, we also investigated the ratio of out of order datagram deliveries.

From an application layer perspective, our multicast traffic was Constant Bit Rate (CBR) with a payload of 4 bytes. For each of the configurations above, we experimented with four multicast flows differentiated by the interval between two successive message transmissions (250, 500, 750ms and 1sec). As a result of 6LoWPAN header compression, the number of actual bytes on link would vary per hop: a message leaving the source has a shorter on-link length than when copied beyond the first hop. Furthermore, TM adds 8 bytes to each datagram in the shape of a HBHO, which also prohibits UDP header compression. As a result, layer two frames varied in size between 35 and 61 bytes. For these reasons, we use the inter-packet interval to refer to the flows, instead of bytes/sec.

Operating the network under very heavy traffic load is intentional; it brings out algorithm advantages and drawbacks, allowing us to draw conclusions on their performance. Additionally, as 6LoWPANs progress towards general-purpose multi-service deployments, protocols are expected to cope with traffic originating from internet hosts and it is not uncommon to encounter applications with higher throughput requirements. For instance, to adapt to these changing requirements, traditional duty cycling algorithms which relied on the assumption of low data rates are now evolving in order to perform efficiently under high and bursty traffic patterns [13]. It is thus becoming increasingly common to conduct WSN protocol evaluation under conditions of heavy network load [13, 30].

5.2 Packet Delivery Ratio and Loss

For each multicast group, we calculate Packet Delivery Ratio (*PDR*) with eq. 2, where N is the number of unique multicast datagrams sent by traffic sources to the group's IPv6 address, M is the number of multicast group members and R_i is the number of unique multicast datagrams sent to this group and received correctly by node i . This equation is valid under the assumption that multicast group membership remains unchanged throughout the multicast flow's entire lifecycle, which holds true in the experiments presented here.

$$PDR = \frac{\sum_{i=1}^M R_i}{N \times M} \quad (2)$$

PDR is an application layer metric and can take values in $[0, 1]$. Packet loss is calculated as $1 - PDR$, thus packet loss 0% is the equivalent of 100% *PDR*, which means that all multicast datagrams were received correctly by all multicast group members.

5.2.1 Line Topology

Investigating the behaviour of TM for different I_{min} values, we observe that packet delivery ratio can vary between perfect (0% loss) and extremely poor. The four graphs in Fig. 4 illustrate the results for four different datagram transmission rates in the line topology over NullRDC. When configured with $I_{min} = 125ms$, TM achieves 100%

Fig. 4 Packet delivery ratio for TM over NullRDC in the line topology. Each graph illustrates results for a different Multicast Traffic rate. Lines correspond to an individual or a group of I_{min} values, as indicated in the legend. Notice the different Y-axis scale for each subfigure.

delivery across all 20 hops regardless of traffic rate. Delivery ratio drops rapidly with higher I_{min} values and with higher transmission rates. Losses occur when node caches are full and a new packet arrives, overwriting an older one before the latter gets copied further down the line. Since there is no path redundancy in this topology, a packet lost in this manner cannot be recovered. In all scenarios, a very heavy packet loss increase is observed when I_{min} becomes higher than the multicast flow's inter-datagram interval. The reason is that, even with constantly resetting trickle timers, ICMPv6 control packet exchange is not nearly quick enough to keep up with the multicast datagram arrival rate. Thus, the loss phenomenon described above is guaranteed to occur. Heavy packet losses mostly happen at the first hop, which acts as a *filter* for the remainder of the line: The first node throttles traffic to a rate which is more tenable for the rest of the network, which is why the decremental trend is a lot smoother beyond the first hop.

We cherry-picked TM's I_{min} value of 125ms, (which is the best choice for NullRDC), compared its packet delivery ratio with SMRF and plotted the results in Fig. 5. We include results with three F_{min} , $Spread$ configurations, each corresponding to a different line in the plot: *Diamonds*: (0, 1) respectively, *Squares*: (31.25, 4) and *Triangles*: (31.25, 8). Results for SMRF were comparable for different multicast traffic rates, we have thus combined all four rates in this single figure. Since SMRF only forwards each datagram once, it was anticipated to demonstrate a higher packet loss rate. Results confirm this and also indicate that losses increase as the forwarding delay increases.

We performed the same measurements over ContikiMAC and results are significantly different, as illustrated in Fig. 6. The first observation is that with ContikiMAC,

Fig. 5 Packet Delivery Ratio over NullRDC (Line Topology).

Fig. 6 Packet Delivery Ratio with ContikiMAC for both algorithms and various configurations (Line Topology).

packet loss rates are higher across the board. This is due to the fact that multicast datagrams are transmitted as broadcast frames at the link layer, which is rather inefficient as discussed in sec. 4.1. Another observation is that over ContikiMAC, lowering TM's *Imin* has an adverse result: delivery ratio decreases instead of increasing. *Imin* values of 125 and 250ms stand out as having underperformed, compared to the remaining

Fig. 7 Packet Delivery Ratio in the tree topologies for different algorithms, over both RDC layers and various traffic rates. (a): Over ContikiMAC. (b): Over NullRDC.

four values which yielded comparable results and are thus averaged out into a single line. SMRF-related results are similar to those observed over NullRDC: i) Packet loss is slightly higher compared to TM and ii) Increased forwarding delay (by increasing F_{min} and $Spread$) results in higher losses.

5.2.2 Tree Topologies

In tree topologies, TM has a comparative advantage since there is path redundancy: a node may receive a multicast datagram from any of its potentially multiple neighbours. We thus anticipated its packet losses to be considerably lower. Multicast traffic transmission rates once again appeared to have minimal impact on the results and are thus averaged out. Based on our experience from the line topology, we singled out TM's two best I_{min} candidates: 125 and 500ms and compared results with SMRF's two best F_{min} , $Spread$ pairs: (0, 1) and (31.25, 4). Results are illustrated in Fig. 7, with the graph on the left corresponding to ContikiMAC and the one on the right to NullRDC.

When operating over NullRDC, the (0, 1) SMRF configuration which performed very well in the line topology now severely under-performed its (31.25, 4) counterpart. The reason is that it is very susceptible to the hidden terminal problem, which was the original motivation behind introducing $Spread$. The results over ContikiMAC also confirm the finding that configuring TM with a sub-optimal I_{min} value can lead to very poor delivery ratios. Since packet trains indirectly mitigate the hidden terminal problem, the performance of SMRF's (0, 1) configuration is comparable to the delivery ratio exhibited when $Spread > 0$. Depending on network density, SMRF can actually deliver more packets than TM, despite the fact that it is less reliable by design. This occurs when the latter is configured with a sub-optimal I_{min} .

Fig. 8 End-to-End delay over NullRDC (Line Topology).

5.3 End to End Delay

5.3.1 Line Topology

As discussed in sec. 4, after our initial analysis we anticipated TM to exhibit low speed compared to the very straightforward SMRF. Figures 8 and 9 illustrate the results over NullRDC and ContikiMAC respectively.

For TM in the case of NullRDC we observe similar results as those for our analysis of Packet Delivery Ratio. Reducing I_{min} yields significant performance improvements: As I_{min} decreases, after resetting their trickle timers nodes exchange cache content information more frequently and inconsistencies are detected earlier, leading to significantly lower hop-by-hop forwarding delay. A similar observation applies to SMRF: Increasing the value of $Spread$ effectively increases forwarding delay. The effect is augmented over twenty hops, leading to longer end to end delays. However, as anticipated by our discussion in sec. 4, SMRF configurations with low $Spread$ (2 and 4) are considerably faster than TM with $I_{min} = 125ms$.

We approximated the three most representative lines from Fig. 8 (*stars*, *circles* and *white triangles*) with linear regression trend lines (not displayed in the diagram for clarity). All three trend lines demonstrated very good fit (co-efficient of determination was very high: $R^2 > 0.998$). The slopes of the three lines were 0.006, 0.049 and 0.105 respectively. These three values represent an approximation of the estimated per hop forwarding delay. In other words, SMRF with F_{min} , $Spread$ (31.25, 2) is about 2.15 times faster per hop than TM with $I_{min} = 125ms$.

Fig. 9 End-to-End delay over ContikiMAC (Line Topology). Lines which would otherwise overlap are summarized into a single one (*triangles and circles*).

For ContikiMAC, TM's end to end delay behaves in a similar fashion to packet delivery ratio with respect to I_{min} values. Reducing I_{min} has a positive impact until the value of 500ms. Any further reductions cause a radical performance degradation, with $I_{min} = 125ms$ severely under performing. The difference between the two algorithms is even more extreme than in the case of NullRDC, as shown in Fig. 9, with SMRF being over 5 times faster than TM, on a hop-by-hop basis. This is caused by the fact that TM relies heavily on link-local multicast messages (link layer broadcasts) to exchange cache content information between nodes. As discussed in section 4.1, link layer broadcasts are fundamentally inefficient with ContikiMAC.

5.3.2 Tree Topologies

We repeated the measurements in tree topologies, with results illustrated in Fig. 10. Network density is a very important factor for both algorithms. As network density increases, RPL's DODAG tends to become *shallow* and *wide*: Maximum path length decreases, while the number of direct ancestors for a single node and the number of leaf nodes increase. As a result, with SMRF datagrams have to traverse fewer hops to reach their destinations. Similarly, TM has been designed to be density-aware [17]: High density and path redundancy result in fewer message exchanges before a datagram can reach all its intended recipients. In Fig. 11, the Y axis illustrates the number of nodes observed to be reachable after a number of hops, as a percentage of the total number of nodes across all experiments for this network density. Notice how, for both algorithms, the maximum observed hop count decreases as network density increases. The choice of duty cycling did not have significant impact so the figure combines results from both RDC algorithms.

Results regarding TM's I_{min} and end to end delay confirm the previous findings: Low values of I_{min} perform very well over NullRDC, while over ContikiMAC the

Fig. 10 End to end delay for different algorithm configurations, over both RDC layers. (a): Over NullIRDC. (b): Over ContikiMAC.

Fig. 11 Hop count distribution for tree topologies.

optimal choice is a value of 500ms. In all scenarios, irrespective of duty cycling and multicast traffic injection rate, SMRF outperforms TM significantly.

In the case of TM over ContikiMAC in a sparse topology, increasing traffic rate decreases end to end delay. This may seem confusing in the first instance but there is a very reasonable explanation: Increased traffic rate becomes the most frequent trigger for trickle timer resets, thus reducing per hop forwarding delay. In more dense topologies, maximum hop count is lower and mitigates this phenomenon.

5.4 Out-of-Order Arrivals

As discussed in sec. 4, SMRF's operation does not entangle datagram ordering. In all our experiments, all messages delivered by SMRF were in the correct order. Thus, in this section we present evaluation results pertaining to TM only.

Fig. 12 Packets received Out-of-Order as a percentage of total packets received (Line Topology). Over ContikiMAC different I_{min} values have minor impact and are thus combined into a single line.

In Fig. 12 we present out of order datagram arrivals as a percentage of total transmitted datagrams, with each sub-figure corresponding to a different multicast traffic rate. Interestingly, in the case of ContikiMAC, changes in I_{min} do not appear to influence packet ordering significantly and are therefore combined into a single line. Beyond the fourth hop, out of order datagram delivery ratio fluctuates between about 35 and 45% for all traffic rates. Compared to NullRDC, packet order is considerably worse with ContikiMAC due to higher error rate on a hop-by-hop basis. At each trickle timer reset, a node will copy and forward multiple datagrams in a row. However, the algorithm offers no guarantees that datagrams forming this *batch* will be transmitted in the correct order. Furthermore, some of them may get lost and get re-transmitted at a subsequent pass, thus arriving late. Conversely, with NullRDC, hop-by-hop transmission is more reliable and as a result TM delivers fewer datagrams out of order.

For slow traffic rates (top two sub-figures) over NullRDC, low I_{min} values (125 or 250ms) both achieve 0% out of order arrivals (*triangles* in the graph coincide with *squares*). This happens because the algorithm has a chance to copy and forward packets before the queue in each node's cache can build up. With a very good hop-by-hop successful transmission rate, need for re-transmissions is infrequent. As I_{min} increases, cache queues start building up and the batch phenomenon described above occurs, increasing out of order deliveries.

Experimentation in tree topologies (Fig. 13) confirms the findings: The percentage of datagrams getting delivered out-of-order is lower over NullRDC than over ContikiMAC, with low I_{min} values performing better than their high counterparts.

Fig. 13 Out-of-Order datagram delivery in tree topologies.

Over ContikiMAC, the choice of I_{min} does not play a significant part, while the phenomenon is mitigated by increasing network density (as a result of the fact that each datagram has to cross fewer hops before reaching all intended recipients).

5.5 Energy Consumption

Through the facilities provided by Contiki's energy consumption estimation module (energest) [11, 12], we measured the time each node spent in each of the following three states over the duration of each experiment: i) MCU active, ii) RF listening / receiving, iii) RF transmitting. Since we are simulating sky motes, we then converted these time values to estimated energy consumption based on typical datasheet power levels at an operating voltage of 3.0V [41].

NullRDC keeps radio transceivers always on (no duty cycling). As a result, the majority of energy is consumed during idle listening or packet reception, with other components contributing insignificantly. For this reason, we only consider ContikiMAC for the evaluation of the two algorithms in terms of energy consumption.

5.5.1 Line Topology

Fig. 14a illustrates average (per-node) energy consumption for each of the two algorithms under different multicast traffic rates. Observe how consumption attributed to radio reception remains relatively constant across different experiments. The reason is that Radio on/off cycles are controlled by the duty cycling algorithm and are thus unrelated to the behaviour of upper layers in the stack. The main difference between

Fig. 14 Energy Consumptions over ContikiMAC for both algorithms in the line topology. (a): Average per node. (b): Per node average, normalised for the number of received packets.

the two algorithms is caused by radio transmissions, with TM consuming more energy in this state due to its periodic ICMPv6 control datagram exchange and due to the fact that each node may end up forwarding the same cached datagram multiple times (until all its neighbours have received it or until it gets replaced by a newer one in the node's cache). In the case of SMRF, radio reception and radio transmission contribute to total consumption at a ratio of about 1:1. Consumption attributed to micro-controller activity is also higher in the case of TM, providing yet another indication of the algorithm's increased complexity compared to SMRF.

Values displayed in this figure are averages across all nodes in the network. On a hop-by-hop basis, total energy consumption decreases linearly with distance from the traffic source, with nodes close to the source demonstrating an approximate 20% higher consumption than the average, while nodes near the end of the line consume as little as 20% less than the displayed average.

Contrary to our anticipation, network-wide energy consumption decreases as the multicast traffic rate increases, a side-effect of increased packet loss rate. Fig. 14b depicts the same results normalised by the number of received datagrams for each scenario. In this context, the line drawing plots average energy consumption per successfully delivered datagram. This normalisation brings out the anticipated (incremental) trend with increasing multicast traffic rate.

5.5.2 Tree Topologies

Previous conclusions are confirmed under tree topologies, with results illustrated in Fig. 15. Average, per-node energy consumption is lower with SMRF, with radio transmissions being the most significant factor in the case of TM. Consumption due to listening and reception fluctuates slightly more than in the line topology, while network density has a positive effect in total consumption. The effect of network density is more significant in the case of TM, which was anticipated since the algorithm is *density-aware* by design [17].

Fig. 15 Energy Consumption for various network densities (Tree Topologies).

6 Hardware Experiments

In order to evaluate the validity of simulated results, we conducted similar measurements on a hardware testbed formed by 11 Sensinode N740 NanoSensors, each one equipped with an IEEE 802.15.4 RF transceiver. Similar to simulated experiments, one of the nodes acted as multicast traffic source and RPL root, while the remaining nodes would join the same multicast group and act as multicast forwarders and traffic sinks. Our nodes were installed in two different rooms, with enough distance and physical obstacles (full height concrete walls and closed wooden/glass doors) to guarantee a multi-hop topology at link-layer. Thus, the traffic source could only communicate with 5 of the 10 participating nodes.

Under the assumption that radio links are symmetric, the network density of this deployment is $ND = 0.72$ ($|E| = 40, |V| = 10$ in eq. 1, sec. 5). We used devices of identical specification, all transmitting with the same power and configured with the same receiver sensitivity level, thus the assumption is not unrealistic. Deviations can still be caused by spatial interference and by differences between devices introduced during manufacturing.

From a network layer perspective, the topology would alter from time to time and from experiment to experiment due to decisions made by the RPL routing protocol. On most occasions, all nodes were within up to three hops away from the source, with distances of four or five hops observed on very few occasions.

In this deployment, we evaluated packet delivery ratio and out of order packet arrival for the configurations listed in Table 3. The port of the Contiki OS for our hardware platform [32] does not support ContikiMAC yet. As a result all experiments were conducted over NullRDC.

In Fig. 16 we illustrate average datagram delivery ratio for both algorithms under different configurations and traffic rates, with results being a good match of simulated observations. For each experiment permutation, averages were calculated over all nodes and all iterations. Configuring TM with $I_{min} = 500ms$ severely under performs all other scenarios and exhibits heavy packet losses under most traffic rates. As was

Table 3 Hardware Testbed Configuration

<i>Nodes</i>	11 Sensinode N740 NanoSensors
<i>PHY and MAC</i>	IEEE 802.15.4 with CSMA
<i>Duty Cycling</i>	NullRDC
<i>Iterations</i>	10 for each parameter permutation
<i>Message Size</i>	4 app. layer bytes (variable size on link)
<i>Traffic Interval</i>	$\in \{62.5, 125, 250, 500, 1000\} ms$
<i>TM Params</i>	$(Imin, Imax) \in \{(125ms, 11), (500ms, 9)\}$ $k = 3$
<i>SMRF Params</i>	$Fmin = 0ms, Spread = 1$ $Fmin = 31.25ms, Spread = 4$

Fig. 16 Datagram delivery ratio during hardware experiments.

the case during simulations, TM with $Imin = 125ms$ exhibits higher delivery ratios than SMRF in most cases.

The results are inverted when the traffic rate is very high (1 message per 62.5ms). We observe that TM's packet loss increases abruptly when traffic inter-packet interval is lower than the value of $Imin$. The reason for this behaviour is that TM's trickle timers are not refreshing frequently enough. As a result, node caches can overflow before all messages have been forwarded, resulting in dropped datagrams.

As discussed in section 3.2, TM can re-arrange packet order resulting in out of order datagram deliveries. In Fig. 17, we illustrate the percentage of datagrams delivered out of order by TM for each experimental configuration. SMRF delivered all packets in order during all hardware experiments and is thus not displayed in the figure. The most notable observation is that incorrect packet order frequency peaks when the traffic rate is equal to $Imin$, with observed frequencies on either side of that value being significantly lower.

Fig. 17 Out-of-Order datagram delivery for TM during hardware experiments. With SMRF, all delivered datagrams are in the correct order.

Fig. 18 Code sizes for both algorithms and different hardware platforms in a typical configuration with support for 1 multicast group and TM cache size of 6 messages.

7 Code Sizes and Memory Requirements

For both algorithms, we evaluated code size when compiled for three different hardware platforms: i) Sensinode N740 NanoSensors (cc2430 System-on-Chip [4], 8bit 8051-based MCU, 8 MB RAM) with the Small Device C Compiler (SDCC), ii) Tmote Sky (16bit MSP430 MCU, 10 MB RAM) with the msp430-gcc toolchain, iii) 8bit AVR MCU (avr-gcc).

Fig. 18 illustrates Contiki image sizes when configured to use each of the two algorithms and when built without multicast support. In both our implementations, changes to the TCP/IP core are minimal and this is reflected in the diagram for all platforms (white bar portions). In order for SMRF to operate, it requires RPL group

management, as specified in the RPL RFC. Implementing this functionality for Contiki results in an increase of the RPL engine code size. This is only necessary for correct SMRF operation and is automatically disabled when building with TM support.

As discussed in section 3.2, TM's design is of high complexity compared to SMRF. Each of the two algorithms is implemented as a separate code module, with respective code sizes depicted with the label *Multicast Support* (bars with black fill in the chart). The size of the module implementing TM is about ten times larger for all platform/compiler combinations. For instance, when building with SDCC, TM's code size is 12028 bytes whereas SMRF's module only occupies 718 bytes of code. With msp430-gcc, the respective sizes are 3676 and 326 bytes in the .data segment.

We also investigate algorithm memory requirements and scalability with their operational parameters. For TM, we evaluate scalability with message cache size and for SMRF we research into the impact of an increase in the number of supported multicast groups. For SMRF, each multicast group requires additional storage space in the node's routing table. The memory requirement increase incurred by each additional group is reflected in Fig. 19 for two different hardware platforms (lines with diamonds and squares). The increase equals 24 bytes per group for both platforms. The vertical offset between the two lines is caused by two factors: i) platform differences in terms of low level drivers and configuration and ii) different hardware architectures and compiler optimization capabilities.

The same figure illustrates TM's memory requirements as message cache size increases. We observe that the increase is very abrupt, with each additional message requiring an additional 256 bytes of RAM. This amount is necessary in order to cache the entire message headers and payload, alongside local control information required by the algorithm. Without altering Contiki's configuration for the Sensinode platform, images with TM support will only fit available memory when cache size is lower than 7 messages. For the Sky platform, memory restrictions are satisfied when the cache is configured to hold fewer than 9 messages.

8 Conclusions

In this work we have disclosed design and implementation details of the Stateless Multicast RPL Forwarding (SMRF) algorithm for IPv6-based wireless sensor networks. We have also presented the outcomes of an in-depth comparison between SMRF and the Multicast Forwarding Using Trickle (TM) algorithm. We presented evaluation results obtained by simulations as well as from experiments conducted on a multi-hop hardware testbed. The latter result set validated simulated findings.

We have demonstrated that TM's performance and energy consumption are very sensitive to changes in the value of its configuration parameter I_{min} , with the optimal depending on the choice of underlying duty cycling algorithm. In the case of deployments without duty cycling, decreasing the value of I_{min} results in better performance. In the more realistic scenario of duty-cycled networks we observe that, in order to minimize delay and packet loss I_{min} should be configured to a value higher than the CCI. Our experiments were conducted with a CCI value of 125ms, in which case $I_{min} = 500ms$ is the optimal configuration, performing slightly better than when

Fig. 19 Memory requirements and scalability.

Table 4 Algorithm Selection Criteria and Optimal Configuration

Criterion	TM	SMRF
Reliability	✓	
Low Delay		✓
Code Size	High	Low
Delivery Order		✓
Energy Efficiency		✓
Configuration		
NullRDC (no duty cycling)	Low I_{min} e.g. $I_{min} = 125ms$	$F_{min} = 0ms$ (lower delay) $F_{min} > 0ms$ (lower loss)
ContikiMAC (duty cycling)	$3 * CCI \leq I_{min} \leq 4 * CCI$ e.g. $(CCI, I_{min}) = (125ms, 500ms)$	$F_{min} > 0ms$

$I_{min} = 375ms$. Lower values have a negative impact on packet loss, while higher values increase the ratio of packets delivered out of order. From our experiments we also observed that the values of I_{max} and k do not have a significant impact on performance.

On the other hand, SMRF is less susceptible to variances of this nature, has lower end-to-end delay and is more energy efficient in exchange for an occasional slight drop in reliability. SMRF is also a lot less complex and has lower memory requirements, making it suitable for severely constrained devices.

Ultimately, the choice of multicast forwarding algorithm should be based on the anticipated usage of a sensor deployment. In Table 4 we highlight a summary of selection criteria and recommended configuration values for both algorithms.

The sources for both implementations are available on github, as a fork of Contiki's source tree⁵ and are distributed under the terms of the 3-clause BSD license.

⁵ <https://github.com/g-oikonomou/contiki-sensinode/tree/mcast-forward>

References

1. Bhattacharyya (editor), S.: An Overview of Source-Specific Multicast (SSM). RFC 3569 (2003)
2. Butt, T.A., Phillips, I., Guan, L., Oikonomou, G.: TRENDY: An Adaptive and Context-Aware Service Discovery Protocol for 6LoWPANs. In: Proc. Third International Workshop on the Web of Things (WoT 2012). Newcastle, UK (2012)
3. Carzaniga, A., Khazaei, K., Kuhn, F.: Oblivious low-congestion multicast routing in wireless networks. In: Proc. Thirteenth International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2012) (2012)
4. A True System on Chip solution for 2.4 GHz IEEE 802.15.4 / ZigBee®. CC2430 Data Sheet (rev. 2.1) (2007)
5. Cheshire, S., Krochmal, M.: DNS-based service discovery. RFC 6763 (2013)
6. Cheshire, S., Krochmal, M.: Multicast DNS. RFC 6762 (2013)
7. Clausen, T., Herberg, U.: Comparative Study of RPL-Enabled Optimized Broadcast in Wireless Sensor Networks. In: Proc. Sixth International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP 2010). Brisbane, Australia (2010)
8. Das, S.M., Pucha, H., Hu, Y.C.: Distributed hashing for scalable multicast in wireless ad hoc networks. IEEE Transactions on Parallel and Distributed Systems **19**(3), 347–362 (2008)
9. Dunkels, A.: The ContikiMAC Radio Duty Cycling Protocol. Tech. Rep. T2011:13, Swedish Institute of Computer Science (2011)
10. Dunkels, A., Mottola, L., Tsiftes, N., Österlind, F., Eriksson, J., Finne, N.: The Announcement Layer: Beacon Coordination for the SensorNet Stack. In: Proc. European Conference on Wireless Sensor Networks (EWSN) (2011)
11. Dunkels, A., Österlind, F., Tsiftes, N., He, Z.: Demo abstract: Software-based sensor node energy estimation. In: Proc. Fifth ACM Conference on Networked Embedded Sensor Systems (SenSys 2007). Sydney, Australia (2007)
12. Dunkels, A., Österlind, F., Tsiftes, N., He, Z.: Software-based on-line energy estimation for sensor nodes. In: Proc. Fourth Workshop on Embedded Networked Sensors (Emnets IV). Cork, Ireland (2007)
13. Duquennoy, S., Österlind, F., Dunkels, A.: Lossy links, low power, high throughput. In: Proc. 9th ACM Conference on Embedded Networked Sensor Systems (SenSys 2011) (2011)
14. Feng, C.H., Zhang, Y., Demirkol, I., Heinzelman, W.B.: Stateless multicast protocol for ad hoc networks. IEEE Transactions on Mobile Computing **11**(2), 240–253 (2012)
15. Flury, R., Wattenhofer, R.: MLS: an efficient location service for mobile ad hoc networks. In: Proc. 7th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC), pp. 226–237 (2006)
16. Hui, J., Culler, D.: The dynamic behavior of a data dissemination protocol for network programming at scale. In: Proc. 2nd international conference on Embedded networked sensor systems (SenSys), pp. 81–94 (2004)
17. Hui, J., Kelsey, R.: Multicast forwarding using trickle. Internet Draft (version 01) (2012). (draft-ietf-roll-trickle-mcast-01)
18. Hui, J., Vasseur, J.P.: The Routing Protocol for Low-Power and Lossy Networks (RPL) Option for Carrying RPL Information in Data-Plane Datagrams. RFC 6553 (2012)
19. Hui, J.W., Culler, D.E.: IP is dead, long live IP for wireless sensor networks. In: Proc. 6th ACM conference on Embedded network sensor systems (SenSys '08), pp. 15–28 (2008)
20. Hui (editor), J., Thubert, P.: Compression format for IPv6 Datagrams over IEEE 802.15.4-Based Networks. RFC 6282 (2011)
21. Klauck, R., Kirsche, M.: Bonjour contiki: A case study of a DNS-based discovery service for the internet of things. In: Proc. 11th International IEEE Conference on Ad-Hoc Networks and Wireless (ADHOC-NOW 2012), *Lecture Notes in Computer Science (LNCS)*, vol. 7363, pp. 317–330. Springer (2012)
22. Klauck, R., Kirsche, M.: Enhanced DNS Message Compression - Optimizing mDNS/DNS-SD for the Use in 6LoWPANs. In: Proc. 9th International Workshop on Sensor Networks and Systems for Pervasive Computing (PerSeNS 2013) (2013)
23. Koutsounikolas, D., Das, S.M., Hu, Y.C., Stojmenovic, I.: Hierarchical geographic multicast routing for wireless sensor networks. Wireless Networks **16**(2), 449–466 (2010)
24. Levis, P., Brewer, E., Culler, D., Gay, D., Madden, S., Patel, N., Polastre, J., Shenker, S., Szewczyk, R., Woo, A.: The emergence of a networking primitive in wireless sensor networks. Commun. ACM **51**(7), 99–106 (2008)

25. Levis, P., Clausen, T.H., Hui, J., Gnawali, O., Ko, J.: The trickle algorithm. RFC 6206 (2011)
26. Levis, P., Patel, N., Culler, D., Shenker, S.: Trickle: A self-regulating algorithm for code propagation and maintenance in wireless sensor networks. In: Proc. First USENIX/ACM Symposium on Networked Systems Design and Implementation (NSDI), pp. 15–28 (2004)
27. Li, J., Jannotti, J., De Couto, D.S.J., Karger, D.R., Morris, R.: A scalable location service for geographic ad hoc routing. In: Proc. 6th annual international conference on Mobile computing and networking, MobiCom '00, pp. 120–130 (2000)
28. Lin, K., Levis, P.: Data discovery and dissemination with DIP. In: Proc. 7th international conference on Information processing in sensor networks, IPSN '08, pp. 433–444. Washington, DC, USA (2008)
29. Lynn, K., Sturek, D.: Extended multicast dns. Internet Draft (version 01) (2012). (draft-lynn-homenet-site-mdns-01)
30. Michopoulos, V., Guan, L., Oikonomou, G., Phillips, I.: DCCC6: Duty Cycle-Aware Congestion Control for 6LoWPAN Networks. In: Proc. 2012 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops). Lugano, Switzerland (2012)
31. Montenegro, G., Kushalnagar, N., Hui, J.W., Culler, D.E.: Transmission of IPv6 packets over IEEE 802.15.4 networks. RFC 4944 (2007)
32. Oikonomou, G., Phillips, I.: Experiences from Porting the Contiki Operating System to a Popular Hardware Platform. In: Proc. 2011 International Conference on Distributed Computing in Sensor Systems and Workshops (DCOSS). Barcelona, Spain (2011)
33. Oikonomou, G., Phillips, I.: Stateless Multicast Forwarding with RPL in 6LoWPAN Sensor Networks. In: Proc. 2012 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops). Lugano, Switzerland (2012)
34. Okura, A., Ihara, T., Miura, A.: Bam: branch aggregation multicast for wireless sensor networks. In: Mobile Adhoc and Sensor Systems Conference, 2005. IEEE International Conference on, pp. 354–363 (2005)
35. Österlind, F., Dunkels, A.: Approaching the maximum 802.15.4 multi-hop throughput. In: Proc. Fifth ACM Workshop on Embedded Networked Sensors (HotEmNets 2008) (2008)
36. Royer, E.M., Perkins, C.E.: Multicast operation of the ad-hoc on-demand distance vector routing protocol. In: Proc. of the 5th annual ACM/IEEE international conference on Mobile computing and networking (MobiCom '99), pp. 207–218 (1999)
37. Sá Silva, J., Camilo, T., Pinto, P., Ruivo, R., Rodrigues, A., Gaudêncio, F., Boavida, F.: Multicast and ip multicast support in wireless sensor networks. *Journal of Networks* **3**(3), 19–26 (2008)
38. Sanchez, J.A., Marin-Perez, R., Ruiz, P.M.: Beacon-less geographic multicast routing in a real-world wireless sensor network testbed. *Wireless Networks* **18**(5), 565–578 (2012)
39. Sanchez, J.A., Ruiz, P.M., Liu, J., Stojmenovic, I.: Bandwidth-efficient geographic multicast routing protocol for wireless sensor networks. *IEEE Sensors Journal* **7**(5), 627–636 (2007)
40. Shelby, Z., Hartke, K., Bormann, C.: Constrained application protocol (CoAP). Internet Draft (2013). (draft-ietf-core-coap-15)
41. Tmote sky: Ultra low power ieee 802.15.4 compliant wireless sensor module. Moteiv Corporation (2006). Tmote Sky Datasheet Revision 1.0.2
42. Song, S., Choi, B.Y., Kim, D.: MR. BIN: Multicast Routing with Branch Information Nodes for Wireless Sensor Networks. In: Proc. 19th International Conference on Computer Communications and Networks (ICCCN 2010), pp. 1–6 (2010)
43. Song, S., Kim, D., Choi, B.Y.: Agsmr: Adaptive geo-source multicast routing for wireless sensor networks. In: Proc. Wireless Algorithms, Systems, and Applications, *Lecture Notes in Computer Science*, vol. 5682, pp. 200–209. Springer Berlin / Heidelberg (2009)
44. Winter (editor), T., Thubert (editor), P., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, J.P., Alexander, R.: RPL: IPv6 Routing Protocol for Low power and Lossy Networks. RFC 6550 (2012)