BMFA: Bi-Directional Multicast Forwarding Algorithm for RPL-based 6LoWPANs

Georgios Z. Papadopoulos^{1,2}, Andreas Georgallides¹, Theo Tryfonas¹, and George Oikonomou¹

¹ Faculty of Engineering, University of Bristol, UK,
[a.georgallides,theo.tryfonas,g.oikonomou]@bristol.ac.uk

² IRISA, Télécom Bretagne, Institut Mines-Télécom, France,
georgios.papadopoulos@telecom-bretagne.eu

Abstract. In scenarios involving point-to-multipoint network traffic, transmitting to each destination individually with unicast may lead to poor utilisation of network bandwidth, excessive energy consumption caused by the high number of packets and suffers from low scalability as the number of destinations increases. An alternative approach, would be to use network-layer multicast, where packets are transmitted to multiple destinations simultaneously. In doing so, applications adopting a one-to-many communication paradigm may improve their energy efficiency and bandwidth utilisation. In this paper, we present Bi-directional Multicast Forwarding Algorithm (BMFA), a novel RPL-based multicast forwarding mechanism. BMFA improves its pre-predecessor SMRF in that it allows multicast traffic to travel both upwards as well as downwards in an RPL tree. At the same time, it retains SMRF's low latency and very low energy consumption characteristics. Our performance evaluation results, conducted using the Contiki operating system, show that BMFA outperforms its rival Trickle Multicast / Multicast Protocol for Low power and Lossy Networks (TM / MPL), in terms of reducing both delay and energy consumption.

Key words: Internet of Things, 6LoWPAN, Wireless Sensor Networks, IPv6 Multicast, Trickle.

1 Introduction

In environmental monitoring scenarios, it is expected that networks will be formed by a potentially very high number of sensor nodes and therefore scalability is an essential requirement. In cases when sensor devices are powered by batteries, it is impractical or outright untenable to replace batteries very frequently due to high management cost and possibly hard-to-reach installation locations. Thus, long battery life is important. Low energy consumption may also be considered important in deployments with mains-powered devices, in order to reduce financial cost, but also in order to comply with national and international regulations where applicable.

In scenarios involving point-to-multipoint network traffic, transmitting to each destination individually using unicast at the network layer may lead to

poor utilisation of network bandwidth, excessive energy consumption caused by the high number of packets and suffers from low scalability as the number of destinations increases. An alternative approach is to use network-layer multicast, where packets are transmitted to a set of destinations simultaneously. This can lead to energy-efficiency improvements for applications that require one-to-many communications. Examples of such applications include service discovery and network management.

In this paper, we present the design and implementation of a new multicast forwarding algorithm for 6LoWPANs, namely, Bi-Directional Multicast Forwarding Algorithm (BMFA). We demonstrate that BMFA can achieve very low energy consumption and therefore meet the requirements of aforementioned use-cases, ultimately increasing the lifetime of a smart object deployment.

2 Background and Related Work

A number of works in the area of multicast for Internet of Things (IoT) and Wireless Sensor Networks (WSNs) have been proposed. A large number of multicast forwarding solutions encountered in current literature are based on geographic routing, such as those discussed in [1], [2]. However, most of those approaches have certain characteristics and make assumptions which render them unsuitable for IPv6-based deployments. For instance, many of them assume that, for every multicast message, the sender is aware of the addresses or unique identifiers of all intended destinations. Additionally, some efforts suffer from poor scalability while others rely on unrealistically large network packets. Finally, they are only applicable where the source as well as all destinations are within the boundaries of the same WSN deployment.

Multicast Protocol for Low-Power and Lossy Networks (MPL) is the standard for multicast forwarding and group management in Low-Power and Lossy Networks (LLNs) [3]. MPL is independent of the protocol used for unicast routing. It will thus operate in Routing Protocol for LLNs (RPL)-based [4] 6LoWPANs, as well as in deployments where RPL is not in use. In a nutshell, MPL routers maintain a cache of multicast datagrams they have seen. Neighbours exchange information about the content of their caches by using ICMPv6 messages. If one node detects that one of its neighbours has not received a multicast datagram that the former has in its cache, it will schedule subsequent forwarding of the datagram(s) in question. The exchange of ICMPv6 messages is governed by trickle [5], thus reducing network control traffic when multicast traffic is not present, but reacting quickly to the arrival of new multicast datagrams.

SMRF [6, 7] is a multicast forwarding algorithm based on the topology information provided by RPL. Nodes participating in an RPL network exchange topology information in order to build a Destination Oriented Directed Acyclic Graph (DODAG) and construct their routing tables. When RPL's "Storing with multicast support" Mode of Operation (MOP) is used, a node can join a multicast group by advertising its address in an outgoing Destination Advertisement

Fig. 1: The BMFA algorithm.

Object (DAO) message. Upon reception of such a message from one of its children, the parent node registers the multicast address in its routing table; then the same address is advertised by the parent in its own DAO messages. When a multicast packet destined to that address arrives, it will be forwarded downwards the tree until it reaches the recipient nodes.

3 BMFA Operation

SMRF is very lightweight but this comes at the cost of a severe limitation: it is only capable of forwarding traffic “downwards” in the RPL DODAG tree [6, 7]. In this paper, we present a multicast forwarding algorithm called BFMA, an extended version of SMRF scheme. BFMA’s primary design goal is to alleviate this limitation (i.e., allow both upwards and downwards traffic) while at the same time maintaining SMRF’s lightweight and energy-efficient nature. In order to support bi-directional traffic and avoid routing loops, BMFA uses the 20 flow-label bits of the IPv6 header. BMFA also uses information provided by RPL’s group membership scheme. Its operation is the following (Fig. 1):

- A node will accept an incoming multicast datagram if and only if the digest value in the Flow Label inside the IPv6 header does not contain the digest value of its own Link Local address and if the datagram’s link layer source address is the link layer address of either the node’s preferred RPL parent or the link layer address of one of its children.

Topology parameters	Value
Topology	$360 \times 200\ m$
Number of nodes	21 fixed sensors (20 sources)
Node spacing	$x = 6m / y = 8m$
Simulation parameters	Value
Duration	68 <i>minutes</i>
Data collection scheme	CBR & VBR
TM	lmin in [125, 500, 700] <i>ms</i>
BMFA	Spread in [2, 4]
Routing model	RPL [4]
Number of hops	Multihop (5 hops maximum)
MAC model	CSMA
Duty cycling	ContikiMAC (CCI 125 <i>ms</i>)
PHY	IEEE 802.15.4
Hardware parameters	Value
Antenna model	Omnidirectional
Radio propagation	2.4 <i>GHz</i>
Transmission range	TX: 50 <i>m</i> , Interference: 60 <i>m</i> (UDGM)

Table 1: Simulation setup.

- If the message gets accepted and if the node is member of the multicast group, then the message will get delivered up to the network stack locally.
- If the message gets accepted the packet’s Flow Label field is updated to the digest value of the packet sender’s link layer source address.
- If the message gets accepted and if there is an entry for the datagram’s multicast destination address in the routing table, meaning that a node under us in the RPL structure is a group member, the packet will get forwarded.

4 Performance evaluation

We have implemented BMFA for the Contiki open-source operating system for the Internet of Things. Contiki also features an implementation of an earlier version of the MPL algorithm, called Trickle Multicast (TM). For this evaluation we therefore compare BMFA with Contiki’s implementation of TM, using the Cooja simulator [8]. Our setup consists of 21 simulated nodes, each one of them operating as a multicast traffic source, a group member, or a simple traffic forwarder. Depending on the experiment, all nodes would support either BMFA or TM. Our experiments aim to investigate performance changes of those two algorithms based on different configurations and under different network traffic patterns. Table 1 summarising important simulation parameters.

4.1 Network Delay

We investigate network delay as a factor of traffic bit rate and network density (defined in [7]). As can be observed from Fig. 2, TM does not perform as it was expected based on the configured parameters. For instance, TM configured with $lmin = 750ms$ leads to the lowest end-to-end delay across the board (for different traffic bit rates per network density). Moreover, end-to-end delay is expected to

Fig. 2: BMFA vs TM end-to-end delay performance.

be lower as I_{min} decreases, while our simulation results present the opposite. Lastly, high delays were anticipated under TM since it uses link layer broadcasts, which are fundamentally inefficient under ContikiMAC, as demonstrated in [7].

On the other hand, when using BMFA in low-density topologies (i.e., up to 0.35), the end-to-end delay declines slightly as the inter-packet delay increases, while when the bit rate decreases (Variable Bit Rate with 1 – 2s inter-packet delay) overall delay reaches its maximum. This is caused by packets spending more time in a node's cache when the bit rate increases; packets do not get transmitted until all packets preceding them are forwarded first. Furthermore, an inter-packet delay of more than 1s leads to the opposite results, since all cached packets get forwarded before the next one arrives. On the other hand, for high densities (i.e., 0.71) the delay continues its descending trend as the inter-packet delay increases. Based on the fact that for high network densities a node is expected to be selected as preferred parent from a greater number of nodes (RPL's DODAG becomes shallow and wide), more packets are expected to wait into its cache until they get forwarded; in this case packets transmitted with VBR arrive neither too soon nor too late to the recipient nodes, resulting to even better results. To summarize the performance of the two algorithms, BMFA outperforms TM by at least five times under most configurations.

Fig. 3: BMFA vs TM average node energy consumptions.

4.2 Energy Consumption

Through the facilities provided by Contiki's energy consumption estimation module, (**Energest** [9]), we measured the time each node spent in different power consumption states (MCU active; RF listening / receiving, RF transmitting). Since we are simulating Texas Instruments MSP430F5438 experimenter board nodes, we converted these time values to estimated energy consumption based on typical datasheet power consumption levels at an operating voltage of 3.0V.

For TM it can be observed (Fig. 3) that under higher network densities, less energy is consumed due to the fact that agreement between all nodes can be achieved with fewer ICMPv6 control message exchanges. This can be also observed by the fact that as density increases, less energy is required for transmitting than for listening. For BMFA, irrespective of network density, as the inter-packet delay between the transmitted packets increases, overall energy consumption decreases since fewer packets are forwarded. In the case of a very dense network (0.71) and for high bit rates, we can see that energy consumption with BMFA approaches the one observed with TM. This happens because nodes are consuming too much energy by keeping the radio on as a result of picking up transmissions from their large number of neighbours, despite the fact that they only forward packets received only from their children or preferred parent. By comparing the two algorithms we can see that BMFA is more energy efficient than TM since it forwards each packet only once and there is no additional ICMPv6 control message exchange. Moreover, we must highlight that the energy consumption attributed to CPU activity indicates TM's higher algorithmic complexity.

5 Conclusion

In this paper, we have presented BMFA, a multicast forwarding mechanism that minimises network-wide energy consumption. We have implemented BMFA for the Contiki OS and have undertaken a thorough performance evaluation using the Cooja simulator. Our results show a typical trade-off between network performance, energy consumption and reliability. More specifically, our results show that BMFA outperforms TM, in terms of reducing the end-to-end delay, design complexity, code size and energy consumption, while on the other hand, TM severely outperforms BMFA in terms of reliability. Finally, note that as documented in [6, 7], the performance of multicast forwarding for all engines discussed in this paper is heavily dependent on the underlying MAC layer. For a more detailed description of possible optimizations see [6, 7].

Acknowledgements

The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 609094.

References

1. A. Carzaniga, K. Khazaei, and F. Kuhn, "Oblivious low-congestion multicast routing in wireless networks," in *Proceedings of the 13th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2012, pp. 155–164.
2. C. H. Feng, I. D. Y. Zhang, and W. B. Heinzelman, "Stateless Multicast Protocol for Ad Hoc Networks," *IEEE Transactions on Mobile Computing*, vol. 11, no. 2, pp. 240–253, 2012.
3. J. Hui and R. Kelsey, "Multicast Protocol for Low-Power and Lossy Networks (MPL)," RFC 7731, Feb. 2016.
4. T. Winter, P. Thubert, A. Brandt, J. Hui, R. Kelsey, P. Levis, K. Pister, R. Struik, J. Vasseur, and A. R., "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks," RFC 6550, 2012.
5. P. Levis, T. Clausen, J. Hui, O. Gnawali, and J. Ko, "The Trickle Algorithm," RFC 6206, March 2011.
6. G. Oikonomou and I. Philips, "Stateless Multicast Forwarding with RPL in 6LoWPAN Sensor Networks," in *Proceedings of the IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops)*, 2012, pp. 272–277.
7. G. Oikonomou, I. Philips, and T. Tryfonas, "IPv6 Multicast Forwarding in RPL-Based Wireless Sensor Networks," *Springer Wireless Personal Communications*, vol. 73, no. 3, pp. 1089–1116, 2013.
8. F. Österlind, A. Dunkels, J. Eriksson, N. Finne, and T. Voigt, "Cross-Level Sensor Network Simulation with COOJA," in *Proceedings of the 31st Annual IEEE International Conference on Local Computer Networks (LCN)*, 2006.
9. A. Dunkels, F. Österlind, N. Tsiftes, and Z. He, "Software-based On-line Energy Estimation for Sensor Nodes," in *Proceedings of the 4th ACN Workshop on Embedded Networked Sensors (EmNets)*, 2007, pp. 28–32.