Poster: Building a Stairway to Centralised WSN Control

Pablo Corbalán, Victor Cionca, Ramona Marfievici, Donna O'Shea, and Dirk Pesch
Nimbus Centre for Embedded Systems Research
Cork Institute of Technology, Cork, Ireland
{firstname.lastname}@cit.ie

Abstract

We present SMOG, a mechanism to build and maintain a highly accurate centralised full network topology model in wireless sensor networks. We evaluate its scalability, accuracy, and reactivity in simulations and a testbed. Our findings show SMOG achieves high accuracy with low overhead.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless Communication

General Terms

Design, Algorithms, Experimentation, Performance

Keywords

Wireless Sensor Network, Centralised, Model, RPL, NDP

1 Introduction

Centralised network management in wireless sensor networks (WSNs) has many advantages in terms of improved network lifetime, reliability, and delay [4]. Centralised approaches have also been applied to routing in WSNs [2] and have been defined by ISA100.11a and WirelessHART to build time-slotted channel hopping (TSCH) schedules and disseminate routes in the network. Recently, the Software-defined Networking (SDN) paradigm has emerged to simplify control and management in other network domains by centralising network control and decoupling control and data planes, encouraging us to consider centralised solutions also in sensor networks. SDN also facilitates multi-tenancy and network function virtualisation, features that are useful in many sensor network applications. However, to the best of our knowledge, some of the referenced protocols only manage to build partial network topology models, therefore limiting the network-wide decisions that can be taken. Furthermore, no deep analysis of the cost of creating complete network models has been provided.

International Conference on Embedded Wireless Systems and Networks (EWSN) 2016
15–17 February, Graz, Austria

ISBN: 978-0-9949886-0-7

Contribution. We propose SMOG, a mechanism to build and maintain a centralised full network topology model using probabilistic data structures. We consider different modes of operation to report neighbourhood information and show preliminary results on SMOG's scalability, accuracy, and reactivity in simulations and a 100-node testbed.

2 SMOG Design

Model and Mechanisms. SMOG builds a full network topology model as a directed graph, based on neighbourhood information collected at the sink through SMOG messages. Nodes discover and maintain their neighbours via lower layer protocols. RPL is used for neighbour discovery and to build upward routes to the sink, and IPv6 NDP to detect neighbour unreachability. Neighbourhood information is stored in the neighbour cache. Each node sends SMOG messages that contain: the number of neighbours, the RPL rank and preferred parent MAC address, and a Bloom filter (BF) [1] that contains the MAC addresses of a node's neighbours. An insert event in the neighbour cache triggers an insert in the BF, while a removal triggers the BF's complete recomputation. Because the probabilistic nature of BFs allows false positives (FPs), we designed a FP discovery mechanism based on RPL rank and neighbour adjacency information, not presented here due to space limitations. Note that the appearance of FPs leads to a probabilistic model.

Modes of Operation (MOPs). An important dimension of SMOG is its reactivity in reporting neighbour events. To this end, we propose three modes of operation: *i) eventful*: an event triggers a message 1 to 5 s after it occurred. If another event occurs before the message is sent, the previous message is suppressed and a new SMOG message is rescheduled; *ii) periodic*: a message is sent during the second half of a static interval (5 min); *iii) stateful*: an event triggers the same behaviour as in *eventful* but with larger intervals (10 to 15 s). In the absence of events, SMOG messages are sent during the second half of a dynamic interval (2 to 20 min). To limit overhead, if no events occur the interval doubles, eventually reaching a maximum. When neighbour events occur, the interval halves, eventually reaching a minimum.

Inconsistencies. Another important dimension of SMOG is its accuracy. We use three types of inconsistencies to assess it: *Missing Node/Edge*: a node/link that exists in the real topology and is not contained in the model; *FP Edge*: a link contained in the model does not exist in the real topology.

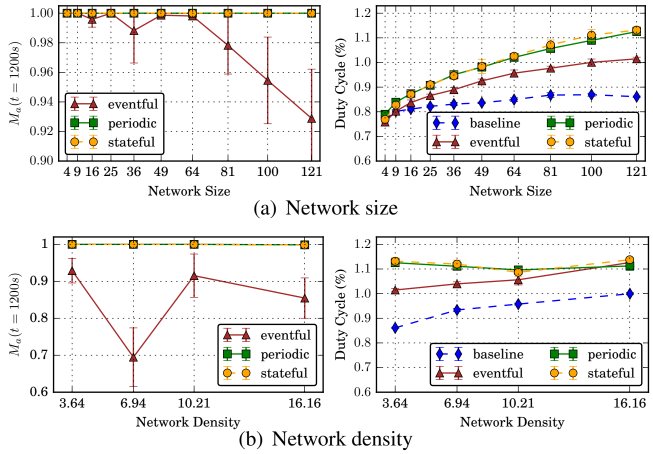

(a) Network size

(b) Network density

Figure 1. Simulations: Model accuracy at $t = 1200$ s (left) and duty cycle (right).

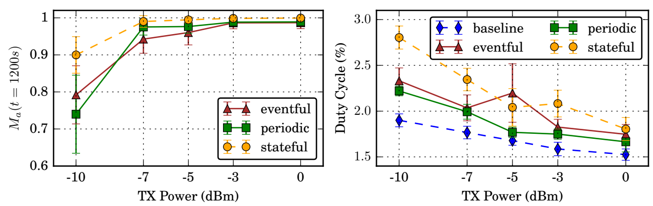

Figure 2. Testbed: Model accuracy at $t = 1200$ s (left) and duty cycle (right).

3 Evaluation

To evaluate SMOG we ran an extensive set of experiments in Cooja, Contiki's network simulator, focusing on scalability, and in Indriya [3], a 100-node testbed, focusing on effectiveness under realistic channel conditions. The key metrics we looked at are: model accuracy ($M_a(t)$), computed as the difference between the real network topology and SMOG's graph adjacency matrices, reactivity, i.e., inconsistency duration, and overhead, quantified by measuring the radio duty cycle and compared to a baseline, i.e., the duty cycle of a null application running RPL and NDP.

Experimental settings. We use Contiki's network stack with 6LoWPAN, RPL, NDP, and ContikiMAC with default settings. SMOG sends its messages on top of UDP. Neighbour reachability was confirmed with link layer ACKs. All experiments were carried out for 20 minutes, with a maximum number of 40 neighbours, 256-bit Bloom filters, and 8 hash functions. Each experiment was repeated five times.

Simulation. We simulated different square grid network sizes, from 4 to 121 nodes, and different network densities, from 3.64 to 16.16 on average, for the 121-node network. We used UDGM radio model and placed the sink, i.e., RPL root, in a corner in all experiments. Fig.1 shows SMOG scales, at least, up to 121 nodes with different network densities with *periodic* and *stateful*, achieving high accuracy with low overhead over the baseline. The accuracy achieved by *eventful* appears to decrease with network size and density, although having the lowest overhead. In the simulations, neighbour events only occurred while the network was being built and no further events occurred after network convergence. The network congestion created by RPL when building the routing topology, decreased *eventful*'s PDR and, consequently, its accuracy. On the other hand, the interval reports sent by

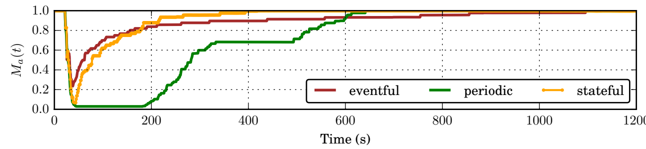

Figure 3. Testbed: Model accuracy in time at -3 dBm.

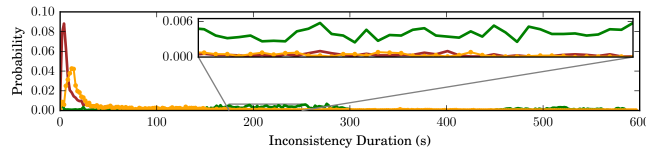

Figure 4. Testbed: Inconsistency duration probability distribution for all MOPs at -3 dBm.

periodic and *stateful*, allowed them to achieve high accuracy with a higher overhead than *eventful*. The impact of FPs on the accuracy was low due to the combination of BF size, number of hash functions, and network densities analysed.

Testbed. We ran experiments in Indriya with five transmission (TX) powers (from -10 to 0 dBm) resulting in networks with different densities (from 6.55 to 11.27 on average). Node 1 acted as sink. Fig.2 shows different trends, in terms of accuracy and overhead, than in simulations, arguably due to different network and channel conditions. In the testbed, dynamic channel conditions produced neighbour events after network convergence, increasing *eventful*'s accuracy gradually as shown in Fig.3. These events also increased *eventful* and *stateful*'s overhead w.r.t. the baseline. Contrary to simulations, *periodic*, unresponsive to neighbour events, had the lowest overhead, increasing its accuracy at lower pace. *Stateful* achieves always the highest accuracy at the highest overhead cost. At -10 dBm, the poor results can be explained by the low connectivity in the testbed. Fig. 4 shows *eventful* is the most reactive, followed closely by *stateful*. *Periodic* presents broader peaks at multiples of 1.5×2.5 min, as per definition. The second lower peak is the result of the inconsistencies that could not be resolved with the first report, but were solved in the second. Overall, the testbed results show a trade-off between reactivity and overhead.

4 Conclusions and Future Work

We presented SMOG, a mechanism to build a highly accurate centralised full network topology model. We tested SMOG in simulations and a testbed, showing its high accuracy with low overhead under different network conditions. Despite the encouraging results, however, future work is required to validate SMOG over a broader range of scenarios, i.e., under churn. These additional experiments are already on our agenda, and will enable us to ascertain to what extent SMOG is affected by other variables.

Acknowledgements. The authors thank Timofei Istomin for helping with testbed experiments. This work was funded by Science Foundation Ireland (SFI) under grant 13/IA/1885.

5 References

[1] B. H. Bloom. Space/Time Trade-offs in Hash Coding with Allowable Errors. *Commun. ACM*, 13(7):422–426, July 1970.

[2] Dawson-Haggerty et al. Hydro: A Hybrid Routing Protocol for Low-Power and Lossy Networks. In *Proc. of SmartGridComm*, 2010.

[3] M. Doddavenkatappa et al. Indriya: A Low-Cost, 3D Wireless Sensor Network Testbed. In *Proc. of TridentCom*. 2011.

[4] M. Zimmerling et al. pTunes: Runtime Parameter Adaptation for Low-power MAC Protocols. In *Proc. of IPSN*, 2012.