Routing in Large Realistic PLC Smart-Grids

Yehuda Ben-Shimol
Communication Systems Engineering Dept.,
Ben-Gurion University, Beer Sheva, ISRAEL
Email:benshimo@bgu.ac.il

Yehiel Zohar
Communication Systems Engineering Dept.,
Ben-Gurion University, Beer Sheva, ISRAEL
Email:yehielz@post.bgu.ac.il

Abstract*—Networking in power line communication (PLC) smart-grids (PLC-SGs) is a major challenge due to the unique characteristics of the physical-, link- and network layers. The RPL (Routing Protocol for Low-Power and Lossy Networks) routing protocol is considered as one of the potential candidates for routing for PLC-SGs, however, the behavior and performance of RPL is not yet scrutinized for its suitability as a proper routing solution for large and realistic PLC-SG networks. In this work we investigate the performance of the RPL protocol in large PLC-SGs, configured from hundreds to thousands of network elements. Field measurements are used to set the channel and neighborhood characteristics. We consider electrical meter reading and adaptive billing as potential applications using the routing layer. We incorporated several improvements for RPL's features and configuration parameters, in order to address the constraints imposed by the requirements of practical PLC-SGs. Our results are based on extensive simulations, and show that in spite of these improvements RPL may not be suitable routing solution for large PLC-SGs, operating under realistic conditions.

I. INTRODUCTION

Operators of electrical power lines are trying to remotely read and monitor electric power meters at any given time to encourage a more efficient electricity usage, using adaptive billing as a function of time and/or type of consumers. Several meter reads per minute is a common requirement. More applications and services are being planned for the smart-grid, such as electrical power control and remote monitoring. The smart grid can also serve as the underlying network for the "smart home" for devices that are already connected to the power-line. Therefore, the demand for higher data transmission rates and shorter response time is ever increasing.

Using power lines to carry information imposes several constraints that need to be addressed: 1) A unique noise profile which may lead to scenarios where only half of the spectral bandwidth is available for communication for less than half of the time. This severe noise conditions, which need to be resolved by a dedicated MAC protocol, often causes high loss of information or even total unavailability of the communication channel due to the direct impact of the noise profile on the performance of the MAC protocol. As a result, the routing mechanisms may not function properly, or efficiently, especially when the noise turns links to become unidirectional rather than bidirectional. 2) A need to transmit with rates ranging from several $Kbps$ up to $1Mbps$. Low transmission rates are used to overcome severe noise conditions in order to reduce the bit error-rate (BER). This also complicates the decision of choosing the right transmission rate as part of determining the routing path to be used by the routing protocol. 3) A PLC-SG network is expected to include up to several thousands of network elements (NEs), therefore, the topology may change frequently, which, in turn, may result with many failures and re-establishments of communication channels and/or their qualities. The dynamics of such a network is usually very high and challenging. 4) The specific nature of the underlying electrical network allows transmissions that pass through the neighborhood transformer to propagate through the medium-voltage lines to other areas of the network, turning any "innocent" broadcast transmission into a "broadcast storm". Although broadcast is usually related to a data link layer activity, many routing protocols rely on it as part of their algorithmic solutions. The wide regional effect of many broadcast transmissions reduces the performance of the routing protocol, and therefore must be considered by solutions at the network and even higher layers of the communication stack. These conditions impose a complex set of constraints to be handled by the routing protocol.

Routing protocols may be classified as proactive or reactive (there are also hybrid solutions). Proactive protocols calculate and maintain forwarding rules (tables) in advance. This implies that the routing protocol identifies and reacts to changes in parameters that influence the forwarding information. On the other hand, reactive routing protocols are triggered by the requirement to deliver actual data to calculate the routing path.

The goal of the paper is to investigate and explore the behavior of RPL, a proactive routing protocol, under conditions that resembles large and realistic

---

This work was partially supported by the ISG consortium. http://isg.co.il.

PLC-SGs. By “large” we mean networks constructed from at least hundreds of NEs as is common in scenarios of electrical meter reading in European cities. For the “realistic” part, a set of field experiments over a PLC network of 1024 NEs was used to obtain real physical parameters of the links. The measurements allow us to construct networks with links having realistic physical parameters that affect the communication through them. We implemented RPL following the definitions in [2], using the OMNET++ network simulator together with its INET framework [3], and used extensive simulations to evaluate RPL’s performance under the unique challenges imposed by the underlying PLC network.

The rest of the paper is constructed as follows: section II describes the RPL protocol. Section III states the research question at hand and describes the system model that is used to investigate the RPL protocol. Section IV presents simulation results for various parameters. Summary and conclusions are given in section V.

II. ROUTING IN PLC NETWORKS

Routing in PLC networks is an evolving area that covers various solution approaches, resembling the development of routing protocols for lossy and low power wireless networks (LLNs). Two of the most considered routing protocols for PLC networks are RPL [1], [2] and LOADng [4], for which there is already an on-going research on their performance in LLNs (see [1], [5] and references within). RPL is considered as “more general” since it supports multiple traffic profiles (P2P, P2MP, MP2P) as opposed to the LOADng that supports only P2P traffic [5]. To the best of our knowledge, the published results regarding the performance evaluation of the two routing protocols in PLC networks are limited in the sense that the unique characteristics of the PLC channels are not considered, and that the underlying MAC protocol practically neglects them. The realistic characteristics of PLC channels greatly impact on the behavior and performance of both MAC and routing protocols. Moreover, the objective function (OF) embedded in standard routing protocols usually tries only to minimize the hop count. The channel characteristics also break the assumption that the primary cause of channel disturbances is white noise, and that the received signal power is mostly determined by the distance between the receiver and transmitter. Existing simulation-based studies were performed for relatively small networks that are suitable for the limited environment of “home networking.” Small network configurations that are composed of only a few tens of NEs are not sufficient to expose many of the limitations of routing protocols in “real life scenarios”. An example of “real life" phenomenon is the propagation of communication signals through the medium-voltage wires interfering with the communication activity in adjacent (and sometimes also remote) neighborhoods. This, in turn, has an impact on topology changes and may load the network with routing management activity.

RPL is a distance-vector protocol that constructs a Destination Oriented Directed Acyclic Graph (DODAG) according to a specified OF, metrics and constraints. RPL is proactive, so each NE constructs a forwarding table to other NEs in the network, without any necessity of a data message to be sent. RPL uses three main types of messages to construct the DODAG, which are variations of ICMPv6 messages: DODAG Information Object (DIO) that is used to construct and maintain the upward routes; DODAG Information Solicitation (DIS) that triggers neighboring NEs to send their current state information (DIO) and Destination Advertisement Object (DAO) which is used to construct the downward routes by advertising prefixes of reachable NEs in a sub-tree.

The process of constructing a DODAG is initiated by its root (concentrator), which broadcasts the first DIO. Each of the neighboring NEs receiving the DIO message decides whether to join the DODAG, based on its policy. In case of a join decision, the NE computes its rank, with respect to the message originator (the root), and broadcasts it to its neighbors. The DIO message contains, inter alia, the transmitter’s rank, so each NE is able to compute the link’s step of rank (Sp) and obtain its own rank, in an additive manner. This process converges as each NE has at least one preferred parent towards the root, although there still might be changes as a result of new updates that suggest smaller ranks, with respect to the root. At the end of this stage, there is an optimized upward route from each NE to the root. In the second stage, each NE advertises its reachability to NEs that are part of its sub-tree, in order to build downward routes. Routes advertisements are published using a DAO message that contains, inter alia, the address and next-hop of the publishing NE. There are two modes of operation for this stage: *storing* and *non-storing* mode. In general, storing mode defines that every NE holds reachability information of its sub-tree, while in the non-storing mode only the root holds a source-routing table to the entire DODAG. In this study, we only examine the non-storing mode, where the whole traffic passes and source-routed through the root. Thus, the overhead of managing and processing traffic is more massive and challenging while storage constraints are relieved for all NEs but the root.

The *Trickle algorithm* [6] is used for establishing eventual consistency in a network, by sending DIO messages in lower frequency when there is no change in the network topology or in ranking of the NEs. In cases of abrupt topology changes, the algorithm solicits the creation and publication of DIO messages more frequently, until the network returns into consistency in manner of topology changes. Moreover, the Trickle mechanism allows configuring boundaries on the rate of DIO messages creation, and avoids network saturation.

One must consider the characteristics of the PLC channel on the performance of the routing protocol. One of the characteristics of PLC channels is the inability to sense the carrier due to very low signal to interference and noise ratio (SINR) levels. The MAC protocol cannot rely on carrier sense and collisions are much more frequent than the case were carrier sensing is possible. This might result with overload of control messages as a response to critical topology changes. RPL does not support asymmetric links, and this is of main concern in PLC-SGs. The effect of asymmetric links is noticeable when examining how downward routes are being built. The NE's decision to reach the root (in the upward direction) using a specific parent is based only on the quality of the downlink. Thus, a parent selection is in many cases erroneous. RPL should quickly construct a connected network (even if not optimal). It should also converge fast and as close as possible to an optimal DAG.

A good routing protocol should not be "too hysteric" due to small local changes (e.g., edge/node failure, edge/node join, change of edge weight), or due too small changes in link qualities, especially if the changes are temporal. This means that a NE should be busy with routing maintenance activity only a small fraction of the time in order to be available for data delivery. This is the main question of this work - to investigate the validity of RPL for touting in PLC-SGs.

III. SYSTEM MODEL

We represent the PLC-SG as a directed network $G = (V, E, W)$ with $V = \{v_1, \ldots, v_N\}$ representing the set of N nodes (NEs), $E = \{(i,j)|\ i,j = 1, \ldots, N, i \neq j\}$ is the set of directed links, each is assigned a weight (cost) $w_{i,j} \geq 0$. A directed path from v_i to v_j is denoted as $P_{v_i,v_j} = \{(v_i, n_1), (n_1, n_2), \ldots, (n_{k-1}, v_j)\}$ where $n_1, \ldots, n_{k-1} \in V$ are intermediate nodes (or relays); the weight of path P_{v_i,v_j} is $W(P_{v_i,v_j}) = w_{v_i,n_1} + w_{n_1,n_2} + \ldots + w_{n_{k-1},v_j}$.

A. Physical and MAC layers

A field experiment using a dedicated PLC-SG was used only for gathering statistical data on the communication activity. The collected measurements provided a large set of average parameters of the communication links between the NEs. This information helped us to implement a dedicated MAC layer that uses a wide range of transmission rates, according to the link properties it has measured during the field experiment. The noise power levels varied between $20dB$ and $60dB$. This range is severe, however it was obtained from field measurements and represents scenarios of practical operational PLCs.

A unicast message to a neighbor is transmitted with a rate that is determined by the SINR and BER over the channel. 11 transmission rates of $2^i \cdot 100bps,\ i \in \{0, 1, \ldots, 10\}$ are used for SINR ranges that are determined by

$$(-30 + 3j)dB, \quad j \in \{0, 1, \ldots, 10\}. \qquad (1)$$

That is, the best transmission conditions are for SINR levels of $0dB$ and above with transmission rate of $102.4Kbps$, and worst transmission conditions are for $-30dB \leq \text{SINR} < -27dB$ with transmission rate of $100bps$. Links with SINR < -30dB are considered as disconnected. Links experiencing SINR levels that are close to the boundaries determined by (1) are affected by small noise levels. To avoid jitter between adjacent transmission we added around these boundaries a Hysteresis filter of $0.5dB$ and we have found it to be an important improvement for the stability of RPL. Our hysteresis filter is implemented at the MAC layer, unlike [7] which presents a hysteresis mechanism for RPL's OF. For a broadcast transmission, a node calculates the median of the known transmission rates to all of its neighbors and uses it for the transmission rate of the broadcast. This guarantees that at least half of the neighbors are able to receive the broadcast message, considering the physical conditions of the channel, while keeping the transmission rate relatively high (by excluding the nodes having the worst SINR levels). One needs to keep in mind that RPL uses broadcast extensively, thus this mechanism performs better than constraining the transmission rate to be of the worst neighbor.

B. Modification of RPL

The OF described in the formal documentations of RPL aims at minimizing the hop count between the root (concentrator) and the node (NE). As described earlier, minimal hop count is by no way optimal, due to other factors of channel behavior and traffic load. Careful investigation (together with our industrial colleagues) of the parameters effecting the performance of the network layer and below, indicated that we need to resolve some problematic situations, such as bottlenecks and bad selection of preferred parents to pass through towards the root. The standard OF was modified so it considers multiple factors such as: signal attenuation, expected number of retransmissions until success (ETX), available storage size for routing information, physical distance etc. The weight (cost) of each link was expressed in terms of time, and is required to be additive. This time measure represents the effect of each of the factors on the transmission time of a message using that specific link. For example, the higher the ETX value is, the longer the expected time required to successfully transmit a message over that link. This way, we may "punish" links of high ETX or a slow transmission rate by increasing their weight. Preventing bottlenecks is achieved by limiting the greediness of every node to minimize the number of hops to the root at all costs. For the present work we use a modified OF that considers: hop count, SINR (provided by the physical layer) and the transmission time. The resulting cost of the directed link (i, j) is

$$w_{i,j} = W_1 \cdot N_{hops} + W_2 \cdot ETX(\overline{SNR}) + W_3 \cdot D_{i,j} \tag{2}$$

where N_{hops} is the number of hops between node i and the root, $ETX(\overline{SNR})$ is the average number of transmission retrials as a function of the average SINR sensed by node i. $D_{i,j}$ is the transmission time (which is approximately equal to the inverse of the transmission rate, neglecting the propagation time) from node i to node j. W_1, W_2, W_3 are non-negative coefficients allowing us to calibrate the OF and convert the units of the physical factors into the relevant time units.

We have also added the TPL a suspension mechanism to better handle the churn of parent ranking as a result of the received noise. Using such a suspension factor α that prevents frequent changes of the computed ranks after receiving a DIO message, helps keeping the topology more stable. That is, NE i computing the rank $r_{i,j}$ after receiving the k'th ($k > 1$) DIO message from parent j using

$$r_{i,j}(k) = \alpha \cdot r_{i,j}(k-1) + (1-\alpha) \cdot mr_{i,j} \tag{3}$$

where $mr_{i,j}$ is the parent rank calculated directly from the information contained in the DIO message and the measured SINR. Using experiments, not presented in this paper due to lack of sufficient writing space, we concluded that $\alpha \simeq 0.7$ provides the best performance. This value complies with similar results from practical field measurements done by Mobix, for their proprietary routing protocol [8].

C. Simulated topology

Models that use geographical positions of the NEs usually ignore the fact that two nearby NEs may not communicate directly if they are connected to different electrical phase wires. Conclusions made from simulated topologies that are generated by considering communication parameters that only rely on nodes positions are not presenting true behavior of a PLC-SG accurate enough. Our alternative approach was to use a set of field experiments operated on a real low-voltage PLC network. These experiments resulted with a very detailed set of parameters for each of the communication links. We used this dataset to construct a "random PLC-SG network" by randomly selecting a subset of vertices (484 in the present case), together with their corresponding link properties. This allows us to simulate networks that resemble real PLC-SGs more accurately. The number of NEs in the simulation relates to the maximal number of NEs that might be accessed via a single network concentrator when using a MAC with similar capabilities to the one used in our model (for PLC networks using a slower MAC with a single channel, a common practice is to connect up to 200-250 NEs to a single concentrator).

D. Collected statistics

During the construction of the tree, we gathered a detailed statistics on the activity in the MAC and network layers for each of the 484 NEs. In this work we present statistics of the following parameters: 1) The elapsed time until the RPL protocol constructs a connected network - this answers the question of how long it takes for the network to become operational for actual data delivery. 2) The elapsed time until the RPL protocol converges to its minimal weighted routing tree. 3) Stability of the minimal DODAG topology. 4) The percentage of time that a NE dedicates to the construction and maintenance of the routing tree (i.e., overhead of management traffic). 5) The network activity level required to maintain a routing tree under realistic operating conditions.

IV. RESULTS

Without the modifications mentioned above, RPL did not converge to its final DODAG for large networks. Using the modified RPL protocol on the simulated topology resulted with the DODAG shown in Fig. 1. We have verified, by offline centralized algorithm (using the physical conditions taken from the simulation run time) that the presented topology is a minimal weighted DODAG.

The first statistics we refer to is the progress of connecting NEs to bidirectional communication with the network concentrator. The initial state of all NEs is *inactive*, so both joining the network and sending a DAO update to the concentrator is required prior to considering a NE as being in state *connected*. Figure 2 shows the percentage of connected NEs as a function of time. The results presented in Fig. 2 show that the network becomes fully operational after approximately 2,000 seconds. This behavior is typical in practical PLC networks of similar size that were observed in other simulations.

RPL frequently tries to improve the DODAG by distributing updated DIO messages. A better

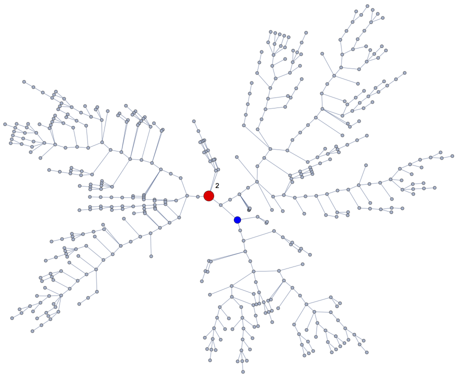

Figure 1. Routing tree of a 484 NEs PLC network. NE 2 (red) is the concentrator initiating the RPL instance.

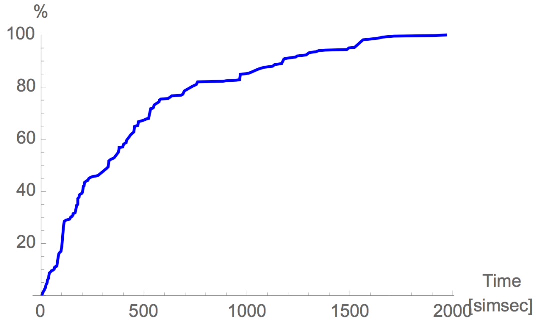

Figure 2. Percentage of connected NEs vs. time

DODAG is one having a lower total weight, i.e., the sum of the ranks of all the NEs in network, which are computed, using (2) and (3). Fig. 3 describes the network's weight as a function of time, starting with the first occurrence where all the NEs already have a finite rank. Each NE with initial state *inactive* is assigned an infinite rank, and all NEs obtain a finite rank (although not minimal) after several tens of seconds. This does not imply that in this stage the network is already connected, since a bidirectional NE-concentrator connectivity is required for all NEs for a network to be fully connected.

Observing the results shown in Fig. 3, one may divide the evolution of network's cost into two parts. The first part (grayed area) describes the process of *fast convergence*, where RPL improves the network to achieve a DODAG of a minimal weight. The second period (to the right of the red dashed line) illustrates the *maintenance phase* which is mainly characterized by small fluctuations of the weight of the DODAG. These fluctuations are caused by noise that is followed by changes in link qualities which affect the rank of the NEs (using (3)). We note again that without the hysteresis window, the maintenance phase suffers from much larger fluctuations.

Less fluctuations implies a more consistent topology of the DODAG. We measure the DODAG consistency by analyzing the decisions made by the NEs due to the DIO messages they receive. Each DIO message received by a NE may cause a change in its rank or a change in the selection of its preferred parent in the DODAG. Thus, following the number of DIO messages causing such a change may serve as a measure of the consistency of the topology. In large PLC-SGs, as discussed here, it is reasonable to always experience a certain level of inconsistency due to noise in the channel. We already showed that the total weight of the DODAG fluctuates in Fig. 3. Figure 4 shows the total number of DIO messages, every second, that are followed by a change in the rank or the parent of each NE. Although many changes are observable, even in the maintenance phase, we manage to stabilize the resulted DODAG structure by calibrating the parameters of the OF and α (see expressions (2) and (3), respectively).

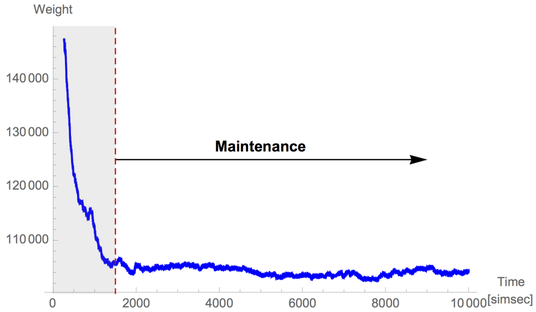

Figure 3. Total weight of the PLC-SG vs. time

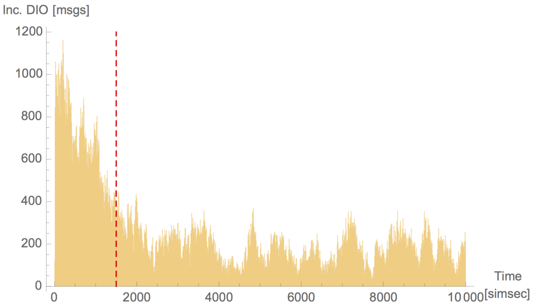

Figure 4. Number of inconsistent routing decisions vs. time

Having a consistent connected network is desired, however, for any practical purpose it is also desired that the NEs will dedicate a small fraction of their resources for maintenance activity of the topology and free as much communication resources as possible for transmission of actual data. Each connected NE may be in one of two states: *idle* or *busy*, where *busy* means transmitting or receiving. Since the analysis in this work is for the construction and maintenance of the routing DODAG only, an *idle* state may be dedicated later for transmitting actual data. For this reason, we measured the percentage of NEs in the *busy* state as a function of time. The blue curve in Figure 5 shows the percentage

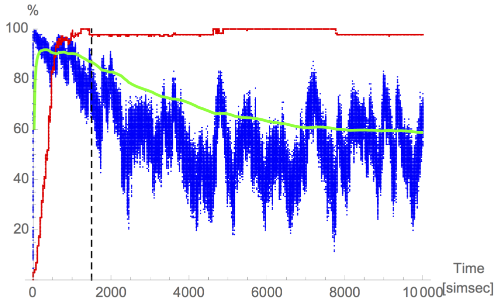

Figure 5. Percentage of busy NEs (blue), the average percentage of busy NEs (green) and of optimally ranked NEs (red).

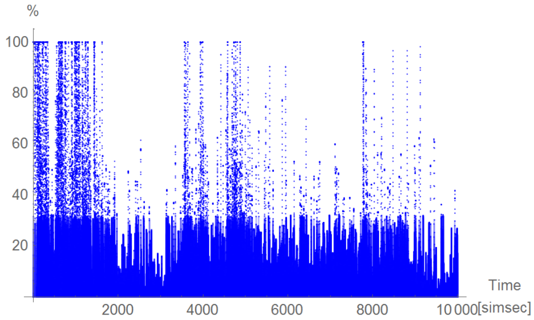

Figure 6. Activity level of a specific NE the PLC network

of busy NEs sampled every $50ms$ (the number of events is huge and cannot be displayed otherwise). The sampling interval was chosen with accordance to the longest transmission time that might be used, following the noise levels and the MAC protocol in use. The red curve shows the percentage of NEs having an optimal rank (which was computed outside the simulation). It is evident that the number of busy NEs is high. The average number of busy NEs converges towards a value of approximately 60%, which is presented by the green line. Eliminating the transient behavior of the fast convergence area, the percentage of busy NEs is decreased to 52%. This shows that in any given time, approximately half of the NEs are busy in maintaining the DODAG. This leads to a serious doubt in the applicability of the RPL protocol for routing in large PLC networks under realistic channel conditions. One immediate implication is the possibility to use a MAC protocol with multiple channels. This will allow to increase the number of NEs transmitting simultaneously using different orthogonal channels.

The results presented in Fig. 5 may be misleading in the sense that the number of busy NEs does not necessarily relate to the fraction of time that these NEs are available for data transmission. Fig. 6 shows the activity level for the previous second of one node that is 2 hops away from the concentrator (the blue node in Fig. 1). It is evident that there are periods of times where this NE would not be available for data transmission due to the activity of maintaining the DODAG. The most prominent character of this plot is the relatively low values of load that a NE experiences. For a NE that is close to the concentrator this is crucial since many other NEs are members of its subtree in the DODAG.

V. SUMMARY AND CONCLUSIONS

We have discussed the applicability of the RPL protocol for routing in large and realistic PLC-SGs. We have pointed out several criteria that any good routing protocol must comply with, and examined the performance of RPL against these criteria. We have found that under realistic conditions, that is, severe noise conditions and hundreds of nodes, RPL (as defined in [2]) does not perform as expected. A part of the network was not reachable from the concentrator and a DODAG was not created.

While trying to improve this faulty behavior we incorporated several mechanisms: a modified OF that considers multiple network parameters besides the formal hop count, a hysteresis decision window at the MAC and a relaxation mechanism for the ranking of the NEs. We have found that these mechanisms improved RPL's performance substantially: all NEs were reachable within a reasonable time, and the DODAG quickly converged to its minimal weight. The fluctuations in the network weight were also reduced substantially. However, our investigations were focused on the level of NEs' activity required to construct and maintain the DODAG. We have found that in any given time, about half of the NEs were busy with DODAG maintenance. Closer investigation of specific NEs activity showed that a large number of them were almost always busy with routing maintenance, leaving no resources for actual data delivery.

Our findings raise a doubt in the applicability of RPL as a preferred routing solution for large and realistic PLC-SGs.

REFERENCES

[1] E. Ancillotty, R. Bruno and M. Conti,"The role of the RPL routing protocol for smart grid communication," *IEEE Comm. Mag.*, Vol. 51, No. 1, pp. 75–83, 2013.

[2] T. Winter (editor) et al., "RPL: IPv6 routing protocol for low power and lossy networks," RFC 6550, 2012.

[3] http://www.omnetpp.org/ and http://inet.omnetpp.org/

[4] T. Clausen (editor) et al., "LOADng," https://tools.ietf.org/html/draft-clausen-lln-loadng-12.

[5] D. Sharma and A. N. Shukla, "A Comparative study of the routing protocols LOAD and RPL in low-power and lossy networks (LLN)," *Sch. J. Eng. Tech.*, Vol. 2, No. 1, pp. 85–87, 2014.

[6] P. Levis (editor) et al., "The trickle algorithm," https://tools.ietf.org/html/rfc6206.

[7] O. Gnawali and P. Levis, "The Minimum Rank with Hysteresis Objective Function," RFC6719, 2012.

[8] http://www.mobix.com