

Effect of Downlink Traffic on Performance of LoRaWAN LPWA Networks: Empirical Study

Konstantin Mikhaylov, Juha Petäjäjärvi and Ari Pouttu
Centre for Wireless Communications
Oulu, Finland
{firstname.lastname}@oulu.fi

Abstract—Today the LoRaWAN is among the most widely adopted Low Power Wide Area Network (LPWAN) technologies. In contrast to the other alternatives, LoRaWAN enables deploying both continent-wide public networks, and private networks composing one or several gateways. In this paper, we empirically investigate how does the presence of the downlink traffic affects the performance of uplink for LoRaWAN operating in 868 MHz EU bands. Our results show that in real-life the downlink transmissions can compromise performance of uplink and must be accounted for when planning a network. Also we demonstrate the effects that the selection of secondary frequency channels and data rates have on the performance of a LoRaWAN network. The reported results reveal new dependences and provide the ground truth reference for the future analytical works. For these reasons, the paper can be of interest to both practitioners, planning and deploying LoRaWAN networks, and the researchers performing analytical work on LoRaWAN.

Index Terms— LoRaWAN, LPWAN, downlink, uplink, performance, packet error rate, measurements.

I. INTRODUCTION

A great number of the Internet of Things (IoT) devices is deployed in our living environment already today and many more will appear in the years to come. One of the few commonalities between these devices is the wireless communication, which they utilize.

Albeit the landscape of the wireless communication technologies today is excessively diverse [1], the ones, which are known as the Low Power Wide Area Networks (LPWANs) [2][3], stand out of the mass. These solutions were developed with the ambitious goal of enabling low cost communication for massive and widespread deployments of the resource-limited autonomous machines.

Among dozens of the LPWAN technologies existing today, LoRaWAN [2]-[4] is among the most popular ones. This is also one of the very few technologies, which supports deployment of the operator-controlled continent-wide “public” networks, and of “private” networks controlled by micro-operators or individuals. The technology has been in massive roll out from 2015 and by now is present in over 100 countries [4].

Due to this deployment flexibility and the specifics of its technical solutions (e.g., operation in ISM bands

and the use of ALOHA-like media access in multiple frequency channels [3],[5]) since very early days the problem of LoRaWAN scalability has drawn significant attention. The maximum capacity of a single-gateway LoRaWAN network was assessed and the fundamental throughput limits were determined in [6]. The further studies [7]-[11] complemented these initial results. Notably, if the early works implied orthogonality of signals with different spreading factors (SFs) and impossibility of receiving packets under collisions, the recent ones account for the inter-SF interference and capture effect [12].

The early works dealing with LoRaWAN were primarily focusing the uplink-only networks. The effect of the downlink on the LoRaWAN network’s performance got attention only recently in [11] and [13]. The results reported in these studies show that LoRaWAN downlink performance is affected by the duty cycle restrictions on one hand and the amount of downlink traffic on the other. In the case of extensive downlink traffic, the acknowledgements get lost due to collisions, causing end devices (EDs) to deduce that their uplink frames are not delivered. Importantly, the majority of papers dealing with the performance of LoRaWAN networks published, including [11] and [13], use analytical methods and there is a clear lack of the experimental validation of these results and effects.

In this paper, we address this omission and report the results of an extensive measurement campaign, during which we empirically investigated how the performance of a LoRaWAN network is affected by the selection of communication parameters on one hand, and the presence of downlink traffic on the other.

II. EXPERIMENT SETUP

Due to the limited space, we refrain from detailing the basics of the technical solutions underlying LoRaWAN. For these, a reader can refer to, e.g., [6][10]-[13]. The details of the LoRa physical layer were reported, e.g., in [14] and [15].

A diagram sketching the structure of our experimental setup and a photo of the experimental environment are presented in Fig. 1. To prevent any influence of commercial LoRaWAN deployments or of the other radio systems, all the measurements were conducted in the 11x6,5x5,5 meter anechoic chamber.





Fig. 1. Experiment setup in the anechoic chamber

Execution of the measurements in isolated environment also allowed to alleviate the limitations regarding the on-air duty cycle imposed by the frequency regulations. In our experiments we used two types of devices:

- Firstly, five to ten EDs built using the CWC modular hardware platform [16] and equipped with RN2483 LoRaWAN transceivers have been deployed.
- Secondly, a Linux-based MultiConnect® Conduit™ commercial LoRaWAN gateway (GW) has been installed. The LoRaWAN network server (NS) was deployed right on the GW and configured to accept the frames from all the EDs involved in the experiments. A special supervisory software application was instrumented in Node-RED and deployed on top of the NS. The application logged all the received uplink packets and their parameters for further analysis and could generate traffic in downlink in response to uplink for the selected EDs.

In total up to eight 868MHz band frequency channels have been utilized in the measurements – the three default channels (868.1 MHz, 868.3 MHz and 868.5 MHz), and up to five secondary channels (869.1 MHz, 869.3 MHz, ... 869.9 MHz). The EDs were configured to follow the duty cycle regulations in the default channels, but had no restrictions for operating in the secondary channels. The GW followed the duty cycle restrictions in all the channels used.

TABLE I. PARAMETERS OF THE EXPERIMENTS

Phase	Nodes	Downlink	DR	Secondary channels
I	5	yes/no	unique ²	unique ⁴ /shared
II	10	yes/no	unique per pair ³	unique per pair ⁵ /shared
	10	yes/no	DR1	unique per pair ⁵ /shared
	10	yes/no	DR3	unique per pair ⁵ /shared
	10	yes/no	DR5	unique per pair ⁵ /shared
	10	partial ¹	unique per pair ³	unique per pair ⁵ /shared

¹ - only nodes 1,2,6 and 7 (here and further node numbers match ones in Fig. 1)

² - node 0 – DR1, node 1 – DR2, ..., node 4 – DR5

³ - nodes 0 and 5 – DR1, nodes 1 and 6 – DR2, ..., nodes 4 and 9 – DR5

⁴ - node 0 – 869.1 MHz, node 1 – 869.3 MHz, ..., node 4 – 869.9 MHz

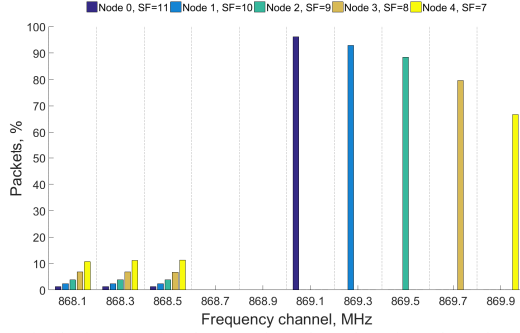
⁵ - nodes 0 and 5 – 869.1 MHz, ..., nodes 4 and 9 – 869.9 MHz

All the EDs and the NS were configured to enable activation by personalization. All uplink transmissions were sent with no ACK requested. The adaptive data rate (ADR) mechanism has been disabled. The transmit power for all the nodes was set to 14 dBm. For receive window 1 (RW1) the same SF and channel as for uplink was used, for RW2 – the SF12 and 869.5 MHz channel.

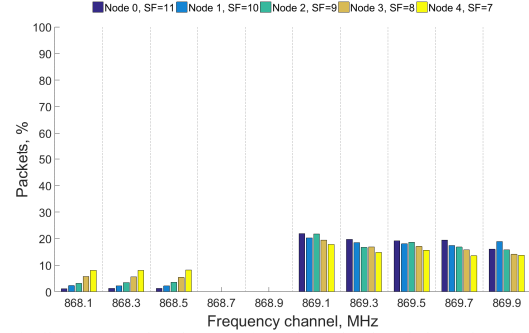
Once powered, each ED configured its LoRaWAN transceiver and started generating radio frames and sending them in uplink. A random delay (0.5 to 1.5 seconds) was introduced between each two sequential packets to prevent any synchronization. Even though this traffic rate is hardly conventional for a typical LPWAN application, this enabled us to artificially increase the volume of our test network. Namely, a single test ED could imitate 88-700¹ traditional nodes. The packets transmitted by EDs included 36 bytes of application payload, which contained the 4-byte software counters for the transmitted uplink and the received downlink packets. These counters, extracted from the log of the NS, were used to estimate the packet error rate (PER). Additionally to the payload, for each uplink packet the radio signal strength indicator (RSSI), signal noise ratio (SNR) estimate, used frequency channel (CH) and other service data provided by the GW were recorded. Note, that given the small distance between the EDs and the GW (see Fig. 1), a 60 dB attenuator was added between the antenna and the radio input of the GW to prevent any receiver saturation.

The measurements were handled in two phases (see Table 1). In the first phase only five EDs were used. In the second phase the number of EDs was increased to ten. The first phase was used to validate the correctness of the system's behavior in general and estimate the performance under minimal interferences. The second phase enabled to study the performance of LoRaWAN under interferences. In total, four different modes (with/without downlink, with unique/shared channels – see Table 1) were tested in the first phase and 18 modes – in the second. Each of these modes was run for over 75 minutes non-stop and involved transmission of five to twenty thousand radio packets. The collected data from the experiments was validated and statistically processed using MATLAB scripts.


¹ for SF7 and SF12, respectively, when compared to 0.1% duty cycle limited network sending packets of 36 bytes



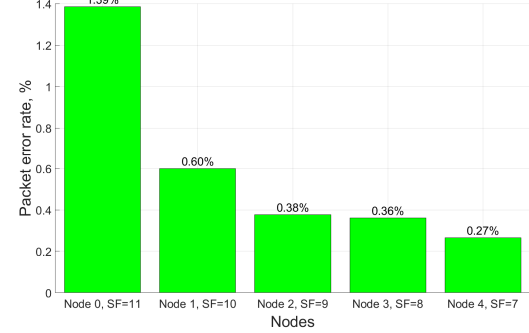
(a) Distribution of packets between freq. channels (unique channels & DRs)



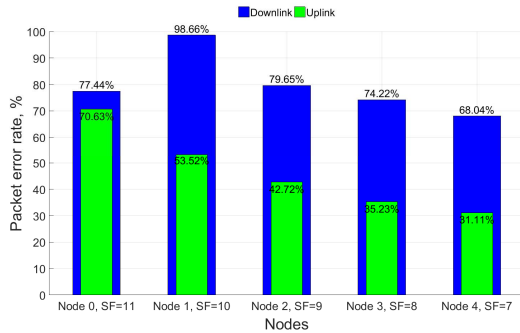
(b) Distribution of packets between freq. channels (shared channels, unique DR)



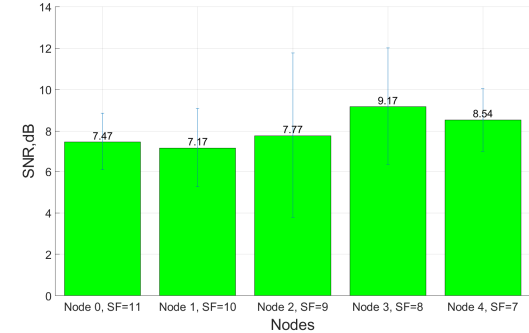
(c) PER for uplink (unique channels & DRs, no downlink)



(d) PER for uplink (shared channels & unique DRs, no downlink)



(e) PER for uplink and downlink (unique channels & DRs)



(f) SNR for uplink (shared channels & unique DRs, no downlink)

Fig. 2. Selected results of measurement phase one involving five end devices

III. RESULTS AND DISCUSSION

A. Measurement results for phase I – five test EDs

The selected results of the first phase are presented in Fig. 2. Figs. 2 (a) and (b) show the distribution of all the uplink packets between the frequency channels when each node had a unique secondary channel and when all secondary channels were shared, respectively. Figs. 2 (c) and (d) reveal the PER for uplink-only communication for the same cases. Fig. 2 (e) illustrates uplink and downlink PER. The latter was measured as the number of downlink packets received by an ED, divided by the total number of uplink packets sent by this ED. Finally, Fig. 2 (f) depicts the mean and the standard deviation of signal noise ratio (SNR) for the uplink packets from the different EDs. The following conclusions can be drawn from the presented results:

1. Even though all EDs transmitted with different SFs and in various secondary channels – some packets were lost. The portion of lost packets was below 1%.
2. If EDs share secondary channels – more packets are lost. The EDs operating at lower SFs are affected less, than the ones using lower SFs.
3. The downlink traffic of the GW is limited – only a small portion of all uplink frames was responded to.
4. Presence of downlink traffic causes significant increase of uplink PER.
5. Distribution of the packets between the frequency channels depends on the data rate (DR) used.

Note, that during the measurements for all EDs the mean SNR of the received packets was above 7 dB and the mean RSSI - between -90 and -95 dBm, thus, theoretically, enabling error-free communication.

B. Measurement results for phase II – ten test EDs

The selected results of measurement phase two are depicted in Fig. 3. Figs. 3 (a) and (b) illustrate the PER for uplink-only case, (c) and (d) – the PER for uplink and downlink, (e) and (f) – the portions of downlink packets sent in RW1 and RW2. Figs. 3 (a), (c) and (e) address the case when each secondary channel is used by only two EDs, and Figs. 3 (b), (d) and (f) - when all the secondary channels are shared among all EDs. Finally, Figs. 3 (g) and (h) reveal the PER in the case when the downlink is sent only to two EDs. The following observations can be drawn from these results:

1. The presence of a second ED operating in the same secondary channel(s) and using same DR increases PER. This increase is higher for nodes with high SFs.
2. Increase of the ED number and of downlink traffic results in increase in both uplink and downlink PER. This is especially notable for high SFs.
3. Under heavy downlink traffic the use of RWs is uneven. The transmissions with high SFs tend to get responded in RW1, while the low SF – in RW2.
4. The downlink transmissions negatively affect also the uplink transmissions with other DRs, including the ones done in other frequency channels.

Result 4 may need some further clarification. In the case illustrated in Fig. 3 (g) the GW was responding only to nodes 2 and 3 (see Fig. 3 (h)), which use SFs 8 and 9, and operate in secondary channels 869.5 MHz and 869.7 MHz, respectively. Comparing Figs 3 (a) and 3 (g) one can see that, e.g., for nodes 1 and 5 the PER has increased from 42% to over 75%. Meanwhile, from Fig 2 (a) (similar charts built for the case of ten nodes also confirm this) one can see that for SF7 over 90% of the packets are sent in secondary channels, which differ from the ones, where downlink packets were sent.

C. Discussion

As one can see from the presented results, a LoRaWAN gateway is limited in respect to its downlink (i.e., sending acknowledgments, commands or the application data). Among the major reasons for this is the restriction for the GWs duty cycle by the frequency regulations. In this respect, the results of our measurements confirm that of [6],[11] and [13].

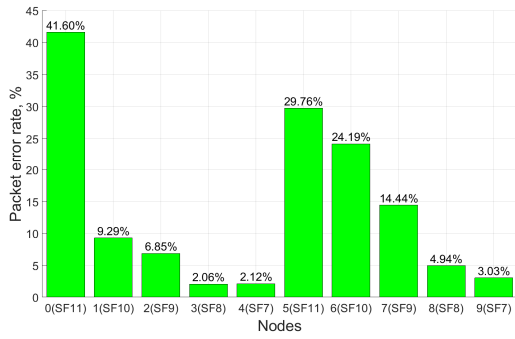
The higher PER for the EDs operating in the shared channels with different DRs can be attributed to the inter-SF interference reported earlier, e.g., in [12]. The higher PER for nodes operating with low DRs is due to the longer on-air time, which increases the probability of collisions. Meanwhile, the fact that, e.g., in the case depicted in Fig. 3 (a) for SF12 (i.e., a 2.3 second long frame and 0.5 to 1.5 second random delay between the frames) only 35% uplink packets for two nodes were lost confirms the presence of the capture effect. The capture effect, along with the different channel budgets coming from the difference in distances between various EDs and the GW (see Fig. 1), as well as due to

the difference of performance of the individual EDs, explains the variation of the PER for different EDs with same DRs. The observed effect of DR of uplink transmission on the distribution of the downlink between the different RWs is due to the difference in the back-off times for the frequency channels for RW1 and RW2, originating from the different on-air times of downlink transmissions and the duty cycle restrictions.


Nonetheless, the most notable result of our experiments is the observed strong effect of the downlink on the performance of uplink. The reason after this is the fact that RW1 in LoRaWAN typically uses the same SF and the same frequency channel, as the previous uplink. Therefore, a GW's downlink transmission in RW1 can interfere with an uplink transmission of another ED selecting this channel. Meanwhile, RW2 is opened in a pre-defined channel and uses a pre-specified SF. Depending on the frequency plan of LPWAN, the channel used for RW2 may also be assigned for uplink. In this case, downlink transmissions may interfere with uplink also for RW2.

In the both described cases, since the power of a downlink transmission significantly (i.e., 60 to 120 dB) exceeds that of the incoming uplink, this is hard to correctly receive uplink under downlink interference. Note also, that since RW2 often uses higher SF than that of RW1, this is often more beneficial energy and time-wise for a GW to use RW1. The substantially shorter on-air time for transmitting a response in RW1 also enables sending more frames even in a channel with more conservative duty cycle restrictions. This is the reason why, e.g., in the case illustrated in Fig. 4e, almost 75% of all the downlink frames have been transmitted by the GW in RW1. To the best of our knowledge, the described effects have neither been reported, nor modelled in the previous studies, which makes comparison of our obtained results with the ones reported there impossible. We are also unaware of analytical models describing these effects.


The highlighted problem of interference between downlink and uplink can be approached in several different ways. First, at the GW side, one can assign RW2 to a dedicated channel and prevent use of RW1. Such functionality is available in particular commercial LoRaWAN GWs. The downsides of this are the reduction of radio resources available to the GW for conveying downlink, the need of using constant SF for downlink, and the potential interferences between the neighboring GWs. Second, one can consider introduction of some form of listen-before-talk mechanism at the GW, the EDs, or both of these. Third, one can increase the density of the GWs, or even introduce in the network the GWs, which would not convey any downlink. Alternatively the GWs receive and transmit chains can be separated in space and use different antennas. The efficiency of these solutions and the possibility of others requires further investigation.




(a) PER for uplink (unique channel & DR per node pair, no downlink)




(b) PER for uplink (shared channels & unique DR per node pair, no downlink)



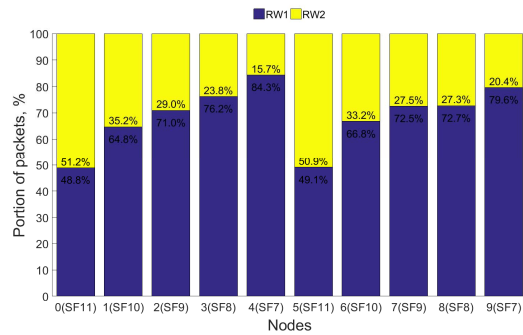
(c) PER for uplink and downlink (unique channel & DR per pair)



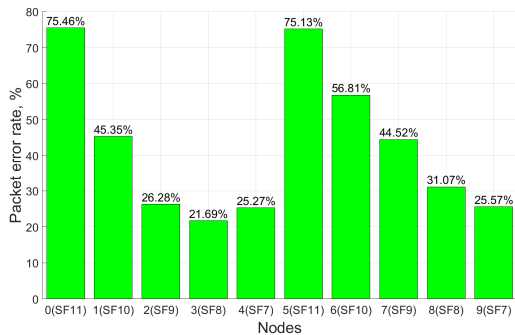
(d) PER for uplink and downlink (shared channels & DR)




(e) Distribution of downlink between RWs (unique channel & DR per pair)



(f) Distribution of downlink between RWs (shared channels & unique DRs per pair)



(g) PER for uplink (unique channel & DR per pair, downlink for EDs 2 and 3)



(h) PER for downlink (unique channel & DR per pair, downlink for EDs 2 and 3)

Fig. 4. Selected results of measurement phase two involving ten end devices

IV. CONCLUSIONS

In the current paper we empirically characterized the performance of a LoRa wide area network, paying special attention to the effect of downlink communication on the network performance. To the best of our knowledge, the current paper is the first one addressing this problem by means of extensive practical measurements with real-life hardware. In this respect, our results compliment the ones presented in the recently published analytical works (e.g., [11] and [13]).

Among the most important results of our measurements is the observation of the strong negative effect, which the downlink communication has on the performance of uplink. Our experiments have shown that even the uplink transmissions encoded with other DRs and done in other frequency channels may get affected. Due to the limited space, in the current paper we do not provide any analytic model of this effect, which obviously has to be developed. This, together with the detailed investigation of the different ways of addressing this problem, both in theory and in practice, we plan to address in our further studies.

The other notable results of this work include: 1) the observed effects of the ED's SF on the distribution of the packets between the frequency channels, and 2) the effect of the uplink DR (i.e., SF) on the selection of the RW in downlink. To the best of our knowledge, these effects have neither been observed, nor reported yet. The possible reasons for these effects have also been discussed, albeit further analytic studies are required to confirm that we have identified them all.

Finally, the results of this study (Figs. 2 (a) and 2 (c)) indirectly prove that a LoRaWAN technology can implement multiple service classes (i.e., priorities) by assigning a separate secondary frequency channel to each priority level. Using this method, the EDs with different levels of priority will compete mostly between themselves. This may be important for increasing the reliability of LoRaWAN communication and for fighting the in-network interferences, and may enable novel applications and use cases with more challenging requirements in respect to the reliability of communication.

ACKNOWLEDGMENT

The work was undertaken in the context of the P2P-SMARTEST project (<http://www.p2psmartest-h2020.eu/>), an Innovation Action funded by the H2020 Programme, contract number 646469, and the LPWAN-evolution activity funded by the University of Oulu. Also this research has been financially supported by Academy of Finland 6Genesis Flagship (grant 318927) program.

REFERENCES

- [1] S. Andreev *et al.*, "Understanding the IoT connectivity landscape: a contemporary M2M radio technology roadmap," *IEEE Com. Mag.*, vol. 53, no. 9, pp. 32-40, Sep. 2015.
- [2] U. Raza *et al.*, "Low Power Wide Area Networks: An Overview," *IEEE Com. Surv. Tut.*, vol. 19, no. 2, pp. 855-873, 2nd q. 2017.
- [3] M. Centenaro *et al.*, "Long-range communications in unlicensed bands: the rising stars in the IoT and smart city scenarios," *IEEE Wireless Com.*, vol. 23, no. 5, pp. 60-67, Oct. 2016.
- [4] A. Weissberger, "LoRaWAN gains momentum: NEC LoRaWAN server + New Zealand nationwide network," IEEE ComSoc Technology Blog (24.11.2017) [Online]. Available: <http://techblog.comsoc.org/category/lpwans>
- [5] J. Petäjälä *et al.*, "On the Coverage of LPWANs: Range Evaluation and Channel Attenuation Model for LoRa Technology", in *Proc. Int. Conf. Intelligent Transport*, 2015, pp. 1-6.
- [6] K. Mikhaylov *et al.*, "Analysis of the Capacity and Scalability of the LoRa Wide Area Network Technology", in *Proc. Eur. Wireless*, 2016, pp. 1-6.
- [7] B. Reynders *et al.*, "Range and coexistence analysis of long range unlicensed communication," in *Proc. 23rd Int. Conf. Telecomm.*, 2016, pp. 1-6.
- [8] O. Georgiou and U. Raza, "Low Power Wide Area Network Analysis: Can LoRa Scale?," *IEEE Wireless Com. Letters*, vol. 6, no. 2, pp. 162-165, Apr. 2017.
- [9] M. C. Bor *et al.*, "Do LoRa Low-Power Wide-Area Networks Scale?," in *Proc. 19th ACM Int. Conf. Modeling, Analysis Simulation Wireless Mob. Syst.*, 2016, pp. 59-67.
- [10] F. Adelantado *et al.*, "Understanding the limits of LoRaWAN", *IEEE Commun. Mag.*, vol. 55, no. 9, pp. 34-40, Sep. 2017.
- [11] F. Van den Abeele *et al.*, "Scalability Analysis of Large-Scale LoRaWAN Networks in ns-3," *IEEE IoT J.*, vol. 4, no. 6, pp. 2186-2198, Dec. 2017.
- [12] K. Mikhaylov *et al.*, "On LoRaWAN Scalability: Empirical Evaluation of Technology Susceptibility to In Network Interference", in *Proc. Eur. Conf. Netw. Commun.*, 2017, pp. 1-6.
- [13] A.-I. Pop *et al.*, "Does Bidirectional Traffic Do More Harm Than Good in LoRaWAN Based LPWA Networks?," [Online]. Available: <https://arxiv.org/abs/1704.04174>
- [14] A. Augustin *et al.*, "A Study of LoRa: Long Range & Low Power Networks for the Internet of Things", in *Sensors*, paper 1466, Sep. 2016.
- [15] Semstech, "AN1200.22: LoRa™ Modulation Basics", [Online]. Available: <https://www.semtech.com/uploads/documents/an1200.22.pdf>
- [16] K. Mikhaylov and J. Petäjälä *et al.*, "Design and Implementation of the Plug&Play enabled Flexible Modular Wireless Sensor and Actuator Network Platform", *Asian J. Control*, vol. 19, no. 5, pp. 1-21, Sept. 2017.