Paper ID AM-SP1343

Energy Analysis of LoRaWAN Technology for Traffic Sensing Applications

Samarth Mathur[1], **Anand Sankar**[2], **Prajwal Prasan**[2], **Bob Iannucci**[2]
1. Carnegie Mellon University, Department of Electrical and Computer Engineering
NASA Ames Research Park, Moffett Field, CA 94035 USA
+1 650 495 4875 – samarth.mathur@sv.cmu.edu
2. Carnegie Mellon University
Department of Electrical and Computer Engineering, USA

Abstract:

Traffic sensing is considered an integral part in the vision of Intelligent Transportation Systems (ITS). Cities will be able to work towards better traffic management solutions and road safety with this data. Building the network infrastructure for traffic sensing will place a tremendous burden on cities in terms of cost and energy. To change this, there is a need for a paradigm shift in connectivity infrastructure. In this paper, we explore the possibility of using LoRaWAN to enable the applications of crowdsourced traffic calming and traffic sensing. We explore the features and inadequacies of LoRaWAN in order to propose an adaptive data aggregation and re-transmission algorithm for relaying traffic data from sensors. To validate the use of LoRaWAN, we conduct an energy analysis of the protocol and determine the effect of parameters such as payload size and data rate on the energy consumption. The analysis leads to the conclusion that the sensor devices can be made ready to last for more than five years on small batteries as would be required for such an application.

KEYWORDS:
Low-Power Wide-Area Networks, LoRaWAN, Traffic Sensing, Smart Cities

I. Introduction

Rapid growth of vehicles due to urbanization has led to a need for traffic calming, an approach for implementing road safety measures and moderating vehicle flow. An essential part of traffic calming is to ascertain the choice of the calming measures. For evaluating alternatives, cities use traffic sensing techniques that can be rapidly deployed such as traffic surveillance cameras, radars, inductive loops and pneumatic road tubes. Data collection from the sensors and their subsequent analysis can end up significantly delaying the evaluation of the calming measures. An effective solution would be one in which the sensors can be rapidly deployed, and information rapidly obtained. To that effect, there is a need to build fast and reliable networks for the sensors. One way to do this is by establishing a minimalist networking infrastructure across the city. This network could then be incrementally built using a bottom-up approach to support our requirements. We call this concept *crowdsourced traffic calming*.

We present the idea of *crowdsourced traffic calming* and *crowdsourced cities* in more detail in a companion paper [1]. For envisioning this concept, pervasive and open networks that provide city-wide coverage are indispensable. It would be impractical and costly for cities to supply power for all traffic sensors, so it is reasonable to assume that the sensors would be energy-constrained. Once deployed, they would have to last for five to ten years when powered using modest batteries. This brings an additional requirement for the networks: to have minimal energy needs.

Being short-range, Wi-Fi and Bluetooth technologies are ineffective in this scenario. Cellular networks would be ideal in terms of coverage. Infrastructure for pervasive cellular coverage of cities is already prevalent. But the cost and energy demands of cellular modems hinder their use. Even Wireless Sensor Networks (WSNs) using mesh networks drain unnecessary power as they have to keep their receivers continually on to transmit data across the mesh. To satisfy these requirements, a device-centric connectivity solution is required which can offer wide-area coverage at a lowest energy-cost. This has led to the emergence of Low-Power Wide-Area Networks (LPWANs) such as LoRaWAN [2], SigFox [3], NB-IoT [4], and WAVIoT [5]. LPWANs offer considerable advantages over mesh networks, which have been discussed by Filho et al [6] and Centenaro et al [7].

For our application, the Long-Range Wide-Area Network, or LoRaWAN is a promising candidate. It targets key Internet of Things requirements such as low cost, low energy consumption and high network capacity. It has an open specification for its protocol, and the network can be customized according to user requirements. To utilize LoRaWAN for our sensor network, there is a need to scrutinize its energy implications on the sensor devices. In our paper, we expand on this problem to determine the energy cost of communication for traffic sensors. In Section II, we give a brief introduction to LoRaWAN and the LoRa modulation scheme in order to theorize expected results. Section III describes our approach to setting up experiments for performing energy analysis. We also lay out an adaptive algorithm for aggregation and transmission of sensor data. The experimental results are detailed in Section IV. In Section V, we analyze the results obtained to determine the estimated lifetime of our sensor device radios. We also present other concerns regarding network capacity and range. The key conclusions drawn are then presented in Section VI along with planned future work.

Contributions: This work presents an analysis of energy consumption characteristics of LoRaWAN for the application of traffic sensing.

- We define an adaptive algorithm for data transmission for our application that can aggregate and re-transmit data.
- We analyze the effect of LoRaWAN data rate on the energy cost per transmission.
- We analyze the effect of varying LoRaWAN payload size on the energy-per-useful-bit for data transmission.
- We analyze the effect of data aggregation and time period of data collection on the lifetime of our end-device.

II. Related Work

A. The LoRa Wide Area Network (LoRaWAN) Protocol

An outline of the device-centric LoRaWAN network architecture is presented in Figure 1. Unlike mesh networks, sensors are typically laid out in a star or star-of-stars topology in LoRaWAN. This results in shorter transmission periods than multi-hop networks. The sensors can remain in a sleep state, only waking up to transmit data. The sensor information is received by gateways which extend coverage across the city. The gateways use IP networks to aggregate this data on a cloud-based network server. The traffic data can then be accessed for monitoring through user-created applications.

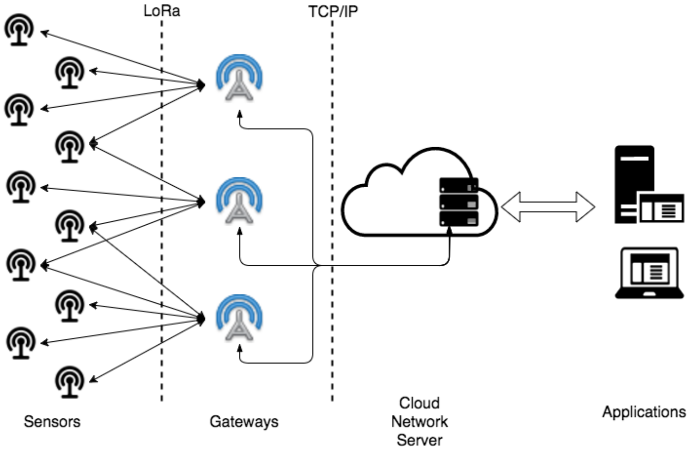

Fig. 1: LoRaWAN Network Architecture

LoRaWAN defines the following classes of end-point devices for varying operations [8]:

1) **Class A**: Class A devices use bi-directional communication, where each end-device has an up-link transmission followed by two short down-link receive windows. Transmission slots are scheduled in an ALOHA network-like manner.

2) **Class B**: In addition to the Class A receive windows, Class B devices open extra receive windows at scheduled times. Class B devices require a time-synchronized beacon from the gateway to open the receive windows at scheduled times.
3) **Class C**: Class C devices have continually open receive windows, only closed when transmitting. They have the lowest latency among all classes, but the highest energy consumption.

Class A devices have the minimum energy requirements among the three classes. In a traffic sensing system, nearly all communication is directed from the sensor to the gateway. Minimal receive slots are needed. Therefore, we select Class A operation for our devices. LoRaWAN also supports message acknowledgements (ACKs) and Class A devices can receive the ACKs in either one of the receive windows. In the event of a transmission failure, the LoRaWAN specification does not implement packet re-transmission. Nevertheless, if there is a need to safeguard data-loss, we can take advantage of the ACKs to develop a re-transmission algorithm.

B. LoRa Technology

LoRaWAN is based on the LoRa modulation scheme which uses chirp spread spectrum (CSS) modulation to enable long-range communication. The scheme generates a chirp signal that progressively varies in frequency as linearly increasing or decreasing cycles. Payload data is chipped at a higher rate, determined by the spreading factor (SF) or processing gain, and modulated onto the chirp signal. The bit-rate R_b of the modulated LoRa signal is then given by

$$R_b = \frac{SF \times CR}{T_s} \tag{1}$$

where

- SF is the spreading factor (7-12)
- CR is the coding rate, which is used for forward error correction (4/5 to 4/8)
- T_s is the symbol time in seconds and defined by

$$T_s = \frac{2^{SF}}{BW} \tag{2}$$

Substituting Equation 2 in Equation 1, we observe that increasing the spreading factor leads to a lower bit-rate of the modulated LoRa signal. Equation 2 also implies that a higher spreading factor has a longer time-on-air for LoRa packets. This would result in higher energy usage for the transceiver. LoRa controls the spreading factor and bandwidth through a parameter called the data rate (DR). Table I lists the primary data rates for the North American region. Other parameters which can be controlled are the payload size of a packet, transmit power for the radio and the error coding rate.

TABLE I: Data Rates (DR) for North America

DR	*SF/Bandwidth*	*Bit Rate*	*Max Payload Size*
0	10 / 125 kHz	980 b/s	19 bytes
1	9 / 125 kHz	1760 b/s	61 bytes
2	8 / 125 kHz	3125 b/s	134 bytes
3	7 / 125 kHz	5470 b/s	250 bytes
4	8 / 500 kHz	12500 b/s	250 bytes

The choice of the spreading factor is a challenging design decision to ensure optimal energy consumption and receiver sensitivity. The receiver sensitivity S in dBm of a radio at room temperature is given by

$$S = -174 + (10 \times \log_{10} BW) + NF + SNR \tag{3}$$

where

- -174 is the term due to thermal noise in dBm/Hz
- BW is the receiver bandwidth in Hz

- NF is the receiver noise figure in dB
- SNR is the minimum output signal-to-noise ratio (in dB) required at the receiver

As the spreading factor is an indication of the processing gain, increasing it allows a lower minimum SNR to be processed at the receiver. Using Equation 3, this leads to a better receiver sensitivity for the radio. An increase in bandwidth decreases the receiver sensitivity, but the data is transferred faster, leading to lower energy consumption. Overall, as the data rate is increased, we would expect to observe a reduction in receiver sensitivity and energy use. Lower energy costs created by increasing the data rate will be accompanied by a reduction in the receiver sensitivity, and consequently, the transmission range of the sensors.

The frequency and payload size of data transmission also impact energy use. The number of transmissions in a day would be directly related to the energy use of the device. Increasing the payload size leads to a higher number of symbols per packet, thereby increasing the energy cost per packet transmitted. However, a packet with a larger payload carries more information than a packet with a small payload. To analyze the effect of payload size, we choose our figure of merit as the energy-per-useful-bit of transmitted data. The energy-per-useful-bit E_b is defined as

$$E_b = \frac{Total\ energy\ consumed}{Total\ bits\ transmitted} \tag{4}$$

We would expect the energy-per-useful-bit to decrease as the payload size increases. This is because the size of packet meta-data such as the preamble, header and error correction information remains nearly constant. Smaller payloads would have higher overhead costs as compared to larger payloads. This makes it desirable for the sensors to send data with larger payload sizes. This observation motivated us to create an algorithm for aggregating data from consecutive time-periods and transmitting the data as a single LoRa packet. This aggregation will result in a higher latency in data transmission, as data would be transmitted only after the aggregation is completed. Lower energy costs created by reducing the frequency of transmission would be accompanied by a higher data latency.

C. Traffic Sensing Using Magnetometers

For gathering traffic data, we use a wireless sensor developed at Carnegie Mellon University to detect vehicles and estimate their speeds and type. The principal components of the wireless sensor are a MEMS magnetometer, a low-power ESP32 microprocessor and a LoRa transceiver. Coleri et al have already demonstrated that traffic statistics such as vehicle count [9], speed and classification [10] can be estimated from vehicle signatures captured with a single magnetic sensor. The sensors are designed to be embedded in raised-pavement Stimsonite markers, allowing straightforward placement and access. The magnetometer uses an interrupt mechanism to wake up the processor when a vehicle is detected. The processor analyzes and stores the captured data from the magnetometer. It then aggregates this data in user-specified time-intervals. These data-packets can then be further aggregated and scheduled for transmission as a single payload over the LoRa transceiver. The LoRa transceiver interacts directly with any LoRaWAN gateways within its range, sending over the packets with traffic information.

III. Approach

A. LoRaWAN Packet Transmission Algorithm

The sensors monitor vehicles in real-time, aggregating data by vehicle type over a finite time-period. One data-point is reported per vehicle type. For our experiments, we defined two vehicle types: small vehicles such as cars and two-wheelers, and large vehicles such as buses and trucks. Sensors report two data-points per time period using LoRa transceivers. An individual sensor data-point would be 3 bytes in size and contain the following values:

- The number of vehicles of that type (2 bytes, up to 65,535 vehicles of a type)
- The average speed of the vehicles (1 byte, average speeds of up to 255 miles/hr)

Figure 2 shows the structure of our LoRa physical layer packet. The packet starts with a preamble sequence for receiver synchronization. A header with packet identification and payload control information is then sent followed by the payload. The header and payload have separate cyclic redundancy check (CRC) sections for error checking. Since there were two data-points, 6 bytes of traffic data were reported for a time-interval. Combined with 4 bytes

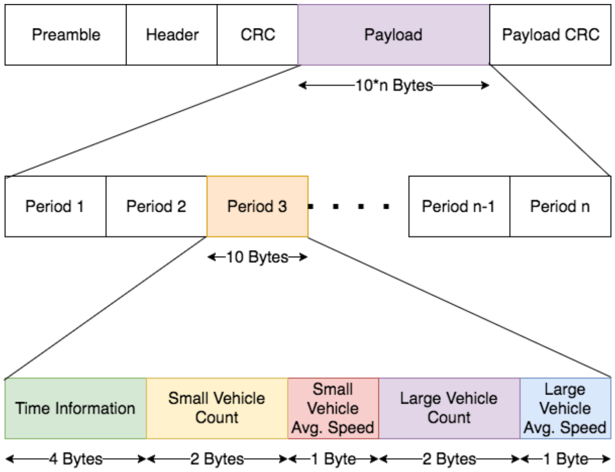

Fig. 2: LoRa Physical Layer Packet Structure

of time and date information, we considered the total transmitted data to be 10 bytes per time-interval. For our experiments, we varied these time-intervals between one and thirty minutes. We also aggregated data from up to five consecutive time-periods, and transmitted the information as a single payload of up to 50 bytes.

It is likely that gateways might experience outages due to power and link failures, sometimes lasting for multiple days. Collision between packets and high interference might also cause transmission failure. To prevent loss of data, a more robust re-transmission scheme is needed. This would require additional buffers for storing and aggregating data. Additionally, user requirements might change over time. For example, transportation officials might be interested in real-time minute-interval data during rush-hour, but might only need one data-point for every five minutes at night. As real-time analysis for night data is not as important, the data could be further aggregated with later intervals by the sensor and transmitted as a single payload. Based on the variability of user requirements, we designed a parameterized algorithm with configurable buffers and time-periods.

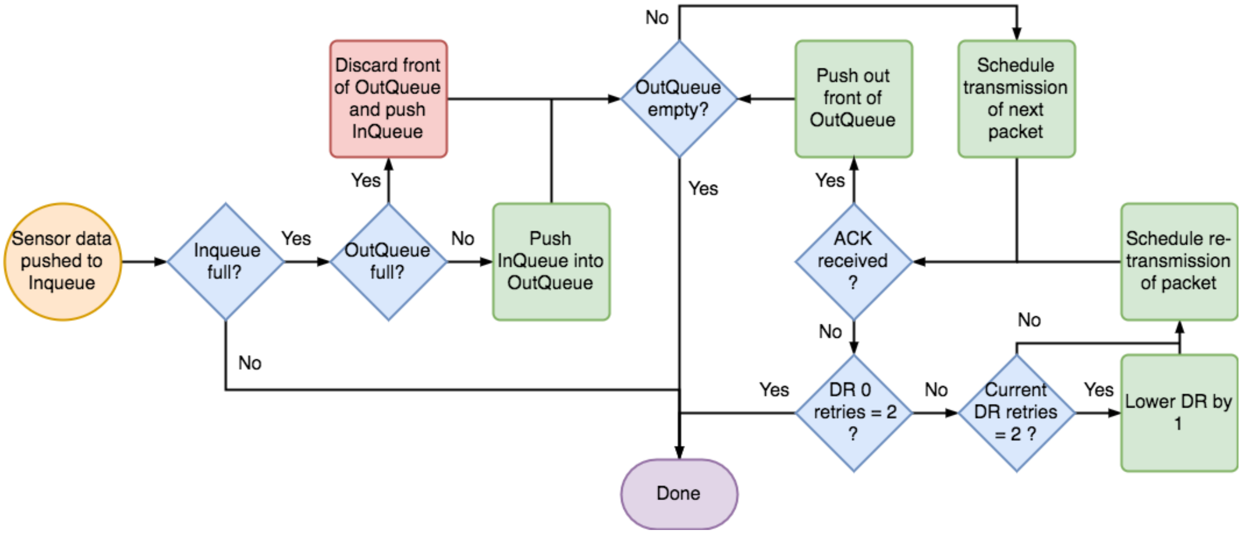

Fig. 3: LoRaWAN Packet Transmission Algorithm

Figure 3 shows a flow chart of the algorithm. Sensors collect data over a user-defined time-interval and send it to the input queue InQueue. The size of the InQueue is customizable and can vary from 10 to 50 bytes. A task waits for the InQueue to fill-up and then pushes the data as a single packet to the output queue OutQueue. The OutQueue can hold up to 5000 bytes of data for managing prolonged outages. The first packet is then read from the OutQueue and transmitted. If an ACK is received from the gateway, the processor discards this packet from the OutQueue, checks for and schedules any packets still present in the OutQueue. If an ACK is not received, the processor re-transmits the packet.

The re-transmission procedure is left up to the end-device designer by the LoRaWAN specification. In our design, the data rate is lowered after every two tries if an ACK is not received. After two tries at the lowest data rate, the processor goes back to sleep, The packet remains in the OutQueue. The processor tries to retransmit this packet when the InQueue is filled again. If the OutQueue is full, the packet with the earliest time-interval is discarded to allow newer data to be recorded. Data is therefore lost only when the OutQueue is full. When no data is to be sent, the LoRa transceiver switches to a sleep mode.

B. Experimental Setup for LoRa Tranceiver Energy Cost

We designed our experiments to study the effect of the data rate (spreading factor and bandwidth) and the payload size on the energy cost and energy-per-useful-bit of the LoRa radio. We used a Semtech SX1276 LoRa transceiver module as our end-device radio and a Multi-Tech MultiConnect Conduit with an MTAC-LORA-915 card as our LoRaWAN gateway. The transmit power was kept at $+14$dBm, and the coding rate was set to $4/5$. We used an Arduino implementation of the LoraMAC-In-C library (LMIC) [11], which is a C implementation of LoRaWAN by IBM. To measure real-time current consumption of the LoRa transceiver, we used a PowerDué: a customized Arduino development board with computation and instrumentation capabilities. For all experiments, the LoRaWAN code to be tested was downloaded and run on the PowerDué. Packet transmission was determined to be successful when data was received at the gateway, which we observed through a simple application server. We also verified it by checking the ACKs at the LoRa transceiver.

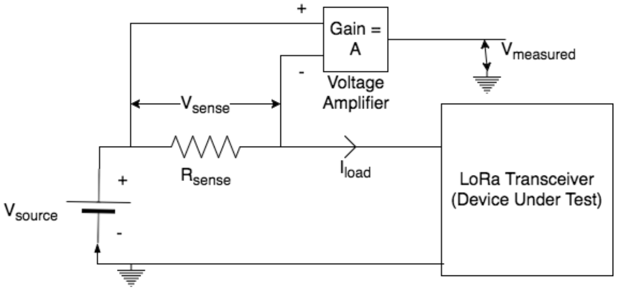

Fig. 4: Setup for measurement of LoRa power consumption

Figure 4 shows the general setup used for power measurement on the PowerDué board. The measured voltage $v_{meas}(t)$ at the output of the setup is an amplified version of the actual voltage across a sense resistor R_{sense} connected in series with the LoRa radio. Knowing the measured value of voltage $v_{meas}(t)$ and the amplifier gain A, we can find the load current $i_{load}(t)$ drawn by the transceiver as

$$i_{load}(t) = \frac{v_{sense}(t)}{R_{sense}} = \frac{v_{meas}(t)}{A \times R_{sense}} \tag{5}$$

The value of R_{sense} should be such that it minimizes its effect on the transceiver load current. For our setup, we use a 1.33Ω resistor as the sense resistor. The transmission power p(t) can then be approximated as

$$p(t) = V_s \times i_{load}(t) = V_s \times \frac{v_{meas}(t)}{A * R_{sense}} \tag{6}$$

where V_s is the 3.3 V source voltage. Based on the power and transmission time, we can calculate the energy cost for different payload sizes and data rates. The energy cost for transmission E_{tx} is given by

$$E_{tx} = \int_{t_1}^{t_2} p(t)dt \tag{7}$$

where $t_2 - t_1$ is the time taken to transmit a LoRa packet. Using the value of E_{tx}, we then calculate the energy-per-useful-bit E_b as

$$E_b = \frac{E_{tx}}{S} \tag{8}$$

where S is the payload size in bits. Using Equation 7, we can also measure the energy costs E_{rx_1}, E_{rx_2} for the two receive windows and E_s for sleep mode respectively. The estimated lifetime of the device is then obtained as

$$Est.\ life\ (days) = \frac{E_{battery}}{((E_t + E_{rx_1} + E_{rx_2}) \times n_t) + E_s} \tag{9}$$

where:

- $E_{battery}$ is the total energy available from the battery
- E_t is the energy consumed per transmission
- E_{rx_1} is the energy cost for opening the first Rx window
- E_{rx_2} is the energy cost of the second Rx window
- n_t is the number of transmissions in a day
- E_s is the energy cost for sleep mode

C. Assumptions

We performed our design and experiments with the following assumptions.

1) In North America, LoRaWAN operates in the 902-928 MHz frequency band. The primary data rates for the region are listed in Table I. Our measurements are performed in this frequency band only.
2) It is assumed that the transceiver remains in sleep mode, only waking up to transmit packets and open the receive windows. The sleep time of the device per day was calculated as

$$T_{sleep} = T_{total} - (T_{tx} + T_{rx_1} + T_{rx_2}) \times n_t \tag{10}$$

   where T_{total} is the total time in a day, T_{tx}, T_{rx_1} and T_{rx_2} are the time taken to transmit and the time taken to open the two receive windows respectively. The SX1276 data sheet states the typical sleep current value for the transceiver as 1.5 μA. We base our sleep mode energy calculations on this value.
3) We assume that the total battery capacity of the end-device to be 2500 mAh, the typical battery capacity for two AA cells. Due to environmental and chemical factors, batteries degrade over time, losing their capacity. However, in our paper, we do not consider this aspect. Out of the 2500 mAh capacity, only 1000 mAh would be available for communication, the remaining being used for sensing and computation.
4) In LoRaWAN, there is an adaptive data rate (ADR) feature available in which the network manager can control the data rates for the end-devices. The nodes which are nearer to the gateway will use higher data rates and lower output powers, while the devices which are farther off will use lower data rates and higher output power. For the purpose of energy analysis, we disable the ADR feature. But it should be noted that this feature can further help to save power and reduce congestion on the network by distributing available channels among the devices.
5) A network trial conducted by Semtech [12] for their transceiver found a packet error rate of 10% for DR 1 at a distance of 1350 meters. The transmit power was set to +13dB. Our lifetime calculations are done for DR 1 and a transmit power of +14dB, we assume a similar worst-case packet error rate of 10%. This means that 90% and 99% of the packets would be successfully transmitted within the first and second tries respectively. The expected error in our calculations should then be within 1%.

IV. Experimental Results

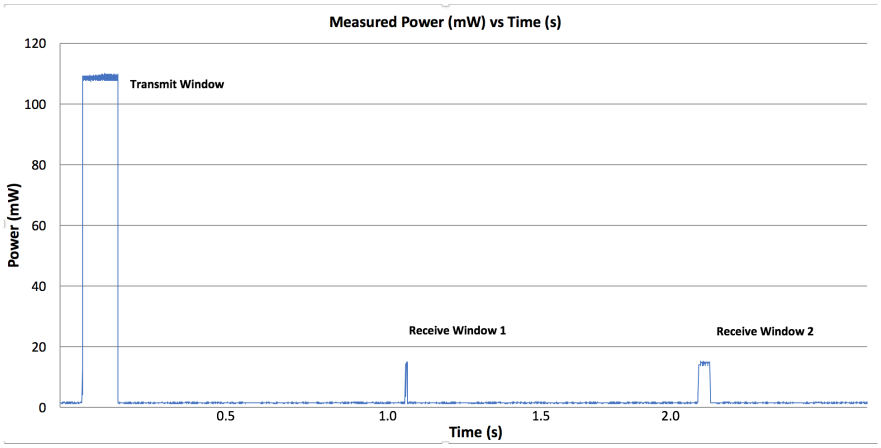

Fig. 5: Transmit and receive windows for DR 2, Payload = 10 bytes

The measured power for a transmission cycle is shown in Figure 5. A large transmit window is observed, followed by two short receive windows. This result follows the profile for normal Class A functioning of the transceiver. The first receive window uses the same data rate as the transmitted packet. The second receive window uses a fixed data rate, DR 1 in our case. The first window is opened one second after a packet is transmitted, and the second window is opened one second after the first one is closed.

TABLE II: Energy consumption per transmission for data rates 0 to 4 and payloads 10 to 50 bytes

Payload (bytes)	*Energy cost per transmission* (mJ)				
	DR 0	*DR 1*	*DR 2*	*DR 3*	*DR 4*
10	42.00	23.23	11.70	6.45	2.97
20	51.21	27.84	15.18	8.25	3.73
30	93.21	32.59	17.36	9.98	4.40
40	102.42	36.93	20.83	11.62	5.34
50	144.42	42.02	23.15	13.35	5.90

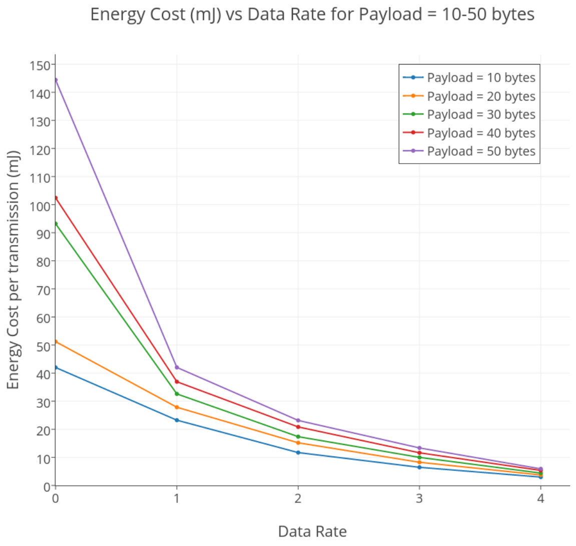

Fig. 6: Energy consumption per transmission for data rates 0 to 4 and payloads 10 to 50 bytes

Table II and Figure 6 show the energy consumption per transmission for data rates of 0 to 4 and payload sizes of 10 to 50 bytes. It can be inferred from the figure that the energy consumption decreases as the data rate increases. For payload sizes, an opposite trend is observed, as higher payload sizes lead to an increase in the energy consumption. This result aligns itself with our theoretical predictions in the previous section.

Table III and Figure 7 show the energy-per useful-bit per transmission. An intriguing observation here is that the energy-per-useful-bit decreases as the payload size increases. It validates our assumption that it would indeed be useful to aggregate data from multiple time-intervals before transmitting it to the gateways. Another unusual observation here is the difference in the DR 0 curve with respect to the curves for other data rates. This is because

TABLE III: Energy per useful bit per transmission for data rates 0 to 4 and payloads 10 to 50 bytes

Payload (bytes)	*Energy per useful bit* (mJ/bit)				
	DR 0	*DR 1*	*DR 2*	*DR 3*	*DR 4*
10	0.525	0.290	0.146	0.080	0.037
20	0.320	0.174	0.094	0.051	0.023
30	0.388	0.135	0.072	0.041	0.018
40	0.320	0.115	0.065	0.036	0.016
50	0.361	0.105	0.057	0.033	0.014

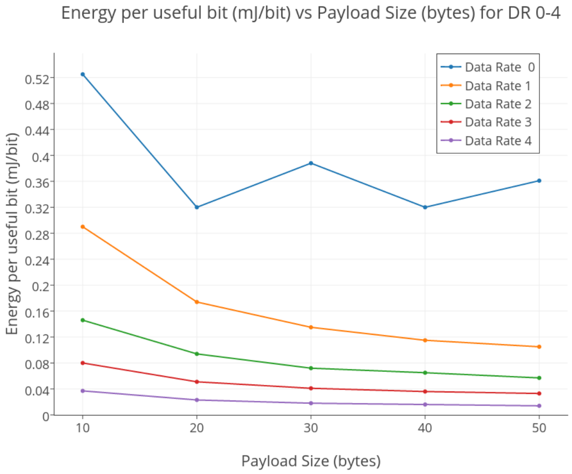

Fig. 7: Energy consumed per useful bit for data rates 0 to 4 and payloads 10 to 50 bytes

for DR 0, the maximum packet size in one transmission is limited to just 19 bytes. To send 40-bytes of data, we would require to transmit two 19-byte and one 2-byte packet.

V. Analysis

A. Lifetime estimates for LoRaWAN devices

Based on the results obtained in the previous section, we can determine that DR 0 is not suited for data aggregation. Therefore, we select the next worst data rate in terms of energy expenditure, which is DR 1 (SF 9, bandwidth 125kHz). We have already detailed the results of the transmission energy cost E_{tx} in the previous section. We similarly calculated the energy per transmission of the two receive windows E_{rx_1} and E_{rx_2}. For DR 1, we got the average values of E_{rx_1} = 0.52 mJ, and E_{rx_2} = 0.60 mJ. The sleep energy E_s was calculated on the basis of the load current value of 1.5μA. Using Equation 9, we calculated the estimated lifetime of the transceiver based on the data-collection time-intervals and number of packets aggregated. It is to be noted that this lifetime estimate is only for the transceiver, and does not include the end-device energy cost due to computation and sensing.

TABLE IV: Estimated lifetime of the LoRa transceiver, time interval 1 to 6 mins, aggregation 1 to 5 packets, and DR = 1

Interval (min)	*Estimated lifetime of device (days)*				
	Agg=1	*Agg=2*	*Agg=3*	*Agg=4*	*Agg=5*
1	306	513	659	777	856
2	609	1017	1304	1534	1687
3	909	1512	1934	2271	2494
4	1205	1999	2550	2989	3278
5	1499	2478	3154	3690	4041
6	1790	2948	3743	4370	4781

Table IV shows the estimated lifetime for the LoRa transceiver given a 1000 mAh battery. It is plotted as a function of time-period of data collection in Figure 8. We can immediately infer that the battery life of the devices is better when packets are aggregated together. The estimated lifetime varies from 306 days in the case of data-collection over one minute time-intervals with no aggregation of multiple intervals to approximately 4800 days in the case of data-collection over six minute-intervals and data aggregation over five time-periods. The lifetime is found to be over five years when data is collected in 3-minute intervals and at least three consecutive packets are aggregated. With 6-minute intervals, an estimated lifetime of over five years can be achieved with no aggregation at all.

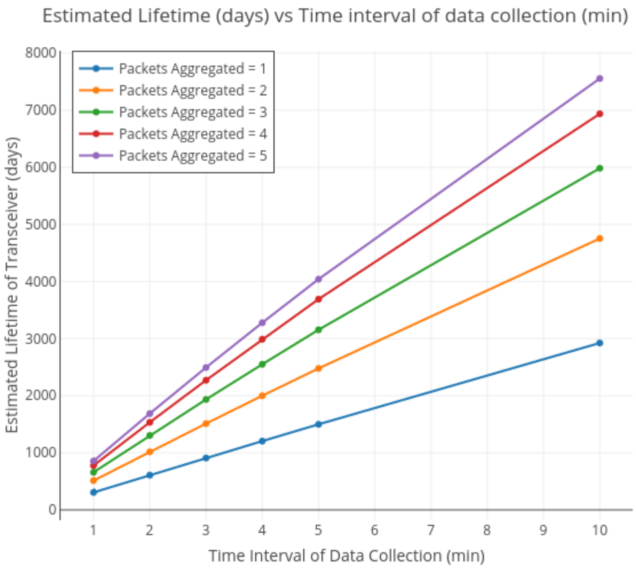

Fig. 8: Lifetime of the device vs time period of collection for aggregation sizes 1 to 5

The trade-off for better energy cost for the device is the loss of real-time data. This trade-off can be avoided to some extent by using variable time-intervals through the day. Traffic sensing could be done for shorter time-intervals, such as two minutes, during rush hour. There would be no aggregation during this time as it will allow the cities to quickly react to traffic congestion and take countermeasures. The monitoring could be done at 6-minute intervals for the rest of the day, with the data being aggregated for every hour. As there is least traffic at night, the sensors could use the largest time period and aggregation sizes to collect data. This can be achieved if we can send weekly or daily schedules to the sensors over LoRaWAN.

B. LoRaWAN range, capacity and performance considerations

The city of Palo Alto has a total area of 67 km^2 [13] and a road network of approximately 320 km [14]. Assuming that we need to monitor traffic per every half mile of road, with sensors in each road direction and lane, we would require sixteen hundred sensors for covering the entire road network of the city. Past work done for the range of a LoRaWAN gateways has found that a single gateway can support devices over a range of 1.2 km for the highest spreading factor with up to +14dB of margin [7]. We also conducted range tests in the city of Palo Alto and found most areas within a mile, or 1.6 km to have adequate coverage [1]. A radial range of 1.2 km per gateway would allow overlapping coverages, so that the gateways could be used to triangulate sensor location using time difference-of-arrival techniques [15], without the end-devices having to expend any energy on localization.

We would thus require around fifteen LoRaWAN gateways to be installed in the city for complete coverage, with each gateway handling about 120 devices. Assuming the worst-case scenario of all end-devices collecting data for one minute time periods with no aggregation, each end-device would be transmitting 10 bytes/min, or 1.33 bits/sec. The gateways would therefore be handling 160 bits/sec of data. The theoretical maximum capacity for DR 1 is 1757 bits/sec [12]. The channel load would therefore be under 10%.

A vital aspect for wireless systems is to consider the placement of sensors. Embedding our sensor device and transceiver module into the physical constraints of Stimsonite markers imposes new challenges on the antenna structure for communication. The pavement itself, a part of the ground plane has a detrimental effect on the directivity of the antenna, increasing the take-off angle. This problem, however, is beyond the scope of this paper and is analyzed in detail in another companion paper [16].

VI. Conclusion

In this paper, we analyzed the usefulness of LoRaWAN as a minimalistic network for crowdsourced traffic calming. We found its implementation to be open and customizable. We described an adaptive algorithm that could be used to transmit sensor data collected over user-defined time-intervals. For better energy-efficiency and preventing data loss, we added support for packet aggregation and packet retransmission. From the experimental results, we observed that lower data rates lead to higher energy expenditure for the sensor devices. However, even with a data rate of 1, we were able to achieve an estimated lifetime of over five years for 6-minute time intervals with a modest 1000 mAh battery. By analyzing the energy-per-useful-bit for transmission of packets of varying payloads, we validated our prediction of better energy efficiency in aggregating information from multiple time-intervals. With these results, we can conclude that LoRaWAN is indeed a rational choice for establishing a crowdsourced traffic calming network.

Such a network will empower cities to monitor traffic in almost real-time, instead of relying on passive information like they do today. It can open up the opportunity of smarter traffic management solutions for cities, with actuators such as signaling also being connected to the network. With the possibility of customization of parameters such as receiver sensitivity, data rates and transmit power, LoRaWAN can also find other applications in the greater realm of smart cities. The rudimentary analysis done for network capacity supports the notion that the network can concurrently support multiple applications. This work can be built up on to estimate the energy costs for communication in those applications. For future work, we would also like to analyze the performance and capacity considerations of LoRaWAN in greater depth.

VII. Acknowledgements

We would like to thank Ervin Teng, Cef Ramirez, João Diogo Falcão and Ritchie Lee for their continued help and support throughout the project. This work was supported in part by the National USDOT University Transportation Center for Safety for which the authors are grateful.

References

[1] B. Iannucci and A. Rowe, “Crowdsourced Smart Cities,” in *ITS World Congress*, Montreal, Quebec, Canada, 2017.
[2] LoRa-Alliance, “LoRaWAN: What is it?” [Online]. Available: https://docs.wixstatic.com/ugd/eccc1a_ed71ea1cd969417493c74e4a13c55685.pdf
[3] B. Moyer, “A New IoT Cellular Network,” White Paper, Electronic Engineering Journal. [Online]. Available: https://www.eejournal.com/article/20141103-sigfox/
[4] Y.-P. E. Wang, X. Lin, A. Adhikary, A. Grovlen, Y. Sui, Y. Blankenship, J. Bergman, and H. S. Razaghi, “A Primer on 3GPP Narrowband Internet of Things,” in *IEEE Communications Magazine*, vol. 55, Campinas, SP, Brazil, 2017, pp. 117–123. [Online]. Available: https://arxiv.org/pdf/1606.04171.pdf
[5] WAVIoT, “What is NB-Fi Protocol.” [Online]. Available: http://waviot.com/technology/what-is-nb-fi.html
[6] H. G. S. Filho, J. P. Filho, and V. L. Moreli, “The adequacy of LoRaWAN on smart grids: A comparison with RF mesh technology,” in *Smart Cities Conference (ISC2), 2016 IEEE International*, Campinas, SP, Brazil, 2016. [Online]. Available: http://ieeexplore.ieee.org/abstract/document/7580783/
[7] M. Centenaro, L. Vangelista, A. Zanella, and M. Zorzi, “Long-range communications in unlicensed bands: the rising stars in the IoT and smart city scenarios,” in *IEEE Wireless Communications*, Italy, 2016. [Online]. Available: https://arxiv.org/abs/1510.00620
[8] LoRa-Alliance, “LoRa Specification 1.0.” [Online]. Available: https://www.lora-alliance.org/portals/0/specs/LoRaWAN%20Specification%201R0.pdf
[9] S. Coleri, S. Y. Cheung, and P. Varaiya, “Sensor Networks for Monitoring Traffic,” in *In Allerton Conference on Communication, Control and Computing*, Berkeley, CA, 2004.
[10] S. Coleri, S. Y. Cheung, P. Varaiya, B. Dundar, S. Ganesh, and C.-W. Tan, “Traffic Measurement and Vehicle Classification with Single Magnetic Sensor,” in *Transportation research record: journal of the transportation research board*, Berkeley, CA, 2005.
[11] M. Kooijman, “Arduino-LMIC Library,” "August" 2016, "Github Repository". [Online]. Available: https://github.com/matthijskooijman/arduino-lmic
[12] Semtech, “LoRa Modulation Basics: AN1200.22,” 2015. [Online]. Available: http://www.semtech.com/images/datasheet/an1200.22.pdf
[13] C. of Palo Alto, “About Palo Alto.” [Online]. Available: http://www.cityofpaloalto.org/visiting/about_palo_alto.asp
[14] E. G. Brown Jr., B. P. Kelly, M. Dougherty, and J. Appleton, “California Public Road Data 2014.” [Online]. Available: http://www.dot.ca.gov/hq/tsip/hpms/hpmslibrary/prd/2014prd/2014PublicRoadData.pdf
[15] Q. Yao, S.-K. Tan, Y. Ge, B.-S. Yeo, and Q. Yin, “An area localization scheme for large wireless sensor networks,” in *Multi-disciplinary Trends in Artificial Intelligence*, Singapore, 2005. [Online]. Available: http://ieeexplore.ieee.org/document/1543864/
[16] A. Tadkase, N. Srinivasan, and B. Iannucci, “Combating Ground Reflection for Wireless Sensors,” in *ITS World Congress*, Montreal, Quebec, Canada, 2017.