LEMoNet: Low Energy Wireless Sensor Network Design for Data Center Monitoring

Chenhe Li
Department of Computing and Software
McMaster University
Hamilton, ON, Canada
email address

Jun Li
Department of Computing and Software
McMaster University
Hamilton, ON, Canada
email address

Mehdi Jafarizadeh
Department of Computing and Software
McMaster University
Hamilton, ON, Canada
email address

Ghada Badawy
Computing Infrastructure Research Center (CIRC)
McMaster Innovation Center
Hamilton, ON, Canada
email address

Rong Zheng
Department of Computing and Software
McMaster University
Hamilton, ON, Canada
email address

***Abstract*—Today's data centers (DCs) consume up to 3% of the energy produced worldwide, much of which is wasted due to over-cooling and under utilization of IT equipment. This wastage in part stems from the lack of real-time visibility of fine-grained thermal distribution in DCs. Wireless sensing is an ideal candidate for DC monitoring as it is cost-effective, facility-friendly, and can be easily re-purposed. In this paper, we develop LEMoNet, a novel low-energy battery operated wireless sensor network design for monitoring DCs. It employs a two-tier network architecture and a multi-mode data exchange protocol to balance the trade-offs between low power consumption and high data reliability. We have evaluated the performance of LEMoNet by deploying custom-designed sensor and gateway nodes in a production DC as well as through extensive simulation studies in networks of various sizes. We show experimentally that LEMoNet achieves an average data yield over 98% in the production DC. It scales well in large and dense networks in large-scale simulations. Under normal operations with one temperature and one humidity reading every thirty seconds, the battery lifetime of LEMoNet sensor nodes is projected to be 14.9 years on a single lithium coin battery.**

***Index Terms*—Data center monitoring, Wireless sensor networks, Bluetooth low energy**

I. Introduction

The computing needs of our society, the majority of which are fulfilled in data centers (DCs), double every five years. Despite increasing productivity, DCs are burdened with high capital and operating costs. These costs arise from widespread industry practice of over-design and over-provisioning as the server utilization is typically between 12% – 15% for on-premises DCs and 65% for cloud DCs [1]. The over-design stems from the fact that DC owners and operators have no real time visibility of the performance and operational conditions of their DCs, thus becoming overly cautious. For instance, high capacity cooling systems are commonplace to mitigate occasional "hot-spots" during peak demand periods but enhancing air-circulation can fix this problem at a fraction of the cost [2]. Traditionally, DC monitoring systems measure only thermal metrics, such as the Information Technology Equipment (ITE) rack inlet temperature and the flow rate of air through Computer Room Air-Conditioning (CRAC) units. These variables are measured with on-board sensors mounted at key locations. Alone, the sensors do not provide sufficient spatial resolution to fully describe the thermal performance of a DC. This may lead to early equipment failure, low energy efficiency and degraded server performance [3]. Moreover, in co-location DCs, where equipment, rack space, and network bandwidth are leased to customers. DC operators have no access to on-board sensors in customers' ITE and are not allowed to connect any monitoring device to them directly.

Wireless sensing is an attractive and cost-effective approach to infrastructure monitoring. It increases the flexibility in sensor placement and allows seamless, automated system reconfiguration. As a result, Wireless Sensor Networks (WSNs) have been investigated as a low-cost solution for finer-grained thermal mapping and thermal forecast in DCs [4]–[7]. However, none of the previous work that utilizes battery-powered sensors was able to achieve a lifetime of more than a couple of months which hinders wider adoption of WSN-based monitoring solutions in DCs. Furthermore, these research prototypes generally rely on special-purpose protocols and hardware, which are often incompatible with existing Data Center Infrastructure Management (DCIM) solutions.

To achieve long battery life and high scalability, we argue for rethinking WSN design to support low duty cycle environment monitoring in DCs. In particular, we recognize that in traditional multi-hop WSNs, significant energy is consumed as the result of protocol overheads, such as topology formation and maintenance, time synchronization, medium access control, and for reliable data transfer. Much of the overhead is due to the *symmetry* among transmitting and receiving devices.

This research was supported by the collaborative research and development grant *NSERC CRDPJ 500720* from the Natural Science and Engineering Research Council of Canada, and Cinnos Mission Critical Incorporated.

ISBN 978-3-903176-16-4

In contrast, for DC monitoring, heterogeneous devices with different power sources, wireless technologies, and compute and storage capabilities can be utilized.

In this paper, we design LEMoNet, a novel low-energy wireless sensor network design for monitoring co-location DCs. LEMoNet follows a two-tier network architecture, where the bottom tier includes wireless sensor nodes and more powerful gateway nodes connected via Bluetooth low energy (BLE); and the top tier consists of gateways and a controller connected via WiFi. The sensor nodes are battery operated and can be deployed and replaced quickly. Gateways are AC powered and collect readings from multiple sensor nodes. Gateways can be seamlessly added or upgraded without any change on the sensor nodes. BLE is chosen as it offers high data rates (compared to Zigbee) and low power consumption for short range communications. More importantly, BLE supports broadcast communication without connection setup. The gateways and the controller communicate using WiFi since each gateway needs to transfer readings from multiple sensor nodes to the controller at a higher data rate. To reduce power consumption, normal sensing data is sent in a best effort manner using broadcast and sensors can be immediately put to sleep upon transmission of data. High reliability is achieved by aggregating packets received by multiple gateways at the controller. This effectively shifts the complexity from sensors to the controller.

LEMoNet utilizes the Bluetooth v5.0 compliant protocol stack, WiFi and MQTT, a lightweight machine-to-machine publish/subscribe messaging protocol commonly adopted in the Internet-of-Things (IoTs). Therefore, it can be implemented on commercial-of-the-shelf (COTS) devices with only software modifications. However, to minimize the power consumption of sensor nodes in different operation modes, we design our own sensor device with low-power sensing modules. As a result, under normal operations, with a single temperature and humidity reading every thirty seconds, the battery lifetime of LEMoNet sensor nodes is projected to be *14.9 years using a single lithium coin battery*.

We have evaluated the proposed system in a production DC containing 56 server racks with up to 42 servers each. A total of 58 sensors are deployed next to the front and back sides of 7 racks. Moreover, to evaluate the scalability of the protocol, we have conducted extensive simulations in dense and large network settings. Experiment results show that LEMoNet can achieve an average Packet Reception Rate (PRR) of over 98% at very low energy cost.

The rest of the paper is organized as follows. In Section II, an overview of the BLE protocol is provided and related work on DC monitoring is reviewed. An overview of LEMoNet design and components is presented in Section III. Details of LEMoNet protocol are described in Section IV. In Section V, we discuss the hardware and simulation implementation followed by performance evaluation in Section VI. Finally, we conclude the paper in Section VII.

TABLE I: Advertising Packet Types

Advertising PDU Type	**Max Data length**	**Allow scan**	**Allow conn.**
ADV_IND	31 bytes	YES	YES
ADV_DIRECT_IND	N/A	NO	YES
ADV_NONCONN_IND	31 bytes	NO	NO
ADV_SCAN_IND	31 bytes	YES	NO

II. Background and Related Work

In this section, we first give an overview of the BLE protocol since the information is crucial in understanding LEMoNet design; then we provide a review of existing research on WSNs for DC monitoring.

A. BLE background

BLE devices can operate in four roles: A *peripheral device* is an advertiser, which is connectable and can operate as a slave in a connection. A *central device* scans for advertisers and can initiate connections. It operates as a master in one or more connections. A *broadcaster* is a non-connectable advertiser. Finally, an *observer* scans for advertising packets but cannot initiate connections. BLE advertising devices do not utilize carrier sensing multiple access, and thus advertising packets are transmitted immediately when data become available. A peripheral device or a broadcaster can broadcast payloads up to 31 bytes using the advertising channels.

There are four types of advertising packets, shown in Table I. A peripheral device sends connectable undirected advertising packets (ADV_IND) to request a connection to a central device. Non-connectable devices advertise non-connectable undirected packets (ADV_NONCONN_IND) to scanning devices. The scannable undirected advertising packet (ADV_SCAN_IND) allows additional information exchange via scan responses. Scanners can be either active or passive. An active scanner can send scan requests to an advertising device to request additional information. A passive scanner, on the other hand, can only receive data from an advertising device. BLE link layer uses a white list for device filtering. The white list contains a set of device MAC addresses for filtering so that only scans and connection requests from devices on the list are processed.

The BLE standard supports star topologies and more recently mesh topologies [8]. Some BLE devices support multiple roles simultaneously. In LEMoNet, only the star topology is used. Sensor nodes are peripheral devices and gateways are central devices.

B. WSNs for DC monitoring

WSN technologies have been investigated as a low-cost candidate solution to finer-grained thermal mapping and accurate hot-spot detection in DCs since they require no additional network and facility infrastructure in an already complicated IT environment. It was reported in [4] that energy saving measures adopted utilizing information collected through a 588-node WSN deployed in a 1,200m^2 area, result in an overall energy reduction of 17% and reduced Power Usage

Efficiency (PUE) from 1.94 to 1.51. The payback time of deploying such a WSN is 3.4 years. In [5], researchers at Microsoft Research designed RACNet, a large-scale sensor network for high-fidelity DC environmental monitoring. RACNet consists of wireless master nodes powered by USB ports on servers or wall power. The wireless master node and several wired sensors form a daisy chain to cover one side of a rack, collecting data at different heights. The master nodes are equipped with IEEE 802.15.4 radios to form a mesh work for data collection. To improve data yield, a Wireless Reliable Acquisition Protocol (WRAP) was designed that uses token passing for network-wide arbitration. Daisy chaining multiple sensors to a master node allows powering the sensors and ensures reliable gathering of sensor data. However, wired connection is restrictive for sensor placement, and placing extra wires along the rack is generally considered undesirable by DC operators.

Rodriguez et al. [6] observed significant temperature variation in different locations in DC using a small-size network consisting of 10 nodes. In [9], Chen et al. deployed temperature and airflow sensors to monitor inlet and outlet server temperatures and CRAC units. All sensors run TinyOS and form a single-hop network using 802.15.4. It has been demonstrated that using the real-time sensor data to calibrate a computational fluid dynamic model, temperature evolution of servers with highly dynamic workloads can be forecast at an average error of 0.52^{o}C, within a duration up to 10 minutes. Cluster area sensor network (CASN) [7] comprises a number of TelosB sensor nodes running TinyOS attached to compute servers or workstations. It verifies servers physical presence through wireless cluster-wide command dissemination, and thus enhance the security of datacenter management. CapNet [10] is a real-time WSN protocol for power capping for data centre management. It uses distributed event detection to eliminate the overhead of regularly polling all nodes in the network. In CapNet, sensor nodes use a single IEEE 802.15.4 channel for communication inside a cluster, where the transmission schedule is slotted and coordinated by a power-capping manager.

All of the above work uses IEEE 802.15.4 in RF transceivers. They primarily aim to validate the utility of WSNs in DCs and/or to achieve high data yield. Battery lifetime of the sensor devices, costs and compatibility with existing DCIM solutions are generally not considered.

III. OVERVIEW

We design and develop LEMoNet, a low energy wireless monitoring Network for DCs. LEMoNet is energy efficient, easy to install and manage within existing IT infrastructure while delivering better visibility into a DC's thermal environment. It also supports real-time alerts in event of outages such as overheating in DCs.

A. Design Considerations

In LEMoNet, there are two key design considerations: 1) what are the suitable radio technologies, and 2) what duty cycle should be supported for DC monitoring. To answer the first question, we compare the power consumption of low-power RF modules commonly used for IoT applications in Table II.

TABLE II: Specification of Popular RF Modules

	ESP8266	**TI cc2520**	**TI cc2640r2f**
Standard	WiFi 802.11b	Zigbee	BLE
Data rate	11 Mbps	250 kbps	1, 2 Mbps
TX power	17 dBm	5 dBm	5 dBm
Sensitivity	-91 dBm	-98 dBm	-97 dBm
Link budget	108 dBm	103 dBm	102 dBm
TX current	170 mA	33.6 mA	9.1 mA
RX current	56 mA	18.5 mA	5.9 mA

From Table II, we see that compared to Zigbee, BLE modules achieve lower power consumption, higher data rate and comparable link budgets. Similar observations were made in [11], where the authors compared the usage of four wireless technologies (WiFi, classical Bluetooth, Zigbee and BLE) for a construction noise monitoring and concluded that BLE is suitable for high-rate sensing. Furthermore, for reasons that will become clearer later, BLE is advantageous as it natively supports connectionless broadcasting data transmissions and has better co-existence with WiFi than Zigbee [12]. Therefore, we adopt BLE for LEMoNet sensor nodes.

To answer the second question, we conduct a measurement study by instrumenting server racks with temperature sensors in a raised-floor high performance DC and a modular enclosed DC with two in-row cooling units. The two DCs are chosen as they represent very different thermal environments and IT workloads. Figure 1 show the temperatures reported by sensors at the top, middle and bottom of one server rack in the cold aisle[1]. In Figure 1a, the cooling units failed around time 18:35 whereas in Figure 1b, one cooling unit was shut down by the researchers around time 18:44. Some server temperature rose up to 35^{o}C, which is beyond the acceptable operating temperature range from 18^{o}C to 27^{o}C recommended by the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) [13]. However, in both cases, it took 2 – 5 minutes to reach beyond 27^{o}C. Prior to the cooling outages, the temperatures remain stable for a long period of time. Therefore, collecting temperature sensor readings at the intervals of 30s would be sufficient for hot-spot detection as well as characterizing temporal fluctuation as the result of interplay between cooling and IT workloads.

B. System Architecture

LEMoNet consists of sensor nodes, gateways and a central controller organized in a two-tier network architecture (Fig. 2). The first tier is between sensor nodes and gateways; and the second tier is between gateways and a central controller. Sensor nodes are battery powered and operate at low duty cycles.

[1]In DCs with air cooling, server racks are lined up in alternating rows cold air intakes facing one way and hot air exhausts facing the other. The rows composed of rack fronts are called cold aisles. Typically, cold aisles face air conditioner output ducts. The rows the heated exhausts pour into are called hot aisles.

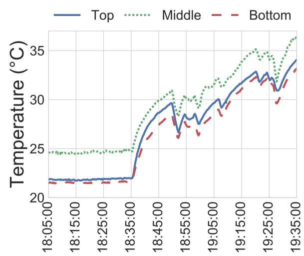

(a) Rack in a Raised Floor DC. The set point was 22^oC. The servers are at high load.

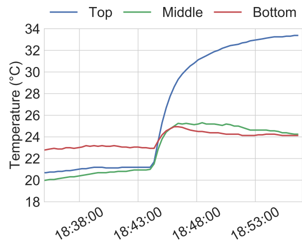

(b) Rack in a Modular DC. The set point was 20^oC. The utilization of all servers are 70%.

Fig. 1: Inlet temperature before and after a CRAC unit outage

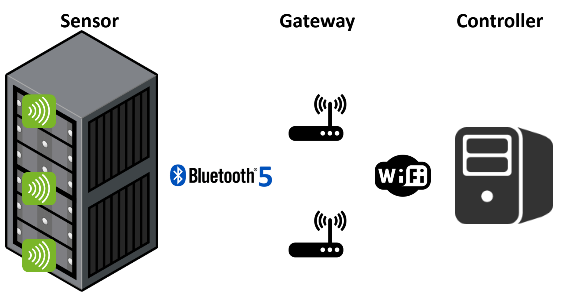

Fig. 2: LEMoNet system architecture

They measure and report DC's environmental conditions (such as temperature, humidity, pressure and air flow) using BLE. Gateways are wall powered and are responsible for collecting data from sensor nodes and forwarding them to a central controller using WiFi. There is no association between sensor nodes and gateways. Any gateway that receives sensor data can forward it to the controller. Multi-packet reception at gateways offers opportunities to improve the reliability of data collection at no extra cost. This architecture makes the network scalable with little management overhead since adding a new sensor node only involves deciding its location and installing it.

IV. The LEMoNet Protocol

LEMoNet builds upon the BLE stack and leverages existing message types in BLE.

A. Sensor to/from Gateway

As illustrated in Figure 1, thermal conditions in DCs tend to change slowly. A majority of sensor data are non-mission critical and do not require high reliability. Broadcast connectionless communication in BLE is thus a good fit for collecting sensor data where a sensor node does not need to be paired to a specific gateway. There are however two scenarios where reliability is highly desirable. First, when sensors detect outage events such as overheating, alerts should be generated as soon as possible. Second, updating configuration parameters and software images on sensor nodes require reliable data transfer. To support the former, we leverage a particular message type in BLE to provide acknowledgment to sensor data without device pairing. For the latter, connection oriented communication in BLE is utilized.

To this end, LEMoNet sensors can operate in three modes, namely, reporting non-urgent data from sensors to a gateway in the normal connectionless (NCL) mode, reporting urgent data from sensors to a gateway in the scannable connectionless (SCL) mode, and the reliable connection (RC) mode for exchange of bulk data. In this study, transmission of the urgent data in SCL model is triggered when temperature measurements are above the overheating threshold (In our case, 27 ^oC in cold aisle). Otherwise, the sensors stay in the NCL mode. The RC mode is activated for parameter updates or over-the-air software update.

NCL Mode: In this mode, a sensor sends ADV_NONCONN_IND advertising packets (Service Type 0) once containing non-urgent sensing data and immediately goes to sleep (Fig. 4a). When a gateway receives the advertising packet, it puts the packet into a non-urgent message queue. Gateways publish non-urgent messages every 1 seconds. Due to the broadcast nature of advertising packets, multiple gateways can receive it. By combining multi-packet receptions at the gateways, the controller can tolerate packet losses at individual gateways.

SCL Mode: A sensor node enters the SCL mode if at least one of its sensor readings falls outside a pre-defined range. This usually indicates abnormal thermal conditions[2]. In the SCL mode, a sensor indicates in its ADV_SCAN_IND advertising packet that the data is urgent (with Service Type field set to 1). Upon reception of the advertising packet, a gateway whose whitelist contains this sensor node sends a scan response message acting as an acknowledgment (Fig. 4b). The gateway then publishes the urgent data immediately to the central controller via MQTT. A sensor node retransmits the data up to 10 times if it fails to receive the scan response message. Upon receiving an acknowledgment or exceeding maximum retransmission attempts, the sensor node goes to sleep.

RC Mode: The RC mode is used for down-link communication from gateways to sensors for commands, configuration parameters and software updates (Fig. 4c). In LEMoNet, all communications are initiated by sensor nodes, and thus such operations can only be done asynchronously. Sensor nodes periodically (e.g., every hour) transmit an ADV_SCAN_IND advertising packet with Service Type set to 2. If an update is needed, the gateway whose white list contains the sensor node responds with a connection request message to pair with the sensor. After pairing is successful, the gateway sends relevant data to the sensor via a data channel. Upon completion of data transfer, the gateway terminates the pairing. The sensor acknowledges the termination and goes to sleep.

Lastly, all packets from sensor nodes are timestamped at the gateways, which are synchronized using the network time protocol (NTP). In a local area network, NTP can achieve time synchronization accuracy on the order of tens of ms, which is sufficient for thermal profiling of DCs.

[2]In this study, we do not consider abnormal readings due to sensor failure.

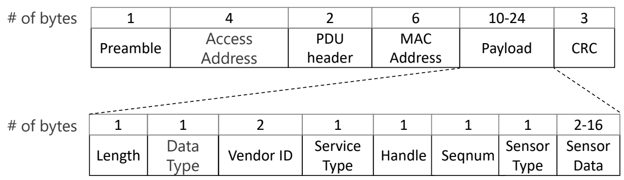

Fig. 3: LEMoNet Advertising Packet Format

B. Gateway to Controller

Gateways communicate with a central controller using WiFi channels 1, 6 and 11. This minimizes overlapping between WiFi and BLE advertising channels. For situations where WiFi coverage is insufficient, a multi-hop mesh network among the gateways can be employed. The MQTT protocol [14], a lightweight machine-to-machine publish/subscribe messaging protocol, is employed at the application layer. In MQTT, gateways publish non-urgent data periodically and urgent data immediately. The central controller subscribes to sensor reading topics. Merging multiple non-urgent data and sending them once reduce the number of packets between gateways and the controller, and limit interference with BLE traffic.

C. Multi-packet reception

In the NCL mode, an advertising packet that contains non-urgent sensor data may be lost. Failures of packet receptions arise from two sources: packet corruption (received with error) and packet loss (packet not received). Traditionally, to improve the reliability of data transfer, two mechanisms can be utilized, namely, packet retransmission and forward error correction (FEC). Both mechanisms incur additional power consumption on the sensors because of either extra transmissions or extra bits in the packets. In LEMoNet, since broadcast sensing data can be received at any gateway in the vicinity, the controller can combine the packets from the gateways and remove duplicated ones using a sequence number field in sensor messages. We have also experimented with combining packets with corrupted bits that fail CRC. However, evaluation studies show the improvement is marginal and thus the mechanism is omitted from this work.

D. LEMoNet Messages

LEMoNet uses standard BLE messages for scan request/response, connection request/response, etc. To support the three service types, we customize the payload of advertising packets to include additional information. The format of advertising packets is shown in Fig. 3. More specifically, in the payload field,

- **Vendor ID** distinguishes LEMoNet nodes from other BLE devices.
- **Service Type** specified the service type (0 for non-urgent sensor data, 1 for urgent sensor data, 2 for commands and software updates).
- The parameter characteristic value **Handle** allows a gateway to directly update sensor parameters after pairing with the sensor node.

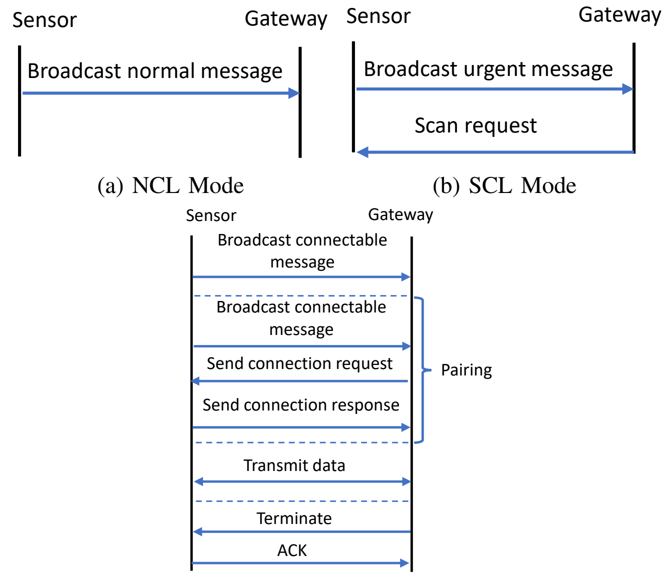

(a) NCL Mode (b) SCL Mode

(c) RC Mode

Fig. 4: Data Transmission Timeline

- **Sequence number** tracks unique sensor data transmitted from each device.
- **Sensor Type** is of one byte. Each set bit indicates that the packet contains sensor data of the respective type.
- **Sensor Data** contains sensor reading. Each reading is of two-byte long.

V. Implementation

In this section, we present the implementation of LEMoNet gateway and sensor hardware for testbed evaluation as well simulated modules for large-scale simulations.

A. Hardware

In choosing hardware components, we need to consider power consumption, response time and resolution requirements following industrial standards and best practices. For example, in DC monitoring, the accuracy of temperature measurement should be at least ± 0.33 °C, and pressure sensors are required to distinguish differential pressures in the range from -2.491 hPa to 4.982 hPa [15]. Consequently, integrated digital sensor chips are selected to build LEMoNet sensor nodes for their ultra-low power consumption, precise measurement results and ease of use. The LEMoNet BLE shield for gateways and sensor nodes are shown in Fig. 5.

1) Sensor Node: The core of LEMoNet sensor node is a HY-40R201PC module [16] featuring a TI CC2640r2f wireless micro-controller. It is compatible with the BLE v5.0 specification. A sensor node contains an SHT-31-DIS temperature and humidity sensor [17] and an LPS25HB pressure sensor [18].

2) Gateway's BLE Shield: For fast prototyping, we develop a BLE shield using an HY-40R201PC module that can be attached to Raspberry Pi 3 (RPI3), a single board computer that supports Linux systems. Existing BLE module on RPI3's can only receive 90% of advertising packets with the BlueZ library [19]. One possible reason is that the WiFi and BLE

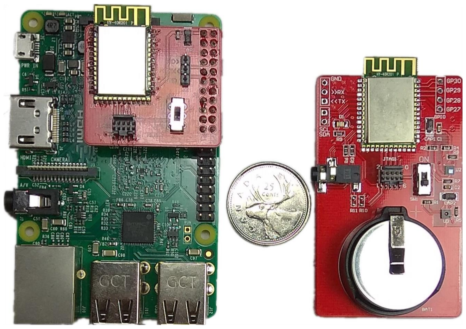

Fig. 5: LEMoNet BLE shield and Sensor

TABLE III: Key Simulation Parameters

Parameter	Value
BER Model [22]	$P_e = \frac{N-1}{N} P_{e(N-1)} + \frac{1}{N} P_{e1}$ ψ = 20ns
AdvInterval	0.1s
ScanInterval	2.5s
Tx Power	0dBm
Channel Model	Path loss exponent α = 4.7 Log normal $N(0.5, 0.25^2)$
Noise floor	-110dBm

transceivers on RPI3 share the same antenna. Such a low packet reception rate necessitates the design of our own BLE shield. The BLE shield communicates with RPI3 via a UART interface. The whitelist of TI CC2640r2f device can only store up to 16 MAC addresses. We populate the whitelist in sensor nodes with the MAC addresses of the gateway's BLE shields.

B. Simulation Modules

To evaluate the scalability of LEMoNet, we have extended an existing implementation of BLE physical (PHY), MAC and Host Controller Interface (HCI) layers in OMNET++ [20] using the MiXiM framework [21].

The key simulation parameters and the bit-error-rate (BER) model are summarized in Table III. In the BER model, P_e is the bit error rate, P_{e1} is the BER of the first bit, and $P_{e(N-1)}$ is the BER for the rest of the bits in each GFSK hop[3]. Delay spread is one of the effective parameters in the calculation of $P_{e(N-1)}$. According to the recommendations in [24] and our measurement study, it was set to $\psi = 20ns$ for DC environment. More details about the GFSK BER model can be found in [22].

VI. Performance Evaluation

In this section, we present performance evaluation results. We start with profiling the power consumption of LEMoNet sensor nodes in different modes and then evaluate the performance of the LEMoNet protocol using a testbed in a production DC and simulations. In DC monitoring, the number of sensors and their placement are based on engineering experience and budget constraints. It is typical to place three sensors at the top, middle and bottom in the front of each server rack and three in the back. For evaluation, we also consider dense deployment scenarios where every server is instrumented with one sensor in the front and one in the back.

[3]Similar to classic Bluetooth (in its basic rate), BLE also uses GFSK modulation in the physical layer [23].

TABLE IV: Current Usage Profile under Different Working Mode

Mode	MCU	Sensors	RX	TX	Current
Standby	OFF	OFF	OFF	OFF	2 μA
Wake-up	ON	ON	OFF	OFF	3.8 mA
Measuring	OFF	ON	OFF	OFF	1.2 mA
Processing	ON	ON	OFF	OFF	3.8 mA
Advertising	ON	OFF	OFF	ON	8.9 mA
Unicast TX	ON	OFF	OFF	ON	8.9 mA
Receiving	ON	OFF	ON	OFF	5.9 mA

A. Power Consumption Profiling

A Power Monitor FTA22D [25] and a Fluke 87V industrial multimeter are used to profile the power consumption of LEMoNet sensor nodes.

As described in section IV, there are three modes of sensor operations, and each of them has different average power consumption. After waking-up, the node sends commands to SHT-31-DIS and LPS25HB sensors to notify them to start sampling data. Afterwards, the processor enters the standby state for 3ms. Next, the processor reads data from sensors and determines the next state. In NCL mode, data are broadcast to gateways and the sensor node will be put in a low-power mode Measured data transmission time varies slightly in SCL mode and NCL mode depending on how long it takes to receive a SCAN_REQUEST message that confirms the reception of the data message in SCL. In RC mode, the sensor node pairs with a gateway to receive firmware upgrades. The transmission time depends on how long it takes to receive a connection request packet form gateways (i.e. pairing) and the size of the data.

Table IV gives the average current when a sensor node is in different working modes. Assume a sensor is powered by a battery with capacity 1000mAH. We can estimate its battery lifetime when operating in the NCL mode by considering the average current in different states. Fig. 6 summarizes the lifetime under different sampling intervals to collect sensor measurements. For instance, 10s means that a sensor node wakes up every 10 seconds to collect and transmit its sensor data via advertising. When a sensor node operates in the SCL or RC modes, it has to stay for longer periods of time and thus its lifetime will be reduced. However, both are likely to be rare occurrences in DCs.

B. Testbed Evaluation

To evaluate the performance of LEMoNet, We deployed 58 sensor nodes and 3 gateway nodes in a high performance computing DC. Fig. 7 shows the floor plan of the data center. A total of 35 sensor nodes were deployed on the top, middle and bottom of 6 racks (in the first three rows) on both the front

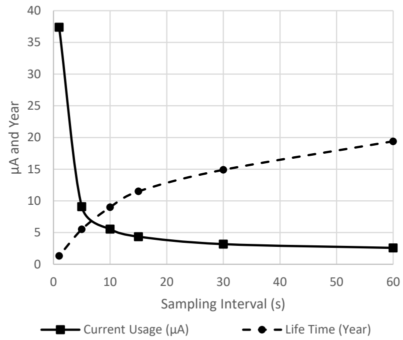

Fig. 6: Current usage in the NCL mode

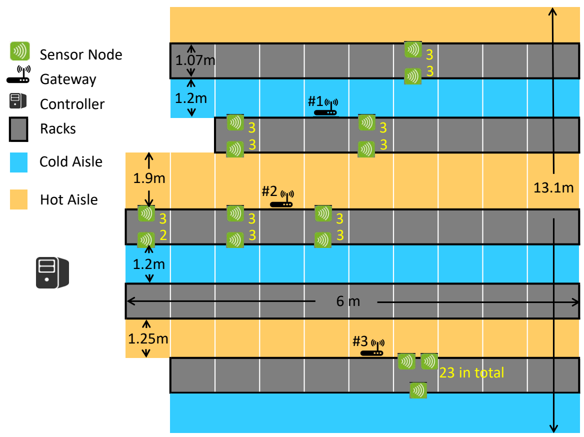

Fig. 7: Placement of 58 sensors and 3 gateways in the production DC

and back sides. The remaining 23 sensors were instrumented on a single rack in the bottom row with 40 servers. Gateways were placed on the top of three racks. We set the transmission power of both sensor nodes and gateways to 5dBm.

1) Packet Reception Rate: In this set of experiments, we evaluate the packet reception rate of the proposed network. The sampling interval is set to 30 seconds. The PRR for all sensor nodes and two sensor nodes are shown in Fig. 8. The two sensor nodes chosen have relatively low received signal strength from all three gateways. The average RSS of Sensor 1 at gateway 1, 2, 3 are respectively, -66.33Bm, -77.7dBm and -71.6dBm; and the average RSS of Sensor 2 at gateway 1, 2, 3 are respectively, -60.1dBm, -90.1dBm, and -74.8dBm. In Fig. 8, in each group, we report the percentage of received packets at each gateway, and at the controller. A couple of observations can be made. First, there exist much diversity in packet reception across different gateways. Although the PRR at individual gateways can be as low as 93.5%, the PRR at the controller is higher than that of the best gateway. This is because the controller by combining correctly received packets from multiple gateways, can attain a very high PRR (at 99.4%).

2) Effect of sampling interval: Next, we investigate the effect of sensor sampling interval. Smaller sampling intervals result in higher traffic loads in the network. Table V reports

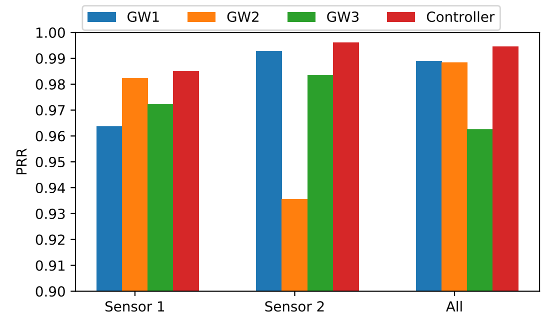

Fig. 8: Packet reception rate. GW1, GW2, GW3 and C are the percentages of received packets at the three gateways and at the controller.

TABLE V: Effects of sampling interval on PRR in the NCL mode

Sampling interval	Min PRR	Mean PRR	Max PRR
10 s	94.1%	98.5%	99.6%
20 s	95.3%	98.6%	99.6%
30 s	97.2%	99.1%	100.0%
60 s	98.4%	99.3%	100.0%

the PRRs at the controller after multi-packet reception. The PRRs are calculated for every 512 packets transmitted by each sensor nodes. As expected, as the sampling interval increases, the traffic load decreases and in general the PRR increases. The average PRR in all cases is above 98.4%. Note this is accomplished at no extra power consumption on the sensor nodes.

3) Number of retransmissions in the SCL mode: In this set of experiments, we fix the sampling interval to be 30 seconds and evaluate the co-existence of NCL and SCL modes. Specifically, three scenarios are evaluated,

- **Normal operation:** The temperature threshold for sensors in cold aisles is set to 24^oC[4]. When the measured temperature exceeds the threshold, sensor nodes switch to the SCL mode to ensure reliable transfer of their measurements.
- **Sensors in cold aisle:** All sensor nodes in cold aisles are in the SCL mode while sensor nodes in hot aisles are in the NCL mode unless an outage occurs. This and the next scenario are used to stress-test the network.
- **All sensors:** Sensor nodes in both hot and cold aisles are in the SCL mode.

Fig. 9 shows the cumulative distribution function of the number of retransmissions a sensor node makes in the SCL mode before receiving a scan request. We find in the first scenario, at most one retransmission is needed to deliver measurement data reliably while in the second and third scenarios, around 90% of data can be delivered reliably with 2 to 3 retransmissions. In the first and second scenarios, the average PRRs for sensor nodes in the NCL mode are 98.6%

[4]High cold aisle temperatures are often indications of potential hot-spots and inefficient cooling.

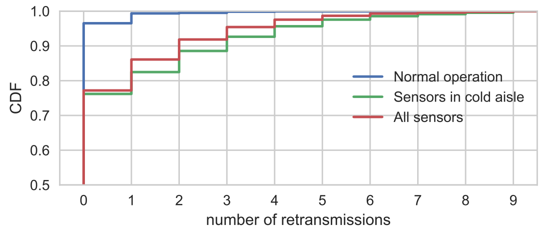

Fig. 9: The number of retransmissions before receiving a scan request

and 99.0%. The PRRs for sensor nodes in the SCL mode are at 100%.

C. Simulation Study

We have conducted simulations to evaluate the effects of key parameters in the LEMoNet protocol and its scalability in dense and large-scale deployments.

1) Effects of sampling interval and the percentage of time in the SCL mode: In the first scenario, a DC of $18m \times 10m$ comprised of 60 racks in 6 rows is monitored by 360 sensors (6 sensors per rack: 3 sensors on the cold aisle, and 3 sensors on the hot aisle). 3 gateways are installed on the top of every other row.

In the first set of experiments, we vary the *sampling interval* among $\{5s, 15s, 30s, 60s, 120s\}$. Sensors enter the SCL mode with the probability of 10%. The left plot in Figure 10(a) shows that except for when *sampling interval* is 5s, the average PRR is 99% or above in the NCL mode. The right plot in Figure 10(a) shows the probabilities of the number of required retransmissions for delivering the packets from sensors in SCL mode correctly. Recall from Section IV, sensors try at most 10 times to transmit a SCL mode packet. It can be seen from Figure 10(a) that this upper-bound is sufficient to avoid any packet loss as in fact at most 5 retries are needed to ensure 100% PRR in the SCL mode. Apart from PRR, we have also assessed the power consumption of the sensor nodes in the first set of experiments. Figure 11(a) depicts how the average power consumption decreases as the sampling interval increases. The non-linearity is consistent with the estimates in Table 6. The power consumption during standby becomes a dominant factor at long sampling intervals.

In the second set of experiments, we vary the SCL mode percentage among $\{1\%, 10\%, 20\%, 50\%, 100\%\}$ with the *sampling interval* set to 15s. The results are shown in Figure 10(b). We see in all the cases, the average PRRs in the NCL mode is above 99%. However, the power consumption slightly increases with higher SCL mode percentages (Figure 11(b)), because of the additional power consumption is needed for retransmissions.

2) Scalability Analysis: Next, we evaluate the scalability of LeMoNet in dense and large-scale deployments. In the third set of experiments, the same DC layout is used but more sensors are deployed per rack, varying in $\{10, 20, 30, 40, 50, 60\}$. A total of 6 gateways are deployed with one gateway per rack.

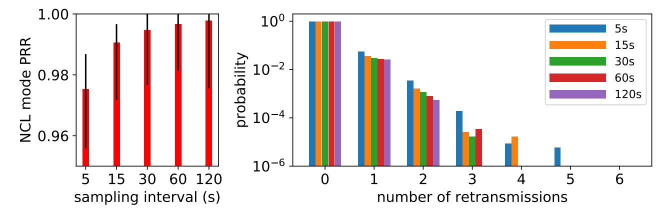

(a) Experiment set 1: Effect of the sampling interval. Sensors per rack = 6 (3 front / 3 back), SCL mode percentage = 10%

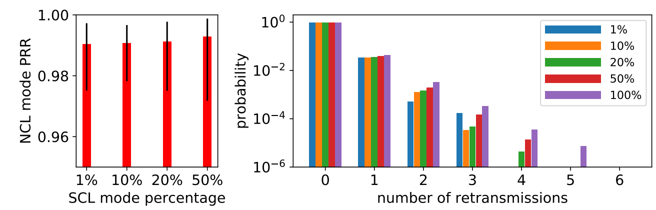

(b) Experiment set 2: Effect of the SCL mode percentage. Sensors per rack = 6 (3 front / 3 back), sampling interval = 15s

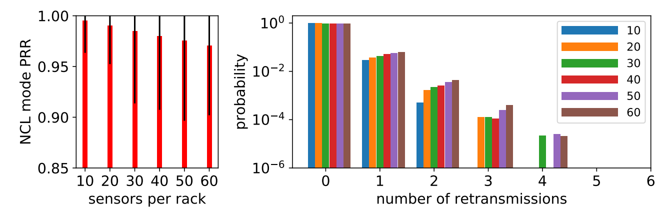

(c) Experiment set 3: Scalability wrt sensor density. SCL mode percentage = 10%, sampling interval = 30s

Fig. 10: Packet reception rates in the NCL mode and the number of retransmissions in the SCL mode

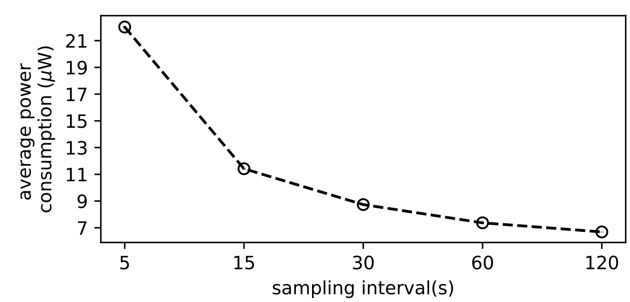

(a) Power consumption in experiment set 1. SCL mode percentage = 10%

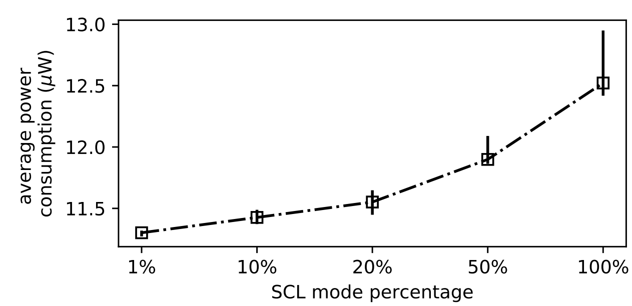

(b) Power consumption in experiment set 2. Sampling intervals = 15s

Fig. 11: Average power consumption of sensors

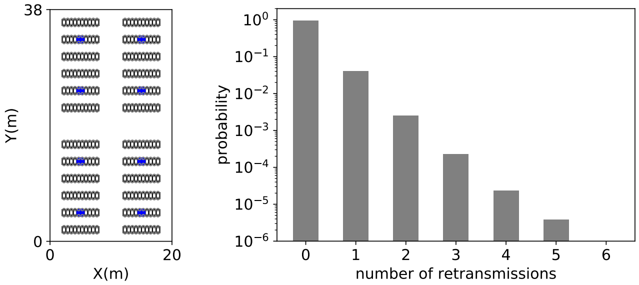

(a) Layout (b) SCL mode retransmissions PDF

Fig. 12: Large scale experiment. SCL mode percentage=10%, sampling interval = 30s

The *SCL mode percentage* and *sampling interval* are set to constant values 10% and 30s, respectively.

The left plot in Figure 10(c) shows that the PRR in the NCL mode slightly drops with the number of sensors approximately linearly. It is the direct effect of the increased number of data packets and subsequently collision probability. However, as shown in the right plot of Figure 10(c), packets in the SCL mode are all delivered reliably.

In the fourth set of experiments, we deploy 1440 sensors in a DC of size $38m \times 20m$ with a total of 12 rows and 240 racks (equivalently, 6 sensors per rack). 8 gateways are deployed as shown Figure 12(a). In the experiments, we set *sampling interval* = 30s and *SCL mode percentage* = 10%. The max, min and average PRR for sensors in the NCL mode are 99.87%, 95.19% and 98.99%, respectively. There is no loss for packets transmitted in the SCL mode and the probabilities of the number of retransmissions are comparable to that in Experiment 1 and 2. From both sets of experiments, we can conclude that LEMoNet scales well to dense and large-scale DCs. However, the number and placement of gateways in these scenarios would affect the PRR of data transmitted in the NCL mode and will be investigated in our future work.

VII. Conclusion

In this paper, we have designed, developed and evaluated LEMoNet for DC monitoring. It employs a two-tier network architecture and a multi-mode data exchange protocol to balance the trade-offs between low power consumption and high data reliability. Experimental and simulation studies demonstrated the high reliability, scalability and power efficiency of LEMoNet.

References

[1] P. Delforge, "Americas data centers are wasting huge amounts of energy," *National Resources Defense Council, vol. Issue Brief*, pp. 14–08, 2014.

[2] R. R. Schmidt, E. Cruz, and M. Iyengar, "Challenges of data center thermal management," *IBM Journal of Research and Development*, vol. 49, no. 4.5, pp. 709–723, 2005.

[3] X. Wang, X. Wang, G. Xing, and C.-X. Lin, "Leveraging thermal dynamics in sensor placement for overheating server component detection," in *Green Computing Conference (IGCC), 2012 International*. IEEE, 2012, pp. 1–10.

[4] R. Mahdavi and W. Tschudi, "Wireless sensor network for improving the energy efficiency of data centers," Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States), Tech. Rep., 2012.

[5] C.-J. M. Liang, J. Liu, L. Luo, A. Terzis, and F. Zhao, "Racnet: a high-fidelity data center sensing network," in *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*. ACM, 2009, pp. 15–28.

[6] M. G. Rodriguez, L. E. O. Uriarte, Y. Jia, K. Yoshii, R. Ross, and P. H. Beckman, "Wireless sensor network for data-center environmental monitoring," in *Sensing Technology (ICST), 2011 Fifth International Conference on*. IEEE, 2011, pp. 533–537.

[7] K. Hong, S. Yang, Z. Ma, and L. Gu, "A synergy of the wireless sensor network and the data center system," in *Mobile Ad-Hoc and Sensor Systems (MASS), 2013 IEEE 10th International Conference on*. IEEE, 2013, pp. 263–271.

[8] M. W. Group, "Bluetooth® mesh profile specification 1.0," 2017.

[9] J. Chen, R. Tan, Y. Wang, G. Xing, X. Wang, X. Wang, B. Punch, and D. Colbry, "A high-fidelity temperature distribution forecasting system for data centers," in *Real-Time Systems Symposium (RTSS), 2012 IEEE 33rd*. IEEE, 2012, pp. 215–224.

[10] A. Saifullah, S. Sankar, J. Liu, C. Lu, R. Chandra, and B. Priyantha, "Capnet: A real-time wireless management network for data center power capping," in *Real-Time Systems Symposium (RTSS), 2014 IEEE*. IEEE, 2014, pp. 334–345.

[11] J. Hughes, J. Yan, and K. Soga, "Development of wireless sensor network using bluetooth low energy (BLE) for construction noise monitoring," *International Journal of Smart Sensing and Intelligent Systems*, vol. 8, no. 2, pp. 1379–1405, 2015.

[12] Y. Wang, Q. Wang, Z. Zeng, G. Zheng, and R. Zheng, "Wicop: Engineering wifi temporal white-spaces for safe operations of wireless body area networks in medical applications," in *2011 IEEE 32nd Real-Time Systems Symposium (RTSS 2011)*. IEEE, 2011, pp. 170–179.

[13] ASHRAE. (2016, Jun.) Data center power equipment thermal guidelines and best practices. [Online]. Available: https://tc0909.ashraetcs.org/..._2016_REVISED.pdf

[14] U. Hunkeler, H. L. Truong, and A. Stanford-Clark, "Mqtt-sa publish/subscribe protocol for wireless sensor networks," in *Communication systems software and middleware and workshops, 2008. comsware 2008. 3rd international conference on*. IEEE, 2008, pp. 791–798.

[15] P. Rod Mahdavi, "Data center energy efficiency measurement assessment kit guide and specification," 2014.

[16] L. SHENGRUN Rechnology Co. (2017, Nov.) Hy-40r201 bluetooth ble 5.0 module specifications. [Online]. Available: http://www.tuner168.com/uploads/soft/171114/3-1G114161337.pdf

[17] Sensirion. (2017) Sht3x-dis datasheet. [Online]. Available: https://www.sensirion.com

[18] STMicroelectronics. (2016) Lps25hb datasheet. [Online]. Available: http://www.st.com

[19] M. Krasnyansky, "Bluez: Official linux bluetooth protocol stack," 2003.

[20] K. Mikhaylov, "Simulation of network-level performance for bluetooth low energy," in *Personal, Indoor, and Mobile Radio Communication (PIMRC), 2014 IEEE 25th Annual International Symposium on*. IEEE, 2014, pp. 1259–1263.

[21] A. Köpke, M. Swigulski, K. Wessel, D. Willkomm, P. Haneveld, T. E. Parker, O. W. Visser, H. S. Lichte, and S. Valentin, "Simulating wireless and mobile networks in omnet++ the mixim vision," in *Proceedings of the 1st international conference on Simulation tools and techniques for communications, networks and systems & workshops*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2008, p. 71.

[22] H. Liu, V. Venkatesan, C. Nilsen, R. Kyker, and M. E. Magana, "Performance of frequency hopped noncoherent gfsk in correlated rayleigh fading channels," in *Communications, 2003. ICC'03. IEEE International Conference on*, vol. 4. IEEE, 2003, pp. 2779–2783.

[23] J. Masuch, M. Delgado-Restituto *et al.*, *Ultra low power transceiver for wireless body area networks (Appendix: Radio Specifications Imposed by the BLE Standard)*. Springer, 2013.

[24] H. Hashemi and D. Tholl, "Analysis of the rms delay spread of indoor radio propagation channels," in *Communications, 1992. ICC'92, Conference record, SUPERCOMM/ICC'92, Discovering a New World of Communications., IEEE International Conference on*. IEEE, 1992, pp. 875–881.

[25] M. S. Inc. (2014, May) Mobile device power monitor manual. [Online]. Available: http://msoon.github.io/.../LVPMManual.pdf