BLITZ: A Network Architecture for Low Latency and Energy-efficient Event-triggered Wireless Communication

Felix Sutton, Reto Da Forno, Jan Beutel, Lothar Thiele
Computer Engineering and Networks Laboratory, ETH Zurich, Switzerland
{firstname.lastname}@tik.ee.ethz.ch

ABSTRACT

We consider wireless sensing systems where an event must be rapidly communicated from its source to a remote host. Due to the non-deterministic nature of event arrivals, multi-hop dissemination using state-of-the-art radio duty-cycled protocols leads to a trade-off between latency and energy efficiency. In order to circumvent this system design constraint, we propose BLITZ, a network architecture that leverages interference-based flooding to rapidly wake-up the multi-hop network and disseminate the event on-demand. We show that by embracing network flooding in combination with simultaneous transmissions, we can realize low latency event-triggered multi-hop communication without having to sacrifice energy efficiency. We introduce an analytical model to quantify the limits of our approach, present a prototype implementation using a multi-radio wireless sensor platform, and experimentally evaluate BLITZ in a laboratory setting and in an indoor testbed.

KEYWORDS

Event-triggered; interference; network flooding; wake-up receiver

1 INTRODUCTION

Motivation. Early warning systems are an important application domain of wireless sensor networks, whereby a spatially distributed network of resource constrained source nodes detect *events* which are rapidly communicated to a remote host for analysis. In such systems, events arrive according to a non-deterministic arrival rate. Once an event is detected at a source node, it is fully characterized by application-specific sensor *data*. Examples of events and their associated data include the detection of an intruder using an infrared sensor as characterized by the captured image [25], the detection of a gas leak using a gas sensor as characterized by the hydrocarbon response [1], and the detection of a rock wall fracture using an acoustic sensor as characterized by feature extraction [5].

In order for the remote host to react quickly to an event, the network architecture must adhere to the following requirements:

- ***Responsiveness:*** Events and their associated data must disseminate through the network with minimal latency.
- ***Energy efficiency:*** Energy consumption must be minimized so to maximize operational lifetime.

HotWireless'17, October 16, 2017, Snowbird, UT, USA

ACM ISBN 978-1-4503-5140-9/17/10…$15.00
https://doi.org/10.1145/3127882.3127883

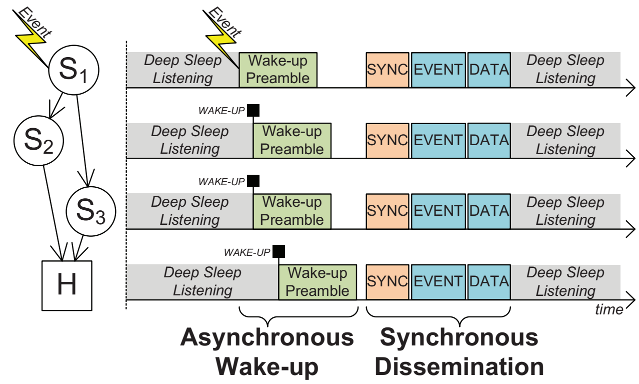

Figure 1: Example multi-hop network and the radio activity of BLITZ using interference-based flooding.

Challenges. The dissemination of events to a remote host may be achieved using *periodic communication*, *e.g.*, using LPL [19], LPP [16] or variants thereof, where nodes communicate according to a specified radio duty-cycle. However, this communication scheme exhibits a fundamental trade-off between latency and energy efficiency. When an event is detected, a node must wait until the beginning of its next communication round before dissemination commences, thereby increasing latency and adversely impacting responsiveness. Furthermore, since all nodes must communicate periodically to maintain network state, precious energy resources are expended irrespective of the detection of an event. One may improve energy efficiency by communicating less frequently, but only at the cost of increasing latency. This fundamental trade-off is a severe design constraint, that until recently, has received little attention in the literature.

Approach. We overcome this fundamental trade-off by facilitating *event-triggered communication*, where nodes in a multi-hop network only communicate when there is an event to disseminate, therefore conserving precious energy resources between event arrivals. We propose BLITZ, a novel network architecture that combines two orthogonal communication primitives using interference-based flooding. As illustrated in Fig. 1, we consider a multi-hop network of resource-constrained *source* nodes S_i and a remote *host* H. During periods of inactivity, *i.e.*, where no events are detected, the entire network resides in a *deep sleep listening* state where energy consumption is at its lowest, while continuously listening to the wireless channel using an ultra-low power wake-up receiver. When an event is detected, BLITZ invokes two communication primitives, namely, *(i)* asynchronous wake-up, and *(ii)* synchronous dissemination. The asynchronous wake-up primitive quickly awakes the network hop-by-hop from the energy conserving state, before the synchronous dissemination primitive synchronizes the network, allocates network bandwidth, and disseminates the event and its associated data to the host. The entire network then returns to the deep sleep listening state until the next event is detected.

The BLITZ network architecture leverages *interference-based flooding*, whereby the asynchronous wake-up and synchronous dissemination primitives employ network flooding, while allowing neighboring nodes to cause interference through simultaneous transmissions. These techniques may appear counterintuitive, since flooding potentially wastes energy through redundant transmissions and simultaneous transmissions may lead to packet corruption. But we will show that the proposed network architecture facilitates event-triggered multi-hop communication without having to trade-off responsiveness and energy efficiency, and without suffering from packet corruption.

Contributions. We make the following contributions:

- We present a novel communication architecture that supports low latency and energy-efficient event-triggered wireless communication using interference-based flooding.
- We introduce an analytical model to quantify the limits of BLITZ compared to a state-of-the-art protocol based on periodic communication.
- We present a prototype of BLITZ and experimentally evaluate its performance with respect to latency in an indoor testbed, and energy efficiency in a laboratory setting.

2 PROTOCOL DESIGN

Overview. In order to motivate the design of BLITZ, we first outline the intuition behind achieving efficient event-triggered wireless multi-hop communication:

> *Sleep as long as possible, quickly wake-up when there is something interesting to share, and promptly arbitrate information transfer.*

We next discuss how we realize this functional behavior through the combination of two orthogonal communication primitives.

2.1 Asynchronous Wake-up

Requirements. In order for all nodes in the network to sleep as long as possible, each node must only wake-up if an event is detected or if a neighbor has an event to disseminate.

Challenges. Due to the non-deterministic nature of event arrivals, nodes do not know in advance when to wake-up for dissemination. A common approach in the literature is to employ a periodic communication scheme, however this leads to a system design trade-off between latency and energy efficiency.

Proposal. An alternative approach is to employ an on-demand wake-up scheme facilitated by wake-up receivers [4]. Wake-up receivers are ultra-low power low-complexity demodulation circuits capable of detecting a wireless signal. The power dissipation of wake-up receivers is low enough to keep them always on, and thus still achieving a node lifetime of several years using a single battery charge. The always-on wake-up receiver enables a node to wake-up its neighbors at any time by transmitting a wake-up preamble, *i.e.*, a burst of a carrier signal. The wake-up receiver detects the wake-up preamble and awakes the node from the deep sleep listening state.

In order to quickly wake-up a multi-hop network, BLITZ floods wake-up preambles. As soon as a node has received a wake-up preamble, it immediately transmits a wake-up preamble to be received by other nodes. A practical limitation of this approach is that the simultaneous transmissions from neighboring nodes may destructively interfere, resulting in either a delayed or a missed wake-up. Rather than attempting to avoid simultaneous transmissions, we instead encourage the interference by randomizing the structure of the wake-up preamble transmission such that the probability of complete destructive interference is reduced, while taking advantage of the superposition of signals from neighboring nodes to improve reliability. Specifically, we leverage carrier frequency randomization, as introduced and evaluated in [22], which represents the wake-up preamble as a sequence of random FSK symbols.

2.2 Synchronous Dissemination

Requirements. Once all nodes in the network are awake, the available network resources must be arbitrated based on *(i)* which nodes have an event to disseminate, and *(ii)* how much bandwidth is required to disseminate each event and its associated data.

Challenges. The key challenge is that when the network awakes from the deep sleep listening state, there is no network state to leverage. Specifically, there is no local or global time synchronization, the network topology is unknown, and the bandwidth requirements between source nodes and the host are yet to be determined.

Proposal. In principle, one may apply any number of multi-hop protocols from the literature to acquire network state and deliver the event and data to the host. However, the time and energy required to acquire the necessary network state will impact overall responsiveness and energy efficiency. We therefore chose to leverage a *synchronous* and *topology-agnostic* protocol. This unique combination of properties takes advantage of a globally-synchronized schedule for rapid bandwidth arbitration without having to spend time and energy to discover the network topology.

It has been shown in [3] that commodity IEEE 802.15.4-compatible transceivers can be used to achieve topology-agnostic global time synchronization by leveraging constructive interference. Through careful time-triggered operation of the transceiver, the transmission of IEEE 802.15.4 symbols can be aligned so that they constructively interfere, thus improving the reliability of packet reception. When constructive interference is combined with network flooding, which is known as a Glossy flood [3] in the literature, one can quickly synchronize a multi-hop network with a fine-grained resolution.

The Event-based Low-power Wireless Bus (eLWB) [23] is a synchronous protocol that leverages time-slotted Glossy floods for the dissemination of events and their data. Since the eLWB combines topology-agnostic synchronization with adaptable bandwidth allocation, we chose to integrate it into the BLITZ network architecture.

3 COMPARATIVE ANALYSIS

We next present an analytical model to quantify the limits of BLITZ. We compare the performance of BLITZ to the eLWB as it is a representative of a class of protocols that are competitive with respect to latency, energy-efficiency and reliability [13].

Model Description. We consider a multi-hop network consisting of source and host nodes. The source detects the arrival of sporadic events, which are characterized by the smallest time between events T_{event}. The source and host leverage a multi-radio platform architecture consisting of a microcontroller interfaced to an IEEE

802.15.4-compatible transceiver and a wake-up receiver. The platform dissipates power P_s when the microcontroller is in sleep mode, P_p when the microcontroller and transceiver are active, and P_{wur} when the microcontroller is in sleep mode and the wake-up receiver is active. It is assumed that the power dissipation for the transmission and reception of IEEE 802.15.4 packets, and the transmission of a wake-up preamble, are equal.

We consider the following performance metrics:

- Best-case and worst-case latency, L_{min} and L_{max}, is the minimum and maximum time between the arrival of a sporadic event at the source and the reception of the event and its associated data at the host, respectively.
- Energy efficiency is represented by the average energy consumption per event E_{avg}, as evaluated at the source.

eLWB. The interaction between source and host using the eLWB is illustrated in Fig. 2. The eLWB supports three types of rounds, namely SYNC, EVENT, and DATA rounds. The eLWB maintains synchronization using periodic SYNC rounds with duration T_S, according to a constant period $T_{round} \leq T_{event}$. In this model, the SYNC round period is constrained to ensure there is at least one opportunity to disseminate between the smallest event inter-arrival time T_{event}. When an event is detected, the source node requests the host to schedule an EVENT round with duration T_E. During the EVENT round, the source disseminates the event to the host and requests bandwidth for the associated data. The host then schedules a DATA round with duration T_D for data dissemination. A detailed description of the round structure and the time separation between rounds, δ_{SE} and δ_{ED}, can be found in [23].

The eLWB rounds are implemented using time-slotted Glossy floods [23], whereby the transceiver sequentially receives a packet and immediately retransmit the packet a fixed number of times. Since the number of retransmissions is constant for all nodes in the network, the transceiver on time for each round is equal for all nodes in the network. However, due to the multi-hop propagation of Glossy floods, nodes residing $h > 1$ hops away from the host must listen for a duration of $(h-1)T_{hop}$ until the round commences.

The best-case latency occurs when an event arrives just as a SYNC round begins, while the worst-case latency will be T_{round} longer, as represented in equations (1) and (2), respectively. Since the SYNC round period is independent of event arrivals, there will be on average $\eta = \frac{T_{event}}{T_{round}}$ rounds per event, leading to an average energy consumption per event according to equation (3).

$$L_{min}^{\text{eLWB}} = (h-1)T_{hop} + T_S + \delta_{SE} + T_E + \delta_{ED} + T_D \tag{1}$$

$$L_{max}^{\text{eLWB}} = L_{min}^{\text{eLWB}} + T_{round} \tag{2}$$

$$E_{avg}^{\text{eLWB}} = \Big(\eta(h-1)T_{hop} + \eta T_S + T_E + T_D\Big) P_p + \Big(T_{event} - (\eta(h-1)T_{hop} + \eta T_S + T_E + T_D)\Big) P_s \tag{3}$$

BLITZ. Fig. 3 illustrates the interaction between source and host using the proposed BLITZ network architecture. Once an event is detected, the asynchronous wake-up primitive awakes the multi-hop network from the deep sleep listening state by flooding a wake-up preamble of duration T_W. Once the preamble transmission is complete, the node must wait a guard time of T_G until the synchronous dissemination primitive commences. The guard time is the sum of the time to initialize the transceiver T_{init} and to receive a wake-up preamble T_{wake} accumulated over $h \geq 1$ hops.

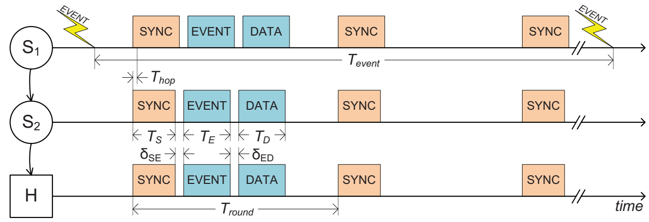

Figure 2: eLWB event and data dissemination.

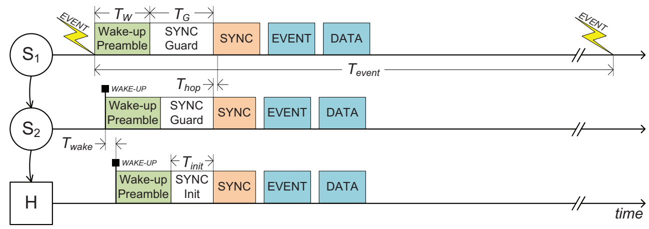

Figure 3: BLITZ event and data dissemination.

Table 1: Parameterization of analytical model.

Parameter	Value	Parameter	Value	Parameter	Value
P_s	8.5 μW	P_{wur}	16 μW	P_p	50 mW
T_S	22 ms	T_E	18 ms	T_D	22 ms
T_W	1.4 ms	T_G	17.5 ms	T_{hop}	1.3 ms
T_{wake}	0.75 ms	δ_{SE}	6 ms	δ_{ED}	12 ms

Since the asynchronous wake-up primitive starts immediately after an event is detected, the best-case and worst-case latency for BLITZ are identical. As expressed in (4), the latency is equal to the best-case latency for eLWB plus the overhead for the asynchronous wake-up primitive. The average energy per event for BLITZ is given by equation (5).

$$L_{min}^{\text{BLITZ}} = L_{max}^{\text{BLITZ}} = T_W + T_G + L_{min}^{eLWB} \tag{4}$$

$$E_{avg}^{\text{BLITZ}} = \Big(T_W + T_G + (h-1)T_{hop} + T_S + T_E + T_D\Big) P_p + \Big(T_{event} - (T_W + T_G + (h-1)T_{hop} + T_S + T_E + T_D)\Big) P_{wur} \tag{5}$$

Results. Using measurements from a prototype implementation, we parameterize the model according to the values listed in Table 1 for $h = 10$ hops. The analysis indicates that the latency for eLWB varies between 91.7 ms and 91.7 ms + T_{round}, while the latency for BLITZ is constant at 110.6 ms. While the eLWB may achieve lower latency than BLITZ in some specific cases, the eLWB worst-case latency will increase as T_{round} is increased so to reduce energy consumption. BLITZ instead achieves a constant low latency of 110.6 ms, while only dissipating 16 μW during periods of inactivity.

We next investigate if the overhead associated with the BLITZ asynchronous wake-up primitive improves energy efficiency. Given an application-specific worst-case latency, L_{max}, we evaluate the maximum eLWB round period according to equation (2), and then determine the T_{event} where the average energy per event for eLWB and BLITZ are equal. Fig. 4 illustrates the resulting non-linear protocol partitioning consisting of four unique regions. The first region represents where L_{max} is less than BLITZ can support, the

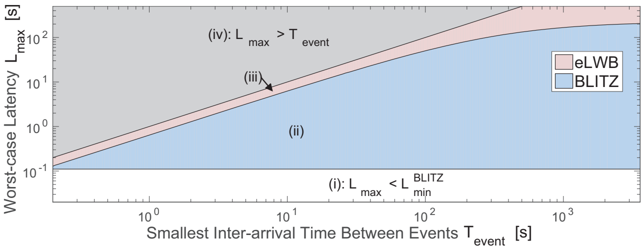

Figure 4: Results of the comparative analysis.

second region where BLITZ is more energy-efficient, the third region where eLWB is more energy-efficient, and the fourth region where $L_{max} > T_{event}$, and is therefore infeasible under the model assumptions. The results highlight the superiority of BLITZ over the eLWB in terms of latency and energy efficiency for a wide-range of application scenarios. The eLWB is preferred only when the constraint on worst-case latency is relaxed, which results in the periodic eLWB SYNC rounds consuming less energy than keeping the wake-up receiver always on.

4 EXPERIMENTAL EVALUATION

We next present a prototype of BLITZ and experimentally evaluate it in terms of latency and energy efficiency.

Prototype. We follow the event-triggered design methodology presented in [23] to realize a wireless sensing platform, as depicted in Fig. 5, that is responsive and energy-efficient by design. The platform architecture maps sensing and communication tasks onto dedicated microcontrollers, and interconnects them using an interface with predictable run-time behavior [21]. The application microcontroller implements asynchronous wake-up using the hardware presented in [22], while the communication microcontroller implements synchronous dissemination using an IEEE 802.15.4-compatible transceiver.

4.1 End-to-end Latency

Setup. We deployed a total of 16 prototypes into the FlockLab [11] indoor testbed, which is configured with fine-grained tracing capabilities [12]. The location of each node is depicted in Fig. 5. Node 23 was configured as the host, while all remaining nodes were configured as source nodes. A total of 100 events were generated at node 33 with T_{event} = 5 s. Once the host receives an event and its associated data, it schedules a second SYNC round that contains a broadcast message informing all source nodes to return to the deep sleep listen state.

We evaluate the performance of BLITZ using the following metrics: *(i) wake-up reception rate (WRR)* is the ratio of the number of nodes that wake-up compared to the number of events generated, *(ii) wake-up delay* is the time between an event generated at node 33 and the reception of the wake-up preamble, and *(iii) mean latency* is the average time between an event generated at node 33 and received at the host, and similarly, the average time between the host initiating a broadcast and when it is received by a source node.

Results. The results of the testbed evaluation are shown in Fig. 6. The asynchronous wake-up was 100% successful for all nodes in the network, except for nodes 16, 22, and 28. While this can be attributed to poor RF propagation, it is important to note that despite these nodes not participating in the all synchronous dissemination

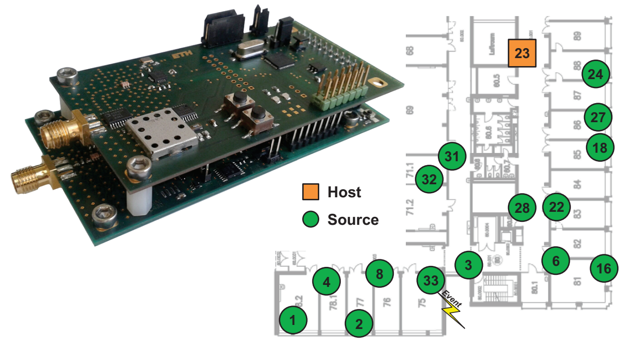

Figure 5: BLITZ prototype and testbed deployment.

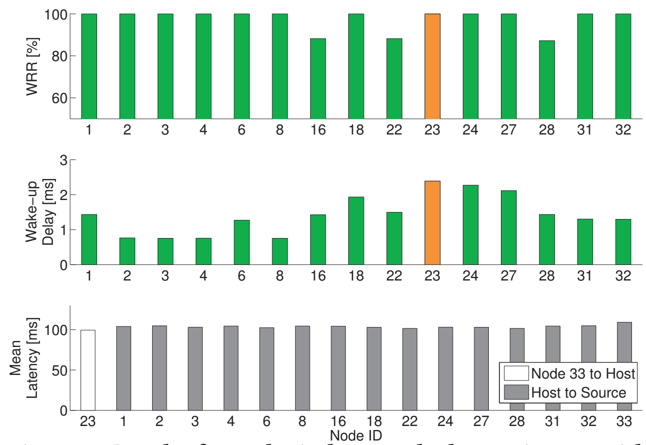

Figure 6: Results from the indoor testbed experiment, with the host at node 23.

occasions, there was sufficient network connectivity to disseminate all generated events to the host. The wake-up delay was measured between 0.75 ms and 2.4 ms, which is consistent with the parameterization of the wake-up receiver, and the latency associated with network flooding, *i.e.*, wake-up delay increases with hop count. The mean latency from node 33 to the host is 99.4 ms. This is very close to the latency of 98.3 ms as determined by the analytical model presented in Sec. 3 when parameterized with h = 4 hops. The mean latency from the host to source averaged across all nodes is 103.8 ms. We conclude that the BLITZ prototype achieves a low end-to-end latency consistent with the analytical model presented in Sec. 3.

We now compare the BLITZ latency performance with the closest work in the literature, the ROD-SAN [24] radio-on-demand architecture, where a wake-up delay of 300 ms and a minimum dissemination delay of 600 ms per hop are reported. While the experimental setup differs significantly in terms of network architecture, platform design and deployment conditions, we highlight that BLITZ supports a wake-up delay on the order of milliseconds and a mean latency of 99.4 ms through a 4-hop network.

4.2 Energy Efficiency

Setup. Using the RocketLogger [20] precision measurement device, we measured the power profile of the BLITZ protoype during the dissemination of an event and its associated data. The prototype was supplied with 2.5 V while having its on-board low-dropout regulator bypassed.

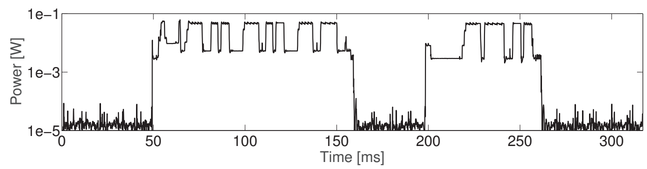

Figure 7: Power profile of source node using BLITZ.

Results. As illustrated in Fig. 7, the source node dissipated approximately 16 μW during periods of inactivity, and approximately 50 mW during wake-up preamble transmission and synchronous dissemination. Assuming a sporadic event arrival with T_{event} = 120 s, the presented prototype would operate for at least one year on a low-capacity coin cell battery.

5 DISCUSSION

Due to the limitations of wake-up receivers, as surveyed in [2], the latency and energy efficiency of BLITZ may be adversely impacted. Since wake-up receivers exhibit a lower sensitivity compared to IEEE 802.15.4-compatible transceivers, the per-hop range for successful asynchronous wake-up is lower than for synchronous dissemination. This may lead to some nodes not awaking from the deep sleep listening state, potentially disconnecting the host from the network and therefore increasing end-to-end latency. One possible solution is to increase the spatial diversity of wake-up receivers by deploying additional source nodes that only implement the asynchronous wake-up primitive.

Furthermore, wake-up receivers are susceptible to interference, meaning a wake-up preamble may be received without a neighbor transmitting a wake-up preamble. This will cause the entire network to awake from the deep sleep listening state only to return to sleep after a timeout. We intend to address this waste of precious energy resources in future work using pattern classification techniques.

6 RELATED WORK

As surveyed in [6], several protocols have been proposed for event-triggered wireless communication including, SIFT [7] and Alert [17]. However, since these works are based on variants of periodic communication protocols, they all exhibit a fundamental trade-off between latency and energy efficiency. To the best of the authors knowledge, we are the first to present a network architecture for event-triggered communication that leverages interference-based flooding that achieves low latency and energy efficiency.

Wireless communication technologies such as backscatter [10, 14] and RFID [8] are well suited to a range of wireless sensing applications. However, the limited range and reliance on the existence of high-powered RF signals to facilitate communication limits their adoption, particularly when high spatial-diversity and per-hop link ranges greater than ten meters are desired.

Since the introduction of wake-up receivers, only a few wireless protocols in the literature, most notably WUR-MAC [15], FLOOD-WUP [18] and GWR-MAC [9], leverage wake-up receivers for efficient multi-hop dissemination. However, these works make simplifying assumptions on network topology and arbitration of network bandwidth, and lack experimental evaluation. Specifically, line or star topologies are typically assumed, which circumvents the challenges associated with the rapid wake-up of dense multi-hop topologies. Furthermore, the allocation of network bandwidth is typically contention-based, whereby a random back-off mechanism exacerbates latency and severely hinders energy efficiency. We instead tackle these challenges with BLITZ, and experimentally evaluate a prototype implementation.

7 CONCLUSIONS

This paper presents the design, analysis, and experimental evaluation of BLITZ, a network architecture for efficient event-triggered multi-hop wireless communication. BLITZ leverages interference-based flooding where all nodes participate in the dissemination of an event and its data using network flooding, while embracing the interference caused by simultaneous transmissions. We experimentally evaluate the performance of a prototype implementation that exhibits a mean latency of 99.4 ms across 4-hops, while dissipating only 16 μW during periods of inactivity.

ACKNOWLEDGMENTS

This work was scientifically evaluated by the SNSF and financed by the Swiss Confederation and Nano-Tera.ch.

REFERENCES

[1] Fabien Chraim et al. 2016. Wireless gas leak detection and localization. *IEEE Transactions on Industrial Informatics* 12, 2 (2016), 768–779.
[2] Ilker Demirkol et al. 2009. Wake-up receivers for wireless sensor networks: benefits and challenges. *Wireless Communications, IEEE* (2009).
[3] Federico Ferrari et al. 2011. Efficient network flooding and time synchronization with Glossy. In *IPSN*. IEEE.
[4] GU Gamm et al. 2012. Low-power sensor node with addressable wake-up on-demand capability. *International Journal of Sensor Networks* 11, 1 (2012), 48–56.
[5] L Girard et al. 2012. A custom acoustic emission monitoring system for harsh environments: application to freezing-induced damage in alpine rock-walls. *Geoscientific Instrumentation, Methods and Data Systems Discussions* (2012).
[6] David C Harrison et al. 2016. Rare event detection and propagation in wireless sensor networks. *Comput. Surveys* 48, 4 (2016), 58.
[7] Kyle Jamieson et al. 2006. Sift: A MAC protocol for event-driven wireless sensor networks. In *EWSN*. Springer, 260–275.
[8] R. Jurdak et al. 2010. Radio Sleep Mode Optimization in Wireless Sensor Networks. *IEEE Transactions on Mobile Computing* (2010).
[9] Heikki Karvonen et al. 2014. A generic wake-up radio based MAC protocol for energy efficient short range communication. In *PIMRC*. IEEE, 2173–2177.
[10] Bryce Kellogg et al. 2015. Wi-fi backscatter: Internet connectivity for rf-powered devices. *ACM SIGCOMM Computer Communication Review* 44, 4 (2015), 607–618.
[11] Roman Lim et al. 2013. FlockLab: A testbed for distributed, synchronized tracing and profiling of wireless embedded systems. In *IPSN*. ACM.
[12] Roman Lim et al. 2015. A Testbed for Fine-Grained Tracing of Time Sensitive Behavior in Wireless Sensor Networks. In *Local Computer Networks Workshops*.
[13] Roman Lim et al. 2017. Competition: Robust Flooding using Back-to-Back Synchronous Transmissions with Channel-Hopping. In *EWSN*.
[14] Vincent Liu et al. 2013. Ambient backscatter: wireless communication out of thin air. *ACM SIGCOMM Computer Communication Review* 43, 4 (2013), 39–50.
[15] S Mahlknecht et al. 2009. WUR-MAC: energy efficient wakeup receiver based MAC protocol. In *IFAC*.
[16] R Musaloiu-E et al. 2008. Koala: Ultra-Low Power Data Retrieval in Wireless Sensor Networks. In *IPSN*. IEEE.
[17] Vinod Namboodiri et al. 2008. Alert: An adaptive low-latency event-driven MAC protocol for wireless sensor networks. In *IPSN*. IEEE, 159–170.
[18] Chiara Petrioli et al. 2014. A Novel Wake-Up Receiver with Addressing Capability for Wireless Sensor Nodes. In *DCOSS*.
[19] Joseph Polastre et al. 2004. Versatile low power media access for wireless sensor networks. In *SenSys*. ACM.
[20] Lukas Sigrist et al. 2017. Measurement and Validation of Energy Harvesting IoT Devices. In *DATE*.
[21] Felix Sutton et al. 2015. Bolt: A stateful processor interconnect. In *SenSys*. ACM.
[22] Felix Sutton et al. 2015. Zippy: On-Demand Network Flooding. In *SenSys*. ACM.
[23] Felix Sutton et al. 2017. The Design of a Responsive and Energy-efficient Event-triggered Wireless Sensing System. In *EWSN*. ACM.
[24] Hiroyuki Yomo et al. 2015. Radio-On-Demand Sensor and Actuator Networks (ROD-SAN): System Design and Field Trial. In *GLOBECOM*. IEEE.
[25] P. Zappi et al. 2010. Tracking Motion Direction and Distance With Pyroelectric IR Sensors. *IEEE Sensors Journal* (2010).