

Received May 2, 2019, accepted May 28, 2019, date of publication May 30, 2019, date of current version August 2, 2019.

Digital Object Identifier 10.1109/ACCESS.2019.2920100

Atomic-SDN: Is Synchronous Flooding the Solution to Software-Defined Networking in IoT?

MICHAEL BADDELEY¹, USMAN RAZA², ALEKSANDAR STANOEVI², GEORGE OIKONOMOU¹, REZA NEJABATI¹, MAHESH SOORIYABANDARA², AND DIMITRA SIMEONIDOU¹

¹Department of Electronic and Electrical Engineering, University of Bristol, Bristol BS8 1TH, U.K.

²Toshiba Research Europe Ltd., Bristol BS1 4ND, U.K.

Corresponding author: Michael Baddeley (m.baddeley@bristol.ac.uk)

This work was supported in part by the Engineering and Physical Sciences Research Council (EPSRC) Centre for Doctoral Training (CDT) in Communications under Grant EP/I028153/1, and in part by Toshiba Research Europe Ltd.

ABSTRACT The adoption of software defined networking (SDN) within traditional networks has provided operators the ability to manage diverse resources and easily reconfigure networks as requirements change. Recent research has extended this concept to IEEE 802.15.4 low-power wireless networks, which form a key component of the Internet of Things (IoT). However, the multiple traffic patterns necessary for SDN control makes it difficult to apply this approach to these highly challenging environments. This paper presents Atomic-SDN, a highly reliable and low-latency solution for SDN in low-power wireless. Atomic-SDN introduces a novel synchronous flooding (SF) architecture capable of dynamically configuring SF protocols to satisfy complex SDN control requirements, and draws from the authors' previous experiences in the IEEE EWSN Dependability Competition: where SF solutions have consistently outperformed other entries. Using this approach, Atomic-SDN presents considerable performance gains over other SDN implementations for low-power IoT networks. We evaluate the Atomic-SDN through simulation and experimentation, and show how utilizing SF techniques provides latency and reliability guarantees to SDN control operations as the local mesh scales. We compare the Atomic-SDN against the other SDN implementations based on the IEEE 802.15.4 network stack, and establish that the Atomic-SDN improves SDN control by orders-of-magnitude across latency, reliability, and energy-efficiency metrics.

INDEX TERMS SDN, IoT, cyber physical systems, WSN, Sensor networks, synchronous flooding, concurrent transmissions, low power wireless, control systems, industrial IoT.

I. INTRODUCTION

A. CONTEXT AND CHALLENGE

By providing an abstracted network view and virtualized network functions, Software Defined Networking (SDN) allows services to be centrally programmed onto functionally agnostic hardware. The ability to reconfigure the network as needed, quickly install new protocols, or slice network resources across applications and tenants, allows networks to adapt to changing requirements or shifts in business needs.

Implementation of this concept has been immensely successful in data-centre and radio-access networks [1],

The associate editor coordinating the review of this manuscript and approving it for publication was Ting Yang.

where reliable and low-latency links support communication between a logically centralized controller (or a number of controllers) and the rest of the network. This ensures that control tasks, such as installing flowtable rules, cause minimum disruption to the data forwarding operations.

Extending the benefits of SDN to the wireless domain and to low-power sensor networks has gained popularity in recent years [2]. Consequently, a number of SDN architectures have been implemented for IEEE 802.15.4 low-power wireless networks. This has provided a model for SDN control and configuration across a low-power multi-hop mesh. However the shared nature of the underlying wireless medium, a need to transmit data over multi-hops, and stringent resource constraints pose significant challenges not present in the conventional wired SDN networks.

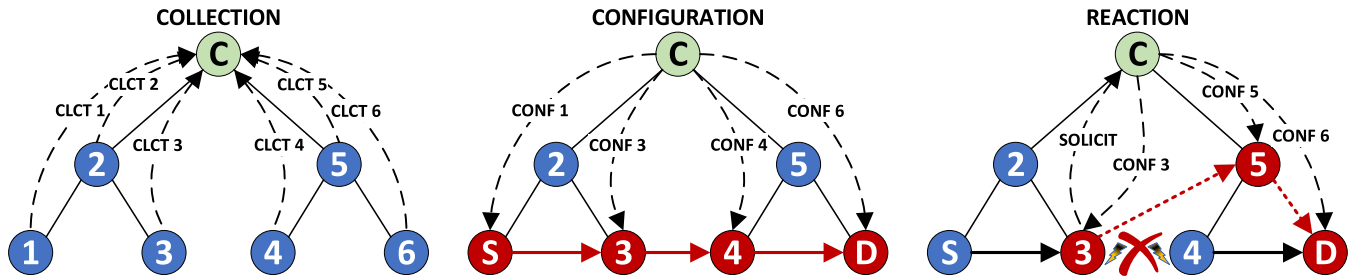


FIGURE 1. Core SDN services within a low-power mesh network: *Collection* (CLCT), *Configuration* (CONF), and *Reaction* (SOLICIT + CONF). Nodes that need to receive instruction from the SDN controller are marked in red. Nodes S and D mark the source and destination nodes for a point to point link across the mesh.

Firstly, additional SDN overhead increases contention over the shared wireless medium as well as competition with existing network protocols. Recent efforts have attempted to address this issue through optimization of the SDN protocols, reduction in message frequency, and prioritization or dedication of network resources [3]–[6]. Secondly, the multi-hop mesh topology prevalent in low-power wireless networks introduces delay and unreliability at each hop. This motivates the need for an SDN protocol that supports an ultra-fast hop-by-hop forwarding scheme and diversity techniques to achieve a very high reliability. Lastly, SDN requires frequent back-and-forth communication between the controller(s) and network nodes. The flow of this traffic follows a variety of different patterns including *many-to-one*, *one-to-many* and *one-to-one* communication. Unfortunately, the standard protocols such as the IPv6 Routing Protocol for Low Power and Lossy Networks (RPL), employed by several wireless SDN architectures, provide less-than-optimal performance for the plurality of these traffic patterns.

A radically different design is required, supporting all the traffic patterns to provide the following three essential core SDN services, which are presented in Figure 1.

- *Collection (many-to-one)*: allow the controller to gather node information and infer the current network state.
- *Configuration (one-to-many)*: allow the controller to provide instruction to all nodes within the mesh, or a subset thereof.
- *Reaction (one-to-one + one-to-many)*: allow nodes to query or alert the controller for instructions on how to react to new input or events.

This new approach ultimately needs to remove the complexities of mapping traditional SDN architecture to the currently available protocols in low-power wireless; where fundamental challenges arise from the controller not only having to communicate reliably with all nodes, but that each individual operation (for example, to set a path between two nodes) can mean the replication of control messages across multiple nodes in order to correctly configure the network.

This paper therefore proposes utilizing Synchronous Flooding (SF) as a basis for SDN control in IEEE 802.15.4 low-power wireless networks, and draws on the authors’ extensive experience in implementing SF solutions

for the International Conference on Embedded Wireless Systems and Networks (EWSN) Dependability Competition [7], [8], where a version of Atomic-SDN placed 2nd for both *collection* and *dissemination* categories in 2019 [9].

B. MOTIVATION

Over the past few years SF has been shown to be extremely capable in delivering fast, reliable communications in low-power wireless networks, and solutions based on this technique have consistently placed within the EWSN Dependability Competition [10]–[14]; achieving high reliability, low latency, and increased energy efficiency over standard approaches. These results support the case for using SF as a platform for SDN control in low-power wireless, the arguments for which are as follows:

Firstly, the broadcast nature of SF supports one-to-all communication within a single flood. This renders them inherently stateless, without the need for network topology. Currently, multi-hop mesh networks rely on underlying protocols such as Routing for Low-Power and Lossy Networks (RPL), which builds a Directed Acyclic Graph (DAG). This distributed routing protocol is typically used to funnel data from sensor networks towards a single border router, where they are later processed. However, *one-to-many downwards* communication is a common issue in RPL networks. An example of this challenge is shown in the *configuration* scenario in Figure 1. In this case, the SDN controller wishes to set a Point-to-Point (P2P) link from S→D across multiple branches of the RPL DAG. The tree-like topology forces the controller to navigate multiple branches to reach all destinations, resulting in packet duplication as it individually transmits to each child. This is particularly relevant in RPL non-storing mode which doesn’t support multicast forwarding, although recent efforts attempt to address this [15]. This issue isn’t specific to SDN in low-power wireless, however the complex requirements of SDN control means it is a highly visible and present issue for SDN implementations based on IEEE 802.15.4 networks.

Secondly, SF protocols benefit from minimal latency and extremely high reliability. They are able to aggressively and concurrently propagate across the entire network within short time period. Nodes don’t need to worry about interfering




FIGURE 2. Atomic-SDN uses time-sliced SF control to maximize network resource utilization during control periods. SDN collection (CLT), configuration (CFG), and reaction (REACT) opportunities are preceded by a indication (IND) flood.

with their neighbors and wait for their next transmit opportunity, but are able to immediately relay a packet on a different frequency. This approach, as employed in Atomic-SDN, is shown later in Figure 11. Additionally, the inherent nature of flooding means that the network benefits from a great deal of spatial diversity. In many scenarios this allows messages to skirt around interference hot-spots without re-transmitting, saving on energy and decreasing latency. This combination of frequency and spatial diversity is particularly relevant when operating over the 2.4GHz band which, in close proximity to external IEEE 802.11 devices, can pose significant interference to low-power wireless networks.

Finally, the time-synchronized nature of SF allows SDN control to be decoupled from other network processes. Currently, as neighboring nodes within a low-power wireless network share a single link, any additional control messaging increases contention over scarce resources, causing increased delay and reduced reliability for other control protocols (such as RPL and 6LoWPAN) and application data. Current approaches have tried to mitigate the burden of this overhead through reduction in the number of control messages, the use of source routing headers, and optimization of the protocol [4], [5], [16]. However, this limits the effectiveness of the SDN architecture: sacrificing responsiveness and fine-grain configurability for performance. SF provides a highly reliable means of communicating to all network nodes, and completely isolating this overhead to free-up network resources.

C. APPROACH

Atomic-SDN introduces a novel architecture to allow dynamic configuration of SF protocols in response to changing SDN requirements, and slices network resources in time to isolate the SDN overhead from other network processes and temporally decouple the SDN control plane. This allows configuration of other low-power wireless network layers: such as IEEE 802.15.4, Bluetooth Low Energy (BLE), or 6TiSCH.

Figures 2 and 3 show a high-level overview of this concept. Atomic-SDN designates these flooding periods as SDN control ‘opportunities’, where each opportunity is preceded by an indicator (IND) flood that informs the network of the type of SDN function that will follow (if any), and how that function will be configured: the role of each node, duration of each flood, etc. Using SF in this way allows Atomic-SDN to ensure maximum utilization of network resources within any given control period, with minimal latency and extremely

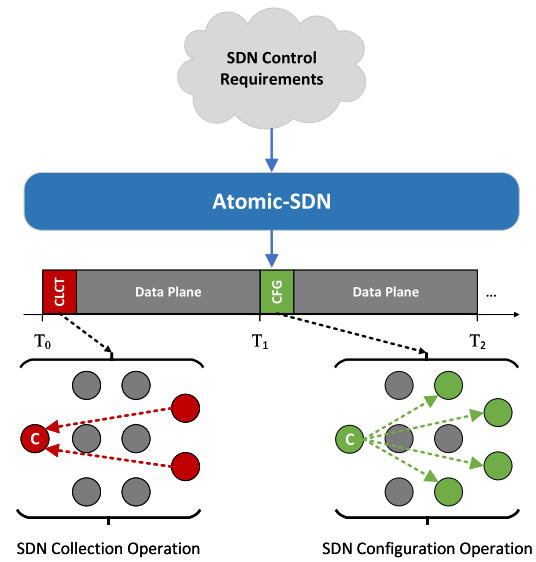


FIGURE 3. High-level overview of the Atomic-SDN approach in Figure 2.

high reliability. Although other flooding solutions match SF in terms of reliability, such as the Asynchronous Flooding (AF) approach used in Bluetooth Mesh, these tend to have far higher latencies in a Radio Duty Cycling (RDC) environment [17].

Although SF in and of itself is not novel, its application in solving the problem of SDN control in low-power wireless has not yet been explored. We show that, by utilizing SF to service the multiple traffic patterns required in SDN control, reliable and scalable SDN for low-power wireless networks can be achieved within the local mesh network (up to 100 nodes). As control messages are rapidly propagated across the entire network and reliably received in a single flood by all participating nodes, this dramatically reduces the burden of SDN control overhead that has frustrated current approaches.

D. CONTRIBUTION

This paper makes the following contributions:

- We propose SF as a mechanism for facilitating SDN control in low-power wireless networks.
- We devise a flexible middleware system for the design, instantiation, and scheduling of SF protocols.
- We apply this solution to the challenge of SDN control in low-power wireless networks, and present

Atomic-SDN: a scalable solution that offers considerable improvements in reliability, latency, and energy efficiency over current architectures.

- We implement Atomic-SDN in Contiki for motes with TI MSP430F1611 Microcontroller, and CC2420 radio.
- We evaluate Atomic-SDN against other SDN implementations for IEEE 802.15.4, through simulation on emulated hardware.
- We evaluate Atomic-SDN on a real-world testbed, and show that it can provide considerable reliability and latency improvements over current approaches to SDN control in low-power mesh networks.

E. OUTLINE

The remainder of this paper is presented as follows: in §II we provide a brief overview of related works exploring SDN in low-power IEEE 802.15.4 mesh networks, and provide necessary background on information on Synchronous Flooding (SF) and the concept of Concurrent Transmissions (CTs); we present Atomic-SDN in §III, where we cover design aspects; §IV characterizes Atomic-SDN performance through analysis of theoretical bounds on latency; in §V we evaluate Atomic-SDN through simulation and compare it to non-flood based SDN architectures for low-power wireless networks; and we conclude in §VI.

II. BACKGROUND MATERIAL AND RELATED WORKS

This section introduces current approaches examining SDN in low-power wireless networks, as well as necessary background material on Concurrent Transmissions (CT) and Synchronous Flooding (SF).

A. SDN IN LOW-POWER WIRELESS NETWORKS

Recent research has considered how to extend SDN control to low-power wireless networks. Whereas traditional SDN concepts have been successfully applied to other networking environments, such as data centers and optical, the constraints of low-power wireless networks (IEEE 802.15.4 in particular) pose considerable challenges to centralized control architectures. We provide a brief outline of current approaches, which have been covered in detail in recent surveys [2], [18], [19], and highlight how they attempt to overcome the challenges of implementing SDN within a constrained environment.

Sensor OpenFlow [20] argues for the use of SDN in sensor networks, proposing a custom low power protocol based on the traditional southbound protocol for SDN, OpenFlow [21]. The authors highlight the difficulties of implementing Out-Of-Band (OOB) control plane communication within a sensor network and attempt to mitigate SDN overhead through the introduction of Control Message Quenching (CMQ) [22], whereby retransmissions of SDN control messages from individual nodes are throttled in order to give the controller time to respond to the initial asynchronous request for instruction.

SDWN (Software Defined Wireless Networks) [3] provides an architectural framework and highlights novel uses for SDN in low-power wireless sensor networks. Specifically,

the authors introduce the idea of using SDN flowtables to configure in-network data aggregation and Radio Duty-Cycling, allowing the programmatic installation of rules which can help reduce the number of transmissions and improve the energy consumption of individual nodes. In addition, a form of Protocol Oblivious Forwarding (POF) [23] is proposed to reduce memory footprint, allowing flowtables to match on byte arrays within the packet, rather than needing multiple rules for specific packet types.

SDN-WISE [4] builds on architectural concepts introduced in SDWN, as well as introducing stateful flowtables: essentially turning the flowtables into a Finite State Machine (FSM). This allows simple controller logic to be ‘programmed’ into the nodes, where they can perform certain actions under one state, whilst performing a different set of actions when in another. For example, this could be used to allow nodes to run their SDN flowtable actions in a low-energy mode.

CORAL-SDN [5] reduces the effect of overhead generated by other control protocols on the SDN stack, and uses a mechanism to reduce RPL control messages in a IPv6 based IEEE 802.15.4 network as nodes initialize and associate with the SDN controller. This frees up resources for the SDN protocol, improving its scalability. However reducing the frequency of RPL control messages may cause issues when trying to maintain end-to-end links between the controller and the edges of the network, particularly in interfered or dynamic networks.

Additionally, SDN concepts are included in recent standardisation efforts from the IETF 6TiSCH Working Group (WG) [24], which aims to incorporate elements of SDN within its proposals for centralized scheduling mechanisms. 6TiSCH is engaged in developing scheduling processes for IEEE 802.15.4-2015 TSCH. This amendment allowed the creation of channel hopping schedules but did not define how these schedules should be properly configured or maintained. However, 6TiSCH foregoes traditional SDN elements such as flowtables and focuses the centralized allocation of TSCH slotframe resources (the TSCH channel/time slots) within the network.

B. CONCURRENT TRANSMISSIONS AND SYNCHRONOUS FLOODING

As previously stated, SF solutions have consistently beaten other approaches in reliability and latency metrics at the IEEE EWSN Dependability Competition. A comprehensive review of SF based protocols can be found in a recent survey and tutorial of CT in IEEE 802.15.4 networks [25]. Through a novel architecture, Atomic-SDN can dynamically configure and support multiple SF protocols from a single framework, allowing SDN control data to be propagated across entire network with minimal latency, while benefiting from high reliability due to the temporal and spatial diversity inherent in broadcast flooding communication protocols. In comparison to contemporary SDN solutions for low-power wireless, this considerably reduces the amount of traffic




FIGURE 4. Synchronous flooding using back-to-back transmissions in a 3-hop network (based on the schedule in Figure 5). Blue indicates a transmission, whilst green indicates a reception.

between nodes and the controller and allows Atomic-SDN to provide a scalable SDN solution in the local mesh (up to 100 nodes).

1) TEMPORAL DISPLACEMENT

The authors of Glossy [26] first proposed the use of CT to achieve highly reliable *one-to-many* communication within multi-hop low-power mesh networks. They found that as long as the maximum temporal displacement between concurrently transmitted signals (of the same data) was less than half a microsecond ($\Delta_{max} > 0.5\mu s$), then that data can be reliably demodulated without the transmitted signals interfering with one another.

2) MANAGING CLOCK DRIFT

However, the radio-driven nature results in clock drift, meaning synchronization can not be reliability maintained over multiple transmissions. Consequently, the authors proposed interleaving transmission (T_x) and reception (R_x) slots so that successful receptions could help correct this drift and align the next slot. Recent approaches have demonstrated techniques to estimate this drift [7], [12], and that slot interleaving is not necessary. This allows nodes to repeatedly T_x after the first reception so that data is forwarded at every slot, meaning the time taken to fully propagate the packet across the network is substantially reduced. This back-to-back T_x approach is utilized as the flooding primitive in Atomic-SDN, demonstrated in Figures 4 and 5.

3) FLOOD OPERATION

With reference to this back-to-back T_x approach, each flooding period is partitioned into slots. Both the maximum number of slots (MAX SLOTS) and maximum number of transmissions (MAX TX) are statically configured at the start of each flooding round. At the start of each flooding round the initiating (source) node transmits a packet, and repeatedly transmits on every slot until MAX TX. All other nodes have set their radios to receive. The receiving node then relays the packet on next slot, concurrently transmitting with all other forwarding nodes.

MAX SLOTS is used to calculate the maximum flooding time, while MAX TX is the number of times a node

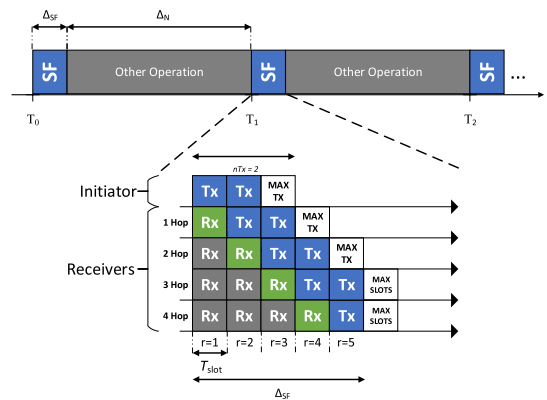


FIGURE 5. Synchronous flooding protocol used in Atomic-SDN. Back-to-back transmissions flood the network with minimal latency.

concurrently transmits after the first reception. Factors such as external interference, poor connectivity, and the network hop distance need to be taken in to consideration when these variables are set. Increasing them allows for greater reliability, and at a minimum the number of slots needs to equal the hop distance of the network, while minimising them allows for lower end-to-end latency in protocols with multiple flooding periods, as well as reducing energy consumption. In essence, these values are a trade-off between latency and allowing greater temporal and frequency diversity (if paired with slot-by-slot channel hopping [12]).

4) TIME SYNCHRONISATION

The length of each timeslot (T_{slot}) is determined by the time needed to transmit the packet (T_{tx}) (i.e. made up of the preamble, SFD, MPDU, and packet data), a software delay (T_{sw}) introduced by the micro-controller, a radio calibration delay (T_{cal}), and a processing delay incurred by the receiver radio (T_{rs}) incurred by the hardware, s.t.

$$T_{slot} = T_{tx} + T_{sw} + T_{cal} + T_{rs}$$

Non-initiating nodes listen for flood transmissions. When they successfully receive a packet, a relay counter in the header indicates how many hops (and consequently how many slots) have elapsed. Nodes combine this with their knowledge of T_{slot} to calculate the reference time of the

initiator and synchronize to that node. After synchronization, they relay the packet on the next timeslot; alongside any other neighbours who also received within that slot. Nodes at the next hop will repeat the process, and so on, until nodes have either reached MAX TX or the flooding period (Δ_{SF}) has elapsed (calculated from T_{slot} and MAX SLOTS).

Once synchronized, the initiating node effectively acts as a timesync for the network, allowing non-initiating nodes to duty-cycle their radio. The benefit of this approach, key to the operation of Atomic-SDN, is that it allows the protocol to be temporally decoupled from normal network operation; allowing it to be run alongside other control and application protocols or, as in the case of Atomic-SDN, be used to regularly configure those protocols.

5) MULTIPLE INITIATORS

Subsequent studies to the original Glossy paper have shown that the receiver is able to reliably demodulate multiple concurrent transmissions of the same data not necessarily because of so-called constructive interference, as the authors first considered, but likely as a result of transmissions being demodulated as non-coherent Minimum-Shift Keying (MSK), as well as Direct Sequence Spread Spectrum (DSSS) minimizing the error rate [27]–[29].

When multiple transmitters concurrently transmit different data, then the technique relies upon the Capture Effect [30]. This is found within IEEE 802.15.4 radios and refers to the phenomenon that the strongest signal out of multiple co-channel signals will be demodulated. It occurs either if one of the signals is around 3dB stronger, although this depends on the particular hardware and modulation schemes used, or if one of the signals is received significantly earlier than the other competing signals. Although the signals may still interfere, there is a high probability that one of the transmissions will be demodulated.

This property has therefore been used to great effect in *many-to-one* SF data collection protocols in IEEE 802.15.4 [31], allowing multiple initiators to participate in a shared flood. However, there is doubt as to how well these protocols would perform on other physical layers. Indeed, recently the authors of [32] have experimentally demonstrated that when nodes transmit different data and CTs are applied to Bluetooth physical layers (which do not experience capture effects as significantly as IEEE 802.15.4), then reliability drops significantly.

III. ATOMIC-SDN DESIGN

Atomic-SDN has been designed to tackle the issues faced by current approaches to SDN in low-power wireless networks. It implements the three core functions necessary for SDN control, as well as providing association with the SDN controller. Moreover, it facilitates these functions *as quickly as possible, as reliably as possible*, and maintains *scalability* in the local mesh. These functions and their associated traffic patterns are detailed below.

- **Collection (*many-to-one*):** Nodes need to be able to update the controller of their local and neighborhood state, so that the controller can make informed decisions when configuring the network.
- **Configuration (*one-to-many/one-to-all*):** The controller needs to be able to configure multiple nodes within the network, either to set data flows across the mesh, or to independently provide instruction to a number of nodes.
- **Reaction (*many-to-one/one-to-many*):** Nodes need to be able to react to unexpected flows or events by soliciting the controller for instruction, and quickly receiving a response.
- **Association (*many-to-one/one-to-all*):** Nodes need to be able to join the controller and be configured with initial instructions and network settings.

Atomic-SDN moves away from previous approaches addressing the challenge of SDN architecture in low-power wireless networks. Rather than layering the SDN architecture on top of standard asynchronous or synchronous Layer-2 protocols in the IEEE 802.15.4 networking stack (such as RPL), Atomic-SDN adopts Synchronous Flooding as the mechanism for communication between the SDN controller and nodes within the multi-hop mesh network.

Indeed, SF is increasingly seen as the ‘go-to’ solution for low-latency control in low-power wireless networks, particularly when applications require highly-robust communication for unpredictable and opportunistic traffic patterns. This view is supported by the consistent and continued success of SF solutions in the IEEE EWSN Dependability Competition [10]–[14], which benchmarks protocols on reliability, latency, and energy efficiency across multi-hop networks.

A. GENERAL APPROACH

Atomic-SDN provides periodic SDN control *opportunities*, where an initial *indicator* (IND) flood instructs all network nodes as to the type of SDN service that will follow (if any), as well as maintaining time synchronization across the mesh (as shown in Figures 2 and 9). This allows Atomic-SDN to separate SDN control from the data plane, and slice the network across time so that control messages are no longer in contention with other protocols (such as RPL, 6LoWPAN, or application-layer). Due to the broadcast nature of SF multiple nodes can be quickly and reliably serviced in a single flood, without replicating messages across multiple topology branches. This provides performance improvements orders-of-magnitude over current approaches to SDN in low-power wireless sensor networks.

However, to implement the different core SDN services within a multi-hop mesh network, multiple traffic patterns must be supported (*one-to-all, one-to-many, many-to-one, one-to-one*). Crucially, the plurality of these patterns are not supported by a single SF primitive or protocol and, as such, multiple protocols are needed to fulfil all required communication types. Unfortunately, the complex and low-level nature

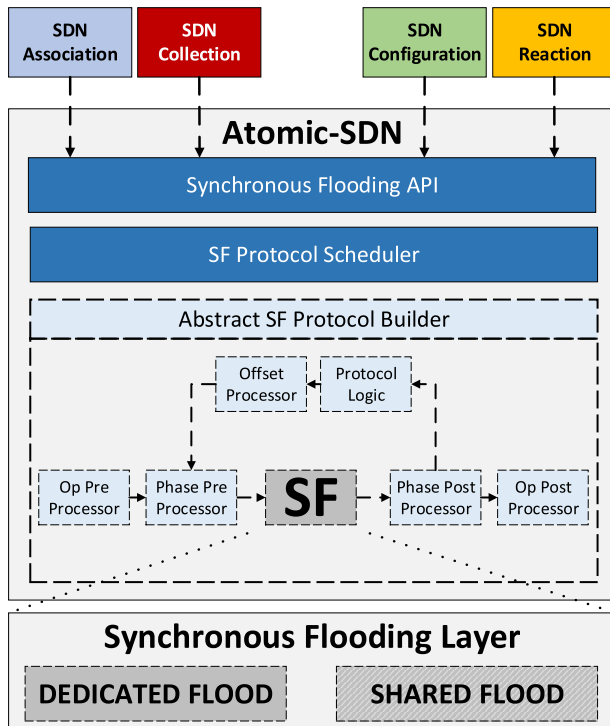


FIGURE 6. Atomic-SDN architecture.

of SF implementations has meant that, to date, there has been no unified framework allowing multiple SF protocols (such as Glossy [26], Chaos [33], LWB [34], or CRYSTAL [31]) to coexist within a single architecture.

Atomic-SDN solves this issue by introducing a novel SF architecture that allows the construction of complex, higher-level communication by applying pre and post logic functions on top of SF primitives. In this manner, different flooding protocols can be configured, instantiated, and scheduled, as the SDN control requirements change; allowing Atomic-SDN to adapt the SF protocol to the SDN service dictated by the controller, and meet application Quality-of-Service (QoS) requirements.

Figure 6 shows an overview of this architecture. By applying this configurable logic on top of generic flood primitives, SF protocols can be dynamically reconfigured at each SDN control period. The basic approach is as follows (going upwards from the lower layers):

- SF layer manages the lower level time synchronisation and concurrent transmissions.
- Floods are packaged into generic dedicated (single initiator) or shared (multiple initiators) flooding primitives.
- These primitives are configured with *offsets*, *guards*, and *logic blocks* to create logical *phases*.
- Phases are linked and scheduled to create an SF protocol.
- SF protocols are mapped to SDN functions, and tailored to their current service requirements, to create an SDN control *opportunity*.
- The SDN controller periodically initiates a required SDN operation during a scheduled SF control slot.




FIGURE 7. Atomic-SDN data collection and data dissemination protocols.

B. ATOMIC-SDN FLOODING OPERATIONS

To achieve the core SDN functions; *collection*, *configuration*, *reaction*), as well as network *association*; Atomic-SDN needs to perform three distinct traffic patterns:

- Single source to all destinations (*one-to-all*)
- Single source to a subset of destinations (*one-to-many*)
- Multiple sources to a single destination (*many-to-one*)

Atomic-SDN provides two SF protocols, (*collection* and *dissemination* shown in Figure 7), which can be used individually or in conjunction, to fulfil these patterns. Each schedule of dedicated or shared floods repeats until the SDN opportunity is complete.

The first two traffic patterns (*one-to-all* and *one-to-many*) can be achieved through an SF dissemination flooding protocol. In its most simple case, this allows the controller to rapidly and reliably communicate information to the entire network within a single flood, allowing SDN to bypass the packet duplication issues inherent in other SDN architectures for low-power wireless. The flood is then propagated across the network as nodes successfully receive the packet and start to relay the transmission. If a node is designated as a destination (designated in the same manner as before), it will read the packet data after the flood has ended, otherwise it will act as a forwarder.

The third pattern (*many-to-one*) is more complex, as previously discussed in §II. In SF collection protocols based on shared flood phases, multiple sources will compete as initiators. In each flood, only one source will successfully be received by the destination. Therefore, competing nodes that were not successful must continue to re-transmit until they are acknowledged in an ACK flood. As nodes are acknowledged they will switch their role from source to forwarder, and help with future transmission phases. This continues until all source nodes have had their transmissions acknowledged, which is indicated by a STOP consisting of one or more empty T_x floods plus a NACK phase.

C. ABSTRACT PROTOCOL BUILDER

The success of SF for control solutions in low-power wireless is rooted in the mechanism’s ability to provide low-latency and high reliability even under extremely adverse conditions. As such, there have been a number of attempts to take the core flooding principle, and tailor it to diverse application requirements in order to facilitate protocols for *one-to-all* communication [26], data collection [31], *many-to-many* communication [34], network consensus [33], [35], and interference management [36].

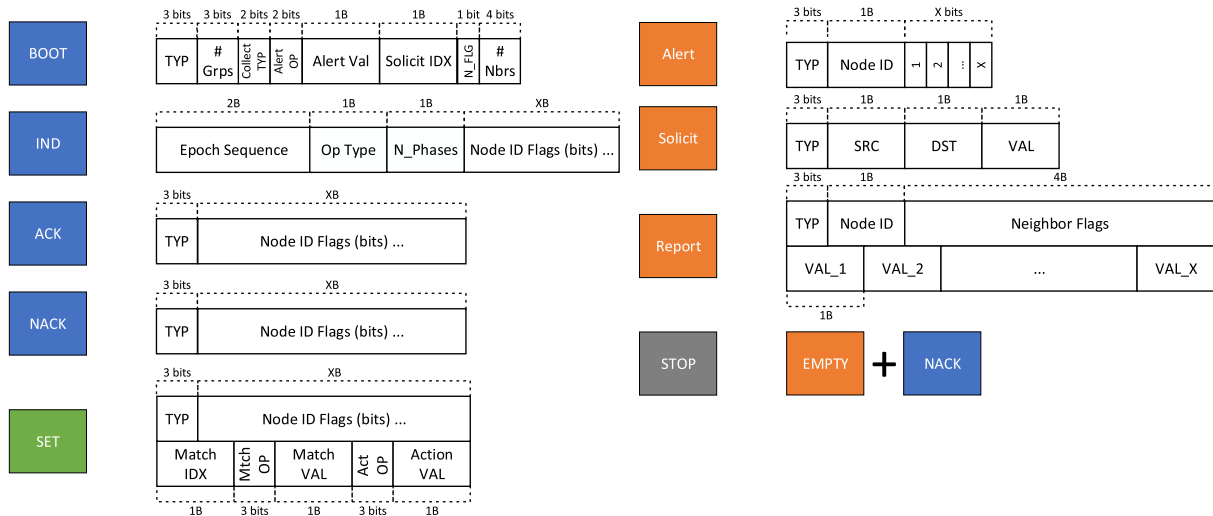


FIGURE 8. Atomic-SDN phase types built from the SF protocols outlined in Figure 7. These phases are chained together to create higher-level functionality in the form of an SDN ‘opportunity’. From left to right: one-to-all phases (blue), many-to-one phases (orange), one-to-many phases (green), and a STOP phase (grey).

Each of these protocols satisfies a specific set of application requirements. However, to fully implement SDN a number of different traffic patterns need to be supported, and achieving this therefore requires multiple SF protocols. Yet the underlying low-level implementation of proposed SF protocols have, to date, varied significantly; co-existence of multiple protocols within a single stack is therefore particularly challenging despite being based on the same basic underlying mechanism.

To address this issue, Atomic-SDN implements an Abstract Protocol Builder (APB) middleware layer (as shown in Figure 6), which uses generic flooding primitives attached with configurable protocol-specific logic to allow flexible construction of complex high-level synchronous flooding protocols. This mechanism is currently used within Atomic-SDN to implement the data collection and data dissemination protocols in Figure 7; however, the abstract nature of the APB means that it can be easily extended to implement any SF based protocol in order to suit additional traffic patterns or requirements.

1) FLOOD PRIMITIVES

In Atomic-SDN, generic ‘Flood Primitives’ are defined as a single flood as shown previously in Figure 5, configured with a MAX TX number of transmission slots, each with duration T_{slot} . If a node is able to successfully complete all MAX TX transmissions it will exit the flood process, otherwise it will exit at Δ_{SF} , the time taken for all transmission slots to elapse. Flood primitives are currently implemented as a one-to-all back-to-back transmission flood, however any lower synchronous flooding layer could conceivably be used, such as the Glossy interleaved $R_x T_x$ approach [26], or a consensus primitive such as Chaos [33].

2) PHASES

Phases are the building blocks of Atomic-SDN, allowing higher-level SDN functionality to be realized by chaining multiple phases into a series of logic decisions. Each phase is a self-contained unit consisting of a flood primitive configured with MAX TX transmissions and duration Δ_{SF} , combined with an associated data structure and the concrete implementation of the following abstract functions, as shown in Figure 6, which define phase behavior based on the current node role:

- Pre and post processing logic.
- Guard to allow for drift and processing in other nodes.
- Offset from initial phase reference.

By defining these functions, phases can be configured to perform a specific, self-contained role, whilst propagating the associated phase packet types shown in Figure 8. Multiple phases can then be chained together in order to build up higher level processes, known as opportunities, allowing full protocols to be implemented through the combination of a number of simple blocks.

3) OPPORTUNITIES

Atomic-SDN defines the concept of SDN opportunities, whereby the controller regularly and synchronously initiates a period of SDN control across the network. These are shown in Figure 10, where highlighted phases blocks are repeated until the opportunity is complete; either through a predefined number, or through a STOP phase. The type of opportunity is chosen by the controller prior to the flooding period, where the opportunity logic is constructed through the combination of a number of phase types, along with pre and post processing logic. Prior to execution, every opportunity is announced by the controller through a special one-to-many IND phase. This phase instructs the network as to what type of SDN control opportunity to expect (if any), the number of phases in

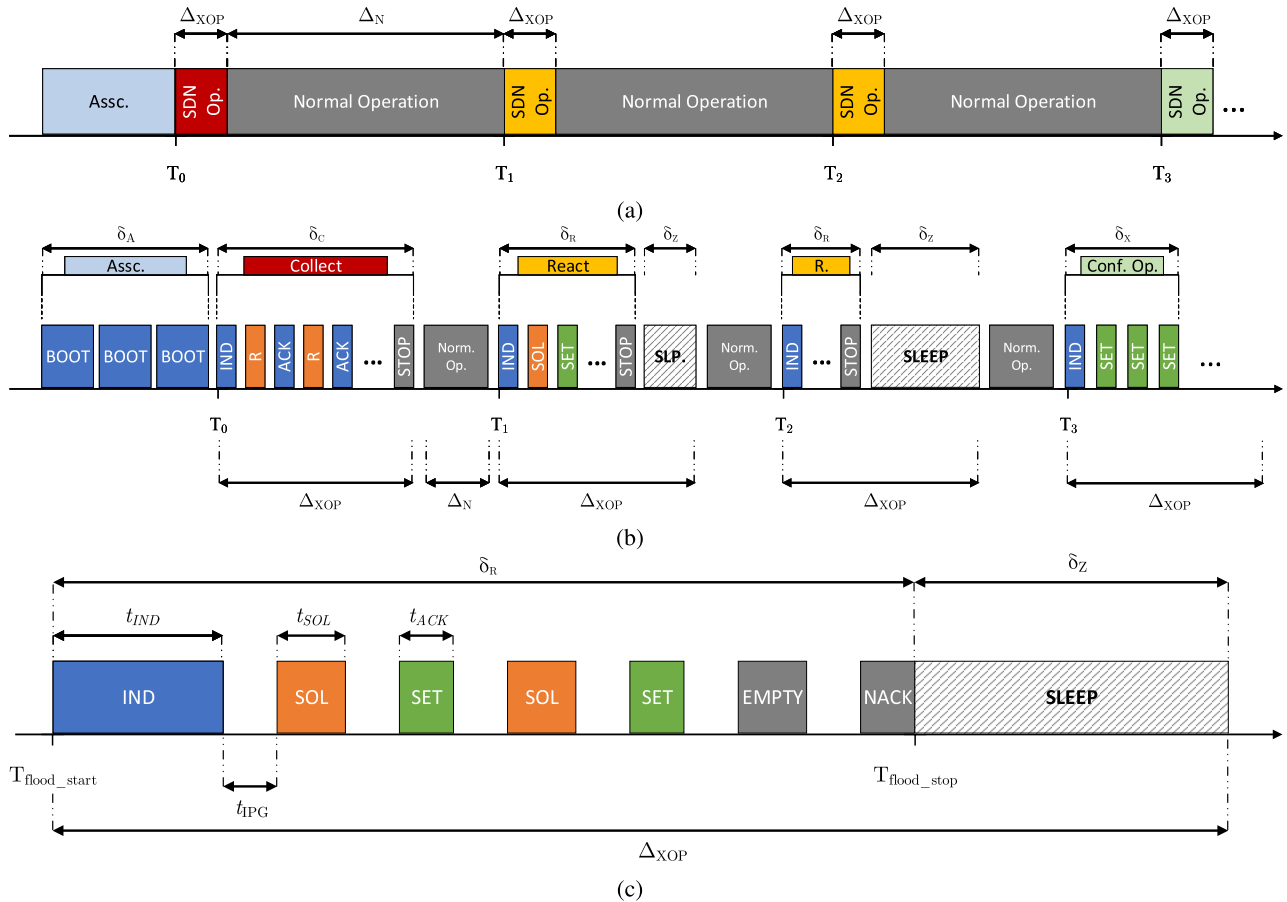


FIGURE 9. Example Atomic-SDN schedule at epoch, opportunity, and phase level. Opportunity timings are detailed in §IV. (a) Atomic-SDN epoch schedule. (b) Atomic-SDN opportunity schedule. (c) Atomic-SDN phase schedule (REACT Opportunity).

TABLE 1. Description of phases shown in Figure 8.

Phase Type	Description
BOOT	Current SDN settings for network association.
IND	Indicate which <i>opportunity</i> type will follow (if any).
ACK	Acknowledge receptions at the controller.
NACK	Acknowledge no receptions at the controller.
SET	Configure an entry in the SDN flowtable.
ALERT	Alert the controller that an event has been triggered.
SOLICIT	Solicit the controller for instruction.
REPORT	Report state information to the controller.
STOP	End <i>opportunity</i> before the allotted time (Δ_{XOP}).

that opportunity, and distributes the current epoch sequence number. Additionally the IND phase includes a variable length array of mapped Node ID flags. Used in conjunction with the current opportunity type, these flags indicate the role of each node within the flood (*source, destination, or forwarder*).

4) EPOCHS

We define an ‘Epoch’ as the period of time between regularly scheduled SDN control opportunities, with periodicity T_i , where a trade-off is considered when setting the epoch length, and consequently the frequency of SDN opportunities. As synchronous flooding periods in Atomic-SDN inherently block other processes, a longer epoch allows a greater amount of time to be devoted to normal network operation; whether that is application processes, other low-power wireless protocols, or to allow nodes to sleep and therefore conserve energy.

D. SCHEDULING

Atomic-SDN operates a two-stage scheduling process, as highlighted in Figure 9. Firstly, self-contained flood ‘Phases’ are chained together within a short period to allow the construction of higher-level SDN functionality. Then, at a macro level, these flooding periods are scheduled periodically to provide regular SDN ‘opportunities’, as well as maintaining tight time synchronization across all nodes.

1) HIGH-LEVEL ‘OPPORTUNITY’ SCHEDULING

One of the core principles behind Atomic-SDN is the separation in time of control processes from normal

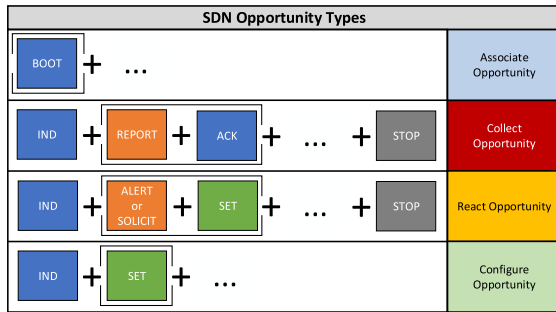


FIGURE 10. Atomic-SDN control *opportunities* built from the phase types defined in Figure 8. Highlighted phases are repeated until the end of the opportunity.

network operation. By slicing control independently from normal operation, the controller is able to define a short period of time in which it is able to communicate with and instruct associated nodes in the local mesh.

With each control period serving a single SDN control function, this necessitates some decision making and scheduling from the controller: choosing what type of control opportunity to initiate at the start of each epoch, and instructing nodes when to quit the control period and resume normal network operation. This scheduling process is shown in Figure 9a which shows a high-level timeline of Atomic-SDN. Once the type of SDN control opportunity is chosen, a mandatory *indicator* (IND) phase is scheduled at the start of the control period. This *one-to-many* phase allows the controller to propagate the opportunity type (if any) to the rest of the network, as well as assigning nodes’ roles (*source*, *destination*, or *forwarder*) and distributing any additional information, such as maximum length of the control period.

2) LOW-LEVEL ‘PHASE’ SCHEDULING

As described in §III-C, each flood is packaged into self-contained ‘phases’ which accomplish specific functions within a larger SDN opportunity. After receiving the IND phase propagated at the start of each opportunity, nodes within the mesh participate in a pre-defined schedule mapped to the SDN opportunity defined within the IND, where the schedule consists of a number of distinct phases of one or more types, and each phase (in and of itself) has its own low-level slot schedule. The phase schedule is configured depending on the the current phase type, where guard times, offsets, and protocol logic are determined from the current node role within the context of the larger SDN opportunity.

E. CHANNEL HOPPING AND NETWORK ASSOCIATION

Network association is achieved through BOOT and IND phases. BOOT phases distribute the current SDN configuration to joining nodes (match/action information for flowtables, what information should be included in *collect* opportunities, etc.). IND phases are scheduled every epoch and, as well as containing SDN opportunity information, allow nodes to re-associate themselves if they have de-synchronized from the network.

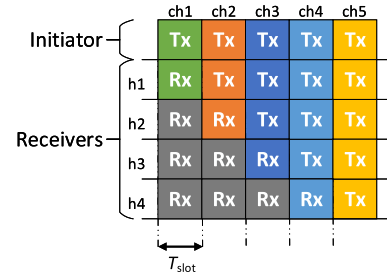


FIGURE 11. Per-slot channel hopping in an Atomic-SDN flood.

Figure 11 shows how Atomic-SDN employs per-slot channel hopping as proposed in [12]. In every IND or BOOT phase the Atomic-SDN controller distributes the current epoch sequence number, which is used to generate a pseudo-random channel hopping sequence for the network. Once known, nodes increment this number every epoch, meaning that if they miss an IND phase due to their duty-cycling or from interference they will retain knowledge of the hopping sequence. At SF primitive slot, nodes concurrently hop to the next channel in this sequence. Known association channels are seeded into every second channel in the sequence (for example, {*ch2*, *ch3*}), so that when a node is trying to join (or re-join) the network, it merely has to listen on one of these known channels for a long enough period until it hears a transmission and can re-synchronize to the controller.

This mechanism, combined with the spatial diversity of SF, allows Atomic-SDN to survive extremely high levels of interference compared to other SDN architectures for low-power wireless networks, providing reliable network control where other solutions would struggle. It has successfully implemented and tested within competition scenarios [9], and is evaluated further in §V.

IV. CHARACTERIZATION OF THE LATENCY BOUNDS

As each SDN opportunity is temporally scheduled, it is possible to characterize Atomic-SDN in terms of the lower latency bounds needed to complete SDN control opportunities (the bounds referenced Figure 9). This is derived through the following information: *the number of participating active nodes (n)*, *the duration of each phase type (t_x)*, and *the duration of the period between phases (t_{IPG})*.

The time taken in each Atomic-SDN control slice is shown in Equation 1, where (Δ_{XOP}) is the time spent on the SDN opportunity (δ_{SDN-OP}) plus a variable sleep period (δ_Z) up to a maximum control bound dictated by the control periodicity.

$$\Delta_{XOP} = \delta_{SDN-OP} + \delta_Z \tag{1}$$

(Δ_{SDN-OP}) varies according to the SDN opportunity being run, and depends on the the underlying SF protocol used to facilitate that function.

A. CONFIGURATION

The configuration opportunity is a *one-to-many* process consisting of an indication phase followed by *n* number of SET

phases and Inter Phase Gaps (IPG), where n depends on the number of individual configuration messages that need to be sent in order to accomplish a specific function within the network. This means that the latency bounds of a configuration opportunity are dependent on the complexity of the function, rather than the scale of the network.

$$\delta_X(n) = t_{IND} + n * (t_{IPG} + t_{SET}) \quad (2)$$

B. COLLECT

The collect opportunity is a *many-to-one* process consisting of an indication phase followed by n number of REP/ACK phase pairs and the IPG, alongside the *stop* overhead of 2 empty REP/ACK phase pairs.

$$\delta_C(n) = t_{IND} + (n + 2) * (t_{IPG} + t_{REP} + t_{ACK}) \quad (3)$$

C. REACT

The react opportunity is a *one-to-one* or *many-to-one* process consisting of an indication phase followed by n number of SOL/SET phase pairs and the IPG, alongside the *stop* overhead of 2 empty SOL/SET phase pairs.

$$\delta_R(n) = t_{IND} + (n + 2) * (t_{IPG} + t_{SOL} + t_{SET}) \quad (4)$$

V. EVALUATION

In this section we evaluate Atomic-SDN and compare its performance against μ SDN [16], a publicly available SDN architecture for IPv6 enabled IEEE 802.15.4 networks; and SDN-WISE [4], a SDN implementation for IEEE 802.15.4 based on the RIME communication stack [37]. Both are implemented in Contiki [38], the same low power Operating System (OS) on which Atomic-SDN is built. Both were chosen evaluation candidates as they can run on top of multiple different IEEE 802.15.4 MAC layers, and therefore provide multiple baselines to compare Atomic-SDN performance. This section demonstrates that, by utilizing SF to create periodic SDN control slices, Atomic-SDN displays considerable performance gains across all metrics in comparison to other SDN architectures for low-power wireless. Furthermore, this mechanism is only possible due to the novel SF framework developed for Atomic-SDN, which allows multiple SF protocols to be configured and instantiated in order to satisfy the plurality of traffic patterns necessary for full SDN control. Figures [12-16] summarize our results.

A. SIMULATION AND TESTBED SETUP

All simulation configuration settings are outlined in Table 2. Simulations were performed using the Cooja simulator and hardware emulator for Contiki OS. Cooja emulates TelosB motes that use the TI MSP430F1611 CPU and CC2420 radio, which is the target platform hardware required by the Atomic-SDN to run the lower layer SF code and is also compatible with SDN-WISE, as well as emulated EXP5438 motes (TI MSP430F5438 CPU and CC2420 radio) used by μ SDN. Additionally, Cooja provides a simulated Multipath Ray-tracer Medium (MRM) radio environment

TABLE 2. Simulation parameters.

Parameter	Setting
Duration	1h
Topology	Grid
Radio Medium	Multipath Ray-tracer Medium (MRM)
Link PRR	~90%
Transmission Range	~300m
SNR Reception Threshold	-100dB
Capture Effect Preamble	64 μ s
Capture Effect Threshold	3dB
SDN Collect Period	60s
SDN Flowtable Lifetime	300s
Application Data Period	60→75s




FIGURE 12. Time taken to complete each Atomic-SDN opportunity as the local mesh scales.

that allows rays to be combined at the receiver (necessary to simulate SF). All simulations were performed across a grid topology, on all nodes, with nodes placed at 300m intervals.

As the Atomic-SDN controller needs to keep track of all network nodes, a maximum of 70 nodes were used in simulations due to the memory constraints of the emulated TelosB hardware. The performance evaluation simulations were run over a 1h period, with an SDN opportunity frequency of 1 second. At the start of each opportunity the type of SDN control scenario (either *collection*, *configuration*, or *reaction*) was set in a round-robin process.

Atomic-SDN is evaluated against μ SDN and SDN-WISE, which were configured to adopt two separate MAC scenarios: firstly using ContikiMAC, an energy saving MAC layer, and secondly using *always-on* Carrier Sense Multiple Access (CSMA). In all simulations the controller collects state information from all nodes every 60s, and node flowtable entries have a 300s lifetime. Additionally, each simulation runs a data collection application where nodes attempt to send application data to a sink node at a variable rate of 60→75s.

Testbed experimentation was performed on a 19 node TelosB testbed. The nodes are located over two floors

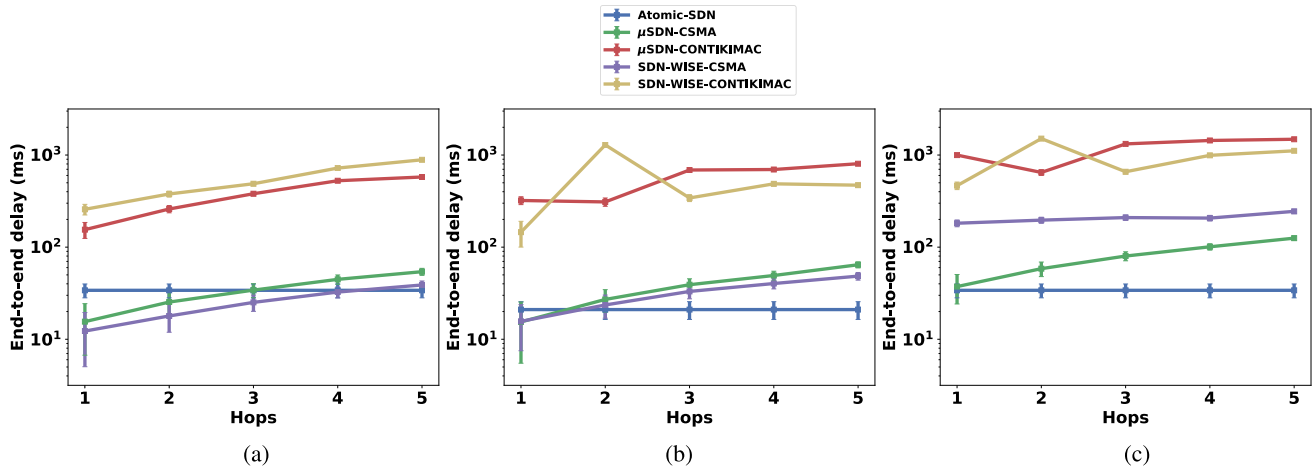


FIGURE 13. Mean collection, configuration, and reaction delays versus hop distance in 30 node network when an individual node participates. Atomic-SDN exhibits similar latencies to CSMA-based μ SDN and SDN-WISE, however it additionally retains consistent delay across all hop counts. (a) Mean collect delay. (b) Mean configure delay. (c) Mean react delay.

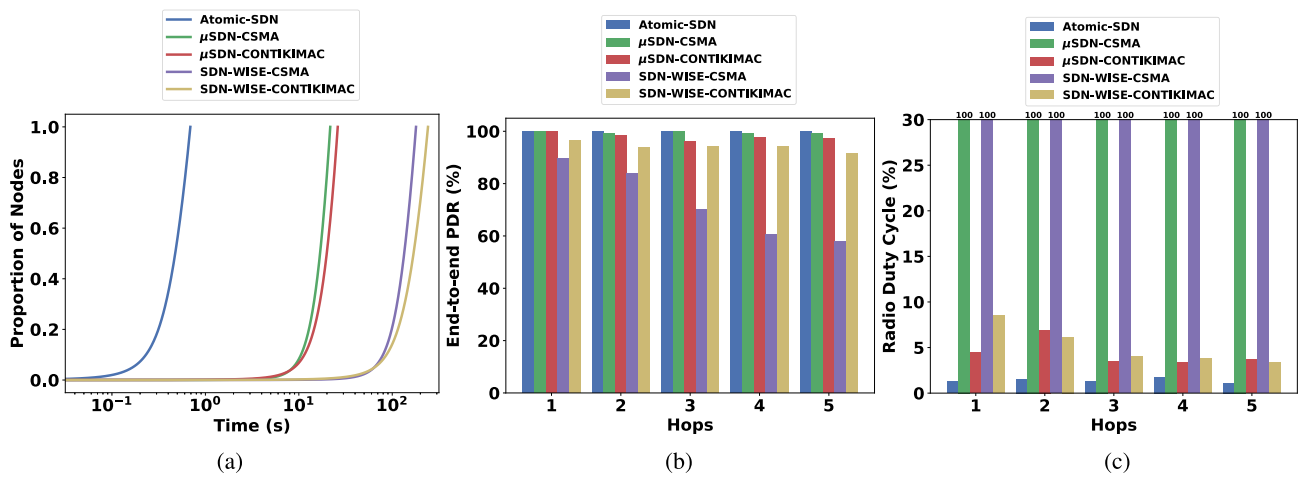


FIGURE 14. (a) Time taken to complete an SDN *react* operation concurrently for all nodes in a 30 node network, as well as (b) end-to-end Packet Delivery Ratio (PDR), and (c) Radio Duty Cycling (RDC) versus hop distance from the controller. As CSMA based μ SDN and SDN-WISE are always-on, they exhibit 100% RDC, denoted at the top of the plot.

and exhibit a mixture of dense clusters as well as remote multi-hop branches. As with the simulations, these experiments were also run over 1h periods. However, the focus was solely on examining Atomic-SDN through analysis and evaluation of the SDN *react* operation in an artificially harsh environment, where reception losses of up to 75% are injected directly into the SF layer.

B. SIMULATION: SDN SCALABILITY IN THE MESH

A key challenge for SDN in IEEE 802.15.4 low-power wireless has been to maintain scalability as the mesh grows from a handful, to hundreds of nodes within a local cluster. As the number of nodes increases, the controller needs to appropriate a greater proportion of limited network resources to SDN control. However, in Atomic-SDN the minimum latency bounds discussed in §IV establish that certain guarantees can be made concerning the resources needed to complete an SDN function when the number of participating nodes is known.

We explore this by evaluating the completion time for the three SDN opportunity types (*collection*, *configuration*, and *reaction*) in networks of increasing size: running simulations for 30, 40, 50, 60, and 70 node mesh networks (limited to 70 nodes due to hardware memory constraints on the Atomic-SDN controller node). Figure 12 shows these results, and demonstrates how the use of SF protocols means that the time to complete all SDN opportunity types (for all network nodes) increases linearly with network size, regardless of hop count.

In each SDN opportunity type we have assumed a worst-case scenario where the controller needs to interact with each node independently. However the number of nodes participating in each SDN opportunity, necessary to fulfill the requirements of higher-level application functions (virtually located at a centralized controller), would likely be a smaller subset of all nodes and therefore incur lower delay.



FIGURE 15. 19 Node TelosB Testbed. Initiating nodes are marked in green, and the SDN controller in blue.




FIGURE 16. PDR overlayed with mean round-trip latency for SDN *react* operations, injected with probabilistic reception misses, in a 19 node testbed. Atomic-SDN maintains high reliability even during 75% injection of receive misses.

The SDN *configure* opportunity is a *one-to-all* process where, after the initial indication phase, each node is configured in turn, in a scheduled fashion. In a 70 node network this allows the configuration of all network nodes within 800ms, assuming each node requires a separate configuration message. However, this time could be substantially reduced (to tens of milliseconds) if a configuration message is relevant to all, or a subset of nodes; allowing the SDN controller to configure multiple nodes in a single flooding phase.

Both the *collect* and *react* opportunities utilize the same underlying SF protocol, and therefore exhibit equivalent delay. In this two phase protocol, the competition between nodes to successfully transmit their data to the controller means that the minimum bound on the completion time is dictated by the number of nodes that need to communicate with the controller. This is a worst-case scenario where it is assumed that all nodes try to perform the SDN operation at exactly the same time, which inevitably causes contention and retransmissions. It does not necessarily follow that this would be the case in a real-world situation. However, these times are still considerably less than the time it takes to complete the same SDN operations in current SDN architectures for low-power wireless networks [4], [5], [39], [40].

Yet, despite Atomic-SDN demonstrating considerable scalability in comparison to other low-power wireless SDN

architectures, there are still questions surrounding the scalability of Concurrent Transmissions when there are 100s or 1000s of nodes and a large number of hops [41]; particularly in extremely dense networks where the number of concurrent transmitters could be significant. Although the authors of [42] propose mechanisms for managing these issues and employing CT based protocols across large networks, there has yet to be experimental evaluation of large-scale SF in real-world scenarios.

C. SIMULATION: PERFORMANCE COMPARISON

We compare Atomic-SDN against current approaches to SDN in IEEE 802.15.4 using μ SDN and SDN-WISE in two configurations: always-on CSMA and the duty-cycled ContikiMAC. Simulations were performed in Cooja using emulated target hardware, and Atomic-SDN performance gains are evaluated in terms of *latency*, *reliability*, and *energy efficiency*. The simulated network was limited to 30 nodes in order to accommodate SDN-WISE which, although specified as able to support longer routing headers, does not have this feature available in the current SDN-WISE Contiki implementation.

1) 'BEST-CASE' - INDIVIDUAL NODE PARTICIPATES

We first consider a 'best-case' scenario where just a single node participates in the SDN control operation (i.e. there is no competing SDN control traffic from other nodes). Table 3 averages these results and Figure 13 shows mean delay versus hop distance from the SDN controller for an individual node performing one of three SDN operations: *collection*, *configuration*, and *reaction*. In each case, Atomic-SDN maintains consistent latency over all distances, due to the minimum bounds on latency inherent in SF protocols. However, as expanded upon in §IV, this bound is affected by the underlying SF protocol which supports it (which can require N number of phases). In this scenario, where there is only one individual node communicating with the controller, it is possible for CSMA based architectures to achieve better latency results than Atomic-SDN at lower hop distances depending on the configuration of the lower layer SF primitive (guard times, number of slots, etc.). However, in CSMA, the radio is always on. Compared to the duty-cycled ContikiMAC configurations of μ SDN and SDN-WISE, Atomic-SDN maintains

TABLE 3. Mean latency, Packet Delivery Ratio (PDR), and Radio Duty Cycle (RDC) for a single node to perform each SDN opportunity type in a 30 node network.

Architecture	Latency (ms)	PDR (%)	RDC (%)
Atomic-SDN	34.0	100.00	1.34
μ SDN-CSMA	33.94	99.58	100
μ SDN-ContikiMAC	340.42	96.45	3.76
SDN-WISE-CSMA	25.75	68.49	100
SDN-WISE-ContikiMAC	544.23	93.84	5.15

minimal latency bound by the length of the flood, regardless of hop distance.

2) 'WORST-CASE' - ALL NODES PARTICIPATE

We next evaluate performance when considering a 'worst-case' scenario where all nodes in the network need to participate in the SDN control operation. These results are presented in Figure 14.

The SDN *react* operation is used to benchmark the time taken for all network nodes to concurrently solicit and receive instruction a single from the controller. Figure 14a shows the effectiveness of SF protocols in comparison current SDN implementations for low-power wireless, where nodes need to perform two or three-way handshakes across multiple Layer 2/3 links. Not only is Atomic-SDN able to perform this operation on all network nodes within milliseconds, but this is orders-of-magnitude faster than the non-SF approaches, which can take seconds or even minutes.

Figure 14b shows Atomic-SDN achieves 100% reliability compared to both the CSMA and ContikiMAC configurations of μ SDN and SDN-WISE. As μ SDN implements end-to-end acknowledgements for SDN control traffic, it presents with a higher Packet Delivery Ratio (PDR) in comparison to SDN-WISE (particularly the CSMA configuration) which has no transportation layer guarantees. Additionally, ContikiMAC causes high channel utilization through packet retransmissions. Although this improves the overall PDR in the case of SDN-WISE, as packets have a higher chance of surviving each hop, it also causes considerable contention on links that experience high traffic loads. This therefore results in a drop in PDR for μ SDN nodes at greater hop distances as they contend with nodes nearer the controller.

Finally, Figure 14c shows the Radio Duty Cycling (RDC) at each hop. ContikiMAC configurations show reduced energy efficiency at lower hop counts as nodes need to serve messages from their children; whilst energy for Atomic-SDN increases at higher hop counts as nodes closer to the controller receive in the first few transmission slots on an ACK, and so spend less time participating in the flood. CSMA configurations, which don't perform any duty-cycling, display high PDR across all hops. Barring any contention, nodes should always be able to receive transmissions as the radio is always on. In comparison, the use of SF in Atomic-SDN means it can benefit from a highly reliable MAC, whilst retaining the low-energy operation of duty-cycled approaches. By using the APB to construct multiple SF protocols tailored for each SDN task Atomic-SDN demonstrates near perfect reliability when collecting state information from all nodes. Although these results are based on simulation and 100% reliability cannot be guaranteed, multiple studies and extensive experimental evaluations (as covered in a recent survey [25]) have demonstrated up to 99.99% reliability is achievable using SF protocols, even under heavy interference.

D. TESTBED: SDN CONTROL UNDER INTERFERENCE

A 19 node testbed was used to evaluate the potential of Atomic-SDN to provide reliable SDN control in high interference scenarios, reaching all mesh nodes irrespective of link quality. The layout of this testbed is shown in Figure 15: nodes are located over two floors and it presents a number of interesting features such as a dense cluster, isolated multi-hop paths, and non Line-of-Sight (LOS) lossy links.

The SDN controller was configured to repeatedly initiate an SDN *reaction* operation from all nodes with minimal periodicity (i.e. back-to-back) over 1h periods. Forces reception (R_x) drops were injected into the SF layer at each node in order to simulate the effects of interference, stepping from 0%, 25%, 50% to 75% probability. Latency and PDR results for each drop rate are presented in Figure 16, and show how Atomic-SDN is able to maintain near 100% probability even with 75% reception misses. However, it shows that as the probability of missing a reception increases, this also increases the average time taken for each node to complete the SDN operation. As receptions are missed, the underlying SF protocol ensures nodes will re-transmit their controller solicitation data until they hear an acknowledgement.

Not only do these testbed results demonstrate the resilience of the SF approach to SDN control pioneered in Atomic-SDN (and also shown through success at the IEEE EWSN Dependability Competition), but overcoming this issue of resilience and reliability is particularly relevant in both Advanced Metering Infrastructure (AMI) and industrial control scenarios: where nodes may suffer from lossy links due to harsh environments, or multipath effects due to surrounding industrial machinery and assets. By maintaining a reliable and low-latency link to an SDN controller, based on a stateless and topology agnostic control mechanism, Atomic-SDN can provide programmable SDN control without assuming network stability.

VI. CONCLUSION

Attempts to apply SDN concepts within low-power wireless have consistently been met with the same problem: that implementing centralised control architectures for constrained, low-power mesh networks generates complex and considerable overhead that the underlying physical and MAC layers struggle to handle. A number of approaches have managed to reduce this overhead, but a truly responsive and dynamic SDN architecture needs to be able maintain a timely view of the network state, handle network changes within milliseconds rather than seconds or minutes, and be resilient to interference. This is particularly relevant if SDN research in low-power wireless is to move beyond sensor networks and start addressing other aspects of IoT such as Industrial Wireless Control for Cyber Physical Systems, where strict latency and reliability requirements cannot be met by current SDN control solutions.

Synchronous Flooding, a radically different approach compared to current standards, can provide this. This paper

has introduced Atomic-SDN, a unique solution for SDN in low-power wireless networks that utilizes SF to provide highly-reliable SDN control with minimal latency and extremely high reliability, without affecting application-layer traffic or other control processes. By facilitating the propagation of control messages across the network in a flood, within dedicated control timeslots, Atomic-SDN allows the SDN layer to operate without knowledge of topology, as well as benefiting from the spatial and temporal diversity inherent within flooding protocols. Furthermore, Atomic-SDN has demonstrated resilience to extremely high levels of interference, not only through the testbed experimentation in this paper, but also in extensive benchmarking against WiFi interference at the 2019 IEEE EWSN Dependability Competition.

REFERENCES

- [1] D. Kreutz, F. M. V. Ramos, P. E. Verissimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig, "Software-defined networking: A comprehensive survey," *Proc. IEEE*, vol. 103, no. 1, pp. 14–76, Jan. 2015.
- [2] K. Sood, S. Yu, and Y. Xiang, "Software-defined wireless networking opportunities and challenges for Internet-of-Things: A review," *IEEE Internet Things J.*, vol. 3, no. 4, pp. 453–463, Aug. 2016.
- [3] S. Costanzo, L. Galluccio, G. Morabito, and S. Palazzo, "Software defined wireless networks: Unbridling SDNs," in *Proc. Eur. Workshop Softw. Defined Netw.*, Oct. 2012, pp. 1–6.
- [4] L. Galluccio, S. Milardo, G. Morabito, and S. Palazzo, "SDN-WISE: Design, prototyping and experimentation of a stateful SDN solution for Wireless Sensor networks," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, Apr./May 2015, pp. 513–521.
- [5] T. Theodorou and L. Mamas, "CORAL-SDN: A software-defined networking solution for the Internet of Things," in *Proc. IEEE Conf. New. Function Virtualization Softw. Defined Netw. (NFV-SDN)*, Nov. 2017, pp. 1–2.
- [6] M. Baddeley, R. Nejabati, G. Oikonomou, S. Gormus, M. Sooriyabandara, and D. Simeonidou, "Isolating SDN control traffic with layer-2 slicing in 6TiSCH industrial IoT networks," in *Proc. IEEE Conf. New. Function Virtualization Softw. Defined Netw. (NFV-SDN)*, Nov. 2017, pp. 247–251.
- [7] U. Raza, Y. Jin, A. Stanoev, M. Baddeley, and M. Sooriyabandara, "Competition: CROWN-concurrent receptions in wireless sensor and actuator networks," in *Proc. Int. Conf. Embedded Wireless Syst. Netw.*, Feb. 2018, pp. 223–224.
- [8] U. Raza, Y. Jin, and M. Sooriyabandara, "Competition: Synchronous transmissions based flooding for dependable Internet of Things," in *Proc. Int. Conf. Embedded Wireless Syst. Netw.*, Feb. 2017, pp. 278–279.
- [9] M. Baddeley, A. Stanoev, U. Raza, Y. Jin, and M. Sooriyabandara, "Competition: Adaptive software defined scheduling of low power wireless networks," in *Proc. Int. Conf. Embedded Wireless Syst. Netw.*, Feb. 2019, pp. 1–2.
- [10] A. Escobar, F. Moreno, B. Saez, A. J. Cabrera, J. Garcia-Jimenez, F. J. Cruz, U. Ruiz, A. Corona, J. Klaue, and D. Tati, "Competition: Big-bangbus," in *Proc. Int. Conf. Embedded Wireless Syst. Netw.*, Feb. 2018, pp. 213–214.
- [11] M. Trobinger, T. Istomin, A. L. Murphy, and G. P. Picco, "Competition: Crystal clear: Making interference transparent," in *Proc. Int. Conf. Embedded Wireless Syst. Netw.*, Feb. 2018, pp. 217–218.
- [12] R. Lim, R. Da Forno, F. Sutton, and L. Thiele, "Competition: Robust flooding using back-to-back synchronous transmissions with channel-hopping," in *Proc. Eur. Conf. Wireless Sensor Netw. (EWSN)*, Feb. 2017, pp. 270–271.
- [13] A. Escobar, J. Garcia-Jimenez, F. J. Cruz, J. Klaue, A. Corona, and D. Tati, "Competition: Redfixhop with channel hopping," in *Proc. Int. Conf. Embedded Wireless Syst. Netw.*, Feb. 2017, pp. 264–265.
- [14] B. A. Nahas and O. Landsiedel, "Competition: Towards low-power wireless networking that survives interference with minimal latency," in *Proc. Eur. Conf. Wireless Sensor Netw. (EWSN)*, Feb. 2017, pp. 268–269.
- [15] O. Bergmann, C. Bormann, S. Gerdes, and H. Chen, *Constrained-Cast: Source-Routed Multicast for RPL*, document draft-ietf-roll-ccast-01, Internet-Draft, Internet Engineering Task Force, Oct. 2017.
- [16] M. Baddeley, *μSDN: A Low-Overhead SDN Stack and Embedded SDN Controller for Contiki*. [Online]. Available: <https://github.com/mbaddeley/usdn>
- [17] S. Labs. (2018). *Bluetooth Mesh Network Performance*. [Online]. Available: <https://www.silabs.com/documents/login/application-notes/an1137-bluetooth-mesh-network-performance.pdf>
- [18] I. T. Haque and N. Abu-Ghazaleh, "Wireless software defined networking: A survey and taxonomy," *IEEE Commun. Surveys Tuts.*, vol. 18, no. 4, pp. 2713–2737, 2nd Quart., 2016.
- [19] S. Bera, S. Misra, and A. V. Vasilakos, "Software-defined networking for Internet of Things: A survey," *IEEE Internet Things J.*, vol. 4, no. 6, pp. 1994–2008, Dec. 2017.
- [20] T. Luo, H.-P. Tan, and T. Q. S. Quek, "Sensor openflow: Enabling software-defined wireless sensor networks," *IEEE Commun. Lett.*, vol. 16, no. 11, pp. 1896–1899, Nov. 2012.
- [21] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "Openflow: Enabling innovation in campus networks," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 69–74, Apr. 2008.
- [22] T. Luo, H.-P. Tan, P. C. Quan, Y. W. Law, and J. Jin, "Enhancing responsiveness and scalability for openflow networks via control-message quenching," in *Proc. Int. Conf. ICT Converg. (ICTC)*, Oct. 2012, pp. 348–353.
- [23] H. Song, "Protocol-oblivious forwarding: Unleash the power of sdn through a future-proof forwarding plane," in *Proc. 2nd ACM SIGCOMM Workshop Hot Topics Softw. Defined Netw.*, Aug. 2013, pp. 127–132.
- [24] P. Thubert, *An Architecture for IPv6 Over the TSCH Mode of IEEE 802.15.4*, document draft-ietf-6tisch-architecture-20, Internet-Draft, Internet Engineering Task Force, Mar. 2019.
- [25] T. Chang, T. Watteyne, X. Vilajosana, and P. H. Gomes, "Constructive interference in 802.15.4: A tutorial," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 1, pp. 217–237, 1st Quart., 2018.
- [26] F. Ferrari, M. Zimmerling, L. Thiele, and O. Saukh, "Efficient network flooding and time synchronization with glossy," in *Proc. 10th ACM/IEEE Int. Conf. Inf. Process. Sensor Netw.*, Apr. 2011, pp. 73–84.
- [27] V. S. Rao, M. Koppal, R. V. Prasad, T. V. Prabhakar, C. Sarkar, and I. Niemegeers, "Murphy loves CI: Unfolding and improving constructive interference in WSNs," in *Proc. IEEE 35th Annu. IEEE Int. Conf. Comput. Commun. (INFOCOM)*, Apr. 2016, pp. 1–9.
- [28] C. H. Liao, Y. Katsumata, M. Suzuki, and H. Morikawa, "Revisiting the so-called constructive interference in concurrent transmission," in *Proc. IEEE 41st Conf. Local Comput. Netw. (LCN)*, Nov. 2016, pp. 280–288.
- [29] A. Escobar-Molero, "Improving reliability and latency of wireless sensor networks using concurrent transmissions," *At-Automatisierungstechnik*, vol. 67, no. 1, pp. 42–50, Jan. 2019.
- [30] K. Leentvaar and J. Flint, "The capture effect in fm receivers," *IEEE Trans. Commun.*, vol. COMM-24, no. 5, pp. 531–539, May 1976.
- [31] T. Istomin, A. L. Murphy, G. P. Picco, and U. Raza, "Data prediction + synchronous transmissions = ultra-low power wireless sensor networks," in *Proc. 14th ACM Conf. Embedded Netw. Sensor Syst. CD-ROM*, Nov. 2016, pp. 83–95.
- [32] B. A. Nahas, S. Duquennoy, and O. Landsiedel, "Concurrent transmissions for multi-hop bluetooth 5," in *Proc. Int. Conf. Embedded Wireless Syst. Netw.*, 2019, pp. 1–12.
- [33] O. Landsiedel, F. Ferrari, and M. Zimmerling, "Chaos: Versatile and efficient all-to-all data sharing and in-network processing at scale," in *Proc. 11th ACM Conf. Embedded Networked Sensor Syst.*, Nov. 2013, Art. no. 1.
- [34] F. Ferrari, M. Zimmerling, L. Mottola, and L. Thiele, "Low-power wireless bus," in *Proc. 10th ACM Conf. Embedded Netw. Sensor Syst.*, Nov. 2012, pp. 1–14.
- [35] B. A. Nahas, S. Duquennoy, and O. Landsiedel, "Network-wide consensus utilizing the capture effect in low-power wireless networks," in *Proc. Int. Conf. Embedded Networked Sensor Syst.*, Nov. 2017, Art. no. 1.
- [36] T. Istomin, M. Trobinger, A. L. Murphy, and G. P. Picco, "Interference-resilient ultra-low power aperiodic data collection," in *Proc. 17th ACM/IEEE Int. Conf. Inf. Process. Sensor Netw.*, Apr. 2018, pp. 84–95.
- [37] A. Dunkels, "RIME—A lightweight layered communication stack for sensor networks," in *Proc. Eur. Conf. Wireless Sensor Netw. (EWSN)*, Jan. 2007. [Online]. Available: <https://scholar.google.com/scholar?q=Rime%20-%20a%20lightweight%20layered%20communication%20stack%20for%20sensor%20networks>

- [38] A. Dunkels, B. Gronvall, and T. Voigt, "Contiki—A lightweight and flexible operating system for tiny networked sensor," in *Proc. 29th IEEE Annu. Int. Conf. Local Comput. Netw.*, Nov. 2004, pp. 455–462.
- [39] E. Muncio, J. Marquez-Barja, S. Latré, and S. Vissicchio, "Whisper: Programmable and flexible control on industrial IoT networks," *Sensors*, vol. 18, no. 11, p. 4048, Nov. 2018.
- [40] M. Baddeley, R. Nejabati, G. Oikonomou, M. Sooriyabandara, and D. Simeonidou, "Evolving SDN for low-power IoT networks," in *Proc. 4th IEEE Conf. Netw. Softwarization Workshops (NetSoft)*, Jun. 2018, pp. 71–79.
- [41] C. Noda and C. M. Pérez-Penichet, B. Seeber, M. Zennaro, M. Alves, and A. Moreira, "On the scalability of constructive interference in low-power wireless networks," in *Proc. Eur. Conf. Wireless Sensor Netw.*, 2015, pp. 250–257.
- [42] Y. Wang, Y. He, X. Mao, Y. Liu, and X. Li, "Exploiting constructive interference for scalable flooding in wireless networks," *IEEE/ACM Trans. Netw.*, vol. 21, no. 6, pp. 1880–1889, Dec. 2013.



of Things, and cyber physical systems.

MICHAEL BADDELEY received the M.Eng. degree in computer and electronic systems from the University of Strathclyde, in 2010. He is currently pursuing the Ph.D. degree with the University of Bristol. He worked for a number of years as a Research and Software Engineer for Thales, both in U.K. and in Germany. He is also a Research Engineer for Toshiba Research Europe Ltd., Bristol, U.K. His current research interests include software defined networking, the Internet



of Things, and cyber physical systems.

USMAN RAZA received the Ph.D. degree in information and communications technologies from the University of Trento, Italy. He is currently a Principal Research Engineer with Toshiba Research Europe Ltd., Bristol, UK. During his multiple roles in both academia and industry, he has filed patents, published papers and book chapters, participated as an Invited Panelist and a Speaker, and chaired scientific events. His current research interests include cyber-physical systems, industrial wireless systems, and low power wide area networks. He was a recipient of the Endeavour Research Fellowship at the University of New South Wales, Australia, and an offer of the William J. Fulbright Scholarship from the U.S. Department of State. He received the IEEE Communications Society Heinrich Hertz Award, in 2019, the Mark Weiser Best Paper Award, in 2012, and the SenseApp best paper award, in 2014. He has successfully bid for very competitive and prestigious European Horizon 2020 projects.

of Things, and cyber physical systems.



ALEKSANDAR STANOEV received the M.Eng. degree (Hons.) from the University of Bristol, in 2018. He is currently a Research Engineer with Toshiba Research Europe Ltd., Bristol, U.K. His current research interests include low-power wireless MAC protocols and the embedded IoT hardware design.



of research that led to one of the largest smart meter wireless mesh networks in the world. He is currently the Managing Director of Toshiba Research Europe Ltd., Bristol, U.K., and also an Honorary Visiting Professor with the School of Engineering, Cardiff University. He has more than 80 peer-reviewed publications, several book chapters, and holds more than 20 patent applications. He is involved in standardization activities within IETF, IEEE, and ETSI, and is a Fellow of IET.

GEORGE OIKONOMOU has ported Contiki/Contiki-NG for numerous hardware platforms. He is currently a Lecturer with the Electrical and Electronic Engineering Department, University of Bristol. He is also the Developer of Contiki-NG's 6LoWPAN multicast forwarding support. His current research interests include energy-efficient technologies for networks of severely constrained, wireless embedded devices and the IoT. He has a strong track record of delivering deployable and high-TRL research outputs. He is a Co-Founder, a Steering Group Member, and a Maintainer of the Contiki-NG open-source embedded operating system for the Internet of Things.



of research that led to one of the largest smart meter wireless mesh networks in the world. He is currently the Managing Director of Toshiba Research Europe Ltd., Bristol, U.K., and also an Honorary Visiting Professor with the School of Engineering, Cardiff University. He has more than 80 peer-reviewed publications, several book chapters, and holds more than 20 patent applications. He is involved in standardization activities within IETF, IEEE, and ETSI, and is a Fellow of IET.

REZA NEJABATI received the M.Sc. degree (Hons.) in telecommunication and information systems and the Ph.D. degree in electronic systems engineering from the University of Essex. He joined the University of Bristol from the University of Essex, where he was a Lecturer with the School of Computer Science and Electronic Engineering. He has made pioneering contributions to the fields of optical networking, data center networks, software defined networking, network virtualization, and network function virtualization. He has authored or coauthored more than 200 peer reviewed papers and several standardization documents. His current research interest includes disruptive new Internet technologies with focus on high-speed networks. He is a Fellow of the U.K. Higher Education Academy.



of research that led to one of the largest smart meter wireless mesh networks in the world. He is currently the Managing Director of Toshiba Research Europe Ltd., Bristol, U.K., and also an Honorary Visiting Professor with the School of Engineering, Cardiff University. He has more than 80 peer-reviewed publications, several book chapters, and holds more than 20 patent applications. He is involved in standardization activities within IETF, IEEE, and ETSI, and is a Fellow of IET.

MAHESH SOORIYABANDARA received the B.Sc. degree in electronic and electrical engineering from the University of Peradeniya, Sri Lanka, and the Ph.D. degree in engineering from the University of Aberdeen, U.K. In 2012, he received the Toshiba Research Europe's Distinguished Researcher Award and Toshiba Corporate R&D Prize for his contribution to smart grid communications. In 2013, he received Toshiba Corporation's Corporate Technology Prize in recognition



of research that led to one of the largest smart meter wireless mesh networks in the world. He is currently the Managing Director of Toshiba Research Europe Ltd., Bristol, U.K., and also an Honorary Visiting Professor with the School of Engineering, Cardiff University. He has more than 80 peer-reviewed publications, several book chapters, and holds more than 20 patent applications. He is involved in standardization activities within IETF, IEEE, and ETSI, and is a Fellow of IET.

DIMITRA SIMEONIDOU is currently a Professor with the High Performance Networks Group (HPN), University of Bristol, Bristol, U.K., where she moved with her team of 20 researchers, in 2012 after 15 years with the University of Essex, Colchester, U.K. She has authored or coauthored more than 400 publications, many of which have received best paper awards, 11 patents, and several major standards contributions. She has led numerous collaborative research projects, and is currently managing a portfolio of 16 active projects funded by EPSRC, EC, and industry, including large transnational collaborations with the USA, Brazil, and Japan. She has chaired a number of international conferences (including a TPC Chair of the European Conference on Optical Communication 2013 and a Subcommittee Chair of the Optical Fiber Communication Conference 2015) and committees across these technical fields.

...