

Connecting Low-Power and Lossy Networks to the Internet

JeongGil Ko and Andreas Terzis, Johns Hopkins University

Stephen Dawson-Haggerty and David E. Culler, University of California at Berkeley

Jonathan W. Hui, Cisco Systems, Inc.

Philip Levis, Stanford University

ABSTRACT

Many applications, ranging from wireless healthcare to energy metering on the smart grid, have emerged from a decade of research in wireless sensor networks. However, the lack of an IP-based network architecture precluded sensor networks from interoperating with the Internet, limiting their real-world impact. Given this disconnect, the IETF chartered the 6LoWPAN and RoLL working groups to specify standards at various layers of the protocol stack with the goal of connecting low-power and lossy networks to the Internet. We present the standards proposed by these working groups, and describe how the research community actively participates in this process by influencing their design and providing open source implementations.

VISION OF LOW-POWER AND LOSSY NETWORKS

For nearly a decade, wireless sensor network research and development largely assumed that the Internet architecture was ill suited for sensor network applications. Many in the field argued that Internet protocols were impractical for the resource constrained devices that were being embedded in the physical world; that the end-to-end architecture was inappropriate for the localized algorithms and in-network processing required to achieve robustness and scalability; and that an architecture designed to accommodate a wide range of applications was unnecessary as sensor networks would be tailored to specific target applications.

As a result, academic and commercial efforts during this period did not assume the traditional interfaces and layers of network architecture. This freedom to develop new protocols without regard to an existing architecture allowed sensor network researchers to quickly invent and evaluate new protocols, as well as gain experience of their use in the field. As the community progressed, new network abstractions, such as routing

state distribution and collection-based routing, emerged [1].

Ironically, this freedom of thought meant that innumerable good ideas were sequestered in disjoint, non-interoperable stovepipes with little opportunity to impact the real world. Furthermore, without an IP-based architecture, these systems require an application-layer gateway to communicate with other systems or even the wider Internet. Application gateways are complex to design and manage since they perform significant functional and semantic translation and carry application-layer state.

In the past couple of years, the push to develop and deploy the smart grid has placed a pressing need to deploy networks that have an unprecedented scale when compared to existing IP networks (e.g., 10 million endpoints within a single network targeted by energy utilities), allow multivendor interoperability, and utilize low-cost communication devices. Sensor networks are considered a natural fit for smart grid applications such as automated metering infrastructure (AMI) and home area networking (HAN).

To address this need, the Internet Engineering Task Force (IETF) created Working Groups (WGs) to standardize protocols for constrained networks.

6LoWPAN: The IPv6 in Low-Power Wireless Personal Area Networks Working Group was chartered to standardize necessary adaptations of IPv6 for networks that use the IEEE 802.15.4 physical (PHY) layer, and has defined how to carry IPv6 datagrams over IEEE 802.15.4 links and perform necessary configuration functions to form and maintain an IPv6 subnet.

RoLL: IETF chartered the Routing over Low-Power and Lossy Links Working Group to standardize a link-independent IPv6 routing protocol for resource constrained devices. The Working Group has defined the RPL routing protocol, which builds a robust topology over lossy links with minimal state requirements.

In this article we introduce the standardization work proposed by the two WGs discussed

above. We also present how the findings from a decade of research in wireless sensor networks have influenced the standardization work of the IETF, and finally describe efforts from the research community to integrate these standards in sensor network research.

6LOWPAN: IPv6 OVER LOW-POWER WIRELESS PERSONAL AREA NETWORKS

IPv6 is in many ways a good fit for sensing and monitoring applications: its nearly infinite address space enables a future with ever more ubiquitous computation. Moreover, integration with the Internet architecture promotes flexible, principled optimization and innovation at each layer in the stack.

In support of this vision, the IETF chartered in 2005 the 6LoWPAN working group to standardize the use of IPv6 over IEEE 802.15.4 radios [2], which have vastly different characteristics from previous link technologies such as Ethernet or WiFi. For example, the maximum transmission unit (MTU) of 802.15.4 is only 127 bytes, far smaller than the 1500 bytes supported by Ethernet and WiFi. Additionally, 802.15.4 does not provide a full-broadcast domain where all nodes are capable of receiving messages from all other nodes using a single physical transmission. Instead, 802.15.4 networks are composed of overlapping broadcast domains, where a radio's neighbor set is defined by those nodes reachable with a single transmission. To address these issues, the 6LoWPAN Working Group has focused on two work items:

- How to carry IPv6 datagrams in 802.15.4 frames
- How to perform necessary IPv6 neighbor discovery functions (e.g., address resolution, duplicate address detection) in a network with overlapping broadcast domains

The remainder of this section presents the respective techniques in detail.

FRAGMENTATION

The 6LoWPAN WG first addressed how to carry IPv6 datagrams in IEEE 802.15.4 frames and published RFC 4944 in 2008 [3]. IPv6 specifies that the link must support an MTU of no less than 1280 bytes, requiring 6LoWPAN to define a link-layer fragmentation and reassembly mechanism. The fragmentation mechanism is simple and only provides the ability to encode a datagram using multiple link frames. It does not include end-to-end recovery of lost fragments, expecting that link-layer acknowledgments will provide sufficient delivery success rates. The fragmentation header is 4–5 bytes and contains three fields: *datagram size*, indicating the size of the datagram being fragmented, *datagram tag*, identifying all fragments for a particular datagram, and *datagram offset*, which indicates the fragment's location within the datagram. By including the datagram size in all fragments, a node can immediately allocate an appropriate reassembly buffer even if fragments are delivered out of order.

HEADER COMPRESSION

Traditional header compression techniques used in IP networks are flow-based, observing which portions of the header change rarely across packets within a flow and eliding those portions when possible. Flow-based techniques are commonly used over point-to-point links, and work well because the same compressor and decompressor is used for the lifetime of the flow. In sensor networks, however, the path of a flow may change frequently due to time-varying dynamics of low-power wireless communications. If traditional flow-based techniques were used, changing a next-hop route would require migrating compression state to the new route, limiting the effectiveness of header compression.

As a result, the WG developed a new compression format that does not require per-flow state [4]. This new format statelessly compacts headers in two ways: first, by removing redundant information across the link, network, and transport layers;¹ second, by assuming common values for header fields and defining compact forms of those values.²

The header compression format also allows stateful compression for arbitrary IPv6 address prefixes and is effective for two reasons. First, nodes within the sensor network's subnet all share the same subnet ID. Second, many applications involve all nodes communicating with a single common destination (e.g., data collection traffic). In both cases, prefix information used for compression is common to all nodes and not specific to individual flows. To support stateful compression, each node maintains a context table, each entry containing an IPv6 prefix. Nodes configure the context table while registering themselves using the 6LoWPAN Neighbor Discovery protocol, presented in the following subsection.

NEIGHBOR DISCOVERY AND AUTO-CONFIGURATION

The IPv6 Neighbor Discovery (ND) protocol defined in RFC 4861 allows nodes to perform address resolution, discover neighboring routers, determine unreachability, and perform duplicate address detection [5]. The challenge with the existing IPv6 ND protocol is that it utilizes multicast communication and assumes that the link provides a single broadcast domain. Because the IEEE 802.15.4 link layer does not support multicast, link-layer broadcast is used to deliver IPv6 multicast messages. Broadcast transmissions are unreliable because they cannot be acknowledged at the link layer. Furthermore, broadcast transmissions can be expensive, especially when nodes duty-cycle their radios to reduce power. As a result, the 6LoWPAN WG set out to optimize IPv6 ND by minimizing its reliance on multicast [6].

Like IPv6 ND, 6LoWPAN ND also uses Router Advertisement (RA) and Router Solicitation (RS) messages to allow nodes to discover neighboring routers. RS messages are sent using multicast, but this is the only use of multicast within 6LoWPAN ND. Two key assumptions are made to support this: first, link-local IPv6 addresses are derived from the link-layer (medi-

The header compression format also allows stateful compression for arbitrary IPv6 address prefixes and is effective for two reasons. First, nodes within the sensor network's subnet all share the same subnet ID. Second, many applications involve all nodes communicating with a single common destination.

¹ IPv6 payload length is derived from lower layers, the Interface ID of IPv6 addresses may be derived from link-layer addresses, and IPv6 version is always 6 by virtue of using 6LoWPAN.

² Traffic class and flow label are 0; hop limit is 1, 64, or 255; link-local IPv6 addresses are common; and commonly used multicast addresses only differ in a few bytes.

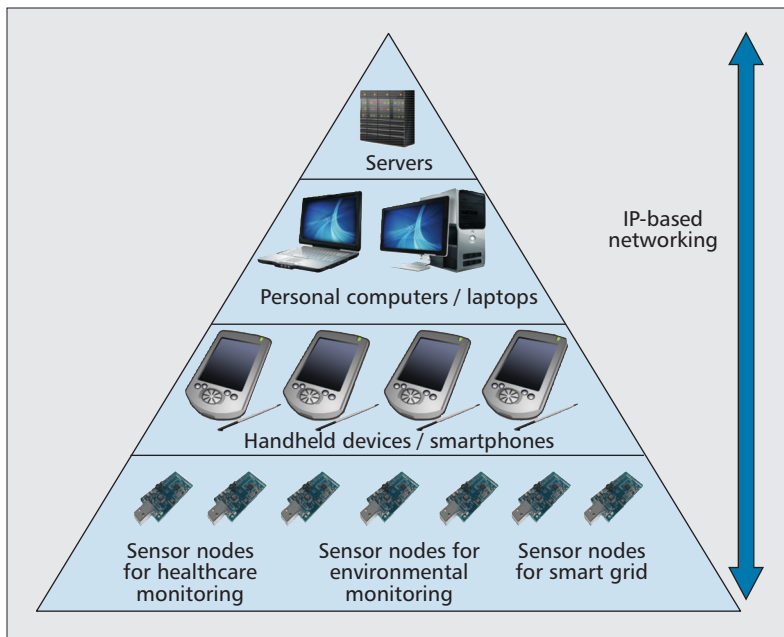


Figure 1. Low-power and lossy networks connect the Internet of Things to the existing IP-based network architecture.

um access control, MAC) addresses, so 6LoWPAN nodes do not need to use multicast Neighbor Solicitation (NS) messages to resolve an unknown link-layer address; the address can be constructed directly from the IPv6 address. Second, addresses other than link-local addresses are assumed not to be on-link; communication with these nodes must proceed through a router. As a result of these simplifications, 6LoWPAN nodes do not need to cache packets or store per-neighbor state while on-link determination and address resolution proceeds; a packet is either sent directly to the destination (if it is link-local) or otherwise to a router.

The address registration mechanism conveniently provides address unreachability detection and duplicate address detection. Whereas IPv6 ND treats the neighbor table as a cache, nodes using 6LoWPAN ND must treat the neighbor table as a registry. Because address registration requests are acknowledged, a node can infer unreachability from the lack of acknowledgments. Finally, routers can optionally perform duplicate address detection by forwarding address registration requests to a central server to verify whether or not the address is already in use.

RPL: IPV6 ROUTING PROTOCOL FOR LOW-POWER AND LOSSY NETWORKS

While the work from the 6LoWPAN WG opened the possibility of using IPv6 in IEEE 802.15.4 networks, standardizing a routing protocol was outside the scope of that working group. This led to the creation of the Routing over Low-power and Lossy network (RoLL) working group in 2008. As its first goal, the RoLL working group set out to study the routing requirements for a broad range of applications. These requirements were the distillation from a decade

of academic research in wireless sensor networks and experience in deploying such networks. This process resulted in four RFCs describing routing requirements for four different target environments: RFC 5548 for urban environments [7], RFC 5673 for industrial low-power networks [8], RFC 5826 for home automation [9], and RFC 5867 for building automation [10]. The documents include requirements for latency, reliability, and network management capabilities.

The WG defined their target network architecture as a “low-power and lossy network (LLN),” defined by a set of unique characteristics. For example, LLN devices can maintain very little state and are optimized for saving energy. Furthermore, traffic patterns in LLNs can be more complex than unicast flows and finally, such traffic must be carried over links that can forward only small frames. Since these LLN characteristics impose unique routing requirements existing protocols may be inappropriate for such networks. To test their suitability in this context, the WG evaluated well-known routing protocols such as Open Shortest Path First (OSPF), Intermediate System to Intermediate System (IS-IS), Ad Hoc On Demand Vector (AODV), and Optimized Link State Routing (OLSR), against a set of core metrics: routing table scalability, packet loss response, control overhead, and the ability to impose routing constraints. The WG concluded that no current IETF routing protocol could provide acceptable performance in the unique conditions that characterize LLNs. As a result, the WG developed the IPv6 Routing Protocol for LLNs, known as RPL [11].

RPL supports three traffic patterns: multi-point-to-point traffic (MP2P), point-to-multi-point traffic (P2MP), and point-to-point (P2P) traffic

RPL’s design adopted results from the wireless sensor network research community’s decade of research. Next, we present the key mechanisms which heavily influenced the design of the RPL protocol.

Routing state propagation: Conventional link-state routing protocols propagate state using scoped floods; most existing distance-vector protocols typically use periodic routing beacons to maintain freshness. However, both mechanisms have inherent limitations which become especially problematic in large, low-bandwidth networks. Instead, RPL uses the Trickle [9] algorithm for scalable state propagation. At its core, Trickle is an epidemic protocol which reacts quickly to changes in routing state and tapers off as the rate of state changes decreases.

Spatial diversity: In an environment where nodes are likely to fail due to environmental causes and interference can cause previously good links to rapidly become unusable, nodes should be able to quickly switch away from unavailable routes. Thus, RPL makes use of *spatial diversity* as a technique to achieve reliability: a router can maintain multiple potential parents towards a destination instead of a single one.

Expressive link and node metrics: Existing routing protocols define static link costs; the routing protocol then responds to links breaking or becoming available. On the other hand, LLN links exhibit

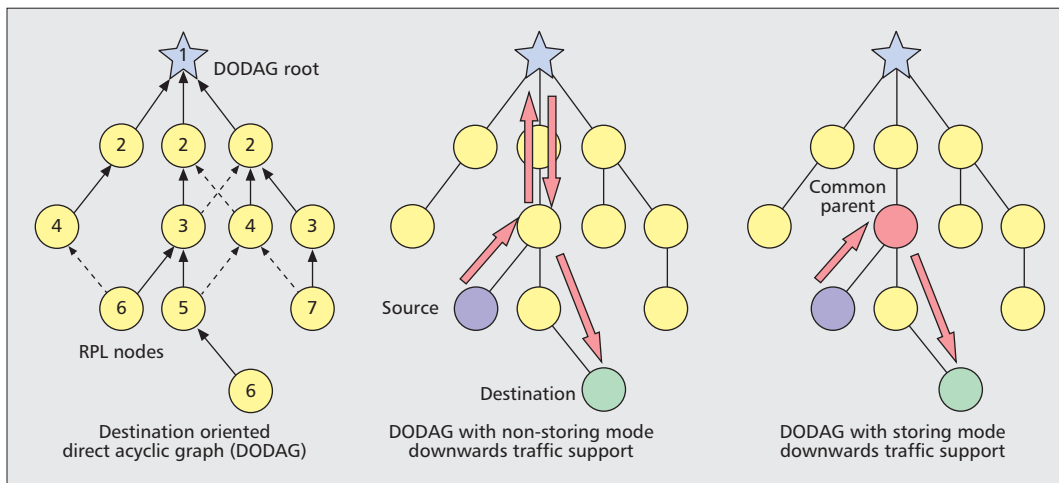


Figure 2. RPL nodes that form a directed acyclic graph (DAG) rooted at a destination node support multi-point-to-point traffic (left). The solid arrows point to a node's preferred parent; the dotted arrows point to other nodes included in the parent set. Sample rank values for each node are also presented. The DODAG, also supports point-to-multipoint traffic and point-to-point traffic in either a non-storing mode (center), in which the root attaches source routing headers to data packets, or storing mode (right), in which each node maintains the routing state for its descendants.

significant temporal variation across multiple dimensions (e.g., reliability). In order to represent this variability, RPL includes a flexible framework for incorporating dynamic link metrics such as ETX (Estimated number of Transmissions) [12].

Of the three traffic patterns, multipoint-to-point traffic is the dominant traffic type for current LLN applications: typically, a small set of destinations act as edge routers connecting the LLN to the Internet. To support this traffic pattern, RPL constructs a destination oriented directed acyclic graph (DODAG) and uses this graph to route data traffic. This approach was based on a broad literature in *Collection routing* for sensor networks [13, 14] and the fact that tree routing would result in minimal routing states kept at the LLN nodes. One of the mechanisms that RPL borrows is the Trickle algorithm to regulate the transmission of DODAG information objects (DIOs), used to maintain the DODAG. DIO messages include objective functions (OFs) and are used to compute the rank value for each node.³

RPL pays special attention to the changes in the rank value of a node. Essentially, a node can only move “closer” to the destination (i.e., decrease its rank value) without first “poisoning” its own routes, since such an action can never lead to a routing loop. This rule allows RPL to provide reliable collection of data from a large set of nodes to a single destination while optimizing for metrics defined by the OF: the WG has defined standard metrics for link quality, latency, and throughput.

RPL also specifies local and global repair mechanisms for recomputing routes when an inconsistency is detected or based on administrative decisions. Local repair means detaching a node's sub-DODAG by increasing its rank value. Once a root initiates a global repair event all the nodes in the DODAG recompute their rank values and reconfigure their parent sets. The left side of Fig. 2 provides an example topology of a

RPL network. Solid arrows represent each node's preferred parent (determined from the node's neighbors and their rank values) while dotted arrows point to the other nodes in the parent set.

In order to support routing to various destinations within the DODAG which are the root, RPL uses a second message type: the Destination Advertisement Object (DAO). RPL supports scenarios where in-network nodes do not have enough memory to store routes to all possible destinations. In this case, the DAO messages, which contain information on the desired parent set of a destination node, are propagated up the DODAG until they reach the root. The root gathers DAOs from all nodes in the DODAG, and uses them to construct “down” routes to various destinations. Data to these advertised destinations is forwarded along a DODAG until it reaches the root, which then attaches a source routing header and sends it back down the DAG. Alternatively, nodes in the DODAG may store next-hops to downstream destinations. However, a key design simplification was not supporting “mixed-mode operation” in which storing and non-storing nodes coexist, since this was still considered a research issue; thus, all nodes in a DODAG must either store or not store routes. The center and right DODAGs presented Fig. 2 illustrate examples of non-storing and storing modes, respectively.

Security is an important design consideration for LLNs since they are often part of the critical infrastructure. Given that link-layer security mechanisms are not always sufficient to guarantee the integrity of routing messages, RPL defines an optional⁴ cryptographic mode of operation in which the advanced encryption standard (AES) is used for message authentication; RPL also supports RSA signatures for checking the integrity of routing messages. More concise elliptic curve cryptography (ECC)-based signatures were also suggested but these remain a topic of future research: in the LLN setting it

RPL pays special attention to the changes in the rank value of a node. Essentially, a node can only move “closer” to the destination (i.e., decrease its rank value) without first “poisoning” its own routes, since such an action can never lead to a routing loop.

³ OFs specify the metrics and constraints used to compute the routing path and parent node set. The rank of an RPL node is a scalar representation of the node's location within a DODAG. Nodes use these values to determine the set of parent nodes closer to the DODAG root. The rank value is derived from the metrics specified in the OFs.

⁴ Since minimizing the implementation complexity is critical for LLNs, the security mechanisms are specified as optional in RPL.

The 6LoWPAN and RoLL working groups leveraged results from a decade of research by the wireless sensor network research community. In turn, the research community is making efforts to implement and distribute the proposed standards.

Implementation	ROM (B)	RAM (B)
6LoWPAN header compression	7222	1581
TinyRPL — non-storing mode	9166	308
TinyRPL — storing mode	8428	1454
TinyRPL — no downward routes	6990	264

Table 1. Memory usage of 6LoWPAN header compression and TinyRPL in tinynos2.x for telosB nodes. Both implementations require a small amount of memory, making them viable for low-cost resource-constrained wireless devices. The memory usage of the non-root nodes is shown for TinyRPL's non-storing mode. Each node stores up to 30 unique destinations in storing mode. The header compression implementation includes a 1500-byte buffer for reassembling fragmented datagrams.

is desirable that a signature be as short as possible while still providing strong cryptographic guarantees.

IMPLEMENTATIONS OF 6LoWPAN AND RPL

As discussed above, the 6LoWPAN and RoLL WGs leveraged results from a decade of research by the wireless sensor network research community. In turn, the research community is making efforts to implement and distribute the proposed standards to encourage the development of standards-based wireless sensing systems and provide feedback based on real implementations to the standards community. As a first step in distributing these standards, several widely used open-source software platforms including, TinyOS, OpenWSN, and Contiki [15], have released reference implementations of the 6LoWPAN and RoLL standards.

Along with these open implementations, multiple commercial implementations have helped drive the validation of the new IETF standards for constrained devices. The IP for Smart Objects (IPSO) Alliance, a marketing organization whose goal is to drive the adoption of IP in commercial markets, has held multiple interoperability events among its members to validate the specifications described in the drafts from the two WGs. The ZigBee Alliance, an organization focused on developing networking standards for home area networks has also been holding interoperability events among its members as part of its certification process.

BLIP 2.0 AND TINYRPL: IMPLEMENTATIONS IN TINYOS

TinyOS, a widely-used, open source software platform for wireless sensor network research, is one of the first platforms to provide an open implementation of the 6LoWPAN and RoLL

standards. Specifically, blip and TinyRPL (available at <http://code.google.com/p/tinyos-main/>) are open source implementations of the respective standards in TinyOS under the new BSD license.

First released in September 2008, blip is the Berkeley Low-Power IP stack. The project's goal is to provide TinyOS users with a protocol suite that users can build real-world applications. blip takes a practical approach to standards: where they are in good shape, follow them, but where the standards are not fully defined, develop custom solutions. It also provides a framework for researchers to develop new IP-based protocols while maximizing code-sharing.

The IETF 6LoWPAN and RoLL groups have made significant progress towards workable standards, and blip-2.0 is the effort within the project to provide implementations of the latest header-compression standards, as well as a more flexible API and framework. It currently supports, the latest header compression draft (6lowpan-hc).

Along with blip-2.0, TinyRPL was released in September 2010, as a catalyst for the use of IETF routing standards in the TinyOS research community. While TinyRPL does not implement the entire RPL standard,⁵ it provides a usable building block for sensor network researchers.

Both blip and TinyRPL implementations require only a small amount of resources (e.g., ROM, RAM). Table 1 summarizes the memory usage of the two implementations. Note that the memory usage of RPL varies depending on the type of downward routing used. Nevertheless, considering that the widely used MSP430F1611 microcontroller has 48 kbytes of ROM and 10 kbytes of RAM [16], and that other microcontrollers offer similar amounts of memory, the small footprint of the two implementations make them suitable for resource constrained platforms. As an example, blip has already been deployed in several real-world deployments to show the benefits of IP-based addressing in wireless sensor network systems and deployments that use TinyRPL are soon to follow.

SUMMARY AND OUTLOOK

Motivated by the need to support the upcoming automated metering infrastructure and home area networking applications, and energized by a decade of research in wireless sensor networks, the IETF has started standardizing protocols that underlie the emerging Internet of Things.

In this article we review protocols from the 6LoWPAN and RoLL WGs that define mechanisms to adapt IPv6 for 802.15.4 links, including encapsulation of IPv6 packets and auto-configuration functions, and routing protocols for resource constrained devices. These efforts, however, are only the beginning.

Work has already started in the Constrained RESTful Environments (CoRE) Working Group to define an application-layer protocol for constrained devices and we expect additional work in routing, transport, management, and security. In general, it will be important to recognize that the presence of structure and constraints is as

⁵ The current version of TinyRPL does not support the security mechanisms and supports a limited number of objective functions and routing metrics (e.g., MRHOF with ETX).

much an opportunity for profound innovation as their absence. Many problems that have been studied in a manner devoid of context appear anew in a particular context, with a new set of trade-offs and criteria. It becomes possible to innovate in one area while utilizing the best available technology everywhere else, rather than what happens to be available in a local code base, and then to make clear assessments of the impact of a particular innovation. Considering the impact of empowering billions of new inter-networked devices, the emerging IETF protocol stack for low-power and lossy networks should be fertile ground for continued research in wireless sensor networks.

REFERENCES

- [1] P. Levis *et al.*, "The Emergence of A Networking Primitive in Wireless Sensor Networks," *Commun. ACM*, vol. 51, no. 7, 2008, pp. 99–106.
- [2] IEEE 802.15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (LR-WPANs), available at <http://standards.ieee.org/getieee802/download/802.15.4-2003.pdf>.
- [3] G. Montenegro *et al.*, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks," RFC 4944, Sept. 2007.
- [4] J. Hui and P. Thubert, "Compression Format for IPv6 Datagrams in 6LoWPAN Networks," Internet draft, 2010.
- [5] T. Narten *et al.*, "Neighbor Discovery for IPv6 (IPv6)," RFC 4861, Sept. 2007.
- [6] Z. Shelby, S. Chakrabarti, and E. Nordmark, "Neighbor Discovery Optimization for Low-Power and Lossy Networks," Internet draft, 2010.
- [7] M. Dohler, T. Watteyne, and T. Winter, "Routing Requirements for Urban Low-Power and Lossy Networks," RFC 5548, May 2009.
- [8] K. Pister, P. Thubert, and S. Dwars, "Industrial Routing Requirements in Low-Power and Lossy Networks," RFC 5673, Oct. 2009.
- [9] A. Brandt, J. Buron, and G. Porcu, "Home Automation Routing Requirements in Low-Power and Lossy Networks," RFC 5826, Apr. 2010.
- [10] J. Martocci *et al.*, "Building Automation Routing Requirements in Low-Power and Lossy Networks," RFC 5867, June 2010.
- [11] T. Winter, P. Thubert, and RPL Author Team, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks," Internet draft, 2010, work in progress.
- [12] D. S. J. De Couto *et al.*, "A High-Throughput Path Metric for Multi-Hop Wireless Routing," *Proc. 9th ACM Int'l. Conf. Mobile Computing and Networking*, Sept. 2003.
- [13] P. Buonadonna and Gilman Tolle, "MultihopLQI," available at <http://www.tinyos.net/tinyos1.x/tos/lib/MultiHopLQI>, 2004.
- [14] O. Gnawali *et al.*, "Collection Tree Protocol," *Proc. 7th ACM Conf. Embedded Networked Sensor Sys.*, Nov. 2009, pp. 1–14.
- [15] A. Dunkels, B. Grönvall, and T. Voigt, "Contiki, A Lightweight and Flexible Operating System for Tiny Networked Sensors," *Proc. 1st IEEE Wksp. Embedded Networked Sensors*, Tampa, FL, Nov. 2004.
- [16] Texas Instruments, MSP430x1xx Family User's Guide (Rev.F), <http://www.ti.com/lit/pdf/slau049f>, 2006.

BIOGRAPHIES

JEONGGIL KO (jgko@cs.jhu.edu) received his B.Eng. degree in computer science and engineering from Korea University and M.S.E. degree in computer science from Johns Hopkins University, where he is currently working toward a Ph.D. degree. His research interests include wireless medical sensing systems, low-power embedded network system and protocol design, and the deployment of such embedded systems to real environments. He is a recipient of the Abel Wolman Fellowship awarded by the Whiting School of Engineering at Johns Hopkins University.

STEPHEN DAWSON-HAGGERTY is a computer science graduate student at the University of California (UC) Berkeley. He works on enabling ubiquitous access to data collected from the world, moving from the embedded sensor tier into the cloud. He has an undergraduate degree from Harvard University and a Master's from UC Berkeley.

JONATHAN W. HUI received his B.S. degree in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, Pennsylvania, in 2003, and M.S. and Ph.D. degrees in computer science from UC Berkeley in 2005 and 2008, respectively. For his Ph.D. dissertation, he did seminal work on IPv6 in constrained networks, having designed, implemented, and evaluated the world's first IPv6/6LoWPAN network for low-power wireless radios. At the same time, he helped start Arch Rock Corporation, where he was lead engineer and led the development of the PhyNet low-power wireless platform, the world's first commercial IPv6/6LoWPAN product and the first of its kind to receive the IPv6 Ready-Phase 2 (Gold) designation. Cisco acquired Arch Rock in September 2010. Currently, he is a lead engineer within the Smart Grid Business Unit at Cisco, where he drives the design and development of mesh networking technology for smart grid solutions. He participates in the IETF and has authored standards specifications related to IPv6 in low-power wireless networks. He also participates in the IP for Smart Objects Alliance to promote the use of IP for smart objects.

DAVID E. CULLER received his B.A. from UC Berkeley in 1980, and his M.S. and Ph.D. from MIT in 1985 and 1989, respectively. He joined the EECS faculty in 1989 and is the founding director of Intel Research, UC Berkeley, and associate chair of the EECS Department. He is a member of the National Academy of Engineering, an ACM Fellow, and was selected in *Scientific American's* Top 50 Researchers and *Technology Review's* 10 Technologies that Will Change the World. He was awarded the NSF Presidential Young Investigator and the Presidential Faculty Fellowship. His research addresses networks of small embedded wireless devices, planetary-scale Internet services, parallel computer architecture, parallel programming languages, and high-performance communication.

PHILIP LEVIS is an assistant professor of computer science and electrical engineering at Stanford University. He received an Sc.B. in computer science and biology from Brown University in 1999, an M.S. in computer science from Colorado University Boulder in 2001, and a Ph.D. in computer science from UC Berkeley in 2005. In 2008 he received an NSF CAREER award and was named a Microsoft New Faculty Fellow. He is a co-author of the Internet draft describing RPL.

ANDREAS TERZIS received his Ph.D. degree in computer science from UC Los Angeles. He is an associate professor in the Department of Computer Science at Johns Hopkins University, where he leads the Hopkins interNetworking Research Group (HiNRG). He is a recipient of the National Science Foundation (NSF) CAREER award. His research

Motivated by the need to support the upcoming automated metering infrastructure and home area networking applications, and energized by a decade of research in wireless sensor networks, the IETF has started standardizing protocols that underlie the emerging Internet of Things.