



Extending IP to Low-Power, Wireless Personal Area Networks

Extending IP to low-power, wireless personal area networks (LoWPANs) was once considered impractical because these networks are highly constrained and must operate unattended for multiyear lifetimes on modest batteries. Many vendors embraced proprietary protocols, assuming that IP was too resource-intensive to be scaled down to operate on the microcontrollers and low-power wireless links used in LoWPAN settings. However, 6LoWPAN radically alters the calculation by introducing an adaptation layer that enables efficient IPv6 communication over IEEE 802.15.4 LoWPAN links.

Several leading radio manufacturers have implemented IEEE 802.15.4, which specifies a wireless link for low-power personal area networks (LoWPANs). 802.15.4 is widely used in embedded applications, such as environmental monitoring to improve agricultural yields, structural monitoring to track building and bridge integrity, and industrial control to provide more sense points and control points at lower cost. These applications generally require numerous low-cost nodes communicating over multiple hops to cover a large geographical area, and they must operate unattended for years on modest batteries. Such requirements target a very different set of applications than do WPAN technologies

such as Bluetooth, which eliminate wiring for headsets, game controllers, and personal devices. Accordingly, 802.15.4's capabilities are more limited than other WPANs and WLANs – they have small frame sizes, low bandwidth, and low transmit power. Additionally, the microcontrollers typically coupled with LoWPAN radios have limited memory and compute power. These constraints led many LoWPAN vendors to embrace proprietary protocols and link-only solutions (such as ZigBee), presuming that IP was too memory- and bandwidth-intensive for them to scale it down as necessary.

6LoWPAN introduces an adaptation layer between the IP stack's link and network layers to enable efficient

Jonathan W. Hui
Arch Rock

David E. Culler
University of California, Berkeley

transmission of IPv6 datagrams over 802.15.4 links, dramatically reducing IP overhead.¹ The adaptation layer is an IETF proposed standard and provides header compression to reduce transmission overhead, fragmentation to support the IPv6 minimum *maximum transmission unit* (MTU) requirement, and support for layer-two forwarding to deliver an IPv6 datagram over multiple radio hops. 6LoWPAN achieves low overhead by coupling traditional protocol layers; it uses information in the link and adaptation layers to compress network- and transport-layer headers. Drawing on IPv6 extension headers, it employs the *header stacking* principle to separate orthogonal concepts and keep the header small and easy to parse.

Here, we discuss key 6LoWPAN concepts to demonstrate how it enables efficient support for IPv6 over 802.15.4 links.

IPv6 over IEEE 802.15.4

The IPv6 protocol is designed to supersede IPv4 and enable the Internet to scale for decades to come. To overcome dwindling unallocated address space – and in anticipation that networked appliances and instruments will vastly outnumber conventional computer hosts – IPv6 expands the IP address space from 32 to 128 bits. Recognizing the growth in link bandwidth, IPv6 increases the minimum MTU requirement from 576 to 1,280 bytes. To simplify routers and increase performance, IPv6 implements fragmentation at the endpoints, rather than in intermediate routers. To increase protocol efficiency and eliminate the need for ad hoc link-level services to bootstrap a subnet, IPv6 includes scoped multicast as an integral part of its architecture. Core IPv6 components, such as Neighbor Discovery (ND),² use link-local scoped multicast for address resolution, duplicate address detection (DAD), and router discovery. Stateless address autoconfiguration (SAA)³ simplifies configuration and management of IPv6 devices by enabling nodes to assign themselves meaningful addresses.

IPv6 also reflects the advances in link technologies the Internet uses. Ethernet has prevailed as the dominant link, and its throughput has increased at an extraordinary rate. Current WLAN technologies, such as Wi-Fi, mirror Ethernet capabilities by supporting similarly sized MTUs and high link rates. Both links operate in the context of ample power and highly capable devices.

WPAN technologies, on the other hand, operate with lower power. IEEE 802.15.4 was designed specifically for long-lived application domains that require numerous low-cost nodes, and these constraints limit the capability of LoWPAN links and the microcontrollers to which they're attached. Throughput is limited to 250 kbps in the 2.4-GHz band and 20 or 40 kbps in other bands. The frame length is limited to 128 bytes to ensure reasonably low packet error when bit-error rates are non-negligible and reflects microcontrollers' limited buffering capabilities. 802.15.4 defines short 16-bit link addresses, in addition to IEEE EUI-64 addresses, to reduce header overhead and memory requirements. Communication range is short (tens of meters) because transmission power increases polynomially with range. Unlike most typical WPAN and WLAN installations, LoWPANs communicate over multiple hops. Finally, the associated microcontrollers typically have about 8 Kbytes of data RAM and 64 Kbytes program ROM.

Due to these resource constraints and LoWPANs' multihop nature, supporting IPv6 over LoWPAN networks presents several challenges. First, IPv6 datagrams aren't a natural fit for LoWPANs. Low throughput, limited buffering, and frames that are one-tenth the size of the IPv6 minimum MTU requirement make datagram fragmentation and compression a necessity for efficient operation. For example, link headers can limit effective link payload to 81 bytes, making the IPv6 (40 bytes), User Datagram Protocol (UDP; 8 bytes), and TCP (20 bytes) headers seem exceedingly large. Second, because 802.15.4 is both low-power and low-throughput, it's more prone to spurious interference, link failures, dynamic link qualities, and asymmetric links. Such characteristics require the network layer to be responsive and adaptive while remaining energy efficient, and they affect all aspects of networking, including fragmentation, compression, forwarding, and routing. Third, a LoWPAN's expected topology is a mesh of short-range connections. This negates the assumption that the link is a single broadcast domain on which a core of IP architectural components – such as IPv6 ND and SAA – relies. The IETF 6LoWPAN working group addressed these issues with RFC 4944.¹ In the remainder of this article, we provide a basic overview of RFC 4944 and touch on the issues that remain to be addressed.

6LoWPAN Adaptation Layer

The 6LoWPAN format¹ defines how IPv6 communication is carried in 802.15.4 frames and specifies the adaptation layer's key elements. 6LoWPAN has three primary elements:

- *Header compression.* IPv6 header fields are eliminated from a packet when the adaptation layer can derive them from the link-level information carried in the 802.15.4 frame or based on simple assumptions of shared context.
- *Fragmentation.* IPv6 packets are fragmented into multiple link-level frames to accommodate the IPv6 minimum MTU requirement.
- *Layer-two forwarding.* To support layer-two forwarding of IPv6 datagrams, the adaptation layer can carry link-level addresses for the ends of an IP hop. Alternatively, the IP stack might accomplish intra-PAN routing via layer-three forwarding, in which each 802.15.4 radio hop is an IP hop.

The key concept applied throughout the 6LoWPAN adaptation layer is that it uses stateless compression to elide adaptation-, network-, and transport-layer header fields – compressing all three layers down to a few bytes, combined.⁴ We can see that it's possible to compress header fields to a few bits when we observe that they often carry common values, reserving an escape value for when less-common ones appear. Common values occur due to frequent use of a subset of IPv6 functionality (such as UDP, TCP, and Internet Control Message Protocol version 6 [ICMPv6] as Next Header values) and simple assumptions of shared context (for example, a common global routing prefix for the entire LoWPAN). 6LoWPAN also elides redundant header information across protocol layers (for instance, IPv6 length fields and IPv6 addresses are derived from lower-layer headers).

Traditional IP header compression techniques are stateful and generally focus on optimizing individual flows over a highly constrained link.⁵ These methods assume that the compressor and decompressor are in direct and exclusive communication and compress both network- and transport-layer headers together. They optimize for long-lived flows by exploiting redundancies across packets within a flow over time, requiring the endpoints to initially send packets un-

compressed. Flow-based compression techniques are poorly suited for LoWPANs. Traffic in many LoWPAN applications is driven by infrequent readings or notifications, rather than long-lived flows. Communication over multiple hops requires hop-by-hop compression and decompression and per-flow state at each intermediate node. Many LoWPAN routing protocols obtain receiver diversity via rerouting, which would require state migration and reduce compression effectiveness. In contrast, stateless compression in 6LoWPAN doesn't require any per-flow state and lets routing protocols dynamically choose routes without affecting compression efficiency. Looking at 6LoWPAN's specifics, we can see how extensively it employs stateless compression.

Because 802.15.4 is both low-power and low-throughput, it's more prone to spurious interference, link failures, dynamic link qualities, and asymmetric links.

Encapsulation Header Format

6LoWPAN uses header stacking to keep orthogonal concepts separate and enforce a well-defined method for expressing its capabilities. Analogous to IPv6 extension headers, 6LoWPAN expresses each capability in a self-contained subheader: *mesh addressing*, *fragmentation*, and *header compression*. Mesh addressing supports layer-two forwarding, and fragmentation supports the IPv6 minimum MTU requirement. 6LoWPAN identifies all header formats using *dispatch* subheaders, including uncompressed IPv6, 6LoWPAN compressed IPv6, and other adaptation-layer headers. A 6LoWPAN Not-A-LoWPAN (NALP) dispatch enables it to coexist with other protocols. The header stack is simple to parse and supports stateless compression. The fragmentation header is elided for small datagrams, indicating that a single frame carries the entire payload. Similarly, the mesh header is elided when 6LoWPAN frames are delivered over a single radio hop, so the path source and destination are identical to those in the link-layer header. Figure 1 shows typical header stacks and details for each subheader.

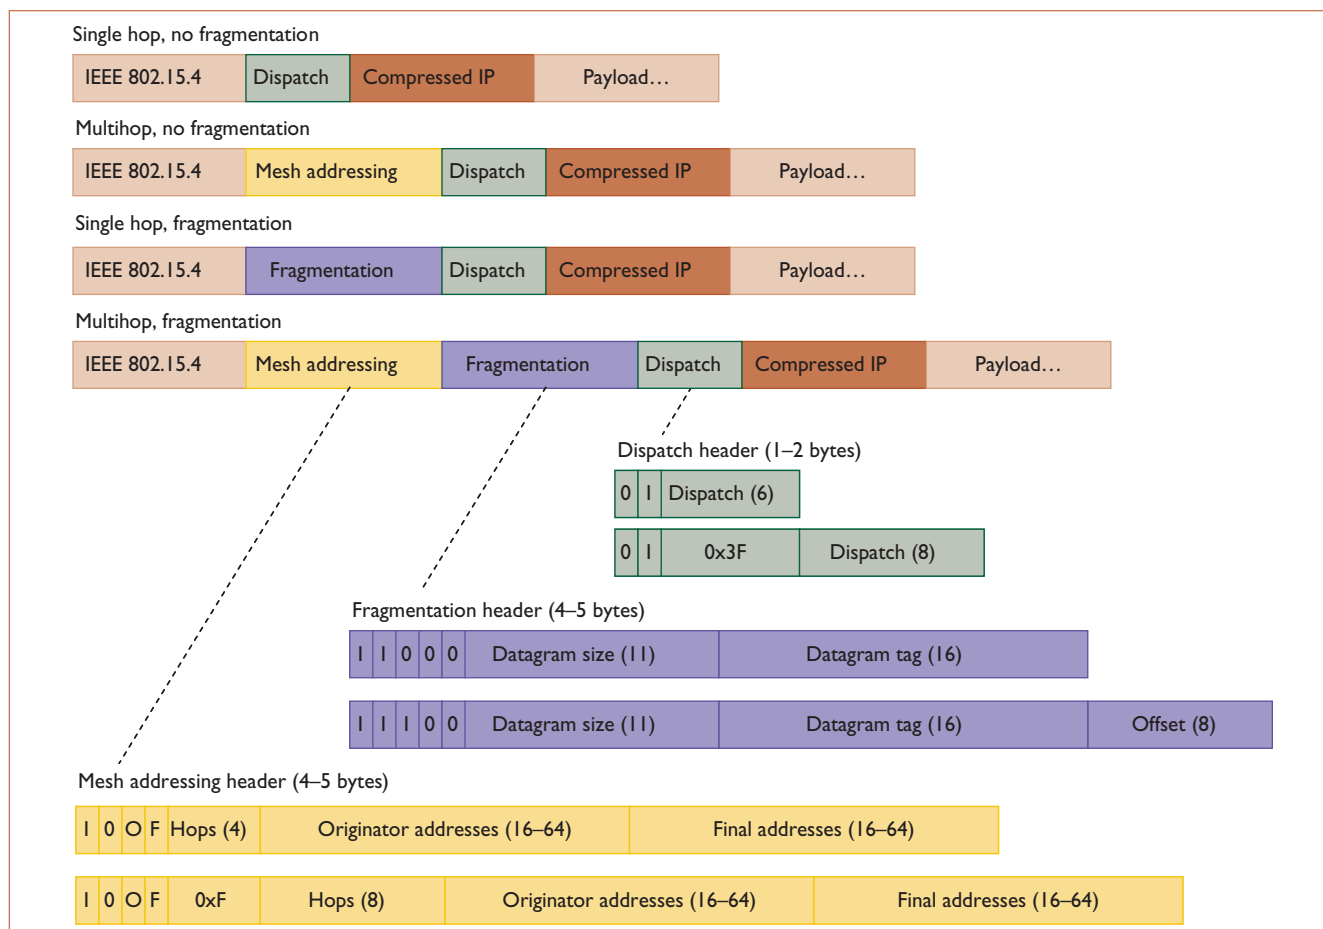


Figure 1. 6LoWPAN headers. 6LoWPAN uses header stacking to keep orthogonal concepts separate and naturally supports compact headers by eliding headers that are unused. It defines a mesh addressing header to support layer-two forwarding, a fragmentation header when the IPv6 datagram is too large to fit in a single 802.15.4 frame, and header compression to reduce IPv6 header overhead.

Network- and Transport-Layer Header Compression

Stateless compression elides fields in network- and transport-layer headers. The 6LoWPAN format defines HC1, a compression scheme optimized for link-local IPv6 communication. HC1 is identified by an encoding byte following the Compressed IPv6 dispatch header, and it operates on fields in the upper-layer headers. 6LoWPAN elides some fields by assuming commonly used values. For example, it compresses the 64-bit network prefix for both source and destination addresses to a single bit each when they carry the well-known link-local prefix. 6LoWPAN compresses the Next Header field to two bits whenever the packet uses UDP, TCP, or ICMPv6. Furthermore, 6LoWPAN compresses Traffic Class and Flow Label to a single bit when their values are both zero. Each compressed form has reserved values that indicate that the

fields are carried inline for use when they don't match the elided case. 6LoWPAN elides other fields by exploiting cross-layer redundancy. It can derive Payload Length – which is always elided – from the 802.15.4 frame or 6LoWPAN fragmentation header. The 64-bit interface identifier (IID) for both source and destination addresses are elided if the destination can derive them from the corresponding link-layer address in the 802.15.4 or mesh addressing header. Finally, 6LoWPAN always elides Version by communicating via IPv6. Hops Left is the only field always carried inline. Fully compressed, the HC1 encoding reduces the IPv6 header to two bytes.

A bit in the HC1 encoding indicates transport-layer compression, which is currently defined only for UDP. 6LoWPAN might compress the upper 12 bits of both Source and Destination Ports to a single bit each when either carry

a predefined value, allowing significant compression when communicating within a 16-port range. 6LoWPAN uses a third bit to elide Payload Length. Checksum is the only field that UDP messages must carry fully inline. Fully compressed, the UDP header is reduced to four bytes.

HC is an alternative IPv6 header compression scheme for 6LoWPAN proposed in a separate Internet draft.⁶ HC generalizes HC1 to support compression of arbitrary network prefixes. It recognizes that all nodes in a LoWPAN are likely to have a *common routing prefix* (CRP) and exploits this shared context to maintain significant compression when communicating over multiple IP hops. By defining two different dispatch values, HC can support both a link-local prefix and a CRP. HC also supports a 16-bit compressed form for both source and destination addresses. The 16-bit form enables greater compression than HC1 when the IID is derivable from the short link address. HC also uses the 16-bit form to compress well-known multicast addresses and carry the four-bit scope inline and maps the 112-bit group ID down to nine bits. HC can compress the IPv6 header to two bytes.

Adaptation-Layer Evaluation

By exploiting commonly occurring shared context, 6LoWPAN eliminates many aspects of IPv6 overhead. In the best case, nodes send small IPv6 datagrams over a single 802.15.4 hop. Because no mesh addressing or fragmentation headers are required, 6LoWPAN can fully compress the IPv6 header to two bytes, and the overhead added to a raw 802.15.4 frame is only seven bytes (one dispatch, two HC1, and four UDP). Expanding this capability to cover multiple hops requires additional addressing information, either in the mesh addressing header or the compressed IPv6 header. Using 16-bit addresses increases the overhead to 12 and 11 bytes, respectively. When communicating with IP devices outside the LoWPAN, packets might carry a full IPv6 address inline. Note that HC always lets 6LoWPAN compress at least one address when using the link-local prefix or CRP. Comparatively, ZigBee has a seven-byte header for communicating over a single hop and a 15-byte header when communicating over multiple hops, which is equal or larger to 6LoWPAN's compressed UDP/IPv6 header. But unlike 6LoWPAN, ZigBee provides no mechanisms to communicate end to end with arbitrary IP devices.

To demonstrate 6LoWPAN header compression's benefits, we compute the energy cost of communicating a variable-sized data payload carried in the ideal network header (zero bytes), 6LoWPAN compressed headers (for link-local, intra-PAN, and internetwork), and an uncompressed IPv6 header. The computation includes the 192- μ s transmit/receive turnaround time and transmission times for the preamble, start-of-frame delimiter, 802.15.4 link headers with short source and destination addressing modes, a nine-byte 802.15.4 security header, 6LoWPAN headers, IPv6/UDP headers, and a variable data payload. This analysis also includes transmission costs for 802.15.4 link acknowledgments. We took all parameters from the CC2420 data sheet using the standard 250 kbps data rate with 0 dBm transmit power. Figure 2 shows significant cost reduction compared to uncompressed IPv6, especially when the compression removes the need for fragmentation. Decreasing the header size saves energy per packet, increases available payload, and potentially eliminates fragmentation costs. Overhead for 6LoWPAN communication with fully compressed headers is negligible compared to raw 802.15.4 operations. Typical IPv6 communication will present a mix of internetwork and link-local communication, depending on the traffic characteristics. Although we don't consider MAC-layer costs, these are orthogonal to whether 6LoWPAN frames carry IPv6 datagrams.

IPv6/6LoWPAN Architecture

The 6LoWPAN format specification defines how fragmentation, compression, and layer-two forwarding are represented in an 802.15.4 frame. However, the implementation of those capabilities is out of that document's scope. 6LoWPAN's dependencies on the specific operations defined in the 802.15.4 MAC are minimal, supporting essentially any MAC protocol that provides the 802.15.4 frame format. Similarly, the 6LoWPAN format doesn't specify how IPv6 capabilities, such as ND and SAA, are orchestrated to configure the LoWPAN to be consistent with the adaptation layer. Next, we outline IPv6 over 6LoWPAN's key architectural issues.

IEEE 802.15.4 in Practice

802.15.4 presents several pragmatic issues that have significant architectural impact beyond the 6LoWPAN adaptation layer. Whereas in con-

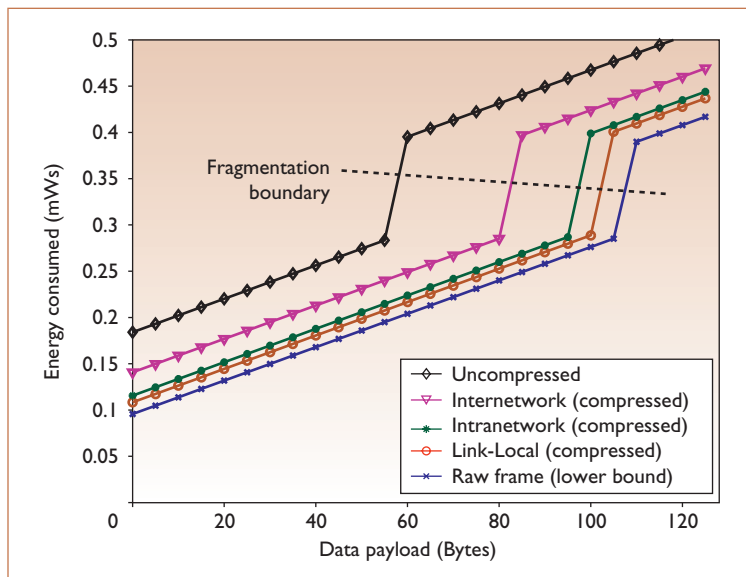


Figure 2. Energy cost for transmitting IPv6/6LoWPAN packets on the CC2420 radio. The 6LoWPAN adaptation layer adds minimal header overhead while providing effective IPv6 and transport-layer header compression, which make transmission of IPv6 datagrams over 802.15.4 cost effective. The header overhead is comparable to other networking technologies that don't provide end-to-end IP interoperability.

ventional WPAN settings, the user typically adjusts device and host placement so that the link between them is adequate, in typical LoWPAN settings a network of many devices is embedded in a physical environment at particular, meaningful locations. Network protocols must deal with the many exigencies that arise. Multihop routing extends range and helps avoid obstacles. Thus, a LoWPAN network isn't typically a single broadcast domain. Moreover, the link quality between any node pair is often complex and time-varying due to environmental factors. Hop-by-hop retransmission schemes help make lossy 802.15.4 links viable for multihop communication, but alone they aren't sufficient. Links that are reasonably good on average – say, with 90 percent packet-reception reliability – will often experience bursts of loss due to changes in the noise floor and spurious interference. Routing can overcome such bursts when forwarding datagrams by selecting an alternate path. In effect, routing can exploit receiver diversity by dynamically selecting from multiple next-hop candidates. To deal with these link challenges, the network layer requires extra visibility into detailed link behavior to build and maintain effective routing structures.

Many LoWPAN applications have significant

device mobility within the LoWPAN, giving rise to time-varying connectivity relationships, in addition to variations induced by changing environmental factors. For example, package tracking might involve numerous devices moving among a set of stationary ones. This isn't IP mobility in the traditional sense because nodes might remain in close physical proximity and be connected within the LoWPAN. However, such variations require that the routing topology adapt to connectivity changes.

The 802.15.4 specification defines only a limited set of power-management mechanisms for edge devices and no power management for forwarding devices. Consequently, most commercial implementations and industrial standards built on 802.15.4 forego the defined power-management mechanisms when defining routing protocols. To conserve energy, nodes must duty cycle the radio, but doing so requires both transmitter and receiver to coordinate when and how to communicate. Common mechanisms for this involve sampled listening techniques, in which the receiver periodically listens for lengthened transmissions, or scheduling techniques, which involve time synchronization between nodes. 6LoWPAN, so far, avoids requiring particular MAC features. When adapting IPv6 components to operate over 802.15.4 links, we should exercise similar care regarding dependencies on the specific underlying MAC protocol.

Mesh Under vs. Route Over

Two important architectural issues for IPv6 over LoWPAN are how link-level factors inform routing and at what layer datagram forwarding occurs within the LoWPAN. Traditionally, IP routing occurs at the network layer in a manner largely independent from the underlying links that implement the individual hops. 6LoWPAN, in its role as an adaptation between the link (layer two) and the network (layer three), can support routing at either layer.

In a *mesh under* organization, the network stack performs no IP routing within the LoWPAN; instead, the adaptation layer seeks to mask the lack of a full broadcast at the physical level by transparently routing and forwarding packets within the LoWPAN. By emulating a full broadcast link, it potentially provides compatibility with IPv6 protocols that expect such communication behavior. The challenge

is that logical link emulation is significantly more complex in LoWPANs than in traditional infrastructure-based 802.11 topologies. Mesh topologies require forwarding over multiple radio hops, and link-local multicast must deliver packets to all nodes in the entire LoWPAN. Many mechanisms that exist to form, maintain, and diagnose IP routing must also be recreated at the link layer for meshing to operate reliably.

Alternatively, *route over* performs routing at the IP layer, with each node serving as an IP router. We can view it as a collection of overlapping link-local scopes, with each link-local domain defined by the inherent radio connectivity. Unlike mesh under, route over supports layer three forwarding mechanisms within the LoWPAN that can utilize network-layer capabilities defined by IP, such as IPv6 routing or hop-by-hop option headers and ICMPv6 for configuration and management. Route over also lets IP routing protocols span different link technologies, enabling better integration into more capable networks. It also lets IP-based protocols constrain IP communication to local radio coverage, rather than an entire LoWPAN.

Issues of link- versus network-layer routing aren't unique to 6LoWPAN – they arise in Frame Relay, Asynchronous Transfer Mode (ATM), switched Ethernet, 802.11 meshing, and, to some extent, VPNs. For example, we've experienced similar challenges with IP over ATM, in which independent link-level routing makes it difficult to optimize IP routes end-to-end. Additionally, two independent routing layers can have unintended interactions, especially when reacting to changes in link state. Instead, IP over ATM found much better success with Multiprotocol Label Switching (MPLS) in a route over configuration in which routing occurs at layer three but forwarding occurs at layer two. MPLS might provide useful guidance as 6LoWPAN matures.

Addressing, Autoconfiguration, and ND

Using SAA, each host generates a link-local IPv6 unicast address from its IEEE EUI-64 address, 16-bit address, or both. In mesh under, the link-local scope covers an entire LoWPAN, possibly over multiple hops, and a link-local address is sufficient for communication within the LoWPAN, whereas a routable address is required to communicate outside. In route over, a link-local address is sufficient to communicate

Related Links

- The Internet Engineering Task Force (IETF), which is concerned with the evolution of the Internet architecture and the smooth operation of the Internet: www.ietf.org
- Homepage for the 6LoWPAN working group within the IETF, concerned with supporting IPv6 over IEEE 802.15.4 links: www.ietf.org/html.charters/6lowpan-charter.html
- Homepage for the ROLL working group within the IETF, concerned with addressing routing over low-power and lossy links: www.ietf.org/html.charters/roll-charter.html
- Homepage for the IEEE 802.15.4 task group, concerned with the development of wireless personal area networks: www.ieee802.org/15/pub/TG4.html

with nodes in direct radio communication, but a routable address is required to communicate with devices that are multiple radio hops away.

For all unicast addresses, regardless of their scope, it's cost effective to derive them from the 802.15.4 link address. 6LoWPAN's binding between link, adaptation, and IP headers enables 6LoWPAN to elide IP addresses derived from link addresses and removes the need for address resolution. Similarly, autoconfiguration should configure interface addressing using a CRP so that 6LoWPAN can elide the prefix. 6LoWPAN can use the short link address to derive IP addresses – allowing reduced header overhead when using the mesh addressing header or HC – and maintain privacy by using a token with local scope. Because link addresses must be unique within the LoWPAN, a mapping between IP and link addresses removes the need for DAD, which requires expensive floods or additional mechanisms to maintain uniqueness. The network might assign short link addresses at the link layer or use DHCPv6 when carried in the IPv6 address.

ND also lets a node discover neighbors, maintain reachability information, configure default routes, and propagate configuration parameters. ND is currently defined only for operation on a single link. Although it can potentially run unmodified with mesh under, it presents several issues. ND extensively uses link-local multicasts, which must be implemented in mesh under using expensive floods. It also uses *neighbor unreachability detection* (NUD) to determine reachability to other nodes within link-local scope, but placing multiple link hops in between can cause unexpected timing interactions between layers if the link is incapable of

rerouting quickly enough. These issues become more significant as the LoWPAN's size increases. Instead, the existing ND protocol might actually be better suited to route over. Because the link-local scope is defined by the radio communication range, a link-local multicast reduces to a link-layer broadcast, and IP neighbors are defined by neighbors within physical connectivity. Expensive floods aren't needed to support ND advertisements, and NUD traverses only a single link hop. However, we must make modifications to ND to support its ability to auto-configure a collection of nodes connected over multiple IP hops.

Routing

Limited memory and communication capabilities constrain the routing state at each node as well as the routing information that might be communicated. These restrictions preclude using protocols that rely on complete link-state information. Traditional distance vector mobile ad hoc networks (manet) protocols are also ill-suited because they assume a high rate of mobility for all nodes in the network, whereas LoWPAN nodes are generally more stationary. Consequently, manet protocols use frequent floods to discover and maintain routes. Caches used to optimize communication only trade memory for communication. In addition, most of these protocols exchange route maintenance information at rates that far outpace typical LoWPAN communication and react to link fading with expensive route-repair actions. Instead, LoWPAN routing protocols must operate using incomplete information and tolerate some inconsistency. Interestingly, we're returning to scalability issues similar to those encountered with the early Internet, but this time in a wireless setting. The new Routing over Low Power and Lossy Links (ROLL) working group within the IETF routing directorate will soon address these challenges.

Although a full treatment of routing protocol design is beyond this article's scope, we've outlined several underlying design issues. Indeed, we can even apply the 6LoWPAN philosophy of elision based on shared context to routing-table structure. IP routing protocols typically require 32 bytes to store the IPv6 addresses for the destination and next hop, which can easily fill hundreds of bytes of memory. Although routing table length is traditionally reduced through route summarization, doing so binds

host addresses to the underlying topology. This is problematic when link qualities can change due to environmental effects or node mobility. 6LoWPAN supports a different alternative by reducing the route table width. Because 6LoWPAN establishes a mapping between link and IP addresses, intra-PAN routing essentially operates on link addresses and can reduce 32 bytes of addressing to four bytes.

6 LoWPAN only opens the possibility of supporting IPv6 over 802.15.4 links. As we've seen, the ad hoc network topology and strict resource constraints have implications for core pieces of the IPv6 architecture. Although we've seen IP solutions for ad hoc networks and others for extremely resource-constrained devices, little has been done to complete the IPv6 story for low-power, wireless networks such as IEEE 802.15.4. We've completed a high-quality IPv6 network stack based on 6LoWPAN, including ICMPv6, stateless and DHCPv6 autoconfiguration, forwarding, routing, and UDP and TCP transport. It's only 24 Kbytes of ROM and 3 Kbytes of RAM in size. We plan to use this implementation to aid in developing and standardizing mechanisms required to complete the 6LoWPAN story.

We'll also assess broader architecture concepts in the context of LoWPAN extended IP networks, including proxies, the domain name system, and application-layer protocols. LoWPANs are poised to form the next tier of the Internet, finally bringing physical information and control into our broad computing infrastructure. LoWPANs have the potential to connect the next billion hosts to the Internet. □

Acknowledgments

This work is supported in part by the US National Science Foundation under CRI:0435454 and NeTS-NR:0435454.

References

1. G. Montenegro et al., *Transmission of IPv6 Packets over IEEE 802.15.4 Networks*, IETF RFC 4944, Sept. 2007; <http://tools.ietf.org/html/rfc4944>.
2. T. Narten et al., *Neighbor Discovery for IP version 6 (IPv6)*, IETF RFC 4861, Sept. 2007, <http://tools.ietf.org/html/rfc4861>.
3. S. Thomson, T. Narten, and T. Jinmei, *IPv6 Stateless Address Autoconfiguration*, IETF RFC 4862, Sept. 2007; <http://tools.ietf.org/html/rfc4862>.

Related Work in Low-Power Wireless Personal Area Networks

RFC 4944 gives a complete description of 6LoWPAN,¹ the packet format standardized by the IETF to enable IPv6 communication over low-power, wireless personal area networks (LoWPANs). Support for IP in resource-constrained environments has a long history, including over telephone modems that gave rise to Point-to-Point Protocol, Dynamic Host Configuration Protocol for autoconfiguration, and header compression. 6LoWPAN differs in how it exploits shared context, frequently occurring simple cases, and cross-layer redundancy to vastly reduce header overhead when communicating over a dynamic, multihop topology. 6LoWPAN builds on prior work with stateless IP header compression.² Many efforts have addressed links in which multihop forwarding is required, including frame relay and Asynchronous Transfer Mode. 6LoWPAN is unique in that it also addresses severe resource constraints.

IEEE 802.15.1 (Bluetooth) is another wireless link technology that falls under the WPAN classification. Intended to serve as a cable-replacement technology, Bluetooth supports relatively high throughput for a limited number of nodes within a small range. IEEE 802.15.3 pushes WPAN capabilities further, with greater throughput and support for more nodes. Although both are intended for battery operation, they only target lifetimes of several days to several weeks. In contrast, 802.15.4 is intended for low data-rate applications in which numerous nodes must be low-cost and have multiyear lifetimes on modest batteries. The 802.15.4 standard supports up to 64,000 nodes within a PAN compared to a small handful with other WPAN links. 802.15.4 has also reduced complexity, intended to function with eight-bit microcontrollers providing 8 Kbytes of RAM or less. Although IP over Bluetooth using the Bluetooth Network Encapsulation Protocol has been around for several years, it's typically used to provide a point-to-point connection over a single radio hop.

Researchers have developed numerous mesh network layers over 802.15.4, as open source projects (such as TinyOS and

micro-IP (uIP), industrial forums (ZigBee and WirelessHART), or proprietary offerings (Dust Networks, Sencicast, and Millennial Net). Each has defined its own set of incompatible packet formats tied to particular MAC features, routing algorithms, and addressing. Many address only the individual 802.15.4 subnet, leaving all further communication protocols to be defined via ad hoc gateways. 6LoWPAN potentially lets us unify this disparate activity and enable embedded 802.15.4 devices to be incorporated into Ethernet, Wi-Fi, General Packet Radio Service, and other environments within a uniform IP framework. uIP and other embedded TCP/IP stacks provide IP host functionality and are widely used in wired and powered settings. However, almost no embedded IP stacks directly address the issues related to supporting IP over low-power mesh topologies in LoWPANs.

Within the IETF, the mobile ad hocs networks (manet) working group and related research activities' tremendous effort has been devoted to reactive and proactive routing protocols for mobile devices. This work has assumed capable, high-bandwidth links and powerful hosts with high, random mobility. As such, it used conventional IP datagrams and frame formats and hasn't attend to the impact of resource constraints. Work in the IETF Ad-Hoc Network Autoconfiguration (AUTOCONF) working group is devoted to developing solutions for stateless address autoconfiguration and Neighbor Discovery in settings in which IP connectivity is naturally viewed as a collection of overlapping partial broadcast domains. The Routing Over Low-Power and Lossy Links (ROLL) working group was recently chartered to address routing in LoWPANs.

References

1. G. Montenegro et al., *Transmission of IPv6 Packets over IEEE 802.15.4 Networks*, IETF RFC 4944, Sept. 2007; <http://tools.ietf.org/html/rfc4944>.
2. C. Westphal and R. Koodli, "Stateless IP Header Compression," *Proc. IEEE Int'l Conf. Comm. (ICC 05)*, vol. 5, 2005, pp. 3236–3241.

4. C. Westphal and R. Koodli, "Stateless IP Header Compression," *Proc. IEEE Int'l Conf. Comm. (ICC 05)*, vol. 5, 2005, pp. 3236–3241.
5. V. Jacobson, *Compressing TCP/IP Headers for Low-speed Serial Links*, IETF RFC 1144, Feb. 1990; <http://tools.ietf.org/html/rfc1144>.
6. J. Hui and D. Culler, "Compression Format for IPv6 Datagrams in 6LoWPAN Networks," draft-hui-6lowpan-hc-00, June 2007; <http://tools.ietf.org/html/draft-hui-6lowpan-hc-00>.

Jonathan W. Hui is a lead engineer at Arch Rock. His research interests are in network architectures, protocols and algorithms, and system design and implementation for low-power wireless networks. Hui has a BS in elec-

trical and computer engineering from Carnegie Mellon University, and an MS and a PhD in computer science from the University of California, Berkeley. Contact him at jhui@archrock.com.

David E. Culler is a professor of computer science at the University of California, Berkeley, and CTO of Arch Rock. His research interests include wireless embedded networks, efficient operating systems, and parallel computer architecture. Culler has a BA in math from UC Berkeley, and an MS and a PhD in computer science from the Massachusetts Institute of Technology. He is a member of the National Academy of Engineering and a fellow of the ACM and the IEEE. Contact him at culler@cs.berkeley.edu.