INTERNET Ø: INTERDEVICE INTERNETWORKING

End-to-End Modulation for Embedded Networks

Neil Gershenfeld and Danny Cohen

The Internet may be the most complex system ever engineered; from the first host in 1969, it's grown to comprise more than 1 billion routable host addresses [1]. Its future expansion may be more dramatic still due to the demand to extend the Internet from people to things [2], but the frontiers of high-speed networking have receded further and further from the requirements of small, cheap, slow devices. These things need the Internet's original insights, rather than their current implementation; this is being done in the IØ initiative.

The demand for networking embedded devices has led to a proliferation of standards and protocols, including X10, HomePlug, LonWorks, BACnet. CEBus, Fieldbus, ModBus, CAN, Lin, I^2C, SPI, SSI, ASI, USB, EPC, IrDA, Bluetooth, 802.15.4, and ZigBee. While each of these has been optimized for a particular domain, all are encountering many of the same issues that the Internet faced as it grew, including inadequate address space, the need for naming and routing across networks, and mutual incompatibility. This situation is in fact analogous to the early days of the Internet itself.

Early packet-switched networks, including ARPANET, PRNET, and SATNET either relied on complex protocol converters at their

 8755-3996/06/$20.00 ©2006 IEEE

interfaces or couldn't connect at all. The simple, profound solution to this problem was to agree on a representation of a packet that is independent of the network that carries it [3], [4]. This was enshrined in the end-to-end principle for internetworking: the function of a network should be defined by what is connected to it rather than imposed by the construction of the network [5]. E-mail, instant messaging, and the Web could all be invented without requiring agreement on changes to the Internet's infrastructure. While strict adherence to this principle has been challenged by the demands placed on the Internet [6], it has served as a valuable design guide as the Internet has grown far beyond its initially-anticipated applications.

The demand for networking embedded devices has led to a proliferation of incompatible standards and protocols.

IØ can be understood as extending the end-to-end principle from computers to devices. It grew out of a series of testbeds (see Figure 1). In Figure 1(d), for example, the lights and switches could be plugged anywhere into a modular track system. Each device communicated with IP packets so that its function could be determined by the logical configuration of the network rather than the physical wiring, and each contained a Web server so that its state could be seen and changed remotely over the Internet as well as locally. These devices could be configured by a user demonstrating their operation, for example by pushing a button on a light and then operating a switch to be associated with the light, rather than requiring the use of an external computer.

1. Internet Ø testbeds. (a) A networked bathroom shelf for managing a senior's medication, from the White House/Smithsonian Museum Millennium technology demonstrations. (b) Furniture for information navigation, at New York's Museum of Modern Art UnPrivate House show. (c) An interactive stage for the Flying Karamazov Brothers. (d) A programmable building, the Media House in Barcelona. (Photos courtesy of Neil Gershenfeld.)

Networking lights and switches isn't an arbitrary demonstration; it has serious implications for the economics of building construction. In the United States, the construction industry is a trillion-dollar-per-year business; just the payroll in 2002 was US$235 billion [7]. The cost of drawing, following, checking, and later revising wiring diagrams could be replaced by simply servicing the building's infrastructure with energy and information if its configuration could be determined by the occupants, but the benefit in cost and convenience would be lost if the installation must be done by a skilled network engineer and supported by an IT department. Likewise, the architecture of a conventional industrial control system is fixed by a controller that must be expensively modified to add a new component; in a networked peer-to-peer system, a sensor could be directly read by a local display, a control processor, and a remote server. Embedded networks also have significant energy implications; residential and commercial buildings were responsible for roughly 40% of the source energy use in the United States in 2004 [8]. More efficient buildings have been observed to recover at least 40% of that [9], but the cost and complexity of installing the required sensors and control systems has been an obstacle to their widespread adoption.

IØ brings the Internet down to embedded devices in a way that remains compatible with everything above it.

The name IØ emerged (initially as a joke) from the testbed installations to contrast the technological requirements of networked infrastructure from those of the high-speed Internet2 network [10]. These attributes were expressed in eight guiding principles; none of these alone is new, but their intersection is.

- *IP to leaf nodes:* Because IØ can reduce the cost of IP connectivity from tens of dollars to dollars (or less) per node, there isn't an economic reason to switch to something else for the last hop to a device, and retaining IP brings the Internet's interoperability and scalability directly to embedded devices rather than requiring the configuration of gateways for protocol conversion used by alternate standards for embedded networks.
- *No performance numbers:* One of the most dramatic aspects of the growth of the Internet has been in speed, progressing by six orders of magnitude from the ARPANET's original 56 kb/s to today's 40 Gb/s OC-768 backbone. This was enabled by the absence of performance numbers in the Internet specifications, allowing IP packets to travel over transports that weren't imagined in 1969. Many of the more recent alternatives for embedded networking simultaneously specify a logical protocol, its physical representation, and the allowable data rates. While this specificity eases implementation for the originally-intended application, it has the consequence of embedding technological assumptions that constrain future growth.
- *Compiled standards:* Fitting IP into a light switch or thermostat requires simplifying both the hardware and the software; even though silicon scaling enables increasing integration, device complexity still imposes an overhead in the cost of design, fabrication, packaging, processing, and power consumption. An IP protocol stack can fit in a few hundred or thousand bytes of microcode by jointly implementing the parts used by a particular application rather than the norm of separately writing each layer and then imposing the overhead of interlayer message passing. Layering is a useful abstraction that provides modularity for future developments [11], but this generality does not need to be retained once it is built into an embedded device [12], analogous to the compilation of high-level code for execution in a target processor.
- *Open standards:* This shouldn't need comment, but does; along with their technological motivations, competing standards for embedded networking have also been driven by proprietary concerns. For an IØ device to be able to join the rest of the Internet, its specification must be (and is) an open one.
- *Peers don't need servers:* In the Barcelona installation the lights and switches stored pointers for their associations so that servers could add value to the network but weren't required to run the house, much as a search engine helps organize the Internet but doesn't operate it. Many of the alternative approaches for embedded networking impose the need for an external server for two devices to be able to interact; along with bringing IP to leaf nodes, those devices should have the resources required to independently implement their functionality.
- *Physical identity*: A networked light switch can have multiple names: its physical location (the switch by the door), an address associated with the network it is connected to (192.168.1.101), a name on that network (myswitch.myhome.mynetwork), and possibly also a persistent hardware address that's independent of the network (00:0B:5D:8E:87:2D). For peers to not need servers, a device must be able to generate and associate these names by physical interaction with the device rather than requiring remote operations. For example, in the Barcelona installation, devices generated randomized IP addresses [13] to avoid the need for serialization or an address server. Pressing a programming switch on a light and then operating a light switch caused them to exchange these addresses, relating the physical and logical identities without requiring explicit knowledge of the latter.

- *Big bits:* This is where IØ differs most from recent networking practice. The duration of a bit and its speed of propagation define a size. Electromagnetic signals in wired or wireless networks travel on the order of the speed of light, or about 300 m/μs. If a bit is smaller than the size of a network, then it's necessary to impedance-match junctions to eliminate reflections. On the other hand, if a bit is larger than a network, the transient response to it can equilibrate independently of the topology of the network. Modern computer networks operate in the limit of small bits, but for low (and even not-so-low) data-rate devices, using bits that are big enough to settle on the local network eliminates the need for impedance-matched hubs in wired networks or agile transit-receive switching and collision detection in wireless ones.
- *End-to-end modulation:* In the near-field limit for big bits, signals can equilibrate. In the time domain, this corresponds to communicating in impulse responses. Although the transient dynamics depend on the details of the medium, information can be communicated solely in the occurrence of an event rather than its frequency, amplitude, or phase. Much as Morse code can be carried by any medium that supports a transient disturbance (e.g., clicked on a telegraph, banged on a pipe, or flashed from ship to shore), IØ encodes bits in the timing of the onset of an impulse response (a click) so that not just the data in a packet but also its modulation can be carried end-to-end.

End-to-end modulation enables interdevice internetworking.

An IØ IP packet is sent serially as conventional ASCII bytes with serial line IP (SLIP) [14] framing. However, instead of the usual RS-232 voltage levels, an easily-implemented pulse position code is used with two time slots per bit, with a one represented by a click in the first interval and a zero by a click in the second (Figure 2). Clicks in both intervals identify the start and stop bits, self-consistently providing the byte framing, time origin, and data rate.

Figure 3 shows a 240-bit UDP/IP IØ packet, comprising 160 bits for an IPv4 header, [3], [4] 64 bits for the UDP header, [15] and 16 bits for the SLIP framing. This looks like a barcode and in fact, could be used that way and scanned

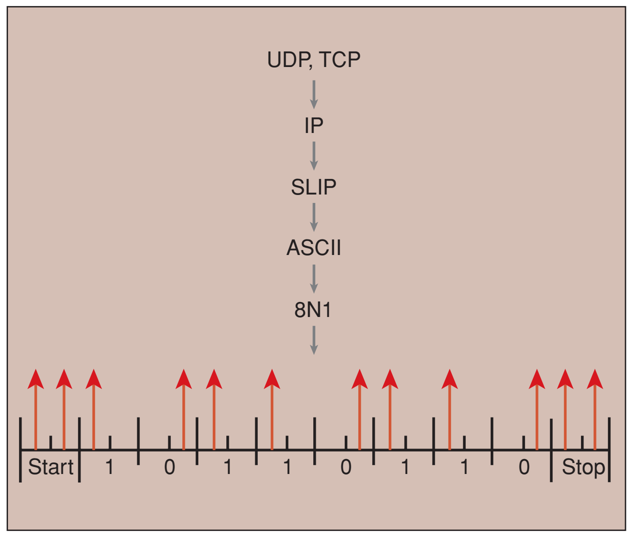

2. An IØ byte.

3. A UDP IØ packet.

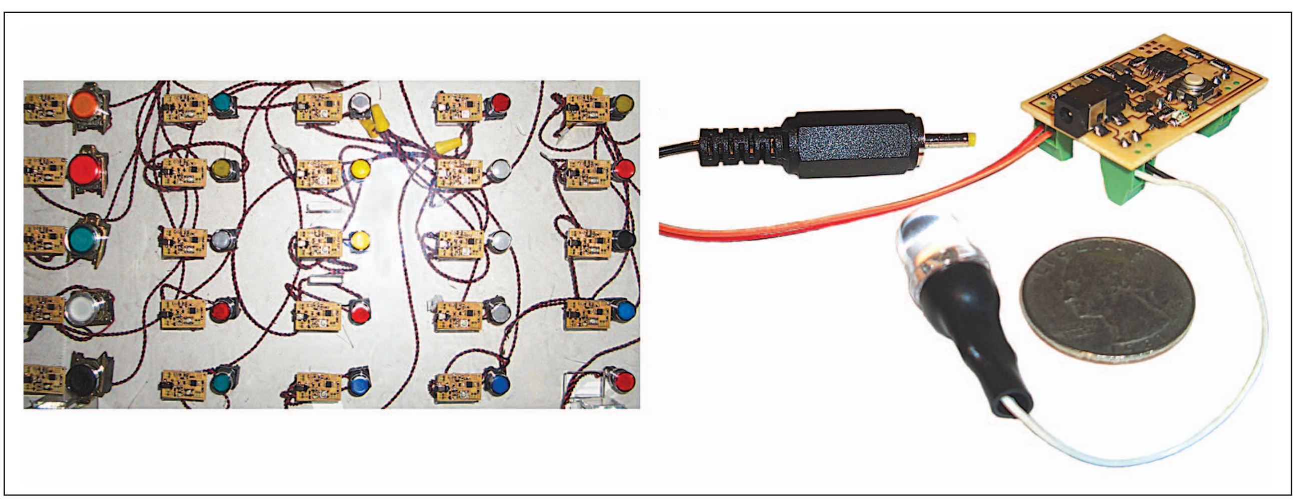

4. A dc powerline IØ node and 25-node control panel network.

optically, then the raw signals could be carried over any other IØ transport, wired or wireless, electromagnetic, acoustic, or optical.

For example, Figure 4 shows a dc powerline IØ implementation targeted at taking advantage of the 24- or 48-V power distribution in industrial and building control systems. The clicks are capacitively coupled over the dc supply, providing both energy and information for a control input and switched load via a 0.65 mm power plug or screw terminal as an alternative to separate networking cables and connectors.

The schematic of this dc powerline interface is shown in Figure 5. A coupling capacitor is connected to the unregulated supply, with a pull-up resistor to the locally regulated supply. A click is generated by pulling the capacitor down with a MOSFET.

For the 0.01 μF coupling capacitor and 1.7 A MOSFET used here, the slew rate is

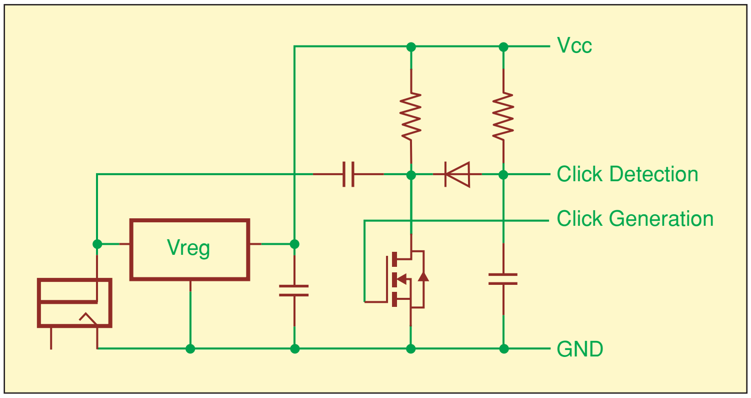

5. A dc powerline IØ interface.

$$\frac{dV}{dt} = \frac{I}{C} = 1.7 \times 10^8 \frac{V}{s},$$

giving an initial click duration on the order of 30 ns at 5 V. As shown in Figure 6, ringing in the wiring extends this to a few hundred ns.

To receive a click, the bypass capacitor is also connected to a diode detector followed by an RC network to stretch the click to a few μs, as shown in Figure 7, so that it can trigger a comparator.

Figure 8 shows a byte (10110000) formed from these clicks. The implementation of reading and writing IØ packets, including click generation and timing recovery, required 224 bytes of microcode for a minimal UDP packet (coincidentally, a byte of code per bit of packet data).

Beyond timing the arrival of individual clicks, the constraints of the IØ encoding can serve as a kind of modulation scheme for noise rejection and channel sharing. A low data-rate device may be able to time the click arrival to a small fraction of the click spacing, a resolution on the order of 100 ns in the dc powerline example above. As shown in Figure 9, these times can be used in a decoding tree to self-consistently reject spurious events that are not compatible with the click framing, and separate interleaved click streams to do statistical time division multiple access (TDMA) for channel sharing. The implementation of collision detection for carrier sense multiple access (CSMA) is also simplified by the relatively small

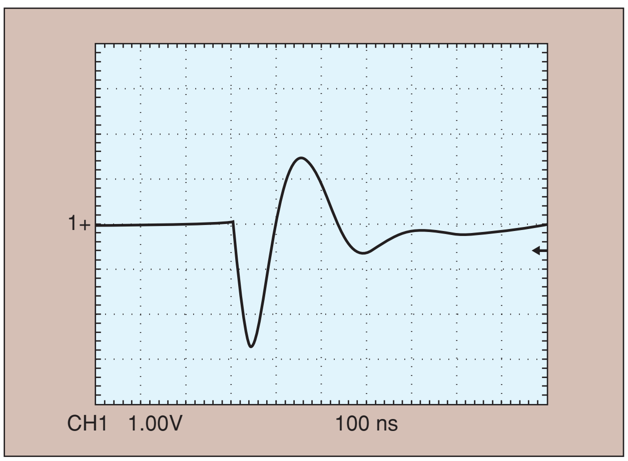

6. A dc powerline click.

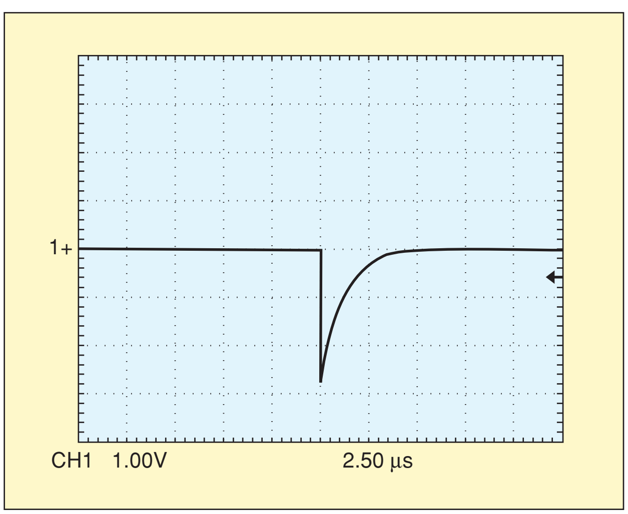

7. Detected click.

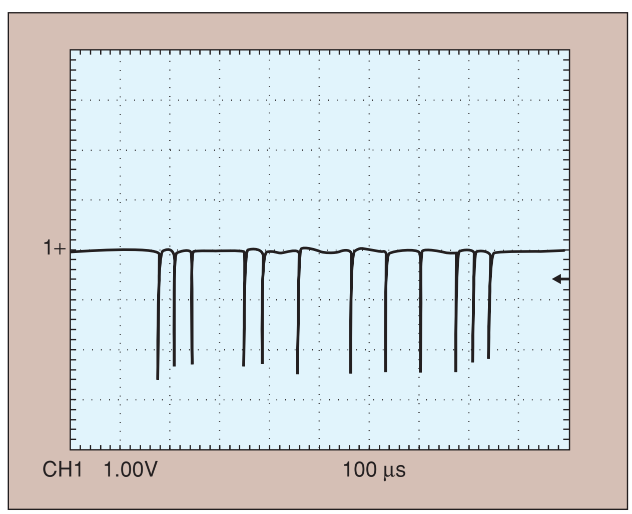

8. An IØ byte (10110000).

fraction of a bit duration that is associated with propagation compared to the much greater fraction during which the impulse response equilibrates and a potential collision can be detected before transmission.

Figure 10 shows the construction of an IØ packet as viewed by a conventional network sniffer and its payload, a Web page, after it has passed through a stateless IØ to RS-232 SLIP bridge (Figure 11). This Web page was sent via UDP (rather than transmission control protocol (TCP) as specified by the HTTP protocol [16]) to eliminate the need for a resource-constrained IØ device to maintain the state of connections for serving Web pages that fit in a single packet, and the need for exchanging three synchronize-acknowledge (SYN-ACK) packets over a bandwidth-constrained network

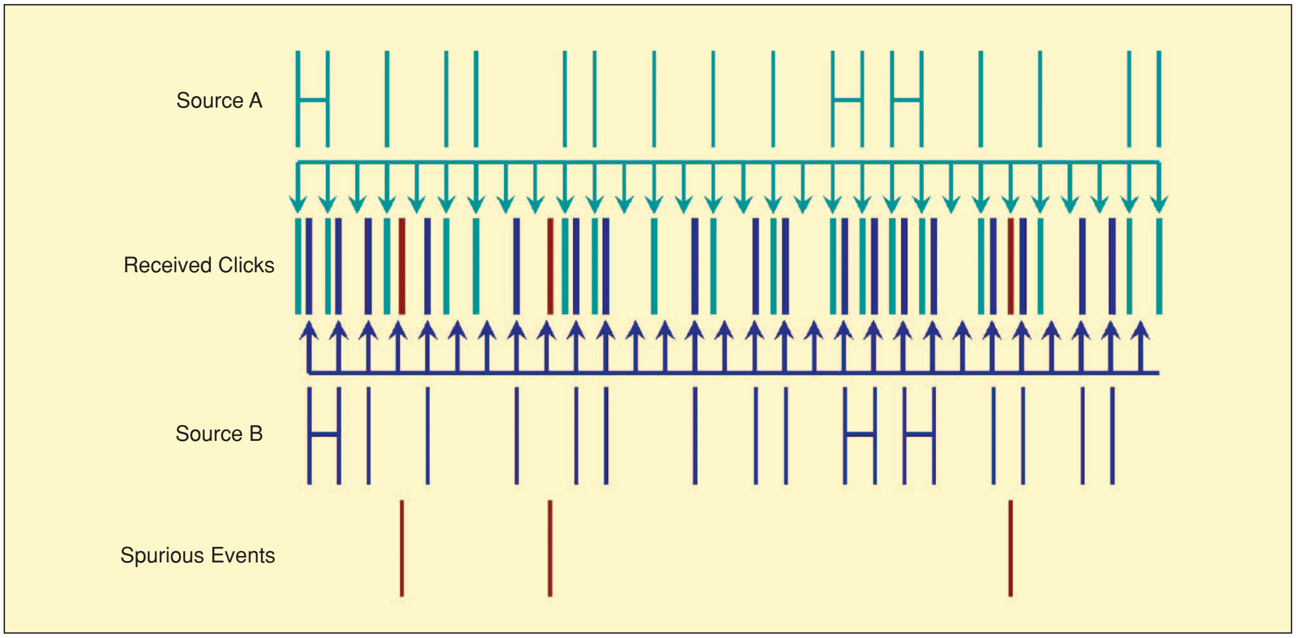

9. Self-consistent IØ decoding (horizontal bars show recovered click and byte framing).

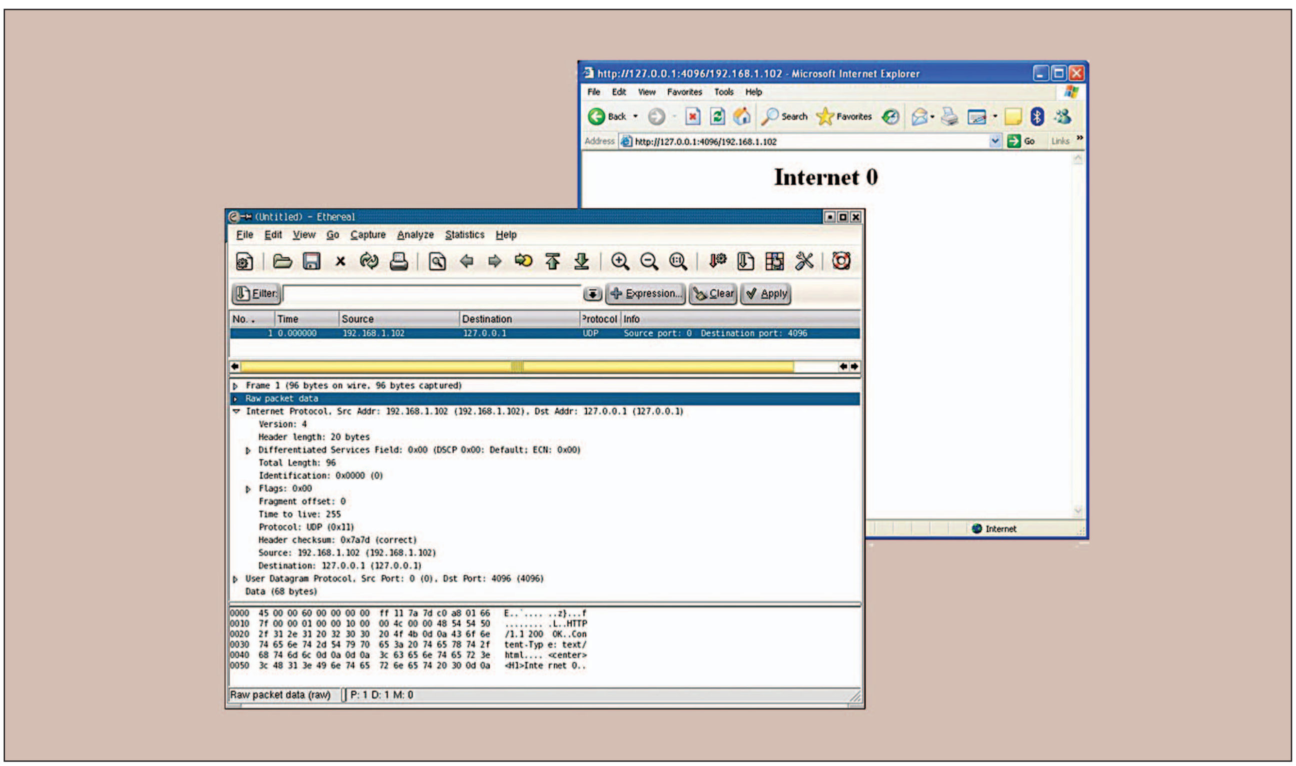

10. A THTP Web page.

before data can be transmitted. This trivial hypertext transfer protocol (THTP) transport (implemented here by a UDP to TCP bridge in the SLIP interface) is analogous to the relationship between FTP (which uses TCP) and trivial file transfer protocol (TFTP) [17] (FTP over UDP).

Information can be communicated solely in the occurrence of an event rather than its frequency, amplitude, or phase.

Because an IP packet is used as the native format in an IØ network, there is no need for address resolution protocol (ARP) traffic to resolve hardware media access control (MAC) addresses [18]; if 128-bit IPv6 addresses are used for global routing, [19] this corresponds to $2^{128}/4\pi\,6378137^2 = 6.6 \times 10^{23}$ available addresses per square meter of the Earth's surface (i.e., roughly Avogadro's number). IØ installations have variously used randomized address self-assignment within subnets, automatic assignment by an address server, or manual address assignment by an installer.

Since the packet representation is unchanged across an IØ network, the physical (PHY), MAC, and network protocol layers effectively merge. This means that it's possible to interconnect networks at the physical link rather than software protocol layers [20] (Figure 12). For example, Figure 11 also shows a stateless bridge between dc powerline and IR optical IØ networks; all it needs to do is generate a click in one of the media in response to receiving one in the other. Just as IP-based internetworking enables the Internet's end-to-end architecture, IØ's interdevice internetworking enables end-to-end modulation, with analogous implications for things.

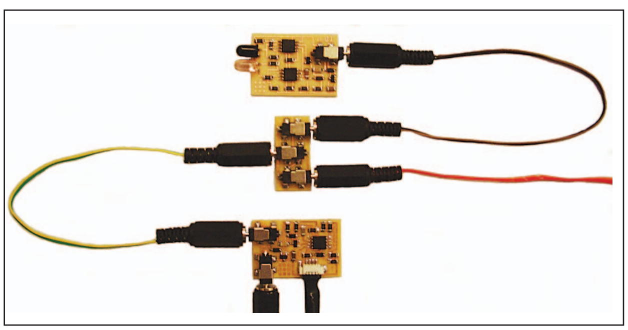

11. Stateless IØ dc powerline RS-232 and optical bridges (bottom and top), and hub (middle).

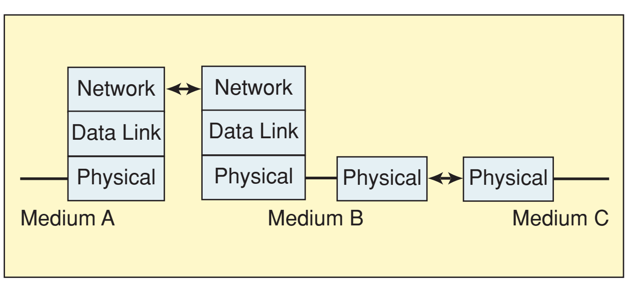

12. Comparison between internetworking (left) and interdevice internetworking (right).

Consider the rollout of the EPC and ISO 18000 standards for RFID [21]. One of the most significant obstacles has been neither the cost of the tags nor the readers but the cost of configuration. With great effort, a 96-bit standard was defined for the contents of a tag, but this provides no guidance for what a tag reader should do with this information, hence the need for configuring readers and middleware to send tag data to servers. And any changes in the use of the tags must be reflected in changes in this installation. As viewed from IØ, however, inductive loading is just one more channel that can be used to time clicks. Done this way, each IPID tag carries a packet that contains its own routing information, subsequent tags can perform different functions, and a reader is reduced to being a bridge or gateway between wireless and wired transports.

Far-field wireless links can be implemented in the same way, by using band-limited impulses as is already done in ultrawideband (UWB) radios [22]. The most important uses of IØ, however, are likely to take advantage of underused communication channels, including visible as well as IR optical transports, ultrasonics for unregulated wireless links that remain confined within an acoustic space, and RF industrial, scientific, and medical (ISM) bands that lack dedicated protocols. The diversity of these options allows the physics of a channel to be associated with its content, such as using a near-field transport for secure key distribution for embedded cryptographic authentication [23]. And they can each be introduced without adding to the further proliferation of incompatible standards. Just as the IP protocol is not optimal for any one thing but is good enough for almost anything, this generality is appropriate when minimizing cost and complexity is more important than maximizing performance for a specific task.

Conversely, IØ is inappropriate when saturating a channel's capacity is more important; it does not replace today's many optimized network transports. And IØ does not specify how a light switch should describe itself to a light bulb; there are already many domain-specific device-independent device-description languages, including UPnP for consumer electronics [24] and BACnet for building automation [25]. These standards sit above IØ once IP connectivity has been established. Most importantly, IØ is not an alternative to today's Internet (call it Internet 1), it extends it. Just as Internet2 speeds up the Internet, IØ brings the Internet down to embedded devices in a way that remains compatible with everything above it.

Although the boundary between IØ and the rest of the Internet can be stateless, there are many possible reasons to want functions in those interfaces, including proxying access, caching data, managing identities, and providing security [26], [27]. Likewise, an IØ bridge between media types need not decode a packet to pass it, but can in order to route it. In each case, these configurations can be introduced based on the needs of an application rather than dictated by boundaries between transport media.

Routing between IØ peers through the existing Internet does implicitly rely on its servers; growth of the former will challenge the capacity of the latter. But there are encouraging hints that the distinction between leaf nodes and central servers can ultimately disappear by deriving and implementing networking protocols as distributed solutions to constrained optimizations [28]. This approach helps explain existing protocols and may eventually replace them with optimal adaptive alternatives, but even then the physical distinction between big and small bits that lies behind end-to-end modulation will still apply to any future successors to IP.

IØ can contribute to the future of networking in one more way: by reducing the cost of not just acquisition and configuration but also experimentation. It's much easier to alter a network of dollar-scale devices than it is to reprogram essential Internet servers. IØ simplifies the development as well as implementation of Internet connectivity, making it accessible to new people as well as new things [29].

ACKNOWLEDGMENTS

We would like to acknowledge the contributions of the students and colleagues who participated in the emergence of IØ, including David Dalrymple, Aaron Rosado, Jon Santiago, Karen Sollins, Rob Beverly, Raffi Krikorian, H. Shrikumar, Matt Hancher, and Pehr Anderson from MIT; Todd Snide from Schneider Electric; Doug Johnson and Arshan Poursohi from Sun Microsystems; Kerry Lynn from Cisco Systems; Terry Riley from New York's Museum of Modern Art; Paul Magid and the Flying Karamazov Brothers; and Vicente Guallart from Metapolis. This work was supported by NSF (CCR-0122419), DTO (0705N7KZ-PB), and the Center for Bits and Atoms.

REFERENCES

[1] X. Meng, Z. Xu, B. Zhang, G. Huston, S. Lu, and L. Zhang, "IPv4 address allocation and the BGP routing table evolution," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 1, pp. 71–80, 2005.

[2] N. Gershenfeld, R. Krikorian, and D. Cohen, "The Internet of things," *Sci. Amer.*, vol. 291, no. 44, pp. 76–81, 2004.

[3] V. Cerf and R. Kahn, "A protocol for packet network intercommunication," *IEEE Trans. Commun.*, vol. 22, pp. 637–648, 1974

[4] Internet Protocol. IETF RFC 791, 1981 [Online]. Available: http://www.ietf.org/rfc/rfc0791.txt

[5] J.H. Saltzer, D.P. Reed, and D.D. Clark, "End-to-end arguments in system design," *ACM Trans. Comput. Syst.,* vol. 2, no. 4, pp. 277–288, 1984.

[6] M.S. Blumenthal and D.D. Clark, "Rethinking the design of the Internet: the end-to-end arguments vs. the brave new world," *ACM Trans. Internet Technol.*, vol. 1, no. 1, pp. 70–109, 2001.

[7] U.S. Census Bureau, "Economic Census," 2002.

[8] *Indicators of Energy Intensity in the United States*, U.S. Department of Energy. [Online]. Available: http://intensityindicators.pnl.gov

[9] B.V. Neida and T. Hicks, "Building performance defined: The Energy Star National Energy Performance Rating System," U.S. Environmental Protection Agency, 2001 [Online]. Available: http://www.energystar.gov

[10] Internet2 [Online]. Available: http://www.internet2.edu/

[11] A.S. Tannenbaum, *Computer Network,* 4th ed. Englewood Cliffs, NJ: Prentice Hall, 2003.

[12] D. Clark and D. Tennenhouse, "Architectural considerations for a new generation of protocols," in *Proc. ACM SIGCOMM*, 1990, pp. 200–208.

[13] Dynamic Configuration of IPv4 Link-Local Addresses, IETF RFC 3927, 2005 [Online]. Available: http://www.ietf.org/rfc/rfc3927.txt

[14] A Nonstandard For Transmission Of IP Datagrams Over Serial Lines: SLIP, IETF RFC 1055, 1988 [Online]. Available: http://www.ietf.org/rfc/rfc1055.txt

[15] User Datagram Protocol, IETF RFC 768, 1980 [Online]. Available: http://www.ietf.org/rfc/rfc0768.txt

[16] Hypertext Transfer Protocol—HTTP/1.1, IETF RFC 2616, 1999 [Online]. Available: http://www.ietf.org/rfc/rfc2616.txt

[17] The TFTP Protocol, IETF RFC 1350, 1992 [Online]. Available: http://www.ietf.org/rfc/rfc1350.txt

[18] An Ethernet Address Resolution Protocol, IETF RFC 826, 1982 [Online]. Available: http://www.ietf.org/rfc/rfc0826.txt

[19] Internet Protocol, Version 6 (Ipv6), IETF RFC 2460, 1998 [Online]. Available: http://www.ietf.org/rfc/rfc2460.txt

[20] Information technology—Open Systems Interconnection—Basic Reference Model: The Basic Model, ISO/IEC 7498-1, 1994.

[21] Information technology—Radio frequency identification for item management—Part 1: Reference architecture and definition of parameters to be standardized, ISO/IEC 18000-1, 2004.

[22] M-G.D. Benedetto and G. Giancola, *Understanding Ultra Wide Band Radio Fundamentals*. Englewood Cliffs, NJ: Prentice Hall, 2004.

[23] F.-X. Standaert, G. Piret, N. Gershenfeld, and J.-J. Quisquater, "SEA: A scalable encryption algorithm for small embedded applications," *Springer Lecture Notes Comput. Sci.*, vol. 3928, pp. 222–236, 2006.

[24] "The universal plug and play forum" [Online]. Available: http://www.upnp.org/

[25] "A data communication protocol for building automation and control networks" [Online]. Available: http://www.bacnet.org/

[26] K.R. Sollins, "Designing for scale and differentiation," *ACM SIGCOMM Comput. Commun. Rev.*, vol 33, no. 4, pp. 267–276, 2003.

[27] Design Issues in Internet ø Federation, K. Sollins and J. Li, 2006 [Online]. Available: http://dspace.mit.edu

[28] J. Wang, L. Li, S.H. Low, and J.C. Doyle, "Cross-layer optimization in TCP/IP networks," *IEEE/ACM Trans. Networking*, vol. 13, no. 3, pp. 582–268, 2005.

[29] N. Gershenfeld, *FAB: The Coming Revolution on Your Desktop—From Personal Computers to Personal Fabrication*. New York: Basic Books, 2005.

Neil Gershenfeld is with the MIT Center for Bits and Atoms, Cambridge, MA. *Danny Cohen* is with Sun Microsystems, Menlo Park, CA. CD ■