CHAPTER 7

Examples and Applications

Lots of information about Internet of Things applications has been published to date, but virtually all these examples assume a continuation of *current* networking architecture models. Specifically, IPv6 extended to the very edge of the network, with end devices powerful enough (in terms of processor, memory, etc.) to run an IP protocol stack. But as has been described in preceding chapters, this architecture is unsuitable for the "next wave" of IoT end devices to be brought onto the network. They will simply be too cheap, too numerous, too hard to manage, and too varied to support the traditional networking model.

Another incorrect assumption made about the future of the Internet of Things is that the *data* models will remain much the same as today: well-defined, one-to-one relationships between IP–equipped end devices and big data servers at the core of a network accessed over the "cloud." But this traditional approach cannot *fully exploit* the potential richness and power of the IoT for a number of reasons:

- Data handling and storage at the big data servers
- Impracticalities of end-to-end control loops
- Inability to exploit a publish/subscribe world made up of neighborhoods and affinities of end devices

This chapter will first explore the impact each of these issues has on potential new IoT applications and then provide specific examples of new potential IoT applications.

Controlling the Cacophony

Machine-to-machine data interchanges are currently tedious: all raw, device-specific data must first be sanitized and then formatted to conform to big data representation schemes based on application programming interfaces (APIs) such as Representational State Transfer (REST) or Simple Object Access Protocol (SOAP). Next, the IP-stack-based transceivers on current end devices must send the data without collision with other IP device traffic, often using Carrier Sense Multiple Access with Collision Avoidance/Detection (CSMA/CA or CSMA/CD). This is a lot of work for a simple temperature sensor, with its restricted and terse purpose-built vocabulary.

The users of big data are interested in an integration of small data: the device-abstracted, protocol-abstracted information streams. The onus of converting sensor raw data to big, data-friendly "small" data cannot be easily delegated to every end device as has been imagined to date for the Internet of Things. Managing diverse device driver interfaces and their specific interfaces and protocols rapidly spins out of control at the scope of the emerging IoT.

As the number of edge devices proliferates, the network effects of the traffic generated by billions of publishers and subscribers overwhelms the processor and memory processor enhancements enabled by Moore's Law (which is linear), as shown in Figure 7-1. Recall that machine-to-machine communities and their interactions are more akin to social networks; in other words, they are Metcalfe's Law or $O(n^2)$ (Order-*n*-Squared)–based. The data processing, storage, and networking requirements for cloud-based IoT analysis and control services will not be able to keep up with the deluge of small data emanating from the edges of the network. (They can barely keep pace, even in today's simplified and managed end-to-end thin-client IP applications.)

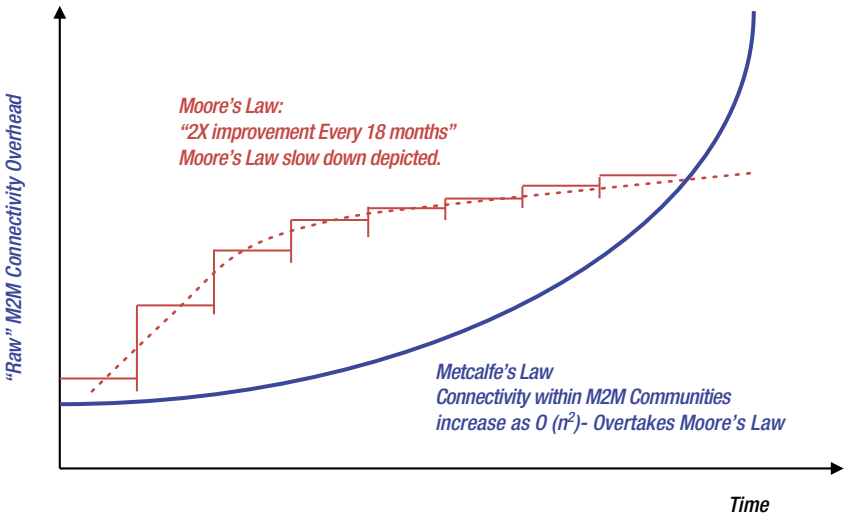

Figure 7-1. *Much of the current thinking on the Internet of Things assumes that constant hardware improvements (due to Moore's Law) will allow traditional networking schemes to be extended to the IoT. But in fact, the machine social network will grow much faster (Metcalfe's Law) and will require a more specialized architecture*

This will be true either for chirp-based networks or legacy IP end devices; the amounts of data are simply too great. For this reason, the emerging IoT architecture removes the overhead of the task of aggregating and transporting data from *both* the end devices *and* the big data servers. It instead segregates it within propagator nodes that can be deployed near the edges of the network (refer to Chapter 4).

Intelligence Near the Edge

The emerging Internet of Things architecture also provides for migration of intelligence toward the edge of the network in the form of publishing agents within the propagator nodes and/or distributed integrator functions. Through these capabilities, IoT applications may rely on these distributed intelligences to manage the conversion of chirp data streams to and from end devices such as sensors and actuators to small data flows that are more easily consumed by the big data integrator functions. This process will enable the rapid proliferation of a dizzying variety of applications using very simple, low-cost, or intermittently available end devices that are simply not possible with traditional IP networking schemes.

Incorporating Legacy Devices

An added benefit of this architecture is that applications requiring more-sophisticated end devices that *do* justify the cost and complexity of IP on board (video surveillance, for example) may also use the same architecture, easing the load on big data servers and making possible the extended publish/subscribe network of neighborhoods and affinities (see the following sidebar). The core objective is to encourage and manage a more equitable *division* of labor, one that only improves with time, as devices at the edge are permitted to be simpler in function. Simpler devices will rapidly proliferate at the edge once a supporting network infrastructure is in place that can both manage chirp streams on behalf of the end devices *and* create small data flows suitable for the benefit of big data integrator functions. See the following "Nailing a License Plate to a Stump" sidebar.

NAILING A LICENSE PLATE TO A STUMP

Many of today's Internet of Things commentators have hailed the address expansion incorporated within IPv6 as the solution for the IoT. And it is certainly mathematically true that IPv6 creates more than 340 undecillion (more than 3.4×10^{38}) potential addresses, which some have said is enough to assign one to every atom on the surface of the earth. (For logistical reasons, the practical limit is likely much less.) But compared with the roughly 4 billion unique addresses possible with IPv4, that is indeed a substantial improvement and certainly sufficient to address every potential end device in any conceivable Internet of Things.

But this analysis confuses addresses with functionality. It would certainly be possible to nail an automobile license plate to a tree stump (see Figure 7-2), which would make the stump uniquely identifiable and thus potentially addressable. But it does not magically enable the stump to drive away on the highway like a car. It is obviously missing the horsepower (a motor), means of transportation (wheels), and intelligence (a driver) to make any usefulness on the highway impossible.

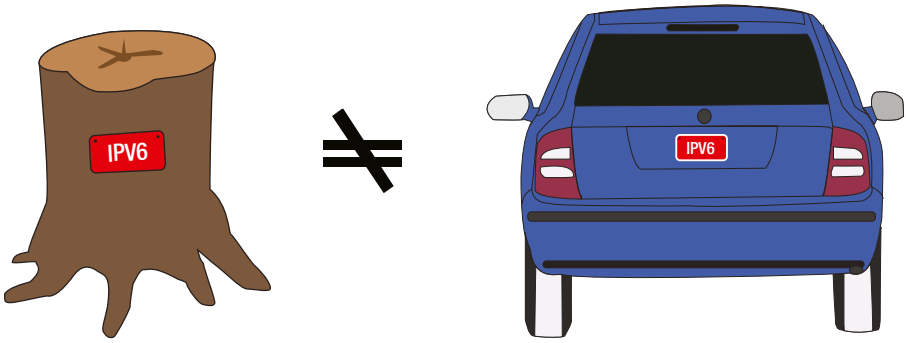

Figure 7-2. *Addressing is not performance*

In the same way, the capability to address an end device sensor or actuator is only a small part of the issue in the IoT. Without burdening the end device with horsepower (memory and processor), means of transportation (IP stack), and intelligence (central management and oversight), its data cannot make it to the "information superhighway," either.

Thus, the IPv6 address space alone doesn't solve the essential application problem in the IoT: enabling the connection of billions of end devices that are too simple to support full networking. The new emerging architecture of the Internet of Things creates the simple chirp structure that allows for the development of applications without demanding untenable requirements at the end devices.

Staying in the Loop(s)

One of the key challenges of extending legacy IP architectures to the Internet of Things is the inherent constraint created by using a protocol originally developed for host-to-host communications (peer-to-peer, by definition) to the very different and inherently asymmetrical world of the IoT. One of the impacts of this legacy on IoT applications is the difficulty of managing control loops over long distances and via the nondeterministic global Internet. Unlike a host-to-host interaction, IoT end devices and actuators often have very little or no intelligence of their own, so the task of managing them would fall to integrator functions accessed via some sort of round–trip control loop over a long-distance link.

Round-trip control via IP and the global Internet is an impractical means of controlling simple end devices at the extreme edges of the Internet of Things, especially because some may be only intermittently connected. Instead, localized control through distributed intelligence in nearby propagator nodes allows autonomous or semiautonomous control via on-board integrator functions or publishing agents (see Figure 6-2).

Given the delay and jitter (variation in delay) inherent in the global Internet, the existing IP network is a cumbersome and ultimately impractical solution for control of myriad simple end devices, as shown previously. But (as described in Chapter 6) the

emerging IoT allows the control loops to be *decoupled* and thus become *isochronous*. An efficient lower-level local control loop may be in place between the propagator node and end device, whereas occasional updates and exceptions are communicated upstream to subscribed integrator functions. In turn, "tuning" and configuration messages may occasionally be received from integrator functions for implementation at the local end device actuator.

Okay on their Own

This multilevel control will allow for IoT applications that may function substantially autonomously in real time most of the time, including for extended periods when out of communication with a distant integrator function. This option of a much more rapid response from a local publishing agent or integrator function is a key feature of the emerging Internet of Things for autonomous and semiautonomous (advise and consent) tasks.

This distribution of intelligence throughout the emerging IoT architecture bodes well for its future. More autonomy means less supervisory control and less drain on resources required for round-tripping. The predictive elements (integrator functions) become more seasoned at being proactive, and reactive elements give way to proactive behavior. The overall system evolves to be more predictive, lean, and agile.

All the World Is a Subscription

Another legacy limitation of the host-to-host nature of IPv6 is that connections are inherently point-to-point between known devices. (Routers are required to create and manage these relationships.) This creates isolated "silos" of data, in that there are separate sets of end devices deployed for different functions. So in contrast with the emerging IOT, they may not be able to contribute their information to an integrator function, even if the combination would provide much more powerful information.

As described in Chapter 5, the emerging Internet of Things architecture is not limited by the concept of preset device-to-device relationships. Instead, integrator functions will create information neighborhoods made up of a wide variety of small data flows forwarded by propagator functions from many chirp data streams. The lowly chirp-enabled sensor is now a participant of the connected world, without changing its fundamentally simple function of publishing a specific category of real-time, raw, simply formatted data.

Exploring Affinities

Many new classes of IoT applications will be possible, in which integrator functions seek out potentially useful information by examining affinities between many publishing sources (see Figure 7-3). An example of this is temperature, pressure, and vibration; or small data streams that seem to vary in relationship to one another, or in relationships to an Internet data source, such as weather reports. Again, this would be more difficult in the traditional IP point-to-point environment, in which different device types tend to be segregated from one another.

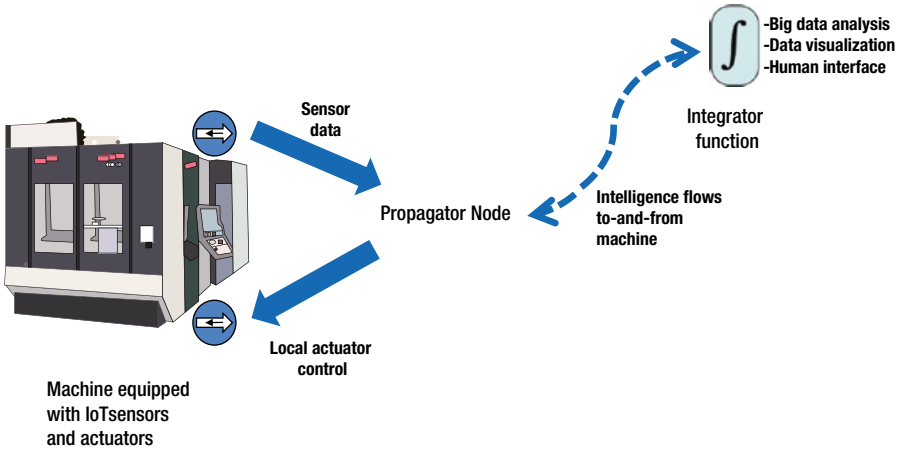

Figure 7-3. *Unlike traditional networks, many important and illuminating relationships may be unknown at the time of installation of an Internet of Things application. But over time, integrator functions may expand their "neighborhood" of information sources by exploring other data streams that share some affinities with the existing neighborhood of data sources. These new sources may be included for a time to "test" their usefulness, and may be later dropped or replaced and new sources explored*

In a world in which the data emanating from many IoT applications may be marked as public small data streams by their owners, the potential exists for incredible insights and efficiencies of scale as integrator functions build extensive subscriptions. The key aspect setting these applications apart from legacy Internet of Things applications built on traditional IPv6 networking is that the relationships between end devices and integrator functions may be unknown at the outset. Instead, they are built and refined over time by the integrator functions. A larger social network for data exchange emerges. The data streams will span the gamut: chirp sensory data, changing subscriber patterns, preferred data routing paths on specific days, and so on. End devices, propagator nodes, and the publishing agents within will "belong" to multiple "information social networks" informed by neighborhoods of subscribed data.

Social Machines

The information social networks will be free to grow to quite large sizes simply because machines are not constrained by Dunbar's Number (which theoretically limits the maximum interactions that humans can actively support to fewer than 200). Freed from legacy-style, predefined peer-to-peer interactions with thousands (or millions) of end devices by the distribution of networking intelligence and decoupling of control loops, integrator functions will be able to digest unprecedented amounts of distilled and directed data (see Figure 7-4).

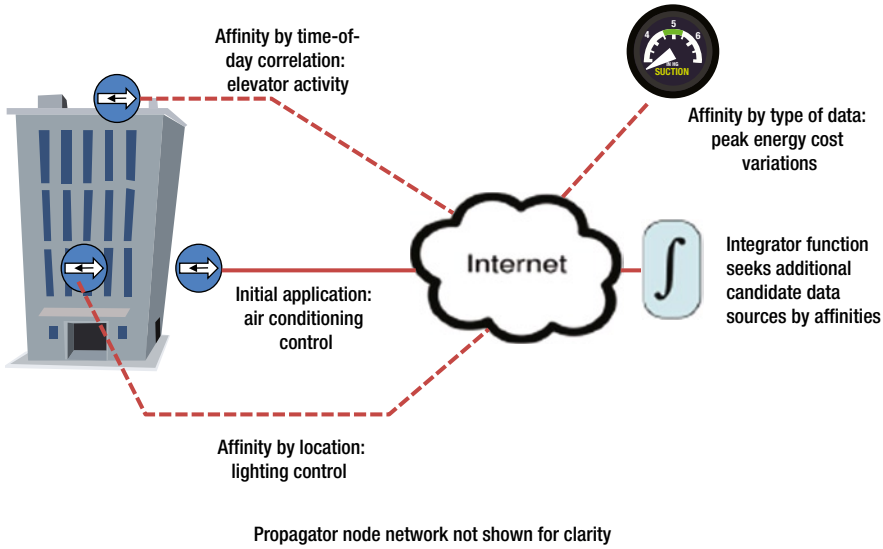

Figure 7-4. *Distributed networking capabilities in an ever-expanding mesh of propagator nodes will provide more and more refinement in pruning and tuning of chirp data streams to create more efficient small data flows. This growing efficiency will allow integrator functions to analyze more end points with less processing of individual data packets. The IoT will become more useful as the architecture expands*

From the machine-to-machine IoT perspective, intelligence is inferences drawn over time from multiple and diverse data sources. The proliferation over time of more and varied chirp–enabled end devices and propagator nodes will continue to expand the available universe of potentially interesting data streams. As more propagator nodes are added with the growth of the Internet of Things, the overall application data exchange flow rates will continue to improve linearly through the proactive use of pruning/aggregation/exception handling, as shown previously. But the Metcalfe's Law network effect of the *information* will be growing even more rapidly in an $O(n^2)$ relationship.

Agriculture

Managing an agricultural enterprise is a difficult, multivariate endeavor; many man-made and natural factors are in play. In Figure 7-5, a lower–level local control loop applied by a distributed integrator function autonomously "manages" the actuators controlling the irrigation system valves (when they turn on or off, based on local moisture level sensors). This isochronous loop monitors and controls the amount of water applied locally within specific zones, avoiding over- or under-watering. But there is no need for round-trip control by a distant higher-level integrator function or for burdening a distant integrator function with a continuous stream of data from local moisture sensors.

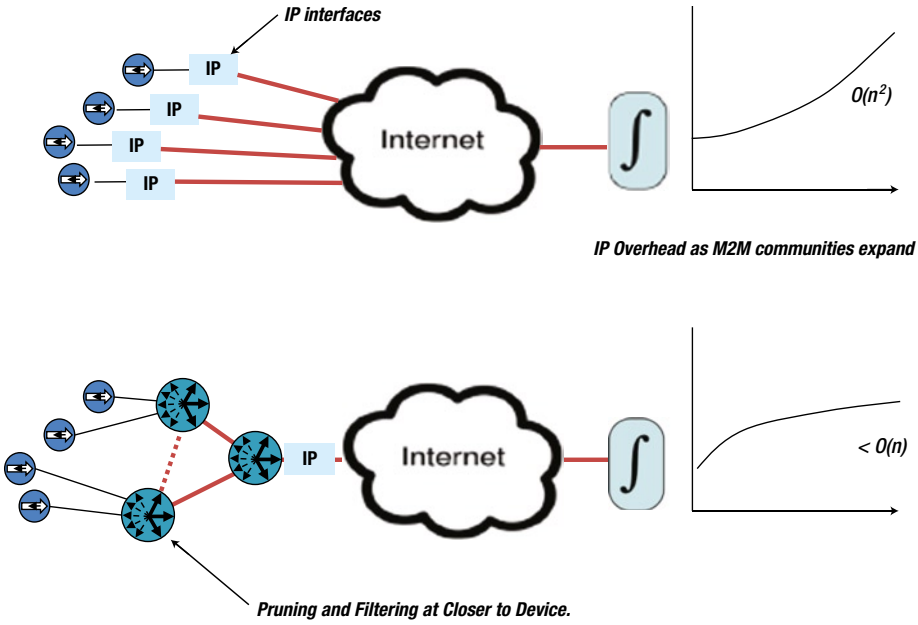

Figure 7-5. *A wide variety of sensors and other data sources combine to optimize yields and profitability from a farm field. Although lower-level control loops might monitor and manage irrigation through distributed integrator functions, additional integrator functions deployed at a "higher" level may take a wider variety of information into their analysis*

In the less-than-perfect world, however, patterns of water *absorption* by the crops are not easily discernible by these lower-level control loops. An airborne drone equipped with appropriate sensors (such as infrared) may be deployed to scan the corn field and collect a more global view of the terrain and where more water may be needed. The drone provides this information through its wireless interfaces to a smartphone or other general-purpose processor running an integrator function operating at a "higher" level than the moisture-sensor-irrigation-valve control loop. The integrator function correlates this to the current sprinkler map and fine-tunes it to ensure more even water distribution. Farmers may also be provided with suggestions regarding changing the terrain to provide slopes for more efficient irrigation. A few weeks later, the drone conducts another survey. Over time, the lower control loop, in conjunction with the upper control loop, generates a more comprehensive view of its region of interest.

The cost of one sophisticated but "remote" sensor (the drone) may be lower than implanting multiple simpler moisture sensors. The control loop is still being closed, and the drone and sensor ensemble is more modular, reusable, and upgradeable. Economies of scale kick in. The drone may be used by a farming community, a shared resource. If the control loop is being monitored weekly, the same drone can be used to close multiple control loops in adjoining farms. A swarm of such drones can be used to cover large areas in a low-cost, scalable manner.

The integrator functions may also discover and subscribe to a variety of other data streams and sources to create a richer combination of information. Weather forecasts, spot produce prices, the current cost of transportation, the availability and cost of contract field workers, and many other factors may be taken into account. Some of these other data streams will not be generated through the farmers' own efforts, but made available by others with public markers on the data streams. Trend analysis of these factors over time may point to the ideal moment to harvest the crop for maximum profits. The community of farmers, through shared resources and the integration of many data sources (some unforeseen at the time the application was first deployed), can compete more effectively with those in other regions.

SHOW ME YOUR DATA; I'LL SHOW YOU MINE

The previous farming example suggests that the group shares information of similar type for a similar goal: all are farmers in a specific area. But because the emerging Internet of Things architecture is fundamentally based on a publish/subscribe model, the creators and consumers of data streams may not be always have as much (or anything) in common. For example, a restaurant owner might want to know about foot traffic in a nearby shopping mall in order to target ads for video kiosks or instant coupons on social media. A trucking company might want to know about unusual traffic patterns created by an accident and detected by in-pavement or video-intersection sensors in order to reroute their fleet.

These and many more as-yet-unimagined opportunities may exist for sharing of data streams that are already being created. A nonmonetary "exchange" market place might emerge—or even one based on market pricing or auction models. Because the chirp protocol is category-based and publish-friendly, chirp streams and small data flows from nonaligned organizations can be acted upon. A key enabler of these potential exchanges is that the entire IoT architecture is oriented toward a publish/subscribe model rather than defined peer-to-peer relationships, even at the lowest levels. The chirps from the simplest sensor can be shared with an unlimited number of integrator functions without any change or reconfiguration required.

Home Health Care

The agricultural example described cooperative use of a population of Internet of Things sensors and actuators by defining information neighborhoods of related elements and seeking out affinities of potentially related and pertinent information. But other IoT applications will be more restrictive in their deployment and operation. The need for secure, private, and purpose-built communications will proliferate within local machine-to-machine communities.

In Figure 7-6, a small private IoT community of "vital signs" sensors feeds into a local integrator function for analysis and pattern matching near the network edge (the patient). Proactive monitoring and management of medication dosage and/or additional home care is also logically delegated to a private patient's home via additional sensors or inputs from a caregiver's smartphone.

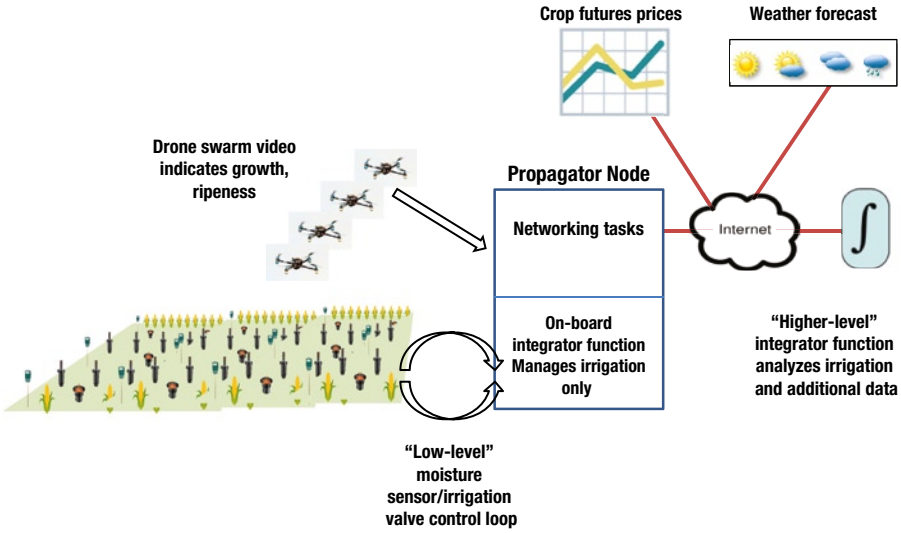

Figure 7-6. *A home heath IoT application might use a variety of sensors and other inputs, including wearables and ingestibles, to create a complete picture of a convalescing patient's condition. Local alarms may be triggered for particular threshold readings or combinations of conditions and events. Periodic regular reports and occasional exceptions may be forwarded to off-site medical personnel for emergent response or long-term analysis*

Because individual sensors need not be burdened with the processor, power, and memory overhead required to support IPv6, they might be smaller, lighter, cheaper, and less invasive, which could include wearable and ingestible form factors. With the local analysis enabled by the emerging Internet of Things architecture, readings from many sensors may be considered together, along with variables such as room temperature and time of day, allowing a more-sophisticated combined analysis rather than simply alarming on one boundary condition.

This private information neighborhood becomes adaptive and self-learning to provide the first-tier reporting and response initiation autonomously. If the patient's heart rate or breathing becomes erratic, the patient and caregiver will know immediately based on alarms and other feedback devices triggered by the local integrator function. Notice of the exception condition would also be transmitted to distant medical personnel, who are made aware of it immediately. The stimulus-response is more proactive, potentially averting or reducing the severity of an event. Round-trip communication to a distant integrator function is now restricted to escalated issues only.

Safe and Efficient Process Control

Natural resource– and commodities-processing enterprises such as oil refineries present a demanding "analog" application. Maintaining liquid flow rates, temperatures, and pressures may be critical to ensuring higher yields of end products. Keeping values within tolerances may help avoid leaks and spills, mitigating environmental impacts, and government fines; not to mention worker and public safety. Environmental monitoring sensors (air, water, vibration, etc.) can also help keep the plant operating within required specifications.

In these types of applications, the more data that can be gathered at more points, along with autonomous or semiautonomous feedback loops allowing for control of actuators such as valves and vents, the better. Chirp-based sensors can be smaller, cheaper, more rugged, and demand less power than traditional IP-based devices, allowing them to be deployed in greater numbers and with less management and technical support. Redundancy through sheer number of sensors is a corollary benefit.

As with other applications, lower-level control loops might allow near-instantaneous response to local factors, such as actuating a valve to reduce the flow of ingredients to moderate a chemical reaction that is exceeding norms, with only exceptions sent "up the line" for additional monitoring and analysis. This would be much more efficient than requiring a round–trip data exchange for small adjustments.

A "carpet" of moisture sensors below key pipes and junctions might detect leaks at their earliest stages, long before they would be otherwise noticed. Footfall, wireless, or infrared sensors might help track personnel to ensure safe practices and operations, as well as to allow rapid response and rescue in case of an emergency.

Another advantage of deploying a wide variety of sensors in large numbers is the capability to analyze data flows from many devices. A neighborhood of interest might include liquid detectors, temperature monitors, and vibration sensors. Combinations of changing readings might pinpoint a future maintenance problem such as a worn bearing that is leaking slightly, a bit hotter than normal, and creating a small vibration in the equipment. Recognizing this state earlier allows work to be scheduled without excessive downtime, even when no individual sensor type showed out-of-tolerance readings on its own.

Better Perimeter Security and Surveillance

Facilities are only as secure as their most vulnerable access point. One way to increase security is again to increase the number of points being monitored. A field of footfall sensor "motes" is impractical if each must be burdened with the overhead of a full IP networking stack. It would be impractical to wire them all, and a combination of solar and battery power might be enough to drive simple chirp logic (see Figure 7-7).

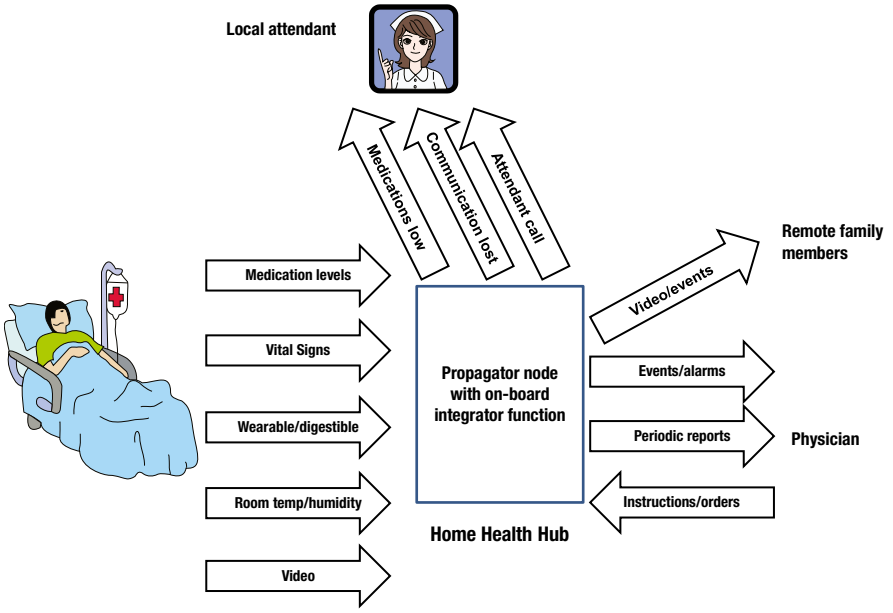

Figure 7-7. *In the emerging IoT architecture, networking demands on end device sensors are minimal, enabling many more of them (both in type and sheer number) to be "seeded" into the environment to improve installation security*

As "swarming" algorithms become more sophisticated, very simple airborne or ground drones might help patrol vulnerable areas of the perimeter. Autonomous local coordination of drone travel (along specific routes, in response to detected potential breaches, etc.) could be combined with the propagation of alert and alarm messages as dictated by the situation.

Similarly, video camera "swarms," operating in coordinated manner, could track/follow persons of interest as necessary. A camera swarm might collectively focus its attention to look for particular patterns or people. The cameras may be stationary, but through handoffs to others in the shared network, they still effectively provide ubiquitous surveillance coverage. In places where cameras are not deployed, mobile units with cameras will provide the needed continuity. Video surveillance will operate seamlessly as mobile and stationary cameras are employed as members of a collective intelligence community.

As discussed in Chapter 5, because integrator functions are IP-based, they may incorporate native IP data streams from more sophisticated cameras and sensors, combining these with small data flows aggregated by propagator nodes from chirp device streams. A single point of analysis and control thus manages both legacy and emerging devices. They might include biohazard, radiation, and other threat sensors.

Faster Factory Floors

With the increasing automation of the factory floor, the autonomous or semiautonomous lower–level control and feedback loops made possible through distributed intelligence within the Internet of Things may allow for higher production and better use of human resources. If integrator functions can handle lower-level adjustments and controls of operating machinery, human eyes and minds may be freed for longer-term analysis and optimization, based on exception and historical data collected at a higher level.

Machine autodiagnosis, parts supplies and quality, temperature and emission sensing all can be combined with video analysis of production lines and conveyers to maximize efficiency. As with some other applications, a key benefit of the emerging Internet of Things is the potential small size and cost of chirp–based end devices, allowing for much broader usage.

For example, industrial robots in factory automation are increasingly equipped with force and vision sensing for adaptive motion control. They are capable of stopping upon encountering an obstacle without damaging themselves or the obstacle. Factory floor environments that once needed to be rigidly structured (to ensure that “dumb” robots operated safely) are now more flexible in their designs with sensor-guided control.

Mobile robots, as in Automatic Guided Vehicles (AGVs) previously were required to move on preset paths, following lines inlaid or painted on the floor. More AGVs now use location markers on passageways and real=time data from other AGVs to collaboratively determine collision-free trajectories in factories with no markings on the floor. Sensor–driven path planning in real time in untaught factory floor environments is now practical; it was unthinkable only a decade ago. As more IoT sensor end devices become part of smart buildings, the character of industrial robots will continue to become more adaptive to changes in the environment. This will significantly reduce the cost of preplanned factory automation infrastructure.

True Home Automation

A new class of home and enterprise Access Points (APs) will be developed with the appropriate end device chirp transceiver built in, as shown in Figure 7-8. These will support both legacy Wi-Fi (IP) and chirp communications, and will typically include an IoT propagator node and (often) a publishing agent or an integrator function. These ambidextrous devices will appear as two logically distinct devices, even if they are using the same transceivers (for example, 2.4GHz unlicensed band radios). Thus each of these chirp-aware APs in the house, part of a mesh network, can provide access to all publishers and subscribers within the home legacy and Internet of Things communities. Each node and its agents can be regulated by a supervisory control system, which can move agents, remove them, update them, and so on.

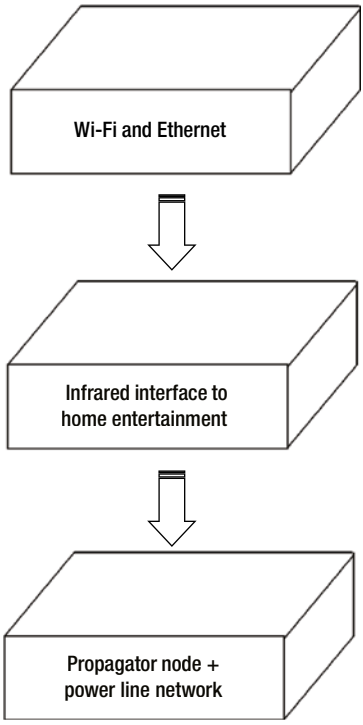

Figure 7-8. *In home and small office environments, attractive modular packaging will allow consumers to "build up" combinations of needed functions based on a "base" propagator node with IP functionality mated to additional transceivers to serve new and legacy devices. These "stacks" would often include a local integrator function serving as the home automation hub. This hub would be accessed wirelessly by an app on a smartphone, tablet, or PC*

In actual packaging, propagator nodes may be stackable, as shown previously, supporting multiple interfaces and their disparate tree-like networks (e.g., Wi-Fi and chirp infrared). Device-specific agents would reside on the propagator node networks, specific to one type of transceiver interface and sensor type. This would encompass a tight low-level interface with language and protocol *specific* to the device and its function. Thus a temperature sensor need "know" nothing more than how to transmit its temperature over an IR link. If no transmission is received, its agent "knows" that something is amiss, not the device. Further, to simplify matters, only the publishing agent needs to know how to parse and read meaning into the terse chirp stream, pruning, aggregating, and forwarding small data flows toward integrator functions as appropriate.

Local home automation monitoring and control will take the form of an on-board integrator function. This might be managed by a front panel or (more likely) a smartphone/tablet/PC app and would provide an extensible means of interacting with all the devices in the home, whether chirp-based or legacy IP. This could easily expand to include alarm and home entertainment functions. There will also exist "translator"

modules to permit non-Wi-Fi/non-chirp devices (such as TV remotes) to be incorporated as data sources or control points (much as a universal TV remote today). One of these modules will likely be a replacement for the home AC main plug. A simple chirp-based interface could provide information about energy usage and allow remote power-on/power-off.

Legacy translators such as these will be important for many years until Internet of Things–enabled devices replace current technology. Given that the life of some large appliances is 15 years (as opposed to 2–3 years for electronic technology), transition technologies are needed (see Figure 7-9).

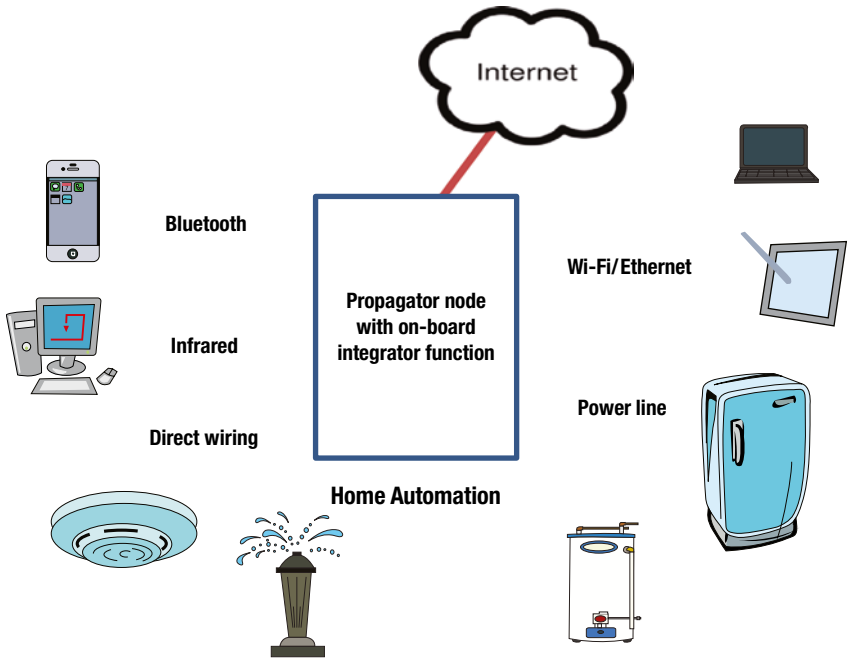

Figure 7-9. *The Internet of Things home hub brings together security, entertainment, activity monitoring, environmental comfort, energy usage and other interests within a single intelligent framework. Simple chirp protocols will suffice for the overwhelming majority of end devices in the home environment, but legacy IP and non-IP (e.g., TV remote) streams may also be supported*

This will bring about *true* home automation, with a variety of appliances, sensors, and other devices smoothly coordinated via a single point of intelligence, as discussed previously. The fabled Internet of Things toaster need not be burdened with a processor, memory, and an IP stack—a simple chirp interface will do. Physical interfaces may be varied, as noted in Chapter 2: Infrared, Bluetooth, Power Line, and other interfaces may all come together at the propagator node. Clusters of simple chirp devices, many not even yet imagined, will connect via these interfaces. Integrator functions will interpolate events

and data, detecting movement in the house and adjusting heating and cooling zones, for example, or turning off lights in unused rooms.

IoT end devices will also be able to communicate tersely and cogently with external integrator functions, reaching these via the IP interface of the home propagator node and the home's high-speed broadband Internet link. For example, the trashcan might chirp its level of "fullness," which the home network relays to the garbage collection company. Trucks and arrival times are accurately scheduled.

In another example, kitchen appliances might transmit status to a contracted appliance maintenance-and-repair depot, which would in turn schedule a repair visit to cover multiple devices (as shown in Figure 7-10). A list of those repairs is made available to services in that locality, if permitted. The home user can also specify the schedules that suit her, as opposed to the other way around. The combined information is a "request for quote" sent to multiple repair service subscribers. One may be selected. The repairman's visit is scheduled/confirmed by the home user. Alternatively, she may choose to have the parts mailed to her and do it herself. Another set of electrical appliance retailers will present their bids.

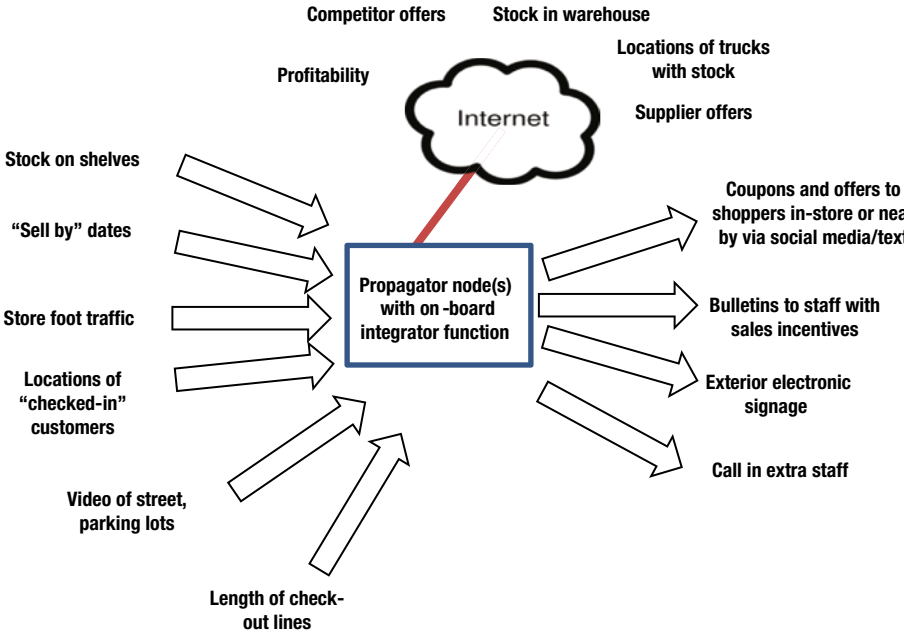

Figure 7-10. *Although many home automation IoT applications will be localized, the potential exists for end devices to update distant contracted organizations about their status and health. Maintenance reminders and service visits may be scheduled in response. Major appliance and equipment OEMs may offer these services to their own customers; others may offer a service supporting many different brands and types of equipment*

The IoT-enabled home may also coordinate with smart meters for gas and (especially) electric utilities to minimize usage during expensive time-of-day billing periods by throttling down or turning off some appliances and scheduling operations with an eye to maximum economies of cost and utility demand, as well as current and expected weather. Cooperative programs with utilities may offer additional price advantages if the utility is allowed to bias these decisions to match its generating capacities.

Wholesale and Retail: Beyond RFID

Dozens of Internet of Things applications have already been suggested and/or are being rolled-out now, both for online merchandisers and brick-and-mortar locations. To date, these applications have often been based on technologies such as Radio-Frequency Identification (RFID) along with IP-based readers and sensors. RFID chips are generally inert until powered-up by a nearby reader, so there may be many applications in which a simple chirp-based device will provide more functionality.

In the competitive world of retailing, well-stocked and properly "fronted" (products aligned to the shelf edge) displays are more enticing to shoppers. Low-cost, chirp-based sensors might be deployed along a shelf edge. Powered by overhead light, they might identify when product displays require attention. Or sensors in the floor or shoppers' carts might trigger coupon deals and other offerings based on location within the store.

Shoppers who have "checked-in" to a specific store on social media might welcome these offerings catered to their location and path through the store. As in an earlier example, foot or parking lot traffic might be used to gauge the type, number, and attractiveness of offers presented. After all, there's no need to offer major discounts at this moment if it is known that the parking lot is currently filling up. This is another area in which a variety of data streams, including factors such as current and expected weather, might be brought together by an integrator function for autonomous or semiautonomous action (see Figure 7-11).

Figure 7-11. *Combining data from many sources, a retailer might make available offers and coupons specifically related to stock on hand, store traffic, weather conditions, and many other factors*

In the back-end and wholesale environments, inexpensive and/or disposable (even biodegradable) shock and temperature sensors might monitor shipping conditions in transit and help track stock in the warehouse. Efficiently rotating stock based on first-in-first-out or expiration dates might be aided by more data from individual containers or cases. (And then there is the oft-told IoT tale of the refrigerator noticing that one is low on milk, seeing that one's location is near the store, and generating a text reminder to pick up a half-gallon on the way home.)

A Broader "Net" in Natural Sciences

As with the security examples noted previously, a larger number and great geographic spread of sensors is important for improving the usefulness of natural science observations. Strain and crack gauges spread over very large areas might allow better monitoring of geological conditions, perhaps leading to prediction capabilities for natural events such as earthquakes and volcanic eruptions. Detection of snow levels, CO^2 emissions (as from a wildfire), air and water pollution, and many other parameters may be much easier with cheaper, lighter, and more-easily-managed end devices. Small, cheap, solar-powered IoT chirp devices might allow scientists to cast broader nets for data than before.

Speaking of nets, wild and farmed animal populations (fish, cattle, birds, and so on) might also be monitored with implantable and/or digestible chirp devices. As with many other applications, the promiscuous forwarding built into the basic chirp architecture and the propagator node might allow recruitment of a very wide-ranging network of nonaligned propagator nodes as a "free" propagation medium for this data. This might allow a collection of data from far afield without the cost of building out a propagator node network specifically for one application. A sparse mesh network of "volunteer"

IP–capable propagator nodes (which are more power-hungry) will provide the networking path, but the underlying lower level of end devices will remain light and terse in their communication protocols.

Living Applications

Emerging IoT applications will be “alive" in the sense of being adaptive, self-healing, self-forming, and largely collaborative. The architectural foundations will together enable unprecedented innovation in the development and deployment of new applications in the Internet of Things:

- Minimal networking requirements for end devices
- Provision for local autonomy of action
- Distributed intelligence to offload both end devices and integrator functions
- A flexible publish/subscribe model creating neighborhoods of information