CHAPTER 3

On the Edge

Although the backbone architectures of networks garner the most attention, the *actual drivers* of network deployments are the devices at the edge. If that statement seems odd, consider desktop architectures such as twisted-pair Ethernet and the near-ubiquitous Wi-Fi, neither of which made great strides until the technologies were embedded in silicon and offered nearly free on every computer and smartphone sold.

This "edge effect" is amplified by the sheer numbers. There are orders of magnitude more end points than networking devices in most networks. From a cost, deployment, and product life cycle standpoint, it's always been true—until the end points are network-ready, a network architecture is only theory.

These factors apply even more directly to the Internet of Things. There will be literally billions of networked end points, eventually dwarfing the world population traditional Internet to date, as shown in Figure 3-1. But unlike any other network deployment, the IoT end points should be extremely *inexpensive*, *autonomous*, and mostly *untouched* by human command and control.

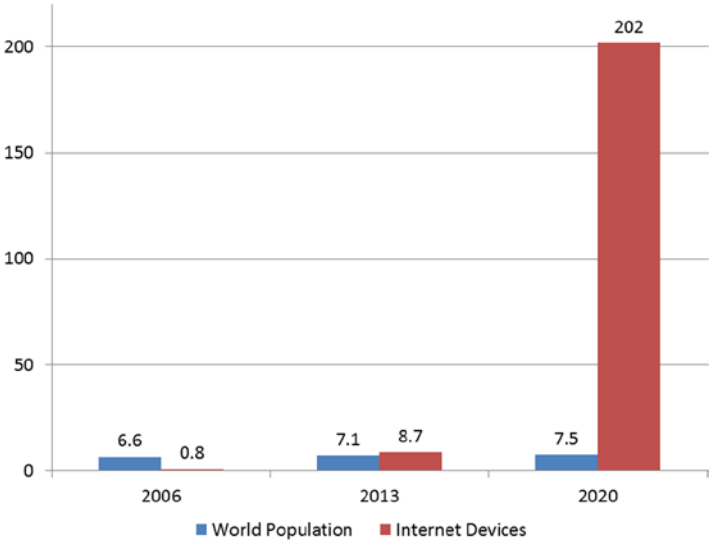

Figure 3-1. *The number of Internet-connected devices exceeded the number of humans in the world around 2009, but the Internet of Things will cause a further exponential increase in the number of devices. Sources: Cisco Systems, International Data Corporation, Population Reference Bureau, U.S. Census World Population Clock, United Nations Department of Economic and Social Affairs*

A World of Different Devices

For many people, the IoT conjures up visions of smartphones, laptops, and similar intelligent, human-oriented devices. But in fact, the statistical bulk of the IoT will consist of relatively simple devices such as pollution sensors, diesel generators, air conditioning systems, building lighting components, and so on. Familiar computing devices, designed for human "high touch," will mostly stay on the traditional Internet, but the Internet of Things will reach far out to the edges of the network to devices that have *never been connected* in the past.

Because these classes of devices have never been connected, there are limited technical models upon which to draw. The connectivity challenges are substantial: there may be limited bandwidth, lossy connections, intermittent links, and power-off periods. In addition, end devices may be mobile or stationary, appearing and disappearing from the network at any time. But the greatest challenges are in the manufacturing, deployment, and management of this vast population of end devices.

As seen in Figure 3-2, IoT devices could be virtually anything that runs on any sort of electricity (or provides or has access to energy such as heat, motion, or light that may be converted to electricity for signaling). IoT-enabled devices may be built in millions of factories and shops across the globe and purchased in millions of different venues. There is no existing (or imagined) technology or business process that could possibly manage this sort of far-flung, uncoordinated global supply chain.

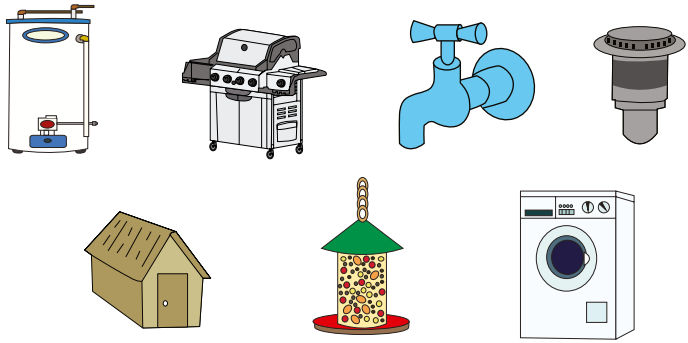

Figure 3-2. *The Internet of Things will include a dizzying variety of end devices, both traditional and obscure*

Intended to be Untended: Some Examples of IoT Systems

Similar to the fish shown in Figure 3-3, IoT devices must act autonomously and independently. It is only from an *external* viewpoint that the devices will appear coordinated. When powered up or otherwise triggered, an IoT device will simply bleat out its data and/or listen for its data. But that sending and receiving will have no bearing on most IoT devices' prime functions.

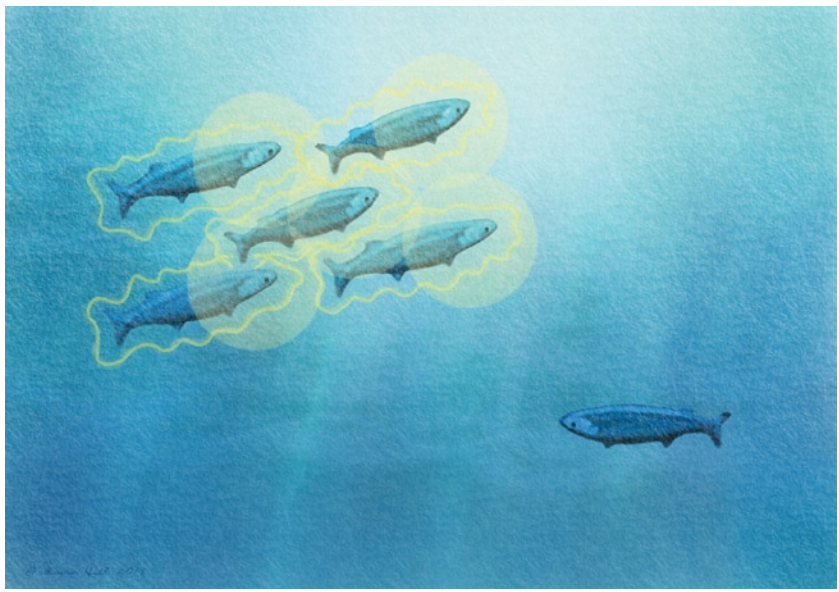

Figure 3-3. *Each fish in a school participates in group movements and behaviors when in contact with others, yet can also exist independently*

For example, streetlights will continue turning on and off with the setting and rising of the sun, regardless of whether their status messages are being received somewhere else. Electrical generators will continue cranking out kilowatts without "knowing" whether their terse broadcast reports on lubricant viscosity are being studied. Because networking on the Internet of Things frontier is so lossy, intermittent, and uncertain, it is important not to hamstring the devices with an end-to-end data assurance requirement.

This grows out of the recognition that the Internet of Things will only *indirectly* interact with humans. The vast majority of the communications will be machine-to-machine: generally end devices and integrator functions exchanging information through lossy and intermittent links, typically through relays (propagator nodes). Humans will interact with the integrator functions, retrieving reports or setting parameters that bias the operation of the remote end devices. Interactions that are real time, mission-critical, or human-oriented will mainly continue to use the traditional Internet and other existing "reliable" networking protocols.

Because the vast majority of IoT end devices will be engineered to operate independently of network connectivity, *individual* data messages are completely *uncritical*, as noted earlier. This allows for end devices that cease sending or receiving when powered off, wireless links that are extremely weak or intermittent, solar-powered end devices and other network elements that literally "go dark," and other realities of networking at the edge.

Temporary and Ad Hoc Devices

In fact, an entire class of IoT end devices may exist only transiently as hastily formed networks. Smart disposable "motes" may be deployed for specific purposes, perhaps sending data only for as long as their limited batteries last. A sensor network of this type might be measuring the pressure change of an intruder's footfall, for example, in a temporary protective alarm ring around a facility. The cost, size, and power savings that come from avoiding the overhead of traditional protocols are substantial and will drive these devices and networks to simpler chirp architectures.

Addressing an Uncertain Frontier

One of the major issues to be addressed when contemplating the Internet of Things is how messages to and from end devices may be addressed. This issue was discussed briefly in the preceding chapter and is covered in more detail in Chapter 6, where the three key IoT addressing concepts are explored: self-classification of end device type with external markers, non-guarantee of absolute end device address uniqueness, and end device address derivation from the environment. These basic concepts will allow the uncoordinated "crowd" of end devices to be scaled into a global Internet of Things.

Reliability Through Numbers

Although much of the Internet of Things will be generally predicated on the fact that individual end device connections will be lossy, intermittent, and unreliable, an interesting phenomenon will be the build-up of *reliable* information from a very large number of *individually unreliable* sources.

As an example, consider strain gauge sensors on a highway bridge (see Figure 3-4). It might be desirable to distribute hundreds or thousands of these at many locations on the bridge. Data from the gauges might be collected wirelessly and propagated to an integrator function monitoring the condition of the bridge. But it would be nearly impossible to wire an external power source to each of these sensors. In this case, it might be more practical to make a significant percentage of these devices "solar-powered," energized by either the sun or existing streetlights on the bridge.

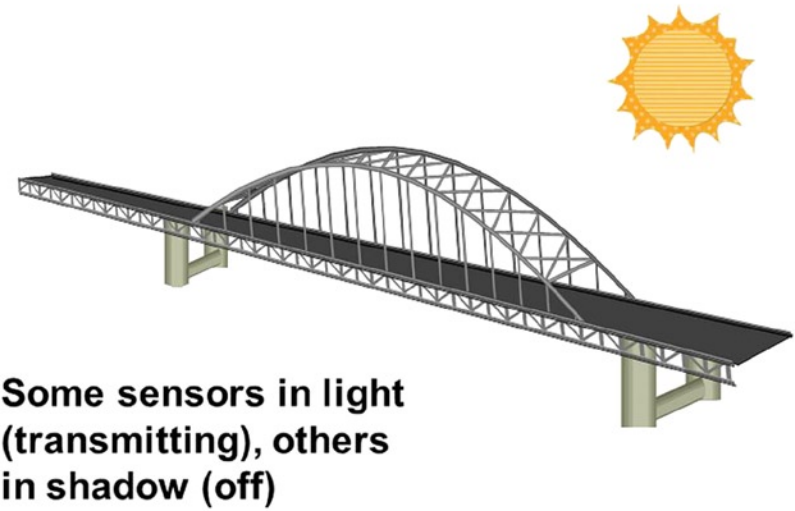

Figure 3-4. *Thousands of individually unreliable solar–powered strain gauges on a bridge effectively create a single reliable integrated data source*

With the movement of the sun through the sky, different sensors might be illuminated at different times of the day. Some cease to broadcast when in shadow, whereas others begin broadcasting their status when the moving sun casts light on them. Still other sensors' broadcasts might be occasionally interrupted by passing vehicles. However, there will always be hundreds of sensors broadcasting, although no single sensor would be guaranteed to be active at any particular moment.

This is over-provisioning through sheer numbers of end devices, creating a *net* consistency and reliability through integration that would be impractical or prohibitively expensive to provide through highly reliable individual sensors. Similarly, integrator functions might analyze and interpolate information from a variety of unrelated devices to detect events or trends such as a power outage.

Meaning from Many

At this point, it might be worthwhile to briefly consider some examples of how information flowing to and from simple end devices is transported and becomes meaningful in the machine-to-machine world of the Internet of Things (IoT applications will be more thoroughly explored in Chapter 7).

The true power and utility of the IoT comes when vast quantities of data from end devices in the form of short chirps are consolidated, analyzed, and integrated to create "small data" streams of rich information. The resulting small data flows percolate "up" and are converted into big data content. This process will be a key driver for the deployment of the IoT (and was the inspiration for this book). End device chirps that are briefly stored and analyzed at integrator functions will allow the development of perspective and some learning from experience.

End Devices in Dedicated Networks

In the example of streetlights mentioned previously, the on-or-off state and/or OK/Fault status being repeatedly transmitted by small modules within each individual streetlight would be collected via one or more propagator nodes. This communication, in the form of chirps, might be wireless or via very low speed data modulated over electrical power cables. Propagator nodes at central points in the street grid receiving these chirps might ignore repeated transmissions (or reduce the number), bundling the data for forwarding to an integrator function. The propagator node may add contextual information *not available* from the end devices, such as time of day, weather, location, and so on.

The combined data would then typically be encapsulated in an IP packet and forwarded by the propagator node toward an integrator function, as described in Chapter 4. This might be via the traditional Internet, a private wide area network (WAN), or some combination.

The integrator function (typically software operating on a general-purpose processer; see Chapter 5) would be receiving chirp data from streetlights across the city. From these "small data" feeds, a big data perspective could be developed based on analysis and integration over time or as a snapshot of status. Individual streetlight failures or faults beyond a previously defined threshold might cause the integrator function to generate an alarm and report for a human operator's action or might even be integrated with scheduling software to add faulty lights to a repair worker's schedule automatically. In this way, data from relatively "dumb" devices becomes a powerful tool for system management.

Expanding to the World

In the preceding example, the network was fairly sequestered. In fact, this might be desirable for security or other proprietary reasons, and the chirp protocol permits this (see Chapter 6). But tremendous potential uses for chirp data from simple end devices arise in broader settings.

The data from a significant portion of end devices will simply be transmitted with generic public markers (see Chapter 2) that allow it to be interpreted by any integrator function with an understanding of that end device classification and type (moisture sensor versus temperature gauge versus strain gauge, and so on).

One of the key opportunities for new meaning extracted from "small data" is that integrator functions may process a wide range of nominally *unrelated* data sources, as opposed to the fixed end-to-end IP conversations typical of most of the traditional Internet. Integrator functions may perform something of a seemingly variable "random walk"—collecting data in a contingent fashion from a wide variety of end devices *anywhere in the world* based on sampling trends and events. The externally self-classified format of chirps allows all of the elements of the IoT to recognize potentially interesting, but previously unknown, data sources.

There could be huge benefit, for example, from integrating data from thousands of wind speed and direction sensors, barometric pressure gauges, and temperature readings, along with public domain weather reports, to create a highly localized footprint of potential tornado formation. These end-device data sources might not be centrally owned and controlled, but an integrator function could easily seek them out and add them to an ever-growing set of inputs.

The capability to handle both proprietary and generic uses of data creates the need for an IoT architecture and chirp protocol that can be public or private (see Chapter 2). Some data streams from end devices will actually be used by multiple unrelated integrator functions, a factor that propagator nodes must take into account when bundling and forwarding end device chirps (see Chapter 4).

Converting States to Chirps

For a large majority of devices on the Internet of Things, only the bare minimum amount of data will be contributed to these higher-level analyses. As noted previously, a simple On/Off state or an "OK/Fault" condition might be the only useful information that the end device may present. Or a simple voltage differential or current reading will be of interest for a moisture sensor, temperature gauge, or similar device.

For simpler devices such as these, the analog-to-digital interface may likewise be very simple. Ideally, integrated silicon chips will be developed, which simply detect the presence of voltage (or a similar condition) and directly create chirps through very simple logic. This obviates the need for processing, memory, or other computing functions within the majority of end devices.

More importantly, this means that there is no *significant redesign* needed for millions of existing un-networked devices, appliances, and machines. Instead, a simple connection to an existing point in end device wiring or circuitry will provide the information needed to create chirps.

Again, this is a departure from the thinking behind much of the Internet and traditional networks, in which the end devices must have all the functionality needed to create digital data (typically in frames or packets). Instead, much of the IoT will function more along the lines of telemetry, in which states and conditions are coded as simply as possible and then broadcast, as shown in Figure 3-5.

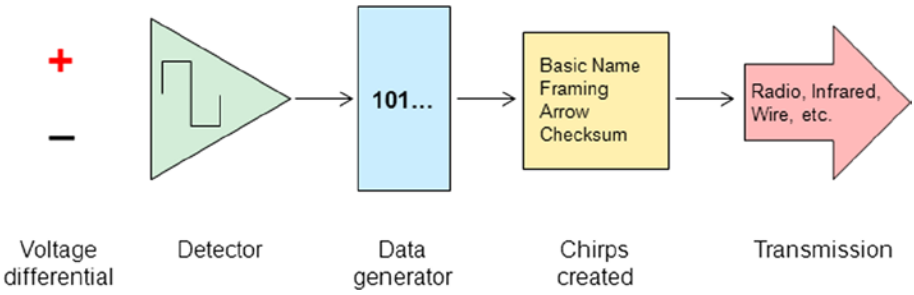

Figure 3-5. *In millions of simple end devices, basic physical states will be converted to chirp payloads. An address, "arrow" of transmission, and checksum are added to this payload to form the complete chirp packet*

It is likely (and perhaps desirable) that some number of standardized chirp formats will be created to handle specific very common states and conditions, such as On/Off, Green/Yellow/Red status states, and so on. A list of suggested potential chirp formats is listed in Appendix A.

"Setting" End Devices

Many of the end device examples explored thus far have been sensors and other devices that will simply broadcast states and conditions and *listen for nothing*. Although this situation may be true for the majority of devices in the Internet of Things, additional billions of devices will be receive-only or bidirectional (sending *and* receiving chirps).

The messages intended for these devices will typically be generated by integrator functions and then propagated "down" (or away from the integrator function) toward end devices by propagator nodes. (Propagator nodes may be directly IP-connected to the general-purpose computer hardware hosting the integrator function, but will more often be reached via the traditional Internet.)

This "direction" of travel is determined by the "arrow" of transmission within the chirp message markers encapsulated in IP, as discussed earlier. Integrator functions will generate these chirps based on human programming, preset alarm conditions, or through routines generated by interactions between integrator functions. At the final destination, the "last" propagator node strips the IP encapsulation and generates native chirps bound for the end device.

As before, many of the targeted end device appliances and actuators will be very simple and thus, the chirps will have very simple payloads. In this way, these end device–bound chirps may resemble the "SetRequest" of Simple Network Management Protocol (SNMP). A key difference with SNMP, however, is that the IoT end device need not acknowledge taking an action directed by the chirp, nor even the reception of an individual chirp. This eliminates a tremendous amount of protocol overhead throughout the network.

As with the end device chirps propagating "up" through the network, these chirps moving "down" will simply be repeated. Because each individual chirp is so tiny, and repeated transmissions may be squelched at the propagator node without clogging

wide area connections, the cost of over-provisioning through repetition is small. In applications for which it may be important that the integrator function have some acknowledgment that a chirp was indeed acted upon, bidirectional (send and receive) end devices can be deployed (see Figure 3-6).

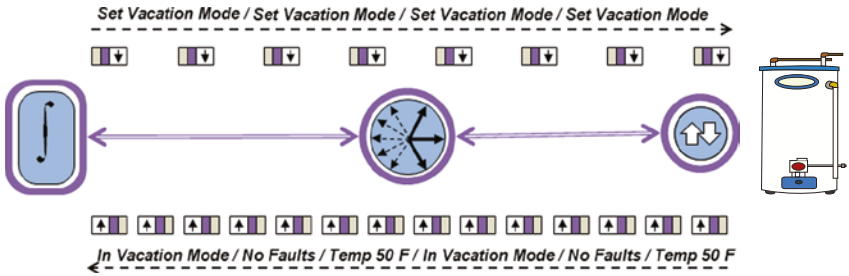

Figure 3-6. *Receive-only and bidirectional end devices receive transmissions from integrator functions via propagator nodes, which handle broadcast bundling and unpacking en route*

Where necessary, these bidirectional end devices may receive chirps (setting a valve position in a process control environment, for example) and also be continually broadcasting chirps that indicate the valve position. In this way, the integrator function need only repeat the command to move the valve—until eventually chirps are received, indicating that the valve is now in the desired position.

The unreliability of any individual transmission suggests that IoT chirp protocols may not be the best choice for real-time and especially critical or dangerous functions. Traditional Internet and other networking protocols will continue to work well in those situations, of course. But for billions and billions of end devices, chirp protocols will provide "good enough" functionality at a vastly reduced cost of bandwidth, processing, memory, and other factors.

The end result may be "neighborhoods" of interest built up by integrator functions consisting of data streams from a combination of IoT chirp protocol end devices (converted to IPv6 by propagators) and more sophisticated end devices communicating via native IPv6. Information extraction and analysis takes place within the integrator function, as described in Chapter 5.

Cornucopia of Connections

Many Internet of Things discussions assume wireless connectivity over traditional networking schemes such as Wi-Fi, Bluetooth, and so on. And it is likely that this will be true, particularly between *propagator nodes* or between *propagator nodes* and *integrator functions*. But connections to end devices may be widely varied—some quite sophisticated and others extremely prosaic. A more detailed look at wireless connectivity will be explored in the extensive sidebar that concludes this chapter titled "Wire-Less vs. Wireless."

Within a residence or enterprise, many end device connections to a collocated propagator node may be via copper media. A few of these may be dedicated wiring, but existing copper wiring infrastructures such as telephone, data, and especially AC power line wiring will often be much more cost-effective. Because a very large number of end devices will be plugged in to AC mains (as will the propagator nodes), there will be a natural opportunity to exploit this in many cases. The amount of IoT data will be low, as noted earlier, so existing AC power line chips and protocols (IEEE 1901 for example) provide more than enough capacity for Internet of Things communications.

With the low data rates and duty cycles of most IoT end devices, other potential existing technologies may also be considered (see Figure 3-7 "Examples of IoT End Devices"). Open-space optical networking techniques such as infrared (IR) may be useful in the home environment, for example. Although IR has mainly been used for remote control of home entertainment and similar devices, networking protocols such as the open-source Linux Infrared Remote Control (LIRC) may present an interesting low-cost alternative for IoT chirp networking (see the "Wire-Less vs. Wireless" sidebar).

Application/ Device	Mode[1]	Chirp Size (bytes)	Frequency/ Minute	Repetition %	Effective Data Rate (kbps)	Typical Transmission[2]
Environmental Sensors (Temperature, vibration, strain, moisture, etc.)	S	4.5	1	90%	<1	Z,B,I,W
Lighting Monitoring and Control	B	6.5	1	90%	<1	P,W
Household Appliances	S	4.5	<1	99%	<1	I,P,Z,B
Inventory / Supply Chain	S	4.5	1	95%	<5	R,Z,W
Video Surveillance (Stand By)[3]	B	4.5	60	99%	5	W,Z,C
"Smart" Signs	R	7.5	1800	90%	22	W,Z,P,I
Home Entertainment Control	B	7.5	600	99%	50	I,W
Process Control (Valves, flow rates, etc.)	B	7.5	60	75%	150	Z,W,C
Industrial Machine Control / Diagnosis	B	7.5	6000	25%	340	W,Z,I,C

[1] S = Send-only device; R = Receive-only device; B = Bidirectional
[2] B = Bluetooth; C = Copper (wired); I = Infrared; P = Power line; R = RFID; W = WiFi; Z = ZigBee. Listed in approximate order of suitability/preference.
[3] In stand by mode, local integrator function is processing video surveillance stream, so only "OK" or "fault status is transmitted by IoT interface. When movement of a particular type is detected locally, device would shift to IP mode in order to transit large frames of full video. Local device may also be triggered to full video mode by command from remote integrator function.

Figure 3-7. *A small sampling of IoT End Device types shows tremendous variety in communications types and effective data rates. Note that these are well below typical Internet data rates*

No matter which connection techniques are used, chirp messages to and from end devices need only reach a propagator node, where they will be bundled, pruned, and retransmitted as needed to move the traffic through the traditional Internet for connections to integrator functions.

Chirp on a Chip

As noted previously, many Internet of Things end devices will have relatively simple information to share or receive, such as simple states or conditions that may be communicated by the presence or absence of voltage or some other simple "signal." When this is combined with the very simple structure of the IoT chirp packet, the potential exists for extremely cost-effective, mass-produced, integrated silicon chips. These could provide state detection, chirp formation and transmission technology (wired, IR, or radio frequency [RF]) in a single small, inexpensive, and low-power package. (Receive-only and bidirectional integrated devices will also exist with slightly different requirements.)

Development and widespread distribution of these "chirp on a hip" components will be critical to the expansion of the Internet of Things because they will make possible connections to millions of different types of relatively inexpensive devices.

"Chirp chips" might be offered in a variety of tiers, defined by the integration of different functions. Global positioning system (GPS) receivers, electromagnetic position indicators, accelerometers, and other indicators of environmental condition might be interesting potential add-ons, as might radio-frequency identification (RFID), as discussed in the section following). But it's likely that a majority of IoT chirp chips will be relatively simple single-function modules optimized for lowest cost, smallest size, and minimal power consumption. Development and integration of chirp chips is discussed in more detail in Chapter 8.

Aftermarket Options

Integrated chirp chips can become available quickly for new purchases of IoT–ready OEM equipment. But billions of devices already exist that users will desire to have connected to the Internet of Things. For these devices, add-on and aftermarket alternatives need to be developed.

For many simple needs, such as power On/Off or Red/Yellow/Green status, a simple module might plug in between the end device and the AC mains. These might communicate via power line or wireless technologies and would require no software or configuration of the end devices. It can be imagined that these might be built into devices such as power bars and surge protectors. (In this case, the device might also function as a propagator node for all the attached end devices.)

Additional packaging options for aftermarket IoT connections in some applications could include small stand-alone devices based on Universal Serial Bus (USB) and other standardized interfaces, especially those that provide power as part of the interface. Because of the simplicity of the chirp networking protocol, add-on aftermarket devices may be very compact and draw little power. IoT devices will not require the high speeds possible over these interfaces (and their associated costs), but these standards may still be useful due to their wide availability in the market.

RFID Integration in the Internet of Things

For some market participants, interest in the Internet of Things has been *primarily* around the spread of RFID capabilities. RFID is based on a small physical device (a "tag" or "label") that broadcasts stored data, such as a serial number and other information. These devices can be self-powered, but more often are temporarily energized by RF fields generated by the receiving device (the "reader" or "interrogator"). RFID is widely used in inventory tracking and asset management, and its applications are constantly expanding.

Some have viewed each individual tag as an end device in the Internet of Things, and there is certainly some similarity in basic capability between a tag and a simple chirp-enabled end device. But it seems more likely that RFID *readers* will actually function as IoT end devices, perhaps combined with a propagator node.

Because typical RFID tags communicate only identification parameters and have no defined interfaces to other signals (such as voltage presence or differential) within the device to which they are attached, they are significantly more limited than a full chirp end device. But interesting potential exists for combinations of RFID information and chirp data to be received by a propagator node, which could bind the information together before forwarding directly to the integrator function (see Figure 3-8).

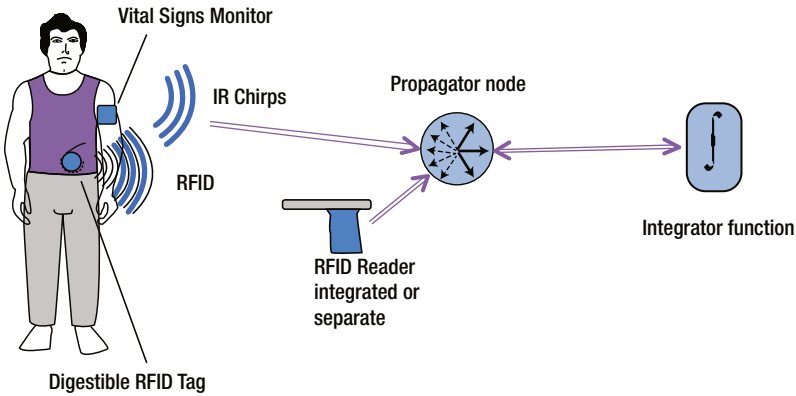

Figure 3-8. *Some applications may combine both RFID and chirp signals to provide both location and state inputs for analysis by an integrator function*

End Devices with Higher Demands

As noted in Chapter 2, relatively simple end devices will predominate numerically within the Internet of Things. But there will still be billions of devices with more demanding communication needs, such as video surveillance systems, teller machines, and telepresence information kiosks, among many others. Many of these have real-time data requirements, high bandwidth needs, and/or human interfaces that make data reliability and bandwidth critical.

For the most part, these devices will therefore remain directly connected to the existing Internet via traditional networking protocols such as TCP/IP. These high-data-need devices will certainly often share the Internet backbone with traffic from propagator nodes comprised of consolidations of chirps to-and from IoT end devices.

But interesting opportunities may exist for combinations of chirp and traditional protocols within a single device. In some cases, simpler status or environmental conditions that are less time-critical might be sent and received via chirps while a high-demand end device is in "StandBy" mode. Then when the device is fully activated (for a human interaction, perhaps), a traditional Internet connection is established for the duration of the high-data-need transaction.

Another potentially interesting application might be to make use of the IoT chirp interface as a back-channel or chording input to the traditional high–bandwidth Internet connection, perhaps in a different frequency or physical domain (see Sidebar "Wire-less vs. Wireless"). Chirp–enabled end devices will likely constitute the vast numerical majority of the Internet of Things, but billions of higher–demand IoT end devices will still comfortably coexist.

The Big Idea: "Small" Data

This chapter has explored the variety of Internet of Things devices in some detail. The only common denominator for IoT-enabled devices may be data—*just a little* for each: tiny squirts and squibs of data—a few bytes reporting moisture content of soil or wind direction or a short instruction to set a valve to a new position. As introduced in Chapter 2 and more fully explained in Chapter 6, these tiny information exchanges are in the form of chirps: simply structured self-classified data packets with minimal overhead.

Individually not impressive or meaningful, these end device chirp data streams become powerful tools when combined and analyzed within integrator functions (see Chapter 5). But first, these myriad chirps must be transported across the Internet of Things networking frontier and (usually) through the traditional Internet. That job falls to the propagator nodes, which will be explored in the next chapter.

WIRE-LESS VS. WIRELESS

Most people picture wireless connectivity when thinking of ways to connect end devices in the IoT. And when thinking of wireless, most consider traditional existing protocols such as Bluetooth, ZigBee, Wi-Fi, and cellular/4G/LTE. Many IoT end devices may indeed be connected using one or more of these protocols (see Figure 3-9), but not only to *these* wireless protocols.

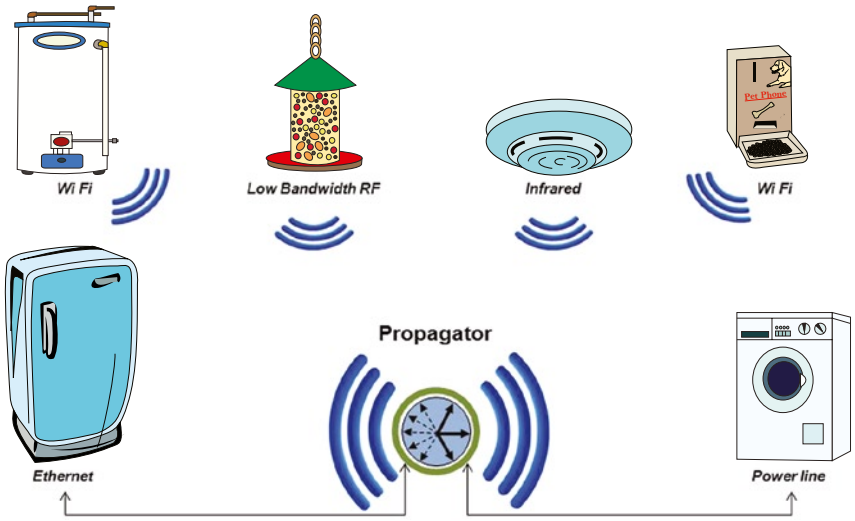

Figure 3-9. *IoT end devices will communicate over various means: optical IR, wireless, power line. Many propagator nodes will be equipped with multiple physical wired and wireless interfaces*

Again, because the total data transmitted to or received from an individual IoT end device is exceedingly small in the vast majority of cases, all those traditional protocols are by definition overkill. Many sensor-type devices will generate only a few bytes of data per hour, for example, and still effectively less after repeated identical transmissions are squelched. Even the lowest ZigBee data rates, for example, are on the order of 20 kilobits per second. This is multiple orders of magnitude greater than the data rate that will be needed for most IoT end devices, although there will be exceptions for other classes of IoT end devices.

Due to these low data rates and duty cycles, the protocol stacks and wireless RF sophistication of standard chips will not be necessary for the typical IoT end device link. Much simpler (read: cheaper) solutions based on simple modulation schemes within existing unlicensed frequencies can therefore be considered.

As noted earlier, these alternatives might include power line, television white spaces frequencies, and open space optical links (IR or visible). The first is obviously potentially attractive for any end device that plugs into AC mains, as long as a propagator node is also plugged into the same building or household somewhere. IR is familiar to most of us in the form of TV and other entertainment system remotes. Wire-less need not be traditional wireless.

Navigating an Already Wireless World

There may be a number of low-cost, unsophisticated wireless modulation schemes developed for the Internet of Things (some possible approaches are suggested in Chapter 6). With such small data rates and duty cycles, very low baud rates are needed, so signaling techniques can be quite simple. It likely goes without saying that virtually all IoT networking must take place in unlicensed frequencies. (It is somewhat contrary to the low cost and simple protocol characteristics of the IoT end device to consider licensed RF bands, although there is nothing in the chirp structure that would preclude this.)

But these new potential wireless IoT solutions will not be deployed in virgin territory—traditional wireless protocols such as Wi-Fi, Bluetooth, and many others are already widely (and unpredictably) deployed using unlicensed RF bands.

Coexistence by Camouflage

Because it will be necessary for IoT wireless signals to coexist with existing protocols such as Wi-Fi, this would seem to demand a traditional wireless protocol stack in every IoT end device and propagator node. Yet this would certainly destroy the low-cost model needed for widespread Internet of Things acceptance and proliferation.

The solution to this problem is based on “hiding in plain sight” within the traditional unlicensed RF environments based on an understanding of their operation. The key will be to exploit time and frequency domain differences. Because IoT chirps are so short and individually uncritical, they may be squeezed into “spaces” that naturally occur when more-sophisticated protocols are in operation (see Figure 3-10).

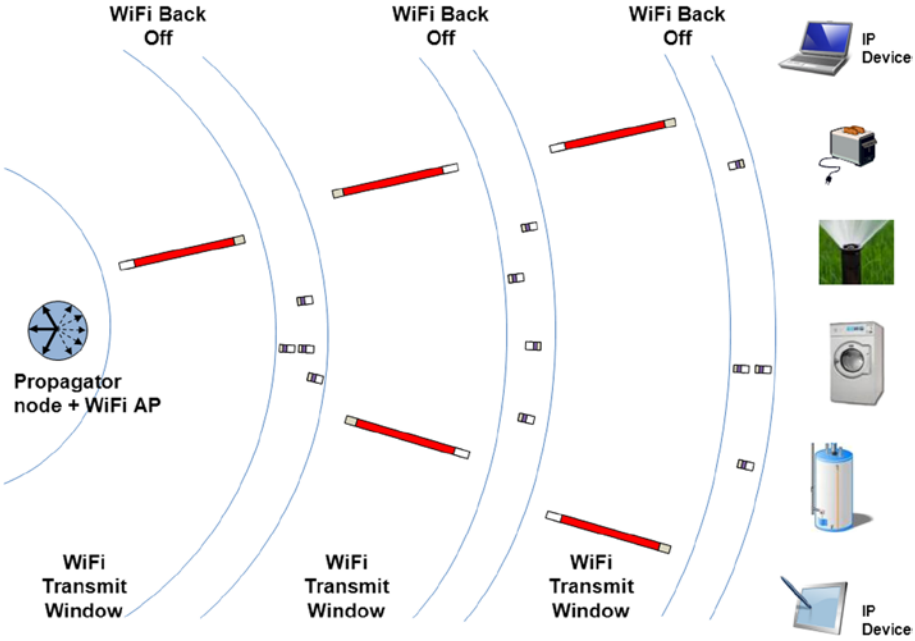

Figure 3-10. *IoT chirps "squeezed" between Wi-Fi send-receive cycles, as fully explained in Chapter 6*

Collisions? Who Cares?

Selecting Wi-Fi as an example, IoT devices may easily operate within the "quiet time" back-offs inherent in Carrier Sense Multiple Access with Collision Detection (CSMA/CD). IoT end devices simply broadcast or listen for their chirps. Because the chirps are very short, there is little statistical likelihood of one colliding with a Wi-Fi packet, even within a fairly busy Wi-Fi network. And even if one collision does occur, that chirp is individually uncritical, and another will likely get through relatively soon. Randomized timing between chirps will also help avoid any "deadly embrace" problems with devices communicating via traditional wireless protocols (see Chapter 2).

The effect on the Wi-Fi network is also minuscule, again because of the very small chirps and low duty cycle of the typical Internet of Things device. So there is no need to burden IoT end devices or the chirp protocol with any collision detection, avoidance, or recovery capabilities. Propagator nodes, on the other hand, may be the appropriate places to incorporate either a full traditional wireless stack or a "listen for a pause" capability to hold transmissions and avoid unnecessary collisions (see Chapter 4). By bundling and pruning IoT chirp broadcasts, the propagator nodes can be "good citizens" within traditional wireless environments.

"Chording" and Baud Rate

When fully considering all the "wire-less" options such as power line and optical signaling, an interesting set of opportunities is presented. Sending or receiving data in multiple domains simultaneously may increase the baud rate (or decrease the potential for collisions and interference). An individual device might send simple chirps via RF and IR to increase the amount of information transferred while remaining "below the radar" in terms of traditional wireless networking in the same environment.

Like a musical chord, sending multiple pieces of information simultaneously in two frequency domains may offer a potential for very rich communications using the very simple chirp protocol. As more fully described in Chapter 6, this allows a much higher baud (signaling or information) rate than is possible within the bandwidth of a single medium.