

WIRELESS SENSORS STREAMLINE DATA DISTRIBUTION

With applications ranging from home automation and remote meter reading to industrial sensor networks, low-power, low-cost wireless devices are set to reshape the control and data-distribution landscape.

Self-organizing, wireless sensor networks have immediate utility in a variety of industrial, medical, consumer and military applications. As a result, a new IEEE wireless standard, IEEE 802.15.4 (Callaway), has been proposed that aims to derive the optimum power, transmission distance and data rate requirements for devices that would best suit this space. With a lower-power profile and lower data rate than Bluetooth, this technology presents a number of interesting hurdles to its implementation, such as battery use and size of device. The networking protocol presents another challenge, encompassing latency, node acquisition time, route discovery and message confirmation. While a full understanding of the physical- and data-link-layer pa-

BY TOD RIEDEL, SOKWOO RHEE AND SHENG LIU




ILLUSTRATION AND DESIGN BY MIRIA RAMM-STERN

parameters of this technology is essential for its proper use, the potential impact depends also on the real-world implementation. Unlike the myriad profiles of the now emerging Bluetooth devices, there are three general application classes that have been derived for these devices: periodical sampling, event driven and “store and forward.” There are also three basic topologies in which this technology can be deployed: star, mesh and star-mesh hybrid. Each of these implementation modes has relative advantages and disadvantages that must be properly understood in order to match the application requirements to the appropriate wireless sensor network.

■ Why now?

That low-power, self-organizing networks are becoming a reality is largely a result of significant advances in microelectromechanical systems, low-power radio and digital circuit design. Wireless sensor networks are now capable of operating with submilliampere power consumption, allowing a 3-volt dc coin battery to power the sensor node for periods of up to five years and beyond, depending on the sampling rate. Such sensor nodes, when integrated with a coin battery, are portable, unobtrusive

and easily designed into small devices. Low-power, low-data-rate applications include aiding digital precision instruments on the factory floor, collecting water and gas meter readings, monitoring shipments through the supply chain

Unlike the myriad profiles of Bluetooth devices, low-power sensor nets have three basic classes.

and reporting on the vital signs of individual wearers. All of these applications share three common requirements: small form factor, long battery life and a robust, efficient network protocol. Proper implementation starts, however, with the basic choice of topology.

■ Topology options

The basic star topology is a single-hop system in which all wireless-sensor nodes communicate bidirectionally with a base station or gateway (Fig. 1a). The base station can be a PC, PDA, dedicated building-control device, embedded Web server or other gateway to a higher-data-rate device. The nodes are identical and the base station serves both to communicate

data and commands among endpoints, and to transfer data to a higher-level system like the Internet. The nodes do not pass data or commands to each other; they use the base station as a coordination point. Among wireless-sensor networking topologies, the star system is the lowest in overall power consumption but is limited by the transmission distance between each node and the base station (Fig. 2). That distance is typically 10 to 30 meters in the ISM band.

Mesh topologies are multihopping systems in which all wireless sensor nodes are identical and communicate directly with each other to hop data to and from the base station and to pass commands to each other (Fig. 1b). A mesh network is also highly fault-tolerant because each sensor node has multiple paths back to the gateway or to other nodes. The multihop system allows for much longer range than a star topology, but consumes more power since nodes need to always “listen” for messages or for changes in the prescribed routes through the mesh.

A star-mesh hybrid seeks to take advantage of the low power and simplicity of the star topology, as well as the extended range and self-healing nature of a mesh network topology (Fig. 1c). A star-mesh hybrid organizes sensor nodes around routers or repeaters which, in turn, organize themselves in a mesh network. The repeaters serve both to extend the range of the network and to provide fault-tolerance. Since wireless-sensor nodes can communicate with multiple routers or repeaters, the network will reconfigure itself around the remaining routers if a repeater fails or if a radio link experiences interference.

■ Form factor

A typical node configuration for wireless sensor networks has two main components: an RF transceiver, which is primarily analog and runs in the high-frequency, 300-MHz to 2.4-GHz ISM bands, and the MCU, which is digital and runs in a relatively low-frequency band in the kilohertz to several-megahertz range. Typically, the RF transceiver must be accompanied by a number of external components such as inductors, capacitors or surface acoustic wave filters. Because these external components are bulky and expensive, it has been challenging to in-

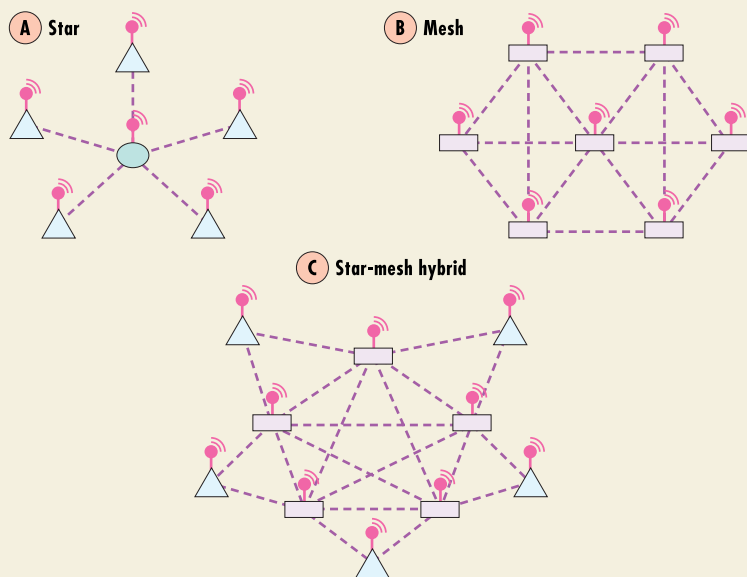


Fig. 1 Wireless sensor nets have three main topologies: star (a), mesh (b) and star-mesh hybrid (c). A star-mesh hybrid organizes sensor nodes around routers which, in turn, organize themselves in a mesh network.

tegrate RF circuitry that can meet the size and cost requirements. With rapid advances in CMOS process technology, several small, low-cost, highly integrated RF transceivers are now available.

Meanwhile, the performance and integration level of off-the-shelf industrial microcontrollers are also improving rapidly. More and more peripheral devices are being embedded in the MCU without significant cost additions.

For example, some microcontrollers come with built-in voltage supervisor/regulators, which have been traditionally considered key external components for the MCU. Many microcontrollers even include an on-chip, low-power, real-time clock and hardware encryption block, both of which help reduce the size and cost of digital circuitry.

The emergence of “combo” chips is even more encouraging. Currently, several companies are introducing the RF transceiver and the MCU in one silicon unit. Previously, this was difficult due to interference and noise issues between RF and digital circuitry. But as CMOS RF technologies improve, it has become possible to design integrated RF-digital chips, further reducing the size and cost.

■ Battery life

One critical advantage of wireless sensor networks is their independence from the wiring constraints and costs of traditional networks. This advantage will not materialize unless an adequate wireless power source is available, so power efficiency is a critical design factor. If the battery must be replaced often (every week or every month), the labor cost for battery replacement will overwhelm the initial wiring cost savings. Therefore, long battery life (typically from five to 10 years) is essential. In addition, since the philosophy of the sensor networks is “wireless anywhere,” the size of a sensor node must also be considered. In many cases, even AA batteries are too bulky to power the sensor node, so coin cell batteries are the only option.

Typically, RF components consume more than 70 percent of the total power in full-operation mode, sometimes consuming even more while receiving than transmitting. The RF components also burn significant amounts of power during switching or waking up. Consequently,

many scenarios must be considered in the power budget.

The power the RF circuitry consumes is highly dependent on the modulation scheme. Wideband RF chips, like those for Bluetooth, consume much more power than typical narrowband radios because of the complex baseband processing. Although wideband radios offer better immunity to interference, for many sensor network applications narrowband radios remain a practical and more power-efficient choice. Currently, several companies offer RF chip solutions that can achieve data rates of up

While wireless sensors offer freedom from wiring constraints, power efficiency now becomes critical.

to 1 Mbit/second with less than -85-dBm sensitivity in receive (Rx) mode, and draw no more than 10 mA of current at 3 Vdc. A few state-of-the-art RF chips have been developed that operate in the 2.4-GHz band with a current drain at 15 mA. Given the advantage of the 2.4-GHz band in terms of worldwide regulatory compliance and coverage, these RF chips are certainly viable candidates for wireless sensor network systems.

Impressive progress has also been made in the area of microcontroller power savings. Until recently, the typical power consumption of 8-bit microcontrollers has been 4 mA per Mips. With advanced chip fabrication processes and new microcontroller architectures, however, this number has recently decreased to 0.5 mA per Mips in certain new devices, helping to reduce the overall power consumption of the wireless nodes.

■ Network protocol

A typical network protocol stack, following the Open System Interconnect model, for self-organizing, wireless sensor networks is shown in Fig. 3. In general, each layer in the reference model can be designed independently, as long as interfaces between layers are consistently defined. To establish a reliable, ad hoc sensor network with a tight power budget, however, all layers in the protocol stack should be designed to meet the

same set of system-level requirements, such as energy constraint, bandwidth efficiency, adaptability and robustness. To achieve a viable solution, design trade-offs must be made at all layers, while taking into account intrinsic limitations of channel capacity, device processing speed and variations in RF link quality.

■ Physical layer

From radio signal path loss models, it is well-known that required output power varies exponentially with radio range, with an exponent of 2 in free space and 4 in cluttered areas. For the same end-to-end distance, forwarding information over multiple links, with limited transmit power for each link, can result in power consumption much lower than directly transmitting signals over one long link. To operate on small batteries for extended periods, sensor networks must employ radios with extremely low transmit and receive power and rely on multihopping for long-range connectivity. Popular radios such as cellular phones, IEEE 802.11 and Bluetooth, with typical current drain at 30 mA or higher, are not suitable here.

As discussed above, low-power chip radios fabricated with advanced CMOS process technology are now available that offer 100-foot line-of-sight range with 10-mA current drain at 3 Vdc. When these radios operate with less than 0.1 percent duty cycle, a coin cell battery with 220-mA-hr capacity can last for more than two years in a suitable environment.

In sensor network applications that rely on cooperative channel sharing and distributed data routing, however, reducing the duty cycle of individual nodes directly impacts network-level performance. Therefore, higher layers in the protocol stack must be carefully designed with this mind in order to support physical-layer implementation with ultralow duty cycles.

■ Data link layer

The data link layer in the protocol stack generally provides two main services: media-access control (MAC) and error control. Among a variety of MAC schemes, carrier sense, multiple access (CSMA) is the most popular in ad hoc sensor networks, mainly due to its ease of

implementation, but more importantly, for its efficacy in exploiting channel reuse in a large-scale network.

With CSMA, a network node always listens to the communicating channel and checks its availability before it starts transmitting a data packet. If the channel is busy, the node backs off for a random period of time before the next attempt. In most cases, such as IEEE 802.11, radios remain in listening mode even during a back-off period. However, radio circuits consume a significant amount of energy even if they're only listening. Therefore, radios should be turned off when each network node is in the back-off period or has no data to broadcast. Both listen period and back-off period are key design parameters for CSMA.

CSMA is well-suited for networks with sporadic traffic, but its performance degrades dramatically when the channel is constantly occupied by long packets or streaming data. To improve the accessibility of busy channels, particularly for critical data packets, a noncontention-based mechanism should be established, in addition to the regular CSMA. Transmission scheduling based on a cen-

tralized beacon has been an effective scheme for contention-free channel access. Beacon-based scheduling can work well in a centralized system with a star topology. For general sensor networks with a decentralized topology, however, transmission scheduling that requires proper synchronization proves

With no explicit contention control built in, an error-control scheme must be used to ensure reliability.

to be highly challenging (Goldsmith). Instead of offering guaranteed time slots, one effective way of improving channel accessibility to essential information is to prioritize data packets; packets with a high priority can occupy the channel with a low probability of collision when all listening nodes with lower-priority packets collectively back off for a longer period of time.

Further improvement on channel access can be achieved when higher-layer protocols are designed with a MAC objective. For instance, certain sensor network applications require periodical sampling of sensor data. If the application layer allows for dynamic adjustment of the sampling interval and phase shift of sampling sequence, the air channel can be effectively shared by a relatively large number of nodes with periodical transmission.

Due to the cost constraint on device hardware, CSMA with collision detection is not feasible in sensor network applications. The alternative, CSMA with collision avoidance, offers an effective approach to contention control. CSMA-CA introduces a considerable amount of overhead in network traffic, however. Without any explicit contention control, an error-control scheme must be incorporated into the data link layer to ensure an adequate transmission success rate. Common error-detection techniques such as cyclic redundancy check implemented with acknowledgment handshake prove to be effective in sensor networks. A flexible combination of data-link-layer acknowledgment (node to node) and network-layer acknowledgment (end to end) can offer an adequate

transmission success rate and achieve the desired energy efficiency.

■ Network layer

The network layer is responsible for route discovery and data packet delivery. In ad hoc sensor networks, where large numbers of nodes are deployed randomly, discovery of multihop routes (self-organization) in a mesh topology is a difficult task. It is equally challenging to maintain and repair routes (self-heal) when nodes are relocated or fail. Numerous distributed-routing algorithms have been developed over the years that support ad hoc, multihop networks. In general, these routing algorithms can be divided into two categories: proactive and reactive (Goldsmith). In a proactive routing protocol, all nodes in the network constantly maintain tables for routes between certain source-destination pairs, regardless of whether these routes are needed.

Proactive routing can deliver data packets faster than reactive routing because no discovery time is needed. The routing table size grows exponentially with the network size, however, and maintaining these tables can quickly become impractical for typical sensor networks employing a high number of nodes. On the other hand, in a reactive routing protocol, routes are discovered based on the demands of source nodes initiating data for specific destinations. Once a route is discovered, the nodes will maintain the route information for a limited period. The routing table size can be relatively small and remains constant in relation to network size, but on-demand route discovery often leads to long latency, making it ineffective for real-time applications.

Most distributed-routing algorithms for ad hoc mobile networks, proactive or reactive, are developed based on a flat network architecture such as the mesh topology. Without any hierarchy, every node in an ad hoc network takes equal responsibility for relaying packets for other nodes. In a flat network implementing a fully distributed routing algorithm, every node not transmitting needs to actively listen to the channel in order to serve as relay for route-through traffic. As a result, power efficiency of distributed-routing algorithms in a mesh struc-

Topology	Power usage	Range
Star	Low	Short
Mesh	High	Long
Star-Mesh hybrid	Low	Long

Fig.2 The star system consumes the least power but its range is shorter.

Application
Presentation
Session
Transport
Network
Data link
Physical

Fig.3 All the layers in the OSI reference model can be designed independently.

ture is intrinsically low. Using the star-mesh hybrid framework, intelligent routing can be developed that achieves high power efficiency, low latency and robust connectivity. Given limited code space for routing table storage on each sensor, reactive routing offers the more compact solution to sensor network applications. The latency issue associated with reactive routing can be effectively resolved by directing traffic to flow to a few nodes that are designated as data collection stations, with each station attracting traffic within a relatively local area.

Maintaining traffic to a local neighborhood is essential to ensuring scalability of ad hoc networks. It is observed (Gupta) that the per-node capacity in an ad hoc network decreases asymptotically as the network size increases. This result is based on the condition that the average path length between the source and destination grows in proportion to the network size. To avoid grinding the per-node capacity to a halt in a large-scale network, all traffic in the network should remain local—that is, the average hop count of data packets should be low compared with the network size.

■ Application class

Among various industrial, building and home applications, the following application classes represent the most common modes of acquiring and propagating sensor data:

- Periodic sampling. For applications where a certain condition or process needs constant monitoring, such as temperature in a conditioned space or pressure in a process pipeline, sensor data is acquired from a number of remote points and forwarded to a data collection center on a periodic basis. The sampling period mainly depends on how fast the condition or process varies and what intrinsic characteristics need to be captured. Since the duty cycle of a remote node varies in proportion to sampling rate, the application layer on the protocol stack should always seek to use a minimal sampling rate while fulfilling the monitoring requirement. In many cases, the dynamics of the condition or process to be monitored can slow down or speed up from time to time. Therefore, if the application layer can adapt its sampling rate to the changing dynamics of the condition

or process, oversampling can be minimized and thus power efficiency of the overall network system can be further improved.

Another critical design issue associated with periodic sampling applications is the phase relation among multiple nodes. If two nodes operate with identical or simi-

Given code-space limitations, reactive routing is the more compact solution for sensor networks.

lar sampling rates, collision between packets from the two nodes is likely to happen repeatedly. It is essential that the application layer detect this repeated collision and introduce a phase shift between the two transmission sequences to avoid further collision.

- Event driven. There are many cases that require the monitoring of one or more crucial variables and transmission occurs only when a threshold is reached. Common examples include fire alarms, door and window sensors, and instruments that are used intermittently. To support event-driven operations with adequate power efficiency and speed of response, the sensor node must be designed so that its power consumption is minimal in the absence of any trigger-

If the application layer can adapt its sampling rate, overall network efficiency can be optimized.

ing event, and the wake-up time is relatively short when the threshold is reached. These design requirements should be accounted for in all layers of the protocol stack.

- Store and forward. In many applications, sensor data can be captured and stored or even processed by a remote node before being transmitted to the basestation. Instead of immediately transmitting every unit of data acquired from a sensor, the aggregation and processing of data by remote nodes can potentially improve overall network performance in both power consumption

and bandwidth efficiency. The application-layer protocol should provide proper application programming interfaces for effective integration of data aggregation and processing algorithms. ■

For more on short-range wireless, see: "Multiband OFDM: Why it Wins for UWB," www.commsdesign.com/design_corner/OEG20030624S006

"Bluetooth Drives Autos Toward Hands-Free Operation," www.commsdesign.com/design_corner/OEG20020304S0067

"Developing Bluetooth Applications In Java," Part 1, www.commsdesign.com/design_corner/OEG20030611S0026, and Part 2, www.commsdesign.com/design_center/OEG20030612S0017

"Switch-Mode GaAs Power Amps Shine in WLAN Designs," www.commsdesign.com/design_center/OEG20030522S0016

"Tutorial On Spread-Spectrum Technology," www.commsdesign.com/design_center/wireless/design_corner/OEG20030506S0029

References:

Callaway, E., P. Gorday, L. Hester, J.A. Gutierrez, M. Naeve, B. Heile and V. Bahl, "Home Networking with IEEE 802.15.4: A Developing Standard for Low-Rate Wireless Personal Area Network," *IEEE Communications Magazine* (August 2002), 70-77.

Goldsmith, A.J. and S.B. Wicker, "Design Challenges for Energy-Constrained Ad-Hoc Wireless Networks," *IEEE Wireless Communications* (August 2002), 8-27.

Gupta, P. and P.R. Kumar, "The Capacity of Wireless Networks," *IEEE Transactions on Information Theory* (March 2000), 388-404.

Tod Riedel (triedel@millennial.net) is the president and co-founder of Millennial Net. Riedel was previously North American director of enterprise consulting services for Compaq Computer Corp. He holds a BS from Cornell University and an MBA from Harvard's Graduate School of Business Administration.

LOW-POWER WIRELESS

Sokwoo Rhee (sokwoo@millennial.net) is chief technology officer and co-founder of Millennial Net. Prior to that, he was a research associate at MIT focusing on wireless biomedical instrumentation. Rhee holds a B.S. degree from Seoul National University, an M.S. degree and a Ph.D. from Massachusetts Institute of Technology.

Sheng Liu (sliu@millennial.net) is vice president of research and co-founder of Millennial Net. Prior to that he was a senior development engineer at Raytheon and a research scientist at MIT, directing industry-sponsored research programs in the areas of controls, robotics, simulations, signal processing, and mechatronics. He has a B.S. degree from National Taiwan University, an M.S. degree from University of California, Irvine, and a Ph.D. from MIT.