Millennial Net

Wireless technologies for industrial manufacturing applications

BY MARTIN HANSSMANN, SOKWOO RHEE, & SHENG LIU

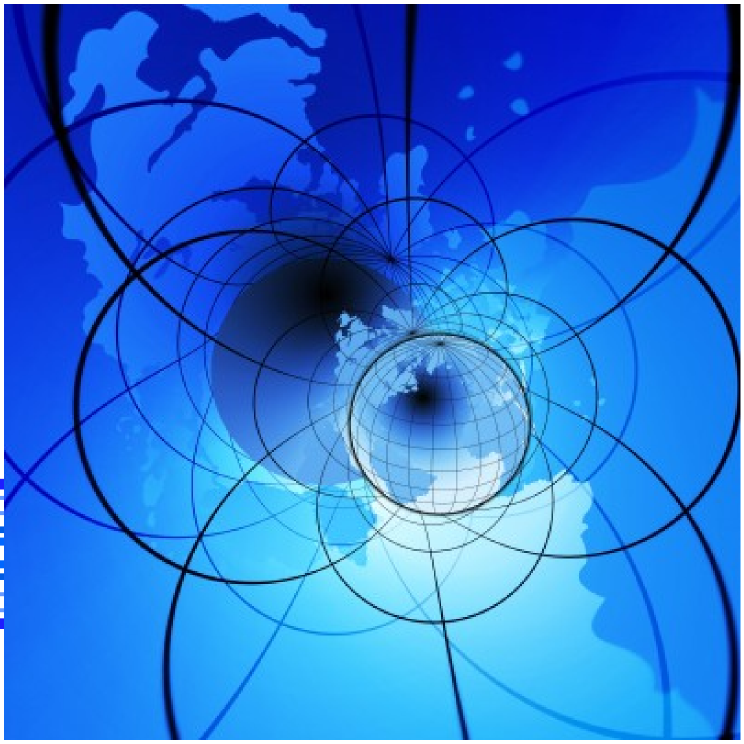

NO WIRING CONSTRAINTS

WIRELESS TECHNOLOGIES HAVE existed within industrial and commercial facilities for many years; however, some recent advances in wireless technologies have opened up new applications. Historically, two-way radios were one of the earliest forms of wireless communications in manufacturing; today, this is often replaced by cellular/mobile phones, pagers, or Personal Digital Assistants (PDAs). However, in most industrial facilities, the use of wireless has extended considerably beyond mobile devices and has become a central component within the overall IT infrastructure and productivity of the company.

Industrial Applications of Wireless Technologies

The use of wireless technologies can be broadly classified into two categories: a) those predominantly used for the transfer of information and b) those implemented predominantly for the use of data acquisition architecture. This article will review current technologies falling into both of these categories and will then focus more intensely on wireless data acquisition strategies.

The robustness, security, attractive cost, and high bandwidth of 802.11 wireless communications (often referred to as Wi-Fi) have resulted in many companies adopting this technology within their back office IT infrastructure; most commonly to provide wireless connectivity to portable computers for use in conference rooms, shop floor areas, and out buildings. The primary objective is to extend the company's existing Ethernet network access to users outside of their wired environment. For many companies, this has provided a relatively low-risk entry into establishing a familiarity with wireless technologies for their IT functions as well as their employees. Because these networks typically form part of the back-office infrastructure with access available to large groups of individuals, their applicability within control architectures is often seen as a risk to the robustness of the overall plant control solution.

WIRELESS MESH NETWORKS ENABLE NUMEROUS EMBEDDED APPLICATIONS TO BE INDEPENDENT OF THE ISSUES OF WIRING COSTS AND PHYSICAL CONSTRAINTS.

While Wi-Fi is not typically part of an industrial controls application, the use of 802.11 technology for long-range point-to-point communications is becoming more common for use in supervisory controls. In this scenario, in addition to a local control architecture serving the needs of the plant, there is an ability to remotely view (and in some cases interact with) the shop floor equipment via dedicated wireless transceivers that link the facility to a remote site. This type of wireless architecture is found in offshore drilling and oil production operations, refineries, distributed mining operations, and others. It offers remote monitoring of all aspects of an operation with significantly fewer resources. An additional key benefit of an operation linked via wireless monitoring is to be able to better optimize the overall throughput through an improved real-time understanding of each of the operational elements. Although the use of public band frequencies such as 900MHz and 2.4GHz is prevalent, it is also possible to use proprietary frequencies that have the benefit of increased security and resistance to interference from other local RF applications. The range of these transceivers can vary from less than one mile to twenty miles or more to accommodate the distance between the various elements of operation. However, because this is a point-to-point technology, environmental variations and other interferences will affect reliability; therefore, the use of these solutions is predominantly limited to secondary monitoring and open-loop applications.

As facilities have become more experienced with wireless technologies within their operations, these technologies have been extended to more integrated applications, including the wireless extension of common bus technologies such as DeviceNet, Modbus, and Fieldbus. Because of the ability to rapidly integrate these wireless bus technologies within an existing facility’s control solution, they provide an attractive alternative to wiring in cases where there is a cluster of sensors remote from the main facility. However, because of the nondeterministic nature of all wireless technologies, the design of a control solution employing this technology must be sufficiently tolerant of environmental and other interference to ensure continued robust operation of the plant in the event of wireless communication disruptions.

Recently, a more robust wireless technology, called the mesh network, has been developed for wireless sensor integration within a facility. Wireless mesh networks enable numerous embedded applications to be independent of the issues of wiring costs and physical constraints. At the same time, the unique nature of wireless mesh networks requires many fundamental challenges to be addressed, with robustness and scalability being the most important issues. In typical wireless mesh network systems, these two factors generally go against each other because of the self-adjusting and nonhierarchical nature of the networks. For example, as the number of nodes increases, the robustness of network becomes harder to guarantee. On the other hand, to make the network more robust, smaller sized networks are preferred. When the issue of network responsiveness is added on top of these two issues, the equation becomes even more complicated.

Using the wireless mesh networking technology, hundreds or thousands of sensors and actuators can be placed without any wiring constraints. Due to the ad hoc nature of wireless mesh networks, the sensor nodes form a network automatically with minimal human intervention. The network is maintained autonomously, healing itself if any damage occurs to the network. The wireless mesh network is reliable and robust because the network “learns” based on its own changes or problems in the topology and adapts itself to situations very quickly, even at the individual packet-transmission level. In essence, a wireless mesh network does not assume any predefined topology or placement. The network is designed to adapt itself to them at the initialization and keep adapting itself continuously.

Much of the technology development in the wireless sensor networking field comes out of the traditional static networking industry (i.e. DSDV, DSR, AODV), which is based on a relatively static environment and stable links. As a result, these developments are challenged to address all of the above-named requirements to the level required for production purposes. By contrast, the research on Persistent Dynamic Routing comes out of the sensor research field that presumes an unstable, rapidly changing environment.

In this article, several important factors that determine the performance of the wireless sensor network are discussed. In most cases, these factors have a conflicting impact on each other, so that trade-off decisions are unavoidable in network design. Among those factors, scalability, reliability, responsiveness, mobility, and power efficiency will be discussed more in detail. In addition, the advantage of Persistent Dynamic Routing will be discussed, and finally, the concept of increasing capacity to improve scalability will be introduced along with several techniques to achieve this.

Challenges in Designing Wireless Sensor Networks

> SCALABILITY REFERS TO THE ABILITY OF THE NETWORK TO GROW, IN TERMS OF THE NUMBER OF NODES, WITHOUT EXCESSIVE OVERHEAD.

Many people believe that if a good wireless sensor node is built, a good wireless sensor network can be built. Although a good sensor node is an important part of a wireless sensor network, this is not whole story. The characteristics of a good sensor network include: scalability, reliability, responsiveness, mobility, and power efficiency.

A sensor node is just a very small part of the wireless sensor network. There are more important design challenges that go into making a network good. To build a good wireless sensor network, all of the above factors must work in harmony. The challenge is that this harmony can be difficult to achieve. The complex inter-relationship between these factors is a balance; if these factors are not managed well, the network can suffer from overhead that negates its applicability in the real world.

Scalability

Scalability refers to the ability of the network to grow, in terms of the number of nodes, without excessive overhead. This is an important real-world requirement where networks must support more than the small handful of nodes typical in a pilot implementation.

This is due to the network overhead that comes with the increased size of the network. In ad hoc networks, the network is formed without any predetermined topology or shape. Therefore, any node wishing to communicate with other nodes should generate more packets than its data packets, i.e., control packets or network overhead. As the size of the network grows, more control packets will be needed to find and keep the routing paths. Moreover, as the network size increases, there is higher chance that communication links get broken in communication paths, which will end up creating more control packets. In a small network, the amount of control packets is almost negligible. But when the network size starts increasing, the overhead increases rapidly. Since the available overall bandwidth is limited, the increase of overhead results in the decrease of usable bandwidth for data transmission. As the network size grows further, there will be very small amount of bandwidth left for application data transmission.

Reliability

Reliability is the ability of the network to ensure reliable data transmission in a state of continuous change of network structure. Typically, there is an inverse relationship between scalability and reliability in ad hoc wireless networks; as the number of nodes in the network increases, the more difficult it becomes to ensure reliability.

This scalability characteristic of ad hoc networks described above imposes an interesting question on the reliability of the network. Since an ad hoc network is designed to automatically adapt itself to a changing environment or interference, it will issue more control packets when it faces dynamics. More dynamics in the environment will increase the number of control packets and, at some point, the network cannot sustain the amount of overhead caused by the dynamics, which will result in less reliability of data transmission. This breaking point will show up earlier in a large-sized network. So, network scalability and reliability are closely coupled and typically they act against each other.

Responsiveness

Responsiveness is the ability of the network to quickly adapt itself to changes in topology. To achieve high responsiveness, an ad hoc network should issue and exchange more control packets, which will naturally result in less scalability and less reliability.

Mobility

Mobility refers to the ability of the network to handle mobile nodes and changeable data paths. Generally, a wireless sensor network that includes a number of mobile nodes should have high responsiveness to deal with the mobility. So, it is not easy to design a large scale and highly mobile wireless sensor network.

Power Efficiency

Power efficiency, the ability of the network to operate at extremely low power levels, also plays an important role in this complex equation. A typical method for designing a low-power wireless sensor network is to reduce the duty cycle of each node. The drawback is that as the wireless sensor node stays longer in sleep mode to save power, there is less chance that the node can communicate with its neighbors. This will decrease the network responsiveness and may also lower reliability due to the lack of the exchange of control packets and delays in packet delivery. In addition, a more complicated synchronization technique will be necessary to keep more nodes in low duty cycle, which may also affect scalability.

Managing the Design Tradeoffs

The complex issue of managing these tradeoffs comes down to how the communication overhead can be minimized while maintaining the network reliability and responsiveness. As explained above, there are many conflicting factors involved in the design of wireless sensor networks, and there are always tradeoffs. When choosing a wireless sensor network for an application, careful consideration of the balance of these factors within the context of the needs of the application is critical.

Addressing the Design Challenges

This article presents two sets of techniques address the real-world design challenges outlined above. 1) Persistent Dynamic Routing (PDR) is a breakthrough set of techniques that addresses all of these requirements to support production-grade wireless sensor network implementations. 2) High-Capacity Wireless Sensor Networks (HC_WSN) present emerging research on the issue. These techniques are not mutually exclusive and can be used together to manage the design tradeoffs of wireless sensor networks.

Persistent Dynamic Routing

PDR provides a mechanism for the network to ensure reliable data transmission without dropping data packets. Combined with the technique of dynamic route discovery, which discovers the best route for packet delivery on the fly, PDR enables a level of scalability and power efficiency that other networking systems cannot achieve.

Almost every existing ad hoc network protocol assumes some level of static status of the network. For example, the route discovery process of AODV assumes there is at least a short duration during which a “snapshot” of the complete route to the destination is possible. The data packet of DSR carries the full route information in itself, which assumes the existence of a “full” route at that moment.

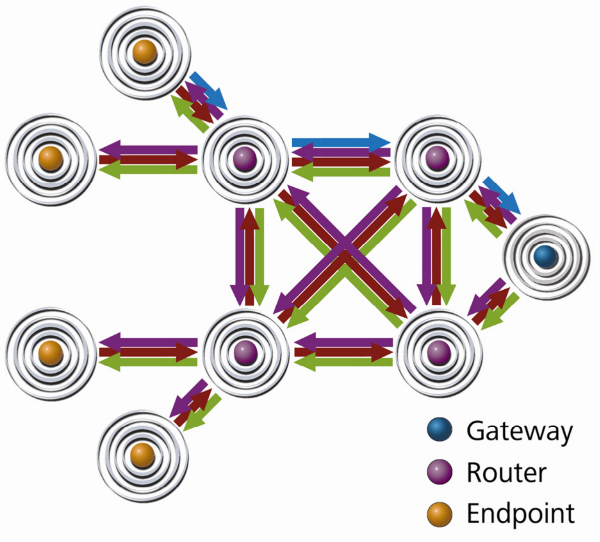

Flooding approaches. Gateway establishes tree structure for dynamic addressing. Route discovery packet broadcasts through the network establishing mesh routing (AODV), which floods the network. Route response packet broadcasts to validate the route. Data packet sent to the gateway.

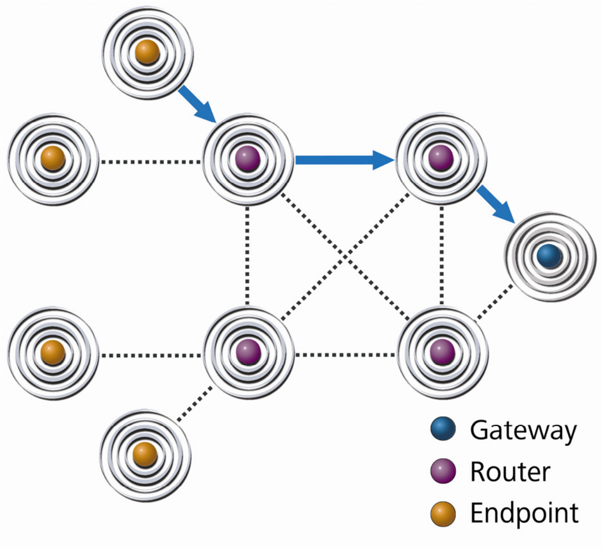

Persistent Dynamic Routing. With PDR, the endpoint sends the data packet, which is dynamically propagated through the network and delivered to the gateway.

In the case of a relatively static network with low fluctuation and interference, this assumption can hold with a reasonable level of success. But in a highly dynamic environment, an assumption of this kind of “quasi-static” status does not hold. In other words, the network may be continuously changing so that it is impossible to establish a full route from the source to the destination at a point in time. In this case, traditional routing algorithms such as AODV or DSR can present difficulties. For example, in an AODV system, the source will keep sending the route discovery packet but will not get a definite route response from the destination, which will result in continuous flooding of the network with route discovery packets. As a result, the data packet will not even be sent into the network since route discovery process is incomplete. More route discovery packets translate into more overhead. This problem becomes more serious in a large-sized network; since the route discovery is essentially a flooding process, the efficiency of the network will drop significantly.

With PDR, the data packet does not need to wait until the route discovery process grabs a “full” route at a moment in time. A data packet is released and navigates through the network with the best knowledge it has collected from its neighbors at that moment. It works in a manner similar to the mechanism of navigating a maze without any prior knowledge of the maze. The data packet does not wait until the full route is confirmed; rather, it starts navigating the network with whatever information it has about the destination.

PDR can significantly decrease the overhead of packet delivery in a highly dynamic network because it does not send excessive numbers of route discovery packets nor does it use proactive route updates. Also, the route discovery packet in PDR does not go more than one hop in each discovery process, resulting in less flooding of the network. In practice, flooding is used only once at the very beginning of the network formation and, from then on, route discovery is only done in the local area to collect knowledge on the best route to the destination. This “best knowledge” has no guarantee that it is correct, and the data packet does not “ask” for that kind of guarantee. In this sense, PDR is based on the probabilistic rather than deterministic approach. In a relatively static network, the higher probability that the destination matches the actual deterministic route would give PDR the same level of performance as AODV, if not better. In highly dynamic environments, PDR produces significantly less overhead in packet delivery than the AODV flooding approach.

RELIABILITY IS THE ABILITY OF THE NETWORK TO ENSURE RELIABLE DATA TRANSMISSION IN A STATE OF CONTINUOUS CHANGE OF NETWORK STRUCTURE.

High-Capacity Wireless Sensor Networks

Lee et al [1] pointed out that the average number of packets per data delivery increases as the size of the network increases in AODV. The increase of packets per data delivery is not just proportional to the size of the network; rather, it increases exponentially. When the size of the network grows further, almost no data can be delivered since the number of packets needed per data delivery approaches infinity. This phenomenon is largely due to the ad hoc nature of the wireless sensor network. As the size of the network increases, more links are likely to break between the source and the destinations, and more packet exchanges are necessary to heal and reconstruct the delivery paths. Therefore, the issue is how to minimize the number of packets per data delivery in a large-sized wireless sensor network.

One way to increase the scalability of the network without increasing overhead is to increase the capacity of the network, i.e. to increase the volume of data that can be transmitted through the network. Increasing capacity will enable the network to handle more packets in a given network size. One way of achieving this is to increase the network bandwidth; however, higher bandwidth typically consumes more power, negating one of the critical factors (and benefits) of wireless sensor networks. Therefore, alternative approaches are needed. Several available approaches include multiple gateways, frequency reuse, and data aggregation. Each approach can increase capacity, but comes with a tradeoff.

Multiple Gateways

Typical wireless sensor networks use a single gateway to aggregate the sensor data and pass it to the host application or system. Adding additional gateways to the network can reduce the number of hops the data must take before reaching the gateway. Since each data packet takes fewer hops, more data packets can be transmitted through the network without adding to the total number of hops. In this scheme, the multiple gateways must be connected via some kind of “fatter pipe” such as Ethernet or 802.11.

Frequency Reuse

Another technique is to reuse radio frequencies. Using multiple channels through channel hopping can increase the robustness of the network by increasing resistance to interference and RF noise. But channel hopping or dynamic channel switching can also be used to send more data through the frequency band. The tradeoff here is that more power will be consumed to use multiple frequencies in a network. For example, a node will require a more complex route-discovery/link-recovery process, which will take a longer time and, as a result, burn more power. A longer route-discovery process could also result in a less responsive network. In practice, using multiple frequencies in one network can also require more memory and processing speed in a sensor node.

Data Aggregation

Data aggregation is another method for increasing network capacity without increasing the communication bandwidth. With this technique, multiple data packets are aggregated into one (e.g. at the routers) packet which is then transmitted through the network. Packet aggregation serves to minimize the number of packets propagating through the network, which minimizes the overhead, thus maximizing the throughput of real sensor data. This technique typically results in a drop in responsiveness and increases the delay in packet delivery.

Commercial and Industrial Applications for Wireless Sensor Networks

The use of wireless sensor technology has been proven in a large multitude of applications within the commercial and industrial markets. For the most part, these fall into two categories: 1) condition or asset health monitoring, and 2) secondary monitoring. Their promise of increased robustness, low power, and relatively low cost has led to proof of concept deployments within virtually every sector of industry, including oil and gas, power generation, building automation, mining, medical, and many others. Although the technology has been proven to be capable of meeting the needs of many of these applications, it has often not achieved widespread deployment for several reasons, including the need for 1) business process changes to fully realize the business benefits, 2) incremental training of existing field technicians more familiar with wired technologies, 3) long term cost of ownership data and understanding, and 4) industrial standards allowing for increased co-existence and interoperability with other wireless devices.

A WIRELESS MESH NETWORK DOES NOT ASSUME ANY PREDEFINED TOPOLOGY OR PLACEMENT.

This article will present several examples where wireless mesh networks have achieved greater success in industrial gases monitoring, secondary monitoring, and energy management.

In precision industrial fabrication facilities such as aeronautical manufacturing plants and precision casting foundries, it is necessary to conduct “secondary” data measurements to confirm that fabricated assemblies and sub-assemblies are within specified tolerances. Ideally, this is done on the shop floor directly after completion of a specific fabrication process or step. Historically, precision measurement tools and gages were used to perform such measurements and the results were recorded on a paper-based system for later comparison with the specifications. By implementing a wireless sensor network, these facilities have been able to directly link real-time secondary monitoring to plant operations without compromising data integrity and reliability (Wilkinson [5]). A further benefit is the automatic historical data logging and records management intrinsic to such a centralized monitoring system. This is of particular importance in regulated industries.

The use of wireless sensor networks holds a great deal of promise for asset management applications, e.g., the use of a mesh network to monitor industrial gases. Industrial or specialty gases are widely used within most industrial facilities. The storage of these gases is typically remote from the primary operations area. Because these gases are integral to the manufacturing operation, monitoring their usage and real-time availability is important to overall plant operations. Typically, this has necessitated facilities personnel to manually monitor existing consumptions and then project replenishment needs. The added complexity of aligning replenishment needs with delivery schedules of the gases supplier also needs to be considered. Implementing a wireless mesh network to automatically monitor and report current usage and current cylinder pressure of the gases provides an excellent asset optimization opportunity. To facilitate both facilities personnel (users) of the gases and the gases supplier (vendors) to access this data, it is commonly integrated with an internet-accessible monitoring application. This can be further optimized by the vendor through an automated or semi-automated replenishment function providing both users and vendors an increased efficiency within their respective operations.

Commercial means of monitoring (and controlling) energy consumption have long been available using traditional wired infrastructure. However, in many cases this limits monitoring due to the relatively high costs of implementation in existing facilities or for use with mobile applications. Ideally, monitoring the specific energy consumption of each piece of equipment within a facility can help drive equipment-specific energy reduction strategies as well as an improved allocation of overall costs within an industrial facility. The latter is becoming increasingly important, as the cost of energy has increased to become a significant factor in the overall operational costs. The use of a wireless sensor network provides an ideal platform for such a monitoring application. Because mesh networks can support multiple hundreds of sensors, even large plants can conduct widespread monitoring. In addition to monitoring energy consumption, it is also possible to simultaneously monitor other information, such as temperature, pressure, and other metrics, using low power (battery powered) devices. This provides a comprehensive data source for optimizing plant energy consumption strategies. Simply monitory consumption has proven to lead to significant savings, as demonstrated by a research project by Pacific Northwest national Laboratory [6]. Furthermore, in cases where there is a transient monitoring need, the rapid deployment and non-intrusive nature of wireless mesh networks easily allows them to be re-deployed once the necessary data has been gathered.

Conclusions

Wireless technologies are gaining increased acceptance as a viable technology within the industrial solution space. However, it is important for control system architects to fully understand not only the benefits but also the challenges and limitations of wireless technologies. The complex inter-relationships between the various wireless technology characteristics require careful tradeoffs to be made in the design of the solution. This is especially true for applications that need to handle continuous data streams where high capacity is critical to maintain the scalability and reliability of the network. Although still relatively new, Mesh Networks have proven themselves to have significant advantages in robustness, scalability, and low power consumption over other wireless technologies for use in industrial monitoring applications. Mesh Networks employing Persistent Dynamic Routing address the scalability challenges of wireless sensor networking that can limit implementation in production situations by significantly increasing the overall packet delivery rate and reducing network communication overhead. High Capacity WSN, an approach that will emerge as the important next trend in the wireless sensor networking industry, can be implemented in addition to PDR techniques and will further support the increasing number of applications that require the wireless sensor network to handle greater amounts of data.

References

[1] Lee, SJ., et al, “Scalability study of the ad hoc on-demand distance vector routing protocol,” International Journal of Network Management, Volume 3, Issue 2, (2003) pp 97-114

[2] Goldsmith, A. J., Wicker, S. B., “Design Challenges for Energy-Constrained Ad Hoc Wireless Networks,” IEEE Wireless Communications, August 2002

[3] Perkins, C. E., Royer, E. M., Das, S. R., “Performance Comparison of Two On-demand Routing Protocols for Ad Hoc Networks,” IEEE Personal Communications, February 2001

[4] Abolhasan, M., Wysocki, T., Dutkiewicz, E., “A Review of Routing Protocols for Mobile Ad Hoc Networks,” Ad Hoc Ne Ad Hoc Networks,” Ad Hoc Networks, Volume 2, Issue 1

[5] Wilkinson, General Manager – Starrett Advanced Technology Division. “Data Integrity and Industrial Reliability via Wireless Mesh Networks”

[6] D.J. Hammerstrom, Principal Investigator – Pacific Northwest National Laboratory ” Pacific Northwest GridWiseTM Testbed Demonstration Project”

Contact:
Millennial Net, Inc.
285 Billerica Road
Chelmsford, MA 01834 USA
Tel: +1 978-569-1921
Fax: +1 978-256-3162
www.millennialnelnet.com