Control Over Wireless:
Current Applications and Future Opportunities

Jay Werb, ISA100 Wireless Compliance Institute
Soroush Amidi, Honeywell Process Solutions

Background

This paper describes how systems based on the ISA100.11a [1] wireless standard can be used for control.

When industrial users consider applications for wireless control, they generally cite three key benefits: improved reliability, improved control, and cost savings.

Improved reliability: In some situations where wired connections are exposed to extreme conditions, a wireless replacement can actually be more reliable. In other situations, a redundant wireless connection can serve as a backup for wiring.

Improved control: With introduction of wireless sensing devices, it becomes reasonable to add additional sensing points to a process, for more optimal operation.

Cost savings: Up to 90% of installed cost of conventional measurement technology can be for cable conduit and related construction. New applications are now economically feasible with wireless, and some existing applications can be more widely deployed.

To begin our review, let us examine two case studies of how wireless is typically used today for control.

Wireless Control Example – Refining

Our first case study is from an oil refinery, where wireless is deployed in a fuel/air controller for furnace units. A cracking unit's furnace requires reliable temperature control. In the past, wiring from temperature transmitters was not reliable due to wear and tear caused by the application, resulting in frequent wire replacement, furnace inefficiency, and production scrap. The cracking unit is an important unit of a refinery and close control of the furnace temperature is required to obtain the best conversion rate and maximize the profitability of the plant. Users needed a solution that eliminated wires while sending reliable data that could be used by the fuel/air controller.

The solution involved an ISA100 compliant network and battery powered wireless temperature transmitters. Process data from wireless temperature transmitters feed into the existing

fuel/air controller that controls the temperature of the furnace. So, in this case, wireless data is used as an input for closed loop control. Wireless temperature transmitters provide data to the air/flow controller, which then manages the temperature at the furnace.

Similar applications may be found in steel plants, where wired transmitters have difficulty with melted cables. Steel plants are important users of wireless technology for various applications including control.

Wireless Control Example – Chemical

Our second case study is from a chemical manufacturer. In this application, plant operators need to maintain an appropriate level of bitumen in holding tanks. With improved insight to the tank levels, plant operators are able to more efficiently manage inventory levels. Previously, control relied on manual reads of the bitumen tank levels.

The solution was to install an ISA100 wireless network with wireless level indicators (differential pressure gauges). Level indication is sent continuously to the console operator, who then manages the bitumen level based on data from the wireless transmitters. The main benefit is more efficient bitumen inventory management without over filling.

This is an example of wireless used in an open loop control where a control operator decides when to open and close the valve to control the bitumen level. This example shows the trust of the operator in the data sent by the wireless transmitters.

General Model for Control Over Wireless

Now that we have reviewed some simple case studies and benefits, we will consider a more general model of control over wireless.

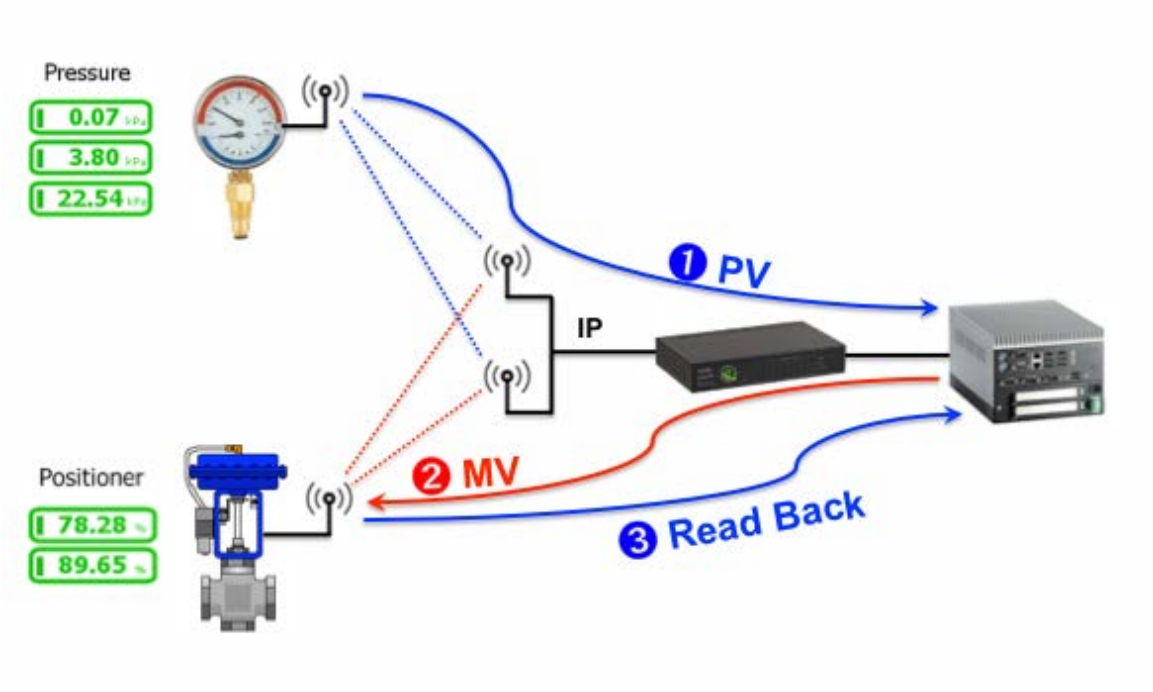

Figure 1: Control Over Wireless

Figure 1 shows a reference physical configuration in a wireless control system. On the left, we have a wireless pressure sensor on the top and a wireless positioner on the bottom. They both have redundant wireless connections that directly communicate to a high performance IP backbone. Once the message is on the IP backbone, it is forwarded quickly and very reliably to a controller via an ISA100.11a gateway. When people start thinking about wireless control, they usually imagine something like this picture.

The IP backbone may be wired, such as Ethernet, or it may itself be wireless, such as a WiFi mesh. The IP backbone is normally a plant-wide resource supporting a variety of applications such as data, voice, and video. In the limited context of this paper, the IP backbone provides a reliable and fast connection between locations in the plant. When the message is on the backbone, it propagates quickly and reliably to wherever it needs to go.

This picture shows one wireless hop between the wireless devices and the high performance IP backbone. This is a typical configuration for a wireless control system where we take advantage of the high-speed backbone. The initial wireless connection to the backbone has been compared to a highway on ramp.

Figure 1 also shows the data flow. A process variable (PV), pressure in this case, is published to the controller. Based on this input and other considerations, the controller sends a manipulated value (MV) to the positioner. The positioner sends the actual position back to the controller. All this needs to happen on the time scale required by the process, such as 1-2 seconds.

In practice, the solution does not have to be 100% wireless. The input device may be wireless and the output device may be wired and vice versa. For example, in the refinery example above, the temperature sensor was wireless and the actuator was wired.

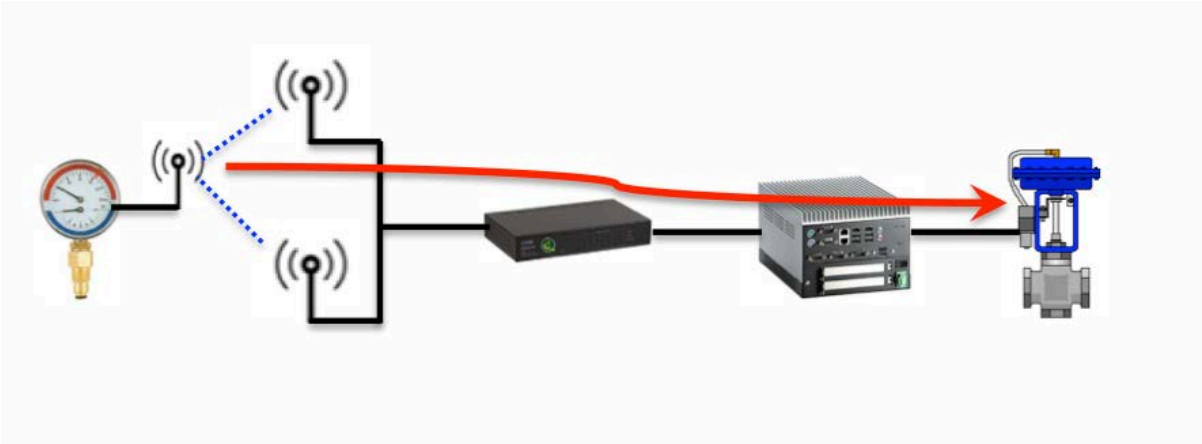

Figure 2: Wireless Sensor, Wired Actuator

A generalized "entry level" architecture is shown in Figure 2. A wireless transmitter periodically publishes data to a controller through an ISA100 network. The controller executes control logic and uses wired fieldbus communication to transmit commands to a positioner. The input is wireless but the output is through a wired connection.

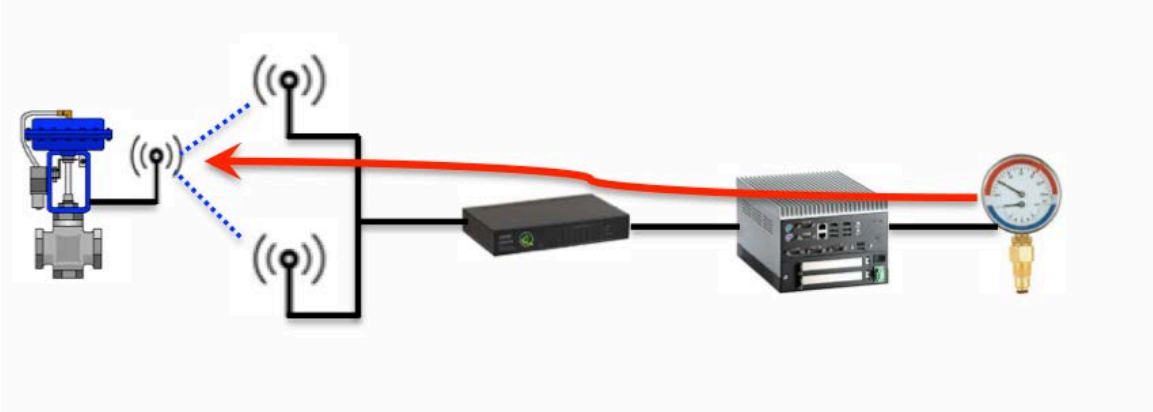

Figure 3: Wired Sensor, Wireless Actuator

Figure 3 shows the reverse. On the right, a wired input device transmits sensor data to a controller, where control logic is executed. Then actuation occurs via a wireless link through an ISA100 network.

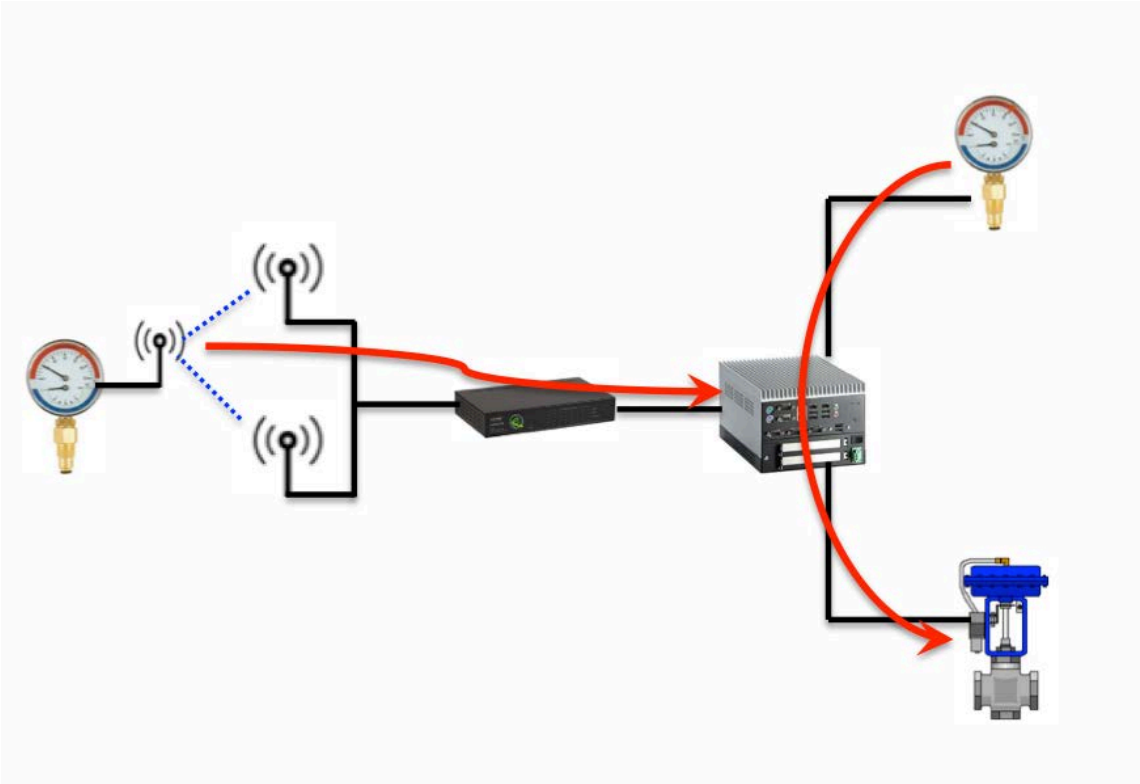

Figure 4: Wired Sensor, Wireless Actuator

Figure 4 shows another combination, with a wireless sensor providing secondary input. On the right, the primary input is wired, and the output is also wired. On the left, a wireless sensor provides a secondary input that would be otherwise unavailable.

The four examples shown in Figures 1-4 illustrate a basic range of realistic ways to use ISA100 wireless technology today in control schemes.

Wireless Control – General Benefits

Three main reasons are generally cited for wireless control: improved reliability, improved control, and cost savings. These benefits apply to both open and closed loop applications, and are summarized in Table 1.

Benefit	Description
1. Improved reliability	Troublesome wired sensors replaced by wireless counterparts. Wireless may serve as a backup for wired technology.
2. Improved control	Add wireless devices to existing processes for more optimal control.
3. Cost savings	Up to 90% of installed cost of conventional measurement technology can be for cable conduit and related construction. New and existing applications are now economically feasible.

Table 1: Major benefits of wireless for control

The first benefit, as previously described in the bitumen holding tank example, is applicable to situations where a strong wireless connection is actually more reliable than a troublesome wired connection. In those cases, one might reasonably replace wired sensors with their wireless counterparts. Alternatively, it might make sense to add a wireless sensor that can serve as a backup if the wired connection fails.

Second, wireless is being used for secondary inputs to optimize an existing wired process. A secondary input may have not been feasible without wireless. With the ability to bring in data wirelessly from even the most remote plant areas, new opportunities to optimize plant processes can be sensibly implemented.

Third, the much lower cost of a wireless installation makes a broad range of applications economically feasible for the first time. That's basically the second use case where a process that was controlled manually (manual reading and then manual action) can be automated (automated reading with manual action, to be followed at some future date by automated action). New applications in a plant can now be cost justified for the first time, and existing wired applications can be more widely deployed in wireless form.

These benefits are in addition to gains from using the same network and system for monitoring applications, which is often the initial application of ISA100.11a wireless technologies.

Wireless Control – Application Classes

The ISA100.11a standard was designed to cover a wide range of applications, summarized in the standard as application classes 0-5. The ISA100 usage classes are shown on Table 2.

Safety	0	Emergency action	Always critical
Control	1	Closed loop Regulatory control	Often critical
	2	Closed loop Supervisory control	Usually non-critical
	3	Open loop control	Human in the loop
Monitoring	4	Alerting	Short-term consequences
	5	Logging Downloading/uploading	No immediate consequences

Table 2: Application Classes from ISA100

The most critical application class is emergency action, Class 0, shown at the top of the chart. The least critical application class is logging, Class 5, shown at the bottom of the chart. This paper focuses on control applications, Classes 1 to 3 which are circled in Table 2.

Control	**1**	**Closed loop Regulatory control**	**Often critical**	**Control of primary actuators High frequency cascades**
	2	**Closed loop Supervisory control**	**Usually non-critical**	**Low frequency cascade loops Multivariable controls Optimizers**
	3	**Open loop control**	**Human in the loop**	**Manual flare Remote opening of security gate Manual pump/valve adjustment**

Table 3: Control Applications Enumerated in ISA100

Table 3 summarizes the control application examples for Classes 1, 2 and 3 as described in the ISA100.11a standard itself.

The examples for open loop control (Class 3), as described in the standard, all involve a user in the loop. An operator manually initiates a flare, and watches the flare. A guard remotely opens a security gate. An operator manually adjusts pump or valve.

Closed loop supervisory control (Class 2) applications are usually non-critical. Examples include low frequency cascade loops, multivariate controls, and optimizers.

Class 1, closed loop regulatory control, is at the top of the ISA100.11a control hierarchy. The standard describes direct control of primary actuators, such as:

- where a host connection is available on demand of 99.99% or more;
- with link outages of more than a half-second intolerable; and
- with demand rates once every 4 seconds.

High frequency cascades are also referenced in the ISA100 standard.

Class 0, safety applications, are not explicitly covered by the standard, but some applications with safety components are commonly considered as candidates for wireless. If safety application parameters are consistent with Classes 1-5 requirements, then ISA100 technology is applicable. For example, gas monitoring and safety showers are commonly cited as candidates for wireless.

Wireless Control – Latency Requirements

What are the latency requirements for control? As an introduction, Figure 5 shows a baseline wireless control scenario.

Figure 5: Baseline Control Scenario

Two devices are shown, often within radio proximity of each other. One device publishes data to the other device at regular intervals, and the other device receives that signal and does something useful with the information. A third device, not shown in the picture, might serve as a controller, and/or control logic may reside in the devices themselves. This is a fundamental configuration layout for control. 1-second latency is frequently mentioned as a baseline requirement for a configuration like this, as shown in Figure 6.

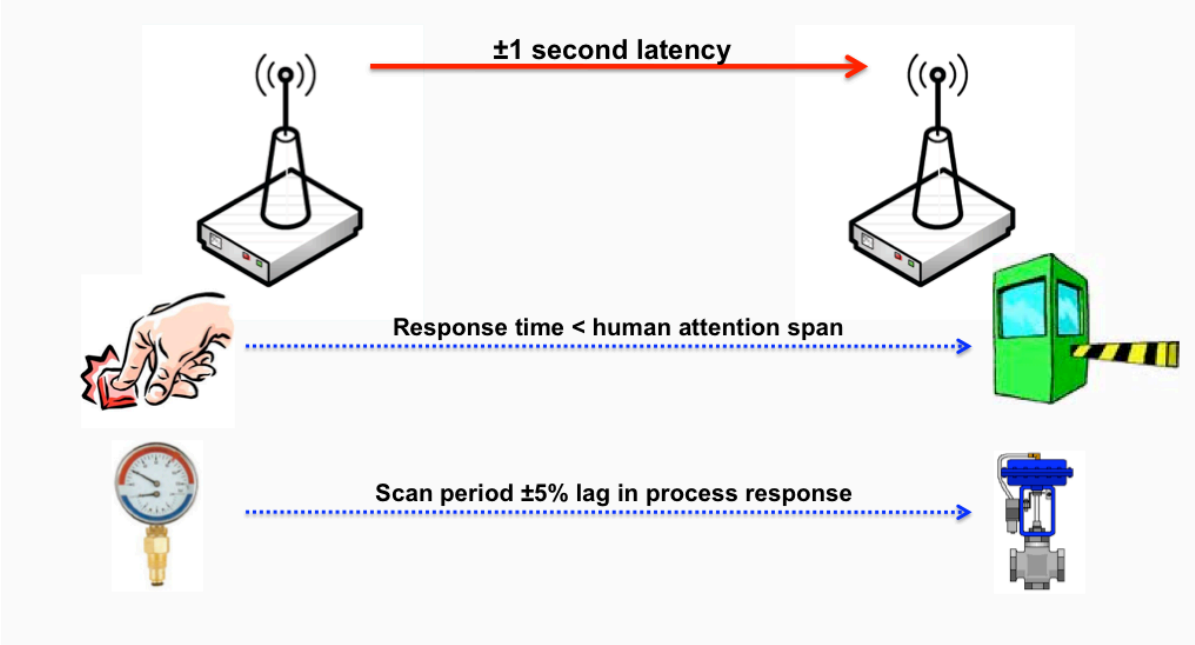

Figure 6: Latency Requirements for Control

Figure 6 suggests that something on the order of 1-second latency, end to end, should be the basic performance reference for control. Some applications can manage with less performance, and a few may need a faster system. That being said, ISA100 was designed under the principle that 1-second latency* covers a reasonable range of applications, without involving painful user trade-offs or asking the user to apply unproven new concepts in plant operation. Maintaining and supporting existing operating principles is one of key benefits of using ISA100 technology.

For example, let's consider a human-in-the-loop of a control application, shown here as a person pressing a button. Normally, users want to see some feedback from their actions. It is well known that when a response time exceeds about 2 seconds, the user's attention will start to drift and/or the user will start to doubt the reliability of the action. So we suggest 1-2 seconds, end-to-end, as a reasonable target for incorporating a human-in-the-loop in a control schema.

As another example, consider how automated control works within a control schema, shown here as a pressure sensor feeding data to a valve controller. Generally, the scan period should be some fraction of the lag in the process response. Figure 6 gives a value of 5% as a rough rule of thumb, with a plus-or-minus to indicate that the actual value depends on the application.

Where did the value of 5% come from? Basically, it's a guideline that is common in actual applications. For example, to prevent excessive valve activity in the presence of higher

* For simplicity of exposition, this paper treats reporting interval and latency as being roughly equivalent. Generally, our concern here is that control inputs and outputs do not remain unknown states for longer than the target "latency".

frequency noise, one might reasonably make the period between scans less than one-tenth of the dead time, or one-twentieth of the lag in the process response. In practice, that kind of thinking tends to result in numbers in the range of 5% shown Figure 6. (More detail can be found in reference [2].)

How does this map to actual systems? When we look at actual ISA100 solutions and systems intended for control, scan rates on the order of one second are common, with system architectures designed to support that.

A reader might observe that these reference performance parameters are not substantially downgraded just because it's wireless. That's the point of course. Control is control, however the data is transmitted. When ISA100 was conceived, the designers had the goal of wireless control very much in mind, and worked from the ground up to make sure it would meet reasonable user expectations for control. That DNA is clearly evidenced in actual ISA100 systems.

For example, one commonly asked question is: Do we need to actually transmit data every scan, even if the data hasn't changed? The answer of course depends on the application. Control systems go to a failure state when data is uncertain. If wired control logic requires current data, why should that change just because the data channel is wireless? Reflecting this line of thinking, we have observed that most ISA100 control solutions are optimized to transmit data at each interval, with the message propagating quickly and very reliably.

Wireless Control – System Requirements

Table 4 shows the main system requirements for wireless control.

1. Rate and Latency	**Publication rates 1-2 seconds** **Controlled latency, ~50% publication rate** **4 Hz publication in constrained configurations**
2. Flexible System Architecture	**Engineered and scalable IP backbone**
3. Mesh Networking	**Interoperable peer-to-peer connections** **Function blocks at the device level** **Battery life is deterministic**
4. Reliability	**Wireless transmission is deterministic** **Wireless transmission is received** **Wireless transmission is accurate**
5. Security	**Wireless transmission has not been hacked**

Table 4: System requirements for wireless control

For the first requirement, we've already discussed sampling rates and latency. For control, the user needs high confidence that messages will be transmitted on time. We list 1-2 second latency as a baseline target for control, to support human-machine interaction. The ISA100 standard was intended to support reporting four times per second, or possibly more, in meticulously designed installations.

Second, if a network can't be architected to meet the user's performance requirements, the rest of this list doesn't matter. Whenever control is involved, an ISA100.11a network should be engineered as an extension of a mature and scalable IP backbone. It is self-evident that future-oriented technology needs to fully leverage the plant's – and society's – current and future investment in mission-critical IP technology. This propensity – to leverage IP technology – is clear in the ISA100.11a standard and in actual products.

Third, true peer-to-peer networking is essential for cases where it is impossible to arrange proximity to an IP backbone, but where high performance machine-to-machine connections are required nonetheless. For time-critical control applications in such configurations, it might not be acceptable to wait for messages to travel all the way to a controller and then all the way back again wirelessly. For those network configurations, control needs to involve a simple and direct wireless connection between remote devices. These connections need to keep running even if mesh network operation is temporarily interrupted. Peer-to-peer support also implies that control logic runs in the devices themselves, using approaches such as function blocks that have been established in the wired world.

Notice that "*deterministic*" battery life is listed as a mesh networking requirement. Devices used for control tend to consume a disproportionate amount of shared network communication and energy resources. It is essential that these devices operate for a well-defined lifetime without maintenance, and without unnecessary network caveats. Just as importantly, the existence of these control devices should not create system-wide maintenance problems by running down the batteries of other devices on the network or by consuming unnecessary network resources. As will be shown in the network architecture section of this paper, designers of ISA100 were careful to eliminate unnecessary transmission of control messages, by applying techniques to ensure that each message will arrive at its destination expeditiously and directly.

Fourth and fifth, the reliability and security requirements are listed separately. The actual technical solutions are best integrated as they are in ISA100.11a. Especially in control, deterministic on-time message delivery is essential. Wireless communication has an unavoidable statistical component, but it is essential that every effort be made to rig the odds favorably, using deterministic techniques. And of course, high confidence is essential that data, when received, has not been garbled or hacked.

ISA100.11a Network Architecture for Control

Figure 7 is an annotated version of a basic reference diagram found in the ISA100.11a-2011 standard.

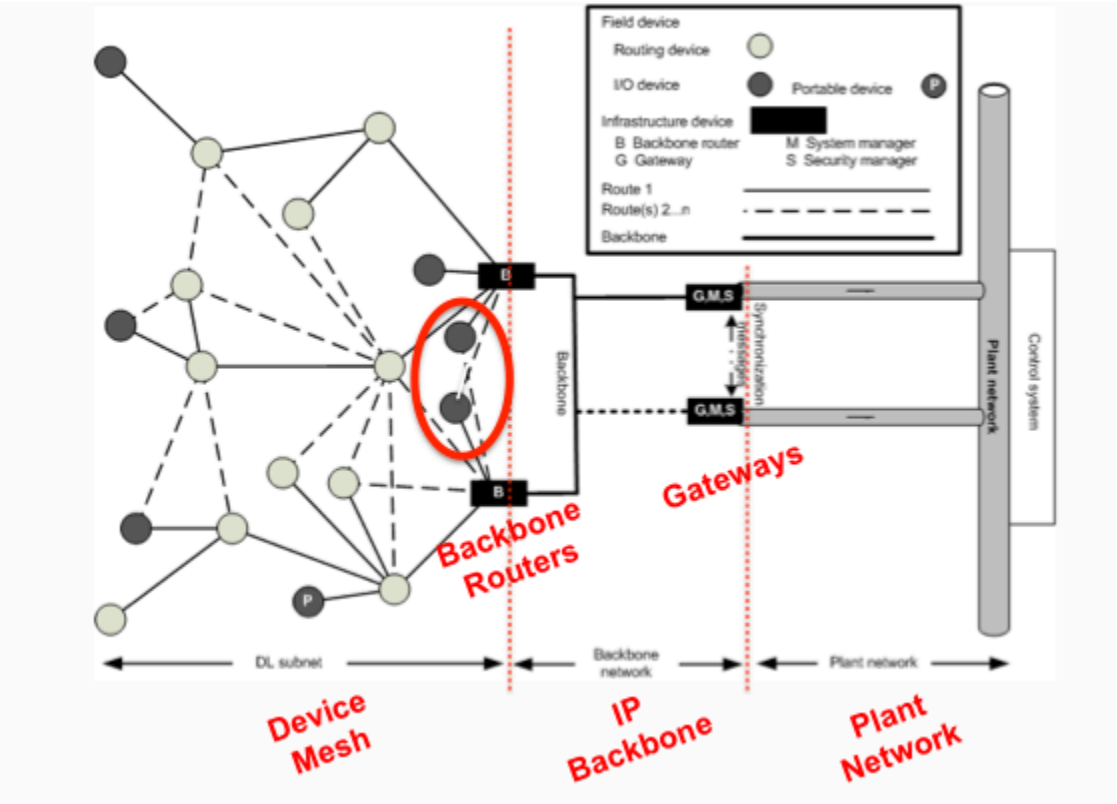

Figure 7: ISA100.11a Network Architecture

The plant network domain, with a control system, is illustrated to the right. The plant network is shown as connected to redundant ISA100 gateways.

The IP Backbone domain is illustrated in the middle. The backbone provides coherent network services to the plant, through access points that are called backbone routers. The backbone might itself be wireless, such as a high performance WiFi mesh.

The Device Mesh domain is shown on the left. This is a mesh network of battery-powered field devices. Performance of the device's connection to plant network tends to degrade as devices are located farther to the left of this diagram, away from the backbone.

In Figure 7, notice that the ISA100.11a standard emphasizes two I/O devices, that are circled in the figure, with redundant connections to multiple backbone routers. ISA100 included various optimizations targeted at connections in that configuration. The mesh network operates behind those devices (to the left in Figure 7), delivering a reasonable level of service but at lower

priority. The circled control devices receive special attention within the system architecture and priority in operation.

The circled control devices employ an ISA100.11a feature known as duocast, which is an optimization method whereby one message is received simultaneously by multiple backbone routers and then acknowledged by each in turn. This provides a very high probability that a message will be received by the backbone on the first transmission. Combined with duocast retries, a wireless ISA100.11a control device can be specifically configured for low latency and high reliability.

Another fundamental ISA100 optimization is the architecture of the network itself. The ISA100 network standard was designed to be a plant-wide resource, not dozens of little networks with hundreds of edge conditions. If the network is configured for control, control devices should be able find at least one direct connection the backbone, and preferably two or more. This allows for controlled latency and reliability, with predicable and consistent battery life. Such connections allow for reporting rates to be on the order of 1-5 seconds, without significantly degrading network performance or battery life of neighboring devices.

Figure 8 represents some of same general concepts as Figure 7, but in a graphic that focuses attention more on control aspects.

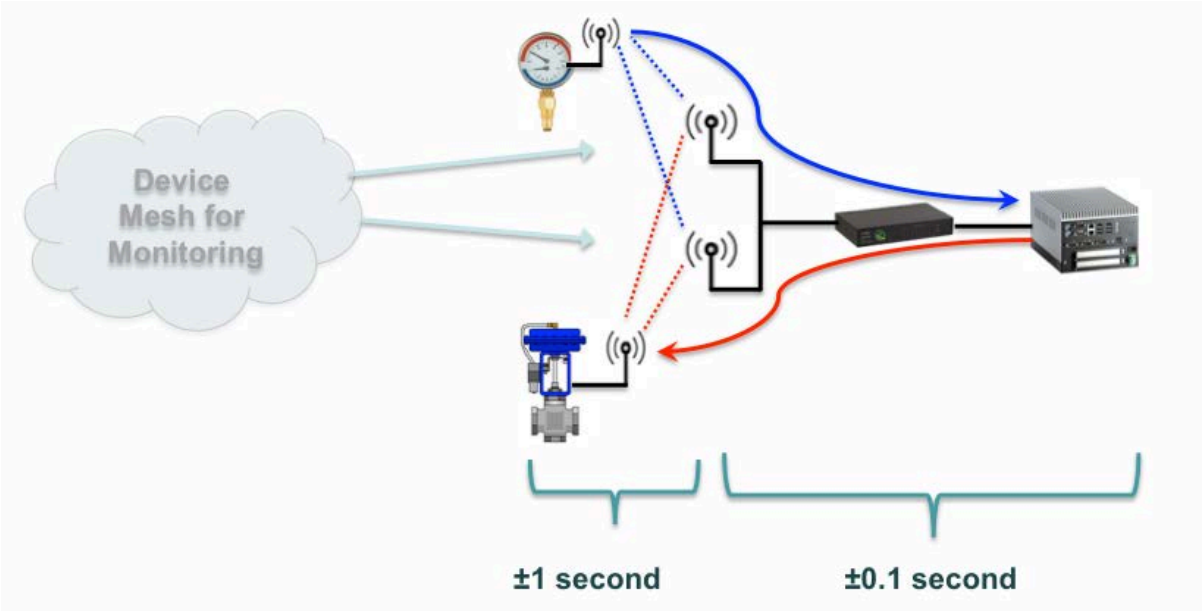

Figure 8: ISA100.11a Network Architecture

In this example, a pressure gauge feeds data to an ISA100 backbone, through a gateway and to a controller. The controller then sends data to a positioner, using the same network. We show both control devices with redundant wireless connections directly to the backbone.

The order-of-magnitude latency is shown at the bottom of the diagram. As mentioned earlier, the backbone network is high-speed and should run relatively quickly, at a fraction of a second. The direct wireless connection to the backbone may involve some significant latency, shown here as 1 second. That latency is intended to include cycle time for the device to assemble the data and then find a workable transmission opportunity to the backbone.

Figure 8 also shows a cloud to the left, indicating that the network is also used for monitoring, at the same time as control. It is important to recognize that real-life installations will have both control and monitoring working simultaneously in the same network. It is essential that control have priority over monitoring. The monitoring function shouldn't interfere with the control function if the network is properly configured for control.

In contrast, Figure 9 shows broadly what happens when the mesh network is involved in control applications.

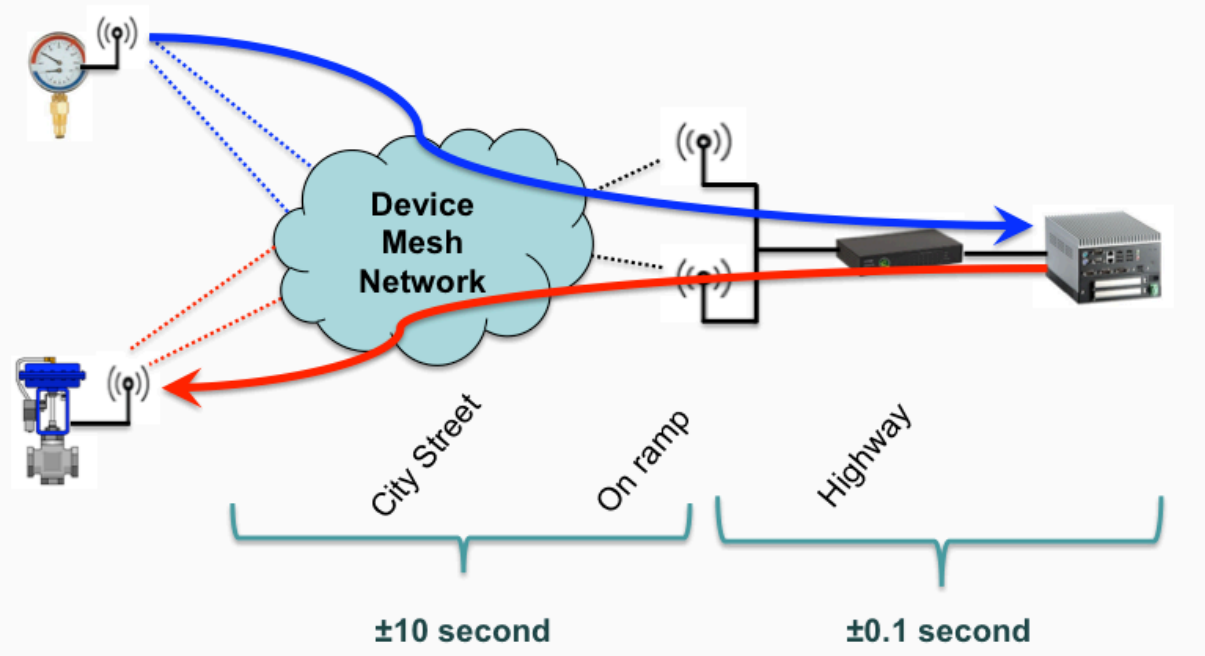

Figure 9: Wireless Control Through the Mesh

Here, control messages need to propagate *all* the way through the mesh to get to the controller, and then propagate *all* the way back along essentially the same path. This complete communication cycle can take a long time, for example 10 seconds in each direction. It is theoretically possible to align all of the planets and shorten that latency, but most battery-powered device mesh networks of this kind are not designed to complete this round trip in anything resembling the previously mentioned target of 1 second, as indicated in Figure 8.

As indicated in Figure 9, the network can be compared to a system of roads. The high speed IP network on the backbone is like a highway. The connection from the device mesh to the backbone is like the on-ramp, with a traffic light that operates during rush hour. In this analogy, the device mesh is like city streets, with a traffic light at each corner. Obviously, we need to avoid "city streets" in order to meet the demands for industrial process control.

There are also good reasons to *avoid* transmitting control data frequently through a battery-powered mesh network. Energy consumption is typically the most important restriction. Devices on the mesh need to sleep most of the time to conserve their batteries. For that reason, we show the performance of the device mesh being in the range of 10 seconds, mostly due to the need to preserve battery power in the routing devices. As a rough rule of thumb, if a mesh needs faster publication rate than about 5 to 10 seconds, then extra energy should be added to the mesh repeaters somehow, such as by providing external DC power to repeaters that carry control messages.

To summarize the core concepts of Figure 9, if an application's control data can be in an unknown state for something like ten seconds, then perhaps it can use a device mesh for control. ISA100.11a fully supports that configuration, as well as the case where approximately 1-second latency is required.

But what if the control application on the mesh network requires fresh data, say every 2 seconds? That is, how can an ISA100.11a solution address the case where a combination of mesh and fast data is required? One solution is peer-to-peer networking, as shown in Figure 10.

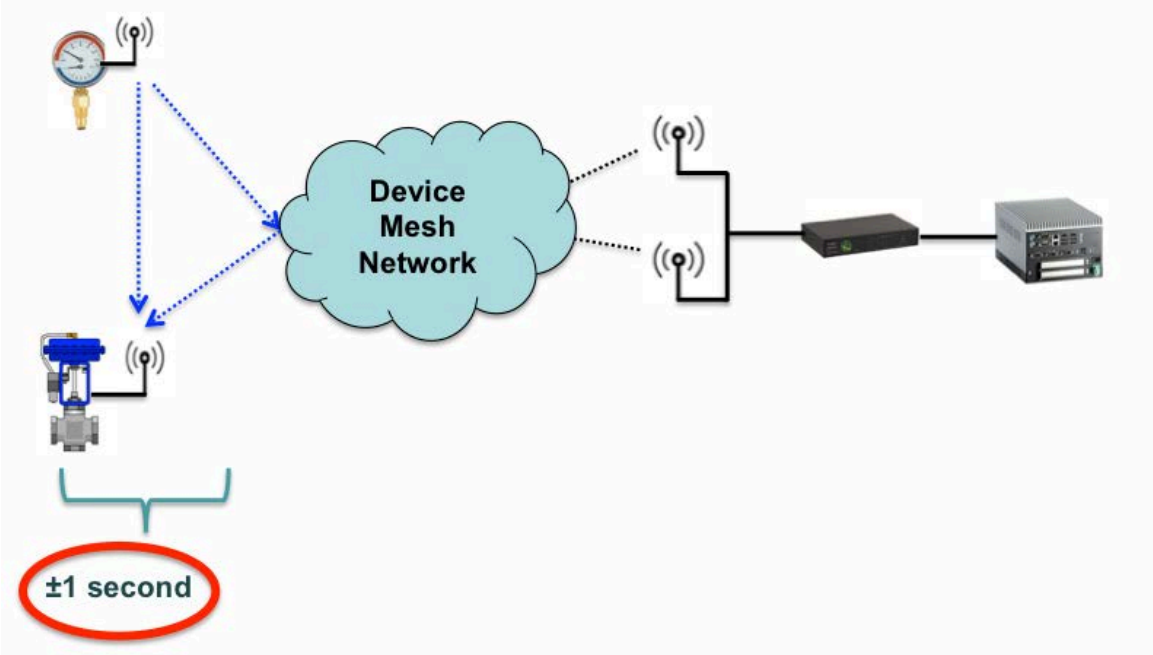

Figure 10: Wireless Control In the Mesh

The input block and output block might both be remote from a controller, but they are typically in physical proximity to each other. In many cases, they can establish a direct wireless connection that is strong and reliable. It is also typically possible to arrange a redundant path through a single mesh hop as shown in the figure. The result can be both efficient and effective. One advantage of this approach is that the connection and control algorithms keep working even if there is an interruption in service in remote components. The ISA100 standard fully supports peer-to-peer connections, including an object model that allows for control logic running in the devices themselves.

Conclusion

When industrial users consider applications for wireless control, they generally cite three key benefits: improved reliability, improved control, and cost savings. These benefits can be achieved using the ISA100.11a architecture, meeting reasonable user expectations for a broad range of control applications. One network can handle monitoring and control applications simultaneously, by giving critical control messages communication priority. A high speed IP backbone, which may itself be wireless, ensures that messages are quickly transmitted to and from controllers. For applications located in remote locations in a plant, true peer-to-peer networking allows for control loops without involving controllers that may take many seconds to access. The result is a system architecture that can coherently deliver mission-critical performance in a useful variety of real-world configurations.

References

1. ISA-100.11a-2011, An ISA Standard, *Wireless systems for industrial automation: Process control and related applications*, available electronically at www.isa.org, DOWN7967.

2. Verhamme, Ignace., *Wireless Control with ISA100.11a*, Honeywell Process Solutions, white paper available online.