cbtnuggets.com

IoT Protocols You Need to Know: ISA100.11a

by Landon D. Foster

9–11분

ISA100.11a is a IoT protocol used by many IoT devices and is most likely to be found in industrial settings, like petroleum refineries and manufacturing plants. As you might understand from its name, ISA100.11a was developed and promoted by the Industrial Society of Automation (ISA). It provides a platform that can be built upon to supplement some controls. It also presents a reliable method of gathering data that is less expensive and maintenance intensive than traditional wired communications.

IoT is supposed to make our lives easier, or more accessible; If you use Alexa or Google Home devices, you have interacted with IoT devices. Chances are, even if you don't have one of these hubs, you have some manner of IoT devices in your home — even if it's just your utility meters. Let's cover some basic IoT protocols.

How Does ISA100.1A Work?

ISA100.11a is based on a PHY (Physical Layer) and MAC layer (1 and 2 of the OSI model) called 802.15.4. You're likely at least passingly familiar with this standard, as it also governs Bluetooth and BLE. 802.15.4 covers a lot of spectrum, including some of the 900MHz band, but is most famously in the 2.4 GHz

band, where it crosses over WiFi.

When your phone is connected to your headphones, or speaker, you're using this PHY. 802.15.4 is considered a LR-PAN or Low Rate Personal Area Network type of PHY. What makes ISA100.11a so useful is that it takes advantage of the lower rate structure to allow it to have a longer reach. ISA100.11a, however, doesn't use the same channel structure as some of you may be familiar with for WiFi. It has 16 channels in the same space, and uses a smaller channel width as well. You can see the channels commonly used in this band below.

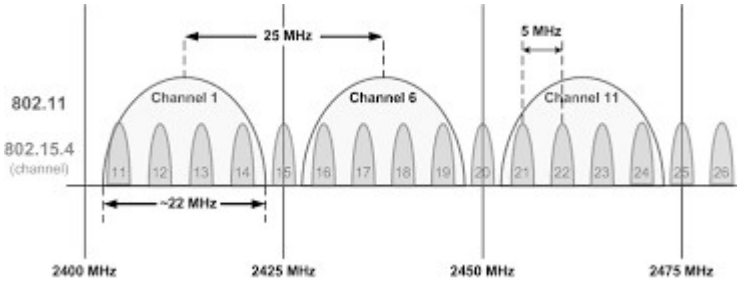

Key Features of 802.15.4 PHY

We're going to take a deep dive into the unique and important parts of the protocol now. This isn't an exhaustive list, of course. We don't have space here to go into all of it, and it gets deep pretty quickly.

TDMA

ISA100.11a uses a version of TDMA (Time Division Multiple Access), a standard method of contention free, or more deterministic access to the medium. In TDMA, all the clocks on the devices that make up the network are rigorously synchronized, and the timing is used to provide slots for the individual devices to transmit without getting in each other's way through collisions.

For example, slots are allocated to device A- Device A is set to report to the system every 60 seconds. It would be allocated the number of slots (10 ms duration) that it needs to send its data on a rotating basis. These slots are "reserved" from the total amount of slots available to be used, but only as many times as it needs. In contrast, a device that needs to send twice the data, or reports twice as often would generally be allocated twice the slots. Time slots can also be allocated dynamically by a network manager or analogous module.

It should be pointed out that this is also used in conjunction in both with frequency hopping, and the two differ by the domain that they vary in. TDMA varies the time domain, when the transmission is sent, whereas DSSS varies the frequency space that the transmission is sent on.

Superframes

A superframe is a unique feature to these types of networks: a collection of timeslots repeating on a cyclic schedule. More plainly, superframs are sets of timeslots that repeat according to a pattern. The number of timeslots included in a particular superframe determines its length, and therefore, how often it repeats, and logically extended, how often a device that uses this superframe may communicate.

Superframes can be conceptualized as a method of organizing communication through "links". Each time slot in a given superframe is dedicated to a specific link, such as device A communicating with Device B. Conversely, from Device B's point of view it is scheduled to listen for device A in the same time period. Multiple superframes can overlap in a network in the same time period. Remember, each device has to be viewed as its own entity, so this does not mean that a device is listening

and talking at the same time.

A device can also participate in multiple superframes, and participation in a superframe is not considered compulsory. As mentioned before, the period of a superframe is dependent on the length of the superframe and is how often that particular superframe repeats. How often the superframe repeats is inversely proportional to the length of the superframe. (EG If the Superframe is Length 2, it will repeat at rate 1. If it is length 1, it will repeat at rate 2, twice as often, because it is half as long.) Readers who are familiar with traditional 802.11 WLANs will notice that this is a stark contrast to the pseudo-random access that WLANs enjoy through the CSMA/CA in the order that is imposed on the network. One should recall that the focus of these networks is not arbitrary user access, but rather instrumentation and controls responding (often) in cycles as long as once a day.

A key learning for operators of these networks is this: shorter superframes mean lower latency and higher bandwidth, but are more battery intensive to the devices themselves and the entire superframe need not be filled with transmissions. Superframes can be filled out with null timeslots to artificially lengthen the superframe and increase the amount of time between individual transmission cycles.

Network Layer Peculiarities

The network layer in ISA100.11a uses the 6LoWPAN model. At the network level, most of the

the unique character of ISA100.11a drops off, and gives way to 6LoWPAN. Addresses assigned to devices are IPv6 analogues, and allows up to 128-Bit addresses to be used. At the high end,

ISA100.11a addressing means that the network is nearly infinite in scalability and eliminates the possibility of duplicate addressing, even in hyper-dense environments. This also means that ISA100.11a allows for connectivity between different devices, even across a linked routing backbone. Backbone routing complicates the process for ISA100.11a and provides a different Schema for different situations.

The network layer in ISA100.11a is also notable for having an intrinsic power-saving mechanism built in.

Using fully 128-bit addresses isn't always practical in a small network. Small here is considered to be less than 1000 devices, compared with a WAN network or similar. As transmitting 128-bits for every address is power consuming and not an efficient use of airtime, the standard allows for 16-bit aliases to be used over the local subnet. This can be conceptualized as the way we use names in the social world — within your family you might use a shorter version of a name, which is generally understood.

Your uncle Rich, for example, within your family is known simply as Rich. In the wider world he is known by his full name which would be cumbersome and unwieldy to use in common conversation but may be required to identify him among a much larger group like the general public. In this example, Lord Rich Richington Pendleton the Third, or similar. Imagine having to use full names every time you address someone instead of pronouns! The system manager assigns 16-bit aliases and 128-bit full addresses in an ISA100.11a network. Each device maintains an address table with these name associations, which is

constructed during the initial join process. When a transmission traverses the network backbone, the network layer of the

backbone or the device performing the "Gate" function of the subnet handles this translation from 16 to 128 bit and the inverse.

Routing

Routing in a ISA100.11a network is handled at the data link layer in most cases, but can be

handled by the network layer in some cases as mentioned previously. Inside a discrete subnet, between devices, the traffic is routed at Layer 2 through the data link layer. However, when the traffic passes a backbone router, or crosses into another subnet, it's both translated into the 128-bit address and handled by the network layer by necessity.

It's notable that this is not considered a problem because this is no longer a WIRELESS medium but instead being transmitted over wireline. Power conservation on wireline is not considered a priority. A third schema is present when the data is moved over a gateway, the formal boundary of the WSN (wireless sensor network). The entire data payload must be de- and re-encapsulated as it's being actively reinterpreted to fit within the plant network and being made ready to be used with the controls or sensor interpretation networks. This combination of Layer2 and Layer 3 routing is referred to as Mesh-Under and Routing-Over.

Final Thoughts

That's a short but deep preview of ISA100.11a. It's much more than this, but anything less than a week class may not be fair to the protocol. That said, the above highlights some major differences between it and other protocols you might run into

more often like WiFI. It is the author's strong conviction that IoT will only grow and that network engineers must be prepared for this growth, or be left behind.