



# **LECTURE 6: AZURE IOT HUB AND THE SENSOR LIFECYCLE**

**Ken Birman**  
**Spring, 2020**

# NEXT “ON BAT”... A 2-LECTURE SEQUENCE

We will be looking at IoT again, but this time with the goal of thinking about how DHT sharding interacts with IoT applications with sensors.


It will stretch over 2 lectures, so you won't see the whole story today.

By the end we'll have learned how big data systems compute on sharded data in a DHT, and then looked at how that approach could also work for large-scale collection of IoT data from sensors.

# THE FUTURE SMART WORLD

## Puzzles...

- Who “builds” this world?
- Will they maintain it properly?
- Can the devices be trusted?



<http://www.plmconnections.com>

# LONG-STANDING ISSUE WITH SENSORS

People have talked about using sensors to create a “smart world” since 1980’s, but it hasn’t been as simple as they imagined!

It is fairly easy to put RFID tags on devices, but those are passive.

In fact “full fledged” IoT with sophisticated sensors and actuators poses a wide range of challenges that we are only starting to appreciate.

# IOT IS EVERYWHERE, BUT POORLY MANAGED

Your Internet router, and networked printer

Cortana/Alexa/Siri/Google Nest

Your TV and home entertainment system

The network-connected microwave, fridge, range.

Smart hot-water heater, and A/C, and room heating units

Smart power meter, to connect them all together

Smart water meter (might even be able to diagnose leaks)

Solar panels on the roof, energy storage batteries in the wall



# IoT DOESN'T NEED TO BE OBVIOUS!



Estel: Italian design firm specializing in smart offices

The technology is subtle but pervasive. Dozens of smart devices

# ... EXAMPLES OF IOT IN THE OFFICE

Room occupancy, temperature, humidity sensors and sector control

Sensor to detect exterior light, actuator to control lights & window shades

Desktop microphone for conferencing

Smart copier/scanner with network-enabled functionality

The elevator system

The espresso machine that automatically orders new coffee packs

Door locks that check ID cards

# ... EVEN THE ELECTRIC POWER GRID IS SMART

Most of the world's bulk electric power systems are becoming smart

This is IoT on a “grand scale” and covers more than just power: coal/gas delivery, scheduling of power plants, maybe even water delivery, too.

But this means that the power grid will need to keep a close eye on everything using electric power, or generating it. More IoT!



# ... THE LIST REALLY IS ENDLESS

Smart farm

Smart city

Smart highway

Smart emergency first-response....



# MORE PUZZLES: CONTEXTUALIZATION

How do IoT devices know which room they are in?

- Alexa, adjust the shades to block the glare on my display
- Siri, use active noise cancellation to block that street noise
- Cortana, find me a nearby conference room we can book for an hour.

... in addition to the IoT devices themselves we will need increasingly detailed “environmental maps” for everything, down to individual rooms!

# ... CONTEXTUALIZATION IS A CLOUD COMPUTING TASK WE MIGHT DO IN A MICROSERVICE!

We have a database of knowledge about the residents of a home

We also have IoT sensors and knowledge of where they are located.

Occupant says “Hey Cortana, can you block this glare?”

You do a query to figure out that the sunlight from a particular window in the TV area is reflecting onto the occupant, who is on the couch there...

# WHO KEEPS THIS STUFF SECURE + ROBUST?

Even if every light bulb “could” have a computer in it, why would this benefit anyone, and who would make sure the broken ones are replaced?

How can we protect privacy and ensure that these things are secure?  
What costs could be incurred for violations?

What if a sensor malfunctions? Can we figure out that it needs repair?

# SITUATION TODAY?

Very poorly managed, huge numbers of IoT devices yet very little attention to software upgrades, network security issues raised.

There are network-enabled printers that turned out to have entire spy computing systems embedded in them, to retain copies of everything.

Largest “zombie/bot” population? By one estimate, it may be Internet Wifi routers with default password settings!



# AZURE IOT TODAY: AIMING FOR A MINIMAL BUT ADEQUATE LAUNCH POINT.

Microsoft has focused on IoT for corporate customers with huge numbers of smart devices, and little control over them.

And within that first step, they focus on management of the “fleet” of sensor and actuator devices:

- Unmanaged sensors are a danger and a nightmare to the “owner”
- Seems like a necessary first step, in any case
- Can we “secure” the IoT devices, and make them “trustworthy”?

# KEY ARCHITECTURAL ELEMENTS?

Microsoft product: Azure IoT Hub, IoT Edge and Intelligent Edge

- First, the hub handles secure registration of devices and status tracking
- Next, it automates software upgrades
- It deals with issues of intermittent connectivity
- For devices that can be controlled from the cloud, it creates a “model” to enable you to perform those control actions

# AZURE IOT: DATABASE OF SENSORS.

The first step centers on secure registration of devices. The Azure IoT Hub manages a scalable database of sensors and associated data.

The enterprise owner also records information such as:

- Device make and model,
- Software revision level, battery lifetime, when it was last serviced
- Where it is located, role it plays (information for contextualization)
- Additional application-specific information or “knowledge”




# THIS MAKES IT AN ACTIVE DATABASE!

In a normal database the data tuples are just plain old data objects.

In Azure IoT Hub, the objects in the database are intended to hold meta-data on behalf of the “real” sensors and actuators, and to keep this data in sync with the actual sensor or actuator *when connectivity is possible*.

In effect, we now have meta-data describing the sensor combined with live properties (like battery level, photos cached, filter settings) that are wired to the actual device and change in real-time!

# DEVICE SECURITY



Azure Sphere Development Kit

The level of security for today's network-enabled IoT devices is poor to non-existent, making them way too easy to hack or disable.

So Microsoft has a new product aimed at sensor *manufacturers*. The Azure Sphere is a special low-power security chip that embodies a hardware root of trust and low-power cryptographically protected HTTPS.


With Azure Sphere, device manufactures can secure existing sensor products, and the resulting sensors will interoperate with Azure IoT hub.

# IOT PROGRAMMING VIA TEMPLATES

Like much of the cloud, Azure IoT offers “recipes” that developers download and then customize.

Here is an example from a scenario that they “story-board” on the Azure IoT Hub website.

This one relates to smart manufacturing



# EVERY AZURE IOT DEVICE HAS A “PROXY”

Many devices have limited network connectivity and won't always be online.

So in Azure IoT Hub, every device has a cloud-hosted “representative”: a software agent that can respond to device operations 24x7, and then will push updates (like new software revisions) when an opportunity arises.

The agent can also schedule maintenance operations.

# PROXY PROGRAMMING

In this Azure proxy mode, you can send information to a device even if the device is currently disconnected! The proxy is always available.

For example, a firmware update or patch, or new device configuration.

But obviously the action can't occur until the device connects. So there is always a back-and-forth: Event "to" the device, and later, an event "back". Applications will need to work in this very asynchronous way.

# QUALITY OF SERVICE ISSUES

For many devices, network quality is also an issue.

Over time, a network link might be unavailable, or available but slow, or temporarily very fast (at a high price).

In normal networks we don't think about this much. But for IoT, our applications may need to be dynamically responsive as conditions change.

# THE QUALITY OF THE SENSOR ITSELF IS ALSO A SERIOUS CONCERN

With large numbers of sensors we often get redundancy

Wood and Marzullo explored this in a system called Meta here at Cornell.

Key idea: first, that sensors have “range of accuracy”, and second that time also has a range of accuracy. Finally, that by leveraging this insight, we can actually identify and correct for many kinds of inaccuracy!

# CLOCKS ALSO HAVE LIMITED ACCURACY

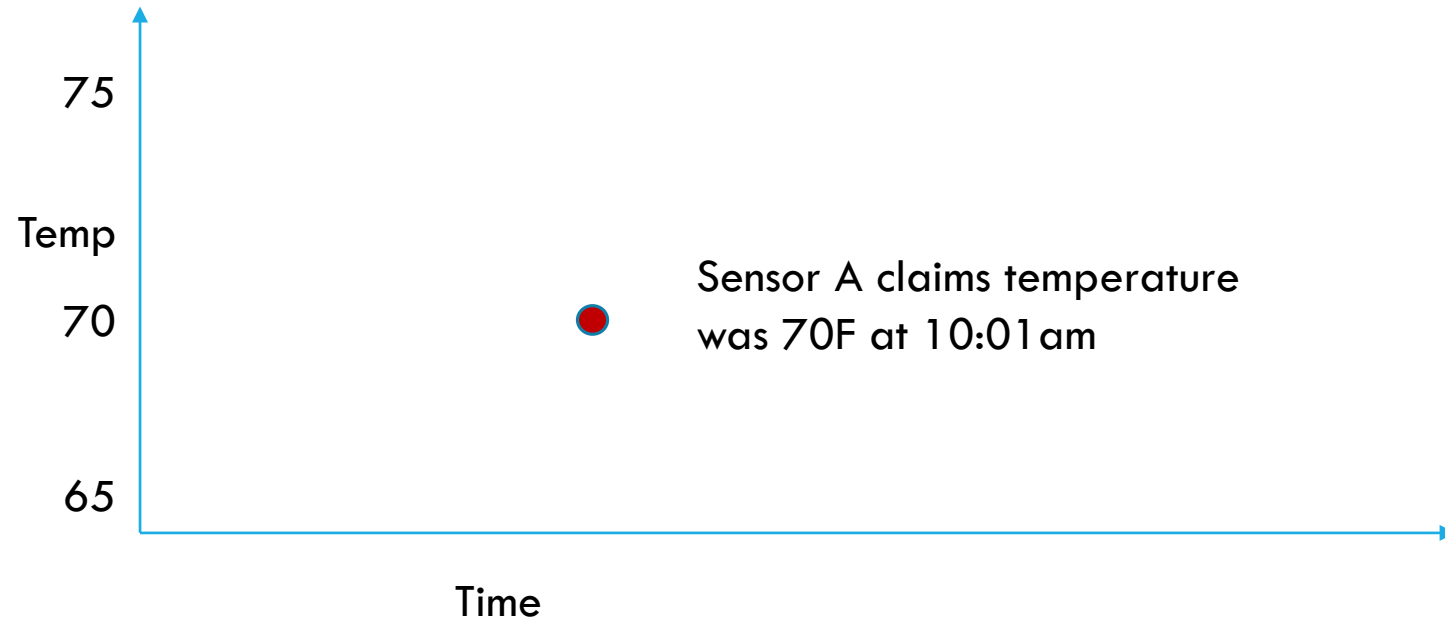
There are standard protocols for synchronizing them.

But at any instant in time, clocks could still have drifted from the exact time.

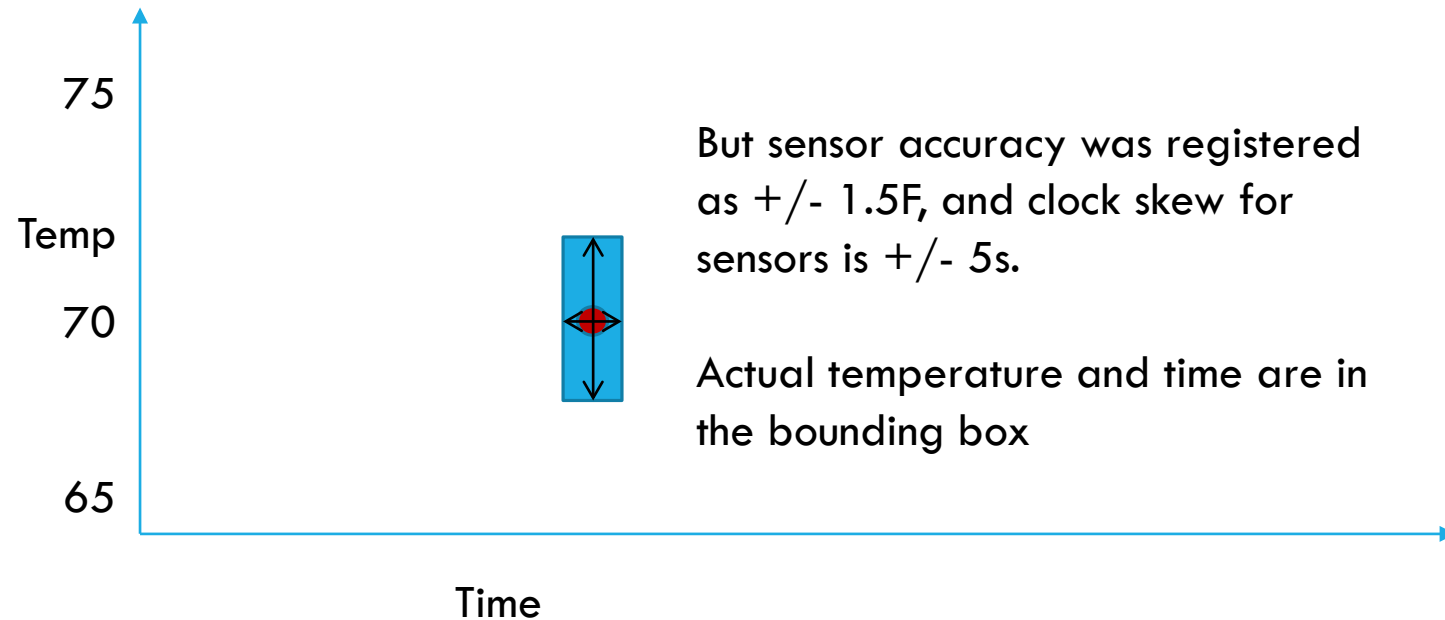
The term clock *skew* is an estimate of the worst-case drift between two non-faulty clocks, if you were to read them at the same instant in time.




# DEEP DIVE: HOW DOES SENSOR ACCURACY IMPACT THE WAY WE MIGHT USE THEM?




# DEEP DIVE: HOW DOES SENSOR ACCURACY IMPACT THE WAY WE MIGHT USE THEM?




# DEEP DIVE: HOW DOES SENSOR ACCURACY IMPACT THE WAY WE MIGHT USE THEM?



# DEEP DIVE: HOW DOES SENSOR ACCURACY IMPACT THE WAY WE MIGHT USE THEM?



# DEEP DIVE: HOW DOES SENSOR ACCURACY IMPACT THE WAY WE MIGHT USE THEM?



Knowledge of temperature *trends* could give us a further way to improve the data! Also, by now we can see that we need to schedule service on the yellow sensor, or remove it entirely.

# META: “RULES” FOR COMPUTING WITH IMPRECISE SENSORS

Suppose we tell the smart system to “turn on the office A/C if the temperature rises above 71F”

Does this mean “if the temperature might be more than 71?” Or “if the temperature is definitely more than 71?”

The correct action is very different! In our example, the temperature is probably 70.1, but *could* be as high as 71.6

# “MIGHT” VERSUS “DEFINITELY”

***Might*** is a way to talk about “any value in the bounding box.”

- If the temperature is  $70.1 \pm 1.5$  at time  $10:03 \pm 5s$ , it “might” have any value in the temperature range, or the time range.

***Definitely*** would mean that the allowed range has “no intersection” with the bounding box

- In our example, the temperature is definitely more than 69F

# TEMPORAL ACCURACY CONCERNS

Fred Schneider likes to distinguish three properties:

- Precision: How close is a clock to other non-faulty clocks?
- Accuracy: How close is the clock to “true” time?
- Drift: How quickly is a clock losing precision or accuracy?

Recall: For a pair of clocks, the “skew” is their difference if measured at the same instant in true time. Drift is the “rate at which skew increases.”

Precision and accuracy are *both* forms of skew.



# PATRIOT MISSILE DEFENSE IN THE IRAQ WAR

The Patriot anti-missile system was implemented as a two-processor system.

A radar unit, on the ground, computed the trajectory of an incoming threat and tells the anti-missile when and where to intercept the target. Then the anti-missile launches and makes a pinpoint rendezvous.

Would you prefer for the two clocks (the radar, and the defensive missile) to have the best possible precision, or accuracy?

# WHAT ACTUALLY HAPPENED?



Iraq launched Scud missiles at various targets (“flying garbage cans filled with dynamite (and maybe worse)”).

Israel fired on dozens of Scuds. Time after time, the Patriot missiles kept missing: they reached the right spot, but just a few milliseconds late!

The anti-missile would punch through the fuselage of the Scud, but miss the warhead, and then both would fall to earth... and explode on the ground



# WHY WAS THE PATRIOT LATE?

Later they discovered that the clocks in the radar computer system were drifting very slightly – they were slightly fast.

After 3 or 5 days, this might add up to a few milliseconds of inaccuracy.

So when the radar said to the patriot: “intercept a thread that will be at location  $(x,y,z)$  at time  $T$ , the time “ $T$ ” was a few milliseconds late! Since then, the issue has been fixed (IronDome works really well).

# OTHER SIMILAR CASES THAT COULD ARISE

A smart highway tells two cars to follow some pattern: car A can enter the passing lane and pass car B. Car B should stay to the right until car A has passed by. If this is given by times, the same question arises.

Or two drones scanning some field, taking commands from a control system.

In Fred's terminology: We need a high degree of precision for temporal coordination in such cases, but this is complicated by the tendency to use absolute time of day when computers interact with one-another.

# MORE PUZZLES: “INFORMATION FLOW”



Because the IoT world will be large and geographically distributed, information needs to be moved around to carry out intelligent tasks. Yet customers may insist that we limit and control the flow of private data.

# PRIVACY



Suppose that the cloud learns sensitive information

- If Siri/Alexa/Cortana can hear your commands, it can listen to you.
- Who owns that information it gathered?
- Who should be permitted to listen to it?

It would be nice to say “all my private stuff stays in my house” but only the cloud really has the resources to understand everyone at large scale

# SECURITY



Beyond personal privacy, companies developing new products are at risk of competitors spying on them!

They benefit from the technology, but if their IP gets stolen, they lose it all!

So there are questions of protection not of private data, but also of company IP. But if you own your private data, how will the system figure out which remarks are private and which are related to your job?

# FAULT-TOLERANCE

With so many moving parts, some will fail!



In fact failures can occur at every level

- A sensor could go offline
- It could stay online but send confused/incorrect data, or bad timestamps
- The network connection could freeze up or break
- The cloud servers could shut down or be reconfigured



# EXAMPLE: LANDLORD'S MANAGEMENT TOOL

To see how these can come into conflict, consider a solution aimed at the landlord of a large residential complex.

It might have hundreds of units (apartments or condos)

Landlord needs to operate this intelligently

# AVOIDING HARM TO A BUILDING



In fact a major goal for smart buildings is to anticipate problems and fix them, perhaps even before they become serious.

Think about an apartment complex with a forgetful tenant in a drought area. If he starts the shower running but forgets and goes shopping, huge water waste can occur. Someone should shut the water off!

But then there is a privacy issue raised: will this reveal that the tenant has early senile dementia? He could be harmed if this became known.

# SIMILAR ISSUES

Window left open and yet heat is on.

Stove was left on, or lights, or TV... but tenants aren't home. Microwave was set to 30 minutes instead of 30 seconds.

Tenant falls and can't get to a phone or call for help.

Fire in the area. First-responders need to know where people are located.

# IOT CAN HELP!

The challenge is to build trustworthy solutions that won't violate privacy and yet will help in cases like these.

The answer is to create systems that safeguard your data by keeping it close to where they gather it (which is why Azure IoT Edge is designed to even live right in the home).

When data does get sent to the cloud, it should be for a specific task. Then, when finished with that task, “leave no trace behind”.

# SOCIETAL BALANCING ACT!



Those of us who will build these systems need to decide what to do! And worse, we can't even trust the sensors themselves!

The law is far behind the technology curve. Don't expect to find answers in courtrooms for a long time.

Yet because companies ultimately have a revenue and profit motive, if you focus purely on corporate goals, you might do something evil.

# DOES EXISTING CLOUD ENCOUNTER THESE ISSUES?

A few of the same issues arise in any cloud-scale system.

So in fact our challenge isn't really to invent everything from scratch, and some of these policy issues are familiar ones..

The actual challenge is more subtle: we need to learn to “repurpose” what the cloud is already able to do, in these new roles and ways.

# BIG PART OF THE ISSUE



Leakage

At very large scale, the cloud “mixes” tasks on behalf of different users (“multi-tenancy”), for cost-efficiency reasons.

For example, AWS Dynamo and Azure Cosmos both hold data for huge numbers of applications. The same is true for the global file system.

Of course they keep track of which data is owned by which user, and you don’t normally “see” other people’s data. But information could *leak*.

# SUMMARY

Tackling the IoT opportunity will be a huge opportunity, but also a huge investment and an enormous amount of work.

Today's cloud is incredibly powerful and this existing powerful solution will make a big difference.

But getting from here to there will still require new technology development.