



nLine Inc.



## Watching the Grid:

### Utility-Independent Measurements of Electricity Reliability in Accra, Ghana

**Noah Klugman\***, Joshua Adkins\*, Emily Paszkiewicz (UC Berkeley / nLine Inc.);  
Molly Hickman (Virginia Tech / nLine Inc.);  
Matthew Podolsky (UC Berkeley);  
Jay Taneja (University of Massachusetts Amherst);  
Prabal Dutta (UC Berkeley / nLine Inc.)

IPSN' 2021  
May 21, 2021, virtual

My name is Noah Klugman. I'll be talking about our work watching the grid, utility-independent measurements of electricity reliability in Accra Ghana

Limited measurements make it hard to fix power quality issues, which are too common

Power quality issues are common around the world to the point where they are often holding back economic development. However, many utility are not able to sufficiently measure the performance of the grid making any performance issues hard to fix. Missing data on reliability impacts different stakeholders in different ways.

## Limited measurements make it hard to fix power quality issues, which are too common



### **Utilities**

Goal: to provide reliable electricity to customers

Challenge: don't know when or where the grid is failing, making it difficult to diagnose problems and perform prompt maintenance

For example, the lack of reliability measurements hurts utilities, who have the goal of providing reliable electricity to customers but don't know where or when the grid is failing, when measurements don't exist.

## Limited measurements make it hard to fix power quality issues, which are too common



### **Utilities**

Goal: to provide reliable electricity to customers

Challenge: don't know when or where the grid is failing, making it difficult to diagnose problems and perform prompt maintenance



### **Regulators**

Goal: to ensure utilities are operating fairly

Challenge: don't have adequate, utility-independent data on grid performance

Missing data on reliability impacts regulators. A primary goal of a regulator is to ensure that the utility is actually providing the quality of service that they are advertising to their customers. However, without good grid reliability data regulators are forced to depend on self reported reliability metrics from the utility, which, while not inherently bad, is a real conflict of interest.

# Limited measurements make it hard to fix power quality issues, which are too common



## Utilities

Goal: to provide reliable electricity to customers

Challenge: don't know when or where the grid is failing, making it difficult to diagnose problems and perform prompt maintenance



## Regulators

Goal: to ensure utilities are operating fairly

Challenge: don't have adequate, utility-independent data on grid performance



## Ratepayers

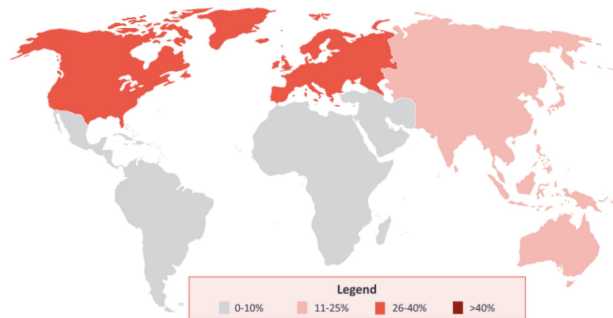
Goal: to pay the least for the highest quality electricity

Challenge: don't have historical data on quality of service in an area; have to be involved in grid maintenance by manually reporting failures

And a lack of grid reliability information is hard for ratepayers, who are just trying to pay the least they can for the best quality electricity they can access. However, because they don't have access to quality of service data, they may accidentally decide to invest in building a house or opening a business in an area where power is poor, a potentially disastrous economic decision.

Existing grid sensors are often prohibitively expensive, and are rarely deployed in many countries

Global Smart Meter (Energy and Water)  
Penetration by Region (2019)



Source data: IoT Analytics Research 2019 <https://iot-analytics.com/smart-meter-market-2019-global-penetration-reached-14-percent/>

This graph sort of blew my mind, here is the global smart meter penetration by region, from 2019, looking at both energy and water meters. Sorry for the confusing colors, here the regions in gray are actually the least covered, with only 0-10% penetration. So when you think of smart meters as a clear solution that has existed for many years, you really don't see that as the environment that much of the world is experiencing.

## Key claim

We can detect small and large power outages and power quality issues by aggregating measurements from networked sensors installed at outlets in households and businesses at the edge of the grid, enabling a utility-independent, high-resolution, and low-cost system well suited for deployment in under-instrumented areas.

This brings us to our key claim that we can detect small and large power outages and power quality issues by aggregating measurements from networked sensors installed at outlets in households and businesses at the edge of the grid, enabling a utility-independent, high-resolution, and low-cost system well suited for deployment in under-instrumented areas.

A single slide introduction to the grid

So before we jump in I want to give a single slide introduction to the grid

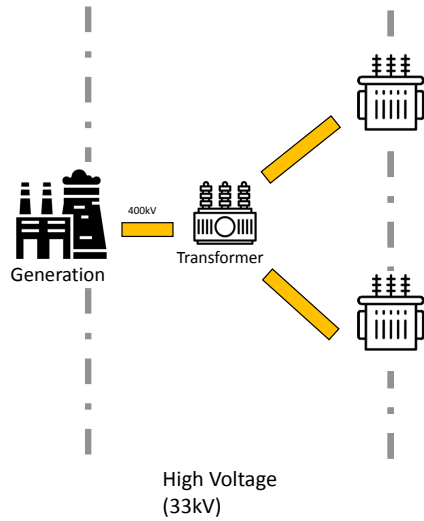


## A single slide introduction to the grid




We start with generation, these are the large power plants, from

## A single slide introduction to the grid



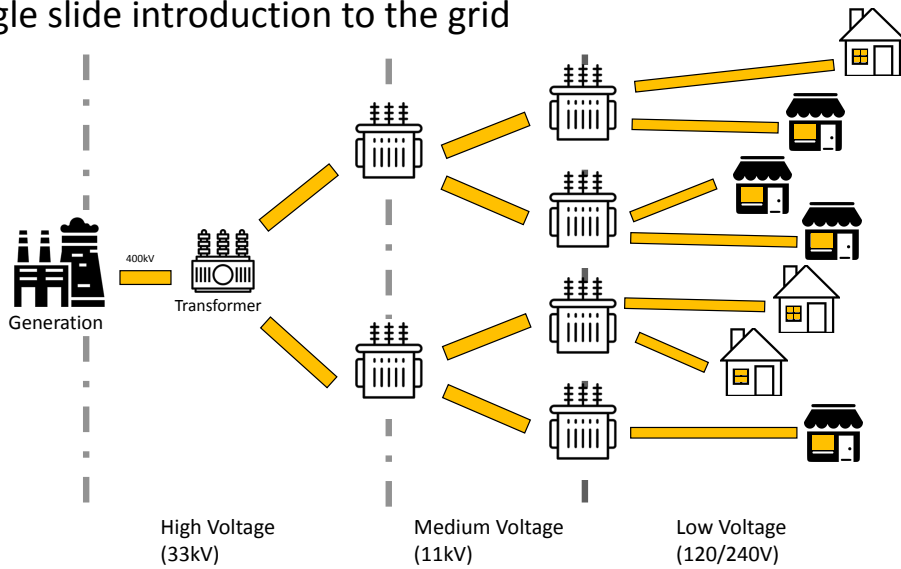
which emanate high voltage lines which are stepped down in the high voltage network to

# A single slide introduction to the grid




medium voltage lines, which are again

## A single slide introduction to the grid



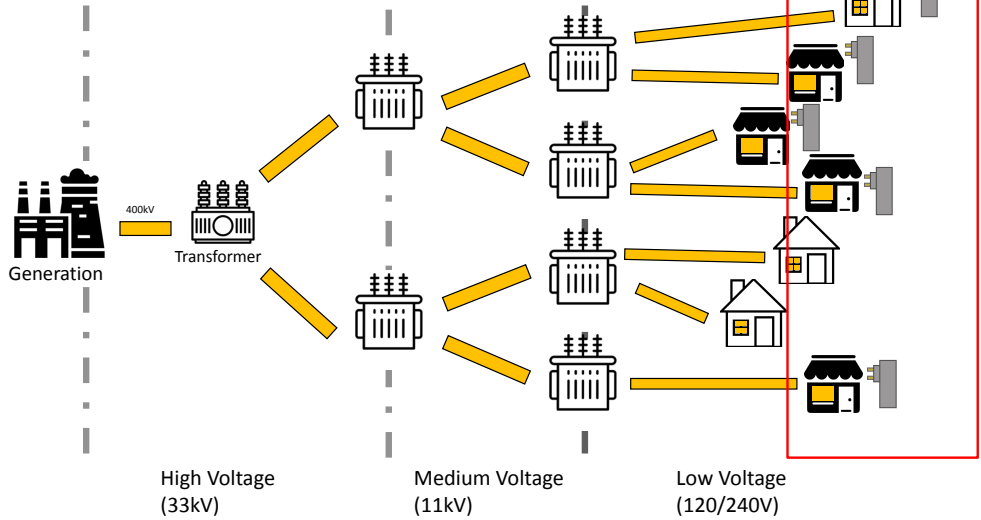
are stepped down in the low voltage part of the network which is where the lines enter customer households and businesses.

A high voltage outage occurs!



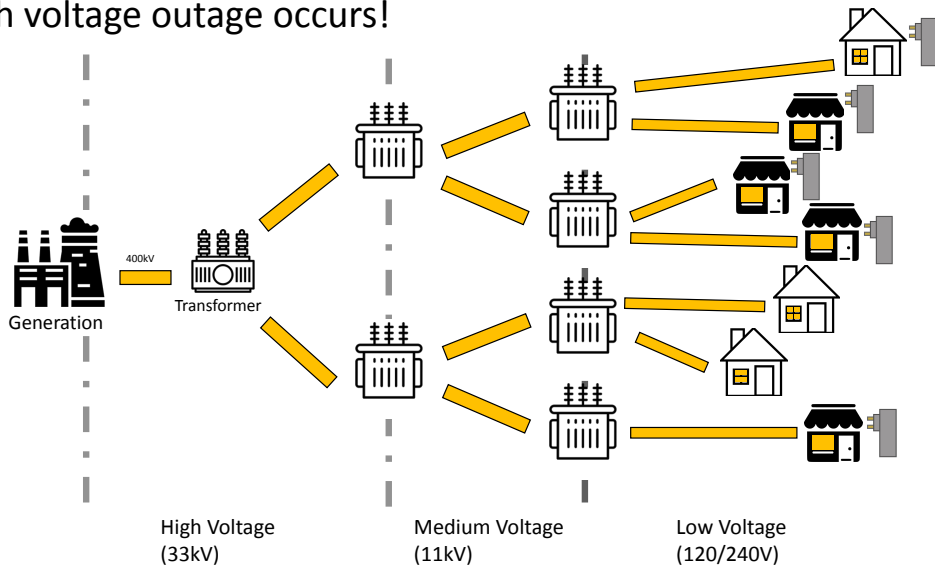
Lets talk through an example of how we may detect an outage. Lets say a high voltage outage occurs and You might have missed it but we

A high voltage outage occurs!




added some sensors to the low voltage part of the grid over here. These are our sensors that we call PowerWatch, they plug in at outlets in each of these structures. Anyways,

A high voltage outage occurs!



So when a high voltage outage occurs,

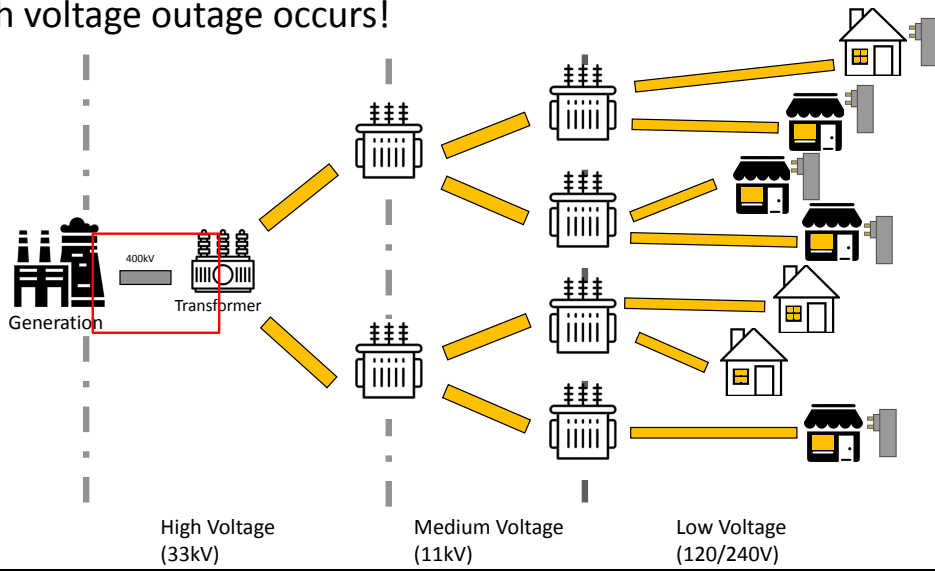
A high voltage outage occurs!



you see

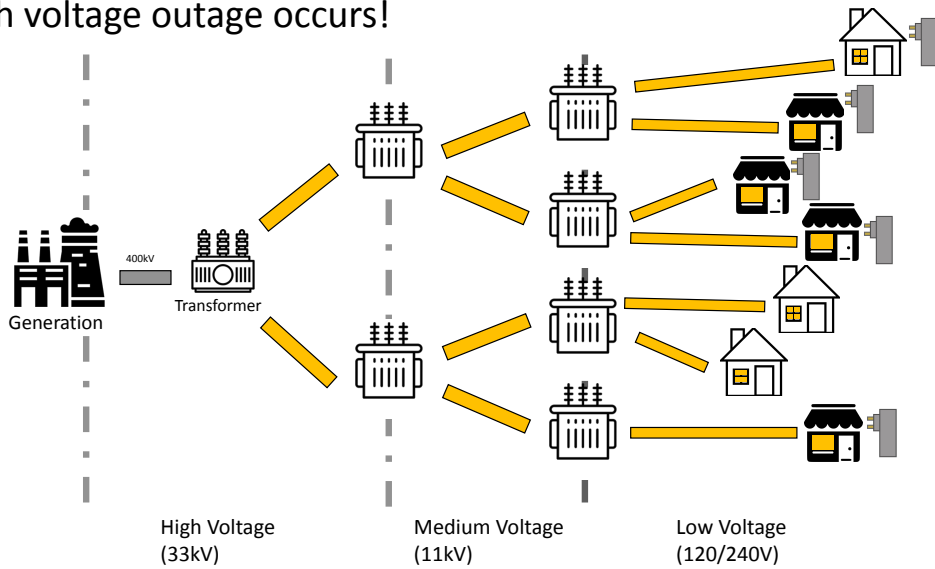


A high voltage outage occurs!



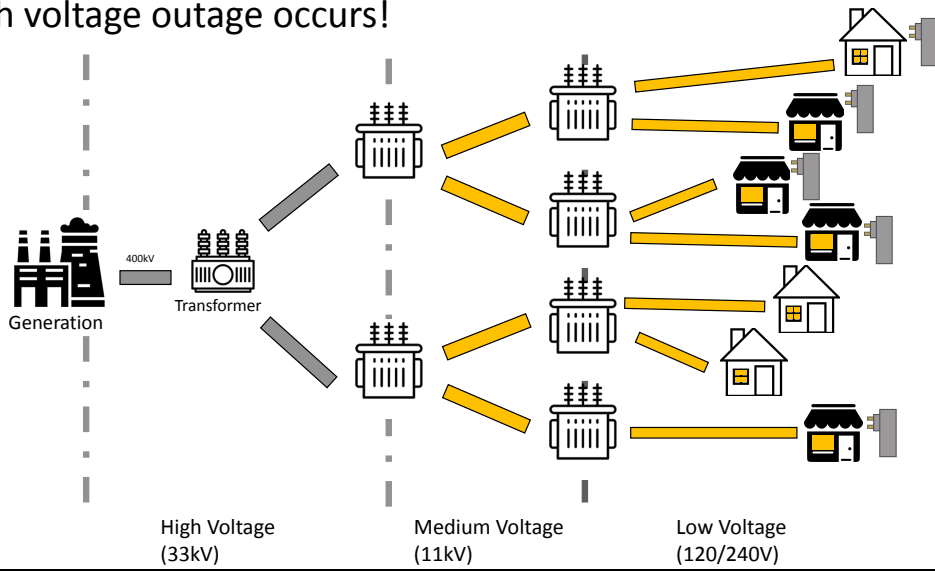
the line in the high voltage area past generation go dark,

A high voltage outage occurs!




and the consequence this has on the network is predictable right?

A high voltage outage occurs!



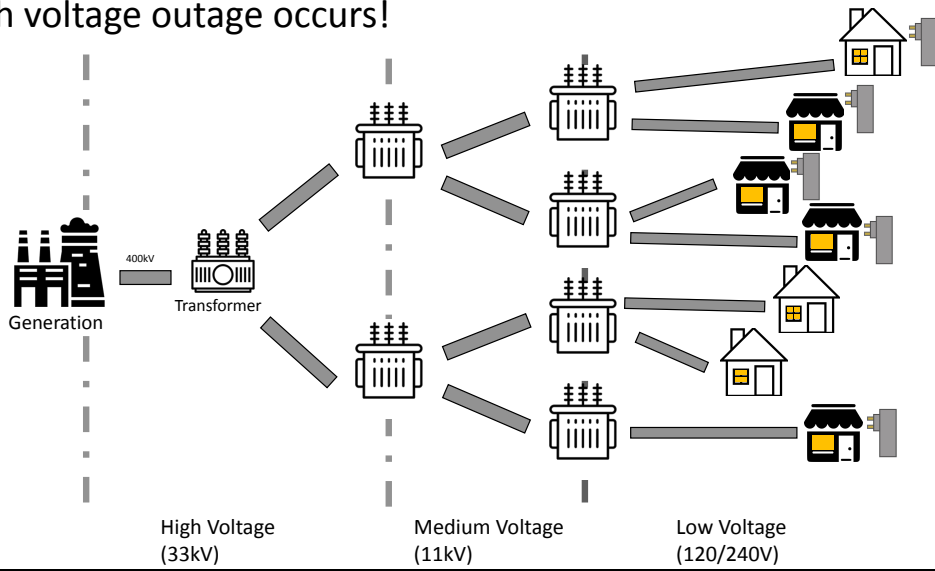
So this percolates down, so every single edge

A high voltage outage occurs!



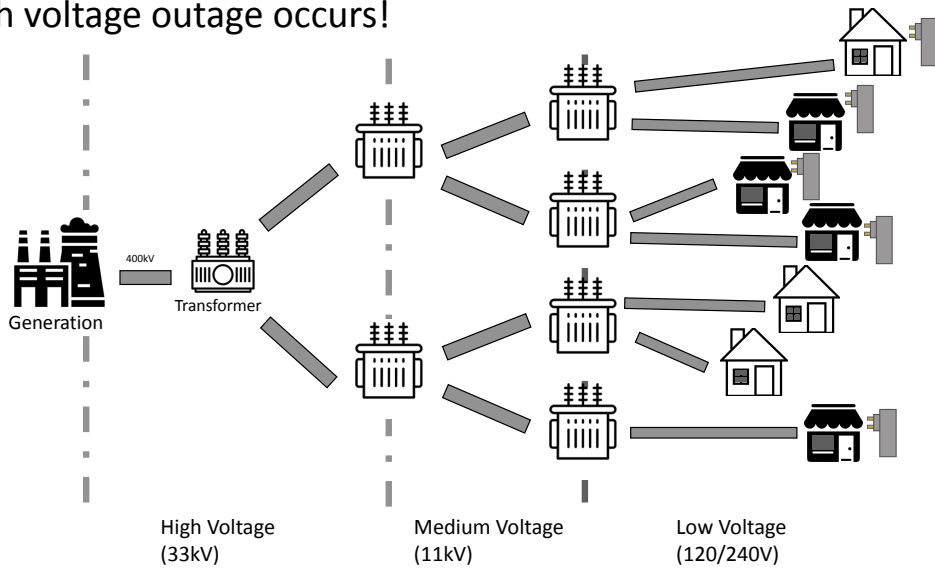
node will experience this

A high voltage outage occurs!



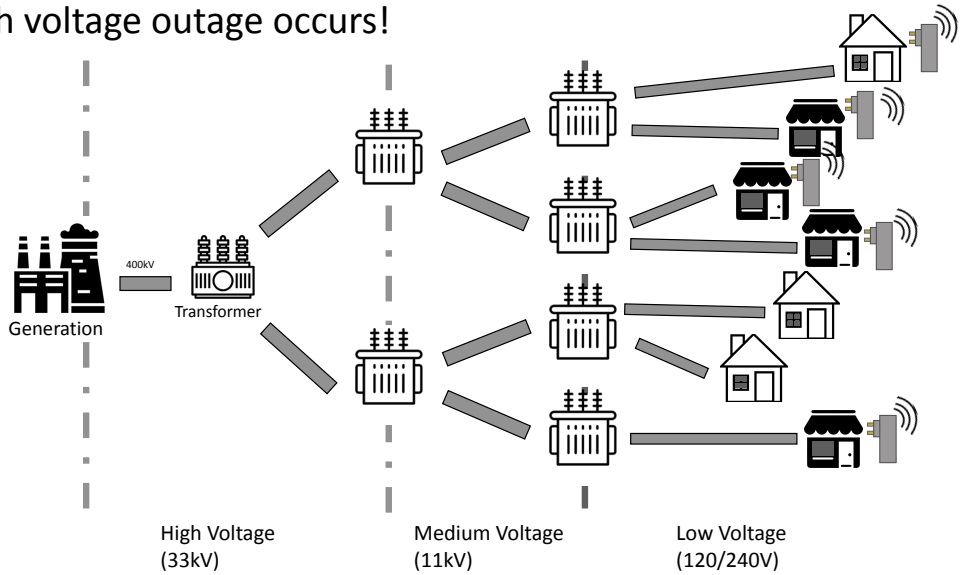
outage roughly together in time.

A high voltage outage occurs!




All going out together. If you think about our deployed PowerWatch sensors, what will happen is that each one of these will experience the outage,

A high voltage outage occurs!



or a loss of power, at the same time at each of the outlets they are installed at.


A high voltage outage occurs!



And what this will look like,




A high voltage outage occurs!




For our high voltage outage in the aggregate is each of these sensors will send the report that they lost power, and their location, and the timestamp,

A high voltage outage occurs!




through the cellular network which we use as our data backhaul in part because of its wide availability around the world,

A high voltage outage occurs!



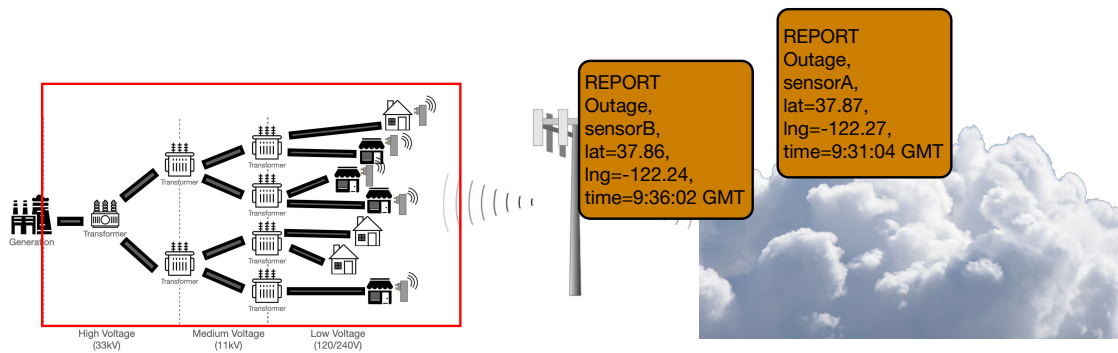
and this will reach our cloud. We will get this information

A high voltage outage occurs!



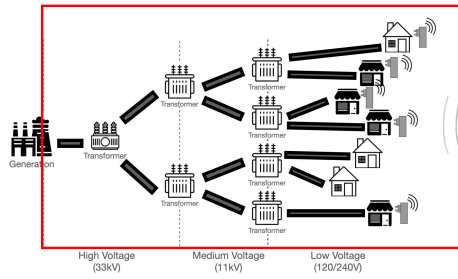
and what we do

# A high voltage outage occurs!



is that we actually

# A high voltage outage occurs!




REPORT  
Outage,  
sensorB,  
lat=37.86,  
lng=-122.24,  
time=9:36:02 GMT

REPORT  
Outage,  
sensorA,  
lat=37.87,  
lng=-122.27,  
time=9:31:04 GMT

REPORT  
Outage,  
sensorD,  
lat=37.82,  
lng=-122.25,  
time=9:31:12 GMT

look for

# A high voltage outage occurs!



REPORT  
Outage,  
sensorB,  
lat=37.86,  
lng=-122.24,  
time=9:36:02 GMT


REPORT  
Outage,  
sensorA,  
lat=37.87,  
lng=-122.27,  
time=9:31:04 GMT

REPORT  
Outage,  
sensorD,  
lat=37.82,  
lng=-122.25,  
time=9:31:12 GMT

REPORT  
Outage,  
sensorD,  
lat=37.92,  
lng=-122.25,  
time=9:33:12 GMT

a space-time cluster.

## A high voltage outage occurs!



REPORT  
Outage,  
sensorC,  
lat=37.84,  
lng=-122.24,  
time=9:31:10 GMT

REPORT  
Outage,  
sensorD,  
lat=37.92,  
lng=-122.25,  
time=9:33:12 GMT

REPORT  
Outage,  
sensorA,  
lat=37.87,  
lng=-122.27,  
time=9:31:04 GMT

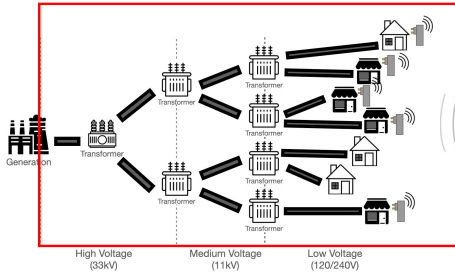
REPORT  
Outage,  
sensorB,  
lat=37.86,  
lng=-122.24,  
time=9:36:02 GMT

REPORT  
Outage,  
sensorD,  
lat=37.82,  
lng=-122.25,  
time=9:31:12 GMT

These sensors experienced the same outage, so the time bound should be pretty narrow, and they are all on the same network, so the space bound should be narrow as well.



# A high voltage outage occurs!



REPORT  
Outage,  
sensorC,  
lat=37.84,  
lng=-122.24,  
time=9:31:10 GMT

REPORT  
Outage,  
sensorD,  
lat=37.92,  
lng=-122.25,  
time=9:33:12 GMT

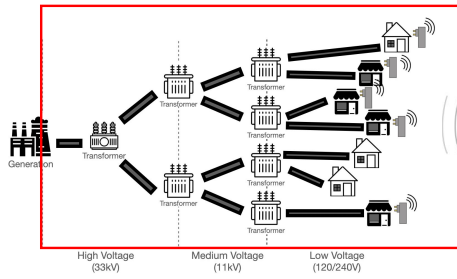
REPORT  
Outage,  
sensorA,  
lat=37.87,  
lng=-122.27,  
time=9:31:04 GMT

REPORT  
Outage,  
sensorB,  
lat=37.86,  
lng=-122.24,  
time=9:36:02 GMT

REPORT  
Outage,  
sensorD,  
lat=37.82,  
lng=-122.25,  
time=9:31:12 GMT

nd in this case, with our high voltage outage, we would see this cluster

## A high voltage outage occurs!



REPORT  
Outage,  
sensorC,  
lat=37.84,  
lng=-122.24,  
time=9:31:10 GMT

REPORT  
Outage,  
sensorD,  
lat=37.92,  
lng=-122.25,  
time=9:33:12 GMT

REPORT  
Outage,  
sensorA,  
lat=37.87,  
lng=-122.27,  
time=9:31:04 GMT

REPORT  
Outage,  
sensorB,  
lat=37.86,  
lng=-122.24,  
time=9:36:02 GMT

REPORT  
Outage,  
sensorD,  
lat=37.82,  
lng=-122.25,  
time=9:31:12 GMT

**Space-time cluster**  
**VALID OUTAGE**

And call it valid outage. The cluster gives us confidence that this was a real outage, not just sensors randomly reporting noise.

## Outline

1. Validating PowerWatch in the lab
2. Deploying PowerWatch at scale in Accra, Ghana
3. Evaluating data returned without ground truth
4. Early learnings about the grid

So that is the basic sensing methodology here. For the rest of the talk what we will do is talk through how we validated our sensor, called PowerWatch in the lab, our scaled deployment, how we evaluated our deployment quality without ground truth, because again we are in areas where this measurement is not being taken, and then finally we will touch on some early learning we can present on the grid in Accra.

## Outline

1. **Validating PowerWatch in the lab**
2. Deploying PowerWatch at scale in Accra, Ghana
3. Evaluating data returned without ground truth
4. Early learnings about the grid

Ok, so lets explore PowerWatch in the lab

## PowerWatch hardware and performance

- Power loss/restoration (~seconds accuracy)
- Grid Voltage (every ~2 minutes, ~3V error)
- Grid Period/Frequency (every ~2 minutes, 10 us precision)
- GPS Location
- 3G Cellular
- Battery backup



We were able to implement PowerWatch well within the bounds of our application accuracy. We timestamp power loss and restoration within seconds, syncing our time with an RTC and over network time. We sample grid voltage and grid frequency every two minutes and both of those are sampled within the accuracy requirements given to us by the folks using our data. Like I mentioned earlier, we are also taking GPS location, which helps with deployment management, using 3g cellular as our backhaul, and the PowerWatch sensor has a battery backup that lets us stay alive during extended power outages.

## PowerWatch: distributed sensors coupled with cloud-based analytics



The cloud is fairly straightforward. Each sensor is coupled with our analytics, first we go off to a postures and timescaledb database, and then we do our space-time clustering with apache spark which we are running from datanricks.

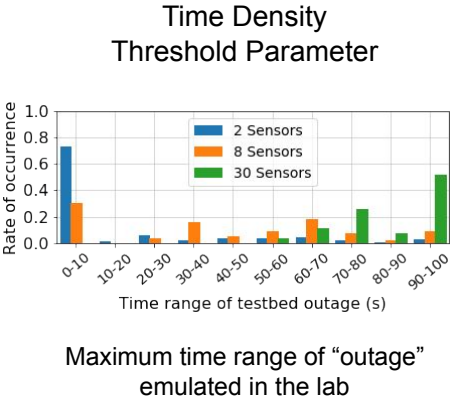
## Core data processing algorithms

PowerWatch uses STDBSCAN (Space-Time Density-Based clustering) to identify outages.

For our core data processing we use STDBSCAN which is a space-time density-based clustering algorithm that requires two different parameters both of which we learn, one in the lab and on in deployment.

# Core data processing algorithms

PowerWatch uses STDBSCAN (Space-Time Density-Based clustering) to identify outages.




The time density parameter is learned by simulating outages and looking at the largest time variance between sensors in the outage, which we find to be 100 seconds



# Core data processing algorithms


PowerWatch uses STDBSCAN (Space-Time Density-Based clustering) to identify outages.

Time Density  
Threshold Parameter



Maximum time range of "outage"  
emulated in the lab

Space Density  
Threshold Parameter



Max distance to nearest  
deployment site (excluding outliers)

We learn our space density threshold from our deployment by calculating the maximum distance between each of our deployment sites as the maximum distance that sensors could be in space.

## Outline

1. **Validating PowerWatch in the lab**
2. Deploying PowerWatch at scale in Accra, Ghana
3. Evaluating data returned (without ground truth)
4. Early learnings about the grid


Ok, so we presented our implementation and now

## Outline

1. Validating PowerWatch in the lab
2. **Deploying PowerWatch at scale in Accra, Ghana**
3. Evaluating data returned (without ground truth)
4. Early learnings about the grid

Will move on to how well PowerWatch performed in the field


## Design allows for measurement without relying on the utility for deployment or management



Stepping back, again, one of the unique things about PowerWatch is that it lets us to take these measurements without relying on a utility for deployment or management. This let us deploy in sites that we selected ourselves based on characteristics we were most interested in, namely areas where new transformers were scheduled to be injected in the area, and it was as easy as training a field staff member to knock on doors, find a participant, and the sensor was installed as soon as it was plugged into an outlet.

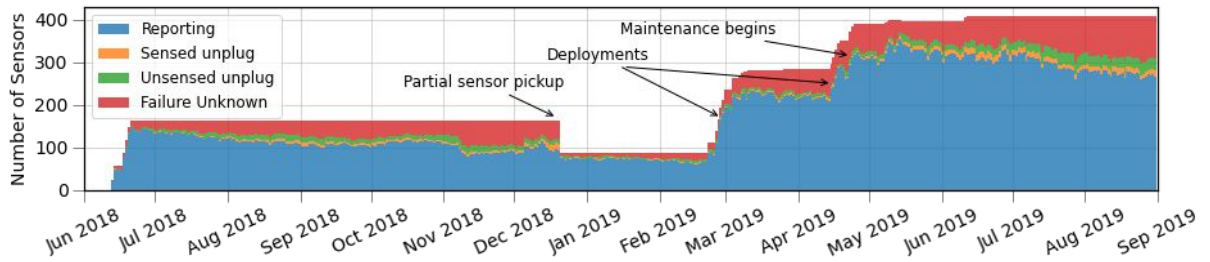
## We deploy 462 PowerWatch sensors in homes and businesses in Accra, Ghana for over a year

- Between June 2018 - December 2019 we enrolled in Accra:
  - 462 PowerWatch participants
  - Participants were clustered at 152 transformers



using this methodology, we were able to enroll 462 participants in 152 sites, in each site we placed 3 sensors.


## The sensor performs well enough in the field to produce a meaningful data stream



And the sensors performed well enough over this time period to produce a meaningful data stream. Of course we are expecting a certain number of failures. I think what is really interesting in this graph is the impact that field officers revisiting the sites and asking people to keep the sensors plugged in, and doing this daily maintenance and management work, really did increase the quality and number of sensors that were functioning as expected.

## Cellular network supported our deployment

- Mean DRR of 97.4%
- 95% of sensors have DRR over 95%
- Local queuing helps correct for cellular issues



We can also look at how well the cellular network supported our deployment, we had a very high data reception rate. We achieved this with some local queuing, so when the network was down we would keep data local and retry, but with that in mind we did achieve a 97 percent pdd, indicating that we collected the quantity of data we were expecting.

## Outline

1. Validating PowerWatch in the lab
2. **Deploying PowerWatch at scale in Accra, Ghana**
3. Evaluating data returned (without ground truth)
4. Early learnings about the grid

As interesting than whether the entire implementation of the deployment worked,



## Outline

1. Validating PowerWatch in the lab
2. Deploying PowerWatch at scale in Accra, Ghana
3. **Evaluating data returned (without ground truth)**
4. Early learnings about the grid

are the questions of whether the data collected is representative of the energy reliability of Accra. To get there we first have to ask ourselves is whether we

Question 1: Are we detecting outages?

are actually detecting outages or are we just seeing sort of random noise across a distributed population of sensors.

# We can sanity check our deployment with big outages

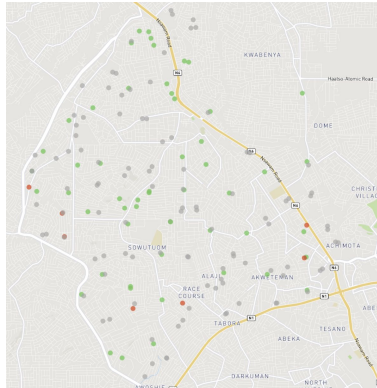
The screenshot shows the GhanaWeb website interface. At the top, the logo 'GhanaWeb' is visible along with social media icons and a search bar. Below the logo, a breadcrumb trail reads 'You are here: Home → News → 2019 03 13 → Article 729968'. A navigation menu includes categories like News, Sports, Business, Entertainment, World, and TV / Radio. The main content area features a sidebar with 'Latest News' and a list of categories. The article title is 'PDS explains cause of Tuesday night's widespread dumsor', dated 'Wednesday, 13 March 2019'. The article text explains that technical challenges at the Power Distribution Services Ghana (PDS) caused power fluctuations in Accra, leading to a 'dumsor' (blackout). It mentions that PDS recently took over electricity distribution from the ECG and that power supply to certain areas has been curtailed.

So the first easy way to check this is that we can sanity check against outages that appear on the news, as these as ground truth. So here is an example of an outage, which was a high voltage outage like we saw in my earlier example, that took place on March 13th, 2019.

## A sensed high-voltage outage

- Gray dots = Sensors not reporting anything
- Green dots = Sensors reporting power
- Red dots = Sensors reporting outage

No Outage



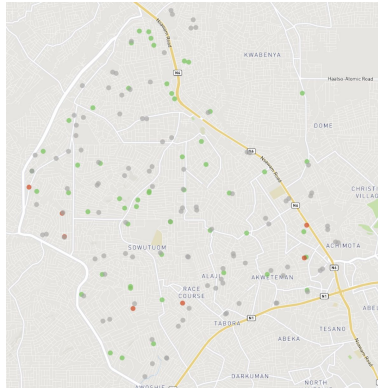
6:38 pm on March 12, 2019

Look at how this shows up in our system, here you see a snapshot of our system at 6:38 on March 12. The green dots are visualizing sensors that are reporting they have power.

## A sensed high-voltage outage

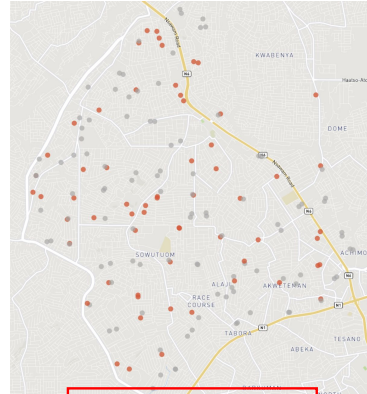
- Gray dots = Sensors not reporting anything
- Green dots = Sensors reporting power
- Red dots = Sensors reporting outage

No Outage



6:38 pm on March 12, 2019


High Voltage Outage



6:42 pm on March 12, 2019


You see all those green dots, just a couple minutes later, have turned to red dots, PowerWatch sensors reporting an outage, just a couple minutes later. This shows a large outage, which is expected from a high voltage outage. But obviously we are not going to be able to validate everything this way.

We find spatial and temporal patterns in the data unlikely to occur for reasons other than power outages or restorations



And that is because we see a ton of clusters. Again, the intuition is that people are not calling each other to decide to unplug these at the same time. So we have to use other methods to see the validity of these clusters.

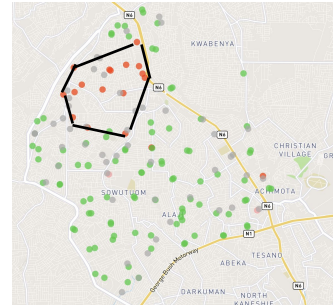
## Outages occur in dense spatial clusters as expected



One method is to ask whether outages occur in dense clusters as we would expect. This is a medium voltage outage here.

## Outages occur in dense spatial clusters as expected

- In an outage cluster very few powered sensors are within the convex hull of an outage
  - Holds regardless of the outage size




We find that in every outage cluster, very few powered sensors are within the convex hull of the outage. And this holds true regardless of the size of the outage.



# Outages occur in dense spatial clusters as expected

- In an outage cluster very few powered sensors are within the convex hull of an outage
  - Holds regardless of the outage size



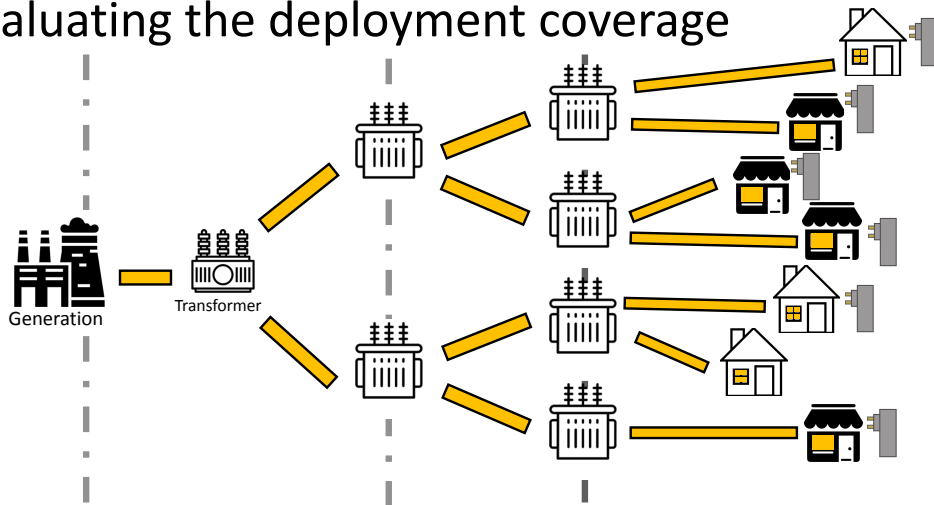
Outage Size	Mean # Powered Inner Sensors/Outage	Mean % Powered Inner Sensors/Outage	Max # Powered Inner Sensors/Outage	Max % Powered Inner Sensors/Outage
<b>3-10 Sensors</b>	0.03	0.33%	2	20%
<b>10-30 Sensors</b>	0.09	0.51%	2	11.76%
<b>30+ Sensors</b>	0.31	0.60%	2	4.65%

In this table you can see that outages even of 30+ sensors there was a maximum of 2 sensors powered in the convex hull. If there was random noise we wouldn't expect this relationship.

Question 2: How representative is our sample?

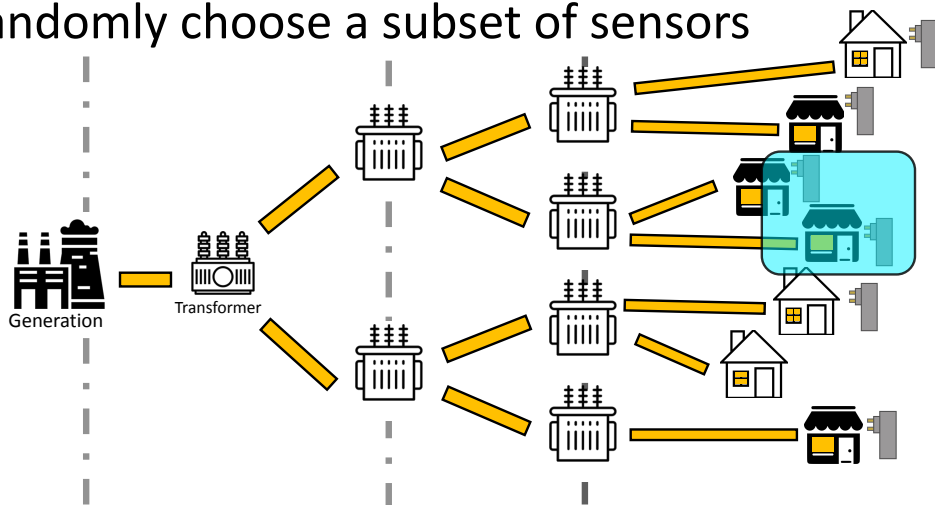
The second question we ask now that we are confident we are seeing outages is do we have enough sensors to accurately estimate reliability?

## Evaluating the deployment coverage



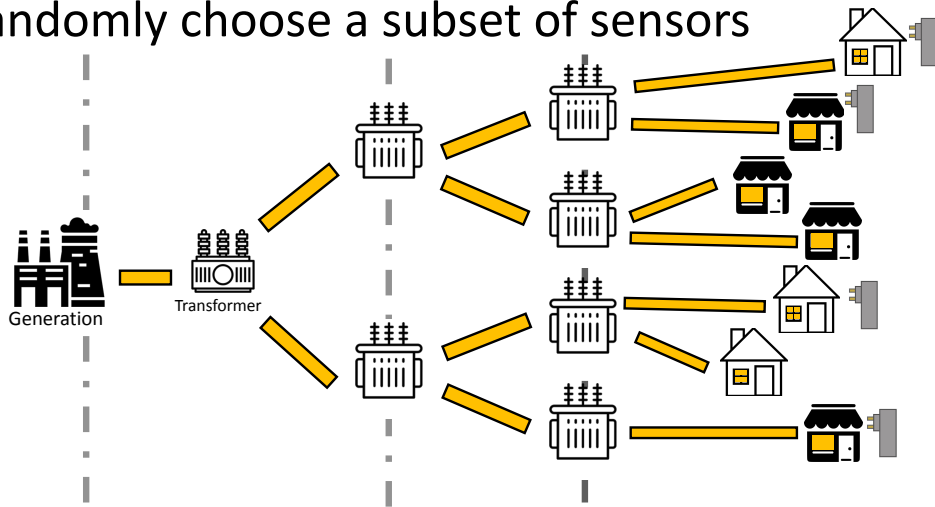
And one way that we could evaluate this when ground truth isn't present is that we can actually run an experiment, which we have been calling a drop out study, where we

Randomly choose a subset of sensors




randomly choose a subset of sensors to remove and then

## Randomly choose a subset of sensors




we play back the outages that we would have detected with just the sensor subset, and see whether if information is lost.

# Playback an outage with sensor subset



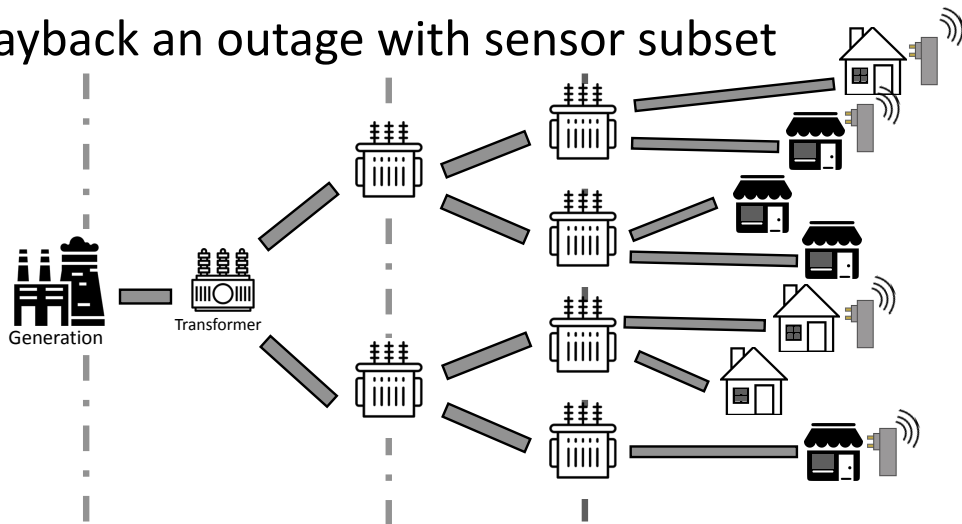
So here our outage occurs

# Playback an outage with sensor subset



And each sensor from our subset stops charging

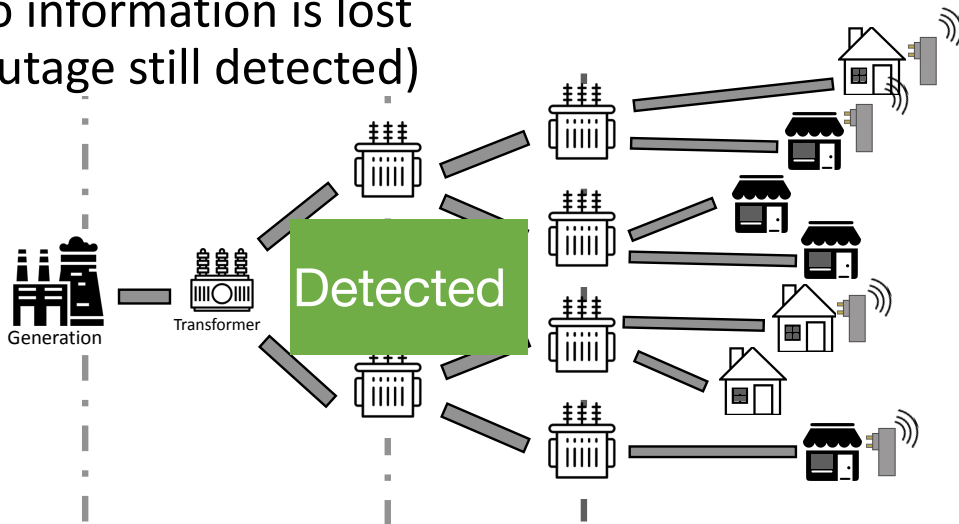
## Playback an outage with sensor subset



So here in this example, which is a high voltage outage, no information is lost. Every sensor detected this outage and a space time

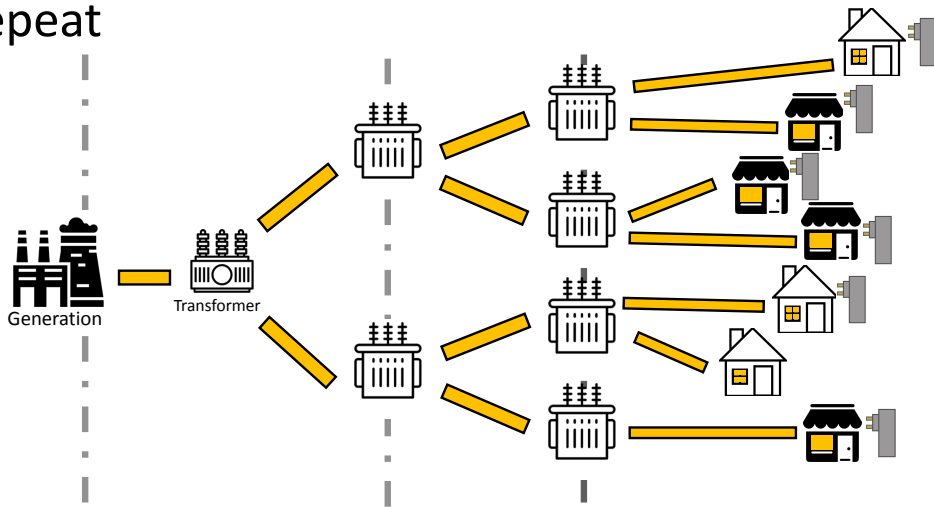


No information is lost  
(outage still detected)



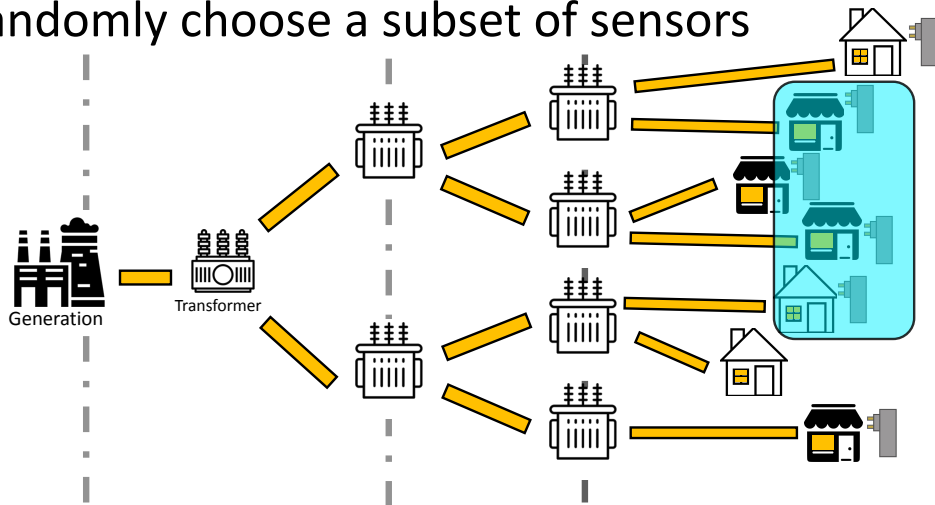
cluster could be formed, even with the smaller subset of sensors.

Repeat



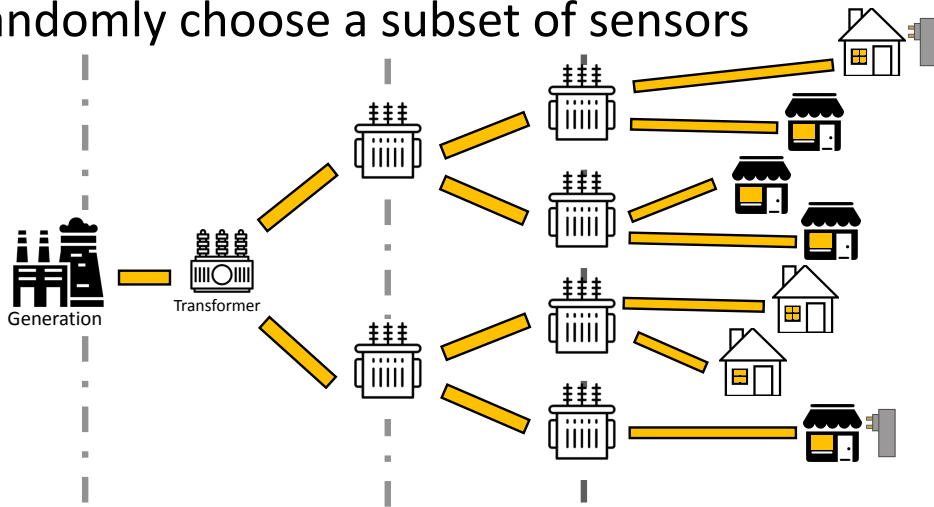
Eventually though, as we *drop out more and more sensors*, we will start to see instances where when we play back an outage, information is lost.

Randomly choose a subset of sensors




We randomly choose

Randomly choose a subset of sensors




a subset of sensors [to drop out]

# Playback an outage with sensor subset




Play back an outage. For example, this medium voltage outage is a bit different,

# Playback an outage with sensor subset




Happening here on the bottom

# Playback an outage with sensor subset



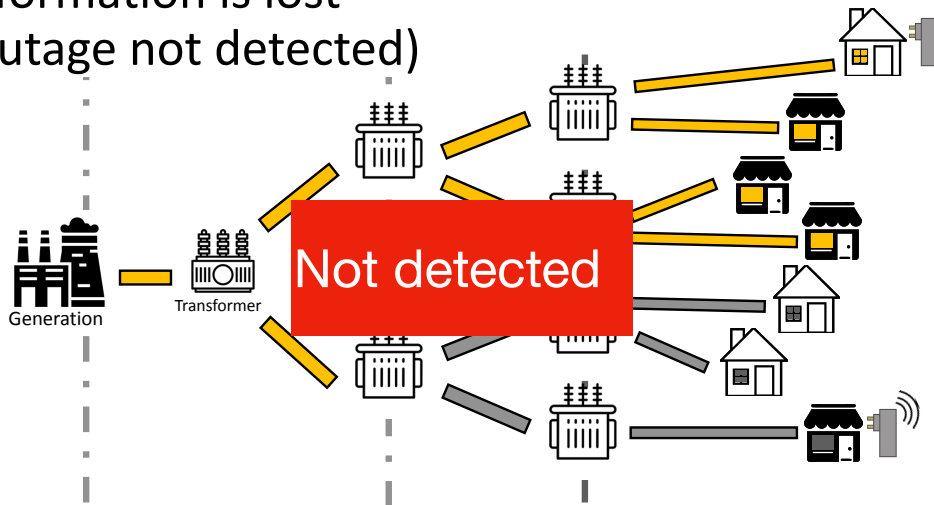
## Playback an outage with sensor subset



And now with our subset of sensors, for this outage only a single sensor will go out, where before two would have gone out.



Information is lost  
(outage not detected)

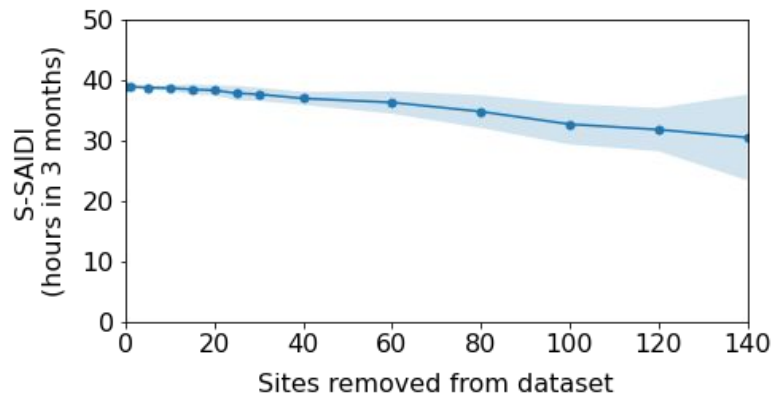


And because we don't have enough measurements to get a space time cluster we miss this outage in our experiment, where we would have caught it otherwise, leading to information loss. Our intuition is that if we have sufficient coverage, we can remove quite a few sensors before we start seeing dramatic drops in our predicted outages

Information contributed by each site is often redundant,  
suggesting sufficient deployment density

Now we ran this many times and we find is that information contributed by each site was rudent. They  
way we show this is here

Information contributed by each site is often redundant, suggesting sufficient deployment density



Where on the y axis we see S-SAIDI which is a metric we define that captures the hours of power outages experienced by **[the average] customer [in our sample]**, in this case over 3 months. On the x axis we see the number of sensors that we simulate removing from our dataset for that run. If we didn't have sufficient density, we would expect large variation, or large amounts of information loss. Instead, what we can see after running this many times is that the information contributed by each site is often redundant, the bounds don't really grow large until we are removing nearly 1/4 of our sensors, which is great as this suggests that we have sufficient deployment density to be capturing the majority of larger outages [1]. Now of course there could be many single house outages that we miss from our sample all together - but this gives us a good sense that even adding many more sensors probably wouldn't drastically impact our measurement of larger outages

Or... "a metric we define that captures the hours of power outages experienced by **[the average] customer [in our sample]**, in this case over 3 months" → the average number of hours without power experienced by the customers in our sample in a period of time - in this case, 3 months.

Should also mention the total number of sites that were reporting/deployed in this period. 150 I think?  
To put x axis in perspective

[1] We have a better reason to believe we capture most MV outages and ALL HV outages (the combinatorics method, i.e. the chances of deploying at 150 sites and not being "under" every substation are vanishingly small).

The main message of THIS graph is that our deployment adequately samples from the range of reliability experiences at our sites - basically, we haven't chanced to sample too much from sites with abnormally good or abnormally bad reliability.

## Outline

1. Validating PowerWatch in the lab
2. Deploying PowerWatch at scale in Accra, Ghana
3. **Evaluating data returned (without ground truth)**
4. Early learnings about the grid

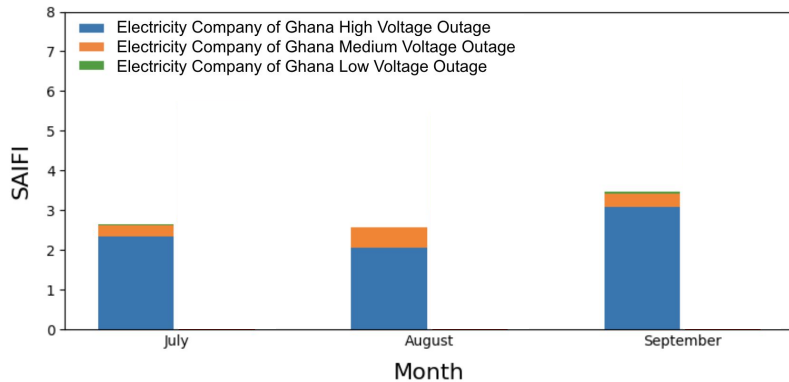
Ok, so with the bit of time we have left,

## Outline

1. Validating PowerWatch in the lab
2. Deploying PowerWatch at scale in Accra, Ghana
3. Evaluating data returned (without ground truth)
4. **Early learnings about the grid**


I can share three of our early learnings from the grid collected by our system.

## We detect more outages than are reported



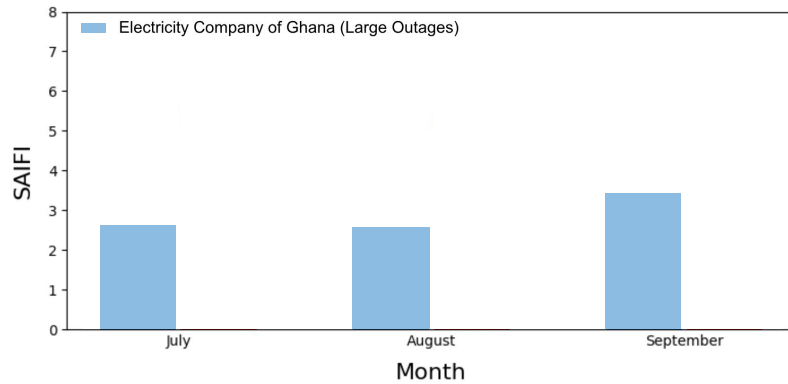
So here we are looking at SAIFI on the y axis, which is the frequency of outages experienced by the customer. So what this is saying is that for the month of July 2019 the average customer experienced a bit less than 3 outages. And this data was provided to us by the electricity company of Ghana, or ECG, the national utility that runs the grid we were studying. ECG has automated measurements on the high voltage network and on a lot of the medium voltage. This allows them to desegregate this data into high voltage, medium voltage, and low voltage outages which are colored here. Now what is important to note here is this little bit of green here, the low voltage, are shown to really contribute nothing at all to the overall outages seen by the utility.

## We detect more outages than are reported



So here we are looking at SAIFI on the y axis, which is the frequency of outages experienced by the customer. So what this is saying is that for the month of July 2019 the average customer experienced a bit less than 3 outages. And this data was provided to us by the electricity company of Ghana, or ECG, the national utility that runs the grid we were studying. ECG has automated measurements on the high voltage network and on a lot of the medium voltage. This allows them to desegregate this data into high voltage, medium voltage, and low voltage outages which are colored here. Now what is important to note here is this little bit of green here, the low voltage, are shown to really contribute nothing at all to the overall outages seen by the utility.


## We detect more outages than are reported



Now to compare this against powerwatch, which doesn't have a perfect method of differentiating outages between high and medium voltages, we group the medium and high voltage ECG outages into "large outages".




## We detect more outages than are reported




And when we look at the large outages experienced by powerwatch, which we define as clusters that are greater than 10 sensors, we see parity largely, which is exciting! This gives us insight that our sparsely deployed low cost sensors might be experiencing the same set of outages that the utility is sensing from their Scada system, but at a fraction of the cost.

## We detect more outages than are reported



However this story gets a bit more interesting when we bring back in the low voltage outages, remember the small sliver here.


## We detect more outages than are reported



When we look at the powerwatch detected small outages, we see a drastically different story, at times showing a nearly 2x total increase of SAIFI, which is a very different estimate. Again, this is a research system, and we can't claim it's a bullet proof result, but it does provide some evidence that there is more to look into here, and the larger number does align with the more common perception that reliability continues to need to be improved in Accra.


## We observe many hours outside nominal voltage

- ECG defines acceptable range as +/- 10% of 230 volts.
- We find **44.9%** percent of measurements below and **3.5%** above this range.




We are also able to observe many hours where our sensors observe voltage of that are well outside of the nominal voltage range, + or - 10% of 230 voltage. This has huge effects on appliance health and the overall usability of the grid.

In much of the world, grids are not instrumented at the customer level




So to wrap up the problem that powerwatch was designed to solve is that in much of the world, grids are not instrumented at the customer level.

In much of the world, grids are not instrumented at the customer level



This is the normal view of automated metering, which sits at the high and medium voltages.

# Our system sees more



As opposed to our system, which by going allway to the leaf node sees much more, and at a fraction of the cost.

## Systems like PowerWatch still face barriers to scale even though the data has real-world value

- Resulted in multiple publications on meta-systems required for scale
- Contracted to scale deployment to 1500 sensors by Government of Ghana
- Acting as primary source of data for monitoring and evaluation of entire Ghana compact



PowerWatch didn't scale easily. Actually, the bulk of the publications leading up to this deployment describes meta systems we developed to support scaled deployments, and generally there is a lot of remaining critical work to be done to make systems like powerwatch more easily implemented. Similarly, many of the most interesting challenges with this work was how to validate our data stream without ground truth, and the techniques we propose nod towards future work as well. However, we have had some continued success, and we recently scaled our deployment to 1500 PowerWatch sensors by the government of Ghana. Our new deployment is now acting as a primary data source for the Government of Ghana and the US Government, as well as other firms, each of whom are involved in monitoring and evaluating a large investment that is improving the grid. And it is being used by lots of other researchers, its been this wonderful experiencing of taking and packaging our insights so that now we know how to take the measurements it leads to the positive impact that motivated us originally.



## Watching the Grid: Utility-Independent Measurements of Electricity Reliability in Accra, Ghana

- Sparsely-deployed sensors at imprecisely-selected households and businesses allow us to observe more outages and voltage quality problems than operating utility.
- We achieve coverage at a fraction of the cost of utility grade sensors, creating a financially-viable path toward monitoring for the most resource-constrained utilities.
- We demonstrate the first true utility-independent direct measurement methodology, freeing regulators and independent evaluators from reliance on the utilities they are auditing.
- We collect the largest high-resolution, open-source dataset on the experience of electricity consumers in a low- or middle-income country ([github.com/lab11/powerwatch-ipsn2021](https://github.com/lab11/powerwatch-ipsn2021)).

Noah Klugman @naklugman 



nLine Inc.



# What does it mean to “Measure Reliability”?

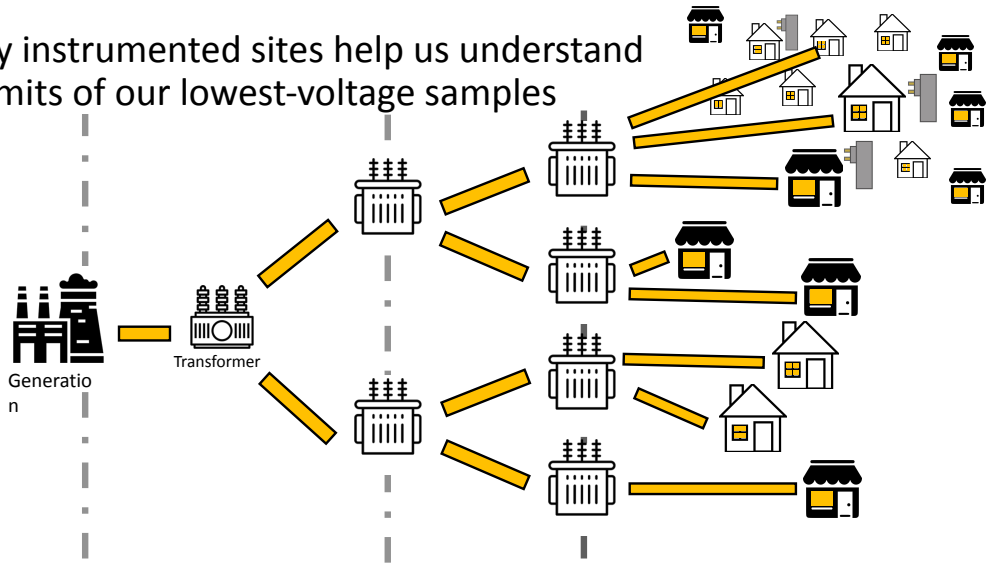
System Average Interruption **Frequency** Index (SAIFI)

$$\text{SAIFI} = \frac{\text{Total Number of Customer Disruptions}}{\text{Total Number of Customers}}$$

System Average Interruption **Duration** Index (SAIDI)


$$\text{SAIDI} = \frac{\text{Sum of All Customer Disruption Durations}}{\text{Total Number of Customers}}$$

Highly instrumented sites help us understand the limits of our lowest-voltage samples




To answer that we performed an experiment in the field where we highly instrumented two sites putting many many more sensors in each site than we normally would.

Highly instrumented sites help us understand the limits of our lowest-voltage samples



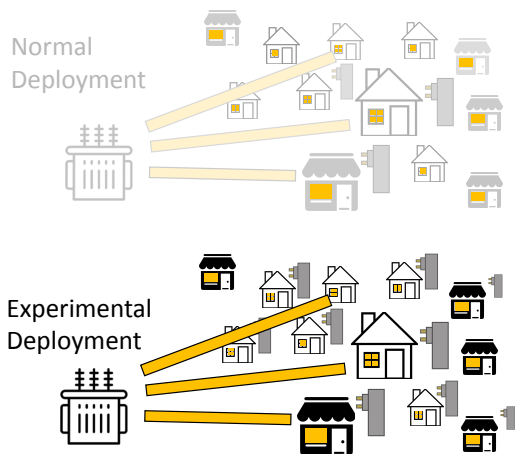
So where we would have normally deployed 3 sensors in a site

## Highly instrumented sites help us understand the limits of our lowest-voltage samples




We now

## Highly instrumented sites help us understand the limits of our lowest-voltage samples



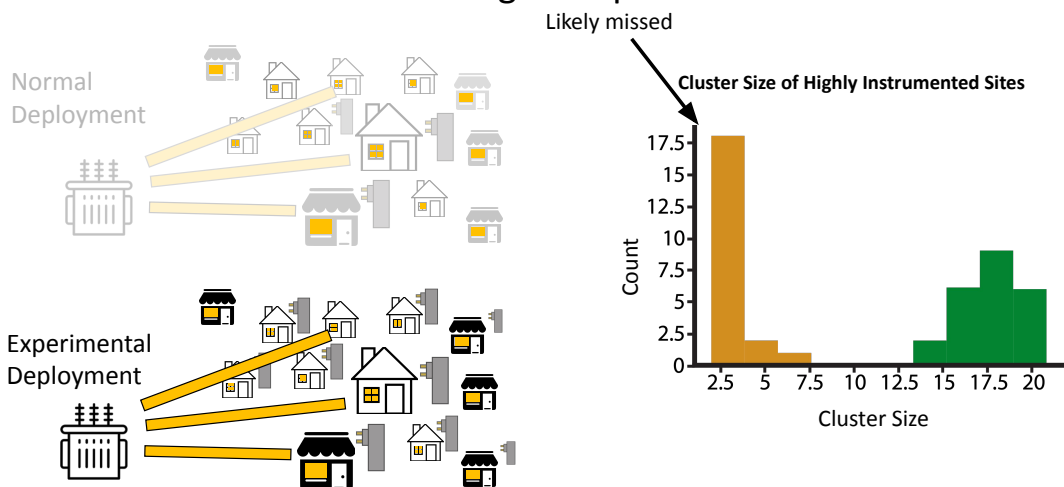
Deploy around 25. Here you can see the graphic of that experiment.

## Highly instrumented sites help us understand the limits of our lowest-voltage samples



And what we find is that in these two sites we see clusters of outages that really are of two different modes.

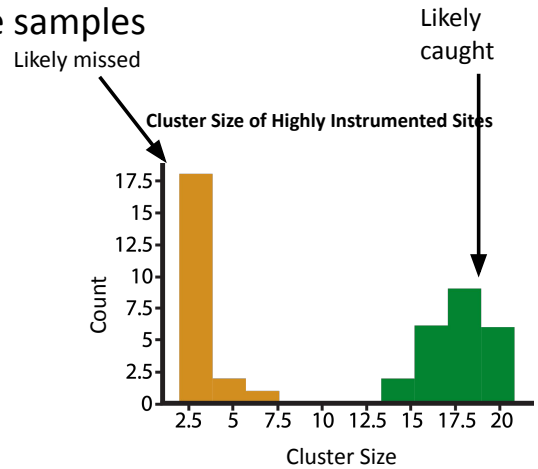
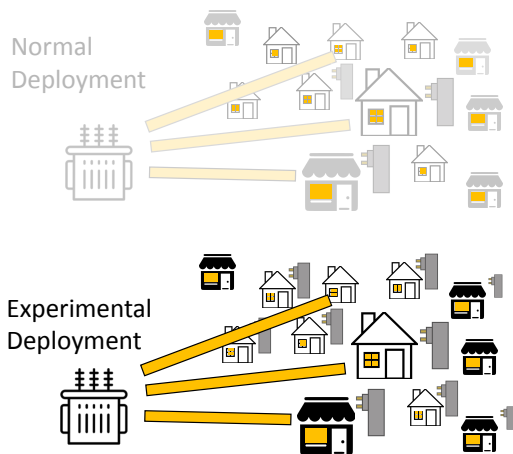
## Highly instrumented sites help us understand the limits of our lowest-voltage samples



The first of these very small outages, which are impacting a small number of houses, we would likely miss unless we had happen to have a sensor in that house.




## Highly instrumented sites help us understand the limits of our lowest-voltage samples



More optimistically though there is these second mode, where clusters are fairly large. And the fact that there are regular outages stemming from common issues in the low voltage network, which would lead to these clusters, and that even if we are placing sensors fairly randomly within a site, we are likely to still be detecting these types of low-voltage outages.

# Co-reporting rates and voltage correlations scores are higher for sensors under the same infrastructure

- Can expose some knowledge of grid infrastructure
  - 4.5% average likelihood of previously reporting with other sensors
  - Goes up by 3x when known to be under the same infrastructure




	Whole Population	Under same feeder	Under same transformer
<b>Percent co-reporting outages</b>	4.5%	11.22%	11.64%

Finally, we are able to start looking at whether we can identify the topology of the grid. We look at sensors that are known to be under the same infrastructure, and find that these are more likely to experience outages and voltage correlation, as we would expect, although reconstructing the network topology from these relationships remains as future work.

# Testbed time with improved sensor hardware/firmware

Older Sensors



Newer Sensors

