WHITE PAPER

Aclara RF Phase Detect

How Aclara is using smart meters, RF communications, and mathematics to automatically calculate the phase of your meters.

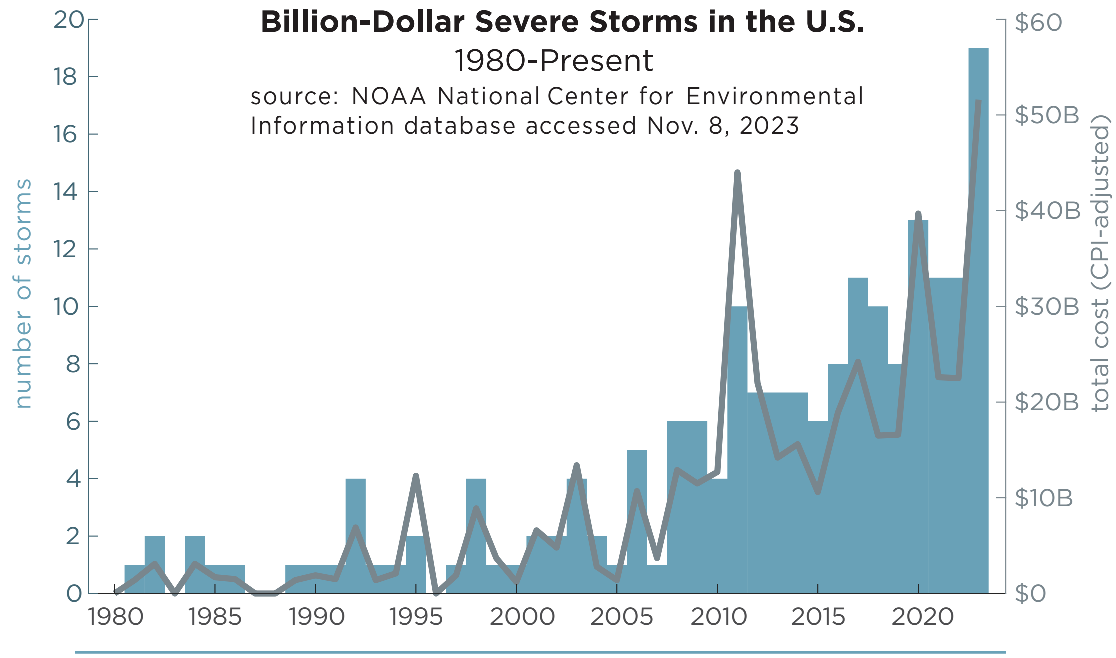

Figure 1

There's little disagreement that it's getting harder and harder for electricity providers to keep the lights on. 2023 set a record for infrastructure damage in the U.S. due to severe storms. This was no anomaly. Storms have been increasing in both quantity and strength for several decades now. Warmer air temperatures combined with warmer ocean temperatures leads to more water in the air. That's more energy that can be released when conditions are right. If this trend continues, and there is no reason to think it won't, 2023 will soon look like an average year.

In this environment it is difficult to keep an accurate model of the distribution network. In the aftermath of a storm the distribution utility's highest priority is to restore power as quickly as possible. The lineworkers on the front lines aren't so concerned with matching the network reality with the topology map in the engineering office, particularly if a new configuration means they can restore power more quickly.

Let's suppose on Day 0, your network model – the map your engineering team keeps of the distribution network wiring in all its detail – is a perfect representation of the truth of your actual network. You know with 100% confidence the wiring phase of every one of your electric meters. Then a severe storm comes through and knocks out power to many of your customers. Your line crews do their best to restore power but when the dust has settled and all your customers are turned back on, you no longer have that 100% confidence in your model. Then, before you can correct the mistakes, another storm comes through and, again, the wiring changes. After several iterations of this, you can see how that formerly perfect model has been degraded. It becomes a battle between your engineering modeling team and nature. Our own informal survey of utilities is that in practice the network model is about 85% accurate.

But so what if the model is only 85% accurate? Why should you care about the 15% of meters that are mapped incorrectly? The most pressing reason is that it causes your load balancing to be imperfect which leads to efficiency loss. The more severe the mismatch the higher the loss. In effect that 15% error is eating into your margin.

There is a more critical problem, though, that until recently seemed to be in the distant future: There are a lot of new distribution analysis and automation applications emerging that require the network model to be nearly 100% accurate. Some of these are outage management techniques, designed to combat the difficulties in maintaining and restoring power in the changing climate. Many others offer solutions to problems brought on by the increasing prevalence of electric vehicles and distributed energy resources.

If you've been living with the 15% model error until now, chances are you will soon find yourself needing to correct it.

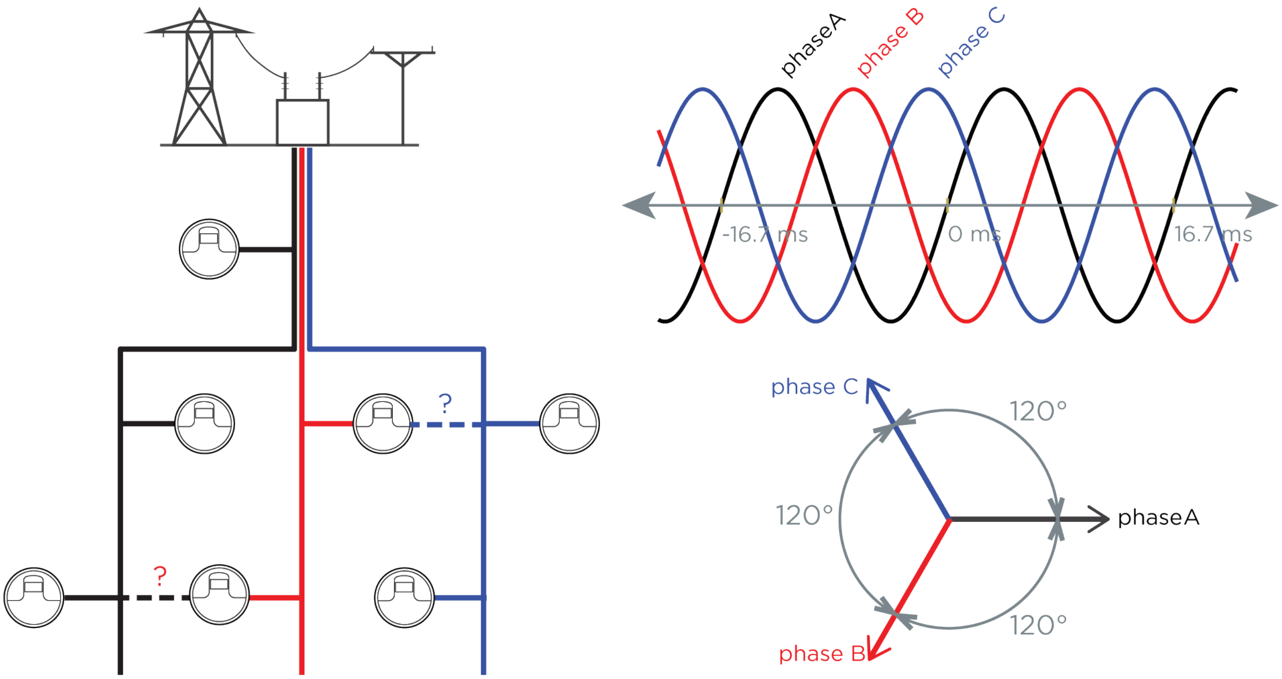

Figure 2: The left-side image depicts a small, 12-load three-phase feeder. A substation feeds three meters connected to a three-phase medium voltage line. The feeder splits into three laterals, each single phase. Dashed lines indicates that the engineering model for the network show incorrect connections. The rightside image shows the voltage waveforms for each of the three laterals (phases) and their corresponding phasor representation. These images demonstrate that if you know the phase of a meter, you know the lateral that it is on, and vice versa.

Nearly all of these applications assume that the topology of the network is known.

If you've been living with the 15% model error until now, chances are you will soon find yourself needing to correct it. This could be an expensive proposition. Manually checking the wiring phase on each meter requires visiting each meter with a line crew. Just doing this once is in most cases prohibitively expensive. If you are required to do it after each storm, it's downright impossible. Wouldn't it be nice if there were a way to use your existing capital to do it automatically?

Automatic phase detection of electric meters can be approached in one of two ways: either as a topology estimation problem or as an angle estimation problem. This is illustrated in Fig. 2 (left side image), where a very small, 7-load three-phase feeder is depicted. The feeder splits into three laterals, each single phase. The defining characteristic of a three-phase network is that each of these laterals is fed by a 60 Hz voltage waveform that is 120 degrees out of phase with the other two laterals. The voltage waveforms for each of the three laterals (henceforth simply referred to as "phases") are demonstrated on the right-side image of **Fig. 2.**

You can see from this depiction that if you know the phase of a meter, you know the lateral that it is on and vice-versa. This applies to all the meters on the feeder except for those on the three-phase section of the line near the feeder. So, in most cases, the problems of angle estimation or topology estimation are the same. We have chosen to address the problem as an angle estimation problem. In this way, we can directly observe which wiring phase the meters are connected to, and we can discern phase on three phase sections.

UNPRECEDENTED PRECISION

Phase detect will become your go-to topology mapping tool

Our patented technology uses Aclara's point-to-multipoint communications network to implement a time-synchronized measurement across all your network's smart meters. By establishing a common time across all meters, the relative angles of each of the voltage waveforms can be measured directly. We have found that in practice this is done with surprising precision: far less than 1° error.

Consider the same simple distribution network as in Fig. 2, redrawn in **Fig. 3**. Let's say this network uses Aclara RF, which uses a modem on every meter that can communicate with collectors, typically mounted to poles. The collectors feed the data they collect from the meters to the cloud, so each meter only needs to talk to one collector. In practice, though, most meters are able to communicate with at least two collectors, often more.

The phase detection process kicks off by having

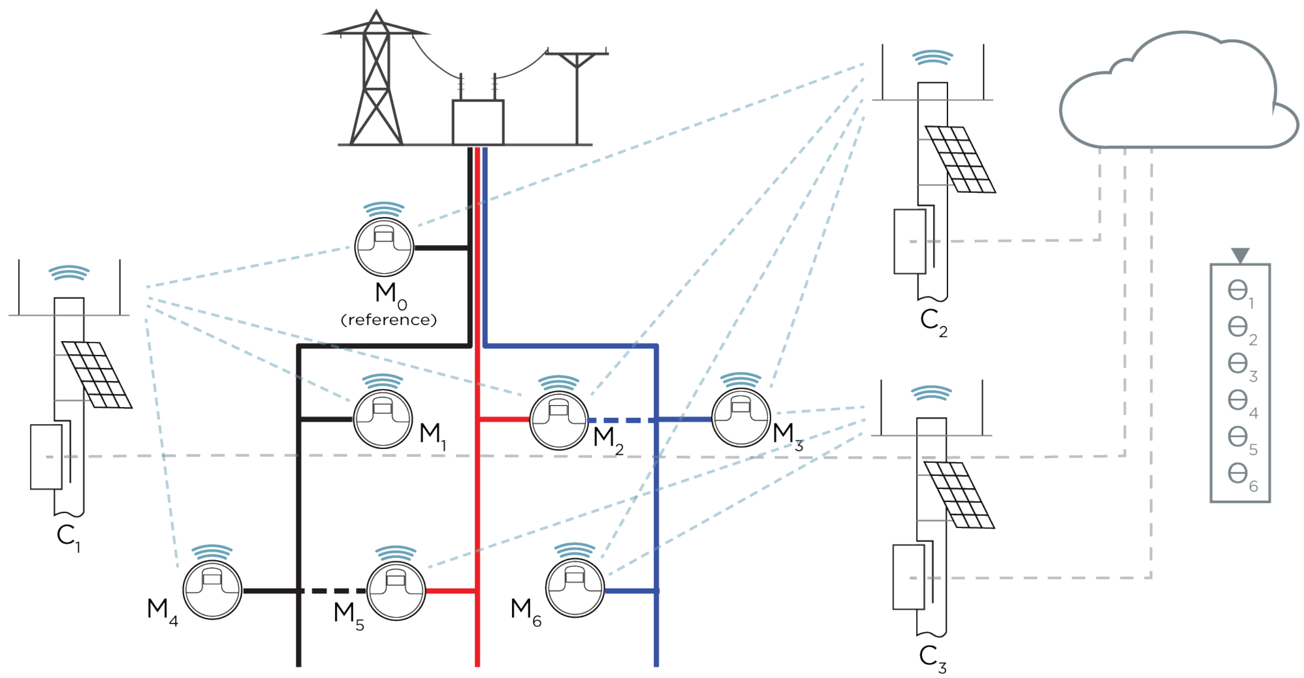

Figure 3: The same distribution network of Fig. 1. This illustrates how a modem on each meter communicates with the pole-mounted collector unit, which feeds data into the cloud. Phase detection begins when each collector transmits a signal, or beacon.

each collector transmit a signal, called a beacon, one-by-one. When a meter receives a beacon from any collector it makes a measurement of its voltage waveform. When each collector has transmitted exactly one beacon every meter should then have at least one measurement set in its memory. Many will have more than that. All the measurements from every meter are transmitted back to the cloud via the collectors. These data are then processed by an algorithm and the phase of each meter is provided. The calculation happens shortly after the data are collected and uses only the measurements that were collected from the beacons just transmitted. Meter phases can be determined quickly and on demand.

We have given each meter a name: M with a unique subscript. Similarly, we will call the collectors C with a unique numerical subscript. We connect two nodes in this diagram by an edge if that meter responded to a beacon from that collector. So, for example, Meter 4 responded to beacons from collector 1. Meter 0 responded to beacons from collector 1 and collector 2. Meter 6 responded to collector 3 and collector 2.

Every meter, except for meter 0, has an unknown phase angle, designated as Θ, associated with it. We will say that meter 0 has a known phase of A and that its voltage waveform angle is therefore 0 degrees. Henceforth, we will call this the reference meter. This is the one meter on your network for which you are certain of the phase. The problem for the cloud calculation is then to compute the other six phase angles from the beacon data.

Let's rearrange Fig. 3 so that all the collector nodes are in one column and all the meters are in another column. We will keep the edges the same. This is shown in **Fig. 4**. This arrangement should look familiar to students of artificial intelligence for it resembles a neural net. While we don't use a neural network specifically, we do use a technique very closely related to it called belief propagation. This, along with neural nets and other graphical processing methods, are part of an emerging set of graphical processing algorithms that have shown to be extremely adept at solving problems robustly where large amounts of interlocking data and unknowns are involved. This fits our problem perfectly because in reality there will be hundreds of collectors and tens, if not hundreds of thousands, or even millions of meters.

The actual representation would be extremely large, far larger than we could depict here, with even more edges crossing each other in a big mess. While we as humans would have a hard time visualizing it, it turns out to be very easy to depict programmatically in a computer algorithm.

Let's look at how belief propagation works. Next to each meter node we will write out initial knowledge of that meter's phase, as shown in **Fig. 4.** For the reference meter we know that it's on Phase A. Every other meter has an unknown phase so we will put a "?" next to each of them.

Each meter node will send a message to all the collectors that it is connected to. This is not a real message in the sense that meters transmit to collectors in the communication network. Rather, it is a mathematical message within the algorithm that abstract versions of the meters and collectors communicate with each other. It's very similar to how a neural net works, if you imagine each node is a neuron and sends synapses to other neurons that it is connected to. If we are certain of a meter's phase, we send a message that relays that certainty. If we are not, then the message communicates the nature of the confusion.

After all meter nodes have sent their messages to the collector nodes, a calculation is done on each collector node to best resolve the ambiguity between the meter nodes it has received messages from. Corrections to each meter node's current phase estimates are then sent back to the meter nodes, again as mathematical messages. The correction calculations are performed on each meter node using the messages from multiple collectors. Often there will still be some ambiguity at each of the meter nodes. However, it will be less than the ambiguity that existed before the process started. That ambiguity can be further reduced by repeating the process: the meter nodes transmit messages to the collector nodes, and the collector nodes back to the meter nodes. After a few iterations of this, the ambiguity will be gone and a solution will have been arrived at. We can write the estimated phase of each meter beneath the initial estimate.

It's not much of a simplification to say that the process of learning in the human mind is that of creating new connections between neurons. The more connections a neural net has, the more robust it performs. The same phenomenon is at work here. What makes this method work so well is that each meter node is connected to multiple collector nodes. Recall that these connections symbolize communication in real life between the meter and the collector. It is the point-to-multipoint feature of the Aclara RF communication network combined with multiple redundancies that provides this robustness. In fielded deployments, we've found that this method fails less than 1% of the time.

PRACTICAL ROBUSTNESS

Phase detect works under some of the harshest conditions you will experience

As we've already noted, our own informal surveys suggest that your existing phase models are probably already at least 85% correct or more. This is not the scenario we've outlined in the discussion above. Recall that in our belief propagation graph (**Fig. 4)** under initial estimate, our reference meter is phase "A" but all the others are "?". By this we have modeled a scenario in which we have complete uncertainty about every meter except the reference meter. It's as if we are saying that there is an equal probability of these meters being on Phase A, B, or C. This is a far cry from the 85% certainty you likely have in your existing phase maps.

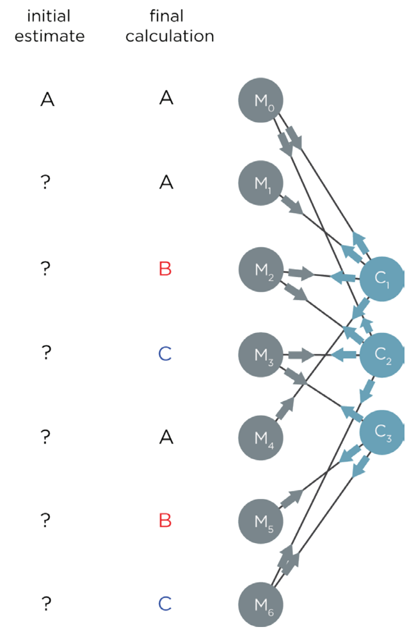

Figure 4: Here, the network from Fig. 2 is shown in a simplified view. In this configuration, it resembles an artificial intelligence neural net. While the arrangement is similar, the technique used for phase detection is called belief propagation. It is similar to other graphical processing algorithms that solve problems using interlocking data sets.

From a mathematical perspective, this means there is already a great deal of information in your existing topology that we are ignoring. The challenge is determining which 15% of your smart meters are mapped incorrectly and what the correct phasing should be. If we can only label our initial estimates as "A", "B", "C", or "?", how can we indicate partial certainty?

Fortunately, belief propagation permits a solution. In the graph shown in **Fig. 5**, we have changed the initial estimate for each meter node to a set of probabilities. Meter 1, for example, is estimated to be on Phase B with 85% certainty. If the network model is in error, though, we say that there is an equal chance of it being on Phase A or C: 7.5% each. We have indicated this as a bar graph where the height of the bars for each phase indicates our initial certainty in the network model's reported phase for that meter.

One of the most elegant features of belief propagation is that it permits the passing of messages containing probabilistic data. We may use the same iterative message passing algorithm described in the previous section, but with probabilistic messages. After several iterations, the algorithm converges to a solution that blends the data collected from the beacons with the existing network model in the most harmonious way. By bringing in information from the existing model, we have further improved the precision of the algorithm!

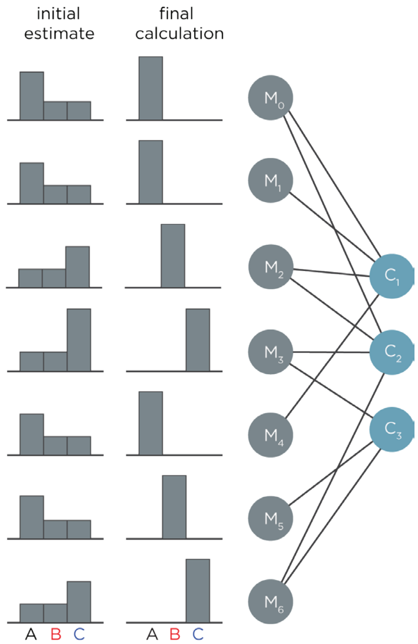

Figure 5: This depiction of belief propagation shows the technique's ability to permit messages containing probabilistic data. Here, the Initial Estimate for each Meter node (M) is now a set of probabilities. The height of the bar shown indicates the initial and final certainty in the reported and computed phase for that meter.

You may have noticed that even the reference meter has been permitted some ambiguity. In fact, it is initialized with the same uncertainty that all the other meters have been initialized with. The not-so-obvious implication of this is that we really don't need reference meters at all. This is a surprising outcome that we did not anticipate in design but rather discovered by accident when a reference meter stopped reporting data during a test of the algorithm, and we were able to produce accurate meter phases regardless.

There are other advantages as well. If the algorithm is having a difficult time computing a solution or some source of ambiguity exists, the final calculation will show that uncertainty. This is reported to you as a user by a confidence figure returned with each meter's phase calculation. Often that ambiguity can be used to track down other systemic problems, whether they arise from meter or communications issues, or actual distribution network issues that may otherwise have remained hidden.

The not-so-obvious implication of this is that we really don't need reference meters at all.

OPPORTUNITIES FOR GROWTH

Phase detect generates data that is unique among AMI networks

Underneath the hood, each of the meters is measuring the phase angle of the voltage sinusoids relative to each other at the instant a beacon is received. This is demonstrated in **Fig. 6**. For model correction it is only necessary to measure this to within about 15 degrees, but our analyses of deployments at existing utility networks have shown accuracy far better than that. Transmission engineers have understood the value of simultaneous phase angle measurements – or synchrophasors – for some time now. However, this is just starting to show value in distribution networks, largely due to the proliferation of EVs and IBRs. Our study of reported phase angels from trial deployments has given us some insight into just how valuable this data could be for distribution networks.

In **Fig. 7**, we have plotted the phase angles reported by our phase detect product for a set of about 400 meters that all tuned out to be on Phase A. For this experiment, the system was configured to calculate phase once per day, so we have one value for each meter each day over a period of one month. It is important to point out here that prior to this it was not possible to make simultaneous measurements across a population of smart meters. While our original purpose was to make synchronous measurements for the purpose of calculating phase, it made sense to take a closer look at the data, and see what it could possibly tell us.

The first thing that caught our attention was how spread out the angles from the entire population are over a single day: as much as 20° on some days. Conventional wisdom is that there should not be so much excursion from 0° on distribution networks. However, this is a wisdom that was born in analysis and simulation of pristine distribution networks that exist only on paper. Actual networks have imperfections and pathologies that are difficult to model and, all too often, are only anticipated by the teams that maintain them and not so much by academics. It seems there is value in looking at meter phase angles after all.

The plot in **Fig. 8** is the same, except that the phase angles from each meter are color coded to indicate similar behavior. We produced this by running a simple time-series clustering algorithm. Four distinct groups of meters emerged, each identified by a unique phase angle walk over the one-month observation period.

We had access to the coordinates of each of the 400 meters in this study. In F**ig. 9**, we have indicated the position of each on a map with color matching the behavioral group each belongs to in the time-series phase plot. It is obvious that meters are clustered together spatially by their phase behavior. It turns out that each of these clusters corresponded to different

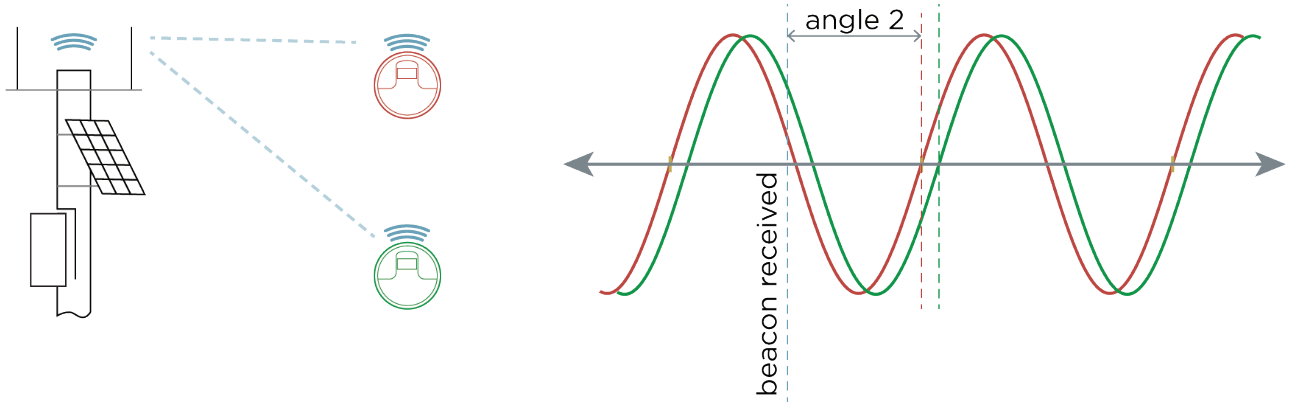

Figure 6: At the instant a meter receives a beacon it measures the phase angle of the voltage waveform. After processing of these angles from all meters and beacons the phase angle relative to the reference meter has been shown to be accurate to within 1°.

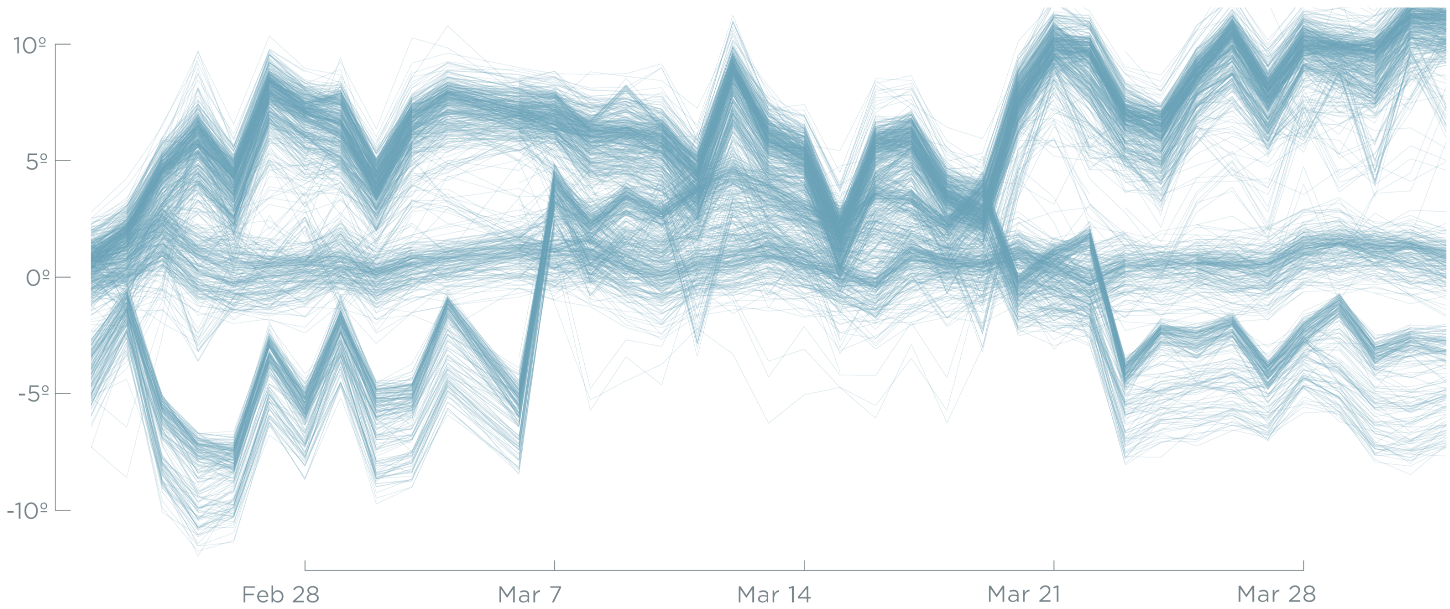

Figure 7: Plot of the phase angles recorded by Phase Detect for 400 meters over the course of 30 days. Each meter received one phase calculation per day. The plot showed that the angles from across the meter population could vary as much as 20 degrees over the course of one day.

substations. In fact, what we have stumbled across is a way of identifying substation connectivity from this data in addition to phase connectivity.

The explanation for this derives from the similarity of voltage waveforms between meters that are located near each other topologically. Think of it this way: two meters on the same distribution transformer will have nearly identical voltage waveforms – and thus nearly identical voltage phase signatures over time. Those two meters will have similar waveforms to meters down the street on the same feeder, but not as similar. A meter on the same feeder a mile away even less in common.

Direct measurement of the phase angle of the voltage waveform at each meter is a completely new measurement that has not been fully considered until now because it was largely considered impractical if not impossible. Hubbell is actively engaged in research and development to exploit this signal for full topology identification. This includes both meter-to-feeder mapping and meter-to-transformer mapping. Based on results similar to those presented here, we believe this to be very possible.

The world is changing around us and, at times, it can seem as if it is getting ahead of us. We believe in the power of human ingenuity to adapt to those changes. The mathematics show that it is possible to stay abreast of drastic changes in your network topology using data from your existing assets, albeit in ways that you maybe haven't thought of before. The results confirm the mathematics. Hubbell, by its commitment to applied research, is determined to make these advanced mathematical techniques accessible to you and give you the tools you need to adapt and survive.

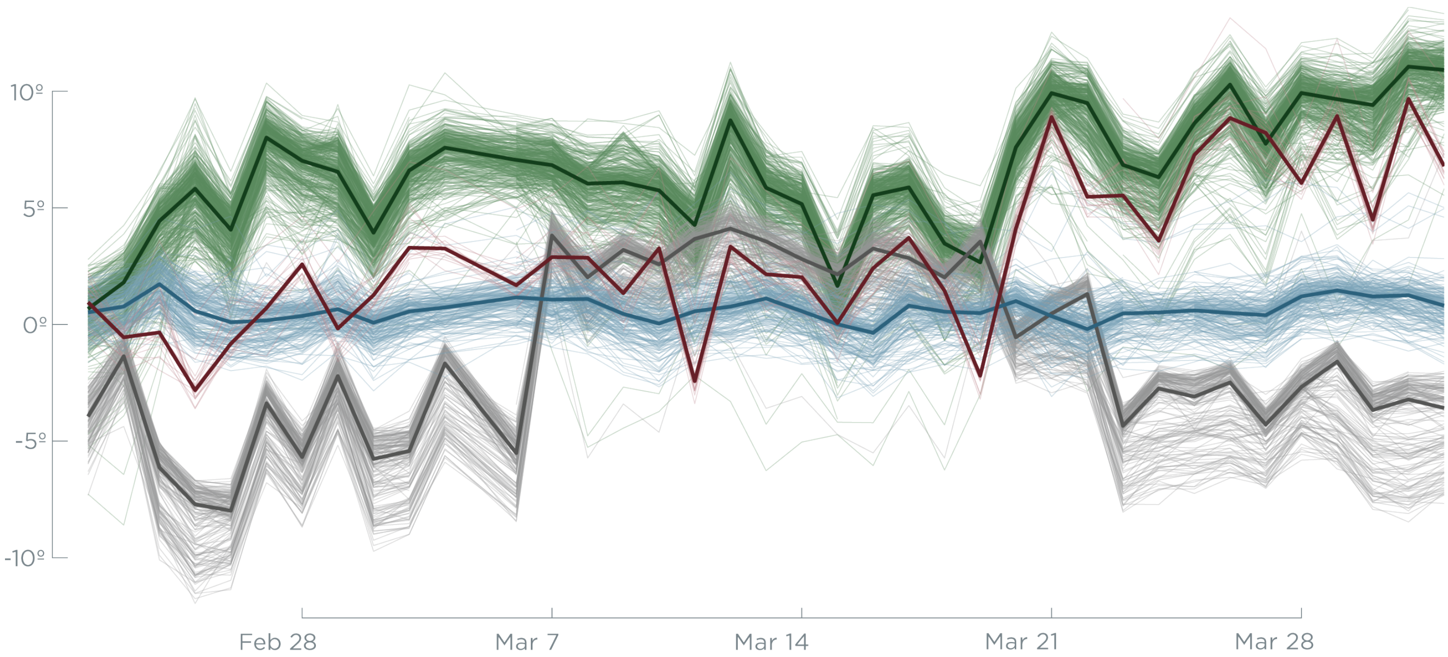

Figure 8: The same plot of meters as Fig. 7 is shown, except the phase angles are color coded here to show similar behavior patterns. Bold lines are the mean of the phase angle for each group. Four distinct behavioral groups emerged from this data.

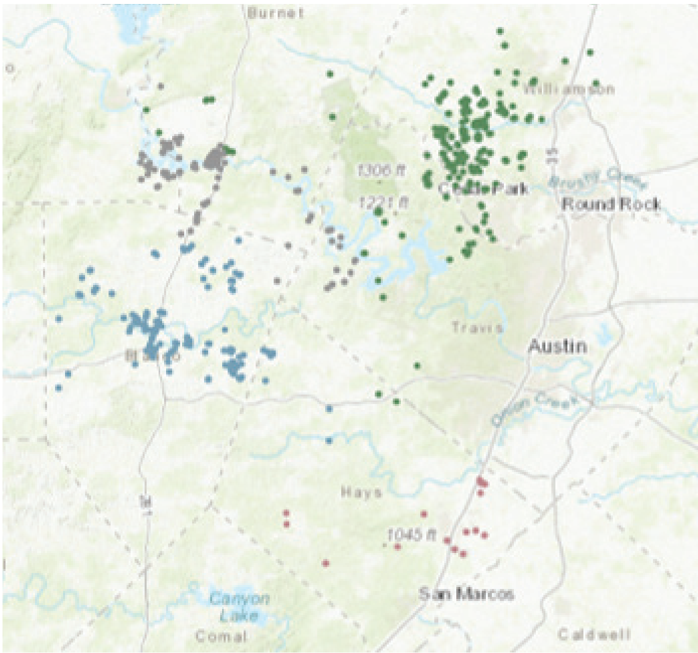

Figure 9: Using the coordinates of the 400 meters in the study, the map shows the position of each meter in its corresponding phase behavior color from Fig. 7. The map revealed that meters with similar phase behavior were clustered together spatially. Each color cluster corresponded to a different substation, revealing new possibilities for using the phase connectivity data set to identify substation connectivity.

About the Author

David Rieken, D.Sc

The leader of Hubbell.R, an industrial research laboratory within Hubbell that is committed to solving problems in electrification through applied analytical solutions. He has conducted smart grid research at Hubbell (formerly Aclara) for 18 years. Prior to that he performed research at MIT Lincoln Laboratory, Boeing, and General Dynamics. He is passionate about bringing the same quality those research laboratories are known for to bear on problems in distribution automation.

A proud member of the Hubbell family.

Aclara.com | 800-297-2728 | info@aclara.com

ACL-ACLM-WP-EN-01908 | 04/2024