21, rue d'Artois, F-75008 PARIS

http://www.cigre.org

CIGRE US National Committee

2019 Grid of the Future Symposium

Improving the Reliability of the Distribution Grid Using a Distributed Restoration System with Dynamic Leader Assignment

M. QUINLAN, Y. SHARON
S&C Electric Company
USA

SUMMARY

Fault location, isolation, and service restoration (FLISR) is the process of operating switching devices to isolate an unhealthy section and restore power to deenergized but otherwise healthy sections of a network from nearby circuits. Automated FLISR systems are able to restore power to customers much quicker, improving customer satisfaction and reducing metrics related to outage time. Automated FLISR can be performed through a variety of methods. One of these is through distributed intelligence, consisting of network switching devices communicating to a nearby device designated as a leader, which can coordinate the restoration process. This method maximizes restoration speed by limiting the scope of restoration and communications to the immediate area. Presented here is an implementation of distributed intelligence where leadership is assigned dynamically to an open tie device. This change grants additional benefits in the speed and reliability of the restoration.

KEYWORDS

Distribution Automation, Advanced DA/DMS Applications, Smart Grid Systems Architecture, Performance Optimization

michael.quinlan@sandc.com

I. INTRODUCTION

Transmission systems are built as a mesh network with multiple paths between every node. While lines can be disconnected because of a fault, if the system is operated correctly, a single line that disconnects does not result in loss of service to any customer. Distribution systems, on the other hand, maintain a radial topology with a single path between every customer and the single source that serves this customer. Any section that gets disconnected because of a fault will result in loss of service to all the customers on this section as well as all sections downstream of it. Those customers on the faulted section must wait until the fault is removed before power can be restored to that section. As for the customers downstream of the faulted section, while a momentary outage is inevitable, they can be restored sooner if connected to adjacent feeders with sufficient reserve capacity, as seen in Figure 1.

The controlling of distribution feeder switches in response to an outage event typically follows a strategy known as fault location, isolation and service restoration (FLISR). When a fault occurs, or loss of the transmission feed for example, switches can be operated to isolate the affected area and to restore power to the healthy sections when capacity is available from neighboring circuits.

With a non-automated distribution grid, a utility is notified when customers call in to report outages. A picture of the area that is affected can be determined after a few customers call in from different points on the feeder. At this point the utility dispatches a truck to isolate the area around the event by manually operating the switches. This takes time as the switches are sometimes many miles apart from one another. During this time, the command center can determine where excess capacity is available, and how healthy sections of the feeder can be restored. Once determined, the truck can begin the switching operations to restore the rest of the customers. Finally, at this point the crew of the truck can go back and work to repair the cause of the event itself. This time-intensive process can be even further complicated at night and under storm conditions. The costs associated with sending out these crews can become quite large, as well as the costs to the customers who go without power during the whole process. For example, a feeder with 2,000 customers divided to 6 sections which has a fault in one of its upstream sections, will result in 1,666 customers, in unfaulted sections, losing power. Through this manual process of restoration, the unfaulted sections will get restored within approximately 100 minutes. According to [1] the cost to the customers is in excess of $275,000.

An automated FLISR application can perform many of these tasks much faster, reducing the time and number of customers impacted by a fault by automatically isolating the trouble area and transferring sections to adjacent circuits to restore service. In addition, the fault isolation feature of the technology can help crews locate the trouble spots more quickly, resulting in shorter outage durations for the customers impacted by the faulted section. In the same scenario as above, but with an automated FLISR application restoring the unfaulted teams in less than a minute, the cost to customers associated with the fault event is around $112,000. The business case for automating FLISR has resulted in many utilities incorporating automation into their system.

In Section III a list of current FLISR methods is discussed, along with the benefits and drawbacks. In Section IV an overview of a distributed intelligence approach is given. An original scheme of distributed intelligence based FLISR is described in Section V, followed by examples in Section VI.

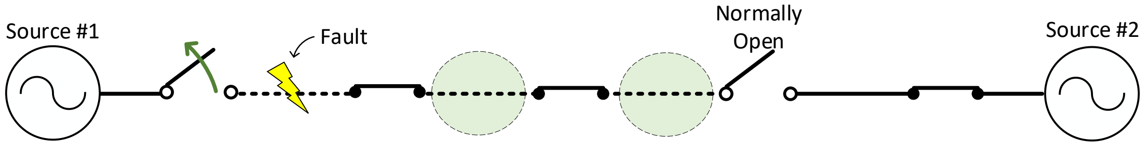

Figure 1: A fault appears on a radial network and is cleared by an upstream device. Three sections are unpowered as a result; one contains a fault, and the two circled sections are healthy and ready to be restored by the alternate source.

II. NOMENCLATURE

Section: Electrically connected segment of feeder bounded by automated switches, including laterals and customer loads.

Connected Region: Electrically connected group of sections bounded by open switches.

III. CURRENT APPROACHES

Current automatic implementations of FLISR can be separated into a few different categories: Loop-Restoration, Substation Automation, Local Heuristic Automation, and Advanced Distribution Management Systems (ADMS). Each of the categories has strengths and typical applications or topologies in which they are applied. When choosing a restoration scheme, many different aspects need to be considered which will affect the cost and benefits of the installed system, including the following:

Setup Effort: The cost of new hardware: computing equipment, additional field devices, upgrades to existing devices and communication infrastructure. The cost of engineering effort of configuring the system, and software licensing.

Maintenance Effort: The cost of performing periodic software updates, maintaining and upgrading the communication system, updating the system models as the network changes (such as load growth), and reoccurring licensing fees.

Reliability: The overall probability that the system restores successfully, given the individual probabilities of various contingencies (e.g. communication loss due to contingency or interference, inoperability of a protective, computing, or communication device), vulnerability to cybersecurity threats.

Optimality: How much load the system restores as percentage of how much load can potentially be restored. Other optimality criteria consider critical load separately from non-critical load, and number of switching operations.

Adaptability: The ability to maintain the level of performance as the system changes, including load variation or configuration into an alternate state.

Scalability: A system whose cost is high regardless of the size of the network it is deployed on is less scalable then one whose cost is more proportional to size of the network.

Speed of restoration: The amount of time customers will experience an outage. Investor-owned are penalized or rewarded by their regulators as a function of the frequency and duration of sustained outage their customers are subject to. In North America an outage is considered sustained if it lasts more than 5 minutes [2]. Therefore, the benefit in restoring in less than 5 minutes. However, several utilities started tracking sustained outages as short as 1 minute, and even momentary outages, in order to improve their customer's experience. Shorter restoration times may also allow some distributed generators to ride through an event. Therefore, the desire is to restore as quickly as possible.

Loop-Restoration

One of the simplest implementations of FLISR, loop-restoration can restore unpowered sections of a feeder without communication nor human intervention. In the scenario where a single tie connects two feeders, the tie is programmed to close after voltage is lost on one side (due to fault event or loss of transmission feed, for example). In the case where voltage is lost is due to a device locking out after a fault, the fault will become energized from the other direction, but then will quickly be isolated by the protection scheme. With a proper protection scheme, this will result in only the two reclosers nearest the fault ending up locked-out. All other sections of the two original feeders are powered.

The main advantage of this approach is the ability to restore quickly, and without the use of a communication system. No communication system leads to an easy installation and maintenance; possibly only requiring a firmware upgrade of existing equipment [3]. Also, because there is no reliance on communication, it is considered highly reliable. However, this can only be used in simple systems where there is a single tie switch connecting feeders, and the adjacent feeder always has enough reserve capacity to reenergize the unfaulted, unpowered sections. Additionally, it may be undesirable to knowingly reenergize the fault from the other direction, causing more stress on the system and increasing hazardous conditions.

Substation Automation

When the network is more complex than loop-restoration can handle, automation can be performed by logic residing in the substation. This entity can monitor and control the switches in the feeder, and based on a pre-determined response, can reconfigure the network to restore power. This entity acts as an expert system which must be configured with the specific response to *every* situation it may face.

With a view of the whole feeder (and possibly neighboring feeders), the network can be restored from multiple sources and perform partial restorations when the capacity is too small from any given source [4]. This solution involves a large amount of set-up to program the expert system. Therefore, this leads to a system that cannot handle new configurations that were not designed for originally. Additionally, since all the switches communicate to, and all logic resides in, the same device, a single point of failure exists for the feeder(s) being controlled.

Local Heuristic Automation

When the number of possible reconfigurations becomes too numerous to program into an expert system, a distributed intelligence approach can be used. This allows the devices to coordinate with each other, through local communication, to reenergize one unpowered section at a time. The system can base its decision on real time loading values to judge whether a neighboring feeder has enough excess capacity to pick up the present load in each of the sections.

This solution is adaptable to changes in the system, such as a new topology or loads that gradually increase over time. The solution is also resilient to the loss of a device. If any of the devices cannot operate or communicate, the rest of the feeder is not prevented from restoring by utilizing the devices which are still healthy. This system requires local peer to peer communications to function and does require limited topology information to be loaded into each device [5]. As no expert system has to be programmed, and the topology may be loaded from existing system models, the setup time is reduced. However, since the system is being solved locally, one section at a time, the network may fall into local optimal solutions. This may result in some sections being dropped which would be connected using an approach with a larger view of the system.

Advanced Distribution Management Systems (ADMS)

A more recent FLISR solution, is for the logic to reside in the control center. Communications from the devices across multiple feeders report back to the control center which consolidates the information and can solve for a new state after an event occurs. The system can then command each device to operate directly. As the control center has a picture of the current state of the entire network, it may be able to solve for a more optimal solution. Similar to the previous scheme, this solution is flexible to changes in loads and other operating conditions, since it is periodically receiving device conditions and measurements.

An ADMS system provides many other benefits that go beyond FLISR. This, and the fact that configuration and software updates are all done centrally, make it an appealing proposition. With the necessary communication, computing, and system modeling infrastructure, the initial cost of adding an ADMS system is in the range of several millions of dollars [6][7]. However, setting up and maintaining the necessary infrastructure, which may include GIS, SCADA, OMS, Distribution Automation devices, substation equipment, and AMI [6], all connected together, can be an order of magnitude higher [7].

ADMS is also slower than the others, because of the communication latency from the devices in the field to the control center, and back. The resulting restoration time can therefore be up to 5 minutes [8]. Additionally, there is potentially a single point of failure for the entire system. If the controller in the control center, or the communication to it, is compromised or lost, no restoration will be performed throughout the system. The need to route the signal between the control center and the field devices through several communication devices (routers, switches, radios, etc.) may also reduce the reliability.

IV. DISTRIBUTED INTELLIGENCE APPROACH

Another way to perform FLISR, is the distributed intelligence approach. Similarly, to the local heuristic approach, the logic for reconfiguring a feeder resides in multiple local locations (such as the switching devices themselves). However, instead of restoring one section at a time, multiple sections can be restored at the same time. The centralized approach can generate an optimal solution for the system by considering all devices, but it comes at the cost of long computation and communication times. The local heuristic approach has low communication requirements and low computational needs; however, it may find itself falling short of an optimal solution. The distributed intelligence approach increases the scope of the heuristic approach to perform nearly as good as the centralized approach in terms of optimality, while only increasing communication and computational burden slightly. This approach is also faster than both the centralized approach, because communications are still local to the feeder, and the local heuristic approach, because fewer switching operations are performed to arrive at the same solution. Finally, like the local heuristic and the centralized approaches, it can be adaptive to an ever-changing system condition, and with self-discovery, can substantially reduce the setup effort.

The structure of a distributed FLISR approach can be separated into components as seen in Figure 2: data propagation, protection operation, isolation, restoration planning, and restoration execution. The progression between these components in response to a fault event is depicted in Figure 3. This fault location and isolation is performed local to the event, and the restoration is orchestrated by a leader device. A leader is responsible for coordinating the closing of open devices, which prevents source overloading and inadvertent connection of two sources which can occur when devices take uncoordinated action. A leader is also responsible for calculating the new desired topology, for which it will be restoring to. The domain assigned to a leader device can vary, ranging from a group of feeders, to a single feeder, to a single section. The more information the device calculating the restoration has, the better the outcome, but this comes at the cost of the amount of data needed to be collected.

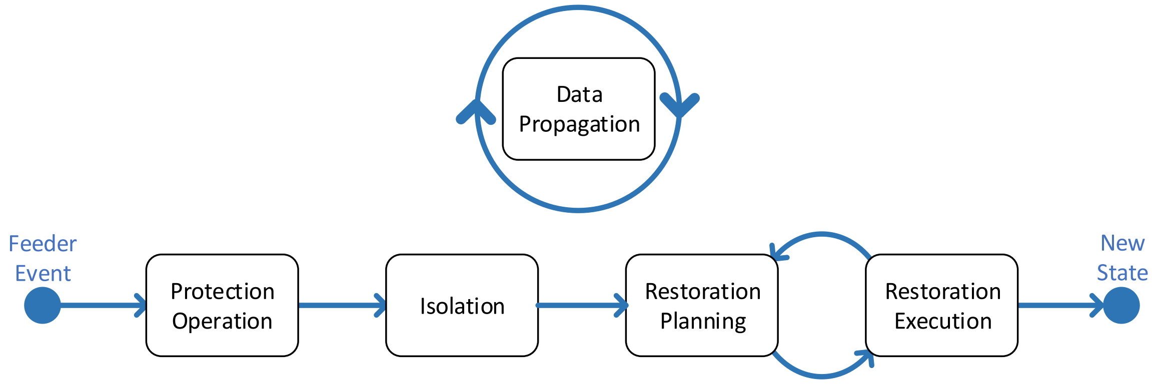

Figure 2: Diagram of distributed intelligence approach.

Data Propagation

During the setup of a device, information on the neighboring device(s) must be entered to allow the topology to be discovered. This information, along with the state of the device and measurements taken by the device, are sent to the leader to enable the leader to perform restoration. Leaders can be elected during this self-discovery period through a number of leader selection schemes, examples of such can be found in [9] and [10]. Note that this data can be passed around periodically or after an event occurs. Sending the data around periodically will result in the restoration calculation using the data that was recorded up to one data period ago. Since the load does not change quickly, this is typically acceptable. Sending the data after the event allows the leader to use the most up to date information, however, this comes at the cost of speed of the restoration. If the communication channel is slow, sending all the necessary data all at once to the leader may overload the channel. Step #1 of Figure 3 is shown with the data being passed around periodically.

Protection Operation

Fault location can be performed through a variety of methods, which end with the device immediately upstream of the fault opening. Once a device has opened and cleared the fault, this will trigger the remaining components to be run. Common methods of locating the fault and clearing it are as follows:

<u>Time-based Coordination</u>	Devices are programmed to open when a fault is detected with increasing speed, as they get further from the substation. The benefit of time-based coordination is that it does not need communications, and only a single device needs to operate. However, only so many devices can coordinate in a row, and the further upstream the fault is, the longer the fault is present and stressing the system. Additionally, as the system gets reconfigured into new topologies, it becomes difficult to remain coordinated.
<u>Communication-based</u>	Communications are enabled between devices and their immediate upstream/downstream neighboring devices. Using a permissive or blocking scheme, the devices can coordinate to only open the one just upstream the fault. This allows coordination of virtually endless devices, however requires the effort to set up and maintain the communication network [11].
<u>Circuit Testing</u>	Each device is programmed to operate at the same time when a fault is detected and will reclose once shortly after voltage appears upstream again. The device is then set to operate quickly immediately after it closes, then more slowly once no fault appears. This method also allows coordination to occur on a large number of devices, but without communications [12]. This subjects the customers upstream to a momentary outage that the previous two methods do not, as well as incurs more operations on the devices.
<u>Sectionalizing</u>	The most upstream device on a feeder will operate for a fault, and the devices on the feeder, which in this scenario typically do not have fault interrupting capability, will sequentially operate (from the most downstream, up) depending on the number of reclose attempts of the upstream device. This solution will isolate the fault without communications but may cause more operations than necessary and can only segment the feeder into a few sections.

Isolation

Upon a device opening upstream of a fault, the devices downstream of the fault should be opened to isolate the fault. This can occur once the upstream device locks-out, or during an open interval of a reclose sequence if quicker restoration is desired. If performed before lock-out, however, the fault may be found to be temporary, and some switching operations may have been performed unnecessarily. The downstream device operation can be initiated through messaging or a timing-based method following detection of loss-of-voltage. Upon performing isolation, a message will be sent to the leader indicating that a fault has been isolated by the device and restoration can begin.

Restoration Planning

When a leader receives information on an isolated fault, having data on all the devices in its domain, the leader will calculate the best way to restore the de-energized sections. This scope of this document will not cover the details of this calculation; more information can be found in [13] or [14]. Given the loads in each section and capacities of the sources in neighboring feeders, the logic will determine what switch positions and sequence of operation will allow for the optimal sections to be restored.

Restoration Execution

Once the new configuration is identified the leader may begin to operate devices. For example, to prevent forming temporary closed loops, the leader can close open devices first, then close the tie devices. In order to close the tie devices, the leader will contact the neighboring leader, with the requested amount of load the tie will need to pick up. If the request is accepted, the neighboring leader will reserve that capacity to be used for restoration. If any of the requests are denied, the original leader will update its network data, and restoration planning is started again. In the situation where a leader denies a request, devices that operated before the request was denied may have been operated unnecessarily, increasing the restoration time and possibly decreasing asset life. To prevent this, leaders can send tie closing requests before any operations occur. Then, upon receiving all confirmations, begin device operations in the original order. If any device denies the operation command, or times-out, the leader goes back to restoration planning with the troublesome device in a fixed position. Upon completing all the device operations, the process is complete, and the new system configuration is reached. As data is continuously passed around, the system will be ready to respond to any new events that may happen in this configuration.

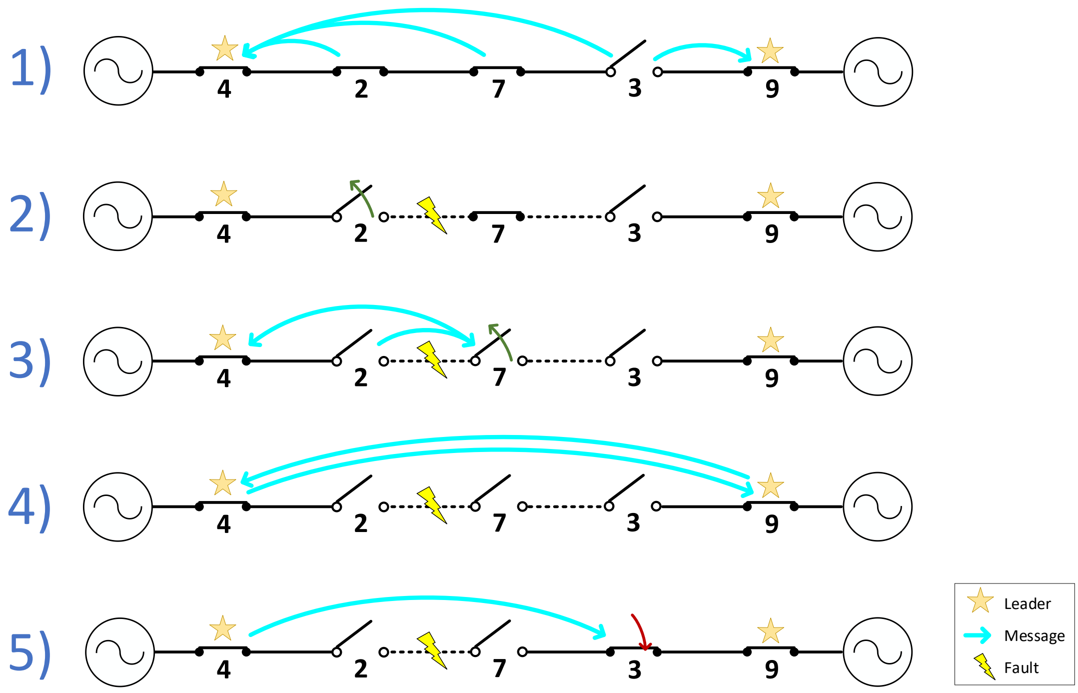

Figure 3: An example of the process of restoration under the distributed intelligence approach on a radial feeder connected to an adjacent feeder by a normally open tie point. Here a leader device is designated by the star and is device #4 in this example for the left feeder, and device #9 for the right.

a) Data propagation: In step 1 of Figure 3, data is periodically passed from each device to the leader of its respective feeder. This allows the leader to have a vision of the current conditions of the system.

b) Protection Operation: In step 2, a fault appears in the section between device #2 and #7, wherein device #2 opens to clear the fault, resulting in de-energizing the feeder downstream of the fault.

c) Isolation: In step 3, device #7 now opens to isolate the fault from the rest of the system. The trigger for the device to open is a message received from the device which operated on protection. The leader is then contacted to alert it to the fault, and to trigger restoration planning to begin.

d) Restoration Planning: Device #4 now performs a calculation, using its previously gather data on the feeder, to determine how to reconfigure the system. It determines that device #3 should close.

e) Restoration Execution: Before closing device #3, the leader must coordinate with the leader of the other feeder connected to the tie. In step 4, the leaders communicate and decide that it is acceptable to close the tie. Now device #4 is clear to close all devices, and commands device #3 to close. This results in a new system configuration, where all healthy de-energized sections are restored.

V. LEADER IN TIE

Here a new scheme of distributed intelligence is proposed, following these principles and procedures:

Principals

- Each connected region will have only one leader
- Each leader has a complete view of its connected region, and a limited view of adjacent feeders
- Leadership is assigned dynamically in response to system changes
- Leaders are assigned to open-tie switches

Procedures

- A single leader for each connected region is elected during steady state. This leader, whether predetermined or discovered using a leader election scheme, will be one of the open-tie devices at the boundary of the connected region. Criteria for selecting which tie, if there are multiple ties, can be based on: highest reserved capacity, the closest in terms of message hops, etc.
- Periodic data is sent to all open-tie devices within a feeder. This ensures that every open-tie in the feeder has the topology and current state of the feeder and can potentially act as the leader.
- As a result, every tie maintains full information of the two connected regions it is a part of. The tie will pass equivalent information about one side to the other when passing data. This allows devices in the feeder to have a view of the neighboring feeders, without passing large amounts of data around. The equivalent data will represent the neighboring feeder as if a source was on the immediate other side of the tie device.
- Upon isolating a fault, a device will elect a new leader from one of the ties downstream of itself (in the same unpowered connected region). The isolating device will wait for an acknowledgement from the new leader. If no acknowledgement is received, the isolating device will repeat the process for the next available tie, until there are no more (once all ties are exhausted, restoration cannot take place).
- This leader will serve both to calculate restoration, and to coordinate closing operation through handshaking with an adjacent connected feeder leader, if the elected tie is not also the leader of that adjacent feeder.
- Leader will contact other leaders through common tie devices, as it may not know who the leader is on the other side of an open-tie device.
- Upon completion of restoration, the leader will trigger a new leadership election.

In practice, this scheme will have devices pass around data on the feeder, and equivalents of neighboring feeders to its neighbors during steady state. This data will arrive at all the tie devices, allowing each one to assume the role of leader if needed. Once a device has isolated downstream of a fault, it should determine the tie downstream of itself as the leader. If multiple ties exist, one will be chosen based on some deterministic ranking. It should send a message to that leader indicating an isolated fault, and that restoration should be performed. Note that the original leader of the feeder is not the best candidate to receive this message, as it may no longer reside within the connected region, and does not provide additional benefits. The new leader should then make the determination that restoration can commence and run a calculation to determine the best new state of the feeder. It will coordinate the closing of ties (including itself) with the leaders of the neighboring feeder sections. Once confirmed, it will command devices to operate to reach the new state.

Several restoration processes can happen simultaneously, in the case where two or more switches are downstream of the fault. However, as the original system was radial, these must now be in separate connected regions. A leader will be elected for each of these areas and complete the restoration process independently. In case two of these areas connect to the same adjacent feeder and try to restore power from that feeder, the leader of the adjacent feeder will coordinate the restoration between the two restoring leaders.

When determining the leader device to perform the new configuration calculation and coordinating with neighboring feeders, any device within the section can perform the task. Following the principles above, and in particular selecting a tie device to be leader, has the following significant benefits compared to other implementations of the general distributed intelligence approach:

Reduced Setup Effort – Device's on the feeder do not need any user input to assign roles, such as a leader. Since the devices elect dynamic leaders, the setup effort is reduced.

Improved Speed – In a scenario where only one tie exists in a section, that tie is elected leader, and the number of messages needed to perform the restoration is reduced compared to if another non-tie device was selected. The leader no longer needs a message to contact the closing tie for permission nor to command it to close. This leads to short restoration times overall.

Increased Reliability – In the same scenario as Figure 3, there are three device's which need to be communicating in order to perform the restoration. The isolating device, the leader device, and the tie device which will close. If the leader is selected as the tie device, this reduces the number of critical devices, improving the reliability.

Increased Scalability – The data sent to the leaders (and other devices) only contains information on the current feeder, and equivalent information on neighboring feeders. Thus, regardless of how many feeders are in the network, the requirement from the communication infrastructure remains the same.

VI. EXAMPLES

Example 1

The example in Figure 4 compares the response to the same fault event using a distributed intelligence scheme when (a) the leader is set to be the closed device at the head of the feeders, and (b) the tie device.

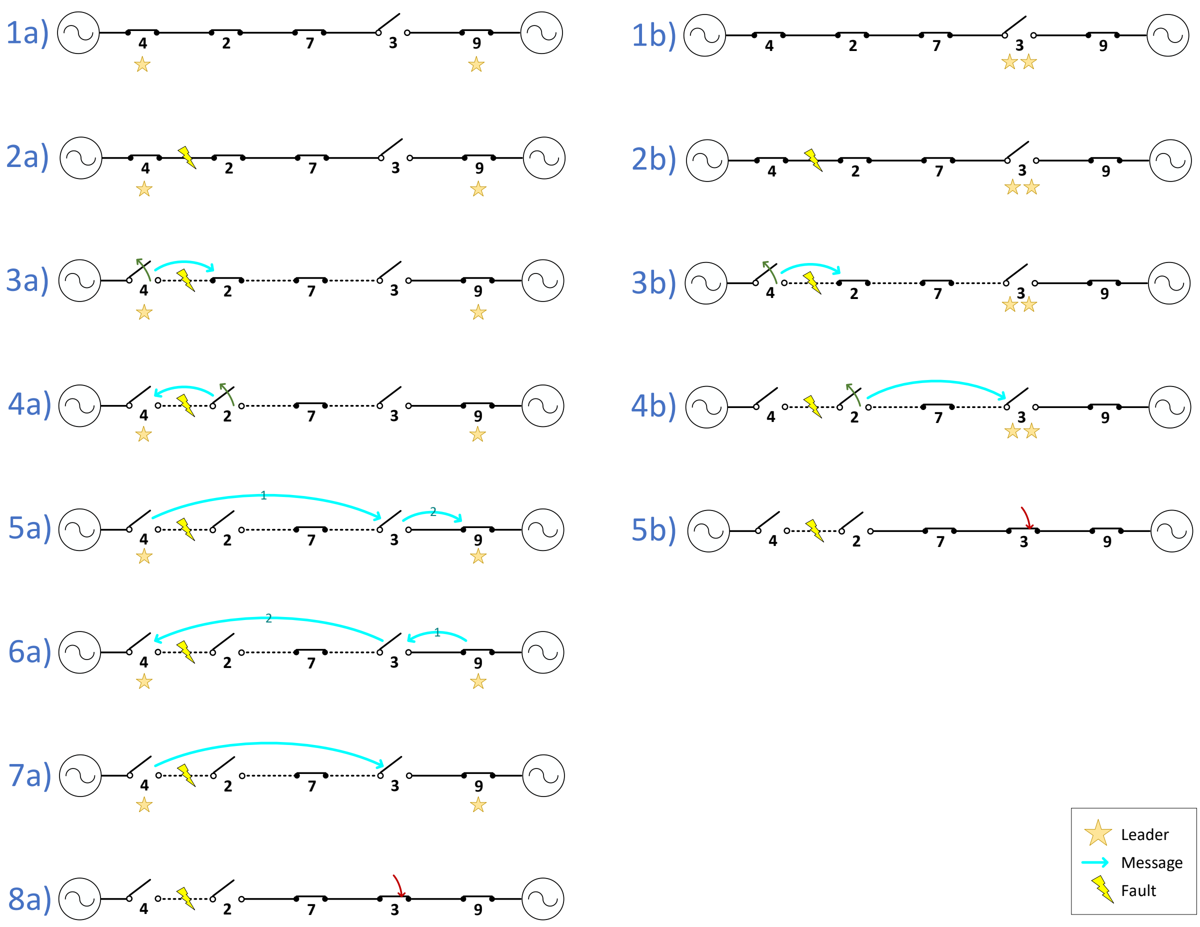

Figure 4: Diagram showing results from arbitrary leader vs tie-based leader. In scenario 'a' on the left, the restoration took place using 7 messages in 8 steps. In scenario 'b', where the leader was in the tie device, the restoration used 2 messages in 5 steps.

1) Single line diagram of two feeders, 5 switches total. One normally open point in the middle.
2) Fault appears in the section between device #4 and #2
3) Time-based protection opens device #4, de-energizing the fault and downstream sections.
4) Isolation opens device #2 which contacts its leader
    a. In scenario 'a' on the left, leader #7 is contacted
    b. In scenario 'b', the isolating device determines that the tie, device #3, should be leader and is contacted
5) The leader determines the new desired state and
    a. Contacts the tie, to ask for permission from the neighboring leader to close
    b. Gains permission from itself (as the tie and leader of the other feeder) and closes, completing the restoration.
6) The leader receives permission to close the tie (scenario 'a' only)
7) Leader commands device #3 to close (scenario 'a' only)
8) Device #3 closes, completing the restoration (scenario 'a' only)

Example 2

In the following example we demonstrate that dynamically electing leaders does not result in undesirable configurations, such as closed loops or overload conditions, even in more complex scenarios. Specifically, we consider a scenario where multiple sources are lost simultaneously, for example, because they are all connected to the same sub-transmission where a fault occurs.

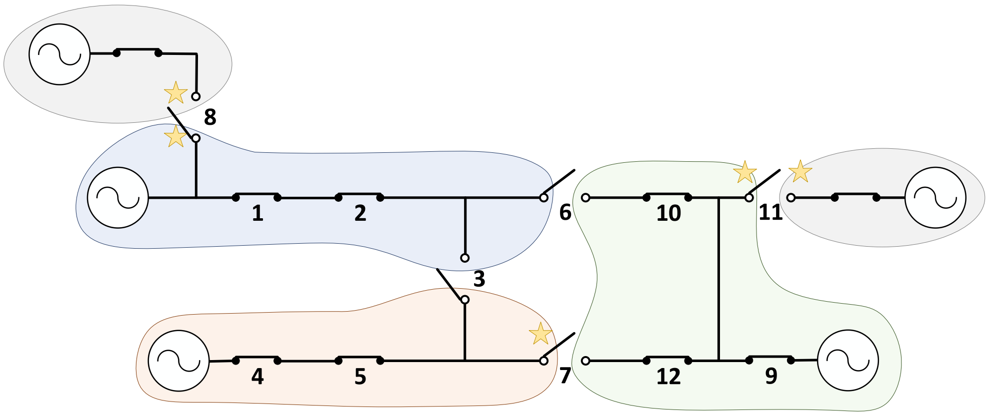

Figure 5: Network in steady state, with 5 feeders covering the area encircled. Stars are shown above open tie devices, on the side of which they are leaders of the feeder.

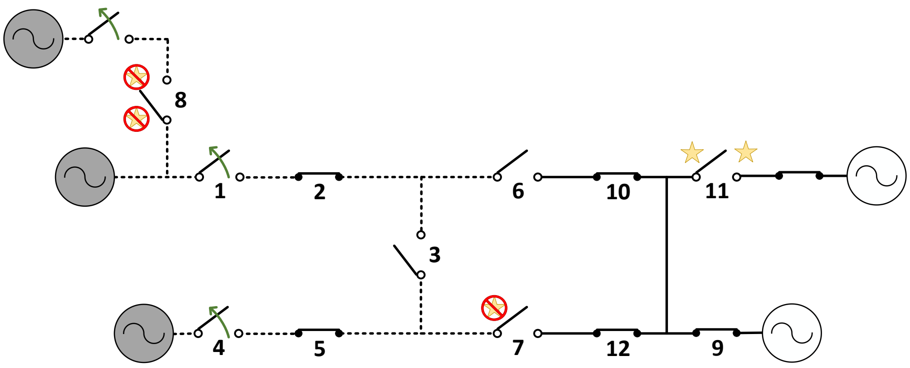

Figure 6: In this scenario the three sources on the left, colored grey, are lost simultaneously. Here the switches connected to the sources (devices #1 & #4) are set to open upon a loss of the source. Additionally, the leaders in each of the de-energized feeders sensed a loss of voltage in the system, and no longer consider their feeder available for restoring neighboring feeders. This is represented by the removal of the star icon.

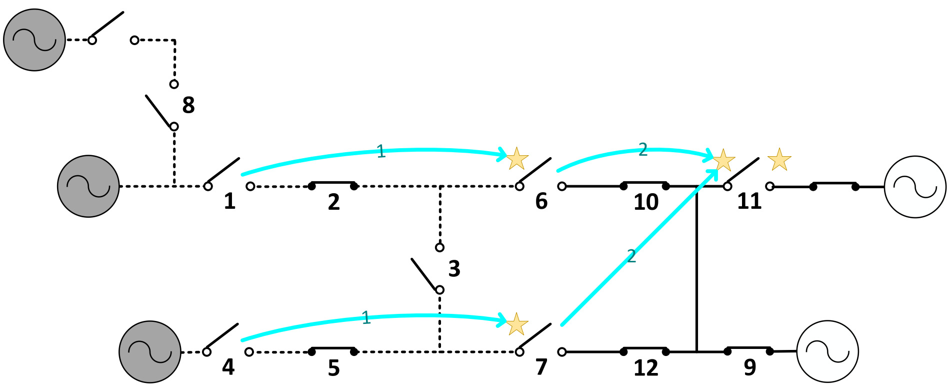

Figure 7: Devices #1 & #4 now inform the ties in their de-energized feeder that the source is isolated, and restoration can begin. Device #6 & #7 independently decide to close themselves to power their respective feeders. They both ask for permission from the leader of the powered feeder to close.

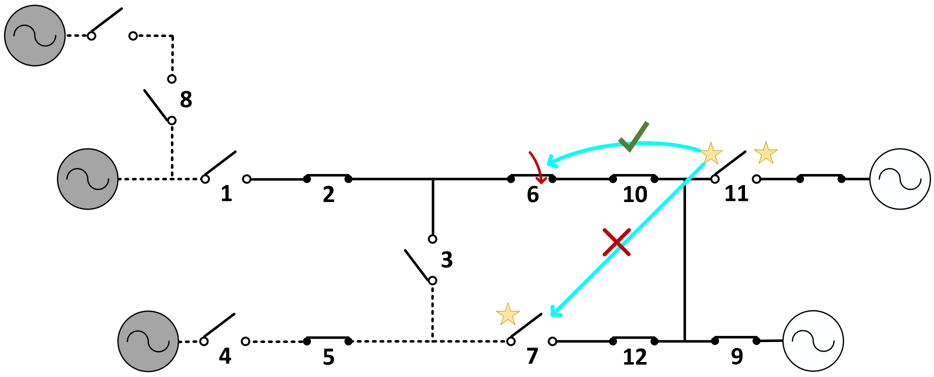

Figure 8: The leader of the powered feeder only approves one of the ties to close. It is a race condition in this example for which tie will be allowed to close, as the source may be overloaded if both are allowed to close. Here, device #6 is approved to close and does; powering the sections to its left. Device #7 is denied and remains open.

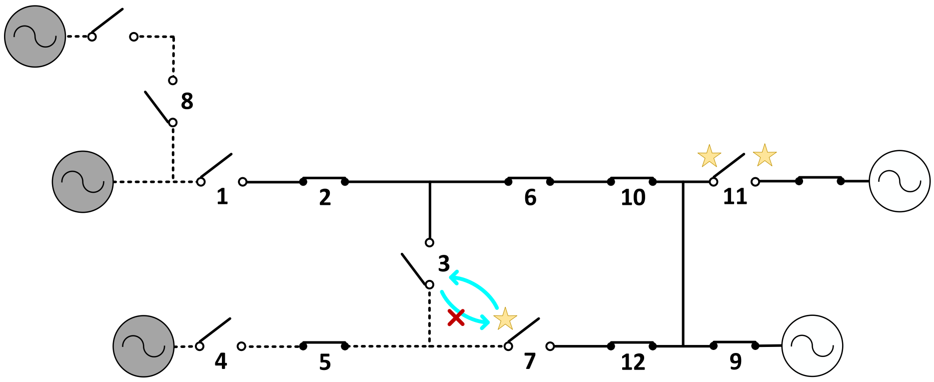

Figure 9: Device #7 knows that device #3 is a tie and attempts to get permission to close it as an alternate solution. The permission is denied by the tie, as it sensed a recent voltage loss in its feeder. Device #7 can periodically attempt to get permission for itself or device #3 to close. Once steady state has been arrived at again, and enough capacity is available, one of the devices can close.

I. CONCLUSION

This paper presented an approach to perform FLISR using the distributed intelligence approach. This approach uses local decision making and communications to locate and isolate a faulted segment. A leader of the unpowered sections then performs restoration, coordinating with leaders of neighboring feeders. Distributed Intelligence provides benefits over others in scalability, reliability, and speed, while most often reaching the optimal solution. The specific scheme described here entails dynamically electing the leaders to perform the restoration, specifically electing tie devices within the connected section that is restoring. This provides further benefits of speed and reliability to the FLISR process.

BIBLIOGRAPHY

[1] M. Sullivan, J. Schellenberg and M. Blundell, "Updated value of service reliability estimates for electric utility customers in the united states." Lawrence Berkeley National Laboratory - LBNL Report #: LBNL-6941E, January 2015

[2] State of California. Reliability Standards. (July 2019, https://www.cpuc.ca.gov/General.aspx?id=4965).

[3] Loop Automation: Improving energy availability and reducing power outages. (Schneider Electric USA - Document Number: 0107BR1401, July 2014).

[4] SEL DNA: Distribution Network Automation. (Schweitzer Engineering Laboratories, Inc., 2017, https://selinc.com/solutions/distribution-network-automation/capabilities).

[5] IntelliTeam® II Automatic Restoration System. (S&C Electric Company, Descriptive Bulletin 1042-32, June 2016, https://www.engerati.com/system/files/intelliteamr_ii_automatic_restoration_system_1042-32-060616.pdf).

[6] Electric Infrastructure and Operations. (Orange and Rockland Utilities, INC, January 2018, http://documents.dps.ny.gov/public/Common/ViewDoc.aspx?DocRefId={38E832C4-4EFC-4D44-9192-5BC619180B12}).

[7] Petition of Duquesne Light Company for Approval to Modify its Smart Meter Procurement and Installation Plan. (Pennsylvania Public Utility Commission, P-2015-2497267, 2015, http://www.puc.pa.gov/pcdocs/1485531.docx).

[8] Evolve at the Pace of Change: Optimizing the Distribution Grid. (General Electric, July 2018, https://www.ge.com/content/dam/gepower-pw/global/en_US/documents/software/adms/Distribution%20Optimization%20White%20Paper%20_07.13.18%20.pdf).

[9] R. G. Gallager, P. A. Humblet and P. M. Spira, "A distributed algorithm for minimum-weight spanning trees" (ACM Transactions on Programming Languages and Systems, vol. 5(1), January 1983, pages 66–77)

[10] E. Korach, S. Kutten and S. Moran. "A Modular Technique for the Design of Efficient Distributed Leader Finding Algorithms" (ACM Transactions on Programming Languages and Systems. vol. 12(1), January 1990, pages 84-101)

[11] M. Yalla, *et al.*, "Application of peer-to-peer communication for protective relaying" (IEEE Transactions on Power Delivery, vol. 17(2), April 2002, pages 446-451)

[12] J. Gates, "ComEd advances system reliability" (T&D World, September 2013, https://www.tdworld.com/distribution/comed-advances-system-reliability)

[13] R. A. Jabr, R. Singh and B. C. Pal, "Minimum loss network reconfiguration using mixed-integer convex programming" (IEEE Transactions on Power Systems, vol. 27(2) , May 2012, pages 1106-1115)

[14] D. Elizalde, D. Staszesky and M. Meisinger, "Use of distributed intelligence for reliability improvement using minimum available distribution assets" (2006 IEEE/PES Transmission & Distribution Conference and Exposition: Latin America, 2006, pages 1-6)