Grid Fault Recovery and Resilience: Applying Structured Energy and Microgrids

William Cox, *Member, IEEE*, and Toby Considine, *Non-Member*

***Abstract*-- We show how transactive techniques and Structured Energy address fault containment, resilience, and recovery.**

By applying the structured energy description of a grid as a topology of microgrids, including nesting, union, and intersection, we demonstrate a new way to understand fault resilience, containment, and recovery. Factoring operations management into microgrid management and cross-microgrid management reduces the complexity of each.

We show how structured energy enables flexible changes in configuration as energy needs and supplies evolve, fail, and are restored, thus improving resilience in the face of failures.

We suggest architectural approaches to grid configuration to improve fault resilience, including following geographic distribution of energy resources with distributed and transactive operation of those resources.

***Index Terms*—microgrids, power system faults, fault tolerance, redundancy, fault tolerant systems topology, aggregation, transactive operation**

I. NOMENCLATURE

We use the following terms in this paper:

Topology
Energy Grid
Microgrid
Consumption-Only Microgrid
Supply-Only Microgrid
Transactive Operation
Structured Energy
Aggregating and disaggregating microgrids

II. INTRODUCTION

THE focus of this paper is on how to apply well-tested technique from other areas and architectural understanding from Structured Energy [1] to electrical grid fault containment, resilience, and recovery.

As we saw in the United States in the aftermath of Hurricane Sandy, energy resources that were available could not be used, and the recovery of broad geographic areas of the power grid took weeks rather than hours or days. Structured Energy show ways to mitigate such losses and delays in a consistent, automatable framework.

The recursive nature of Structured Energy and related transactive operation allows application at different granularity, from facility components to broad grid areas at collections of distribution grids and substations.

We use the approach of [2] in taking microgrids as the fundamental structure of the smart grid, and aggregate, disaggregate, and change configuration of a set of microgrids applying the architectural understandings of structured energy.

We consider actions of microgrids with respect to inter-microgrid energy flows, rather than how microgrids are implemented and deployed. This is analogous to Service-Oriented Architectures [3], where *what* is requested is key, rather than *how* it is accomplished.

III. BACKGROUND

In this section we briefly summarize relevant aspects of Structured Energy and Transactive Operation [1].

A. Structured Energy

Structured Energy defines a graph of microgrids, each with its own balance of supply and demand, where relevant. We extend common definitions of Microgrids (e.g. [4]) to encompass *consumption-only* and *supply-only* microgrids.

In that paper we showed how microgrids form a topology over their components. This understanding allows us to describe dynamic evolution and configuration where we can *aggregate* microgrids[1] to create larger ones and can *disaggregate* a microgrid to define sub-microgrids.

Different sequences and subjects for aggregation (and the mirror disaggregation) can define a set of microgrids with the same components. We call these *aggregation* (or *disaggregation*) *paths*, which we will generally simply call *paths.* We use the existence of different paths to restore service as well as to reconfigure to extend service and integrate additional microgrids.

Not all paths are feasible; both electrical connectivity and communication connectivity is required for transactive operation and for aggregations.

We extend this discussion of relevant aspects of Structured Energy as applied to fault resilience in Sections 7 and 8.

B. Transactive Operation

Transactive operation [1] uses transactive techniques to

W. T. Cox is with Cox Software Architects LLC, Summit NJ 07901 USA (e-mail: wtcox@coxsoftwarearchitects.com).

T. Considine is with TC 9, Inc, Pittsboro NC USA (e-mail: toby.considine@gmail.com).

[1] The base for the inductive definition is that a single component can be treated as a consumption-only or supply-only microgrid.

balance supply and demand operationally as well as through forward or future contracted use and supply. Transactive operation therefore uses actors that can actively balance business needs, energy costs, and availability to meet their goals.

We use transactive operation to aggregate and disaggregate microgrids: the behavior of each component and each microgrid is determined by the availability of energy at cost-effective prices, when considered in the business context of the participant.

Transactive techniques range from pure market-based interaction (see e.g. [5] and [6]) to transactive controls [7] to fiat prices that manage expected response through the economic demand and supply curves (see e.g. [8]). For this paper, independent operation using transactive techniques and principles is assumed; specific transactive mechanisms are of limited relevance.

IV. DEFINITIONS

In this section we state definitions specific to this paper.

A. *Fault*

We define a fault as failure of a particular microgrid or component to fulfill its obligations, including but not limited to supply or consumption of energy. We do not address partial failure, e.g. a restriction in possible consumption or supply.

B. *Fault Containment*

A fault is contained if the rest of the system, broadly considered, can function effectively in areas not affected by the fault.

Fault tolerance encompasses any or all of fault containment, resilience, and recovery.

C. *Fault Resilience*

Fault resilience describes "…the tendency or ability to spring back and … recover to normality after a disturbance." More broadly and *a propos* our purposes, "… [an] enhanced ability to deal with the unexpected." [9]

In this paper we address the ability of a grid comprised of multiple microgrids and components to provide energy services in the face of failures, whether caused by natural disturbance, component or equipment failure, or other causes.

V. A STRUCTURED ENERGY VIEW OF A GRID

A. *A Grid is a Directed Graph*

As in [1] we consider the *paths* in aggregating or disaggregating a grid to define a directed graph for a particular configuration, with the edges from the containing to the contained microgrids. See **Figure 1**.

We will say that *looking up* in such a Structured Energy graph is (from a particular microgrid) looking toward containing microgrids; conversely *looking down* is looking toward contained microgrids.[2]

In Figure 2 looking up from microgrid M is toward the containing microgrid G at the top. From M we look down to the contained microgrids including S. Note that G likely also participates in a broader grid—the relationship applies at any scale.

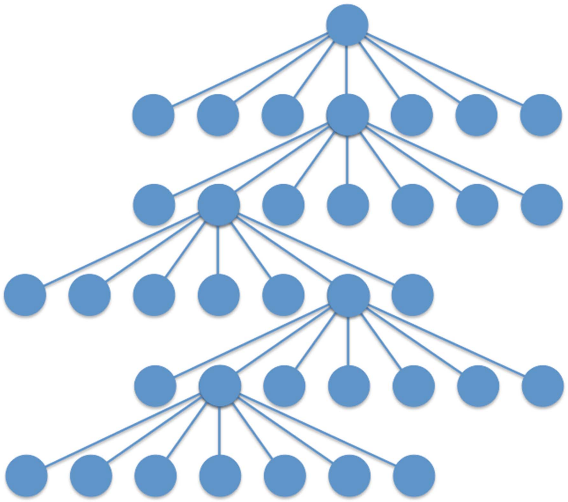

Figure 1 Structured Energy View of a Grid

B. *Aggregation, Disaggregation, and Reconfiguration*

Considering a directed graph such as that in **Figure 1**, aggregation with respect to a microgrid M reflects how one built or configured the transitive closure looking down in the graph. Disaggregation likewise addresses all feasible paths that could have been taken to assemble a specific microgrid M, of which the current graph is one.

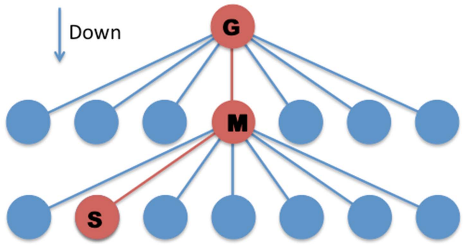

Figure 2 Microgrid M looks up to Grid G and down to smaller (contained) microgrid S

Note that our figures show a specific configuration—we speak of disaggregating and re-aggregating or *reconfiguring* a particular microgrid as identified by the components in the transitive closure of the graph, and address a variety of fault and recovery situations. This reconfiguring in the absence of faults is merely a different organization of the components and containing microgrids for a specific set of components.

VI. FAULTS AND GRIDS

We simplify the discussion by treating a single fault within a specific grid (or, interchangeably, microgrid) G in order to make the algorithms more clear.

In **Figure 3** a fault is detected with respect to microgrid M (indicated by a heavy black circle inside the ellipse), and the

---

[2] Again, a single component, for closure, is considered a consumption-only or supply-only microgrid.

other red microgrids are those actually or potentially affected by the fault.

In this paper we assume that G is stable prior to the fault, the single fault occurs, and recovery is applied to the consequences of that fault. See Section XII.

Faults may be of many types; we describe two classes. Communication failure may cause a microgrid to be unable to participate in its containing microgrid. Similarly an energy line failure may prevent exchange of the net difference of energy in a structured environment. [1] [10]

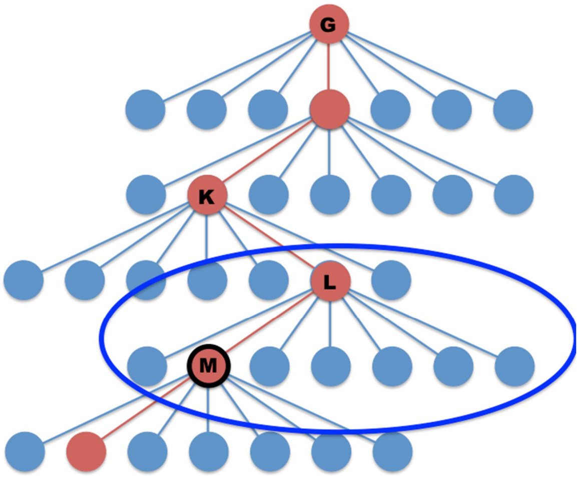

Figure 3 Fault Showing potentially affected microgrids in red. The failed microgrid M is marked with a heavy black circle.

VII. An Algorithm for Fault Containment

Again consider **Figure 3**, where microgrid M is the detected point of failure. Other red microgrids are affected by that failure—the containing microgrid has lost a component, and so forth. The cause of the failure of M might be in the red microgrid below; we will return to this issue.

Containment of faults dictates that the nodes beyond the marked layer in the graph are affected as little as possible. So microgrid L contains the fault if it can operate without one of its components, which depends on that containing microgrid's energy balance and operating range.

If L cannot operate, we progress up the graph until we find a containing microgrid that can operate normally, say first to L then to microgrid K. We have thus contained the fault in L by finding a containing microgrid K that can operate normally. But there's another side to containment—microgrids and their components within the containing microgrid L.

By assumption L cannot operate with M failed; this reflects on the fragility of the design or configuration for the specific configuration and paths; this is the subject of another paper in preparation.

By assumption, the microgrids looking down from L can and will operate independently unless, of course, they have also failed. But our assumption is that there is a single failure.

For continued operation, the primary concern would be whether a give microgrid has sufficient consumers and suppliers to balance, and to what extent. We suggest that the six siblings of M can continue to operate subject to such operational constraints.

Now consider the microgrids next down from M. M cannot operate due to the fault. But can the microgrids in the lowest level continue to function? Why has M failed?

In the event of coordination failure (due to communication or failed or partially failed cross-links) the lowest level microgrids can continue to function on their own, perhaps suboptimally—else why participate in the aggregation M—but perhaps a single failure in microgrid N has driven the failure of M. In that case, the remaining blue microgrids on the lowest level can continue to operate independently (or perhaps with degraded coordination through the aggregation M).

In this manner we limit the failures to the three microgrids that have actually failed, and allow the remaining 11 microgrids to continue to operate. We discuss architectural and structural issues in Section IX.

We have built up and down from stable microgrids to limit the scope of the fault. In the next section we address recovery.

VIII. An Algorithm for Fault Recovery

We suggest that a primary objective is to maintain energy supply as broadly as practicable in the face of failure.[3]

We have maintained operation, perhaps independent of the original grid G, for all but the actual failed grids. This is a significant step forward, as a common implementation without structured energy would likely address the stranded consumers by repairing inward until full service is restored. Our approach is more self-organizing, and can rapidly build both in and out from the fault to restore energy connectivity.

How can we recover from the fault in M? In a Structured Energy approach, we either disaggregate (separate down) or aggregate (build up) for each set of microgrids. **Figure 1** shows one grouping of component microgrids into larger ones; there are, subject to communication and energy flows, potentially many such groupings.

Consider microgrid L in **Figure 3**. The relevant components, recalling that microgrids form a topology over the set of components, are those presently in the transitive closure of the constituents of L.

While there may be exactly one path, as shown in the Figure, with care in design and architecture to define redundant pathways, there may be many ways to disaggregate downward toward the components, that is, to the next level of microgrids.

So we use the microgrids defined, and look at alternate aggregations that can in effect route around the failure. In **Figure 3** we are concerned with re-integrating the microgrids that are siblings of M. These might be connectable to K, to one of L's siblings, and so forth on up the graph, or disaggregated at a lower level and then re-aggregated.

The feasible alternate connections can be computed in advance, making the complexity of restructuring essential the order of a constant for a given failure. Graph exploration algorithms are well known, and the exploration can be pruned by actual feasible connections—and the design of redundant con-

---

[3] Of course, "as broadly as possible" has many definitions.

nections can be guided by this simple algorithmic approach to fault recovery.

The re-integration into G of disconnected microgrids involves three steps:

1. Consider alternate aggregations and determine which one(s) allow reintegration subject to business criteria such as number of connected grids, number of components, customers, etc.
2. Select an alternate aggregation that represents an improvement on the situation with the fault.
3. Reintegrate the component microgrids of L and M with the respective un-faulted microgrids.

In essence, this algorithm allows simpler use of redundant pathways, guidance on defining and implementing redundant pathways, and permits pre-computation of alternate configurations in a straightforward way.

We note that energy connectivity is most critical to the potential restructuring. Information connectivity is relatively simpler and less expensive, and completed standards in deployment such as OASIS Energy Interoperation [11], OASIS Energy Market Information Exchange [12], and the OpenADR2 Profiles [13] of both can rapidly and easily aggregate microgrids with a simple change to a single URI.

IX. ARCHITECTING FOR FAULT RESILIENCE

Building fault resilience through structured energy leads to a number of architectural understandings.

First, distributed resources must be distributed and diverse, both geographically and in control and management.

Creating a parallel graph with Distributed Energy Resources (DER) managed at a high level allows for greater failure possibilities, and for less resilience as we've defined. The microgrid connection of DER at a low level in the graph permits use during disconnected operation as well as for limiting the scope of failure.

In other words, if DER is attached only to a high level grid, the same failure modes apply, but DER can (where still connected) be used only for resilience at the high level. DER should be distributed through the various levels of a reconfigurable grid.

Such *distributed* energy resources must be not only geographically distributed, but distributed in management and transactive operation, in order to contribute to recovery and support resilience. Focus only on geographical distribution under single management or control leads to an inability to support users of energy during fault recovery.

Second, multiple connections for both communication and energy flows increases the number of possible structures below a particular microgrid. It's a truism that multiple connectivity, as e.g. in packet store-and-forward networks, allows greater reliability. By creating multiple paths for energy transfer among components the entire structure can be more reliable. We have shown specifically how to take advantage of the reliability improvement from redundant connection.

Third, it reduces the risk to the larger grid of introducing new technologies and interactions. Structured energy limits the interactions between components of a microgrid to other components within the same microgrid. Whatever the diversity within a microgrid, it is hidden from the interaction of that microgrid as a component within the next level within the structure. Containment and distribution apply not just to energy resources but also to technology evolution.

X. DISCUSSION

The approach and techniques described here are well known and well tested in other areas such as in network fault detection and data communications path management. The application of these techniques to power distribution fault resilience, and to energy surety is new.

The algorithms we have described take no account of the underlying technology of a component, nor would they be improved by doing so. A component may need to understand its own processes and capabilities to participate, but its peers within a microgrid need not. The complexity and diversity within a component are irrelevant to fault detection and recovery within a microgrid. The algorithm relies on the effect of the underlying fault, and how the microgrids recovers from that effect. This is true whether the component implements existing technology, novel technology or is itself a microgrid.

The effect of this non-determinism about the underlying technology is that this approach can be built out without knowing in advance what technologies will be implemented in each component. This approach is pre-adapted for adoption of novel technologies into each component of a microgrid. In particular, this model addresses the growing concerns or consumer choice and adoption associated with distributed energy resources.

The second characteristic of the algorithms is their simplicity. Because it does not require variations based on technology choices in each component, it can be easily implemented using a small amount of code. This approach lends itself well to implementation in application specific integrated circuits (ASICS) so the cost of software on each component can be minimal. This means that not only does this approach scale up, as described above, but it also scales down. Even the smallest component system can use this approach internally—there is no barrier to using this approach even within a residential microgrid, in which the components are appliances, local generation, and local storage.

XI. SUMMARY AND CONCLUSIONS

Microgrids as described by Galvin and by Structured Energy are recursive systems of systems. The external characteristics of a microgrid are the same, whether that microgrid encompasses a home, a factory, a neighborhood, an industrial park, a city, or a region. A grid built from microgrids gains resilience from the diversity of its components.

Such diversity means that systems that manage a grid must be either very complex or very simple.

More complex models would require extensive integration with the introduction of each new component. The complexity of integrating each new technology is at least geometric. Today's grids manage this complexity by limiting diversity of technology. This leads to extensive testing of each new technology and locks out all technology providers but those with the deepest pockets.

This paper describes a simple model, and assumes the principles of service-oriented integration. Because the algorithms provide reliability and resilience without consideration of the underlying technology, the model defines micrograds that can accept diversity of new technologies while requiring minimal integration. New technologies and new component providers can be easily introduced without changing the underlying systems and network management.

XII. FUTURE WORK

The awareness to detect and isolate faults can be used to recover and re-aggregate. In this paper recovery addresses the problem of a microgrid detecting a failure, and creating one or more new stable systems. Recovery more broadly includes re-incorporation of previously failed nodes back into the now stable grid. A more detailed model for such recovery will build on this work.

Our model for recovery is moreover close to a solution for graceful insertion, that is, for adding additional nodes to an existing dynamic system. Perhaps some additional registry of available services would assist in such insertions, but how small can that registry be? Can these techniques reach so far as to incorporate, say, a new fuel cell into an active microgrid without configuration, and without losing the capability of fault detection we have described? Further work is required to develop and describe these modalities.

XIII. BIBLIOGRAPHY

[1] William Cox and Toby Considine, "Structured Energy: Microgrids and Autonomous Transactive Operation," in *Innovative Smart Grid Technologies*, 2013, Available from IEEE and at Authors' web sites.

[2] Toby Considine, William Cox, and Edward Cazalet, "Understanding Microgrids as the Essential Smart Energy Architecture," in *Grid-Interop*, 2012, Available from Grid-Interop and Authors' web sites.

[3] OASIS, "Reference Model for Service Oriented Architecture 1.0," 2006. [Online]. http://docs.oasis-open.org/soa-rm/v1.0/

[4] Galvin Electricity Initiative. Smart Microgrids. [Online]. http://www.galvinpower.org/microgrids

[5] Edward G. Cazalet, "Automated Transactive Energy (TeMIX)," in *Grid-Interop*, 2011.

[6] GridWise Architecture Council. (2013) Transactive Energy Conference. http://www.gridwiseac.org

[7] S Katipamula, D Hatley, D. Hammerstrom, D. Chassin, and R. Pratt, "Transactive Controls: Market-Based GridWise Controls for Building Systems," Pacific Northwest National Laboratory, Technical Report 2006.

[8] Chellury Ram Sastry, William Cox, and Toby Considine, "The Semantics of Price and of Price-Responsive End Nodes," in *Grid-Interop*, 2012, Available from Grid-Interop and Authors' web sites.

[9] Laurenzo Strigini, "Fault Tolerance and Resilience: Meanings, Measures and Assessment," in *State of the Art Report Assessing Measuring, and Benchmarking Resilience Coordination Action*. London, UK: City University, 2009, http://amber.dei.uc.pt/stateart.php.

[10] Toby Considine and William Cox, "Smart Loads and Smart Grids-Creating the Smart Grid Business Case," in *Grid-Interop*, 2009, Available from Grid-Interop and Author's web sites.

[11] OASIS, "Energy Interoperation 1.0," February 18, 2012. http://docs.oasis-open.org/energyinterop/ei/v1.0/energyinterop-v1.0.html

[12] OASIS, "Energy Market Information Exchange Version 1.0," January 11, 2012. [Online]. http://docs.oasis-open.org/emix/emix/v1.0/emix-v1.0.html

[13] OpenADR Alliance. http://www.openadr.org/

XIV. BIOGRAPHIES

William Cox (M2009) is a leader in commercial and open source software definition, specification, design, and development. He is active in the Smart Grid Interoperability Panel and is a member of its Smart Grid Architecture Committee. He is co-chair of the OASIS Energy Interoperation and Energy Market Information Exchange Technical Committees, past Chair of the OASIS Technical Advisory Board, and a member of the WS-Calendar and oBIX Technical Committees. He participates in related IEC and ASHRAE standards work.

He has developed enterprise product architectures for Bell Labs, Unix System Labs, Novell, and BEA, and has done related standards work in OASIS, the Java Community Process, Object Management Group, and IEEE.

He earned a Ph.D. and M.S. in Computer Sciences (minor in Electrical and Computer Engineering) from the University of Wisconsin-Madison.

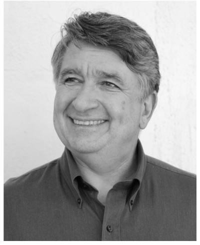

Toby Considine is recognized thought leader in applying IT to energy, physical security, and emergency response. He is a frequent conference speaker and advises companies and consortia on new business models and integration strategies.

He has supported and managed interfaces to and between buildings, cogeneration plants, substations, chilled water plants, and steam and electrical distribution. This work led to his focus on standards-based enterprise interaction with engineered systems in buildings, and to his work in OASIS. Toby is chair of the OASIS oBIX Technical Committee, a web services standard for interface between building systems and e-business, and of the OASIS WS-Calendar Technical Committee.

Toby has been part of national smart grid activities since delivering the plenary report on business and policy at the DOE GridWise Constitutional Convention in 2005. He is a member of the SGIP Smart Grid Architecture Committee, and is active in several Smart Grid Domain Expert Workgroups.

Before coming to UNC, Mr. Considine developed enterprise systems for technology companies, manufacturing plants, architectural firms, and media companies old and new.