Scalability, Resilience, and Complexity Management in Laminar Control of Ultra-Large Scale Systems

Jeffrey Taft, PhD
Connected Energy Networks
Cisco Systems, Inc.

Paul De Martini
Resnick Institute
California Institute of Technology

Abstract

Ultra-large scale control systems that are based on layered optimization decomposition and Network Utility Maximization (NUM) have structural properties and operational modes that give such control systems useful scalability and anti-fragile properties. We denote control architectures based on these principles as Laminar Control and we suggest that the structure so defined has properties such as layer-by-layer segmentation of control signal traffic, abstraction of local state into scalar signal flows, self-similarity of data flow patterns at each layer, and support for islanded operation and re-connection.

While the underlying physical networks (such as power systems) and associated communication networks are not necessarily self-similar, the imposition of the Laminar Control paradigm allows the creation such a structure at the application node level. Consequently, instead of having a set of agents with randomly structured logical connections, the optimization nodes become hubs and furthermore, it may be possible to restrict peer-to-peer communication at any level in the hierarchy to delimited specifiable domains. State determination can be similarly partitioned so that domain state need not be shared globally.

In combination, these properties can invest distributed control networks with three valuable characteristics: scalability of control system real time network communications, resilience of the logical control network, and complexity bounds.

Introduction: Laminar Control Architecture and Data Flows

Distributed and hierarchical control methods have been available for decades. In the case of electric power systems, some of these concepts have been used in portions of the grid, but not as an Ultra-Large Scale (ULS) control. By ULS we mean the concept developed at the Software Engineering Institute to describe extremely large, complex systems with the following characteristics:[1]

- Decentralized data, development, and control
- Inherently conflicting diverse requirements
- Continuous (or at least long time scale) evolution and deployment
- Heterogeneous, inconsistent, and changing elements
- Normal failures (failures are expected as a normal part of operation)

In the power grid domain, certain approaches to wide area grid control have employed a single physical variable that presumes to characterize a key aspect of system state. At the transmission and generation level, system frequency is used for this purpose. It is widely used in incremental area balancing via Area Control Error and Automatic Generator Control (AGC).[2] System frequency has also been proposed as the basis for control of large number of responsive loads not owned by the electric utility.[3]At the distribution level, feeder voltage is used in Volt/VAr control systems, both in those that act centrally and those that are composed of a collection of independent agents.[4] Such methods have enjoyed a degree of success but encounter difficulties in two areas:

1) Such systems can become unstable due to feedback through the grid itself
2) When multiple functions want to use the same infrastructure for possibly competing or even conflicting purposes, as is happening with distributed

---

[1]Peter Feiler, John Goodenough, et al., Ultra-Large-Scale Systems The Software Challenge of the Future, Software Engineering Institute, June 2006

[2] NERC Resources Subcommittee, Balancing and Frequency Control, available online: http://www.nerc.com/docs/oc/rs/NERC%20Balancing%20and%20Frequency%20Control%20040520111.pdf

[3] PNNL Staff, Grid Friendly Controller Helps balance Energy Supply and Demand, available online: http://readthis.pnl.gov/MarketSource/ReadThis/B3099_not_print_quality.pdf

[4]Naveen Venkatesan and S K Solanki, Coordination of Demand Response and Volt/VAR Control Algorithm Using Multi-Agent System, IEEE 2012 Transmission and Distribution Conference and Exposition, 2012 IEEE PES, Orlando, Fla. May 2012.

energy resources,[5] then a single variable is not sufficient to enable proper coordination or federation of the multiple control processes involved

Recently a method based on Network Utility Maximization[6] and Layering for Optimization Decomposition[7] has emerged as a distributed control paradigm for ultra-large scale controls, especially for power grid control.[8] In this approach, a structured network of optimization nodes communicates hierarchically via scalar signals to cooperate in solving a joint optimal control problem. The nodes are arranged in a manner that corresponds to a layered decomposition of the full optimization problem. This layered decomposition can be mapped onto the structure of a physical system such as a power grid to solve the control issues of federation, disaggregation, and constraint fusion while allowing for local "selfish" optimization. This approach is a hybrid of central and decentralized control, made distributed by virtue of the various parts engaging in cooperation to solve a common problem: that of grid control. In the layered decomposition approach Network Utility Maximization is used to provide overall coordination, with most nodes acting as both a coordinator for sub-problems at the next tier below, as well as a sub-problem solver for the coordinator for the next tier above.

The optimization decomposition consists of breaking the problem at each layer into a master problem and set of sub-problems, using iteration to decouple the constraints. The master problem and sub-problems cooperate by exchanging simple signals which, depending on the decomposition type, can be seen as either resource allocation signals (primal decomposition) or prices (dual decomposition). Problems may be decomposed recursively, leading to a multi-layer structure that inherently supports distributed computation. When this type of mapping is done for hierarchical controls, especially for power grids, it leads to a logical data flow structure that has useful characteristics.

---

[5] P. De Martini, DR 2.0:A Future of Customer Response, available on line: http://www.demandresponsesmartgrid.org/Resources/Documents/FINAL_DR%202.0_13.07.08.pdf

[6] Mung Chiang, Steven Low, et al., Layering as Optimization Decomposition: A Mathematical Theory of Network Architectures, Proceedings of the IEEE, Vol. 95, No. 1, January 2007.

[7] Daniel P. Palomar and Mung Chiang, A Tutorial on Decomposition Methods for Network Utility Maximization, IEEE Journal on Selected Areas in Communication, August 2006,pp. 1439-1451.

[8] Jeffrey Taft and Paul De Martini, Ultra-Large Scale Power System Control Architecture A Strategic Framework for Integrating Advanced Grid Functionality., available online: http://www.cisco.com/web/strategy/docs/energy/control_architecture.pdf

There are three principal data flows that matter here:

1. “Vertical” signal exchange among optimization nodes for multi-tier coordination
2. Hub and spoke signal exchange among master and sub-problems nodes for tier level coordination
3. Peer to peer flows within a layer for local cooperation

Figure 1 shows these primary data flows. Note that the first flow in the list is essentially a subset of the second on the list, but it is helpful to consider an entire vertical optimization signal chain as one of the flows. We view this model as having two axes of distributed intelligence and refer to the vertical single flow and optimization chain as coordination distributed intelligence, and the tier hub-and spoke and peer-to peer interactions as application distributed intelligence. The overall vertical flow across all tiers with optimization node operation is denoted as Deep Area Coordination (DAC).

Note that in this method we are explicitly providing a multi-tier and intra-tier coordination mechanism for control. This is a “team” approach, wherein explicit deterministic communication is used to implement coordination, as opposed to swarm or flock behavior, as exhibited in some biological systems and as mimicked in some agent-based software systems.[9] However, the potential advantage of flock behavior in avoiding hazardous regions of operation via trajectory steering can be achieved in the Laminar Control framework in a manner that is well-behaved and manageable by human grid operators who are supervising the overall operation of the control system but are not “inside the control loop.” By setting and adjusting optimization criteria and constraints, humans supervising the control system can guide the system to avoid entering regions of reduced stability margin. Doing this via a few top level control nodes (as opposed to adjusting all of the nodes in the entire system) is roughly the equivalent of steering the trajectory of the dynamic system (the grid) in state space; adjusting goals at a few nodes is essentially the same idea as leading the “flock.”

[9] Felipe Cucker and Steven Smale, The Mathematics of Emergence, Japanese Journal of Mathematics, October 2006. Available online: http://ttic.uchicago.edu/~smale/papers/math-of-emergence.pdf

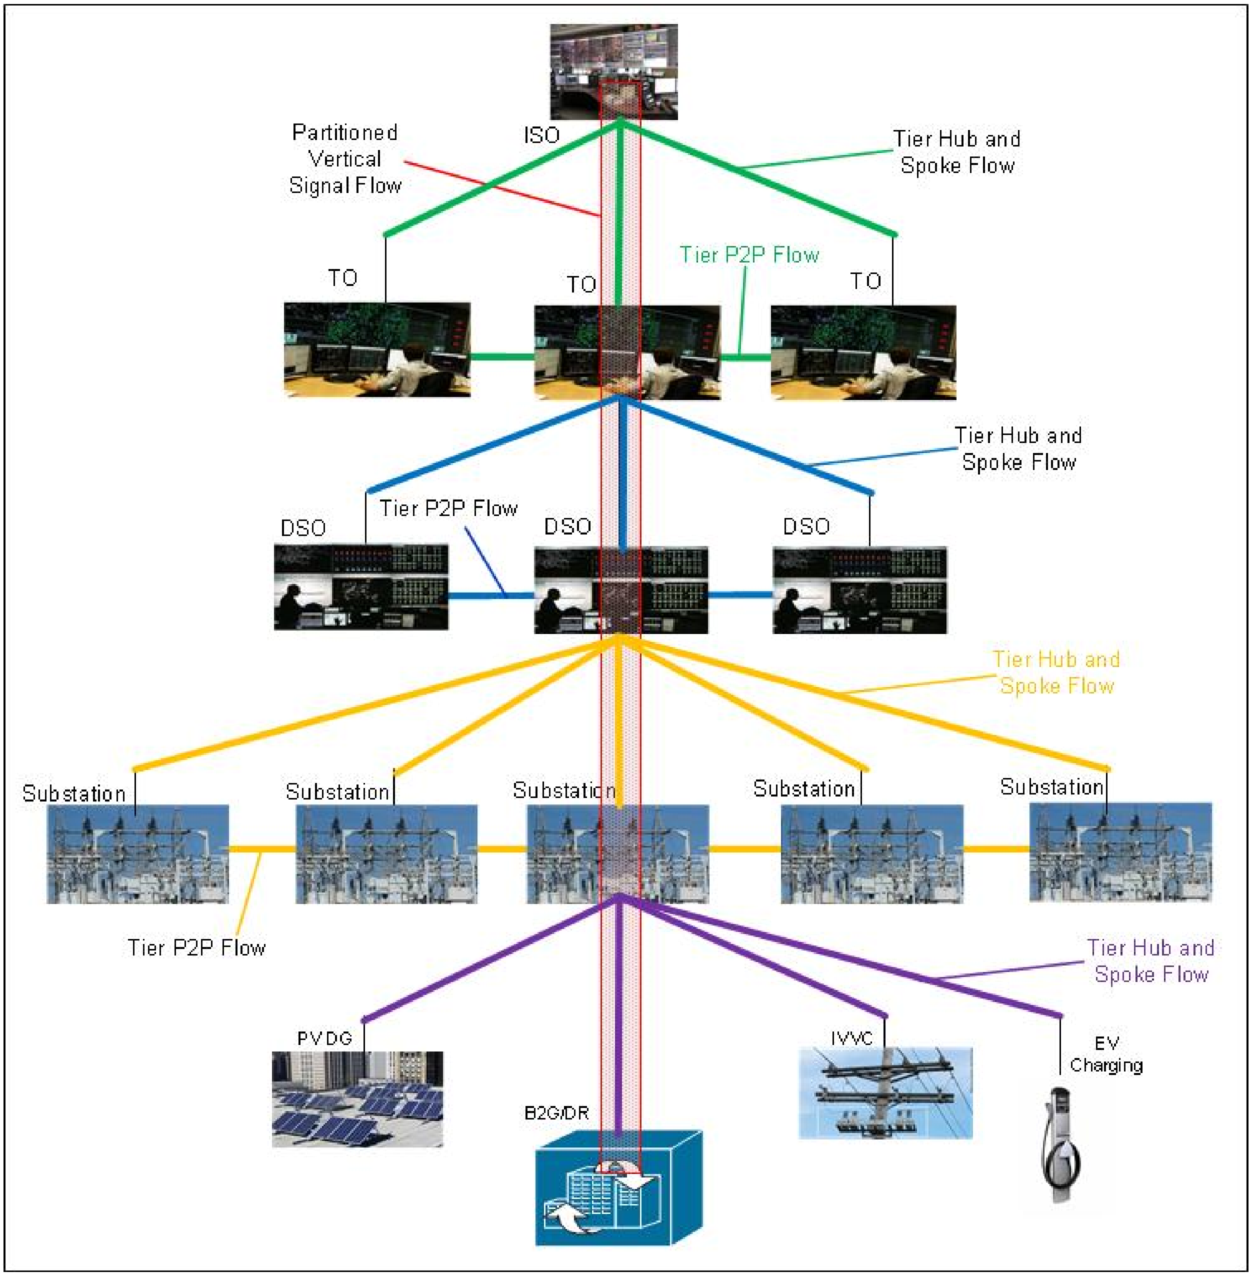

Figure 1 Primary Data Flow Patterns in Laminar Control for Power Grids

Note that we decompose the optimization problem by tiers that match physical power grid tiers. In addition, the elements are broken up at each tier into domains, with the definition of a domain depending on the level in the power grid hierarchy. Domain definition is flexible providing the system designer with freedom to bound complexity. At each domain in the logical control architecture, there is a computational element that solves the optimization problem for that domain, which means that in general each domain optimizer simultaneously acts as a sub-problem solver for the level above and as a master solver for the level below. Adding tier coordination functions to the vertical signal flow leads to the concept of a simple coordinator for each node, where the optimizations are calculated and where orchestration for the sub-tier is handled. Figure 2 illustrates a domain and coordinator structure.

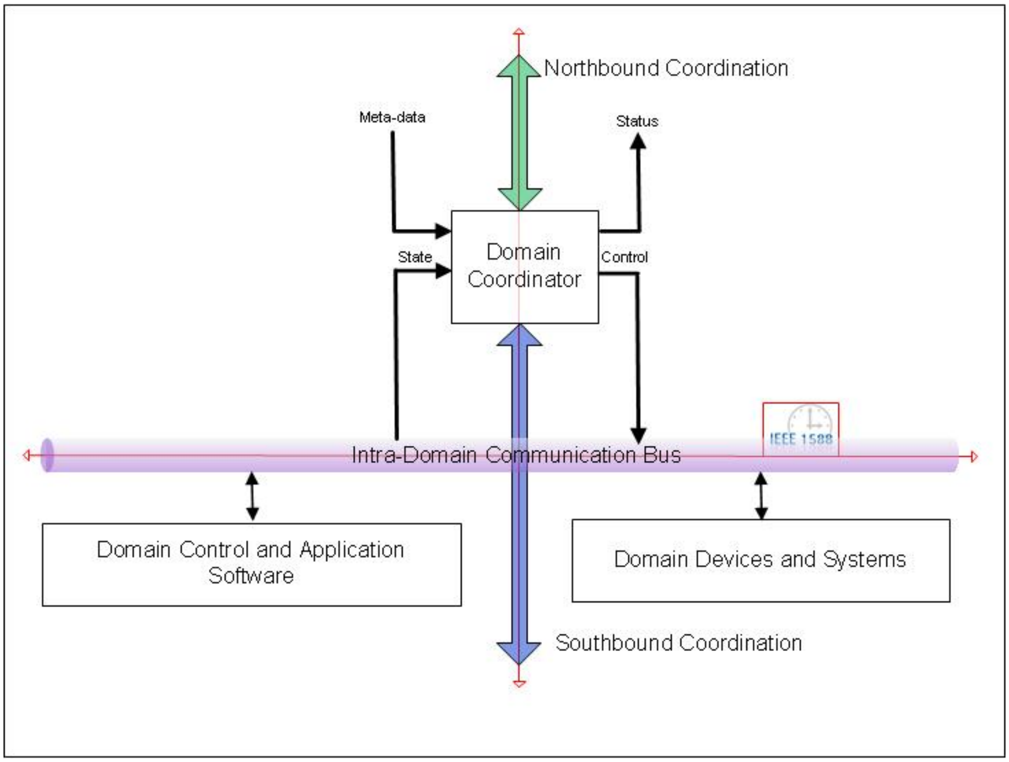

Figure 2 Simplified Domain/Coordinator Structure

The coordinator has a north bound data flow for coordination to the master domain above and a southbound data flow for coordination of the sub-problem domains below. It also connects to the domain data bus to communicate with measurement and control elements in its domain. IEEE 1588 network timing is used throughout the entire architecture to support application level synchronization.

The internal structure of the coordinator itself is also simple as shown in Figure 3. Note that the optimization engine may take many forms and that it may different in each domain or from level to level. It can take the form of a classical steepest decent search tool[10] or a mixed integer nonlinear programming engine, a simple set of equations in the case where the optimization has a closed form solution, or even a market-like mechanism for transactive type controls.

[10] An example would be Newton's Method for finding the minimum point on a curve, such as is often done with state estimation solutions.

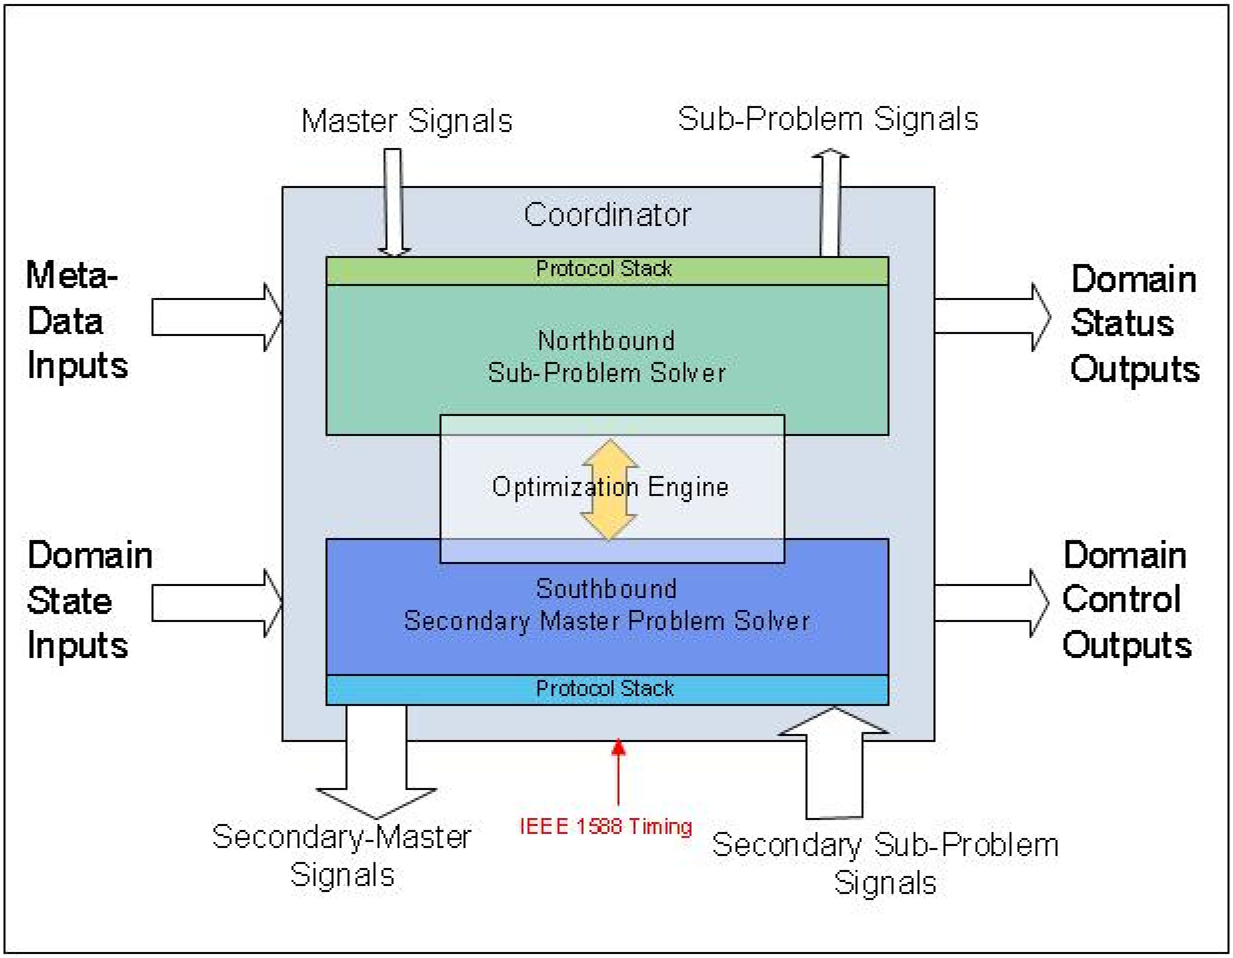

Figure 3 Coordinator Internal Structure

In the following sections we discuss the data flows and other aspects of the Laminar Control model to uncover the key properties that support practical implementations in scalable form.

Decoupling of Laminar Signal Flows by Layer

In a Laminar Control implementation, each optimization layer node except for the very top and the very bottom has two parts: a northbound portion that solves a sub-problem element for the master problem residing one layer up, and a southbound portion that acts as the master for the decomposition of the current layer sub-problem into a new master and set of secondary sub-problems.

One significance of this structure is that signals flowing between southbound master and northbound sub-problem are confined to the two-half layers involved (see Figure 4), meaning that such signals do not aggregate in the communications sense when moving up the optimization node stack. This is therefore an automatic mechanism for preventing the top level data pipe bandwidth requirement from growing without bound as the control system scales upward in size.

At each level, the number of signals involved depends on the number of defined sub-problem elements. It is always possible to define additional domains as necessary, thus controlling the southbound "fan-out" from master to sub-problem set at the cost of increasing slightly the fan-out of the layer above. Note that the fan-in at the top of each layer is one except for the very top node, which has a fan-in of zero.

This degree of design freedom provides another computational advantage: the detection and localization of sensor failures in large scale sensor networks makes use of the theory of random matrices and in particular uses tests on sensor measurement covariance matrices to determine if the largest eigenvalues have Tracy-Widom distributions.[11] As the number of measurement points (sensors) increases, being able to specify the number and the size of domains allows the control engineer to limit the sizes of the matrices involved and therefore to limit the computational requirements. This is another aspect of scalability that is inherent in the structure of Laminar Control.

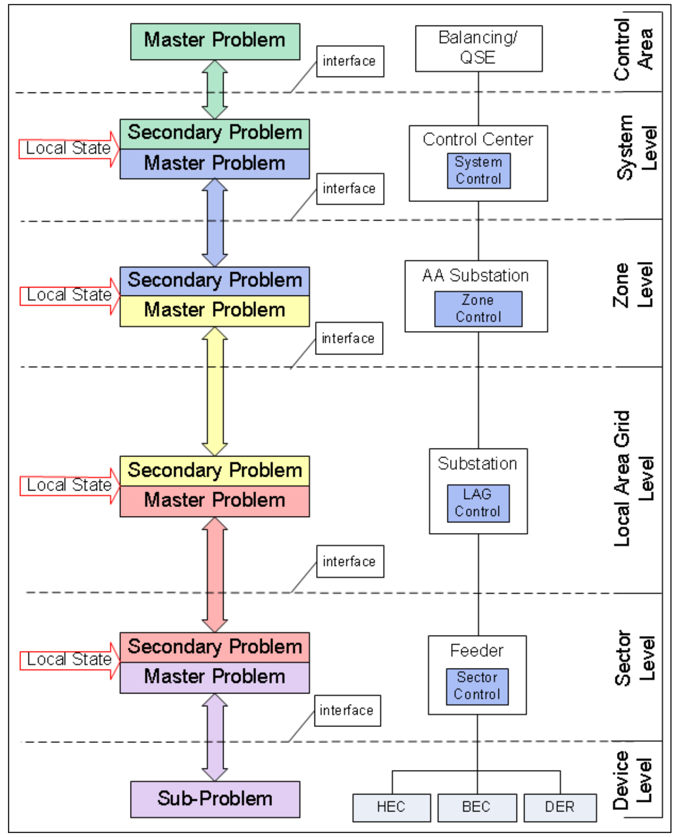

Figure 4 Vertical Coordination Signal Flows

Information Abstraction and Entropy Rate Reduction via Laminar Signals

We suggest that the Laminar Control model has the inherent property of encoding domain grid state in such a way as to accomplish an effective data rate reduction

[11]Romain Couillet and Merouane Debbah, Signal Processing in Large Systems, IEEE Signal Processing Magazine, 24, January 2013, pp. 24-39.

without the need to resort to data compression/decompression processing.[12] To appreciate this, we use an information-theoretic concept to define the term "analytic." Employing the definition of Shannon entropy,[13] we may then define an analytic as a data processing algorithm that reduces entropy for data sets, or entropy rate for data streams.[14] The stream definition is especially useful for real time analytics, as opposed to Big Data/Hadoop offline analytics models.

In the Laminar model, master problems and sub-problems exchange the information necessary to solve the layered optimization problem in a lightweight manner and encode the relevant information on grid state and constraints for the involved domains.[15] Since the master and sub-problems solve the control problem based on state and constraints, this must mean that as the optimization process iterates, the signals converge to represent the information that was relevant to the control problem. In other words, an entropy rate reduction takes place as the optimization process proceeds and therefore the signal generation process constitutes an analytic that performs information abstraction about the control problem, the system state, and the constraints.

We may take advantage of this property in two ways:

1. The amount of data that must be passed from level to level in the control hierarchy can be limited to just the NUM coordination signals, thus greatly limiting the necessary bandwidth of the inter-layer communication links.
2. System state (power state in the case of electric grids) can be determined in a distributed fashion and on a local domain basis, which eliminates the need to determine and distribute state globally, thus reducing computation complexity and bandwidth requirements by employing multiple smaller distributed data acquisition and state calculation processes and limiting state distribution.

---

[12] The formal proof of this property has yet to be developed and is an area of potential research work.

[13] R. Ziemer and W. Tranter, Principles of Communications Systems, Modulation, and Noise, Houghton Mifflin, Boston, 1976.

[14] Jeffrey Taft, Distributed Intelligence for Physical Networks: Sensing, Data and Analytics, Control, and Platforms Part 2: Data and Analytics, Nov 2011. Available online: http://www.cisco.com/web/strategy/docs/energy/data_and_analytics.pdf

[15] The formal proof of this property has yet to be developed and is an area of potential research work.

Both of these consequences are important to obtaining a distributed control system implementation that scales automatically as it is rolled out in an incremental fashion.

Self-Similarity of Laminar Control Networks

Another way to look at Laminar Control networks is that they may be designed to be self-similar. Self-similarity[16] is a geometric property whereby a whole object is similar in shape to any of its parts – an example would be a complex leaf made up of leaflets of the same shape as the overall shape of the complex leaf. This allows us to think about the data flow traffic patterns in a unified way at any level in the power grid hierarchy, which is an aid to selecting and configuring communication network protocols and communication network Quality of Service measures. It also can simplify the implementation of network security policies that make use of traffic models. Self-similarity can be used to achieve network robustness to random connectivity failures, which are in fact a significant concern in the implementation of real distributed control networks.

Another way to think about this property is to view the set of partitioned grid state elements residing in the domains at each tier as a multi-resolution representation of state. Multi-resolution representations have been widely used in many fields, including digital image processing and machine vision,[17] and digital signal processing.[18] We suggest that the Laminar Control structure inherently contains a multi-resolution representation of grid state. This issue is important because it leads to powerful scalability properties that become crucial in advanced grid environments where very large volumes of data must flow from millions to tens of millions of endpoint devices and where communication link failures are common, as is the case with ULS systems.

It is important to understand that the underlying physical network (such as a power grid) and its converged communication network may not be strictly self-similar, meaning that the repeating pattern may not always be completely filled out.

---

[16] John E., Hutchison, Self-Similar Sets. Indiana Univ.Math. J. 30 (1981), pp. 713–747.

[17] James L. Crowley and Arthur C. Sanderson, Multiple Resolution Representation and Probabilistic Matching of 2D Gray-Scale Shape, The Robotics Institute, Carnegie Mellon University, Dec 1984. Pittsburgh, PA. Available online: http://www-prima.inrialpes.fr/jlc/papers/Crowley-Sanderson-PAMI87.pdf

[18] Stephane Mallat, A Theory for Multiresolution Signal Decomposition: The Wavelet Representation, IEEE Transactions on PAMI, Vol. 11, No. 7, July 1989, pp. 674-693.

Consider, for example, a distribution grid from distribution control center to primary substations to feeder circuits. An attempt to calculate the Hausdorff dimension, a measure of self-similarity,[19] would encounter difficulties such as the number of substations not matching the number of feeders emanating from a substation and the number of feeders is not being constant from substation to substation. One might consider some version of approximate self-similarity using an average dimension, but consider instead that we can obtain the key benefits of self-similarity and scale freedom by:

a) relaxing strict adherence to the definition
b) structuring the control network(as opposed to the underlying grid and communications network) so that self-similarity is achieved

Item “b” has already been referenced in the section on Information Abstraction for the purpose of localizing the determination and distribution of system state. Here we see the same issue, but for the purpose of achieving scale freedom of the control network, which when properly applied to network design, leads to inherent robustness to the loss of individual control network node connectivity.[20]

Additionally, this structure facilitates mechanisms to provide hierarchical control that behaves in a manner similar to what Smith describes as “population engineering”,[21] i.e., making changes to a small number of nodes rather than changing all or most of them to effect a change in system behavior. Tier level master problem nodes can influence large numbers of sub-problem nodes using the well-known periodic coordination techniques[22] of having the master adjust sub-problem state estimates to take into account tier interactions, or having the master change sub-problem trajectory directives, or having the master add terms to the sub-problem performance indices. These represent practical methods to implement the “soft control” suggested by Smith to enable the equivalent of flock behavior in a large number of nodes by interaction with a very small number of soft control nodes, the idea being to steer the system away from hazardous regions of operation

---

[19] Dierk Schleicher, Hausdorff Dimension, Its Properties, and Its Surprises, The American Mathematical Monthly, Mathematical Association of America, 114, June-July, 2007, pp. 510-528.

[20] Reuven Cohen, et al., Breakdown of the Internet Under Intentional Attack, Phys. Rev. Letters 86, 3682 (2001).

[21] David M. D. Smith and Neil F. Johnson, Evolution Management in a Complex Adaptive System: Engineering the Future, available online. http://arxiv.org/pdf/cond-mat/0409036.pdf

[22] Robert E. Larson, A Survey of Distributed Control Techniques, First International Conference on Distributed Computing Systems, Huntsville, Alabama, October 1-4, 1979, pp. 245-252.

before they are actually encountered. Thus, using the Laminar Control approach, we can avoid some of the issues of the agent-based approach while still obtaining the benefits of the flock behavior model.

Implications for Ultra-Large Scale Control Implementation

The combination of the properties listed above has significant implications for the design of Laminar Control systems. The benefits fall into several categories described below.

Resilience and Anti-Fragility of the Control Network

The previous section mentioned a form of robustness that derives from network structure. Such robustness manifests itself as graceful degradation (limited loss of functionality or performance) in the face of moderate amounts of stress in the form of random link failures. In addition, in control engineering, there is a class of designs also labeled robust (the H_2-H_∞ controller designs[23]) that are relatively insensitive to uncertainty and variation in the parameters of the model of the system being controlled. Such methods may be viewed as "control hardening" in the sense that they can tolerate a degree of stress but will fail (perhaps catastrophically) when the stress becomes extreme enough.

More recently, the concept of anti-fragility has been introduced and has become linked to the term resilience.[24] In control engineering, resilience has been a topic of attention for some years and has resulted in an evolving series of definitions, including:

"*Resilient control systems are those that tolerate fluctuations via their structure, design parameters, control structure and control parameters*." [25]

More recent definitions also address the element of security by including the response of the system to malicious attack as well as the fluctuations just mentioned.

[23] Eduardo N. Goncalves, et al., Multiobjective Optimization Applied to Robust H_2/H_∞ State Feedback Control Synthesis, Proc. 2004 American Control Conference, Boston, MA, June 30-July2, 2004, p 4619 – 4624.
[24] Nassim Taleb, Anti-Fragile: Things That Gain From Disorder, Random House, New York, 2012.
[25] S. M. Mitchell and M. S. Mannan, "Designing Resilient Engineered Systems", Chemical Engineering Progress102(4), April 2006, pp. 39-45.

The anti-fragile concept which Taleb poses as mostly a philosophical issue suggests that such systems not only tolerate stressful random fluctuations, but may actually improve as a result of encountering them. If we view grid faults as examples of such stress, then Laminar Control systems can exhibit such behavior in the manner in which they can handle feeder section isolation and microgrid islanding.

Understanding how this works requires the introduction of some additional aspects of how Laminar Control is to be implemented in practice. When a section of the Laminar Control tree (be it microgrid or circuit section) becomes isolated or islanded, two additional modes of operation are available:

1) the islanded microgrid or isolated circuit section uses its local optimization criteria and local grid state to operate in a manner that adapts to changing local conditions, rather than just continuing on the basis of the last command/information from a supervisory system that is no longer connected to the island; in effect it becomes a mini-system all its own, performing coordination within the domain of the islanded portion as if it were a complete system
2) the islanded microgrid or circuit section also can seek a different master control node and rejoin the control chain through that node if communication network connectivity permits, in which case the new master will automatically adapt to its expanded set of sub-problem nodes

This means that a customer microgrid or utility microgrid would seamlessly and automatically reconnect to an alternative circuit or generation source as available following an area outage as experienced during recent super storms.

Viewing anti-fragility as a spectrum or matter of degree rather than a binary property, what we can say is that the Laminar Control framework supports the anti-fragile concept somewhat more than do many other control approaches that may exhibit robustness but not the type of adaptivity that would allow a Laminar Control system to manifest superior resilience.

Scalability of the Control Network

By providing an automatic mechanism for information abstraction, and by limiting signal aggregation to each master/sub-problem hub-and-spoke logical sub-network, the Laminar Control model provides a solution to the problems of scalability and incremental rollout. Furthermore, the Laminar Control model allows for the delimiting of state determination and state information distribution, thus providing a way to manage the state computation scaling issue and avoiding the global-scale state distribution problem. State determination in distribution grids is difficult due to the complexity and uncertain knowledge of grid topologies and the fact that distribution grids are often unbalanced, unlike transmission grids. In addition, the number and type of devices attached to distribution grids is constantly increasing as new capabilities penetrate at the distribution level. Finally, as loads become both responsive and transactive, the concept of distribution grid state must be extended to devices and systems which are not owned by the utility but which have significant impact on grid operations. The total number of grid state elements can easily reach into the tens of millions for large grids with significant penetration of advanced capabilities. Under such scenarios, it is crucial that the control architecture have inherent scalability, especially as regards state determination.

Inherent scalability has another benefit: support for incremental rollout. Since utilities typically cannot build whole systems in a green field manner, new technology must be introduced in an incremental manner through rollout programs. The control architecture discussed in this paper inherently supports incremental capital deployment and adaptation of microgrids in several ways:

1) scalability of the network communications is structural
2) ability of the control framework to integrate subsets of the full control system and to integrate new sections automatically as they come online (see the discussion above on islanding and microgrids)
3) ability to implement the control architecture at various tiers in various orders and bring each tier online as elements of that tier become available (this is the multi-tier version of item 2); this means that the control architecture can be built top-down, bottom-up, middle-first, or any combination thereof and the elements will integrate as links are put in place.

Scalability of the control network leads to scalability of the underlying communications network, which is crucial in a distributed intelligence environment, especially one in which many of the data links have limited bandwidth capability.

<u>Complexity Bounding and Complexity Management</u>

Many of the approaches to distributed control for power grids are elegant in theory but represent considerable complexity in operation. Many of these advanced methods have failed to gain traction, leading to energy systems sometimes being characterized as a graveyard of new control theories.[26] Complexity introduces risk, and electric utilities may be reluctant to adopt new control approaches on the basis of new functionality only, unless the risks associated with the new technology are outweighed by its ability to reduce or help manage other risks faced by the utility. In this regard, it is important that the control system not be so complex as to become a risk in itself.

As grid control expands to ultra-large scale, complexity expands exponentially, raising the issue of how a utility may manage the control system itself, as well as ensuring proper control of the grid by such a control system. Consider the multi-agent approaches to grid control. In such approaches, a variety of software agents are defined to implement the control functions. Agents are autonomous software components that may change roles spontaneously, may relocate themselves in a network, or form temporary "teams" to cooperate on a problem. While this is a very elegant concept, such" free-range" behavior does raise a number of issues in terms of management of the agents, and in terms of being able to verify proper operation, or perform after-event diagnoses. The nature and number of interactions in such systems is potentially so large as to be effectively intractable. This leads to effectively unbounded growth of complexity of the control system agent actions and interactions, which represents a severe risk to the utility in not being able to assure effective grid operation. In addition, there is no formal methodology for the specification of agents, or for validation or testing of agent system designs. Specification of agents and agent systems is an art, with little in the way of science

[26] Romeo Ortega, et al., Energy Processing and Control Systems: Joint Past, Common Future, available online: http://www.nester-ru.eu/attachments/057_Ortega%20%20Energy%20Processing%20and%20Control%20Systems.pdf

to aid the designer, although much research continues to be done in this area. Presently agents represent significant risk to the practical user, especially for ultra-large scales at which no agent systems have ever been deployed.

The Laminar Control structure can provide data flow determinism and application synchronization, two keys to managing operational complexity in systems of this type. By applying such constraints, we can regularize operational behavior, thus achieving two ends:

- simplifying the design tasks
- making deviations from expected performance easily detected and analyzed

The same structure that uses self-similarity for scaling also provides a familiar template for optimization signal traffic flows. The hub-and-spoke model is simple and well known in the utility world, and is much easier to debug and diagnose than randomly variable agent-to-agent patterns. By employing the hub-and-spoke pattern at each tier level, we are able to achieve the necessary capability in building block fashion, rather than resorting to ad hoc structures in different tiers and across different control system instances.

Timing can also be used to advantage in managing operational complexity of the control network. Layered optimization decomposition is consistent with several models for timing and synchronization of individual optimizer nodes. Thus the control system designer can select the amount of temporal determinism in accordance with other system tradeoffs, such as the cost of distributing timing signals. Such determinism greatly aids in the commissioning, debugging, monitoring, and diagnosis of the control system.

Conclusion

Public policy has set the US electric industry on a path towards a hybrid power system increasingly powered by customers' distributed resources. Based on EIA and other credible forecasts, installed DER capacity may reach nearly 30% by 2020 – effectively all the incremental capacity installed over this decade. It is clear that new grid operating systems described by EPRI as Grid 3.0[27] are required to enable policy goals.

This Laminar Control approach addresses EPRI's Grid 3.0 requirements by providing both a solution for ultra-large scale controls needed at the expected adoption of distributed energy resources globally and an effective architecture to manage the transition distribution networks from legacy one-way infrastructure to modern multi-way systems. Laminar Control has unique structural properties and available modes of operation that go beyond the issues of federation, disaggregation, and constraint fusion. These properties may make it possible to provide scalability, enhanced network security, resilience, and management of system and operational complexity, all of which are quite valuable in practical implementations across an evolving power system.

Likewise, the modular structure lends itself to discrete capital investments to keep pace with DER adoption. The additional advanced technology/distributed controls overlay on a new modern grid infrastructure (platform) installed when replacing aging infrastructure. Such a core platform includes an enabling field area communications network as well as a transition from analog to digital protection relays and distribution grid designs that bear resemblance to an electrical bus as opposed to traditional large-to-small wire.

This allows utilities to plan investments using a three-pronged approach:

- Address immediate reliability gaps
- Invest in grid modernization as capital budget allows
- Invest in an overlay of advanced technology as DER adoption increases

Modularity, scalability, resilience are key attributes of systems that will necessarily evolve with the changing pace and shape of electric industry transformation. Laminar Control in this context offers utilities, policy makers, technology firms and others a very effective, lower risk architectural approach for the future of distribution.

[27] EPRI staff, Needed: A Grid Operating System to Facilitate Grid Transformation, EPRI, 2011