Decision-Guided Self-Architecting Framework for Integrated Distribution and Energy Management

Jorge Arinez, *Member, IEEE,* Stephan Biller, *Member, IEEE,* Alexander Brodsky,
Daniel Menascé, *Senior Member, IEEE,* Guodong Shao, and João P. Sousa, *Member, IEEE,*

***Abstract*—The future success of the Smart Grid (SG) depends on many technological developments and innovations. The SG will be characterized by decentralized generation, improved network communications, demand response, greater reliability and robustness to failures. The future power grid will also have many interacting components that require new capabilities to efficiently coordinate and manage power and information flows. Though many power distribution engineering and analysis tools have been developed to operate the grid in a real-time manner, there are still areas requiring improved systems integration and decision-making. One such critical area that has great opportunity for improvement is at the interface between Distribution Management System (DMS) operated by distribution utilities and Energy Management System (EMS) operated by end consumers. This paper proposes the development of a set of technologies in the form of Decision-Guided Self-Architecting (DGSA) Framework whose purpose is to provide integrated, real-time optimal decision support in the common domain that intersects the operations of DMS and EMS. This is a key area requiring attention as the overall grid may be sub-optimized unless there is well-integrated and coordinated, optimal decision-making across the various power generating, distributing, and consuming entities of the SG. In addition, this paper also describes the supporting technologies upon which the framework is based and presents two key applications of the DGSA Framework: Energy Management Systems and an Energy Investment Advisor. Also presented are two cases that illustrate various scenarios for how the framework would be implemented and the type of decisions that would be made with such a framework.**

***Index Terms*—Smart Grid, Distributed Management Systems, Energy Management Systems, Self-Architecting Systems.**

I. Introduction

PRESENTLY, there is significant interest and effort underway both at a national and international level to develop next-generation or 'Smart Grid' electrical networks [1], [2], [3], [4]. These efforts are focused on addressing the limitations of the current electricity grid, including key issues such as large transmission and distribution inefficiencies (i.e., 1/3 fuel-to-electricity conversion efficiency and 8% transmission line losses [5]), highly coupled components that are prone to cascading failures, uni-directional communication, and isolated systems with limited integration and connectivity.

J. Arinez and S. Biller are with the General Motors R&D Center, Manufacturing Systems Research Lab, Warren, MI e-mail: jorge.arinez@gm.com and stephan.biller@gm.com, A. Brodsky, D.A. Menasce, and J.P. Sousa are with the Computer Science Department of the Volgenau School of Information Technology and Engineering, George Mason University, Fairfax, VA email: brodsky@gmu.edu, menasce@gmu.edu, jpsousa@cs.gmu.edu and Guodong Shao is with the Manufacturing Systems Integration Division at the National Institute of Standards and Technology, Gaithersburg, MD email: gshao@nist.gov.

Manuscript received October 1, 2010; revised October 22, 2010.

The future Smart Grid (SG) however promises to address most if not all of the present deficiencies of the current grid. One of the key objectives is to better integrate the final electrical consumers and the distributors/generators through the deployment of improved information and communications networks. In principle this will allow real-time visibility into demand requirements, allowing the power supply to be more responsive. Not only does this capability require such aforementioned advanced networks, there will be a need to improve the integration of Distribution Management System (DMS) with Energy Management System (EMS).

To realize the full potential of all that is envisioned for the SG [5], there also needs to be a layer of advanced intelligence in the form of computational tools that can use all of the newly flowing data to make optimal decisions in real-time. Such a suite of advanced optimization tools needs to be more than just a static set of computational algorithms, rather these tools need to evolve, adapt, and modify themselves according to the developments in the structure and technological foundation of the SG.

The SG will require a high degree of systems integration across the entire power generation, transmission, distribution, and consumption chain. Though not necessarily easy, it is technically possible to interface and link the various information systems that will need to communicate with each other across the grid. However, the data integration of the various types of information and control systems used in the power industry alone does not provide the capability nor the guarantee that the system as a whole will be operable in a system optimal fashion. Further, though real-time control capabilities do exist, it is not feasible to dynamically generate an optimization model due to the time-consuming nature of formulating the required model(s) and performing the associated validation and analysis. The only exceptions are highly stable and deterministic power grids.

Thus, it is clear that the interaction among various components and factors of the SG (both utility grids and commercial and industrial (C&I) customer facing microgrids) is very complex. Distributed generation such as solar- or wind-based is inherently intermittent, which introduces significant swings in SG power supply that must be balanced in real or near-real time. On-site storage and aggregation of schedulable demand can be used in the SG to dynamically compensate for sudden reductions in supply or peaks in demand.

Borghetti et al. [6] indicate that advanced DMS have the ability to integrate optimization to perform such a short-term scheduling of resources in power distribution networks. In their work they develop a short-term scheduling approach using a Mixed Integer Linear Programming (MILP) algorithm to optimize electrical generation for distributed resources, the approach considers operation requirements and distribution constraints. Their approach, while successfully applied to a complex network, has all of the resources under the span of control of a single DMS. The problem, though challenging, does not involve the complexities associated with interactions from external EMS, which has its own set of requirements [7] and may not necessarily be easy to model nor control.

Other researchers have developed systems or platforms upon which advanced DMS may be created. Silva et al. [8] develop a web-based DMS application that includes Object-Oriented (OO) models of the components of an electrical distribution network. They develop algorithms to process topologies, estimate state, perform power flow, and short circuit analyses. Casolino et al. [9] present a similar OO development environment based on timed Petri Nets that also integrate a discrete event scheduler. Though this framework has benefits of flexibility and extensibility, it is limited solely to DMS applications and does not have any optimization capabilities.

Despite these tools and frameworks, making optimal real-time decisions across the boundary of DMS and EMS remains elusive and is far beyond the human ability of utility operators and energy managers. Furthermore, current DMS/EMS are not designed to optimally adjust to new grid components, factors, conditions, and faults. Neither are they able to continuously self-architect, configure, or schedule, which are crucial factors for SG efficiency, reliability, and resilience to man or nature-induced system failures and outages.

To bridge these gaps, we propose a Decision-Guided Self-Architecting (DGSA) framework, which may be thought of as a technology platform for developing key decision support tools that can be deployed in various SG applications. The intent of the DGSA framework is to serve as the optimizing 'brain' or computational engine for autonomous yet inter-related smart grids. It is in particular well-suited to work with existing DMS and EMS allowing them to continuously adjust to new grid components, factors, conditions, and faults, and optimally schedule and actuate in a real-time controllable supply and demand of power.

The challenge in developing systems that support these tasks is that each task requires building a mathematical abstraction using a specialized modeling or computational language. Essentially, the same underlying reality must be modeled multiple times using different mathematical abstractions. This makes their development prohibitively expensive. Perhaps most importantly, it has proven to be extremely difficult and expensive to extend the smart grid systems with new technologies, components and factors, which no doubt will be continuously introduced to market over time.

II. Decision-Guided Self-Architecting (DGSA) Framework

The DGSA framework will be designed for use in both utility distribution grids as well as customer-facing microgrids, and to support a range of SG components, such as renewable sources, on-site storage, and distributed generation. The DGSA framework shown in the middle column of Figure 1 consists of two main layers: a Self-Architecting (SA) framework/language and a Decision Guided Management System (DGMS) and Query Language (DGQL). Each layer, though related, is separate and provides unique functionality. The DGSA framework, which is built on a relational database management system (R-DBMS), is also designed as an extensible technology platform on top of which other applications can be constructed. The following sections will discuss two such examples of potential applications, a DGSA-based: (1) EMS and (2) Energy Investment Advisor.

In essence, the DGSA framework is an open, flexible Information Technology (IT) architecture in which existing DMS and EMS (right most column of Figure 1) can be enhanced with database-driven optimization modeling capabilities. These capabilities provide DMS/EMS with real-time decision-making support and intelligence. These advanced capabilities in turn interact with a library of simulation models and tools (left column) for model extraction, validation, and verification.

A. Self-Architecting Framework/Language

The Self-Architecting (SA) framework, which is the upper layer of the overall DGSA framework, captures human-intelligible models of the SG's structure and possible controllable states. Structural views include components (e.g., generation, transformation, use, and control) and lines (e.g., transmission and control), these views are used both at design-time to convey stakeholder intentions, and at operation-time to follow up on the SG status, all the while fed by real-time monitoring.

Applying the SA framework and infrastructure to energy management provides functionality already available in service-oriented autonomic software systems. Specifically, it will enable the distributed, self-aware, highly survivable, self-tuning management of energy systems. It requires three core sets of capabilities:

1) Service-orientation: The foundational step is to endow grid components with a presence in the Internet of Energy Services (IES). This step entails creating generic models of grid components as services, including descriptions of their electric and physical properties, such as generation/demand as a function of time and/or environmental conditions, (e.g., for photovoltaic (PV) generators), operating costs, CO_2 emissions, power factor, transmission efficiency, etc. These models need to be computer-readable and related by ontologies for generality and to facilitate service discovery.

2) Self-awareness: A key challenge is to build the equivalent of 'plug-and-play' functionality for the SG. This challenge entails augmenting grid components, ranging from generators to power lines to appliances, with the capability to advertise their presence once connected to the grid and to monitor their own status. Furthermore, it involves augmenting the grid with service registries to facilitate service discovery across organizational boundaries in a distributed, scalable, and secure fashion while abiding by the terms of commercial contracts.

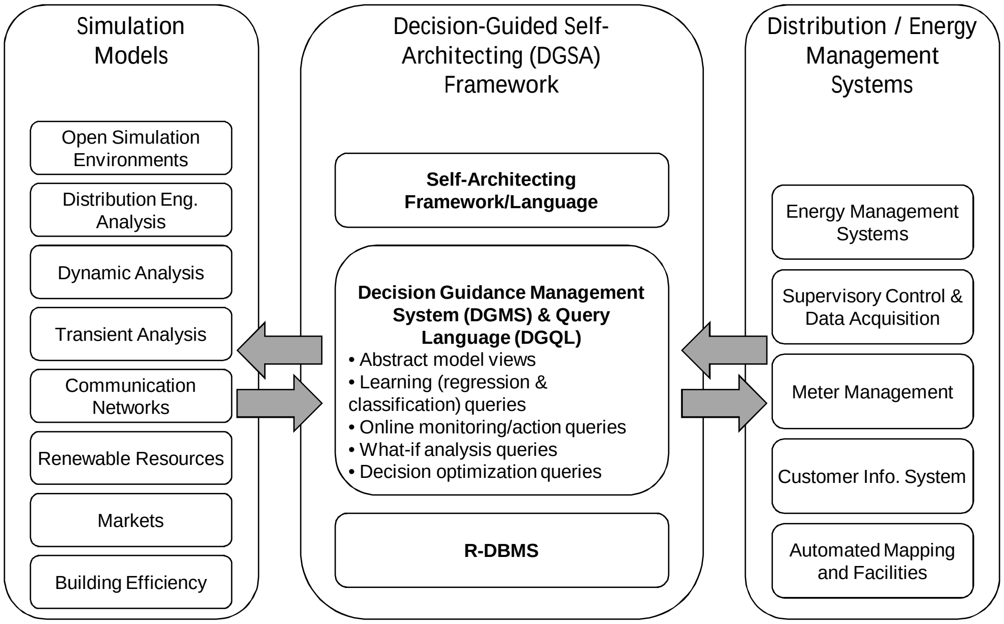

Fig. 1. Decision-Guided Self-Architecting (DGSA) Framework

Monitoring of the services at operation-time keeps the knowledge in the registries up-to-date.

3) Self-architecting: Building on service-orientation and self-awareness, self-architecting takes the effectiveness of autonomic grid management to the next level. Central to self-architecting are models of best-practice patterns combining complementary grid components that jointly work in particularly effective ways, e.g., combinations of PV generators and conventional generators, and/or energy storage components.

One of the benefits of service discovery is that a single organization does not have to acquire and maintain all the components that comprise an effective pattern; components can be discovered across organizations, leveraging self-awareness mechanisms. Patterns are parameterized by quality goals, and the resulting Quality of Service (QoS) can be tuned by changing the number and properties of the components to discover and assemble. Work is still needed to design distributed optimization protocols such that similar to the routing algorithms on the Internet, local optimization decisions take the overall grid towards improved QoS.

The work in self-architecting (SA) systems is motivated by the complexity of making architectural decisions manually in the presence of complex quality tradeoffs: decisions made to optimize certain features may cause adverse effects in other parts of the system. SA accomplishes its goal by leveraging computer-readable models during operation, which capture the system's structure, and fed by probes and gauges, the status of the structural elements.

SA builds on a service-oriented view of systems. Service-orientation adds a level of indirection between provision and use, called service discovery, based on an abstract description of function and quality, this service can be provided by any number of concrete components. This indirection allows the creation of a cross-organization market of services, thus leading to better performing systems. Furthermore, when automated, service discovery plays a key role in making systems autonomic: self-healing, self-adaptive, self-optimizing, and self-protecting.

Prior work in SA research at George Mason University (GMU) [10], [11] was initially applied to software systems and termed 'Self-Architecting Software Systems', or SASSY. Figure 2 provides an overview of the existing SASSY infrastructure, which includes Self-(Re)Architecting, Monitoring Support, and Adaptation Support components. Monitoring Service monitors the QoS values and status of individual service providers and passes that information to Gauge, which computes the utility function for the software system.

If the system's utility falls below a threshold, Gauge informs the Self-(Re)Architecting component to develop a plan for resolving the situation. Self-(Re)Architecting finds an architecture that maximizes a specified utility function that is representative of the desirable objectives. After that, Change Management Service is invoked to effect the changes through adaptation operations provided by the Adaptation Service. Change Management decides how to make the change, and controls and coordinates the dynamic and distributed adaptation of the deployed system.

A key innovation of SASSY is that its Self-(Re)Architecting component works based on models designed by domain experts and system stakeholders, as opposed to computer programmers. These models express the system's goals in terms of activities to be performed and end-to-end QoS goals associated with the stakeholder-defined activity sequences.

B. Decision-Guidance Management System (DGMS) & Query Language (DGQL)

The second layer in the DGSA framework is the DGMS, shown in detail in Figure 3. The work on DGMS at GMU

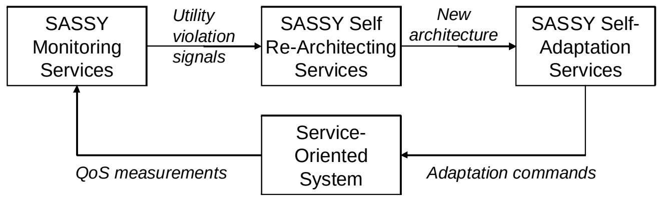

Fig. 2. Self-Architecting Software Systems (SASSY) [10], [11].

includes work on the CCUBE Constraint OO Database System [12], [13], [14], [15], [16] and CoJava which takes Java simulation code and automatically translates it into a constraint optimization problem and solves it [17], [18], [19]. Also, Service-Composition CoJava (SC-CoJava) is an extension of CoJava with composite services to support quick modeling of supply chain [20], [21] followed by Stochastic SC-CoJava, which is an extension of SC-CoJava with two-stage stochastic programming capability [22]. In addition, CoReJava is a CoJava extension, which allows regression analysis of (unknown) parameters in a Java simulation program [23] and the Decision-Guidance Query Language (DGQL) [20], [24].

The DGMS platform allows for the fast iterative development of decision-guidance systems that support (1) construction of learning sets; (2) statistical learning; (3) probabilistic prediction and simulation; and (4) stochastic or deterministic optimization. The domain knowledge for all these tasks is expressed in a parameterized database query language (DGQL), so that the development of models is as simple as the development of database reporting applications. The DGMS engine generates the corresponding mathematical models, such as a MILP or constraint programming (CP) models at run time, and applies a variety of meta-optimization heuristics and commercial optimization solvers. Thus, it combines the extensibility of database application modeling with the benefits of optimization algorithms based on mathematical and constraint programming.

As depicted in Figure 3, DGMS supports functions such as what-if analyses, monitoring and control, statistical learning, and decision optimization. These functions do not need to be manually implemented when required for the management of the SG. Rather, they are automatically derived from abstract model views by the DGMS compiler to describe SG components and factors. These views can describe distribution system operation, sensing and communication, demand response resources, customer behavior, and demand profiles. Other scenarios, ranging from service to critical loads, dispatchable distributed generation and storage, and renewable (nondispatchable) resources to the charging of Plug-in Hybrid Electric Vehicle (PHEV) and Electric Vehicle (EV), can be represented with such database views. The operation of wholesale and retail markets, contractual terms for procurement of electricity, forecasts (of prices and renewable wind and solar conditions), and distribution-level assets and their on-site generation planning and operations are further views that can be represented in the DGMS.

These abstract models and views are written as though one implements a database (DB) reporting application using the Structured Query Language (SQL), which is very simple and intuitive for DB application developers or business people with DB skills. Essentially, each such model is comprised of table schemas that hold the relevant information and SQL views that compute the relevant QoS and business metrics, such as energy consumption and efficiency, and operational costs. They can also be annotated by indicating that some of the table columns are unknown, while another view can be annotated to indicate that the value it computes (e.g., adjusted cost) is to be used as an optimization objective.

Given this information, DGMS will automatically generate, at run time, a formal mathematical programming problem with mathematical equations, inequalities, and the objective function and deploy a mix of algorithms best suited for the problem at hand, e.g., MILP using IBM ILOG CPLEX optimization solver.

Therefore, when a new component or a factor of the SG is introduced, the only requirement is to add a simple SQL-view model for this component, whereas all the learning, optimization and other DGMS functions are automatically implemented with the use of the DGMS compiler. With the appropriate abstract model views, the DGMS layer that supports the DMS/EMS will support decision optimization functions such as:

- Microgrid optimal operational scheduling, including local generation, spot market purchase, on-site storage charge or discharge, and thermal storage
- Microgrid optimal load shedding distribution
- Optimal thermostat settings
- Optimal PHEV and EV charging schedules
- Optimal contractual terms including curtailment level commitment and peak demand limits
- Optimal Return On Investment (ROI) including in renewable generation, storage, local generation, and thermal energy storage

Also, given the abstract model views, DGMS will support the learning functions (parameter calibration and classification) such as:

- Occupancy prediction classification
- Demand forecasting functions
- Heat ventilation and air conditioning (HVAC) temperature/settings functions
- QoS system wide metrics
- Microgrids fault detection classification
- Outage classification
- Online monitoring/action queries
- Occupancy prediction detection
- Demand change detection
- HVAC temperature /settings abnormality detection
- QoS abnormality detection
- Microgrids fault detection
- Outage detection
- Load shedding conditions detection

Finally, the lower layer of the DGSA framework is the Relational Database Management System integrated with DGMS. Here all system static and dynamic data are normalized and stored in R-DBMS. Thus, part of the framework development

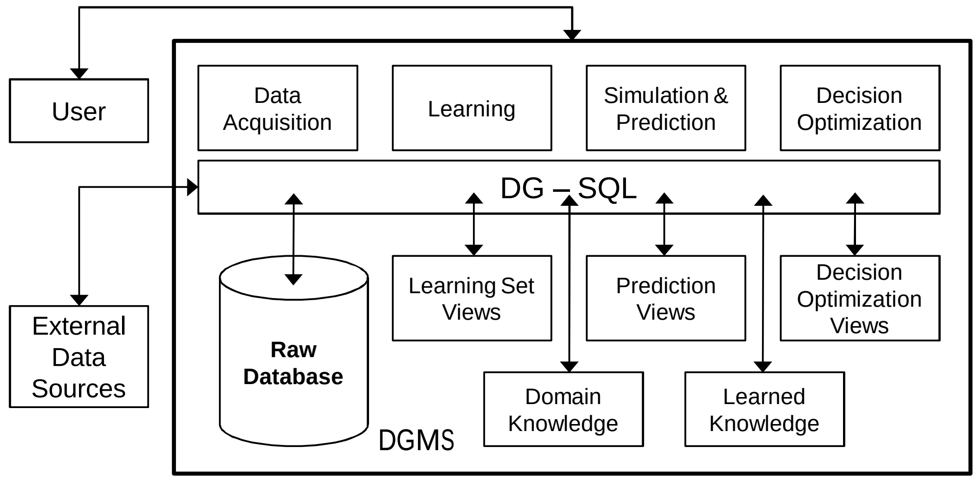

Fig. 3. High-level view of the Decision-Guidance Management System (DGMS)

is to create a flexible, yet normalized (i.e., in Boyce-Codd Normal Form) design of database schema that can be maintained by the database administrator.

III. DGSA Application 1: Energy Management Systems

Most existing C&I customer microgrids already use an EMS, which receives a continuous stream of supervisory control and data acquistion (SCADA) data on the status of power consumption and equipment to take control actions. Today, making operational decisions through EMS, even sub-optimally, may be very complex and unique for various organizations. With the introduction of new 'green' technologies such as on-site renewable generation, on-site storage, energy harvesting, and/or charging stations for electric vehicles, the complexity of making optimal decisions to minimize power consumption, carbon emissions, and cost, grows significantly.

To be able to make optimal operational decisions on energy management, we propose using the DGSA Framework described earlier to extend the capabilities and intelligence of existing EMS, which we call a DGSA-based EMS. The DGSA-based EMS will provide guidance in an iterative process by providing a facility energy manager with a small number of optimal actionable recommendations, and receiving the manager's feedback, constraints, and priorities, until the manager satisfies with a recommendation and decides to execute it.

To create a DGSA-based EMS, there are two development paths that may be taken. First, the database structure of the DGMS shown in Figure 3 may be integrated with that of an existing EMS. This approach allows for tight coupling between EMS functions and the DGSA decision-making capabilities. However, depending on the customer's EMS technology platform, the database structure may not always be available if the EMS is a commercial product with a proprietary database architecture. A second, more flexible approach is that the DGMS receives data directly from the physical devices or as a data fed indirectly from the EMS. Then, the DGMS performs the necessary optimizations and provides recommended actions to a user. Upon acceptance of the actions by the user, the DGMS sends the required control actions to the EMS, which then sends the specific commands to the devices to effect control.

To understand some of the different types of control actions that could be taken by the DGSA-based EMS, consider as an example a facility such as a university campus, large government building, or a manufacturing plant. The decisions to be actuated by an EMS may involve a diverse range of factors and trade-offs that affect total energy consumption, carbon emissions, and operational costs. Another important component of the DGSA-based EMS is a market optimizer designed to promote fair collaboration of microgrid stakeholders to achieve higher savings in energy consumption. The DGSA-based EMS will guide or make the following classes of decisions (that are either recommended to the energy manager or automatically actuated):

1) *Setting the target peak demand and curtailment commitment:* A high peak demand target may be prohibitively expensive, but it could lead to less interruption of power and cheaper per kWh rates; whereas, a low peak demand target may lead to more interruption (and the need to shed load), using more expensive (and emitting) local generation, buying power on the spot market, or any mix of the above.
2) *Control of supply and/or demand in response to unforeseen changes in supply and/or demand:* For example, supply can drop as a result of the interruption of a renewable source (e.g., the sun stops shining), or a curtailment signal from a utility, or demand may increase, e.g., due to cold/hot weather. The response may involve load shedding to decrease consumption, using local generation, on-site storage, buying on the spot market, or any mix thereof.
3) *Scheduling of consumption through aggregation and/or prioritization:* For situations such as power consumption from manufacturing processes, on-site battery charg-

ing, or ice production for cooling, one would like to enable such consumption when the other demand is supposed lower and cheaper (e.g., off-peak hours and below peak demand target). On the other hand, we would like to schedule in such a way to guarantee, with high probability, that we can curtail energy when required. This ”energy option” may generate revenue that exceeds the average per kWh cost. These trade-offs must be optimally made while taking into account business constraints such as deadlines for completion of electrical vehicle charging or production requirements.

4) *Optimal load shedding:* In the case of sites with multiple buildings the key decision pertains to identifying and selecting those which can be interrupted with minimal inconvenience or negative economic impact. For example, various units within an organization may indicate that they are willing to pay more for not being interrupted, or conversely, that they are ready to be interrupted in lieu of high enough financial reward. This requires designing a special market (which is non-intrusive to stakeholders) and identifying an optimization solution to find a fair (equilibrium) price, to determine the best load shedding. The DGSA Framework is particularly well suited to these types of decisions because the DGSA-based EMS can be integrated with the DMS of a local utility providing improved awareness of load shedding opportunities.
5) *Learning occupancy patterns:* Given a stream of information such as that from motion/occupancy sensors and organizational and personal calendars, one can make much better decisions on HVAC thermostat setting or lighting, based on statistical learning of occupancy patterns instead of on static rules. For example, how do we decide on optimal settings of ”occupancy criteria”? This involves the identification of parameters such as X, Y, and Z to frame decisions such as ”if after X time of day, the room sensor has not been activated for Y minutes, we predict that the room will be empty until Z hours the next day”.

IV. DGSA Application 2: Energy Investment Advisor

The development and deployment of energy efficient building technology requires making a range of decisions by stakeholders on optimal mixes of investments and operational energy management. These decisions may be both very complex and unique for various organizations.

Consider an organization that is making investment decisions on energy efficient technologies in their organizations, such as a residential complex, a university campus, a hospital complex, a large government, or a industrial facility. There is a large range of energy management technologies available today such as: improved window insulation, renewable (e.g., solar) power sources, traditional back-up generators, battery storage, highly efficient HVAC systems, energy harvesting solutions, and soon charging stations for EVs. New and more efficient technologies will be continuously introduced in the coming years. What mix of these technologies makes economic and environmental sense for a particular organization now and over time? And what interrelated contractual arrangement should the organization make with power and gas utilities, as well as load curtailment companies (e.g., what should the target peak load and/or curtailment commitment be, and how should utility meters be unified?).

More specifically, how do we assess the energy, renewable energy credits, and cost savings and ROI for a particular mix of investments and contractual terms? And, how do we recommend an optimal mix, e.g., with maximum ROI subject to budget limitations over time? Unfortunately, simple answers such as ”introducing technology X typically saves Y%”, do not work for any non-trivial system, since the technology components are highly inter-dependant and may involve complex interactions among them. To do the assessment, one needs to analyze historical and projected energy/power consumption demand patterns (per device, over space and time). Then, a baseline must be computed, which is what the consumption and costs would be without introducing new technologies. Finally, one needs to assess the consumption with the new technology mix introduced.

The last part is particularly challenging: we need to assess the consumption not per historical power/energy consumption pattern, but for optimally scheduled and configured power/energy demand. This ”optimal” scheduling and configuration problem has many facets. For example, if one introduces power storage and back-up/local generation, the organization could commit to a much higher curtailment (and get significant economic rewards for it). Or, the organization may schedule production of ice during off-peak hours to reduce the consumption during more expansive rates. Or, the organization may schedule an interruptible load (e.g., charging vehicles or cooling or heating water) within the peak demand bounds in order to create an ”energy option.” Or, an organization may decide that significant solar generation (which is highly interruptible) could be compensated with local dispatchable generation. Both the assessment of savings (energy, emissions, cost) for a particular mix of energy efficient technologies and a recommendation for the ”best” mix requires considerable formal modeling, decision optimization, and statistical learning software solutions.

V. DGSA Framework and Standardization for Smart Grid Components

For the DGSA framework to be robust across the various commercial DMS and EMS, standards must be in place to reduce the difficulty and effort to integrate these two types of different but related systems. Development of standards for the SG are well underway since they are critical to enabling interoperable systems and components. Though DMS/EMS technologies and components may be developed by many different companies, SG standards applied to the DGSA framework will enable diverse technologies, systems, and their components to work together to securely exchange meaningful, actionable information and promote consistency in systems management and maintenance. To this end, The

National Institute of Standards and Technology (NIST) has identified a list of standards for SG interoperability [4].

Data communication standards define information models and messages for communication between parties and provide a mechanism to instantiate, format, store, and exchange common, meaningful information. Relevant standards may be developed for systems such as building system (ANSI/ASHRAE 135-2008/ISO 16484-5 BACnet, ANSI/CEA 709, and CEA 852.1 LON Protocol Suite), and a metering model (ANSI C12 Suite, ANSI/CEA 709, and CEA 852.1 LON Protocol Suite). Sample standards for information exchange include common information models that define application-level EMS interfaces and messaging for distribution grid management in the utility space among control center systems (IEC 61968/61970 Suites), a de facto communication protocol used at the distribution and transmission level between control centers and substations (DNP3), a specification that defines messages exchange between utilities and commercial/industrial customers for price-responsive and direct load control (OpenADR), an open standard for data exchange based on a publish/subscribe mechanism (OPC-UA Industrial), and a standard for geographic data exchange (Open Geospatial Consortium Geography Markup Language (GML)).

VI. Proposed Case Studies

A. George Mason University Campus Energy Management

GMU will focus on a case study demonstrating both R&D solutions: DGSA-based EMS and the Energy Investment Advisor. This study will focus on how to significantly reduce power consumption and carbon emissions, lower peak demand, and bring about significant savings in adjusted costs (costs minus SG related revenues such as curtailment). The study will follow the following steps:

1) Develop a methodology and metrics for assessing the performance of the DGSA-based EMS and enacting Energy Investment Advisor actions, based on a formal statistical foundation.
2) Design and deploy a data collection infrastructure, using a relational DBMS capable of collecting both operational data during the demonstration, as well as a range of historical data and SG components specifications.
3) Before running a demonstration, perform analyses on the optimal mix of SG investments, such as solar, local generation, on-site storage, and a mix of controllable contractual terms. This will be done using DGSA Energy Investment Advisor.

B. Energy Management in Automotive Manufacturing

Automotive manufacturing plants represent a wide diversity of facilities with processes that have challenging energy management needs. Automotive manufacturing plants typically have EMSs which provide real-time energy consumption information, however, this data is often at a very high level of abstraction. In this study, the plant selected will have an EMS, which allows monitoring the status of lights, HVACs, exhaust fans, power meters, substations, and flow meters (which include those for natural gas, compressed air, chilled water, and steam). The plant's EMS also provides control functions for states such as occupied/unoccupied or on/off in addition to managing set points. Consequently, the challenge to deploy the DGSA framework and extend the plant's conventional EMS involves defining and acquiring the data collection mechanisms to allow the DGMS to provide effective operational recommendations. Hence, this study will follow these steps:

1) Understand core EMS database structure and how to augment and extend it with the database structure of the DGMS.
2) Develop linked table structures as interfaces to the DGSA framework from the production system database.
3) Integrate the manufacturing process energy data collection with the EMS and DGMS.
4) Perform validation between EMS and utility energy and power consumption data to ensure accuracy of simulations.

VII. Discussion

In summary, energy management and investment applications require predicting behavior of a complex system and making decisions to move the system towards desirable outcomes, such as reducing energy/power consumption and carbon emissions, and saving operational costs, while maintaining desirable comfort level. In such applications, predictions and decisions are to be made in the presence of large amounts of dynamically collected data and learned uncertainty models. There has been extensive research in the areas of operations research, mathematical and constraint programming, machine learning and data mining, and database systems. However, there are no cohesive frameworks, algorithms, and systems that unify the models and computational paradigms of all the components. Without unification of models for the related but different tasks, users are forced to express their knowledge of the underlying domain multiple times using different mathematical abstractions. This makes the development of decision support systems costly and time consuming, and extremely difficult to modify and extend. Moreover, developing decision guidance applications today requires considerable expertise in operations research and mathematical programming that most software application developers do not have.

Thus, the DGSA framework, proposed in this paper, is designed to address the research needs at the interface of real-time decision support tools for integrating DMS and EMS. More specifically, the contributions of this paper are as follows: the introduction of the DGSA framework and how such a framework can help satisfy the objectives of the SG. Second, we review the underlying technologies of Self-Architecting Systems, which allow the optimal run-time system configuration, and DGMS, which allow use of an extensible model library of smart grid components and perform decision optimization, learning, and prediction tasks. Third, we propose two applications of the DGSA Framework: an EMS and the Energy Investment Advisor. Finally, two proposed cases are described that illustrate scenarios for how the framework would be implemented.

REFERENCES

[1] "The Smart Grid: an Introduction," U.S. Department of Energy, Tech. Rep., 2008.

[2] R. Pratt, P. Balducci, C. Gerkensmeyer, S. Katipamula, M. Kintner-Meyer, T. Sanquist, K. Schneider, and T. Secrest, "The Smart Grid: An Estimation of the Energy and CO2 Benefits," Pacific Northwest National Laboratory, U.S. Department of Energy, Tech. Rep. PNNL-19112, Revision 1, January 2010, prepared for the U.S. Department of Energy under Contract DE-AC05-76RL01830.

[3] "Gridwise Architecture Tenets and Illustrations," Pacific Northwest National Laboratory, U.S. Department of Energy, Tech. Rep. PNNL-SA-39480, October 2003.

[4] "NIST Framework and Roadmap for Smart Grid Interoperability Standards, Release 1.0," NIST, Tech. Rep., January 2010, NIST Special Publication 1108.

[5] H. Farhangi, "The Path of the Smart Grid," *Power and Energy Magazine, IEEE*, vol. 8, no. 1, pp. 18 –28, jan. 2010.

[6] A. Borghetti, M. Bosetti, S. Grillo, S. Massucco, C. Nucci, M. Paolone, and F. Silvestro, "Short-term scheduling and control of active distribution systems with high penetration of renewable resources," *Systems Journal, IEEE*, vol. 4, no. 3, pp. 313 –322, Sep. 2010.

[7] J. Arinez and S. Biller, "Integration Requirements for Manufacturing-Based Energy Management Systems," in *IEEE PES Conference on Innovative Smart Grid Technologies*, January 19-21, 2010.

[8] M. Silva, J. Saraiva, and A. Sousa, "A Web Browser based DMS-Distribution Management System," vol. 4, 2000, pp. 2338 –2343.

[9] G. Casolino, M. Landolfi, A. Losi, and M. Russo, "A Software Environment for Developing Object-Oriented Applications in Distribution Management Systems," vol. 2, Nov. 2005, pp. 1457 –1460.

[10] D. A. Menascé, J. M. Ewing, H. Gomaa, S. Malek, and J. P. Sousa, "A framework for utility-based service oriented design in SASSY," in *WOSP/SIPEW*, 2010, pp. 27–36.

[11] D. A. Menascé, J. P. Sousa, S. Malek, and H. Gomaa, "QoS Architectural Patterns for Self-Architecting Software Systems," June 2010, pp. 195–204.

[12] A. Brodsky, J. Jaffar, and M. J. Maher, "Toward practical constraint databases," in *VLDB*, 1993, pp. 567–580.

[13] A. Brodsky and Y. Kornatzky, "The lyric language: querying constraint objects," in *SIGMOD '95: Proceedings of the 1995 ACM SIGMOD international conference on Management of data*. New York, NY, USA: ACM Press, 1995, pp. 35–46.

[14] A. Brodsky, V. E. Segal, J. Chen, and P. A. Exarkhopoulo, "The CCUBE constraint object-oriented database system," in *SIGMOD 1999, Proceedings ACM SIGMOD International Conference on Management of Data, June 1-3, 1999, Philadelphia, Pennsylvania, USA*, A. Delis, C. Faloutsos, and S. Ghandeharizadeh, Eds. ACM Press, 1999, pp. 577–579.

[15] A. Brodsky, J. Jaffar, and M. J. Maher, "Toward Practical Query Evaluation in Constraint Databases," *Constraints, An International Journal*, vol. 2, pp. 41–304, 1997.

[16] A. Brodsky, V. E. Segal, J. Chen, and P. A. Exarkhopoulo, "The CCUBE constraint object-oriented database system," *Constraints*, vol. 2, no. 3/4, pp. 245–277, 1997.

[17] A. Brodsky and H. Nash, "CoJava: Optimization modeling by nondeterministic simulation," in *Principles and Practice of Constraint Programming - CP 2006, 12th International Conference, CP 2006, Nantes, France, September 25-29, 2006, Proceedings*, ser. Lecture Notes in Computer Science, F. Benhamou, Ed., vol. 4204. Springer, 2006, pp. 91–106.

[18] ——, "CoJava: A Unified Language for Simulation and Optimization," in *CP*, 2005, pp. 877–877.

[19] ——, "CoJava: A Unified Language for Simulation and Optimization," in *OOPSLA Companion*, 2005, pp. 194–195.

[20] A. Brodsky and X. S. Wang, "Decision-guidance management systems (DGMS): Seamless integration of data acquisition, learning, prediction and optimization," in *HICSS*, 2008.

[21] M. Al-Nory and A. Brodsky, "Unifying simulation and optimization of strategic sourcing and transportation," in *Winter Simulation Conference*, 2008, pp. 2616–2624.

[22] M. Al-Nory, A. Brodsky, and H. Nash, "A Service Composition Framework for Decision Making under Uncertainty," in *ICEIS*, 2009, pp. 363–375.

[23] A. Brodsky, J. Luo, and H. Nash, "CoReJava: Learning Functions Expressed as Object-Oriented Programs," in *ICMLA*, 2008, pp. 368–375.

[24] A. Brodsky, M. M. Bhot, M. Chandrashekar, N. E. Egge, and X. S. Wang, "A Decisions Query Language (DQL): High-Level Abstraction for Mathematical Programming over Databases," in *SIGMOD Conference*, 2009, pp. 1059–1062.

Jorge Arinez (M'09) received his Ph.D. (2000) and Masters (1995) degrees in Mechanical Engineering from the Massachusetts Institute of Technology and his B.A.Sc. (1993) from the University of Toronto. He is currently a staff researcher in the Manufacturing Systems Research Lab at the General Motors R&D Center. His research interests include the real-time monitoring, control, and quality of manufacturing systems with a focus on energy efficiency and sustainability of plant operations.

Stephan Biller (M'07) received a Dipl.-Ing. degree in electrical engineering from the RWTH Aachen, Germany, a Ph.D. in Industrial Engineering and Management Science from Northwestern University and an MBA from the University of Michigan. He is a GM Technical Fellow and currently a Lab Group Manager in the Manufacturing Systems Research Lab at the General Motors R&D Center. Here he has responsibility for innovations in sustainable plant floor systems and controls and is presently focusing on the digital factory, the real-time information enterprise and the interoperability of the two.

Alexander Brodsky received his Ph.D. (1991) and M.Sc. (1983) in Computer Science, and B.Sc. (1982) in Mathematics and Computer Science, from the Hebrew University of Jerusalem. He is Director, Center for Smart Power Grids, and Associate Professor of Computer Science at George Mason's Volgenau School of Engineering. His research interests include decision-guidance systems, databases, constraint optimization and their application to energy efficiency and smart grid.

Daniel Menasce (SM'05) received his Ph.D. in Computer Science (1978) from the University of California at Los Angeles. He is the Senior Associate Dean and a Professor of Computer Science at George Mason's Volgenau School of Engineering. He is a Fellow of the ACM and a Senior Member of IEEE. His areas of research interest include autonomic computing systems, software performance engineering, analytic models of computer systems.

Guodong Shao is a computer scientist in the Manufacturing Systems Integration Division at the National Institute of Standards and Technology. His research interests include manufacturing simulation, sustainable manufacturing, system integration, and decision guidance applications. He serves on the Executive Board of the Winter Simulation Conference. He holds a Masters Degree from University of Maryland at College Park. He is a Ph.D. candidate in Information Technology at George Mason University.

João P. Sousa (M'09) received his Ph.D. in Computer Science (2005) and Masters in Software Engineering (1995) from Carnegie Mellon University, and his Masters in Applied Mathematics (1994) and his B.Sc./Lic. in Electrical and Computer Engineering (1988) from the Technical University of Lisbon, Portugal. He is currently an Assistant Professor of Computer Science at George Mason's Volgenau School of Engineering and a member of the Center for Smart Power Grids. His research interests include applications of ubiquitous computing to smart spaces, self-configuring and autonomic cyber-physical systems, and energy management.