The Role of Context and Resilient Middleware in Next Generation Smart Grids

Abhishek Dubey, Subhav Pradhan, and
Douglas C. Schmidt
Vanderbilt University
Nashville, TN, USA
{abhishek.dubey, subhav.m.pradhan,
d.schmidt}@vanderbilt.edu

Sebnem Rusitschka, and Monika Sturm
Siemens Corporate Research
Munich, Germany
{sebnem.rusitschka,
monika.sturm}@siemens.com

ABSTRACT

The emerging trends of volatile distributed energy resources and micro-grids are putting pressure on electrical power system infrastructure. This pressure is motivating the integration of digital technology and advanced power-industry practices to improve the management of distributed electricity generation, transmission, and distribution, thereby creating a web of systems. Unlike legacy power system infrastructure, however, this emerging next-generation smart grid should be context-aware and adaptive to enable the creation of applications needed to enhance grid robustness and efficiency. This paper describes key factors that are driving the architecture of smart grids and describes orchestration middleware needed to make the infrastructure resilient. We use an example of adaptive protection logic in smart grid substations as a use case to motivate the need for context-awareness and adaptivity.

CCS Concepts

•**Computer systems organization → Dependable and fault-tolerant systems and networks;** *Resilience;*

Keywords

Smart Grid, Autonomous Resilience, Middleware

1. INTRODUCTION

Emerging trends and challenges. Electrical power grids and operations are software-reliant systems that are under pressure to evolve in response to the growing volatility in their environments and requirements. For example, the energy generation environment is changing due to the advent and adoption of distributed energy sources, such as wind and solar power. These new energy sources are increasingly decentralized and distributed through large energy parks located far from load centers and "prosumers" (who can both produce and consume energy locally). Decentralized energy sources also have more variation (*e.g.*, fluctuating generation capacity) than traditional power sources, such as coal or nuclear power [3]. These changes are motivating significant modifications in legacy power system infrastructure and operations. For example, additional transmission and distribution capacity is needed to support new sources of energy, which in turn requires investments in new physical infrastructure, such as power lines and substations.

Likewise, power system requirements are changing in response to new regulatory models, such as unbundling energy suppliers from grid operators and giving consumers more choices in selecting their energy provider(s). These models are hastening the emergence of *next-generation smart grids* based on *transactive energy systems*, which provide a network environment for distributed energy nodes, as opposed to the hierarchical structure in legacy power grids [1, 12]. Key goals of these regulatory models are to flatten peak loads and incentivize the adoption of new energy sources. Achieving these goals, however, requires software architects and systems engineers to devise more adaptive and context-aware strategies for operating smart grids safely in markets where decentralized sources of energy can be traded dynamically, as well as meet the growing demands of energy producers and consumers [15].

The changes in environments and requirements needed to support next-generation smart grids are motivating significant modifications in legacy power system infrastructure and operations. For example, the management and configuration of physical devices and software architecture inside legacy substations has traditionally been governed by electrical grid specifications, such as IEC61850 [7] . In these legacy systems communication is statically configured along pre-specified channels. Likewise, the system architecture is defined using logical nodes that are similar to component-based software elements [4], though these logical nodes are statically bound to physical devices. This architecture yields tightly coupled functionality that is bound to relatively fixed locations.

Although pre-configuration makes it easier to validate system behavior, reliance on pre-configuration is problematic for next-generation smart grids since it yields systems that are inflexible, costly to evolve, and react inefficiently to dynamic changes. For example, pre-configured provisioning and control of protection devices exacerbates the under-utilization of capacity during non-peak usage. If enough contextual information is available, adaptive relays can be used to dynamically adjust grid protection equipment to handle

M4IoT 2016, December 12-16, 2016, Trento, Italy
© 2016 ACM. ISBN 978-1-4503-4663-4/16/12...$15.00
DOI: http://dx.doi.org/10.1145/3008631.3008632

different load levels [8, 5].

Promising Solution → A context-aware middleware for smart grids. An alternative way to meet the changing environments and requirements discussed above involves digitally enhancing the operation of the legacy power transmission and distribution infrastructure in a manner that enhances software deployment flexibility, yet preserves system resilience. At the heart of these enhancements are context-aware IoT systems. This approach enables the use of software as the "universal integrator" of smart grid components, which is a common trend in modern cyber-physical systems[16].

This paper provides the following contributions to research on context-aware techniques in smart grids that help determine where to deploy software functionality and intelligently protect grid resources:

1. We create a taxonomy that showcases the role of context in guiding the development and operation of electrical power systems. This taxonomy can guide software architect and systems engineer decisions on when and where to flexibly deploy software functionality in next-generation smart grids.
2. We describe how context-aware middleware can be applied to smart grids by showing how to automate the deployment and online parameter reconfiguration decisions of adaptive protection system by using resilient orchestration middleware called CHARIOT [9, 10].

The remainder of this paper is organized as follows: Section 2 examines the role of context in guiding the development, operation, and evolution of software-reliant smart grids; Section 3 describes the intelligent devices that enable context awareness in smart grids; Section 4 gives an overview of the CHARIOT resilient orchestration middleware; Section 5 explains prototype for our adaptive protection case study using context-aware IoT systems and CHARIOT; and Section 6 presents concluding remarks.

2. CONTEXT AND SMART GRIDS

To help structure our discussion, Figure 1 depicts the rates at which key resources and capabilities in power systems change over time (shown on a scale from "slow rate of change" to "fast rate of change" on the Y axis), grouped according to how these resources and capabilities map to following key domains in the electrical power system sector (shown on the X axis):

- The financial domain consists of value assets (*e.g.*, power generation volume and transmission capacities) and transactions for value transfer between market stakeholders.
- The physical domain consists of power system equipment, such as transmission lines, circuit breakers, transformers, and capacitor banks.
- The digital domain consists of various IT resources, such as computing and network hardware and software.

Figure 1 shows that power systems constitute a range of domains (*i.e.*, financial, physical, and digital) in which changes occur at different rates (*i.e.*, ranging from slow to fast). Moreover, changes in all of three domains are now occurring more frequently than in the past, *e.g.* through market liberalization, digitalization of power equipment and power electronics, and the rate of change within digitalization itself (*i.e.* available computing power, computing and architectural paradigms around data-driven applications).

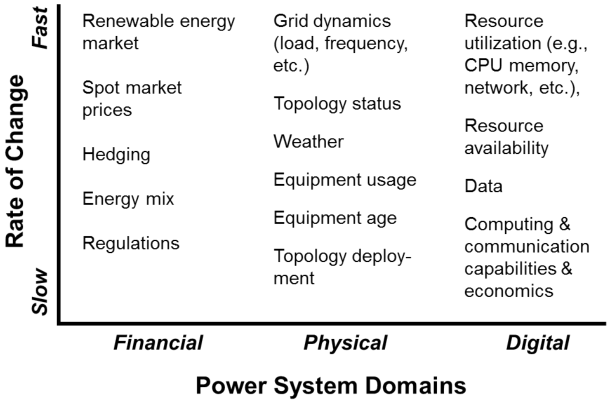

Figure 1: Rates of Change and Volatility of Context in Power System Domains

Next-generation smart grids are thus characterized by increased *volatility*. This volatility stems not only from the inherent perturbations of renewable, decentralized electricity generation, but also through a paradigm shift. In particular, power consumption follows available electricity through energy efficiency, demand response, and dynamic demand, rather than maintaining and operating costly spinning reserve power plants that are less flexible and adaptive with respect to power generation [13].

Developing, operating, and evolving software for next-generation smart grids that is responsive to the types of changes shown in Figure 1 requires a deeper awareness of the role of *context* to guide decisions during design, development, and operation. For example, due to the increased volatility in the financial and physical domains outlined above, the digital domain must compensate for this volatility through greater resilience and adaptivity. Likewise, the digital domain has also become more volatile due to various factors, such as centralized vs. decentralized computing economics, that change over time (*e.g.*, depending on technology innovation rates and market conditions).

3. INTELLIGENT DEVICES THAT ENABLE CONTEXTAWARENESS

IEDs at the edge of a smart grid topology hierarchy shown in Figure 2 decouple computing from the actual physical assets and controls to enable the digitalization of functionality that has historically been tightly coupled to dedicated (electro-)mechanical power protection devices. These IEDs can be reprogrammed to change either their operational logic (*e.g.*, between various protection schemes) or to deploy new functionality (*e.g.*, protection, whereas IED previously performed only monitoring). These new devices thus enable more flexible and adaptive binding of functionality onto different elements throughout a smart grid hierarchy. Decisions on where to deploy this functionality can be guided by a range of contextual information, including

- **Characteristics of the physical grid**, such as fore-

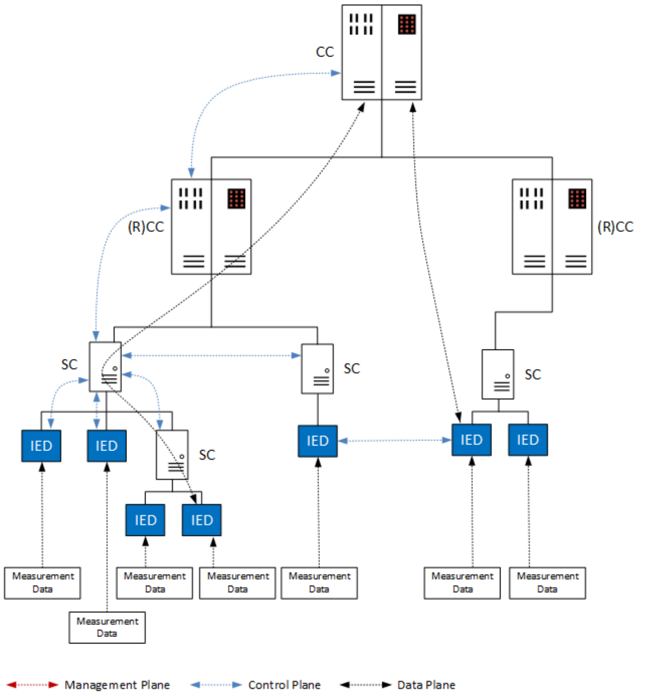

Figure 2: Elements in Electrical Power Systems include Central Control Centers (CC), Regional Control Centers (RCC), Substation Controllers (SC) and Intelligent Electronic Devices (IEDs)

casted or actual power load conditions, which may be handled by migrating functionality from central control centers to IEDs—or vice versa—to perform locally coordinated actions in the system.

- **Characteristics of the underlying information technology (IT) infrastructure**, such as availability or processing load, which may be handled by migrating functionality from IEDs to other IEDs or central control centers—or vice versa—to ensure functionality is performed.

For example, IEDs in a smart grid can be instrumented with sensors *e.g.*, for synchronized phasor measurement also called phasor measurement units (PMUs) to collect data about the health and status of power system in real-time. Online analytics performed at the IED, substation controller, and/or control center elements in a smart grid hierarchy (see Figure 2) can use this data to derive the context in which a smart grid is operating at any moment. Results from such context-aware analytics can then be used to navigate through transient load conditions and dynamically derive system configuration and protection parameters. In turn, these parameters can be used to reconfigure a grid so it adaptively sheds load, while minimizing service disruptions to energy producers, consumers, and traders.

Just as analytics is functionality that can be deployed flexibly and adaptively at various locations within the smart grid hierarchy where and whenever it is needed, other functionality can benefit from such context awareness. As described in Section 1, next-generation smart grids have decentralized volatile generation, new markets, as well as variable and responsive consumption. These capabilities yield varying local/regional characteristics, as well as overall increased dynamics compared to legacy centralized power systems. It is therefore important that monitoring, protection, and control functionality be deployed with adaptive parameters that leverage awareness of the context in which system resources and operations occur.

For example, higher sampling rates may be needed when a power system undergoes transient dynamics. Likewise, adaptive protection mechanisms are needed when abnormal stress exists in the system, but certain lines have additional capacity due to local weather conditions (such as cooler temperatures). As more digital components are introduced into power systems, more sources of potential failure exist. It is therefore essential that these systems can adaptively migrate functionality from one failed IED to a neighboring IED with available and appropriate computing resources and connections.

4. CHARIOT: CONTEXT-DRIVEN ORCHESTRATION MIDDLEWARE

CHARIOT [9, 10] is context-driven orchestration middleware that supports (1) model-based definition of goals that specify applications required by an IoT system, (2) the composition of—and the requirements imposed by—the applications, and (3) the constraints governing the deployment and (re)configuration of applications. CHARIOT uses these services to support initial application deployment, failure avoidance, fault management, and operations management of distributed IoT systems.

Goal-based system description model. CHARIOT use the concept of goal-based system description to describe the IoT system being managed. IoT system goals are defined as high-level capabilities that can be decomposed into smaller sub-capabilities using the concept of *capability decomposition*, as shown in Figure 3. Next, these capabilities are mapped onto logical components that provide the capabilities. These components can have inter-dependencies that are also captured with CHARIOT's goal-based model.

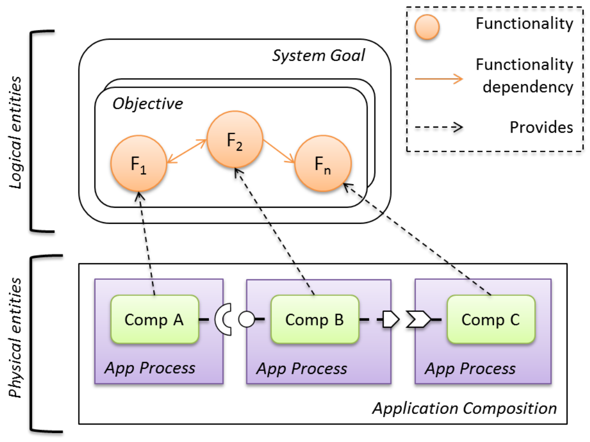

Figure 3: Example of a Goal-based Description of an Application.

Runtime architecture. CHARIOT provides a monitoring and deployment infrastructure, as well as a novel management engine that uses IoT system information stored persistently to formulate *Satisfiability Modulo Theories* (SMT) constraints that encode system properties and application requirements. CHARIOT can therefore use SMT solvers (such as Z3 [2]) to dynamically compute optimal system (re)configuration at run-time. CHARIOT's distributed architecture allows any node in the system to compute the required solution and then automate the resulting deployment to available system resources. Failure detection, group management, and leader election in CHARIOT is provided by Zookeeper [6].

CHARIOT's approach to IoT system management is based on the concepts of *configuration space* and *configuration points*. If a system's state is represented by a configuration point in a configuration space, then any change that invalidates a configuration point requires moving the IoT system from one configuration point to another in the same configuration space. As such, a configuration space includes (1) a goal-based description of different IoT systems, (2) replication constraints corresponding to redundancy patterns associated with different systems, (3) constraints on valid component-to-node deployment mappings, and (4) available resources, including different nodes and their corresponding resources, such as memory, storage, and computing elements.

At its core, CHARIOT encodes the constraints in the form of a matrix of decision variables. A component-to-node (C2N) matrix comprises rows that represent component instances and columns that represent nodes. The size of this matrix is $\alpha \times \beta$, where α is the number of component instances and β is the number of available nodes.

Each element of a C2N matrix is encoded as an integer variable whose value can either be 0 or 1. A value of 0 for an element means that the corresponding component instance (row) is not deployed on the corresponding node (column). Conversely, a value of 1 for an element indicates deployment of the corresponding component instance on the corresponding node.

CHARIOT uses other matrices to encode the resource availability of nodes (updated dynamically using the monitors). Communication resource requirements are encoded using a square matrix $\beta \times \beta$. The placement constraints are therefore written in terms of resources required and net resources available. Additional constraints for redundancy and collocation are also added. The solution of the constraint problem is a feasible placement answer, wherein "placement" means an instance of the C2N matrix.

Using Context and Goal Information for Runtime Management and Reconfiguration. As shown in Figure 4, CHARIOT's *Monitoring Infrastructure* is responsible for detecting changes in the *sensing* phase. After detection and diagnosis, its *Management Engine* determines the goals that are required based on the current context of an IoT system. This context is derived from known global system states. Thereafter, actions needed to reconfigure the system such that changes are handled are calculated in the *planning* phase based on Z3 [2], which is an open source SMT solver. After reconfiguration actions are computed, CHARIOT's *Deployment Infrastructure* is responsible for taking those actions to reconfigure the system in the *acting* phase.

5. CASE STUDY: ADAPTIVE PROTECTION FOR SMART GRIDS

This section presents a case study that expands upon the importance of context discussed in Section 2 by motivating the role of context-aware middleware in adaptive protection for smart grids. We first explain how adaptive protection mechanisms can leverage contextual information to enhance the resilience of next-generation smart grids. We then describe our ongoing work on modeling and managing the protection system assembly with the CHARIOT goal based description concept.

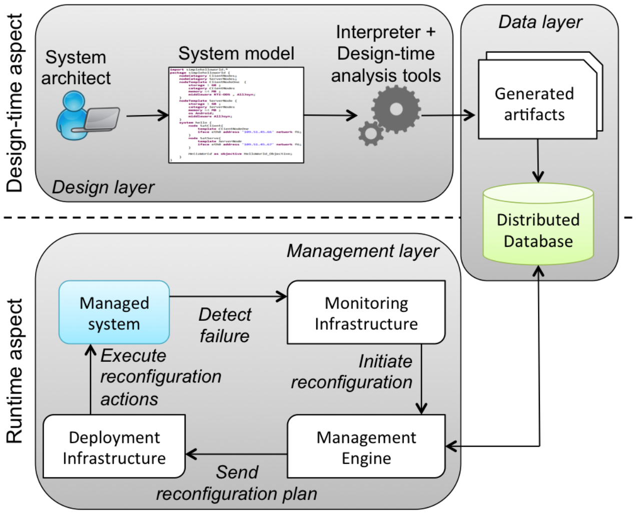

Figure 4: CHARIOT's Self-reconfiguration Mechanisms.

5.1 Adaptive Protection Mechanisms

Adaptive protection in a power system involves the use of adjustable protective relay settings (*e.g.*, current, voltage, feeders, and equipment) that can change in real-time based on signals from local sensors or a central control system [14]. Although these concepts have been studied since the 1990s [11] they are only recently are becoming practical with the increasing penetration of IEC 61850 capable IEDs, IEC 61850 based substation data model, functional and communication description and configuration, as well as the availability of phasor measurement units (PMUs) that deliver a real-time view of wide area power system status and health through GPS-synchronized measurements.

In modern substations that comply with IEC 61850 , substation management via the substation bus (including data management, computing, and communication performed with IEDs) is separated from the so-called process bus, as shown in Figure 5. The process bus handles communicative connection to sensors and actuators, such as circuit breakers, which are directly coupled to physical electrical assets at the substation level, such as transformers and loads. The decoupling of physical assets and actuators from data management and computing, including object-oriented data and communication standardization for substation automation yields more flexible configurability. This configuration flexibility can also enable context-aware reliable adaptivity of specific smart grid capabilities, such as adaptive protection.

The design of adaptive protection scheme may vary greatly, depending on the physical system assets (such as PMUs, static var compensators (SVCs), and power generation types) and the type of protection required (such as distance protection or overcurrent protection) . For example, adaptive

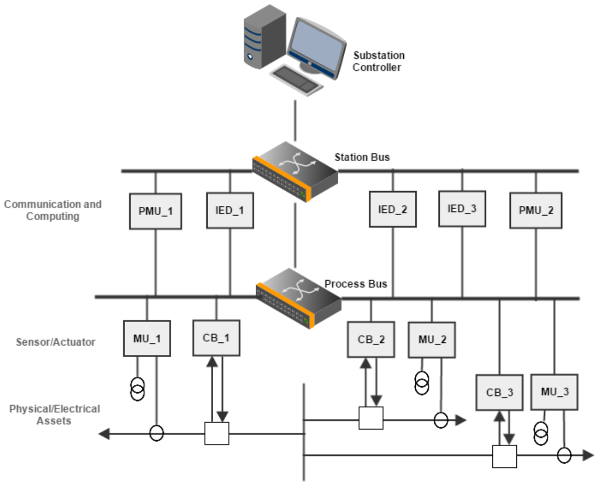

Figure 5: Modern substation configuration can be used for adaptive protection of Smart Grids

distance protection [17] is based on flexibility from an SVC and real-time system measurements from PMUs. SVCs are typically placed along a transmission line to countermeasure voltage drops by providing fast-acting reactive power on the high-voltage transmission line.

The flexibility introduced through the SVC, however, must be accommodated by the ability to *adapt* distance protection settings dynamically online. For example, it may be necessary to determine the actual system state via PMUs, the injected current via the SVC, as well as calculate the apparent impedance of a transmission line. Without this dynamic information, the pre-configured distance relays may underreach or over-reach, depending on whether the shunt injected current is negative or positive and the apparent impedance is either increased or decreased, respectively.

New settings can be determined for multiple protection zones by comparing synchronized voltage measurements at both ends of the transmission line, as well as online calculation of the apparent impedance, which takes into account the synchronized measurement of the shunt injected current. The calculations of adaptive settings are triggered when any changes in the compensation levels are detected. When a fault occurs, this adaptive protection setting accounts for the actual compensation levels and apparent impedance of the line before a signal to circuit breakers are issued to clear the fault.

The adaptive protection scheme described above increases the cost of communication and coordination between the IEDs and availability of real-time data, In addition, conventional adaptive protection mechanisms assume that the configured IEDs are available, *i.e.*, they have sufficient computing capacity to adapt set points and no communication error occurs between the configured distance relays on both sides of the transmission line nor between configured IED and circuit breakers within the substation.

5.2 Adaptive Protection in CHARIOT

This section describes how to apply CHARIOT to model a simplified three bus system with three different protection assemblies. We assume that each bus has a PMU, a relay, and a breaker. Figure 6 shows the three different functionalities (relay, pmu and the breaker) related one of the bus protection system (the other two protection systems are omitted for brevity). The composition shown in the figure describes the communication patterns between the functionalities.

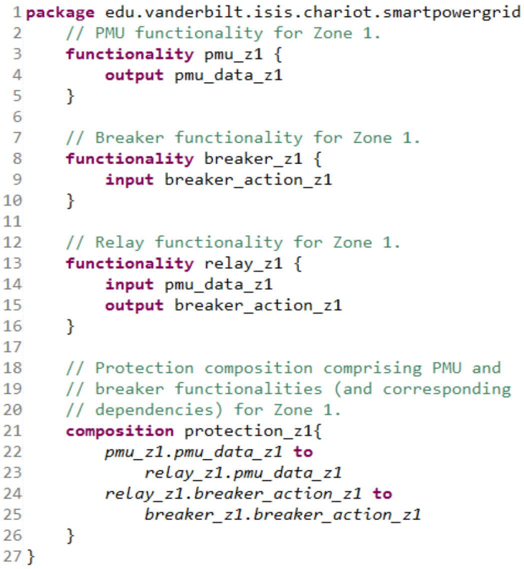

```
package edu.vanderbilt.isis.chariot.smartpowergrid {
    // PMU functionality for Zone 1.
    functionality pmu_z1 {
        output pmu_data_z1
    }

    // Breaker functionality for Zone 1.
    functionality breaker_z1 {
        input breaker_action_z1
    }

    // Relay functionality for Zone 1.
    functionality relay_z1 {
        input pmu_data_z1
        output breaker_action_z1
    }

    // Protection composition comprising PMU and
    // breaker functionalities (and corresponding
    // dependencies) for Zone 1.
    composition protection_z1{
        pmu_z1.pmu_data_z1 to
            relay_z1.pmu_data_z1
        relay_z1.breaker_action_z1 to
            breaker_z1.breaker_action_z1
    }
}
```

Figure 6: Snippet of Functionality and Composition Description in CHARIOT-ML.

Figure 7 shows the goal based breakdown of the three different compositions (one each for a protection zone). It also shows how CHARIOT can describe the replication constraints (replicate per node) to ensure as many PMUs as possible as available in a protection zone. The per-node constraint ensures multiple instances of a PMU are automatically configured if there are multiple PMU devices in a zone. Based on the composition and the system specification, CHARIOT can identify and start the relay instance in either zone 2 or zone 3 if the relay in the protection zone 1 fails. The objective instance (for example, protection_z1 in figure 7) provides the context for the relay component when it is started on a different IED. Note that the IED node model has not been shown in the figure.

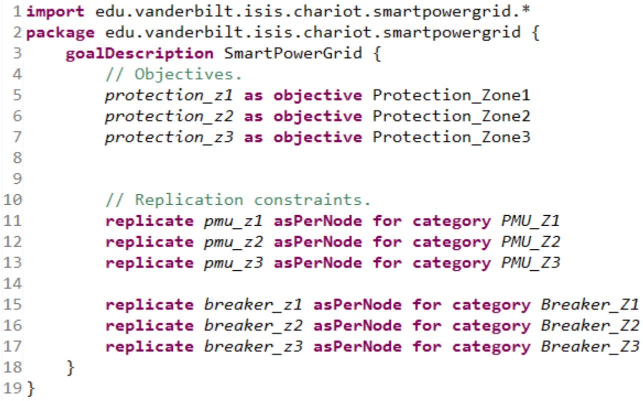

```
import edu.vanderbilt.isis.chariot.smartpowergrid.*
package edu.vanderbilt.isis.chariot.smartpowergrid {
    goalDescription SmartPowerGrid {
        // Objectives.
        protection_z1 as objective Protection_Zone1
        protection_z2 as objective Protection_Zone2
        protection_z3 as objective Protection_Zone3


        // Replication constraints.
        replicate pmu_z1 asPerNode for category PMU_Z1
        replicate pmu_z2 asPerNode for category PMU_Z2
        replicate pmu_z3 asPerNode for category PMU_Z3

        replicate breaker_z1 asPerNode for category Breaker_Z1
        replicate breaker_z2 asPerNode for category Breaker_Z2
        replicate breaker_z3 asPerNode for category Breaker_Z3
    }
}
```

Figure 7: Snippet of System Description in CHARIOT-ML.

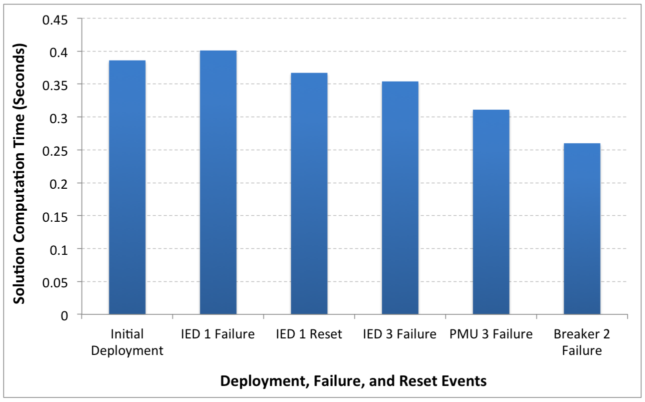

Figure 8: Total Solution Computation Time for Different Events for the Three Bus System (refer to section 5.2).

Figure 8 demonstrates the time taken by CHARIOT middleware to compute solution for different events related to the three bus system. We performed a small-scale (three protection zone) simulation experiment on a single Windows 7 machine with 8 GB memory and 8 cores. As shown in the figure, we consider three kinds of events: (1) initial deployment, (2) failure, and (3) reset. The first two events are self explanatory, the third event corresponds to the action of being able to sucessfully resetting the failed IED node. For the aforementioned small-scale system, the solution computation time is less than 0.4 seconds and we can see that the time to compute a solution decreases as the number of failures increases, *i.e.*, when the number of nodes in the system decreases. This is due to the reduction is search space.

6. CONCLUDING REMARKS

The use of *pre-configured* architectures in legacy power systems tightly couples functionality in relatively fixed locations that makes it hard for these systems to meet changing environments and requirements in the energy sector. In particular, the inflexibility of legacy power system architectures increases the time, effort, and costs associated with developing, operating, and evolving these systems into next-generation smart grids.

To avoid accumulating excessive technical dept it is essential to consider the role of context in power system domains that are incurring greater rates of change and volatility. In particular, without enhanced resilience and adaptivity in the digital domain, the accidental and inherent complexity of smart grids would be magnified to the point that the technical debt and other demands imposed by the financial and physical domains could not be satisfied.

Modeling and managing IoT systems with dynamic orchestration middleware like CHARIOT helps digitally-enhance smart grid operations to more effectively manage risks arising from volatility in the power system domain by flexibly determining where to deploy functionality throughout a grid hierarchy. Meeting these needs motivated our work on context-aware orchestration middleware.

7. REFERENCES

[1] Annual Energy Outlook 2014. http://www.eia.gov/forecasts/aeo/pdf/0383(2014).pdf.

[2] L. M. de Moura and N. Bjørner. Z3: An efficient smt solver. In *TACAS*, pages 337–340, 2008.

[3] A. R. Di Fazio, T. Erseghe, E. Ghiani, M. Murroni, P. Siano, and F. Silvestro. Integration of renewable energy sources, energy storage systems, and electrical vehicles with smart power distribution networks. *Journal of Ambient Intelligence and Humanized Computing*, 4(6):663–671, 2013.

[4] G. T. Heineman and B. T. Councill. *Component-Based Software Engineering: Putting the Pieces Together.* Addison-Wesley, Reading, Massachusetts, 2001.

[5] S.-C. Hsieh, C.-S. Chen, C.-T. Tsai, C.-T. Hsu, and C.-H. Lin. Adaptive relay setting for distribution systems considering operation scenarios of wind generators. *IEEE Transactions on Industry Applications*, 50(2):1356–1363, 2014.

[6] P. Hunt, M. Konar, F. P. Junqueira, and B. Reed. ZooKeeper: Wait-free Coordination for Internet-scale Systems. In *Proceedings of the 2010 USENIX conference on USENIX annual technical conference*, volume 8, pages 11–11, 2010.

[7] R. Mackiewicz. Overview of iec 61850 and benefits. In *2006 IEEE PES Power Systems Conference and Exposition*, pages 623–630. IEEE, 2006.

[8] V. A. Papaspiliotopoulos, G. N. Korres, and N. D. Hatziargyriou. Protection coordination in modern distribution grids integrating optimization techniques with adaptive relay setting. In *PowerTech, 2015 IEEE Eindhoven*, pages 1–6. IEEE, 2015.

[9] S. Pradhan, A. Dubey, T. Levendovszky, P. S. Kumar, W. A. Emfinger, D. Balasubramanian, W. Otte, and G. Karsai. Achieving resilience in distributed software systems via self-reconfiguration. *Journal of Systems and Software*, pages –, 2016.

[10] S. M. Pradhan, A. Dubey, A. Gokhale, and M. Lehofer. Chariot: A domain specific language for extensible cyber-physical systems. In *Proceedings of the Workshop on Domain-Specific Modeling*, DSM 2015, pages 9–16, New York, NY, USA, 2015. ACM.

[11] G. Rockefeller, C. Wagner, J. Linders, K. Hicks, and D. Rizy. Adaptive transmission relaying concepts for improved performance. *IEEE Transactions on Power Delivery*, 3(4):1446–1458, 1988.

[12] E. Santacana, G. Rackliffe, L. Tang, and X. Feng. Getting smart. *IEEE Power and Energy Magazine*, 8(2):41–48, March 2010.

[13] P. Siano. Demand response and smart grids - a survey. *Renewable and Sustainable Energy Reviews*, 30:461–478, 2014.

[14] SmartGrid.Gov. Smart grid definitions of functions.

[15] J. F. Stevens and J. H. Allen. The smart grid: Managing electrical power distribution and use.

[16] J. Sztipanovits. Embedded software and systems: Challenges and approaches. In *Proceedings of the 4th International Workshop on Hybrid Systems: Computation and Control*, HSCC '01, pages 4–, London, UK, UK, 2001. Springer-Verlag.

[17] M. P. Thakre and V. S. Kale. An adaptive approach for three zone operation of digital distance relay with static var compensator using pmu. *International Journal of Electrical Power & Energy Systems*, 77:327–336, 2016.