

Received April 27, 2018, accepted May 31, 2018, date of publication June 27, 2018, date of current version July 25, 2018.

Digital Object Identifier 10.1109/ACCESS.2018.2851186

Leveraging Standards to Create an Open Platform for the Development of Advanced Distribution Applications

RONALD B. MELTON¹, (Senior Member, IEEE),
KEVIN P. SCHNEIDER¹, (Senior Member, IEEE), ERIC LIGHTNER²,
THOMAS E. MCDERMOTT¹, (Fellow, IEEE), POORVA SHARMA¹, (Member, IEEE),
YINGCHEN ZHANG³, (Senior Member, IEEE), FEI DING³, (Member, IEEE),
SUBRAMANIAN VADARI⁴, (Senior Member, IEEE), ROBIN PODMORE⁵, (Life Fellow, IEEE),
ANAMIKA DUBEY⁶, (Member, IEEE), RICHARD W. WIES⁷, (Senior Member, IEEE),
AND ERIC G. STEPHAN¹

¹Pacific Northwest National Laboratory, Richland, WA 99352, USA

²U.S. Department of Energy, Washington, DC 20021, USA

³National Renewable Energy Laboratory, Golden, CO 80401, USA

⁴Modern Grid Solutions, Redmond, WA 98073, USA

⁵Incremental Systems Corporation, Bellevue, WA 98007, USA

⁶Department of Electrical Engineering and Computer Science, Washington State University, Pullman, WA 99164, USA

⁷Department of Electrical Engineering, University of Alaska Fairbanks, Fairbanks, AK 99775, USA

Corresponding author: Ronald B. Melton (ron.melton@pnnl.gov)

This work was supported in part by the U.S. Department of Energy with the Battelle Memorial Institute under Contract DEAC05-76RL01830, in part by the Manager and Operator of Pacific Northwest National Laboratory with Alliance for Sustainable Energy, LLC, under Contract DE-AC36-08GO28308, and in part by the Manager and Operator of the National Renewable Energy Laboratory.

ABSTRACT Modern electric power distribution systems are data rich and include growing numbers of distributed energy resources and distribution automation. To take advantage of distribution automation and manage growing penetrations of distributed energy resources, distribution utilities need applications for planning and operations that use all available data and may incorporate distributed approaches to operate and control. The industry would benefit from distribution management applications based on a common platform that makes systems of each type interchangeable. This paper describes an approach to enabling cost-effective development and deployment of advanced applications for distribution system planning and operations based on development of an open-source, standards-based platform for application development called GridAPPS-D, which leverages data abstractions for application development based on standards, such as the distribution system common information model.

INDEX TERMS Distribution system analysis, dynamic simulation, power modeling, power simulation, smart grid.

I. INTRODUCTION

Modern electric power distribution systems offer an abundance of data, and distributed energy resources (DER) and distribution automation are growing. Distribution utilities would benefit from an improved ability to take action on this data. However, a traditional distribution management system (DMS) is typically isolated from, or at best loosely coupled with, related systems such as distributed energy resource management (DERMS), supervisory control and data acquisition (SCADA), outage management systems (OMS), geographic information systems (GIS), advanced metering infrastructure (AMI), customer information systems (CIS), meter data management systems (MDMS), and fault location,

isolation, and service restoration (FLISR). To take advantage of distribution automation and manage growing penetrations of DER, distribution utilities need an integrated approach, with applications for distribution system planning, operations, management, and control that use all available data and may incorporate distributed approaches to operations and control [1]. The industry refers to software platforms that integrate multiple systems as advanced distribution management systems (ADMS). Unfortunately, distribution utilities have faced high costs to implement and integrate ADMS and related applications [2]. These high costs have limited the deployment of such systems to large utilities. Developing and deploying advanced distribution system planning, operations,

and management applications based on a common platform design would make them more broadly applicable and also interchangeable, which would benefit the entire industry.

This paper describes an approach to enabling cost-effective development and deployment of advanced applications for distribution system planning and operations. This approach is based on development of an open-source, standards-based platform for application development called GridAPPS-D [3], [4]. A key objective for GridAPPS-D is to reduce the time and cost required to develop, integrate, and deploy advanced applications.

An element of this approach is to define data abstractions for application development based on interoperability standards such as the Common Information Model (CIM) [5], [6]. This standards-based approach will support wide adoption of these capabilities and interchangeability of the applications.

Effective application of the standards-based approach is required to achieve the desired outcomes. In particular, in order to support the ability to develop applications that may be run for a variety of distribution system networks, an efficient means of referencing and interacting with the network model must be provided. The paper gives a summary of the challenges inherent to using relational database systems to store network models, followed by a description of how more appropriate database technology in the form of a “triple-store” database [7], [8] is used to overcome these challenges. Consideration of other NoSQL [9] database technology to support compound queries that can associate simulation results with reference network models based on the distribution system common information model (CIM, ISO/IEC 61968 and 61970) [5], [6] is also discussed.

We then describe the expected benefits of the advanced, data rich, data-driven applications, discuss the application of the GridAPPS-D platform to develop such applications, and conclude with a summary.

II. APPROACH – THE USE OF STANDARDS-BASED DATA ABSTRACTIONS, WELL-DEFINED DATA EXCHANGE MECHANISMS, AND TRIPLE-STORE DATABASE TECHNOLOGY

A. GridAPPS-D STRUCTURE

GridAPPS-D has been designed to address the challenges outlined in the previous section. It is designed to provide a services-based platform supporting the development of applications. GridAPPS-D is being implemented as an open-source system, providing a reference architecture and implementation that can be used by others in implementing similar application development tools or in adapting existing systems or creating new ones for operational deployment of applications that comply with our approach.

The conceptual architecture for GridAPPS-D is shown in Fig. 1. The figure depicts the logical functionality and architecture of GridAPPS-D. The application of standards-based data abstractions and data exchange mechanisms

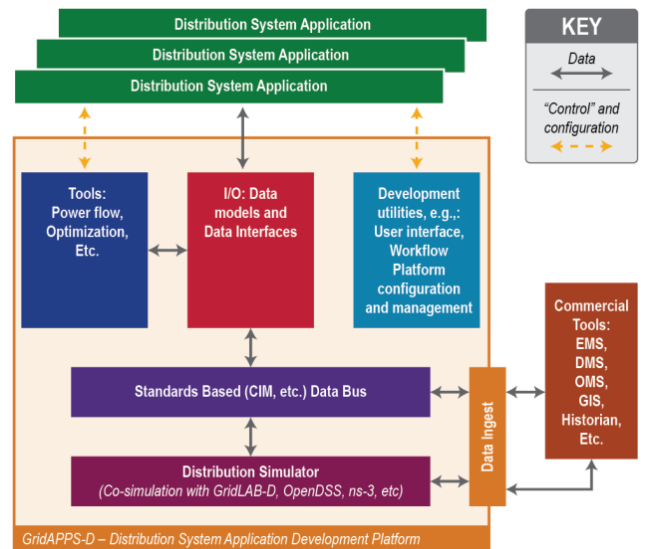


FIGURE 1. GridAPPS-D conceptual architecture.

described in the previous section is included in the I/O and in the data bus.

Other features of the platform include data management of network models, related information, and time-series data. Co-simulation is used so that application developers can simulate a variety of distribution system behaviors. “Data Ingest” provides the ability to exchange data with existing systems such as energy management systems (EMS), distribution management systems (DMS), outage management systems (OMS), geographic information systems (GIS), data historians, and so forth. A provision is included for access to tools such as optimizers. Finally, the ability to configure the platform is supported. This paper focuses on the features that are being implemented to support the use of CIM-based models, which enables the standardization of the application developer’s data model and data exchange. More details on the objectives and functionality of GridAPPS-D are available in its conceptual design [3].

B. STANDARDS-BASED ABSTRACTIONS

A key design objective for the GridAPPS-D platform is to enable development of advanced distribution system planning or operations applications that can be deployed at multiple utilities within different vendor systems and require minimal time and cost to integrate into the user’s environment—in other words, portable applications [3]. To facilitate this, the platform is to be standards-based to the greatest extent possible.

To achieve these objectives, the CIM and similar standards such as MultiSpeak® [10] were used, taking advantage of their information models and data abstractions to define the data abstractions used by application developers. Use of the CIM provides a consistent means of specifying system models and well-defined information models that support messaging within the system. The CIM specification makes use of

the Unified Modeling Language (UML) [11]. UML enables developers to specify, visualize, construct, and document artifacts such as data structures, software interactions, and use cases. The CIM uses UML to produce non-technology-specific specifications. These specifications serve two purposes for CIM: (1) semantically, they provide a human interface that clearly defines transmission and distribution system concepts that can be understood by application developers; (2) syntactically, CIM and UML act as blueprints that provide the basis for actual physical implementations of data structures and software interactions in applications, services, and database systems.

When the data models and abstractions are combined with a well-defined data exchange mechanism, application developers will be able to write portable applications with minimal integration cost as long as the target systems also adopt the data models and data exchange mechanism.

C. THE QUERY CHALLENGE

The use of a specified data exchange mechanism and standards-based data models and abstractions is not sufficient to enable portable applications. One must also consider the mechanism for applying the data models and abstractions to reference data in the target distribution system. Referencing objects and data in the target system is a form of querying the network model and the corresponding messages between network components. For an application to be portable, it must be able to make references that can be acted on for any network model that has components and data operated on by the application.

Historically, CIM-based network models have generally been managed in relational databases and queried through complex Structured Query Language (SQL). Unfortunately, relational structures make logical networks difficult to naturally represent [12], causing additional overhead to manage complex concepts and connectivity. The problem is compounded when highly sophisticated relational (SQL) queries must be developed to perform graph traversal and pattern discovery using large numbers of joins across CIM concepts. The object-class hierarchy of the network model is decomposed into numerous linked relational tables. Special query-building tools are then used to interact with the model using the exchange format queries and exchanged using CIM XML, a variation of the Resource Descriptive Framework schema (RDFS) that depicts class associations and aggregation structures in graph form and simplifies class inheritance by prefixing inherited properties with parent class names [13]. In its native form, client libraries are used to support queries of the CIM XML model.

However, the use of relational database technology is an impediment to enabling development of portable applications, so alternative approaches were explored that support more direct querying and management of network models.

III. ENABLING DEVELOPMENT OF PORTABLE APPLICATIONS

A. MATCHING DATABASE TECHNOLOGY TO THE PROBLEM

As mentioned, GridAPPS-D is designed to support the development of advanced, data-driven distribution system management and/or planning applications that take advantage of the data-rich environment expected in modernized electric power distribution systems. In addition, these applications are intended to be portable, that is, able to be easily deployed on any vendor system or infrastructure that supports the data abstraction, data model, and data exchange mechanism defined by the GridAPPS-D reference architecture and implementation.

The application developer, then, uses the GridAPPS-D data abstractions and messages to make references to components of the network model, data exchanged between components of the network model, and changes to settings of those components. The structure of these references follows the RDFS form of subject-predicate-object—in other words, triples. This form naturally matches the capabilities of triple-store and graph database technology [14]. The use of a triple-store database technology for network models coupled with a hybrid time-series store [15] frees the application developer from having to know in advance the details of the specific network model. In effect, the application is able to adapt itself to the specific network model.

The selection of triple-store technology was based on a formal evaluation of the benefits and performance of three database technologies: conventional relational database management system (RDBMS) technology, an open-source graph database, and an open-source triple-store database. The details of this evaluation are included in [16] and summarized in Table 1. In the table, white indicates that the technology did not meet the requirement, light shading that the requirement was partially met, and dark shading that the requirement was fully met. Triple-store technology was chosen, allowing the CIM XML reference model to be managed and queried in its native format and eliminating the overhead of representing and querying reference model graph data in RDBMS form.

TABLE 1. Strengths of three candidate database technologies.

Requirements	Triple-		
	RDBMS	Store	Graph
Administrative free schema maintenance	Dark	Light	White
Ability to intuitively represent CIM objects	Dark	Dark	White
Schema evolution free solution	Dark	Dark	White
Value type and referential integrity checking	Dark	Dark	White
Indexing	Dark	Dark	White
Supports data management language	Dark	Dark	White
Data partitioning	Dark	Dark	White
Directed graph	Dark	Dark	White
Polymorphism	Dark	Dark	Light
Ease of technology transfer	Dark	Dark	White
Community involvement	Dark	Dark	White
Product maturity	Dark	Dark	White

A key point of the triple-store approach is the ability to directly specify data via the object–attribute–value statements in code, using the SPARQL query language. In contrast, the use of RDBMS technology requires a query-building mechanism that must be regenerated for each individual network model, thus impeding the portability of the applications.




FIGURE 2. Circuit diagram and photo of the common single-phase center-tapped secondary service transformer.

B. TRIPLE-STORE USE CASE – TRANSFORMER

Use of the triple-store approach is illustrated by considering an interaction with a single-phase, center-tapped secondary transformer component, as shown in Fig. 2. There are hundreds (or even thousands) of these on most radial distribution feeders. Compared to the balanced three-phase transformers commonly found in transmission applications of the CIM, this example is more complicated in many respects:

- There is one phase on the primary (A, B, or C) and two on the secondary (s1, s2), plus the neutral (N).
- There are three windings, but none of them are delta or wye. Instead, they are individual phase connected (Iii).
- The secondary center tap is connected to neutral, which creates 240 volts between s1 and s2 for larger loads and DERs. The polarity of s1's winding is opposite from that of s2, and this must be accounted for in the CIM.
- These transformers are purchased and stocked in bulk, so the ratings and impedances are stored in a catalog of data. The catalog increases efficiency and helps to enforce accuracy and consistency in the models.

Tank-level transformer modeling and the AssetInfo pattern in CIM provide the flexibility needed to represent single-phase service transformers. Fig. 3 shows the UML instance diagram of a transformer connected to phase C of a bus (i.e., ConnectivityNode in CIM). The transformer has three windings (i.e., Ends in CIM), each connected to its own terminal, and the TransformerTankEnd instances carry the phasing information for transformer tank modeling. (See [16] for

CIM unbalanced line examples, which use a different phasing pattern). The primary Terminal connects to the primary ConnectivityNode, while both secondary Terminals connect to the secondary ConnectivityNode. The transformer catalog voltages and polarities appear near the bottom of Fig. 3. It is essential that endNumber values for each TransformerTankEnd match those of the corresponding TransformerEndInfo. There are other catalog attributes not shown in Fig. 3, but it is important to note the phaseAngleClock value of 6 for end-Number 3. This reverses the polarity of windings connected to s2 in Fig. 2.

Fig. 4 through Fig. 7 illustrate a four-query process of retrieving all the service transformer data. Each such illustration shows a portion of the CIM UML class diagram to the left, a SPARQL query to the upper right, and a snippet of output to the lower right. A SPARQL query consists of triples that act to filter results from the database. Each triple selects a type (i.e., CIM class), an attribute value, or an association (i.e., a link on the UML diagram). Each color on these illustrations matches a feature of interest on the UML diagram with a SPARQL triple that filters on that feature. SPARQL syntax is similar to that of SQL, but usage for CIM is much simpler because JOIN clauses or supplemental tables are not needed for polymorphic associations. SPARQL queries can be written only using the UML, which serves as the database schema. The order of triples does not matter, except for some performance differences if the database does not automatically optimize queries.

The four query steps are

1. Retrieve the catalog volt-ampere and voltage ratings, with corresponding endNumber, phaseAngleClock, and winding resistance values, as shown in Fig. 4. These are grouped in three rows per tank; a 75-kVA tank is shown in the output. The DESC clause is used here only to sort the output from a large database so that the 75-kVA tank appears first. There are only two CIM classes involved in this query.
2. Retrieve the catalog no-load test data (Fig. 5). This is a shorter query but involves three CIM classes. Normally, the CIM units are SI, but this query actually models the transformer test report, not the transformer itself. The no-load loss units are kW, and the exciting current units are percent. This query is agnostic to the no-load test winding, but it could return more than one row, unexpectedly, if there is more than one no-load test per transformer.
3. Retrieve the catalog short-circuit test data (Fig. 6). This models another test report, so the units of loss are kW and the units of impedance are percent. Because these are three-winding transformers, we expect three pairwise short-circuit test results, which do appear in Fig. 6.
4. Connect the transformer to the circuit with proper phasing, and reference the proper catalog data (Fig. 7). For transformer tank modeling, the CIM Terminal phases attribute must be ignored, and it is crossed out at

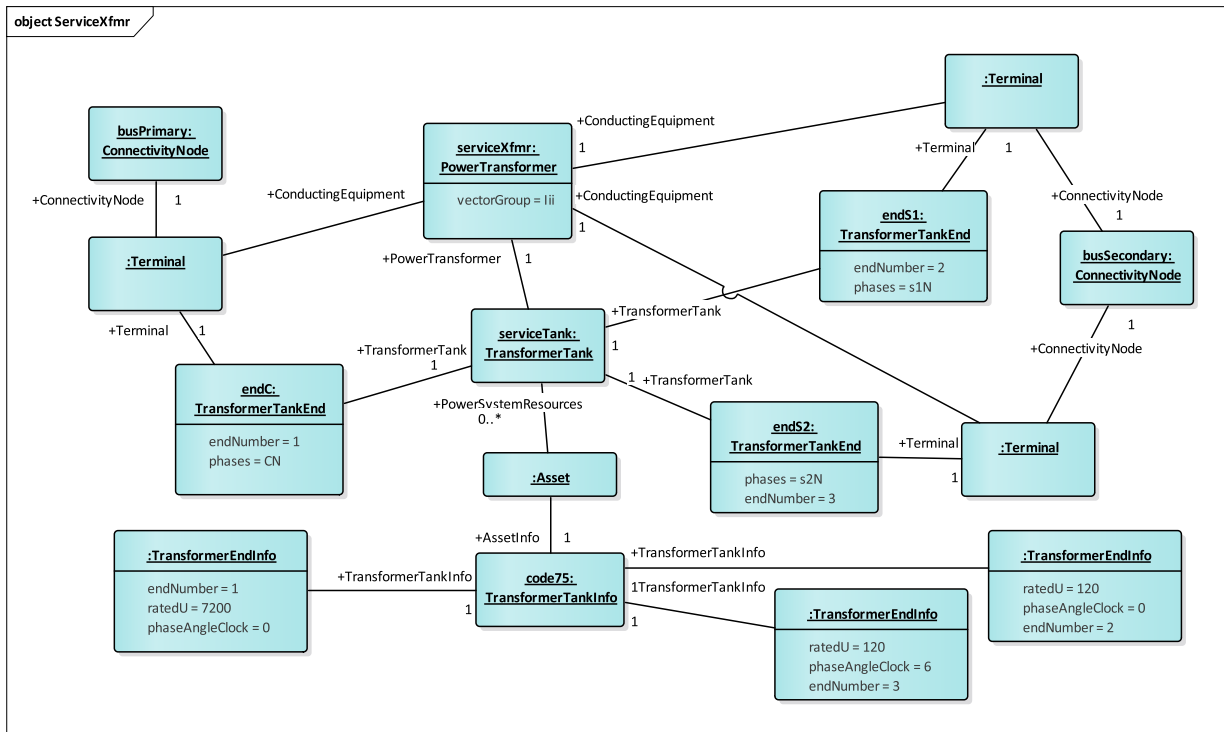


FIGURE 3. UML instance diagram for a service transformer on phase C, referencing catalog data.

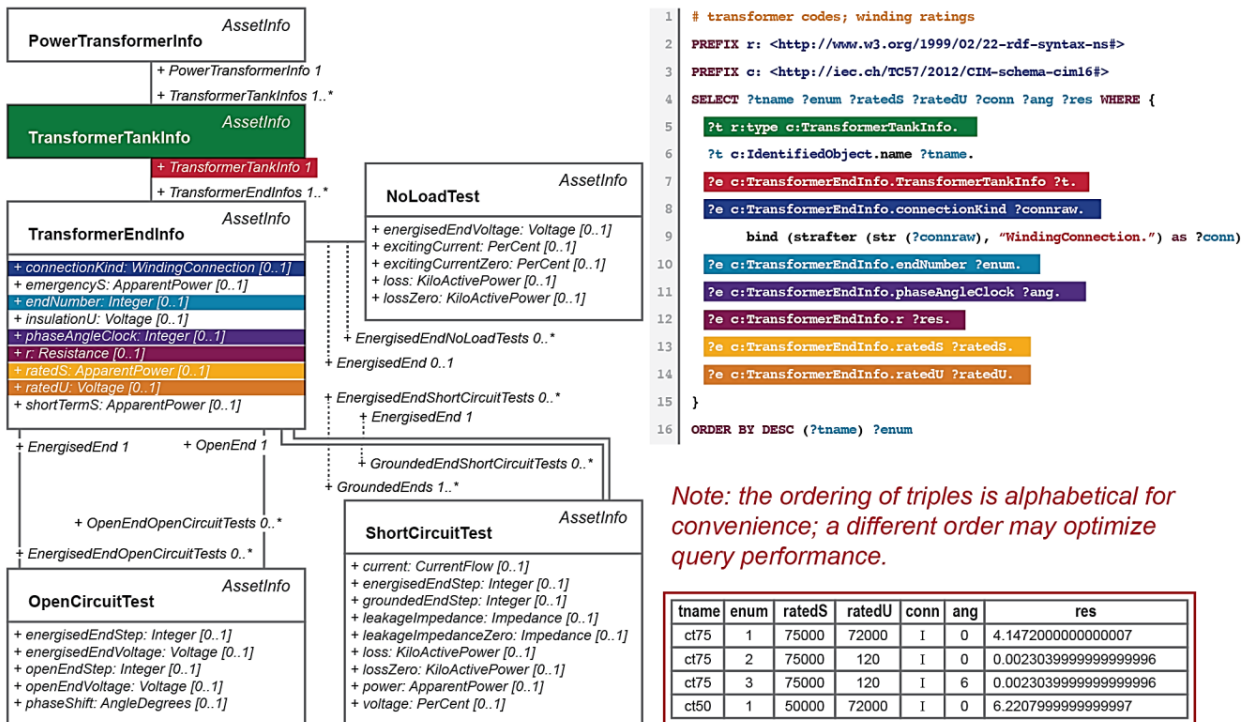


FIGURE 4. Step 1: Cataloging the transformer winding ratings.

the upper left of Fig. 7. Line 17 of the SPARQL matches the instance end with the catalog endNumber, which is essential if we need the correct winding rating

information for this output. In this example, we show how to retrieve the system nominal voltage instead of the winding rated voltage. These can be different;

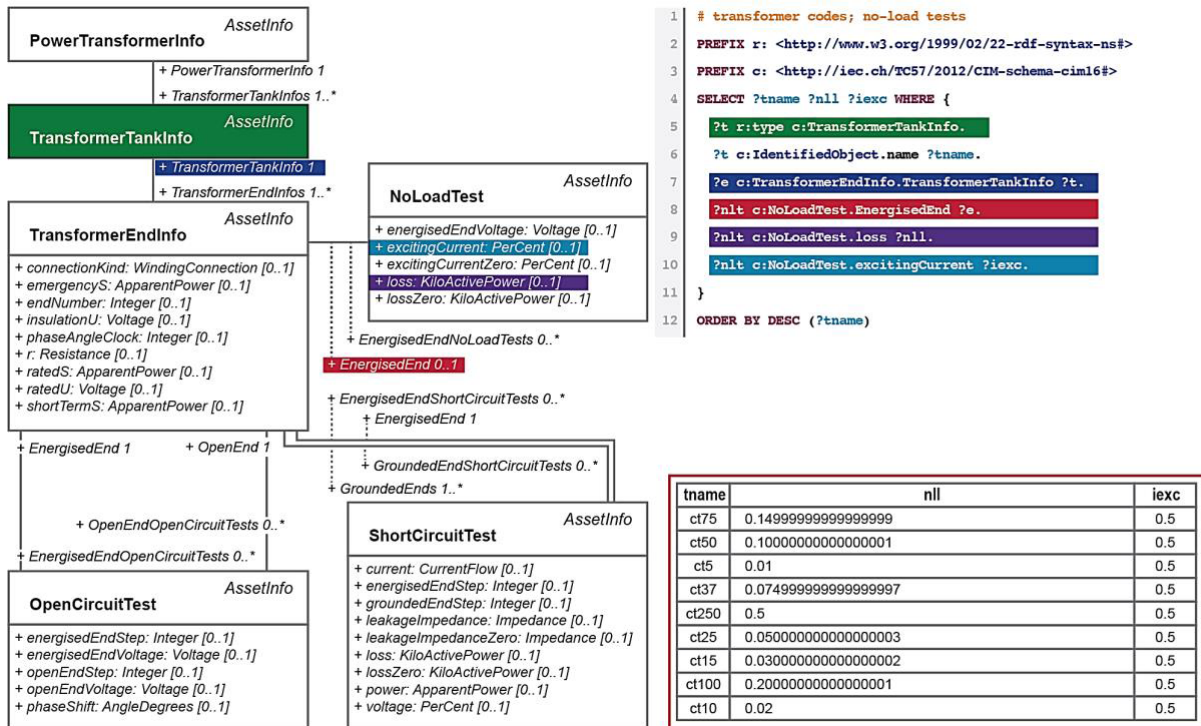


FIGURE 5. Step 2: Retrieving no-load test data for catalog entries.

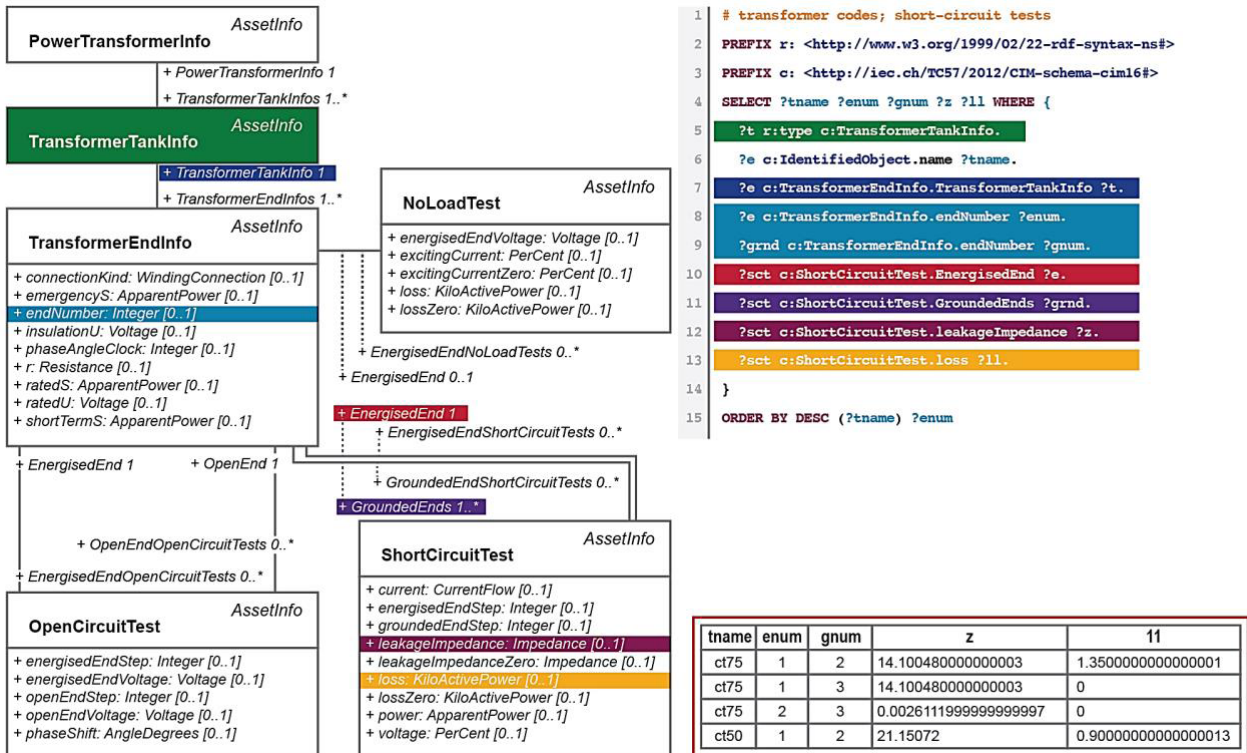


FIGURE 6. Step 3: Retrieving short-circuit test data for catalog entries; both ends have to match.

e.g., a 12.5-kV transformer on a 12.47-kV system. As displayed in Fig. 7, the xfmrcode and enum values point to the correct catalog values for this instance.

This example is probably the most complicated of those used in GridAPPS-D for distribution power flow modeling. The queries could be partitioned differently, depending on the

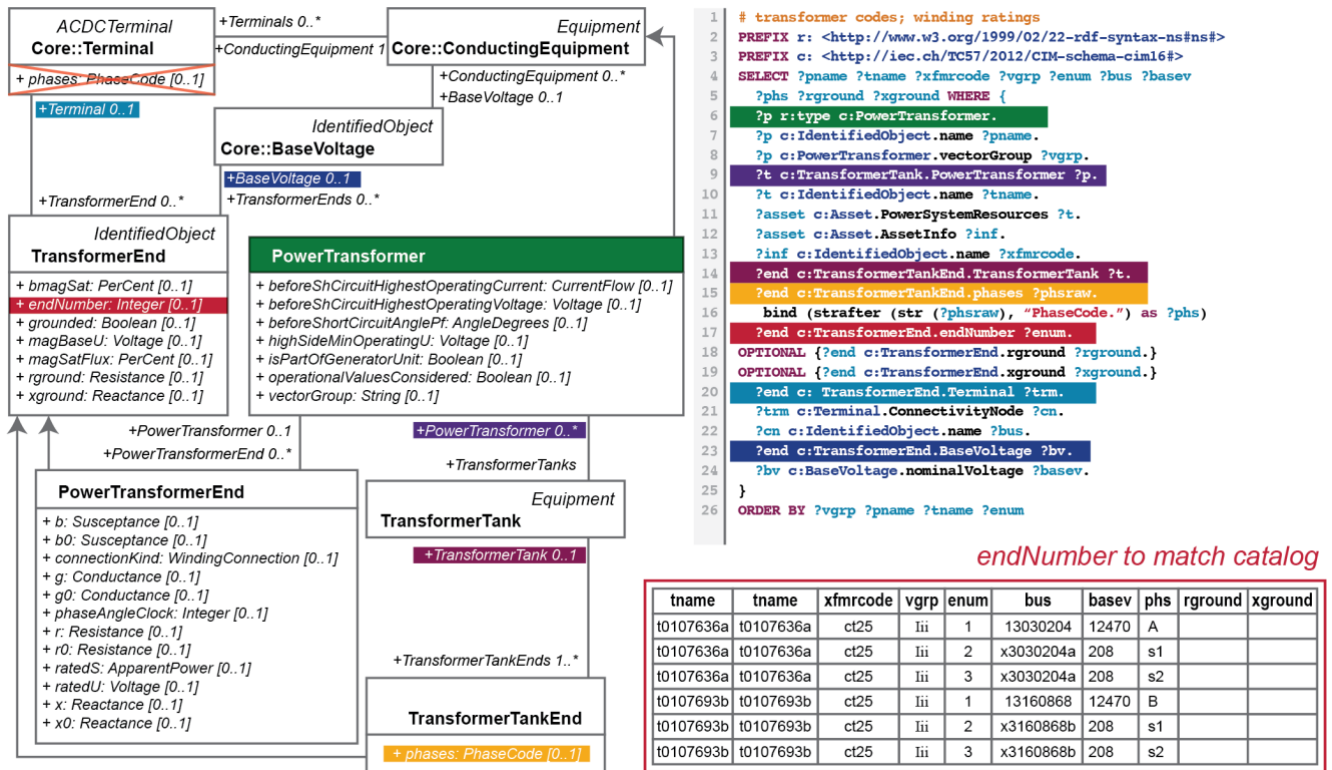


FIGURE 7. Step 4: Connecting transformer to catalog data, phases, and connectivity nodes.

application’s needs. Fig. 4 through Fig. 7 are convenient for the simulators used in GridAPPS-D [17], [18], and for explanation here. A Java component of GridAPPS-D traverses the SPARQL result sets to write the input data for the simulator such as GridLAB-D. Support for multiple simulators is planned.

IV. EVALUATING THE BENEFITS OF ADVANCED APPLICATIONS

In parallel with establishing an approach to developing advanced applications, a way to evaluate the potential benefits of such applications was also considered [19]. The evaluation process was treated as a test of the hypothesis that the advanced applications will provide measurable benefits as compared to conventional applications using limited data.

Twelve applications are being developed for evaluation. These applications are summarized in Table 2. They include “end-use” applications intended for application in utility operations and “service” applications that provide capabilities for use by other applications. The functional requirements for GridAPPS-D explicitly recognize the possibility of both types of applications. The design of the platform supports interaction between applications.

Many applications in Table 2 perform overlapping functions. For example, two perform volt-var control, and several incorporate or adapt load forecasting. This overlap

TABLE 2. Example applications under development.

Title	Type
Circuit visualization and GUI services	Service
State estimator with load forecasting	Service
Model validation after GIS import	Off-line
Volt-var control with reconfiguration	Operations
Responsive (transactive) load control	Operations
Locational pricing forecasting	Operations
Solar forecasting from irradiance	Operations
DERMS optimal power flow	Operations
Restoration with DER	Operations
Volt-var optimization with DER	Operations
Cyber traffic monitor	Operations
Resilient microgrid control with storage	Operations

exemplifies one motivation for the platform: to foster the development of novel applications that use untried methods on real circuits. Naturally, some of these competing applications will prove more widely applicable than others. A utility should be able to mix-and-match applications to suit their needs.

Applications may be developed in C/C++, Java, MATLAB, or Python (most developers are using Python for the Table 2 examples). In addition, several services are provided for application developers:

1. SPARQL query access to the database (Section III).
2. State estimator voltage, current, switch, and tap outputs.

3. Circuit visualization with marking, labeling, and strip chart outputs.
4. Parameter configuration at start-up, and graphical user interface (GUI) forms to collect a limited number of user run time inputs.
5. Test case scripting and logging.
6. Access to time-series data.
7. Several installed toolkits, e.g., the simulators GridLAB-D and OpenDSS as services, and NumPy, SciPy, MySQL, VOLTRON, and MATLAB as run time capabilities.
8. Open-source optimization solvers such as Bonmin, IPOPT, and SCIP.
9. Open-source machine learning package, e.g., LIBSVM.

Note that the tools listed in items 7, 8, and 9 are examples and may or may not be distributed with GridAPPS-D. When not distributed with the platform, the user would have to install those toolkits separately, paying any required license fees, before the dependent application(s) could run. Software dependencies are important considerations for an application developer, in planning the transition from research to production.

V. APPLYING THE PLATFORM

GridAPPS-D currently runs in a Linux virtual machine (VM). Although it can be built from sources, the primary form of distribution is as a set of Docker containers [20]. Users can install the Docker infrastructure on their computer and then download the Docker containers. Several platform usage scenarios are then feasible:

1. Start and run the application through its browser interface. Utilities could use the platform this way to evaluate new applications, or to evaluate applications on their own circuits. The App Hosting Manager allows a user to install and configure new applications to run in the platform, by modifying configuration files but without having to write new code. GridAPPS-D will also be able to ingest any distribution circuit provided in CIM format.
2. Write scripted scenarios and responses using the Test Manager, and run those through GridAPPS-D. This mode can be used for a more rigorous evaluation, and also for operator training.
3. Write a new application, using one of the open-source examples as a template. This mode should provide a faster on-ramp for application developers to develop a standards-compliant product.
4. DMS vendors can use the platform to develop and test their own standards-compliant interfaces. Any GridAPPS-D code may be incorporated into a commercial product, pursuant to its BSD license terms. The goal is for an application to be deployable from one platform to another, simply by moving the program file(s) and updating local configuration files.

VI. SUMMARY

GridAPPS-D [4] is an open-source, standards-based platform built to support the development of a wide range of advanced, data-driven applications for electric power distribution system operations and planning. Through a well-defined data model, abstraction, and exchange mechanism, the platform offers the possibility of portable applications that can be deployed with less time and cost than ones custom-built for each system.

In implementing GridAPPS-D, a new approach to managing standards-based network models and messaging was employed: a combination of database management technologies with triple-store technology at the core. This novel approach has been shown to provide the basis for application portability and for improved management of CIM-based network models.

GridAPPS-D should enable the industry to begin to transition to development of portable applications that can be deployed within any operational environment that supports the GridAPPS-D programming model. This will enable utilities of all sizes to select affordable planning and operations applications that meet their individual needs. As a reference architecture and implementation, GridAPPS-D provides a starting point for vendors and other interested parties to apply the potential of portable applications as they develop new products and services or adapt existing offerings. It also provides a common platform for future research and development of advanced distribution system planning, operations, management, and control applications. The source code and documentation are available at <https://github.com/GRIDAPPSD/GOSS-GridAPPS-D> and <http://gridappsd.readthedocs.io/en/latest>.

REFERENCES

- [1] Y. Agalgaonkar, C. Marinovici, S. Vadari, K. Schneider, and R. Melton, "ADMS state of the industry and gap analysis," Pacific Northwest Nat. Lab., Richland, Washington, DC, USA, Tech. Rep. PNNL-26361, 2016. [Online]. Available: https://www.pnnl.gov/main/publications/external/technical_reports/PNNL-26361.pdf
- [2] U.S. Department of Energy. (Feb. 2015). *Voices of Experience: Insights Into Advanced Distribution Management Systems*. [Online]. Available: <https://www.smartgrid.gov/ADMS.html>
- [3] R. Melton, K. Schneider, T. McDermott, and S. Vadari, "GridAPPS-D Conceptual Design," V1.0. Pacific Northwest Nat. Lab., Washington, DC, USA, Tech. Rep. PNNL-26340, 2017. [Online]. Available: http://www.pnnl.gov/main/publications/external/technical_reports/PNNL-26340.pdf.
- [4] Pacific Northwest National Laboratory. *GridAPPS-D*. Accessed: May 2018. [Online]. Available: <http://gridappsd.readthedocs.io/en/latest>
- [5] *Application Integration at Electric Utilities—System Interfaces for Distribution Management—Part 11: Common Information Model (CIM)*, ed. 1.0, Standard IEC 61968, Jul. 2010.
- [6] *Energy Management System Application Program Interface (EMS-API)—Part 301: Common Information Model (CIM) Base*, ed. 3.0, Standard IEC 61970, Aug. 2011.
- [7] M. F. Hebler, R. Blace, and A. Perez-Lopez, *Semantic Web Programming*. Hoboken, NJ, USA: Wiley, 2011.
- [8] P. R. Paulson, T. D. Gibson, K. L. Schuchardt, and E. G. Stephan, "Provenance Store Evaluation," Pacific Northwest Nat. Lab., Richland, WA, USA, Tech. Rep. PNNL-17237, 2008. [Online]. Available: https://www.pnnl.gov/main/publications/external/technical_reports/PNNL-17237.pdf

- [9] J. H. E. Han and G. J. Le Du, "Survey on NoSQL database," in *Proc. 6th Int. Conf. Pervasive Comput. Appl. (ICPCA)*, 2011, pp. 363–366. [Online]. Available: <http://ieeexplore.ieee.org/document/6106531/>
- [10] G. McNaughton and R. Saint, "How NRECA's multispeak specification supports interoperability of diverse electric grid automation systems," in *Proc. IEEE Power Energy Soc. General Meeting-Convers. Delivery Elect. Energy 21st Century*, Jul. 2008, pp. 1–5. [Online]. Available: http://www.gridwiseac.org/pdfs/forum_papers/118_paper_final.pdf
- [11] Object Management Group. (2017). *Unified Modeling Language, Version 2.5.1*. [Online]. Available: <http://www.omg.org/spec/UML/About-UML/>
- [12] I. Robinson, J. Webber, and E. Eifrem, *Graph Databases*. Newton, MA, USA: O'Reilly Media, 2013.
- [13] A. de Vos, S. E. Widergren, and J. Zhu, "XML for CIM model exchange," in *Proc. PICA Innovative Comput. for Power-Electr. Energy Meets Marketing, 22nd IEEE Power Eng. Soc. Int. Conf. Power Ind. Comput. Appl.*, Sydney, NSW, Australia, May 2001, pp. 31–37. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&number=932315&isnumber=20177>
- [14] J. Huang, D. J. Abadi, and K. Ren, "Scalable SPARQL querying of large RDF graphs," *Proc. VLDB Endowment*, vol. 11, no. 4, pp. 1123–1134, 2011. [Online]. Available: <https://people.csail.mit.edu/jshun/6886-s18/papers/Huang11.pdf>
- [15] T. Elsethagen et al., "Data provenance hybridization supporting extreme-scale scientific workflow applications," in *Proc. New York Sci. Data Summit (NYSDS)*, Aug. 2016, pp. 1–10. [Online]. Available: <http://ieeexplore.ieee.org/document/7747819/>
- [16] T. E. McDermott, E. G. Stephan, and T. D. Gibson, "Alternative database designs for the distribution common information model," in *Proc. IEEE/PES T&D Conf.*, Denver, CO, USA, Apr. 2018.
- [17] R. C. Dugan and T. E. McDermott, "An open source platform for collaborating on smart grid research," in *Proc. IEEE PES General Meeting*, Jul. 2011, pp. 1–7. [Online]. Available: <http://ieeexplore.ieee.org/document/6039829/>
- [18] D. P. Chassin, J. C. Fuller, and N. Djilali, "GridLAB-D: An agent-based simulation framework for smart grids," *J. Appl. Math.*, vol. 2014, May 2014, Art. no. 492320. [Online]. Available: <https://www.hindawi.com/journals/jam/2014/492320/>
- [19] T. F. Sandquist and K. P. Schneider, "GridAPPS-D evaluation framework: A systems engineering approach," Pacific Northwest Nat. Lab., Richland, Washington, DC, USA, Tech. Rep. PNNL-26362, 2017. [Online]. Available: https://www.pnnl.gov/main/publications/external/technical_reports/PNNL-26362.pdf
- [20] Docker. (2018). Docker—Build, Ship and Run Any App, Anywhere. Docker Inc., San Francisco, CA, USA. Accessed: Apr. 3, 2018. [Online]. Available: <https://www.docker.com>



KEVIN P. SCHNEIDER (S'00–M'06–SM'08) received the B.S. degree in physics, and the M.S. and Ph.D. degrees in electrical engineering from the University of Washington. He is currently a Chief Engineer with the Pacific Northwest National Laboratory, Seattle Research Center, Seattle, WA, USA. He is also an Adjunct Faculty Member with Washington State University, an affiliate Associate Professor with the University of Washington, and a licensed Professional Engineer in Washington State. His main areas of research are distribution system analysis and power system operations. He is the Past Chair of the Distribution System Analysis Sub-Committee and the current Vice-Chair of the Analytic Methods for Power Systems Committee.



ERIC LIGHTNER received the B.Sc. degree in electrical engineering from the University of Maryland and the M.Sc. degree in electrical engineering from Johns Hopkins University. He was a Program Manager for advanced technology development with the Department of Energy for the last 27 years, and for the past 14 years has focused on the modernization of America's electric grid. As part of this work, he conceived of and championed the smart grid program and its associated activities with the Department of Energy. He is currently the Director of the Federal Smart Grid Task Force within the Office of Electricity Delivery and Energy Reliability whose mission is to ensure awareness, coordination, and integration of smart grid related activities both within the U.S. Department of Energy and across other federal agencies. In addition to directing the Federal Smart Grid Task Force, he is leading the development of an industry-led Voluntary Code of Conduct for data privacy related to customer energy usage information, managing research in advanced distribution management systems and low-cost sensors, and leading several efforts regarding Smart Grid Customer Engagement.



THOMAS E. MCDERMOTT (S'77–M'81–SM'92–F'14) received the B.S. and M.E. degrees in electric power from the Rensselaer Polytechnic Institute, and the Ph.D. degree in electrical engineering from Virginia Tech. He has a licensed PE in Pennsylvania. He is currently a Chief Engineer in integration, electricity, and infrastructure with the Pacific Northwest National Laboratory, with over 35 years of experience in consulting, software development, and research. His technical interests include simulation, transactive energy, and distribution system analysis.



POORVA SHARMA (M'16) received the B.E. degree in computer science from Washington State University, where she is currently pursuing the M.S. degree. She is currently a Research Engineer with the Data Sciences Group, Pacific Northwest National Laboratory (PNNL). She specializes in developing new software tools in the field of knowledge management, distributed systems, integration systems, and messaging middleware. Her work focuses on developing software platforms for data management and data integration in power grid and building energy research. She worked on projects under PNNL's Future Power Grid Initiative and developed the GridOPTICS Software System. She also worked on VOLTTRON—a platform for distributed sensing and control.



RONALD B. MELTON (S'75–M'80–SM'94) received the B.S.E.E. degree from the University of Washington, Seattle, WA, USA, in 1977, and the M.S. and Ph.D. degrees from the California Institute of Technology, Pasadena, CA, USA, in 1978 and 1981, respectively. From 1980 to 2002, he was with the Pacific Northwest National Laboratory, Richland, WA, USA, in various leadership and management positions and on a variety of scientific and technical problems. From 2002 to 2008, he was with private industry, primarily in system security engineering for nuclear material control and accounting and critical infrastructure systems. In 2008, he returned to the Pacific Northwest National Laboratory, where he is currently involved in grid modernization research and development. He has been actively involved in the IEEE PES Innovative Smart Grid Technologies Conference and is also the Secretary of the IEEE PES Committee on Connected Buildings, Loads, and Consumer Systems.



YINGCHEN ZHANG (M'07–SM'16) received the B.S. degree from Tianjin University, China, in 2003, and the Ph.D. degree from the Virginia Polytechnic Institute and State University in 2010. He is currently a Group Manager and a Principal Researcher with the National Renewable Energy Laboratory and a Visiting Research Assistant Professor with the University of Denver. He has over 10 years of experience in the power industry, in the areas of power system operation and planning, renewable integration, sensing, and monitoring. He has authored or co-authored over 50 peer reviewed publications and holds one U.S. patent. His key areas of expertise lies in predictive analytics for energy systems, advanced energy management systems for future grids, the impact of large-scale integration of renewable energy resources on power system operations, and sensing and analysis for a system with high renewable penetration.



FEI DING (M'14) received the Ph.D. degree in electrical engineering from Case Western Reserve University, USA, in 2014. She joined the National Renewable Energy Laboratory, Power System Engineering Center, as a Research Engineer in 2015. Her research interests include distributed renewable energy grid integration, distribution automation, distribution system operation and planning with grid edge intelligence, and advanced distribution management systems.



SUBRAMANIAN VADARI (S'84–M'88–SM'95) past roles have been the Vice President of infrastructure with the Battelle Memorial Institute, a Partner and a Global Lead for system operations and smart grid, Accenture, and a Lead Engineer with ESCA Corporation (now GE Digital Grid Solutions). He is currently the Founder and the President of Modern Grid Solutions, delivering consulting and training services to a global set of entities (cities, utilities, and vendors) seeking deep

subject matter expertise in developing the next-generation electric energy and other sustainable mechanisms. He is also an affiliate Professor with the University of Washington and a Technical Consultant with New York State Smart Grid Consortium, where he participates in an architecture role in the core REV team. He has published over 50 articles and several blogs on a broad range of subjects ranging from electric system operations to smart grid, big data, and Internet of Things. As a Smart Grid and Smart Cities Expert, he offers sought-after perspectives on the entire electric value chain.



ROBIN PODMORE (M'73–F'96–LF'13) was born in Palmerston North, New Zealand. He received the bachelor's and Ph.D. degrees in electrical engineering from the University of Canterbury, New Zealand, in 1968 and 1973, respectively. In 1973, he was a Post-Doctoral Fellow with the University of Saskatchewan, Canada. From 1974 to 1978, he managed the Power Systems Research Group, Systems Control, Palo Alto, CA, USA. From 1979 to 1990, he was the Director and the Vice President of business development with ESCA Corporation (now GE Grid), Bellevue, WA, USA. In 1990, he founded Incremental Systems Corporation, where he is currently the President. He has been an Industry Leader and the Champion for open energy management systems and affordable and usable Operator Training Simulators. From 2012 to 2016, he served as the Vice President for New Initiatives and Outreach for the IEEE Power and the Energy Society. In 2009, he cofounded the IEEE Smart Village Initiative and is currently the Vice Chair. He is a licensed Professional Engineer with the State of California and an NERC Certified System Operator. He is a member of the National Academy of Engineers.



ANAMIKA DUBEY (M'16) received the M.S.E. and Ph.D. degrees in electrical and computer engineering from The University of Texas at Austin, USA, in 2012 and 2015, respectively. She is currently an Assistant Professor with the School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA, USA. Her research interests include distribution system operation and planning for resilience and efficiency, grid integration of distributed energy resources, and demand response.



RICHARD W. WIES (S'92–M'99–SM'17) received the B.S., M.S., and Ph.D. degrees in electrical engineering from the University of Wyoming, Laramie, in 1992, 1995, and 1999, respectively. Since 1999, he has been with the University of Alaska, Fairbanks, where he is currently an Associate Professor of electrical and computer engineering. His current research interests are focused on the development of distributed generation and load control strategies for improving system efficiency, stability, and renewable energy penetration in remote microgrids. His other research interests include projects related to the efficient, reliable, and safe operation of electric power and energy systems. He is a member of the IEEE PES. He is a licensed Professional Engineer in the State of Alaska.



ERIC G. STEPHAN received the B.A. degree in mathematics and computer science from Eastern Washington University, Cheney, WA, USA, in 1986. He is currently a Data Scientist who is active in the international data standards community. He has been active in the data provenance community for over 12 years. His research interests include scalable science and engineering data architectures, applied semantic Web, and experimental reproducibility.

...