Towards a Decentralised Hierarchical Architecture for Smart Grids

Daniel Rech
Karlsruhe Institute of Technology (KIT)
daniel.rech@student.kit.edu

Andreas Harth
Institute AIFB
Karlsruhe Institute of Technology (KIT)
harth@kit.edu

ABSTRACT

We present a hierarchical distributed communication and control architecture for Smart Grids. The proposed topology consists of multiple layers to allow for robust and flexible data access and resource allocation in large decentralised Smart Grid systems. We introduce a scenario involving different Smart Grid actors, and develop an architecture using Linked Data principles (a set of standards for data access on the web). We propose a simple language to express allocation constraints, and map the resource allocation problem to a constraint satisfaction problem. We further provide initial experimental results for decentralised data access and resource allocation in Smart Grids.

1. INTRODUCTION

Smart Grids are distributed systems consisting of a very large number of actors. The actors vary in their use of the grid: some actors produce energy, some actors consume energy and some do both. The actors behave largely autonomously but have to communicate and work together to balance energy supply and demand. While there have been small deployments of Smart Grids with hundreds of participants (e.g., in German projects MeRegio[1], SmartWatts[2] and eTelligence[3]), the design of suitable architectures to support envisioned Smart Grid functionality on a large scale is a current research topic.

Previous works (e.g., [19, 9]) assume a centralised model in which utilities estimate the consumption based on historic data and send out price signals in regular intervals to entice actors to adjust their consumption. In contrast, we assume a decentralised hierarchical architecture for scalable and flexible communication and coordination between Smart Grid participants. We follow a general trend towards distributed architectures, emerging both in data management research (Dataspaces, [7]) and on the web (Linked Data, [2]),

[1] `http://meregio.de/`
[2] `http://smartwatts.de/`
[3] `http://etelligence.de/`

EnDM 2012, Berlin, Germany Copyright 2012 ACM 978-1-4503-1269-1/12/03...$10.00.

Consumers and producers of energy have different requirements and a management infrastructure for data and functionality has to take into account all of these. A distributed communication system requires the interconnection of originally separate data and functionality in a flexible way. Reading and writing data requires a clearly specified interface, with communication channels, protocols and conventions that each participant has to follow. While we do not specifically discuss the possibilities that semantic technologies offer for data integration, we believe that semantic technologies can help to address data integration issues that naturally arise in distributed environments with hetereogenous actors and system [20].

Our contributions are:

- We show that the hierarchical topology of the power distribution network can also be beneficially applied to the communication network, for both data access and resource allocation.
- We explain how concepts developed in the context of the World Wide Web (such as Representational State Transfer, REST) and the Semantic Web (such as Linked Data) provide flexible and scalable means to manage interconnected data and functionality.
- We map the problem of balancing energy production and consumption to a constrain satisfaction problem.

The remainder of the paper is organised as follows: Section 2 identifies challenges for managing data and functionality in Smart Grids. Section 3 introduces an idealised network topology for both power distribution and communication. Section 4 describes the communication and control interfaces for accessing data and services as well as methods for resource allocation. Section 5 presents initial experiments and explains how the proposed architecture satisfies the requirements. Section 6 relates our approach to existing work, and Section 7 concludes.

2. DATA MANAGEMENT CHALLENGES

In today's energy grid infrastructure, a small number of large power plants generate the bulk of the energy, and adjust power levels during hours of peak consumption. Such a "load-following supply" strategy leads to many inefficiencies, for example, related to high transport costs [10]. Further, regenerative energy sources such as solar and wind require a "load-following demand" strategy: actors have to consume the available energy at the time of generation. Thus, rather than controlling a few large energy producers, Smart Grids have to provide means for controlling a great many small actors, which may consume and produce energy.

The high-level requirements and the push towards decentralised energy generation suggest a decentralised communciation architecture. While data management is an important aspect in Smart Grids, we also need means to manage functionality. In other words, we not only need to be able to query Smart Grid data, but also need the abillity to affect change and control behaviour. Based on these insights, we derived technical requirements for a decentralised infrastructure, which we present in the following. We revisit the requirements as part of the evaluation in Section 5.3.

1. The current state of the network can be queried in an efficient way.
2. A subgrid is managed by an actor in the next higher hierarchy level.
3. Each node has to be able to manage up to several hundred other nodes.
4. Each subgrid can be disconnected and operated autonomously.
5. The network must be resilient and continue to work in case of node failures.

We build on top of well-developed technologies tested in computer systems of the internet. Our architecture mostly follows the design principles of Dataspaces [7], incorporating the basic ideas behind Linked Data [2]. The interfaces follow RESTful architectural patterns [5].

3. SMART GRID TOPOLOGY MODEL

The physical topology of the current power distribution grid is arranged according to voltage [13, 1]. Power plants are connected via a high-voltage network, which feeds into networks of subsequently lower voltage, ultimately leading to household networks with a voltage range between 100 and 240 V. As the current power distribution network is already organised into different subgrids, we propose that a smart communication and control infrastructure follows such a hierarchical partitioning. Our architectural model for the communication infrastructure thus mimics an idealised version of the physical power distribution network.

Definition 1. *A Smart Grid consists of nodes N and interconnects E; nodes represent actors (sensors, actuators) in the grid, and are connected to b (branching factor) other nodes. The levels of a network are denoted by L.*

We assume that several million actors can be part of the complete network, divided into subdivisions with tens to hundreds actors each. In other words, each household comprises tens of devices, a low-voltage node connects tens to hundreds of households to a higher-voltage network, and tens to hundreds of these are connected to form a network one voltage level up. We assume energy flow between those actors on the same level that are connected to the same parent actor. We assume communication flow only between actors in the direct next layer. The children of the root note represent the large power plants and storage facilities; the root node itself is a virtual actor.Figure 1 illustrates the structure of the network model.

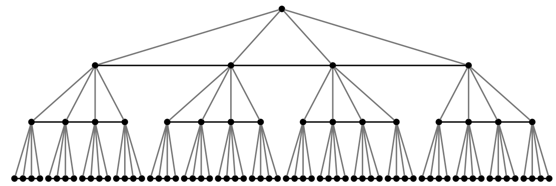

Figure 1: Self-similar network topology of the power grid. Branching factor $b = 4$ and $L = \{0, 1, 2, 3\}$ in this example. Black lines denote energy flow, grey lines denote communication flow.

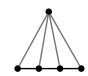

Figure 2: Atomic tree motif of which the entire network is composed of.

For our architecture we exploit the self-similarity of the topology (see work in complex networks [14, 17]), pick a motif (illustrated in Figure 2) and define operations on the motif. The architecture of the entire system is composed of such modular motifs or subgrids. Each subgrid manages its own data and resource allocation, and provides aggregated data to the tier above.

A hierarchical model has the practical benefit that it is possible to deploy computing nodes at already existing interconnection points in the power network. We also reduce unnecessary energy flow between different voltage levels by optimising energy consumption in the subgrids. Further, the hierarchical topology can make the entire system more resilient in face of outages, compared to a completely centralised model which introduces single nodes of failure which affect the entire grid. Finally, the proposed topology aids in partitioning the computational work necessary for communication and resource allocation.

4. COMMUNICATION AND CONTROL

In the following we first define the terms we use for Smart Grid actors. Next, we explain how the scenario maps to a constraint satisfaction problem (CSP) and finally describe the interactions between actors in different levels in the topology. We use a subgrid with branching factor $b = 2$ and levels $L = 0, 1$ to illustrate the definitions; for the example imagine a household (level 0) with an electric car and a solar panel (level 1).

Definition 2. *Let $\mathcal{A}$ be the range of acceptable energy consumption values for an actor and $a \in \mathcal{A}$ be the current energy consumption or energy production value of the actor in watt.*

We assume that each actor has assigned a value for a. To keep our model simple, we assume the possible consumption values to be integer values. We model energy production as negative consumption. We call an actor controllable if the actor has more than one value in its value range $\mathcal{A}$.

An actor might support energy slack as defined by Katz et al. [10]; as a result, an actor supporting energy slack is controllable. A person managing the actor can define a time interval for the task of the actor to complete, for example "have the car charged when I have to go to work in 5 hours". The actor can calculate for itself how much energy is needed to achieve that goal. There may be limitations; for example, a battery cannot be charged faster than at a specific rate or a washing machine cycle should not be interrupted. Based on the specified input, the actor derives an acceptable value range for energy consumption.

We pick the motif in Figure 2 to explain the different parts the system and communication scenarios between Smart Grid actors:

Definition 3. *We call the actor higher in the hierarchy controller, and the actors lower in the hierarchy are called participants. The current energy consumption of a controller is defined by the sum of energy consumptions of its participants. Each controller can also be a controllee for an actor in the next higher layer of the topology.*

We apply these definition to the example: the electric car and the solar panel are participants, and the household is the controller of the subgrid. Assume that the car has to be charged when by the beginning of the working day. A minimum consumption of 1 kW throughout is needed to achieve fully charged batteries. The car's batteries might be charged faster (for example, with 3 kW).The actor is currently attached to the grid and consumes 1 kW. To sum up: $\mathcal{A}_1 = \{1, 2, 3\}, a_1 = 1$.

The solar panel's position towards the sun can be adjusted to influence the produced energy production level. The possible energy production can be varied between 3 kW and 4 kW, currently producing 3 kW. $\mathcal{A}_2 = \{-4, -3\}, a_2 = 3$. The household controls all actors within its subgrid. The range of its energy consumption is defined by the sum of the energy consumption of its participants. In this case: $\mathcal{A}_3 = \{-3, -2, -1, 0\}$.

In the example, the goal is for the subgrid to operate autonomously. No further inflow or outflow of energy should be required and therefore the actor's consumption should equal zero: $a_3 = 0$.

Definition 4. *The controller's task is to set the consumption value of its participants in order to match a given energy consumption value a_c:* $\sum_{n=1}^{N} a_n \stackrel{!}{=} a_c$

The controller can make use of the value range $\mathcal{A}$ to adjust each participant to the energy available within the grid.

The current energy consumption value of each actor of a subgrid can be understood as one variable a_x of the problem. We express the resulting mathematical problem as a finite linear constraint satisfaction problem (CSP) [16]. With the formulation as a CSP, we can apply off-the-shelf CSP solvers to the optimisation problem.

Regarding the example with three actors, the controller can solve the CSP $a_1 + a_2 \stackrel{!}{=} a_3$ because it knows both the range of acceptable consumption values of each participant and

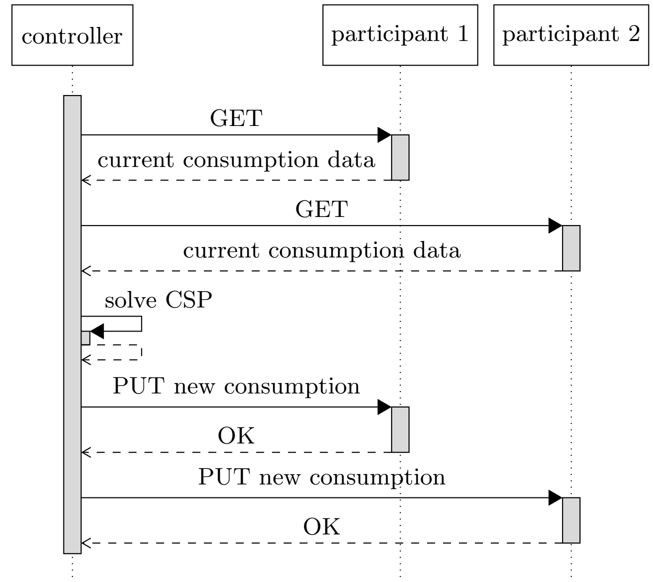

Figure 3: Communication flow of one controller with two participants.

the reference value. The solution to the CSP is $a_1 = 3$ and $a_2 = -3$. With these values the controller can adjust the energy consumption values of the participants. As a result, the subgrid can operate autonomously.

In summary, the communication flow between the controller and the participants follows a sequence of three steps:

1. The controller queries all the participants for their current consumption data.
2. The controller solves the CSP in order to have a configuration for the resource allocation to its participants.
3. The controller adjusts the consumption values of its participants accordingly.

Figure 3 visualises the communication between the the controller and the participants in our example.

5. EVALUATION

In the following we report on the actor interface implementation and initial experiments conducted on a machine with Intel Core 2 Duo 2.53GHz CPU and 4GB RAM. We further explain how the proposed system architecture addresses the requirements we have identified.

5.1 Actor Interface Implementation

We follow the proposed data model for Smart Grid actors from Wagner et al. [20], where an actor is addressable with an IP-address and exposes energy consumption data via an HTTP interface [6]. We however extend the interface with two features:

- The energy consumption of an actor can be adjusted.
- The actor returns a range of acceptable consumption values in addition to the current consumption.

Performing a GET request returns the current energy consumption and the range of acceptable energy consumption values. The actor will respond with RDF-formatted data for content representation. The consumption value of an actor can be modified by sending a PUT request. The controller sends the new energy consumption value in watt to the participant. If the value is within the participant's acceptable value range, the participant will adjust the consumption accordingly. Otherwise, the requests will be refused by returning a HTTP 403 Forbidden response.

Initial experiments show that querying many clients scales in a linear way as expected. We are able to access at least up to several hundreds of sources in parallel; the performance of these concurrent URI lookups is limited by the bandwidth of the network and the number of concurrent threads that the operation system can manage.

Still, the propagation down the tree can cause heavy delays on operations. A controller setting a new energy consumption value for a participant waits for the operation to be finished. The participant might need to solve its own CSP (in a controller role of another subgrid) and delegate setting operations to its participants. As the propagation down the tree can take a long time, we currently experiment with asynchronous approaches.

5.2 CSP Solution Strategy

The solution of the CSP gains complexity with an increasing range of choices an actor offers and with an increasing amount of actors. In the initial experiments we used XCSP 2.1 [15] to express the CSP problem. The finite linear CSP was translated into a SAT problem using the order encoding method by Sugar [18]. The SAT problem was solved by Minisat v2 [3, 4].

The transition from CSP to SAT introduces a limitation regarding possible variables as the translation cost regarding memory grows exponentially. We are able to solve the CSP for 1000 actors with a range of possible values of 5 in a few seconds.

Future work includes further optimisation of the CSP solver performance, and experiments with different topology choices to identify limiting factors for the size of a subgrid.

5.3 Addressing the Requirements

In the following we revisit the requirements identified earlier in the paper and explain how the presented architecture addresses the requirements.

1. The current status of the network can be queried in an efficient way.
   Our experiments show that the status of several thousand nodes can be queried in seconds. How the entire network behaves is subject to future experimentation.
2. A subgrid is managed by an actor in the next higher hierarchy level.
   If we assume that resource allocation takes place in intervals of several minutes, the lookup performance is sufficient to retrieve the current system state even over slow connections. Querying the information can become a problem if triggering one status update propages down the network tree.
3. Each node has to be able to manage up to several hundred other nodes.
   The most costly operation to perform is the solution of the CSP. As expected, the CSP does not scale linearily, which limits the amount of clients an actor with limited hardware can manage. However, first experiments show solvability for a few hundred nodes and a reasonable possible value range.
4. Each subgrid can be disconnected and operated autonomously.
   As each layer can control the managed devices, a connection loss does not need to cause the system to fail. The controlling actor can still operate in a stable way if it can keep consumption and production of energy on a balanced level.
5. The network must be resilient and continue to work in case of node failures.
   Our implementation of the controller can handle the failure of subgrid participants. There are two different types of failure: if an actor has a hardware failure and does not consume energy anymore, the controller can just remove the actor from the calculation. The second type of failure – if an actor fails to report its energy consumption – is more problematic. Given the controller still knows about the operation of the actor, the consumption value might be approximated by using historic data. The issue of securely shutting down actors has to be addressed in further detail. In both situations, the failure of a participant does not mean that the system loses controllability.

6. RELATED WORK

Aggarwal et al. [1] give an overview of the structure of the North American distribution network architecture and derive latency requirements for the future energy grid, concluding that a fiber-based infrastructure is necessary to meet the latency requirements. The conclusion is based on a centralised model for communication, whereas our research focusses on the comparison of different topologies. For future research we plan to include an analysis of latencies and bandwidth constraints for the decentralised approach to manage Smart Grid infrastructure and their influence on stability and controllability of the overall system.

Hauser et al. [8] and Lobo et al. [11] describe how the communication requirements may be met with an IP-based network if additional requirements for latency and security are applied. We add a proposal for a hierarchial topology model of the communication network structure.

Katz et al. [10] explain the challenges of the Smart Grid and introduce desirable features of Smart Grid actors, concluding with the description of a layered architecture and interface specifications to manage energy networks. In contrast, our paper picks up the actor features and proposes a structural model and interface specification. Our focus is on using established web technology and on providing an implementation in order to conduct experiments evaluating different control strategies.

Lu et al. [12] present possible threads to the energy grid and how the hierarchical structure allows to identify the root cause for errors, and provides an understanding about how the topology of the energy network correlates to a successful communication model. Future research includes an in-depth comparison of [12] and our approach regarding resilience of network topologies.

7. CONCLUSION

We presented a distributed architecture for the Smart Grid based on a hierarchical network model. Leveraging the self-similarity of the hierarchical topology, we identified motifs consisting of subtrees in which each node in the network offers standardised data access and control interfaces, leading to what we believe to be a simple yet elegant architecture. Furthermore, we have shown how resource allocation can be performed locally, thus providing a natural fit to energy networks where both energy consumption and production become increasingly decentralised.

We have conducted preliminary experiments showing that a basic Linked Data stack provides enough flexibility to fulfil the outlined requirements while keeping the protocol overhead manageable. Initial results show that a CSP solver can handle the allocation of resources in a simple model. Future work includes an evaluation of more expressive allocation problems.

While we have focussed on the basic architecture for communication and control, several functionalities are missing in our framework, most notably functionalities related to security aspects. For example, we have to assure correct authentification and authorisation before allowing actors to invoke actions which potentially can have adverse effects.

Acknowledgements

This work is supported by the German BMWi Federal Ministry of Economics and Technology (MeRegioMobil, Grant 01ME09005 and iZeus, Grant 01ME12013) and the European Commission's Seventh Framework Programme (PlanetData, Grant 257641). We thank Frank Pallas and the anonymous reviewers for helpful comments.

8. REFERENCES

[1] A. Aggarwal, S. Kunta, and P. K. Verma. A Proposed Communications Infrastructure for the Smart Grid. In *Proceedings of Innovative Smart Grid Technologies (ISGT)*, pages 1–5, 2010.

[2] T. Berners-Lee. Linked Data Design Issues. Technical report, W3C, 2006. http://www.w3.org/DesignIssues/LinkedData.html.

[3] N. Eén and N. Sörensson. An Extensible SAT-Solver. In *Proceedings of 6th International Conference on Theory and Applications of Satisfiability Testing*, pages 333–336. Springer, 2004.

[4] N. Eén and N. Sörensson. Minisat v2. 0 (beta). *Solver description, SAT race*, 2006. http://fmv.jku.at/sat-race-2006/descriptions/27-minisat2.pdf.

[5] R. Fielding. *Architectural Styles and the Design of Network-based Software Architectures.* PhD thesis, University of California, Irvine, 2000.

[6] R. Fielding, J. Gettys, and J. Mogul. RFC 2616: Hypertext Transfer Protocol – HTTP/1.1. http://www.ietf.org/rfc/rfc2616.txt, Jun 1999.

[7] M. Franklin, A. Halevy, and D. Maier. From Databases to Dataspaces: A New Abstraction for Information Management. *SIGMOD Record*, 34:27–33, Dec 2005.

[8] C. H. Hauser, D. E. Bakken, I. Dionysiou, K. H. Gjermundrod, V. Irava, J. Helkey, and A. Bose. Security, Trust, and QoS in Next-generation Control and Communication for Large Power Systems. *International Journal of Critical Infrastructures*, 4(1):3–16, 2008.

[9] C. Hirsch, L. Hillemacher, C. Block, A. Schuller, and D. Möst. Simulations in the Smart Grid Field Study MeRegio. *it - Information Technology*, 52(2):100–106, 2010.

[10] R. H. Katz, D. E. Culler, S. Sanders, S. Alspaugh, Y. Chen, S. Dawson-Haggerty, P. Dutta, M. He, X. Jiang, L. Keys, A. Krioukov, K. Lutz, J. Ortiz, P. Mohan, E. Reutzel, J. Taneja, J. Hsu, and S. Shankar. An Information-centric Energy Infrastructure: The Berkeley View. *Sustainable Computing: Informatics and Systems*, 1(1):7–22, 2011.

[11] F. Lobo, A. Cabello, A. Lopez, D. Mora, and R. Mora. Distribution Network as Communication System. In *SmartGrids for Distribution (CIRED Seminar)*, pages 1 –4, Jun 2008.

[12] N. Lu, P. Du, P. Paulson, F. Greitzer, X. Guo, and M. Hadley. A Multi-layer, Hierarchical Information Management System for the Smart Grid. In *Power and Energy Society General Meeting*, pages 1–8, 2011.

[13] A. M. Miri. *Ausgleichsvorgänge in Elektroenergiesystemen: mathematische Einführung, elektromagnetische und elektromechanische Vorgänge.* Springer, 2000.

[14] E. Ravasz and A.-L. Barabási. Hierarchical Organization in Complex Networks. *Phys. Rev. E*, 67:026112, Feb 2003.

[15] O. Roussel and C. Lecoutre. XML Representation of Constraint Networks: Format XCSP 2.1. *Arxiv preprint arXiv:0902.2362*, 2009.

[16] S. J. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach.* Prentice Hall, 2010.

[17] C. Song, S. Havlin, and H. A. Makse. Self-similarity of Complex Networks. *Nature*, 433(7024):392–395, 2005.

[18] N. Tamura, T. Tanjo, and M. Banbara. System Description of a SAT-based CSP Solver Sugar. *Proceedings of the Third International CSP Solver Competition*, pages 71–75, 2008.

[19] H. Vogt, H. Weiss, P. Spiess, and A. P. Karduck. Market-based Prosumer Participation in the Smart Grid. In *Proceedings of 4th IEEE International Conference on Digital Ecosystems and Technologies (DEST)*, pages 592 –597, Apr 2010.

[20] A. Wagner, S. Speiser, O. Raabe, and A. Harth. Linked Data for a Privacy-aware Smart Grid. In *Informatik 2010: Service Science - Neue Perspektiven für die Informatik, GI Jahrestagung (1)*, pages 449–454, 2010.