# Anticipative Virtual Storage Power Plants

Peter Palensky
Envidatec GmbH
Hamburg, Germany
Email: palensky@ieee.org

Dietmar Bruckner
Institute of Computer Technology
Vienna University of Technology
Email: bruckner@ict.tuwien.ac.at

***Abstract*—Virtual power plants aggregate smaller energy resources to a single "virtual" large one. This is done to appear as a larger player on the market, with one single interface for control, ancillary services and business requests. A virtual storage power plant (VSPP) (also) aggregates loads, coordinating them in a way that it imitates a real storage power plant. This paper discusses how far a VSPP can go into its underlying processes to optimize its performance.**

**Loads, although very simple on the first view, can behave in a pretty complex way. Estimating their current and future state is the goal of an anticipative VSPP. Having insight into the customer processes enables the VSPP to operate more optimized and to trade future shedding potential.**

## I. INTRODUCTION

The electricity system experiences new challenges. A 100+ years old, reliable and optimized infrastructure is now facing the pressure to change its topology, working principles and purpose. In the near future we will need a flexible, bidirectional energy exchange down to the end customer. A customer who by then will be rather called "participant".

Mainly driven by new, decentralized and lightweight ways of generating electricity, customers demand for grid access that allows them to feed back to the grid. Especially renewable energy forms need a good grid due to their stochastic nature. Unfortunately wind and solar radiation does not respect the standard load patterns and need to be supported by other energy forms and an energy grid that absorbs their unpredictable behavior. The expected rise in renewable energy sources makes this need even more acute (Figure 1 gives a relatively conservative outlook from [2]).

One main contribution for a more optimal and more reliable grid operation are "intelligent" loads. This means that loads have certain degrees of freedom (change their schedule, partially shed load for some time, duty cycle, etc.) and that they can interact and cooperate by some means of communication. Before optimization takes place, however, one must understand the nature of the problem: the loads and their properties.

Up to now, most energy/load management systems are configured by trained staff. A tedious and expensive task. It would be great if load management systems could learn the needs and properties of the customer process. Unfortunately these processes are sometimes very complex. This paper is a try to shed some light on this problem and proposes to use

- more data acquisition equipment (embedded nodes that are part of the load management system),
- a set of simple load models, and
- a network that connects these nodes to enable cooperation.

The loads that are supposed to be understood and influenced in a soft and smart way often do not reveal what there "inner driver" is. They appear – at least to the energy consultant – as a black box with an energy consumption as output and several known and unknown inputs. Often they contain sophisticated controls that make the situation even more complex.

## II. LATENT VARIABLES

Expecting that a load with a control module (programmable logic controller, scheduler, thermostat, etc.) behaves intelligently is naive. It is sad reality that most of the commercial loads and controls have no (active) "intelligence" at all. A typical air conditioning system of a supermarket or school is switched on in the morning and switched off in the evening (if at all). There might be a temperature sensor, but it is often overridden, not to speak of $CO_2$ sensors. The same applies to lighting which is typically controlled with a fixed schedule. The only difference one sees is (hopefully) caused by weekends or holidays.

On the other hand, the behavior of loads can be pretty complex. Think of a cooling system in an industrial process (Figure 2). A Wednesday in May might be a typical Wednesday in May: the shift starts at 6:00 a.m., the order situation is typical for May, etc. But if the outside temperature is unusually high, the time constants of the thermal inertia changes. The cooled goods heat up quicker and more energy is needed to

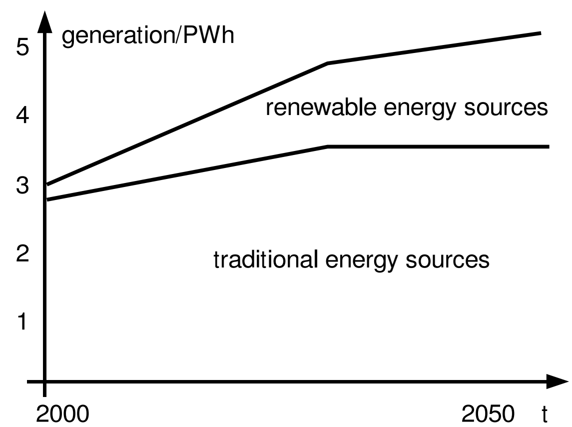


Fig. 1. Renewable energy forms are on the rise.

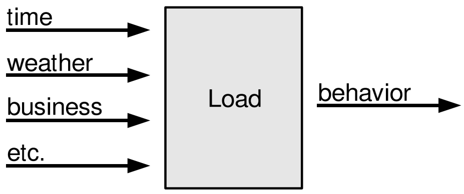


Fig. 2. Examples for inputs to a load process.

cool them. If several of such loads are aggregated on one customer site, it is a pretty complex task to derive the latent variables of the individual loads from the overall consumption chart.

Latent variables are the internal states and energy levels of a consumer (Figure 3). If the consumer follows a deterministic state machine, the state counter might be the latent variable. To outside, only the energy consumption is visible, but it follows the command of its internal variables and triggers that enter the load from outside. A thermal process (cooling, heating) typically has its internal temperature as a latent variable. Its value represents is "virtually" stored energy and its future behavior. The main energy counter at the customer site does not know any of these latent variables and sees only an overall consumption pattern. Often the latent variable is based on a state machine with time and customer behavior as input.

Knowing the latent variables would make energy management much easier. Load behavior could be predicted, schedules could be executed etc. There are three ways of accessing these latent variables:

- Analysis: someone analyses the individual loads and sets up a model for each of them. This is impractical in real life.
- Measurement: the chart measured at the energy counter is analyzed. Deriving individual qualities of such an aggregated set of data is not easy. There are, however, chances that if you add heuristic process knowledge, to get at least an idea of the loads. Google Power Meter [4], for instance, targets exactly this. It is assumed that private homes consists of a standard set of consumers (refrigerator, dryer, etc.) with characteristic load patterns. The aggregated chart is then analyzed to separate it into individual loads. This might work for homes that are alike, but is of no use in an industrial or commercial setting.
- Individual measurement: The way of choice for professional energy information systems and load management. The important loads are equipped with measurement and control nodes that monitor the loads' behavior. Out of their consumption patterns they can derive a simplified model (e.g. a time constant T for a first-order [a.k.a. PT1] behavior or a schedule for lighting).

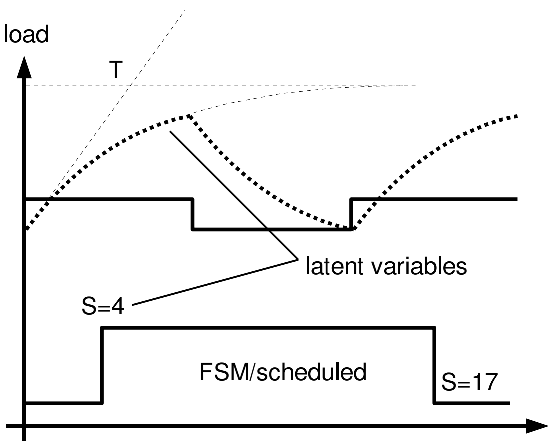


Fig. 3. Loads charts of loads with latent (internal) process variables and scheduled behavior. T is the dominant time constant for an internal (latent) thermal process (upper chart). Visible to outside is just a rectangle shaped load. A scheduled process (lower chart) does not have physical latent variables but follows a finite state machine (FSM) and has a state counter as latent variable.

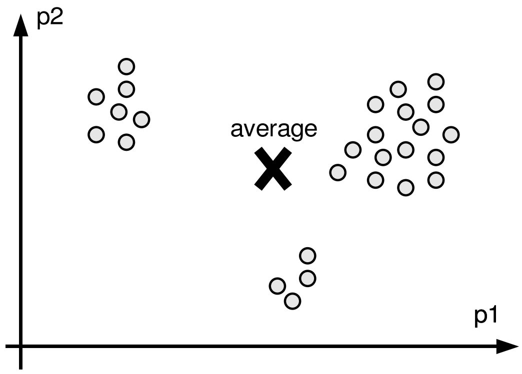


Fig. 4. Samples of system behavior (n process parameters reduced to two dimensions) and the average sample.

It is the latter method that is typically used for industrial and commercial customers. The necessary infrastructure surely costs money, but the gain refinances the investments. Even though (only a few important) individual loads are measured, it is still not easy to derive a model and their latent variables out of their consumption patterns.

## III. Categories of Loads

The authors work with two categories of loads in his simulations [3]:

- Thermal processes of first order
- Scheduled loads (with $< 10$ states)

A load of such a category can then behave in a variety of ways. Schedules might change, time constants might change. etc. It is the constellation of internal process parameters and their current latent variables that result in a specific behavior (i.e. energy consumption).

This leads to a classical problem in "blind" statistics: If you walk beside a cow, you two have an average leg count of three. Monitoring a device for some time with n sensors (a sensor directly corresponds to a know stimulus) leads to a set of samples, each sample being a vector of n elements. The samples form an n-dimensional space, each sample being a

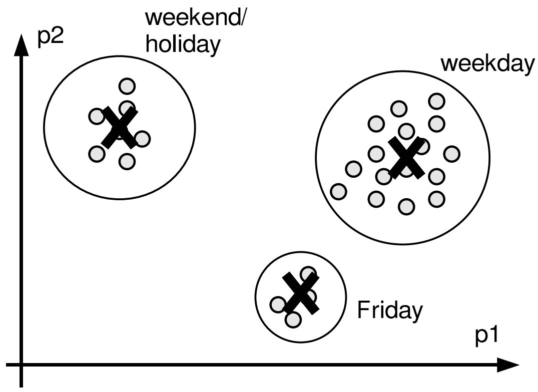


Fig. 5. A schedule might change. The load stays a *scheduled load*, but its timing etc. depends on stimuli like the current day of the week.

point in this space. Figure 4 shows a simplified example: the n dimensions (n process parameters) are reduced to two (p1 and p2).

The average sample lays in the (weighted) center of the samples. Depending on the way of distribution, this average might make no sense at all (think of the legs). Therefore, samples must be categorized, either with heuristic process knowledge or blindly, based on statistical methods. As industrial processes are too complex to be analyzed each time, it is the latter method that must be chosen. The n-dimensional search space is examined for "clouds" of samples, without knowing exactly what these clouds mean (see also [8]). Figure 5 has labeled "clouds", but often it is not possible to find out why a group of samples exists, so they are just named with numbers.

An energy management node has therefore the following tasks:

- Determine which category of load it manages
- Derive the latent variables of the load
- Share its findings (state of "virtual" storage, near-future behavior etc.) with other nodes
- Optimize the collective behavior

The second step works better the more sensors the node applies to the load. If the internal temperature, the state of doors and buttons, etc. is known to the node, it can easier derive the internal values than in the case when it only has the energy consumption pattern.

The latter two steps require communication and cooperation among the nodes.

## IV. Wide area energy management

Aggregating and decomposing energy resources is the principle of so-called "virtual power plants" (VPP, [5]). A number of power generators are grouped and collectively managed. The reason for this is that energy balance is still a process between large players. Entering this market has several barriers, one of them being the size of the power plant. If smaller entities want to participate, they can form such a VPP and appear as one large power plant to the market and to the control system.

A new category of VPP is not only based on distributed generation but also on loads: A "Virtual storage power plant" VSPP. A VSPP aggregates loads and their potential of shedding, shifting and changing their behavior. Unlike existing load aggregates (e.g. demand response via FM broadcasts or power line audio frequency ripple control signals), anticipative VSPPs would allow for planning, as the have – within certain limits – knowledge about the *internals* of their processes.

This is not only "softer" (the individual processes are shed only if their storage characteristics allow for that) but also potentially fair and accountable. Those members of the VSPP that contribute to a certain goal (follow some given load chart, keep frequency stable, etc.) are known to the system and potential rewards can be shared in a fair way.

We approach this new category of power plant with a number of steps.

First studies (like the PROFESY project [1], and the IRON project [7]) and the currently running INSEL project have shown encouraging results. INSEL uses the JEVis system [6] for scheduling large (mid voltage) customers in the city of Hamburg, Germany. The goal is to use historic data for selecting potential future schedules (hoping that the system "can do what it has already done"). A potential improvement of this would lead to an anticipative VSPP:

- Having models of the individual load nodes
- Using latent variables to generate a set of possible future behavior
- Selecting a "good enough" combination of individual future behaviors

The current task is to set up simulations for loads and their respective management nodes [3]. The goal is to derive good load models out of only a few sensors. The ideal case would be the "google method" [4] where the main energy counter is the only sensor input.

## V. Conclusion and further work

This research is still in its concept phase and in its infancy. There are several unresolved issues, mainly scalability and stability questions that must be analyzed and verified by means of simulation.

The system depends on wide area communication infrastructure, which unavoidably leads to the Internet. If network-based control is supposed to be done via this infrastructure, we have to assume certain qualities of service (QoS) that the Internet cannot satisfy. Of particular interest are

- Latency
- Availability
- Information Security

Workarounds for the non-guaranteed latency of the Internet can be based on using side channel information (grid frequency, FM broadcast, etc.) to synchronize the system.

Availability is often solvable with the contract, the VSPP will have. Also traditional power plants do not have an availability of 100%, but have planned and unplanned downtimes. The VSPP surely must satisfy certain maximum ratings for unavailability but a diversified and distributed topology can significantly help to improve that.

Information security is often ignored. The most important aspects of information security are authenticity of the exchanged messages and integrity of the messages. Injection of false and malicious information must be detected, since a “botnet” of mid-voltage power nodes might create some harm if orchestrated in the right way.

Another topic is the prediction and optimization algorithm itself. As there are no resources planned for tuning and configuring the system, machine learning will play a big role. The system is supposed to adopt to changes and to optimize itself without much manual interference. Plausibility and stability, however must always be guaranteed and verifiable by human operators. For this, again, a supporting simulation will be necessary. This is subject to further work.

## REFERENCES


[1] Rausch, T., Palensky, P.: PROFESY: Intelligent Global Energy Management, Proceedings of the 9th IEEE International Conference on Intelligent Engineering Systems INES05, 2005

[2] Union of the Electricity Industry - EURELECTRIC: The Role of Electricity, 2007

[3] Palensky, P., Kupzog, F., Zaidi, A. A., Zhou, K.: A simulation platform for distributed energy optimization algorithms, Proceedings of the 6th IEEE Conference on Industrial Informatics INDIN08, 2008

[4] Google.org Energy Information, http://www.google.org/powermeter/ accessed May 2009

[5] Schulz, C., Roder, G., Kurrat, M.: Virtual power plants with combined heat and power micro-units, 2005 International Conference on Future Power Systems, 2005

[6] Palensky, P.: The JEVis Service Platform - Distributed Energy Data Acquisition and Management, in Zurawski (ed.) *The Industrial Information Technology Handbook*, CRC Press, Boca Taton, Florida, 2005

[7] Stadler, M., Kupzog, F., Palensky, P.: Distributed Energy Resource Allocation and Dispatch: an Economic and Technological Perception, International Journal of Electronic Business Management, Volume 5, Number 3, page 182 196, ISSN 1728-2047, 2007

[8] Sallans, B., Bruckner, D, and Russ, G.: Statistical Model-Based Sensor Diagnostics for Automation Systems, Proceedings of 6th IFAC FET Conference on Fieldbusses and their Applications, p. 239 - 246, 2005