PROFESY: Intelligent Global Energy Management

Thomas Rausch
Institute of Computer Technology
Vienna University of Technology
Gusshausstr. 27-29/E384, 1040 Vienna
Austria
rausch@ict.tuwien.ac.at

Peter Palensky
Institute of Computer Technology
Vienna University of Technology
Gusshausstr. 27-29/E384, 1040 Vienna
Austria
palensky@ieee.org

Abstract – **Global energy management coordinates spatially distributed energy consumers. While local energy management systems use controllers like maximum demand monitors (MDMs), global systems align these localized decisions of MDMs in order to minimize "global" consumption peaks.**
PROFESY is a new idea how to combine and optimize such a network of MDMs. One of the new aspects of PROFESY is its soft interference. Energy management systems naturally take influence on customer processes: they switch off loads, if peaks are about to happen. PROFESY learns and predicts the behavior of individual nodes and softly exploits unused optimization potential. Ideally, the customer's processes are not aware that they are subject to global energy management at all.
The algorithms that were implemented use artificial neural networks and rely on a database which contains historic energy management records and environmental data. The complexity of creating optimized load predictions which again manipulate the power demand of a particular consumer is a challenging task, but the results are encouraging.

I INTRODUCTION

Energy management is probably as old as the deliberate usage of energy itself. Electrical energy, abstractly seen as a resource, is generated out of some primary form of energy (e.g. heat gained from burning fossil fuels), transmitted to the load and finally transformed into some tertiary form of energy (e.g. mechanical). The traditional way of this cycle is production by a small number of power stations, transport over the transmission and distribution grid and consumption in a large number of energy customers. The capacity of production and transport are, as with any resource, limited. These limits might be hard ones (thermal limit of the transmission lines) or soft ones (spot market price of primary energy). Either way, depending on the circumstances, the consumer has to pay a certain price for the consumed energy.
Industry and other large and medium customers often have more complex energy contracts than private homes. While a "small" customer just pays the consumed ("counted") energy, larger ones are also charged for the "way how" they consume: the load chart. Consumption peaks are sometimes extremely expensive, and that is why there are various systems for energy management in use.
Avoiding consumption peaks is a non-trivial tasks. The customer process has to have certain degrees of freedoms, in order to react if a consumption peak is about to happen. Naturally, the only strategy to react is to switch off (or reduce) "unimportant" parts of the energy consuming processes, a method commonly known as load shedding. A popular device for load shedding is the so-called "maximum demand monitor" (MDM). MDMs rely on the fact, that utilities charge their customers by using load charts with a certain measurement period. The average power consumption within these periods leads to the load charts. In Europe, this measurement period is 15 minutes, resulting in 96 measurement samples per day. The fine-grained load behavior within one period is not important, as long as the average load within this period does not exceed some defined value. A (very) simple version of an MDM would permanently monitor the "trend" of the average power (usually the linear extrapolation of the current power consumption to the end of the current measurement period). If it exceeds a certain maximum value, the MDM would switch off consumers with low priorities.
Such MDMs sometimes are networked with the consumers (energy consuming devices, appliances, machines, etc.) via local control networks, use more sophisticated algorithms than the one above, etc. but they are not globally linked. Each MDM controls its own, local site, there is no coordination and cooperation among multiple MDMs. This paper describes how to combine existing MDMs in order to achieve global energy management.
Large energy customers negotiate "net schedules" with their suppliers. If they follow the agreed consumption chart, they stay within their contract. Leaving the schedule would cause high costs. Multi-site customers might be charged based on a collective load chart. This means, the individual load charts of their subsidiaries are combined to one "global" chart. Unfortunately traditional local MDMs do not cooperate and contemporaneous local consumption peaks lead to expensive global peaks.
Global energy management tries to span over the entire number of sites (= "nodes") in order to coordinate the consumption: it acts as a global MDM (Figure 1). Transposing the algorithms and strategies of local MDMs into a global dimension does not always work:

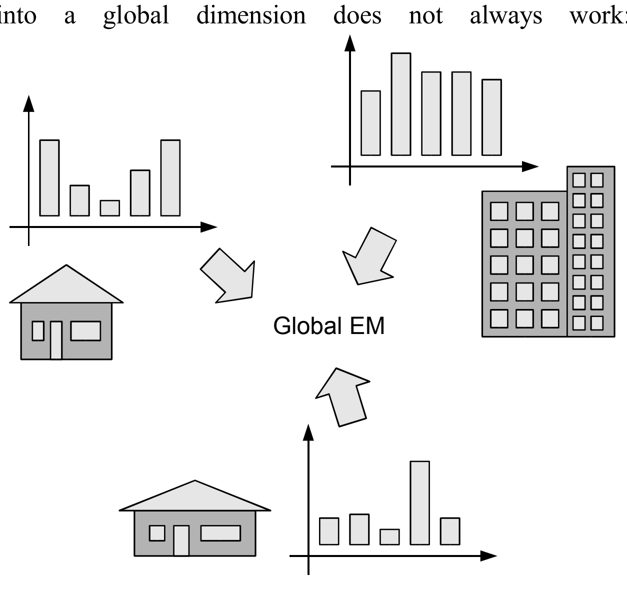

Figure 1
Multiple sites for the global energy managment system

communication between the nodes might be unreliable and/or expensive, the underlying mathematical problem and its algorithms scale badly and configuring such a system is an incredibly complex task.
The presented system PROFESY (PRediction of Optimized load proFiles for Energy management SYstems) takes a new and "gentle" way in global energy management. The key behind PROFESY is profound knowledge about the connected nodes and their properties. Luckily this knowledge is learned and does not need to be configured.
Modern MDMs record a log of their activities. This log at least consists of two parts

- the load charts (e.g. the consumed energy)
- the reserve load charts (the load that theoretically could have been "thrown off")

PROFESY uses MDMs that are connected to the Internet in order to transmit these logs into a database [1]. The collected data is the starting point of PROFESYS' knowledge.
Given a desired "net schedule" for a community of sites, it is the art of global energy management to find out how the individual nodes should behave in order to collectively satisfy the net schedule.

II ENERGY MANAGEMENT

A Local Energy Management

A local energy management system consists of a management station (like an MDM) and one or more controlled energy-consuming devices. The station communicates over a control network (e.g. a LonWorks network [2]) with its managed devices. It logs the consumed energy of all connected devices with a predefined sample rate which is – as already mentioned - 15 minutes in Europe. The resulting 96 values per day are first stored locally and can be periodically synchronized with a central database.
Besides simple logging, the management station also activates and deactivates particular devices, according to their physical circumstances, priorities, and predefined optimization criteria. Due to this control, the management system is able to save energy or at least reduce costs by avoiding peak loads. If the physical process allows the management station to schedule the controlled devices, it can change the time when energy is consumed. Hence, the station is able to modify the schedules and to lower possible peaks.
The degree of "intelligence" of these local management systems varies. Some of them simply log data and do not perform any control at all. Although the term management is not very appropriate for these devices, the provided information about the consumption characteristics is an important contribution to global energy management (Figure 2).
Other local management systems are able to take influence on the consumption of their local devices (heaters, machines, etc.), although they are not prepared to be networked.. These sites, as the previous ones, are "passive" nodes. They show some behavior, offer information, consumption logs, etc. but they do not accept any commands, they cannot be not controlled globally. The global management forecasts their expected load profiles and includes them to its global prediction.
More sophisticated devices can accept load profiles from an external source and attempt to satisfy these charts by managing their devices. From a global viewpoint, these local management systems are very valuable, since the allow a higher flexibility and a greater potential for optimizations. They are the real, the "active" nodes of the global energy management system.
Particular local management stations may even continuously estimate the amount of energy, which may be additionally saved. They for instance "know" that the site could decrease its consumption by 10% for one hour. This reserve is, as explained in the following section, based on virtual energy storages.

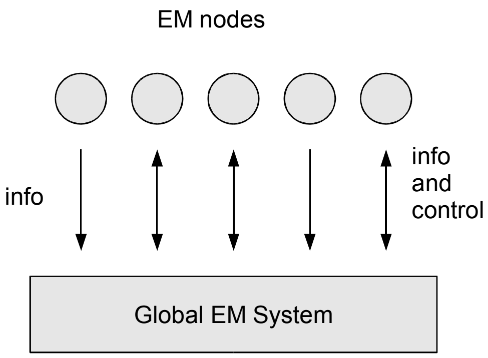

Figure 2
Multiple sites for the global energy managment system

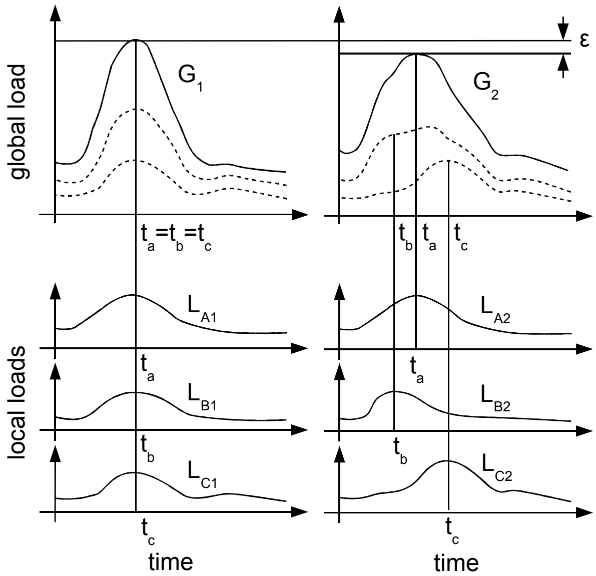

Figure 3
Multiple sites for the global energy managment system

B Global Energy Management

A global consumption peak usually consists of a number of synchronous sub-peaks that sum up to one large peak. Avoiding such a large peak can be done by shifting the underlying sub-peaks. Figure 3 shows 3 local loads L_A, L_B, L_C that result in one global peak.
The right part of the Figure shows – compared to the left one - a smaller global peak: its sub-peaks are not synchronized anymore.. Note that this does not reduce the amount of consumed energy, but only the size of the consumption peak.
Moving sub-peaks is only possible if the physical process offers the already mentioned degrees of freedom. These degrees of freedom are called "virtual energy storages": It

is possible to consume more now, in order to consume less later.
These virtual energy storages fall into two categories, namely
- inert physical processes (usually thermal or mass-flow), or
- logistics.

Filling up such a virtual energy storage means doing something in advance in order to not having to do it later.
PROFESY forecasts and optimizes load profiles in the area of HVAC (Heating, Ventilation and Air-Conditioning) applications and building automation in respect to a global energy management system. HVAC equipment is a perfect example (and in this paper seen as a synonym) for virtual energy storages, because the controlled physical process shows aspects of inertia (a heating system can for instance "store" energy in the temperature of the heated room, an air conditioning system can "store" energy in the percentage of carbon dioxide in the air, etc.).
The term "global" indicates the spatial aspect of the system. A number of buildings, geographically distributed, are maintained by a single operator that wants these buildings to be controlled collectively. According to the local conditions like outside temperature or humidity, the HVAC equipment on site shows a specific behavior. Every individual building that is part of the global optimization is equipped with a local energy management system.
As industrial energy contracts energy are often sold in particular quotas the local management systems try to make full use of the available energy in the given time frame. with a strict schedule. The system would ideally try to consume exactly the amount of energy as contracted. Consuming less than planned or allowed is fine, but exceeding this preallocated contingent is disproportional expensive.
If there is some energy left within such a temporary block, a local energy management system can decide to use this remaining energy to run devices ahead of their normal schedule (i.e. filling some virtual energy storage), in order not to exceed the limits of the next quota.
A local energy management system only evaluates information, which can be accessed on site. It may use environmental information from external data sources like weather services, it does however not coordinate its actions with management stations in other buildings. The optimization potential of local stations is limited.
A global energy management system on the other hand considers all buildings. Similar to a local management system, it tries to avoid the highest peaks. Unfortunately, HVAC equipment in buildings located in the same region or exposed to the same weather conditions usually results in load profiles with akin temporal characteristics, they are synchronzed.
For global energy management, the system forecasts the load profile of a future period for all buildings. By slightly shifting the time when the peak reaches its maximum value, i.e. forcing local management systems to alter the predefined schedules, the sum of all local load profiles can be lowered. An example for this modified operation is an air condition which runs at full power in the morning and is turned off during noon, a classical "peak time" when other devices are inevitably powered on. These modifications, must not decrease the quality of the services (i.e. heating, cooling, etc.) which should be provided by the corresponding devices.

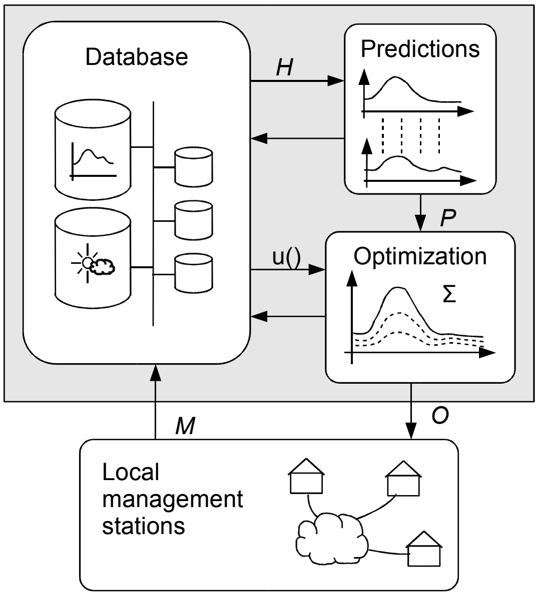

Figure 4
Major parts of the global management system

In addition to spatially bound management decisions, local management stations can also accept external scheduling information, which are received as default load profiles from a global management system. In this case, the local management station tries to control all attached devices in order to consume energy corresponding to this given profile. Due to manual interference, unexpected higher energy demand or unrealistic assumptions about the actual electric power consumption, the local management station may depart from the supplied default profile. These effects decrease the quality of the predictions.
The local management station decides which devices and consumers are to be switched on or off. The global energy management system estimates the optimal amount of consumed energy for the node (the entire building), expresses wishes and recommended profiles, but has no direct influence on the schedules of particular devices.
PROFESY's default load profile have the same resolution as the sample rate of the measurements (15 minutes), although it is not restricted to this sample rate. The applied algorithms are flexible so neither the sample rate nor the duration of the prediction period, i.e. the time for which the prediction is calculated, is restricted.
The prediction of the load profiles is primarily based on historic information like previously recorded load values or historic environmental information (e.g. temperature, thermal insulation). Other valuable information is preprocessed data like the differential change of the temperature to the previous day [3].
Because of the flexibility to alter the schedule of particular devices, multiple possible forecasts can be calculated for the very same basic conditions. Out of this pool of predictions the system selects those profiles that lead – in sum – to the best performance. Hence, the prediction module is supposed to produce a set of very diverse load profiles.
Figure 4 shows the important parts of a global energy management system and the interaction with the management stations on site. The system consists of three major parts: database, prediction module and optimization module.
Historic information H is supplied by the PROFESY

database and used as primary input for the prediction. Beside previously recorded values and historic weather data, the database also provides various other information about the prediction period like calendar information or weather predictions.

The prediction module takes these input values and uses correlations to create one or more forecasts P for the prediction period. In principle, the quality of the forecast is highly depending of the number of data points (i.e. data sources), which contribute to the calculation. It is implementation specific, which data is actually used for the prediction (e.g. if the prediction is created for a Sunday, probably only historic values of Sundays are taken into account).

Every prediction is associated to several probabilities. The “operability” value gives the feasibility that a particular load profile can actually be carried out by the local management system. This parameter can also be seen as “level of realism” towards an successful application of the load chart. An additional “value probability” specifies the likeliness that a single predicted value coincides with the to-be measured value at a particular point in time. For the latter probability, either a confidence interval or a probability function can be used.

The optimization module evaluates these probability values for the calculation of a global minimum. The module searches all possibilities to assign load predictions to buildings and selects that particular set of forecasts with global minimal costs. For large numbers of buildings with various alternative predictions per building the search for a optimal solution with minimal costs is somewhere between tedious and impossible. If the best set of predictions cannot be determined within a reasonable amount of time, alternative algorithms that produce suboptimal, but nonetheless good, solutions must be applied.

The costs of predicted load charts are calculated with a utility function. This is a flexible way to assign manifold costs to the summarized load profile.

Besides the costs, the optimization algorithm also considers probabilistic values for its decision-making. Whether a profile with very low costs, but a minor chance of a successful implementation by the local management station, due to either a low operability or bad value probabilities, is positively evaluated by the optimization modules, depends on configuration of the system. The operator of the global management system must align the ratio of risky, large savings to realistic minor savings. The system, as currently implemented, allows both fine-grained calibration and abstract high-level adjustments.

Finally, a set of optimized load profiles O is identified as (at least nearly) optimal and used to control the local management station on site.

After the operation, the local management stations store the measured real load profiles M into the database. This is primarily used for predicting load profiles in the future. Moreover, these values can be compared with the previously generated forecasts to re-adjust and improve the prediction and optimizing algorithms.

The arrows in Figure 4 which point from the prediction and the optimization module towards the database illustrates the possibility of these modules to store additional meta data for the algorithms or particular datasets into the database. Both modules perpetually evaluate their intermediary results and compare predictions with the actually measured values. In addition, the database may be used to “cache” temporal values, which eliminate tedious recalculations.

III ALGORITHMS

Load prediction is a wide area. Two proofed and popular techniques for creating forecasts are artificial neural networks and regression analysis [4], [5].

The regression model creates a functional correlation with several influencing factors and an error term. Designated regression coefficients are adjusted to approximate previous observation. In [6] various further enhancements for time series analysis are elaborated.

In artificial neural networks (ANNs) a complex mathematical function results from numerous simple functions which are interlinked among each other. ANNs are not explicitly programmed to calculate a particular complex network function. In an initial training phase, pairs of input-output values are applied to the network, which adjusts internal parameters (“weights”). Thus, the network learns even highly complex functions. [7] describes the use of ANNs for load prediction.

[8] compares regression analysis and ANNs and states that not any particular algebraic algorithm, but the accuracy during creation of the models, are relevant for the quality of the prediction. The idea of PROFESY is to create a flexible prediction system, where the quality of the forecasts is permanently increased by continuously adjusting the model parameters. The system should not rely on a model or any other heuristic and detilled information about the physical process itself.

Currently PROFESY uses ANNs for basic prediction. Therefore, an extensive analysis of the input data is essential: the system has to learn. Usually, historical time series are composed of periodic, polynomial and stochastic components [9]. A separate prediction of these parts according to their attributes will likely improve the quality of the forecasts.

PROFESY's predictions are calculated strictly on value bases. So the ANN is trained for the creation of a single sample value. An alternative approach, which may be implemented in future, would be the prediction of larger time frames during a single run of an ANN. This calculation is more complex, but it potentially implies better results, since particular patterns within the profiles may be identified. In the current version, several cycles of training and calculation of the network are necessary for predicting multiple values.

PROFESY is implemented in Java and uses an open source toolkit for the calculation of the ANNs [10]. This software library provides both a flexible application programming interface and a graphical user interface. PROFESY only made use of the programming interface.

The ANN consists of three layers. One input and one output layer with a single hidden layer for interconnection has been identified as best topology for the predictions. The hidden neurons are fully interconnected with their predecessor and with the subsequent layer. The input and the output layers are build from sigmoid neurons while the hidden layer exhibits a WTA (winner-takes-all) functionality.

Since currently only a single value is predicted by PROFESY, there is also only one output neuron present. The number of input nodes of the network on the other hand is dependent on the size of the historic time frame which provides input data for the prediction.

The prediction of future load samples in PROFESY is based on the concept of similar power consumption under similar environmental and operational conditions. This is reflected in the way the ANN is trained. Therefore PROFESY searches the historical databases for similar samples for a particular value. It starts with looking for samples which had been measured on the same week day and time of day as the future sample. The oldest of these samples is regarded as the least influencing parameter to the predicted value. In contrast to this, the latest sample is the most important factor.

Historic load values are treated differently by the implemented algorithms than other input parameters like environmental data. The training of the ANN is performed by input-output-pairs of the following structure.

The oldest recorded sample is used as input parameter for the first training pair of the ANN, while the second eldest is used as output. Some environmental samples can be used as additional inputs. The next training iteration is performed with the oldest and the second eldest samples as input (plus some additional environmental samples) and the antepenultimate sample as output. All subsequent pairs are created similarly. The last input output pair trains the ANN with the latest measured sample, and all preceding samples as input – not to mention additional environmental parameters. After applying all pairs recursively, the weights of the ANN re-adjust and therefore train the network.

The final step when creating a prediction is putting the ANN into calculation mode, i.e. fixing the weights, and to apply all historic samples, together with additional environmental parameters, on the inputs of the trained network. As a result of the calculation of the ANN the output layer provides a prediction corresponding to the values on the input layer of the network.

IV CHALLENGES AND OUTLOOK

A remaining problem in the current version of PROFESY is the mean value dilemma, illustrated in Figure 5. When a particular local system exhibits a special behavior by creating oppositional load characteristics the prediction system may create forecasts which are similar to ordinary average values. Take a machine which has a significant influence on the energy consumption. It is - due to a indeterminable production process - sometimes turned on in the morning, and sometimes in the afternoon: An averaging forecast would create bad results. The algorithms

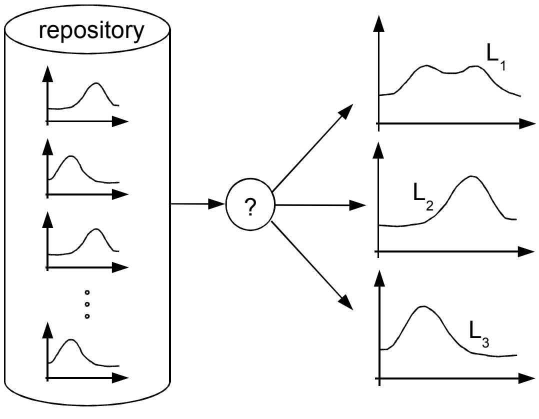

Figure 5
Mean value dilemma

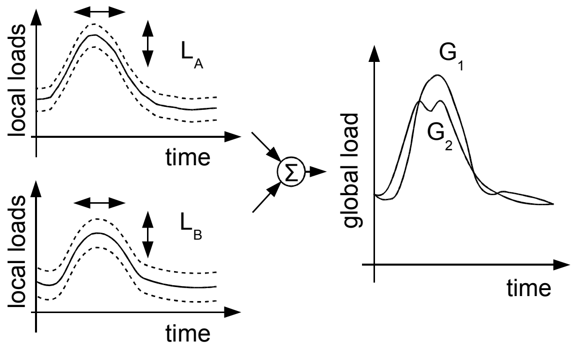

Figure 6
Fuzzy optimization

would have to realize this special behavior of the machine, and not to take a simple average. The historic values, which build the basis of the prediction, must be separated into different conflicting classes.

Currently, only operability values accompany fore casted load values. There are no further probability values available for single samples. Therefore the optimization is quite simple, in particular for a low number of local management systems. Therefore all possible configurations of predicted values are evaluated and the found solution is the global optimum.

In future, the use of probabilistic values will be intensified, in order to account for the uncertainty of the predictions. Due to these probabilities, which accompany not only a complete load profile but every predicted value, the optimization module has to implement more sophisticated algorithms.

Figure 6 shows fuzzy optimization of predictions with probabilistic load values. The thick solid line of the local load forecasts L_A and L_B symbolizes the values of the prediction with the highest plausibility. The dotted lines above and below this value give the highest acceptable variation from the forecast which can be tolerated by a sophisticated optimization algorithm. During optimization the algorithm can rise and lower the prediction for a particular value within these bounds. This flexibility helps in finding better global solutions like the summarized load profiles G_1 and G_2. G_1 is the “trivial” solution. It is the simple arithmetic sum of both local forecasts. G_2 however was created after the optimization algorithm has slightly altered the local predictions. This results in a an earlier peak of L_A, and in an delayed peak of L_B. The summarized load profile G_2 has therefore two peaks. By marginally shifting the expected value of the prediction within a narrow band of load profiles L_A and L_B, the optimization algorithm achieves a better result.

Every modification of the schedule of a particular device by the local management station, due to some global decision, alters the future load consumption of the device.

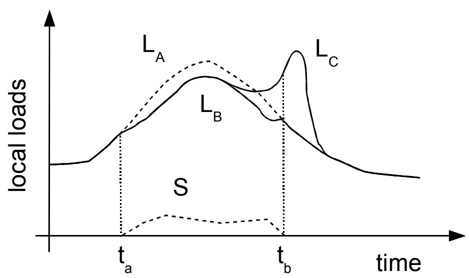

Figure 7
Delayed peak

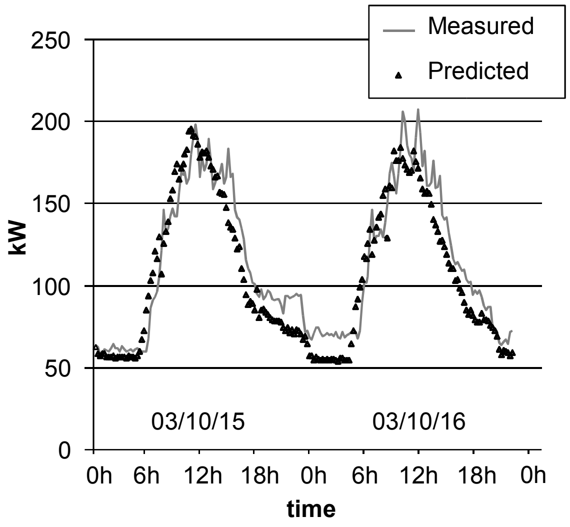

Figure 8
Measured and predicted values

A simple comparison of the predicted and the measured load profile is not possible in this case.

Another problem with modified load charts are unpredictable and delayed side-effects due to previous savings. If the local management station drives a physical process to its very limits, the controlling device might be unable to handle unexpected events, which demand more energy.

In particular, in an early stage of the operation of a global energy management system, there is hardly any historical information available about the behavior of the devices in case of external interference. In the worst case (see Figure 7), a prediction with a low peak (L_B) in respect to some potential savings (S) results in a delayed peak (L_C), which even exceeds the unaltered load profile (L_A).

V CONCLUSION

The basic algorithms described within this paper provide efficient techniques to create accurate load predictions. Figure 8 compares predicted and actually measured load values for an office building in Hamburg for a Wednesday in October (in kW) for these basic algorithms. The prediction takes, besides historical load profiles, also the temperature and the solar radiation during this day into account.

The system predicts a number of possible alternative load charts for each node (i.e. a building with a local energy management system) and chooses the one combination of these alternatives that is feasible and optimal.

This is the big advantage to alternative existing algorithms:, The individual nodes do not have to follow impossible load schedules. They are always asked to do something that they “already did” or that they were once able to fulfill under similar environmental conditions. This “soft” schedules have two advantages:

- they are more feasible than artificially optimitzed ones, and
- the customer process is not disturbed.

PROFESY intelligently combines non-intelligent MDMs, that are already in use. Artificial intelligence and statistics are used to do the trick of gently coordinating the previously uncoordinated sites.

Due to various extensions like probabilistic predictions and fuzzy optimization, existing algorithms must be enhanced. Further research has to be done not only on the algorithms but also on the necessary networks. Issues like scalability and reliability are very important if the system should not only manage tens of houses but maybe thousands.

VI REFERENCES

[1] Peter Palensky, “The JEVis Service Platform - Distributed Energy Data Acquisition and Management”, in: "Richard Zurawski (Ed.): The Industrial Information Technology Handbook", CRC Press, Boca Raton, Florida, 2005

[2] D. Loy, D. Dietrich and H. J. Schweinzer (Eds.), *Open Control Networks, LonWorks/EIA 709 Technology*, Kluwer Academic Publishers, 2001

[3] C. Hüsselmann, W. Kraemer and M. Klaus, “Gasverbrauchsprognose und Energiemanagement mit neuronalen Netzen”, *VDI-Berichte Nr. 1226*, Düsseldorf, VDI-Verlag, 1995, pp. 139-154

[4] L. Frosini and G. Petrecca, “Neural Networks for energy flows prediction in facility systems”, *SMCia/99, 1999 IEEE Midnight-Sun Workshop on Soft Computing Methods in Industrial Applications*, Finland 1999

[5] C. W. Fu and T. T. Nguyen, “Models for long-term Energy Forecasting”, *Power Engineering Society General Meeting*, IEEE , Volume: 1 , 2003, pp. 235-239

[6] G. E. P. Box and G. M. Jenkins, *Time Series Analysis: Forecasting and Control*, Holden-Day, San Francisco, 1976

[7] K. Hufendiek and M. Kaltschmitt, “Einsatz künstlicher neuronaler Netze bei der kurzfristigen Lastprognose”, *VGB-Konferenz „Forschung für die Kraftwerkstechnik 1998*, VGB-TB 233 A, VGB, Essen, Germany, 1998

[8] C. Meisenbach aand S. Gnüchtel, “Lastprognose konventionell oder mit neuronalen Netzen – ein Vergleich”, *VDI -Berichte Nr. 1508*, VDI-Verlag, 1999, pp. 207–218

[9] P. Bretschneider and T Rauschenbach, “Vorhersage musterbasierter stochastischer Signale unter Verwendung von Fuzzy-Technologien”, *Fuzzy control. Proceedings: 11. Workshop des GMA-FA 5.22*, Dortmund, 2001

[10] Joone – Java Object Oriented Neural Engine, www.jooneworld.com.