A System of Systems Approach to the Evolutionary Transformation of Power Management Systems

Jan-Philipp Steghöfer, Gerrit Anders, Florian Siefert, and Wolfgang Reif
Institute for Software & Systems Engineering, Augsburg University, Germany
{steghoefer, anders, siefert, reif}@informatik.uni-augsburg.de

Abstract: Power management systems of the future will be very different from the ones today. They will be complex systems of systems that directly incorporate distributed energy resources into the scheduling processes, react much quicker to changes in power demand and production than today's, and will allow small groups of producers and consumers to participate in the power market. This will lead to a more flexible, robust, and scalable power grid if the technological challenges can be mastered. We propose techniques that support this transition in five evolutionary steps, incorporating proven as well as new ideas. These steps are founded in the belief that agency of system components like power plants, trust as a measure of uncertainty, and self-organisation will be the keys to tackle the challenges at hand. Apart from describing approaches and algorithms that can help us in this matter, we provide preliminary evaluations that show that the techniques are worthwhile and should be further pursued.

1 Power Management as a System of Systems

Systems of systems (SoS) are complex systems composed of entities that can themselves be complex systems. SoS feature five typical characteristics [SC01]: 1) operational independence of individual systems; 2) managerial independence of individual systems; 3) geographic distribution; 4) emergent behaviour; 5) evolutionary development. The first four characteristics are fulfilled by all modern power management systems. Each power plant can be operated independently from the others. Although generators are coupled due to the shared power network and regulatory frameworks, economic operation of the individual plant is – to a certain extent – independent from others. Utilities and companies manage parts of the overall power system independently from each other. Geographical distribution is even increasing with the wide-spread installation of distributed energy resources (DERs) such as biogas plants, solar plants, and wind farms. The relative stability in the network is an emergent behaviour. No single entity of the system can provide this stability in the face of load fluctuations, weather changes, and generator outages.

This paper focuses on the fifth characteristic, evolutionary development. Of course, the state of current, centralised, manually managed power management systems is the result of such a process. We propose, however, to steer this evolution towards escalating autonomy and decentralisation. Such an evolution (see, e.g., [SA00] and [RVRJ12]) is necessary to deal with the growing number of generators, the increasing dependence on unreliable sources, and the increasing ability to control DERs. Extensive studies have been performed on the technological areas that need to be covered by such a transformation and the impact and role the individual technologies will have. We will repeatedly refer to the acatech Future Energy Study [AKM12] in this paper, as it is very recent and comprehensive. It also stresses the importance of increased autonomy of power systems, especially in the context of virtual power plants. The transformation into decentralised, autonomous systems will not be a sudden shift but a gradual evolution that requires both tremendous technological and regulatory changes. It will be complemented by other developments, especially w.r.t. the way power consumers are integrated into the system.

As computer scientists, our focus is on the software side of this process. The first contribution of this paper is an outline of five coarse-grained feature sets that can support an evolution towards a scalable, autonomous system of systems and can be implemented sequentially. These changes can be deployed on a large scale or within individual organisations to make internal management systems more robust, flexible, and cost-effective. The second contribution is an overview of the techniques that have already been developed to implement these feature sets and a discussion of their validity and usefulness. The techniques we propose should not be seen in isolation, but have to be complemented by other technologies and a shift in the legal and regulatory framework (see, e.g., the results of the E-Energy project[1] as well as of several European and international initiatives such as the FP6 project INTEGRAL[2], the European Electricity Grid Initiative[3] and its associated projects such as DISCERN, or the UK's Autonomic Power Systems project).

While current limitations such as missing measuring equipment make the deployment even of the first feature sets problematic, we believe they can contribute to the technologies called for, e.g., in the staged development outlined for virtual power plants in the acatech study [AKM12, Technology Area 12]. Current power management systems, our vision for their future, and related approaches are detailed in Section 2. We propose five feature sets:
1) Agency of power plants and semi-autonomous creation of schedules (Section 3): power plants and organisations work together to create power plant schedules.
2) Trust to manage uncertainty (Section 4): the system handles uncertain predictions and malicious participants using trust values and trust-based scenarios.
3) Hierarchical self-organisation into autonomous virtual power plants (Section 5): the system scales to large numbers of DERs due to a self-organised hierarchy.
4) Trust-aware power markets (Section 6): power plants participate in power markets in their own economic interest. Prices and trade restrictions are based on trust and reputation.
5) Delegation of control with electronic institutions (Section 7): organisations can control the behaviour of their subordinate power plants in a fine-grained way.

Each feature set provides direct benefits to organisations without having later features rolled out in the field as discussed in the respective sections. This paper is focused on power producers, mainly to avoid case distinctions when speaking about producers or consumers. Controllable or stochastic power consumers can easily be integrated into the proposed schemes to increase the flexibility of consumption and enable a mixture of load-led and generation-led operation, e.g., by participating in scheduling. Energy markets

---

[1] `http://www.e-energy.de`

[2] `http://www.integral-eu.com`

[3] `http://www.smartgrids.eu/European-Electricity-Grid-Initiative`

provide a natural interface between these power management applications. Parts of the feature sets have already been implemented and validated in simulations, not only by us but by researchers around the globe. Results of the simulations we performed will be shown in Section 8. The work conducted elsewhere is introduced in the appropriate sections.

2 Power Management Systems – Now and Then

To illustrate the proposed transition towards a self-managing power system, we first outline the current organisation of power management systems and contrast it with the vision of the future system at the end of the transition.

Current power management systems: The current organisation of power management systems is depicted in Figure 1. Big power plants are controlled by electric utilities and other organisations in a flat hierarchy. For each of the big power plants, a schedule is created that postulates the output of the power plant at a given time. Schedules are coarse-grained, providing target values in 15 minute intervals. Small power plants and especially DERs under the control of small cooperatives or individuals produce without external control and feed the power produced into the grid. This lack of control by the electric utilities is compensated by powerful controllable power plants. Current plans are to scale the controllable output further by installing more plants, especially flexible gas-powered ones. Very little measuring equipment is available in the field, so the actual grid status is unclear and predictions that guide energy production are based on intuition, weather forecasts, and historical analysis instead of current, reliable data. Nowadays, they are made by humans or SCADA (Supervisory Control and Data Acquisition) systems in the organisations. Additional power required to maintain the balance between production and consumption is provided by the operating reserve. The continental European grid, e.g., has 3 GW of idle power that can be activated within seconds. To be able to account for imbalances, the power grid is structured into balancing groups where utilities are responsible for equalisation of power consumption and production in each group [UCT09]. Further, entrance requirements to power markets, such as the lower limit of 100 kW for contracts at the European Energy Exchange (EEX), exclude access for small organisations.

Future power management systems: We envision a power grid in which even small power plants can be controlled or participate in a scheduling scheme. For this purpose, the system is structured into hierarchical systems of systems, as depicted in Figure 2. Networked measuring equipment allows to observe the grid status and make decisions based on current conditions. Power plants and consumers are networked, too, and predict future production or consumption. The quality of these predictions is known and can be used to derive possible scenarios for the future to guide scheduling decisions. Groups of power plants are represented by specialised SoS – *Autonomous Virtual Power Plants* (AVPPs) – that create schedules to cover a portion of the load and can flexibly access the power market. AVPPs can represent utilities, small organisations such as cooperatives of farmers that run biofuel plants, or simply a number of DERs that are controlled by the same organisation. Schedules are not only optimised for the economic benefit of large organisations but also take preferences of the individual participants into account. The

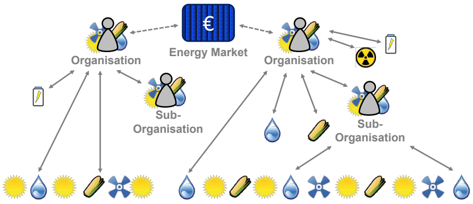

Figure 1: Current structure of power management systems: hierarchies within organisations are flat, some power plants are controlled by organisations directly, and only sufficiently large organisations participate in the power market.

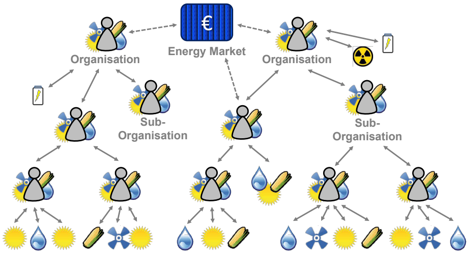

Figure 2: Hierarchical system structure of a future power management system: power plants are structured into SoS represented by AVPPs, decreasing the complexity of control and scheduling. AVPPs participate in the power market and can be part of other AVPPs.

behaviour of AVPPs can be guided by the controlling organisations. Finally, the operating reserve is provided by all controllable power plants in a decentralised scheme.

To make this vision – or a quite similar one – a reality, a number of fundamental changes have to occur. Many of them are legal and political, but some of them are technological. As computer scientists, we are mainly concerned with the latter kind, although an intensive dialogue between the different disciplines, politicians, companies, and consumers is of the essence. The solutions we propose in this paper and elsewhere are all technological and based on the notions of self-organisation and quantified uncertainty with trust values.

Several different approaches to re-structure power management systems have already been proposed. A key concept is the Virtual Power Plant (VPP), usually defined as groups of power plants that are in most cases controlled by a central entity (see, e.g., [LPR09]). Sometimes, membership in a VPP is restricted to certain types of power plants or power plants with predefined properties, such as dispersed generation units and micro-CHPs (combined heat and power units) in [SRK05] or DERs in [BVM$^+$07]. These approaches mainly focus on providing structures to integrate DERs into existing control schemes. Others concentrate on facilitating trading in power markets and distinguish commercial VPPs that participate in power markets and technical VPPs that provide services for the transmission net [PRS07]. None of these approaches combine all the features that we are looking for in the power management system of the future. However, they give important insights into the organisation and functionality of such a system.

3 Agency of Power Plants and Semi-autonomous Scheduling

The first step towards a more autonomous power management system is to enable power plants to proactively participate in the creation of schedules and in maintaining the stability of the grid. Each power plant can then optimise for individual criteria while maintaining global stability and quality. As many decisions require no human intervention, reactions can be faster. Schedules are based on better information due to distributed proposal creation and aggregation based on local weather forecasts, power plant status and load predictions. Together, this gives a strong incentive for small DER operators as well as larger utilities to participate in the scheme: they can better optimise internally and save money in the process while still meeting external goals. Agency of power plants, i.e., the ability to act autonomously in the environment, is the prerequisite for all other changes we propose.

This feature set requires that all power plants – especially DERs – and consumers have access to a communication network, making them accessible from the outside, and allowing the power plants to communicate their current status and predictions of their future output to other parties that create schedules. Merely networking generators to make them accessible from the outside is not enough, however, since the complexity of the system requires them to act autonomously and proactively in the system instead of just providing reactive services that are activated by a central control. Of course, opening power generators like this incurs security and privacy challenges (see, e.g., [Eri10] or [MM09]) and introduces requirements for communication languages and standards to enable semantic and syntactical understanding between the components [AKM12, Technology Area 17].

The superordinate organisations are SoS as defined by [SC01] and are predefined by the electric utility responsible for a certain area. Each of these SoS is responsible for satisfying a specific part of the overall residual load – i.e., the load that has to be satisfied by controllable power plants – by creating a schedule for its subsystems that meets this demand. The schedule creation problem can be solved by a centralised approach (see, e.g., [HLGR06] or [ZCRAC08]) within each organisation as before. With agency of the power plants, however, the schedule is now based on the current status of controllable power plants, predictions by weather-dependent power plants, aand predictions of the future load, enabling more precise and cost-effective schedules. Especially the second point is a major difference to current systems in which weather-dependent power production is predicted centrally or not at all. At this point, predictions are used directly in the scheduling process and no estimation of their accuracy is performed (see Section 4). Thus, a human operator still has to adapt predictions to current events, experiences made in the past, and so forth.

Schedules have to be adapted repeatedly as new data becomes available. Even though predictions for long time horizons are unreliable, they are required to identify if demand has to be covered on the market or surpluses can be sold. The scheduling of the residual load is based on models of the controllable power plants and their current status. The models are formulated with constraints, describing the physical properties of the plants, such as minimal and maximal output, rate of change, as well as shutdown and warmup times. Another factor is the production cost of power, given as a function of the output. Status information includes current output and whether plants are on or in hot or cold standby. A constraint optimiser calculates the optimal output for each power plant so that the combination of outputs is as close to the predicted demand as possible and as cheap as possible. Controllable power plants adopt this schedule and try to fulfil the assigned output as well as possible. To detect divergences from the target output, sealed and secure measuring equipment, similar to smart meters now deployed in households, has to be available at each power station. If these units report deviation from the schedule over a secure network, accounting systems can automatically reduce compensation.

While centralised scheduling allows for very good solutions, has been investigated by the research community, and is in active use in the industry for large controllable power plants, it suffers from a number of drawbacks. Foremost, scheduling is still centralised and requires models of the individual power plants. If power plants are not owned by the company that controls them, this causes conflicts regarding privacy. Small cooperatives or individuals have no interest in disclosing all information necessary to create adequate schedules. As an alternative, power plants could cooperatively formulate valid schedules without disclosing their internal models. Of course, such a solution means increased communication and sub-optimal solutions, but it captures the reality of distributed ownership much better than centralised approaches.

We thus propose a market-based solution [ASSR13b]. In principle, the agents representing power plants allocate a part of the predicted demand by selling or buying energy to or from the organisation that controls them. This is done in an iterative process. Before this process is started, the residual load is calculated. In all iterations, the organisation announces an auction in which a part of the overall residual load is sent to the organisation's controllable power plants in conjunction with a call for proposals. Power plants that want to sell or buy energy to or from the organisation respond with a proposal that includes a target output curve (i.e., a prediction of its future output) as well as the cost. The organisation gathers all proposals and completes the iteration by accepting the best proposal and modifying the remaining residual load accordingly by subtracting the accepted output curve. The other proposals are rejected. The organisation determines the best proposal in a multi-stage decision process based on heuristics and a combination of a filtering and two sorting stages in which only a subset of all proposals passes on to the next stage. In the filtering stage, invalid proposals, i.e., those exceeding the maximum contribution of the corresponding power plant, are rejected. Valid proposals are then sorted by the smoothness of the remaining residual load curve that would result in case a specific proposal was accepted. As it is assumed that a smoother remaining residual load is easier and cheaper to be satisfied, this step improves the algorithm's performance and the organisation's economic efficiency. In the second sorting stage, proposals are sorted by their price-performance ratio and the fraction of the residual load they satisfy within the regarded time frame. The best ranked proposal wins the auction. Having adjusted the remaining residual load according to the accepted proposal, the next iteration begins.

The algorithm terminates as soon as the predicted residual load is sufficiently satisfied. For each power plant, the combination of all its accepted proposals forms the schedule which is a contract with its organisation: the power plant has to comply with its prediction in exchange for remuneration. Because residual load predictions can change from one time step to another, schedules can be revised in following time steps. As only controllable power plants take part in the auction, other consumers or power plants implicitly enter into contracts by predicting their future consumption or output. This mechanism does not need any information about the participating power plants so that privacy is respected.

Using this algorithm, the agents make self-interested decisions (see Section 4). The proposals are conceived by solving a constraint satisfaction optimisation problem which is similar to the one solved by the centralised scheduling approach. However, since agents can change their behaviour and strategy, the optimisation problem has variable objectives.

The solution process can be performed by an off-the-shelf constraint solver. As shown in Section 8, the algorithm achieves results comparable with the centralised approach. In contrast to the approaches presented in [KVM10] and [WLHK06], our market-based scheduling approach does not only regard a single time step but generates long-term schedules.

Regardless of the scheduling mechanism used, all controllable power plants have the duty to participate in a decentralised, autonomous grid stabilisation scheme [AHS$^+$12]. Instead of the current reserve that is provided by large power plants, each producer contributes a small amount of its output to grid stabilisation, thus alleviating the need for costly reserves and distributing the cost on all grid participants. In addition, these reserves can not only be used to stabilise the grid in case load and production are not in balance globally, they can also be used to balance voltage band deviations locally. In Section 8, we show that the algorithms introduced in [AHS$^+$12] fare very well. Both the Honey Bee Algorithm and the Schooling Fish Algorithm are able to deal with fluctuations in the network frequency.

4 Trust to Manage Uncertainty

The second feature set gives agents the ability to measure inaccuracies in predictions as well as deviations from expected behaviour of other agents. This allows making decisions based on past experiences and forming a model of the environment's behaviour, enabling the creation of more robust schedules and dealing with unreliable agents.

Power management systems are inherently mission-critical. Their failure has massive consequences for large numbers of people, industries, and public service. It is thus of utmost importance that the power grid is stable and available at all times. Most of the engineering work in the power systems therefore aims for redundancy and reliability. However, with the transfer of control from human operators to distributed agents, there is a risk that misdeterminations or unforeseen emergent behaviour can bring the system into jeopardy. This risk is mainly associated with the quality of predictions and scheduling decisions based on them. Therefore, the objective of an autonomous power management system must be to increase predictability and minimise this risk. One of the reasons the power system is associated with so much risk is that it is an open system. Power producers and consumers – especially DERs and prosumers such as electric vehicles – are volatile, i.e., they can enter and leave the system at any point in time. Coincidentally, their benevolence can not be assumed. As agents participating in the power grid are not controlled by a single entity and can have different implementations and objectives, it is not clear whether or not they will participate in a way that is sensible for the overall system. Even if agents behave benevolently, the uncertainties imposed by the environment need to be captured. Errors caused by these "external" uncertainties are sometimes referred to as *systematic errors* (e.g., due to badly predicted weather) and *residual errors* (e.g., due to the conditions of a power plant) [CRK$^+$11]. The acatech study refers to this as "trustworthiness" of the information [AKM12, p. 86].

We propose to use *trust* as a central concept to measure this uncertainty and to improve the information's trustworthiness. Trust is based on experiences from the past. A trust model uses these experiences to create a trust value. A trust value can, among other things, give an estimation of the accuracy of predictions, i.e., their credibility, as well as of the reliability of a system participant. With respect to autonomous schedule creation, the following sources of uncertainty exist: accuracy of weather-dependent power plants' *predicted output* (credibility); accuracy of the *predicted demand* (credibility); accuracy of *providing power as scheduled* (credibility); and *availability* of controllable power plants (reliability). The main focus in scheduling is on credibility. A credibility value is derived by evaluating existing experiences capturing the difference between actual and predicted or scheduled power or demand. This value can then be used to calculate an expected deviation from a predicted or scheduled power or demand. Now, semi-autonomous scheduling as described in Section 3 can be further automated: instead of relying on human operators to provide expected deviations, the system can obtain those itself from trust values. By combining predictions and expected deviations, the system can schedule its operation based on the participants' expected behaviour[4].

In our market-based scheduling approach, credibility values are used to assess the risk that a contract is violated. First, this means that proposals of credible controllable power plants are preferred. Second, the maximum amount of energy an agent is allowed to produce per time step is limited with respect to its credibility. Third, depending on the credibility of contractual partners and the contracts, AVPPs schedule an expected amount of reserve power to compensate discrepancies between power production and consumption. Fourth, satisfaction of the residual load is calculated on the basis of the expected output which depends on the proposed output and the proposer's credibility. While the third measure allows to deal with uncertainties at runtime, the other measures incentivise agents to behave more credible. Privacy is still respected because the maximum output is the only information the organisation needs to have about its subordinate power plants.

Although the use of trust values significantly increases a system's ability to deal with uncertainties, a more adequate trust model is needed to precisely measure uncertainties in open systems. That is because simple trust values can only mirror, e.g., an agent's mean deviation from predictions. Since this mean behaviour likely differs from the agent's actual behaviour, a trust value is insufficient to describe the stochastic process underlying an agent's behaviour. There might be, e.g., a solar plant that either makes very accurate or inaccurate power predictions. Based on a corresponding trust value that represents the solar plant's average prediction quality, one would expect this power plant to make rather moderate power predictions, which actually is not the case. Moreover, there are situations in which it is beneficial that an AVPP can adapt itself to more than one possible development of its environment. For instance, if an AVPP knows a number of possible future developments of the power consumption, it will be able to schedule reserve power that allows to satisfy each scenario. Our idea to resolve these issues is to approximate an agent's underlying stochastic process by a number of so-called trust-based scenarios [ASSR13a]. Just like a trust value, they stem from experiences and are determined at runtime. Instead of a single value, however, there are multiple scenarios, each with a probability of occurrence.

---

[4] This refers to forecasting systems – Technology Area 10 in [AKM12]. While accurate forecasts are important, there will always be uncertainty involved which becomes quantfiable with the techniques proposed here.

5 Hierarchies of Autonomous Virtual Power Plants

For the third feature set, we propose to introduce a hierarchy of SoS that self-organises to adapt to a changing environment. Grouping power plants allows to decompose the scheduling task and thus satisfy the power demand in a robust fashion.

The autonomous creation of power plant schedules is an NP-hard problem. Not only is a centralised solution computationally expensive, it also requires a lot of communication to propagate the information required to articulate the models. In the system described so far, the organisations acquire the necessary information and create the schedules. The decentralised approach proposed earlier requires a lot of communication between the power plants controlled by an organisation. This leads to performance bottlenecks that become more severe as the system grows. Thus, to increase scalability of the system, we propose a hierarchical system structure as depicted in Figure 2. Such a structure will also increase efficiency in coping with uncertainties and untrustworthy agents as well as complying with schedules, market contracts etc. Parts of the SoS hierarchy will be predefined to represent the existing structure of utilities, grid operators, or network topology. However, within that organisational structure, a self-organisation process can guide the adaptive creation of hierarchy levels and group power plants in Autonomous Virtual Power Plants (AVPPs). Each AVPP constitutes a system of systems. AVPPs reflect the structure of the power grid and can be located at different voltage levels. Internally, they consist of a mix of generators and energy sources and their main objective is to locally balancing of demand and supply.

In a previous paper [ASSR12], we introduced SPADA that can be used to create partitions of the power plants controlled by one organisation. Each partition is represented by an intermediary – an AVPP – under the control of the organisation. Each AVPP in turn controls the power plants assigned to it and takes up the responsibilities of the organisation. That means that the AVPP collects data and creates the power plant schedules. This additional hierarchy level already increases scalability as organisations now control AVPPs which in turn control the power plants. As an AVPP controls multiple power plants, the organisations need to consider less entities than before. Possible alternatives to SPADA – such as coalition formation and clustering – can, e.g., be found in [HL04].

SPADA is a decentralised algorithm that solves the set partitioning problem (SPP). It partitions a set $\mathcal{A} = \{a_1, \ldots, a_n\}$ into $k \leq n$ pairwise disjoint subsets, i.e., partitions, that exhibit application-specific properties. For power management systems, SPADA uses criteria such as credibility and reliability of the agents to form suitable partitions. The goal is to achieve a good mixture of trustworthy and untrustworthy agents in each partition so that uncertainties can be dealt with locally and each partition is able to compensate for bad predictions on its own as well as a mixture of different energy sources and generator types to enable local balancing of demand and supply.

While a flat hierarchy increases scalability, it is not sufficient for the large number of power plants installed. In Bavaria, a largely rural German federal state with some large urban centres, there are more than 420,000 DERs installed [Ene13]. Such a number of power plants calls for a more hierarchical structure where AVPPs can in turn be controlled by AVPPs. We have therefore extended the SPADA algorithm with a hierarchical partitioning control

(HiSPADA) [SBA$^+$13]. It consists of a control loop that is able to introduce new hierarchy levels by introducing additional AVPPs as intermediaries. The main driver for the introduction of new AVPPs are scheduling times. Whenever an AVPP exceeds a time threshold for the sequential tasks of collecting the necessary data from power plants, calculating their schedules, and disseminating the schedules, it creates a new intermediary level for its child agents. For this purpose, HiSPADA initialises SPADA on the power plants assigned to the AVPP with an additional constraint that a minimum of two partitions has to be created. The partitions SPADA creates based on the mentioned mixture goals are represented by new AVPPs that are in turn controlled by the initiating agent. Likewise, if scheduling times become very short, an AVPP can dissolve itself, relinquishing control of its power plants to a superordinate AVPP, thus making scheduling more flexible for the superordinate AVPP as it has more power plants available. This way, the system self-organises towards a compromise between scheduling times, partition size, and flexibility.

Existing hierarchy levels can also be reorganised such that partitions conform to the originally defined partitioning criteria. Such action can become necessary in case credibility, reliability, or other attributes of the power plants change or the load curve profile the AVPPs have to fulfil changes. Predefined hierarchies are taken into consideration by HiSPADA, and it is possible to use HiSPADA to find suitable sub-hierarchies for predefined organisational entities. As each AVPP acts as a SoS, it is possible to use trust values to assess the quality of each AVPP. The evaluation in Section 8 shows that HiSPADA reduces the scheduling time significantly even for small system sizes.

6 Trust-Aware Power Markets

The introduction of trust-aware power markets in the fourth feature set allows agents to trade even small amounts of power in a secure system in which uncertainties are mediated by the trust mechanisms of the market place.

In a hierarchical SoS as described in Section 5, there is a need for exchanging energy within an organisation on different levels to balance power production and consumption, and thus to increase the system's robustness and efficiency. We propose that AVPPs controlled by an organisation govern an AVPP-internal power market where subordinate power plants and AVPPs can trade power intra-organisationally. An AVPP that needs additional power to satisfy future consumption could thus buy it within its organisation. As the organisation controls this internal market directly, smaller power plants can be granted access to the market, resulting in an increased number of participants and contracts.

When trading energy on power markets, agents have to cope with uncertainties similar to those which arise in the context of scheduling (see Section 3 and Section 4). These uncertainties are mainly due to information asymmetry [BP02], meaning that the parties have only incomplete knowledge about each other and different information about the product. In current over-the-counter markets, human operators therefore ultimately decide on the trustworthiness of trading partners when concluding contracts. Furthermore, a complex legal system ensures that the interests of all traders are safeguarded and incentivises traders to comply with their contracts. This legal system is especially important in situations in which, due to the trading procedure, market participants do not have information about or influence on which agents could or will become their contractual partners. This is the case, e.g., in uniform price auctions that are, e.g., used in the EEX's day-ahead market.

An AVPP-internal market is therefore equipped with a social system based on trust and reputation, complementing the existing legal system. Since trust and reputation stem from established contracts and their fulfilment, such a system provides mechanisms to identify untrustworthy and uncooperative agents and to effectively lower their utility by sanctioning misbehaviour and limiting their access to the market. If market participants are informed of how misbehaviour is sanctioned and the sanctions equate to a reduction in benefit, it is in the agents' best interest to behave benevolently and adhere to the rules of the internal market place. For market participants to be able to take advantage of trust and reputation values, the internal market implements a first-price sealed-bid auction in which bidders have information about the identity of the agent that requested the market to start the auction. While the literature presents various approaches for preventing strategic misbehaviour and gambling through pricing mechanisms (e.g., [CRK$^+$11, VRV$^+$10]), agents also have to be able to identify and cope with basically benevolent agents that unintentionally show non-beneficial behaviour (see Section 4). Trust-based techniques allow participants to deal with both kinds of non-beneficial behaviour. These have to be integrated into the market's social system, the trading procedures, as well as the market participants' decisions.

With respect to the internal market, misbehaviour is, in principle, sanctioned directly in the short- as well as indirectly in the long-term. Direct sanctions are imposed by the organisation, for example, in the form of punitive fines. Indirect sanctions are imposed by the organisation by decreasing reputation values on the one hand, and by contractual partners by decreasing trust values on the other hand. Combined with reputation- and trust-based decisions, these measures decrease a misbehaving agent's utility. Because the risk associated with a transaction is related to the amount of traded goods, the organisation limits access to specific market products or restricts the tradeable volume within a specific time frame dependent on the agents' reputation value. Further, an organisation specifies a minimum reputation value as a prerequisite to take part in the internal market.

In a trust-aware market setting, both contractors are blamed if a contract is not fulfilled. All traders therefore try to lower their risk associated with contracts by preferably concluding contracts with trustworthy agents. Agents that ask the internal market to start a new auction specify a lower bound for a trading partner's reputation value. Bidders also specify a minimum reputation value and decide on the price they are willing to pay or demand depending on an agent's trustworthiness and reputation, resulting in price premiums and discounts for trustworthy sellers and buyers. The internal market's trading algorithm incorporates this information so that contracts are concluded in the interests of all parties. Apart from trust values, it is also possible that agents base their decisions on trust-based scenarios (see Section 4 and [ASSR13a]).

A detailed investigation of AVPP-internal market places as well as the use of trust in an agent-based implementation of the EEX that allows autonomous inter-organisational trading will be subject of a future paper. In the next section, we show that a normative framework can be used to establish AVPP-internal power markets on different levels.

7 Delegation of Control with Electronic Institutions

As the final feature set and a subject of future work, we propose to establish electronic institutions that guide the behaviour of agents in an adaptive way to accommodate changing business goals and impose organisational rules within the hierarchical structure.

Each organisation in the system can establish an electronic institution that embodies business rules and general behavioural guidelines. These guidelines or behavioural *norms* [BPvdT09] influence the decision making process of the individual agents by providing guidance, e.g., on how an agent should behave on the power market, which optimisation goals to use for the scheduling process, or how AVPPs are formed. The norms are observed, changed and sanctioned by a *normative framework.*

Norms are also a way for the organisation to delegate control to individual agents [ASP09]. Organisations usually reserve the right to perform certain actions, e.g., to participate on the power market. Delegating norms can be used to give subordinate entities the possibility to do the same. An AVPP representing a large utility could, e.g., grant access to the power market to an AVPP representing a smaller subsidiary or even an AVPP created by hierarchical self-organisation. This again bolsters scalability and agent autonomy.

8 Evaluation

In this section, we demonstrate results so far obtained with a subset of the presented feature sets and their corresponding algorithms. More precisely, we point out evaluation results concerning autonomous grid stabilisation (see Section 3), power plant scheduling (see Section 3 and Section 4), and the formation of hierarchical AVPPs (see Section 5).

Autonomous decentralised grid stabilisation: In case of an imbalance between energy production and consumption, each controllable power plant contributes to compensate for the power deviation. The decision to what extent a power plant should adjust its output in order to dissolve the mismatch can be made by means of two decentralised algorithms we presented in [AHS$^+$12]. In the paper, we show that controllable power plants that use the Schooling Fish Algorithm (see Section 3) to adapt to changes in the network frequency closely follow a given realistic load curve without any other form of scheduling. The Honey Bee Algorithm, which uses no coordination between power plants, also showed good results but is subject to oscillations in some cases.

Centralised and market-based scheduling: The market-based scheduling algorithm is able to create schedules that allow the system to satisfy the demand even in situations in which untrustworthy agents try to harm the system's stability [ASSR13b]. This is achieved because scheduling decisions are based on trust values. Figure 3 contains data for a centralised scheduling approach and two different variants of market-based scheduling. In the centralised approach, a deviation from the schedule by can not readily be compensated since the culprit is still scheduled as if it was compliant. If trust values are not used, the market-based approach achieves a result similar to the centralised one. However, if trust values are used, the malicious power plant is identified and the capacity provided by other,

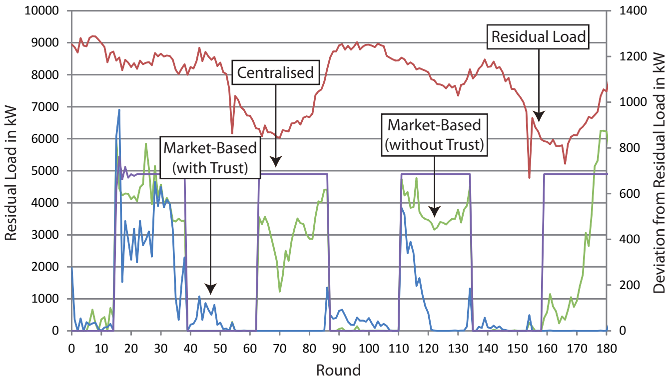

Figure 3: Performance of the market-based scheduling mechanism compared to a centralised approach. A single controllable power plant periodically does not comply with its schedule. If the market-based approach makes trust-aware decisions, it significantly reduces violations caused by the malicious power plant. Results averaged over 100 simulation runs.

more credible power plants is used to restore the original quality of the schedule.

Formation of hierarchical systems of systems: The hierarchical structure of AVPPs is generated by a combination of a decentralised set partitioning algorithm (SPADA) and a hierarchical partitioning control (HiSPADA). When a hierarchy of AVPPs is established, the structure has to follow the system's objectives and thus ensure efficient and stable operation. We showed in [ASSR12] that SPADA's local decisions lead to a partitioning whose quality is within 10% of the solutions found by a particle swarm optimiser. While the structure changes significantly during initial partitioning, subsequent reorganisations due to changes in the environment or in the power plant landscape only require minor changes and are performed locally, i.e., do only involve a small subset of partitions.

HiSPADA's main benefit is a decrease in the maximum sequential scheduling time required, as shown in Table 1 [SBA$^+$13]. The sequential scheduling time is the sum of the runtime of schedule creation over all AVPPs on a path from the root AVPP to a leaf in the hierarchical structure (see Figure 2). Intuitively, after the root node has calculated a schedule for its direct children, these in turn schedule their children. Scheduling is recursively performed on lower levels until the schedules for the physical power plants are created. As each branch performs the scheduling concurrently, no overhead is introduced. Table 1 contains the maximum average scheduling durations on all paths from the root AVPP to the leafs of the hierarchy. It also shows that the hierarchy is relatively shallow for the scenarios used here in which 435 power plants were regarded. Sc. A is the baseline without hierarchies, while Sc. B is a flat hierarchy. Hierarchies are introduced in Sc. C and Sc. D, where Sc. D contains predefined organisations that can not be altered by the algorithm.

In summary, an autonomous power management system implementing hierarchical self-organisation can deal with a vast number of agents by establishing a system structure that supports its objectives and using decentralised control schemes that employ trust values to measure and mitigate uncertainty.

	Scenario A	Scenario B	Scenario C	Scenario D
Max. seq. runtime of scheduling in ms	3624 ± 55	925 ± 638	499 ± 220	484 ± 213
Average depth	–	2	4.99	3.52

Table 1: Evaluation results for scalability and hierarchy depth, averaged over 100 simulation runs.

9 Conclusion and Future Work

In this paper, we propose a transformation of current, centralised, and manually managed power management systems to decentralised, autonomous systems of systems on the basis of coarse-grained feature sets to make these systems more scalable, robust, and flexible. At the heart of the feature sets we propose are known and proven techniques from the fields of multi-agent systems, self-organisation, and trust. They include market-based approaches, trust-based methods to quantify uncertainty, hierarchical self-organisation, as well as normative systems. Preliminary results show that the individual algorithms and mechanisms based on local decisions and techniques for compartmentalisation are promising and that their integration has the potential of yielding the expected benefits.

We are well aware that such a drastic change in the way the power grid is managed will not happen overnight. Indeed, many utilities do not yet have a business case to even start development in this direction. Recent regulatory changes and the increased deployment of networked measurement equipment can, however, increase the pressure on organisations and make the ideas outlined here more achievable and investments more likely. The economic repercussions of the move to smart grids, especially the impact on consumer prices and utilties' bottomlines are still an area of active research (see, e.g., [DCD11]). These issues will become an even stronger focus as the technologies mature, the business cases become more palpable and political pressure increases.

Our future work will be focused on the implementation of three of the concepts put forward in this paper: market-based scheduling, trust-aware power markets, and electronic institutions. In addition, our efforts to achieve a close integration of the concepts and algorithms described here and elsewhere will continue. There are a number of essential additional topics, such as data security and privacy, to which there is always a legal and political as well as a technological side. In the end, a transformation of current power management systems can only be successful when all stakeholders pull their weight.

Acknowledgements This research is partly sponsored by the German Research Foundation (DFG) in the project "OC-Trust" (FOR 1085).

References

[AHS$^+$12] G. Anders, C. Hinrichs, F. Siefert, P. Behrmann, W. Reif, and M. Sonnenschein. On the Influence of Inter-Agent Variation on Multi-Agent Algorithms Solving a Dynamic Task Allocation Problem under Uncertainty. In *2012 Sixth IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*, pages 29–38. IEEE Computer Society, Washington, D.C., 2012.

[AKM12] H.J. Appelrath, H. Kagermann, and C. Mayer, editors. *Future Energy Grid—Migration to the Internet of Energy*. acatech STUDY. acatech – National Academy of Science and Engineering, Munich, 2012.

[ASP09] A. Artikis, M. Sergot, and J. Pitt. Specifying norm-governed computational societies. *ACM Transactions on Computational Logic (TOCL)*, 10(1):1, 2009.

[ASSR12] G. Anders, F. Siefert, J.-P. Steghöfer, and W. Reif. A Decentralized Multi-agent Algorithm for the Set Partitioning Problem. In *PRIMA 2012: Principles and Practice of Multi-Agent Systems*, volume 7455 of *Lecture Notes in Computer Science*, pages 107–121. Springer Berlin / Heidelberg, 2012.

[ASSR13a] G. Anders, F. Siefert, J.-P. Steghöfer, and Wolfgang Reif. Trust-Based Scenarios – Predicting Future Agent Behavior in Open Self-Organizing Systems. In *Proc. of the 7th Int. Workshop on Self-Organizing Systems (IWSOS 2013)*, May 2013.

[ASSR13b] G. Anders, J.-P. Steghöfer, F. Siefert, and W. Reif. A Trust- and Cooperation-Based Solution of a Dynamic Resource Allocation Problem. In *2013 Seventh IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*. IEEE Computer Society, Washington, D.C., 2013.

[BP02] S. Ba and P. Pavlou. Evidence of the effect of trust building technology in electronic markets: price premiums and buyer behavior. *MIS Quarterly*, 26(3):243–268, 2002.

[BPvdT09] G. Boella, G. Pigozzi, and L. van der Torre. Normative Systems in Computer Science – Ten Guidelines for Normative Multiagent Systems. In *Normative Multi-Agent Systems*, number 09121 in Dagstuhl Seminar Proceedings, Dagstuhl, Germany, 2009. Leibniz-Zentrum für Informatik.

[BVM$^+$07] I. Bel, A. Valenti, J. Maire, J. M. Corera, and P. Lang. Innovative Operation with Aggregated Distributed Generation. In *Proc. of the 19th International Conference on Electricity Distribution*, 2007.

[CRK$^+$11] G. Chalkiadakis, V. Robu, R. Kota, A. Rogers, and N. Jennings. Cooperatives of distributed energy resources for efficient virtual power plants. In *The 10th International Conference on Autonomous Agents and Multiagent Systems – Volume 2*, AAMAS '11, pages 787–794, Richland, SC, 2011. IFAAMAS.

[DCD11] L. De Castro and J. Dutra. The economics of the smart grid. In *49th Annual Allerton Conference on Communication, Control, and Computing 2011*, pages 1294–1301, 2011.

[Ene13] Energymap.info. Energymap – Bavaria, 2013. visited April 17, 2013: http://energymap.info/energieregionen/DE/105/111.html.

[Eri10] G.N. Ericsson. Cyber Security and Power System Communication—Essential Parts of a Smart Grid Infrastructure. *Power Delivery, IEEE Transactions on*, 25(3):1501–1507, July 2010.

[HL04] B. Horling and V. Lesser. A survey of multi-agent organizational paradigms. *The Knowledge Engineering Review*, 19(04):281–316, 2004.

[HLGR06] J.S. Heo, K.Y. Lee, and R. Garduno-Ramirez. Multiobjective control of power plants using particle swarm optimization techniques. *Energy Conversion, IEEE Transactions on*, 21(2):552–561, 2006.

[KVM10] K. Kok, G. Venekamp, and P. Macdougall. Market-based control in decentralized electrical power systems. In *First International Workshop on Agent Technologies for Energy Systems, ATES2010, Toronto*, 2010.

[LPR09] P. Lombardi, M. Powalko, and K. Rudion. Optimal operation of a virtual power plant. In *Power Energy Society General Meeting (PES '09)*, pages 1–6. IEEE, 2009.

[MM09] P. McDaniel and S. McLaughlin. Security and Privacy Challenges in the Smart Grid. *Security Privacy, IEEE*, 7(3):75 –77, 2009.

[PRS07] D. Pudjianto, C. Ramsay, and G. Strbac. Virtual power plant and system integration of distributed energy resources. *Renewable Power Generation, IET*, 1(1):10–16, 2007.

[RVRJ12] S. Ramchurn, P. Vytelingum, A. Rogers, and N. Jennings. Putting the "Smarts" into the Smart Grid: A Grand Challenge for Artificial Intelligence. *Communications of the ACM*, 55(4):86–97, 2012.

[SA00] D. Smathers and A. Akhil. Operating Environment and Functional Requirements for Intelligent Distributed Control in the Electric Power Grid. Technical report, Sandia National Laboratories, 2000.

[SBA+13] J.-P. Steghöfer, P. Behrmann, G. Anders, F. Siefert, and W. Reif. HiSPADA: Self-Organising Hierarchies for Large-Scale Multi-Agent Systems. In *Proceedings of the Ninth International Conference on Autonomic and Autonomous Systems (ICAS 2013)*, Lisbon, Portugal, 2013. IARIA.

[SC01] A. Sage and C. Cuppan. On the systems engineering and management of systems of systems and federations of systems. *Information, Knowledge, Systems Management*, 2(4):325–345, 2001.

[SRK05] C. Schulz, G. Roder, and M. Kurrat. Virtual power plants with combined heat and power micro-units. In *2005 International Conference on Future Power Systems*, November 2005.

[UCT09] UCTE. UCTE Operation Handbook – Policy 1: Load-Frequency Control and Performance. Technical Report UCTE OH P1, Union for the Co-ordination of Transmission of Electricity, 2009.

[VRV+10] P. Vytelingum, S.D. Ramchurn, T.D. Voice, A. Rogers, and N.R. Jennings. Trading agents for the smart electricity grid. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, volume 1, pages 897–904. International Foundation for Autonomous Agents and Multiagent Systems, 2010.

[WLHK06] H.F. Wedde, S. Lehnhoff, E. Handschin, and O. Krause. Real-time multi-agent support for decentralized management of electric power. In *Proceedings of the 18th Euromicro Conference on Real-Time Systems*, pages 43–51. IEEE Computer Society Press, 2006.

[ZCRAC08] A. Zafra-Cabeza, M.A. Ridao, I. Alvarado, and E.F. Camacho. Applying risk management to combined heat and power plants. *Power Systems, IEEE Transactions on*, 23(3):938–945, 2008.