Modelling and Decentralised Runtime Control of Self-stabilising Power Micro Grids*

Arnd Hartmanns and Holger Hermanns

Saarland University – Computer Science, Saarbrücken, Germany

Abstract. Electric power production infrastructures around the globe are shifting from centralised, controllable production to decentralised structures based on distributed microgeneration. As the share of renewable energy sources such as wind and solar power increases, electric power production becomes subject to unpredictable and significant fluctuations. This paper reports on formal behavioural models of future power grids with a substantial share of renewable, especially photovoltaic, microgeneration. We give a broad overview of the various system aspects of interest and the corresponding challenges in finding suitable abstractions and developing formal models. We focus on current developments within the German power grid, where enormous growth rates of microgeneration start to induce stability problems of a new kind. We build formal models to investigate runtime control algorithms for photovoltaic microgenerators in terms of grid stability, dependability and fairness. We compare the currently implemented and proposed runtime control strategies to a set of approaches that take up and combine ideas from randomised distributed algorithms widely used in communication protocols today. Our models are specified in Modest, an expressive modelling language for stochastic timed systems with a well-defined semantics. Current tool support for Modest allows the evaluation of the models using simulation as well as model-checking techniques.

1 Introduction

Political and climatical circumstances are causing a shift in electric power production around the world. While large conventional power plants dominated electric power generation up to now, the future will see a drastic increase in the number of distributed microgenerators based on renewable energy sources such as solar and wind power. Electric power grids thus move from a setting in which production was assumed fully controllable so as to always match the uncontrollable, but well-predictable consumption to a setting where the production side becomes uncontrollable, too. External influences such as changing weather conditions can imply drastically higher fluctuations in available electric power.

* This work has been supported by the DFG as part of SFB/TR 14 AVACS, by the DFG/NWO Bilateral Research Program ROCKS, and by the European Union FP7-ICT project MEALS, contract no. 295261.

This problem is amplified by the difficulty of centrally controlling the vast number of geographically distributed microgenerators. New solutions to the problem of matching electricity production and consumption need to be found that are suitable to overcome these new challenges.

The German power grid is a prime example where many of these future challenges are already encountered today. As a consequence of the legal framework enforced by Federal legislation over the last decades, microgenerators of photovoltaic (PV) electric power have been rolled out massively on the rooftops of end user homes all over the country. In spite of a national target growth of 1.5 gigawatt (GW) per year, the total PV generation capacity has increased from 10 GW in 2009 to 25 GW by the end of 2011. The currently estimated actual growth rate is about 1.6 GW per year [9] (despite a target growth of less than 1.1 GW). This growth creates problems, especially in areas with additional microgeneration based on wind or biogas: The Northern German energy provider EWE AG recently reported that the number of emergency situations that required manual intervention to ensure grid stability has grown from less than 1 per week in 2009 to about 1 per day in 2011 [22].

To avoid these situations in the future, improved and better coordinated diagnostic and prediction techniques as well as orchestrated demand-side mechanisms to counter critical grid and/or generation situations are needed. To develop robust and correct mechanisms that do not create unexpected instability, e.g. by introducing oscillatory behaviour, mathematically well-founded models of electric power grids and their components are needed. However, the modelling space is huge, and a precise model reflecting all components in a detailed, physically exact manner will be very complex (if at all possible), and virtually impossible to analyse. Instead, suitable abstractions need to be developed, tailored to the fragments of the system under consideration and the aspects of interest. This will be the topic of the first part of this paper, where we give an overview of the various system aspects of electric power grids, in particular of last-mile micro grids with a significant fraction of microgeneration, and the future challenges faced in such grids (Section 2). We also take a look at the modelling challenges encountered in the study of such systems, surveying different modelling and abstraction approaches suitable for different system aspects and measures (Section 3). One expressive modelling formalism that fits this scenario particularly well is Modest, a modelling language for stochastic timed systems with a formal semantics and good tool support [6, 14], which we will use in the remainder of the paper.

A central issue to ensure the stability of future power grids is proper runtime control for the increasing number of microgenerators. Due to their distributed deployment, decentralised runtime control offers several advantages over centralised management approaches. The ideal is that of a network of independent generators whose control algorithms lead to a self-stabilising system. In the second part of this paper, we thus focus on the study of runtime control algorithms for photovoltaic microgenerators as an example for the modelling and analysis of future power grids. We first introduce the concepts of the currently implemented

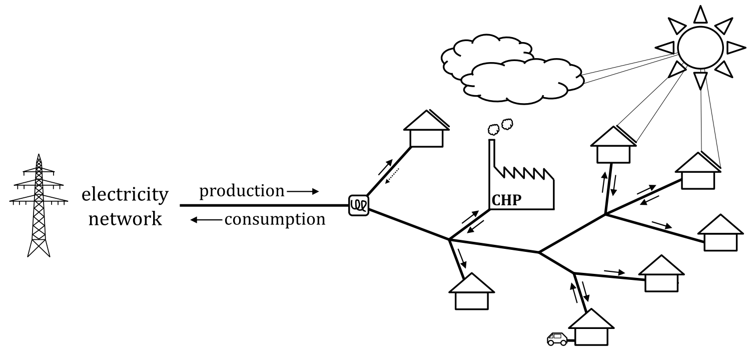

Fig. 1. Power micro grids

and proposed approaches (the former being known to introduce unwanted oscillatory effects) as well as of a potpourri of alternative approaches that take up and combine ideas from communication protocol design (Section 4). We model these approaches in Modest and report on the results of a recent simulation study using these models [3] (Section 5).

Related work. The area of power grid modelling with formal behavioural models is gaining momentum. The most closely related work is likely the recent paper by Chen *et al.* [10], who analyse a multi-player game based on a recently proposed distributed demand-side micro grid management approach [18]. Other tangible work includes the application of probabilistic hybrid automata with distributed control to the power grid domain [21], and work on network calculus in battery buffered households [8].

2 Last Mile Power Micro Grids

The electric power grid is hierarchically structured, with a grid of long distance high voltage lines forming the top layer. At the leaves of the electric power grid hierarchy, we find the low voltage last mile which traditionally connects end consumers to the upper layers. Typically, these last miles have a tree-like structure through which electric power is distributed towards the leaves from a root. This root is a transformer, which constitutes the connection to the upper layer. Since these grids are relatively small (comprising at most a few hundred residential homes or business customers) and have a clear point of separation from the remaining grid, yet may themselves contain multiple independent microgenerators, we call these last miles *power micro grids*. Figure 1 gives a schematic overview of an exemplary power micro grid consisting of seven residential homes and a small industrial customer.

2.1 Elements of Power Micro Grids

A model of a power micro grid needs to take five central aspects into account: (1) the influence of the wide-area power grid it is connected to, (2) the local consumption of electric power, (3, 4) the grid local electric power generation—which can be further divided into (3) potential and (4) actual generation, i.e. the amount of electric power that can be produced in ideal external circumstances (such as weather and time of day), and the amount that is actually produced after control algorithms inside the generators have been applied—and finally (5) the geographic topology and capacities of the cabling inside the micro grid.

Wide-area connection. A power micro grid usually has a single connection point to the wide-area electric power grid. This is a transformer station that converts the network's high voltage to the grid's 400 Volt three-phase current (or 230 V per phase). Traditionally, electric power flows from large conventional 'thermal' power plants through the wide area grid into the micro grids. This infeed is controlled by grid coordinators based on predictions of the local consumption of all the micro grids [17], corrected by runtime observations. Runtime deviations must be corrected in order not to destabilise any grid. Due to the physical limitations related to the power plants in use, only a small fraction of the total generation potential can be employed for runtime adaptation. This is mainly realized with the help of pump-storage plants, where subtracting power is achieved by pumping up water, while adding power is achieved by the reverse, turning water downfall into electric power.

As part of the interconnection to the wide-area grid there is also a safety "fuse", a device that may disconnect the micro grid, intended as a preventative measure for both local events, e.g. to prevent fatal accidents when a cable is damaged during excavation works, as well as interference from the wide area grid, e.g. to prevent excessive infeed that would exceed the electric power flow capacity in the micro grid. With increasing microgeneration inside micro grids, this safety device may actually turn into a problem, for example by disconnecting the micro grid in case local overproduction exceeds the fuse specifications.

Local consumption. At the leaves of the cabling inside the micro grid are residential homes and business customers. In the past they acted only as electric power consumers. The consumption of an individual leaf ultimately depends on a number of factors and decisions by its "inhabitants", yet it roughly follows patterns over the course of a day. Variations may be due to external influences such as temperature, influencing the electric power needed for heating or cooling. As such, consumption is uncontrollable, but predictable within certain error bounds. There is a recent trend to make consumption more controllable via so-called demand-side mechanisms [17], which intend to control the energy consumption of schedulable devices such as off-peak storage heaters and air conditioners. The decisions are to be based on electric power costs or on grid stability conditions.

Generation potential. More and more traditional consumers at the leaves of the micro grid are turning into producer-consumers (a.k.a. "prosumers"). At certain times, they may produce more electric power than they consume. The potential output of the microgenerators installed at these leaves depends first and foremost on the type of generator: Combined heat and power plants (CHP) can essentially operate on demand, independent of external circumstances, while microgenerators based on renewable energy sources such as wind and solar power are inherently dependent on natural phenomena. These vary over time in an uncontrollable manner. Wind turbines show relatively moderate fluctuations since wind intensity usually changes only gradually; the amount of available solar power, however, can change rapidly and significantly when cloud coverage changes quickly.

Actual generation. To avoid grid instability, consumption and production of electric power needs to be matched continuously in real time. The actual electric power emitted into the grid by a locally installed microgenerator may affect this stability. With the further increase of these sources, effective control mechanisms are needed in order to avoid over- or underprovisioning of power. Technically, it is no problem to reduce the output of all relevant types of generators—the problem is to decide when to do so, by which amount, when to switch the generators back on, and by how much. Control algorithms are thus an important aspect of future microgenerators. They are expected to have significant influence on the behaviour of future power micro grids.

Local grid topology. The topology and spatial layout of the micro grid in terms of cable lengths and diameters clearly impacts its behaviour. The grids have been rolled out in the past with the sole perspective of distributing power downstream, i.e. towards the leaves of the last miles. Now there might be upstream power flow in some parts of the grid. It is easy to come up with scenarios where this may result in stability violations (such as excessive voltage) inside the grid that remain unnoticed at the leaves and at the root. The proper reflection of these influences in a way that generalises to arbitrary last miles is very difficult, because it crucially depends on a specific layout.

2.2 Modelling and Abstraction Choices

Since a full model of all individual components of a power micro grid and their precise behaviour is extremely difficult to build and most probably entirely impossible to analyse, the various components have to be represented at appropriate levels of abstraction in a model. These abstractions have to be chosen carefully to make modelling and analysis feasible, yet provide sufficient information to extract reliable answers to the questions of interest from the model.

A first candidate for abstraction is the contribution of the wide-area grid. A detailed modelling of the wide-area grid is clearly out of the scope of a model focussed on just a single power micro grid, while the reverse, the impact of a single micro grid on the behaviour of the entire (e.g. European) electric power grid can be considered negligible. It is thus reasonable to represent the influence of the wide-area grid in the form of a profile, i.e. a deterministic or stochastic function mapping time to the amount of electric power provided. This is an instance of what is called a "load profile", and is itself assumed independent of what happens inside the particular micro grid. In addition, the safety fuse at the root we mentioned does not need to be explicitly modelled; instead it is present in the analysis as part of the characterisation of what "unsafe" or "unstable" states need to be avoided.

When it comes to modelling consumer behaviour, the abstraction level depends on the intended modelling purpose. If the focus is on the effects of consumer behaviour, as in a study of demand-side management mechanisms, a detailed consumer model and the explicit representation of individual consumers are obvious necessities. If this is not the focus, two choices are to be made: Should consumers be represented individually or in aggregated form (i.e. as a load profile), and how detailed does the individual or aggregated consumer model need to be? Modelling consumers individually allows the differentiation of consumer types (e.g. into households and businesses) to be represented directly. These distinctions would only lead to variations of the chosen load profile otherwise. Another fundamental question is whether to use a deterministic, stochastic or nondeterministic model of consumption. While deterministic models are often easier to analyse, they embody the risk of exhibiting or causing spurious oscillations or correlations mainly because they may ignore differences between the participants. A stochastic model typically is a good way to avoid these phenomena by assigning probabilities to different behaviours that are all considered part of the model. When it is not possible to assign probabilities to behaviours, nondeterministic models may capture all possible alternatives, but may often turn out to be hard or impossible to analyse.

The modelling spectrum on the power generation side is similar to that on the consumer side. Given a fixed set of generators of different types, a (deterministic or stochastic) load profile is a good representation of the potential generation. It can represent how the external influences on generation potential vary over time, and since a grid covers only a very restricted geographic area (of maybe 1 km^2), it can be considered constant throughout the geographic dimension, since local differences in wind or cloud cover are negligible at this resolution. With respect to actual generation, a load profile may be a good first step, but hide interesting behaviour that can result from inappropriate control algorithms. For example, the currently deployed control algorithm for PV generators in Germany can lead to oscillating behaviour in times of high potential generation once an unsafe grid state is reached (see Section 4.2). In order to study, for example, whether certain demand-side mechanisms can avoid or buffer these oscillations, one would need at least a simple behavioural model of the actual generation.

Finally, the role played by the grid topology is closely tied to the way the physical aspects of electric power are represented in the model. Intertwined differential equations or calculations with complex numbers are the norm, needed to provide nontrivial answers about frequency and voltage. They are achievable for specific layouts. A common abstraction that helps to provide valid answers on a more abstract level assumes the local grid to behave like a perfect "copper plate", thus eliminating any spacial considerations.

2.3 Properties and Challenges

As the installed microgeneration capacity increases, the effect of power micro grids on the whole network gets more significant. At the same time, as most microgenerators are based on renewable energy sources, the volatility in the micro grids' behaviour becomes an important concern. There are two core objectives of micro grid and microgenerator management: economy and stability, which are deeply intertwined, yet often conflicting interests.

Challenges. In European legislation, an electric power grid has two distinct modes of operation: *emergency operation*, where direct intervention of the grid coordinator is needed to drive the grid to a safe state, possibly impacting service levels on the consumer side, and *normal operation*, where market incentives drive the decisions of the participants. The stability of the grid is a priority concern because reliable distribution is a prerequisite for economic use of energy. However, the most economically beneficial decisions for individual participants may sometimes run counter to the goal of a stable grid. Grid instability is caused by over- or underproduction, respectively under- or overconsumption, i.e. the electric power production does not match the current consumption. It can be stabilised by suitably adjusting production, consumption, or both.

On the production side, the main issue is to avoid overproduction: While some generation technologies such as CHP are perfectly controllable, the upper limit on potential generation of renewable electric power is dependent on natural phenomena; control strategies for these microgenerators can thus only reduce production compared to their genuine potential. On the other hand, the economic interest of microgenerator owners is to feed as much energy into the grid as possible. In this sense, grid stability and production economy are conflicting interests. Control strategies on the production side, whose overriding goal is to ensure grid stability, thus have to be evaluated for efficiency and fairness in the economic sense as well.

In contrast to this, economic interests can be used as a way to guide the consumption side to a behaviour that is beneficial to stability: Over- and underproduction ideally have a direct effect on the price of electricity, which can drive demand in the desired direction. Nevertheless, the study of effective demand-side mechanisms that lead to compensation of production volatility, with or without economic aspects, is an area as widely open for research as the production side.

Properties. We propose the following set of measures to evaluate production control algorithms and demand-side mechanisms, which we collectively call *strategies*, for electric power micro grids:

– *Stability* is the ability of a strategy to keep the grid in a safe state with a minimum of oscillation between safe and unsafe states.
– *Availability* is the overall fraction of time that the grid spends in a safe state.
– *Output* measures the (total or individual, cumulative or averaged) electricity output of the relevant microgenerators, which is usually proportional to the financial rewards of the respective operators.
– *Goodput* relates output to availability: the amount of electric power a generator can add to the grid while the grid is in a safe state.
– *Quality of Service* measures negative impacts on the consumer side. While closely tied to availability, quality of service can also vary while the grid is in a safe state, for example if service reductions are used to achieve safety.
– *Fairness* is the degree to which a strategy manages to distribute adverse consequences equally among the participants. When the grid state does not allow all generators to operate at full power, for example, will each of them be allowed to provide an equal share of the allowed power generation?

3 Formal Modelling Challenges

Power micro grids are complex systems that require expressive modelling formalisms to capture the entirety of their behaviour. Even if only abstracted subsets of a micro grid shall be represented, features such as real-time behaviour and stochastics are necessary, e.g. to model delayed reactions by the grid controller and stochastic load profiles or randomised algorithms. In order to faithfully represent the precise physical behaviour of the electric components together with a discrete control strategy, a versatile modelling formalism is a necessity. A more exhaustive discussion on what kind of modelling features are needed for this problem domain can be found in [16]. However, there is an inherent tradeoff between expressivity and the analysis effort needed to compute results. Every modelling study thus needs to precisely identify the aspects to be included in the model as well as the kinds of properties to be analysed so as to make it possible to select the best matching formalism that is still sufficiently expressive.

3.1 Modest

Modest [6] is a high-level modelling and description language for stochastic timed systems that combines expressive and powerful syntax-level features with a formal semantics in terms of stochastic timed automata (STA). Stochastic timed automata add continuous probability distributions, allowing in particular arbitrarily (e.g. uniformly or exponentially) distributed delays, to probabilistic timed automata (PTA) [19], which themselves can be seen as the orthogonal combination of timed [1] and probabilistic automata [24] (or, equivalently, Markov decision processes [23]). Other special cases of STA are generalised semi-Markov processes (GSMP) [12], which essentially constitute STA without nondeterminism, and both discrete- as well as continuous-time Markov chains (DTMCs and CTMCs). Modest has recently been extended to support the specification of stochastic hybrid automata (SHA) models as well [13].

The key feature of MODEST that makes it attractive for electric power micro grids is that it is built around a *single-formalism, multiple-solution* approach: While expressive enough to specify SHA, most of the various well-known and extensively studied submodels can be easily identified on the syntactic level, and tool support dedicated to these submodels is available [4, 5, 15]. This allows a single language to be used for a wide range of models while benefiting from using restricted formalisms to achieve efficient analysis.

Syntactically, MODEST supports a process algebra-inspired compositional modelling approach. It allows smaller models to be combined into larger, more complex ones, including a parallel composition operator to specify processes or automata that perform their actions independently, subject to the classical interleaving semantics. Actions that are part of the shared alphabet of two or more processes have to be performed by all processes involved in a CSP-style synchronisation. We refer the interested reader to [6] and the MODEST TOOLSET website at

`www.modestchecker.net`

for details concerning the language's design and the semantics of its constructs. The website also contains further documentation, a list of MODEST-related publications as well as examples and case studies. We will use MODEST to build formal models of runtime control strategies for photovoltaic microgenerators in Section 5.

4 Decentralised Runtime Control

A major portion of the photovoltaic (PV) microgeneration capacity is mounted on the rooftops of private households, and is as such connected to the last mile. The often excessive volatility of solar production asks for a highly flexible grid management on this level. For the remainder of this paper, we therefore focus on control strategies for PV microgenerators. As outlined in the previous section, the goal of such a strategy is to reduce actual power output compared to the potential generation whenever this is necessary to maintain grid stability. Otherwise it should allow the output of as much electric power as can be generated.

Let us first take a deeper look at what constitutes a "safe state" for power (micro) grids. There are three fundamental dimensions to stability:

- In Europe, the target **frequency** is 50 Hz. If the frequency leaves the band of 49.8 to 50.2 Hz, this is a serious Europe wide phenomenon.
- In the end customer grid, the downstream customers may witness considerable **voltage** fluctuations because of upstream fluctuations in production and consumption. Deviations of more than 10 % are not tolerable.
- There are individual limits on the **capacity** of grid strands with respect to energy, i.e. the product of voltage and amperage.

The capacity limits are due to the local grid layout and the "fuse" at the connection point to the upper layers. Voltage has a direct linear dependency to production/consumption and is thus a good measure of the grid state. However, voltage changes are local phenomena, entangled with phase drifts in the last mile and intimately tied to the grid topology and the distances and cabling between producers and consumers. Therefore, the frequency is often used instead of voltage as a measure of the grid state, although frequency drifts usually affect the entire European grid and not only a specific last mile and are subject to dampening effects. An approximately linear dependency between production/consumption and frequency is known, albeit being an indirect effect of physical realities. However, it is still considered an appropriate abstraction by domain experts [20, 25]. Roughly, a change in production/consumption of 15 GW approximately corresponds to a 1 Hz change in frequency in the European grid. The currently installed PV generation capacity in all of Germany (see Section 1) thus corresponds to a frequency spread of about 1.7 Hz.

4.1 Centralised vs. Decentralised Control

Photovoltaic microgenerators are difficult to manage. First, this is due to their sheer number, which leads to problems of scalability for any centralised approach. A second problem is their distributed nature: There is currently no measurement, logging and reporting infrastructure in place that enables the collection of accurate and up-to-date information about the state of the grid participants, and there is no communication infrastructure that allows safe remote control. These are two good reasons to consider highly local, decentralised and automatic grid management approaches. Additionally, decentralised approaches that do not need any transmission of information to central coordinators are inherently preferable from a privacy perspective.

The design of a highly local, highly automatic, highly decentralized, and highly flexible grid management is a challenging and pressing problem. It resembles the field of self-stabilising system (SSS) design [11]. SSS are built from a number of homogeneous systems that follow the same algorithmic pattern, with the intention that their joint execution emerges in a stable global behaviour, and can recover from transient disturbances. Compared to the setting usually considered in SSS, there are however some important differences: In a power grid, destabilisation threats must be countered within hard real time bounds. This is usually not guaranteed for SSS. On the other hand, in SSS usually no participant is considered to have knowledge about the global system state, while in a power grid, the participants do in principle have access to a joint source of localized information by measuring amperage, voltage and frequency.

4.2 Current Approaches

About 75% of the PV microgenerators rolled out so far in Germany are non-measured and cannot be remotely controlled. Since 2007, a regulation is in place that enforces a frequency-based distributed control strategy (EN 50438:2007). It stipulates that a microgenerator must shut off once the frequency is observed to overshoot 50.2 Hz. While this was initially meant as a way to stabilise the grid by cutting overproduction, it later surfaced that due to the high amount of PV generation, an almost synchronous distributed decision to take out this portion may induce a sudden frequency drop, followed by the PV generators joining back in, and so on. It hence may lead to critical Europe-wide frequency oscillations.

Due to the obvious problems that widespread use of the current rules may lead to, new requirements are being developed as part of VDE-AR-N 4105 [7]. PV generators will be required to implement the following control scheme:

- As long as the observed frequency is below 50.2 Hz, the generator may increase its output by up to 10 % of the maximum output that it is capable of per minute.
- When the observed frequency crosses the 50.2 Hz mark, the current output of the generator is saved as p_m. When the frequency f is between 50.2 and 51.5 Hz, the generator must reduce its output linearly by 40 % per Hertz relative to p_m, i.e. its output is given by the function

$$output(f) = p_m - 0.4 \cdot p_m \cdot (f - 50.2).$$

- In case the observed frequency exceeds 51.5 Hz, the generator has to be switched off immediately and may only resume production once the frequency has been observed to be below 50.05 Hz for at least one minute.

As we will see (Section 5.3), this relatively complex algorithm is designed to dampen the effect of PV generation spikes and to avoid introducing oscillatory behaviour, but not to actively steer the system towards a safe state where the frequency is below 50.2 Hz.

4.3 Probabilistic Alternatives

If we look at the PV control problem in a more abstract way, it turns out to be remarkably similar to problems solved by communication protocols in computer networks such as the Internet: Limited bandwidth (in our case, capacity of the power grid to accept produced electric power) needs to be shared between a number of hosts (in our case, generators) in a fair way. We thus consider several new control algorithms inspired by concepts from communication protocols, most of which use randomisation to break synchrony and avoid deterministic oscillations:

Additive increase, multiplicative decrease: The first new control algorithm that we study is inspired by the way the Internet's Transmission Control Protocol (TCP) achieves fair usage of limited bandwidth between a number of connections: Bandwidth usage is increased in constant steps (additively), and when a message is lost (taken as an indication of buffer overflows due to congestion), it is reduced by a constant factor (multiplicatively). This *additive-increase, multiplicative-decrease* (AIMD) policy ensures that several users of the same connection eventually converge to using an equal share of the bandwidth. We directly transfer this approach to PV generators: Power output is increased in small constant steps until the frequency is measured to be above 50.2 Hz, at which point the output is multiplied by a constant factor < 1.

Frequency-dependent probabilistic switching: Our hypothesis is that probabilistic strategies may improve stability without requiring fine-grained modifications of the generators' power output as in AIMD. Our next controller thus always switches between full and no power output, but it does so with a certain probability that depends on the current frequency measurement with higher frequencies leading to a higher probability of switching off.

Exponential backoff: Instead of determining the switching probability based on the current system state, we can also unconditionally switch off when the frequency exceeds the allowed value of 50.2 Hz and then wait a probabilistically chosen amount of time before again measuring and potentially switching on.

The precise scheme that we use is *exponential backoff* with *collision detection*. In the computer networks domain, this is commonly employed in CSMA/CD-based (*carrier sense multiple access* with collision detection) medium access protocols such as Ethernet: When one device connected to the shared medium (e.g. the cable) has data to send, it first *senses* the carrier to determine whether another device is currently sending. If not, it sends its data immediately. However, if the channel is occupied or if the sending is interrupted by another device starting to send as well (a collision), it waits a number of time slots before the next try. This number is sampled from a uniform distribution over a range such as $\{1, \ldots, 2^{bc}\}$, where bc, the *backoff counter*, keeps track of the number of collisions and of the number of times that the channel was sensed as occupied when this message should have been sent. The range of possible delays increases exponentially, thus the policy's name; its goal is to use randomisation to prevent two devices from perpetually choosing the same delay and thus always colliding, and to use an exponential increase in the maximum waiting time in order to adapt to the number of devices currently having data to send (again, in order to avoid continuous collisions).

The goals of exponential backoff in network protocols closely match our goals in designing a power generation control scheme: We want all generators to be able to feed power into the grid when it is not "occupied", i.e. when the frequency is below the threshold of 50.2 Hz, and we want to avoid "collisions", i.e. several generators switching on at about the same time and thus creating frequency spikes above that threshold.

Frequency-dependent switching with exponential backoff: An obvious final step is to combine the frequency-dependent probabilities and the randomised delays of the previous two controllers to create one that features randomisation of both switching decisions and waiting times.

We have also considered additional algorithms and variants of those presented above; a full list with detailed explanations can be found in the accompanying technical report [2].

5 Modelling Decentralised Controllers

In order to evaluate the behaviour of the different PV generator control strategies introduced in the previous section, we build MODEST models for power micro grids that use these controllers. As our focus is on the generator control aspect, we represent the other elements of the micro grid listed in Section 2.1 as follows:

- Wide-area connection and local consumption: The influence from the upper layer power grid on our last mile as well as the local consumption within the last mile is modelled as a combined deterministic load profile.
- Generation potential: We model the "worst case" of a maximally sunny day. Each PV generator is assumed to be able to contribute the full amount of power it is capable of (given by a constant MAX) into the grid at any time.
- Local grid topology: We abstract from the physical characteristics of the grid by treating the local connections as a "copper plate" and looking only at the frequency observed. We chose this drastic abstraction due to the reasons outlined in the introduction to Section 4. By treating frequency as a local phenomenon—which it is not—we exaggerate the influence of the individual PV generators. We could easily use voltage as a reference quantity instead since the models are sufficiently abstract.

Since our focus is on effects of overproduction, we only consider the frequency range above 50 Hz, thus representing 50 Hz as frequency value 0 in our model. This value is assumed when all solar generators are switched off and there is no (= zero) influence from the wide-area connection and local consumption. We assume that adding power to the grid has a linear effect on the frequency, so we can describe the grid frequency as the sum of the generator outputs plus the in-feed from the upper layer minus the consumption.

5.1 A Model Template for Power Micro Grids

The detailed models of the control strategies all fit into the same model template shown in Figure 2. The control strategies become part of a `Generator` process, while a `LoadProfile` process represents the wide-area influence and local consumption; the entire system is finally specified as the parallel composition of G instances of `Generator` plus a single `LoadProfile` instance. This template shows a few more noteworthy modelling choices and abstractions:

Each generator repeatedly measures the grid's current frequency, uses this value to decide whether and in which way to modify its own power output, and finally update its output according to this decision. Each of these measure-update cycles takes M time units, with $D \leq M$ time units passing between the measurement and the change of power output. This delay allows us to model decision and reaction times as well as the time it actually takes for the changes made by one generator to be observed by the others. Higher values of D will thus lead to decisions being made on "older" data, while $D = 0$ implies that every change is immediately visible throughout the last mile. We have thus chosen a

```
action init;

const int TIME_BOUND; // analysis time bound
const int G; // number of generators

// Times
const int M; // measure every M time units
const int D; // changes take D time units to take effect (D <= M)

// Frequencies (in Hz above 50.0 Hz)
const real B = 0.3; // frequency when all generators are on full power
const real MAX = B / G; // max output of a generator (contribution to frequency)
const real L = 0.1; // max sum of wide-area influence and local consumption

real input; // background generation (coming from the network), in [0, L]
real[G] output; // generator output, each in [0, MAX]

function real frequency() = input + /* sum of values in output array */;

reward r_availability; der(availability) = frequency() > 0.2 ? 0 : 1;
reward r_output; der(sumoutput) = frequency() - input;
reward r_goodput; der(goodput) = frequency() > 0.2 ? 0 : frequency() - input;

property Availability = Xmax(r_availability / TIME_BOUND | time == TIME_BOUND);
property Output = Xmax(r_output / TIME_BOUND | time == TIME_BOUND);
property Goodput = Xmax(r_goodput / TIME_BOUND | time == TIME_BOUND);

process GeneratorInit(int(0..G) id)
{
   // Generators are initially in a random state
   urgent init {= output[id] = Uniform(0, MAX) =};
   // Each generator "starts" after a random delay in [0, M]
   delay(Uniform(0, M)) Generator(id)
}

process Generator(int(0..G) id)
{
   action measure, update;
   real fm; // frequency measurement
   clock c = 0; // local clock variable

   process Measure()
   {
      measure {= fm = frequency(), c = 0 =}
   }

   /* control algorithm is modelled here */
}

process LoadProfile()
{
   /* load profile is modelled here */
}

par {
   :: GeneratorInit(0)
      /* ... */
   :: GeneratorInit(G - 1)
   :: LoadProfile()
}
```

Fig. 2. A model template for power micro grids

discrete measure-update-wait approach; an alternative is to make the generators reactive, i.e. observe the evolution of the frequency and react when relevant thresholds are crossed.

By use of the `GeneratorInit` process, each generator begins operation after a random, uniformly distributed delay in the range between 0 and M time units; measurements will thus be performed asynchronously. Less realistic, but easier to analyse alternatives would be to have the generators perform their decisions in a fully synchronous manner, or at the same point of time, but in a certain order. However, we have observed that in particular the second alternative generates extreme results (e.g. for fairness) that are clearly artifacts of that abstraction.

5.2 Control Strategy Models

We now explain how to model the control strategies described in sections 4.2 and 4.3 in MODEST to fit into the template introduced above:

Current approaches: We omit the trivial MODEST code for the simple control strategy that turns the generator off when a frequency of at least 50.2 Hz is observed and turns it to full power in all other cases. A direct implementation of the new control scheme according to VDE-AR-N 4105 is shown in Figure 3. The switch between normal and emergency mode is obvious in the model.

Probabilistic Alternatives: Figure 4 shows the model of the AIMD controller. In this case, we chose 10 % of the maximum generator output as the constant value when increasing, and $\frac{2}{3}$ as the decrease factor. The latter has shown to provide a good tradeoff between availability and goodput when we compared our analysis results (see next section) for different reduction factors. The MODEST code for the frequency-dependent probabilistic switching controller is shown in Figure 5. We have chosen a linear function over the range of $[50.0\,\text{Hz}, 50.4\,\text{Hz}]$ for the mapping from measured frequency to switch-off probability. At the critical threshold of 50.2 Hz, the probability of switching off will thus be $\frac{1}{2}$. Finally, the controller based on the exponential backoff approach can be seen in Figure 6; the combination with frequency-dependent switching is not shown because it is just a simple replacement of the `when` conditions in exponential backoff with a probabilistic alternative (`palt`) that uses the chosen probability function.

5.3 A Simulation Study

We have evaluated the different control strategies in a dedicated simulation study [3]. The properties we considered are (as outlined in Section 2.3) stability, availability versus goodput, and fairness. To evaluate stability, we evaluated the frequency traces of exemplary simulation runs with a fixed background load. Figures 7 and 8 show these traces for three of the controllers we studied, namely the on-off controller (Figure 8, left) and the controller according to VDE-AR-N 4105 (right) as well as the combination of the frequency-dependent switching

```
process Generator(int(0..G) id)
{
   real p_m = output[id];
   /* ...template code... */

   process NormalOperation()
   {
      alt
      {
         :: when(fm < 0.2)
            // Increase by 10% of MAX per minute
            update {= output[id] += (0.1 * MAX) / MINUTE,
                      p_m        += (0.1 * MAX) / MINUTE =};
            when urgent(c >= M) Measure()
         :: when(0.2 <= fm && fm < 1.5)
            // 40% gradient
            update {= output[id] = -0.4 * p_m * (fm - 0.2) + p_m =};
            when urgent(c >= M) Measure()
         :: when(1.5 <= fm)
            // Switch off
            EmergencySwitchOff()
      };
      when urgent(c >= D) NormalOperation()
   }

   process EmergencySwitchOff()
   {
      bool waiting;
      clock minute;

      // Switch off
      update {= output[id] = 0, p_m = 0 =};

      // Wait for frequency to be below 50.05 Hz for one minute
      do {
         :: when urgent(c >= M && !(waiting && minute >= MINUTE)) Measure();
            urgent alt {
               :: when(fm <= 0.05 && !waiting)
                  {= waiting = true, minute = 0, c = 0 =}
               :: when(fm <= 0.05 &&  waiting)
                  {= c = 0 =}
               :: when(fm > 0.05)
                  {= waiting = false, c = 0 =}
            }
         :: when urgent(waiting && minute >= MINUTE) break
      }
   }

   Measure();
   when urgent(c >= D) NormalOperation()
}
```

Fig. 3. Model of the controller according to VDE-AR-N 4105

```
process Generator(int(0..G) id)
{
   /* ...template code... */

   Measure();
   when urgent(c >= D) alt {
      :: when(fm <  0.2) {= output[id] = min(MAX, output[id] + 0.1 * MAX) =}
      :: when(fm >= 0.2) {= output[id] *= 2/3 =}
   };
   when urgent(c >= M) Generator(id)
}
```

Fig. 4. Model of additive increase, multiplicative decrease of frequency

```
process Generator(int(0..G) id)
{
   /* ...template code... */

   Measure();
   when urgent(c >= D) update palt {
      :max(0, 0.4 - fm): {= output[id] = MAX =}
      :       fm       : {= output[id] = 0   =}
   };
   when urgent(c >= M) Generator(id)
}
```

Fig. 5. Model of the frequency-dependent probabilistic switching controller

```
process Generator(int(0..G) id)
{
   int bc; // backoff counter
   int backoff; // number of slots to wait till next try
   /* ...template code... */

   process Gen()
   {
      Measure();
      when urgent(c >= D) alt {
         :: when(backoff > 0) update {= backoff- =}
         :: when(backoff == 0) alt {
               :: when(fm <  0.2) {= output[id] = MAX, bc = 0 =}
               :: when(fm >= 0.2) {= output[id] = 0, bc++,
                     backoff = DiscreteUniform(0, (int)pow(2, bc)) =}
            }
      };
      when urgent(c >= M) Gen()
   }

   Gen()
}
```

Fig. 6. Model of the controller with exponential backoff

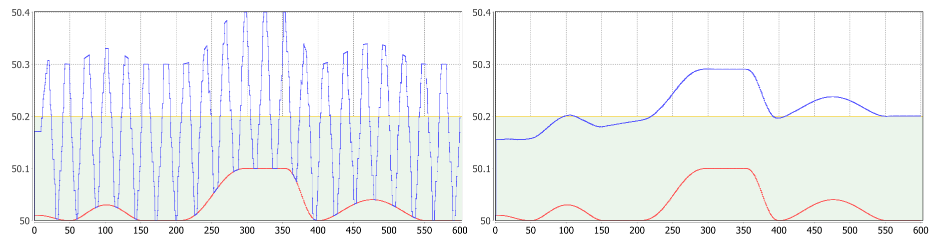

Fig. 7. Behaviour of the 50.2 Hz on-off (left) and VDE-AR-N 4105 controllers [3]

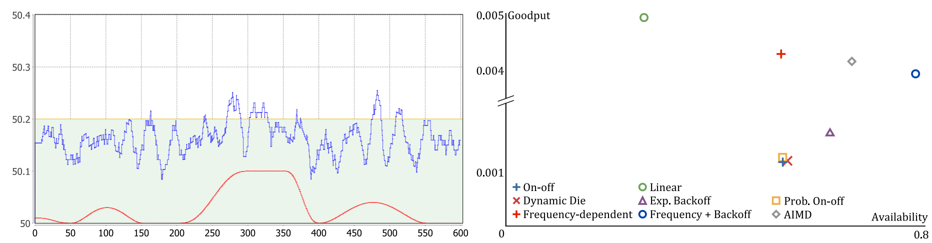

Fig. 8. Frequency-dependent control with backoff (left) and availability vs. goodput [3]

controller with exponential backoff (Figure 8, left). The upper (blue) curves plot the system frequency, while the lower (red) curves show the deterministic background load we used for these runs.

The oscillatory behaviour caused by the current 50.2 Hz on-off controller is clearly visible, as are the different behavioural phases of the new strategy according to VDE-AR-N 4105. The latter clearly avoids oscillations, but does not actively stabilise the grid into a safe state. As expected, the frequency-dependent controller shows a very different behaviour, which appears very erratic, but actually manages to keep the system safe for most of the simulation run. Results for the other newly proposed controllers were mixed: AIMD also works rather well and is at least fairer than its additive-decrease counterpart (as hoped), but exponential backoff alone does not manage to avoid oscillations. The graph on the right-hand side of Figure 8 compares the availability and goodput of the entire set of controllers considered in our simulation study, confirming that the combination of frequency-dependent randomisation with exponential backoff works rather well [3]. It also illustrates that taking inspiration from network protocols for distributed grid operation is indeed a promising direction.

6 Conclusion

This paper has discussed elementary mechanisms for distributed runtime control of power grids facing considerable infeed of renewable energy. We have focussed on the properties and modelling formalisms needed to describe, analyse and manage these systems in a highly flexible, highly automated, and highly decentralized manner. Another system which is highly decentralized, highly flexible and managed in a highly automated way is the Internet. As we have discussed, certain solutions that have been coined as part of Internet protocols can be adapted to serve beneficially in future distributed runtime control of power grids. This benefit might not be restricted to Internet solutions, but might more generally also materialise for some of the genuine Internet design principles, such as:

- Network neutrality and fairness: There is no discrimination in the way the network shares its capacity among its users. Ideally, the net is fair in the sense that if n users are sharing a connection, then on average each user can use about $1/n$-th of the capacity.
- Intelligent edges, dumb core: Putting intelligence into the net itself is much more cost ineffective than placing it at the edges of the networks, i.e. into the end user appliances.
- Distributed design and decentralised control: Distributed, decentralised control is not only a means to assure scalability. It also is a prime principle to protect end user privacy that would be at stake if centralised authorities would collect information for decision making.

There are a number of similarities between the Internet and the power grid, including its excessive size, its hierarchical structure, its organic growth, and its ultimately high dependability. Indeed, it seems to us that this implies a number of very good reasons why the future management of power grids should take strong inspirations from the way the Internet is managed. Our research indicates some first concrete examples of this kind.

Acknowledgments. The authors thank Pascal Berrang (Saarland University) for his assistance in carrying out the empirical studies. Mats Larrson (ABB Schweiz AG), Sebastian Lehnhoff (University of Oldenburg and OFFIS Energie), Martin Ney (Luxea GmbH), Alexandre Oudalov (ABB Schweiz AG), and Holger Wiechmann (EnBW Energie Baden-Württemberg AG) have provided insightful feedback on our findings.

References

1. Alur, R., Dill, D.L.: A theory of timed automata. Theor. Comput. Sci. 126(2), 183–235 (1994)
2. Berrang, P., Bogdoll, J., Hahn, E.M., Hartmanns, A., Hermanns, H.: Dependability results for power grids with decentralized stabilization strategies. Reports of SFB/TR 14 AVACS 83 (2012), ISSN: 1860-9821, `www.avacs.org`
3. Berrang, P., Hartmanns, A., Hermanns, H.: A comparative analysis of decentralized power grid stabilization strategies. In: Winter Simulation Conference (2012), to appear.
4. Bogdoll, J., David, A., Hartmanns, A., Hermanns, H.: mctau: Bridging the gap between Modest and UPPAAL. In: SPIN (2012), to appear.
5. Bogdoll, J., Hartmanns, A., Hermanns, H.: Simulation and statistical model checking for Modestly nondeterministic models. In: MMB/DFT. Lecture Notes in Computer Science, vol. 7201, pp. 249–252. Springer (2012)

6. Bohnenkamp, H.C., D'Argenio, P.R., Hermanns, H., Katoen, J.P.: MoDeST: A compositional modeling formalism for hard and softly timed systems. IEEE Transactions on Software Engineering 32(10), 812–830 (2006)
7. Bömer, J., Burges, K., Zolotarev, P., Lehner, J.: Auswirkungen eines hohen Anteils dezentraler Erzeugungsanlagen auf die Netzstabilität bei Überfrequenz & Entwicklung von Lösungsvorschlägen zu deren Überwindung (2011), study commissioned by EnBW Transportnetze AG, Bundesverband Solarwirtschaft e.V. and Forum Netztechnik/Netzbetrieb im VDE e.V.
8. Boudec, J.Y.L., Tomozei, D.C.: A demand-response calculus with perfect batteries. In: MMB/DFT. LNCS, vol. 7201, pp. 273–287. Springer (2012)
9. Bundesnetzagentur: EEG-Vergütungssätze für Photovoltaikanlagen. `http://www.bundesnetzagentur.de/cln_1931/DE/Sachgebiete/ElektrizitaetGas/ErneuerbareEnergienGesetz/VerguetungssaetzePVAnlagen/VerguetungssaetzePhotovoltaik_Basepage.html` (March 21 2012)
10. Chen, T., Forejt, V., Kwiatkowska, M.Z., Parker, D., Simaitis, A.: Automatic verification of competitive stochastic systems. In: TACAS. LNCS, vol. 7214, pp. 315–330. Springer (2012)
11. Dolev, S.: Self-Stabilization. MIT Press (2000)
12. Haas, P.J., Shedler, G.S.: Regenerative generalized semi-Markov processes. Communications in Statistics. Stochastic Models 3(3), 409–438 (1987)
13. Hahn, E.M., Hartmanns, A., Hermanns, H., Katoen, J.P.: A compositional modelling and analysis framework for stochastic hybrid systems. Formal Methods in System Design (2012), to appear
14. Hartmanns, A.: Model-checking and simulation for stochastic timed systems. In: FMCO. LNCS, vol. 6957, pp. 372–391. Springer (December 2010)
15. Hartmanns, A., Hermanns, H.: A modest approach to checking probabilistic timed automata. In: QEST. pp. 187–196. IEEE Computer Society (2009)
16. Hermanns, H., Wiechmann, H.: Future design challenges for electric energy supply. In: ETFA. pp. 1–8. IEEE (2009)
17. Hermanns, H., Wiechmann, H.: Embedded Systems for Smart Appliances and Energy Managment, Embedded Systems, vol. 3, chap. Demand-Response Managment for Dependable Power Grids. Springer Science+Business Media, New York (2012)
18. Hildmann, H., Saffre, F.: Influence of variable supply and load flexibility on demand-side management. In: EEM'11. pp. 63–68. IEEE Conference Publications (2011)
19. Kwiatkowska, M.Z., Norman, G., Segala, R., Sproston, J.: Automatic verification of real-time systems with discrete probability distributions. Theor. Comput. Sci. 282(1), 101–150 (2002)
20. Lehnhoff, S.: Private communication (2012)
21. Martins, J., Platzer, A., Leite, J.: Statistical model checking for distributed probabilistic-control hybrid automata with smart grid applications. In: ICFEM. LNCS, vol. 6991, pp. 131–146. Springer (2011)
22. Nordwest-Zeitung: EWE spürt Wende deutlich. `http://www.nwzonline.de/Aktuelles/Politik/Hintergrund/NWZ/Artikel/2822057/EWE-sp%FCrt-Wende-deutlich.html` (March 12, 2012)
23. Puterman, M.L.: Markov Decision Processes: Discrete Stochastic Dynamic Programming. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics, John Wiley & Sons Inc., New York (1994)
24. Segala, R.: Modeling and Verification of Randomized Distributed Real-Time Systems. Ph.D. thesis, MIT, Cambridge, MA, USA (1995)
25. Wiechmann, H.: Private communication (2012)