

A COMPARATIVE ANALYSIS OF DECENTRALIZED POWER GRID STABILIZATION STRATEGIES

Arnd Hartmanns
Holger Hermanns
Pascal Berrang

Saarland University – Computer Science
D-66123 Saarbrücken, GERMANY

ABSTRACT

This paper reports on formal behavioral models of power grids with a substantial share of photovoltaic micro-generation. Simulation studies show that the current legislative framework in Germany can induce frequency oscillations. This phenomenon is indeed recognized by the German Federal Network Agency responsible for overseeing the national power grids, and new regulations are currently being identified to counter this phenomenon. We study the currently valid proposal, and compare it with a set of alternative approaches that take up and combine ideas from communication protocol design, such as additive-increase/multiplicative-decrease known from TCP, and exponential backoff used in CSMA variations. We classify these alternatives with respect to their availability and goodput. The models are specified in the modeling language MODEST, and simulated with the help of the *modes* simulator.

1 INTRODUCTION

The electricity markets in Europe, Asia, and the Americas are evolving towards decentralized structures, essentially rooted in political decisions to counter the worldwide climate change. The increase of production based on renewable energy like wind and photovoltaic power generators implies drastically higher fluctuations in available electricity. The problem has two challenging facets, namely power grid stability and economy of power production/supply. The stability of the distribution grids is a priority concern because reliable distribution is a prerequisite for economic use of power, whether or not renewable. This asks for improved and better coordinated diagnostic and prediction techniques, as well as orchestrated demand-side mechanisms to counter critical grid and/or generation situations.

One of the fastest changing power grids is the German one, owed to substantial increases in wind and solar energy production. This is a consequence of the legal framework enforced by Federal legislation over the last decades. It is characterized by an emphasis on microgenerated power, which enjoys priority in the sense that it must be absorbed by the power grid, unless the grid operates in emergency mode; in that case, however, the wasted power must still be monetarized in the accounting as if it had been fed into the grid. For these reasons, microgenerators of photovoltaic (PV) power have been rolled out massively on the rooftops of residential buildings all over the country. In fact, the growth anticipations for PV microgeneration have in the past years been surpassed by a large margin: 7.5 gigawatts (GW) have been installed in 2011, while the German government had announced a target growth rate of 1.5 GW per year in 2009 (Bundesnetzagentur 2012).

This growth creates problems. About 75% of the PV microgenerators rolled out are non-measured and cannot be remotely controlled. Since 2007, a regulation (EN 50438:2007) was in place that enforced a frequency-based distributed control strategy. A too high frequency is an indicator for overproduction of power. The regulation stipulated that a PV microgenerator must shut off once it locally observes the frequency to overshoot 50.2 Hz. While this was initially meant as a way to stabilize 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 very critical, Europe-wide frequency oscillations. As a consequence, new control strategies have been developed, especially by the VDE, the German Association for Electrical, Electronic & Information Technologies (Bömer and Lehner. 2011).

This paper reports on formal behavioral models of these phenomena, focusing on the problems in the last mile to the end customer. The models are specified in the modeling language MODEST (Bohnenkamp et al. 2006), and simulated with the `modes` simulator (Bogdoll, Hartmanns, and Hermanns 2012). We study the oscillatory effect induced by the previous legislative framework and analyze the currently required approach developed by the VDE. We then present a potpourri of alternative strategies that take up and combine ideas from distributed communication protocol design. Among them, we find the concept of additive-increase/multiplicative-decrease known from TCP, the idea of exponential backoff as used in CSMA variations, 802.11e. We explore the properties of these alternatives with respect to stability, availability, goodput and fairness. It turns out that Internet-inspired mechanisms to break synchrony, especially by using randomization, can be considered as one decisive piece in the puzzle of making the power grid future-proof.

Related work. The modeling of power grids based on behavioral models with strictly formal semantics is gaining momentum. The most closely related work is likely the paper by Chen et al. (2012), who analyze a multi-player game based on a recently proposed distributed demand-side micro grid management approach (Hildmann and Saffre 2011). Efficiency aspects of demand side management approaches have also been studied (Tröschel and Appelrath 2009; Lehnhoff, Krause, and Rehtanz 2011). Other tangible work includes the application of probabilistic hybrid automata with distributed control to the power grid domain (Martins, Platzer, and Leite 2011), and work on network calculus in battery buffered end user homes (Boudec and Tomozei 2012).

2 MODELING

The power distribution grid is hierarchically structured, with a grid of long-distance high-voltage lines (380 kV) as the top layer, down to the leaves which traditionally connect end consumers to the upper layers, using 400 V three-phase current (or 230 V per phase). In this paper, we put a certain focus on these leaves, the “last mile”, which is where the masses of PV microgenerators are installed. The power management of these last miles is a particular challenge because of their sheer number, the lack of measurement and reporting infrastructure, and the volatility of the photovoltaic production. These problems call for a highly automatic, decentralized and flexible grid management—a very challenging and pressing problem. We briefly summarize our modeling choices and abstractions here and refer the interested reader to our separate paper on modeling and abstractions for power micro grids (Hartmanns and Hermanns 2012).

2.1 Abstractions

When modeling the last mile of a grid, consisting of electricity consumers (residential buildings and small businesses), local electricity producers such as PV microgenerators on residential rooftops, and a connection to the upper layer of the grid, we treat two key areas in a very abstract fashion:

Physics. We abstract from the detailed physical characteristics of the grid by looking only at the frequency observed. The target frequency of the German power grid is 50 Hz, with frequencies in the range between 49.8 and 50.2 Hz considered stable. 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 PV generators are switched off and there is no (= zero) influence from the upper layer. We treat the grid as a “copper plate” where adding power has a direct 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. Notably, we could equally well use the observed voltage as a reference quantity for the modeling instead of the frequency in our models since the models are sufficiently

```

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

  measure {= fm = frequency(), c = 0 =};
  when urgent (c >= D) update {= output[id] = fm >= 0.2 ? 0 : MAX =};
  when urgent (c >= M) Generator(id)
}

```

Figure 1: 50.2 Hz on-off controller.

abstract. Since frequency changes are a Europe-wide phenomenon and not restricted to a specific last mile, we also exaggerate the influence of each single PV generator. We return to these aspects in Section 5.

Consumption and upper layer in-feed. The influence from the upper layer power grid on our last mile as well as the local consumption is modeled as an abstract (randomized or deterministic) load profile.

2.2 Modest Models

Our models, specified in the modeling language MODEST, consist of a set of $G = 32$ generator processes modeling the PV generators plus a single process representing the abstract background load. 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 updates its output accordingly. Each of these measure-update cycles takes M time units, with $D \leq M$ time units passing between the measurement and the change of 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. For illustration, the declaration of the `Generator` process for the currently deployed control scheme that switches the generator off when 50.2 Hz are exceeded is shown in Figure 1. For more MODEST code, we again refer to Hartmanns and Hermanns (2012).

3 DECENTRALIZED STABILIZATION STRATEGIES

We have modeled the following control strategies for the generator processes:

On-off: The scheme that switches on or off when the 50.2 Hz threshold is crossed, as shown in Figure 1.

Linear: The current requirements (which are part of VDE-AR-N 4105) allow generators to increase their output by up to 10 % of the maximum possible output per minute when the frequency is below 50.2 Hz. When 50.2 Hz are exceeded, the current output is saved as p_m and the generator output is reduced with a 40 % gradient per Hertz, i.e. $output(f) = p_m - 0.4 \cdot p_m \cdot (f - 50.2)$, up to 51.5 Hz. Beyond that, an emergency switch-off takes place.

AIMD: The first new algorithm that we study is inspired by the way the Internet’s Transmission Control Protocol (TCP) achieves fair usage of limited bandwidth through an *additive-increase, multiplicative-decrease* (AIMD) policy: Generator output is increased in constant steps below 50.2 Hz, and if the frequency is measured above 50.2 Hz, the output is scaled down by a constant *factor*.

Probabilistic on-off: After each measurement, our first (very simple) randomized controller proceeds like the on-off controller with probability 0.95 or switches off with probability 0.05 to introduce some disturbance to avoid oscillations.

Dynamic die: We next use a p -sided die, $p \in \mathbb{N}$, instead of fixed probabilities to decide between proceeding like the on-off controller or unconditionally switching off, the latter corresponding to exactly one side of the die. p increases exponentially when the frequency is low, making it more probable for the generator to switch on, and conversely decreases exponentially when the frequency

is above the 50.2 Hz threshold, thus increasing the probability of the generator staying off for some time even when the frequency drops below the threshold again.

Frequency-dependent: Instead of influencing the probability of switching on or off via a stepwise increase or decrease of die size, the frequency-dependent controller makes the probability a function of the currently observed frequency. The probability function has been chosen such that the probability of switching on depends linearly on frequency, being 1 at 50.0 Hz, 0.5 at 50.2 Hz and 0 at 50.4 Hz.

Exponential backoff: The previous three policies all decided probabilistically whether to change the power output or not. In a departure from this, our next controller will unconditionally switch to full power when the current frequency allows and switch off when 50.2 Hz are exceeded, but then wait a probabilistically chosen amount of time before measuring and potentially switching on again. The precise scheme that we use here is *exponential backoff* with *collision detection*, which is commonly employed in CSMA/CD-based medium access protocols such as Ethernet: When one device connected to the shared medium 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 $\{1, \dots, 2^{bc}\}$, where *bc*, the *backoff counter*, keeps track of the number of collisions and of the number times that the channel was sensed as occupied when this message should have been sent. The goal is to use randomization to prevent two devices from perpetually choosing the same delay and thus always colliding.

Frequency + backoff: Our last controller combines the randomization of switching decisions from the frequency-dependent controller with probabilistic waiting times according to an exponential backoff scheme.

4 EVALUATION

We use the `modes` discrete-event simulator for MODEST (which is part of the MODEST TOOLSET, available for download at www.modestchecker.net) to simulate our last mile model with the different control strategies and evaluate the following properties:

Stability: We study how prone the algorithms are to the delay incurred prior to their reactions taking effect. The ability to keep the system in the safe state where frequency is below 50.2 Hz with a minimum of oscillation can be termed as *stability*, i.e. reflecting the absence of frequency oscillations.

Availability: We also look at *availability*, that is the overall fraction of time that the system spends in the safe state below the threshold of 50.2 Hz. We relate this quantity to the goodput, which is the average amount of power a generator can add to the system while keeping it safe.

Fairness: Another important aspect is *fairness*: When the system state does not allow all generators to operate at full power, will each generator be able to provide an equal share of the allowed power generation, or will any of the control schemes continually put certain generators at a disadvantage?

4.1 Stability

Method. In order to find out how prone the different algorithms are to frequency oscillations, we perform two simulation runs per algorithm: one for the case where there is no delay between a generator's frequency measurement and the time at which the effect of its modification of power output can be observed over the entire branch of the grid ($D = 0$), and one for the case where this delay amounts to nine tenths of an entire measurement cycle ($D = 9$, $M = 10$). We then plot and compare the overall system frequency over time. To allow a proper comparison, we use a deterministic background load process. The initial power generation for each generator is randomly determined.

Findings. The simulation plots are shown in figures 2 to 9, with frequency values on the y-axis and simulation time on the x-axis. The safe area below 50.2 Hz is highlighted. The system frequency is the




Figure 2: Behaviour of the 50.2 Hz on-off controller.




Figure 3: Behaviour of the linear controller according to VDE-AR-N 4105.




Figure 4: Behaviour of the AIMD controller.




Figure 5: Behaviour of the probabilistic on-off controller.




Figure 6: Behaviour of the probabilistic on-off controller with dynamic die.




Figure 7: Behaviour of the frequency-dependent probabilistic on-off controller.




Figure 8: Behaviour of the on-off controller with exponential backoff.




Figure 9: Behaviour of the frequency-dependent controller with exponential backoff.

upper curve, while the lower curve is a plot of the background load for reference. The left-hand plots are for the immediate ($D = 0$) case while the right-hand plots show the behavior in the delayed case ($D = 9$).

We first see (Figure 2) that the on-off controller indeed produces extreme frequency oscillations, at least in the delayed setting. Its behavior in the immediate setting, the details of which are not visible due to the scale of the graph, is also very predictable (modulo changes in background load): In a system with G generators, it enters a cycle of length $G + 1$ that starts when a sufficient number of generators is off such that the system frequency is just below 50.2 Hz. The next generator to act that is currently off will then switch on, pushing the frequency above 50.2 Hz and thereby forcing the next generator that is on to switch off. The set of generators that are off thus moves through the total set of generators in a round-robin fashion; assuming that K out of the G generators can be switched on simultaneously without exceeding 50.2 Hz, the frequency will be below that threshold for a fraction $\frac{K+1}{G+1}$ of the total time, giving a simple formula for system availability (for a very simple setting, though).

The linear controller according to VDE-AR-N 4105 is a clear improvement (Figure 3): It exhibits no oscillations or frequency jumps at all, instead mostly following the background load. It completely fails at keeping the frequency in the safe area below 50.2 Hz, however. Such a direct frequency control was most probably not the intention of this controller's designers, though. In fact, if one treats the background load as the cause for the unsafety in this case, one can argue that, if the background load is to a large extent controllable (as is the case when most power is conventionally generated), this will suffice to keep the frequency in the desired range because this controller will *not interfere* with such stabilization attempts.


The AIMD controller is essentially the on-off controller improved to proceed in smaller steps, and this shows in the frequency plots (Figure 4), which show a highly dampened version of the on-off controller's behavior. Indeed, the interesting point about the AIMD controller is not so much its improved stability but the question whether AIMD brings the same fairness into distributed power generation that it brings to TCP in the computer networks setting.

We now come to the family of probabilistic on-off controllers employing various degrees of randomization based on different approaches. Our first observation is that, from the stability perspective, neither the probabilistic on-off controller, nor its variant with dynamic die size, nor the exponential backoff-based controller offer any advantage over the simple deterministic approach (figures 5, 6 and 8). It is particularly surprising that both approaches to randomization (random switching and randomized delays) fail to improve stability (except for a very, very slight reduction in the amplitude of the oscillations when using exponential backoff in the delayed setting). For the randomized delay controllers, this appears to be due to the decisions in the delayed case being effectively 90% synchronous, which means that almost all generators will decide to switch on as soon as the system has become safe for low backoff counter values, which leads to the backoff counter being reset frequently and never even reaching the higher values needed to break synchronicity.

The frequency plots obtained from the two frequency-based on-off controllers, without (Figure 7) and with exponential backoff (Figure 9), tell a very different story: In both cases, the behavior of the system is very different from the other on-off controllers, with the frequency changes in the delayed setting being much closer in magnitude to the AIMD controller. This indicates that it is indeed possible to achieve the effect of fine-grained deterministic control using a suitable number of probabilistic controllers employing an adequate randomization scheme instead. Evaluating the addition of exponential backoff in this case is somewhat difficult: We see that the system appears to remain in the safe area for longer amounts of time when exponential backoff is used, but this effect might potentially be achieved by skewing the probabilities to favor switching off in the non-backoff controller as well. Still, it appears that the controller using exponential backoff copes better with the increase of background load in the middle of the simulation run, so we tend to consider this a beneficial addition in terms of improving stability.

4.2 Availability and Goodput

Method. We obtain results for the system availability by including a property in the model framework that evaluates to the fraction of time that the frequency was below 50.2 Hz in a simulation run. We perform

Figure 10: Availability vs. total goodput for $D = 0$ (left) and $D = 9$ (right).

1000 simulation runs in order to average out the availability numbers over the probabilistic decisions, which in this case also come from using a random background load. Instead of merely showing the availability numbers on their own, we chose to also compute the *goodput* of the generators, i.e. the amount of “useful” power generated, with power being of no use when the system is unavailable. We then show, in Figure 10, availability compared to total goodput for all generators per time unit for the different controllers.


Findings. In terms of availability, four algorithms behave very similarly: The on-off controller, the probabilistic on-off controller, its variant with dynamic die size, and the basic frequency-dependent probabilistic on-off controller. The first three already exhibited almost identical behavior in the plots used to evaluate stability, so their being so close in terms of availability comes as no surprise. The two interesting points about this cluster of algorithms are that 1) the frequency-dependent controller with its vastly different behavior is part of it—however, in the delayed case, it makes this difference shown by providing notably better goodput, and 2) that the on-off controller with exponential backoff is not part of it, showing similar goodput, but increased availability—this being where the reduction in oscillation amplitude compared to the other severely oscillating controllers, as observed in Fig. 8, probably comes in.

The linear controller according to VDE-AR-N 4105 is a clear outlier in the negative sense for availability; as observed before, however, its goal is probably not to actively stabilize the system but to avoid introducing additional instability. The two positive outliers are the AIMD controller and the frequency-dependent controller with exponential backoff. Both manage to improve availability and goodput, with the AIMD controller favoring goodput and the frequency-dependent exponential backoff controller providing higher availability at slightly lower goodput, i.e. more severe but shorter spikes into the unsafe area above 50.2 Hz.

4.3 Fairness

Method. Another aspect worth studying is fairness: we want to find out whether some of the controllers allow certain generators to produce significantly more power than others. Since all generators are identical, any such difference will be due to the random initial power generation. As such, it does not make sense to compute averages over several runs, and we perform one very long simulation run instead, using a scaled version of the usual deterministic load profile. The range of values that we obtain is shown in Figure 11, with the lower (red) area being below the mean output value, the upper (green) area being above the mean, and the black line indicating one standard deviation from the mean in both directions.

Findings. The most unfair controller clearly is the linear one. This is not unexpected: It does not include any radical changes of output, instead trying to progress in small steps. A generator that initially has a high output will thus more or less remain at high output, while a generator that starts low or off does not get a chance to obtain a significant increase in participation. The main point of interest concerning fairness is whether the AIMD approach will result in a fair sharing of the available production possibility in the

Figure 11: Generator output ranges for $D = 0$ (left) and $D = 9$ (right).

same way that it results in a fair sharing of bandwidth when used in TCP. Surprisingly, while it is indeed better than its fairness-ignoring cousin with additive decrease (AIAD), both are still comparatively unfair relative to most other controllers in the delayed setting. This can again be explained by the incremental nature of AIMD, which makes it similar to the linear controller, compared to the very randomized on-off behavior of the remaining controllers, which show a high degree of fairness both for $D = 0$ and for $D = 9$, with the frequency-dependent controller with exponential backoff being a little worse than the others.

4.4 Scaling the Model

One final point of interest is whether the controllers are sensitive to the particular simulation setting of $G = 32$ generators with a total contribution to frequency of at most 0.3 Hz. We thus performed the simulation analysis for stability and availability/goodput for $G \in \{2, 4, 8, 16, 64\}$ as well, in two different settings:

Scaled setting: The contribution to frequency of each generator is scaled such that the total contribution of all generators is the same as for $G = 32$. This will allow us to find out how the number of generators affects the controllers' behavior.

Unscaled setting: The contribution per generator is constant (namely $\frac{0.3 \text{ Hz}}{32} = 0.009375 \text{ Hz}$, as in all models studied in the previous sections), leading to a lower (for $G < 32$) or higher (for $G > 32$) total contribution of all generators. In particular, with a maximum background generation of 0.1 Hz, the system cannot become unsafe for up to 10 generators, but it can reach up to 50.7 Hz for $G = 64$.

4.4.1 Stability

We found that the number of generators itself does not change the system behavior or the performance of the different controllers in terms of stability in a significant way (scaled setting). As an example, consider the behavior of the frequency-dependent controller with exponential backoff: Its behavior in the scaled setting for $G = 16$ and $G = 64$ is shown in Figure 12. Compared to the right-hand plot in Figure 9, the only significant difference is the step size of the changes whenever a generator is switched on or off.

However, once we increase the total frequency contribution (unscaled setting with $G = 64$), the two frequency-dependent controllers start behaving differently. The plot on the left-hand side of Figure 13 shows the behavior of the controller *without* exponential backoff: It quickly starts to induce severe oscillations. This happens because the probability function used by the controller assigns probability 1 to switching




Figure 12: Frequency-dependent controller with backoff for $G = 16$ (left) and 64 , scaled setting, $D = 9$.




Figure 13: Frequency-dependent controller without (left) and with backoff, unscaled setting, $G = 64$, $D = 9$.

off for frequencies beyond 50.4 Hz, which are more likely to be reached at some point when the total contribution of all generators is as high as 0.6 Hz. Once this happens, almost all generators switch off due to the measurement delay of $D = 9$, only to switch on again with a very high probability in the next “round”. The frequency-dependent controller *with* exponential backoff is the only probabilistic one that still works in this setting, albeit not leading to as stable a system as for $G = 32$. Still, as shown on the right-hand side of Figure 13, it manages to reduce oscillations over time. Ongoing experiments using different and non-linear functions to compute the switch-off probability depending on the frequency indicate that these scaling problems can be avoided by using a “better” probability function.

4.4.2 Availability and Goodput

When we look at availability and goodput, we mainly see these findings confirmed: Figure 14 shows the evolution of availability and goodput for changing values of G for the three most interesting controllers. This time, we show the average goodput *per generator* and time unit. In the scaled setting (left-hand side), this implies that a controller that is not affected by changes of G will have constant availability, and goodput inversely proportional to G . We clearly see that the on-off and AIMD controllers show this kind of behavior. The frequency-dependent controller with exponential backoff, however, performs better as G increases: Availability improves, and for twice the number of generators, the goodput is actually larger than half the previous value. This controller’s randomized decisions thus manage to create an averaging effect for a higher number of participants and thereby successfully exploit the potential benefits of randomization.

In the unscaled setting (right-hand side of Figure 14), the results are more difficult to interpret. We first observe that it needs at least 11 generators to reach a system frequency above 50.2 Hz under full background load (22 generators with no background load). For an ideal controller, the goodput per generator should thus be the same for all $G < 11$; after that, an increasing number of generators has to “share” the safe area below 50.2 Hz, thus goodput should again be inversely proportional to G . This appears to be the case for the on-off and AIMD controllers, though both fail to keep availability constant and thus independent

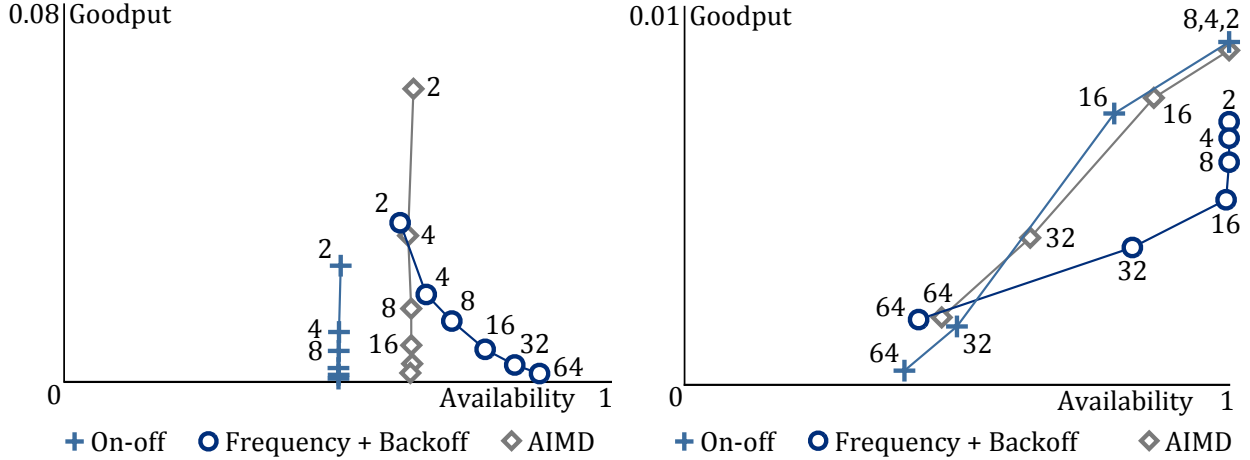


Figure 14: Availability vs. goodput per generator for $D = 9$ in the scaled (left) and unscaled (right) setting.

of the number of generators. The frequency-dependent controller with exponential backoff again behaves differently: While it starts with lower goodput and does not keep this independent of G , it performs better in terms of availability compared to the on-off and AIMD controllers. As observed in Section 4.4.1, it finally starts to break down when $G = 64$. Still, this shows that combining probabilistic switching with randomized delays can indeed lead to a more robust controller; we expect the use of a better probability function and fine-tuning of the backoff procedure to have the potential for significant improvement of the frequency-dependent controllers with and without exponential backoff.

5 CONCLUSION

This paper has introduced a potpourri of decentralized stabilization strategies, and has presented a simulation-based evaluation of their respective properties. The discussion has focussed on the frequency as a single indicator of the grid state, and has assumed a linear impact of PV production on it. The assumption of linearity is an indirect effect of the physical realities. We could also have used the observed voltage as a reference quantity for the modeling instead of the frequency. For low voltage rotating current, the allowed voltage interval is 440 V to 360 V. In this interpretation, the linearity assumption would directly hold, and the analysis results can therefore be transferred to this interpretation right away. It is worth noting that frequency drifts usually affect the entire European grid, and not only a specific last mile. This also means that the influence of a single last mile on the frequency is in fact negligible, and appears amplified in our studies. In contrast, voltage changes are local phenomena, entangled with phase drifts in the last mile. The influence of a single microgenerator on the observed voltage in a last mile is therefore much more substantial, but may actually be skewed by spatial properties.

At the same time, the discussion has focussed on grid stability, not grid economy. The stability of the distribution grids is a priority concern, because reliable distribution is a prerequisite for economic use of energy, whether renewable or not. Of course, the same basic control algorithms can use intraday and spot market prices in addition to frequency or voltage as indicators for the grid state in times when the grid is operating well (which will hopefully be dominating time-wise anyway). These indicators are nowadays easily accessible in residential areas over the Internet (provided there is power to run the residential Internet connection, a grid stability problem) and can be used for decentralized demand-response management. A massive roll-out of such appliances may however make the power grid fall into a similar trap as the current German on-off controllers do, cf. Figure 2. This is because automatic decentralized decisions orchestrated by a central signal may lead to oscillations.

Probably the most tangible contribution of our work is the idea that concepts developed for distributed communication protocols may be attractive candidates for components of future power grid stabilization

strategies. We have emphasized that there is an urgent need for better stabilization strategies in light of the strong growth of PV microgeneration especially in Germany. Decentralized strategies that base their decisions only on the information available are promising, not only because their roll-out does not require new investments in communication infrastructure. Another very attractive feature of the absence of centrally collected information is that privacy concerns are much easier to account for. The Internet-inspired mechanisms to break synchrony, especially by using randomization, appear—at least to us—as one decisive piece in the puzzle of making the power grid future-proof.

ACKNOWLEDGMENTS

The authors are grateful to Holger Wiechmann (EnBW AG), Sebastian Lehnhoff (Universität Oldenburg and OFFIS Energie) and Martin Ney (Luxea GmbH) for inspiring discussions, and to Jonathan Bogdoll and Ernst Moritz Hahn (Saarland University) for their support in preparing the technical report on which this publication is based (Berrang et al. 2012). 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 Seventh Framework Programme under grant agreement no. 295261 (MEALS).

REFERENCES

- Berrang, P., J. Bogdoll, E. M. Hahn, A. Hartmanns, and H. Hermanns. 2012. “Dependability Results for Power Grids with Decentralized Stabilization Strategies”. Reports of SFB/TR 14 AVACS 83. ISSN: 1860-9821, <http://www.avacs.org>.
- Bogdoll, J., A. Hartmanns, and H. Hermanns. 2012. “Simulation and Statistical Model Checking for Modestly Nondeterministic Models”. In *MMB/DFT*, edited by J. B. Schmitt, Volume 7201 of *LNCS*, 249–252: Springer.
- Bohnenkamp, H. C., P. R. D’Argenio, H. Hermanns, and J.-P. Katoen. 2006. “MoDeST: A Compositional Modeling Formalism for Hard and Softly Timed Systems”. *IEEE TSE* 32 (10): 812–830.
- Bömer, J., K. Burges, P. Zolotarev, and J. Lehner. 2011. “Auswirkungen eines hohen Anteils dezentraler Erzeugungsanlagen auf die Netzstabilität bei Überfrequenz & Entwicklung von Lösungsvorschlägen zu deren Überwindung”. Study commissioned by EnBW Transportnetze AG, Bundesverband Solarwirtschaft e.V. and Forum Netztechnik/Netzbetrieb im VDE e.V.
- Boudec, J.-Y. L., and D.-C. Tomozei. 2012. “A Demand-Response Calculus with Perfect Batteries”. In *MMB/DFT*, edited by J. B. Schmitt, Volume 7201 of *LNCS*, 273–287: Springer.
- Bundesnetzagentur 2012. “EEG-Vergütungssätze für Photovoltaikanlagen”. <http://www.modestchecker.net/Link.aspx?id=BNPV2012>. Accessed March 21, 2012.
- Chen, T., V. Forejt, M. Z. Kwiatkowska, D. Parker, and A. Simaitis. 2012. “Automatic Verification of Competitive Stochastic Systems”. In *TACAS*, edited by C. Flanagan and B. König, Volume 7214 of *LNCS*, 315–330: Springer.
- Hartmanns, A., and H. Hermanns. 2012. “Modelling and Decentralised Runtime Control of Self-stabilising Power Micro Grids”. In *ISoLA*, *LNCS*: Springer.
- Hildmann, H., and F. Saffre. 2011. “Influence of variable supply and load flexibility on Demand-Side Management”. In *EEM*, edited by M. Delimar, 63–68: IEEE Conference Publications.
- Lehnhoff, S., O. Krause, and C. Rehtanz. 2011. “Dezentrales autonomes Energiemanagement (Distributed Autonomous Power Management)”. *Automatisierungstechnik* 59 (3): 167–179.
- Martins, J., A. Platzer, and J. Leite. 2011. “Statistical Model Checking for Distributed Probabilistic-Control Hybrid Automata with Smart Grid Applications”. In *ICFEM*, edited by S. Qin and Z. Qiu, Volume 6991 of *LNCS*, 131–146: Springer.
- Tröschel, M., and H.-J. Appelrath. 2009. “Towards Reactive Scheduling for Large-Scale Virtual Power Plants”. In *MATES*, edited by L. Braubach, W. van der Hoek, P. Petta, and A. Pokahr, Volume 5774 of *LNCS*, 141–152: Springer.

AUTHOR BIOGRAPHIES

ARND HARTMANNS studied computer science at Saarland University, Germany, and Linköping University, Sweden, receiving the B.Sc. (2007) and M.Sc. degrees (with honors, 2009) from Saarland University. Currently, he is working as a Ph.D. student at the chair for Dependable Systems and Software at Saarland University. His research interests are modeling tools and formalisms for stochastic timed systems, in particular the MODEST language, and applications thereof. He has been working with MODEST since 2006, and initiated the development of the current MODEST TOOLSET in 2008. His e-mail is arnd@cs.uni-saarland.de.

HOLGER HERMANNS is a full professor in the Department of Computer Science at Saarland University, Germany, holding the chair for Dependable Systems and Software since 2003. He studied at the University of Bordeaux, France, and the University of Erlangen/Nürnberg, Germany, where he received the diploma (with honors) in Computer Science in 1993 and the Doctoral degree (with honors) from the Department of Computer Science in 1998. From 1998 to 2006, he has been with the University of Twente. His research interests include modeling and verification of concurrent systems, resource-aware embedded systems, and compositional performance and dependability evaluation, including dependable energy distribution grids. His e-mail is hermanns@cs.uni-saarland.de.

PASCAL BERRANG is a Bachelor student in computer science at Saarland University since October 2010. He is interested in formal modeling and evaluation of distributed power management systems. His e-mail is berrang@stud.uni-saarland.de.