WIDE-AREA CONTROL SYSTEM FOR BALANCE-ENERGY PROVISION BY ENERGY CONSUMERS

Friederich Kupzog, Peter Palensky

Institute of Computer Technology
Vienna University of Technology
Gusshausstrasse 27-29
1040 Vienna, Austria
+43 1 58801 38411
{kupzog, palensky}@ict.tuwien.ac.at

Abstract: The modern electricity system, changing from the traditional hierarchical to a more and more peer-to-peer oriented structure, needs modern IT and control solutions. While up to now consumers were considered passive players, the grid of the not-so-far future sees them as active grid partners. As a new generation of automated demand response, consumers can react on real-time prices, on grid parameters like frequency or on transport schedules. The critical part of this topic is to find easy, consistent and robust models of the involved processes and the respective control infrastructure. This paper describes a distributed control system that reacts on frequency changes and models consumers as a community of energy storages.

Keywords: Energy management systems, distributed control, decentralized control systems, algorithms, models.

1. INTRODUCTION

One of the largest existing kinds of control networks is the energy transmission and distribution system. Like the Internet, these systems have grown from single, isolated grids to continent-spanning networks. Although their primary task is energy transmission, information exchange is an essential aspect for the operation of the energy system. In the early times of energy grid operations, the problem of transmitting time-critical information was solved elegantly by utilising the network frequency for broadcasting the current energy balance state of the grid. Generation and consumption of energy in the grid has to be balanced at all times, otherwise voltages exceed lower or upper limits, resulting in damage at generation or customer equipment. All generators and loads of the grid participate in a very large and complex control system whose task is to maintain the energy balance by keeping the area control error (ACE) low. Short-term changes in the load situation (within a few seconds) have to be compensated by the generators in real-time (primary control). Only if longer mid-term or long-term discrepancies occur (lasting for minutes or more), secondary or tertiary reserves are activated, that need a considerable activation time until they are fully active. While remote loads are causing the ACE to change, this change has to be detected at the generator's site. Conventional generators have rotating masses, and the rotation speed (frequency) depends on the generator load. Increasing load results in decreasing frequency and vice versa. Thus, the presence of a new remote load in the grid is communicated to the generators by changed impedance resulting in a slightly lower system frequency. Since the system frequency is the same at any place of the connected grid, the deviation from 50 Hz (60 Hz in the U.S.) is a direct and ubiquitous indicator of the ACE. The system frequency as inherent communication channel is used very successfully for communicating the time-critical ACE information in the grid. Less time critical information is traditionally distributed by voice communication, although this approach is more and more replaced by automation networks.

In this paper, a case study for the utilisation of this implicit communication channel is carried out for a special application. The deregulation of European power is accompanied by the preparation of a market for primary control power. At this special energy market, not only energy generators but also energy consumers will be able to offer their ability to modulate their power consumption according to the needs of ACE control. In order to be able to make substantial offers, a large number of widely distributed loads have to be put together and be coordinated. The need for coordination of many small units increases the cost for primary control power provision on the demand side. An always-on real-time communication infrastructure would be far too expensive for this application. Consequently, it is proposed to split the communication requirements in two areas: time-critical coordination realised using the system frequency as implicit communication channel and non-critical administrative communication using available infrastructures such as mobile phone networks or wireless Internet.

2. RELATED WORK

Shedding of electrical loads for the purpose of costs saving and also for electricity grid stabilisation is state of the art. While peak load restriction is a classical application of automation networks on industrial sites where power information of loads in a restricted area is processed, the system frequency is used as one locally measured indicator for the overall grid situation in protection systems that switch off loads in critical grid situations (see e.g. Shokooh et al, 2005). However, this approach deals only with extra-ordinary grid situations where the energy provision to some consumers can be sacrificed (that is, switched off) in order to prevent larger blackouts.

In contrast to that, the approach presented in this paper is intended for normal operating conditions. It originates from the context of the “IRON Project” (Kupzog, 2006), which aims to increase the efficiency of the power system by means of modern information and computer technology with a focus on the demand side. The idea is to utilise so far unused distributed energy storage capacities in inert processes, e.g. heating or cooling applications or material transport. This utilisation is conducted in such a way that the “user comfort” of the energy customer is not negatively influenced. Numerous benefits lye in the utilisation of distributed storages, depending on what that storage system is used for. Here, the application is primary control. The main hindrance for the large-scale realisation of systems incorporating distributed loads is that a very large number of spatially distributed nodes need to be reached by a robust and low-cost control network. Therefore, implicit communication using the power system frequency is considered.

Another approach that focuses on the grid frequency is followed as part of the large U.S. “GridWise” initiative, where research in the area of intelligent loads is also conducted (Trudnowsky et al., 2006). The corresponding GridWise project concentrates on frequency deviation measurement and reaction, but communication here is restricted to the system frequency channel only, resulting in restricted control possibilities. This paper proposes the use of a dual communication strategy in order to increase the manageability of the system. Time critical communication is performed using inherent communication over the power system frequency, and other information is exchanged using existing, low cost and potentially low bandwidth communication infrastructures with low dependability.

3. COMMUNICATION MODELLING

Distributed control systems are typically relying on an explicit communication channel with some given quality of services (QoS). For network based control, hard real time requirements are usually not avoidable, so a general purpose network with best effort transport (as the plain Internet must be considered) can not be used.

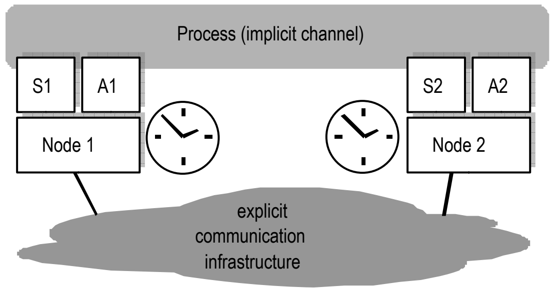

Fig. 3.1: Sensor and actuator nodes of a distributed system can get (coordination and control) information via explicit and implicit communication channels.

One kind of communication channel that is often forgotten is the process itself (Fig. 3.1). Depending on the type of process, important coordination information can be gained out of measurement values. In our case the channel “grid frequency” as even a perfectly simple channel model since it is instantaneously transported. Another source of coordination trigger is time, supposing the local clocks are synchronized.
The control decisions to be done in this system (resource management) depend on process parameters (indicators of the resource) and some rules that ensure fairness, economic operation or safety. The necessary information and the chosen communication channels are therefore:

Information	*Communication channel*
Time	Internet, grid (frequency)
Grid frequency	Grid
Priorities, Schedules	Internet

With this hybrid communication infrastructure it is possible to achieve a real-time-like large area distributed system, without investing in expensive infrastructure.

4. RESOURCE MODELLING

Provided adequate process interfaces (e.g. access to temperature set-points) and communication channels (for coordinating load shifts) are in place, the capability of energy storage in a large group of electrical loads can be utilised without interfering with the safe and uninterrupted operation of these loads. In order to operate such a system consisting of thousands of distributed storages (resources), a consistent model is needed that describes the behaviour or the individual process precisely enough but at the same time is not too complex so that calculations on many models can be processed efficiently. The resource model used here is described in detail in Kupzog and Roesener, 2007 and will subsequently be summarised.

It is assumed that every electrical consumption process can be described as the superposition of potentially volatile power consumption and an energy storage that can be charged and discharged. This energy storage can e.g. be the thermal capacity of a room. Only those loads with a sufficient large amount of storage capability can be utilised in this context (e.g. air conditioning systems). The volatile power consumption originates from the attempts to keep the energy level of the storage (e.g. room temperature) constant. It is the continuous replenishment of energy that is lost due to non-ideal isolation. Although the process cannot feed energy back into the grid and therefore is not a real electrical energy storage, the volatile consumption profile can be estimated for a large number of processes. If this is subtracted from the total load of all considered processes, then positive and also negative power components remain, representing charge and discharge of the storage.

5. CONTROL ALGORITHM

The provision of primary control power is determined by the system frequency f, which can be measured at every point of the system. The system frequency therefore acts as an implicit, unidirectional, real-time and highly reliable communication channel. A system of electrical loads that takes part in primary control has to adhere to a given power characteristic depending on the frequency. The relationship of the power that has to be provided and the network frequency f can be expressed in a function $K(f)$, subsequently referred to as power function. A basic system that fulfils the primary control objective is shown in Fig 5.1. The load or generation of s single unit connected to the power system is modulated with $K(f)$. The shape of $K(F)$ is given e.g. by UCTE, 2004. The term 'modulation' is used to point out the fact that the load (or generator) behaviour is only changed *relatively* to its usual behaviour without primary control activities. The delta character (Δ) will be added to variables to reflect this (see e.g. ΔP in Fig 5.1).

The speciality of the algorithm that shall be discussed here is that is not operating on a single unit but rather coordinating a large number of potentially very small contributors to primary control. However, before dealing with the aspects of distributed contributors, the power function $K(f)$ shall be discussed more in detail since it is the main objective of the system to realise this function as precise as possible.

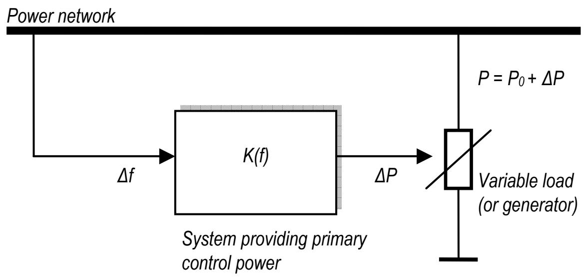

Fig 5.1. Block diagram of a single system providing primary control power. The power dissipation or generation is modulated around a working point P_0 which reflects the normal operation without primary control.

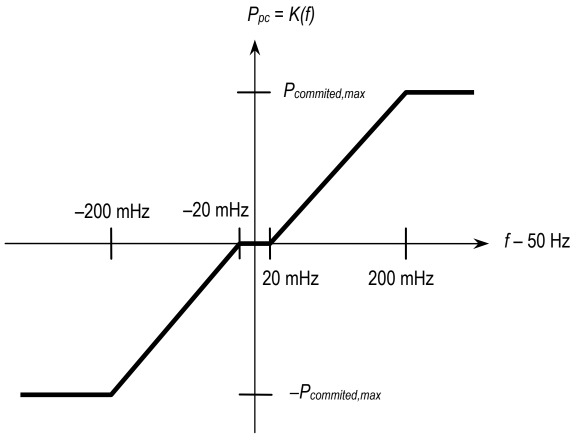

Fig. 5.2. Primary control power as function of the network frequency $K(f)$. Outside the dead band of +/- 20 mHz around the nominal value of 50 Hz the provided control power grows linearly with the frequency deviation up to the maximum of +/- 200 mHz.

As stated in the ground rules of the Union for the Coordination of Transmission of Electricity (UCTE, 2004), the provision of primary control power is linearly increasing with the system frequency deviation from 50 Hz. In order to avoid too frequent control activity, a dead band of +/- 20 mHz has been introduced (UCTE, 2004). As long as the system frequency is within this dead band, no control activities are issued. Only when the frequency exceeds the dead band, primary control power is supplied. A maximum frequency deviation of +/- 200 mHz is allowed; from here on, emergency measures are taken that are not covered by the regulations for primary control. Therefore, the supplied control power is constant outside the +/- 200 mHz band. The resulting frequency-power function is shown in Fig. 5.2. It can be seen that the power function is piecewise linear, monotonously increasing and bounded. These three properties will be used by the algorithm. In fact, the algorithm can work with any function $K(f)$ that has these properties. $K(f)$ determines the power for primary control P_{pc} of a system taking part in the primary control measure. Such a system is characterised by the maximum power it can commit ($P_{committed,max}$, see Fig. 5.2) which itself is an important factor that is reflected in the scaling of $K(f)$.

The differential power ΔP consumed by a system of resources, which is defined relatively to the normal resource operation power, is the sum of all individual differential power amplitudes Δp_i. The system is linear in this regard, not depending on the model of the individual resource. Now, an algorithm is searched that influences the individual resources in such a way that their total differential power ΔP equals (to a certain extend) the required power P_{pc}. These relationships of ΔP, Δp_i, P_{pc} and $K(f)$ are shown in (1).

$$\Delta P = \sum_{\forall i} \Delta p_i \rightarrow P_{pc} = K(f) \tag{1}$$

So far, all considerations were independent of the resource model. Now, in order to be able to find and expression for the individual power Δp_i, the energy storage model discussed in Section 3 will be facilitated. For the design of the distributed control algorithm a linear modelling is of great advantage since this allows the definition of some system-wide sum variables (such as ΔP), that can easily be broken down to individual resources. Two basic possibilities of influencing loads exist. It was found that the fist option, simple (because direct) control of the process power p_i by switching the load off and on is not the preferable solution since it disregards internal process knowledge and has only short-term effects. The more elegant and better solution is to gain influence on process variables that only influence the power indirectly. For thermal applications, this was the process temperature set-point $T_{set,i}$. However, it is possible to generalise this using the energy storage model as outlined in the previous section. Each storage i (either conceptual or real) has an internal energy state s_i that indicates the energy currently stored in the storage. The power consumption of the resource can be described as the power that is needed for compensating the losses of the process, superposed by the power consumed (or released) by charging/discharging the internal energy storage. This is depicted in Fig. 5.3. When the storage energy level is increased, the resource consumes additional power (compared to the normal or average operation). When the energy level is decreased, the resource consumes less power. Additionally, in most cases the charged storage has higher losses than the discharged.

This basic behaviour of a resource can be summarised as in (2) and (3). In (3), the average power is omitted and only the differential power is described:

$$p_i = \frac{ds_i}{dt} + b_i s_i + c_i \tag{2}$$

$$\Delta p_i = p_i - c_i = \frac{ds_i}{dt} + b_i s_i \tag{3}$$

where Δp_i is the differential power consumed by resource i, s_i is the energy stored in resource i and b_i, c_i are a process-specific parameters. For an ideal storage without losses, $b_i = 0$.

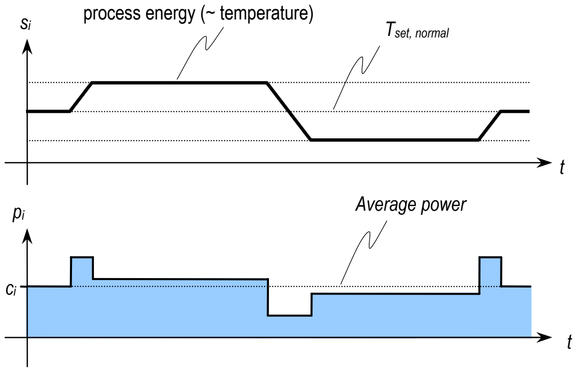

Fig. 5.3: Energy level s_i and consumed power p_i of a DSM resource i according to (2). When the storage energy level is increased, the resource consumes additional power. When the energy level is decreased, the resource consumes less power.

s_i is the process parameter that can be externally controlled. The maximum amount of energy that can be stored in the resource is determined by the value range of s_i. The value ranges of s_i as well as the deviation of s_i are bounded within certain limits, which are process-specific.

5.1. Discussion of the simplified approach with $b_i = 0$

In the following, it will first be assumed that all $b_i = 0$, i.e. all storage processes in the distributed resources are ideal and no additional losses are caused. This simplifies the problem and allows describing the general idea of the approach used here. Later, this simplification will be dropped and the case $b_i > 0$ will be discussed.

Since the DSM resource model is linear, a system-wide storage energy S can easily be defined:

$$S = \sum_{\forall i} s_i \tag{4}$$

This definition is motivated by the search for a relationship between the system frequency f and the individual process energy s_i. For $b_i = 0$, (2), (3) and (4) can easily be combined to (5):

$$\Delta P = \frac{dS}{dt} = \sum_{\forall i} \frac{ds_i}{dt} \overset{!}{=} K(f) \tag{5}$$

By integration of (5), the following is gained for the required system-wise energy state S:

$$S = \int K(f)dt \tag{6}$$

The objective for a collective of distributed storages taking part in primary control is to adhere to the power function $K(f)$. The individual resource has a restricted information horizon. It is not aware of how much power the other resources consume, or what their energy states are. It only can control its own energy state s_i, therefore also its own power consumption, and it can measure the network frequency *f*. Now, (6) gives a simple calculation directive how to calculate the system-wide energy state *S* only from the system frequency. So, *S* is actually known at all places and the remaining question is how to determine the individual storage state s_i from *S* at the place of resource *i* without any knowledge about all other storage states s_j.

Two basic solutions can be applied here. The first, which shall be called the "continuous approach", simply scales down *S* to the individual s_i, so that

$$s_i = k_i \cdot S \text{ with } \sum_{\forall i} k_i = 1 \quad (7)$$

While it is very simple, the drawback of this solution is that s_i continuously follows *S*, resulting in a very fine-grain variation of the resource's energy state. Any small error (e.g. due to quantisation) is multiplied with the potentially very large number of resources and can have major impact on the system performance. Further, all resources would be operated in a synchronous fashion, potentially causing catastrophic resonances among periodic loads such as two-point regulated thermal processes.

The second option, the "discrete approach", allows only discrete values for s_i:

$$s_i = r_i \cdot s_{i,\max} \text{ with } r_i \in \{-1;0;+1\} \quad (8)$$

By doing so, both disadvantages of the continuous approach are avoided. All resources commit either no or their maximum storage potential to the system; and changes between these states are non-synchronous between all resources.
In order to determine when s_i changes its state, the positive value range of *S* is divided into *i* sub-intervals of size $s_{i,\max}$, resulting in a set of interval borders L = $\{l_0, l_2, \ldots l_n\}$. (This could also be done to the negative half of the value range without any differences since both halves are symmetrical.) The current state of r_i is then determined by (9).

$$r_i = \begin{cases} +1 \text{ for } S > l_i \\ -1 \text{ for } S < -l_i \\ 0 \text{ else} \end{cases} \quad (9)$$

$$\text{and } s_i = r_i \cdot s_{i,\max}$$

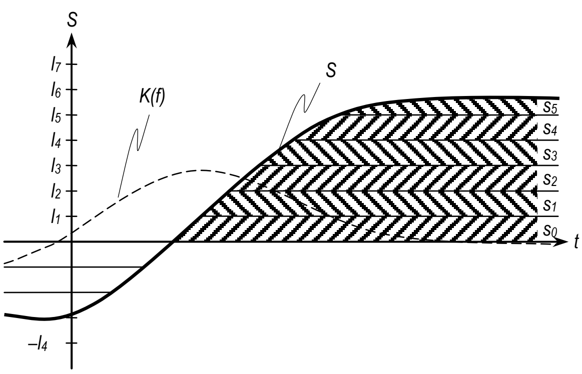

Fig. 5.4. The total value of *S* is constituted by multiple individual resources that commit s_i each. Each individual resource *i* can calculate *S* and knows its activation level l_i. When *S* exceeds this level, the resource becomes active and changes its own energy state s_i accordingly.

In Fig. 5.4 an example is shown how the individual s_i are chosen according to the development of *S* over time. First off all, it can be seen that *S* is calculated by integrating over *K(f)*, which itself is nearly proportional to the system frequency *f*. When *S* exceeds the first activation level l_0, which is zero because it is the first one, the first resource s_0 is activated. After a short while, S exceeds the next activation level and s_1 is activated. This goes on until *S* is no longer increasing. In this example, only six resources become activated. In reality this will usually be many more.
From looking at Fig. 5.4 it becomes intuitively clear that Resource 0 will be active more often than Resource 5. This is due to the non-homogenous probability density function of *S*. The probability density function of *S* plays a strong role for fair workload distribution in the system. Without additional measures, the workload will be distributed unevenly over resources. The most straight-forward solution for this is to re-arrange the activation levels over the value range of *S* from time to time so that heavily used resources will be re-situated in less frequently activated intervals. While the system was designed up to this point for only relying on the system frequency measurement as implicit means of communication, the task of re-arranging the activation levels requires a dedicated communication infrastructure.

5.2. Structure of the resulting system

Taking into account the structural considerations from the previous sections, the following structure for the system of distributed storages providing primary control power is proposed, which is also depicted in Fig. 5.5 (whole system) and Fig. 5.6 (single resource): A set of resources is connected to the power network. Each resource has the capability for storage of the energy $\pm s_{i,\max}$ and its energy state s_i can be changed at any time. All resources measure the system frequency *f* and use this to calculate the required system-wide energy state *S*. It can be assumed that the value of *S* is precisely synchronised and equal in all resources. Technical measures to achieve this are subject to discussion in subsequent sections.

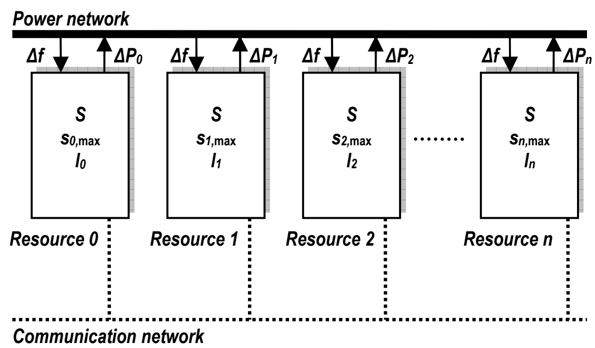

Fig. 5.5: Primary control system consisting of multiple participants coupled by the system frequency *f* and a supporting communication network

Each resource measures f, calculates $K(f)$ and integrates the result over time. Then, it uses its individual activation level l_i to determine its energy state. Changes in the energy state will result in a change of power consumption of the resource, and the system is designed in such a way that all individual changes in power consumption add up to realise the required power characteristic given by $K(f)$.

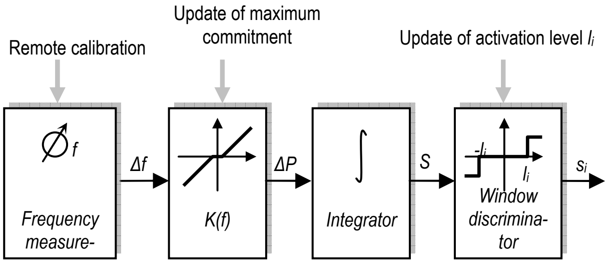

Fig. 5.6: Local calculation of individual set-point offset s_i only on the basis of frequency measurement. Gray arrows indicate possible parameter updates via the communication infrastructure.

This algorithm is executed on each single resource and uses the system frequency as only input. Therefore, a fast and dependable reaction on frequency changes is guaranteed. Nevertheless, there is a need for an additional communication infrastructure that serves for multiple purposes as outlined in Fig. 5.6. However, in contrast to the measurement of f, the information exchange over the dedicated communication channel is not time-critical.

- The precise measurement of the system frequency is crucial for a well-synchronised calculation of S. Since the system frequency is the same in the whole power system, it can be measured at any place. Measurement results at the resource's site can be compared to a high-precision central measurement and a remote calibration of the local measurement units can be performed.
- The number of resources taking part in the system can vary. Therefore, also the committed maximum power for primary control can change. Since this information is part of the power function $K(f)$, it must be possible to update the versions of $K(f)$ stored at individual resources.
- As discussed before, the activation levels have to be re-arranged over the value range of S on a regular basis for maintaining a fair workload distribution in the resources.

5.3. Improved algorithm for the case $b_i > 0$

So far it was assumed that the resource behaviour (power consumption) is only depending on the deviation of s_i ($b_i = 0$ in (3)). In reality though, there are also linear and even non-linear power components. However, non-linear components are assumed to be so small that they can be neglected. Nevertheless, the linear components of the resource power cannot be neglected, which can be concluded from an example shown in Fig. 5.7. In this figure, the system power for ideal resources ($b_i = 0$) and real resources ($b_i > 0$) is compared. The shown trajectories have been calculated on the basis of real frequency data, using (3) as continuous resource model. For the case $b_i > 0$, a realistic value for b has been assumed that was calculated for a household refrigerator ($b = {}^{1}/_{1440\text{s}}$). The result shows qualitatively very clearly that neglecting the linear power component is not possible. Even for relatively small losses in the process, the trajectory is far off the required power function $K(f)$. A solution for overcoming this major obstacle can be found, however this solution shall be motivated by another problem of the current system design.

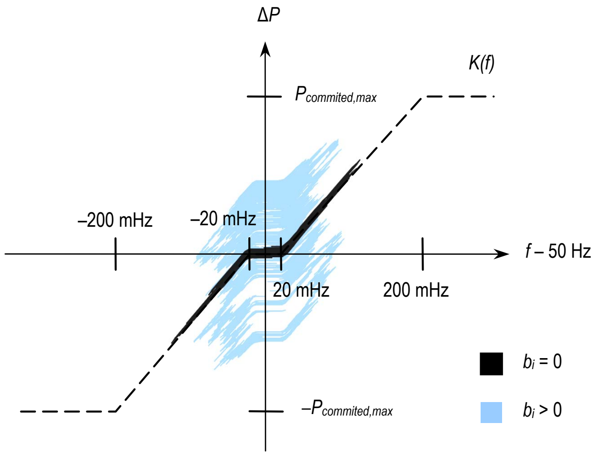

Fig. 5.7: Example trajectories of the overall differential system power ΔP. For $b_i = 0$ the simplified resource model fits well and results in dynamic operation come close to the required $K(f)$. However, with realistic values for b_i the simplified resource model is no longer sufficient.

In the previous discussion the value range of the overall system energy S plays a role. The basic idea of the system is to chose S in such a way that the deviation of S, which is essentially the system power ΔP, meets exactly the requirements of the power function $K(f)$. This is done by integrating $K(f)$ over time. So, the value range of S is actually the value range of the integral

$$S = \int K(f)dt \,. \tag{10}$$

Theoretically, the system frequency f has a perfectly symmetrical probability density. So, it should be guaranteed that the value of the integral crosses zero from time to time and that the integral is not divergent. However, this does not give any guaranties concerning the bounds of S. In fact, even if the probability density function of f is perfectly symmetrical, the slightest measurement offset will cause the integral to become divergent.

A very pragmatic but also efficient way of avoiding the value of the integral to exceed all bounds is to migrate from an ideal integrator to an integrator with 'losses'. This can also be seen as a low pass filter that is transparent for short-term changes (higher frequencies) but filters out the long-term increase or decrease of the integral value (lower frequencies). In an electrical circuit, a capacitor in combination with a resistor would be used, so that the loss in capacitor voltage is proportional to the voltage itself, resulting in an exponential decrease. The ideal integrator can be exchanged by such a component in the hope that the impact of this replacement on the system operation is only minor.

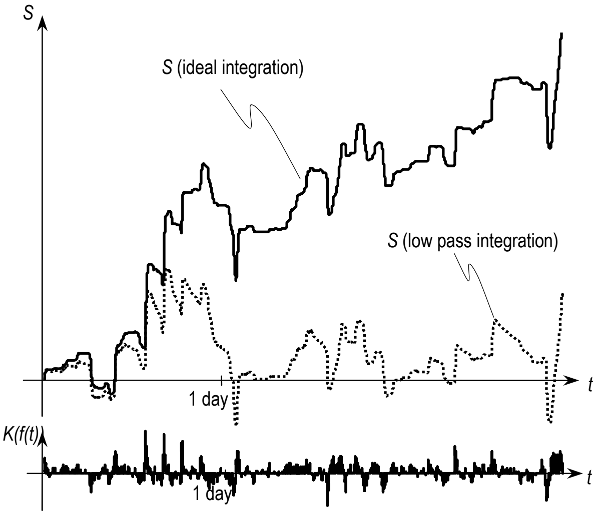

Fig. 5.8: The time-domain comparison of ideally integrated and low pass integrated frequency data measured over 28 hours reveals that the diverging trend of the ideal integration can be successfully avoided.

In order to calculate the impact of a non-ideal integrator, the system shall be examined in the frequency domain. Before doing so, some general matters and conventions should be considered:

- The system is seen as a continuous system and therefore the continuous Fourier transformation is applied. In practice however, the system will probably be implemented in a time-discrete fashion. The findings of the continuous analysis can also be derived using time-discrete analysis methods such as z-transformation.
- In order to avoid the system frequency being mixed up with the frequency domain variable f, the application of the power function on the system frequency $K(f(t))$ will be treated as basic signal $K(t)$ in the time domain and $K_f(f)$ in the frequency domain.
- In general all signals in the time domain will now be written with t as argument (e.g. $S \rightarrow S(t)$, $s_i \rightarrow s_i(t)$) and in the frequency domain with f as argument and subscript (e.g. $S_f(f)$).

At first, the system with ideal integrator shall be considered. By transforming (10) into the frequency domain, the following is achieved:

$$S_f(f) = \frac{K_f(f)}{j2\pi f} + \underbrace{\frac{1}{2}K_f(0)\delta(f)}_{\text{constant component } K_f(0)=0} \qquad (11)$$

The second term falls away since the constant component of the frequency signal is zero. For achieving the total system power, (3) is used and it is assumed that

$$b_i = b \forall i \qquad (12)$$

in order to be able to give a closed form of the solution:

$$\Delta P_f(f) = \underbrace{K_f(f)}_{\text{intended solution}} + \underbrace{b\frac{K_f(f)}{j2\pi f}}_{\text{error term}} \qquad (13)$$

Here, the second term transforms to the integral over $K(t)$ in the time domain, which is an error term that leads to the strong discrepancies between intended and achieved power trajectory as shown in Fig. 5.7. However, for $b = 0$ the result equals the intended solution.

In a second step, the non-ideal integrator shall be considered. While the ideal integrator has the impulse answer $\varepsilon(t)$, now a system with the impulse answer $\varepsilon(t)e^{-t/\tau}$ is chosen, where τ is a parameter that determines how fast the impulse answer approaches zero. In the time domain, this application of a different kind of integrator results in a new version of (10):

$$S(t) = K(t) * \left(\varepsilon(t)e^{-\frac{t}{\tau}}\right) \qquad (14)$$

Here, the operator * stands for the continuous convolution, which is defined by

$$a(t) * b(t) = \int_{t'=-\infty}^{+\infty} a(t)b(t'-t)dt' . \qquad (15)$$

(14) can again be transformed into the frequency domain,

$$S_f(f) = K_f(f)\frac{\tau}{1 + j2\pi\tau f} , \qquad (16)$$

and again (3) can be used to achieve a term for the total system power:

$$\Delta P_f(f) = K_f(f)\frac{j2\pi\tau f}{1+j2\pi\tau f} + bK_f(f)\frac{\tau}{1+j2\pi\tau f} \qquad (17)$$

$$= K_f(f)\frac{j2\pi\tau f + b\tau}{1+j2\pi\tau f} = \underbrace{b\tau}_{=1 \text{ for } \tau=\frac{1}{b}} K_f(f)\underbrace{\frac{1+j2\pi\frac{1}{b}f}{1+j2\pi\tau f}}_{=1 \text{ for } \tau=\frac{1}{b}} \xrightarrow{\tau\rightarrow\frac{1}{b}} K_f(f)$$

It can be seen from (17) that if τ is chosen to equal b^{-1}, then all additional terms fall apart and the system power is exactly equal to the required power function $K_f(f)$. This is a positive outcome, since it was expected that the non-ideal integrator motivated by the divergence problem would have negative impact on the result. The negative effects of non-ideal integrator on one hand and the linear power component on the other hand compensate each other ideally for $\tau = b^{-1}$. This can also qualitatively be explained: the linear power component modelled by the parameter b is the result of non-ideal energy storage. Losses are proportional to the absolute value of energy stored in the resource. Now, in the algorithm to determine the required individual storage state s_i for a given power amplitude, this behaviour reproduced by the non-ideal integrator.

Still, the question remains whether τ can actually be chosen to equal b^{-1}. Since τ is a free parameter of the non-ideal integrator, any value can be assigned to it. But b is just a simplification of the set of individual resource properties b_i (see (12)). Each resource i has its individual value of b_i depending on the quality of its storage isolation. Only for perfect lossless storage, $b_i = 0$. However, τ is a system-wide parameter that cannot be adapted to individual resources since it is used to calculate $S(t)$ which has to be the same in all resources for a precise synchronisation of charging and discharging activities. Consequently, $\Delta P(t) = K(t)$ cannot be exactly achieved, rather a good compromise for τ has to be found. Nevertheless, it has bee shown by (17) that it is possible by careful selection of resources and parameters to come arbitrarily close to the requirement $\Delta P(t) = K(t)$.

6. CONCLUSIONS AND FUTURE WORK

An algorithm for a new distributed automated demand response system was presented. A group of energy consumers react on grid frequency changes and coordinate themselves over a not-necessarily real-time-able communication channel. With this, distributed energy storages can be aggregated. Such storages are already existing in the grid but so far not used. Further, they do not contain dangerous substances, but need smart coordination, typically based on explicit and implicit communication.

Further research is necessary to create a consistent model of implicit and explicit communication channel. The project will conduct a series of field trials in order to evaluate the control algorithms, the reliability of the system and the impact on the energy grid and its applicability for regulation power. This field trial will be supported by further simulations while it will also lead to valuable conclusions for refining the simulation models.

ACKNOWLEDGEMENTS

This project is carried out and financed within the scope of the Austrian Program on Technologies for Sustainable Development, "Energy Systems of tomorrow" – a cooperation of the Austrian Federal Ministry of Transport, Innovation and Technology and the Austrian Research Promotion Agency (FFG).

REFERENCES

Kupzog, F. (2006). Self-controlled Exploitation of Energy Cost saving Potentials by Implementing Distributed Demand Side Management. In Proc. of *4th Int. IEEE Conf. on Industrial Informatics (INDIN 06)*, Singapore 2006.

Kupzog, F. and Roesener, C. (2007). A closer Look on Load Management. In Proc. of 5th Int. IEEE Conf. on Industrial Informatics (INDIN 07), Vienna 2007. (to be published)

Shokooh, F., Dai, J.J., Shokooh, S., Taster, J., Castro, H., Khandelwal, T. and Donner, G. (2005) An intelligent load shedding (ILS) system application in a large industrial facility. Industry Applications Conference, 2005. Fourtieth IAS Annual Meeting. Conference Record of the 2005 IEEE.

Trudnowsky, D., Donelly, M. and Lightner, E. (2006). Power-System Frequency and Stability Control using Decentralized Intelligent Loads. In Proc. of IEEE PES T&D Conference and Exposition, Dallas, Texas, May 2006.

UCTE (2004), Appendix 1: Load-Frequency-Control and Performance. UCTE Operation Handbook. UCTE Brussels, Belgium 2004.