

System Frequency - A resource for Sustainable Electricity

This paper is derived from an invited lecture to the JRC in Petten on 17th November 2011. I am most grateful for the opportunity and for the discussions that ensued, and for subsequent reviews of the text.



This work is licensed under a [Creative Commons Attribution 3.0 Unported License](https://creativecommons.org/licenses/by/3.0/). The file will be available at www.davidhirst.com/electricity, and at www.claverton-energy.com.

Abstract

Alternating Current (A/C) networks need to maintain a constant balance between supply and demand at all times, and the system frequency, visible to all participants, reveals imbalances that must be corrected. Significant power station resources are devoted to correcting imbalances. Millions of refrigerators / freezers, if cheaply adapted, can do the job better. The system frequency then becomes a more powerful and useful indicator for management of slower acting reserves and wider participation by demand and small scale generation in the balancing processes. The potential is for the frequency to fine tune real time prices, and so facilitate wider participation in electricity markets.

Real time prices are not enough to influence the many flexible devices (such as laundry machines) whose value depends upon effective scheduling of their activities. To minimise costs, they need to plan their consumption over longer timescales – hours and days. We need capability to influence these plans to consume at times when wind and renewables are available, as well as when other less flexible demands are low. Flowcost metering is an effective and cheap way to achieve this. By varying the time of consumption, the cost to consumers is minimised, and the efficiency of the system as a whole enhanced.

This changes the best way to use and run the grid, so we need a variety of tools and models to help us plan and develop policies to achieve efficiency. The paper also discusses the necessary models.

1 System Frequency and Imbalance

1.1 System Frequency

The concept of the system frequency pervades the concepts discussed. I recall the twitch of excitement when told by an NGC manager “The frequency is a system-wide signal”. Wow, very few systems have single signals that are truly system-wide.

So let us spend a few minutes contemplating the system frequency and what it can tell us.

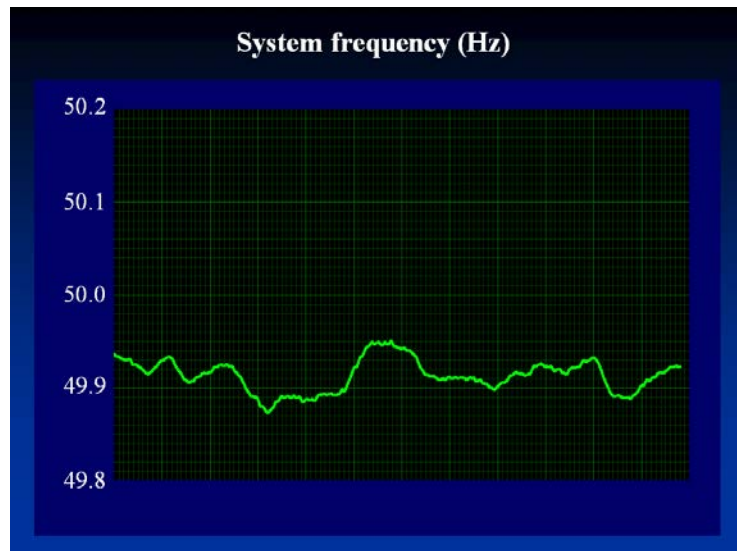
The basics are widely known. Nominally 50Hz (or 60Hz in the US¹), it reflects the spinning speed of the of the large numbers of generators synchronised to each other by multi-phase Alternating Current transmission links across a whole “grid”. Its mathematical model – “the Theory of Synchronous Machines” was developed in the 1930s and remains one of mankind’s great intellectual achievements. It remains the

¹ Also parts of Japan. Japan is unique in having both frequencies in use, and this greatly complicates its management of the rather severe electricity shortages triggered by Fukushima.

foundation of today's profession of electrical engineering, but will not be explored here!

Being system wide, the signal is available to everyone to look at. What you see here is a playback of a past Nordel recording. It represents about 2 minutes of real time.

What the theory says is that the system frequency is an integration of all the electrical activity on the grid. This means is that everything we do with the "mains" interacts with the grid and makes a difference to this system frequency. Mostly, of course, it is too small to be



noticed in the noise of lots of other people doing things. But collectively, all of us doing lots of things changes the system. We see the effect of this in the system frequency, which constantly moves up and down in a seemingly random way. Most of the time, you cannot predict from looking at it, whether its next move will be up or down, or how fast it will move.

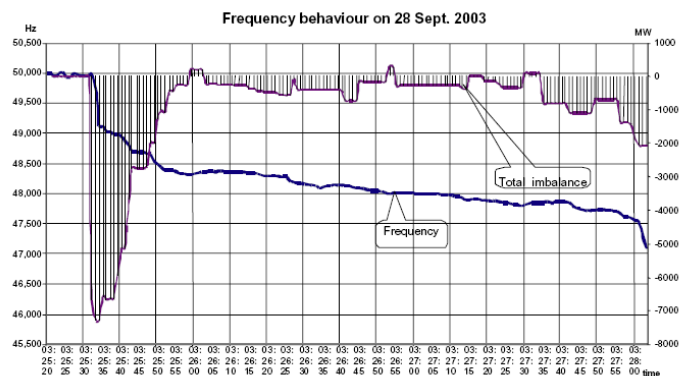
The same signal (near enough²) is what you will see anywhere on the system. For the ENTOS-E system (i.e. most of Europe[2]) the same signal is seen from here, in the North, to Gibraltar in the South, from Trafalgar in the West, to the Dardanelles in the East. It is shared even further away, via A/C links, with ENTSO-E in the process of bringing Turkey into the system.

The frequency shows the instantaneous **balance** of supply and demand, between generation and load across the whole system. When the frequency rises, it is because there is less demand than supply, and when it falls there is less supply than demand. So switching on a light reduces the frequency, and starting up a generator increases it.

1.2 Frequency Response

If the electricity system as a whole does not respond appropriately to changes that occur in the system frequency, and correct the indicated imbalance, it will collapse. As it did in Italy over about 2 minutes (Picture from [3]).

The system is inherently unstable. This is because, when the frequency falls because there is not enough electricity, the generators slow down – the frequency is a measure of their speed of rotation. But a slower generator means that the turbine driving it produce less



² There is some spring in the system, so there can be a few milliseconds delay between different parts of the system. In extreme cases, the springs can oscillate, but exploring this is too specialised a topic for here.

output³. So it slows down further, and, unless this is corrected, the system can descend into collapse. Although load also reduces with a change in frequency, it not by as much, and not enough to compensate.

To maintain stability the system frequency is monitored very carefully by a selection of the synchronised generators, who have been chosen by the “System Operator” to provide “frequency response”: i.e. to be “frequency sensitive”. Their role is to correct imbalances, and they do this by increasing output when the frequency drops (low frequency response), and increasing it when the frequency goes high (high frequency response). They notionally do this in a way defined by the “droop” of the generator, so that, for example, a 1% drop in frequency should produce an increase of (say) 8% in output, and similarly for increases in frequency.

Unfortunately, generators cannot do this instantaneously, so there is a delay between deciding to increase output and it actually happening. They also have to decide whether a change is just temporary noise, and can safely be ignored, or is the beginning of a trend, and so needs response. This is often done with a deadband – a narrow range of frequencies within which no control action is taken.

The acceptability of delay depends on the system: large systems, like ENTSO-E react more slowly, so generators can afford to react slowly; but smaller systems, like Ireland, are far friskier, and so the generators need to react more quickly⁴.

1.3 Cost of Frequency Response

There are cost and operational consequences of providing low frequency response – that is, being willing to increase output at very short notice. First, the genset cannot be running at full output, because it would not then have capacity to increase its output, and this loss of output is costly. Second, it has to have various related services immediately available to support the increase. For coal plant, this usually means all the coal pulverisers, with their big motors, have to be running, so the parasitic losses are greater, overall efficiency is lower and costs increase. Thirdly, it has to be “throttled” so that a reservoir of steam can be held back until it is needed, thus further reducing efficiency. And, fourthly, gensets are significantly constrained in the rate at which they can significantly change output, so it is not feasible to have a single genset on “spinning reserve” duty. The load has to be spread across multiple gensets each able to provide some of the needed reserve, and each facing extra costs as a result.

High frequency response can be easier to provide, after all, it is easy to take your “foot off the throttle”, but also has consequences – largely associated with disposing of the energy within the conversion process that you no longer need⁵.

Gas turbines, particularly Combined Cycle Gas Turbines, also face complex constraints in their ability to respond quickly. I remember hearing that a driver of a gas turbine Le Mans car used to have to “put his foot down” on the accelerator well before he had to break hard for a corner, because of the delay in engine response⁶. But

³ It is more complicated than this, as slowing down releases inertial energy from all the spinning machinery, but this is not enough, on its own, to remove the imbalance.

⁴ Strictly, it is the relationship between a loss and the inertia of the system. Small systems still tend to have large gensets.

⁵ In the UK high frequency response is particularly pricey, although there does not seem to be any fundamental reason for this. Some claim it to be an opportunity for generators to add charges to system costs that others pay?

⁶ Airliners use a lot of extra fuel when landing to keep their engines ready in case they have to abort. If they glided, the engines would be too slow to give power if it became needed.

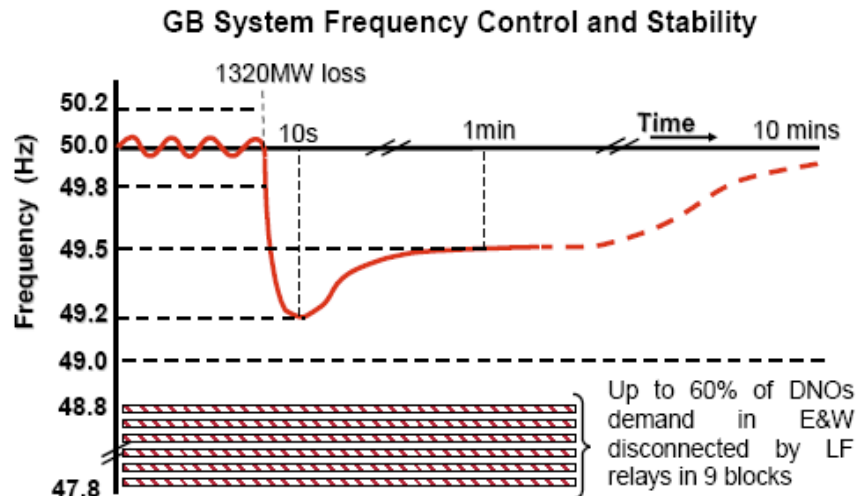
at least he knew when a corner was coming up, System Operators do not always have such solid knowledge.

So, although the system frequency can look random, it is in fact the consequence both of many small random activities, and a smaller number of larger more purposeful, but slower, responses, acting to correct short term imbalances. So we cannot tell whether the frequency is more a result of corrective action, or an indication of system-wide imbalances. The system operator, who has knowledge of how the gensets have responded, and of other actions they have taken can work out the deeper imbalances. At present, for most of us, the frequency looks rather like noise.

1.4 Frequency Events

Occasionally, a large frequency drop occurs [4]. Usually, this is associated with the sudden loss of a big genset, but they can be triggered by loss of an infeed link, or perhaps loss of a load centre. When this happens the system has to replace the lost electricity before the network of gensets in the grid (and hence the system frequency) slow down beyond safe limits. The rate at which the system slows down is largely driven by the inertia of the gensets, so energy in spinning metal is converted to electricity. In the UK, this means that replacement of a loss has to begin within about 5 seconds, and be largely complete within about 10 seconds.

But how much is needed? It is not necessary to completely restore the loss immediately, as the system can revert to a temporary, lower frequency mode, where load is reduced. Since frequency response is costly (and paid for as a system overhead), it is



desirable to minimise it, but not minimise it so far that there is a risk of system collapse. This is one of the more difficult modelling tasks of the System Operator. Broadly, however, a high proportion of the lost generation has to be replaced quickly, and the system restored to an N-1 operational state.

(N minus 1 (N-1) is broadly a state in which the system can survive the failure on any one of N components without significant overall harm).

But we now meet a further time dimension to the issue, the duration of the short term response. In the UK, frequency response is usually expected to last for 30 minutes⁷, with the general expectation that further despatch instructions (or buy sell acceptances) will arrive within that time. The System Operator is expected to be properly aware of the situation, partially from watching the system frequency, but

⁷ To be strictly accurate, this refers to what NGC calls secondary response, and almost everybody else calls primary response. In the UK primary response only needs to last 5 minutes or so, before being replaced by secondary response. Both are automatic, but secondary response is slower. Elsewhere, AGC can be known as secondary response. See section on vocabulary and terminology below

also by monitoring of telemetry from all the frequency sensitive gensets and the load aggregations.

1.5 Hierarchy of Reserves

The core option open to the System Operator is to activate reserves. This may be buying more from existing gensets – asking them to turn up their output; or by ordering further power station starts – usually gensets that are quick to start are inefficient, fairly small (several are required), and expensive; or they may be able to purchase reductions from (big) consumers; or they can ask local distributors to lower voltage, and so reduce demand in their area. In the ultimate case, they may impose blackouts on parts of the system. Indeed, if the frequency is not kept within limits, frequency sensitive relays will cut in, and disconnect chunks of load. If they do so, the hope is that this will rebalance the system, and sometimes it does. System Operator decisions will depend on the overall state of the gensets (perhaps in a market); what they have previously planned to be available; and their assessment of the nature of the failure, most particularly, how long they think the loss will remain.

If the loss is expected to be short term, they may expect to get by with only short term purchases, to tide them over for a few hours until things are fixed. If the loss is longer term, (as will be the case with a scrambled nuclear plant), or even permanent, they may order startup (i.e. buy future output from) plant that will take some hours before it is fully operational.

For System Operators, it is all very complicated, and, crucially, they have to cope with the time delay dimension related to the resources they can call on. They must have systems (and models) that give them a view of the uncertain future, and help them ensure that they have responses to all reasonable contingency events.

There is also political / market dimension. System Operators are solely responsible for the short term, but as they start to look further ahead, responsibility passes towards the market players (the generators and retailers) who sell and buy electricity. If the frequency drop is from a failure of a genset, but the actions (and costs) to recover are taken by everyone else, it seems fair that the owner of a failed generator should compensate them. But how? And if they were contracted for the electricity that the genset will no longer it will no longer be delivering, and what should be the market mechanism for them to pay for this? How far ahead should the system operator plan, and to what extent can the SO expect market players to have capacity available? In other words, how should the market work? Big questions.

Hierarchy of Reserves

- **Fast Reserve** – which may already be spinning but not frequency sensitive
- **Peaking plant** – quick start
- **Capacity Backup** – e.g. warmed up plant, demand reductions, voltage reductions
- **Market readjustment** – longer term replanning

Different operators have different terms and different markets for reserve

Automatic (or Area) Generation Control (AGC)

AGC is a centralised scheme under which a group of gensets within a very large system agree to respond to a centralised area signal that will adjust the set points of their output in a co-ordinated way. The centralised signal is derived from a knowledge of the overall balance of supply and demand within the area, which is acquired from instruments that sense the flows from all gensets and from the various “inerties” to connected areas. This allows the centralised site to work out whether flows are as “scheduled”, and load is as expected. If the area is importing more than scheduled, then the local generation is asked to increase (the AGC signal is a plus), and when the load is less than anticipated, the AGC signal is a minus. The system frequency, which covers a far bigger system than the particular area under control, is a part of the knowledge used to derive the AGC signal. Thus, in following the AGC signal, the generators are “load following”. That is, their output is responding to the variations in load. So the ancillary market concerns load following, rather than frequency response,

This works when the system is large enough for any individual failure to be small enough to have a small impact on the system frequency, and the AGC signal (which changes more slowly than the system frequency) is fast enough to ensure response before the frequency goes outside limits. It is also possible for the gensets to respond directly to frequency, and use the AGC signal to modify the set point of the frequency. When they do this, they may be paid only for load following, and frequency response is bundled in free with it.

The Great Italian Blackout

This arose because the lines carrying electricity to Italy (via Switzerland) broke, so Italy became short of some 6GW. This was beyond their N-1 planning, and a cascade developed that shut down almost the entire system. As a consequence, the system frequency in the rest of Europe rose – suddenly a lot of load had disappeared. However, some of the AGC systems noticed that they were not exporting as much as had been scheduled, and so, following the normal approach, they increased the output of the gensets under their control. This raised even further the system frequency, and the system was rescued only because alert operators managed to manually switch off the AGC before things got further out of hand. Despite being automatic, it still needs constant vigilance.

Terminology and Vocabulary – a confusion for the unwary

This seems a good point to introduce a nuisance complexity: the definitions around frequency (or short term imbalance) management. In different systems, different definitions apply to terms like Primary Response, Secondary Response, Tertiary Response and AGC. Many of the variations concern the time dimension, i.e. how fast and how long it must last, and have pretty rigid definitions of it. Such rigid definitions may be useful in managing specific markets for specific products based on specific technologies, but along with ambiguous terminology, are confusing, do not help system understanding, and inhibit innovation. My preferred vocabulary is: **Frequency Response**, for the first seconds and minutes; **Fast Reserve**, for the early minutes (or **AGC** when it is used); **Backup Capacity** for the first hours; and **Timekeeping**. Each domain can then add time parameters as suits them. It is important to remember that they are an overlapping set, with fuzzy boundaries that move as convenient.

System Time

The frequency response discussed so far is concerned to restore the system frequency to nominal. But if there has been a frequency dip, any clocks driven by the system frequency will have lost a bit of time, and so will be a bit slow. Over time, this may be partially compensated by lifts in frequency, but, in general, the clocks will slow down unless something is done to fix them. This needs extra generation, so that the frequency is raised a bit above nominal for a while, until the clocks catch up. The trouble is that buying this extra generation is an overhead cost to the system operator. Today, it is not clear how much it's accuracy matters. Currently System Operators do have an obligation to maintain system clocks, and tend to minimise their costs by letting the clock slow down during the day, when fixing it needs more generation, and get the clock to catch up at night, when less generation is need to fix the time, and when the available generation is cheaper.

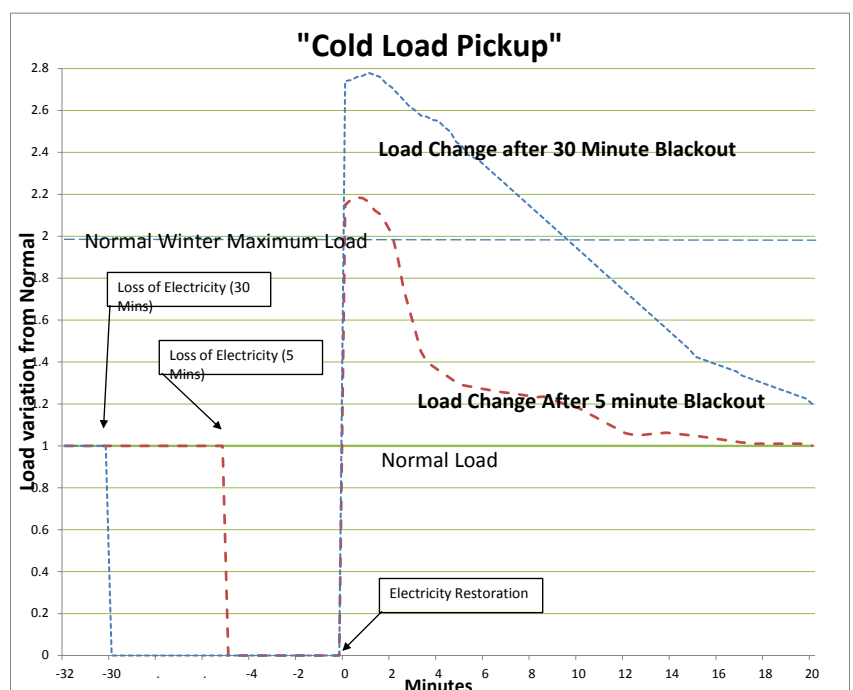
But do we really need it? How important is it today that mains clocks are accurate? While I do not have a clear answer, I can say that, unless there is some sanction against slow clocks, generators will tend to let it drift, and at some point it will make the system dysfunctional – as happened in India (See for example [1]).

2 Demand

So far, we have looked at the system frequency mainly from the supply perspective, aiming to meet a fluctuating, random, but partially predictable demand. But demand is not one thing, it is millions of small demands, aggregated. One vital thing about demands are that they are **diversified**. That is, they are NOT synchronised, but depend upon variations in the timing of individual demands.

It is feasible to draw some conclusions about the statistical distributions from which these probabilities of the timings are drawn, and to derive pretty reliable predictions about some of the parameters of the probability distributions, and how they vary over time of day and seasons. So, within limits, demand can today be pretty precisely predicted. The operation of the system depends critically upon the diversity of timing activity, as, if we all chose to switch on all the devices we use at the same time, this would far exceed the capacity of the system to provide. Without diversity, the result would be instant blackout.

We can see something of this in a phenomenon known as “cold load pickup”, which arises after a blackout. When power is restored, the demand in an area can be several times normal load, as all sorts of devices (for example, heating systems) aim to restore their normal operating conditions. The diagram (derived from[5]) shows the extent of extra demand after outages of 5 minutes and 30 minutes in a sample of Swedish houses. This is one of the reasons



why, after a big blackout, restoration is progressive and protracted. Each area has to get over its cold load pickup before there is capacity to connect another area.

So one thing that gives all systems operators nightmares is the possibility of events that tightly synchronise demand across a wide area. There is no higher need than diversity in demand!

2.1 Demand Heterogeneity

Demand devices are also extraordinarily heterogeneous. There is a quite very broad range of consuming devices and systems. A system operator, looking at the aggregated demands can readily forget that the loads do behave differently. Indeed, it is hard to find a comprehensive and helpful classification scheme.

Fred Schweppe,⁸ classified demand into “power-using devices” and “energy-using devices” in his patent for FAPER – a Frequency Adaptive Power Energy Rescheduler” back in 1979 [8]. Broadly, the services offered by **power-using** devices needed electricity **now**. If there was no electricity, there was no service, so things like lights, TVs, computers (in his day they were minicomputers) were all power-using devices. **Energy-using** devices, on the other hand, had some form of internal energy store, and so could continue to provide their service without continuous use of electricity. So things like immersion heaters (in a tank) could provide hot water at times different from their electricity use. Sometimes heating systems can work this way, at least for a short while. Today a laptop or tablet would count, as battery charging need not coincide with use. The example Fred used was of a high-level water storage tank, which could deliver water at any time, and used electricity only when the pump was replenishing the tank. And fridges, which provide a continuous service of coolth, which I am about to come on to.

I think there is a further helpful classification along a dimension of **priority**. There are lots of electrical things we find useful and helpful, that enhance our lives. Some we could live without, at least in emergency or for a short time. For things like flood lighting, or shop window displays, there may well be circumstances or risks of blackout (or a real time price) where we would wish them switched off. For me, at home, I would have the PC as a higher priority than the TV, but my family might choose otherwise. Some uses, like life support machines, and emergency lighting, are of such high priority we cannot rely only on the public supply, but have local backups. So priorities will be diverse, and depend upon individuals, households, businesses and the public sector. We do not know much about what the price needs to be to persuade us to program them to switch off – we have never tried to find out – but currently we have no sensible way to prioritise our uses, and the all or nothing choice is made for us by the system operator, in choosing which districts to black out when short.

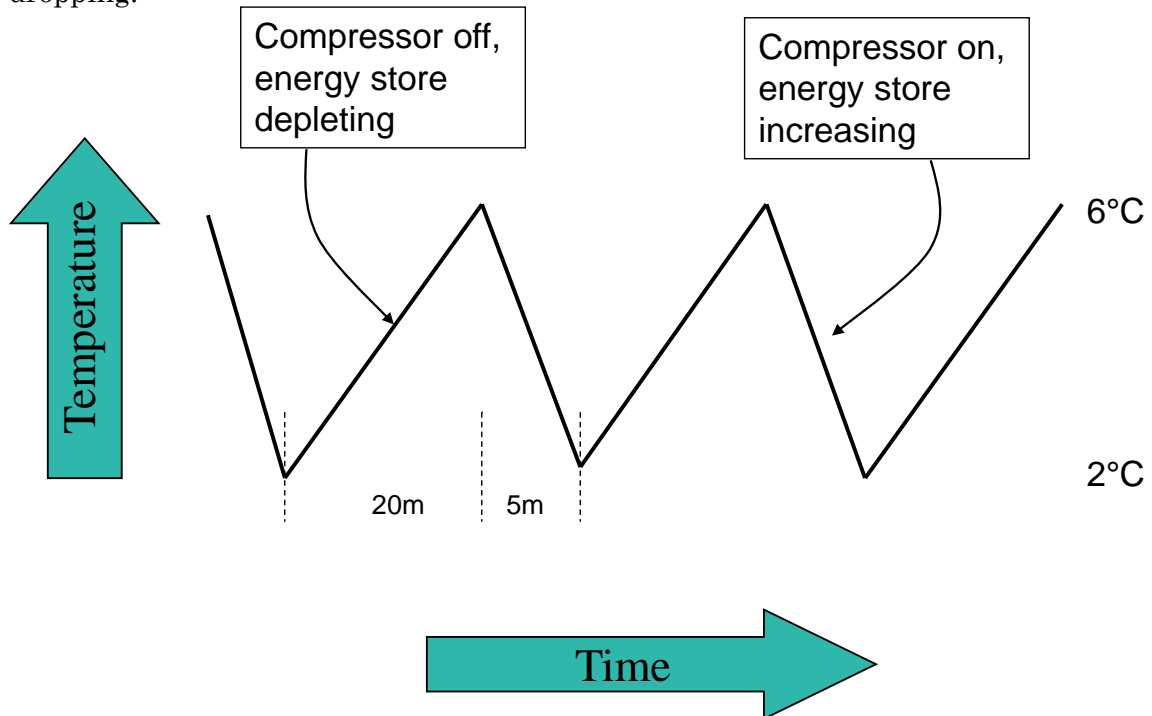
Flexibility is another dimension I will explore further later. There are an increasing range of appliances, exemplified by dishwashers and laundry machines, where, so long as the user deadline is met, the consumption can be planned to suit other criteria. So long as I have clean dishes when I plan to get up, I do not care when the thing runs. The big future load coming up is battery cars.

2.2 Fridges and Freezers

So now to refrigerators and freezers. For my purposes, they are virtually identical, so when I say fridges, it is really shorthand for refrigerators and freezers.

⁸ His book, Spot Pricing of Electricity [6], is a classic, but he is best known for his work on State Estimators [7] – deriving an assessment of the state of the system from available measurements. He died in 1988.


Fridges are energy-using devices, whose paramount duty is to maintain an internal temperature within limits at all times, 24*7, every day of the year. The temperature behaves like a saw tooth. When it reaches the set maximum temperature the cooling unit (which is what consumes electricity) switches on, and the temperature starts dropping.



When the temperature reaches its lower limit, the cooling unit switches off, and the temperature starts its slow climb again, at a speed depending on the insulation, external temperature and how often the door is opened and stuff put in or taken out⁹. The fridge is always delivering its service, but only intermittently consuming electricity. There are few prospects of this changing significantly with technology evolution, but I will skip these arguments here¹⁰.

If the cooling unit switches off before the minimum temperature is reached, no harm is done. Similarly if it switches on before the maximum temperature is reached, no harm is done¹¹. So the device is, within limits, flexible in its consumption pattern.

Domestic refrigeration consumes about 3% of peak demand (about 1.8GW in the UK). They are always on, have flexibility, store variable amounts of energy, and increasingly include computer controllers, so are very cheap to modify – just add software. The controllers have voltage detectors (it is better to switch on and off at a zero voltage point), so they can readily measure frequency. They can switch load very quickly – within about 100ms



⁹ It is actually a bit more complex than this, as there are lags and inertia in different parts of the fridge. But the principle is sound.

¹⁰ The main reason is that the cooling unit has to be seriously oversized to cope with extreme events (freezer makers call one the hunting test, when an entire deer has to be frozen). Cooling units capable of wide variations in output are a long way from being as efficient in normal working as a unit that stops and starts.

¹¹ To be rigidly accurate, there is a very small cost associated with inrush currents when a motor starts. Although visible in real time monitoring, it is almost undetectable in consumption measurements.

– as so can respond far faster than big gensets. As a candidate participants in the electricity system, fridges look good.

There are constraints: temperature limits – and so food preservation – is paramount. Any burden must be shared equitably across all fridges. There are some technical restrictions on the how rapidly cooling units can be switched without harm. There is no scope for communications, except via the frequency, so the control has to be wholly autonomous and independent.

2.3 The business and environmental case

There is great value to be captured. If all domestic UK refrigeration switches off, most of it could stay of for about 15 minutes, so spinning reserve costing around £200m p.a. could be completely replaced by the fridges. The fridge fleet is replaced every 12 years or so, so in 12 years, the transformation could be complete. The value equates to about £20 per fridge, so a bit of software adds £20 in value to every fridge sold. There seems to be a business case.

The is also an environmental case. Estimating models vary, but my estimate is that UK power stations emit around 2 million tons of CO₂ p.a. just to provide frequency response. Extrapolating across Europe, it may be as much as 10 million tons.

Whatever generation we use, we will continue to need frequency response. It would be tragic if we do build a wind based system, and then have to keep fossil fuel plant alive just to provide frequency response. Yet food preservation is one of the most important services that electricity enables, so there is not much prospect of it being made as obsolescent as coal fired power stations¹².

For other domains the numbers vary, and large systems need proportionately less frequency response. Indeed, this is one of the core economic reasons for enlarging systems – doubling the size of the system only marginally increases the need for frequency response, which is largely driven by the biggest single feasible failure. But all systems do need frequency response (even if they call it as something different) and all would benefit from the fridges.



2.4 Risks & Challenges

A key risk is that of **synchronisation**. If the fridges become coordinated and act inappropriately they could do huge harm. So the control has to ensure that this never happens – the natural diversity (all fridge are slightly different) must be preserved at all times.

There is the challenge of changing the operating procedures and, indeed, the thinking of operators when handling a frequency event. System Operators do understand the behaviour of generators, and have sophisticated command systems to order big changes in output. As I put it, they are used to big knobs, with 10s or 100s of MW at their beck and call. Fridges, each averaging say 50W, do not behave in the same way, and do not offer opportunities to intervene. Once the fridges have been manufactured and released “into the wild” they will continue responding to frequency, as they have

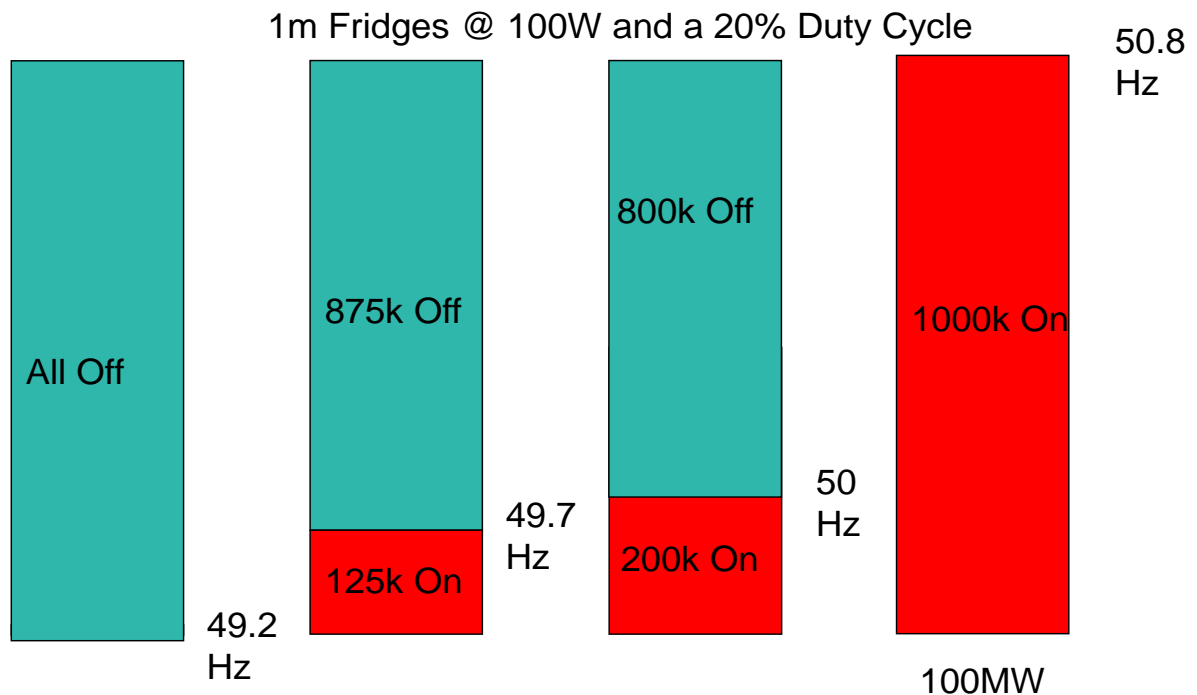
¹² Greater fridge efficiency does reduce the frequency response available, but we have some way to go before it is not worthwhile.

been programmed to do, until they reach the end of their working lives. So System Operations will need to change – becoming less constrained and, I hope, less complex. Planning and scheduling for frequency response will (eventually) become unnecessary, enabling greater focus of the efficiency of the needed reserve. SOs will need to evolve their strategies.

3 Balancing Fridge Control

So what do we need the fridges to do and how should they respond to frequency changes?

One aspect of the fridge control is straightforward, conventional and readily understood. Each fridge gets a frequency to which it is sensitive, and when that frequency is reached, it switches. So if it is on, there will be a low frequency at which it switches off. When it is off there will be a high frequency when it switches on. If these frequencies are spread evenly across the population of (diverse) fridges, then a progressive drop in frequency will progressively trigger a progressive reduction in the



load from the population of fridges, so reducing the imbalance. In these circumstances load will closely follow frequency. As we can follow on this slide.

Actually, there is a complication here. Fridges generally spend more time off than on, so there are fewer fridges available to reduce load than there are to increase load. Low and high frequency response is not symmetrical. Does this really make a difference? No, but there are questions here to explore further!

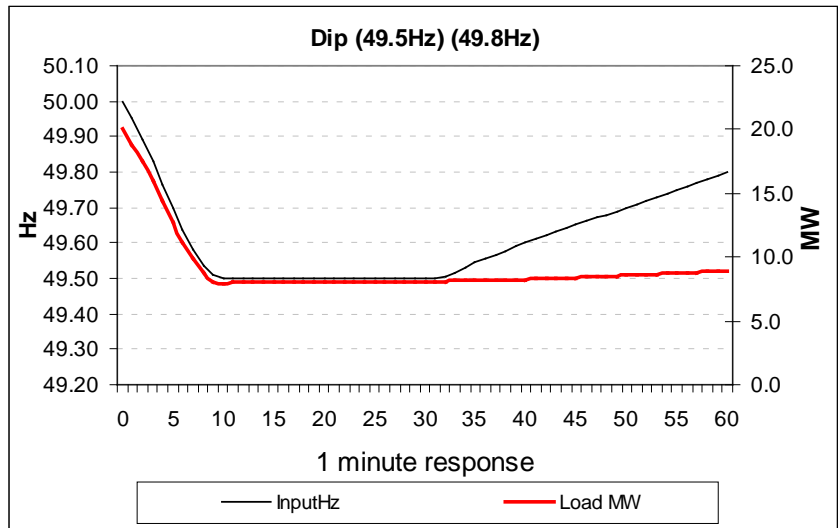
The key challenge, of course, is to ensure that the sensitive frequencies are spread evenly and appropriately across the population of fridges. Three principles are needed: randomness; population dynamism, and controlled recovery.

Including **Randomness** in the control calculations ensures that even two identical fridges will behave differently, so enhancing (and speeding up) the natural tendency to diversity.

Population dynamism recognises that the state of fridges changes over time. There will be fridges that have just switched on, so they are not the best candidates for being asked to switch off. On the other hand, fridges that have been on for some time have

less to lose if they switch off. So the concept is to move the frequency to which each fridge is sensitive as it progresses through its cycle. Fridges that will soon switch off are more sensitive to small drops in frequency. Fridges that have only just switched on will only go off if there is a big drop in frequency. Fridges that hit their temperature limits will ignore the frequency, and switch whatever, as they do today.

Controlled recovery manages the range of frequencies over which the population of fridges is sensitive. It is best explained in the context of a big frequency dip. [slide] As the dip progresses, fridges will progressively switch off, so reducing the load, and (when enough have responded) stop the decline. The frequency will level out. At this point the fridges have done their main job, and it is up to other Fast Reserve to start providing extra generation. The fridge population can keep its consumption low for a while, but progressively fridges will start hitting their temperature limits and switch on again. Some of them will be replaced by fridges that move into a state where they can switch on, but there will be a gradual increase in the overall load. If nothing further happens to restore the system balance, the fridge load will gradually increase to its normal level.



Other generation from many sources will come in – I later discuss how the frequency might drive this – and this will raise the frequency again. Conventional frequency response (while it remains!) will also respond, so the dip may be quite short.

As the frequency rises, we need to repay the energy “borrowed” from the fridges, so we want some of the off fridges to switch on. We can do this by adjusting the fridges perception of what is a high frequency. One convenient way to do this is to use a technique first suggested in 1977 by Stephen Salter [9] for switching night storage heaters. That is to construct a moving average of past frequencies, and use this moving average as the target control frequency. When the measured frequency moves above the moving average, the fridge population will increase its load. So, as the dip recovers, the fridges will start to recover the loaned energy.

If the frequency continues to decline, the fridges will again seek to reduce load.

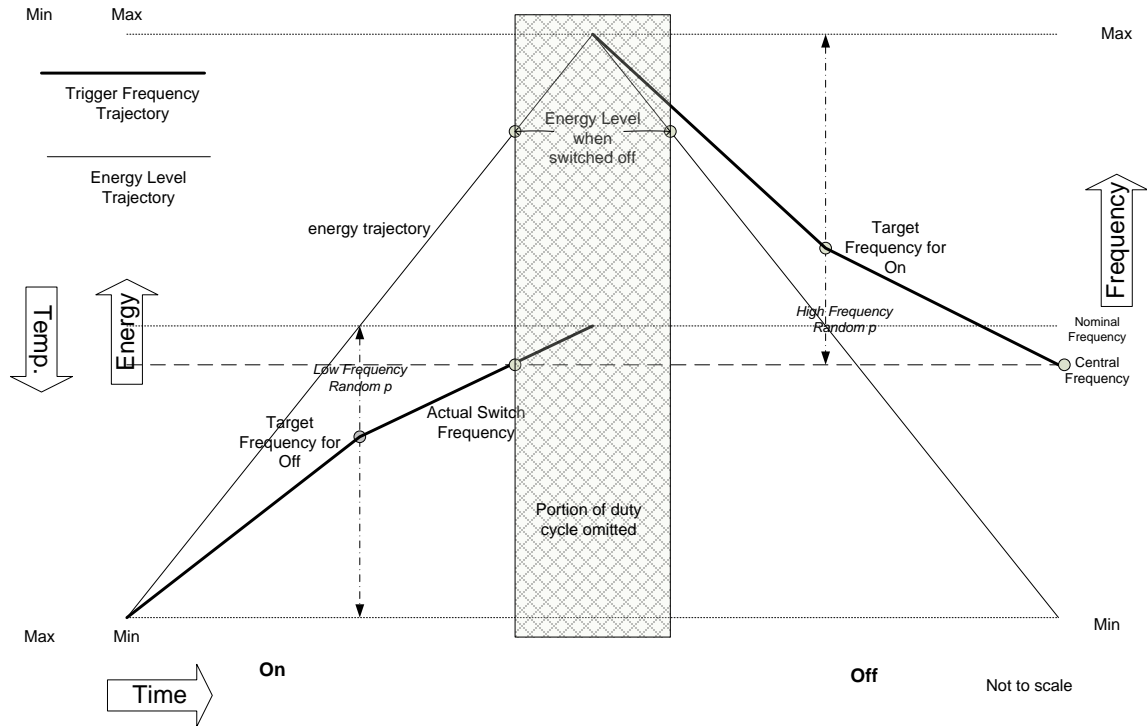
For high frequencies, you just need to think the other way round. I often find this sort of double think hard, but it is important to note that the fridges provides BOTH low frequency response AND high frequency response. In the UK this means the system operator can be billed twice.

All this behaviour is encapsulated in the diagram [overpage from [10]], but I do not propose to explain it here.

3.1 The nature of borrowed energy

I have mentioned energy borrowed from the fridges. What form does this energy take? For an individual fridge, it will only have switched off a bit sooner than normal. So the temperature is higher than it would have been, so the stored energy is less. This will be true for all the fridges that switched off, so the average temperature across the population will have increased, and the energy in them all reduced. Some fridges

(those already off) will have been unaffected, and their temperature (and energy) is as previously expected. Across the whole fleet of these two populations of fridges energy has been given up.



What difference does this make? Well, it means more fridges than normal will want to switch back on again. Their **off** cycle will have been shortened. So their start and stop times will be less diverse than before. Their cycles will be more synchronised. System Operators are likely to worry about this.

We need to restore this lost diversity as soon as possible, and this means getting fridges to start again. Hence the need to start recovering energy as soon as frequency rises again, well before the frequency has been restored to nominal.

Thus the normal system operator assessments of the energy needed to increment the frequency need to be adjusted. They need to be fully aware of how the fridges will behave, and be prepared for the slower recovery of frequency.

3.2 Other questions arising

This control approach raises other questions. What happens if the frequency is not restored? How long can the fridges hold the frequency steady? What happens if there is a “double dip” recession, and the fridges are called upon to serve again before their diversity has been full restored?

I can answer the first of these. It is a pathological case, and, for the system, may well be fatal! If the frequency stays low, more and more fridges will revert to a hunting state. After switching off, their temperature will soon hit limits, and it will switch on again (as soon as it can – there is a minimum off time constraint. As soon as it switches on it will find the frequency below its trigger point, and it will switch off again. All rather like fibrillation, rapid switching, but not really able to do its job effectively. Another analogy arises from overloaded multi-tasking computers. The computer spends its time switching between tasks, but gets little work done before it has to switch again, an phenomena known as thrashing.

Meanwhile the temperature is hovering around its upper limit, so there is no further response available. This is a unhappy state, so it is important that the system

operator does something about it before more than a small minority of fridges get that way.

How long can the fridge response last? Well, it all depends: how deeply the frequency dips; how long the depressed frequency lasts; and, to some extent, on parameters set in the populations of fridges. Broadly, however, since we are dealing with borrowed energy, you can measure it in Wh, or, if you prefer MWh. 1 million fridges can offer about 20MWh, which can be delivered at a peak (but diminishing) rate of around 80MW. Note that it is the fridges that decide this, with the system operator no more than an interested observer!

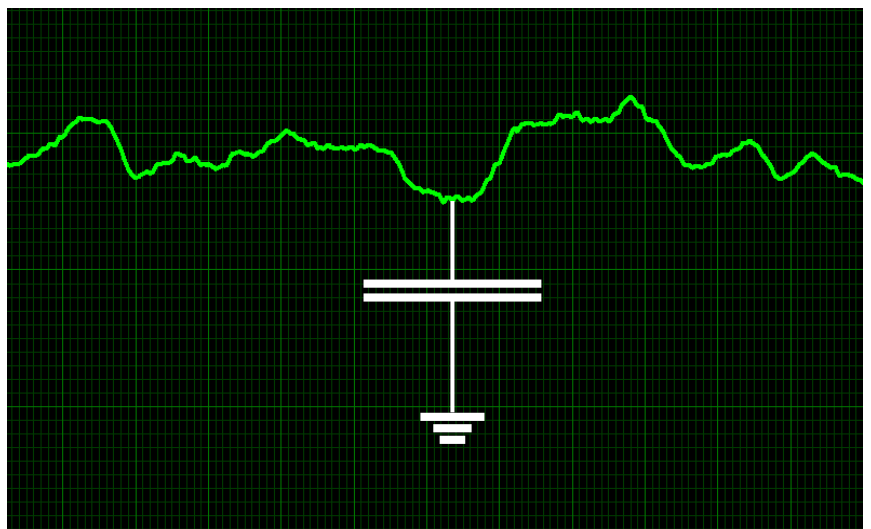
The balancing fridge controlled was conceived for wholly autonomous devices, millions of them, whose collective behaviour was valuable to the electricity system. However, if some communication to the fridges was available, there may be some scope for tuning the fridges to better support the grid as it evolves. This has not so far been tested or modelled. Any such feature would add significant communications costs and complexities.

What happens in a double dip recession? An interesting question, with the answers still incomplete. I would love to really know. The models I built at RLtec were not sophisticated enough to give an detailed solution, and the 100 fridge trial at Electrolux was not deep enough to be definitive. There was good evidence to suggest that, so long as not too much of the Wh storage was used up, the fridges will behave well. There was no evidence of anything disturbing.

3.3 No deadband

But, before getting further into modelling needs, there is another area where the consequences are not as clear as I would like, and that has potentially fascinating implications for system operation, and concerns the accuracy of the analogy with capacitance on the system frequency signal.

Fridges have no deadband. The controller is reacting to the frequency all the time, including when it is close to nominal. So, in normal operation, whenever there is a small variation in frequency some fridges will respond. They switch off early whenever the frequency drops, and switch on early whenever the frequency rises. They will not be very early,



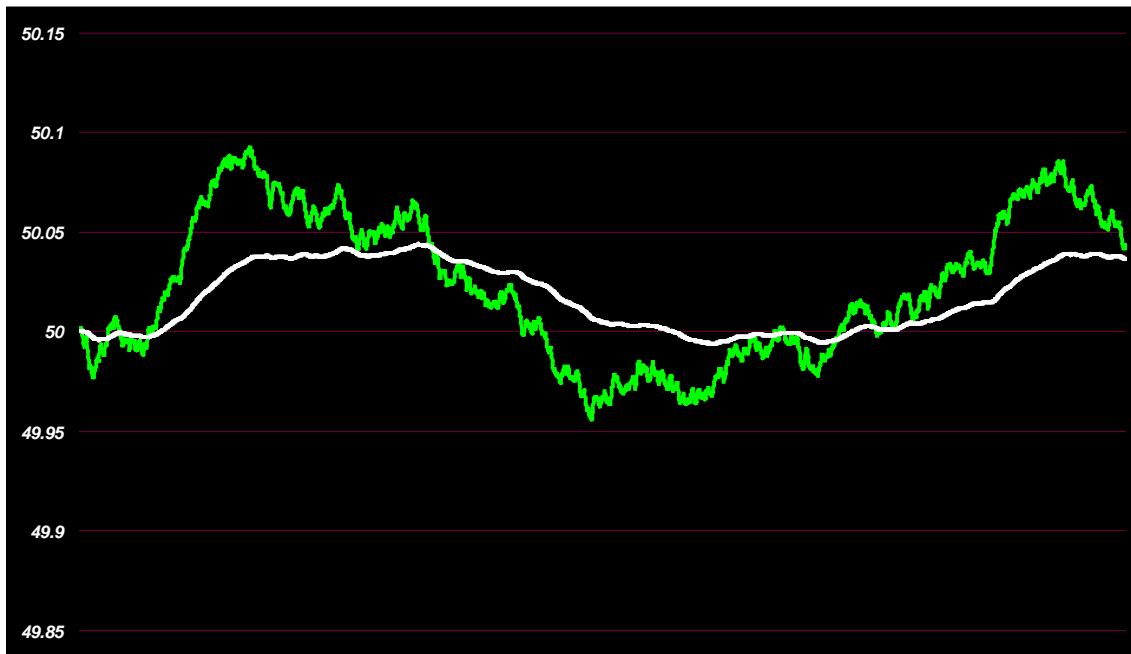
perhaps a few seconds in a cycle of minutes, but they will be constantly making small changes. This is a good thing, for reasons I am about to come on to, but it is not quite intuitively obvious how close to capacitance it really is, and what effects small fluctuations have on the extent of the response.

Consider a small increase in frequency. The fridges closest to normal switching will switch on early. If this is followed by a small decrease in frequency, the fridges closest to switching off will do so. If now, the frequency rises again, the fridges that were most ready to switch have already done so, and the frequency space is depleted until other fridges approach their normal switching points.

3.4 Modelling needs

This means that the extent to which fridges change load in response to frequency depends in part upon the recent past history of the frequency. Variations from a few seconds ago (perhaps minutes) makes a difference to the fridge response delivered.

I can find no reason to believe this matters, but it would be nice to know how great is the impact. I tried to model this as a sliding depletion zone, but this is a simplification of the fridge population behaviour. and the potential complexity of frequency behaviour prevents simplistic or intuitive answers. In this respect, it behaves as an imperfect capacitor, and a straightforward model of the impact – a damped frequency – is not quite perfect.



For me a significant modelling question is: Can one build a computationally simple model of that reliably reflects the aggregate behaviour of many fridges? Put another way, to what extent is it necessary to simulate the complex behaviour of millions of individual (and different) fridges before one can reliably predict the behaviour of the whole.

I once believed that large scale simulation was necessary, but I now believe that, so long as computationally simpler (parameterised) model components are validated by simulation of a few hundred individual (and diverse) fridges, large scale simulation is not essential. It remains an open question.

4 Frequency Tells Price

There is a further but quite profound point about fridges and storage. When you have a system where the frequency is regulated wholly by fridges, then it is possible to build a function that converts the recent history of the frequency into a real time assessment of how many MWh the system is long or short. In effect, one can know how much the capacitance is charged or discharged from nominal. This function is essentially an integration of the variation of the frequency from its nominal value, and so is equivalent to a measure of the variation of the system clock from the actual time.

So far as I can discover, such a reliable assessment of the real time state of what is essentially a market is unique. Most market and trading systems discover whether the system as a whole is in surplus or in scarcity by tracking the price, which emerges

from a complex haggling system among participants. A rising price indicates shortage, and a falling price indicates surplus. In so far as economics has laws, this is a fundamental one. Yet here we do have a reliable mechanical, automatic and built in indication without the haggle and hassle.

So can we turn it around? Can we say that, because the system faces shortage, the price should be higher? When the system is in surplus, the price should be lower.

I think we can, and should. That is, the price at which electricity is bought, and the price paid for its generation should at least be modulated by the frequency. I fear this means a pretty deep rethink of the microeconomics of electricity, and I will offer a brief discussion in a minute or two, after a quick review.

4.1 Fridges in Control

The system frequency is an extraordinary integration of heterogeneous and diverse activity across an entire electricity system, and gives a signal we can all see, and partially interpret. The signal is a core measure of the health of the system, and its reasonable stability is absolutely vital to electricity and so our well-being. Great intellectual and monetary resources are devoted to it. At present its primary control is via large, fossil fuelled power stations, at significant cash and environmental costs.

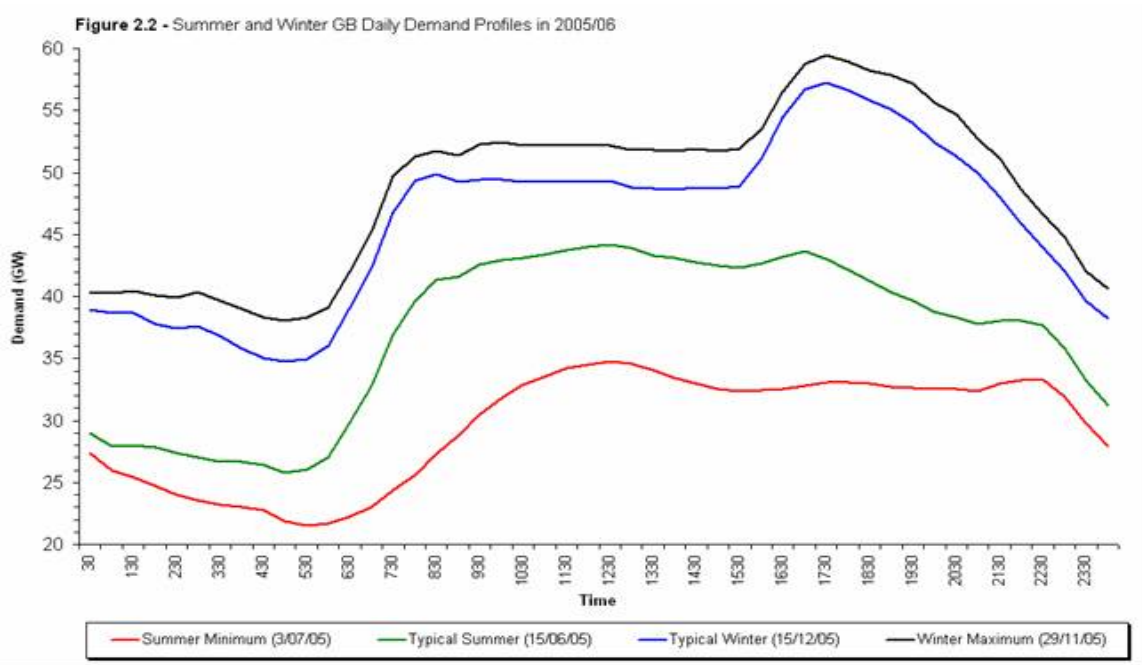
Yet enhanced fridges can do the job much better. They are cheap, environmentally acceptable, always available, highly reliable, and have benign and useful consequences on the future evolution and capacities of the system. They are simple, highly distributed, and give consumers scope for participation and understanding of the electricity system, if they wish.

Yet, at present, their adoption is stalled. There are many reasons for this, not all edifying, but one is the lack of modelling tools that demonstrate the benefits, and support those who need to manage the changing and future system. A significant, but not huge, research project is needed to build prototypes of those tools, and show, beyond reasonable doubt, that the system will work better with smart fridges.

4.2 The pricing of electricity

At this point I considered stopping and asking for questions, but I am pretty confident that, if you have followed me so far, you will ask. “You talk about fridges, but what about other devices: can they also provide frequency response?” To which the short answer is “no”, a full answer needs another lecture, and today we have just time for “well, perhaps, maybe, sometimes, so long as you change the microeconomics of electricity”. So here is a short introduction to the changes we need to contemplate.

I mentioned that fridges mean that, in some fundamental way, the frequency tells the price. What does this tell us about the fundamentals of price? After all, we already know that the price of electricity varies substantially over a day, over the week, and over seasons. This is driven by the current patterns of demand, well known to us.



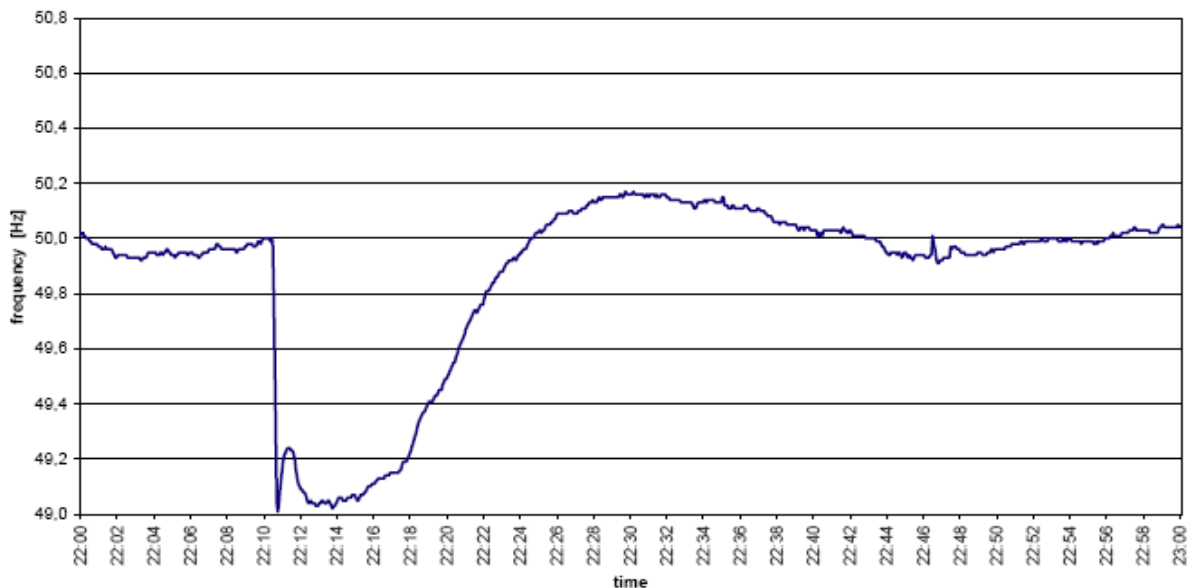
Source: National Grid

I see two significant things to learn. One is that the price varies continuously, as does demand. There are no natural discreet intervals. Thus market structures, by setting prices by the hour or the half hour, are imposing their trading needs on a product that is continuously varying in price. Already, across ENTSO, most frequency disturbances occur close to the hour, when market schedules change. The other is that the real price changes quite fast. To see this, it is useful to share two fairly extreme but still realistic scenarios.

4.3 Future scenarios

One scenario is the scrambling (emergency shut down) of a big nuclear power station. Suddenly, effectively within 10 seconds, the whole balance of the system is upset. For at least a day or two, “baseload” generation that was expected is no longer available. This upset is initially handled by the frequency response, ideally, of course, from fridges. Then generation that was too expensive to run before (was greater than the system marginal price) now has a market. If the plant can come on within about 10 minutes, it is truly a seller’s market, and prices can peak by factors of 10 or more. (In

Frequency Curve "West" on November 4, 2006, 22:00 - 23:00 hours
(Gauging: EnBW Transportnetze)



many cases, this price has been hedged by the System Operator, who may well have paid a fee to have generation physically ready. This physical hedging also usually buys the right to call on the generation at prices below such peaks.) Plant that is slower to start up, but that is cheaper to run may well then edge out the high priced plant, and prices will drop. But the effect on prices will be felt for days ahead. In this scenario, as price based on the fundamentals of electricity can change by an order of magnitude within a few seconds, and has a hangover lasting days. The diagram shows the frequency over the hour following the unzipping of the UCTE network in 2006 [11].

The other scenario is the passing of a storm front. Systems are increasingly incorporating generation, particularly wind, that is subject to the capriciousness of nature, and not to the control knobs of GSMs. When dispersed and diverse, the short variability of wind is very similar to that of demand. That is variations in output do not normally vary very fast, or very far, so can be handled by fridges. One exception to this is when many wind farms are spread across a storm front. Over a very short time, wind generation moves from its maximum, when the wind is as strong as the turbines can stand, to nearly zero, when the wind strength exceeds the design limits of the turbines, and they shut down. This is an extreme event and rarely happens, but does represent the challenge faced by electricity systems in the face of predictable weather trends, the precise timing of which is uncertain. Certainly, within a half hour or less, wind generation can drop by several GW. The half hour before, wind generation is plentiful, so the electricity price is low, but the half hour afterwards, electricity is scarce, so very expensive.

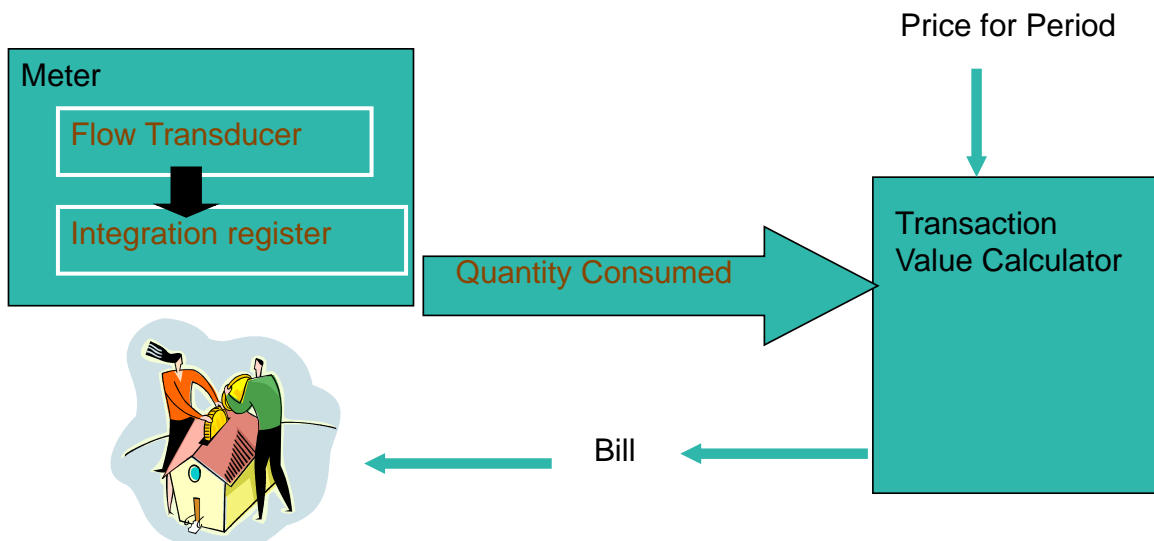
There really is no possibility that frequency **alone** can handle these scenarios.

Nor, indeed, can it handle the more conventional matter of peak and off peak pricing.

4.4 Metering

In my view, the ideal way to address these extreme events, as well as more normal operating modes, requires two (conflicting) features: first, a price that can vary **continuously** (or in very small increments), and second, knowledge of what the price is expected to be over the time ahead, perhaps up to a week. In market terms, **futures**.

There is a huge barrier in the way of the first, in the form of our weights and measures systems, including the so called Smart Meters now being rolled out. They work by incrementing the consumption within a register (sometimes one of several)



and then sending the accumulated result back to a central data processing system.

With simple meters, a register reading is sent back every quarter year or so, perhaps even weekly, and this reading (a few bytes perhaps) is then multiplied by the price, to calculate the bill. In practice there are a rich variety of complications, many of which are undoubtedly designed to confuse customers, and none of which add much by way of value.

If we now want to have the price varying every half hour then we need to send data for every half hour, so there will be 4334 times as much data to send.

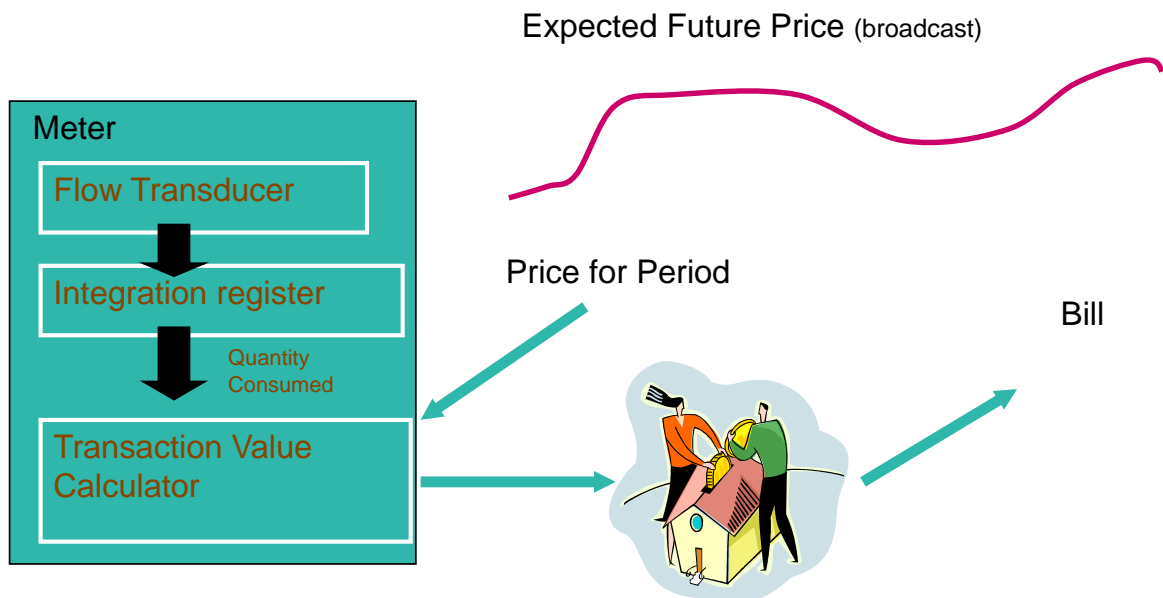
But a half hour is not enough, if we want it by the minute (say) we end up with nearly 1440 times more data.

Personally, I do not think a minute is fine scale enough, but data by the second means nearly 8 million times more data.

All this extra data is of very low value. It offers opportunity to invade privacy, and makes for big complex and wasteful data processing centres, well capable of messing up your bill. It does not help you or your appliances decide what to do.

4.5 Flowcost metering

If, on the other hand, your meter knows the price, it can do the multiplication, and, instead of (or as well as) accumulating your consumption, it accumulates your bill. It really makes no difference how fine grained the pricing intervals are.



How should the meter discover the price? Once convenient and very cheap way is to broadcast the parameters of a future price curve. This is a representation of a price that can vary continuously, and in complex ways, over the time ahead. Updates can be rebroadcast as forecasts become more precise, or circumstances vary. The meter will be aware of the latest broadcast, perhaps an hour or two earlier, and (since it has a clock) it will thus know what the price is **now**. It may be that a new communications system to meters will be useful, but it is not. Indeed, it is optional, and so can be a small addition to the evolving home networking environment.

Why a curve? The most fundamental reason is that, if there were a step change in price, then, sometimes, there will be lots and lots of devices, all waiting for the cheaper period, or all stopping at the start of an expensive period. It would create highly synchronised behaviour, and so put the whole system at risk. I do not believe it was a coincidence that the big unzip of UCTE arose just after the hour. So prices

changes must be gradual, exploiting and preserving the diversity of our demands by encouraging gradual change.¹³

4.6 The Need for Planning

But why worry about the future, what is wrong with a real time price?

Basically, it is not useful enough. We (or our devices) may be willing to sacrifice some consumption now, but for how long? I do not really care if the hot water goes off for a bit, so long as I can have my shower later. If I give up my control of my hot water system, can I really trust the electricity utilities not to leave me cold and dirty when I most need to be fresh and clean?

Almost all the devices that can be flexible in the timing of their consumption need to plan ahead (as indeed do most fossil fuelled generators). A laundry machine is an example. It has many tasks that, once started, cannot readily be interrupted. One a spin task has started, it needs to continue to completion. If a heating cycle has started, any interruption will at least involve a loss of efficiency – heat will be lost before the desired temperature limit is reached. I am no washing machine engineer, but there are other processes where the quality of the wash may be damaged by interruption, and may even damage clothes.

But a full wash involves many tasks, which do not necessarily need to run just after each other. A rinse can perhaps run some time after a wash. What is important, from a user's perspective, is that it completes by the time it is wanted, and does not mess up my clothes. I doubt that you would find a laundry machine control engineer (or salesman) who is willing to pass control of their operations to an electricity utility.

If, however, a laundry machine is aware of the expected price of electricity over the next few hours (or days), it can plan. When the laundry is loaded and the wash program selected, each user can be offered a menu of deadlines, each with an expected price. Broadly, the choice becomes: urgent or cheap. Cheapness will be possible if there are one or more periods where the price of electricity will be low, and the consumption can happen within them. If urgent enough, a user may be willing to pay an exorbitant price, say just after the storm front, to get it done quickly. It is consumer choice, not the utility's, and is closer to the transactions we use for other commodities, when we see the price at the time we buy. It is a fair, and effective way to influence large numbers of consumers to adjust consumption to match available (and renewable) generation.

I have given the example of a laundry machine, as this is perhaps the most complex of demand side appliances, and, if it works for that, it will work for simpler devices, dishwashers, house heating systems, hot water systems and, most fascinating of all, battery cars.

As an aside, the broadcast price curve can also be used for local generation exports. A domestic CHP may well be able to adjust its heating programme to take advantage of higher price periods.

4.7 A sustainable electricity system

What this concept promises is a powerful tool to influence the timing of demand (and small scale generation) over the minutes, hours, and perhaps days ahead. There are several profound implications.

¹³ Another way of thinking about this is that the potential ramp rate of demand – the maximum rate of change – is effectively infinite, and certainly far greater than that of any generation.

is be part of a case for breaking up ENTSO-E into smaller frequency domains, still interconnected, but less tightly, but I will not trouble you with my approach to this today.

And, finally, it gives a possible answer to the question we started with: are devices other than fridges and freezers good for frequency response. I said no, because I hope all these devices will be far too busy optimising their owner's electricity bills and playing a role in optimising the system to waste their capabilities on mere frequency response. They would be pretty poor at it, particularly in comparison with fridges and freezers. Further, if they were reasonably optimised, there would be times when they are not able to provide any response at all, and this will make the systems operators job even harder – not the aim at all.

Right, that is it. A gallop through possibilities for a future electricity system. I hope you found it worthwhile.

David Hirst

Comments and thoughts would be welcome david@davidhirst.com.

PS. Or perhaps for the next lecture

This paper has not been able to touch on the many issues flowcost throws up. The settlement and competition implications; the structure of trading contracts; the opportunity to exploit variable prices; how to allocate between players trading over a shared transport infrastructure; the bundling between fixed infrastructure cost and variable (fuel) costs; the issues from zero marginal cost generation, and how best to influence markets to adopt low carbon generation (FITs ROCs and the like). However, all these issues exist within the current system. What flowcost does do is make more obvious pre-existing contradictions.

I have come to the view that these issues cannot be resolved within the context of a competitive corporate utility free market. The contradictions are too deep, and the need for low carbon behaviour too great to carry on as we are. Given how deregulated free money markets have enabled the banks to do us so much harm, perhaps it is time to come up with a more rational and balanced governance systems. Electricity retailing is a good place to start.

Whatever the outcome, the fridges make it all a bit easier.

References

- [1] Soonee, S.K., S.R. Narasimhan, and V. Pandey. *Significance of Unscheduled Interchange Mechanism in the Indian Electricity Supply Industry*. in ICPODR. 2006: ITBHU http://www.nrlde.org/docs/Significance_of_UI.pdf
- [2] ENTSO-E, 2010 *Interconnected network of ETSO-E*. <https://www.entsoe.eu/index.php?id=77>.
- [3] UCTE, 2004 *Final Report of the Investigation Committee on the 28th September 2003 Blackout in Italy*, Union for the Coordination of the Transmission of Electricity www.ucte.org
- [4] NGC, 2008 *Report of the investigation into the automatic demand disconnection following multiple generation losses and the demand control response that occurred on the 27th May 2008*, NGC <http://www.nationalgrid.com/NR/rdonlyres/D680C70A-F73D-4484-BA54-95656534B52D/26917/PublicReportIssue1.pdf>
- [5] Agneholm, E., 1999 *Cold Load Pick-up*. <http://www.elkraft.chalmers.se/Publikationer/EKS.publ/Abstract/old/AgneholmD.abstract.html>.
- [6] Schweppe, F.C., et al., *Spot Pricing of Electricity*. 1988 Kluwer Academic Publishers.

- [7] Schweppe, F.C., *Power System Static State Estimation, Parts I, II & III*. IEEE Transactions on Power Apparatus and Systems., 1970. DOI: Vol 89: p. pp120-135.
- [8] Schweppe, F.C., *Frequency Adaptive, Power–Energy Re-Scheduler.*, 1979, Application No: US Patent Number 4,317,049,No.
- [9] Salter, S.H. *Off-Peak Load Control based on Voltage and Frequency Trends* in *Third International Conference on Metering, Apparatus and Tariffs for Electricity Supply*. 1977
- [10] Hirst, D.R., *ResponsiveLoad Controller (RL++)*, 2005, Application No: GB 00511361.8,No.
- [11] UCTE, 2007 *Final Report. System Disturbance on 4 November 2006*, UCTE [http://www.ucte.org/ library/otherreports/Final-Report-20070130.pdf](http://www.ucte.org/library/otherreports/Final-Report-20070130.pdf)