

Stabilization of grid frequency through dynamic demand control

Joe Short, David G. Infield, *Senior Member, IEEE*, and Leon L. Freris

Abstract—Frequency stability in electricity networks is essential to the maintenance of supply quality and security. This paper investigates whether a degree of built-in frequency stability could be provided by incorporating dynamic demand control (DDC) into certain consumer appliances. Such devices would monitor system frequency (a universally available indicator of supply-demand imbalance) and switch the appliance on or off accordingly, striking a compromise between the needs of the appliance and the grid. A simplified computer model of a power grid was created incorporating aggregate generator inertia, governor action and load-frequency dependence plus refrigerators with dynamic demand controllers. Simulation modelling studies were carried out to investigate the system’s response to a sudden loss of generation, and to fluctuating wind power. The studies indicated a significant delay in frequency-fall and a reduced dependence on rapidly deployable backup generation.

Index Terms—dynamic demand, dynamic demand control, demand-side management, frequency control, power system stability, intelligent load, refrigerators, integration of renewable power generation

I. INTRODUCTION

A. The need for frequency stability

FREQUENCY DRIFT, upwards or downwards, in a power system is the main indicator of the momentary imbalance between generation and demand. If, at any instant, power demand (taken in this paper to be active power only) exceeds supply, then the system frequency falls. Conversely, if power supply exceeds demand, frequency rises. The system frequency fluctuates continuously in response to the changing demand and due to the practical impossibility of generation being controlled to instantaneously track all changes in demand.

Frequency control of a power system endeavours to match supply as closely as possible to the time varying demand. Given the uncertainty in demand forecasting, this is achieved by ensuring that there is a sufficient quantity of spinning reserve generation on the bars. This generation provides frequency response by altering its output according to the system frequency. Response occurs in two phases. Primary response acts within tens of seconds to halt the decline (or rise) in frequency; secondary response then acts to restore the frequency to near nominal (50Hz in the UK) within tens of minutes. In the UK privatised network, system operators

pay considerable premiums to power generators who provide response [1].

With the recent power blackouts in the eastern United States, London and Italy, system stability is becoming an increasingly important issue to governments who are ultimately accountable for security of supply. Additionally, the UK Government’s announcement [2] to encourage 6 GW of off-shore wind capacity by 2010 indicates that a 20% penetration of variable renewable generation in the UK within twenty years is conceivable. To accommodate this level of variable generation, it has been predicted that balancing costs could rise significantly [3]. This paper examines the possibility of providing the required frequency stability through dynamic demand control thus avoiding the considerably more expensive alternative of additional backup generation.

B. Frequency-responsive loads

In the UK, the National Grid company (NG) has introduced a special tariff which rewards large-scale consumers who agree to provide a limited form of frequency response. Called Frequency Response by Demand Management (FCDM), the scheme involves the placing of certain large loads behind frequency-sensitive relays which isolate the load when the frequency falls below a pre-set level, often 49.7 Hz [4].

Oak Ridge National Laboratory in the United States is carrying out research for the US Department of Energy on the use of controllable loads to increase system security. The proposal is for a pager-controlled switch, attached to consumers such as air conditioners, which could be used to provide a load-reduction response within ten minutes [5].

Pacific Northwest National Laboratory, in a project supported by the US Department of Energy, is developing the ‘Grid Friendly Controller’ which can detect impending grid instability by monitoring extremely low-frequency signatures. In the future, such devices may be fitted to domestic and industrial consumers to provide load-shedding at times of excessive power-system stress [6].

The use of system frequency as an input signal to a load-controller was patented in 1979 in the United States. Called a Frequency Adaptive Power-Energy Re-scheduler (FAPER), the concept can be applied to any electrical consumer which needs electrical energy to function but which is not critically dependent on when that energy is supplied [7].

Econnect Ltd, a UK-based company which specialises in the integration of renewable electricity, has developed an “intelligent load controller” for use on small power grids that have very high penetrations of variable renewable generation. Frequency control is a challenging task on such “micro-grids”.

Joe Short is with the not-for-profit organisation, Dynamic Demand, The Hub, 5 Torrens Street, London, EC1V 1NQ (email: joe.short@dynamicDemand.co.uk)

D. G. Infield and L. L. Freris are with CREST (Centre for Renewable Energy Systems Technology), Department of Electronic and Electrical Engineering, Loughborough University, Loughborough, LE113TU, UK (email:D.G.Infield@lboro.ac.uk).

The Econnect devices monitor system frequency and use fuzzy logic to decide when to switch resistive loads (such as space and water heating) in order to maintain frequency stability.

ResponsiveLoad Ltd, a UK firm, is developing a frequency-dependent load controller similar to the FAPER but which uses various frequency limits to affect the probability of switching. In this way, the controller can move into different modes of operation, depending on the grid frequency at the time [8].

II. SCOPE

The purpose of the work described in this paper is to investigate the effects on a power grid of a large aggregation of frequency-responsive loads similar to the FAPER. The control of such loads is hereafter referred to as “dynamic demand control” or DDC. Examples of appliances that might be suitable for DDC are refrigerators, freezers, air conditioners, water heaters, some pumps, ovens and heating systems.

This study specifically investigates the suitability of domestic refrigerators for use as dynamically controlled loads. Refrigerators are on in all seasons, throughout the day and night, and therefore are available to participate in frequency control at all times. The total energy demand on the UK grid from domestic refrigeration has been estimated as 16.7 TWh per year [9] which amounts to an average load of 1.9 GW. This does not include industrial or commercial refrigeration. The refrigeration load is dependent on ambient temperature, winter load being approximately two thirds of that in summer. Daytime load is also slightly higher than that in the night time. Refrigerators are designed to handle considerable switching as they typically have an on/off cycle of the order of 15 minutes to 1 hour depending on characteristics and contents. Any additional switching caused by frequency control should not therefore present a problem.

In order to investigate the properties of refrigerators as dynamically controlled loads, this paper endeavours to answer as far as possible the following questions:

- Does the load controller have an undesirable effect on the proper operation of the appliance?
- What is the aggregated effect of a very large number of DDC loads on system frequency?
- Do they benefit stability?
- Are there any undesirable emergent properties?
- Do DDC loads provide a function similar to spinning reserve? Could they replace some spinning reserve?
- What benefits might DDC loads have on a system with a high penetration of wind power?

III. METHODOLOGY

A. Simulating the grid

A computer model of a power grid was created incorporating aggregate generator inertia, governor action and load-frequency dependence. Details of this simulation are shown in Appendices A, B and C. The grid simulation was first tested (without any dynamically controlled loads) to check that its behaviour agreed with past studies and data on system operation. Fig. 1 shows the system primary response to a step-change increase in demand from 30GW to 31GW. Immediately




Fig. 1. Simulated response to a step-change in demand from 30GW to 31GW at $T=0$

after the demand increase, the frequency begins to decline with energy being extracted from the spinning reserve and from other frequency sensitive loads on the system (released demand). This continues until the power deficit is zero. Due to governor delay, the system overshoots. A steady state is reached after 20 to 30 seconds when the system frequency settles to a lower value due to the assumed average 4% governor droop. This system behaviour matched closely that detailed in previous work [10].

Real power systems incorporate secondary response which returns the frequency to within normal limits within ten minutes or so. Secondary response was modelled using a simple proportional controller which continually altered the set-point of the generator’s governor by an amount proportional to the frequency error. The gain of this additional controller was set so that frequency was returned to nearly 50 Hz within about 15 minutes, a realistic response time [11].

B. Simulating a refrigerator

Detailed measurements were made of the internal temperatures of a typical commercial domestic refrigerator (the Hotpoint Iced Diamond RSB20) over several hours using four thermocouples and data-logging software. The temperature behaviour over a five hour period is shown in Fig. 2. A simplified mathematical model of the refrigerator was then created. The model comprised several thermally coupled masses as shown in Fig. 3. The method of difference equations was used to calculate heat flows between the thermal masses. Details of the thermal modelling are shown in Appendix D. Values for all the masses and areas of contact were estimated. Estimates for U-values were obtained by taking reasonable initial values (such as for ice, air or insulation) and then tuning them to achieve a match between the simulated temperatures and those of the real refrigerator. The resulting simulated temperature characteristics in Fig. 4 match adequately the experimental results in Fig. 2.

1) *Thermal storage capacity*: In order to achieve behaviour similar to the measured refrigerator, a mass of 2Kg was assumed for the freezer contents and a mass of 3Kg for the freezer box itself. These masses (and the specific heat capacities of the materials) are a crucial factor in determining how much load can be deferred.




Fig. 2. Temperature characteristics of a domestic refrigerator: Hotpoint Iced Diamond RSB20. (Accurate to within $\pm 1.3^{\circ}\text{C}$)




Fig. 3. Simplified thermal model of a refrigerator. Boxes show thermal masses. Arrows show heat flow between the masses

2) *Simulating large numbers of refrigerators:* For the purposes of this study, 1000 individual refrigerators were modelled separately. Each modelled appliance was randomised by altering every parameter to within $\pm 20\%$. This enabled the simulation of a real situation where each appliance would be at a different temperature, and at a different stage in its duty cycling. Duty-cycle lengths differed slightly due to the randomised thermal masses and U-values. During each program cycle, the temperature flows for each refrigerator was recalculated each was switched “on” or “off” according to the calculated “air space” temperature and current system frequency. Trials using greater than 1000 individual models were conducted and no significant difference in the results was found.




Fig. 4. Simulation output of an unmodified refrigerator (i.e. one with a standard thermostat)




Fig. 5. Normal operating strategy of a refrigerator




Fig. 6. Operating strategy of a dynamically controlled refrigerator

Every program cycle, the demand from all of the “on” refrigerators were added together and scaled up by a constant to represent a nominal demand of 1320 MW. This figure was chosen as it is equal to the maximum loss of generation that UK grid operators are required to plan for. (It is also below the 1900 MW of total domestic refrigeration that is in principle available for dynamic demand control).

This study assumes all refrigerators act like a Hotpoint Iced Diamond RSB20 - a relatively small “fridge-freezer”. In reality, much of the 1900MW of domestic refrigeration load will comprise deep freeze units which can be assumed to have greater thermal inertia than fridge-freezers. The load-deferment capability of the simulated system is likely therefore to be conservative.

C. *Simulating the DDC controller*

In a normal refrigerator, the control system (a thermostat) switches the compressor on and off to keep the temperature (usually the air temperature inside the main refrigerator compartment) within certain limits as shown in Fig. 5. In order for a large aggregation of dynamically controlled refrigerators to act as a frequency-dependent load, this strategy has to be altered to respond also to grid frequency.

Fig. 6 shows the strategy adopted for this study, where the allowed temperature-range varies linearly with system frequency. In other words, the thermostat control system is modified so that the two switching temperatures are decreased by an amount proportional to the current frequency deviation from 50Hz. This results in a situation where behaviour is identical to ordinary thermostats while the system frequency is exactly 50Hz. However, if frequency starts to fall, refrigerators will begin switching off, *starting with the coldest*, until frequency-fall is halted by the reduced demand.




Fig. 7. Simulated system frequency resulting from a sudden loss of 1320MW of generation which is then restored between $T=30$ minutes and $T=40$ minutes




Fig. 8. Simulated system frequency resulting from a sudden 1320MW loss of generation at $T=0$ hours which is restored during the ten minutes following $T=0.5$ hours and a "paying back" of this energy starting at $T=1$ hour

IV. RESULTS AND ANALYSIS

A. Response to a sudden loss of generation

The DDC refrigerators were added to the model along with a steady 36000MW load and a steady supply of 37320MW. No spinning reserve was provided on the DDC system. A traditional grid was also created for comparison, on which the refrigerators were replaced by a continuous 1320MW load so as to provide the same aggregate load. The traditional grid was provided with just enough spare spinning reserve to handle a loss of generation of 1320MW. Both systems were subjected to a sudden loss of generation of 1320MW at $T=0$ which was then restored over a ten-minute period starting at $T=0.5$ hours. (This was designed to reflect a realistic situation in which warm-steam backup generation might be available at 30 minutes' notice.) Fig. 7 shows the resulting system frequency. In the traditional system, frequency dropped dramatically at $T=0$ when the generation was lost. It was then stabilised and restored by primary and secondary response. At $T=0.5$ hours, the backup generation started to load up. (This caused a temporary rise in frequency due to a delay in high-frequency secondary response.) On the DDC system, where no spinning reserve was available, frequency-fall was halted instead by a rapid reduction in refrigeration demand. Frequency then began to fall slowly as refrigerators became warmer and started demanding more power. After the lost generation was restored, the frequency stabilised but remained below normal because the average temperature inside the refrigerators was by then higher than normal.

1) "Paying back" the energy: Overall, the DDC system was supplied with around 770MWh less energy than the traditional system. In reality, grid operators would need to "pay back" this energy in order to return both the frequency and average temperatures to nominal. A further simulation was therefore carried out in which an additional 500MW generator was instructed to come on the bars at $T=1$ hour, again increasing its power to maximum over a ten-minute period. This generator was instructed to stay loaded long enough to provide the 770MWh deficit. Fig. 8 shows the resulting system frequency. Fig. 9 shows the total demand on both systems throughout the event. As can be seen, the dynamically controlled refrigerators effectively deferred 770MWh of




Fig. 9. Total demand for the same simulation as above. In total 770MWh has been deferred

demand. This allowed the system to be restored using 500MW of extra stand-by generation (responding after 1 hour) instead of 1320MW of spinning reserve (responding in real time). This is expected to be a cheaper option characterised by a reduction in CO_2 emissions, but more research is required to quantify this.

It can also be argued that the use of DDC may cause the temperature-cycling in appliances to become synchronised (especially after a serious frequency excursion). And that this could have the overall effect of reducing the diversity of domestic loads with potentially negative consequences for power system operation such as large pick-ups in demand following significant load-shifting events. Though no loss-of-diversity issues were found in these simulations (Fig. 9 for example shows no significant pick-up), this too should be the focus of further research. This study did not investigate, for example, the ability of the appliances to respond to two major events one after the other.

2) What was the effect on the refrigerators?: The deferment of generated energy required an increase in refrigerator temperatures. Fig. 10 shows the maximum temperatures reached by any refrigerator in the simulation. As can be seen, freezer contents increased by 2.5°C following the half-hour power deficit.

3) Food safety: If dynamic demand were to be used in reality, a guarantee would be needed that temperatures would




Fig. 10. Maximum temperatures inside the main compartments of the simulated refrigerators, resulting from a deferring of 770MWh. Each point charted shows the maximum temperature found in any of the refrigerators.

not stray beyond limits required for food safety. As has been shown, in dynamic demand control, the refrigerator upper and lower temperature limits are coupled to the system frequency. When the frequency decreases, the upper and lower temperatures increase proportionally. The grid frequency variation is subject to tight legislation which limits long-term excursions. By correctly setting the constant of proportionality, (the gradient of the control function shown in Fig. 6), it could be arranged that, as long as the grid frequency was within legal bounds, the refrigerator temperatures would also be within acceptable bounds for food safety.

4) *An alternative strategy:* It has been shown that dynamic demand control could open the door to a possible alternative strategy for dealing with power imbalances. Rather than have all the necessary reserve spinning on the system, it might be possible to have more standby capacity off the bars. The dynamic demand control provides a delay in frequency fall during which time more generation could be scheduled.

B. Effect of DDC in a power system with a large wind power input

The fast response available from DDC is likely to be of particular benefit to power systems with large inputs of wind power or other time-variable renewable energy sources. In order to investigate this potential, a simulation of a scenario with a substantial wind power input was undertaken.

Wind speed data from 23 UK sites were used in the simulation [12]. A 50-hour period containing considerable variability was chosen, though this sample period did not happen to coincide with the time of highest wind speed. The data sets were put through a computer algorithm which superimposed realistic turbulence (short-term wind speed variation) onto the ten-minute wind speed data. For details of this procedure, see Appendix E. The final wind speed records were averaged into one-minute bins. It was further assumed that the variation in wind speed and physical separation of the wind turbines in each site would smooth second-to-second variations in power. The modelled power system would therefore only experience variations of power over the order of a minute.




Fig. 11. Simulated wind power using measured wind speed data from 23 UK sites from 00:00 30/08/1991

For each site, the power output was calculated on the assumption that a windfarm comprising 150 4MW variable speed wind turbines was present at each site. A purely cubic power-windspeed relationship was assumed with a cut-in wind speed of 2 m/s, a rated wind speed of 15m/s and a furling speed of 25m/s. The output power from the 23 sites were added together to give a total peak generation capacity of 13.8GW. The maximum and minimum power reached in the 50 hours of the simulation was 4.6GW and 0.6GW respectively. During this period, the largest sustained drop in wind power occurred between hours 37 and 42 when 3.5GW of wind was lost in 4.5 hours. The fastest short-term drop occurred during hour 36 when the output fell by 0.8GW in 1 minute. The synthesised wind power trace is shown in Fig. 11.

A simulation was set up with the above wind generation connected, along with 2000MW spinning reserve (providing both primary and secondary response), 1320MW of dynamically controlled refrigeration and a constant 36000MW load. Enough base generation was provided such that the spinning reserve was half-loaded when the wind power was at its average for the 50-hour period.

As before, a control grid was created with no DDC. However, a level of spinning reserve had to be chosen in order for the two systems to be compared. First the DDC simulation was run to measure the standard deviation in grid frequency over the entire 50 hours. This was found to be 0.13Hz and provided a very rough guide to the level of frequency control. The non-DDC system was then given just enough spinning reserve to achieve the same standard deviation in frequency. This amount turned out to be 3170MW. The resulting grid frequency from the simulation for the two systems is shown in Fig. 12.

1) *Smoothing of grid frequency:* Fig. 13 shows a detail of the first ten hours and illustrates how the DDC system considerably reduced the variation in system frequency even though it was operating with substantially less reserve. This is because the simulated controllers reacted more quickly than the generator governor to changes in frequency.

Fig. 14 depicts an important difference between the two systems in terms of the time variation of power output from the spinning reserve generators. The fast-acting controllers handled most of the short-term smoothing, allowing the spinning




Fig. 12. System frequency. Simulation output with wind on the system. Compares 1320MW of DDC plus 2000MW of reserve (black) with no DDC but 3170MW of reserve (grey). Standard deviation in grid frequency was the same for the two systems (0.13Hz)




Fig. 13. System frequency. Detail of first ten hours

reserve generator to change its output much more gradually. In a real situation, this might provide considerable benefit in terms of wear and tear on the generator.

2) *Sudden drop in wind power:* At $T=40$ hours, the wind power had declined dramatically and reached a trough. As can be seen from Fig. 15, both system frequencies fell below the operational limit of 49.8Hz. In a real-life situation, either more spinning reserve would need to be on the system, or some fast-responding (and therefore possibly more costly or inefficient) generation would have to be rapidly scheduled. One option for the UK might include bringing on line the pumped storage




Fig. 14. Spinning reserve output. Detail of first ten hours




Fig. 15. System frequency following a prolonged decline in wind power

facility at Dinorwig which can reach full output of 1.7GW in 16 seconds [13]. But this may compromise the facility's ability to deliver other deferment services.

However, the system with dynamic demand control provided a considerable breathing time. Frequency did not fall below the operational limit until nearly two hours after the non-DDC system. A team of engineers operating the power network may therefore be given a wider choice of generation with which to balance the system including slower acting (and therefore possibly cheaper and more efficient) options. Also the delay may allow generation to be scheduled more cost effectively through the electricity market which operates a "gate closure" time of half an hour in advance of any particular generation slot.

V. CONCLUSIONS

It has been shown that an aggregation of a large number of dynamically controlled loads has the potential to provide significant added frequency stability to power networks, both at times of sudden increase in demand (or loss of generation) and during times of fluctuating wind power.

The devices, if incorporated on a real system, could offer some of the services currently provided by spinning reserve. It is possible that dynamically controlled loads could even be used to replace some of the spinning reserve. The amount they could replace depends on the extent to which low-magnitude but long-term frequency excursions can be tolerated, and on the amount of slower-acting backup generation available. The simulations in the paper show that this amount would be of the order of the total size of the dynamically controlled loads connected.

The use of dynamic demand control may provide a more cost effective strategy for scheduling power because a significant delay (for periods of an hour or more) in the fall of frequency at times of imbalance is introduced. This delay is dependent on the imbalance on the system and the total load under dynamic control. Dynamic demand control has the potential to allow power systems operators to widen their choice of backup generation and to include generators which take more time to bring on line.

When operating with fluctuating wind power, it has been shown that dynamically controlled loads have the potential to offer considerable frequency smoothing and a significant

reduction in governor activity of spinning reserve generators. Dynamic demand control has also been shown to have the potential to allow the system to “ride through” short-term drops in wind power.

The potential demand that could be operated under dynamic control is considerable. In this study, 1320MW was assumed for technical reasons, though the likely amount is much higher. It is worth noting that deep freeze units were not considered in this study, nor were industrial and commercial refrigeration, though there is no reason in principle why these could not be operated in the same way. Also, other loads, such as air conditioning or water heating, could provide similar services. In principle, the potential is several GW in the UK. Similar considerations apply to other countries, although the air-conditioning load may well be larger as a proportion of total load, giving added flexibility.

Deliberately increasing the thermal storage capacity of refrigerators by, for example, introducing phase-change materials into their structure could multiply the potential of dynamic control by a large factor.

A. Economic viability of DDC

The reduction of the amount of spinning reserve needed through the use of dynamically controlled loads may result in considerable cost savings. In the UK for example, the National Grid Company spends approximately £80M per year on frequency response [14]. It is reasonable to assume that this amount could be steadily reduced with the onset of dynamically controlled loads. If so, it should be possible in principle to make alterations to the market in frequency response to produce savings and to divert some of these savings towards incentives for the incorporation of dynamic demand control into appliances.

It is likely that for this long-term approach to be adopted the input of Government and Regulators will be required, at least at the outset. The potential public benefits of dynamic demand control are clear: a more cost effective power system with greater frequency-stability and a capability to absorb larger amounts of variable power from variable renewable generation.

APPENDIX

A. Modelling a generator

A generator connected to the grid is designed to respond to a drop in grid frequency by increasing its output. If the generator is part-loaded (on spinning reserve) this governor action will help stabilise the system. Typically, generator governors are designed with a 4% “droop” characteristic [10]. This means the generator output will increase by 100% for a 4% drop in frequency. The generator’s nominal output can be altered by changing the “set point”, i.e. the frequency at which the generator will output zero MW.

The actual dynamics of real generator sets are highly complex and differ considerably from power station to power station [15]. A truly accurate model would need to comprise many types of generator, each incorporating a control system with several time constants. However, it has been shown that a governor with a droop characteristic can be adequately

modelled as a proportional controller [16]. In this simulation, the first step is to calculate the generator’s target power output, P_{TAR} , using the droop characteristic:

$$P_{TAR} = \left(\frac{f_{SP} - f}{0.04 \times f_{NOM}} \right) P_{MAX} \quad (1)$$

where f_{SP} is the generator’s set point in Hz, f is the current grid frequency, f_{NOM} is the nominal grid frequency (50 Hz for the UK) and P_{MAX} is the generator’s capacity in MW. The next step is proportionally to reduce the error between P_{TAR} and the actual output, P , using:

$$P_{(t+dT)} = P_{(t)} + (P_{TAR} - P_{(t)}) \cdot G \cdot dT \quad (2)$$

where G is the governor gain. A realistic value for G was found to be 0.3 as this resulted in a settling time of the order of 15 to 20 seconds after a step-change in load, a reasonable value [10].

For the simulation, the above equations were incorporated into a “power station” C++ class, any number of which could be connected to the grid.

For simplicity, the total amount of spinning reserve on the system is modelled by a single governor-controlled generator of sufficient size. Also, the total amount of base generation is modelled by an additional very large generator, but on fixed full output.

B. Modelling released demand

Many loads on the grid consist of rotating machines. Hence there is a built in frequency-dependence caused by the fact that these machines slow down as the frequency drops, and thus consume less power. It has been found empirically that the total active power demand decreases by 1-2% for a 1% fall in frequency depending on a load damping constant, D [16]. (N.B., it is logical to assume that D is slowly decreasing over years as more and more mechanical systems are connected behind power electronic interfaces.) This change in power is called “released demand”. It is treated in the simulation as an injection of active power, P_R , calculated using:

$$P_R = -D \cdot P_L \cdot \left(\frac{f - f_{NOM}}{f_{NOM}} \right) \quad (3)$$

where D is assumed here to be 1.0 and P_L is the total load that would exist if there were no built in frequency dependence.

C. Modelling the grid’s inertial energy store

As already stated, the grid frequency falls as all the spinning machines on the system begin to slow down. In effect, the demand deficit is being met by extracting energy from the rotational inertia of all the generators (and spinning loads). The fall in frequency will continue until the demand deficit is met by a combination of released demand and increased generation due to governor response.

In the simulation, all the inertia is assumed to be stored in a single flywheel of moment of inertia, I , rotating at grid

frequency, $\omega = 2\pi f$ rad s⁻¹. The total energy stored is therefore:

$$KE = \frac{1}{2} \cdot I \cdot (2\pi f)^2 \quad (4)$$

The inertial storage capacity of a power system is measured by an *inertial constant*, H , which is the number of full-output seconds of energy stored (assuming nominal frequency). H typically varies between 2 and 8 seconds [16]. Several simulations were carried out with H values varying from 2s to 8s and no significant difference in the apparent effectiveness of the dynamic demand devices was found. This was because, although the settling-time of the system changed slightly, the final frequency for a given power imbalance remained the same. For this study, H is assumed to be 4s throughout. I for the system is calculated once at the start of the run:

$$I = \frac{2 \cdot P_{GMAX} \cdot H}{\omega_{NOM}^2} \quad (5)$$

where P_{GMAX} is the total generation capacity.

For each step of the simulation, the total power surplus, P_S is then calculated:

$$P_S = P_G + P_R - P_L \quad (6)$$

where P_G is the total generation, P_R is the released demand and P_L is the load. Clearly, P_S is the power going into the inertial energy store. Given that for each simulation time slice, dT , energy must be conserved, then:

$$KE_{(t+dT)} = KE_{(t)} + P_S \cdot dT \quad (7)$$

Hence:

$$\frac{1}{2} \cdot I \cdot \omega_{(t+dT)}^2 = \frac{1}{2} \cdot I \cdot \omega_{(t)}^2 + P_S \cdot dT \quad (8)$$

which provides a difference equation for calculating the new frequency for each step of the simulation:

$$\omega_{(t+dT)} = \sqrt{\omega_{(t)}^2 + \frac{2 \cdot P_S \cdot dT}{I}} \quad (9)$$

The simulator then goes through the following steps for each time slice, dT :

- Calculate P_L by summing the connected loads
- Calculate P_G by summing the total generation
- Calculate P_R using Equation 3
- Calculate P_S using Equation 6
- Calculate the new ω (and f) using Equation 9

D. Thermal simulation

The simulation runs through each thermal link calculating the energy flow, dE , from Mass 1 to Mass 2 for the given time slice, dT using:

$$dE = UA(T_1 - T_2)dT \quad (10)$$

$$E_1 = E_1 - dE \quad (11)$$

TABLE I
THERMAL ELEMENT CHARACTERISTICS

Name	Mass (Kg)	Temp (°C)	SHC (J/Kg °C)
Freezer box	3	-15.7	450
Freezer contents	2	-15.7	3000
Fridge air space	0.5	3.5	1000
Fridge contents	2	3.5	2000
Room	∞	25	n/a

TABLE II
HEAT TRANSFER COEFFICIENTS FOR THERMAL PATHS

Link	Mass 1	Mass 2	Area (m ²)	U-Value (W/m ² K)
0	Freezer box	Freezer contents	0.15	12.5
1	Freezer box	Fridge air space	0.225	6
2	Fridge air space	Fridge contents	0.35	12.5
3	Fridge air space	Room	2	0.6

$$E_2 = E_2 + dE \quad (12)$$

where U is the link's U-Value, A is its area, T_1 and T_2 are the temperatures of its connected thermal masses and E_1 and E_2 their stored energies. In the next cycle, new temperatures are calculated for each mass using:

$$T_n = \frac{E_n}{S_n m_n} \quad (13)$$

where E_n is the heat energy stored by mass n , S_n the specific heat capacity (J Kg⁻¹ °C⁻¹) and m_n the mass (Kg).

The parameters chosen are shown in Table I and Table II. These were then randomised to simulate the effect of many realistic refrigerators.

E. Adding synthetic turbulence to measured wind speed

The computer program [17] used existing algorithms [18] to add realistic turbulence to the averaged wind speed data. The technical information reproduced here was provided by the program's author.

The program's input was a time series of wind speed data measured at longer time intervals, e.g. 10 minutes, 20 minutes or 1 hour (10 minutes in this case). The program user chooses a shorter time scale for interpolation, in this case 1 second.

The program first works out the autocorrelation value at the shorter (1 second) time scale, assuming a Kaimal turbulence spectrum, and a turbulence intensity. The turbulence intensity was either measured (time varying), or assumed to be 16.85% if not available (a value taken from the average of others).

The program then performs a random walk with a Gaussian probability density function, using the calculated turbulent autocorrelation, the turbulence intensity, and generated random numbers—1 point per second in the total time span of the measured data. The data points in the random walk are scaled by the measured wind speeds. The scaling factor is interpolated between adjacent measured points. Thus the random walk is turned into a smoothly changing wind speed data stream that follows the measured wind speed data.

REFERENCES

- [1] *Response Prices and Curves*, National Grid Transco, October 2002, accessed online at www.nationalgrid.com/uk/indinfo/.
- [2] DTI, "Hewitt announces biggest ever expansion in renewable energy," Press release, July 2003.
- [3] G. Strbac, "Quantifying the system costs of additional renewables in 2020," Manchester Centre for Electrical Energy, UMIST, Tech. Rep., October 2002, report to the Department of Trade and Industry.
- [4] A. Malins, "Demand side developments," ser. Operations Forum. National Grid Transco, March 2003.
- [5] B. Kirby and M. Ally, "Spinning reserve from supervisory thermostat control," ser. Transmission Reliability Research Review. Washington D.C.: US Department of Energy, December 2002.
- [6] PNNL, "Grid friendly controller helps balance energy supply and demand," Promotional flier from Pacific Northwest National Laboratory accessed online at <http://gridwise.pnl.gov/> in February 2006, 2002.
- [7] F. C. Schweppe, *US Patent No 4317049*, Massachusetts Institute of Technology.
- [8] D. Hirst, *UK Patent No GB2361118*, ResponsiveLoad Ltd.
- [9] DTI, "UK energy sector indicators," 2003. [Online]. Available: <http://www.dti.gov.uk/energy/inform/indicators2005.pdf>
- [10] *The Grid System - Power and Frequency Control*, National Grid Company, 1991, open Learning Materials.
- [11] M. Arthur, "Frequency response," ser. Operations Forum. National Grid Transco, November 2002.
- [12] Data compiled by J. Barton, Centre for Renewable Energy Systems Technology, Loughborough University.
- [13] Inside Dinorwig website, <http://www.fhc.co.uk/dinorwig/d2.htm> (accessed 6 September, 2003).
- [14] Ofgem, "NGC system operator incentive scheme from April 2005," December 2004, see page 58 for ancillary services costs.
- [15] O. I. Elgerd, *Electric Energy Systems Theory: An Introduction*. McGraw-Hill, 1971, p. 325.
- [16] P. Kundur, *Power System Stability and Control*. McGraw-Hill, 1993, pp. 581-592.
- [17] Program written by J. Barton, Centre for Renewable Energy Systems Technology, Loughborough University.
- [18] Energy Research Unit at Rutherford Appleton Laboratory, Engineering Design Tools for Wind Diesel Systems, final report for CEC Contract JOUR-0078. Volume 8 - Logistic Package: Program Documentation.

PLACE
PHOTO
HERE

David Infield (M 2003, SM 2005) was born in Paris on the 30th March 1954.

He was brought up and educated in England, gaining a BA in Mathematics and Physics from the University of Lancaster and a PhD in Applied Mathematics from the University of Kent at Canterbury. He worked first for the Building Services Research and Information Association, Bracknell, UK, on solar thermal system design, and then for the Rutherford Appleton Laboratory, from 1982 until 1993, on wind energy systems and electricity supply modelling. He is currently Director of CREST (Centre for Renewable Energy Systems Technology) and Professor of Renewable Energy Systems with the Department of Electronic and Electrical Engineering at Loughborough University, UK. He is a senior member of the IEEE.

PLACE
PHOTO
HERE

Leon Freris was born in Athens, Greece

on the 8th February 1934. He gained a BSc, an MSc and a PhD in Electrical Engineering from the University of London and a DIC from Imperial College. In 1960 he joined Imperial College, London, as a lecturer and subsequently became the head of the power systems section in the department of Electrical Engineering. In 1995 he retired from Imperial College and joined CREST (Centre for Renewable Energy Systems Technology) at Loughborough University, UK as a visiting Professor in Renewable Energy Systems. His research interest ranged from HVDC transmission, and power system optimization to renewable energy technologies and particularly the electrical aspects of wind power systems. He is a founding member of the British Wind Energy Association and a Fellow of the Institution of Electrical Engineers.

PLACE
PHOTO
HERE

Joe Short was born in London, UK on the 19th December 1970. He graduated in Physics/Philosophy at the University of York. After being trained in news agency journalism at Reuters, and 10 years working on environmental communications with advocacy groups such as Friends of the Earth, Joe gained an MProf in Sustainable Development with the charity Forum for the Future. In 2003 he gained an MSc in Renewable Energy Systems Technology at Loughborough University where he began investigating dynamic demand control. In 2004, Joe founded

Dynamic Demand (www.dynamand.co.uk), a not-for-profit coalition dedicated to promoting the technology as a potential carbon-saving measure amongst policy makers and companies in the UK.