A Probabilistic Method of Modelling Energy Storage in Electricity Systems with Intermittent Renewable Energy

by

John Barton

Doctoral Thesis
Submitted in partial fulfilment of the requirements
for the award of

Doctor of Philosophy of Loughborough University
28th February 2007

Loughborough University

Thesis Access Form

Copy No…………1..………………….**Location**………………………………………………….………..……...…

Author…………...John Barton……………………………………………………………………………..…….

Title….A Probabilistic Method of Modelling Energy Storage in Electricity Systems……………….. with Intermittent Renewable Energy

Status of access OPEN / RESTRICTED / CONFIDENTIAL

Moratorium Period:…………………………………years,ending…………../…………200……………………….

Conditions of access approved by (CAPITALS):…………………………………………………………………

Director of Research (Signature)…………………………………………...…………………………………...

Department of………Electronic and Electrical Engineering…………………………………………………

Author's Declaration: *I agree the following conditions:*

OPEN access work shall be made available (in the University and externally) and reproduced as necessary at the discretion of the University Librarian or Head of Department. It may also be copied by the British Library in microfilm or other form for supply to requesting libraries or individuals, subject to an indication of intended use for non-publishing purposes in the following form, placed on the copy and on any covering document or label.
*The statement itself shall apply to **ALL** copies:*

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

Restricted/confidential work: All access and any photocopying shall be strictly subject to written permission from the University Head of Department and any external sponsor, if any.

Author's signature…………………………………….Date……………………………...…………..……...

users declaration: for signature during any Moratorium period (Not Open work): ***I undertake to uphold the above conditions:***			
Date	Name (CAPITALS)	Signature	Address

CERTIFICATE OF ORIGINALITY

This is to certify that I am responsible for the work submitted in this thesis, that the original work is my own except as specified in acknowledgments or in footnotes, and that neither the thesis nor the original work contained therein has been submitted to this or any other institution for a higher degree.

……………………………………………. (Signed)

……………………………………………. (Date)

Abstract

A novel probabilistic method has been developed for modelling the operation of energy storage in electricity systems with significant amounts of wind and solar powered generation. This method is based on a spectral analysis of the variations of wind speed and solar irradiance together with profiles of electrical demand. The method has been embodied in two Matlab computer programs:

Wind power only: This program models wind power on any time scale from seconds to years, with limited modelling of demand profiles. This program is only capable of modelling stand-alone systems, or systems in which the electrical demand is replaced by a weak grid connection with limited export capacity.

24-hours: This program models wind power, solar PV power and electrical demand, including seasonal and diurnal effects of each. However, this program only models store cycle times (variations within a time scale) of 24 hours. This program is capable of modelling local electrical demand at the same time as a grid connection with import or export capacity and a backup generator.

Each of these programs has been validated by comparing its results with those from a time step program, making four Matlab programs in total. All four programs calculate the power flows to and from the store, satisfied demand, unsatisfied demand and curtailed power. The programs also predict the fractions of time that the store spends full, empty, filling or emptying.

The results obtained are promising. Probabilistic program results agree well with time step results over a wide range of input data and time scales. The probabilistic method needs further refinement, but can be used to perform initial modelling and feasibility studies for renewable energy systems. The probabilistic method has the advantage that the required input data is less, and the computer run time is reduced, compared to the time step method.

Keywords: Energy storage, Modelling, Wind power generation, Photovoltaics, Stand alone power systems, Interconnected power systems, Voltage control, Embedded generation

Table of Contents

Dedicated to the two billion people who at the start of the twenty-first century still lack access to electrical power

1 Introduction

1.1 Background

Electrical energy, unlike most commodities, cannot be directly stored. It must therefore be generated as it is used and electrical power must be continuously balanced to match load with generation. In most conventional power systems, this is achieved by adjusting the rate of generation. Redundancy in the transmission network, together with spinning reserve and standby generating capacity accommodate increases or decreases in demand, or a failure of any one power station, substation or transmission line (Halliday 1988, Watson, Landberg & Halliday 1994).

All electricity systems use some forms of energy storage, including inter-seasonal energy storage, to accommodate variations in electricity demand. In conventional systems this is achieved by storing the fuel used for primary generation. This may be in the form of coal, oil, gas, nuclear fissile material, or hydroelectric dams with inter-seasonal storage of water.

1.2 Renewable Sources of Energy

Renewable sources of energy are likely to provide an increasing proportion of electrical power generation during the twenty-first century. The change will be driven by concerns of fossil fuel resource depletion, fuel supply security, the risks and political unpopularity of nuclear power, and above all the threat of global warming caused by carbon dioxide from the burning of fossil fuels.

Studies by the Royal Commission on Environmental Pollution (RCEP 2004) conclude that the UK and other industrialised countries must reduce their emissions of carbon dioxide by 60% by the year 2050 if atmospheric carbon dioxide levels are to be stabilised at 550 parts per million and catastrophic climate change is to be avoided. At the same time, approximately one third of the world's population, 2 billion people, still do not have access to electrical power or other affordable or adequate energy sources (UNDP, UNDESA & WEC 2000). If energy supplies are to be maintained at a comfortable level, and poorer countries are to improve their standards of living then carbon neutral and low carbon energy sources must be extensively used. All RCEP scenarios show a strongly increasing role for renewable energy during the 21st century (RCEP 2004).

There is also a growing consensus amongst many geologists that oil reserve depletion is likely to cause higher oil prices over the next decade or two (ASPO 2005). This will force a move from conventional, 'easy' oil to unconventional sources such as tar sands, greater use of natural gas in the medium term, but eventually to an economy that does not rely on fossil fuels at all. Some experts think this will take the form of a 'hydrogen economy' (Rifkin 2002).

The renewable sources of energy: biomass, water power, solar power, wind power and geothermal energy currently form only a small percentage of world primary energy supply; 13.8% in the year 2000 (International Energy Agency 2002). Most of this renewable energy was used in the form of traditional biomass for cooking and heating in Africa, Asia and South America. The second and third largest renewable energy sources are hydroelectric power and geothermal power respectively.

1.3 Intermittent Renewable Sources of Energy

Two 'new ' sources of renewable energy, wind (at 0.026%) and solar power (at 0.039%) currently supply a tiny fraction of world energy, but have grown much more rapidly than other sources between the years 1971 and 2000. Wind power has grown at 52.1% and solar power has grown at 32.6% (International Energy Agency 2002).

If we assume a continued compound growth rate for each of these sources of 30%, then their combined share of total energy supply would grow from 0.065% to 100% by the year 2028! Projected growth rates are uncertain, but these sources are likely to be very significant in the next few decades. Biomass energy is limited by the land area available to grow energy crops, and by the size of biomass waste streams and agricultural residues. Hydroelectric power and geothermal power are limited by the availability of suitable sites. Wind and solar power are the only renewable sources that could in principle supply all the world's energy needs.

If wind and solar power are to supply a large proportion of primary energy, most of that energy will be captured in the form of electricity by wind turbines, photovoltaics and solar thermal electricity, because electricity is such a versatile energy form. But unlike biomass and hydroelectric power that can be dispatched to supply energy as required, wind power and solar power are intermittent and variable. They are only available at the times and in quantities that weather conditions permit, but as stated above, electrical energy has to be generated as it is used.

Conventional electricity systems can accommodate at least 5% of wind powered generation, and probably more, with few adaptations (Milborrow 2002). The variability of small amounts of wind or solar power can be accommodated in the same way as variations in electrical load: by varying the electrical generation from fossil fuel generators. If renewable sources are to play their part in a 60% reduction in carbon dioxide emissions, the fraction of intermittent renewable generation will increase well above 5%. This will require a radical new approach to electrical power systems, including novel means of storing energy.

1.4 The Value of Energy Storage

Energy storage may perform many different tasks in an electricity system, on many different time-scales (Barton, Infield 2004). Individually, each task has some monetary value but an energy store that performed just one task would probably not be cost effective. Energy stores that perform two or more functions are more likely to be worthwhile. The following study attempts to add up the total value of a store, based on all the tasks that a store may perform and all the possible revenue streams. It is based on an approach used in a previous technical paper (Bryden, Macfarlane 2000). The resulting equations 1.1 and 1.2 are not very useful in their present form, but it are a starting point for the computer models of this PhD, and a reference or checklist for future specific studies of energy systems.

First, here is a list of parameter names and their meanings, just for this section of the thesis:

Name	**Nomenclature**
Capital costs of fossil fuel generation	CF
Capital costs of renewable generation	CR
Capital costs of storage	CS
Operating costs of fossil fuel generation	OF
Operating costs of renewable generation	OR
Operating costs of storage	OS
Fossil fuel costs	F
Energy Losses of storage	LS
Cost of standing reserve (Standby generation), excluding fuel costs	SB
Cost of spinning reserve (Generators held at part power), excluding fuel costs	SR
Cost of governor control (Frequency control), excluding fuel costs	GC
Cost of stops and starts of fossil-fired generation (wear and tear and operating)	SS
Cost of transmission line repair	TR
Cost of Transmission line investment	TI
Cost of Transmission losses	TL
Cost of support during line faults	LF
Cost of local voltage support	LV
Cost of reactive power	Q
Cost of black start	BS
Electricity used	E

Subscript Meaning	**Nomenclature**
System with fossil and renewable Generation	FR
System with fossil and renewable Generation and energy storage	FRS
Total amount	T

The net cost of electricity without storage is:

$$UC_{FR}=(CF_{FR}+CR_{FR}+OF_{FR}+OR_{FR}+F_{FR}+SB_{FR}+SR_{FR}+GC_{FR}+SS_{FR}+TR_{FR}+TI_{FR}+LF_{FR}+LV_{FR}+TL_{FR}+Q_{FR}+BS_{FR})/E_T \quad (1.1)$$

The net cost of electricity with storage is:

$$UC_{FRS}=(CF_{FRS}+CR_{FRS}+CS_{FRS}+OF_{FRS}+OR_{FRS}+OS_{FRS}+F_{FRS}+LS_{FRS}+SB_{FRS}+SR_{FRS}+GC_{FS}+SS_{FRS}+TR_{FRS}+TI_{FRS}+LF_{FRS}+LV_{FRS}+TL_{FRS}+Q_{FRS}+BS_{FRS})/E_T \quad (1.2)$$

There are three completely new terms when using storage: the capital costs of storage, CS, the operating costs of storage, OS and the electrical losses of storage, LS. These obviously add to the cost of electricity.

The total amount of electricity consumed by customers, E_T is assumed to remain the same when storage is used. For a given amount of renewable generation, the costs of that renewable generation, CR and OR are also unchanged. All other costs are reduced by the use of storage on the system because storage performs the following functions:

1. Energy storage is used for peak-lopping and so reduces the total amount of fossil fuelled generation required, and reduces its capital cost, CF. The remaining fossil fuelled plant is operated at improved load factor.
2. Less fossil fuelled plant is required (see previous point), and therefore operation and maintenance costs, OF are lower.

3. More of the renewable energy generated electricity is used and less is wasted at times of over-supply. Fossil fuel costs are therefore reduced.
4. The fossil fuel costs, F are reduced because the merit order of the remaining fossil fuelled generation is improved. The energy storage tends to replace low merit order, peaking plant, with high fuel costs and low capital costs.
5. The fossil fuel costs, F are further reduced because less plant is operated at part power (spinning reserve), held at standby (burning some fuel but not generating) or held on governor control (part power, continuously changing power output).
6. Other operation and maintenance costs of standby generation, SB and spinning reserve generation, SR are reduced because less of these reserves are required.
7. When energy storage is used for short-term frequency control, less fossil fuelled plant needs to be held on governor control. Its wear and tear cost is therefore lower.
8. Energy storage is used to level the load and level the power output from renewable sources. The remaining fossil-fuelled plant is therefore cycled less. Its wear-and-tear and operating costs due to starting and stopping, SS are reduced.
9. If energy storage is placed at the right locations, it can reduce the total amount of electrical energy that is transmitted from one location to another, or sometimes just make the transmitted power more constant. Here, transmission is taken to include the distribution network. Energy storage in the distribution network could be particularly effective in this function. The strength and number of transmission lines can be reduced without exceeding voltage or thermal limits, and so their investment costs, TI, and repair costs, TR can be reduced. Alternatively their transmission losses, TL can be reduced.
10. Energy storage can support local voltage during line faults, and so maintain a more reliable supply, reducing costs to consumers and electricity companies.
11. Energy storage can be used to control local line voltage by the timely injection or absorption of active and/or reactive power. It can therefore

reduce the cost of other equipment used for that purpose, such as static VAR compensators.

12. Energy storage can continuously supply reactive power, Q (using its power interface electronics). Reactive power would otherwise have to be supplied by synchronous machines or static VAR compensators.
13. In the event of a major power failure across a large proportion of the grid, energy storage may supply enough power to get the large generators operating again. This ancillary service is known as black start, BS.

The economic value of each of these functions is hard to quantify without a computer model of the electricity system. It may never be possible to identify each of these values, because each component of an electricity system is inter-connected and inter-dependent. A system without energy storage would be optimised differently from a system with energy storage, since every investment decision affects the economics of every other decision.

However, even a simple analysis (Bryden, Macfarlane 2000) does appear to show that renewable sources of generation can work synergistically with energy storage. Together they bring a much larger benefit in terms of fossil fuel savings and capital cost savings of other generation than either renewables or energy storage alone.

1.5 System Modelling

When a developer designs an electricity system including intermittent renewable energy, he or she wants to know how well it will work, its life expectancy, its operation and maintenance costs, and will probably want to optimise the system for maximum usefulness and minimum cost. The developer will therefore need to model the system, initially as a feasibility study with only ‘ball-park’ figures for the sizes of components. As a project progresses, detailed design studies become necessary. These include comparisons of specific equipment options and an optimisation of the control strategy.

The feasibility study is particularly important in optimising the cost effectiveness of the system. Energy storage options always become more expensive as the time scale of that storage increases because the energy rating of a store increases with the time-scale of charging and discharging cycles (Barton, Infield 2004). An optimum system may therefore have a relatively small store size in energy capacity terms, and a large renewable energy power rating, but only a mathematical model of the system can indicate the best balance in component sizes.

Energy systems including wind power, solar power, loads and energy storage can be extremely complex to model. Wind speed, solar irradiance and loads all vary on all time scales from seconds to years. The state-of-charge of an energy store depends on the history of energy supply and demand and its own operating characteristics. When a system includes some backup generation, for example a diesel generator, or a grid connection or both, the modelling becomes even more complex.

This thesis describes a novel, probabilistic method of modelling the behaviour of an energy store placed in an electricity grid at the point of supply of wind and solar power. Such an energy store would be placed there in order to minimise the curtailment of renewable energy at times of surplus and/or minimise the unsatisfied electricity demand at times of deficit, because the local grid lacks the capacity to absorb excess renewable energy at all times and/or to supply load deficits at all times. The modelling is used to answer the questions:

How much energy is still curtailed?
How much electrical demand remains unsatisfied?
How are these quantities reduced by the presence of the store?
What fractions of time does the store spend full, empty, filling and emptying? In other words, what fraction of time is the store useful, i.e. neither full nor empty?

There are three principal methods described in technical literature for modelling renewable energy systems: time step, Markov chain and probabilistic. These methods will be described below, with examples, drawing out the strengths and weaknesses of each. Finally, the probabilistic method used in this thesis will be described with its strengths and weaknesses.

1.6 Time Step Simulations

When designing an electricity system, or adding intermittent renewable sources to an existing system, a time-step simulation method is usually used to model the new system and predict its performance, especially when the system includes some form of energy storage. Time-step methods are tried and trusted, easy to code and they capture time-dependent effects very well, even when the system behaviour has significant daily, weekly or seasonal periodicity. They can be as complex as necessary to model all aspects of a system. Time-step simulations may always be the best modelling method to use at the detailed design stage, especially when devising the control software of an energy management control system.

A time-step model would step through a simulation period, e.g. at least one year in steps of, for example, one minute, ten minutes or one hour. The time step method would calculate the surplus or deficit power in each step, calculate the power imported or exported via the grid connection and any remaining surplus or deficit sent to or required from the energy store. Given an initial state of charge (SOC) of the store, the time step method would iterate through the simulation period calculating the state of charge of the store. When the store is full, it can no longer absorb any more power and when it is empty it can no longer supply any more power. Characteristics of the store can be added to the time step model, for example no-load losses, round-trip losses and control strategies. The time step model predicts the remaining curtailed energy, the remaining unsatisfied demand, and the fractions of time that the store spends empty or full and therefore not being used.

Several very sophisticated modelling tools have been written and used for renewable energy in electricity systems. Two of particular interest are Hybrid2 (Barley, Winn 1996, Iqbal 2003, Mills, Al-Hallaj 2004, Panickar, Islam & Bleijs 1999, Panickar, Islam & Nayar 1998) and HOMER (Chun Che Fung, Rattanongphisat & Nayar 2002, Iqbal 2003, Iqbal 2004). Hybrid2 is a joint project between the University of Massachusetts (UMass) and the U.S. National Renewable Energy Laboratory (NREL), and funded by the U.S. Department of Energy. HOMER is written by NREL, is available with a free licence, and is downloadable from NREL. Both of these programs typically model one year of system operation in hourly time steps. They can use either real or synthetic weather data. Both programs model DC-connected components on a DC bus and AC components on an AC bus, with power conversion devices (rectifier and / or inverter) to transfer electricity between the two buses. Both also have a sophisticated solar model that estimates the beam and diffuse components of solar radiation, and calculates the effect of each on a solar panel of a given tilt and azimuthal orientation.

Other time-step simulation programs in the literature include (ARES)-II, standing for ‘Autonomous Renewable Energy Systems’ (Celik 2002a, Morgan, Marshall & Brinkworth 1997) This program models stand-alone systems and predicts battery terminal voltages. Another simulation model is the ‘National Grid Model’ of the UK transmission grid (Halliday 1988, Watson, Ter-Gazarian 1996).

1.6.1 Hybrid2

Hybrid2’s theory manual is available on-line (Manwell et al. 1998). Hybrid2 is a logistical and first-stage design model. It models the long-term performance of a system and aids component sizing decisions, but does not model voltage control or power quality, or aid the design of individual components. For that, a dynamic model would be required. Hybrid2 can model wind power, solar power, loads, energy storage devices, AC and DC transmission systems, power converters of various kinds, diesel generators, load management systems and dump loads. However, hybrid2 can apparently only model stand-alone systems and mini-grids, not large grids or grid-connected systems.

Hybrid2 interpolates between hourly time steps using probabilistic methods, but only for loads and wind power, not for solar power. A normal distribution is usually used for these short-term variations, unless variations can go mainly only in one direction, in which case a Weibull distribution is used. The variation of wind speed within an hour is adjusted for the number of turbines in a wind farm and their spatial separation. Hybrid2 uses the Von Karman spectrum of turbulence for wind speed variations within one hour. Hybrid2 uses temperature when available to make corrections to air density and PV performance. The wind turbine model includes a turbulence model and corrections for height above ground level. Hybrid2 will even fill in gaps in the data using a Markov chain method for wind speed data and interpolation of clearness index for solar data. Ground reflectance is included in the solar model.

1.6.2 HOMER

A lot of documentation about HOMER is also available on-line, for example: (Lambert, Lilienthal 2003) HOMER is primarily an economic optimisation model for feasibility studies, to complement Hybrid2. The theory behind HOMER is less readily available, but HOMER also uses hourly time steps over a period of a year. It accounts for seasonal and diurnal variations in wind speed and loads, and can use real data or synthetic data based on its own data. HOMER cannot model voltage constraints. However, HOMER is available on a free licence and has a very user-friendly graphical user interface.

HOMER apparently does no calculations of variations within each hourly time step. It merely calculates an ‘operating reserve’, equivalent to a spinning reserve, but applies to small-scale systems that may or may not have any spinning generation equipment. This reserve changes over time and is calculated as a percentage of the hourly load, wind and solar power flows. The default values are 10% of load, 25% of solar power and 50% of wind power in each hour. To some extent, HOMER relies on the fact that in practice, solar powered systems always use large battery stores, backup generation or grid connection.

In its solar model, HOMER can construct its own synthetic irradiance data for any longitude and latitude based on monthly average irradiances or monthly average clearness indices.

1.6.3 Drawbacks of Time-Step Methods

Time step methods require large amounts of weather and load data. At minimum they require a year's worth of data to model seasonal effects. If the time step is one hour, then 8760 values of each input variable are required. However, an hourly time step fails to accurately model turbulent variations in wind speed, passing clouds and rapid changes in load. For small-scale electricity systems consisting, for example, of one wind turbine, one array of PV modules and supplying a few households, these short-term variations are more significant in percentage terms than in a larger system. In a larger system with more sources and loads, especially different types of renewable energy or a variety of load patterns, changes in one place or device are mitigated by aggregation.

The amount of data required is inversely proportional to the time step. If a shorter time step is used, for example of 1 minute, then 12.6 million values are required per variable per year. To fully capture the turbulent variations in wind speed or transient effects in the load, a one second time step may be required. Then 757 million values would be required per variable per year. However, weather data is typically measured as hourly averages, and as 10-minute averages at most. The British Atmospheric Data Centre supplies data mainly as hourly averages, (National Environmental Research Council 2004) .

Time step methods can also take a lot of time to calculate. The computer run time is also proportional to the number of time steps, and therefore inversely proportional to the time step length. With an hourly time step, computer run time is quite short, for example HOMER, (Lambert, Lilienthal 2003). But when the time step is one minute or less, a model can take minutes to run, for example the time step model created for the study presented in this thesis.

In addition, time step methods may suffer from inaccuracy due to inter-annual variability in weather. If the chosen year of weather data is not representative of the long-term average, or if there is a significant proportion of years for which that data is not representative (e.g. once every 5 years cloud is greater or wind speeds are much lower), then the time step model will fail to adequately model the performance of a real system.

In some cases, synthetic data can be useful, but the temporal characteristics of synthetic data are less likely to be correct. The properties of wind speed data from a typical synthetic wind generation algorithm are examined in section 2.4.12. Synthetic weather data is often generated using Markov chain methods, see below.

1.7 Markov Chain Methods

A Markov process is one where only the current state of the process can influence where it goes next. The process has no memory of the past (Norris 1997). Ergodic random probabilities also influence the future state. A Markov chain is a discretised form of such a process; a system can occupy any one of a finite number of states. The probability that the system will change from one state, i to any other state, j is given by a transition matrix of probabilities, P_{ij}. This is a standard mathematical method for modelling stochastic variation where the state of a system can be defined by a relatively small number of variables, the 'state variables'.

Markov chains are often used to generate synthetic wind speed data: (Biga, Rosa 1981, Muselli et al. 2001, Poggi et al. 2000) or solar irradiance data: (Masters et al. 2000, Torre, Poggi & Louche 2001, Wan-Kai Pang, Forster & Troutt 2001). In these cases, there is only one variable, the wind speed or the solar irradiance.

Sometimes just one Markov chain is used but in other cases two Markov chains may be necessary, especially when considering variations over short and long time scales in the same data string, for example (Infield et al. 1994). In this synthetic wind algorithm, one Markov chain is used to create a series of hourly average wind speeds with a Weibull distribution, and another Markov chain is used to create a Gaussian distribution of wind speeds within each hour period, varying according to localised wind speed turbulence. Such a method may be described as a higher order Markov chain (MacDonald, Zucchini 1997).

Synthetic data created by Markov chains is often used in time step simulations, for example in the HOMER time-step modelling software (Lambert, Lilienthal 2003). Some simulations create their own wind power and solar PV power data simultaneously in the same model using Markov transition matrices, for example (Castro et al. 1996).

At least one attempt has been made to model a whole energy system using a Markov method (Frean 1983). The method looks useful at first sight, but to model the complexity of a real system with seasonal and diurnal effects would be prohibitively difficult. The number of state variables would be large, the transition matrices would be very large and their values would be difficult to evaluate.

Finally, (Baumgaertner 1995) describes a Markov chain method for predicting the likely behaviour of wind and solar power up to one hour ahead in order to optimise the decision making of a system controller.

1.8 Probabilistic Methods

Probabilistic methods can be used either on their own or in conjunction with time stepping and Markov chain methods. Simulation models that use time-step or Markov chain methods often use standard statistical distributions (e.g. Weibull, Gaussian and beta distributions) to model variations in solar power, wind power and loads, as described below. These probability distributions are used to formulate the probabilities within Markov transition matrices (Celik

2002b, Celik, Marshall 1998, Infield et al. 1994, Poggi et al. 2000, Wan-Kai Pang, Forster & Troutt 2001), and to model the short-term variations within each step of a time-step method(Lambert, Lilienthal 2003, Manwell et al. 1998).

Simulation models that use only probabilistic methods are relatively rare, especially when the system includes a significant energy storage component. Four are described in section 1.8.5.

1.8.1 Probability Distributions in Wind Speed Modelling

Distributions of hourly or ten-minute average wind speeds are most commonly modelled as two-parameter Weibull distributions, for example (Celik 2002b, Celik, Marshall 1998, Chadee, Sharma 2001, Infield et al. 1994, Stevens, Smulders 1979, Wan-Kai Pang, Forster & Troutt 2001). In other situations, lognormal distributions have been found to provide a good fit to wind speed data:(Garcia et al. 1998, Luna, Church 1974, Shaw, McCartney 1985).

Occasionally, a three-parameter Weibull model is used (Van der Auwera, de Meyer & Malet 1980). Alternatively, a special case of the Weibull distribution is the Rayleigh distribution (a Weibull distribution with a shape factor of 2) (Tuller, Brett 1985). This has some physical justification, being derived from a bi-variate normal distribution (Hassan, Sykes 1990, Tuller, Brett 1984, Tuller, Brett 1985). But when compared with other distributions, the two-parameter Weibull distribution most often gives the best fit to wind speed data (Garcia et al. 1998, Tuller, Brett 1985). Wind speed modelling in this thesis concentrates on the two-parameter Weibull distribution for long-term wind speed variations, and on the normal distribution for short-term turbulent variations and for long-term distribution of longer averages e.g. the distribution of monthly average wind speeds. The log-normal distribution is discussed briefly in section 2.6.3.3 but is not used in the probabilistic method.

1.8.2 Probability Distributions in Solar Irradiance Modelling

Sometimes, global horizontal irradiance is analysed and modelled directly (Akuffo, Brew-Hammond 1993, Giraud, Salameh 2001, Rahman, Khallat &

Salameh 1988, Sahin, Sen 1998, Trabea 2000). In most models, the global irradiance is normalised by the extra-terrestrial horizontal irradiance, and referred to as the clearness or clearness index, K_T. The clearness index remains more constant with time of day and season, but still varies with solar elevation angle. As elevation angle decreases, the atmospheric path length of sunlight (the airmass) increases. This reduces the clearness index as recognised by many analyses for example (Gonzalez, Calbo 1999, Jurado, Caridad & Ruiz 1995, Skartveit, Olseth & Tuft 1998, Suehrcke, McCormick 1989, Tovar et al. 1999, Tovar et al. 2001, Tovar, Olmo & Alados-Arboledas 1998).

In addition, the beam fraction of radiation is often distinguished from the diffuse and albedo (reflected) fractions in order to predict the energy capture of an inclined solar panel, (Babatunde, Aro 1995, Gonzalez, Calbo 1999, Gordon, Reddy 1989, Hollands, Huget 1983, Hove 2000, Ideriah 1992, Kudish, Ianetz 1996, Lam, Li 1996, Lopez, Rubio & Batles 2000, Skartveit, Olseth & Tuft 1998, Suehrcke, McCormick 1989, Tiris, Tiris 1998, Tovar et al. 1999, Trabea 2000, Ulgen, Hepbasli 2002, Unozawa, Otani & Kurokawa 2001).

Solar parameters (irradiance, clearness index or other) are often modelled using beta distributions: (Graham, Hollands 1990, Mefti, Bouroubi & Adane 2003, Rahman, Khallat & Salameh 1988, Sahin, Sen 1998, Sulaiman et al. 1999, Youcef Ettoumi et al. 2002). In some papers, other distributions have been used to model solar parameters, sometimes only after a transformation has been applied to the data. A normal distribution has been used by (Aguiar, Collares-Pereira 1992a, Aguiar, Collares-Pereira 1992b, Amato et al. 1985, Jurado, Caridad & Ruiz 1995, Loutfi, Khtira 1992), a bi-exponential distribution by (Ibanez, Beckman & Klein 2002, Ibanez, Rosell & Beckman 2003), a Weibull distribution by (Rahman, Khallat & Salameh 1988) and even a shifted negative binomial distribution by (Poggi et al. 2000). However, beta distributions are the most commonly used and usually give the best fit to the data (Rahman, Khallat & Salameh 1988).

Beta distributions have the useful property that all the possible values lie within a finite range. Thus a beta distribution can represent clearness distributions form very overcast skies, to very clear skies and all conditions in between, while excluding negative values and improbably high values of irradiance. Beta distributions can have a single maximum between the minimum and maximum ends of the range, a maximum at one end or the other, or maxima at both ends of the distribution at the same time. Therefore beta distributions can produce a wide variety of distribution shapes. Sometimes, a single beta distribution can represent the whole range. An alternative approach is to model the solar parameter as the sum of two beta distributions, one for cloudy conditions and another for clear skies, for example (Youcef Ettoumi et al. 2002). This reflects the fact that the sky tends to be either cloudy or clear, spending relatively little time changing from one condition to the other. Solar irradiance therefore often exhibits a bimodal character, (Ibanez, Beckman & Klein 2002, Ibanez, Rosell & Beckman 2003, Jurado, Caridad & Ruiz 1995, Suehrcke, McCormick 1989, Tovar et al. 1999, Tovar, Olmo & Alados-Arboledas 1998). This bimodal nature will be used in the probabilistic modelling method presented in this thesis, section 2.7.

1.8.3 Probability Distributions in Load Modelling

Electricity demand can vary with time of day, day of the week, weather conditions, random factors and even with television schedules. As will be seen in section 2.4.4, demand from individual consumers varies substantially with large step changes as appliances are switched on or off. When demand from many consumers is aggregated together, demand follows a mostly smooth profile over time, with relatively small differences between one day and the same time on another similar day. Electricity distribution companies have sophisticated forecasting techniques (Mandal et al. 2004) that predict the demand from each type of consumer, for example domestic, commercial and industrial.

Electrical demand is often modelled as a deterministic profile over each day (Chun Che Fung, Rattanongphisat & Nayar 2002, Giraud, Salameh 2001, Habibi 2001, Hove 2000, Lachs, Sutanto 1995, National Grid Company plc

2003). Occasionally, models treat demand as a stochastic variable (Yakin 1984).

Synthetic load data can be created in a similar way to synthetic wind speed data (Infield et al. 1994). In this data generator, load has a Gaussian distribution within each hour and a Weibull distribution from one hour to the next. It would be a simple modification to superimpose this onto a daily profile to create synthetic load data for a small group of consumers. The Weibull parameter can be given any shape factor, so that the load distribution can be as skewed as desired.

1.8.4 Probabilistic Models Within Time Step Models

As stated above, some time-step models include a probabilistic calculation of variations within each time step, for example the variation of demand or wind power within each hour (Lambert, Lilienthal 2003, Manwell et al. 1998). This is a relatively minor use of probabilistic methods that enables time-step methods to use longer time steps and therefore run more quickly with less input data. The probabilistic calculation merely determines the size of the spinning reserve or 'operating reserve' (Lambert, Lilienthal 2003).

The method presented in this thesis goes much further, using only probabilistic methods together with seasonal and diurnal profiles.

1.8.5 Purely Probabilistic System Models

Some previous attempts have been made to use purely probabilistic methods for modelling renewable energy (Barton, Infield 2004, Khallat, Rahman 1986, Swift-Hook, Ter-Gazarian 1994, Ter-Gazarian, Kagan 1992). However, these are very limited in scope. (Swift-Hook, Ter-Gazarian 1994) only examines the level of penetration of intermittent renewable energy in an electricity system at which energy storage becomes necessary. (Ter-Gazarian, Kagan 1992) only presents a very simple model based on 3 times of day, and assumes that the renewable generation is never available at the times of maximum electrical demand. In (Khallat, Rahman 1986), Khallat and Rahman use probabilistic methods to predict the hourly capacity factors of photovoltaic arrays at a given site. However, this method does not help to size the energy store. (Barton, Infield 2004) presents a novel method of calculating energy flows to and from an energy store, but makes only a crude attempt to size the store.

This thesis will use the methods of that previous paper, (Barton, Infield 2004) but extending the method to include an improved calculation of store size and a more sophisticated spectral analysis of the intermittency of renewable sources. Even the probabilistic method presented in this thesis is not a purely probabilistic method, because it uses a time step simulation of load variations within a day as part of the calculation of store capacity, section 2.11. This is an example of a time step method within a probabilistic method, in contrast to a probabilistic method within a time step method.

1.9 The Novel Probabilistic Method Presented in This Thesis

Henceforth, this novel probabilistic method will be referred to as 'the probabilistic method'.

The probabilistic method simulates electrical power systems with large fractions of intermittent renewable generation and energy storage. It does not require any time series of weather data. Instead it uses spectral manipulation done on Fourier transforms of time series, probability density functions, and average profiles of diurnal and seasonal variations. The input data required is considerably reduced, although the computer code is longer and more complex than a time-step code. This probabilistic method typically takes less computer time to run and is now more practical to run inside an optimisation routine. Some loss of accuracy may be apparent because of the necessary simplifying assumptions, but it is hoped that the probabilistic method may one day provide a fast and practical tool for feasibility studies, early system design, and some investment decisions.

A primary input to the probabilistic method is the working time period of the store. From this, the method calculates the size of the store using spectral methods. The time period can be thought of as a store working time period, or a typical time for the store to complete one cycle of charging to full and discharging to empty. In reality, in an electricity grid with variable wind power, solar power and loads, the energy store will perform charge and discharge cycles on many different time scales, but the calculated store size is a weighted average of the energy capacity required to absorb all power variations up to and including the chosen time period.

Two versions of the probabilistic method have been written into Matlab computer programs: One works on any time scale but models wind power only with a constant or simply varying electricity load. This will be referred to as the ‘wind only’ program. The second version is currently only written for an energy storage period of 24 hours but models variations in wind power, solar power and electricity demand. This will be referred to as the ’24-hour’ program.

The probabilistic method steps through all possible types of period, treating each period as a statistically independent event. Each period has a different combination of average wind power, average solar power and average electricity demand. The probabilistic method calculates the same outputs as the time step method: the curtailed energy, the unsatisfied demand, and the fractions of time that the store spends full, empty, filling and emptying.

1.9.1 An Overview of the Wind-Only Probabilistic Program

The data flows and calculations performed in the wind only program are shown in figures 1.1 and 1.2.

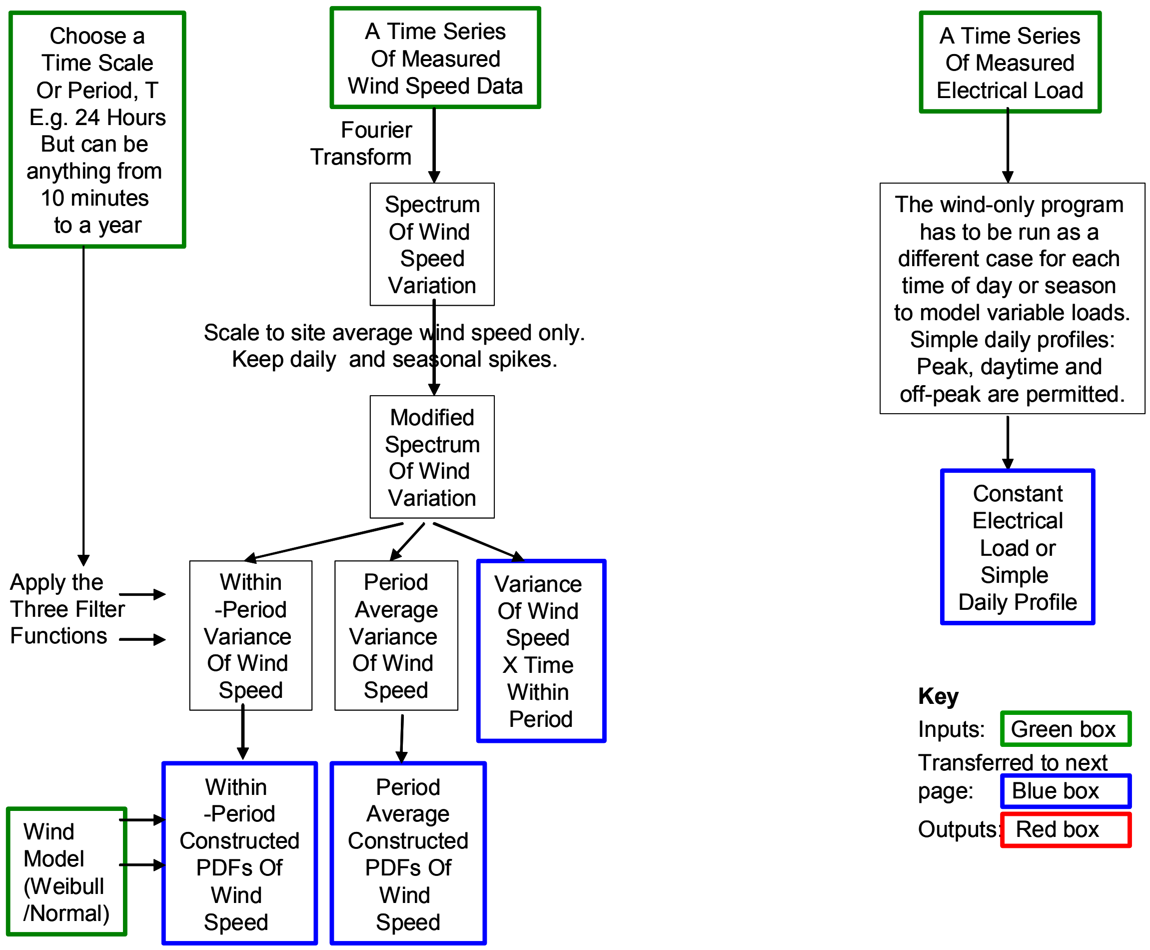

Figure 1.1. Page 1 of the data flows and calculations performed in the wind-only probabilistic computer program

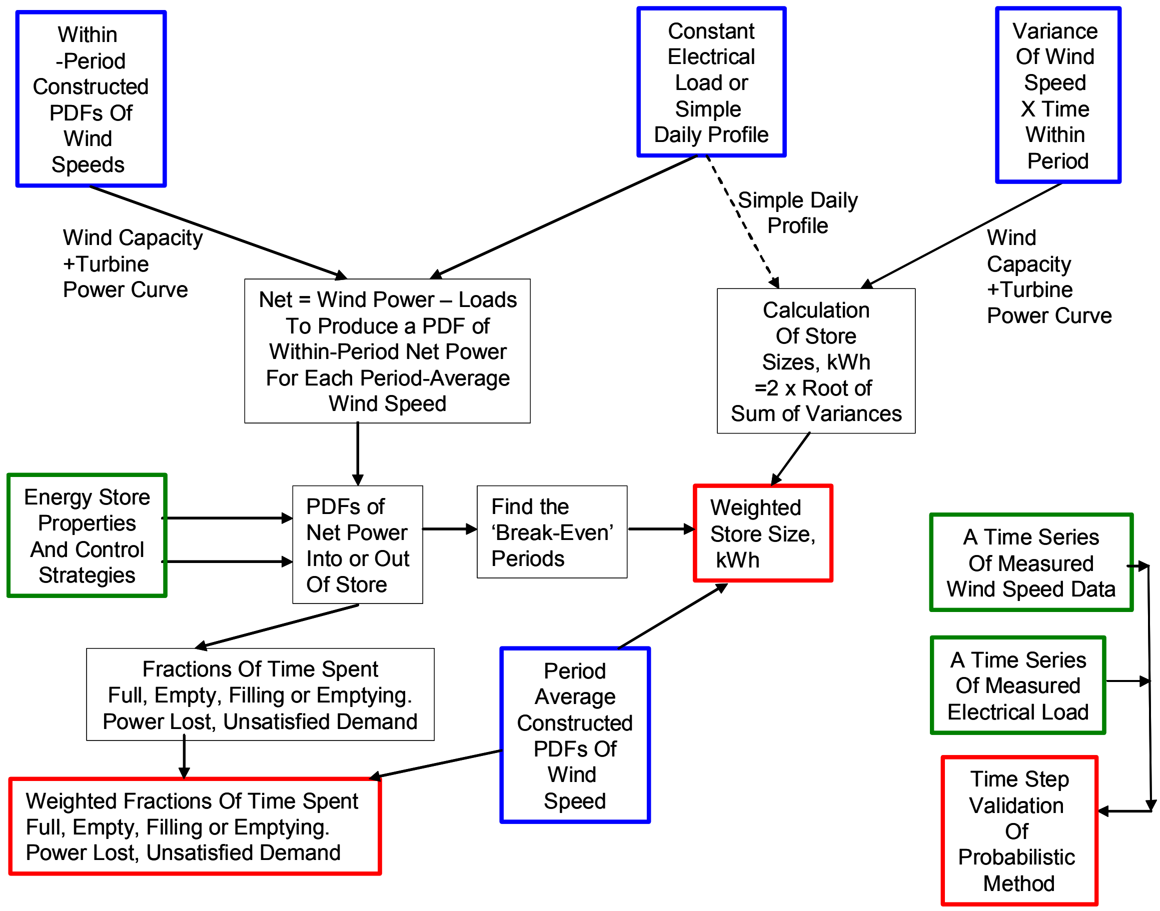

Figure 1.2. Page 2 of the data flows and calculations performed in the wind-only probabilistic computer program

The method is easiest to see for the wind-only version. The wind power spectral density, referred to as the spectrum of wind speed variation is filtered to predict a variance of period-average wind speeds, for example a variance of daily average wind speeds, and also a variance of within-period wind speeds, for example the variance of wind speeds within one day for a given daily average wind speed. A third filter function predicts the variance in wind speed x time over the storage period, a quantity that together with the wind turbine power curve is used to calculate the energy capacity of the store. The capacity of the store is only calculated in the break-even periods, i.e. those when the average wind power is approximately equal to the electricity demand plus the store losses, because these are the periods when the store is most heavily used and is most likely to float between full and empty. In other periods, the store either spends significant fractions of time either full or empty, and floating close to one limit or the other.

Probability density functions (PDFs) of instantaneous wind speed are constructed for each period-average wind speed, and from these PDFs of wind power are calculated. In the simple version of the wind-only program, the PDF of net power is calculated by simply subtracting a constant electrical demand from the wind power. The curtailed energy and unsatisfied demand are calculated by integrating over all periods, i.e. all period-average wind speed conditions. The fractions of time spent full, empty, filling and emptying are calculated using the method described in appendix B.

Another version of the wind-only program has also been written that can be run with time-varying demand using 3 different times of day (peak, daytime mid-rate and off peak) and two seasons (summer and winter). However, this program has to assume that wind speed variations have no periodic dependence on time of day or season. This second version of the wind-only program assumes that the averaging time scale is at least 24 hours but shorter than 6 months. In this case, the PDF of net power is calculated by a simple convolution of the wind power PDF with the electrical load PDF.

These two versions of the wind-only program are essentially the same as the modelling method used to prepare a previous paper, (Barton, Infield 2004), but with a more sophisticated calculation of store capacity.

1.9.2 An Overview of the 24-Hour Probabilistic Program

The data flows and calculations performed in the 24-hour program are shown in figures 1.3 and 1.4.

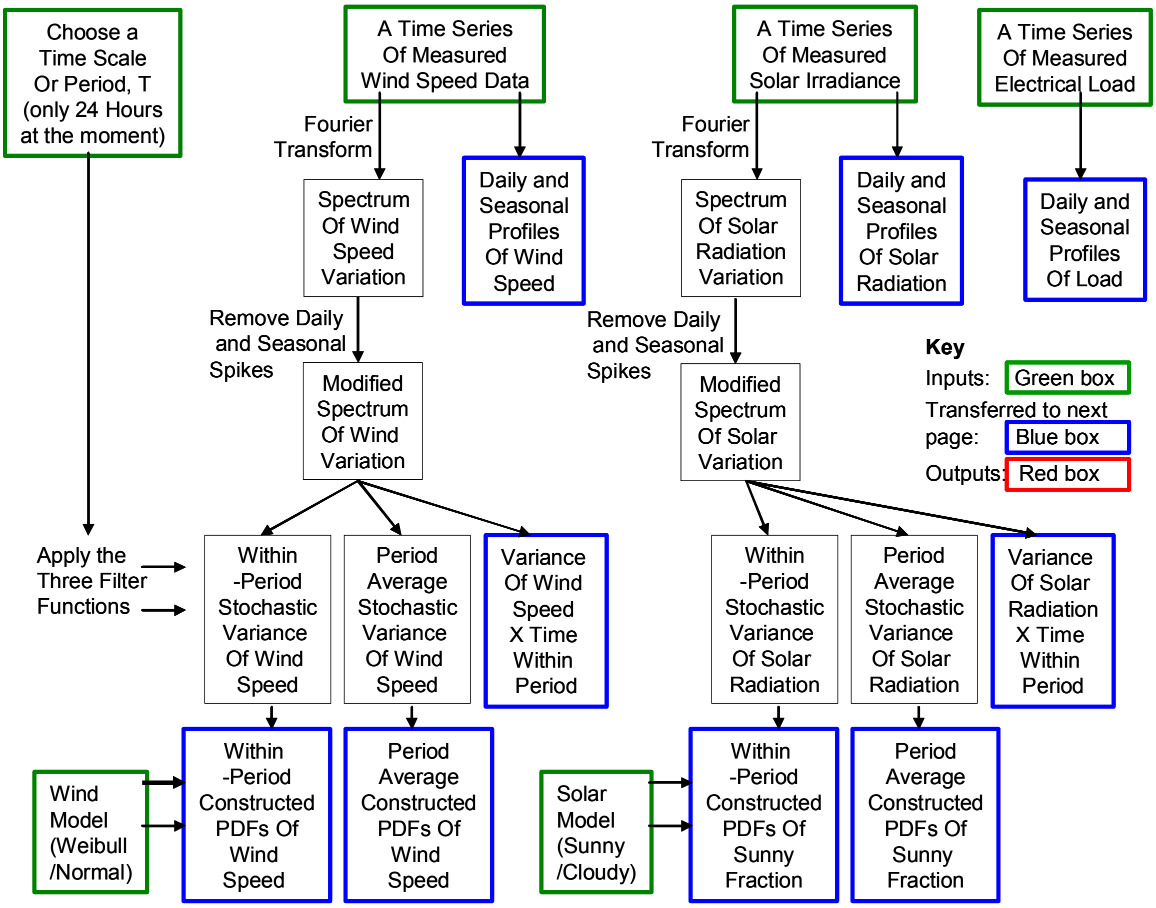

Figure 1.3. Page 1 of the data flows and calculations performed in the 24-hour probabilistic computer program

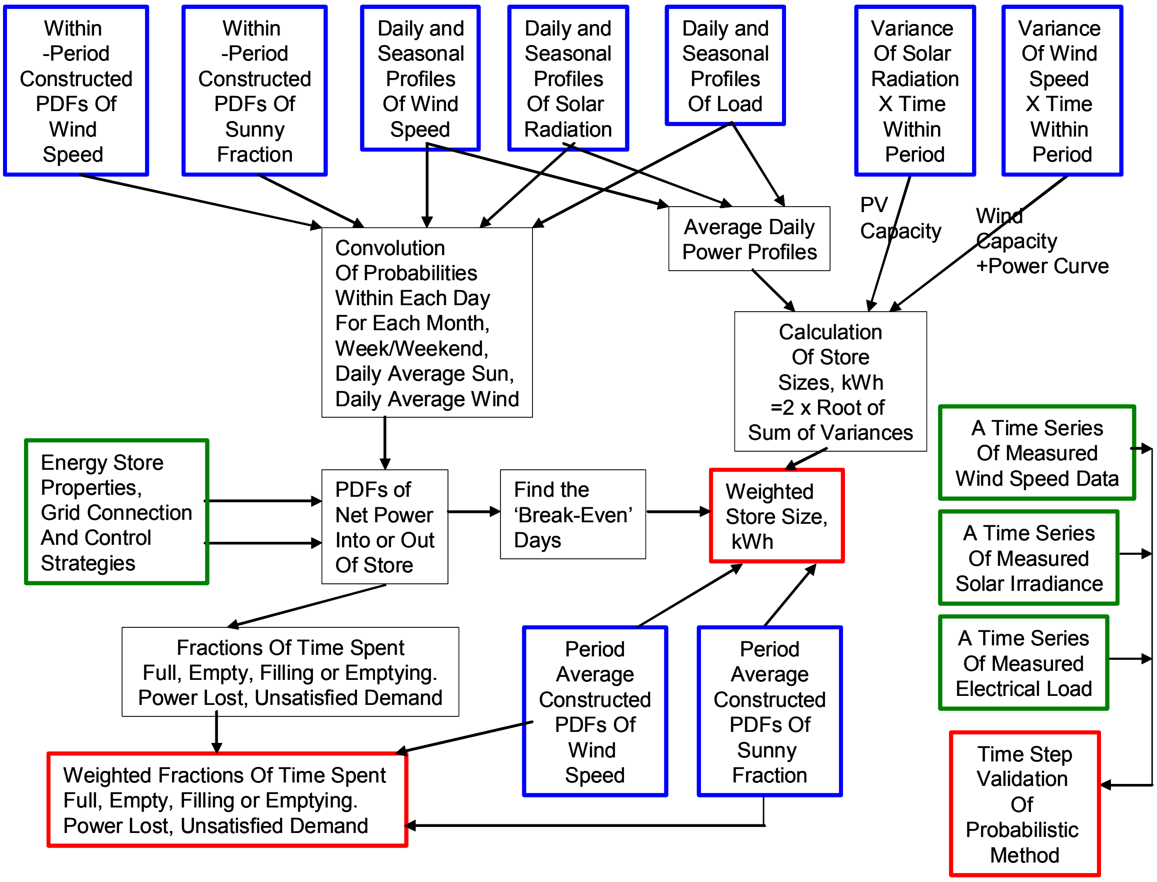

Figure 1.4. Page 2 of the data flows and calculations performed in the 24-hour probabilistic computer program

The calculations performed in the 24-hour program are qualitatively the same as those in the wind-only program, but now including solar power and more diurnal and seasonal effects. Fourier Transforms, filter functions and constructed PDFs are now also applied to variations in solar power. Electricity demand changes with hour of the day, weekday or holiday and month of the year. In addition, due to the highly seasonal and diurnal nature of electricity demand and especially of solar power, daily and seasonal profiles of wind power, solar power and electricity demand are also calculated. In order to treat diurnal and seasonal effects correctly and separately from stochastic variations, the daily and seasonal spikes are first removed from the spectra of wind speed and solar irradiance. In order to simplify the program, electricity demand is treated as a purely deterministic, repeating profile, which it is to a first approximation. Nevertheless, the complexity has increased significantly from the wind-only case: The wind-only program has to step through every period-average wind speed, but the 24-hour program has to step through every combination of daily average wind speed, daily average solar power and daily electricity demand profile. The PDFs of net power are no longer the wind power PDFs with electricity demand subtracted; they are now a convolution of wind power PDF, solar power PDF and electricity demand PDF. The program has to step through a 7-level nest of for loops, as described in section 2.10.3 and shown in fig. 2.74.

1.9.3 Advantages of the Probabilistic Method

If the number of conditions modelled (computer program loops) is less then the number of time steps of the equivalent time step method, then the probabilistic method has a good chance of saving computer run time. This advantage applies when compared to a time step method with very short time steps, for example 1 minute or less, requiring a great number of time steps to model a long total period.

The reduced computer run time makes it easier to run the probabilistic method inside an optimisation routine and thus optimise the sizes of system components during a feasibility study.

The probabilistic method reduces the amount of data that has to be gathered and stored before the model can be run.

When a site has insufficient measured wind speed and solar irradiance data, the conventional ways to model renewable energy output at that site are by using weather data from elsewhere or by using synthetic weather data. Weather data from other sites may not have the same average values and may not have the same diurnal and seasonal patterns. Synthetic data rarely has the same spectral variation, as described in section 2.4.12, and will therefore give an incorrect estimate of the variations in power output and an incorrect prediction of the performance of an energy store. The probabilistic method opens up the possibility of combining an appropriate spectrum of stochastic variations with appropriate seasonal and diurnal profiles to give an accurate model of the system operation at the feasibility study stage without having to embark on a long data measurement campaign.

1.9.4 Limitations of the Probabilistic Method

The probabilistic method requires a longer computer program code. The time step method can be coded in just a few lines, but the probabilistic method requires a long and complex code to integrate the spectra and to construct PDFs of wind and solar power. The programming complexity and computer run time increase exponentially with the number of different forms of variation that are added: wind, solar and variable loads together are much more complex than each variable on its own. Each new source of variation adds at least one, if not two, more ‘for’ loops into the nest of ‘for’ loops of the computer program. As a result of the increased programming complexity, the version of probabilistic program presented here that models variable loads, solar and wind power can only model energy stores that operate over a time scale of 24 hours. In contrast, it was relatively easy to use the probabilistic method to model wind power over all and any time scale. Highly seasonal and diurnal variations create great programming complexity. Some of the potential savings in computer run time are therefore lost.

The probabilistic method is not so easy to validate. In this thesis, the probabilistic method is validated by comparison with a time step method. The time step method, however, can by validated by hand calculation of a few time steps.

The probabilistic method may give less accurate results than the time step method, because simplifying assumptions have to be made in the construction of the PDFs and in the calculation of times spent full or empty. In the time step method, the PDFs are not required since time series of data are used.

The probabilistic method as it stands can only calculate the size of an energy store and its behaviour by starting with the period of operation (the typical time scale of charge and discharge cycles). It cannot start with a store size and predict the period of operation except by iteration.

The probabilistic method requires spectral information and seasonal and diurnal data that is not currently available for all locations. These data must first be calculated from time series of data measured at a variety of locations.

The 24-hour version of the probabilistic model, as currently written into a computer program, would break down if the fraction of time that clouds are present in the sky depended significantly on the solar elevation. It would give the program difficulty in using the variation in solar power calculated from the PSD to predict the fractions of time that passing clouds obscure the sun, and therefore the variations in solar power, see section 2.7.5.

The 24-hour version of the model would also break down if there were a significant correlation or anti-correlation between solar power and wind power beyond seasonal and diurnal variations of each; or a significant stochastic variation of electricity demand; or a significant correlation between electricity demand and either solar power or wind power, beyond diurnal and seasonal effects. In practice, however, the data used In this thesis suggests that such significant correlations or stochastic variations in electricity demand do not exist in the UK.

1.9.5 Information Used by the Probabilistic Method

The following two tables list all the inputs required by each version of the method, together with their probable source in a real modelling exercise. The tables are based on the computer program listings.

Table 1.1 Data Requirements of the Wind-Only Probabilistic Method:

Data	**Probable Source of Data**
Maximum charge and discharge power capacity of the store	Chosen and optimised by the modeller
Round-trip efficiency of the store, including no-load or parasitic losses	Depends on the physical characteristics or electrochemistry of the energy store
Wind turbine power rating	Chosen and optimised by modeller
Grid import and export power capacities	The local electricity distribution network operator (DNO)
Wind Power Spectral Density (PSD) function	Created by fast Fourier transform (FFT) of some representative wind speed data from a location with a similar climate. Height of seasonal and diurnal spikes may be adjusted according to local climate, if known.
Original sampling period of the data from which the FFT was run. Needed to calculate the Nyquist frequency.	From the original wind speed time series, or from the highest frequency of the PSD function
The store operating period, T, or typical charge-discharge cycle time that the store is designed for.	Chosen and optimised by modeller
The frequency increment of the wind PSD	From the PSD function data
Site mean wind speed	Estimated from topology modelling, from a wind atlas or by measure-correlate-predict using data from a nearby weather station
Wind turbine power curve	From potential suppliers, or a generic power curve, since most modern turbines have very similar power curves for a given wind class.

Table 1.2 Data Requirements of the 24-Hour Probabilistic Method:

Data	**Probable Source of Data**
Latitude and longitude of the modelled site, together with time zone for predictions of solar power	From a map of the site or GPS measurements
Maximum charge and discharge power capacity of the store	Chosen and optimised by the modeller
Round-trip efficiency of the store, and separately, the no-load or parasitic losses of the store	Depends on the physical characteristics or electrochemistry of the energy store
Grid import and export power capacities	The local electricity distribution network operator (DNO)
Original sampling period of the data from which the FFT was run. Needed to calculate the Nyquist frequency.	From the original wind speed time series, or from the highest frequency of the PSD function.
The store operating period, T, or typical charge-discharge cycle time that the store is designed for.	Chosen and optimised by the modeller
Wind turbine power rating	Chosen and optimised by modeller
Solar PV power rating	Chosen and optimised by modeller
Backup or local generator power rating	Chosen and optimised by modeller if new, or name-plate rating or measured if existing
Site mean wind speed	Estimated from topology modelling, from a wind atlas or by measure-correlate-predict using data from a nearby weather station

Wind Power Spectral Density (PSD) function	Created by fast Fourier transform (FFT) of some representative wind speed data from a location with an approximately similar climate. Heights of seasonal and diurnal spikes are irrelevant since seasonal and diurnal effects are handled separately.
Solar Power Spectral Density (PSD) function	Created by fast Fourier transform (FFT) of some representative solar radiation data from a location with an approximately similar climate. Heights of seasonal and diurnal spikes are irrelevant since seasonal and diurnal effects are handled separately.
Long-term average solar irradiance at the site where the solar PSD was measured, in order to correctly scale the PSD	Calculated at the same time as the solar PSD
Average wind speeds in each month, daytime and at night, to create the correct daily and seasonal pattern, but not for the average level	Measured at a site with a very similar climate in terms of seasonal patterns e.g. monsoon rains, latitude, land mass size and proximity to a coast
Average solar radiation in each month, from which the program calculates the fraction of time that the weather is sunny, and the total solar irradiance at the modelled site	Measured at a site with a very similar climate in terms of seasonal patterns e.g. monsoon rains, latitude, land mass size and proximity to a coast

Sunny and cloudy attenuation factors and scale factors. These determine how much solar radiation is absorbed or reflected by the atmosphere at a given solar elevation angle, when weather is sunny and when it is cloudy.	Sunny and cloudy attenuation factors are calculated from solar radiation data at a location with similar levels of airborne pollution and dust, from positions of peaks in graphs of clearness index PDF. Sunny scale factor is also calculated from positions of sunny peaks, but cloudy scale factor is calculated to give the correct total variance in solar irradiance as calculated from the integration of the whole solar PSD function including seasonal and diurnal spikes.
Average electrical loads for each hour of the day, weekday or weekend (or holiday), for each month of the year	From DNO or national grid data measured at a location with similar culture, climate and latitude and scaled to the size of the local network.
A calendar of a typical year to tell the program which days are weekdays and which are weekend or holidays	From any location with a similar culture in terms of weekly cycle and holidays.
The frequency increment of the wind PSD	From the PSD function data
Wind turbine power curve	From potential suppliers, or a generic power curve, since most modern turbines have very similar power curves for a given wind class.

More work is required to determine which of the above input parameters are almost universal (making the probabilistic method very useful and easier to use) and which ones change significantly from location to location, making the probabilistic method more problematic.

1.9.6 Spectral Methods

Spectral methods are familiar to engineers in the field of signal processing, vibration analysis, communication systems, control theory and others. The effect of electronic circuits and systems of mechanical inertia, springs and dampers can be understood in terms of their spectral response, as low-pass filters or high-pass filters for example. Any signal or data stream can be represented in the frequency domain by a spectrum, a power spectral density (PSD), calculated by applying a Fourier Transform to the signal. One well known spectrum has been constructed in this way from series of wind speed data; that is Van Der Hoven's spectrum measured at Brookhaven National Laboratory and first published in 1957 (Van Der Hoven, I. 1957). This work has been referenced and presented in several subsequent books, for example (Hassan, Sykes 1990, Spera 1994) and is shown here in fig. 1.5.

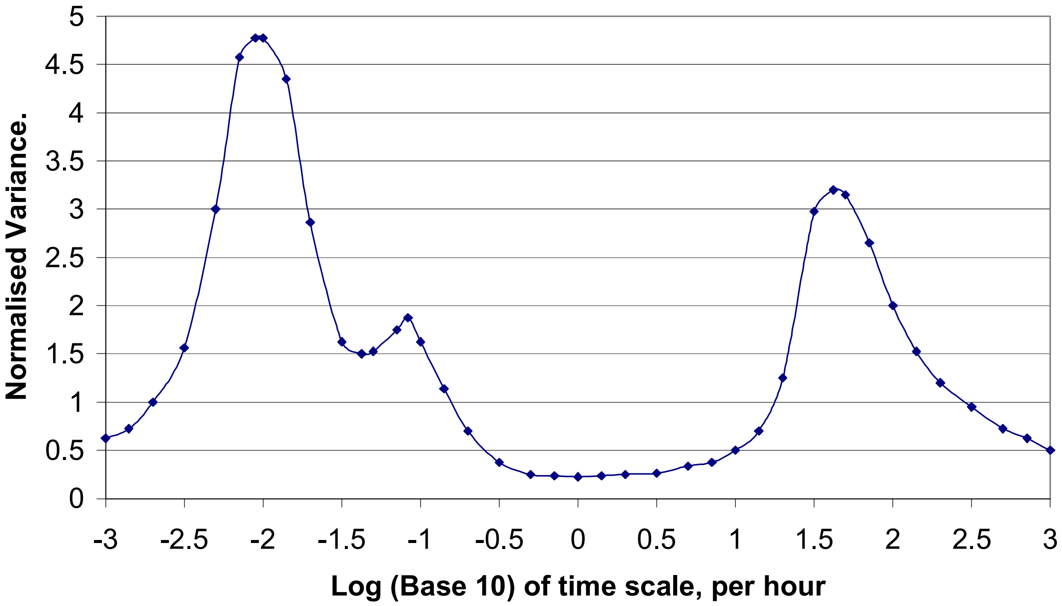

Figure 1.5 The spectrum of horizontal wind speed at Brookhaven National Laboratory, as published in (Van Der Hoven, I. 1957)

To the knowledge of the author, this and other wind speed power spectra have previously been used for illustrative purposes but rarely for quantitative calculations. Two rare examples are (Bossanyi, Anderson 1984, Infield 1990).

The probabilistic method presented in this thesis uses power spectra of horizontal wind speed and global solar irradiance to calculate the variation of wind power and solar power within a period, and the probability distribution functions of period-averaged power. The probabilistic method also uses power spectra to calculate the required size of an energy store, as described in sections 2.5.5 and 2.11.

1.9.6.1 Power Spectral Density and Variance

In this thesis, it is assumed that the area under the PSD curves (autospectral density function), created using a Fast Fourier Transform, represent the variance of the time series from which they were created. The justification for this is taken from (Bendat, Piersol 1993) pages 48, 51, 52 and 55, reproduced below for reference. First, the total area under the curve is the mean square value of the time series, i.e. the variance of the data plus the square of the mean of the data. The autocorrelation at zero time interval is given by:

$$R_{XX}(0) = \int_0^\infty G_{XX}(f)df = \psi_X^2 = \sigma_X^2 + \mu_X^2 \qquad (1.3)$$

Where R_{XX} is the autocorrelation function

G_{XX} is the one-sided spectral density function

ψ_X^2 is the mean square value of the data

σ is the standard deviation of the data

μ is the mean value of the data

Secondly, the mean value of the data appears in the spectral density as a delta function at $f = 0$ with an area of μ_X^2. Therefore, if the variation at zero frequency is excluded from any integration of the spectrum, as it has been in the probabilistic method, then the area under the remaining spectrum must represent only the variance, σ_X^2.

Finally, all relationships and properties of spectra developed from continuous spectra apply equally well to Fourier Transforms, by the Wiener-Khinchin relationship. Therefore, excluding the first (zero frequency) term, the integrated area under a Fourier Transform is also equal to the variance of the time series.

Of course, when a Fourier Transform spectrum has been calculated from one time series, for example a wind speed time series with a mean on 5.1m/s, and the computer model is of a site with an average wind of 8m/s, then the spectrum must be multiplied by the square of the ratio of mean wind speeds, e.g. $(8 / 5.1)^2$ in order to maintain a constant ratio of standard deviation to mean.

1.9.6.2 Why Only Variance and Not Higher Order Moments of Variation?

The probabilistic method uses the variances calculated from the PSDs, and combines variances by addition. Higher order moments of variation, for example skewness and kurtosis, may convey additional information but would be very difficult to manipulate and are not available from the PSDs. In the model, variances are added, assuming that the added variables are randomly and independently distributed, for example random variations of wind power and solar power are assumed to be independent of each other. In any case, higher order moments are not required. When needed, PDFs have been constructed from the calculated mean and variance, adding in knowledge about the characteristics of the system. For example, wind speed distributions are constructed knowing that wind speeds follow a Weibull or normal distribution; PDFs of solar irradiance are constructed from the solar elevation in each hour of the day combined with the mean and variance of solar irradiation within a given day, knowing that solar radiation exhibits a bimodal distribution, either sunny or cloudy, and can never take values less then zero.

1.9.7 Assumptions of the Probabilistic Method

The probabilistic method makes the following assumptions:

1. Wind speed distributions are modelled either as Weibull or normal distributions. When the standard deviation is large compared to the

mean, a Weibull distribution is used, up to a maximum Weibull shape factor of 3.0. Otherwise, for smaller standard deviations, a normal distribution is used.

2. Solar radiation follows a bivalent distribution for a given solar elevation angle: The sky is either sunny or cloudy, with a unique value of radiation for each state, for a given solar elevation.
3. The numerically integrated area under a PSD created from a FFT is equal to the variance of the parameter used to calculate the FFT, with the exception of the first term in the PSD which is equal to the square of the mean value.
4. Wind power and solar power variations ore independent of each other and of electricity demand, except for diurnal and seasonal variations of each.
5. The central limit theorem is applied to calculate the means and standard deviations of net power from the means and standard deviations of each independent variable: stochastic wind power variations, stochastic solar power variations, diurnal and seasonal variations. The central limit theorem is also invoked as a reason for using average profiles of electricity consumption and ignoring minute-to-minute variations in electrical demand from individual households.
6. The standard deviations of wind speed variations within a given time period are assumed to be proportional to the average wind speed over that time period. This is an extension of the concept of turbulence intensity to all time scales.
7. Solar power is assumed to be proportional to global horizontal irradiance and proportional to rated power of the PV device. The PV would produce its rated power at 1000 Watts per square metre of global irradiance. Average PV power output would be increased by tilting the PV module towards the equator, but PV power output would be reduced by heating of the modules. It is assumed that these two effects will approximately cancel in predicting total PV power output. The effects of inclination and module temperature are therefore ignored.

8. The energy store can effectively cope with all power variations up to a certain length of time, that is the store operating time or typical charge-discharge cycle time, T, but cannot accommodate longer-term variations.
9. The fraction of time that the weather is sunny (the ‘sunny fraction’ parameter) is independent of solar elevation angle. However, some variation in month-to-month average sunny fractions is permitted provided there is no overall correlation with solar elevation angle over the whole year.
10. Solar radiation is absorbed or reflected by the earth’s atmosphere according to linear attenuation coefficients, one for sunny conditions and another for cloudy conditions: The solar radiation on the earth’s surface is inversely proportional to the exponent of the thickness of atmosphere through which it travels.
11. Within a given month, the daily average sunny fractions, are distributed according to a beta distribution.
12. The size of the local electricity load and its distribution network is appropriate to a village, a region within a town, or larger. It is large enough to accommodate at least one large modern wind turbine, and is large enough that short-term variations in demand from individual consumers is aggregated out.
13. The size of the electricity system is small enough that it can be modelled as occupying just one time zone and that weather effects on wind and solar power are not completely aggregated out.
14. In the program written to embody the 24-hour version of the probabilistic method, a typical UK calendar year has been used.
15. An energy store is treated as a bin into which energy is put and from which energy is drawn some time later. Nothing is known or assumed about the physics or chemistry of the energy store. The only parameters used by the model are the round-trip electrical efficiency, the parasitic loss, the power ratings of charging or discharging the store, and the typical cycle time, T of the store. The store is like a low-pass filter, smoothing out all variations in power up to the typical cycle time, but no further. Otherwise the probabilistic method (and its time-

step validation) treat the store as a ‘black box’. These assumptions about the store are re-examined in more detail in section 2.9.2.

16. The electricity demand is assumed to be completely predictable; it is a direct function of time of day, weekday or holiday, and month and is assumed to have very similar daily profiles to the total national electricity demand: in thesis to the UK national grid electricity demand.
17. The inefficiency loss of the store (due to its finite efficiency) is accounted as energy enters the store, not as it leaves.
18. Within any store time period, T, the level of backup generation employed, grid import power or grid export power is assumed to be constant. This makes the 24-hour version of the method easier to code.
19. The 24-hour program has to assume that the typical cycle time, T is exactly 24 hours. The simple version of the wind-only program can use any store time scale, but the second version of the wind-only program can only model time scales of between 24 hours and 6 months.
20. In a period, T in which average supply exceeds demand plus system losses, the store spends some time filling, some time emptying and some time full but no time empty. Conversely, in a period in which average supply is less than demand plus losses, the store spends some time filling, some time emptying and some time empty but no time full. The probabilistic method uses these assumptions together with the calculated PDFs of wind and solar power, and some further assumptions to predict the total fractions of time spent full, empty, filling and emptying. For a full explanation, see appendix B.
21. The store control strategy number 2 assumes perfect weather forecasting over the operating period of the store. Strategies 1 and 3 do not require any knowledge of future weather but merely respond to variations in supply and demand of electricity.

1.9.8 Programming Tools

Both the probabilistic and time-step programs presented in this thesis were written in Matlab. Matlab is specifically designed to manipulate matrices, using very concise syntax (implied ‘for’ loops), and runs quickly on matrices. Its special functions are also very useful, e.g. the fast Fourier transform (fft).

2 Methodology

2.1 Introduction to the Methodology

This Chapter describes in detail the probabilistic method developed for modelling hybrid electricity systems with significant fractions of wind and/or solar power. The philosophy of the method is to start with a time period, T, that is the characteristic period of an energy store. The energy store smoothes out fluctuations in energy supply and demand on short time scales, up to this period, T but is not able to smooth out fluctuations on longer time scales.

The primary input data are the power spectral density (PSD) functions representing the time varying renewable energy inputs: the power spectrum of wind speed variation, for example Van Der Hoven's spectrum (Van Der Hoven, I. 1957) and an analogous spectrum of solar irradiance variations, together with diurnal and seasonal profiles of wind speed, solar global irradiance and electricity demand. These spectra and profiles are prepared from real measured data, and may be modified by knowledge of the system to be modelled, but the original time series data itself is not used by the probabilistic method. Thus a significant reduction in input data is achieved compared to time-step methods.

Since there are so few such hybrid electricity systems in existence, and none with the same wind and solar climate or electricity demand as the measured data used in this study, the probabilistic method has been validated using a time-step model. The weather data was measured at a different location, about 100 miles away from the electricity demand data, and over a different period of time. Nevertheless, the time-step model uses the original measured wind speed data, solar irradiance data and electricity demand data with suitable adjustments for average wind speed and seasonal alignment.

Section 2.2 describes the sources of data, the times and locations of measurements and how these measurements were prepared for use in the models. Some preliminary analysis of the data is presented.

Section 2.3 briefly describes Fourier transforms, and explains why a Fourier transform was chosen for modelling of electricity systems with renewable energy and energy storage. In the probabilistic method, it is important to capture the essentials of the variation in solar irradiance, wind speed and electrical loads in the most compact form, and to effectively analyse this to predict the behaviour of a specified system.

Section 2.4 evaluates and compares the importance of different sources of variation in renewable energy supply and load on different time scales. Cyclical or deterministic variations are distinguished from stochastic or random fluctuations. The Fourier spectra of wind speed variations, solar irradiance variations and demand variations are presented. As will be seen, the percentage variation in load is much smaller and more periodic than that of wind or solar variations. This section also discusses the covariance or independence of different stochastic variables. For example, minute-to-minute and day-to-day variations in solar irradiance are treated as independent of wind speed variations. This section examines how valid this assumption is.

Section 2.5 describes the principle of spectrum integration used in the probabilistic method. This section is central to the method. An energy store is modelled as a series of three filter functions applied to the spectra of wind speed and solar irradiance variations. Equations are presented for each filter function. The area under each filtered spectrum represents a statistical variance that is used to model the behaviour of an energy system incorporating wind and/or solar power.

A first equation calculates the variance of period-average values of wind speed or solar irradiance. A second equation calculates the variation of wind speed or solar irradiance within a given time period. These two equations extend previously published work by (Infield 1990) and (Bossanyi, Anderson

1984) and are used to evaluate the charging and discharging power ratings of an energy store and the required power rating of a backup power source and/or dump load. A third equation is entirely new to this thesis and evaluates the size (energy rating) of a store used to accommodate variations in wind power or solar power within a given time frame.

Section 2.6 describes the wind power model. The wind speed variation is largely stochastic with small but significant diurnal and seasonal variations. The profiles of monthly average daytime and night-time wind speeds are presented. The standard deviations of wind speed variations within a given time period are assumed to be proportional to the average wind speed over that time period. This is an extension of the concept of turbulence intensity to all time scales, not just variations within 10-minutes or one hour. PDFs of period-average wind speed and wind speeds within each period are constructed from the calculated standard deviations. The wind speeds are converted into values of wind power using a generic wind turbine power curve. The wind power model is relatively straightforward.

Section 2.7 describes the solar power model. The spectrum of variation of solar irradiance has much larger diurnal and seasonal spikes than the wind spectrum. The solar power model must take into account both the variation in cloud cover and the variation in solar elevation. In addition, solar clearness index (the fraction of solar irradiance that penetrates the atmosphere from space) tends to have a bimodal probability distribution. This means that the beam component of solar radiation tends to be either present or absent, with very little middle ground. Thus the solar power model is much more complicated than the wind power model, and section 2.7 is the longest of chapter 2.

Section 2.8 presents the model of electricity demand. Demand is modelled as if it is an entirely deterministic function of time of day, type of day and month. As will be seen, the stochastic variations in demand are small compared to the deterministic variations, and small compared to the wind and solar power variations when wind and solar penetrations are high enough to justify energy storage.

Section 2.9 presents the energy store model. A rule of conservation of energy is applied over the chosen time scale of the store, for example one day. Energy entering the store is equal to energy leaving the store to satisfy demand, plus energy losses. The round-trip efficiency is a single, invariant number between 0% and 100%. The standing loss (self discharge) is a constant power rate. The store has two other characteristics: the maximum rate at which the store can be charged and the maximum rate at which the store can be discharged. Thus the model's input data is a probability distribution of net power (solar power + wind power – demand – losses) and the model's output is the energy lost due to the store being full or the unsatisfied demand due to the store being empty, the energy lost due to the store charge rate being insufficient and the unsatisfied demand due to the store discharge rate being insufficient.

Section 2.10 describes the process of convolution of probabilities: wind power, solar power and demand probabilities in each period, as performed in the 24-hour computer program. Each input probability density function (PDF) is converted to a PDF of power so that all variables have the same units (kW), in order to make convolution of PDFs possible. The solar irradiance PDF is converted into a PDF of solar power and the wind speed PDF is converted into a PDF of wind power. Most of the convolutions are performed using specially written subroutines, because the PDFs are often 'sparse matrices' (containing many zero values), making the standard Matlab convolution subroutine too slow. For example, the PDF of instantaneous solar power contains just two non-zero values for each hour of the day: the probability of 'sunny' or beam radiation and the probability of 'cloudy' or no beam radiation. Furthermore, at night there is only one value of solar power, i.e. zero!

Since the wind, solar and demand PDFs all have significant diurnal and seasonal variations, the convolution is performed separately for each hour of the day, then all the hourly PDFs are combined to give the total daily PDF.

One convolution is performed for the distribution of power variation within one period (day), and another convolution is performed for the combination of different daily averages. The first type of convolution is the one that takes most calculation time since it is performed many more times.

Section 2.11 describes the calculation of store size (energy capacity). The store size results from the application of the third filter function to the wind and solar variation spectra, as described in section 2.5. The store size calculation adds up the components of variance in state-of-charge resulting from each source of random variation (stochastic solar and stochastic wind power variations) together with an average conversion factor for each variable to give a variance of power x time, equal to a variance of energy or state-of-charge.

The variances of state-of-charge above are calculated for wind power and solar power, based on their stochastic variation alone. Since the wind, solar and demand PDFs all have significant diurnal and seasonal variations, the average diurnal variation is treated as a third component of variance, and is added to the first two.

Section 2.12 describes the modelling of grid import, grid export, and backup generation together with three different control strategies available in the 24-hour computer program. Optionally, the modeller can incorporate a grid link to the local electricity system, and/or a backup generator to improve security of supply while minimising capital expenditure on renewable energy supply and energy storage equipment. The ability to export to grid also minimises energy curtailment and potentially gains revenue. However, these options add to the system cost, reduce its self-reliance and increase its complexity.

The grid import, grid export and backup generation options also present the model with control decisions. Power surplus can be directed either to the store or to the grid. Power deficit can be met either from the store or from grid import or from backup generation. The probabilistic method gives the program user three control options to choose from:

1. Keep the store as full as possible at all times, minimising power cuts.
2. Keep the store as empty as possible, minimising power curtailment.
3. Let the store float with perfect weather forecasting, to minimise both power cuts and curtailment, while maintaining a constant import rate, export rate and backup generation power within each typical store period, T.

Section 2.13 describes the two Matlab programs that incorporate probabilistic methods. The first one models only wind power variations with a constant electrical demand, or demand that varies in a very simple way, but can model a store of any size operating over any chosen time scale. This first program is a simplified version, without diurnal and seasonal variations in wind speed, and will be referred to as the 'wind-only' program.

The second program models any combination of wind power, solar power and time-varying electricity demand, including diurnal and seasonal effects, but is only applicable to a typical store cycling time, T, of one day (24 hours). This second program can also model solar power without wind power or vice versa, and can model constant or varying loads. The second program will be referred to as the '24-hour' program.

The probabilistic method also calculates the proportion of time that the store spends full, empty, filling and emptying. These calculations are described in appendix B. Since the probabilistic method has lost the temporal relationships between high and low net powers (store charging and discharging), the calculations of time spent full and empty use informed guesses of a typical charge and discharge cycle, based on the PDF of net power entering or leaving the store in each period.

2.2 Data Gathered

Wind speed, solar irradiation and electricity load data have been gathered and analysed. Temperature data was also obtained but is not used. The relative importance of the variations of loads, wind and solar power are explored in section 2.4

2.2.1 Weather Data From the Rutherford Appleton Laboratory (RAL)

Wind speed and solar irradiation data was made available from the CCLRC Energy Research Unit (ERU) Test Site at the Rutherford Appleton Laboratory. ERU has conducted research into the exploitation of wind power (Infield et al. 1994, Palutikof et al. 1989, Watson, Ter-Gazarian 1996), and the need for energy storage (Infield 1984). The ERU weather station has been collecting data almost continuously since the late 1980s. Unlike most data available from the British Atmospheric Data Centre (BADC), the RAL data is stored as one-minute averages. This is a high enough frequency to capture much of the wind speed turbulence and the variation of solar irradiance due to passing clouds.

Over time, the number of weather measurements has increased, and their quality has improved. Between May 1986 and April 1994, wind speed was measured at two heights and wind direction at one height. From April 1994, atmospheric pressure and two temperature measurements were added. Then from November 1994, two more anemometers and another wind vane were added. Finally, on the 30th March 1998, a pyranometer was added measuring global horizontal solar irradiance.

The measured data contains a few gaps and bad values, but the studies in this thesis use the most complete of the data sets. All gaps have been carefully filled with data spliced in from nearby in the record. Virtually every month had two or three gaps of a few minutes each where the recording tapes were changed. Some months also had longer gaps of an hour or two, or even up to half the month (May and June 1999, May 2001, December 2001, and January 2002). Where the gaps were significant and the weather had changed during the gap, it was filled with data that had similar start and finish values. Where a gap occurred in the temperature or solar record, the gap was filled with data from very nearby on the same day, or similar hours on another day just a few days before or after the gap. In this way, the diurnal patterns of solar irradiance and temperature were preserved. In some cases, the wind speed record exhibited a significant diurnal variation, especially in the summer. In those cases, the wind speed gap was also filled with data from similar times of day.

Sometimes the whole record was missing for a period of time (a problem with the recording mechanism), and sometimes just one data stream (a problem with just one instrument). Gaps from one instrument sometimes manifest themselves as '-999' or some similar number to indicate a known problem. Such data gaps from single-instrument faults were filled by data from other instruments measured during the same period of time. Elsewhere, data sets simply had impossible or highly unlikely values, for example zero wind speed amongst much higher values: When the measured wind speed dropped abruptly from about 8m/s to 0, and then jumped back to 8m/s some time later, this was also recognised as a gap and replaced by other data.

Two data sets have been compiled from the weather data measured at RAL. Each data set is four years long:

2.2.1.1 Four Years Data of Wind Speed Only

One four-year set of wind speed data has been constructed from the wind speed measured between 1994 and 1998 inclusive. There was no useful data between July 1994 and February 1995. Therefore, the first part of 1994 has been spliced onto the second part of 1995, with the join occurring at 28th May. (Weather patterns at the end of May, 1994 appear to be quite similar to weather patterns at the end of May 1995). Thus five calendar years have yielded four years of useful data. Whenever possible, the wind speed data was taken from the anemometer 18.7m above ground level on tower 2, recorded as 'RLT2S19'. When this was unavailable, the second and third choices were the anemometers at 18m above ground level on towers 3 or 4.

Failing these, data was used from the anemometer at 17m above ground level on tower 2.

The resulting four-year data set of clean wind data consists of 2,102,064 records. This is just a few hours short of four years. This data has been used to construct a power spectrum of wind speed variation but has been little used in the probabilistic method or its time step validation presented in this thesis.

2.2.1.2 Four Years Data of Wind Speed and Solar Irradiance

Another four-year data set of wind speed data, together with simultaneous solar irradiance measurements, has been constructed from the data measured between January1999 and December 2002 inclusive. Again, the preferred anemometer was the one at 18.7m above ground level on tower 2. There is only one pyranometer for the period. It is located 3 m above ground level on the roof of building R64.

This second set of data has been used as one of the inputs to a generic power spectrum of wind speed variations, a power spectrum of solar irradiance variations and for most time step validation presented in this thesis.

Since the data set includes solar data, it is more important to capture the diurnal periodicity. More care was therefore taken to make each day’s worth of data exactly the right length, i.e. 1440 records long. The data set of cleaned weather data is 2,103,840 minutes long. This is exactly four years long, including 1 leap day in year 2000. The preparation of the solar irradiance data is described in more detail in section 2.7.1.

2.2.1.3 Physical Location of the Weather Measurements at RAL

The Rutherford Appleton Laboratory (RAL) is located near Didcot, Oxfordshire. Weather measurement instruments include anemometers, a pyranometer and temperature measurements. Depending on wind direction, some of the anemometers are in the lee of buildings or trees.

When the solar elevation is very low, in the morning and evening, the pyranometer can be shaded. However, the instrument is located sufficiently distant from buildings and trees that the effect on measured irradiance is small.

The weather station at RAL is located at latitude 51.57° North and longitude -1.31° East (+1.31° West). This corresponds to a Great Britain grid reference of 447,000 East and 186,000 North. The location is estimated to within one km and to within 0.01° of longitude or latitude from a map of the area.

2.2.2 Wind Data From an East Coast Offshore Site

Some offshore wind speed measurements have been obtained from a location off the east coast of the UK. The exact location and the long-term average wind speed cannot be disclosed for reasons of commercial sensitivity. Data was obtained for a period of almost 2032 days (5.56 years) at 10-minute intervals. This data is useful for calculating a second spectrum for wind speed variation in order to gauge inter-site variation.

2.2.3 Wind Data From the British Atmospheric Data Centre, BADC

Wind speed data has been downloaded for several locations around Great Britain and offshore islands: Cottesmore in Rutland, Butt of Lewis, St Mary's in the Isles of Scilly and Ronaldsway on the Isle of Man. This data was all measured between the years of 1993 and 2000 inclusive as hourly average values and forms part of the 'ESAWIND' database. All the data streams contained a few gaps that had to be filled for the purposes of this project. Gaps were filled with similar periods of wind speed data from elsewhere in the data record, or with concurrent data from another site.

Each cleaned-up BADC data set contained a whole number of years. The Cottesmore data runs from January 1994 to December 1999 inclusive. The Butt Of Lewis data runs from January 1995 to December 1997 inclusive. The Isles of Scilly data runs from March 1997 to February 2000 inclusive. The Ronaldsway data also runs from March 1997 to February 2000 inclusive.

Thus all the BADC data is very useful for creating spectra of wind speed variation.

The characteristic differences between the wind measurement sites (RAL and all the BADC sites) are discusses in section 2.4.11. Some are maritime sites and others are inland; some are on the eastern side of Britain and others on the west coast. These differences have an effect on the levels of wind speed variation, particularly the diurnal and seasonal variations.

2.2.4 Load Data From Leicester

The Author's thanks go to Central Networks (distribution network operator) for use of their load data. Electrical load data was measured at all branches of two primary substations, 'Red Cross' and 'Braunstone' in the south of the City of Leicester, for the year from February 2001 to January 2002 inclusive.

It was decided to use data from just one substation feeder, 'Braunstone LOC-A', rather than the whole substation, because a single branch is easier to check for data gaps, and because the load on one branch is similar in magnitude to the generation capacity of one large wind turbine; the annual average load on this feeder was 1083.225kW. This size of feeder aggregates the demand of individual premises but would still show local effects, for example the effect of local changes in weather conditions, as would be experienced in medium-scale stand-alone systems. This feeder probably supplies several hundred premises that are a mixture of industrial, commercial and residential premises. 'Braunstone LOC-A' was also chosen because its load profiles closely match the national grid load profiles, within a day, weekly and seasonally, see section 2.8.3.

This single-feeder demand data still had gaps, sometimes of just half an hour but sometimes as long as two days. The half-hour gaps could be filled by linear interpolation, but longer gaps had to be filled with data from the same time of day, in days of a similar type in the same month. Weekdays were filled with data from weekdays, Saturdays from Saturdays and Sundays from Sundays.

2.2.5 British Gas Electrical Demand Data

British Gas plc made available measurements of the electrical load of 14 houses in various parts of the UK, supplied as part of the Solar Cities project. A small-scale analysis was done on data from 53 days between Tuesday 12th December 1995 and Friday 2nd February 1996 inclusive. The electrical demand was measured at minute intervals. The British Gas data was not used in the probabilistic method or in its validation, but is presented here is for illustration and justification of the applicability of daily demand profiles in systems of the order of 1MW average demand. The British Gas data illustrates the effect of aggregation of many consumers.

2.3 The Fourier Transforms

2.3.1 Data Compression of a Fourier Transform

The object of using Fourier Transforms in this project is to capture the essential characteristics of a long time-series in a compact form.

Each data set of processed weather from the Rutherford Appleton Laboratory consists of 4 years of data records at a frequency of 1 minute. There are up to 2103840 records for each data stream of wind speed and solar irradiance. The spectrum resulting from each fast Fourier Transform is just 136 pairs of numbers. The logarithmic frequency interval was 0.05 in base 10 logarithms, or a factor of 1.12 from one frequency bin to the next.
The first of each pair is the logarithm of the centre of the frequency bin and the second of each pair is the power density in that frequency bin. This represents a data reduction factor compared to the original time series of:

$$\frac{2103840}{2\times 136} = 7735 \qquad (2.1)$$

The wind and solar power spectra that were used in most of the calculations presented in this thesis consist of these 136 frequencies.

At Rutherford Appleton Laboratory, 3 days of wind speed measurements have also been recorded at an averaging frequency of 5 seconds. This has also been used to create a wind speed power spectrum extending to shorter time scales. It has been possible to splice this extra spectrum onto the end of the main wind speed spectrum to create a spectrum from 4 years down to 10 seconds (the Nyquist frequency for a 5-second sampling period). The Fourier transform method thus lends itself to combining data from several sources for use in one model.

2.3.2 Limitations of Fourier Transforms

Inevitably, some information is lost when the number of data points is reduced so radically. In the case of Fourier Transforms the temporal relationships are lost. Nothing useful can be deduced from the phase information of the Fourier transform. For example, the spectrum does not tell us how the magnitude of the turbulent, short-term variation changes with the magnitude of the weather dependent, long-term wind speed. The spectrum does not even tell us that the wind speed cannot go negative! Other assumptions have to be made in a practical model, and those assumptions have to be based on the physics of the real system.

When modelling wind speed variations, the probabilistic method assumes that the standard deviation of variations within a period is always proportional to the period-average wind speed. This is an extension of turbulence theory to all time scales. Other assumptions were tried, see section 2.6.5 and produced no improvement in the predictions.

When modelling solar irradiance variations, the probabilistic method simply assumes that the sky cloud state is bivalent – either sunny (beam radiation present) or cloudy (no beam radiation present). The clearness indices are chosen to give the correct total variance and the correct monthly and hourly average radiation. The solar model also makes extensive use of sun-earth geometry to calculate the solar elevation and therefore the extra-terrestrial global horizontal irradiance and the air mass.

2.3.3 The Matlab Fast Fourier Transform Subroutine

As part of this project, various data have been processed using the standard Fast Fourier Transform function within Matlab, called 'fft'. In its raw form, the 'fft' function is not very useful. Its output is a data set of complex numbers, with one number for each frequency from zero up to the sampling frequency, in steps of the total smallest possible frequency interval, i.e. 1/total period. Thus the number of output complex numbers is equal to the total number of data samples.

The first number in the output is the zero frequency or D.C. component, equal to the total of all numbers in the data series. This is discarded.

Since the Nyquist frequency is half the sampling frequency, the second half of the output data set is redundant. The magnitudes are a mirror image of the first half, reflected at the Nyquist frequency. The second half of the data set is therefore also discarded.

The real and imaginary parts of each complex number represent the sine and cosine components of the data series. The methods presented here do not need to know the relative magnitudes of the sine and cosine components. All that is required is the magnitude of each frequency component, equal to the modulus of each complex number. In fact, we require the component of statistical variance, or ‘power’ represented by each frequency component. This is equal to each modulus squared. In order to make the total power representative of the total statistical variance of the data series, each power component has to be normalised by dividing it by the total number of data points in the original data stream.

The phase angle of each frequency component is the argument of each complex number. The phase angles have been plotted of Fourier transform results from several wind speed time series, and found to be apparently completely random. They therefore convey no useful information to a probabilistic modelling method.

The Fourier transform is calculated on equal increments of frequency. The magnitudes of the low frequency components tend to be greater than those of the high frequency components, sometimes by several orders of magnitude. The total variance represented by the higher frequencies is similar to that of the lower frequencies, but the high frequency variance is distributed between more frequencies. In order to obtain a useful, smooth power spectrum on a logarithmic axis of frequency, the magnitudes of the spectral components must be binned according to a logarithmic frequency scale. The higher frequency bins have many Fourier transform components, each of small magnitude, whereas the lower frequency bins have just a few components of larger power. At the lowest frequency end of the spectrum, there may be just one or no components in each bin. In these cases, the spectral amplitude in each bin must be interpolated between adjacent Fourier transform points.

Because half the data set was discarded, the resulting spectrum only represents half the total variance of the original wind speed time series. When calculating the variance represented by an interval of frequencies, the amplitude of the spectrum must be doubled again, (Bendat, Piersol 1993), pages 10 to 12. A Matlab programme has been written that performs all the above conversions to create a smooth spectrum, the integral area of which is equal to half the total variance of the original data stream, within numerical accuracy. The amplitude of this spectrum must then be doubled to give the correct total variance.

The raw Fourier transform output contains as many frequencies as the number of data points in the original time series. Since this consists of several thousand or even millions of time steps, the Fourier transform output has been binned into intervals of 0.05 on the logarithmic time scale in base 10 to create smooth spectra. Thus each data point in the binned spectrum represents variation with a frequency 1.12 times as big as the next slowest, and 0.89 times the next fastest. In the Fourier transform, the high frequencies are closely spaced whereas the low frequencies are widely spaced on the logarithmic scale. Therefore high frequency bins contain a sum of many amplitudes but low frequency bins contain just a few, or even just one or two. Therefore, when binned, the high frequency, right hand end of each spectrum is smooth but the low frequency, left hand end is rough with random scatter.

This process of normalisation and binning has been written into a Matlab program written around the Matlab fast Fourier transform function. This program has been adapted and applied to time series of wind speed data and solar data from Rutherford Appleton Laboratory (RAL) to create the spectra that form part of the input data of the probabilistic method. The program has also been applied to other wind speed data from around the UK and to synthetic time series of wind speed data, see section 2.4.11. The program has even been applied to time series of electrical demand, see section 2.4.5, although the resulting spectrum is not used in the probabilistic method.

2.4 Sources of Variation

2.4.1 Relative Importance of the Different Sources of Variation

Power variations originating from wind power, solar power and electrical demand have been compared. The purpose of this work was to support the assumptions and simplifications made in the probabilistic modelling method. Some representative time series of data from each of the variables have been processed through a fast Fourier transform to create power spectral densities, PSDs.

2.4.2 Magnitudes of Variation in the Original Data

The wind data used in this comparison was RAL wind data, 1994 to1998, whereas the solar data was measured at the same location between 1999 and 2002 inclusive, see section 2.2.1.

Some electrical demand data was measured at a substation feeder in the city of Leicester between February 2001 and January 2002 inclusive. However, this data is only measured as half-hour averages. Electrical demand data was therefore also taken from the National Grid, (National Grid Company plc 2003), where data is given for four example days at one-minute intervals. This National Grid data was only used to create a spectrum of short-term variation, to splice with the spectrum from the Leicester data where the Leicester data could not provide a spectrum. Table 2.1 lists the means and standard deviations calculated from the original time series of data.

Table 2.1 A Simple Comparison of Variation in the Raw Data

	Solar Data RAL 1999 To 2002	Wind Data RAL 1994 To 1998	Demand Data Leicester 2001 To 2002
Mean	118.8 W/m^2	5.09 m/s	1083 kW
Standard Deviation	201.8 W/m^2	2.61 m/s	262 kW

The standard deviation of solar power is almost twice as big as its long-term average. At first sight this seems ridiculous, but the distribution of solar power, fig. 2.1, is very highly skewed. There is a very high probability of zero or near zero irradiance and a low probability of high irradiance. Obviously for almost half the time, the sky is dark and irradiance is zero. In the UK, the inter-seasonal variation is also very strong: The average June irradiance is typically 10 times bigger than the December irradiance. Then the solar elevation angle changes during the day, and cloud cover varies significantly from minute to minute and from day to day.

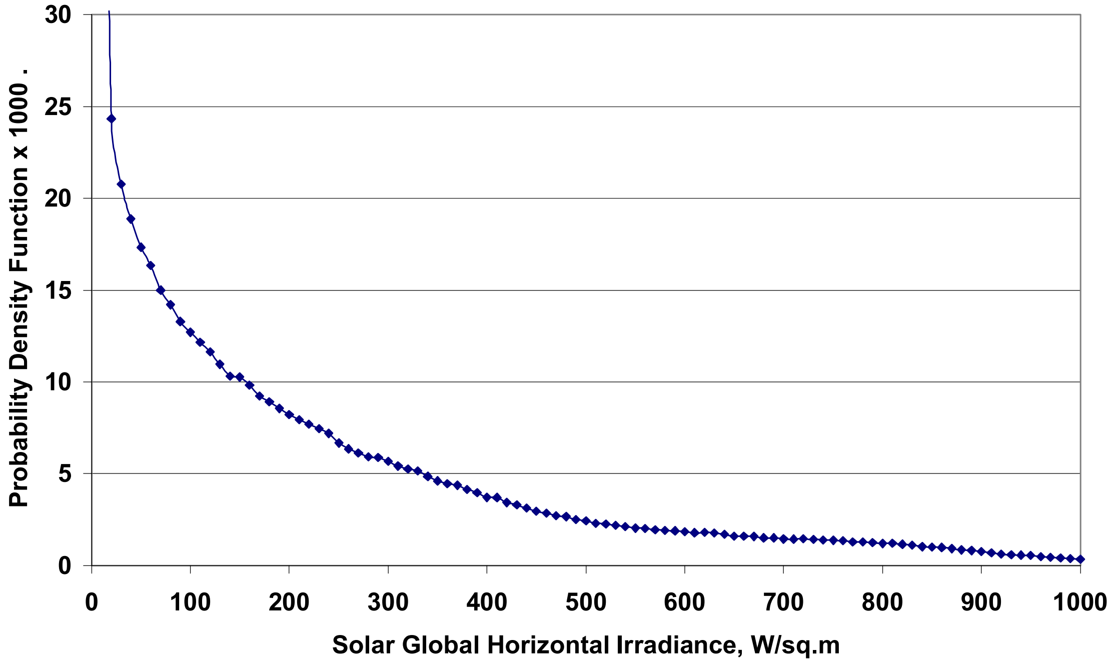

Figure 2.1 Probability Distribution of Solar irradiance at Rutherford Appleton Laboratory From 1999 to 2002

Wind speed data has a smaller variation, having a standard deviation that is just over half of the mean value. The probability distribution of wind speed at RAL is a fairly typical distribution for a temperate climate, and can be approximated by a Weibull distribution with a shape factor of 2.04 and a scale factor of 5.75m/s, fig. 2.2. This Weibull distribution has the same mean and standard deviation as the measured distribution but a slightly different shape, as discussed further in section 2.6.2.1.

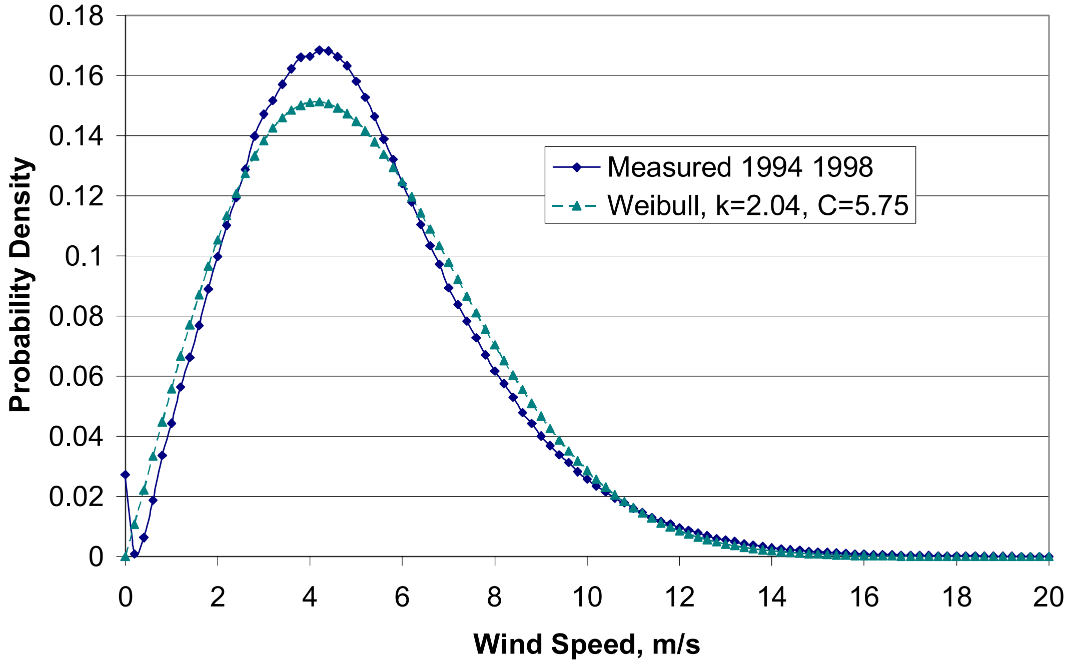

Figure 2.2 Probability Distribution of Wind Speeds at Rutherford Appleton Laboratory From 1994 to 1998

Electrical demand data has quite a small variation, having a standard deviation that is only one quarter of its mean value. This is reflected in the graph of demand probability distribution, fig. 2.3. The smallest recorded value is about half the mean, and the largest recorded value is only about twice the mean.

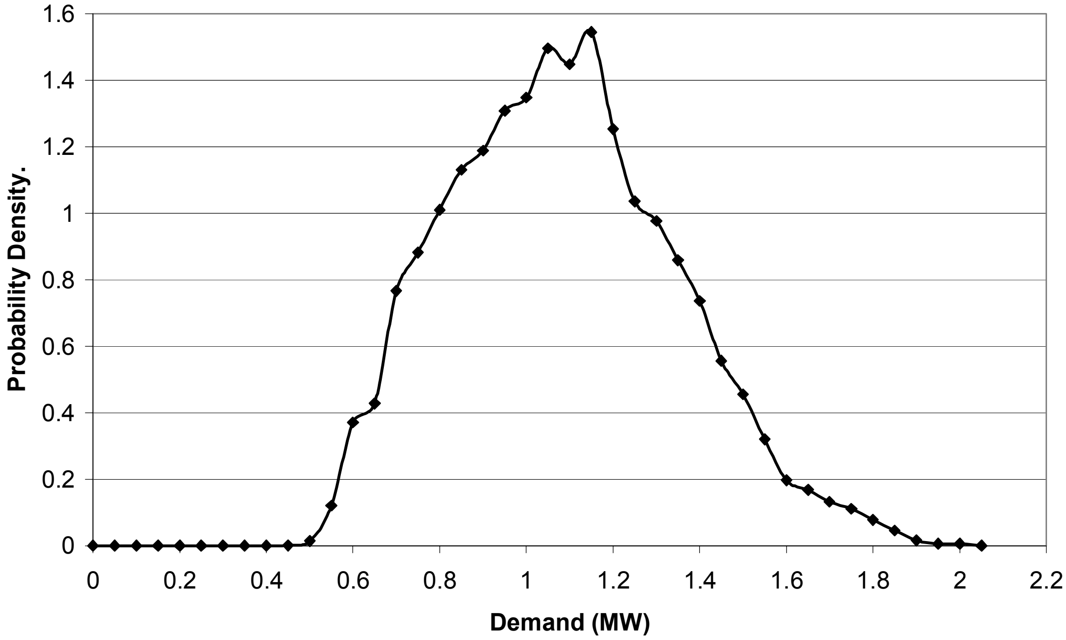

Figure 2.3 Probability Distribution of Demand at Feeder 'LOC-A' in Leicester From 2001 to 2002

2.4.3 Effect of the Wind Turbine Power Curve

The power output from a photovoltaic module is approximately proportional to the radiation it receives, but the power output from a wind turbine is not proportional to the wind speed. The theoretical power available from the wind is proportional to the swept area of the turbine multiplied by the cube of wind speed, and the actual power depends on the power curve of the wind turbine.

A standard, typical turbine curve has been used in this thesis, fig. 2.4 and has also been used in previous publications, (Barton, Infield 2004). This curve has been constructed from an average of three commercially available 1 MW wind turbines (Bundesverband WindEnergie e.V. 2000). Fig. 2.4 shows how a 5^{th} order polynomial has been fitted to the data. Fuller details of the calculation of the turbine power curve are given in section 2.6.4.

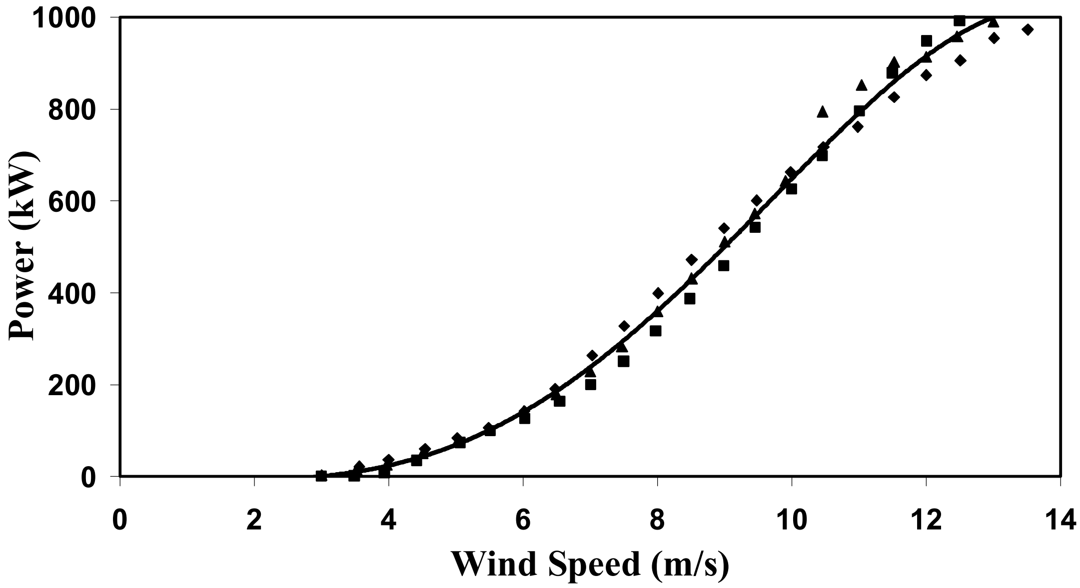

Figure 2.4 Generic Wind Turbine Power Curve

A four-year time series of wind power values has been calculated from the 1994 to 1998 wind speed data, using this turbine power curve, but note that the wind speeds were first scaled from an average of 5.09m/s to 8m/s, which is a more realistic value of average wind speed for a wind farm development. The resulting probability distribution of the wind power is shown in fig. 2.5.

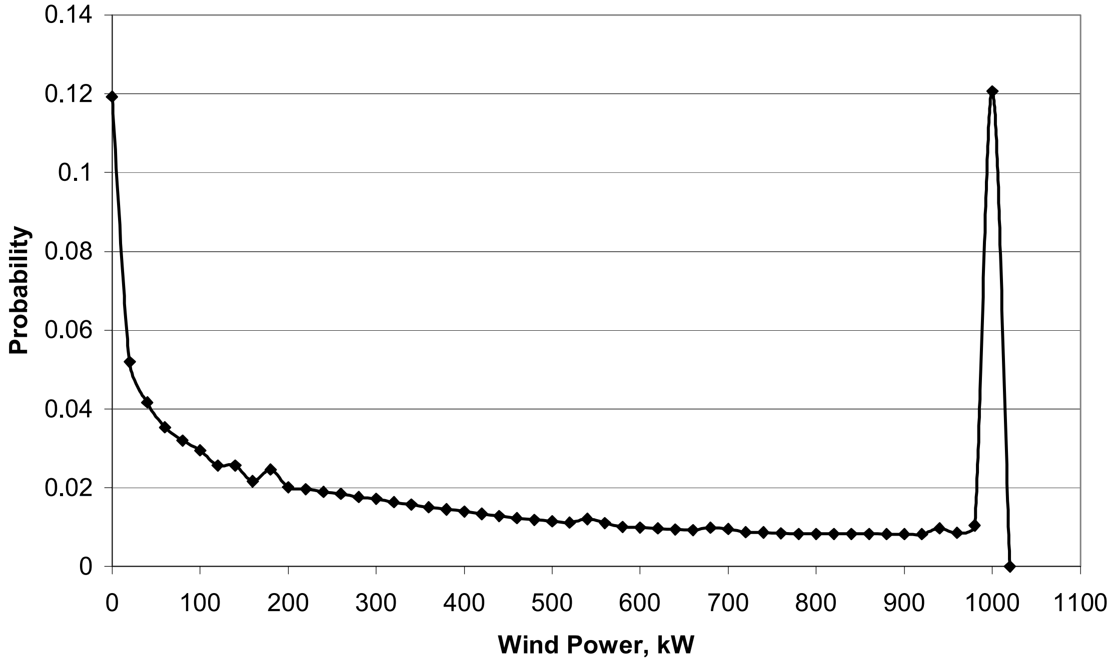

Figure 2.5 Probability Distribution of Wind Powers Based on a 1MW Wind Turbine and Data From Rutherford Appleton Laboratory, 1994 to 1998

The average calculated power is 393 kW and the standard deviation is 353 kW. Thus the standard deviation of wind power is almost as large as its mean value. The effect of the wind turbine power curve has been to increase the variability of the wind power relative to the variability of the wind speed data. This is not surprising, given that the turbine power curve is quite steep around the scaled mean value of 8 m/s. The probability graph also shows that there is a spike of high probability at zero power (12%), and another spike of high probability at turbine rated power output (also 12%), with a broad probability distribution between.

The original wind speed time series and the wind power time series have both been put through a fast Fourier transform program. The resulting spectra are shown in fig. 2.6:

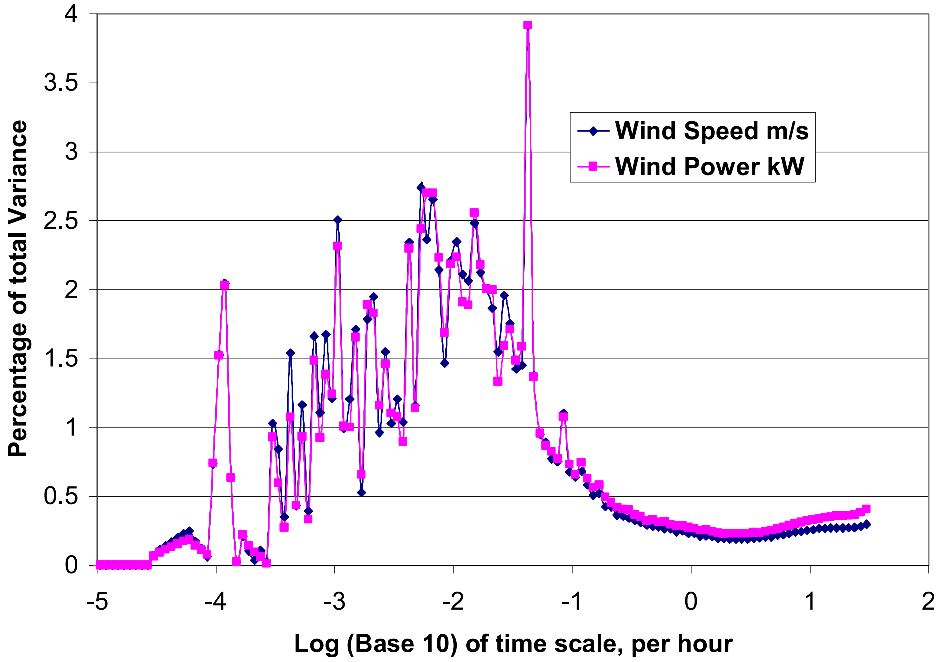

Figure 2.6 Normalised spectra from wind speed and wind power time series

This graph shows that the wind speed scaling and the wind turbine power curve make only a small difference to the resulting (normalised) spectra. High frequency components are slightly increased in significance and some low frequency components are decreased in significance, but the effect is very subtle. The relative importance of different regions of the spectrum is unaffected.

Elsewhere in this thesis, the Fourier transform of the original wind speed time series was used as the basis of the probabilistic modelling method.

2.4.4 Aggregation of Demand - British Gas Data

An example subset of the British Gas data taken from 3pm to 4pm on Tuesday 12th December 1995 is shown in fig. 2.7:

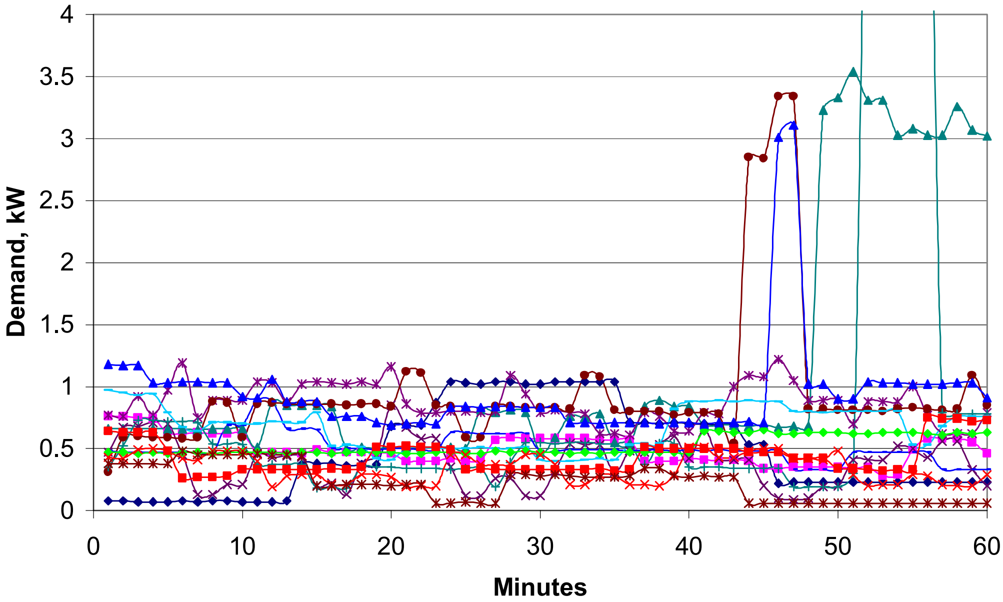

Figure 2.7 Electrical load profiles from individual houses

Each house's demand varies from near zero to several kilowatts, with the peak demand being off the top of the scale at over 6kW, but the average load per house is only 0.64kW. This enormous variation is made up of minute-to-minute variation of each house and variation between houses. The minute-to-minute standard deviation of demand for each house is typically 52% of its hourly average demand, and the standard deviation of house-to-house variation of hourly averages is 45% of the average of all houses.

When demand is averaged over all 14 houses, fig. 2.8, the time series is much smoother. The minute-to-minute standard deviation of total demand within one hour is only 17%. Even with just 14 houses, the random fluctuations due to individual households have been largely smoothed out, and the profile within an hour period is smoother. This justifies the use of hourly averages when modelling an electricity system with average loads of the order of 1 megawatt or more.

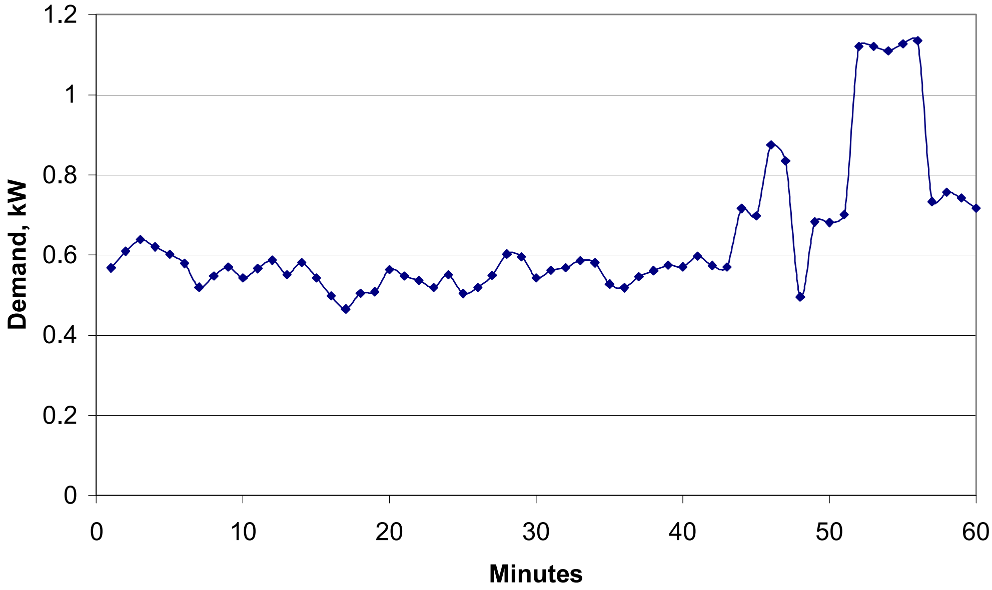

Figure 2.8 Average electrical load profile of 14 houses

A typical substation load of 1MW (excluding a single large load to an industrial premises) would comprise of 1000 houses or more, or a mixture of several hundred commercial, small industrial and residential premises.

As a worked example: If all the houses have a minute-to-minute standard deviation of demand that is 50% of their hourly average value, and all the houses have a similar magnitude of hourly average demand, then we can apply the central limit theorem to calculate the net minute-to-minute standard deviation of demand for a feeder supplying approximately 1000 houses or other premises: Net standard deviation = $50\%/\sqrt{1000} = 1.6\%$. Similar arguments apply to instantaneous house-to-house variation.

2.4.5 Demand Spectrum and Profiles - Leicester Braunstone LOC-A Data

The data from the Loc-A feeder of the Braunstone substation, as described in section 2.2.4 has been used to create a power spectrum of load variations, fig. 2.9 and also profiles of demand for each hour of the day, for weekdays separately from weekend days, and for each month of the year, sections 2.8.3 and 2.8.4.

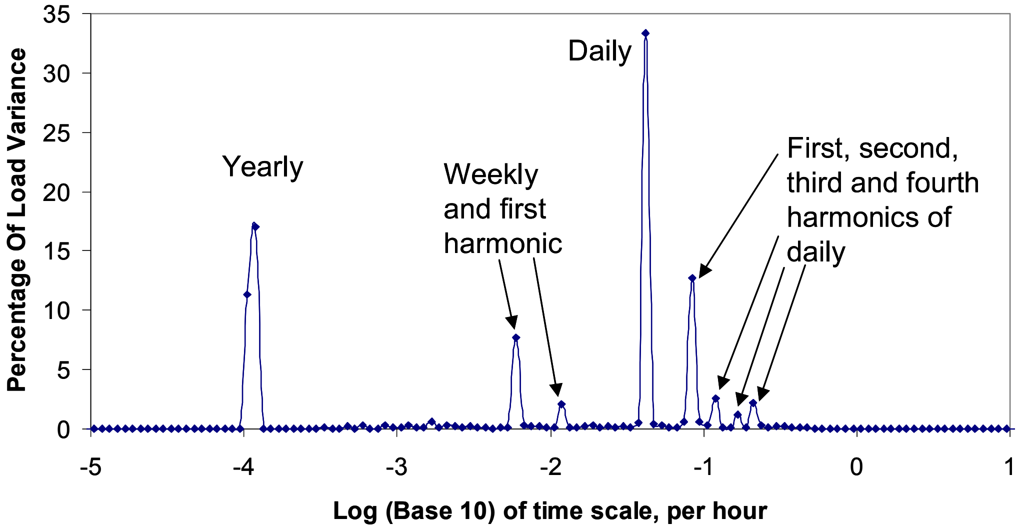

Figure 2.9 Power spectrum of electrical demand at Braunstone Loc-A feeder

Fig. 2.9 shows that the demand spectrum is dominated by periodic variations that appear as spikes in the spectrum.

2.4.6 Spectra of Power Variations

The spectrum of variation (not to be confused with the solar electromagnetic spectrum) has also been prepared by a Fourier transform from the four years of irradiance data measured at RAL, section 2.2.1.2 and is shown in fig. 2.10 and fig. 2.11 below.

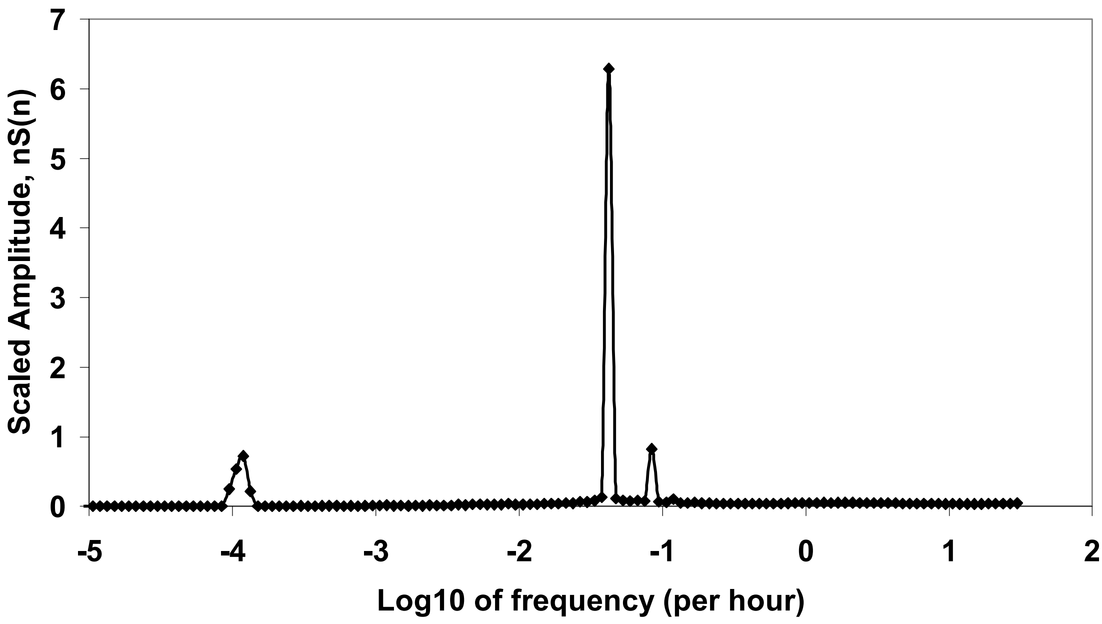

Figure 2.10 Solar power spectrum

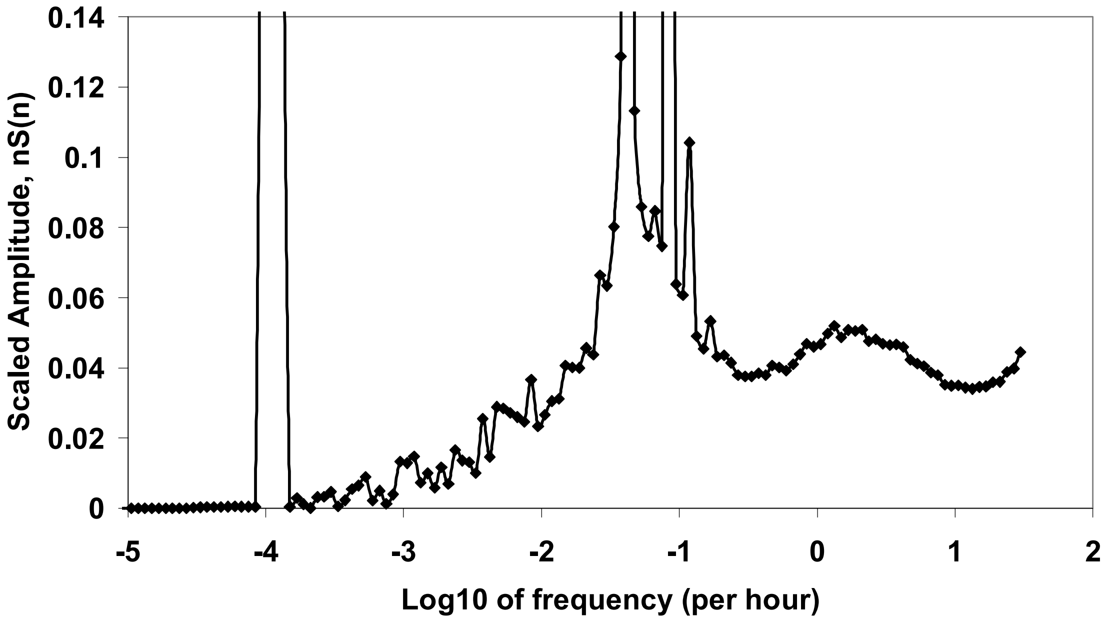

Figure 2.11 Solar power spectrum with expanded vertical scale to show broad spectrum

Spectra of wind power, solar power and electrical demand have been normalised to the same total variance (total 100%) and compared in fig. 2.12.

The wind power spectrum shown is the one resulting from the application of the generic wind turbine power curve onto wind speed data that has been scaled to an average of 8m/s. Solar power is assumed to be proportional to horizontal global irradiance. As discussed in section 2.7.15, the effects of inclined PV modules, and PV module temperature are difficult to model and have been left for future study.

Electrical demand is simply put into a Fourier transform to create a demand power spectrum. Again, the National Grid data was used to create the high frequency portion of the spectrum, for minute-to-minute variations within one hour.

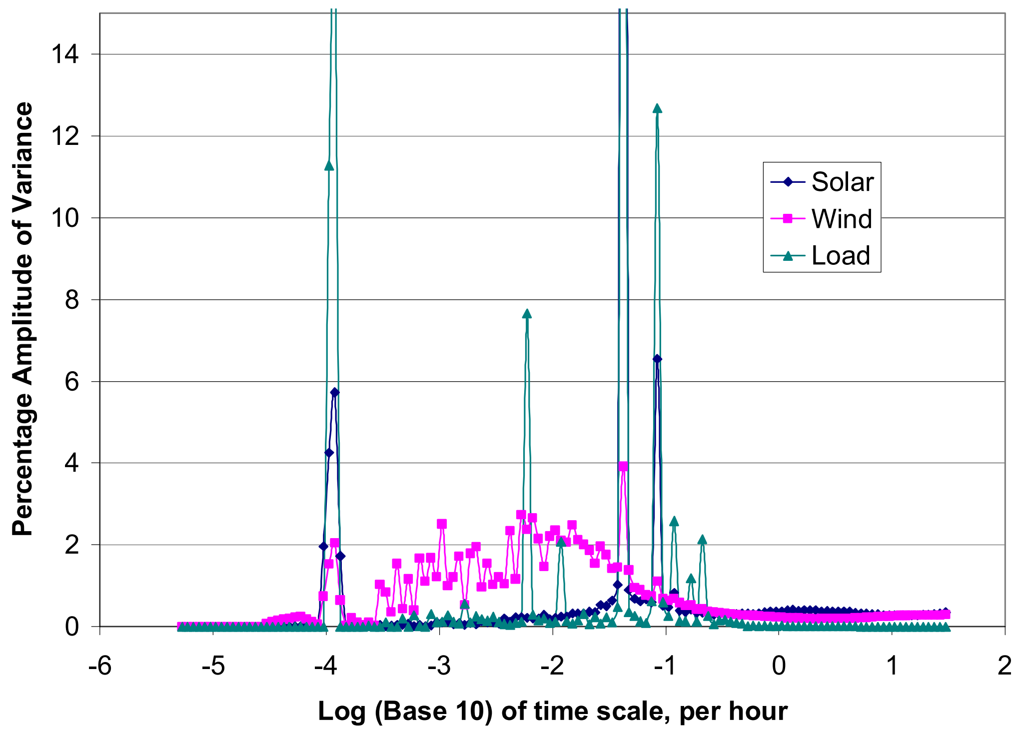

Figure 2.12 Variation spectra of solar power, wind power and electrical demand. Each spectrum is shown normalised by its own total variance.

These spectra show how much variation occurs in a given time frame, or section of the time spectrum, from minutes to hours, days and months. The power spectra also show how much variation is stochastic (apparently random, producing a broad spectrum) and how much is periodic (deterministic, producing discrete spikes in the spectrum). The area under the power spectra curves can be integrated to give the components of variance contained in each part of the spectra, see table 2.2.

Table 2.2 Components of Variance in Each Part of the Power Spectra of Solar, Wind and Load Variations

Spectral Range	Solar power Percentage Variance	Wind Power Percentage Variance	Load Percentage Variance
> 1 Year	0.03%	1.25%	Not Measured
Seasonal	13.67%	4.93%	28.24%
1 Year to 1 Week	2.26%	32.09%	3.50%
1 Week Peak	None	None	7.66%
First Weekly Harmonic	None	None	2.07%
1 Week to 1 Day	6.15%	33.98%	2.71%
Diurnal Peak	49.94%	3.92%	33.37%
First Diurnal Harmonic	6.55%	1.08%	12.68%
Second Diurnal Harmonic	0.83%	Negligible	2.58%
Third Diurnal Harmonic	Negligible	Negligible	1.18%
Fourth Diurnal Harmonic	Negligible	Negligible	2.14%
1 Day to 1 Hour	10.51%	14.02%	3.72%
1 Hour to 2 Minutes	10.08%	8.74%	0.15%
Total	**100%**	**100%**	**100%**
Total Periodic Components	70.98%	9.92%	89.92%
Total Stochastic Components	29.02%	90.08%	10.08%

The vertically expanded solar power spectrum, fig. 2.11 shows a broad spectrum of variation from months down to minutes, with a clearly defined ‘hump’ of variation centred at periods of 30 minutes due to passing clouds.

The variation from day to day is due to passing weather systems and is smaller in magnitude. The full solar spectrum, fig. 2.10 also contains several very large spikes at discrete frequencies corresponding with seasonal variation, diurnal variation and harmonics of diurnal variation. The height of the diurnal spike is so large that its peak is far above the top of the combined figure, fig. 2.12; to make the scale fit this peak would have made all broadband variation too small to be visible at the bottom of the graph. The total variance represented by the seasonal spike, diurnal spikes and diurnal harmonics represent 71% of the total solar variance (the diurnal spike alone is almost 50% of solar variation). The stochastic variation represents just 29% of solar variance. At a lower latitude (nearer the equator) the seasonal periodic component of solar variance would be smaller but the diurnal component would remain very large.

The wind power spectrum also shows a spectrum of variation from months to minutes, with a broad hump of variation centred around one week, fig. 2.12. Most of the stochastic variation in wind power occurs over long time periods – longer than one hour, indicating that it is due to passing weather systems and not due to turbulence. Turbulence intensity is therefore quite modest at the measurement height of 18 metres at RAL. The wind power spectrum contains several spikes at discrete frequencies corresponding with seasonal and diurnal variation, fig. 2.12, but these spikes are much smaller than the corresponding spikes for solar power. Wind power variation is mostly random. 90% of its variance is stochastic and only 10% is periodic. Nevertheless, periodic wind power variations are significant. At RAL, the wind power spectrum seasonal spike contains almost as much variance as the diurnal spike and its harmonic.

So we have seen that solar power variation is mainly periodic but wind power variation is mainly stochastic, and that the random variation of solar power occurs on shorter time scales than most wind power variation. Solar power is therefore considered to be more variable than wind power, but more predictable, when designing stand-alone electricity systems, for example (Allen, Todd 1995).

As seen above, the spectrum of demand variation is almost entirely periodic: The yearly variation spike has no visible harmonics because it is so close to a perfect sinusoid. Weekly and daily variations have visible harmonic spikes. The total periodic variation represents approximately 90% of all demand variation while random fluctuations represent just 10%.

The areas under each line in fig. 2.12 reflect the total percentage variation of each set of data. The standard deviation of wind power is 90% of its mean, the standard deviation of solar power is 170% of its mean and the standard deviation of demand is only 24% of its mean.

2.4.7 Comparison of Absolute Variance in an Example Power System

A hypothetical, simple power system has been examined in which the total demand is exactly balanced, on average, by equal amounts of wind power and solar power. Thus the long-term average wind power (0.5MW) and the long-term average solar power (0.5MW) were each assumed to be equal to half the long-term average demand (1MW). All system losses and energy storage losses are neglected. Variances in system power due to stochastic and periodic components of supply and demand are compared in table 2.3.

Table 2.3 Components of Variance in an Example System Incorporating Solar Power, Wind Power and Varying Demand

Quantity	Solar power Variance	Wind Power Variance	Load Variance
Power Capacity (MW)	0.5	0.5	1.0
Total Variance Calculated from Time Series $(MW)^2$	0.7215	0.2019	0.0585
Total Variances Calculated from Spectrum $(MW)^2$	0.7251	0.2017	0.0574
Total Variances as a Percentage of Total System Variance	73.67%	20.49%	5.83%
Periodic Components Calculated from Spectrum $(MW)^2$	0.5147	0.0200	0.0517
Periodic Components as a Percentage of Total System Variance	52.30%	2.03%	5.25%
Stochastic Components Calculated from Spectrum $(MW)^2$	0.2104	0.1818	0.0058
Stochastic Components as a Percentage of Total System Variance	21.38%	18.47%	0.59%

The total variances calculated from the spectra agree with the variances calculated from the time series, to an accuracy of about 1%. The other results show the characteristics of the different sources of variation, and their relative importance.

The periodic solar power variation is responsible for just over half (52%) of the total system variance on its own. The second largest source of variance is the stochastic variance in solar power, at 21% of total variance. Thus the solar power variation is responsible for 74% of system variance.

The wind power stochastic variance is almost as big as the solar stochastic variance, at 18% of total variance. The wind power periodic variance is only 2% of total variance. Thus the wind power variance is much smaller than the solar power variance, at least for the weather data measured at this site at RAL, and is responsible for only 20% of the system variance.

The periodic component of demand variance is 5.25% of system variance and the stochastic component of demand variance is only 0.59%. Thus the demand variance is much smaller even than the wind power variance, at less than 6% of total system variance. This is despite the average demand being twice the average solar power and twice the average wind power.

Because the variance in the demand is so small, and the stochastic component is only a small fraction of the periodic component, the demand may be considered to be entirely periodic to a good approximation. The Leicester substation demand data was only measured as half-hour averages, and only covers a fraction of the time that the weather data covers, in a different part of the country. However, as the electrical load is almost entirely determined by the time of day, day of the week and time of year, it can be modelled as a series of hourly profiles. It does not matter that the demand is not concurrent or collocated with the weather data.

The relatively large variations of the intermittent renewable sources of energy, especially of solar power, indicate that the challenge of matching demand with supply in such an electricity system is likely to be an order of magnitude greater than in a system with conventional generation and variable load, even if that generation were constant and inflexible. However, much depends on the seasonal and diurnal correlation or anti-correlation of supply and demand, as described in the following sub-sections.

2.4.8 Dependence and Independence of Sources of Variation

In practice, some of the periodic variations reinforce each other, for example solar power with lighting loads. Other variations are complementary, for example solar power with air-conditioning loads, or winter wind power with winter heating and lighting loads. It is important to understand where sources of variation reinforce and where they complement one another and cancel out. This thesis examines the UK environment, where the weather is highly seasonal, but where winter space heating and lighting loads are much more important than summer air-conditioning.

In the probabilistic method, correlation and anti-correlation are modelled as periodic effects. Solar irradiance, wind speed and electricity demand all vary with month of the year and with hour of the day. This makes the modelling relatively simple: daily and seasonal profiles are used for each variable and simply added or subtracted but no correlation coefficients are required.

There may be some residual correlations or anti-correlations not modelled by periodic effects, but these are left for further work. For example, the probabilistic method assumes that instantaneous wind speed is independent of passing clouds (a reasonable assumption) but also that daily average wind speed is independent of daily average solar irradiance within each month. This last assumption depends on the pattern of cloud cover with high and low pressure weather systems and is less obvious.

2.4.9 Seasonal Effects

The 24-hour probabilistic computer program includes models of wind power, solar power and variable electrical demand and performs a separate calculation for each month. The program could have just used 4 seasons, but the electrical demand and wind speed variation apparently lag behind the solar variation by a month or two for the case of England and Wales. The normalised monthly average wind speeds, solar irradiances and demands are shown in fig. 2.13.

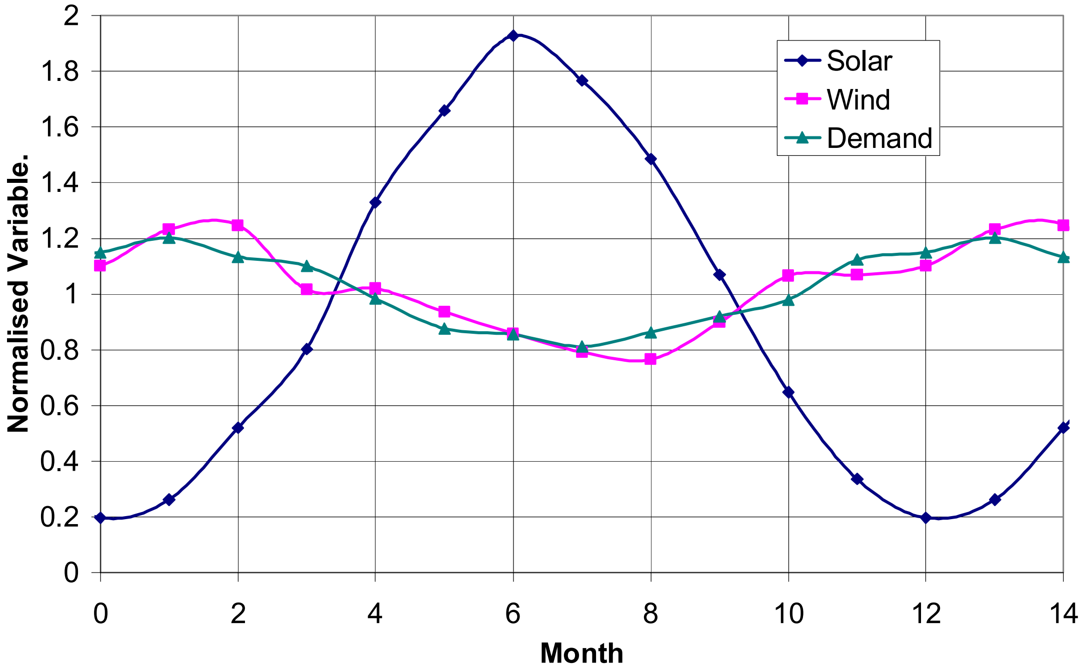

Figure 2.13 Seasonal variations of solar irradiance, wind speeds and electrical demand, normalised by their own mean values.

In fig. 2.13, month ‘0’ represents December, ‘1’ represents January and so on to ‘14’ representing February the following year. This figure shows that solar irradiance is at maximum in June and minimum in December, as expected. However, demand is at minimum in July and maximum in January, probably driven by heating loads, (McSharry, Bouwman & Bloemhof 2005). Wind speed, and therefore wind power, is at minimum in August and maximum in February. Thus the wind power profile lags the demand profile by one month. The solar profile leads the demand profile by one month and is seasonally out of phase with demand and wind power.

2.4.10 Implications for the Probabilistic Modelling Method

Solar power has such a large variance that both periodic and stochastic variations need to be included. Periodic variations are modelled on an hour-by-hour basis for each month of the year.

Wind power also has such a large variance that both periodic and stochastic variations are included. However, the periodic variations of wind power are smaller than solar power, so they are simply modelled as separate daytime and night time averages for each month of the year.

Electrical demand has a smaller total variance, and its variation is almost entirely periodic. The demand is therefore modelled as a profile of demand on an hourly basis for weekdays and weekend days, for each moth of the year. The stochastic variation of demand is completely neglected.

2.4.11 Alternative Wind Speed Spectra from UK Data

The general application and therefore usefulness of the probabilistic method depends on the similarity of the wind power spectrum at different locations.

Some other wind speed data series have therefore also been gathered, as described in sections 2.2.2 and 2.2.3, and used in Fast Fourier Transforms to create wind speed spectra. These are compared with the RAL spectrum to see how the shape of the wind speed spectrum depends on geographical location, fig. 2.14.

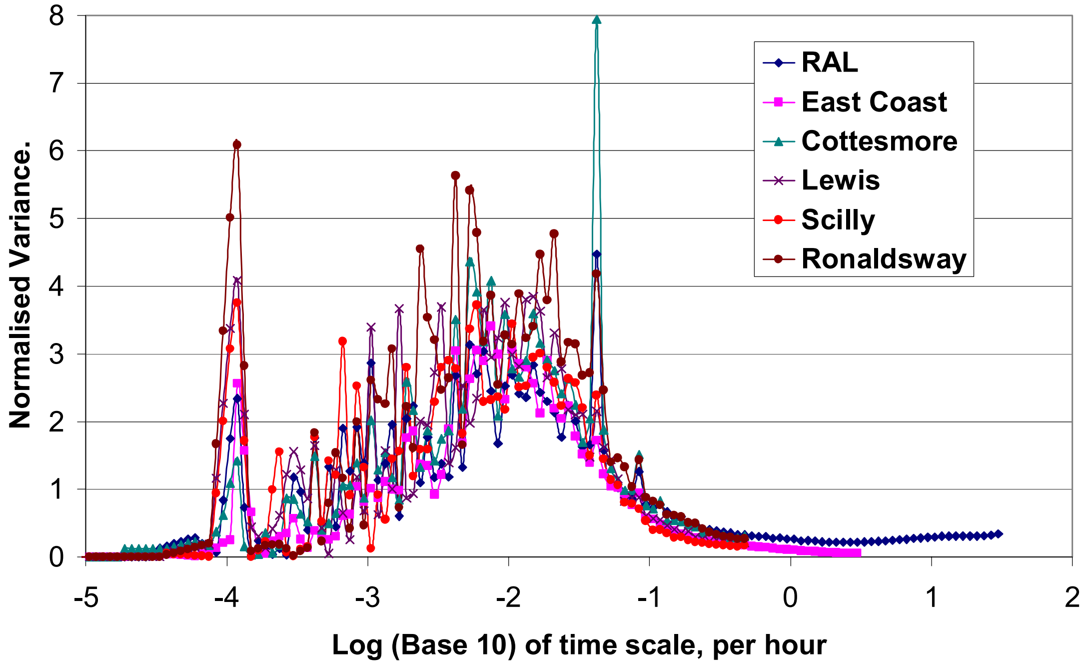

Figure 2.14 Measured wind speed spectra compared, all normalised to 10m/s average wind speed

The above spectra have been normalised by the square of the long-term mean wind speeds and then multiplied by 100 (equivalent to a long-term mean wind speed of 10m/s) to produce a useful vertical scale when plotted. All the spectra have qualitatively very similar spectra. At the left-hand end of the graph is a large spike at –3.94 on the x-axis (= $-\log_{10}$ of 8766 hours) representing annual variation in wind speed. Since this variation is an almost perfect sine wave, this spike has no visible harmonics for these sites in the British Isles. Moving to the right (to shorter time scales), all spectra show a broad hump of weather-related variation. This contains quite a lot of scatter, probably because the data does not span enough years to produce a smooth curve. The humps all have a maximum at about –2 on the x-axis (about 100 hours or 4 days), corresponding to the typical period of passing weather systems. All the spectra also exhibit a spike at –1.38 (= $-\log_{10}$ of 24 hours) on the x-axis corresponding to a period of one day, or the diurnal variation, although the magnitude of this component varies considerably. Then further to the right, many spectra also show the first harmonic of the diurnal spike. This may represent non-sinusoidal diurnal variation or winds that reverse during the day, for example sea breezes. The curves to the right-hand side of the hump are smoother than to the left because each frequency component completes many cycles during the total data period, and each data point on the curve is the sum of many Fast Fourier Transform coefficients. Only one spectrum, that from Rutherford Appleton, extends far enough to show the classic spectral gap of the Van Der Hoven spectrum and at least part of the turbulent spectrum to the right of the gap. The RAL spectrum alone was measured down to a frequency of 5 seconds, see section 2.3.1.

The following sections briefly describe the sources of other wind speed data and compare the remaining differences in their spectra.

2.4.11.1 East Coast

This spectrum is derived from offshore wind speed data, gathered for a proposed wind farm off the East coast of the UK. The exact location and average wind speed have been withheld for reasons of commercial confidentiality.

This spectrum exhibits a relatively small total variance of wind speed in relation to the square of average wind speed, fig. 2.15. The diurnal variation (the spike at -1.4 on the x scale) is the smallest of all the spectra. This is to be expected, since the temperature of the surface of the sea remains very constant between day and night, and thermally driven convection will therefore change little with time of day.

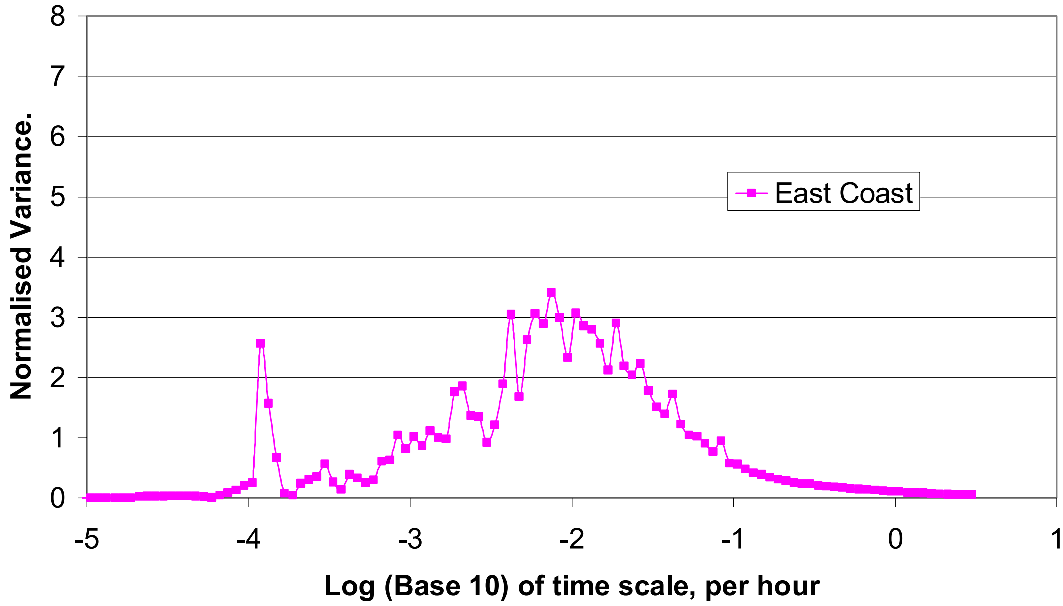

Figure 2.15 Wind speed power spectrum from an east coast offshore site, normalised to 10m/s average wind speed

2.4.11.2 Cottesmore

Cottesmore is located in Rutland, in the English Midlands. Wind speed data was measured on an hourly-average basis for a period of 6 years.

In contrast to the east coast site, this spectrum has the highest diurnal spike of variation, even higher than that of the RAL site, fig. 2.16. Cottesmore is an inland site, many miles from any coast, so thermally driven winds are expected to change a lot between day and night. However, the seasonal variation spike at Cottesmore is the lowest of all the spectra. The reason for this is not known.

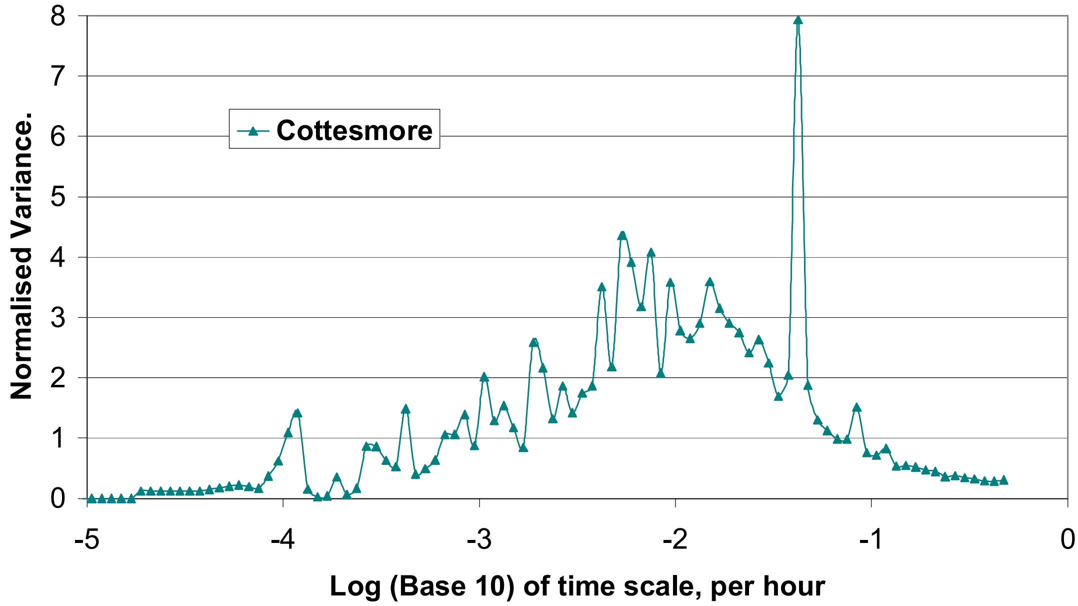

Figure 2.16 Wind speed power spectrum from Cottesmore, normalised to 10m/s average wind speed

2.4.11.3 Butt Of Lewis

This is the most northerly site studied in this thesis, and is on an island off the west coast of Scotland. Hourly average wind speed data was recorded for a period of 3 years.

This maritime location gives the spectrum a relatively small diurnal spike of variation. The seasonal spike of variation is large, fig. 2.17.

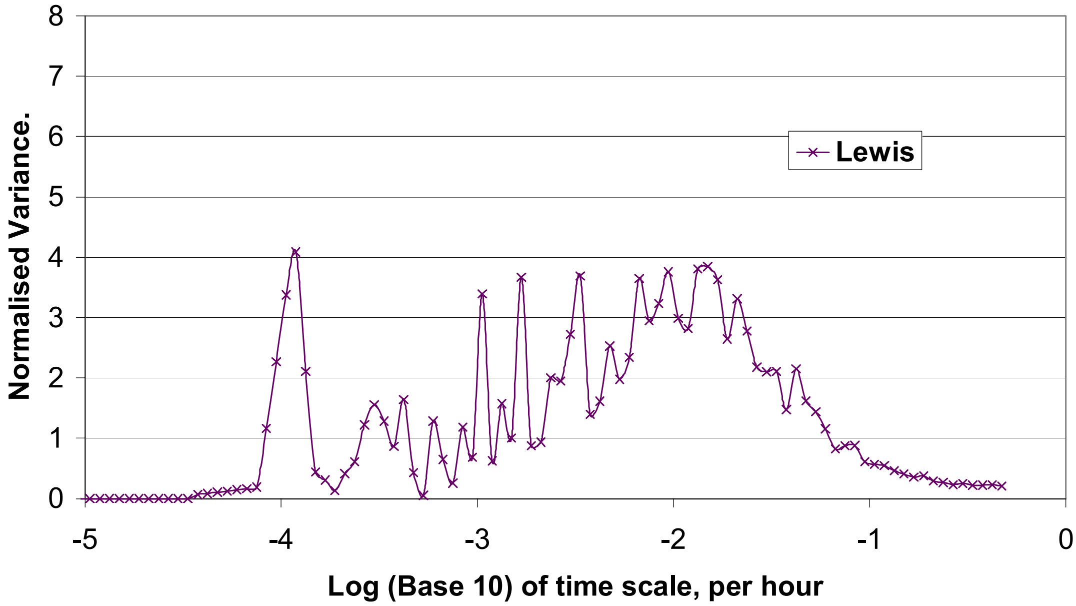

Figure 2.17 Wind speed power spectrum from the Butt of Lewis, normalised to 10m/s average wind speed

2.4.11.4 Scilly Isles

This data was measured at St. Mary's on the Isles of Scilly, off the south-western tip of the UK. Hourly average wind speed was measured for a period of 3 years.

The resulting spectrum has a relatively small diurnal spike, as expected in a maritime location. The seasonal spike of variation is large, fig. 2.18.

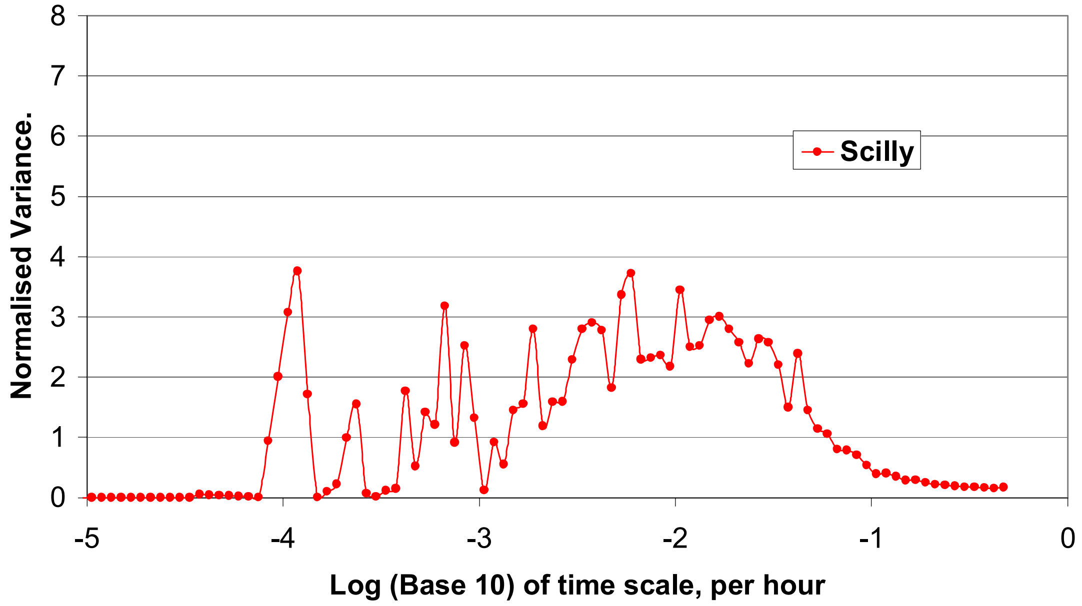

Figure 2.18 Wind speed power spectrum from the Isles of Scilly, normalised to 10m/s average wind speed

2.4.11.5 Ronaldsway

Ronaldsway is the location of an airport on the south-eastern side of the Isle of Man. Hourly average wind speeds were recorded there for a period of 3 years.

This site appears to have the largest total wind speed variance of all sites studied compared to its long-term average wind speed, fig. 2.19. It also has the largest seasonal spike of variation and a moderately large diurnal spike. It is a fairly maritime site, being in the middle of the Irish Sea, but the Isle of Man itself may be large enough and high enough to cause some local diurnal weather effects.

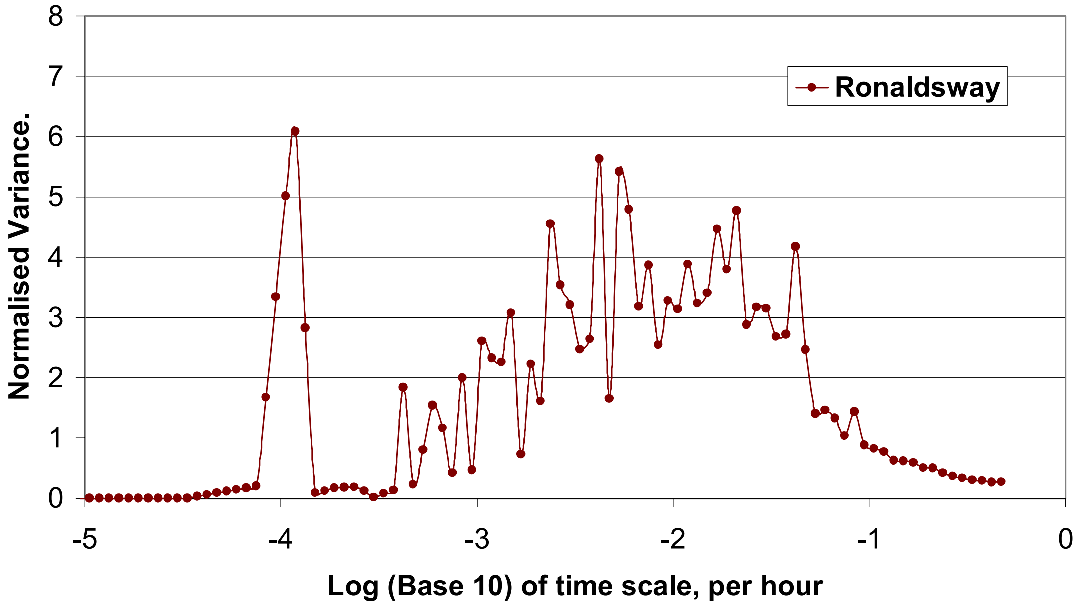

Figure 2.19 Wind speed power spectrum from Ronaldsway, normalised to 10m/s average wind speed

2.4.11.6 Turbulence Intensity

Where calculated, the turbulent portion of the spectra varies in size, dependent on the surface roughness of the surrounding terrain and the height of the anemometer. According to theory, turbulence intensity, *I* varies with height and surface roughness according to the approximate formula:

$$I = \frac{\sigma_u}{\overline{U}} = \frac{1}{\ln\left(z/z_0\right)} \qquad (2.2)$$

Where: σ_u=Standard deviation of axial wind speed

$\overline{U}$ =Axial wind speed

z=Height above ground of wind speed measurement

z_0=Surface roughness length

See (Hassan, Sykes 1990). So a higher anemometer or a smaller surface roughness produces a smaller turbulence intensity.

The probabilistic method presented in this thesis assumes just one turbulence intensity and turbulent spectrum; that calculated from the RAL data measured at 5 seconds interval. The effect of turbulence on wind power would be reduced if aggregated over several wind turbines, (Schlez 2000), especially if those turbines were located over an area wider than the integral length scale of the local wind turbulence, according to Taylor's Frozen Wake Hypothesis, (Townsend 1979).

2.4.11.7 Conclusions From the Study of Different Wind Speed Spectra

Maritime sites, especially on the western side of the British Isles, appear to have the largest seasonal spikes. Landlocked sites, especially sites on the eastern side of Britain, appear to have larger diurnal spikes but smaller seasonal spikes. All the available British data has been normalised by average wind speed and averaged to produce one generic spectrum, see fig. 2.20.

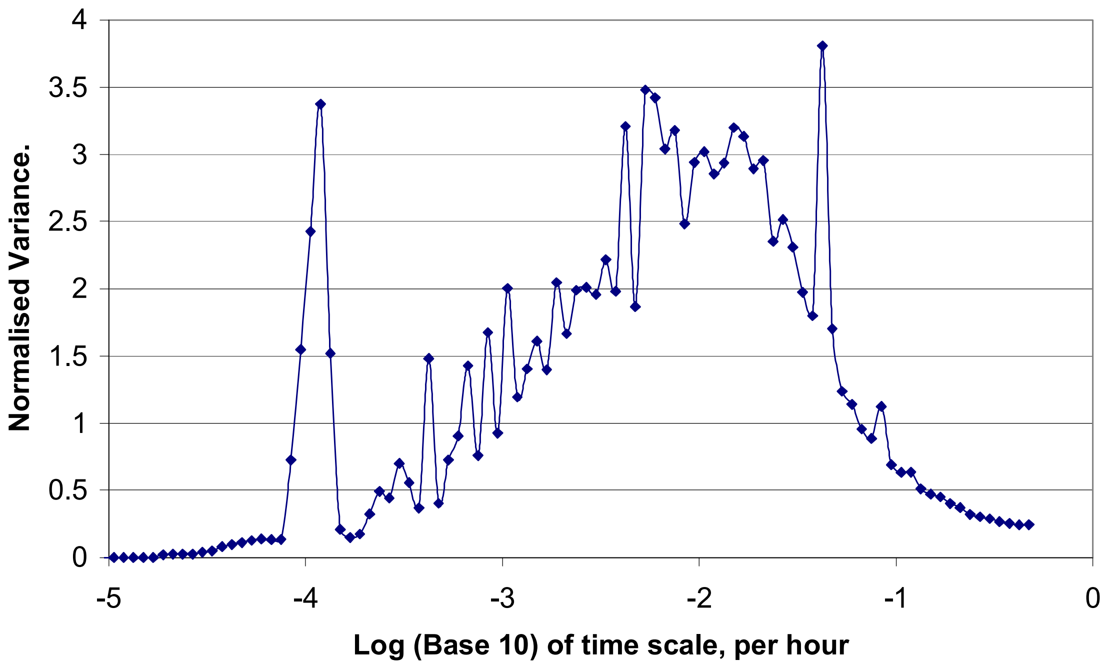

Figure 2.20 Average spectrum from many British sites, normalised to 10m/s average wind speed

This average spectrum is clearly much smoother than any one spectrum over the central portion, with sharp seasonal and diurnal spikes. This spectrum has also been extended to very short timescales by grafting on the turbulent spectrum as measured at RAL, fig. 2.21. This extended spectrum is the one used in all the probabilistic calculations shown in the results, chapter 3 unless otherwise stated.

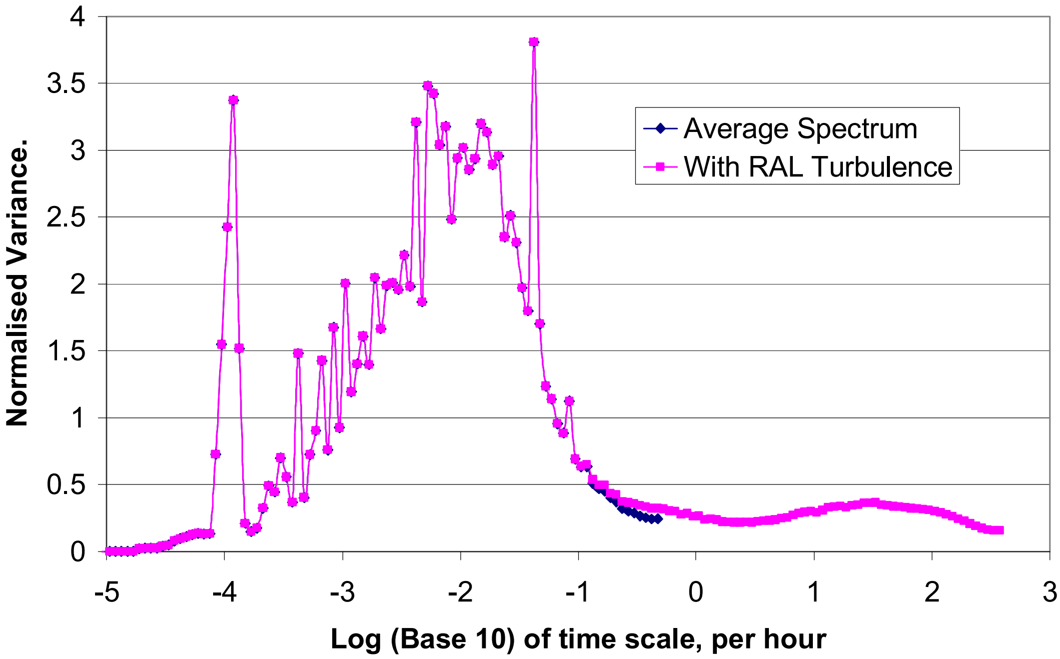

Figure 2.21 Generic wind spectrum used in most probabilistic calculations. The variance has been normalised to an average wind speed of 10m/s

Further work is desirable to investigate how wind speed power spectra change with location around the world, particularly in different climate types, for example equatorial, tropical, desert, central continental and polar climates.

2.4.12 Synthetic Wind Speed Data

A spectrum has also been created from synthetic wind speed data, using a random walk algorithm called ‘Pseudowind’ (Infield et al. 1994), fig. 2.22. The algorithm is sophisticated, employing a short-term generator for a Gaussian distribution of turbulence within a longer-term generator for a Weibull distribution of hourly average wind speeds. The spectrum in fig. 2.22 is actually created by merging 3 spectra from 3 different time series of synthetic wind speed data. The first is 200 years of hourly averages, the second is 10 years of minute averages and the third is 2 months of 1-second interval data. In this way, a smooth spectrum was constructed over all time scales.

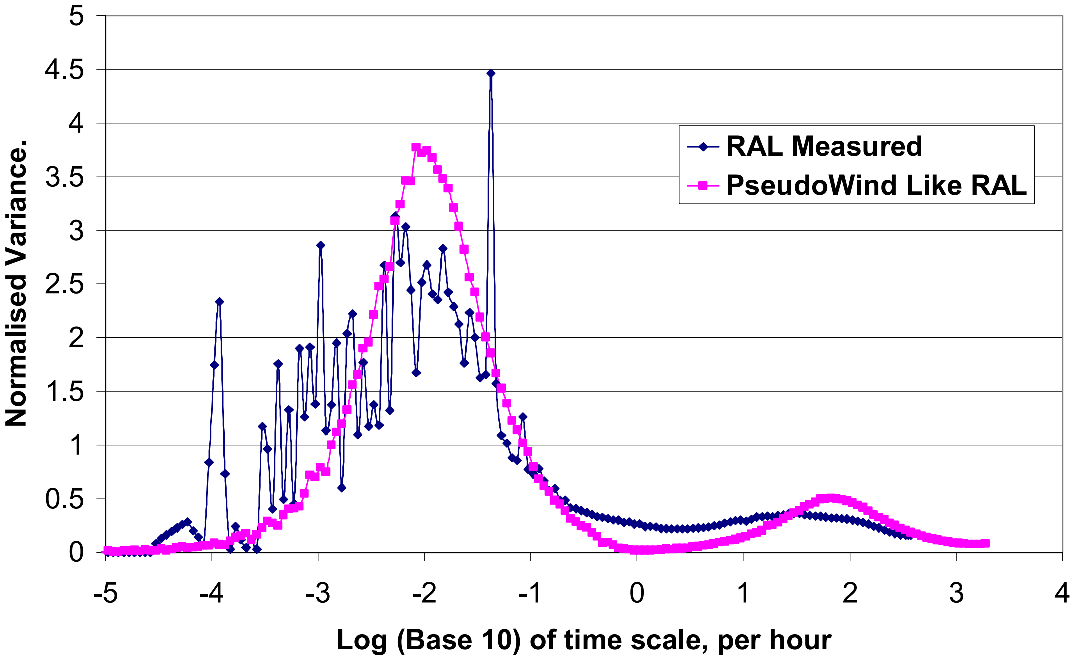

Figure 2.22 Wind speed spectrum from synthetic wind speed data compared with that measured at RAL, normalised to 10m/s average wind speed.

The spectrum created from synthetic wind speed data contains approximately the same total variance as the spectrum from measured data, but it is obviously a different shape. The synthetic wind spectrum has two clear peaks representing weather variation and turbulent variation, separated by a gap of zero amplitude. It has no peaks for diurnal or seasonal variations. In contrast, the spectrum from measured data has two broader, lower peaks of weather and turbulent variation that merge into one another. The measured data spectrum has significantly more variation at time scales of a month or more (at x values of about –3) and significant spikes of diurnal and seasonal variation.

These differences provide justification for using real weather data, or spectral information from it, rather than synthetic wind data. Even a time step simulation using this synthetic data will not capture all the characteristics of real wind speed data. A much more sophisticated synthetic wind speed generator would be required to do this.

2.5 Spectrum Integration Used in the Probabilistic Method

The probabilistic method applies filter functions to the spectra described in section 2.4, and then integrates the resulting spectra to calculate useful quantities. These are used to predict the behaviour of an energy store together with the necessary energy rating of that store. The filter functions have been used in several Matlab programs and have been presented in two conference papers: (Barton, Infield 2005a, Barton, Infield 2005b).

The probabilistic method has been applied to variations in wind power and solar power. If in a given application the stochastic components of variation in electricity demand were considered important, then the method could also be applied to variations in demand. For simplicity, the method is described below for variations in wind speed and wind power. The equations are written for just one frequency component of the wind speed variation spectrum. The description refers to periods of 24 hours, but would work equally well for other time periods, from seconds or minutes up to weeks or months.

2.5.1 The Time Period of the Store and Filter Functions

The probabilistic method assumes that the energy store can effectively cope with power variations within a certain length of time, T, but cannot cope with longer-term variations. The program user specifies this time period, T, for example one day. The method uses a low-pass filter function to calculate the variance associated with the probability distribution of average wind speeds over the time period T, for example daily average wind speeds, section 2.5.2. It uses a second, high-pass filter function to calculate the variance for the distribution of wind speeds within each time period T, for example the distribution of wind speeds expected within any 24-hour period, section 2.5.3. Finally, the method uses a third, state-of-charge filter function to calculate the required store energy capacity for this length of time, section 2.5.5. The filter functions are illustrated in fig. 2.23. The low and high pass filter functions are to scale, but the state-of-charge filter function is shown with arbitrary magnitude.

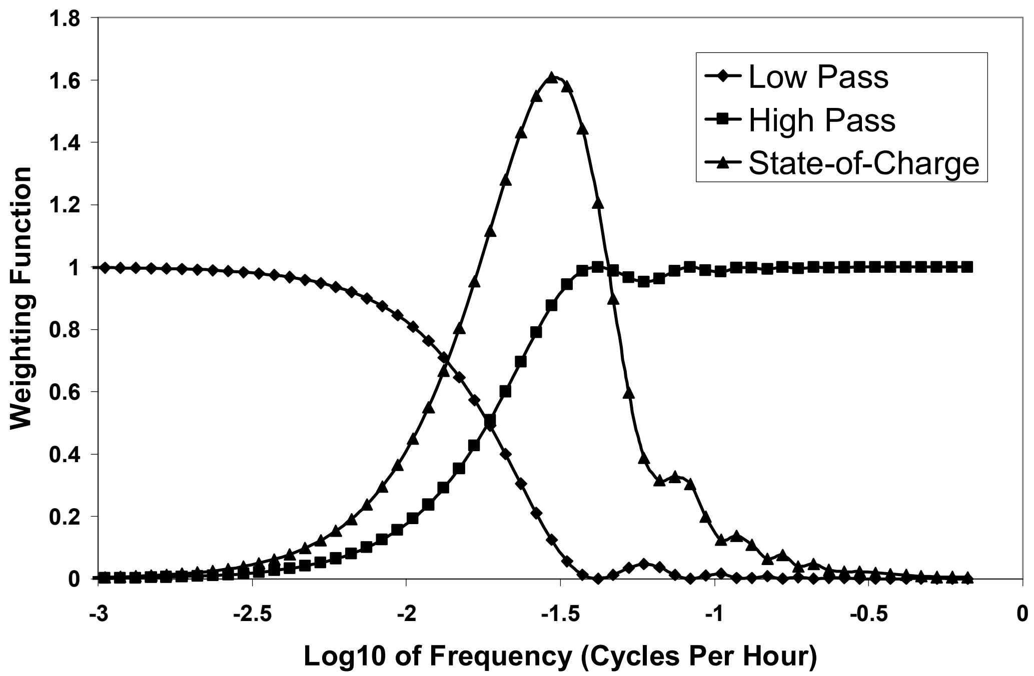

Figure 2.23 Filter functions for a store period of 24 hours

2.5.2 Distribution of Period-Average wind Speeds – The Low Pass Filter Function

For a given spectral frequency component, ω_i with amplitude, A_i and phase angle φ_i, the instantaneous component of wind speed variation is:

$U_i = A_i \sin(\omega_i t + \phi_i)$. Averaging this quantity over the store period, T gives the contribution from this frequency to the period-average wind speed:

$$\overline{U_i} = \frac{A_i}{\omega_i T}\left[\cos\phi_i - \cos(\omega_i T + \phi_i)\right] \tag{2.3}$$

If eq. 2.3 were integrated over all ϕ_i from 0 to 2π, the result would be the average value of each frequency component, which would sum to zero and be meaningless. Instead, eq. 2.3 is squared to give the contribution of each frequency component to the variance in period-average wind speed. When integrated over all φ_i from 0 to 2π to give the average contribution, the resulting integral is:

$$V_{1i} = \left(\frac{A_i}{\omega_i T}\right)^2 \left[1 - \cos(\omega_i T)\right] \tag{2.4}$$

Eq. 2.4 is a low-pass filter, the same as that used by (Infield 1990) for storage modelling, but without the empirical scaling factor of 2.4. The filter for a 24-hour store is shown fig. 2.23. Low frequency components (small ω) have a relatively large effect on the period average wind speed whereas high frequency components (large ω) have a small effect on the period average. Low frequency components remain nearly constant throughout a time period, T. High frequency components of wind variation complete many cycles during the time, T, and time spent above the mean is approximately balanced by time spent below.

2.5.3 Distribution of Wind Speeds Within Each Period – The High Pass Filter function

The variance of wind speed within a period, T is calculated in a similar way, but this time, the important quantity is the difference between the instantaneous value, U_i and the period-average, $\overline{U_i}$:

$$U_i - \overline{U_i} = A_i \sin\left(\omega_i t + \phi_i\right) - \frac{A_i}{\omega_i T}\left[\cos\phi_i - \cos\left(\omega_i T + \phi_i\right)\right] \qquad (2.5)$$

If eq. 2.5 is squared, integrated over time, T and integrated again over all φ_i, then the component of variance within period T results:

$$V_{2i} = \frac{A_i^2}{2} - \left(\frac{A_i}{\omega_i T}\right)^2 \left[1 - \cos\left(\omega_i T\right)\right] \qquad (2.6)$$

This integral, eq. 2.6 represents a high pass filter, previously derived by (Bossanyi, Anderson 1984) and also shown in fig. 2.23. It is actually the complement of the low pass filter function in the sense that summing the two time series resulting from application of the two filters to a given time series results in the original series. We can see that adding the within-period (high pass) variance, V_{2i}, eq. 2.6 to the period-average (low pass) variance, V_{1i}, eq. 2.5 gives $A_i^2/2$. This is the total variance in the wind speed contributed by one frequency component, and confirms that the equations are correct.

Low frequencies (small ω) have a small effect on the high pass (within-period) filter function whereas high frequencies (large ω) have a relatively large effect, fig. 2.23.

2.5.4 Application of the Low Pass and High Pass Filters

The low pass and high pass filters are applied to the wind speed variation spectrum to create two filtered spectra, one for variations of period (24-hour) average wind speed and another for variations within a period of 24 hours, see fig. 2.24.

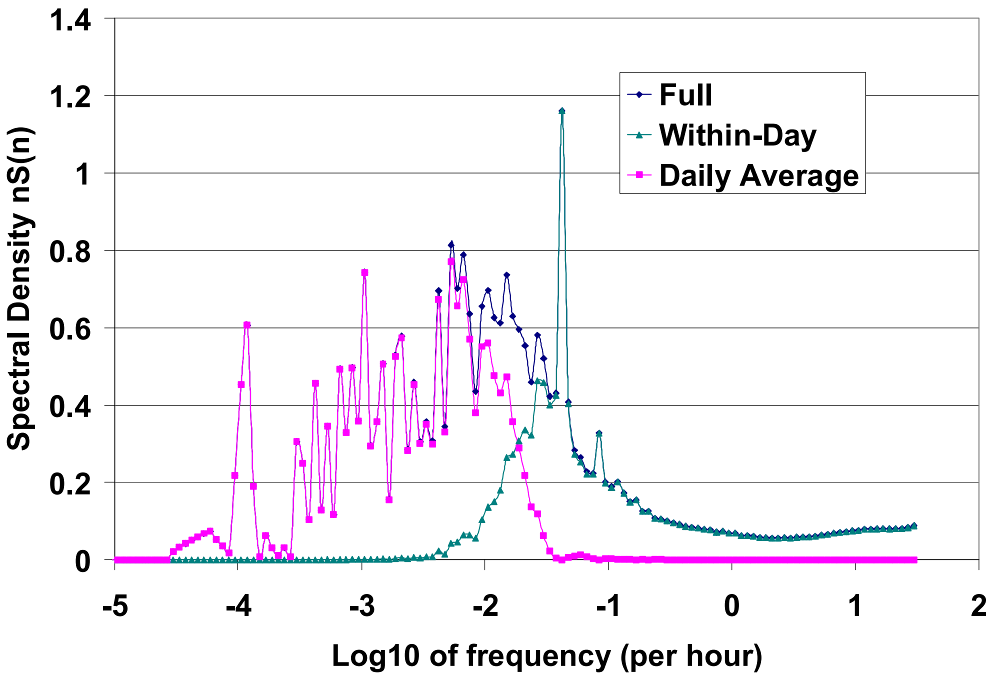

Figure 2.24 The 1994 to 1998 RAL Wind Variation Spectrum and Filtered Spectra. Average wind speed was 5.09 m/s.

The high pass filtered spectrum is integrated to calculate a within-period variance. This is equal to 6.10 $(m/s)^2$ using the RAL wind variation spectrum and a period of 24 hours, when scaled to a daily average wind speed of 8m/s. This corresponds to a standard deviation of 2.47 m/s. This information is used to construct a within-period wind speed probability density function (PDF), see fig. 2.25. The details of this calculation are given in sections 2.6.2 and 2.6.3.

Similarly, the low pass filtered spectrum is integrated to calculate a period average variance. This is equal to 9.36 $(m/s)^2$ for a long-term average wind speed of 8 m/s, the RAL wind variation spectrum, and a period of 24 hours, corresponding to a standard deviation of 3.06 m/s. This information is used to construct a period-average wind speed PDF, see fig. 2.25. Again, the details of this calculation are given in sections 2.6.2 and 2.6.3.

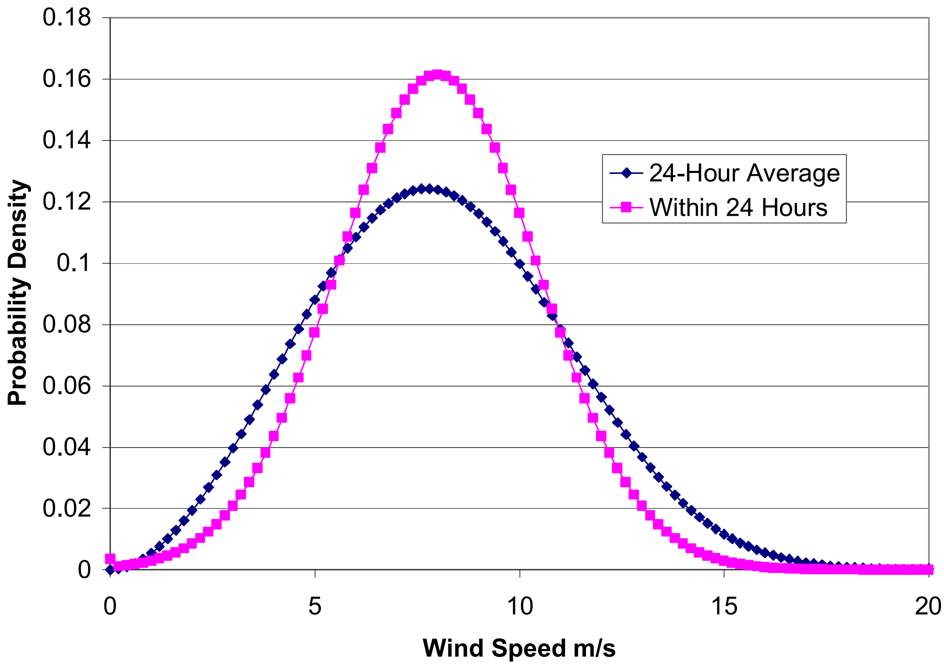

Figure 2.25 Wind Speed Distribution Functions for a 24-Hour Time Period, for an average wind speeds of 8m/s

2.5.5 Variance of State-Of-Charge of the Energy Store

The energy store only smoothes out variations within a period, T. Any longer term variations resulting in net energy imbalances over period T must be accommodated by other sources of generation, power cuts or power curtailment. The following method describes how the energy capacity of such a store is calculated using a third filter function.

This third filter function calculates the variance in state-of-charge of a store associated with a particular frequency component, ω_i. The net power being added to or subtracted from a store at time τ after the start of the period, T is equal to the difference between the instantaneous wind speed, U_i and the period average, $\overline{U_i}$, all multiplied by the effective gradient of the wind turbine power curve, *K*:

$$Net\_Power = K\left(U_i - \overline{U}\right) = K\left\{A_i \sin\left(\omega_i \tau + \phi_i\right) - \frac{A_i}{\omega_i T}\left[\cos\phi_i - \cos\left(\omega_i T + \phi_i\right)\right]\right\}$$

(2.7)

Integrating eq. 2.7 with respect to time, up to time *t* after the start of the period gives the net change in instantaneous state-of-charge, E_i of an energy store due to one frequency component:

$$E_i = \frac{A_i K}{\omega_i}\left\{\cos\phi_i - \cos\left(\omega_i t + \phi_i\right) - \frac{t}{T}\left[\cos\phi_i - \cos\left(\omega_i T + \phi_i\right)\right]\right\} \qquad (2.8)$$

It can be seen that this net energy is zero at time t=0 and at time t=T, but potentially non-zero at other times. The above net energy E_i is squared, integrated with respect to phase angle, ϕ_i, and integrated again with respect to time, t, to find the component of variance in state-of-charge during the time period, *T*. Without showing all the steps of the integration, the average variance in state-of-charge is:

$$V_{SOC} = \frac{1}{T}\frac{1}{2\pi}\int_{\phi=0}^{2\pi}\int_{t=0}^{T} E_i^2 \, dt \, d\phi$$

$$= \frac{1}{T}\int_{t=0}^{T}\frac{A_i^2 K^2}{\omega_i^2}\left\{1 + \frac{t^2}{T^2}\left[1-\cos\left(\omega_i T\right)\right] - \cos\left(\omega_i t\right) - \frac{t}{T}\left[1 - \cos\left(\omega_i T\right) - \cos\left(\omega_i t\right) + \cos\left(\omega_i t - \omega_i T\right)\right]\right\} dt$$

$$= \frac{A_i^2 K^2}{\omega_i^2}\left\{\frac{5}{6} + \frac{1}{6}\cos\left(\omega_i T\right) + \frac{2}{\omega_i^2 T^2}\left[\cos\left(\omega_i T\right) - 1\right]\right\} \qquad (2.9)$$

Eq. 2.9 forms a third weighting function, fig. 2.23, to be applied to the wind variation spectrum, fig. 2.21. Integration of the resulting filtered spectrum gives the variance of net energy. Low frequency components have little effect on the store, since their magnitude varies little during period, T. High frequency components also have little effect on the store, since they complete many cycles during period, T, so each cycle accumulates and discharges very little energy. Only frequency components close to the period of the store have a significant effect on the state-of-charge.

The function tends to zero for very small values of $\omega_f T$, but also tends to zero for very large values of $\omega_f T$. It has its global maximum value when $\omega_f T$=4.58. This corresponds to a frequency component, ω_f , that completes almost ¾ of a cycle during the store time period, T, fig. 2.23.

2.5.6 Calculation of the Turbine Power Curve Gradient, K

As seen above, the variance of state-of-charge requires a wind turbine power curve gradient, K. This is not the average over the whole range of the power curve, but only that around the period-average wind speed. The wind turbine power curve, fig. 2.4 can be combined with the within-period wind speed PDF, e.g. fig. 2.25 to create a within-period PDF of wind power, e.g. fig. 2.26.

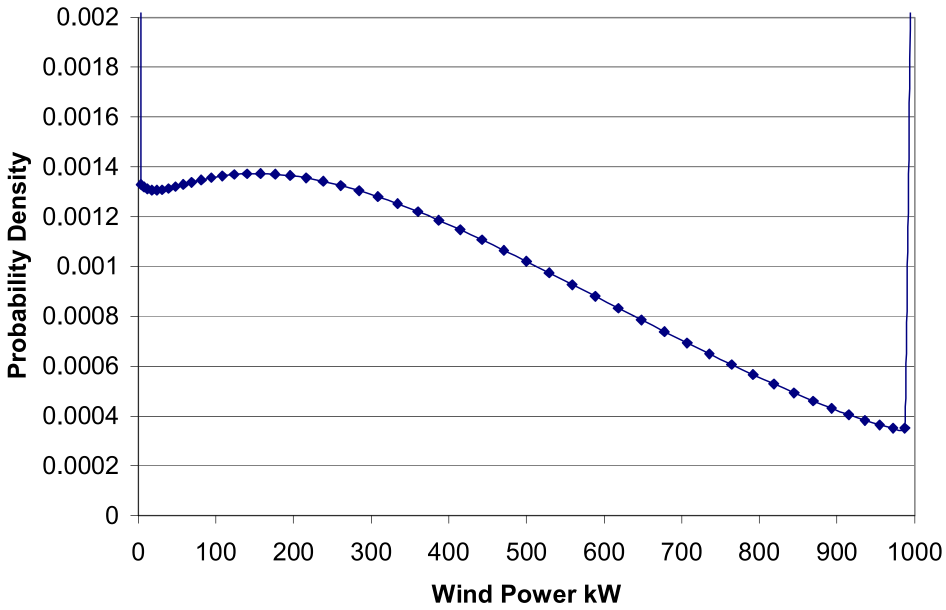

Figure 2.26 Probability Density Function of Wind Power from a 1MW turbine for a Daily Mean Wind Speed of 8m/s

In practice, all that is required is the within-period standard deviation of wind power, σ_P and the within-period standard deviation of wind speed, σ_U . The ratio of these is the effective gradient, $K = \sigma_P / \sigma_U$. (2.10)

Let the PDF of wind speed within one day be $f(U)$, such that $\int_0^\infty f(U)\,dU = 1$

Mean wind speed within a day $\overline{U}_{DAY} = \int_0^\infty U \times f(U)\,dU$ (2.11)

And variance of wind speed within a day $\sigma_U^2 = \int_0^\infty U^2 \times f(U)\,dU - \overline{U}_{DAY}^{\,2}$

Then the wind turbine power output is $P(U)$, as given by the turbine power curve.

Mean power output over one day $\overline{P}_{DAY} = \int_0^\infty P(U) \times f(U)\,dU$ (2.12)

And variance of wind speed within a day $\sigma_P^2 = \int_0^\infty \left[P(U)\right]^2 \times f(U)\,dU - \overline{P}_{DAY}^{\,2}$

(2.13)

Section 2.11.3 describes the wind conditions for which these PDFs of wind speed and wind power must be calculated.

In the case of solar power, K is equal to the photovoltaic module power capacity under standard test conditions, modified by efficiency factors e.g. inverter efficiency and reflectance effects. The probabilistic model neglects all tilt angle effects, PV module temperature effects and variable inverter effects, section 2.7.15.

For a given time period, the energy capacity of the store is calculated using the above equations and the method described in section 2.11.

2.6 The Wind Power Model

There are four aspects of the wind power model used in the probabilistic method and described in the following sections: the seasonal and diurnal variations in wind speed, section 2.6.1; constructing curves of probability density functions (PDFs) of wind speed, 2.6.2 and 2.6.3; conversion of wind speeds into wind power using a turbine power curve, section 2.6.4 and the magnitude of short term variations in wind speed, section 2.6.5.

2.6.1 Seasonal and Diurnal Variations in Wind speed

In a Northern European wind climate, average wind speeds are higher in winter than in summer. The percentage difference in average wind speeds may be small, but the percentage difference in wind power is larger because the available wind power per perpendicular area (wind turbine swept area) is proportional to the cube of wind speed. Data from the measurement site at Rutherford Appleton Laboratory (RAL) shows that January wind speeds are about 60% higher than July wind speeds, fig. 2.27 and fig. 2.28.

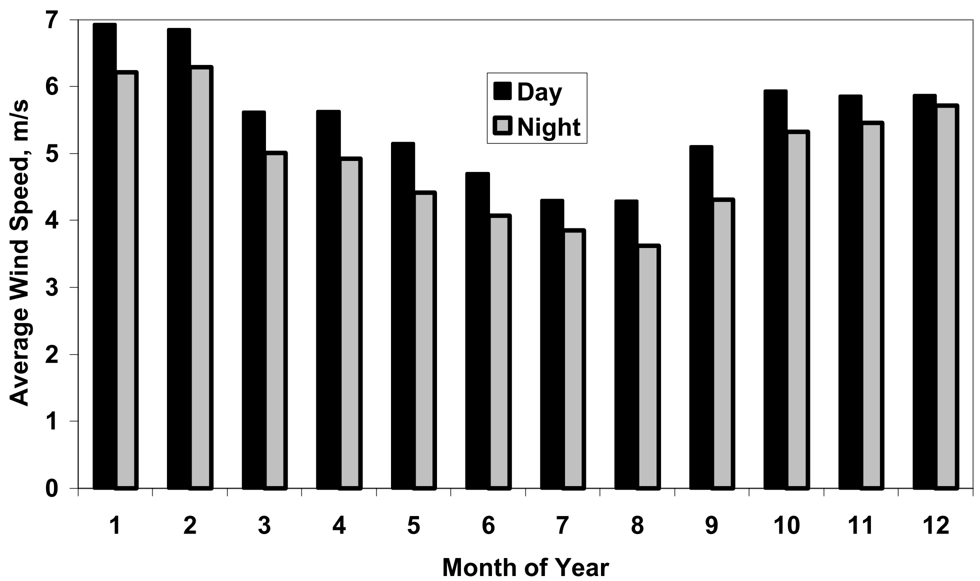

Figure 2.27 Monthly Average Wind Speeds at RAL

At RAL, the power available in the wind during the winter months (December, January and February) is twice that available during the summer months

(June, July and August). The wind speed profile lags the seasonal variation in solar irradiance, so that the highest wind speeds occur in January and February but the lowest wind speeds occur in July and August.

Wind speeds are also higher in the daytime than at night, fig. 2.27 and fig. 2.28. The percentage difference is larger in the spring and summer than in the autumn and winter.

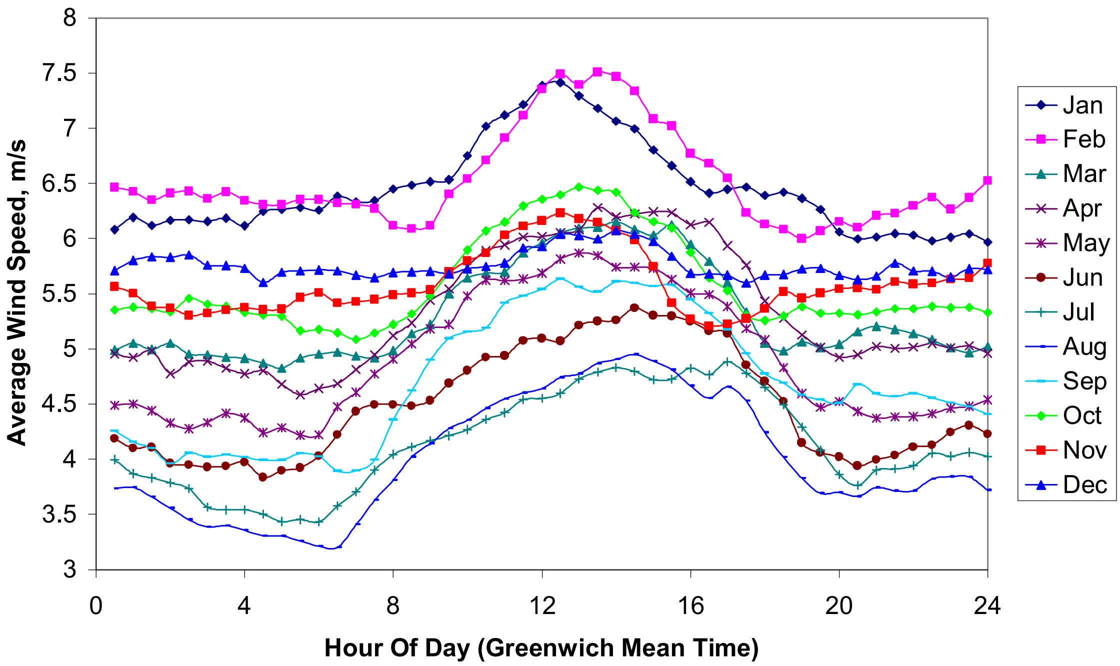

Figure 2.28 Average Daily Wind Speed Profiles for Each Month at RAL

Again, the wind speed profile lags the solar irradiance profile. The maximum average wind speeds occur in the early afternoon and the minimum wind speeds usually occur in the early hours of the morning.

These features make wind powered electricity generation particularly suitable for Northern Europe, where winter lighting and heating loads are higher than in summer, and where daytime and evening electricity loads are higher than night time loads. The weather patterns at RAL are likely to be quite typical of sites in Northern Europe.

The seasonal and diurnal dependence of wind power, together with the seasonal and diurnal dependence of solar power and electrical load, made it essential for the probabilistic model to include seasonal and diurnal variations despite the extra complexity involved, see section 2.10. The full probabilistic method uses different average wind speeds for day and night and for each month of the year, making 2 x 12 = 24 different average wind speeds. The model could have split the year into just 4 seasons, but such a coarse model may not adequately capture the subtleties of variation, for example the way the wind speed variation lags the solar irradiance seasonally, section 2.4.9. The difference between day and night was judged to be small enough that one average wind speed for daylight hours and another for night time was sufficient. Future work to re-evaluate these decisions may be useful.

2.6.2 Wind Speed Distributions Measured at the Rutherford Appleton Laboratory

2.6.2.1 The Long-Term Distribution

The wind speed data measured at Rutherford Appleton Laboratory (RAL) is recorded as one-minute averages, so the RAL data covers a greater range of the wind variation spectrum than most wind speed data and includes some of the turbulent variation, fig. 2.6. The standard deviation of the RAL data should therefore be slightly larger than for ten-minute or one-hour averages. Two RAL wind speed distributions, fig. 2.29 and fig. 2.30 were each measured over periods of 4 years. The shapes of the distributions are quite typical of long-term wind speed distributions in the UK. As stated in section 2.2.1, the 1994 to 1998 data has been used to create one power spectrum of wind speed variations and for one time step validation while the 1999 to 2002 data has been used as one of the inputs to a generic wind speed spectrum and in most of the time-step validations.

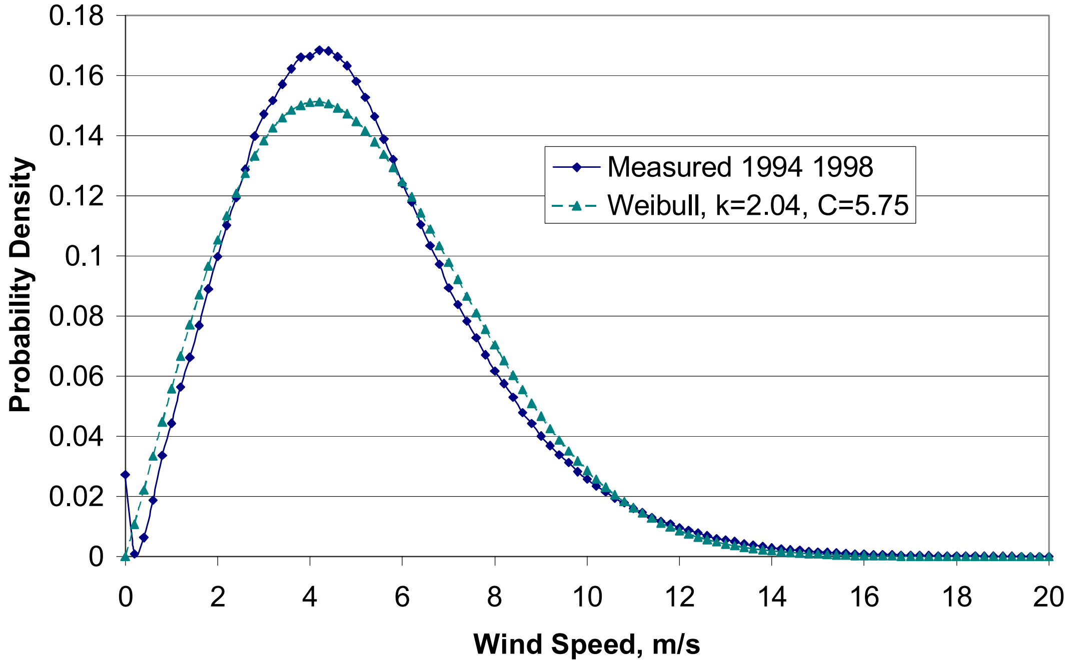

Figure 2.29 PDF of One-Minute Wind Speed Data Measured at Rutherford Appleton Laboratory between 1994 and 1998

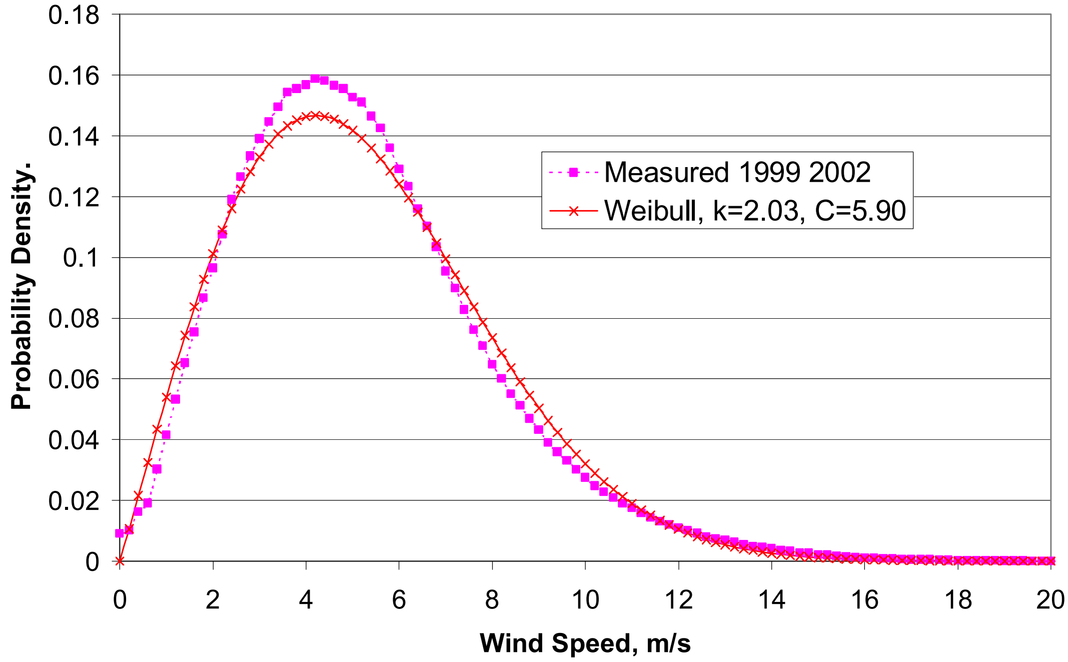

Figure 2.30 PDF of One-Minute Wind Speed Data Measured at Rutherford Appleton Laboratory between 1999 and 2002

Weibull distributions have been added to compare with the measured PDFs. The Weibull distribution with a shape factor of 2.04 and a scale factor of 5.75

has the same mean and standard deviation as the 1994 to 1998 data, fig. 2.29, while the Weibull distribution with a shape factor of 2.03 and a scale factor of 5.90 has the same mean and standard deviation as the 1999 to 2002 data, fig. 2.30. The Weibull shape and scale factors were determined by trial and error to match the mean and standard deviation, while the values calculated using the approximations of section 2.6.3.1 below produced very similar results. The Weibull distributions obviously do not quite fit the measured PDFs exactly. The measured PDFs have higher kurtosis. They have sharper peaks and longer, fatter tails at higher wind speeds than the fitted Weibull curves.

The relative merits of various statistical distributions, including Weibull distributions, used to model wind speed are discussed in section 2.6.3 and an alternative possible statistical distribution is discussed in appendix C.

2.6.2.2 Period-Average Wind Speeds

If the sample-averaging period gets longer, the calculated long-term mean wind speed remains unchanged but the standard deviation of the distribution is reduced because the variations represent a smaller range of the wind variation spectrum, fig. 2.21. The following plots, figs. 2.31 to 2.35 show how the wind speed distributions change with sampling period. Plots are based on the 1994 to 1998 data set.

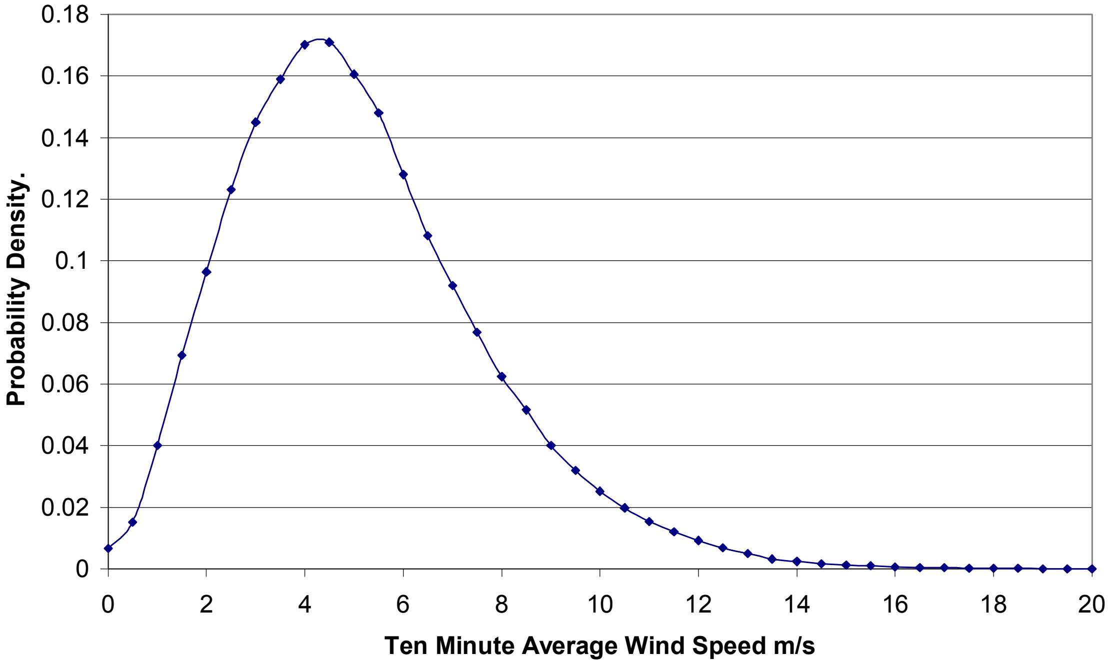

Figure 2.31 PDF of 10-Minute Average Winds Speeds Measured at RAL

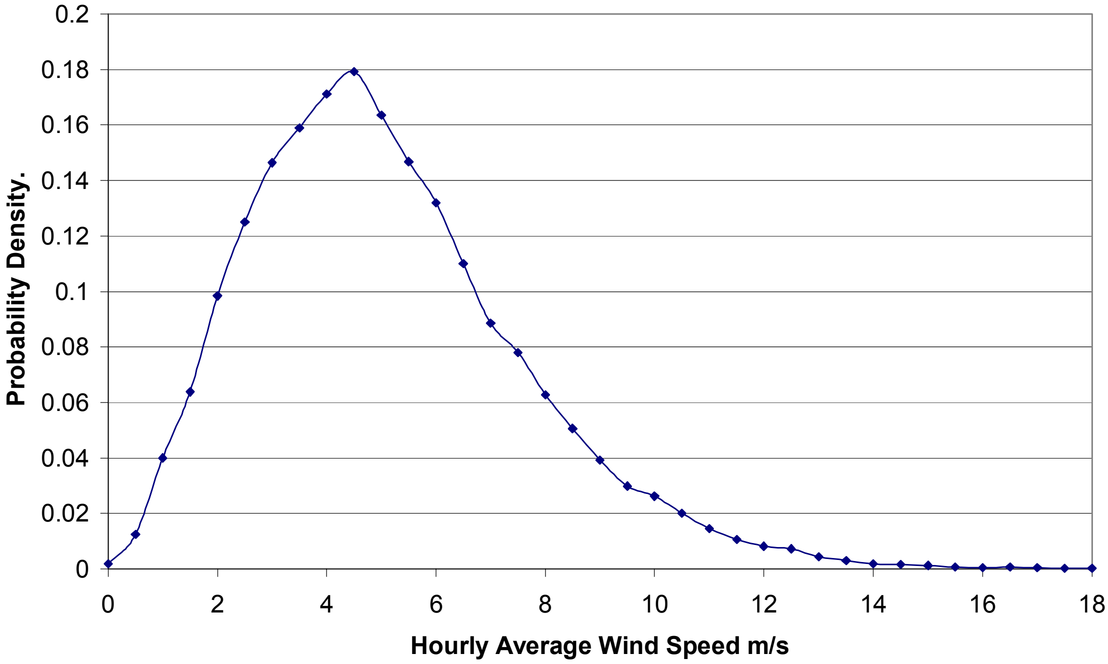

Figure 2.32 PDF of Hourly Average Winds Speeds Measured at RAL

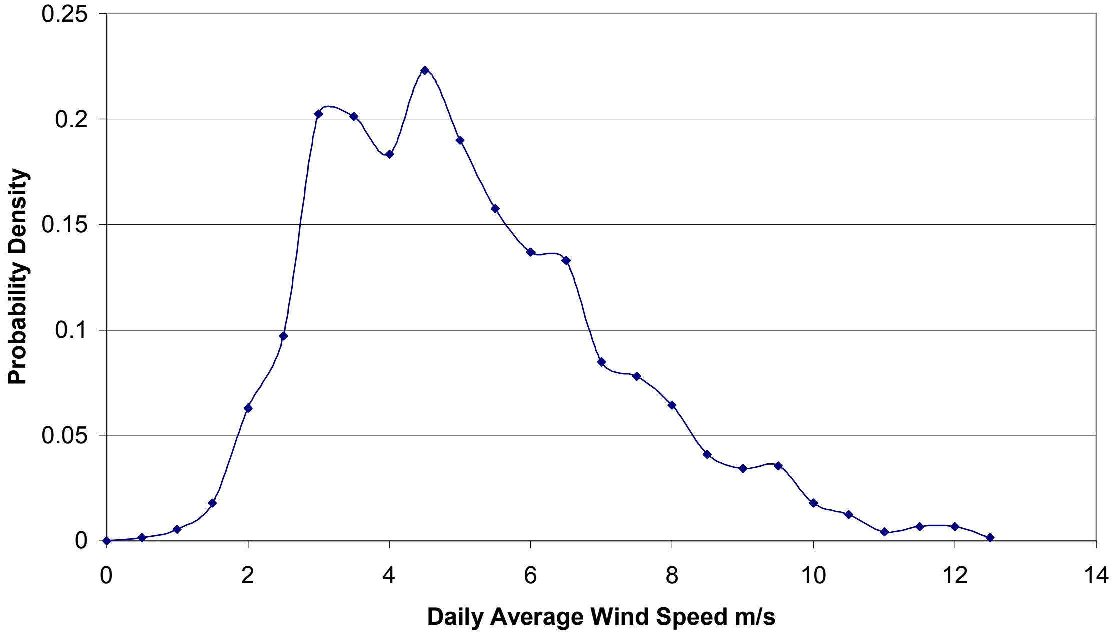

Figure 2.33 PDF of Daily Average Winds Speeds Measured at RAL

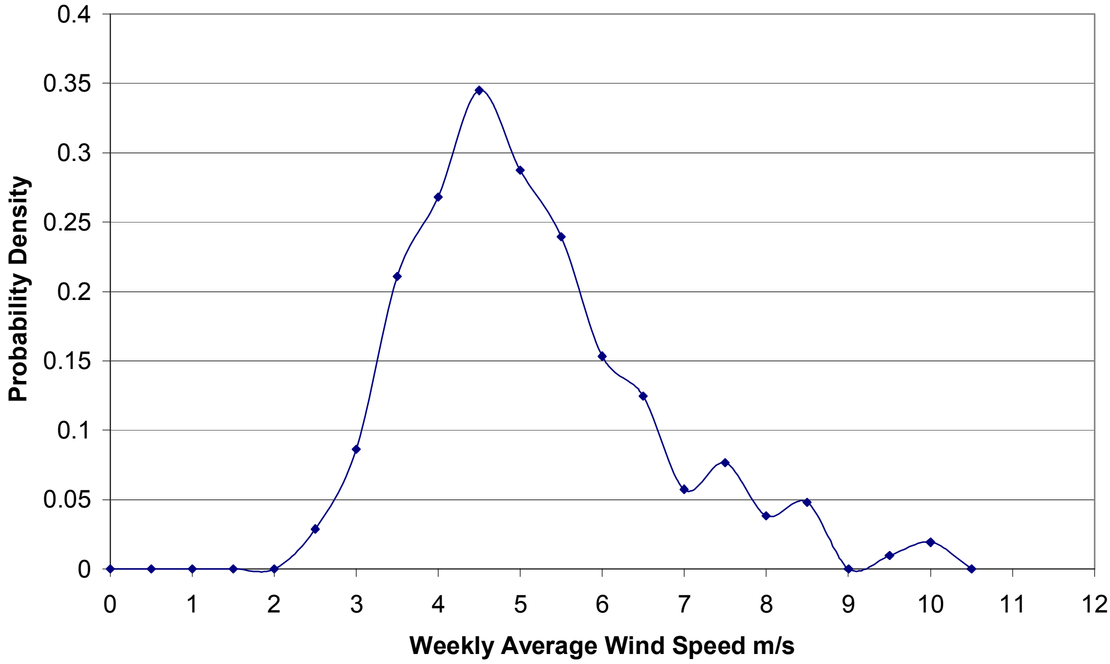

Figure 2.34 PDF of weekly Average Winds Speeds Measured at RAL

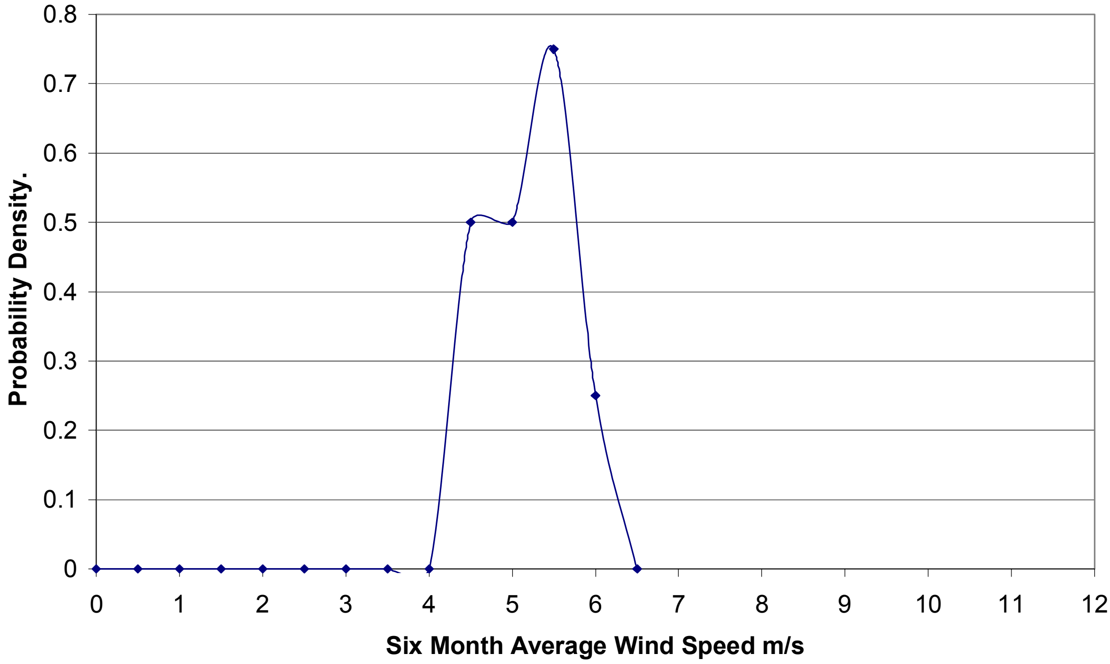

Figure 2.35 PDF of Six-Monthly Average Winds Speeds Measured at RAL

As the averaging period gets longer, the distribution gets narrower and (arguably) more symmetrical like a normal distribution, although the reduced sample size makes the distributions more noisy and rough, figs. 2.33 to 2.35.

2.6.2.3 Variations Within Each Period

As the averaging period gets longer (with a fixed one minute sampling rate), the variation in wind speed within each period gets larger and the distributions move from being approximately Gaussian towards a Weibull shape, as seen in the following plots, figs. 2.36 to 2.40. Let us consider averaging periods where the period-average wind speed was 2.5m/s, 5m/s or 7.5m/s:

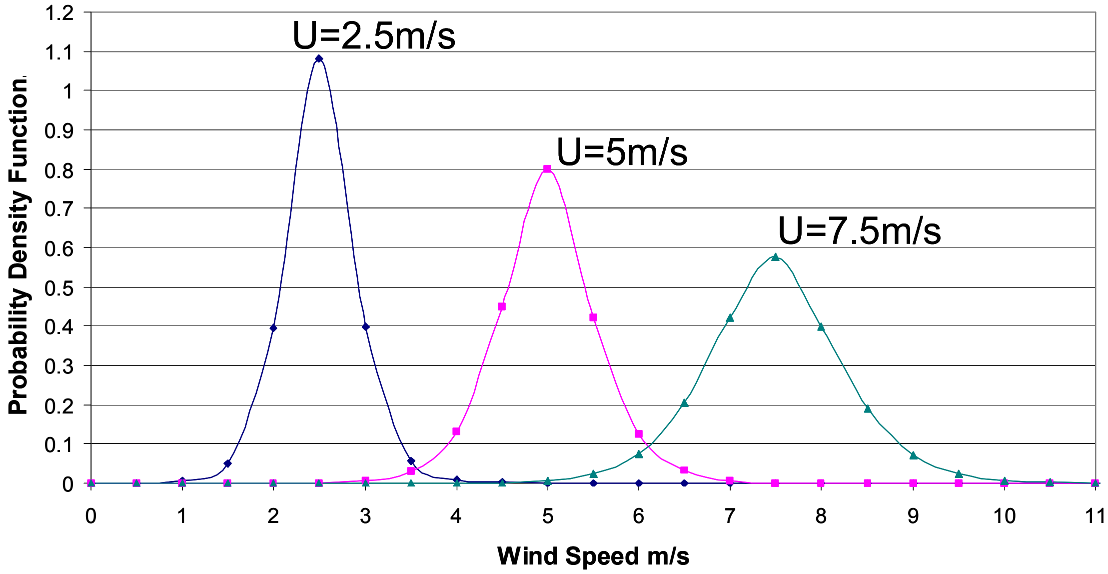

Figure 2.36 PDFs of wind speed within ten-minute periods measured at RAL

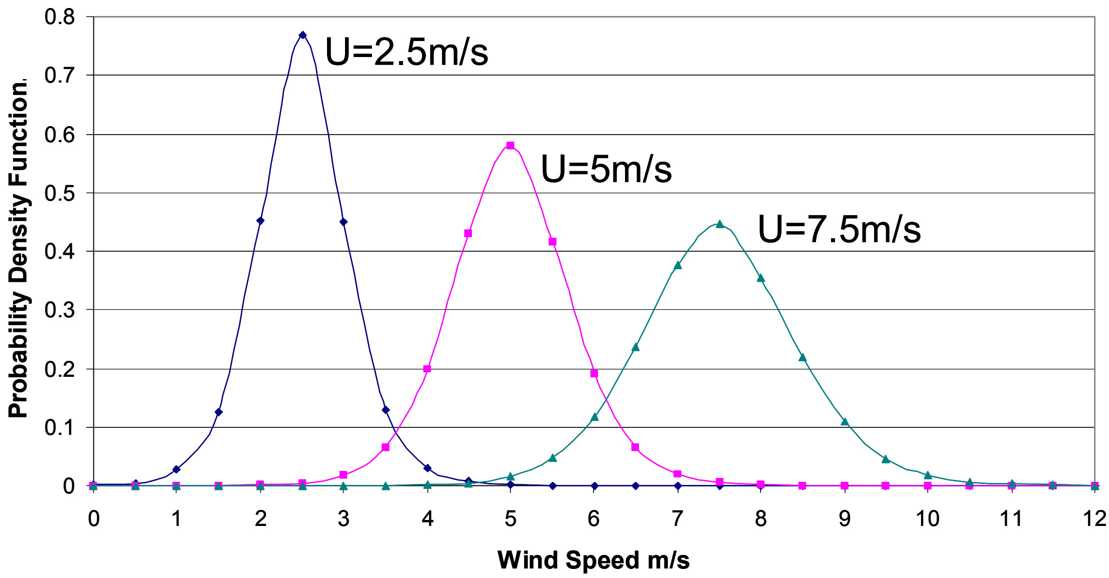

Figure 2.37 PDFs of wind speed within one-hour periods measured at RAL

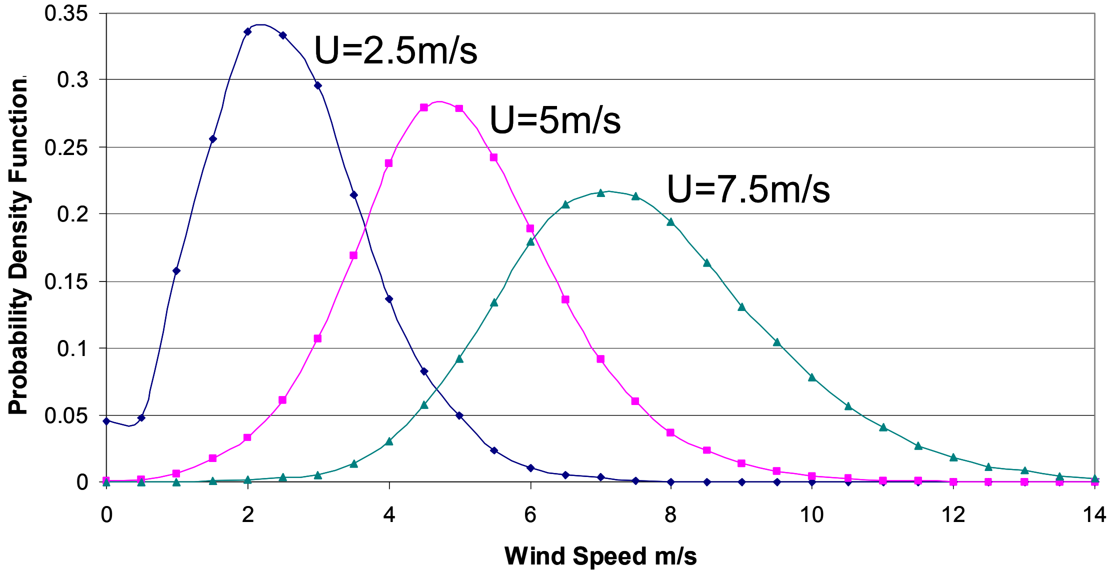

Figure 2.38 PDFs of wind speed within one-day periods measured at RAL

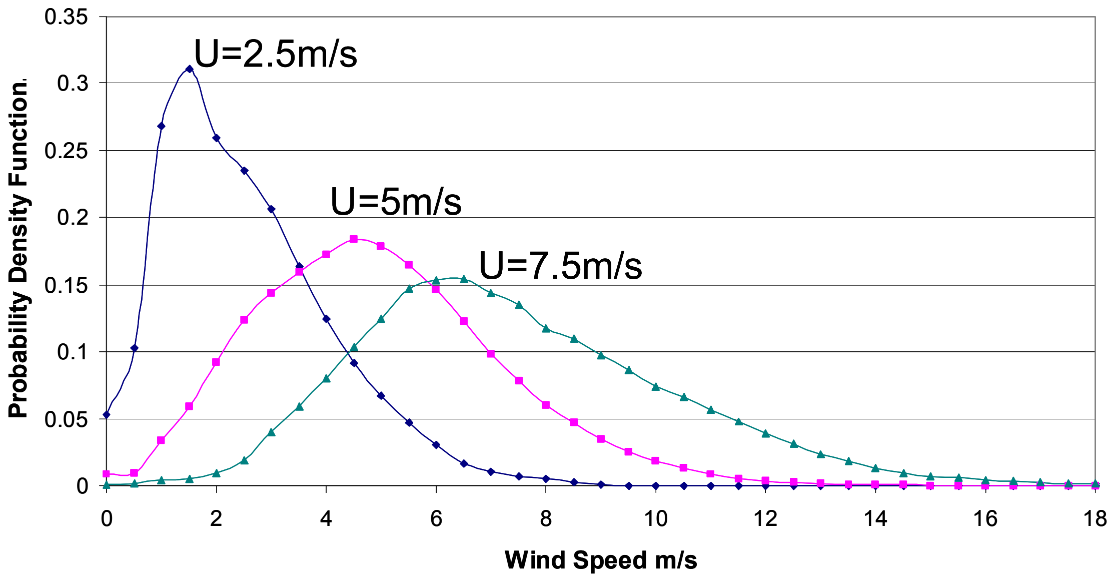

Figure 2.39 PDFs of wind speed within one-week periods measured at RAL

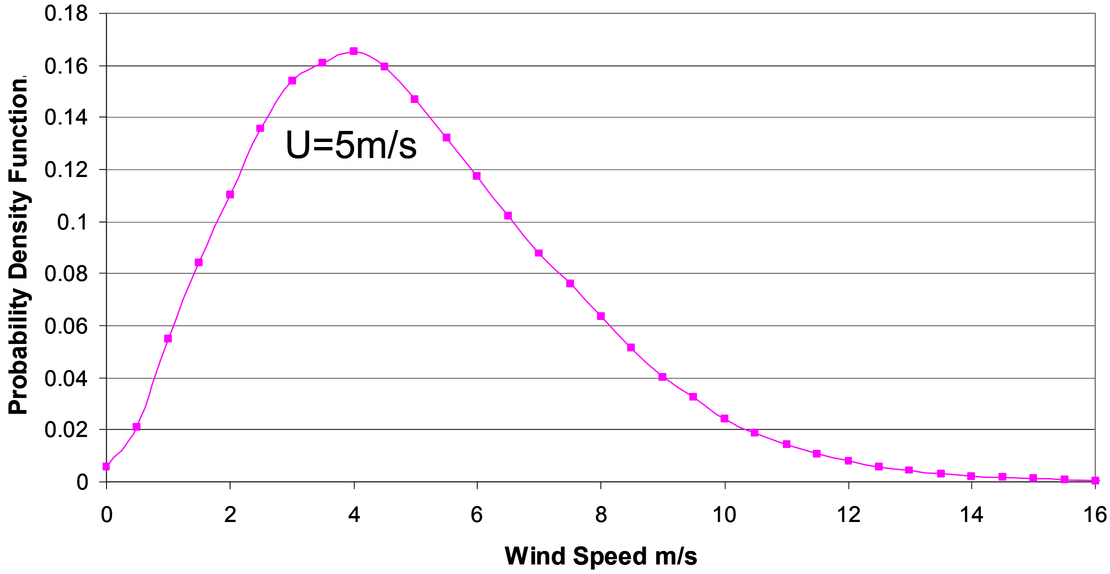

Figure 2.40 PDFs of wind speed within six-month periods measured at RAL

The plots show that the distributions of wind speed within a time period get wider as that time period increases. Therefore the standard deviation of wind speed within a period increases. Short time periods, e.g. one minute have narrow, almost symmetrical distributions, fig. 2.36. Periods of one hour, fig. 2.37, one day, fig. 2.38, and one week, fig. 2.39 have progressively wider PDFs. A period of six months or more has a PDF that is very similar to the long-term wind speed PDF, fig. 2.40. Note that only one curve is shown for wind speeds within six months, one for which the average wind speed was 5m/s. This is because all the six-month averages are tightly grouped around 5m/s; none were as low as 2.5m/s or as high as 7.5m/s.

As the time period increases, the shapes of the PDFs also change. As the time period increases, the shapes change from symmetrical bell-shaped distributions, fig. 2.36 to skewed distributions with a long tail of higher wind speeds, fig. 2.40. The next section, 2.6.3 discusses which statistical distributions fit the data best.

The above figs. 2.36 to 2.39 also show to a reasonable approximation that the standard deviation of wind speed within each period is proportional to the period-average wind speed. This is an assumption made in the probabilistic method, but the relationship between period-average wind speed and within-period standard deviation will be discussed further in section 2.6.5.

2.6.3 Statistical Distributions for Modelling Wind Speed Probability Density Functions

Two primary inputs to the probabilistic method are a wind speed power spectrum, see section 2.4.11.7, fig. 2.21 and a solar irradiance power spectrum, see section 2.4.6, fig. 2.10 and fig. 2.11. Section 2.5 shows how filter functions are applied to these power spectra. The filtered spectra are then integrated to give variances of wind speed and solar irradiance, both of the distributions within a store cycle period, T, and the distribution of the period-average values.

The original wind speed distributions have been discarded as far as the probabilistic method is concerned. But in order to model probability distributions of wind power, approximations to those original distributions must be obtained. The probabilistic method therefore constructs wind speed distributions, using assumed generic statistical distribution functions, from knowledge of only the mean wind speed and the standard deviation (the square root of variance). This is done for both wind speed variations within a store operating period, and period-average wind speeds.

The relative merits of the different statistical distributions are discussed below, concentrating on those typically used for wind speed modelling.

2.6.3.1 Weibull Distributions

The most widely used distribution used to fit long-term distributions of hourly or 10-minute averaged wind speed data is the Weibull distribution, usually with a shape factor between 1.6 and 2.4, see (Celik 2002b, Celik, Marshall 1998, Chadee, Sharma 2001, Garcia et al. 1998, Infield et al. 1994, Stevens, Smulders 1979, Tuller, Brett 1984, Tuller, Brett 1985, Wan-Kai Pang, Forster & Troutt 2001).

(Garcia et al. 1998) found that Weibull distributions fitted hourly-average wind speed data better than log-normal distributions did. (Tuller, Brett 1985) found that Weibull distributions fitted wind speed data better than Rayleigh distributions. This is not surprising since a Rayleigh distribution is a special case of a Weibull distribution with the shape factor restricted to 2.0.

A Weibull distribution can be constructed to fit a wind speed distribution, knowing its mean, $\overline{U}$ and standard deviation, σ (Hassan, Sykes 1990). The Weibull shape factor, k is given approximately by:

$$k = \left(\frac{\sigma}{\overline{U}}\right)^{-1.086} \tag{2.14}$$

And the Weibull scale factor, C is given by:

$$C = \frac{\overline{U}}{\Gamma\left(1 + 1/k\right)} \tag{2.15}$$

Where Γ is the gamma function.

Then the Weibull PDF is $P(U) = \left(\frac{k}{C}\right)\left(\frac{U}{C}\right)^{k-1} \exp\left[-\left(\frac{U}{C}\right)^{k}\right]$ (2.16)

When the standard deviation of a wind speed distribution is large compared to the mean ($\sigma/\overline{U}$ greater than about 0.36) then the corresponding Weibull distribution has a shape factor of less than 3. Such a Weibull distribution has a strong positive skewness. It has a long tail to the right, and is usually a good fit for a long-term distribution of wind speeds, e.g. fig. 2.30. Weibull distributions fit well to long-term PDFs of wind speed averaged over any relatively short period, e.g. one minute, one hour, fig. 2.41, or even one day, fig. 2.42.

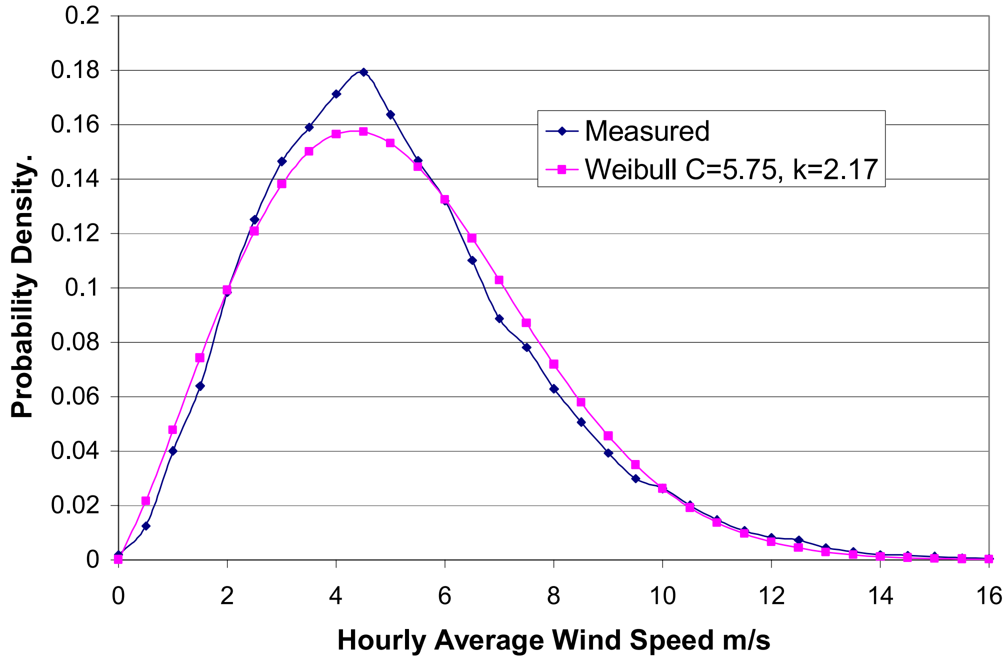

Figure 2.41 PDF of Hourly Average Winds Speeds Measured at RAL Compared with a Weibull Distribution

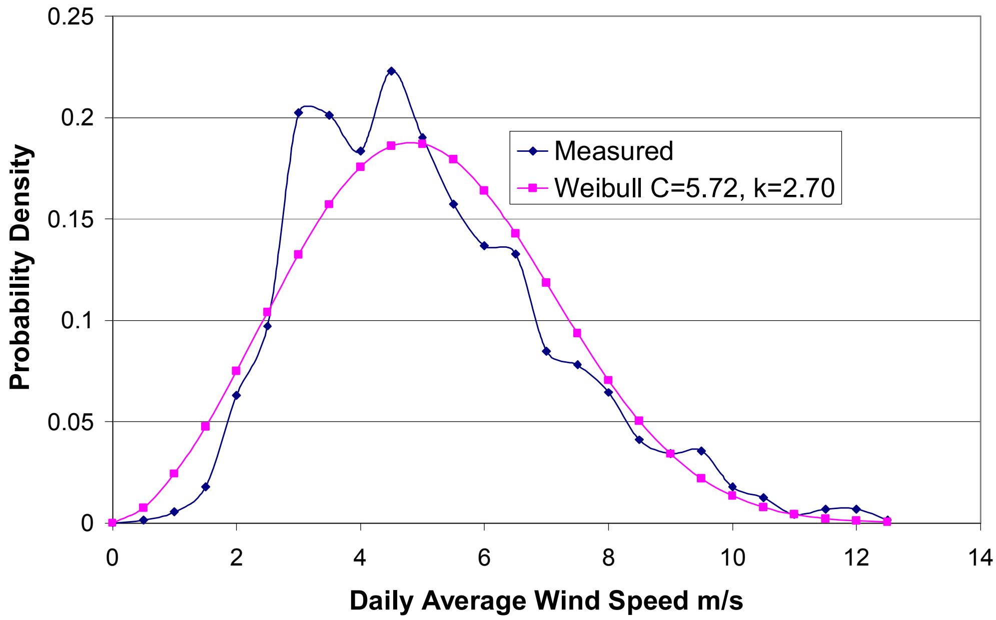

Figure 2.42 PDF of Daily Average Winds Speeds Measured at RAL Compared with a Weibull Distribution

Weibull distributions also fit well to wind speed distributions within a period of time, provided that period of time is long enough, for example six months, fig. 2.43, down to one week, fig. 2.44.

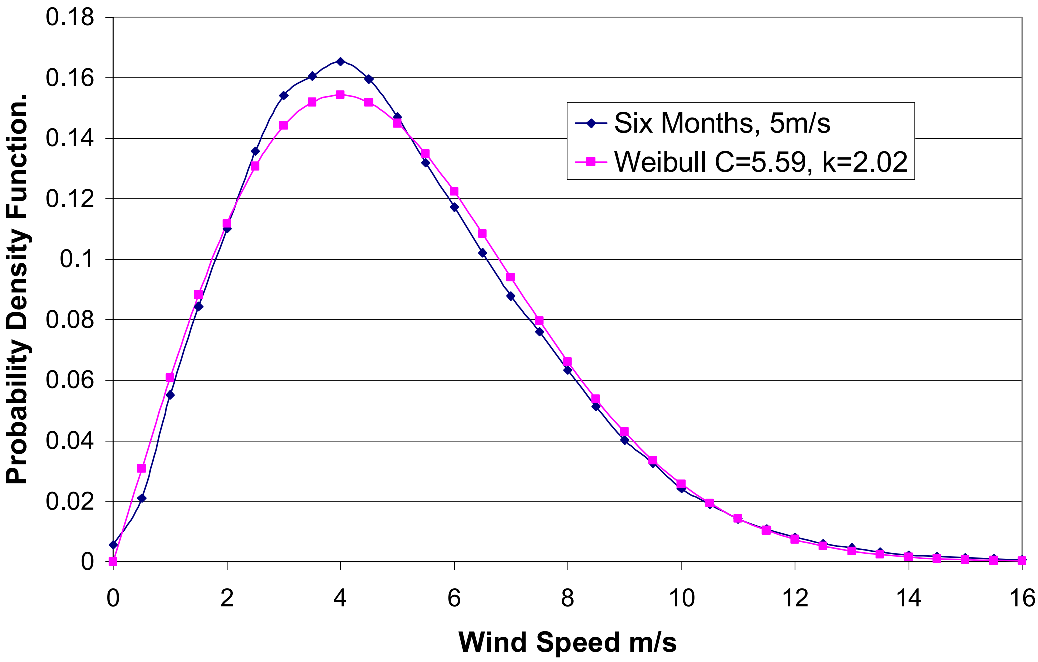

Figure 2.43 PDF of Winds Speeds at RAL Within Six Months Compared with a Weibull Distribution

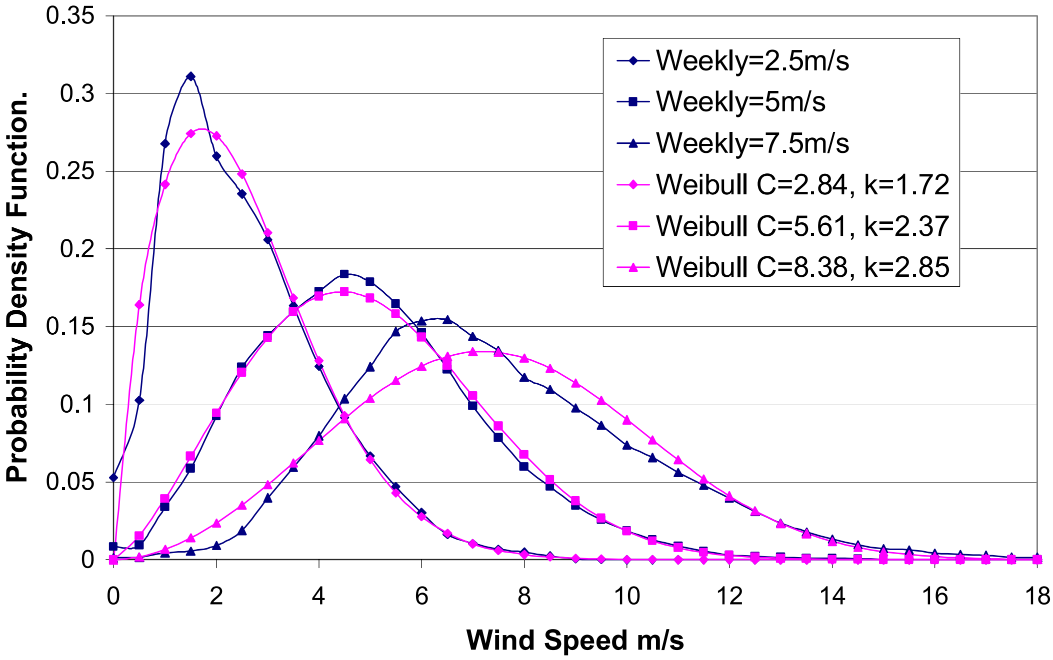

Figure 2.44 PDF of Winds Speeds at RAL Within a Week Compared with Weibull Distributions of the same mean and standard deviation

2.6.3.2 Normal Distributions

When fitting a distribution to wind speed turbulence data (the variations within each 10-minute time period), fig. 2.36, a normal distribution is usually used, (Hassan, Sykes 1990). In the probabilistic method described here, normal distributions are used to model many other wind speed variations too.

Weibull distributions do not tend to fit wind speed distributions when the standard deviation is small compared to the mean, producing a Weibull shape factor greater than 3. This occurs when considering a smaller range of the wind speed power spectrum, for example the long-term distribution of weekly mean or 6-monthly mean wind speeds, figs. 2.34 and 2.35 or the distribution of wind speeds within 10 minutes, one hour or a day, figs. 2.36 to 2.38. Figures 2.45 and 2.46 compare some of these example wind speed PDFs with Weibull and normal distributions that have the same values of mean and standard deviation.

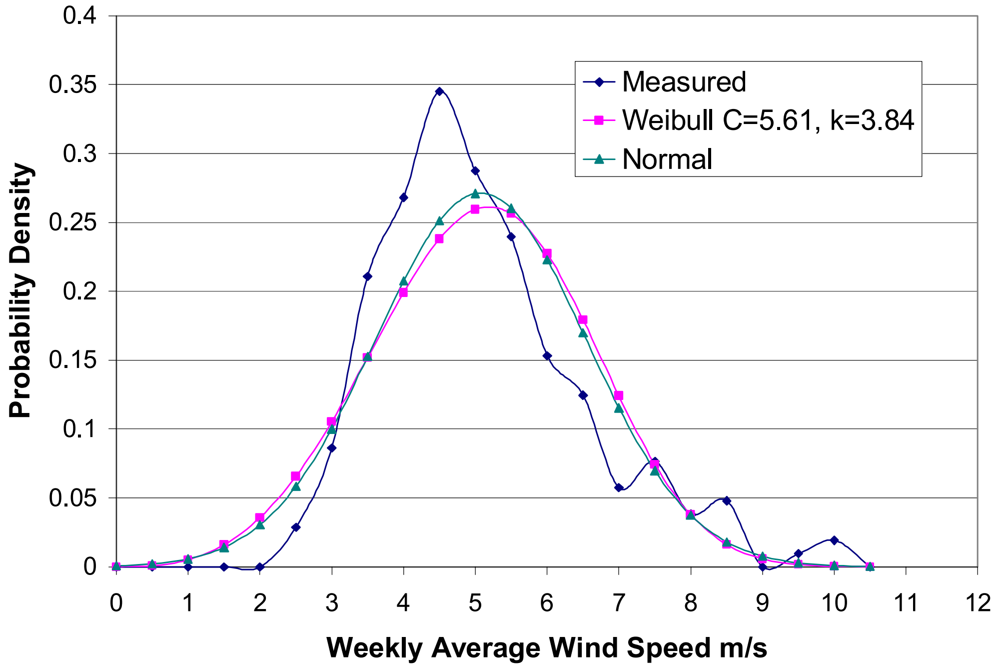

Figure 2.45 Weekly Mean Winds Speeds Measured at RAL Compared with Weibull and Normal Distributions

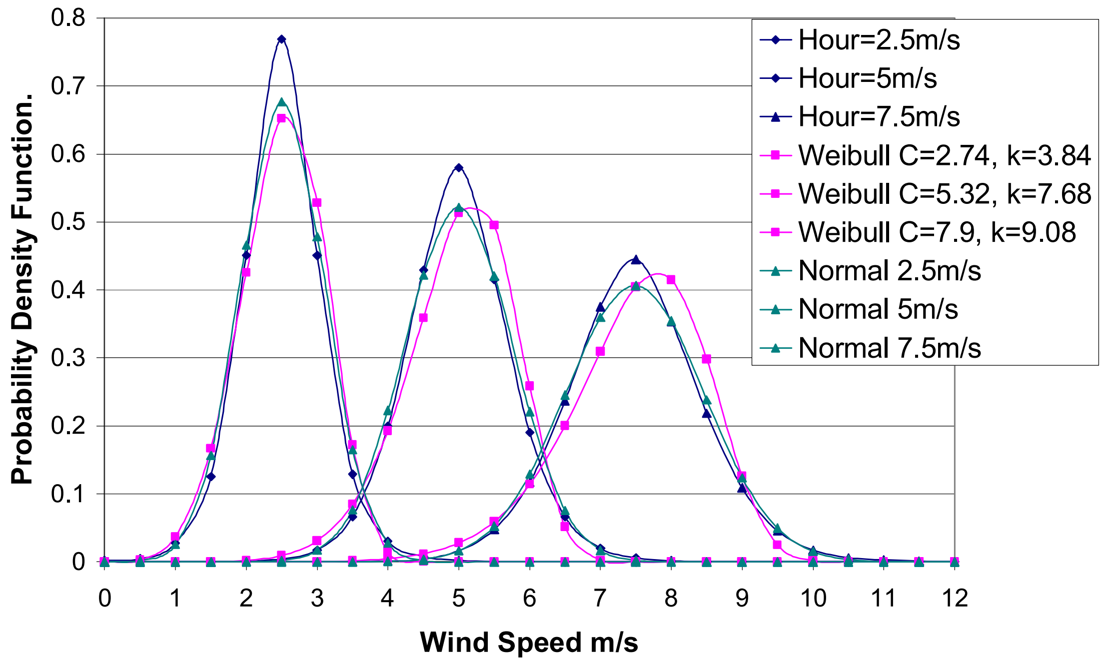

Figure 2.46 PDFs of Wind Speed Within One Hour at RAL Compared with Weibull and Normal Distributions

Wind speed distributions tend to have positive or neutral skewness, even when the standard deviation is small. However, Weibull distributions have negative skewness when the shape factor is greater than about 3.6. The skewness becomes small when the shape factor is more than 3. Fig. 2.45 and fig. 2.46 above show that normal distributions tend to fit wind speed distributions better than Weibull distributions when the standard deviation is small compared to the mean, especially when considering the distribution of wind speeds within a short period of time, fig. 2.46.

In the probabilistic method, a cross-over point was chosen at a Weibull shape factor of 3 ($\sigma/\overline{U}$ of 0.3636). At this point, Weibull and normal distributions with the same mean and standard deviation have very similar shapes, fig 2.47.

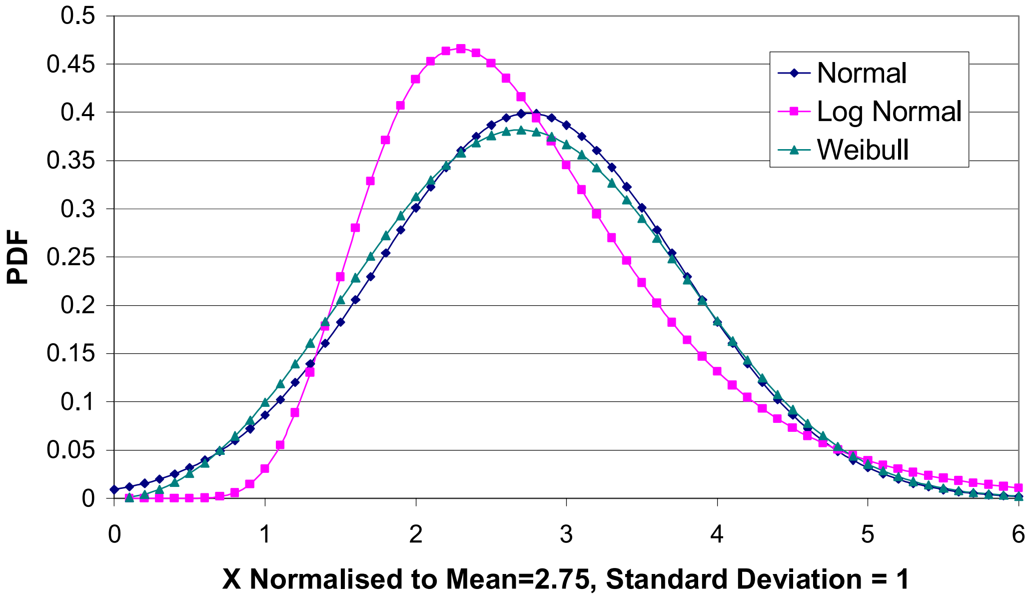

Figure 2.47 Normal, Log-Normal and Weibull Distributions with identical means and standard deviations

In the probabilistic methods presented in this thesis, Weibull distributions are only used when their shape factor is less than 3. Otherwise a normal, Gaussian distribution is used.

2.6.3.3 Log-Normal Distributions

Some long-term wind speed distributions, especially ones with a large positive skewness, are better fitted by a log-normal distribution, (Luna, Church 1974, Shaw, McCartney 1985). Log-normal distributions always have a positive skewness, even when the ratio of standard deviation to mean is small. The skewness of a log-normal distribution is always higher than that of a Weibull distribution, except for extreme cases where the standard deviation is greater than the mean.

The wind speed distributions measured at RAL all have a positive skewness, but generally appear to be closer to a Weibull distribution than a Log-normal. Log-normal distributions were therefore not used in the probabilistic method.

2.6.4 The Wind Turbine Power Curve

The power produced by a wind turbine or wind farm is not proportional to the wind speed, or even the cube of wind speed, but is governed by a turbine power curve. All wind turbines have a cut-in wind speed below which they do not generate any electricity, because the power from the wind is insufficient to overcome mechanical and electrical losses in the turbine. Wind turbines also have a rated wind speed, at which the generator delivers its rated or design electrical power. Above the rated wind speed, a turbine may continue to generate the rated electrical power, or the power output may vary slightly with wind speed, depending on whether the turbine is pitch regulated, stall regulated or yaw regulated, and on other design characteristics. Between cut-in and rated wind speed, the power curve is often close to a cubic curve, because the power available per swept area is proportional to the cube of wind speed.

Most wind turbines also have a cut-out or furling wind speed, above which the turbine stops rotating and stops generating electricity, in order to protect the turbine from damage caused by excessively high winds.

Despite the differences in design, the power curves for different large wind turbines of a given rated power look remarkably similar, fig. 2.48. The turbine power curve used in the probabilistic method is a composite power curve from three commercial wind turbines, see table 2.4 The data was taken from a catalogue of wind turbines published by the German wind energy agency in 2000, (Bundesverband WindEnergie e.V. 2000).

Table 2.4 Technical Data from 3 Commercial Wind Turbines, Each of 1MW Rated Power

Manufacturer:	**NEG Micon**	**BWU**	**Enercon**
Model	NM 1000-250/60	1000/57	E-58
Rotor diameter, m	60	57	58
Rotation speed, RPM	18rpm above 250kW and 12rpm below 250kW, twin speed	22.9rpm above 250kW and 15.3rpm below 250kW, twin speed	10 to 24rpm, variable speed (pitch controlled)
Rotor mass, kg	23,000	18,000	21,000
Cut-in wind speed, m/s	3	3	2.5
Rated wind speed, m/s	13	13	13
Furling wind speed, m/s	20	25	None
Survival wind speed, m/s	54	55.8	59.5
Wind speed, m/s	**Power Curve Values, kW**		
2	0	0	0
3	1.7	0	2.6
4	36.7	7.15	25.55
5	84.0	73.38	75.63
6	142.1	125.4	142.69

7	263.0	199.2	228.82
8	399.2	315.8	360.0
9	540.7	457.9	511.76
10	662.6	625.1	645.0
11	761.3	795.4	852.46
12	873.7	948.0	913.54
13	954.4	1004.0	989.72
14	1020.4	1026.0	980.33
15	1037.8	1020.0	1011.99

These power curves have been plotted in fig. 2.48 and a curve fitted through all the data.

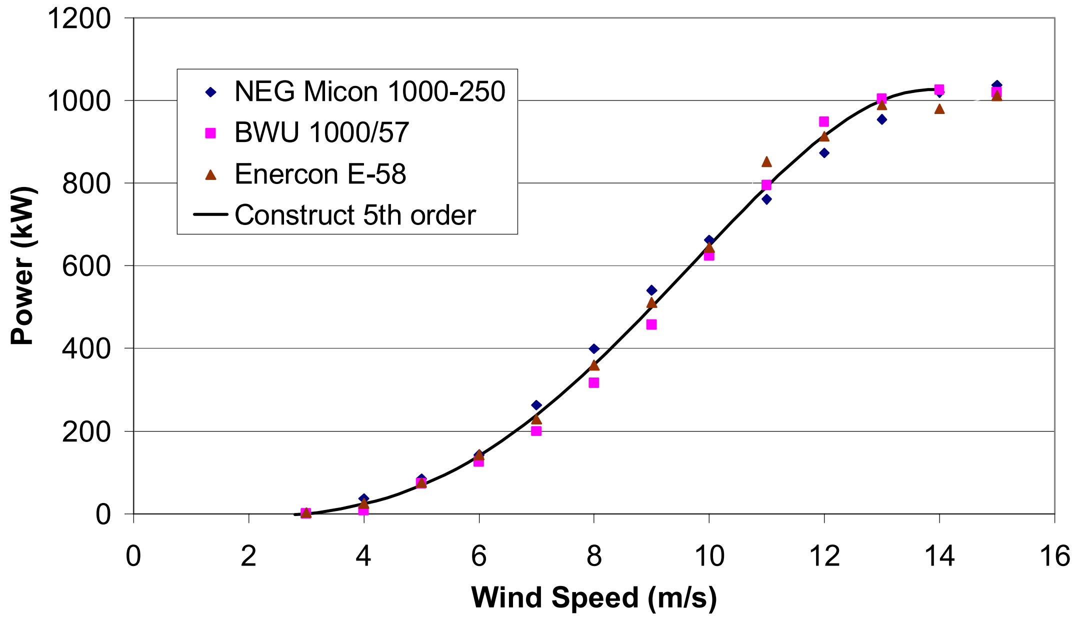

Figure 2.48 Wind turbine power curves from 1MW machines and the probabilistic method's generic turbine curve fitted through them

The fitted curve is a fifth-order polynomial between cut-in and rated wind speed:

$Power = 0$ when $U < 3.0002$

$$Power = 0.0015U^5 - 0.1967U^4 + 4.0058U^3 - 19.268U^2 + 43.607U - 50$$
$$\text{when} \quad 3.0002 \le U \le 12.9944$$

$$Power = 1000 \qquad \text{when} \quad 12.9944 < U \le 25.0$$

$$Power = 0 \qquad \text{when} \quad U > 25.0 \tag{2.17}$$

Although this data is now at least four years old, the similarity of these three turbine curves and the maturity of the technology suggest that turbine power curves will not change radically from this generic turbine power curve.

2.6.5 Magnitude of Short-Term Variations

The probabilistic method always assumes that the standard deviation of wind speeds within any given time period, e.g. one hour, one day, one week are proportional to the hourly, daily, or weekly average wind speed respectively. This is a standard assumption for turbulent variations in wind speed, (Hassan, Sykes 1990); turbulent variations are proportional to the 10-minute average wind speed and the constant of proportionality is the turbulence intensity. However, this assumption is not obvious for longer time scales.

When real wind speed data is analysed, the standard deviation of wind speed within a period are not necessarily proportional to the average wind speed over that period, especially when the averaging period is one day or more in length, fig. 2.49. There is a considerable random scatter in the size of the standard deviation.

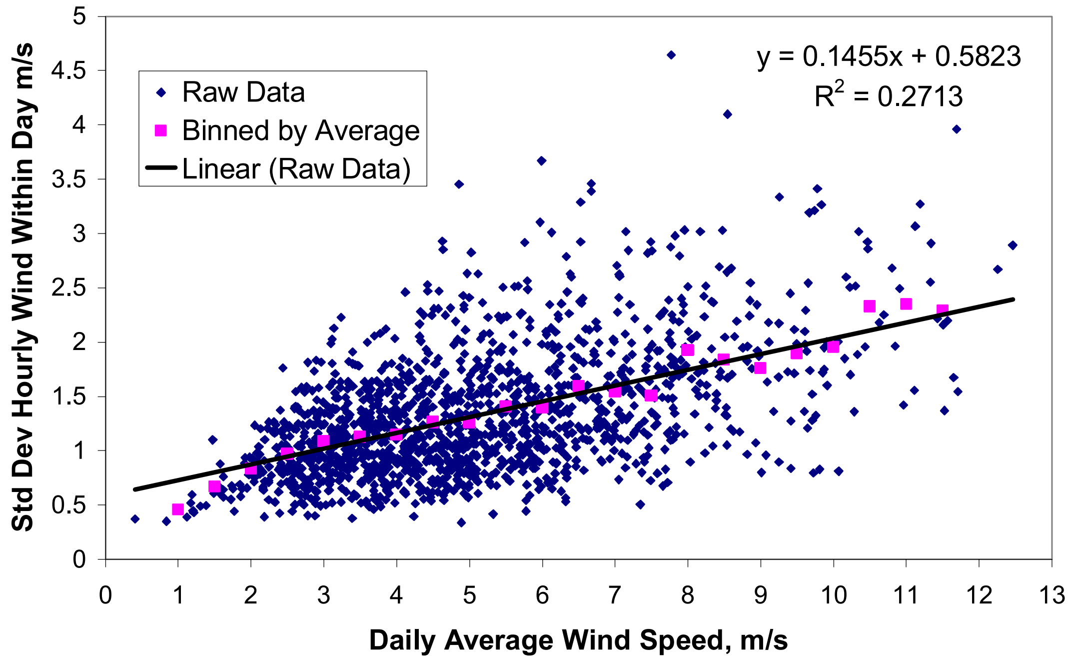

Figure 2.49 Daily average wind speeds plotted against standard deviation of hourly wind speeds within those days, at Rutherford Appleton Laboratory, 1994 to 1998

If the data is binned according to the daily average wind speed, then a non-proportional trend is visible. If a linear trend line is fitted to the original data (or the binned data) this also has a non-zero y-intercept, although the line provides a good fit to the binned averages over a wide range of daily averages from 2m/s to 11.5m/s. When the daily average wind speed is zero, the trend line suggests that the standard deviation is apparently not zero. The linear trend line equation is also quoted on fig. 2.49, together with its coefficient of correlation against the raw (un-binned) data.

Calculations were therefore done to discover whether the predictions of energy store performance were improved by modification to the standard deviation of wind speed within a period. For example, the prediction of the time that the store was empty or full might be improved. However, no improvement was seen.

Further investigation shows that the correlation between standard deviation within a period and the period-average value depends on which parameter is assumed to be the independent variable and on which is assumed to be the dependent variable. If the data is binned according to intervals of standard deviation, and the period average is plotted as the dependent variable, then the correlation is reversed, fig. 2.50. According to this new trend line, if the standard deviation were zero, then the period average wind speed would be positive and non-zero.

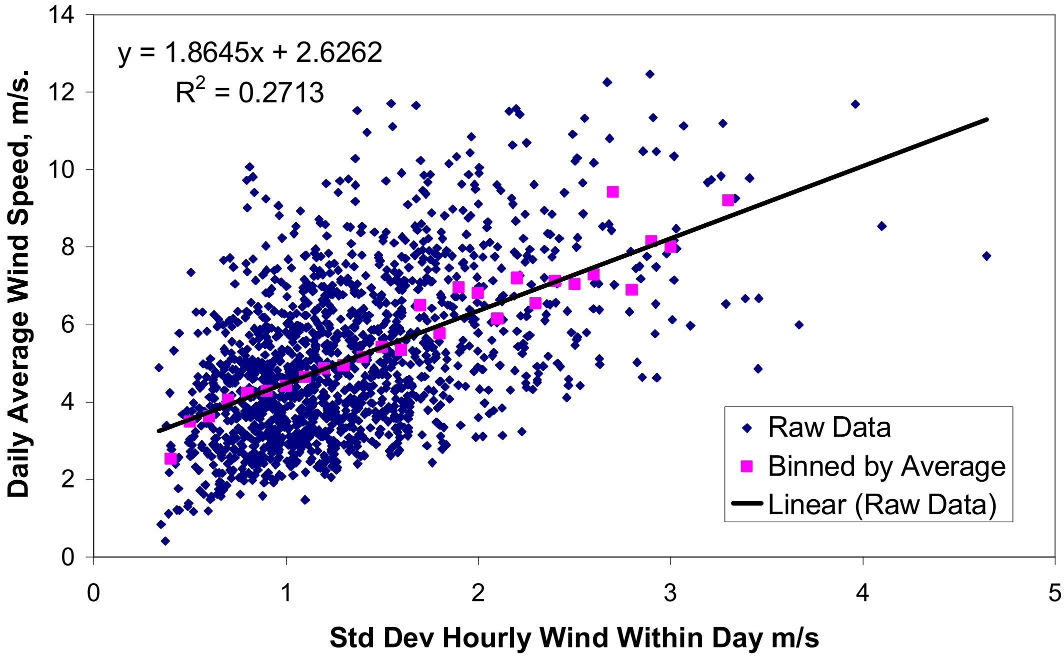

Figure 2.50 Standard deviation of hourly wind speeds within a day plotted against the daily average wind speed of those days, at Rutherford Appleton Laboratory, 1994 to 1998

Again, the binned data is plotted and the new trend line together with its equation and coefficient of correlation are quoted on fig. 2.50.

Similar results were obtained for other sampling periods and averaging periods, for example 10-minutes within a week, or days within a month. The above study shows that although there is a lot of scatter in the standard deviations of variation within a period, no better model has emerged than the original, proportional model. The non-zero y-intercepts shown in figures 2.49

and 2.50 appear to be merely products of the statistical analysis method and do not convey any extra, useful insights into properties of the original data. The probabilistic method therefore uses the assumption of proportionality for standard deviations of wind speed within a given period.

2.7 Solar Power Model

2.7.1 Preparation of Cleaned Solar Irradiance Data

Total solar irradiance on the horizontal plane was measured at Rutherford Appleton Laboratory (RAL) from January 1999 to December 2002 inclusive. The measurements were taken as 1-minute averages and are concurrent with the second data set of four years of wind speed measurements at RAL, see section 2.2.1.2. The data record contained some gaps and bad data that were filled with data copied from elsewhere in the same month, or occasionally from adjacent months or from the same month in another year. The replacement data was always exactly the same length as the gap, so that the total number of data records in each month was made correct. The replacement data was also taken from almost the same times of day, so that the diurnal pattern would be preserved, and the replacement data was a 'best guess' of the missing data. Many months contained just 3 gaps of less than 20 minutes each, created when the recording tapes were changed. Some other months contained many more gaps of just a few minutes each, and a few months contained gaps of several hours or days each.

2.7.2 The Solar Power Spectrum

This solar spectrum, fig. 2.10 and fig. 2.11, contains information about the time-varying nature of solar radiation, and is not to be confused with the solar electromagnetic spectrum from infrared radiation to ultraviolet rays. It is a time variation spectrum, analogous to the wind power spectrum, fig. 2.21 or the Van Der Hoven spectrum, fig. 1.5 (Hassan, Sykes 1990, Spera 1994, Van Der Hoven, I. 1957).

The processed solar radiation data contains exactly 2,103,840 records. That is the number of minutes in 4 years, including one leap year (year 2000). The cleaned data was first input to a Matlab program that performed a fast Fourier transform and created a scaled spectrum of solar irradiance variation, i.e. a 'power' spectrum. The data spans all time scales from 4 years down to 2 minutes, the Nyquist frequency. This produced a smooth spectrum up to time scales of about 1 week, fig. 2.11 and a rough spectrum at longer time scales.

The spectrum is dominated by specific frequencies, fig. 2.10, due to diurnal and seasonal variations in solar irradiance, but stochastic (broadband) variations are also present and significant, fig. 2.11 and see section 2.4.6. The ratio of total standard deviation to mean is larger than for wind speed or electrical load variations.

The fastest variations are related to passing clouds and are located at the right-hand end of the spectrum, fig. 2.11, of the order of 10^{-1} hours. Synoptic variations, due to passing weather systems are in the centre of the spectrum, fig. 2.11, of the order of 10^{2} hours. Interseasonal and inter-annual variations are at the far left-hand end of the spectrum, fig. 2.10, of the order of 10^{5} hours. Not surprisingly, the largest spike is at a frequency of 24 hours, fig. 2.10. Harmonics of 24 hours: 12 hours, 8hours, 6 hours, 4.8 hours, and 4 hours are also clearly visible, fig. 2.11. At the left-hand end of the spectrum, fig. 2.10 there is a spike at 1 year representing the seasonal variation in solar radiation. The seasonal variation is so close to being a pure sine wave that harmonics of annual variation are not visible as separate spikes.

2.7.3 Solar Power Probability Density Function

If the instantaneous solar irradiance is simply put into bins and plotted, then the resulting probability distribution function is not very useful for modelling. The graph is dominated by times of zero radiation (at night) and periods of very low radiation, with a long, lumpy tail at higher values of irradiance, fig. 2.1. This causes the standard deviation of solar irradiance to be larger than its mean, see section 2.4. The lower values of radiation are sometimes due to the sun being low in the sky and at other times due to cloudy conditions, but this graph does not tell us the relative importance of these two effects.

2.7.4 Clearness Index

A more useful approach is to plot clearness index. The position of the sun in the sky is a precise function at each moment in time of sun-earth geometry and orbits. Published equations can give the elevation and azimuthal angle of the sun, (Duffie, Beckman 1974). The elevation of the sun gives the theoretical solar radiation on a horizontal surface, fig. 2.51, given a standard extra-terrestrial solar radiation of 1367W/m^2 and neglecting atmospheric absorption.

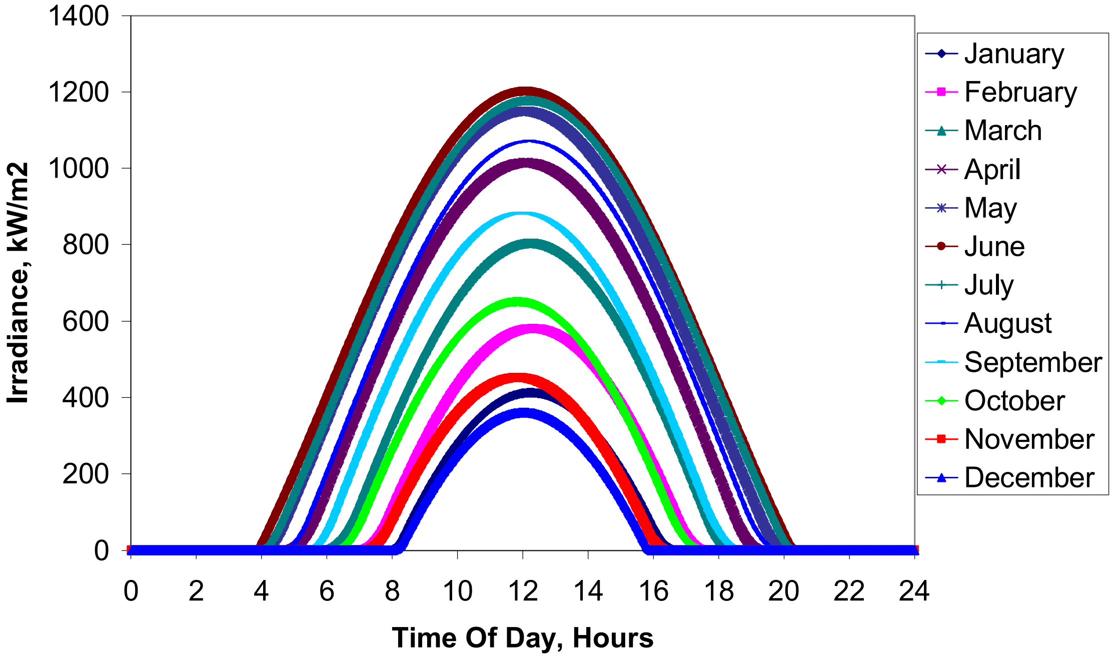

Figure 2.51 Extra-terrestrial global horizontal solar radiation as a function of time of day, averaged for each month

The measured global horizontal radiation has been plotted in a similar way, fig. 2.52. The data was first binned by month of the year, then by minute of the day. Each data point represents an average of 4 years worth of data for that minute of the day, and that month of the year, typically 30 days x 4 years = 120 measurements per point.

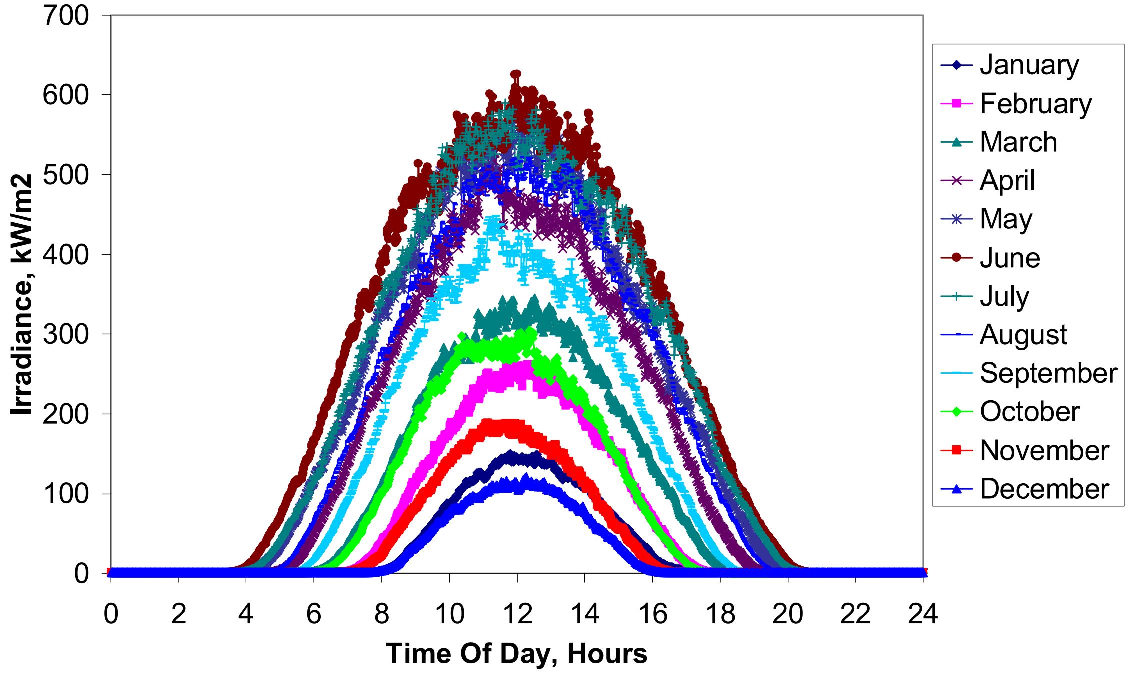

Figure 2.52 Measured global horizontal radiation as a function of time of day at Rutherford Appleton Laboratory, 1999 to 2002

2.7.4.1 The Equation of Time

Both the extra-terrestrial and the actual irradiance graphs clearly show the seasonal variations in solar elevation, sunrise and sunset times. They also show the effect of the 'equation of time', that is the way solar noon moves earlier or later depending on the time of year. January, February, March, July and August solar irradiances are shifted later in the day, but May, September, October, November and December solar irradiances are shifted earlier in the day (Duffie, Beckman 1974), fig. 2.53

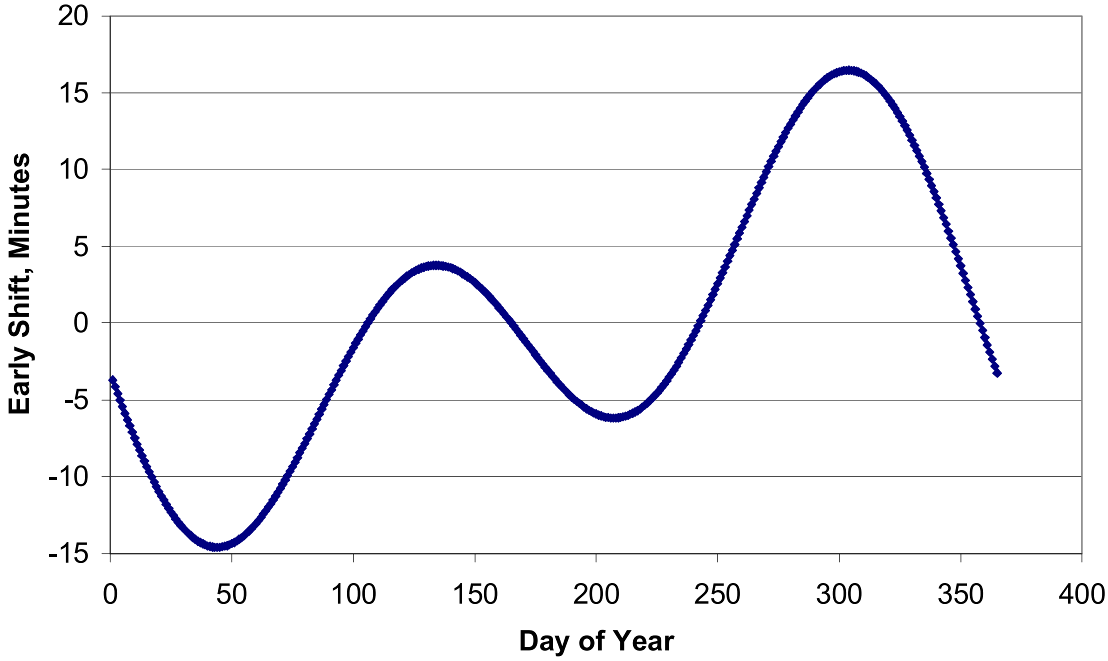

Figure 2.53 The solar equation of time

2.7.4.2 Typical Days in Each Month

As described below, the probabilistic method calculates the average solar radiation for each month of the year. The method performs this calculation by considering the solar elevation on 12 typical days, (Duffie, Beckman 1974) that have extra-terrestrial solar radiations closest to the monthly averages. These typical days are 17th January, 16th February, 16th March, 15th April, 15th May, 11th June, 17th July, 16th August, 15th September, 15th October, 14th November and 10th December.

2.7.4.3 Measured Clearness Index

By comparing the actual global horizontal radiation with the extra-terrestrial global horizontal radiation at the same instant of time, a time series of clearness index was calculated. This is only meaningful when the sun's elevation is more than about 5.7°. This elevation angle corresponds to an air mass of 10; i.e. sine of the elevation angle is 0.1. When the elevation angle is less than 5.7°, the beam component of radiation is very low and the horizontal radiation is dominated by diffuse radiation, some of it coming from sunlight bouncing around the atmosphere from beyond the horizon. At dawn and dusk for example, the calculated clearness index becomes infinite. Therefore, the clearness index has only been used for correlations and empirical formulas where the sun's elevation was greater than 5.7°.

The use of clearness index does not completely remove the effect of solar elevation, fig. 2.54. Clearness index displays a clear positive trend with solar elevation. When the elevation angle is less than 5.7°, the trend is lost in random scatter.

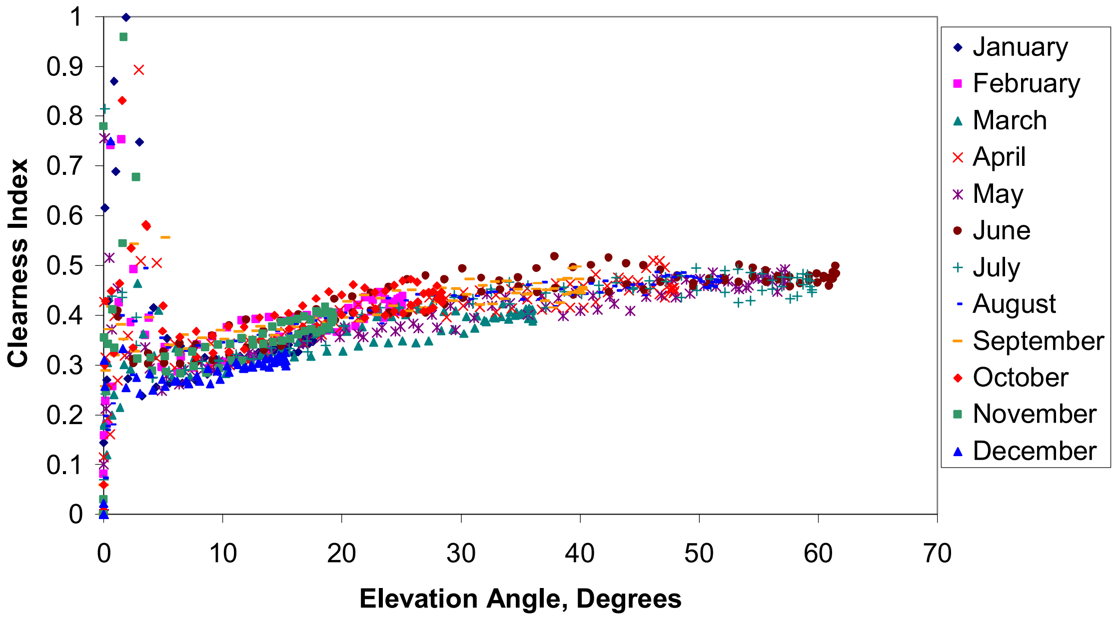

Figure 2.54 Overall clearness index vs. solar elevation

Some seasonal and diurnal dependence is also noticeable. For example, March days tend to be darker than average, especially in the afternoon. In contrast, October days tend to be lighter than average, especially in the morning. The causes of these anomalies are not known. However, the graph clearly shows a broad trend of increasing clearness index with solar elevation.

The clearness index calculated above could have been put through a fast Fourier transform to produce an alternative solar spectrum, reducing the dominance of the daily and seasonal spikes. However, this approach is problematic, since it is not clear how to treat the clearness index at night. If night-time and low sun elevation were simply chopped out of the time series, then winter days would be shorter, leading to speeded-up weather effects, whereas summer days would be longer leading to slowed weather effects. If night time clearness index were treated as constant, then winter variations would be smaller than summer variations. In any case, clearness index still has a strong correlation with solar elevation, fig. 2.54. In clear skies, low elevation means greater air mass and more attenuated beam radiation. In cloudy skies, low elevation means even greater attenuation of the diffuse radiation. In partially cloudy skies, low elevation may also mean a greater probability that a cloud will stop the beam radiation. The fast Fourier transform, fig. 2.10, was therefore performed on absolute solar radiation.

2.7.4.4 Models of Clearness Index in Clear and Cloudy Skies

An overall probability distribution function of clearness index, binned by clearness index is shown in fig. 2.55. All data where the sine of elevation angle is less than 0.1 (solar elevation less than 5.7°) has been removed.

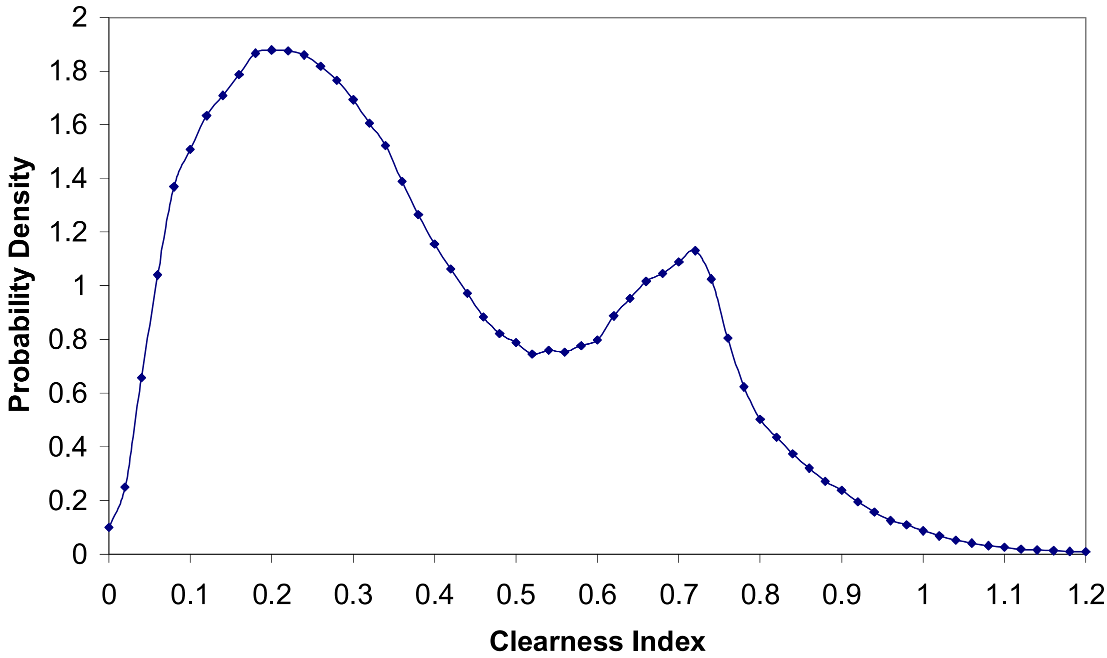

Figure 2.55 Probability distribution of overall clearness index measured at RAL from 1999 to 2002

This graph shows a double peak of clearness index: A broad peak centred at about 0.22, and a smaller, sharper peak centred at about 0.72. It is reasonable to assume that these peaks correspond to cloudy conditions and sunny conditions respectively. The graph means that at any instant of time, the probability of cloudy conditions (no beam radiation at falling on the ground) is high, and the probability of sunny conditions (beam radiation at ground level) is a little lower. The probability of semi-sunny conditions, in the middle of the graph, is low, and the probability of very high clearness indices, at the right hand end of the graph is even lower. Common experience confirms that the sun tends to be either ‘in’ or ‘out’ for long periods at a time, and changes from one state to the other relatively quickly.

As seen above in fig. 2.54, the clearness index also depends on the sun’s elevation. In fig. 2.56, the data has therefore been binned into solar elevation, and then into clearness index. Each line in fig. 2.56 represents a different range of solar elevations, calculated as sine of elevation angle, from 0 to 0.9 in steps of 0.1. The probabilities have been scaled by the solar elevation, mainly to separate the lines and make them clearly visible, but also to weight them according to their relative importance in total annual solar energy. For each value of solar elevation, the clearness index values are binned into intervals of 0.02 as before.

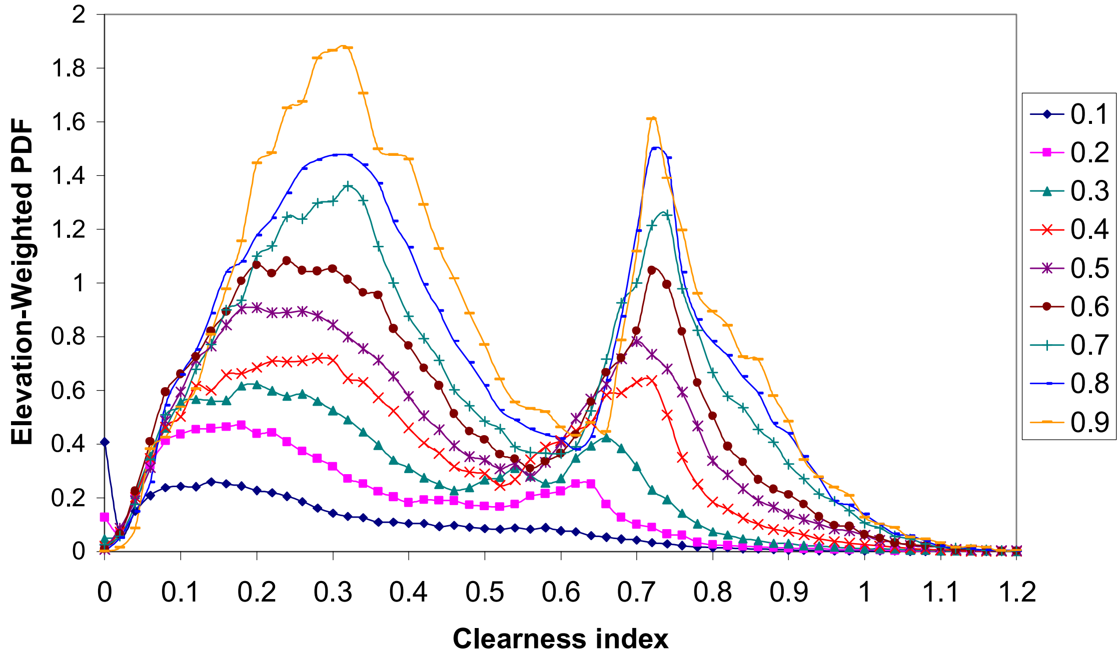

Figure 2.56 Probability density functions of clearness index for each solar elevation, measured at RAL from 1999 to 2002

Fig. 2.56 shows that the highest solar elevation is 0.9. Actually, the latitude of Rutherford Appleton Laboratory is 51.57° north. So on 21st June, at solar midday, the sun's elevation angle is (90° - 51.57°) + 23.45° = 61.88°. The sine of elevation is then 0.882, which gets rounded up to 0.90. the '0.9' line therefore represents a relatively small amount of data measured when the solar elevation was close to its annual maximum value.

Because the heights of the lines have been scaled, the relative heights of the curves are not important, but the positions of the peaks indicate the attenuation of the atmosphere in cloudy and clear conditions. These, and the mean values of clearness index for each value of elevation, have been used in the solar power models of the probabilistic method.

The position of the cloudy peak starts at about 0.29 for a solar elevation of 0.8 or 0.9, dropping to 0.17 for a solar elevation of 0.2. The positions are approximate, since there is some scatter in the graphs of fig. 2.56, and the cloudy peak is very rounded. The position of the sunny peak starts at about

0.74 for a solar elevation of 0.8 or 0.9, dropping to about 0.62 for a solar elevation of 0.2. Again, the positions are approximate, but better defined than the cloudy peaks.

The graphs seem to show that at high solar elevations, the probability of beam radiation is almost as high as the probability of cloud cover; the sunny peak is as high as the cloudy peak. Then with reducing solar elevation, the sunny peak declines in height and disappears entirely. This would suggest that the probability of beam radiation declines with reducing solar elevation.

However, the graphs may also be interpreted as showing that the cloudy and sunny peaks both decline in height, get broader and merge together with reducing solar elevation. As will be seen later, the probabilistic method assumes no change in probability of sunny conditions with changing solar elevation. A ‘sunny fraction’ parameter is calculated that appears to be almost independent of solar elevation, see section 2.7.5.

The positions of the peaks are plotted on a logarithmic scale in fig. 2.57. For this graph, the solar elevation sine was binned into smaller intervals of 0.05 for greater resolution, but the graph is plotted in terms of air mass. Fig. 2.57 anticipates that the solar irradiance is attenuated by the thickness of atmosphere through which it passes, i.e. the air mass:

$$AirMass, M = \frac{1}{Elevation} = \frac{1}{\sin\left(ElevationAngle\right)} \qquad (2.18)$$

The equations quoted on the graph are those of the trend lines fitted through the sunny peak data, the cloudy peak data and the overall average data. Each has its own correlation coefficient, R^2 showing the suitability of a linear fit.

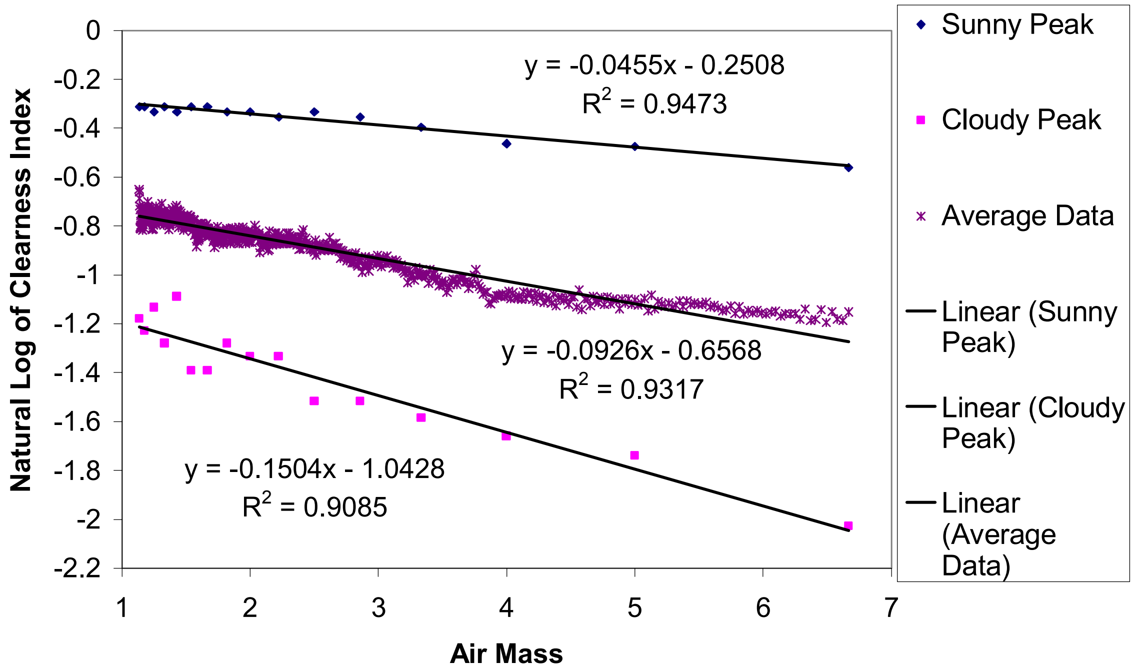

Figure 2.57 Clearness index peak positions as a function of air mass

The model assumes that the solar power is attenuated according to the following law:

$$ClearnessIndex, CI = K\exp(-\alpha M) \tag{2.19}$$

Where: K=Scale factor

α=Attenuation factor

Rearranging: $\ln(CI) = \ln(K) - \alpha M$ (2.20)

Three lines are plotted onto fig. 2.57: one for sunny conditions, another for cloudy conditions, and a third for the average of all conditions at each solar elevation (air mass). A best straight line has been fitted to each set of data. The gradient of each straight line is -α and the y-intercept is ln(K).

For beam radiation the gradient is –0.0455, so the beam radiation attenuation factor, α_S is 0.0455. The beam radiation y-intercept is =–0.2508. The beam radiation scale factor, K_S is therefore 0.7782. So the probabilistic model for sunny conditions is:

$$ClearnessIndex, CI = 0.7782\exp(-0.0455M) \qquad (2.21)$$

The cloudy gradient is –0.1504 so the cloudy attenuation factor, α_C is 0.1504. The cloudy y-intercept is -1.0428. The graph would therefore give a cloudy scale factor, K_C of 0.3525, but the cloudy scale factor was reduced to K_C=0.2729 in order to preserve the correct total variance in solar irradiance and so make the probabilistic model work, as will be seen later, in section 2.7.5. For a given solar elevation, the whole PDF of solar clearness index is replaced by just two discrete values, fig. 2.59. So the probabilistic model for cloudy conditions is:

$$ClearnessIndex, CI = 0.2729\exp(-0.1504M) \qquad (2.22)$$

One can see that the cloudy-sky clearness index is much more dependent on air mass than the sunny-sky clearness index; the slope of the cloudy clearness index is steeper. This is simply because clouds attenuate solar radiation faster than a clear atmosphere.

2.7.5 The Probabilistic Model of Solar Radiation

The probabilistic model used in this thesis uses a ‘sunny fraction’ parameter to describe the probability of beam radiation or cloudy conditions at any instant of time. The cloud cover is assumed to be completely bivalent, transmitting either beam radiation plus some diffuse radiation according to eq. 2.21, or cloudy, diffuse radiation only with no beam radiation, according to eq. 2.22. The probability of the presence of beam radiation is the sunny fraction, f and the probability of no beam radiation is 1-f.

Fig. 2.56 suggests that the sunny fraction might reduce with increasing air mass. However, when the average global solar radiation is plotted on the same graph as the sunny and cloudy peak positions, fig. 2.57, the average line is almost half way between the sunny and cloudy lines, with approximately half the gradient, for all values of air mass. The probabilistic method therefore makes a simplifying assumption that the sunny fraction is independent of time of day (solar elevation). Instead, the sunny fraction is assumed to vary only with month of the year, i=1 to 12 and with the daily average solar radiation, j=1 to N values of daily average, with j=1 representing completely cloudy conditions to j=N representing completely clear skies all day.

In practice, the average level of cloud cover at RAL is very constant through the year. Just one long-term value of sunny fraction could have been used, but the sunny fraction was allowed to vary through the year in order to make the model as generally applicable as possible. Many other parts of the world have dry seasons and rainy seasons, and these may produce radically different sunny fractions at different times of the year.

The model sunny and cloudy clearness attenuation factors, α_S and α_C, are given directly by the graph of peak positions, fig. 2.56 and eq. 2.21 and eq. 2.22. The sunny scale factor is also given by fig. 2.57. This leaves two global quantities to calculate: the cloudy scale factor, K_C and the overall average sunny fraction, $\overline{f}$. We also have two known quantities that can help us calculate these two unknowns simultaneously: the long-term average solar radiation and the long-term variance in solar radiation. The probabilistic method makes two simplifying approximations in order to calculate the cloudy scale factor:

1. The sunny fraction is independent of the solar elevation within a day
2. The month-to-month variation in sunny fraction is effectively independent of the average solar elevation in each month. This assumption certainly applies where the sunny fraction varies little throughout the year and/or the solar elevation varies little throughout the year.

Returning to the power spectrum of solar radiation variation, fig. 2.10 and fig. 2.11, the method preserves the total solar variation due to all factors: diurnal, seasonal and stochastic. For the weather measurement site at RAL, the average solar irradiance is 118.8W/m^2 and the total standard deviation in solar irradiance is 201.8 W/m^2 calculated from the original data or 202.3 W/m^2 calculated by integrating the area under spectrum and taking the square root. The probabilistic method calculates the values of cloudy scale factor and average sunny fraction that preserves these quantities.

At first sight, this method might appear to give ill-conditioned simultaneous equations; the stochastic variation is small compared to the periodic variation, and we are trying to calculate the effect of the stochastic variance using the whole spectrum. However, the sizes of the diurnal and seasonal spectral spikes are proportional to the average solar irradiance, and we have already set the scale factor, under sunny conditions, K_S. The simultaneous equations work well, because the long-term average solar irradiance, $\overline{I}$ is well defined. The following notes should be observed for accurate results:

1. Note that if the latitude of the modelled electricity system is significantly different from the latitude of the measurements from which the spectrum is produced, then the calculation of K_C and $\overline{f}$ would have to be done using the latitude at which the solar spectrum was measured and calculated, not the latitude of the modelled location. Otherwise, the variance represented by the diurnal and seasonal spikes in the spectrum would distort the calculated values of K_C and $\overline{f}$.

2. Note that the value of long-term average solar irradiance, $\overline{I}$ to be used here is the one associated with the spectrum, not the one for the site being modelled. The monthly average solar irradiances, I_i are used later to give the correct average solar irradiance for the modelled site, within each month and in total.

3. Note that the assumption of sunny fraction being independent of solar elevation probably breaks down at high latitudes if the cloud cover is very seasonal. The model would then correctly predict the total solar irradiance and the average irradiance within each month, but would

over-predict or under-predict the variation within each month. This problem is left for future work.

Let us consider one hour, k within a day of cloud type j, within month i.

In general, the expected value of instantaneous solar irradiance, I_{ijk} is:

$$I_{ijk} = I_0 \left[f_{ij} K_S \frac{\exp(-\alpha_S M_{ik})}{M_{ik}} + (1 - f_{ij}) K_C \frac{\exp(-\alpha_C M_{ik})}{M_{ik}} \right] \tag{2.23}$$

Where I_0 is the global, normal extra-terrestrial solar irradiance, 1367W/m^2

f_{ij} is the sunny fraction for that month and day cloud type

The daily average solar radiation is:

$$I_{ij} = I_0 \frac{1}{24} \sum_{k=1}^{24} \left[f_{ij} K_S \frac{\exp(-\alpha_S M_{ik})}{M_{ik}} + (1 - f_{ij}) K_C \frac{\exp(-\alpha_C M_{ik})}{M_{ik}} \right]$$

$$= I_0 f_{ij} K_S \frac{1}{24} \sum_{k=1}^{24} \left[\frac{\exp(-\alpha_S M_{ik})}{M_{ik}} \right] + I_0 (1 - f_{ij}) K_C \frac{1}{24} \sum_{k=1}^{24} \left[\frac{\exp(-\alpha_C M_{ik})}{M_{ik}} \right] \tag{2.24}$$

2.7.6 Calculation of Long-Term Average Sunny Fraction

Now we assume that the total variance, V_{TOT} of the solar spectrum represents not only the variation in solar elevation (periodic effects) but also the variation due to changing cloud cover (largely a stochastic effect). Let $\overline{f}$ be the long-term annual sunny fraction. For the moment we assume this is independent of season or time of day.

So the long-term average solar radiation is:

$$\overline{I} = \frac{1}{12} \sum_{i=1}^{12} \frac{1}{24} \sum_{k=1}^{24} \left[I_0 \overline{f} \frac{K_S}{M_{ik}} \exp(-\alpha_S M_{ik}) + I_0 (1 - \overline{f}) \frac{K_C}{M_{ik}} \exp(-\alpha_C M_{ik}) \right]$$

$$= \frac{I_0 \overline{f} K_S}{12 \times 24} \sum_{i=1}^{12} \sum_{k=1}^{24} \left[\frac{\exp(-\alpha_S M_{ik})}{M_{ik}} \right] + \frac{I_0 (1 - \overline{f}) K_C}{12 \times 24} \sum_{i=1}^{12} \sum_{k=1}^{24} \left[\frac{\exp(-\alpha_C M_{ik})}{M_{ik}} \right] \tag{2.25}$$

Total variance in solar radiation, $V_{TOT} = \frac{1}{12 \times N \times 24} \sum I_{ijk}^2 - \left(\bar{I}\right)^2$

$$= \frac{\bar{f}}{12} \sum_{i=1}^{12} \frac{1}{24} \sum_{k=1}^{24} \left[I_0 K_S \frac{\exp\left(-\alpha_S M_{ik}\right)}{M_{ik}} \right]^2 + \frac{\left(1-\bar{f}\right)}{12} \sum_{i=1}^{12} \frac{1}{24} \sum_{k=1}^{24} \left[I_0 K_C \frac{\exp\left(-\alpha_C M_{ik}\right)}{M_{ik}} \right]^2 - \left(\bar{I}\right)^2$$

$$= \frac{I_0^2 \bar{f} K_S{}^2}{12 \times 24} \sum_{i=1}^{12} \sum_{k=1}^{24} \left[\frac{\exp\left(-\alpha_S M_{ik}\right)}{M_{ik}} \right]^2 + \frac{I_0^2 \left(1-\bar{f}\right) K_C{}^2}{12 \times 24} \sum_{i=1}^{12} \sum_{k=1}^{24} \left[\frac{\exp\left(-\alpha_C M_{ik}\right)}{M_{ik}} \right]^2 - \left(\bar{I}\right)^2$$

(2.26)

Now let us use some working variables to simplify the equations above:

Let $$X = \frac{I_0}{12 \times 24} \sum_{i=1}^{12} \sum_{k=1}^{24} \left[\frac{\exp\left(-\alpha_S M_{ik}\right)}{M_{ik}} \right] \tag{2.27}$$

$$Y = \frac{I_0}{12 \times 24} \sum_{i=1}^{12} \sum_{k=1}^{24} \left[\frac{\exp\left(-\alpha_C M_{ik}\right)}{M_{ik}} \right] \tag{2.28}$$

$$P = \frac{I_0^2}{12 \times 24} \sum_{i=1}^{12} \sum_{k=1}^{24} \left[\frac{\exp\left(-\alpha_S M_{ik}\right)}{M_{ik}} \right]^2 \tag{2.29}$$

$$Q = \frac{I_0^2}{12 \times 24} \sum_{i=1}^{12} \sum_{k=1}^{24} \left[\frac{\exp\left(-\alpha_C M_{ik}\right)}{M_{ik}} \right]^2 \tag{2.30}$$

Then eq. 2.25 becomes:

$$\bar{I} = \bar{f} X K_S + \left(1-\bar{f}\right) Y K_C \tag{2.31}$$

And eq. 2.26 becomes:

$$V_{TOT} = \bar{f} P K_S{}^2 + \left(1-\bar{f}\right) Q K_C{}^2 - \left(\bar{I}\right)^2 \tag{2.32}$$

Rearranging eq. 2.31 we get:

$$K_C = \frac{\bar{I} - K_S X \bar{f}}{\left(1-\bar{f}\right)Y} \tag{2.33}$$

And combining with eq. 2.32 we get:

$$V_{TOT} = \bar{f} P K_S^{\,2} + \frac{\left(1-K_S X \bar{f}\right)^2}{\left(1-\bar{f}\right)Y^2} Q - \left(\bar{I}\right)^2 \tag{2.34}$$

$$\Rightarrow \left(V_{TOT} + \bar{I}^2\right)\left(1-\bar{f}\right) = P K_S^{\,2} \bar{f}\left(1-\bar{f}\right) + \frac{\left(1-K_S X \bar{f}\right)^2}{Y^2} Q$$

Expanding and rearranging terms, we get a quadratic equation in $\bar{f}$:

$$\bar{f}^2 \left[K_S^{\,2}\left(P - \frac{X^2 Q}{Y^2}\right)\right] + \bar{f}\left[2\bar{I}K_S \frac{XQ}{Y^2} - V_{TOT} - \bar{I}^2 - K_S^{\,2} P\right] + \left[V_{TOT} + \bar{I}^2 - \frac{\bar{I}^2 Q}{Y^2}\right] = 0 \tag{2.35}$$

We can now use the standard solution to a quadratic equation:

Let: $$A = K_S^{\,2}\left(P - \frac{X^2 Q}{Y^2}\right) \tag{2.36}$$

$$B = 2\bar{I}K_S \frac{XQ}{Y^2} - V_{TOT} - \bar{I}^2 - K_S^{\,2} P \tag{2.37}$$

$$C = V_{TOT} + \bar{I}^2 - \frac{\bar{I}^2 Q}{Y^2} \tag{2.38}$$

This can be solved to give two possible real values for $\bar{f}$. For the RAL solar data, *A* = -9964, *B* = 15695 and *C* = 10518. Since *A* is negative, only the negative square root gives a positive value of $\bar{f}$, and this is the value used in the model. It is assumed that the signs of A, B and C will not change for all sensible solar data, and thus the negative square root should always be chosen when the coefficients are defined as above.

The cloudy scale factor K_C can now be calculated using eq. 2.33. For the RAL data, K_C is 0.2729 when K_S is 0.7782, as quoted above in eqs. 2.21 and 2.22.

These values have been used to create ideal sunny and cloudy clearness index models, as shown in fig. 2.58.

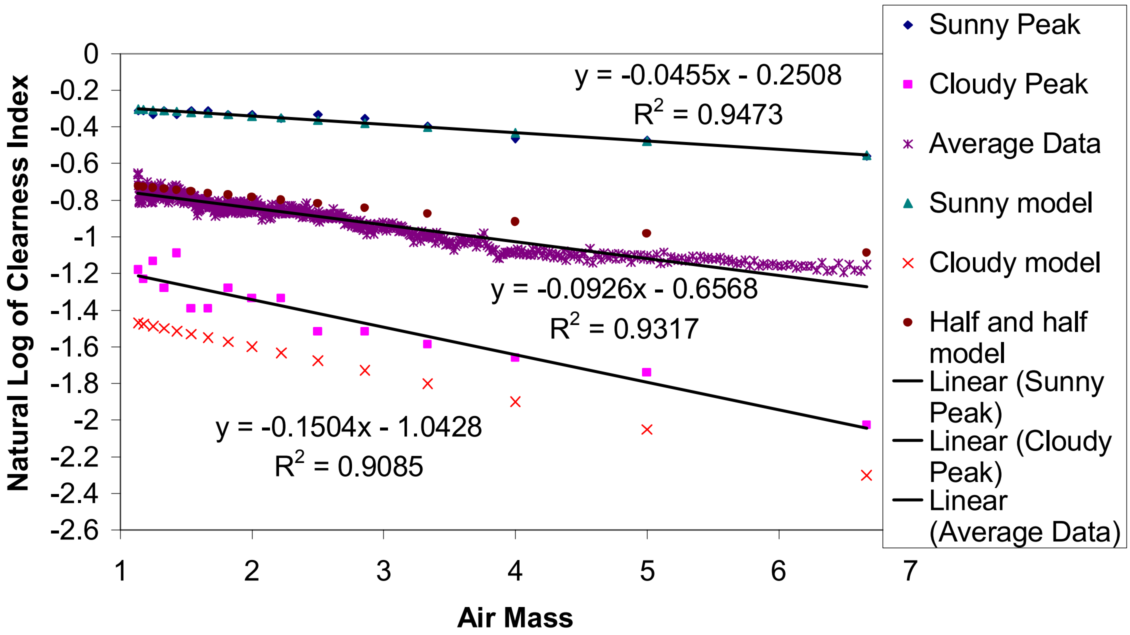

Figure 2.58 Clearness index peak positions as a function of air mass, measured and modelled

The sunny model is very close to the locus of measured peaks of sunny radiation. However, the cloudy model is at significantly lower values of clearness index than its measured peaks. This is because the probabilistic method assumes a purely bivalent behaviour, that is a binomial distribution with n=1, also called a Bernoulli PDF (Hastings, N. A. J., Peacock 1975). In order to preserve total variance of solar radiation, and account for the variation of clearness within the 'sunny' peak and within the 'cloudy' peak, the model values must be further apart than the measured peaks. This is illustrated in fig. 2.59 for the bin of data centred on a solar elevation of 36.87° (sine elevation = 0.6):

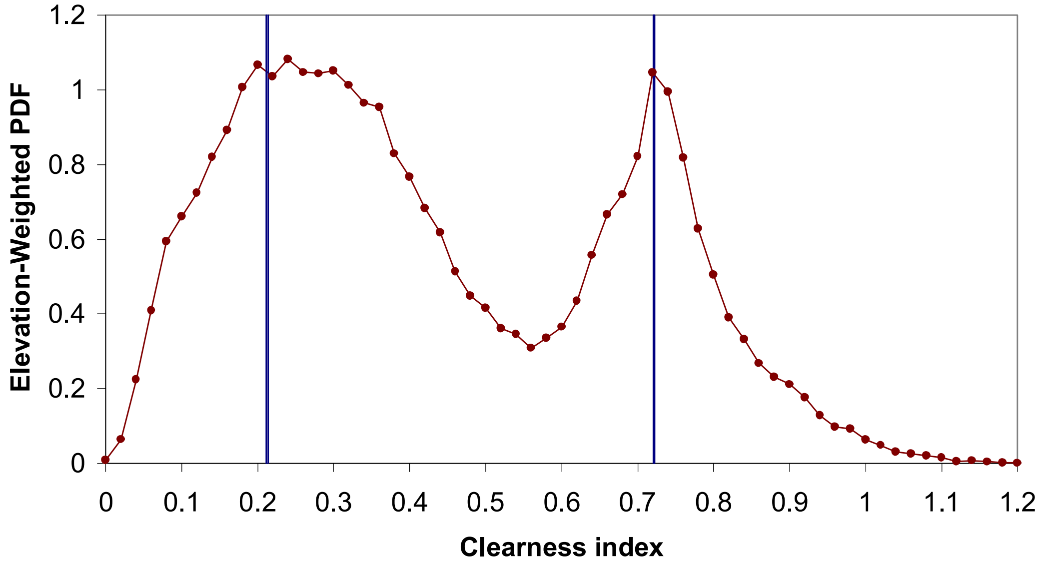

Figure 2.59 Model values of clearness index compared to the actual measured distribution at a solar elevation of 36.87°

2.7.7 Calculation of Monthly Average Sunny Fractions

Now we need to translate the known monthly average solar irradiances, I_i into monthly average sunny fractions, f_i using the long-term, annual average sunny fraction, $\overline{f}$ together with the sunny and cloudy scale factors, K_S and K_C.

The monthly average irradiance is given by an equation similar to eq. 2.24 above, as the monthly average sunny fraction does not change with hour of the day:

$$I_i = I_0 f_i K_S \frac{1}{24}\sum_{k=1}^{24}\left[\frac{\exp\left(-\alpha_S M_{ik}\right)}{M_{ik}}\right] + I_0\left(1-f_i\right)K_C \frac{1}{24}\sum_{k=1}^{24}\left[\frac{\exp\left(-\alpha_C M_{ik}\right)}{M_{ik}}\right] \quad (2.39)$$

Rearranging this, we get an expression for the monthly sunny fraction:

$$I_i - I_0 \frac{K_C}{24}\sum_{k=1}^{24}\left[\frac{\exp\left(-\alpha_C M_{ik}\right)}{M_{ik}}\right] = I_0 f_i \frac{K_S}{24}\sum_{k=1}^{24}\left[\frac{\exp\left(-\alpha_S M_{ik}\right)}{M_{ik}}\right] - I_0 f_i \frac{K_C}{24}\sum_{k=1}^{24}\left[\frac{\exp\left(-\alpha_C M_{ik}\right)}{M_{ik}}\right]$$

$$f_i = \frac{\frac{I_i}{I_0} - \frac{K_C}{24}\sum_{k=1}^{24}\left[\frac{\exp(-\alpha_C M_{ik})}{M_{ik}}\right]}{\frac{K_S}{24}\sum_{k=1}^{24}\left[\frac{\exp(-\alpha_S M_{ik})}{M_{ik}}\right] - \frac{K_C}{24}\sum_{k=1}^{24}\left[\frac{\exp(-\alpha_C M_{ik})}{M_{ik}}\right]} \tag{2.40}$$

2.7.8 Variations of Clearness Index Within a Period Vs. Period Average Clearness Index

Referring to the solar variation spectrum, fig. 2.10, solar power varies on all time scales. For a given time scale, e.g. one day, some of this variation occurs within each period, i.e. changing light levels within a day. The remaining variation occurs between days, i.e. light days and dark days.

It is not immediately obvious how the variation within a period depends on the period average irradiance. By using only the power spectrum, the temporal information of the original irradiance data has been lost. Assumptions have to be made that fit the original measured data and reflect the physical reasons for the variation:

In the case of wind speed data, the short-term wind speed variations are roughly proportional to the period average wind speed, section 2.6.5.

In the case of solar irradiance, short-term variation is due to changing cloud cover and cloud thickness, as well as the changing position of the sun in the sky. The sun's position can be directly calculated from time-of-day, day of year and sun-earth geometry but the effect of cloud cover and cloud thickness has to be inferred from variations in clearness index.

2.7.9 Measured Variations in Clearness Average Over Various Time Scales

The following series of plots show how the short-term variation (as variance) in sunny fraction depends on the average sunny fraction over the period in question. The data used is that measured at RAL from 1999 to 2002, and is the same data from which the power spectrum, fig. 2.10 and fig. 2.11 was produced. Measurements taken when the sun was low in the sky (sine of elevation angle less than 0.1, or air mass greater than 10) were ignored and removed from the data set.

The sunny fraction in each minute was calculated using the equation below, which is a rearrangement of eq. 2.23, applied to any instant of time:

$$f = \frac{\dfrac{I}{I_0} - K_C \dfrac{\exp(-\alpha_C M)}{M}}{K_S \dfrac{\exp(-\alpha_S M)}{M} - K_C \dfrac{\exp(-\alpha_C M)}{M}} \qquad (2.41)$$

2.7.9.1 Variations Within Ten Minutes

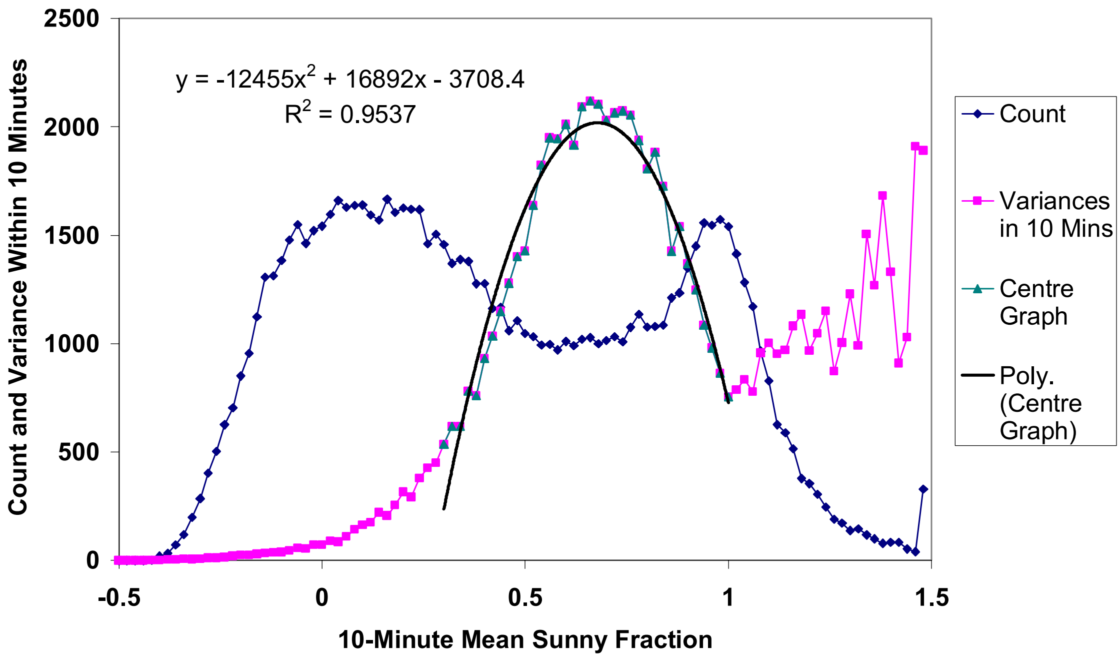

Figure 2.60 One-minute samples of sunny fraction within ten-minute periods

Fig. 2.60 shows the calculated variance in sunny fraction (scaled) within a 10-minute period vs. the statistical count of each 10-minute average. This shows that where the probability density is high, the variation within the period follows a smooth curve, and that for the central portion this curve is approximately quadratic, as shown. The formula and regression coefficient are quoted on fig. 2.60.

The highlighted central portion with a sunny fraction between 0.3 and 1.0 contains 46.0% of the data. The lower end contains 43.0% and the upper end, with very erratic data contains just 11.0%. These very high and very variable sunny fractions are probably due to cloud-side reflection (Durisch, Bulgheroni 1999, Laird, Harshvardhan 1997). The clearness index is significantly higher than in a completely cloudless sky. Similar effects are visible in the hourly and four-hourly data below.

2.7.9.2 Variations Within One Hour

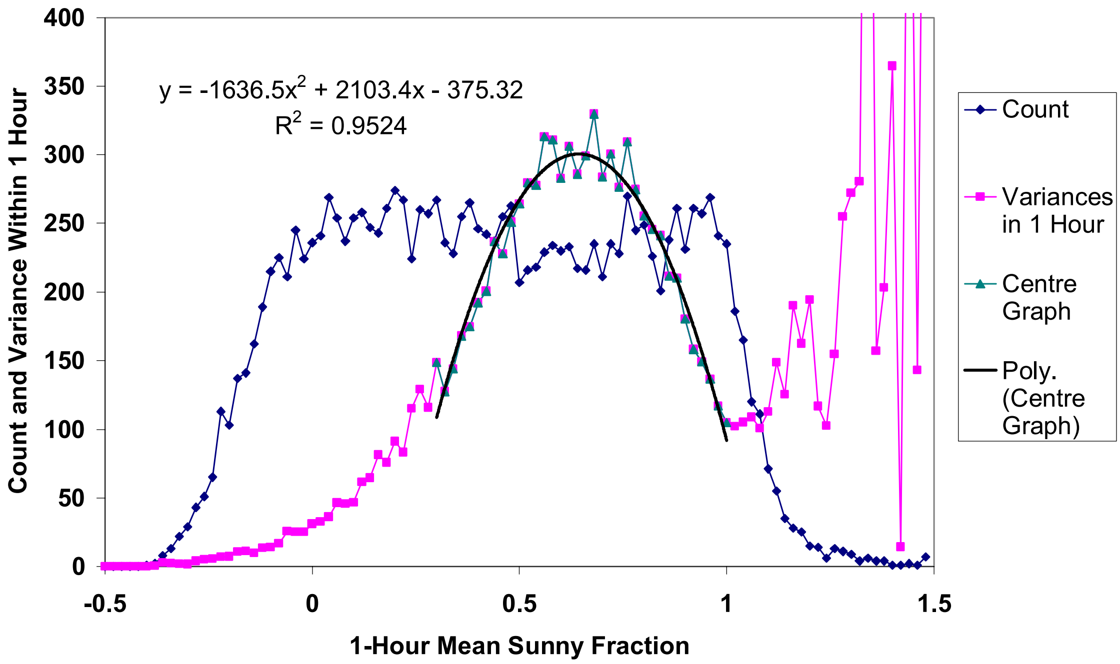

Figure 2.61 One-minute samples of sunny fraction within one-hour periods

Fig. 2.61 shows the calculated variance in sunny fraction (scaled) within a 1-hour period vs. the statistical count of each 1-hour average. Again, where the probability density is high, the variation within the period follows a smooth curve, and for the central portion this curve is approximately quadratic, as shown. The formula and regression coefficient are quoted on fig. 2.61.

The highlighted central portion with a sunny fraction between 0.3 and 1.0 contains 55.5% of the data. The lower end contains 38.7% and the upper end, with erratic data contains just 5.8%.

2.7.9.3 Variations Within Four Hours

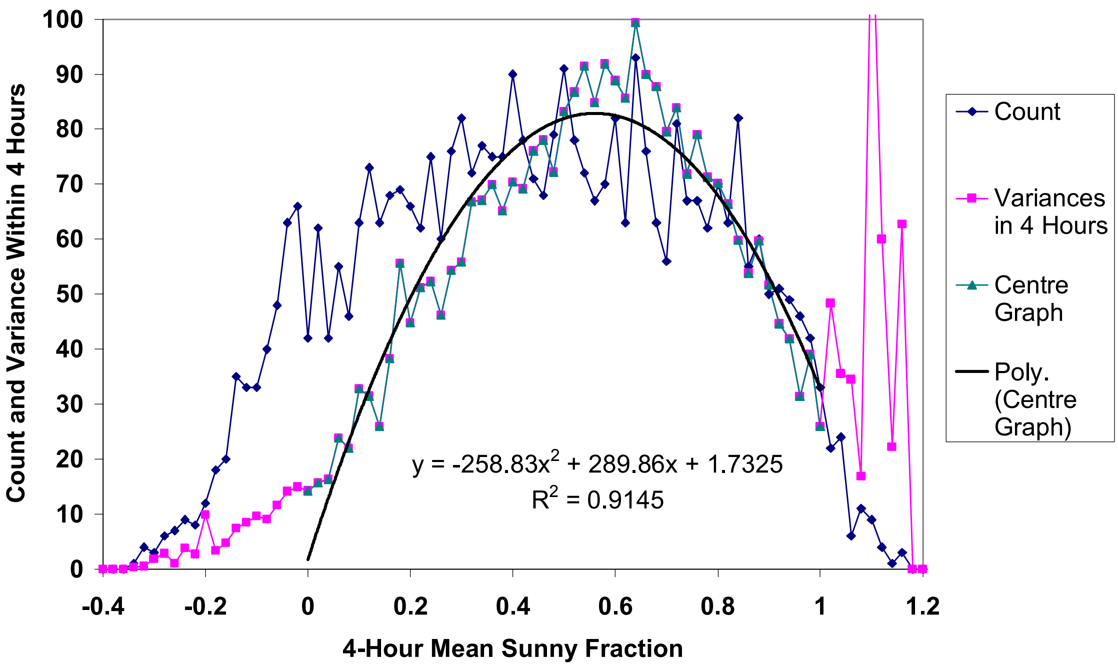

Figure 2.62 One-minute samples of sunny fraction within four-hour periods

Fig. 2.62 shows the calculated variance in sunny fraction (scaled) within a 4-hour period vs. the statistical count of each 4-hour average. Again, where the probability density is high, the variation within the period follows a smooth curve, and for the central portion this curve is approximately quadratic, as shown. The formula and regression coefficient are quoted on fig. 2.62.

The highlighted central portion with a sunny fraction between 0 and 1.0 contains 87.4% of the data. The lower end contains 10.5% and the upper end, with erratic data contains just 2.1%.

2.7.9.4 Variations Within One Day

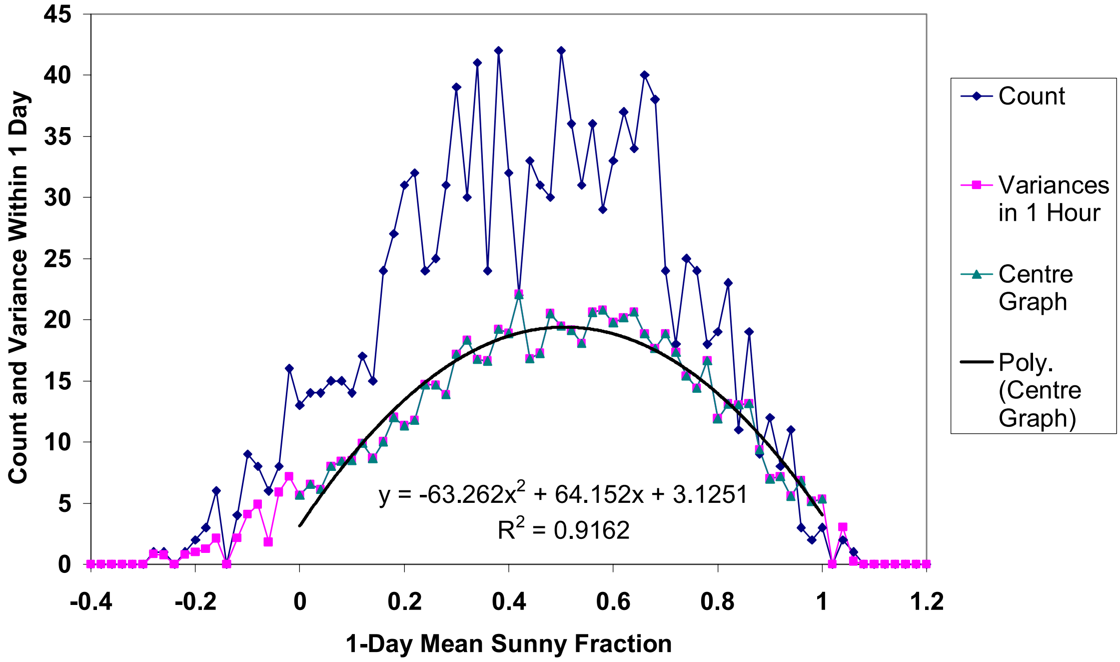

Figure 2.63 One-minute samples of sunny fraction within one-day periods

Fig. 2.63 shows the calculated variance in sunny fraction (scaled) within a 12-hour period vs. the statistical count of each 12-hour average. Again, where the probability density is high, the variation within the period follows a smooth curve, and for the central portion this curve is approximately quadratic, as shown. The formula and regression coefficient are quoted on fig. 2.63.

The highlighted central portion with a sunny fraction between 0 and 1.0 contains 94.7% of the data. The lower end contains 5.1% and the upper end, with erratic data contains just 0.2%.

The averaging period was actually 720 minutes or 12 hours since the average day length is 12 hours.

2.7.9.5 Conclusion of the Study of Variations Within Periods

At all time scales, the short-term variation in sunny fraction is greatest when the period-average clearness index is in the middle of its range at about 0.5. The short-term variation is lower at the low and high ends of the range. This qualitatively confirms the results of (Graham, Hollands 1990) except that Graham and Hollands have fitted a sine curve rather than a quadratic, and have plotted standard deviation vs. atmospheric transmittance rather than variance vs. sunny fraction.

The central quadratic portion of the data (the portion that could reasonably fitted by a quadratic equation) is in the sunny fraction range from 0 or 0.3 to 1.0. These values correspond well with the positions of the peaks in the overall PDF of clearness index, fig. 2.55, these peaks corresponding to cloudy and sunny conditions respectively. This adds further evidence that when a period (10-minutes, hour, day etc.) is completely cloudy or completely sunny, the variation within the period is a minimum, and that most within-period variations are due to cloud cover changes between cloudy and sunny.

2.7.10 Variations in Clearness Index with Changing Solar Elevation

The above plots group together data from all solar elevations, but the average clearness index (not sunny fraction) depends on solar elevation. The following plots, fig. 2.64 to fig. 2.67 group the data by solar elevation, but only for a time scale of one day. Each plot contains data from a group of four months, grouped by average solar elevation. Thus May, June and July are the sunniest months, April, August and September are the second sunniest months, then February, March and October. Finally, the darkest months of the year are January, November and December. The solar elevation is described by the sine of the elevation angle, from 0.2 to 0.9. Obviously, in May, June and July, the elevation can range from 0.2 to 0.9, but in January, November and December, the highest elevation is only 0.4.

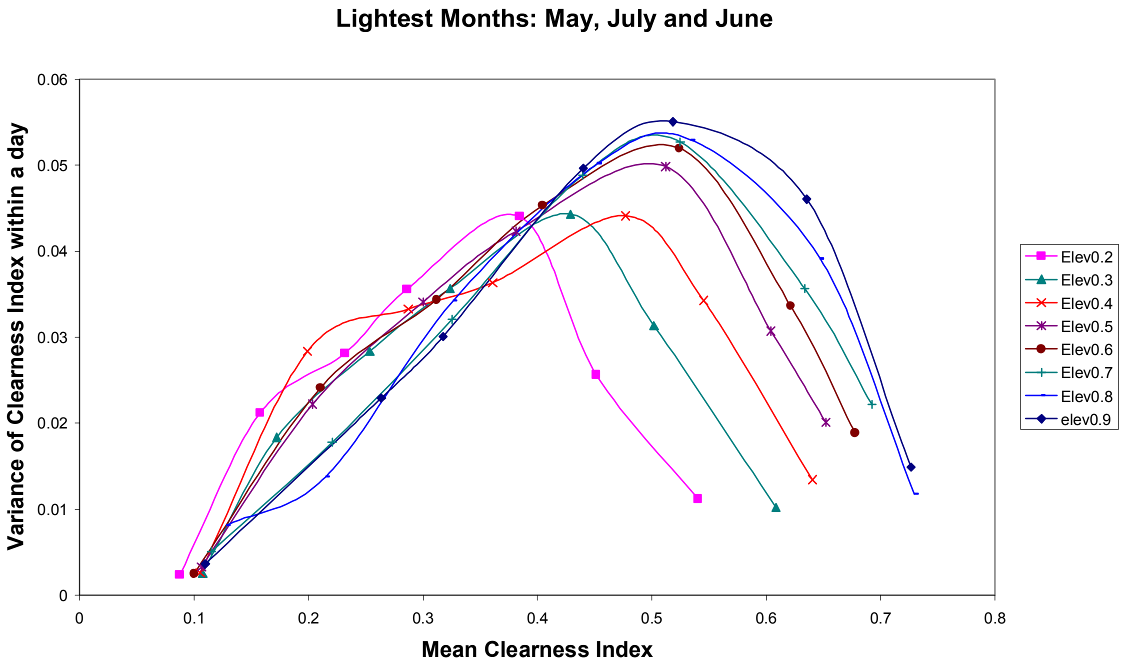

Figure 2.64 Variances of one-minute samples of clearness index within one-day periods, grouped by solar elevation and daily average clearness index in, May June and July

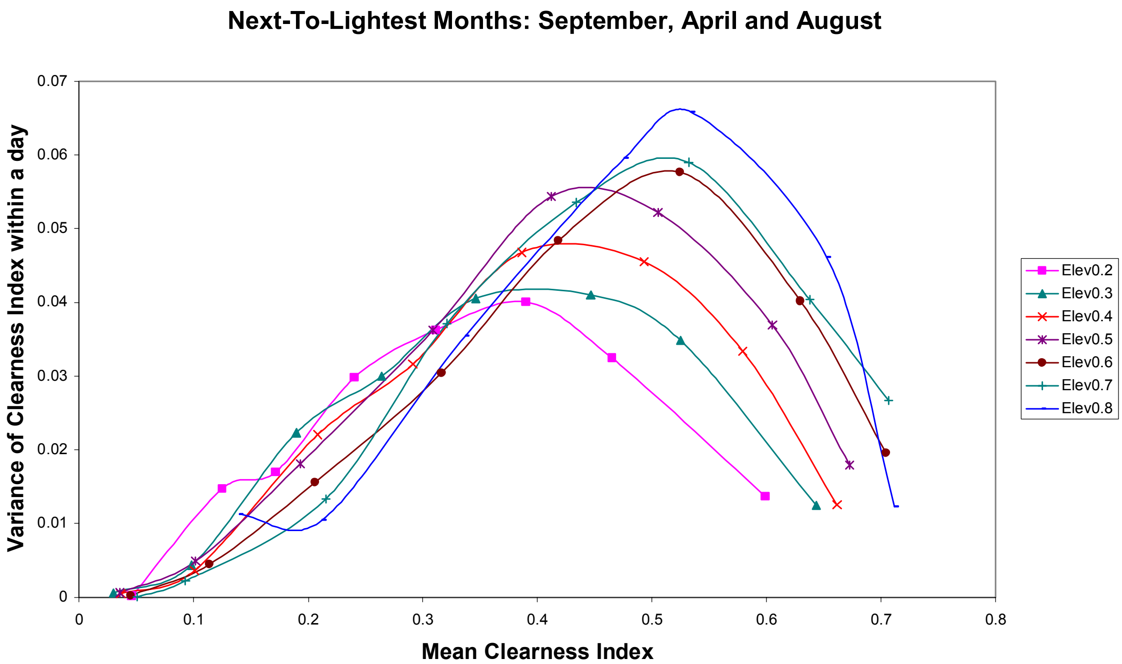

Figure 2.65 Variances of one-minute samples of clearness index within one-day periods, grouped by solar elevation and daily average clearness index in September, April and August

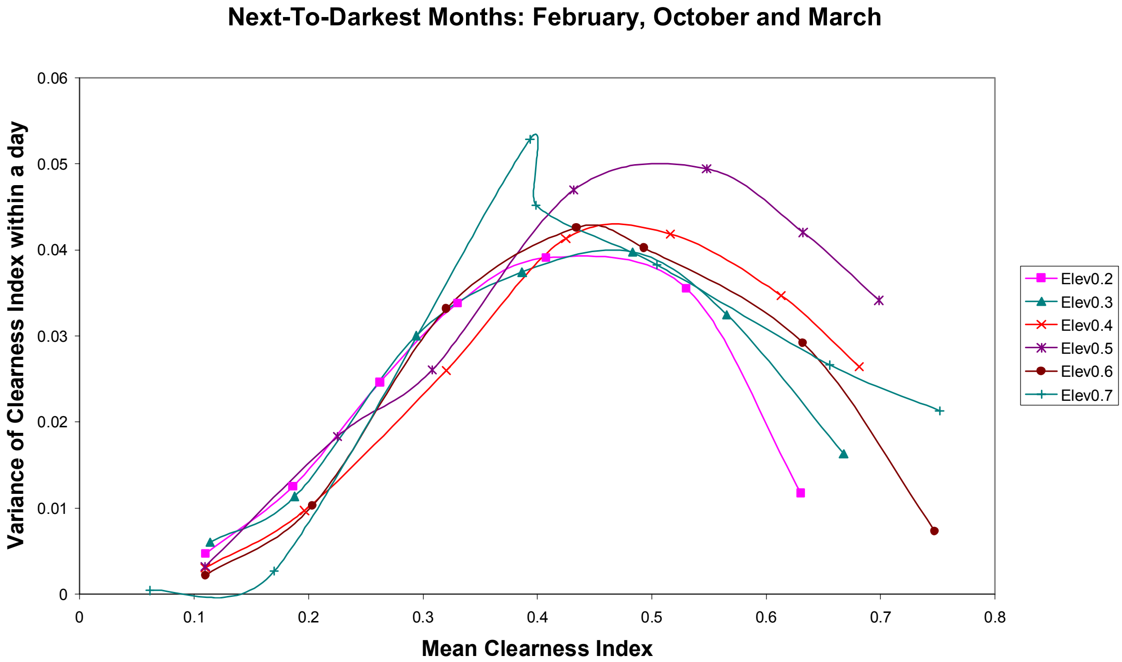

Figure 2.66 Variances of one-minute samples of clearness index within one-day periods, grouped by solar elevation and daily average clearness index in February, October and March

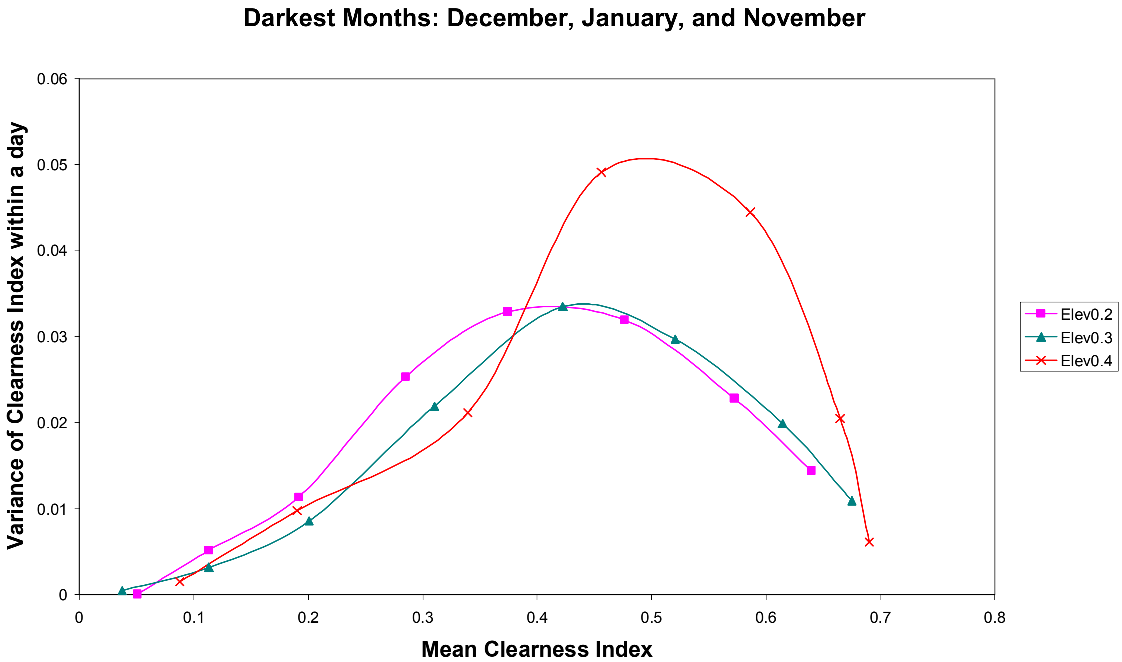

Figure 2.67 Variances of one-minute samples of clearness index within one-day periods, grouped by solar elevation and daily average clearness index in December, January and November

This second group of plots all confirm that the short-term standard deviation of clearness index is a maximum when the daily average is in the middle of the range, and lower at the bottom and top ends of the range. These graphs are also approximately quadratic in shape. A quadratic fit is important to the probabilistic method, as will be seen in section 2.7.11 below.

Increasing solar elevation simply shifts the graph to the right, and reducing solar elevation shifts the graph to the left. This is consistent with fig. 2.58. Increasing air mass increases the attenuation of solar radiation and reduces the average clearness index.

2.7.11 Calculation of Total Variance Due to Changing Cloud Cover

All the stochastic solar variance, represented by the area under the broadband spectrum excluding the periodic spikes, fig. 2.11 has to be due to a changing cloud cover (changing sunny fraction). All other terms are dependent only on solar elevation, which is periodic. Looking at all time scales from minutes up to months, this total variance is given by a summation over all months, i and all hours of the day, k. The variance in irradiance due to cloud cover in each hour of the day, in each month of the year is given by:

$$V_{ik} = f_i \left[I_0 K_S \frac{\exp(-\alpha_S M_{ik})}{M_{ik}} \right]^2 + (1 - f_i) \left[I_0 K_C \frac{\exp(-\alpha_C M_{ik})}{M_{ik}} \right]^2 - I_{ik}^2 \qquad (2.42)$$

And the average radiation in that hour in that month is:

$$I_{ik} = f_i \left[I_0 K_S \frac{\exp(-\alpha_S M_{ik})}{M_{ik}} \right] + (1 - f_i) \left[I_0 K_C \frac{\exp(-\alpha_C M_{ik})}{M_{ik}} \right] \qquad (2.43)$$

To simplify the equations, let us use working variables:

Sunny clearness index, $S_i = I_0 K_S \dfrac{\exp(-\alpha_S M_{ik})}{M_{ik}}$ (2.44)

And cloudy clearness index, $T_i = I_0 K_C \dfrac{\exp(-\alpha_C M_{ik})}{M_{ik}}$ (2.45)

So eq. 2.42 becomes: $V_{ik} = f_i S_{ik}{}^2 + (1 - f_i) T_{ik}{}^2 - I_{ik}^2$ (2.46)

And eq. 2.43 becomes: $I_{ik} = f_i S_{ik} + (1 - f_i) T_{ik}$ (2.47)

Substituting for I_{ik} in eq. 2.46:

$$V_{ik} = f_i S_{ik}{}^2 + (1 - f_i) T_{ik}{}^2 - \left[f_i S_{ik} + (1 - f_i) T_{ik} \right]^2$$

$$= f_i S_{ik}{}^2 + (1 - f_i) T_{ik}{}^2 - f_i^2 S_{ik}{}^2 - 2 f_i (1 - f_i) S_{ik} T_{ik} - (1 - f_i)^2 T_{ik}{}^2$$

$$= f_i S_{ik}{}^2 + T_{ik}{}^2 - f_i T_{ik}{}^2 - f_i^2 S_{ik}{}^2 - 2 f_i S_{ik} T_{ik} + 2 f_i^2 S_{ik} T_{ik} - T_{ik}{}^2 + 2 f_i T_{ik}{}^2 - f_i^2 T_{ik}{}^2$$

Cancelling terms and then factorising:

$$V_{ik} = f_i S_{ik}{}^2 - f_i^2 S_{ik}{}^2 - 2 f_i S_{ik} T_{ik} + 2 f_i^2 S_{ik} T_{ik} + f_i T_{ik}{}^2 - f_i^2 T_{ik}{}^2$$

$$V_{ik} = f_i (1 - f_i)(S_{ik} - T_{ik})^2 \qquad (2.48)$$

Now S_{ik} and T_{ik} depend only on the solar elevation, which is determined by the hour and the month.

Note that this bivalent model fits very well, at least qualitatively with the graphs of sections 2.7.9 and 2.7.10. Eq. 2.48 is a quadratic dependence of the monthly variance, V_{ik} on the monthly average sunny fraction, f_i.

The dependence of S_{ik} and T_{ik} on solar elevation also shifts the graphs to the right with increasing solar elevation, just as the graphs of section 2.7.10 show. The monthly averages of stochastic variances are then:

$$V_i = f_i\left(1-f_i\right)\frac{1}{24}\sum_{k=1}^{24}\left(S_{ik}-T_{ik}\right)^2 \tag{2.49}$$

And the annual average stochastic variance is approximately:

$$V_{Stochastic\_Simple} = \frac{1}{12}\sum_{i=1}^{12} f_i\left(1-f_i\right)\frac{1}{24}\sum_{k=1}^{24}\left(S_{ik}-T_{ik}\right)^2 \tag{2.50}$$

Using the model developed here for RAL data, the above quantity was calculated to be 13666 $(W/m^2)^2$.

The actual formula used in the computer program is adjusted slightly for the different lengths of each month. As will be seen later, real calendar months were used because a real calendar was needed for the load data, to give the correct balance of holidays and working days. The accurate formula as used by 24-hour probabilistic program is:

$$V_{Stochastic} = \sum_{i=1}^{12}\frac{Days_i}{365} f_i\left(1-f_i\right)\frac{1}{24}\sum_{k=1}^{24}\left(S_{ik}-T_{ik}\right)^2 \tag{2.51}$$

Where: $Days_i$ is the number of days in month i.

Using the model developed here for RAL data, and the real month lengths, the stochastic solar variance was calculated to be 13716 $(W/m^2)^2$.

These quantities compare fairly well with the total stochastic variance calculated from integrating the broad solar power spectrum, with the spikes removed, see section 2.4.6. This spectrum integration was calculated to be 12661 $(W/m^2)^2$. The spectrum integration is not precise, since the position of the bases of the daily spikes are only an estimate. It is therefore not clear how much variance belongs to the daily periodic variation (and its harmonics) and how much to the broad spectrum of variation.

2.7.12 Model of Cloudy Periods, Sunny Periods and Mixtures

The simplified model used within the probabilistic method assumes that at a given solar elevation, all variation in clearness index is due to the presence or absence of clouds, and that at any instant of time the sky is either cloudy (no beam radiation) or sunny (full beam radiation and no clouds). Furthermore, periods (e.g. 10-minutes, hours, days or weeks) can be characterised according to the percentage of time that the weather is cloudy or sunny. Henceforth, for description purposes the assumed period of time will be one day, as used in the 24-hour Matlab probabilistic program, section 2.13.

On completely cloudy days, the weather is uniformly cloudy and the clearness index follows the lower line, the cloudy model in fig. 2.58, and eq. 2.22. No account is taken of variation in cloud thickness. Cloudy days therefore have constant sunny fraction (equal to 0) and zero variation in irradiance due to passing clouds. The total variation in clearness index will be low, and will result only from changes in solar elevation and therefore changes in air mass.

On completely sunny days, the weather is uniformly clear and sunny with a clearness that follows the upper line, the sunny model in fig. 2.58, and eq. 2.21. Sunny days therefore also have constant sunny fraction (but this time equal to 1.0) and zero variation in irradiance due to passing clouds. The total variation in clearness index will again be low, resulting only from changes in solar elevation and therefore changes in air mass.

For days that are a mixture of sunny and cloudy conditions, the probabilistic method requires the stochastic cloud-cover variance to be split into two components: short-term (variations within a day) and long-term (variation in daily averages). This is done using the filter functions of sections 2.5.2 and 2.5.3. For the solar spectrum, this is illustrated in fig. 2.68. The left-hand portion of the spectrum determines the stochastic variation of daily averages (the ‘Daily Average’ line). The integrated area under this curve gives a variance of 2000.3 $(W/m^2)^2$. The right-hand, short-term portion of the spectrum (the ‘Within Day’ line) determines the stochastic variation of solar irradiance within a day. The integrated area under this second curve gives a variance of 10661 $(W/m^2)^2$. The sum of the short-term variance and the long-term variance is 12661, as previously calculated by the integration of the whole spectrum, after removing the spikes.

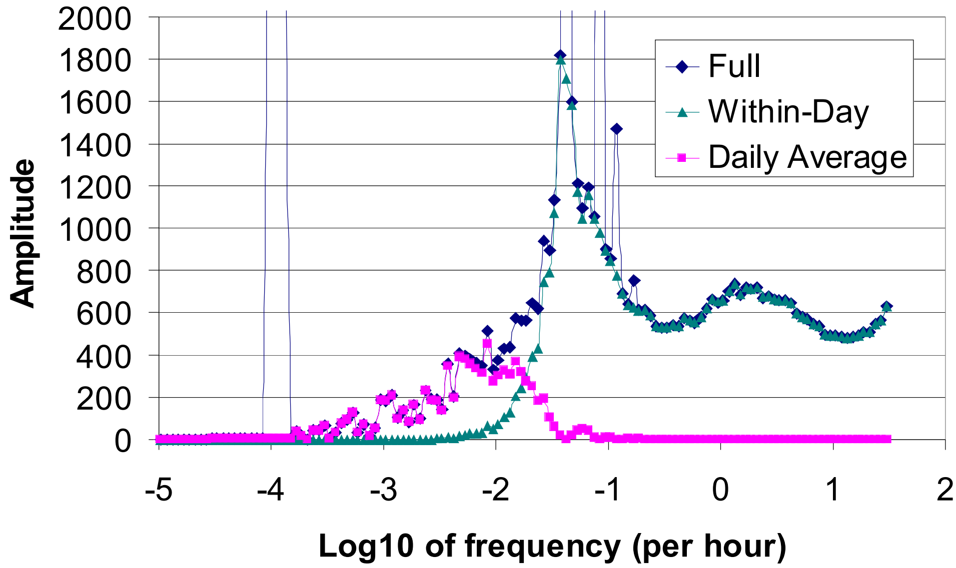

Figure 2.68 Filter functions applied to the solar variation spectrum as measured at Rutherford Appleton Laboratory, 1999 to 2002

Note that these calculations of stochastic variance are based on the solar spectrum after the spikes have been removed. The spikes represent periodic variation in solar elevation and periodic variation in cloud cover. Their removal leaves only the random variation due to changing cloud cover.

The integral of the left-hand, long-term portion of the spectrum = 2000.3$(W/m^2)^2$ has been used to construct probability distributions of daily average sunny fractions, as described in section 2.7.13.2.

2.7.13 Beta Distributions for Daily Average Sunny Fractions

The probabilistic method assumes that within a given month, i, the daily average sunny fractions, f_{ij} are distributed according to a beta distribution.

2.7.13.1 General Characteristics of Beta Distributions

Beta distributions have often been used for modelling probability distributions of solar irradiance and clearness index: (Graham, Hollands 1990, Mefti, Bouroubi & Adane 2003, Rahman, Khallat & Salameh 1988, Sahin, Sen 1998, Sulaiman et al. 1999, Youcef Ettoumi et al. 2002). The properties of beta distributions: formulas for their mean, variance and mode are quoted from (Patel 1976)

One advantage of beta distributions is that they have lower and upper bounded values, and can have a wide variety of shapes depending on the mean and standard deviation. Since the sunny fractions, f_{ij} have a minimum value of 0 and a maximum value of 1, we can use the normalised form of the beta distribution.

Another advantage is that the mean and variance (hence standard deviation) of a distribution are all easy to calculate from the two parameters, a and b. The beta function probability distribution function is:

$$pdf(x) = \frac{x^{a-1}(1-x)^{b-1}}{\mathrm{B}(a,b)} \qquad (2.52)$$

The mean is: $$m = \frac{a}{a+b} \qquad (2.53)$$

And the variance is: $$v = \frac{ab}{(a+b+1)(a+b)^2} \qquad (2.54)$$

What is more, if we know the mean and variance, we can calculate the two parameters, a and b, and so construct a beta distribution with any desired mean and standard deviation. Simple algebraic manipulation yields the following equations:

Rearranging eq. 2.53: $$b = \frac{a(1-m)}{m} \qquad (2.55)$$

And substituting into (2.54): $$v = \frac{\frac{a^2(1-m)}{m}}{\left(\frac{a+m}{m}\right)\left(\frac{a}{m}\right)^2} = \frac{(1-m)m^2}{a+m}$$

So: $$a = \frac{(1-m)m^2}{v} - m \qquad (2.56)$$

The value of b can then be calculated using eq. 2.55.

2.7.13.2 Beta Distributions of Sunny Fraction for Real Solar Data (RAL Data)

The above method has been used to construct beta distributions, fig. 2.69 of daily average sunny fraction, f_{ij} for each month, i of the year to represent the means and standard deviations of daily average sunny fractions at RAL. Note that these probability distributions only approximately reflect the shapes of probability distributions of sunny fraction or clearness index in those months. However, they do preserve the mean and standard deviation values of sunny fraction, and therefore clearness index in a simple but practical methodology. The monthly mean average values of daily sunny fraction are identical to the monthly overall mean sunny fractions, f_i and are calculated directly from the monthly averages of solar radiation. The standard deviations of daily sunny fraction within each month are derived from the spectrum integration, section 2.5.4, such that the total stochastic variance, $V_{Stochastic}$ is preserved.

The long-term (period average) portion of the variance, i.e. the variances in daily average sunny fraction, f_{ij} within each month i are a fixed proportion of the total variance of f within each month $= f_i(1-f_i)$. That proportion is simply the ratio of long-term (period average) stochastic variance to total stochastic variance. In the case of the RAL data model, this ratio is 2000.3 / 13716. This fixed proportion reflects the fact that the total solar power spectrum conveys no information as to which months may have a more variable sunny fraction and which months may have a more consistent sunny fraction.

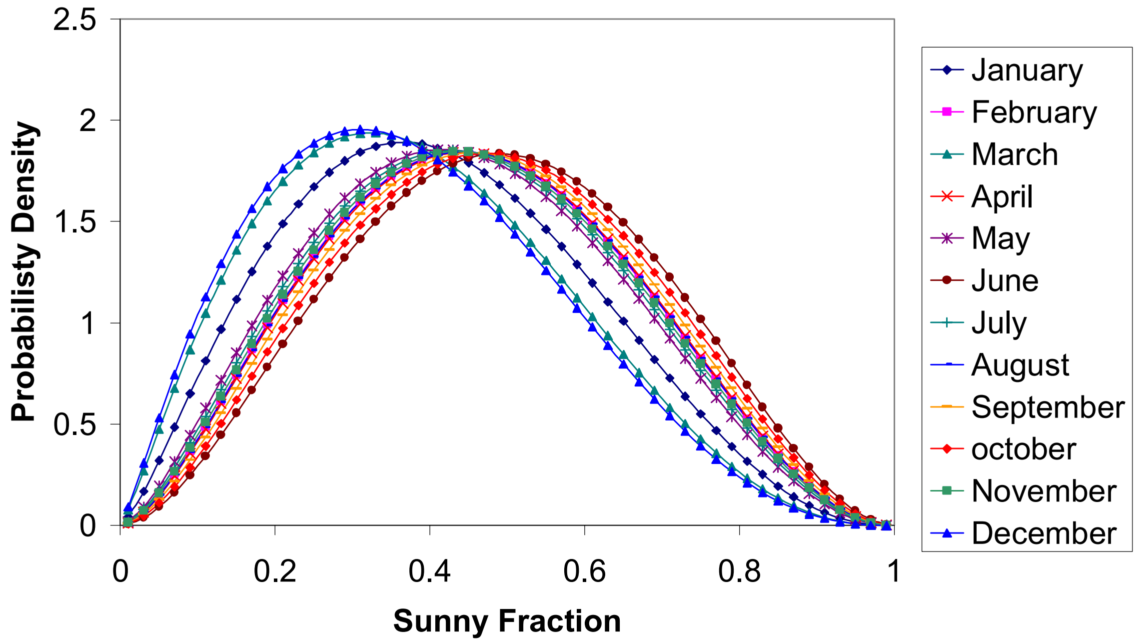

Figure 2.69 Beta distributions of daily sunny fraction for RAL constructed from 4 years of solar irradiance measurements, 1999 to 2002

Table 2.5 presents the same information, the parameters of the beta distributions, as a table. This confirms that the values of standard deviation in daily average sunny fraction in each month are all very consistent.

Table 2.5 Means, Standard Deviations, ‘a’s and ‘b’s of the Beta Distributions of Sunny Fraction for Each Month

Month	Mean Sunny Fraction	Standard Deviation of Daily Sunny Fractions	Beta distribution parameter ‘a’	Beta distribution parameter ‘b’
January	0.4122	0.1893	2.3763	3.3886
February	0.4617	0.1917	2.6616	3.1033
March	0.3846	0.1870	2.2171	3.5478
April	0.4640	0.1917	2.6750	3.0899
May	0.4457	0.1911	2.5692	3.1957
June	0.4927	0.1922	2.8403	2.9245
July	0.4526	0.1914	2.6092	3.1557
August	0.4624	0.1917	2.6657	3.0992
September	0.4720	0.1919	2.7211	3.0438
October	0.4814	0.1921	2.7750	2.9899
November	0.4579	0.1916	2.6399	3.1250
December	0.3761	0.1862	2.1681	3.5968

The values of monthly mean sunny fraction are seen to change with the month. The lowest value is 0.3761 for the month of December, while the highest value is 0.4927 for the month of June. These extremes suggest there is some seasonal dependence, in which months of low solar elevation have low sunny fraction and months of high solar elevation have high sunny fraction. However, February and July have almost average sunny fractions, March has the second lowest sunny fraction, and October has the second highest sunny fraction. Any seasonal dependence is therefore not very strong.

Months with low average sunny fraction have their peak shifted to the left on the graph of PDFs, fig. 2.69, corresponding to larger *b* and smaller *a*, but months of high average sunny fraction have their peak shifted to the right on the graph of PDFs, corresponding to larger *a* and smaller *b*. In practice, *b* is always slightly greater than *a* since the monthly mean sunny fractions are all less than 0.5.

2.7.14 Preservation of Total Stochastic Variance

The probabilistic method uses long-term beta distributions within each month to model the distribution of short-term bivalent distributions. To support the validity of this method, the following argument proves that the total stochastic variance due to changing cloud cover is preserved.

As seen in eq. 2.50 and eq. 2.51, we have expressions for the total stochastic variance in solar irradiance (due to changing cloud cover). These can be thought of as a variance of the sunny fraction, multiplied by a solar elevation-dependent factor. Let us consider a generic average, x (representing monthly average sunny fraction f_i or overall sunny fraction $\overline{f}$) in a bivalent distribution with values of 0 or 1. It is difficult to prove that total variance is preserved specifically for the case of a beta distribution, but it is possible to prove for the general case of a distribution of bivalent distributions:

Each short-term bivalent distribution has a mean value of x , a minimum of 0 and a maximum of 1. The probability of '1' is x and the probability of '0' is $1-x$. The variance within each bivalent distribution is then:

$$v = x \times 1^2 + (1-x) \times 0^2 - x^2 = x - x^2 = x(1-x) \qquad (2.57)$$

The value of x slowly varies, from day to day. Suppose the means of the bivalent distributions have a long-term probability density function (PDF), $p(x)$, such that $\int_{x=0}^{1} p(x)dx = 1$

The long-term mean value of x is the same as the long-term mean value of the long-term PDF: $\overline{x} = \int_{x=0}^{1} xp(x)dx$ (2.58)

The weighted total of the variances within each short-term bivalent distribution, using eq. 2.57 is:

$$v_1 = \int_{x=0}^{1} x(1-x)\,p(x)dx \tag{2.59}$$

And the long-term variance (i.e. the variance in short-term distribution means) is:

$$v_2 = \int_{x=0}^{1} x^2 p(x)dx - \left(\bar{x}\right)^2 \tag{2.60}$$

So the total variance is the sum of short-term and long-term variances: $v_{TOT} = v_1 + v_2 = \int_{x=0}^{1} x(1-x)\,p(x)dx + \int_{x=0}^{1} x^2 p(x)dx - \left(\bar{x}\right)^2$

$$= \int_{x=0}^{1} xp(x)\,dx - \int_{x=0}^{1} x^2 p(x)dx + \int_{x=0}^{1} x^2 p(x)dx - \left(\bar{x}\right)^2$$

$$= \int_{x=0}^{1} xp(x)\,dx - \left(\bar{x}\right)^2 \quad = \bar{x} - \left(\bar{x}\right)^2 \qquad = \bar{x}\left(1-\bar{x}\right) \qquad \text{QED} \tag{2.61}$$

This is the variance of the overall bivalent distribution of mean value $\bar{x}$.

2.7.15 Conversion of Solar Irradiances into Electrical Power

The probabilistic model tries to keep the modelling of solar power as simple as possible. Therefore, the electrical power output of a photovoltaic (PV) device is assumed to be proportional to the global horizontal solar irradiance.

In practice, PV power depends on a number of factors: the PV technology chosen, the tilt and orientation of the PV module, the electromagnetic spectrum of solar radiation, the temperature of the PV module, the intensity of solar irradiance, the power capacity of the inverter or other energy conversion device etc. Each factor affects or interacts with every other one. For example a higher irradiance may cause over-heating of the module. A change in temperature may improve or reduce the PV performance, depending on which technology is chosen. The tilt and orientation of the module affect the relative amounts of beam and diffuse radiation received by the module, and the time of day and year that most radiation is received. If the inverter is over-sized, its standing losses will be significant. However, if the inverter is under-sized, it will be over-loaded at times of high solar power, and solar energy will be lost.

To model all these effects accurately, we would need to model variations in ambient temperature and the heat balance of the solar module knowing its position on the roof and all heat transfer properties. We would need to know the percentage of radiation contained in beam and diffuse components, and construct the geometric interaction of the sun's position in the sky together with the orientation of the module. We would need an accurate model of the module performance and of the inverter. All these effects are the subjects of various academic papers and ongoing research. They are all complex and usually non-linear.

The biggest effects appear to be the tilt and orientation of the module and the temperature of the module. It is anticipated that in the next decade or two, most installed PV modules will be roof mounted and will be positioned flat to the roof. In the northern hemisphere, the modules will therefore be at various tilt angles and all orientations from east, through south-east, south and south-west to west. Most will be approximately facing towards the equator at some moderate tilt angle, tending to increase the solar energy falling on the module. However, most solar energy is received in summer and tends to make the module even hotter than the already elevated ambient temperature. In most PV technologies this tends to reduce the energy captured. As a simple approximation, it was assumed that these two effects will cancel each other.

The probabilistic model program created for this thesis is designed for a scale of at least a few hundred kW of power. This is the typical size of one feeder of an electricity substation, section 2.2.4, and is also the typical power capacity of one large, modern wind turbine. At this size, the system would typically supply electricity to many premises of various types, from domestic to commercial and small industrial. The buildings of those premises would have various shapes and sizes of roof at a variety of tilt and orientation angles. If solar PV modules are fitted to several different roofs, the effect of orientation and tilt angles will aggregate and cancel out, bringing us back to an assumption of electricity output proportional to global solar horizontal radiation.

Given that the whole probabilistic method is a fast and approximate one, appropriate to feasibility studies, it was decided to simply assume that power captured is;

$$Power(kW_e) = \frac{Irradiance\left(W/m^2\right)}{1000} \times PowerRating(kWp) \tag{2.62}$$

2.7.16 How Well Does the Model Match the Real Global Horizontal Solar Irradiance Data?

Fig. 2.70 compares the actual and model-predicted annual probability distribution of global horizontal solar irradiance. The actual distribution was built from measurements of one-minute averages.

The theoretical distributions were built from hourly values calculated using the probabilistic method. That is the average sunny radiation, the average cloudy radiation and the average monthly sunny fraction. Although the hourly distribution of solar irradiance is purely bivalent (two values only in a binomial PDF with n=1), when these hourly distributions are added together for all 24 hours of the day and all 12 months, they produce an almost smooth probability distribution:

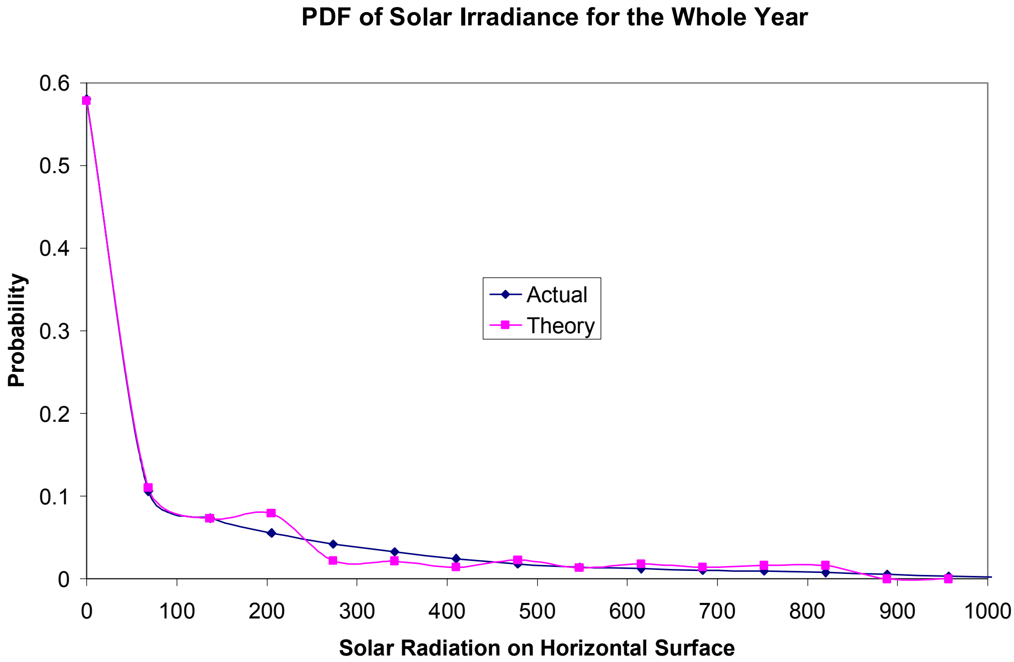

Figure 2.70 Probability distribution of solar power for the whole year, measured at RAL and modelled

Fig. 2.71 below shows that the mean and standard deviations of solar irradiance within each month are modelled very well

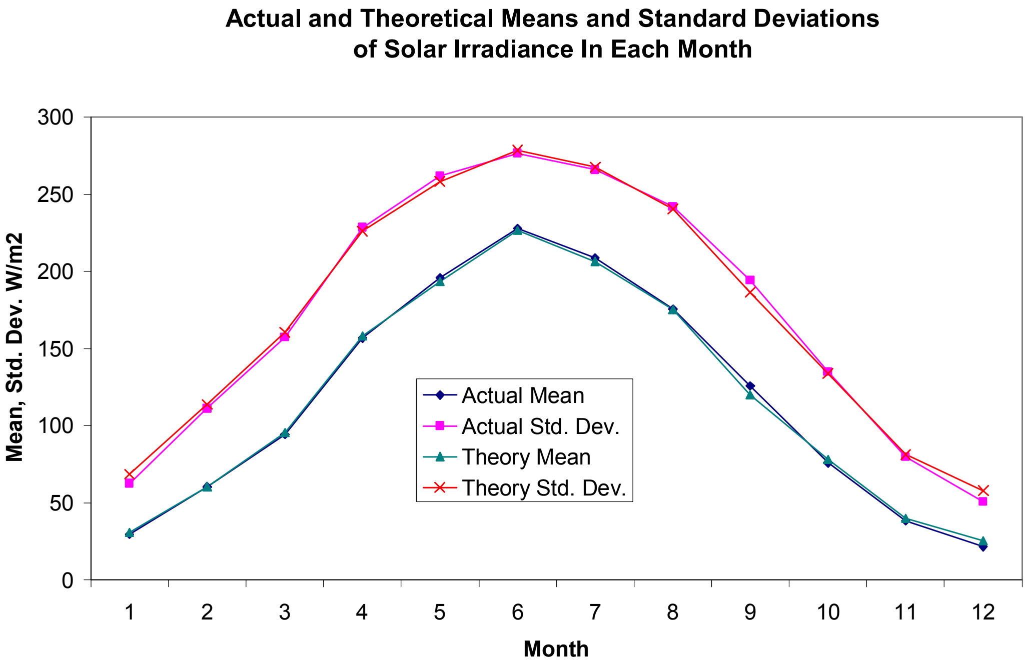

Figure 2.71 Means and standard deviations of solar power through the year, actual and modelled

2.7.17 For Further Work - Improved Predictions of Solar Energy Capture

There are various published models of solar irradiance that can convert a global solar radiation into a radiation on an inclined surface. In future refinements of the probabilistic or time-step models presented in this thesis, it may be possible and desirable to include a calculation of irradiation on an inclined plane.

Many papers present methods of estimating the diffuse fraction or beam components from the global horizontal clearness index, e.g. (Babatunde, Aro 1995, Gonzalez, Calbo 1999, Ideriah 1992, Lam, Li 1996, Lopez, Rubio & Batles 2000, Suehrcke, McCormick 1989, Tiris, Tiris 1998, Trabea 2000, Unozawa, Otani & Kurokawa 2001). Thus the solar radiation could be separated out into beam and diffuse components. The radiation falling on an inclined surface could be better estimated by treating these two components separately.

Alternatively, at least one model estimates the radiation on an inclined surface directly from the horizontal radiation, (Olmo et al. 1999)

2.8 Load Model

2.8.1 Introduction to the Load Model

This section describes how the electrical load model has been constructed for the probabilistic method. One of the key assumptions of this probabilistic method is the size of the electrical power system. It is appropriate to a village, a region within a town, or larger. The system could accommodate the electricity produced by at least one large modern wind turbine.

Unlike the solar power and wind power models, the load is modelled as a deterministic function of time of day, type of day and month of the year. Sub-section 2.8.2 re-visits the reasons for this assumption, based on the magnitude of variation in electrical load and the proportion of variation that is deterministic and periodic.

Sub-section 2.8.3 describes the derivation of the load profiles and compares them with National Grid data.

Sub-section 2.8.4 describes the division of the model year into months and into weekdays and holidays.

2.8.2 Magnitudes of Variation of Electrical Load

Section 2.4.4 shows that as the number of consumers on an electricity network increases, the variation due to each individual consumer rapidly diminishes to a negligible level. If the number of consumers is of the order of 1000, and the standard deviation of variation in demand between consumers is about 50%, then by the central limit theorem predicts that the standard deviation in the mean of a sample of 400 consumers will be about $50/\sqrt{400} = 50/20$ = 2.5%. This is relatively small compared to the actual hour-to-hour standard deviation in total demand of 24% seen in one substation feeder that probably supplies about that number of consumers.

Section 2.4.6 shows that variations in electrical load at the substation feeder level are smaller in percentage terms than those of wind power or solar power (24% vs. 90% and 170% respectively) and that variations in electrical load are dominated by periodic effects: daily, weekly and seasonal cycles. The periodic component of standard deviation in electrical load on the chosen substation feeder is 23% of the long-term average whereas the stochastic component is only 7.6%.

Because the stochastic minute-to-minute aggregated variations in demand are small compared to the periodic variations, and compared to either solar or wind power variations, these stochastic variations were considered negligible for the purposes of the probabilistic method. Only the periodic variations in demand were included in the probabilistic method.

2.8.3 Load Profiles

As stated in section 2.2.4, substation feeder Braunstone LOC-A appears to have a load profile close to the national average, as published in (National Grid Company plc 2003). The load data is recorded as half-hour averages and the number and variety of electrical loads is such that the total load changes smoothly from one half-hour to the next. The loads from the LOC-A feeder have been compared with the national grid average profiles for the whole of England and Wales, (National Grid Company plc 2003), for both summer and winter typical days, fig. 2.72 and fig. 2.73. The Braunstone data shows slightly more variation between day and night and a slightly larger evening peak of load than the national average. The Braunstone sub-station also shows a larger difference between summer and winter than the national grid average, as would be expected due to more limited aggregation.

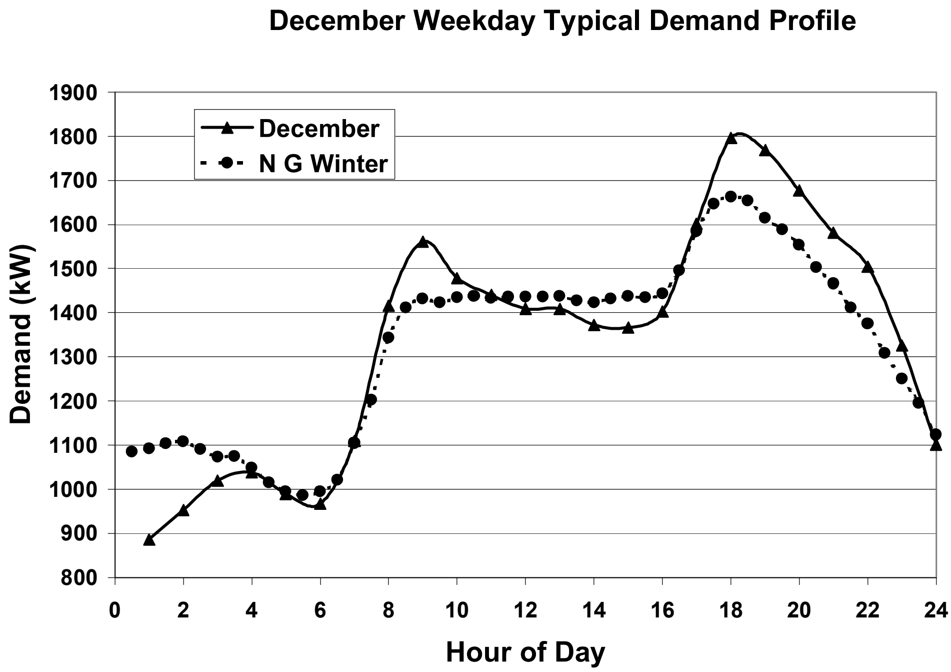

Figure 2.72 December weekday demand profiles

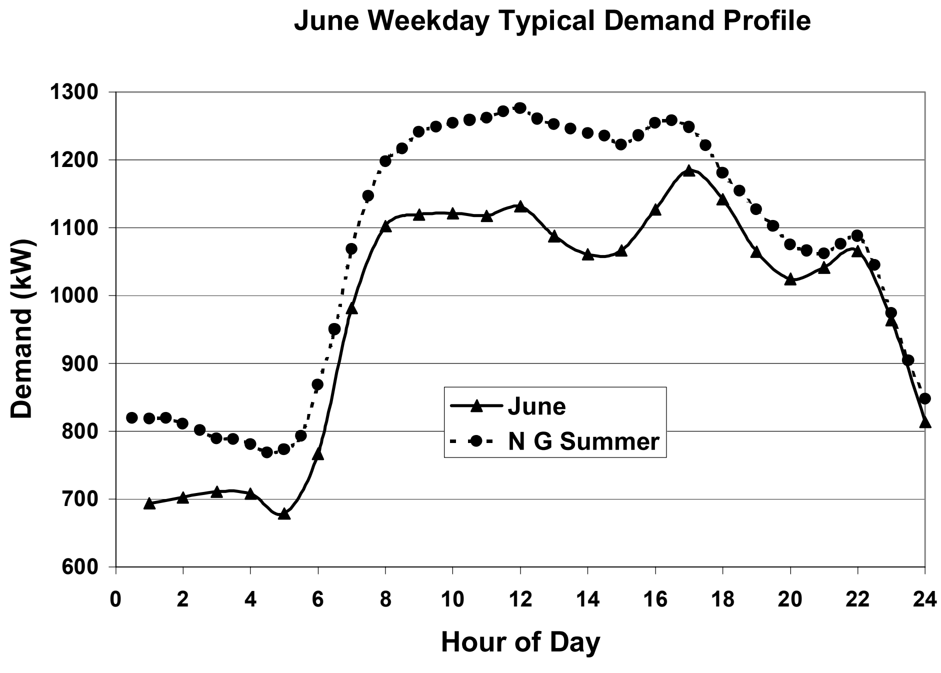

Figure 2.73 June weekday demand profiles

In order to minimise the data required, and to minimise computer run time, the probabilistic model uses hourly time steps of electricity demand. Each pair of half-hour averages have been added together to give the energy consumed in each hour period of the probabilistic model, starting with the first two. The load data is all in, or has been converted to Greenwich Mean Time (GMT). The load profiles have been scaled to an annual average of 1000kW before the probabilistic program user applies his or her chosen load scale factor.

2.8.4 The Model Calendar

The 24-hour probabilistic program uses load profiles calculated for weekdays and weekend days, for each month of the year. The months modelled are the calendar months of the year, with the real number of days in each month (without leap year adjustment, 28 days in February). Months were chosen because the electricity load is out of phase with the seasonal variation in wind power and with the seasonal variation in solar irradiance, see section 2.4.9.

The probabilistic model includes a typical annual pattern of weekends and holidays. In the probabilistic method, Weekdays are treated separately from weekends or holidays because the electricity demand is higher on weekdays. In practice, Saturdays have higher electricity demand than Sundays or bank holidays, but this effect was considered too small to worry about. The important factor was to recognise that different days have different levels of demand and to make the range of demands about right. The probabilistic method uses the UK holiday calendar and resulting profiles of electricity demand, although this method could be adapted to any country of the world.

In each month, at least 2 out of every 7 days are modelled as weekend days. The exact number of weekend days is rounded up or down to the nearest whole number. However, some adjustment was needed, as will be seen later. Starting in January, New Year's bank holiday adds one holiday to January and reduces the number of weekdays by one. February has exactly 8 weekend days and 20 weekdays. Easter sometimes falls in March but more often in April. Occasionally Easter spans the end of March and the beginning of April. The probabilistic model assumes three extra holidays at Easter, one being in March and two in April. May has two bank holidays: one on the first Monday and one on the last Monday. June and July have no bank holidays, but August has one. September, October and November have no bank holidays. Finally, the whole period between Christmas and New Year has relatively low electricity demand. The probabilistic method treats this period as all holiday and December therefore has five extra holidays.

A further correction is needed because so many months have 9 weekend days (rounded up from 2/7 of 30 = 8.57 or 2/7 of 31 = 8.86). To achieve the correct total number of holidays, April loses one of its weekend days, as do September and November. The final numbers of weekdays and weekend days in each month is shown in table 2.6:

Table 2.6 Working days and Holidays in Each Month as Modelled in the Probabilistic 24-Hour Program

Month	Weekends and holidays	Weekdays
January	10	21
February	8	20
March	10	21
April	10	20
May	11	20
June	9	21
July	9	22
August	10	21
September	8	22
October	9	22
November	8	22
December	14	17
Total	116	249

2.9 Energy Store Model

2.9.1 Introduction to the Energy Store Model

This section sets out the assumptions made about an energy store and the mathematical model developed.

The probabilistic method and its time-step validation method both treat a store as a bin into which surplus energy is put and from which energy is retrieved to satisfy an energy deficit some time later. Nothing is known or assumed about the physics or chemistry of the energy store. The only parameters used by the model are the round-trip electrical efficiency, the parasitic loss, the power ratings of charging or discharging the store, and the typical cycle time of the store (probabilistic method) or the energy capacity of the store (time-step validation). Otherwise the probabilistic and time-step methods treat the store as a 'black box'. The effects of round-trip efficiency, parasitic loss and power ratings on net power flows to and from the store are shown in section 2.10.7.

2.9.2 Assumptions Made and Comparison with Real Energy Stores

1. The round trip efficiency of the store is constant except for a constant parasitic loss, see below. The ratio: $\frac{Electrical\_Energy\_Out}{Electrical\_Energy\_In}$ is constant, regardless of the actual rate at which the store is charged or discharged. Many real stores, such as batteries, tend to get more efficient as charge and discharge rates are reduced but less efficient as charge and discharge rates are increased.
2. The parasitic loss (if present) is a constant power consumption and applies regardless of the state of charge of the store, even when the store is not being actively used or when it is empty. This assumption makes the calculations of the probabilistic method much easier than if the parasitic loss were variable. In contrast, real stores have a self-discharge rate that may drop to zero when the store is empty, and electronic power conversion losses that may be zero when the store is not being used.

3. The maximum store charging rate is constant, regardless of how full the store is becoming. A real store, such as a battery, often has a reduced charging rate as the store gets nearly full.
4. The maximum store discharging rate is constant, regardless of how empty the store is becoming. A real store, such as a battery, often has a reduced discharging rate as the store gets nearly empty.
5. The probabilistic method assumes that an energy store has a fixed cycle time or operating period, and that energy cannot be carried forward from one period to the next. The state-of-charge at the end of one period is the same as at the beginning of that period. In practice, real stores operate in continuous time and can store energy for as long as needed. When power variability is low, a store can float between full and empty for a very long time. The state-of-charge at the end of a period is rarely the same as at the beginning.
6. The probabilistic method assumes that a store cannot be full at one time within an operating period and empty at another time within the same period. Thus operating periods are divided into ones of net surplus, in which the store is sometimes full but never empty, and other periods of net deficit in which the store is sometimes empty but never full. Again, in practice, real stores operate in continuous time, and when the power variability is high, the store may be full then empty, or vice versa in a very short period of time.

2.9.3 Typical Parameter Values

The charge and discharge rates of an energy store depend on the physical size of that store, and the parasitic loss also depends to some extent on the embodiment of the store, but round-trip efficiency and typical cycle time are mainly dependent on the chosen technology. Table 2.7 lists some approximate data first published in a paper by this author, (Barton, Infield 2004) and updated in places for clarification.

Table 2.7 Typical Properties of Energy Storage Technologies

Technology	Round-trip efficiency, η	Typical discharge time	Typical time scale, period of operation, T
Super-conducting magnetic energy storage (SMES)	90% (but high power consumption of refrigeration)	10^{-4} to 10^{-3} hours, = seconds	10^{-4} to 10^{-3} hours, = seconds
Super-capacitor	86%	10^{-4} to 10^{-2} hours, = seconds to minutes	10^{-4} to 10^{-2} hours, = seconds to minutes
High-speed flywheel	89%	10^{-4} to 10^{-2} hours, = seconds to minutes	10^{-4} to 10^{-2} hours, = seconds to minutes
Traditional lead-acid batteries	63%	1 to 5 hours	3 to 24 hours
Zinc-bromide batteries	70%	2 to 5 hours	6 to 24 hours
Nickel-cadmium batteries	72%	1 to 10 hours	3 hours to 3 days
Sodium-sulphur batteries	87%	4 to 8 hours	12 to 24 hours
Nickel metal hydride batteries	64%	1 to 4 hours	3 to 24 hours
Flow cells, e.g. Regenesys and vanadium	75%	2 to 12 hours	6 to 24 hours
Hydrogen as a compressed gas	32%	12 hours+	1 day or much longer

2.10 Combination of Probabilities

2.10.1 Systems with Variable Wind Power Only (The First Probabilistic Program)

If an electrical system consists only of a constant load, wind powered generation and a backup supply or a weak grid connection, only the wind power is variable. This simplification facilitates a straightforward implementation of the probabilistic modelling method. The methods of section 2.5 can be applied to the whole wind power spectrum, both its stochastic and periodic components of wind speed variation. The model need not concern itself with correlation or anti-correlation with any other source of variability.

2.10.2 Systems Including Wind Power, Solar Power and Variable Load (The 24-Hour Probabilistic Program)

The task of modelling an electrical system with wind power, solar power and a variable load using a probabilistic method is considerably more complex. Wind power may sometimes correlate positively with solar power e.g. higher wind speeds occur during the daytime, or wind power may sometimes correlate negatively with solar power e.g. higher wind speeds occur during the winter than in the summer. Then the correlation of both wind power and solar power with electrical demand has to be considered.

The 24-hour probabilistic modelling program solves most of this complexity by considering periodic effects separately from stochastic or random effects. Once periodic effects have been removed, the remaining stochastic variation of solar power is shown to have very little correlation with wind power. The variation in electrical demand is almost entirely periodic and therefore shows the minimum of correlation with stochastic variations of either solar or wind power. In any case, the largest single driver of variation in electrical demand (after daily and seasonal effects) is ambient temperature (McSharry, Bouwman & Bloemhof 2005) and not solar irradiance or wind speed. This periodic approach works well but leads to a considerable complexity in computer programming and much longer computer run times compared to the wind-only probabilistic program.

2.10.3 Program Loop Nesting

The program structure of the probability combination subroutine consists of 7 nested 'for' loops, illustrated in fig. 2.74. The code is laid out in pseudo Matlab format. The comment lines are written in green and begin with a '%' to distinguish them from executable code:

```
For Month = 1 to 12
    %Calendar month of the year
                                                  "Slow Loops"
    For Day_Type = 1 to 2
        %Weekday or weekend

        For Slow_Solar_Count = 1 to N1
            %Daily average 'sunny fraction' interval

            For Slow_Wind_Count = 1 to N2
                %Daily average wind speed interval

Store Timescale T = 24 hours  ——————————————————————————————
                For Hour = 1 to 24
                    %Hour of the day.               "Fast Loops"
                    NB daylight treated differently from night time

                    For Fast_Solar_Count = 1 to 2
                        %Instantaneously Cloudy or sunny
                        (Matlab implied loop)

                        For Fast_Wind_Count = 1 to N3
                            %Instantaneous Wind speed interval
                            (Matlab implied loop)
```

Figure 2.74 Nested for loops of the 24-hour probabilistic program

2.10.4 Coding Optimisation

The operation of these 'for' loops is the main reason for the increased computer run time compared to the wind-only, first computer program. However, efficient and well-organised programming has mitigated this increase.

Firstly, Matlab enables ‘implied’ loops by simultaneously performing mathematical operations on all elements (or a range of elements) within arrays and matrices. The operations on arrays and matrices considerably reduce the computer run time compared to the explicit loops necessary in some other computing languages. The implied loops were used at the inner two nested loops.

Secondly, the program does as much pre-processing as possible in other subroutines. For example: matrices of probabilities for wind speed intervals and sunny fraction intervals are pre-calculated. Sun-earth geometry and hence solar elevation and air mass are already calculated. Typical numbers of week and weekend days in each month (based on the British calendar of bank holidays) are pre-calculated, together with their levels of electrical demand. The ‘for’ loops then only have to look up the relevant line of each matrix, each line being called several times in the course of the program. For example, each month of the year has a long-term average daytime wind speed and a long-term average night-time wind speed. From these, probability density functions (PDFs) of long-term wind speed are constructed in a separate subroutine. Each is an array of wind speed intervals with associated probabilities. Each array is called many times during the course of the program, for every possible value of daily sunny fraction, and for both weekdays and weekend days.

Thirdly, the program works mainly in power intervals, not wind speed intervals, at least in the inner loops. PDFs of wind power, solar power and electrical demand are calculated for each hour of the day. These PDFs are directly convoluted for each hour of the day and each daily weather condition. The calculation of the wind power PDFs require the wind power intervals to be converted into wind speed intervals. The constructed wind speed PDFs, section 2.6.3 can then be used to calculate probabilities for each wind power interval. This is achieved via an inverted wind turbine power curve. The turbine power curve only has to be inverted once for each system variant, then treated as a look-up table to save computation time.

Fourthly, meaningless combinations or ones with zero probability are not calculated. For example, during the night-time hours, there is only one possible value of solar power, i.e. zero! Then during the day, at any one instant of time, the solar power can only have two possible values according to the solar power model, section 2.7, i.e. sunny or cloudy.

Fifthly, the number of possible power intervals is reduced compared to the wind-only program. The total number of power levels between minimum net power (max. demand and min. solar or wind) and maximum net power (min. demand, max. solar and max. wind) is only 51. Further refinements to the program may find a more optimum balance between program run time and calculation accuracy.

2.10.5 The Energy Store and Time Scale, T

As outlined in spectrum Integration, section 2.5.1, the probabilistic method uses a time-scale, T as its starting point. In the case of the first computer program (wind only), this time scale is arbitrary and chosen by the user to be anywhere between minutes and months. However, for the second computer program, T is fixed at 24 hours. In future, similar programs could be written to work on other time scales, but this has not yet been done. Referring to the nested 'for' loops in fig. 2.74 above, T occurs inside the first four loops (month, day type, slow solar count and slow wind count) but outside the last three loops (hour, fast solar count and fast wind count), fig. 2.74.

The four outer loops are therefore designated 'slow' loops, slower than 24 hours, and the three inner loops are designated 'fast' loops, faster than 24 hours. Many of the program variable names have been chosen to reflect this convention.

2.10.6 Calculation of System Net Power PDF – Convolution of Probabilities

The electrical demand is treated as a deterministic function of hour of the day, i.e. a load profile, but solar power and wind power vary both randomly and with time of day. The inner two nested loops in fig. 2.74 are used to convolute the wind and solar power probabilities to create PDFs of net power in each hour, i.e. solar + wind - demand. The third loop, the hour-of-day loop, averages all the net power PDFs to build daily PDFs of net power.

Note that the convolution of solar and wind probabilities could have been performed using the specific Matlab 'conv' command. However, in the solar power model of this method, the instantaneous solar power can have only two possible values: sunny and cloudy. The convolution of solar and wind power is therefore done more quickly by a simple addition of shifted wind power PDF arrays.

2.10.7 Calculation of Store Power PDF

The net useful power entering or leaving the store also depends on the finite charging and discharging rates of the store, the finite round-trip efficiency of the store and any standing (parasitic) loss of the store, see section 2.9. The net power of the electrical system is converted to net store power as illustrated by fig. 2.75.

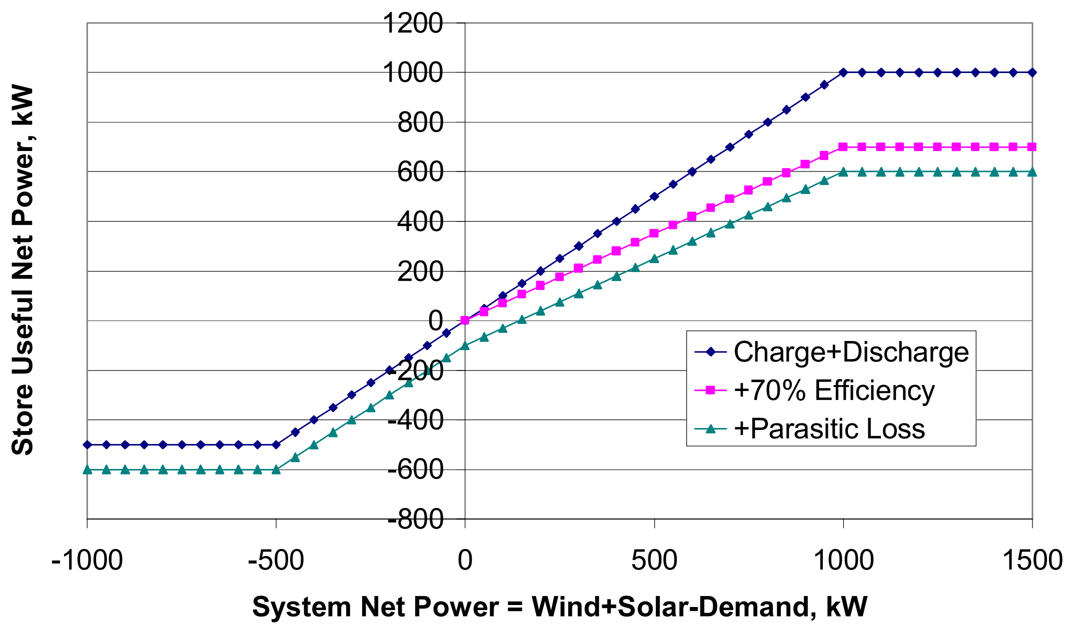

Figure 2.75 Conversion of electrical system power to store power

In the example system of fig. 2.75, the maximum discharge rate is 500kW and the maximum charge rate is 1000kW. Negative powers are therefore first truncated to –500kW and positive powers are truncated to +1000kW.

When the net power is positive, the round-trip efficiency of 70% applies, since efficiency losses are accounted as energy enters the store, not as it leaves. The positive portion of the graph is therefore kinked and the maximum positive power is limited to +700kW

Finally, the store net power is adjusted downward by the parasitic loss of 100kW. This applies to both the positive and negative portions of the graph. The minimum power is reduced from –500 to –600kW and the maximum power is reduced from +700kW to +600kW.

In the probabilistic programs, the electrical system net power PDF is modelled as an array of net power, e.g. –1000kW to +1500kW in steps of 50kW, each with an associated probability in a separate array. The PDF of store net power is calculated by simply converting the power levels using the process illustrated above. The power levels change while the associated probabilities remain unchanged.

2.10.8 Effect of the PDF of Instantaneous Net Store Power on System Behaviour

An example daily PDF of instantaneous net power to or from a store, fig. 2.76 is shown below for illustration only. One daily PDF is constructed for each iteration of the outer four (slow) loops. That is one net PDF for each combination of daily average wind speed, daily average sunny fraction, weekday and weekend day, in each month.

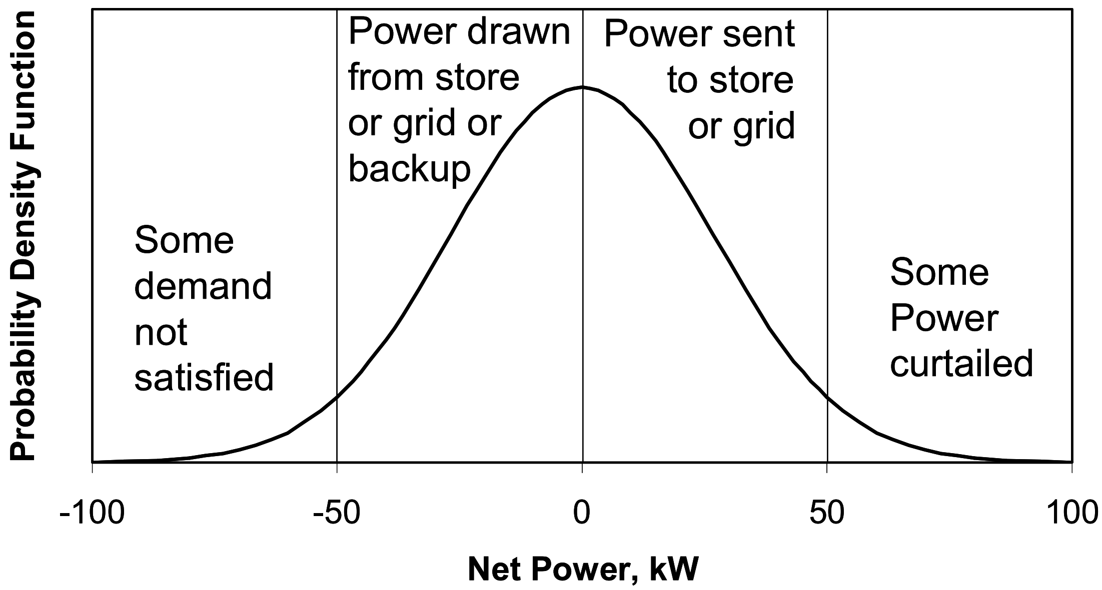

Figure 2.76 Example daily PDF of net instantaneous power

In the above illustration, the average net power is zero, the charging power capacity of the store+grid is 50kW and the discharging capacity of the store+grid+backup power is 50kW. Depending on the control strategy chosen by the user, section 2.12.4, when there is a positive net power, i.e. a surplus of power, some power is sent firstly to the store and secondly to the grid or vice versa. Similarly, when there is a negative net power, i.e. a power deficit, that deficit is met firstly by the store and secondly by the grid and backup power or firstly by grid and backup and secondly by the store. If the positive net power is too large some power has to be curtailed and if the negative net power is too large then some demand will not be satisfied (a power cut results)

The probabilistic programs calculate the total expected powers entering and leaving the store during each 24-hour period by integrating the store net power PDF (adjusted for the store round-trip efficiency, any store parasitic losses and the finite charging and discharging rates of the store).

Periods of 24 hours are thus divided into ones where the net power to the store is positive (more power is sent to the store than is drawn from it over 24 hours), ones where the net power to the store is negative (more power is drawn from the store than is sent to it over 24 hours) and ones where the store power is neutral (net surplus or deficit is balanced by the grid and/or backup generation). These three cases are treated slightly differently.

2.10.9 When the Daily Average Net Power to the Store is Positive

Excess power may occur because the net power exceeds the charge power rating of the store, or because more energy is sent to the store than is drawn from the store over a period of 24 hours (after losses are accounted), or because of a combination of these situations.

When the net store power is positive over a 24-hour period, the probabilistic method assumes that during that period, the store spends some time full, some time emptying, some time re-filling but none empty, fig. 2.77. Note that the probabilistic method knows nothing of the time sequence of power variations. These figures are for illustration only.

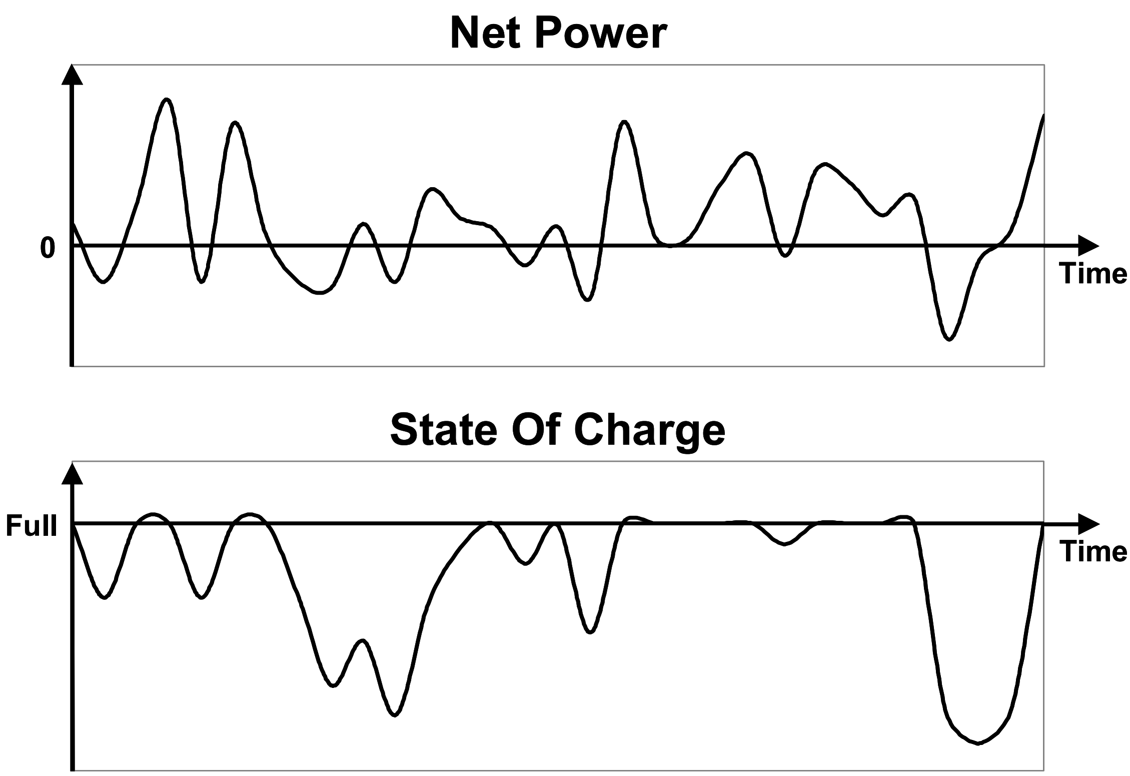

Figure 2.77 Behaviour of a store in a period when its average net power is positive

The probabilistic method also assumes, on average, that the state-of-charge of a store is the same at the end of a period as at the beginning. Therefore, the net excess power flowing into the store has to be balanced by energy curtailment or extra energy exported to the grid. The choice between these two options is decided by the store control option, section 2.12.4.

The flows of power to and from the store, from sources and to demand are calculated in a probability matrix, a simplified example of which is explained in section 2.10.13.

2.10.10 When the Daily Average Net Power to the Store is Negative

Conversely, a power deficit may occur because the net power drawn from the store exceeds the discharge power rating of the store, or because more energy is drawn from the store than is sent to the store (after losses are accounted), or because of a combination of these situations.

When the net store power is negative over a 24-hour period, the probabilistic method assumes that during that period, the store spends some time empty, some time filling, some time emptying but none full, fig. 2.78.

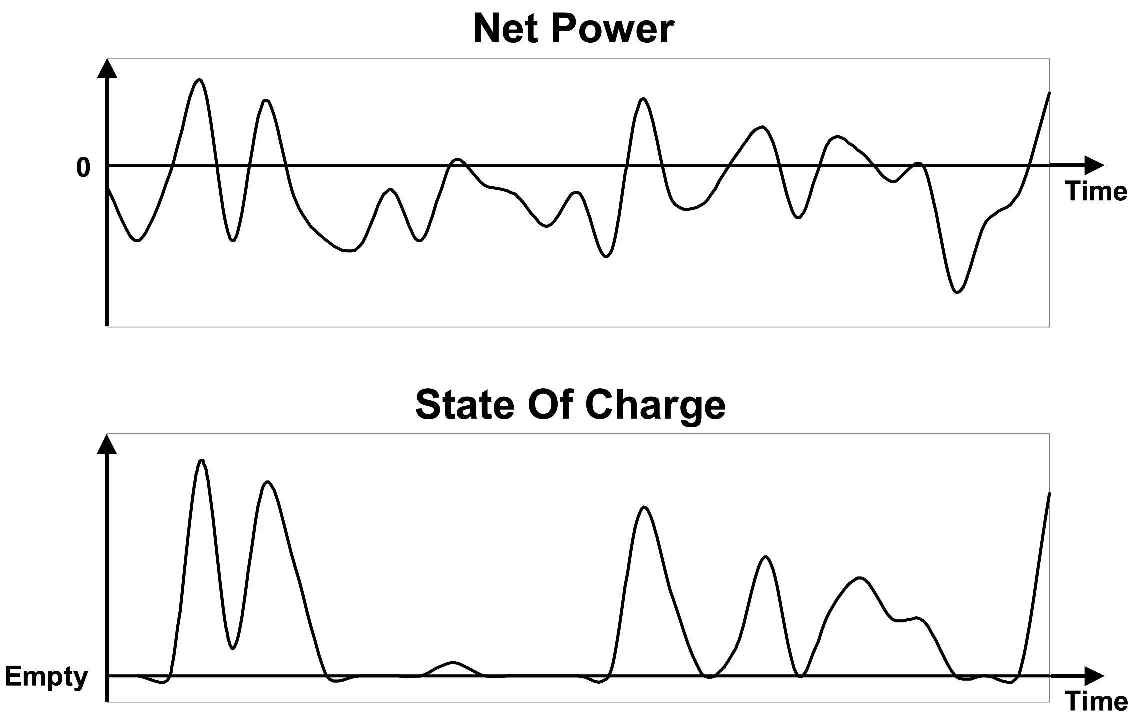

Figure 2.78 Behaviour of a store in a period when its average net power is negative

Again, on average, that the state-of-charge of a store is the same at the end of a period as at the beginning. Therefore, the net deficit of power in the store has to be balanced by imports from the grid, by backup generation or by power cuts (unsatisfied demand). The choice between these options is decided by the store operating procedure, section 2.12.4. Again, see section 2.10.13 for an explanation of the probability matrix.

2.10.11 When the Daily Average Net Power to the Store is Neutral

Any net store surplus or deficit over 24 hours is balanced by the grid and by backup generation. The probabilistic method assumes that the store spends some time filling and some time emptying but none full or empty. This does not necessarily prevent all power cuts or all power curtailment: There may be times when the instantaneous power deficit exceeds the discharge rate of the store together with the grid import and backup generation capacities. There may also be times when the system power surplus exceeds the store charge rate together with the grid export capacity. The power flows are still calculated in a probability matrix.

2.10.12 Store Operating Strategies

In practice, three different operating options can be modelled by the 24-hour probabilistic program, see section 2.12.4:

1. Keep the store as full as possible. Export to grid only when store is full or store charge rate is exceeded. Use grid import and backup generation whenever net power is negative.
2. Use the grid and backup generation at a constant rate to keep the store as balanced as possible
3. Keep the store as empty as possible. Import from grid or use backup generation only when store is empty or store discharge rate is exceeded. Export to grid as much as possible when net power is positive.

2.10.13 The Probability Matrix

At the heart of the probabilistic method is a spreadsheet or matrix of power levels and associated probabilities (Barton, Infield 2004). In the case of the first program (wind variation only) the matrix only has two dimensions: long-term variations in wind speed and short-term variations in wind speed. The 24-hour program (wind, solar and demand all varying) has more dimensions but the principle is the same. Each of the short-term variations (fast, inner ‘for’ loops in fig. 2.74) are represented by columns of the matrix and each of the long-term variations (slow, outer ‘for’ loops in fig. 2.74) are represented by rows of the matrix.

A very simplified example of a matrix is shown here in table 2.8, similar to the one published in (Barton, Infield 2004). This matrix shows just wind power variations as modelled in the wind-only probabilistic program. There is no grid connection, representing a stand-alone system with backup generation. The store operating period, T is 24 hours in this example (the original published example used one hour) and the matrix consists of just 3 rows and 3 columns.

Real matrices would consist of many more rows and columns for greater accuracy.

In this example, the long-term average wind speed is 8m/s and the standard deviation of wind speed variation within each day is 4m/s. This is like a turbulence intensity of 50%, but applies over 1 day, not just to variations within an hour. The distribution of daily average wind speeds is modelled by one PDF. Each row represents a different interval of daily average wind speed. Let us call the mid-value of this interval U_1. In the example, U_1 takes three possible values: 6m/s (representing daily-average wind speeds from 0m/s to 7m/s); 8m/s (representing daily-average wind speeds from 7m/s to 9m/s); and 10m/s (representing daily-average wind speeds above 9m/s). Each element within each row represents a different interval of instantaneous wind speed (second by second) higher or lower than U_1. The instantaneous wind speed $= U_1 + U_2$, where U_2 is the interval mid-value of the instantaneous variation. The matrix has been constructed so that each column represents a fixed ratio of U_2/U_1 (like a turbulence intensity, but in this case the sample period is 24 hours). In the example spreadsheet of Table 2.8, U_2 takes values of –1, 0 and +1 times the standard deviation of variation within one day. Since one standard deviation is 50%, the absolute wind speed, $U_1 + U_2$ takes values of 0.5, 1.0 and 1.5 times the daily average wind speed.

E.g. when the daily average is 6m/s, U_2 takes values of–3m/s, 0m/s and +3m/s, giving absolute wind speeds of 3m/s, 6m/s and 9m/s. These values of $U_1 + U_2$ actually represent ranges of wind speeds: the first column represents values of wind speed less than the mean minus half a standard deviation ($U_1 + U_2 < 4.5$m/s); the middle column represents wind speeds from mean minus half a standard deviation to mean plus a half standard deviation ($4.5 < U_1 + U_2 < 7.5$); the final column represents wind speeds greater than mean plus half a standard deviation ($7.5 < U_1 + U_2$). Each element in the matrix is also associated with a wind turbine power output calculated using the turbine power curve of section 2.6.4 at a wind speed of $U_1 + U_2$, see table 2.8.

For each time scale considered, all variations in wind speed slower than the time scale in question are described by one probability distribution function, $p(U_1)$, (one value of probability for each row of the matrix). In the case of a one-hour store, this would be a Weibull distribution, perhaps even a Rayleigh distribution (a Weibull with a shape factor of 2). In the case of a 24-hour store this would also be a Weibull distribution but with a larger shape factor e.g. about 2.9 corresponding to smaller standard deviation.

Variations faster than the time scale are described by a second probability distribution of wind speeds within the time period, $p(U_2)$, (one value of probability for each column of the matrix). In the case of a one-hour store, this second distribution describes turbulent variations in wind speed, and is best represented by a Gaussian distribution. However, in the case of a daily store, e.g. table 2.8, this second distribution would include turbulent variation and some weather system (synoptic) variation up to periods of one day and would be better represented by another Weibull distribution of high shape factor, e.g. about 2.4. The wind speed distributions are described in more detail in section 2.6.3.

In the spreadsheets, energy can be ‘taken’ from points with high wind power output and ‘given’ to other points in the same row with lower power. This represents the action of a store absorbing transient surpluses and delivering this as useful energy a short period later. Energy cannot be exchanged between rows because this would require longer-term energy storage than is being modelled. The precise redistribution is not important and is not calculated, but the spreadsheets do calculate the total amount of power transferred, and the total amount of extra energy used to supply the electrical demand, see table 2.8. For the example below, the maximum charging rate of the store is 500kW and the discharging rate of the store is 350kW. The electrical demand is a constant 400kW. The round-trip efficiency of the store is 70%, appropriate to an electrochemical storage technology and no standing losses are assumed. In practice, standing losses are accounted as a simple addition to the electrical demand. Let us now work through each daily-average case in table 2.8:

Row 1: When the daily average wind speed, U_1 is 6m/s, the wind turbine is only generating for a portion of the time. When the wind turbine does generate a surplus, this is easily used up at other times in the day. In the first column, $U_1 + U_2$ = 3m/s and the turbine is not generating. The store would be able to discharge up to 350kW (limited by the discharge rate of the store) if that energy were available. In the second column, $U_1 + U_2$ =6m/s and the turbine is generating 140kW. This leaves a deficit of 260kW compared to the demand of 400kW. The deficit is less than the store discharge capacity, so the store would be able to discharge 260kW, if the energy were available. In the third column, $U_1 + U_2$ = 9m/s and there is a small surplus of 100kW. This can all be absorbed by the store, but the useful stored power is only 70kW because the efficiency of the store is only 70%.

The probability-weighted average stored power is 22kW and the probability-weighted average spare discharge capacity is 207kW. The useful extra power delivered to the load is the minimum of these, i.e. 22kW, limited by the wind energy available during the 24-hour period. This leaves zero curtailed power but 201kW of demand to be met by backup generation (averaged over the 24 hours).

Row 2: When the daily average wind speed, U_1 is 8m/s, the times of power deficit are almost exactly balanced by times of power surplus. In the first column, wind speed, $U_1 + U_2$ = 4m/s, the wind power is 24kW and the power deficit is 376kW. However, the spare discharge capacity is again limited to 350kW by the discharge capacity of the store. In column 2, the wind speed is 8m/s and the wind power is 360kW, leaving a deficit of 40kW. This is well within the discharge capacity of the store. In column 3, the wind speed is 12m/s and the wind power is 915kW. This produces a surplus of 515kW of which only 500kW can be absorbed by the store, limited by the charging rate.

However, the stored power is only 350kW because the efficiency of the store is only 70%.

The raw probability-weighted surplus is 28kW larger than the raw probability-weighted deficit, but when the store power ratings and efficiency are taken into account, the stored power is 15kW less than the spare discharge capacity. The average extra power delivered to the load is therefore limited by the stored energy, equal to 109kW. This leaves 5kW of curtailed power and 23kW to be met by backup generation (averaged over the 24 hours).

Row 3: When the daily average wind speed, U_1 is 10m/s, the wind power exceeds the demand most of the time. The probability-weighted average power surplus is 280kW. In column 1, the wind speed, $U_1 + U_2$ = 5m/s and the turbine power is 69kW leaving a deficit of 331kW. This is within the discharge capacity of the store and can be supplied by the store if sufficient energy is available. In column 2, the wind speed is 10m/s and the wind power is 648kW. This produces a surplus of 248kW, all of which could be absorbed by the store, but only 174kW of which could be stored because of the 70% efficiency factor. In column 3, the wind speed is 15m/s (above the rated wind speed of the turbine) and the wind power is 1000kW. This creates a surplus of 600kW, only 500kW of which can be absorbed by the store, and only 350kW can be usefully stored due to the 70% efficiency of the store.

In this row, the probability-weighted stored power is 175kW but the spare discharge capacity is only 103kW. The extra power that can be delivered to the load is therefore limited by the deficit to 103kW. This time, no power is required from backup generation, but an average of 133kW is curtailed, calculated as follows: The useful extra power is 103kW, so the total power absorbed by the store is 103kW divided by 70% = 147kW. The difference between this and the surplus of 280kW is 133kW. Then the power lost due to the finite efficiency of the store is (100%-70%) times 147kW = 44kW.

The spreadsheets include the round trip efficiency of an appropriate energy storage technology. For example, the spreadsheet for one-hour storage assumes an efficiency of 90% appropriate to a high-speed flywheel system, whilst the Daily spreadsheet assumes an efficiency of 70% appropriate to an electrochemical store, for example the Regenesys flow cell system.

Table 2.8 Illustration of a Daily Probability Matrix in Operation – Wind Only Model with Constant Demand of 400kW

Wind Turbine Power Capacity =1000kW	Mean - 1 Std. Dev.	Mean	Mean + 1 Std. Dev.	Probability Weighted Average
Column Probabilities	0.31	0.38	0.31	
Row 1: U_1=6m/s, Probability = 0.43				
U_1+U_2	3.0	6.0	9.0	6.0
Wind Power, kW	0	140	500	208
Surplus Power, kW	0	0	100	31
Stored Power, kW	0	0	70	22
Deficit Power, kW	400	260	0	223
Spare Discharge Capacity, kW	350	260	0	207
Extra Power to Load, kW				22
Required Backup, kW				201
Curtailed Power, kW				0
Store Inefficiency, kW				9

Row 2: U_1=8m/s, Probability = 0.36				
U1+U2	4.0	8.0	12.0	8.0
Wind Power, kW	24	360	915	428
Surplus Power, kW	0	0	515	160
Power Stored, kW	0	0	350	109
Deficit Power, kW	376	40	0	132
Extra Discharge Capacity, kW	350	40	0	124
Extra Power to Load, kW				109
Required Backup, kW				23
Curtailed Power, kW				5
Store Inefficiency, kW				46
Row 3: U1=12m/s, Probability = 0.21				
U_1+U_2	5.0	10.0	15.0	10.0
Wind Power, kW	69	648	1000	578
Surplus Power, kW	0	248	600	280
Power Stored, kW	0	174	350	175
Deficit Power, kW	331	0	0	103
Extra Discharge Capacity, kW	331	0	0	103
Extra Power to load, kW				103
Required Backup, kW				0
Curtailed Power, kW				133
Store Inefficiency, kW				44

2.11 Calculation of Energy Capacity

This section sets out the method of calculating the energy capacity of a store in a given electricity system with a given store operating period, T. As described in the introduction to the methodology, section 2.1, the starting point for the probabilistic modelling method is the typical cycle time or operating period, T of the energy store. The store smoothes short-term variations in power that occur within a time-scale of T, but not longer-term variations with longer cycle times.

The probabilistic method uses the third filter function, section 2.5.5 to calculate the standard deviation of accumulated energy within a period, T originating from each frequency component, ω_i of wind speed variation or solar irradiance variation. The third filter function is used to create a third filtered spectrum of wind speed variations, fig. 2.79 and a third filtered spectrum of solar irradiance variations, fig. 2.80. The area under each of these filtered functions represents a variance in state-of-charge of an energy store with appropriate scaling, e.g. the square of the average gradient of the wind turbine power curve, K, see section 2.5.6. The spikes are first truncated from each spectrum to remove the variance associated with periodic variations, because the state-of-charge variance associated with periodicity is dealt with separately, section 2.11.4.1.

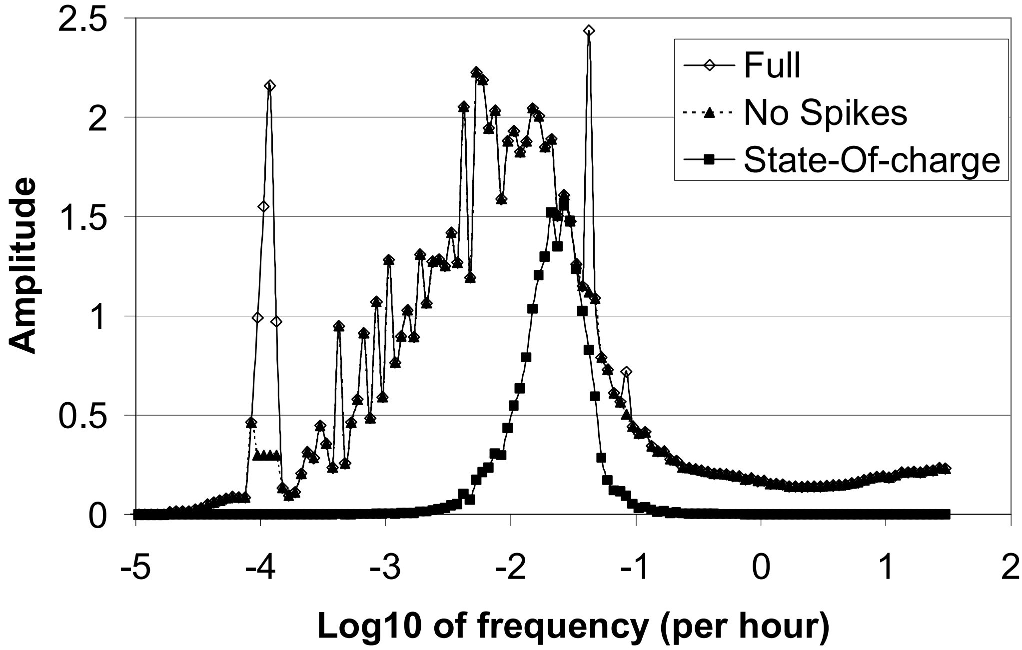

Figure 2.79 Third filter function applied to the generic wind speed variation spectrum, scaled to an average wind speed of 8m/s

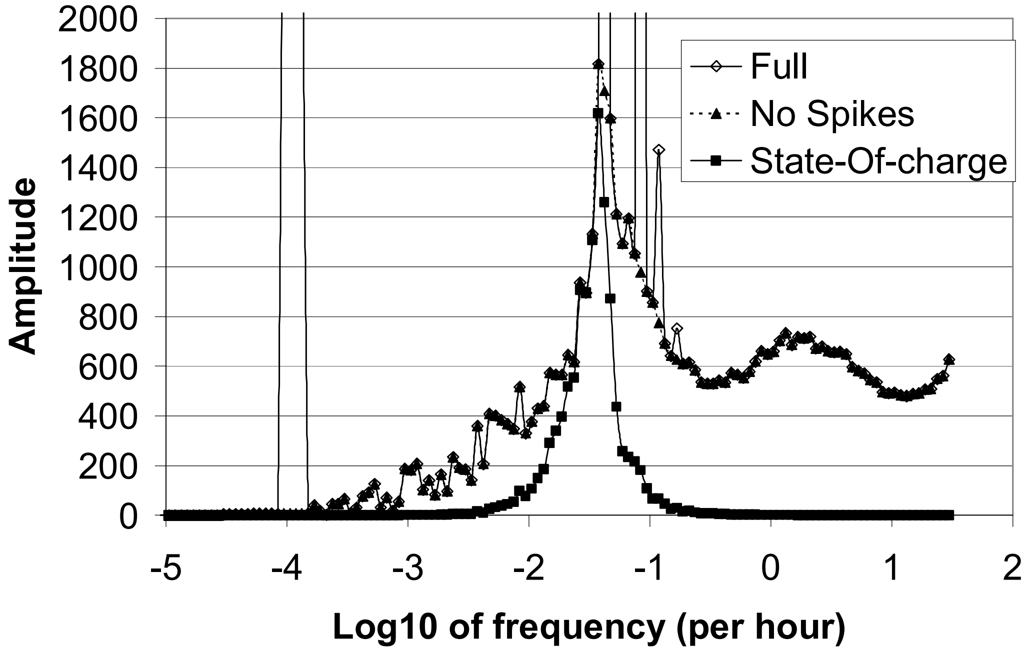

Figure 2.80 Third filter function applied to the solar irradiance variation spectrum, as measured at Rutherford Appleton Laboratory

2.11.1 Balanced Power Conditions for Store Size Calculation

For a given electricity system, and chosen period of time, T, the required size of store is a maximum in periods in which the energy entering the store is exactly balanced by the energy leaving the store, i.e. balanced conditions. A very simple example follows to illustrate why this is so.

Suppose that the net power into or out of a store follows a square wave, $s(t)$ with a mark-space ratio equal to 1.0 and a period, τ . Suppose also that the average value of net power is offset from zero by a value, m :

$$s(t) = m + \Delta \quad \text{when } t < \frac{\tau}{2}$$

$$s(t) = m - \Delta \quad \text{when } t > \frac{\tau}{2} \qquad (2.63)$$

For example, if m =+2kW, Δ=5kW and τ =8 hours, the following square wave results, fig. 2.81:

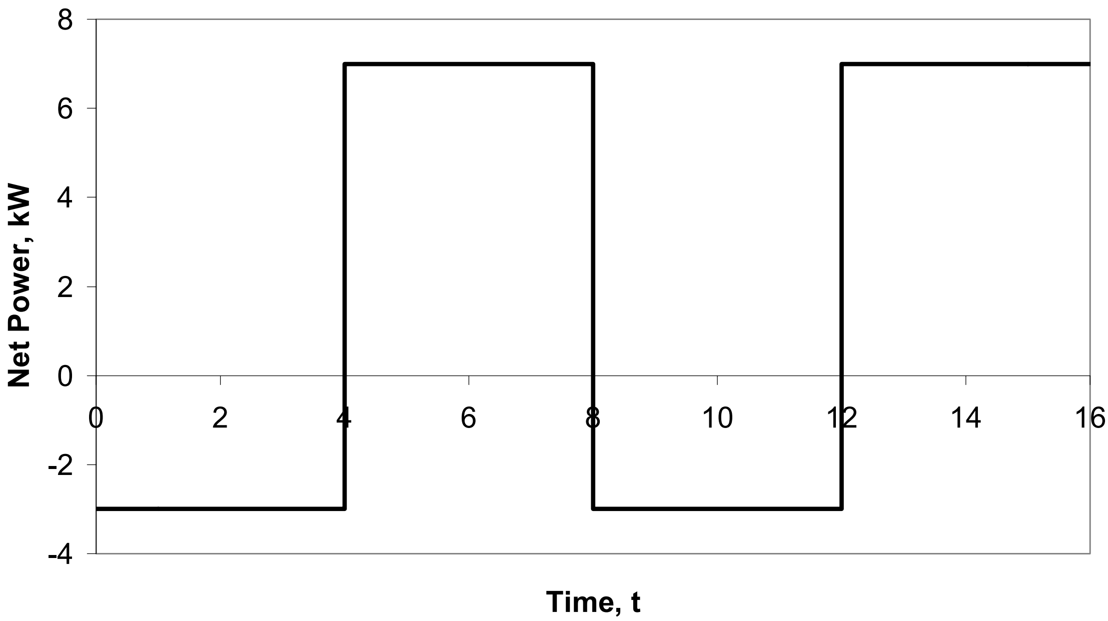

Figure 2.81 Square wave graph of net power into and out of the store

If m is positive (average net power entering the store) then the state of charge would ratchet up until the store is full. Then the state-of-charge will dip below full and recover on each cycle of net power, fig. 2.82:

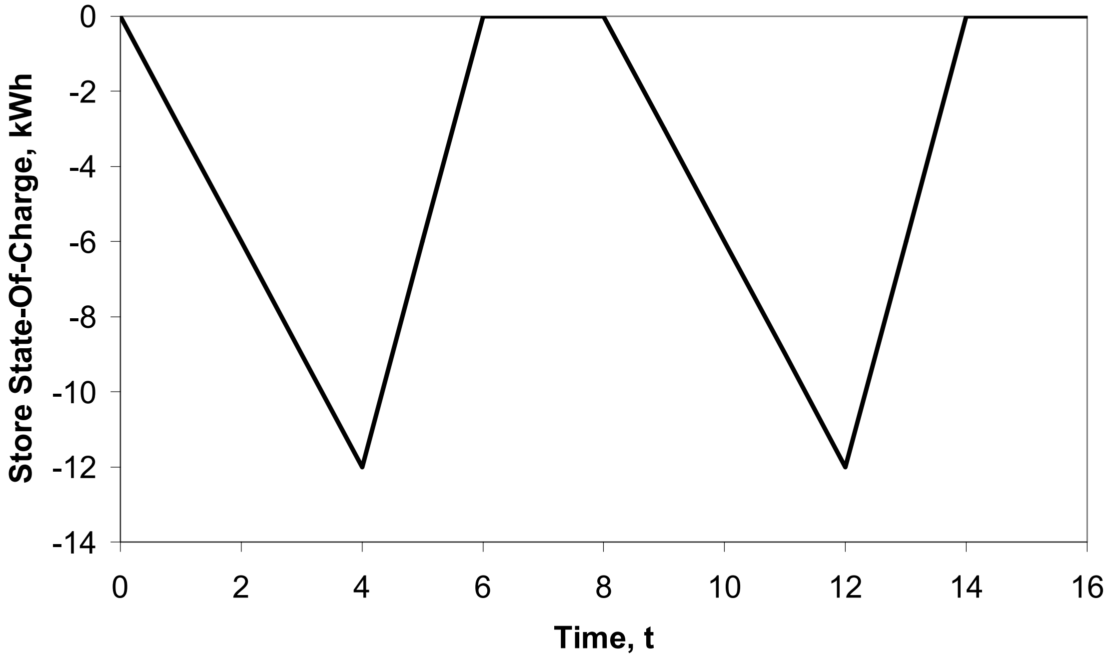

Figure 2.82 State of charge resulting from a square wave of positive net power

The mean value of net power is greater than zero. Therefore the power into the store in the positive half of the square wave is greater than the power drawn from the store in the negative half of the cycle. This causes the upward slopes of state-of-charge (SOC) to be steeper than the downward slopes, leaving a portion of each cycle when the store is full. The maximum excursion from full is given by integrating the net power function, $s(t)$ over the negative half of each cycle, so the negative change in SOC $=(\Delta-m)\frac{\tau}{2}$. (2.64)

In the numerical example above, the maximum excursion from full is 12kWh.

Now let us consider the case where m is negative (average net power leaving the store). Then the state of charge would ratchet down until the store is empty. The state-of-charge will then rise above empty and drop again on each cycle of net power, fig. 2.83:

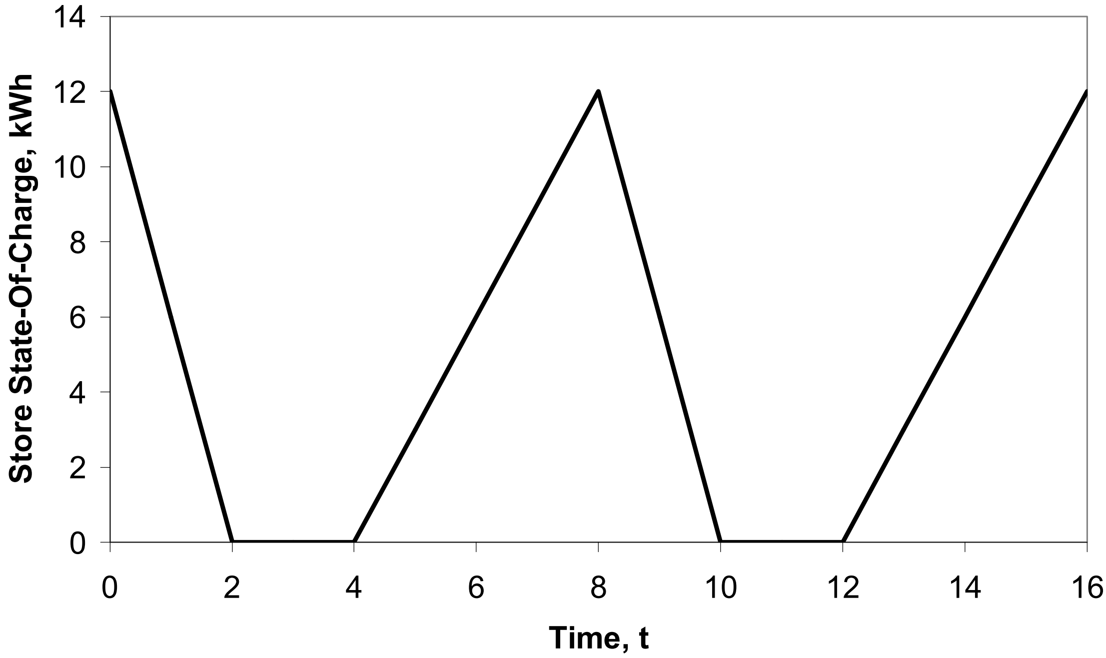

Figure 2.83 State of charge resulting from a square wave of negative net power

The mean value of net power is less than zero. Therefore the power into the store in the positive half of the square wave is less than the power drawn from the store in the negative half of the cycle. This causes the downward slopes of state-of-charge (SOC) to be steeper than the upward slopes, leaving a portion of each cycle when the store is empty. The maximum excursion from full is given by integrating the net power function, $s(t)$ over the positive half of each cycle, so the positive change in SOC$=(\Delta+m)\frac{\tau}{2}$ (2.65)

In the numerical example above, now with m =-2kW, Δ=5kW and τ =8 hours, the maximum excursion from empty is 12kWh.

Whether the net power is positive or negative, the range of state-of-charge is always: $=(\Delta-|m|)\frac{\tau}{2}$, and for a given magnitude and period of power variations, the range of state-of-charge is a maximum when m =0, i.e. when the energy entering the store is equal to the energy leaving the store. Similar analysis would apply for any other shape of power variations, e.g. triangle wave, sinusoid etc. The largest range in state-of-charge is observed, and therefore the largest store size is required, when the store energy is balanced.

2.11.2 Locating the Periods in which Store Energy is Balanced

In the wind-only probabilistic computer program, only wind speed varies while electrical demand remains constant and solar power is absent. This means there is only one long-term 'slow' loop, fig. 2.74, corresponding to varying period-average wind speed. There exists only one iteration of that slow loop in which the store power is most nearly balanced. This is the one where the period-average wind speed is such that wind power balances demand + store losses. This is also the iteration where the store energy capacity should be calculated.

However, in the 24-hour probabilistic program, daily average wind speed changes, daily average demand changes (according to weekdays or holidays and with month of the year), and daily average solar power changes (with daily sunny fraction and as solar elevations change through the year), fig. 2.74. Therefore, there are many different cases in which the store energy will be balanced: This may occur on a low-demand day with moderate wind power and low solar power, or on a high-demand day with high wind power and moderate solar power, or on a moderate-demand day with low wind power and high solar power. The possible combinations are large in number.

Imagining the varying wind power as a continuous line, the wind-only program merely has to find the point on this line where supply = demand + losses, but the 24-hour program has to find a surface in multi-dimensional space (four dimensions in this case). The various parameters of month, day type, daily average wind speed and daily average sunny fraction each represent one dimension in this space. Month and day type are examples of discontinuous variation, with just a few discrete possible values, but daily average wind speed and daily sunny fraction are examples of continuous variation, having a large number of possible values, determined only by the bin sizes used in the program.

The 24-hour probabilistic program solves the problem by considering each two-dimensional map of wind speed and sunny fraction within each discrete case of day type and month.

2.11.2.1 A Worked Example

For example, table 2.9 is derived for weekend days in June in an electrical system with an average load of 500kW, a wind turbine capacity of 500kW and a solar PV installed capacity of 1MW. The store has 100% round-trip efficiency and zero parasitic loss. Ignoring for a moment the improbably high level of solar PV capacity, and the perfectly efficient store, the table shows which days have a net store surplus and which days have a net store deficit:

Table 2.9 The Boundary Between Energy Surplus and Energy Deficit

	Increasing Daily Average Wind Speed →																									
↓ Increasing Daily Sunny Fraction	D	D	D	D	D	D	D	**S**	S	S	S	S	S	S	S	S	S	**S**	D	D	D	D	D	D	D	D
	D	D	D	D	D	D	D	S	S	S	S	S	S	S	S	S	S	S	D	D	D	D	D	D	D	D
	D	D	D	D	D	D	D	S	S	S	S	S	S	S	S	S	S	S	**D**	D	D	D	D	D	D	D
	D	D	D	D	D	D	D	S	S	S	S	S	S	S	S	S	S	S	**S**	D	D	D	D	D	D	D
	D	D	D	D	D	D	**D**	S	S	S	S	S	S	S	S	S	S	S	S	**D**	D	D	D	D	D	D
	D	D	D	D	D	D	**S**	S	S	S	S	S	S	S	S	S	S	S	S	**S**	D	D	D	D	D	D
	D	D	D	D	D	D	S	S	S	S	S	S	S	S	S	S	S	S	S	S	**D**	D	D	D	D	D
	D	D	D	D	D	D	S	S	S	S	S	S	S	S	S	S	S	S	S	S	**S**	**D**	D	D	D	D
	D	D	D	D	D	D	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	**S**	**D**	D	D	D
	D	D	D	D	D	**D**	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	**S**	D	D	D
	D	D	D	D	D	**S**	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	**D**	D	D
	D	D	D	D	D	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	**S**	**D**	D
	D	D	D	D	D	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	**S**	**D**
	D	D	D	D	**D**	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	**S**
	D	D	D	D	**S**	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S
	D	D	D	D	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S
	D	D	D	**D**	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S

In table 2.9, the areas marked with ‘D’ have a net energy deficit and areas marked with an ‘S’ have a net energy surplus. The left-hand boundary represents zero wind speed and zero wind power. Even when the daily sunny fraction is 1.0 (the bottom left hand corner) the solar power is still insufficient to supply the daily average electrical demand.

The middle of table 2.9 represents moderate wind speeds and useful wind power, creating an energy surplus when averaged over the day. Moving from left to right, the boundary between deficit, 'D' and surplus, 'S' occurs first at the bottom of table 2.9, where solar power is high and wind power is low, and later at the top of table 2.9 where solar power is low but wind power is higher.

The right-hand side of table 2.9 represents very high daily average wind speeds that cause the turbine to be furled for part of the day, causing lower wind power again. The top-right corner of table 2.9 therefore has a net energy deficit and is full of 'D's. This creates another boundary between surplus and deficit. Moving from left to right, this second boundary occurs first at the top of table 2.9, where wind power is still high and solar power is low, and second at the middle of the table, where wind power is lower and solar power is moderately high. The boundary between 'S' and 'D' disappears off the right-hand edge of table 2.9 before it reaches the bottom edge. If the table extended to even higher daily average wind speeds, then the furling of the turbine would again cause virtually zero wind power and some power deficit even at highest solar power.

The size of the store is calculated at the points closest to the 'D' to 'S' boundary, in the squares highlighted in red, bold font. These are the squares where a change of one discrete bin increment in either sunny fraction or daily average wind speed would be enough to change the store state from 'S' to 'D' or vice versa. The highlighted squares are either an 'S' surrounded by 3 'D's, or a 'D' surrounded by 3 'S's. These squares represent the daily weather conditions that are the very closest to the boundary.

Note that this boundary between surplus and deficit changes into a broad balanced region when a grid connection and backup generation are available, under control option 2; see section 2.12.4.2. The principles of the calculation remain unchanged.

2.11.3 Energy Capacity for the Wind-Only Case – The First Probabilistic Program

In the case of wind power alone, the energy capacity is calculated from the filtered wind spectrum, after application of the third filter function, see section 2.5.5, eq. 2.9. The filtered spectrum is weighted by 1/(frequency squared) and integrated to give a variance of wind speed x time, with dimensions of (wind speed x time) squared. This quantity is rather meaningless on its own, but when multiplied by the average effective gradient of the wind turbine power curve, K, eq. 2.66, this yields the variance in state-of-charge within a period, e.g. one day. The resulting quantity has dimensions of (power x time) squared, or energy squared, as expected.

The effective gradient of the wind turbine power curve is not immediately available or obvious. The gradient depends which part of the curve we look at. For example, the bottom end of the curve is less steep than the middle or top end of the working range, fig. 2.4. Above the rated wind speed, the gradient is effectively zero. The gradient also depends on the range over which we take the average, since the turbine power curve is a curve not a straight line. The probabilistic method avoids these problems by calculating the standard deviation of wind speed within the store period, e.g. standard deviation of wind speeds within a day, and the standard deviation of wind powers resulting from the PDF of wind speeds within the same day. The standard deviation of within-day wind speeds is already calculated as part of the construction of wind speed distributions, see section 2.6.3. The standard deviation of wind powers has to be calculated by applying the wind turbine power curve to the resulting PDF of within-day wind speeds. Then:

$$Effective\_Turbine\_Curve\_Gradient, K = \frac{Std\_Dev(Turbine\_Power)}{Std\_Dev(Wind\_Speed)} \quad (2.66)$$

The details of how eq. 2.66 is evaluated are described in section 2.5.6.

2.11.4 Solar Power, Wind Power and Demand - The 24-Hour Probabilistic Program

The energy capacity due to varying wind power is calculated as before, but for the case of solar power, there is no power curve to worry about: The power output is proportional to the solar irradiance multiplied by the solar power capacity. The gradient, K in the eq. 2.9 is simply the solar power capacity in $kW/\left(W/m^2\right)$. This should approximately be the power capacity in MW under standard test conditions of 1000W/m2, 25°C cell temperature and an air mass of 1.5.

The solar power varies through the day according to the solar elevation. One function applies to sunny conditions and another for cloudy conditions; see sections 2.7.4 through 2.7.6 and fig. 2.58. Since the solar elevation is a function of sun-earth geometry, and the average sunny fraction is assumed to be constant through the day, the PDF of solar power within a day can easily be directly constructed. The variance of solar power is thus directly calculated for each bin of daily-average sunny fraction.

In the probabilistic method, the electrical demand is assumed to be a direct function of time of day and day type, see section 2.8. The variance of demand within a day can therefore be directly calculated.

Unfortunately, the total variation in state-of-charge cannot be calculated simply by summing the variances of wind, solar and demand variances, because within a day, wind power, solar power and demand are far from independent. All vary with time of day. Even wind speeds are higher during the day than at night. The probabilistic method therefore calculates the average wind power, the average solar power and the demand during each hour of the day and calculates the average net system power = wind power + solar power – demand, for each hour. The method steps through each hour of the day, calculating the average daily periodic accumulated energy surplus or deficit. This periodic variation is not a sine wave, or even a random combination of sine waves, but a strange and ‘knobbly’ shaped wave specific to that combination of daily weather conditions, day type and month. This time, the probabilistic method calculates the difference between the minimum and maximum values of state-of-charge through the day. The difference between minimum and maximum is treated as equal to twice the 'periodic' standard deviation in state-of-charge. The following analysis shows why this is so:

2.11.4.1 Daily Periodic Variations

Referring to section 2.5.5, the effect on variance of state-of-charge of a store due to one component of spectral variation, ω_i with mean-to-peak amplitude, A_i is given by:

$$\frac{A_i^2 K^2}{\omega_i^2}\left\{\frac{5}{6}+\frac{1}{6}\cos\left(\omega_i T\right)+\frac{2}{\omega_i^2 T^2}\left[\cos\left(\omega_i T\right)-1\right]\right\} \qquad (2.67)$$

If we have just one component of frequency, ω_T and that frequency completes exactly one cycle within the store period, T, then $\omega_T T = 2\pi$ the above formula becomes $\frac{A_T^2 K^2}{\omega_T^2}\left\{\frac{5}{6}+\frac{1}{6}\cos\left(2\pi\right)+\frac{2}{4\pi^2}\left[\cos\left(2\pi\right)-1\right]\right\}$ (2.68)

and this reduces to $\frac{A_T^2 K^2}{\omega_T^2}$. The standard deviation of state-of-charge is

therefore $Std.\_Dev. = \frac{A_T K}{\omega_T}$ (2.69)

Now if we consider this one spectral component, ω_T and calculate the accumulated energy surplus or deficit through one cycle of period, T, starting at zero energy and with zero phase angle:

$Power\_Component = A_T K \sin\left(\omega_T t\right)$ at time, t after the start of the period.

Then at time, τ after the start,

$$Accumulated\_Energy = \int_{t=0}^{t=\tau} A_T K \sin\left(\omega_T t\right) dt = \frac{A_T K}{\omega_T}\left[-\cos\left(\omega_T t\right)\right]_{t=0}^{t=\tau} = \frac{A_T K}{\omega_T}\left[1-\cos\left(\omega_T \tau\right)\right]$$

(2.70)

This energy has a minimum value of zero at $t = 0$ and $t = T$, and a maximum value of $\frac{2A_T K}{\omega_T}$ at $t = T/2$.

Thus the minimum-to-maximum range of state-of-charge is $\frac{2A_T K}{\omega_T}$

and half this range is $\frac{A_T K}{\omega_T}$, as above. (2.71)

Thus for a single frequency component, the standard deviation in state-of-charge is equivalent to half the minimum-to-maximum range in accumulated energy. That is why the value of 'standard deviation' for the daily periodic variation in energy is half the minimum-to-maximum range, as used in the probabilistic modelling programs. This becomes important when combining the size of store calculated to cope with periodic variations with the sizes of store calculated to cope with stochastic variations, section 2.11.4.3.

Note: The coding of the equations in the probabilistic Matlab programs actually contain two factors of 2. The first is because a Fourier transform results in a power spectrum, the magnitude of which represents only half the total variance of the original time series, (Bendat, Piersol 1993), pages 10 to 12. The second factor of 2 is because the actual amplitude of each frequency component, A is $\sqrt{2}$ times its root-mean-square value.

2.11.4.2 Adjustment for Finite Power Ratings, Finite Efficiency and Parasitic Losses in the Store

The user inputs to the probabilistic programs include the maximum charging power and the maximum discharging power of the energy store. The instantaneous net power of the electricity system (= wind power + solar power – demand – parasitic losses) may at times be in greater surplus than the maximum charging power of the store, or may be in greater deficit than the maximum discharging power of the store. In these circumstances, some energy is curtailed or some additional electrical demand must be supplied by backup power, grid connection or is not satisfied at all. The actual standard deviation of state-of-charge will be reduced accordingly. The probabilistic method does not have or use any temporal relationships to directly calculate the degree of reduction in required energy capacity. Instead, the probabilistic program assumes that the standard deviation of state-of-charge is proportional to the actual useful power entering or leaving the store, not the external power available.

The charging power includes the inefficiency loss of the store, which is accounted as energy enters the store, not as it leaves. At times of surplus, the charging power is therefore greater than the useful energy stored, unless the store efficiency is 100%. The standard deviation of state-of-charge is also reduced by the finite efficiency of the store.

The parasitic loss is effectively an additional electrical demand, and shifts the actual store power in the negative direction.

All these effects are accounted for by multiplying the variances of solar power state-of-charge, wind power state-of-charge and periodic state-of-charge by a further adjustment ratio:

$$Adjustment\_Ratio = \frac{Variance\_Of\_Net\_Useful\_Store\_Power}{Variance\_Of\_Net\_Electrical\_System\_Power} \quad (2.72)$$

The probabilistic method already calculates the PDF of net useful power entering or leaving the store, as part of the accounting of power flows, sections 2.10.6 and 2.10.7. The variance of net useful store power is calculated directly from this PDF.

The variance of net electrical system power for a given hour has to be calculated separately as the sum of variances of periodic power, solar power and wind power within the day. As a check on the accuracy of the method, when the store efficiency is 100%, when the parasitic loss is zero, and when the charge and discharge rates are so large as to not restrict the operation of the energy store at all, then the variance of net electrical system power is identical to the variance of net useful store power, within numerical error.

2.11.4.3 Combination of State-Of-Charge Variances

Now the total variance in state-of-charge within each day is the sum of the three components. The standard deviation in state-of-charge is the square root of this:

$$SOC\_Std\_Dev = \sqrt{(Solar\_Variance) + (Wind\_Variance) + (Periodic\_Variance)} \tag{2.73}$$

For each condition of month and week day or holiday (24 different combinations), each square represents a combination of daily wind speed and daily sunny fraction, each with its associated probability, based on the PDF of daily sunny fractions for that month and the PDF of daily average wind speeds for that month. As stated in section 2.4.8, the daily sunny fraction and the daily average wind speed for each month are assumed to be independently distributed. The individual probabilities of sunny fraction and wind speed are therefore simply multiplied together. The resulting probabilities are used to weight the energy store sizes for each square, to give a weighted average store size for that month and day type. Taking the example above of weekend days in June, the probabilities of each square are shown in table 2.10.

Table 2.10 Probabilities of Each Combination of Daily Average Wind Speed and Daily Sunny Fraction for Weekend Days in June, x 10,000

	Increasing Daily Average Wind Speed →																								
← Increasing Daily Sunny Fraction	0	0	0	0	0	0	0	**0**	0	0	0	0	0	0	0	0	0	**0**	0	0	0	0	0	0	0	0
	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	**0**	0	0	0	0	0	0	0
	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	**0**	0	0	0	0	0	0	0
	0	0	0	0	0	0	**82**	0	0	0	0	0	0	0	0	0	0	0	0	**0**	0	0	0	0	0	0
	0	0	0	0	0	0	**115**	0	0	0	0	0	0	0	0	0	0	0	0	**0**	0	0	0	0	0	0
	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	**0**	0	0	0	0	0
	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	**0**	**0**	0	0	0	0
	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	**0**	**0**	0	0	0
	0	0	0	0	0	**229**	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	**0**	0	0	0
	0	0	0	0	0	**224**	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	**0**	0	0
	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	**0**	**0**	0
	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	**0**	**0**
	0	0	0	0	**117**	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	**0**
	0	0	0	0	**71**	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	0	0	0	**1**	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

The probabilities have been multiplied by 10,000 and then rounded to the nearest integer to turn them into meaningful numbers. Only the boundary points are shown, i.e. the combinations of wind speed and solar irradiance where the store size is calculated. These squares are highlighted in bold, red text. Note that only the left-hand boundary has significant associated probabilities. These are the squares representing combinations with moderately high daily average wind speeds.

The right hand boundary represents times when the wind turbine is furled for part of the day due to excessive wind speeds and has very low probabilities. In fact, for the purposes of calculating the store size, the right hand boundary can usually be ignored. It is included for completeness in the probabilistic method in case the user would want to model a case with very high wind speeds compared to the rated wind speed of the turbine.

2.11.5 Calculation of Store Energy Capacity From Variance in State-Of-Charge

The standard deviation in state-of-charge of a store is equal to the square root of the variance. The energy capacity of the store is then calculated as +/- 1.0 standard deviations, i.e. the allowed energy range is 2 standard deviations.

This value of 2 standard deviations was found empirically. It has no direct analytical basis, although it is an elegant mathematical value: The excursion of state-of-charge can go positive or negative with equal probability, when the electrical system is balanced, with amplitude equal to the standard deviation of state-of-charge. This value was chosen because it produces good agreement between the probabilistic method and its time-step validation, on any time scale. The probabilistic method is run first using an assumed time-period of storage. The probabilistic method yields a standard deviation of state-of-charge and hence a store energy capacity equal to twice this value. This store energy capacity is then used in a time-step validation of the same system, and the results give excellent agreement with the probabilistic method, for example in energy delivered, energy curtailed, unsatisfied demand and proportions of time spent full and empty.

To give a very simple example, suppose an energy store has a working period of 24 hours, in an electricity system with a turbine that has a completely linear power curve, i.e. power in kW is equal to wind speed in m/s. Suppose the site average wind speed is 8m/s. Using the scaled RAL wind power spectrum produced from the years 1999 to 2002, the variance in state-of-charge is calculated to be 136.63 (m/s-hours)^2. The standard deviation is 11.69 m/s-hours. Then using an example turbine power curve gradient of 1kW/(m/s), the required energy range of the store is 2 standard deviations, equal to 23.38kWh. The average power from the wind turbine is 8kW. Therefore the capacity of the store represents almost 3 hours of wind power production (24kWh), but the store can smooth out variations within a 24-hour period. Most power passes straight from the wind turbine to the load, and only a small fraction has to pass into or out of the store.

Moving back to our earlier example of a real wind turbine power curve and solar power with an average demand of 500kW, the calculated store sizes are listed in the table2.11 for weekend days in June. Only the left-hand portion of the table is shown.

Table 2.11 Store Energy Capacities in kWh Calculated for 24-hour Storage in Weekend Days in June for an Example System

	Increasing Daily Average Wind Speed→													
↓ Increasing Daily Sunny Fraction	726	726	733	843	1077	1408	1683	**1796**	1766	1661	1548	1496	1554	1711
	1009	1009	1014	1101	1295	1589	1843	1948	1917	1817	1709	1658	1705	1842
	1239	1239	1244	1322	1494	1759	1993	2091	2062	1967	1866	1813	1849	1968
	1438	1438	1443	1515	1675	1921	2142	2233	2204	2112	2014	1960	1985	2092
	1619	1619	1624	1691	1841	2074	**2284**	2370	2341	2251	2155	2099	2119	2214
	1789	1789	1793	1857	1997	2219	**2420**	2503	2473	2385	2290	2235	2250	2333
	1949	1949	1953	2016	2149	2358	2551	2631	2600	2514	2424	2369	2377	2449
	2100	2100	2104	2166	2295	2496	2679	2755	2726	2644	2555	2499	2501	2564
	2245	2245	2250	2311	2436	2631	2809	2881	2852	2771	2684	2626	2622	2677
	2385	2385	2389	2450	2573	**2763**	2935	3005	2975	2895	2809	2750	2741	2788
	2521	2521	2525	2585	2706	**2891**	3059	3128	3097	3018	2932	2871	2859	2898
	2652	2652	2657	2717	2836	3017	3182	3248	3216	3138	3053	2991	2974	3007
	2781	2781	2785	2845	2963	3140	3303	3368	3335	3256	3171	3108	3088	3115
	2907	2907	2911	2971	**3087**	3262	3423	3485	3452	3373	3289	3224	3200	3222
	3030	3030	3035	3095	**3210**	3382	3541	3601	3568	3488	3404	3339	3312	3328
	3152	3152	3156	3216	3330	3501	3657	3717	3683	3603	3518	3452	3422	3433
	3271	3271	3275	**3336**	3449	3619	3773	3830	3797	3717	3631	3564	3531	3538

Only the values highlighted in red, bold font are the ones closest to the break-even boundary and are the store sizes used to calculate the weighted store size for weekend days in June. Weighting these store sizes by the probabilities in table 2.10 above, we get a store energy capacity of 2787kWh. This would be enough to supply the long-term average demand of 500kW for only 5.6 hours, but of course most of the power goes directly from source to demand without having to charge or discharge the store.

2.11.6 Overall Energy Store Size

In the wind-only program, only one energy store size is calculated. However, the 24-hour program creates 24 different store sizes; that is one for each month and day type. These store sizes are weighted and averaged to produce an overall store energy capacity. This weighting is in proportion to the total amount of time that the electricity system spends close to the break-even boundary in a whole modelled year. The monthly store sizes are first weighted according to the total of the highlighted probabilities in table 2.10. This is a good indication of the total time spent close to the break-even boundary in each day type and in each month. These weightings are shown for the example case above in the table 2.12:

Table 2.12. Indication of Time Spent Close to the Break-Even Boundary, Probabilities x 10,000

Month	Weekend Days	Week Days
January	0	0
February	71	0
March	216	9
April	485	439
May	622	562
June	839	634
July	833	577
August	549	341
September	344	268
October	180	164
November	7	0
December	0	0

Table 2.12 shows that there are many days in the summer when supply almost equals demand, but few days in winter, especially on weekdays. This is because electrical demand is higher in winter and solar power is much lower; even on windy winter days, 500kW of wind power is insufficient to meet the daily average demand.

Secondly, the store sizes are weighted according to the number of days in the year of each type, e.g. there are 9 weekend days in June. The total number of weekend days and weekdays in each month is shown in table 2.6 in section 2.8.4. The two weighting matrices are multiplied together, element-by-element to give the overall weighting factors for each day type.

For the example above, the resulting yearly overall store size is calculated to be 2548.6 kWh.

2.11.7 Validation

The calculated store sizes have been used as inputs to time stepping models of the same electrical systems. The results, section 3, show good agreement between the probabilistic method and the time step model run in this way. This constitutes the validation of the store size calculation and of the probabilistic method in general.

2.12 Import, Export and Backup Generation

The second probabilistic program (the 24-hour program), and its time-step validation, allow both grid-connected and backup generator operation to be modelled. When there is a power surplus, electricity may be exported via the grid connection as an alternative to charging the store. When there is a power deficit, electricity may be supplied from the grid connection or from a backup generator. If not required, these options are turned off by setting their power capacities to zero.

2.12.1 Grid Connection

The grid connection is modelled as a maximum power capacity of grid import and a maximum power capacity of grid export. The import need not have the same capacity as the export.

The grid electricity also has an associated price in money, e.g. pence per kWh. The price is primarily a function of local electrical demand, which is assumed to be approximately proportional to the total demand on the grid. The price dependence is a simple mathematical formula first derived from a combination of the demand-duration curve for England and Wales, (Burdon 1998) and the price-duration curve for England and Wales, (Milborrow 2000). This makes the crude assumption that price is always a direct function of demand only. Furthermore, the derived demand-price relationship was fitted by a simple mathematical curve. Thus the financial data inputs to the modelling programs are very approximate and any financial results are similarly approximate. Nevertheless, they demonstrate the principle of including costs in the probabilistic modelling method, and therefore the possibility of performing an economic optimisation on the electricity system and its energy storage.

In the models, the grid prices are further modified by a price mark-up of imported electricity and a price mark-down of exported electricity, to reflect grid pricing, transmission losses, and imperfect market pricing.

2.12.2 Backup Generation

This is modelled as an infinitely variable backup generating power, up to a user-defined maximum value. At times of power deficit, grid import is always used before backup generation.

The electricity produced by backup power has a flat cost per unit, e.g. pence per kWh. In the computer programs written for this thesis, the chosen cost-per-unit of backup generation was nearly always greater than the cost-per-unit of grid electricity.

2.12.3 An Alternative Method of Modelling a Grid Connection or Backup Generation

In some weak electricity grids, the amount of intermittent renewable energy that can be connected is limited by the amount of locally generated electricity that the grid can absorb. The grid may be sufficient however, for any likely local electricity demand. In these cases, the grid connection can be modelled as the local demand. In the modelling programs, the local demand is replaced by the capacity of the local grid to absorb renewable energy, and this capacity may vary with time of day, day type and month of the year. There is then no need to model a separate grid connection, for example (Barton, Infield 2004)

Alternatively, the grid may be able to absorb the entire renewable energy supply at all times, but may not always be able to satisfy the local demand on its own. In that case, the local demand is replaced by the net local demand subtracting the grid import capacity. Again, there would be no need to model a separate grid connection.

2.12.4 Control Options for the 24-Hour Probabilistic Program and its Validation Program

Without a grid connection or backup generation, the energy store system has no real choice in its behaviour. It charges when there is a local surplus of electricity and discharges when there is a local deficit of electricity (net local power is adjusted for store losses in all cases).

However, a grid connection or backup generation gives the store control system a choice: Should any surplus electricity be directed to the store or to grid export? Should any deficit in electricity be satisfied by discharging the store, or by importing from the grid, or by running the backup generator? In a real system, the choice depends on the relative costs and opportunities of the options available.

2.12.4.1 Control Option 1

In a system that is in danger of not satisfying all its demand at all times, a sensible control system may cause the store to remain as full as possible. The backup generator and grid import will be used to charge the store whenever their capacity is not required to supply the demand directly. This will certainly be the case when the consequences of unsatisfied demand are severe. A simple example of such a system is an uninterruptible power supply (UPS).

In the 24-hour probabilistic program, this type of control is called ‘control option 1‘. The store stays as full as possible and the store size is calculated on the boundary of: Wind power + solar power + max. grid import + backup generation = demand + losses, averaged over 24 hours. This is where the store is considered to be ‘balanced’, section 2.11.2. Electricity is only exported to the grid when the store is full, or when the store cannot absorb the entire electricity surplus.

2.12.4.2 Control Option 2

This control option attempts to combine the best of options 1 (above) and 3 (below). The store’s management system uses perfect weather forecasting to calculate the net energy surplus or deficit to the store over the following 24 hours. The management system then uses grid export, grid import and backup generation appropriately to balance the power flows to or from the store. Provided the net surplus or deficit is not too large, the store neither goes completely empty nor completely full. The risk of power cuts and the risk of power curtailment are both minimised at the same time.

The 24-hour probabilistic program models this situation by constructing PDFs of wind power and solar power within each 24-hour period, given the 24-hour average conditions. Net PDFs of power to or from the store, including the store losses, are constructed for each set of weather conditions, for each day type and month of the year. In order to simplify the control system and the program software, the grid import or export power is held constant within each 24-hour period. Backup generation is only used if grid import is insufficient to balance the net store power flows, in which case backup generation is also held constant during each 24-hour period. If the grid import capacity and backup generation are insufficient to balance a deficit of electricity over 24 hours, some power cuts will be unavoidable. In this case, the store reverts to 'control option 1' and tries to stay as full as possible at all times in order to minimise those power cuts. If however, the grid export capacity is insufficient to balance a surplus of electricity, some energy curtailment will be unavoidable. Then the store reverts to 'control option 3' and tries to stay as empty as possible at all times in order to minimise that energy curtailment.

The 24-hour time-step validation program simulates this by looking ahead 24 hours and anticipating the net power surplus or deficit, together with the current state-of-charge of the store. The time-step program uses perfect weather forecasting based on the real time series of weather data. It re-evaluates the necessary grid export, grid import and backup generation if necessary in order to achieve a balanced store over the following 24 hours. It re-evaluates the situation once every hour, and so adjusts the grid import or export power and the backup generation once per hour. Note that the time-step program always aims to have a store that is half full (i.e. half way between full and empty) in 24-hours time and adjusts the grid and backup generation accordingly.

The store size in control option 2 is calculated on the entire region in which the electricity surplus or deficit is balanced by the grid and backup generation; see section 2.11.2.1 and table 2.9. Option 2 produces a much broader range of conditions in which the energy capacity of the store comes into play.

As stated earlier, in control option 2, the store is used to make the grid import and export and the backup generation as constant as possible. This is perhaps an artificial situation; a real store management system is more likely to vary the grid import and export power and backup generation in response to the local demand, the state-of-charge of the store and time-varying electricity prices. This strategy would make more use of the flexibility of the grid and of backup generation and therefore make less use of the store and require a smaller store size. Nevertheless, the assumption of constant grid import/export and backup generation is a good starting point, especially if local demand is approximately proportional to national demand and therefore the price of grid electricity varies monotonically with both local and national demand. – If the local demand for electricity increases at certain times of the day, there is no point in increasing the import or decreasing the export of power if the whole network is experiencing the same effect at the same time and so the price of electricity will also be higher at those times. Similarly, if the whole grid is experiencing similar weather effects at the same time of day, and if intermittent renewable energy is evenly distributed throughout the grid, then there is no point in changing the grid import/export in response to local renewable generation. A constant grid import or grid export power will tend to minimise the requirement for and therefore the cost of central electricity generation and backup generation.

2.12.4.3 Control Option 3

In a system that is in danger of curtailing excess renewable energy but can generally supply the local demand, a sensible control system may cause the store to remain as empty as possible. The store will be discharged whenever possible to export electricity to the grid. An example of such a system would be a weak electricity grid in which the level of intermittent renewable generation that can be connected is limited by the voltage rise, for example (Barton, Infield 2004).

In the 24-hour probabilistic program, this type of control is called 'control option 3'. The store stays as empty as possible and the store size is calculated on the boundary of: Wind power + solar power = demand + max.

grid export capacity + losses, averaged over 24 hours. This is where the store is considered to be ‘balanced’, section 2.11.2. Electricity is imported from the grid only when the store is empty, or when the electricity deficit exceeds the store discharge power rating of the store. Backup generation is then only used as a last resort, when grid import is insufficient to accommodate the deficit.

2.13 Matlab Programs

The probabilistic method has been embodied in two computer programs, each validated by a time-step program:

2.13.1 The Wind-Only Probabilistic Program

This works on any time scale from minutes to years but only models variable wind-generated electricity with electrical demand and no solar power. Optionally, this program can be run with time-varying demand using 3 different times of day (peak, mid-rate and off peak) and two seasons (summer and winter). However, this program has to assume that wind speed variations have no periodic dependence on time of day or season. If run with a time-varying demand, the program assumes that the averaging time scale is at least 24 hours but shorter than 6 months.

2.13.2 The 24-Hour Probabilistic Program

This models variable wind power, variable solar power and truly time-varying electrical demand but only over a time-scale of 24 hours. It models energy stores that smooth variations within each period of 24 hours but not day-to-day variations in power.

Future work may enable a program to model variable wind power, solar power and truly time-varying demand on all time scales but such a program was not achievable within the time scale of this PhD project.

Both the probabilistic programs use a spectral approach to wind power variations. The probabilistic 24-hour program also uses a spectral approach to solar power variations. Both use the spectrum integration equations of section 2.5 and both use the same generic turbine power curve, section 2.6.4.

2.13.3 Time Step Validation Programs

Each probabilistic program has been validated by a time step program with exactly the same modelling functionality. For example, the time step programs use the same wind turbine power curve, the same energy store assumptions;

the 24-hour time step program has the same control options 1, 2 and 3. One important difference between the probabilistic and time step models is that the probabilistic programs use a storage time scale as an input, whereas the time step programs take an energy storage capacity in kWh. This storage capacity has to be first calculated by the appropriate probabilistic program modelling the same electricity system.

3 Results and Discussion

As described in section 2.11.7, the probabilistic model results are validated by comparison with time stepping programs. Section 3.1 describes the results of the wind-only programs that model storage on all time scales. Section 3.2 describes the results of the 24-hour modelling programs that model wind power, solar power, variable demand and energy storage with cycle times of 24 hours.

3.1 The Wind Only Programs

The following graphs show the effect of changing various input parameters while the store operating time scale, T is varied between a few seconds and several years.

3.1.1 Datum Case

The datum case for all results of the wind only program is an energy store of 100% efficiency and unrestricted charge and discharge capacities. That is, the maximum charge rate is greater than the largest possible power surplus and the maximum discharge rate is greater than the largest possible power deficit. The electricity system connected to the store consists of 1MW of wind power capacity (typical of one large modern wind turbine) and 400kW of constant demand. The average wind speed is 8m/s (typical of a good wind site) and the standard turbine power curve is used, see section 2.6.4. This curve has a cut-in wind speed of 3m/s and a rated wind speed of 13m/s. The constant demand of 400kW is modelled as a grid connection that allows up to 400kW of wind power output to be absorbed by the grid.

The wind speed power spectrum used in most runs of the probabilistic program is a generic average of spectra derived from various measurement sites around the British Isles. The wind speed data used in the time step program is a scaled wind speed time series measured at Rutherford Appleton Laboratory (RAL) between 1999 and 2002 inclusive, unless otherwise stated.

All the wind speeds have been scaled such that the time-series average is 8m/s.

Figs. 3.1 to 3.7 show the results of the datum comparison between the probabilistic and time step programs over time scales from 5 seconds (0.0014 hours) to 4 years (35,000 hours). Most of the results are plotted against store size in kWh rather than time scale, but the store size was derived from the time scale by the probabilistic program.

In addition, the probabilistic program has been run using a wind speed spectrum derived purely from RAL data measured between 1994 and 1998, and the stepping program has been run using the wind speed time series measured between 1994 and 1998. These extra results are also shown on the datum plots, figs. 3.1 to 3.7 in order to illustrate the improved accuracy or otherwise resulting from the use of a spectrum derived from the same time series as used in the time step program.

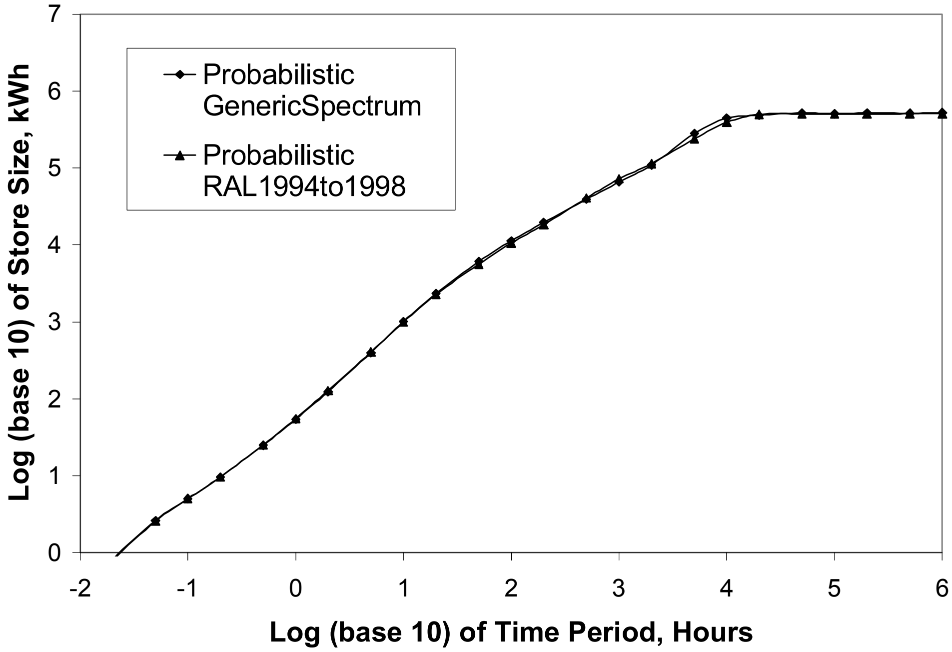

Figure 3.1 Calculated Store Size vs. Store Cycle Time

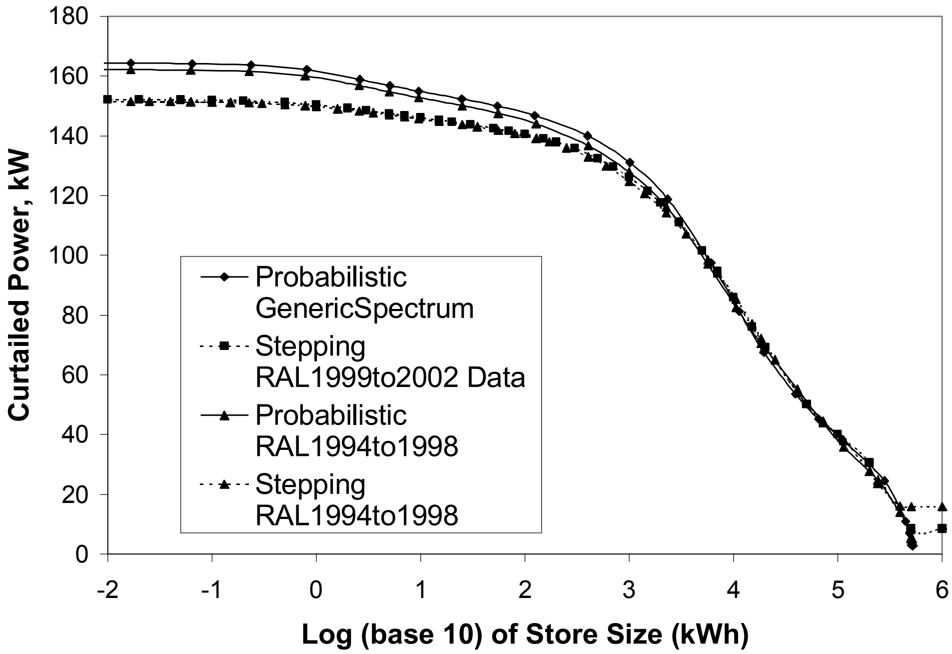

Figure 3.2 Average power curtailed due to store being full

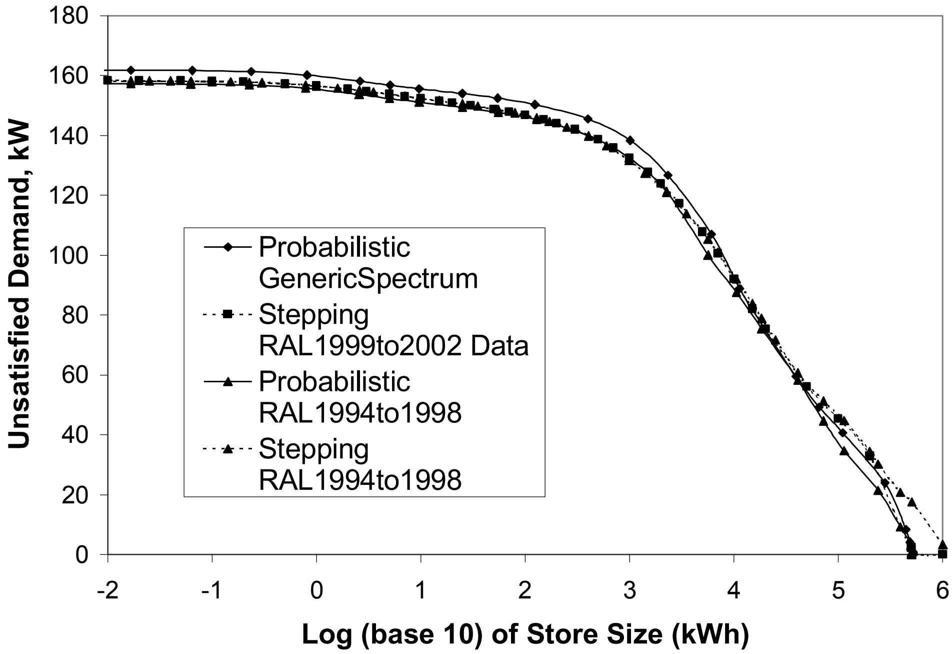

Figure 3.3 Average demand unsatisfied due to store being empty

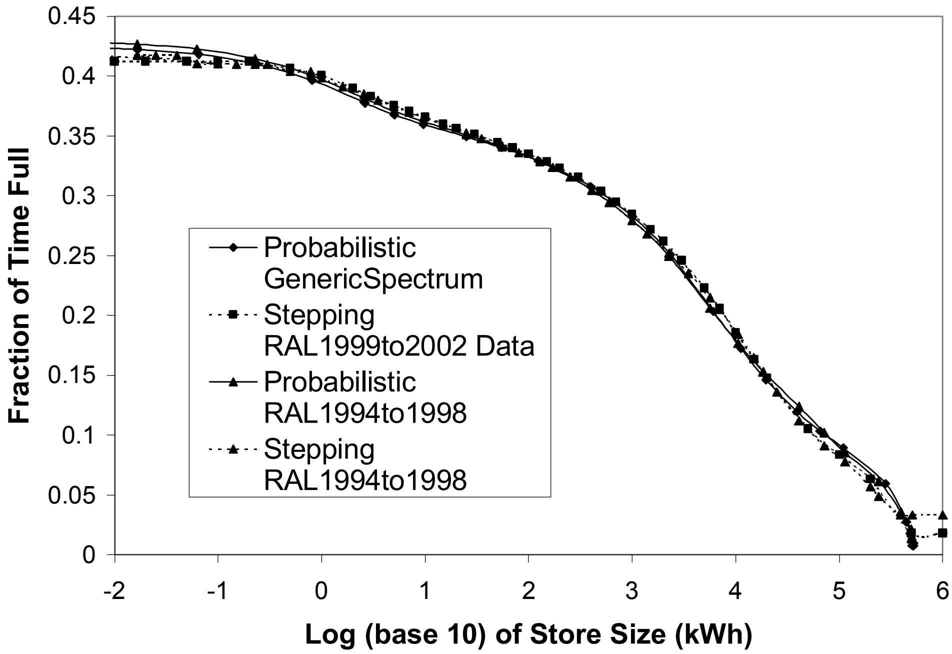

Figure 3.4 Fraction of time that the store spends full

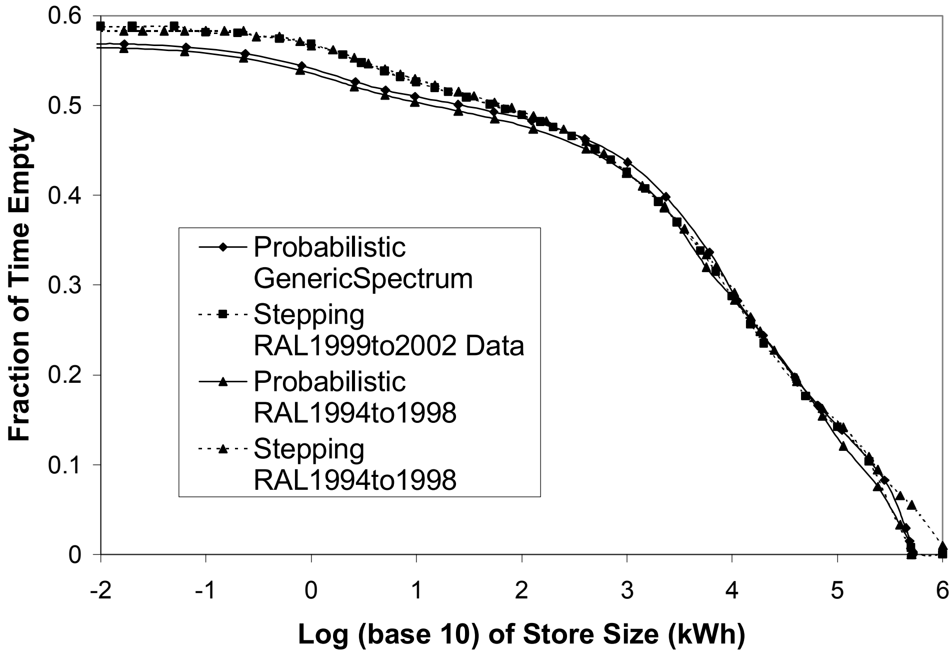

Figure 3.5 Fraction of time that the store spends empty

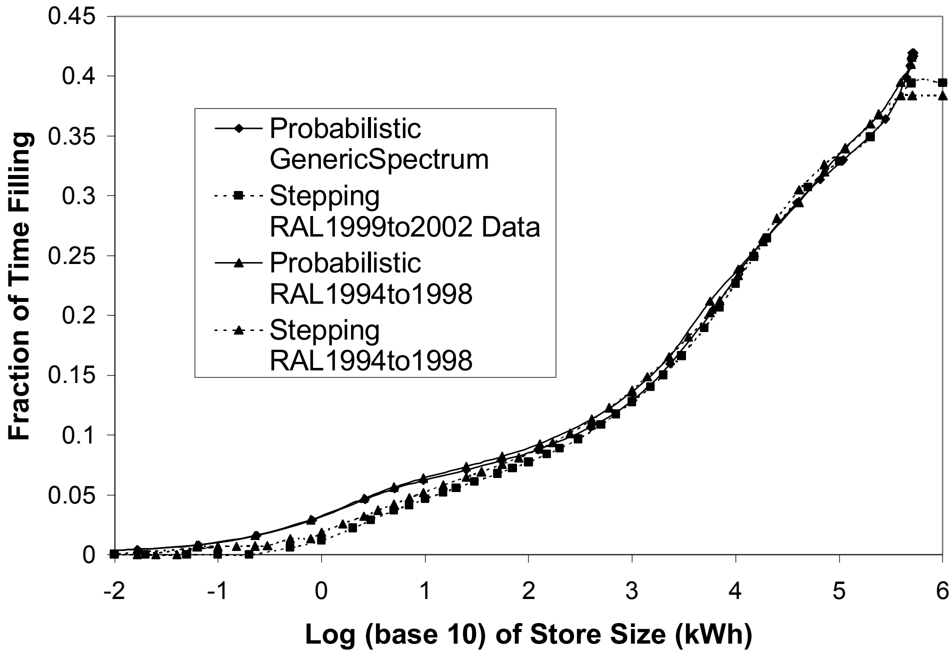

Figure 3.6 Fraction of time that the store spends filling

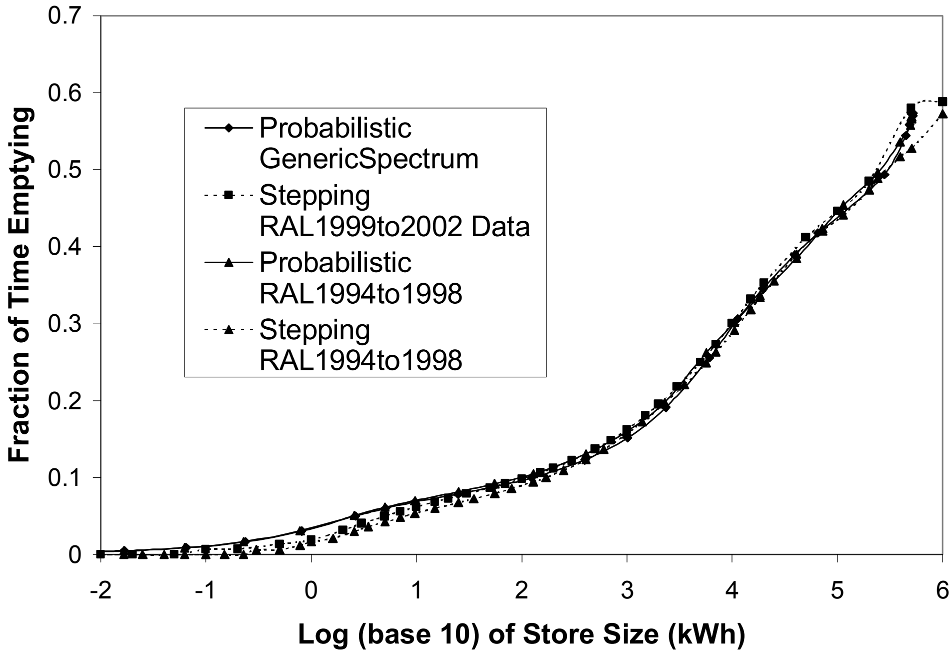

Figure 3.7 Fraction of time that the store spends emptying

In the probabilistic results, the graphs show that very little difference is made by using the RAL 1994 to 1998 spectrum instead of the generic wind spectrum. This gives confidence that the spectrum of wind speed varies little throughout the British Isles and that the generic spectrum is generally applicable, throughout Britain at least.

In the stepping results, very little difference is made by using the 1994 to 1998 time series instead of the 1999 to 2002 time series. Note though that both time series have been scaled to the same average wind speed of 8m/s.

3.1.2 Varying Demand

Figs. 3.8 to 3.14 illustrate the effect of varying the demand (or the maximum wind powered generation that a grid can accommodate). The demand is first set at approximately half the average supply, i.e. 200kW and then at 1.5 times average supply, i.e. 600kW. In these figures, the results are compared with the datum case in which average supply is approximately balanced with demand, which was a constant 400kW. From now on, all probabilistic results use the generic wind spectrum as their starting point and all time step results use the 1999 to 2002 wind speed time series.

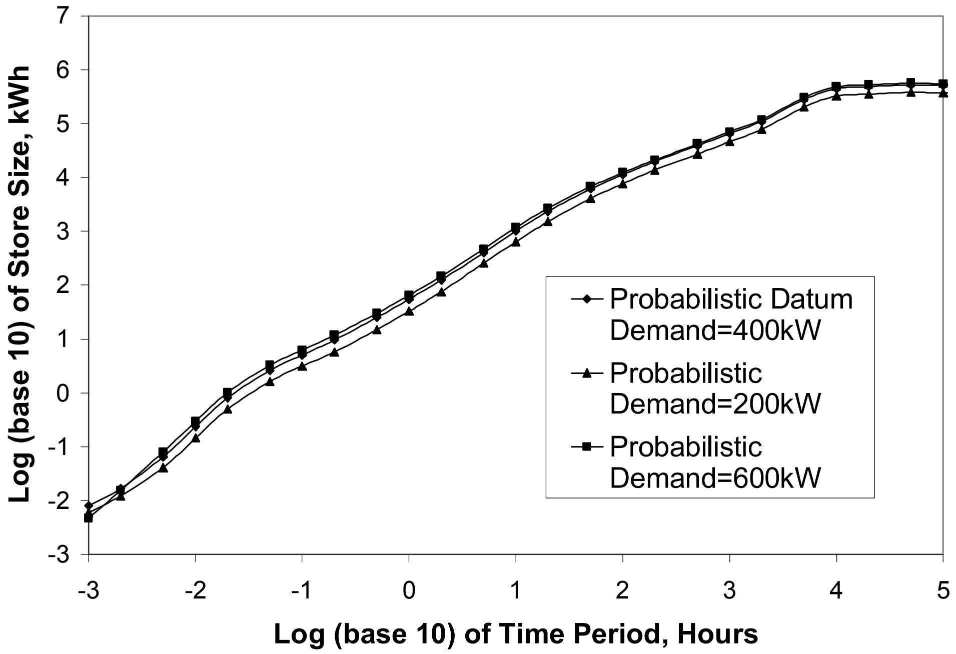

Figure 3.8 Calculated Store Size vs. Store Cycle Time

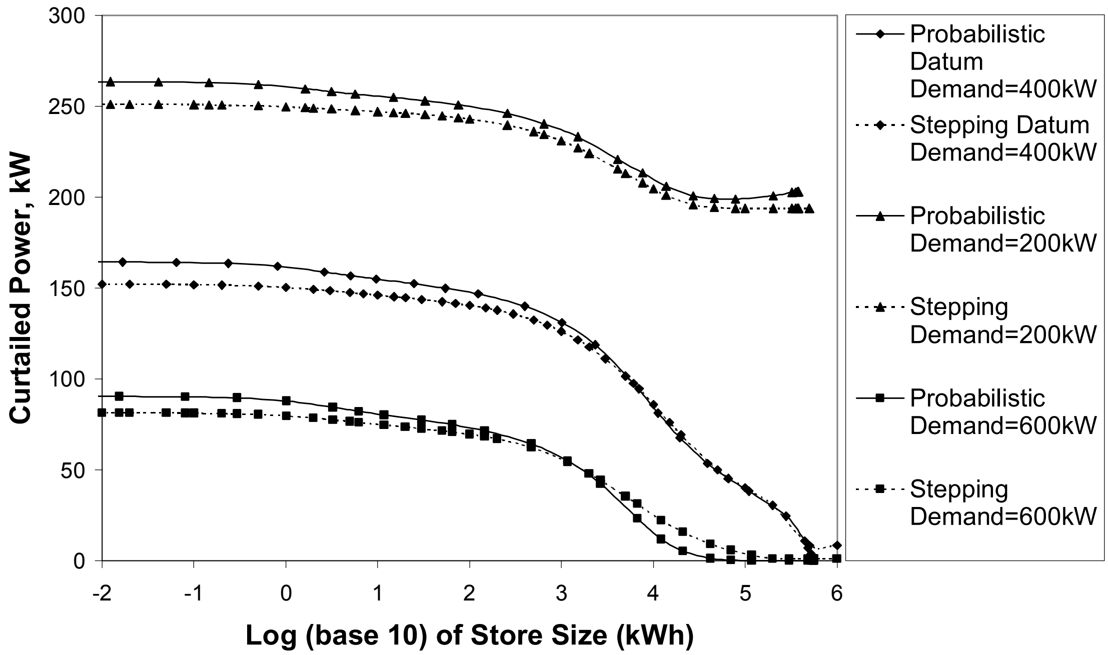

Figure 3.9 Average power curtailed due to store being full

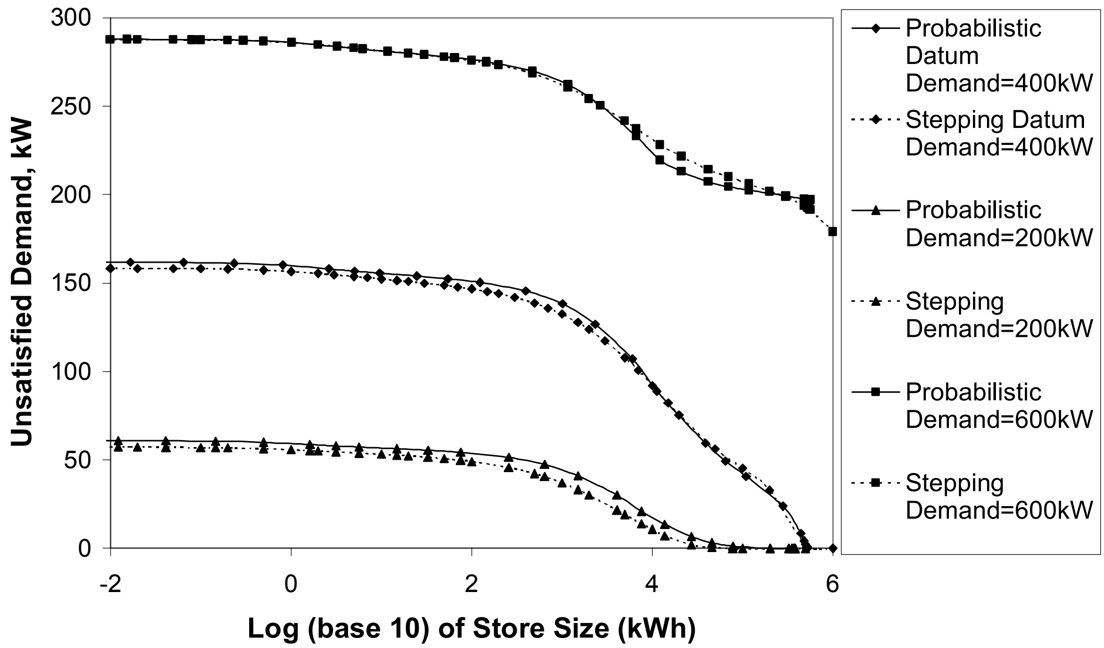

Figure 3.10 Average demand unsatisfied due to store being empty

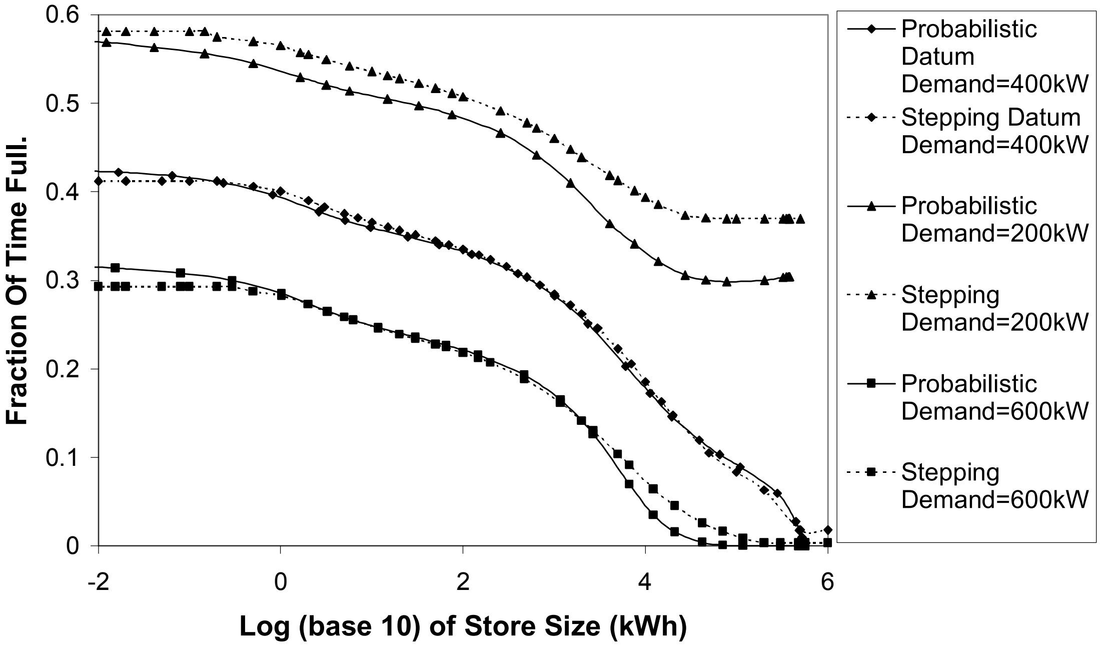

Figure 3.11 Fraction of time that the store spends full

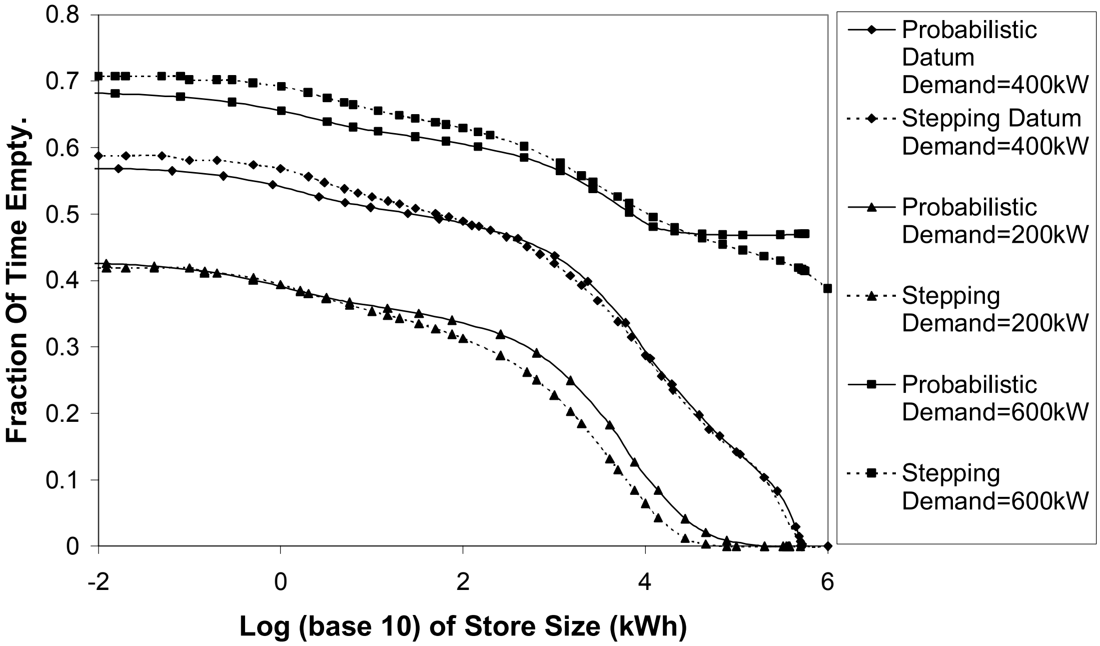

Figure 3.12 Fraction of time that the store spends empty

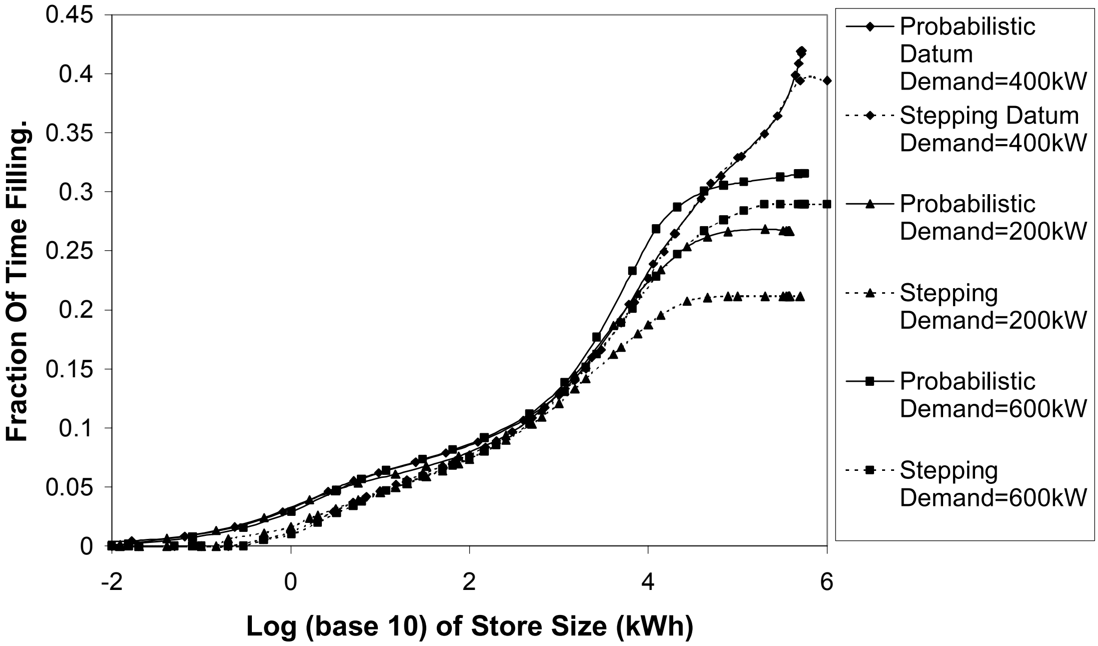

Figure 3.13 Fraction of time that the store spends filling

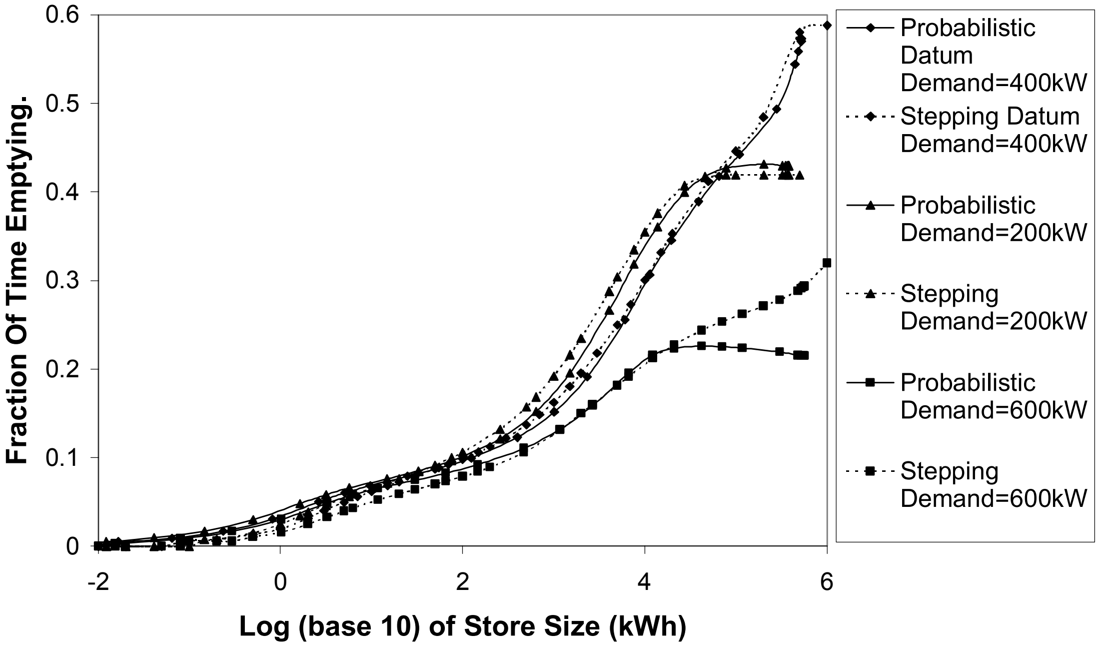

Figure 3.14 Fraction of time that the store spends emptying

When demand is half the supply, the store spends a lot of time full, fig. 3.11 even when the store is large in size. In every case, there are times when some power is always curtailed due to the store being full, fig. 3.9.

Conversely, when demand is 50% larger than the average supply, the store spends a lot of time empty, fig. 3.12, even when the store is large. In every case, there are times when some demand is left unsatisfied due to the store being empty, fig. 3.10.

The different levels of demand make only a small change to the store size, when plotted on the logarithmic scale, fig. 3.8.

3.1.3 Varying Charge and Discharge Rates

In figs. 3.15 to 3.23, the first set of lines represents the datum case again. The second set of lines represents a store with its maximum discharge rate restricted to 200kW while the maximum charge rate is effectively unrestricted at 1000kW. The third set of lines represents a store with its maximum charge rate restricted to 200kW while its maximum discharge rate is unrestricted at 1000kW. In each case, the turbine power capacity was 1MW, the average turbine power was approximately 400kW and the demand was a constant 400kW.

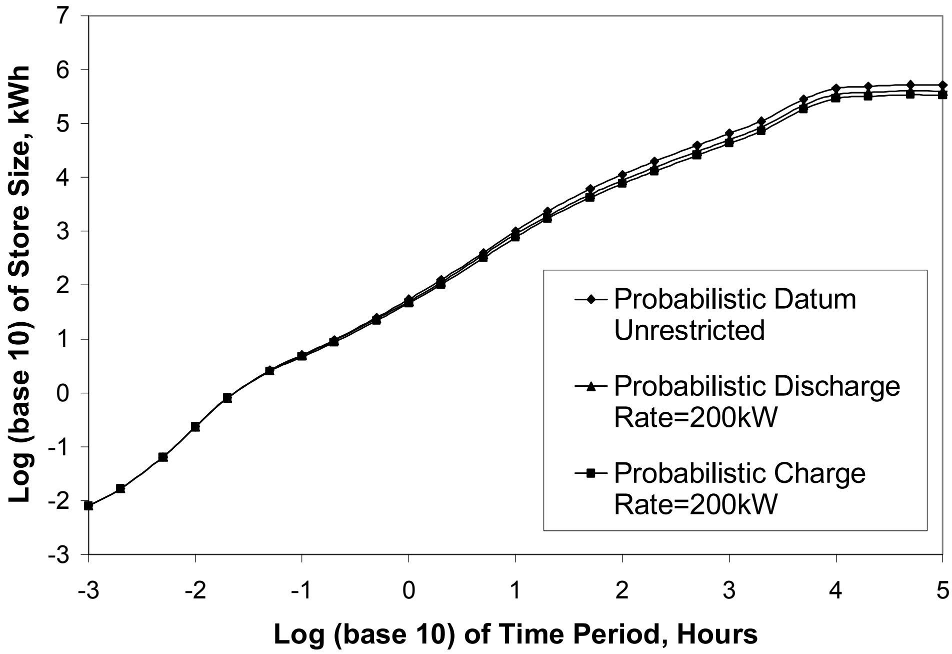

Figure 3.15 Calculated Store Size vs. Store Cycle Time

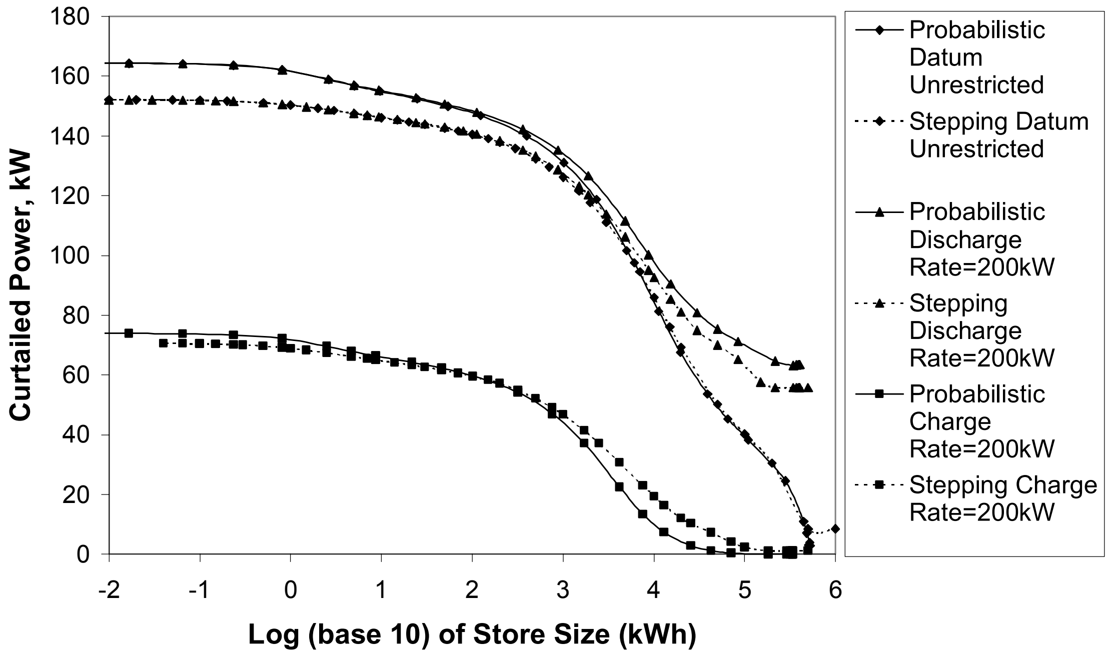

Figure 3.16 Average power curtailed due to store being full

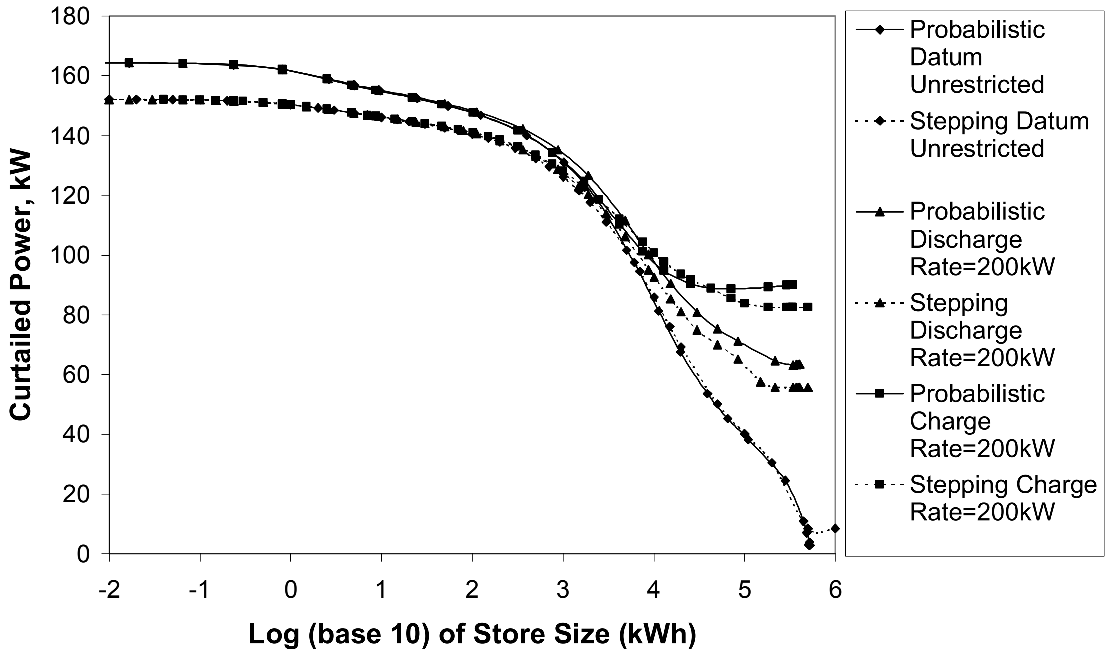

Figure 3.17 Average total power curtailed

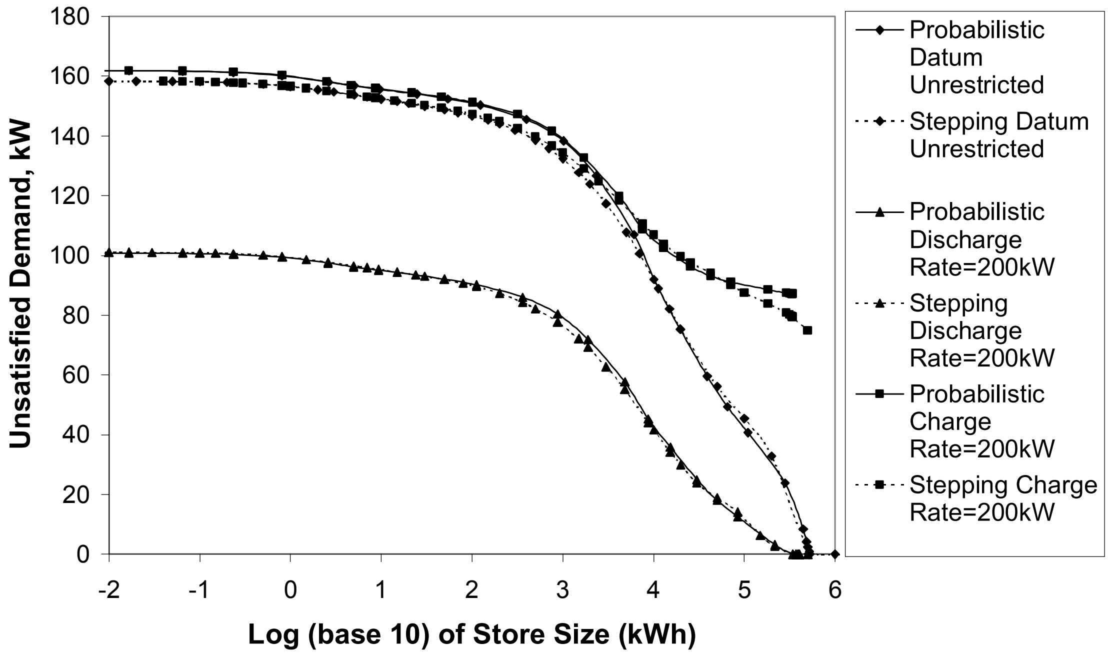

Figure 3.18 Average demand unsatisfied due to store being empty

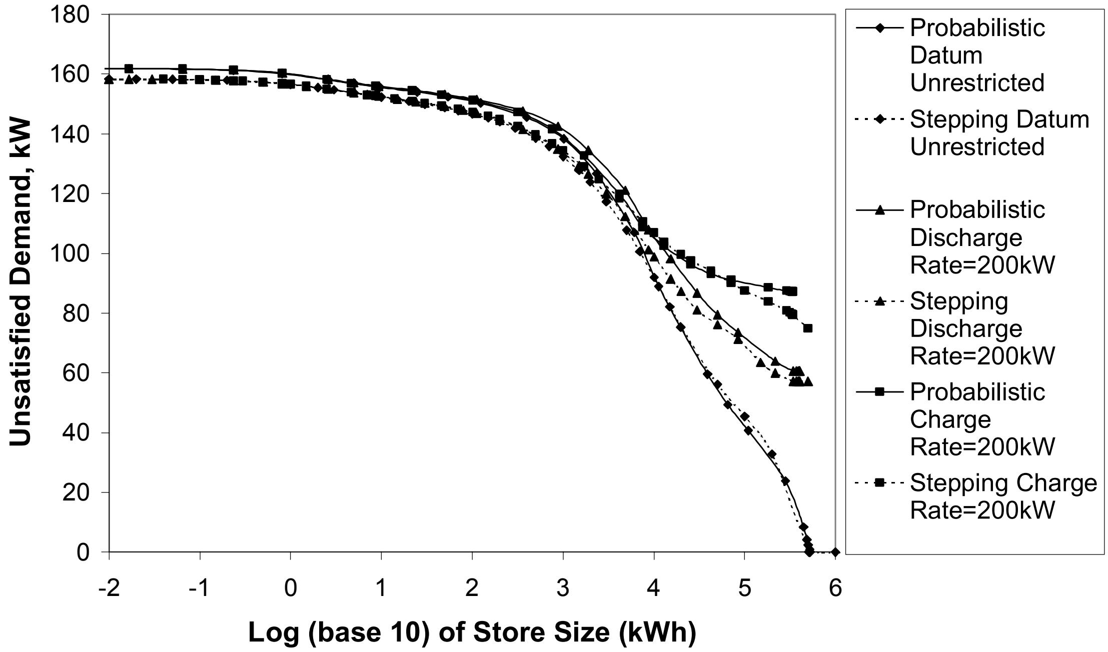

Figure 3.19 Average total unsatisfied demand

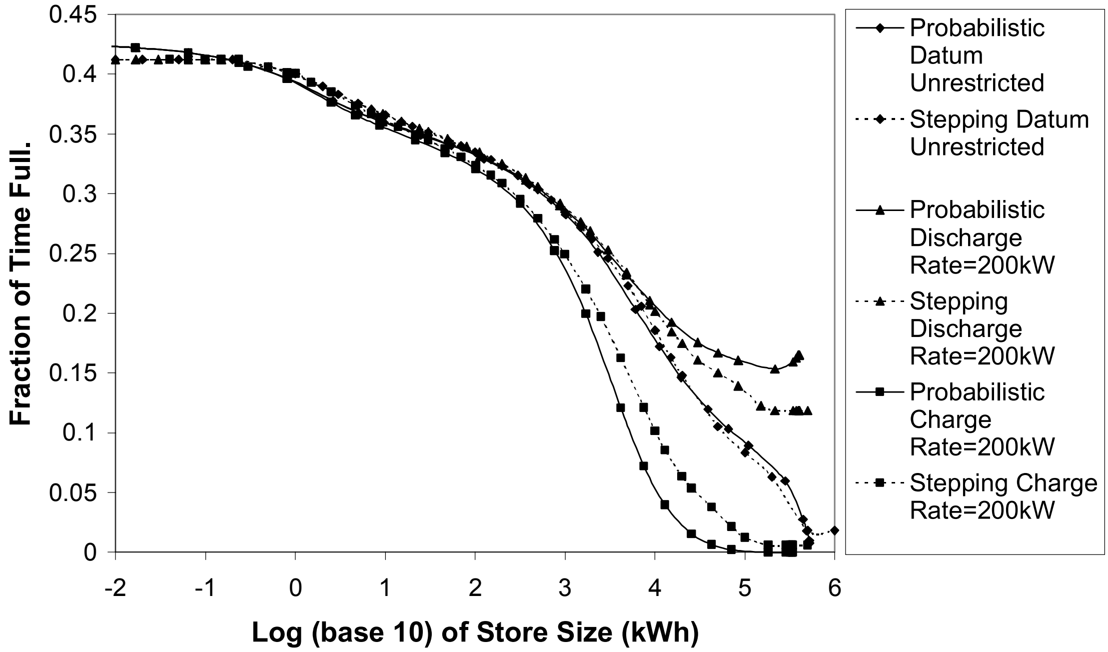

Figure 3.20 Fraction of time that the store spends full

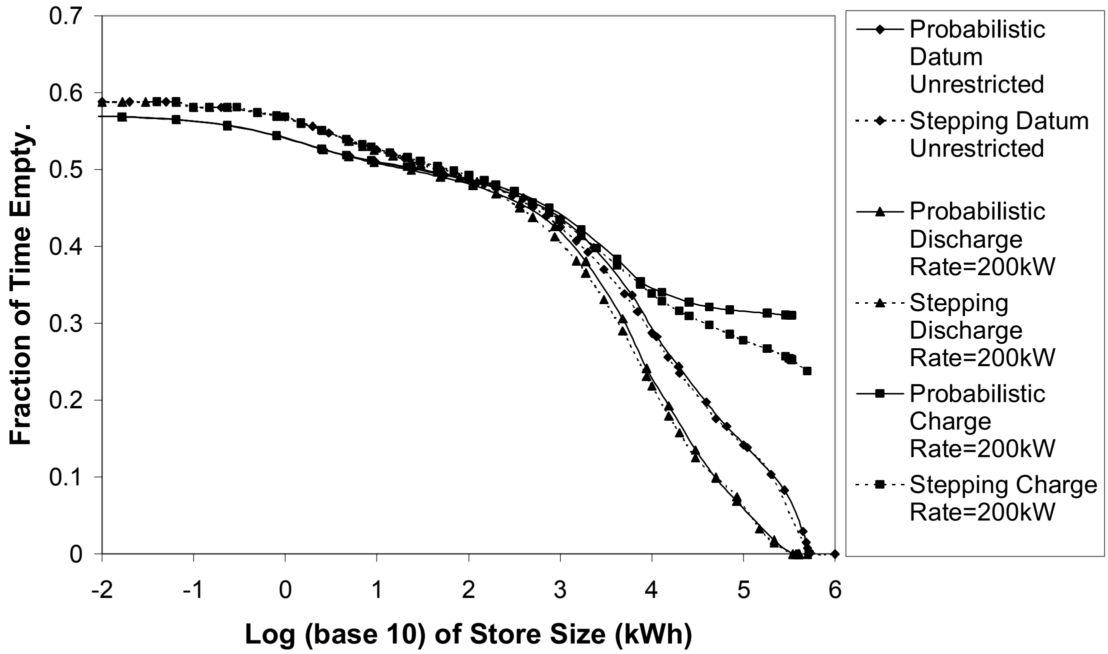

Figure 3.21 Fraction of time that the store spends empty

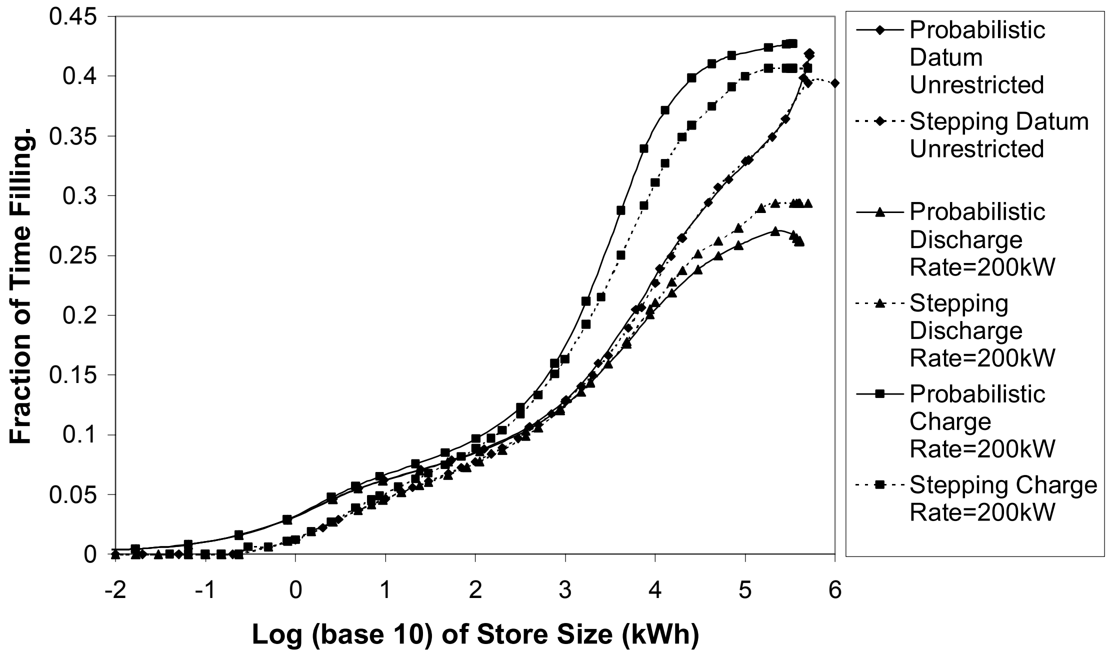

Figure 3.22 Fraction of time that the store spends filling

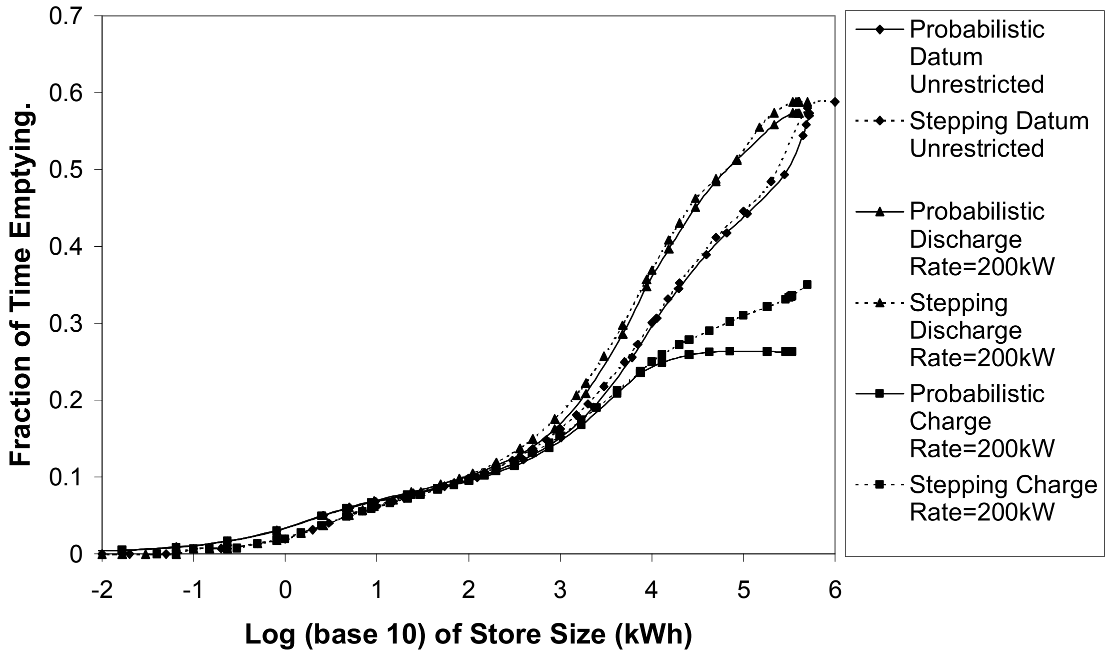

Figure 3.23 Fraction of time that the store spends emptying

When the charge rate is restricted to 200kW, less power enters the store on average than in the datum, unrestricted case. Therefore, the store spends an increased fraction of time filling, at a slower rate, fig. 3.22 and a reduced fraction of time full, fig. 3.20. Because the store is, on average, less full than the unrestricted case, it also empties more quickly, spends less time emptying, fig. 3.23 and more time empty, fig. 3.21.

Conversely, when the charge rate is unrestricted but the discharge rate is restricted to 200kW, less power leaves the store than in the datum case. The store spends more time emptying at a slower rate, fig. 3.23 and a reduced fraction of time empty, fig. 3.21. Because the store is on average more full than in the unrestricted case, it also fills more quickly, spends less time filling, fig. 3.22 and more time full, fig. 3.20.

However, all these effects only become apparent in stores with relatively long cycle times and large store sizes. At shorter time scale and smaller stores, the store spends most of its time either full or empty, and the altered charge rate or discharge rate does not substantially change this pattern.

A restricted charge or discharge rate reduces the energy entering and leaving the store in a given cycle. This reduces the required store sizes, fig. 3.15 but only when the cycle time is long. At short time scales, the variation in wind power within the cycle time, T, and therefore the size of the power surplus or deficit, is generally smaller than the maximum charge or discharge capacity. Longer time scales are associated with greater variations in power and therefore greater required store charge and discharge rates.

Looking at power flows, when the charge rate is restricted to 200kW, an average of 81kW (time step model) or approximately 90kW (probabilistic model) are curtailed due to the power surplus exceeding the charge rate, regardless of the energy capacity of the store. The restricted charge rate reduces the power curtailed due to the store being full, fig. 3.16. It increases the total curtailed power, but only at large store sizes, fig. 3.17.

Similarly, when the discharge rate is restricted to 200kW, an average 57.2kW (time step model) or approximately 62kW (probabilistic model) of demand are left unsatisfied because the power deficit exceeds the discharge rate, regardless of the energy capacity of the store. This reduces the unsatisfied demand due to the store being empty, fig. 3.18. It increases the total unsatisfied demand, but only at large store sizes, fig. 3.19.

3.1.4 Varying Store Efficiency

In this thesis, the store efficiency is defined as the ratio of energy available from the store to energy needed to charge the store. In these first, wind-only models, parasitic losses are ignored. The store efficiency has been set to 100% (datum case), 70% (typical of a chemical energy store) and 30% (more typical of a hydrogen energy storage system). The results are shown in figs. 3.24 to 3.31.

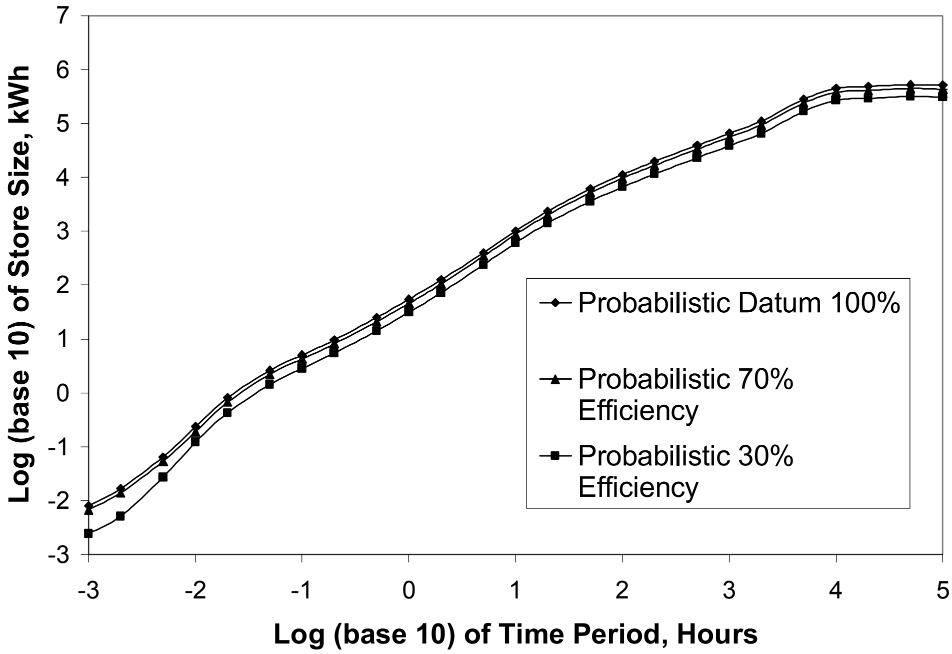

Figure 3.24 Calculated Store Size vs. Store Cycle Time

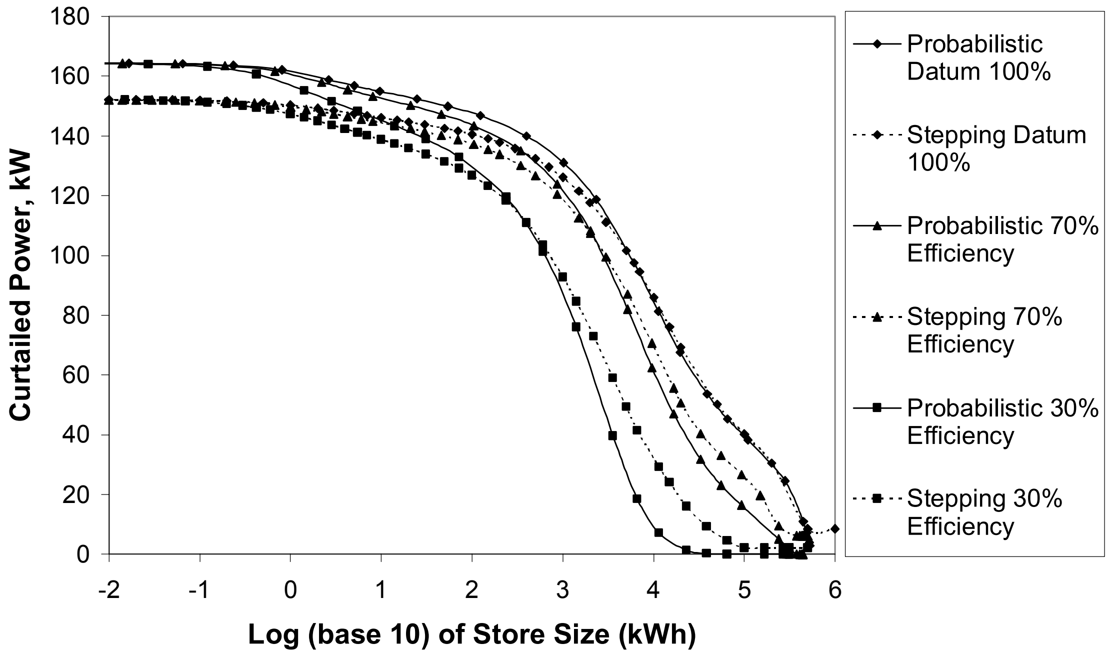

Figure 3.25 Average power curtailed due to store being full

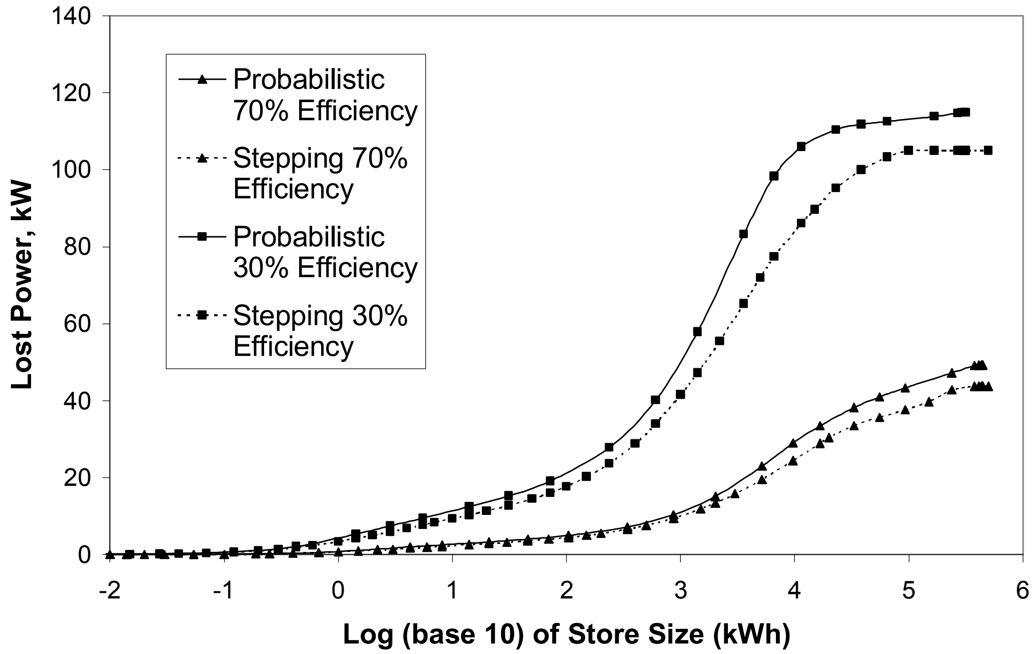

Figure 3.26 Average power lost to store inefficiency

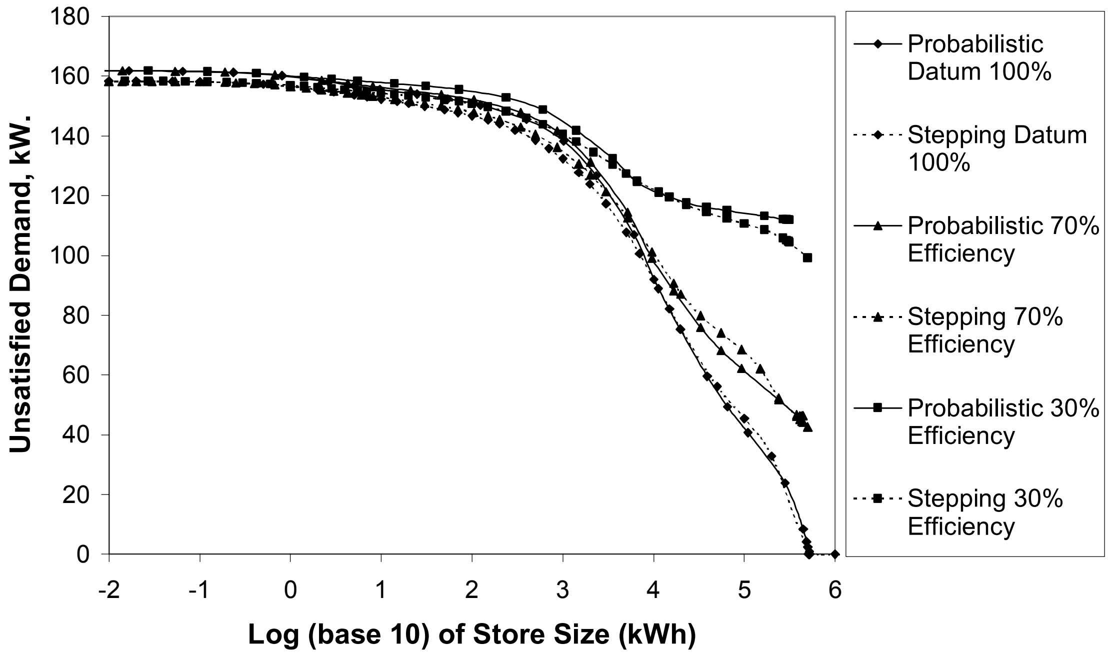

Figure 3.27 Average demand unsatisfied due to store being empty

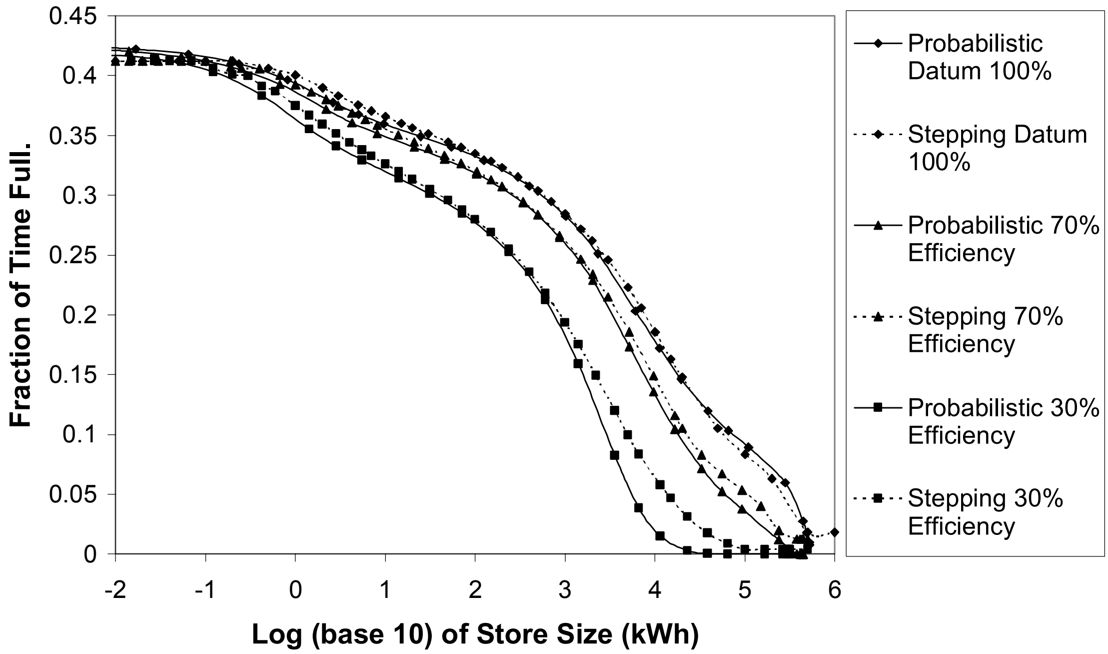

Figure 3.28 Fraction of time that the store spends full

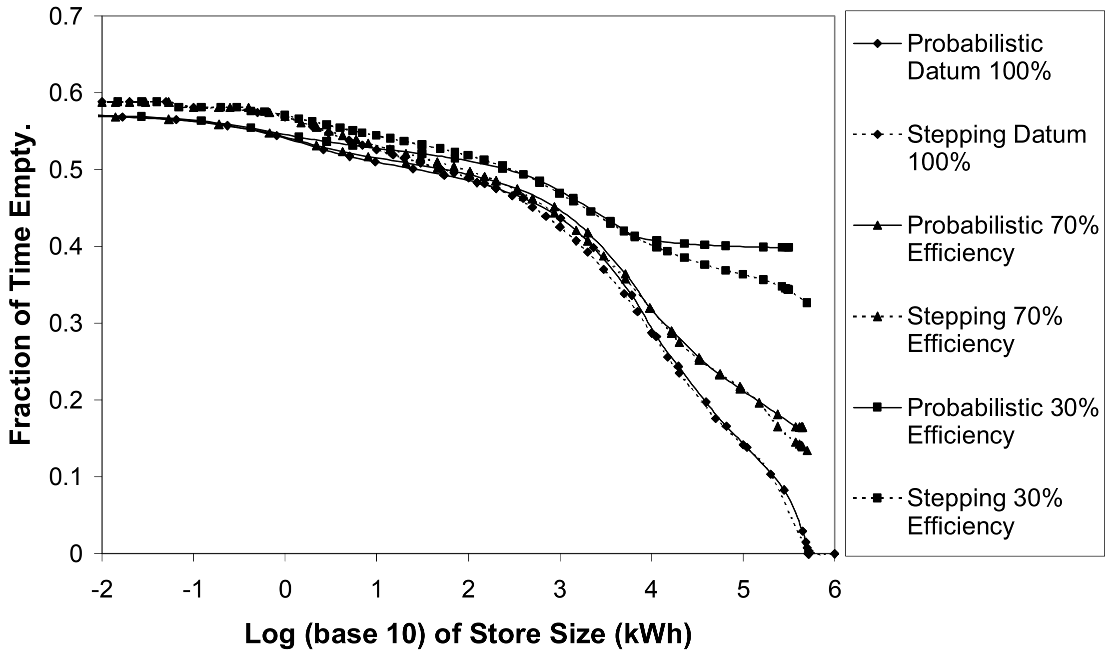

Figure 3.29 Fraction of time that the store spends empty

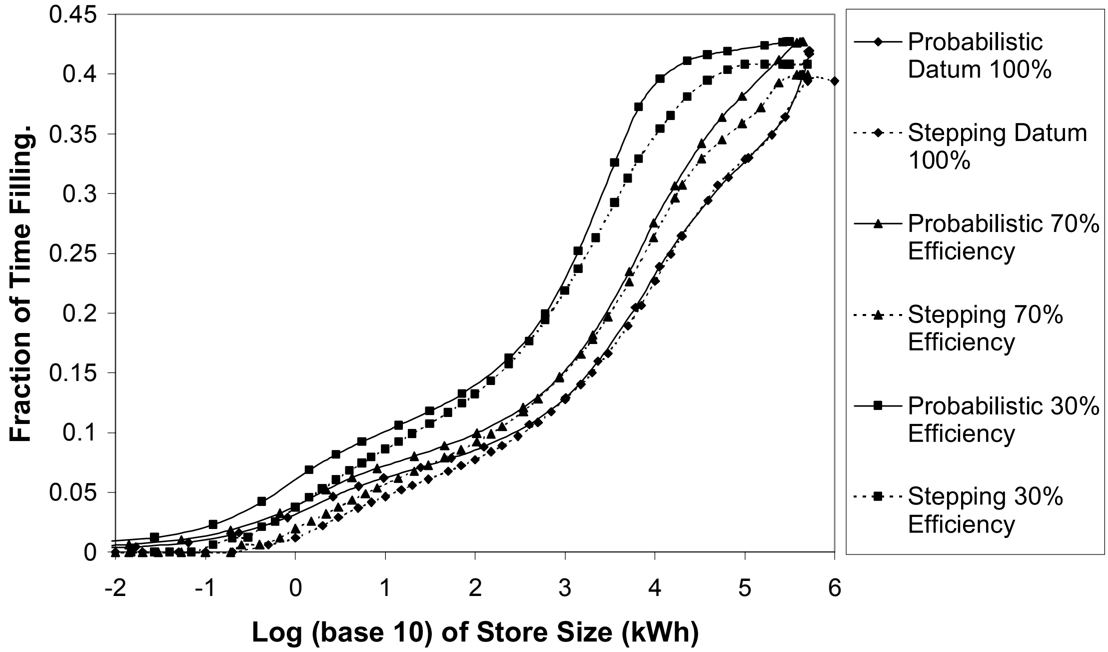

Figure 3.30 Fraction of time that the store spends filling

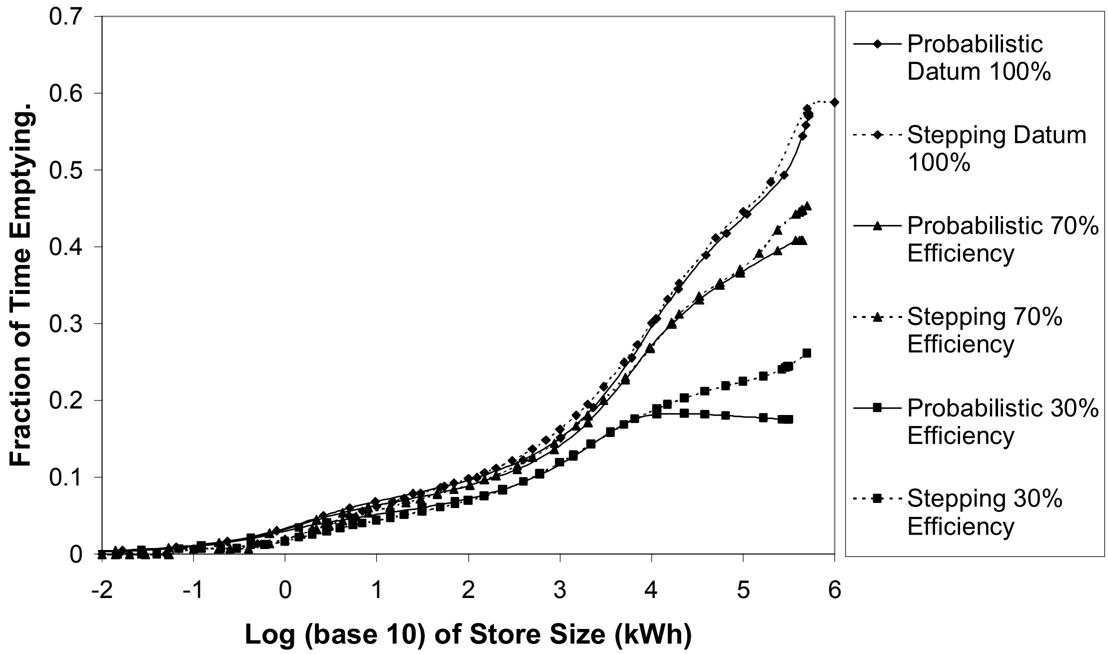

Figure 3.31 Fraction of time that the store spends emptying

When the store efficiency is reduced to 70% or 30%, less power enters the store on average than in the datum, 100% efficient case. Note that inefficiency is assumed to remove energy before it enters the store. Therefore, the store spends an increased fraction of time filling, at a slower rate, fig. 3.30 and a reduced fraction of time full, fig. 3.28. Because the store is, on average, less full than the unrestricted case, it also empties more quickly, spends less time emptying, fig. 3.31 and more time empty, fig. 3.29. As with maximum charge rates, all these effects only become apparent in stores with relatively long cycle times and large store sizes. At shorter time scale and smaller stores, the store spends most of its time either full or empty, and the reduced efficiency does not substantially change this pattern.

A reduced efficiency reduces the effective store energy capacity at a given time scale, fig. 3.24. A similar percentage reduction occurs on all time scales of storage, as the useful power flows out of the store are reduced by the inefficiency.

Note that the power lost to inefficiency, fig. 3.26 is all accounted as the store is charged, not as it is discharged. How the losses are accounted makes no difference to the actual fractions of time spent full, empty, filling or emptying, nor to the physical size of the store; it is merely an accounting effect. The energy capacity of the store is reduced because it is calculated on the energy that can be drawn from it, not the energy that can be directed into the store.

The reduced efficiencies cause the store to be less full on average than the datum case. There is less energy available to supply the demand at times of power deficit, fig. 3.27 and less energy is curtailed due to the store being full, fig. 3.25.

3.1.5 Comparison of Probabilistic with Time Step Methods

All the above graphs show good agreement between the probabilistic and time-step models, especially in the central portions of the graphs. We must exclude the very longest time scales in which the store is so large that it does not have time to get full or empty over the 4-year period of the time step model. We must also exclude the very shortest time scales where spectrum aliasing effects become significant and where the store is so small that it may move from completely full to completely empty (or vice versa) in a single time step.

The probabilistic method results often agree better with the time step method when predicting curtailed power and unsatisfied demand, e.g. figs. 3.9 and 3.10, than when predicting times spent full, empty, filling and emptying, e.g. figs. 3.11 and 3.13, especially at large store sizes. This may be because the probabilistic method first calculates average power flows from the wind turbine(s), to the load, to and from the store and curtailed power. It then uses these as inputs to its estimates of time fractions spent full and empty etc. Thus the calculations of time fractions are one step more removed from input data than the power flow calculations, and incur the errors of that extra calculation step.

In the above figures, the probabilistic method agrees very well with the time step method in the datum case, but less well when power supply and power demand are not balanced, either by system component sizing or by the effect of store parameters. In modelling real electricity systems, this should not be a problem, as system components are usually sized to meet the loads and are therefore similar in magnitude.

The probabilistic method always shows similar trends to the time step method, even if it does not agree in absolute level. When studying the effect of a change in an electricity system, e.g. a small increase in demand, increase in store discharge capacity or a reduced store efficiency, the probabilistic method is therefore likely to give a good estimate of changes to system performance in a back-to-back calculation, even if absolute levels of performance are less well predicted.

Finally, the calculations of curtailed power, figs. 3.2, 3.9, 3.16, 3.17 and 3.25 show a relatively large gap between probabilistic and time step methods when the store size is small. Both methods used the same average wind speed, 8m/s and exactly the same wind turbine power curve. However, the probabilistic calculation used synthesized wind speed PDFs based on Weibull or normal distributions whereas the time step method used the real wind speed time series. The probabilistic method therefore gave a different calculated wind power PDF and different average wind power from the time step method. The time step method predicted an average wind power of 394kW whereas the probabilistic method predicted a wind power between 390kW and 403kW depending on the cycle time of the store. Appendix C discusses further the modelling of wind speed PDFs.

3.2 Time-Varying Wind Power, Solar Power and Demand – The 24-Hour Programs

These programs can only be run for store cycle times, T of 24 hours. The following graphs show results for this one cycle time, but with continuous variation in other system and store parameters. These programs also model variable store parasitic (standing) losses, grid connections and backup generation.

The datum case is again an energy store of 100% efficiency with unrestricted charge and discharge capacities, and the wind power assumptions are as before. In contrast, the electrical demand varies with time of day, weekdays and weekends, and with month of the year, see section 2.8, but the average demand is still 400kW unless otherwise stated. The datum case is a stand-alone system with no grid connection, no backup generation and no solar power although all these variables are explored in the following simulations. The datum case does not require a control strategy, see section 2.12.4 but when a grid connection or backup generation are modelled, all three control options are modelled.

3.2.1 Varying Demand

Figs. 3.32 to 3.39 show the effect of varying demand while all other parameters are held constant.

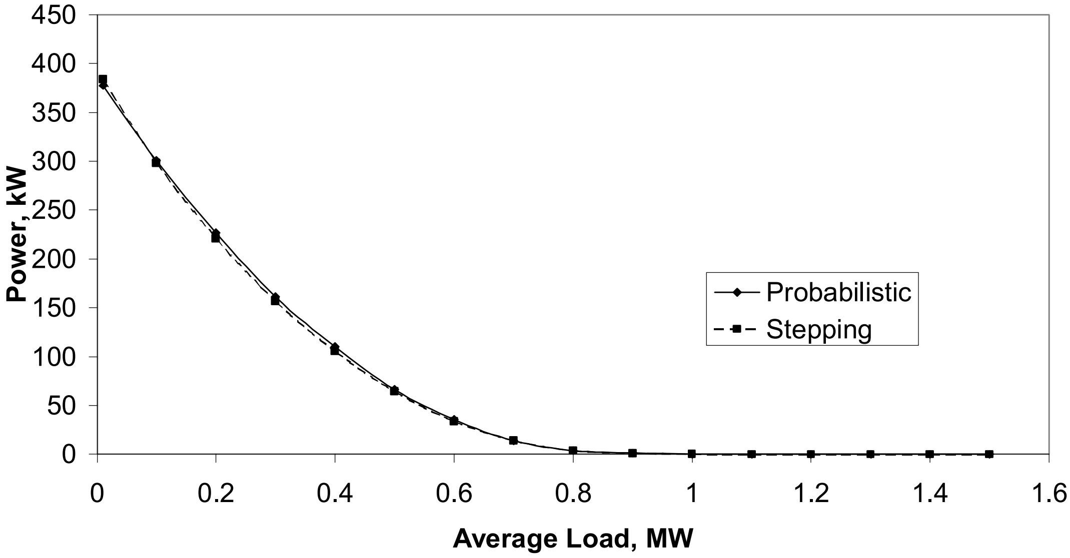

Figure 3.32 Curtailed Power Due To Finite Store Being Full

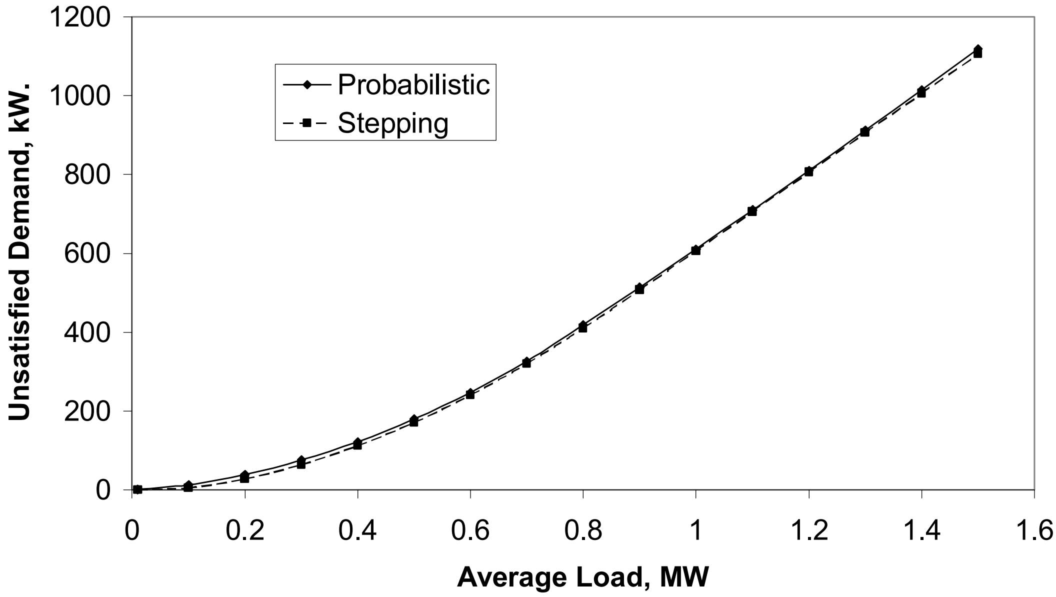

Figure 3.33 Unsatisfied Demand Due to Store Being Empty

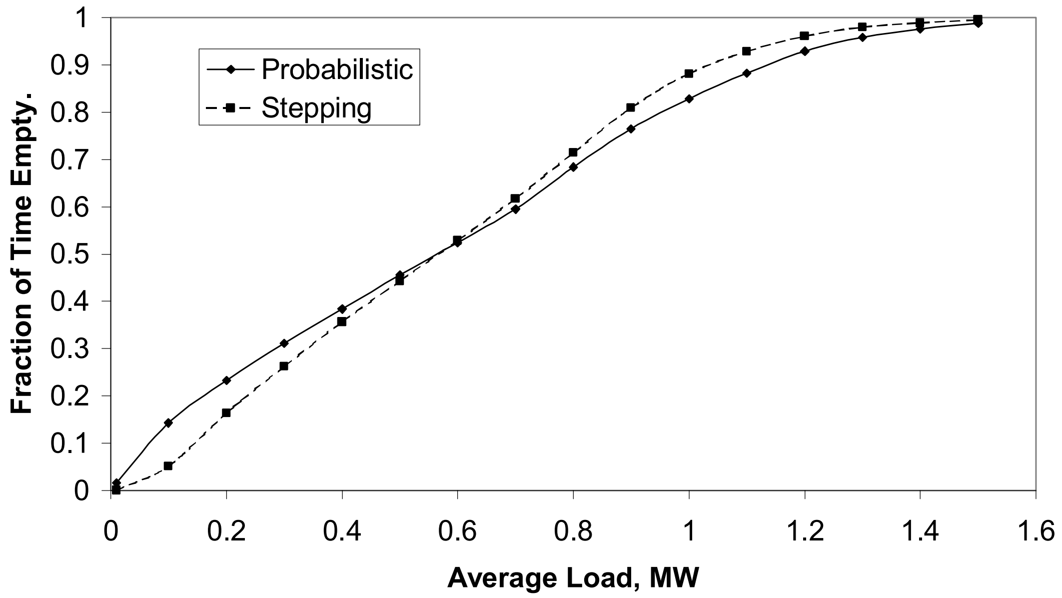

Figure 3.34 Fraction of Time that the Store is Empty

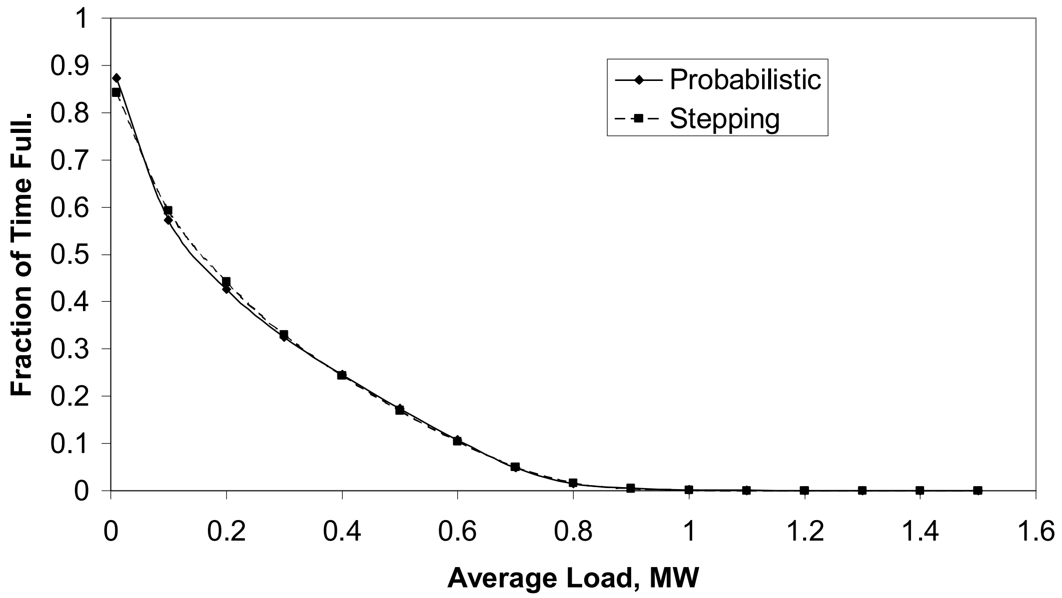

Figure 3.35 Fraction of Time that the Store is Full

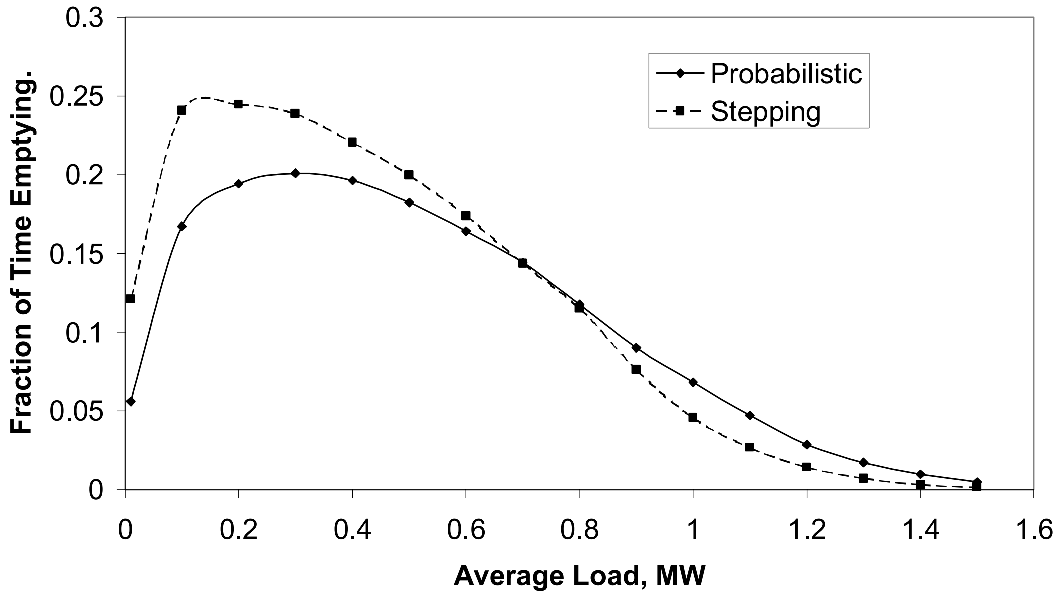

Figure 3.36 Fraction of Time that the Store is Emptying

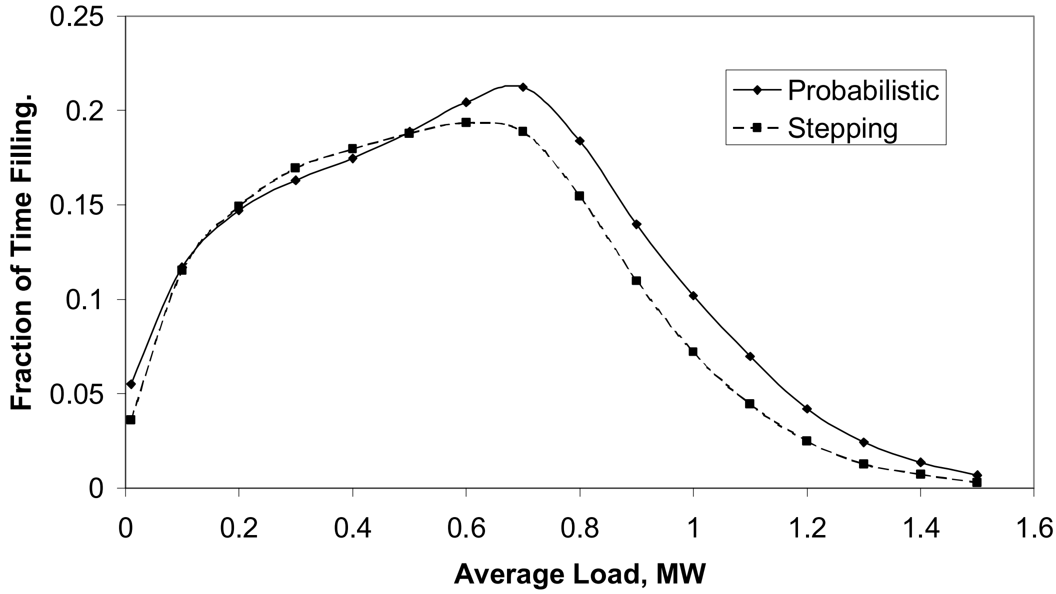

Figure 3.37 Fraction of Time that the Store is Filling

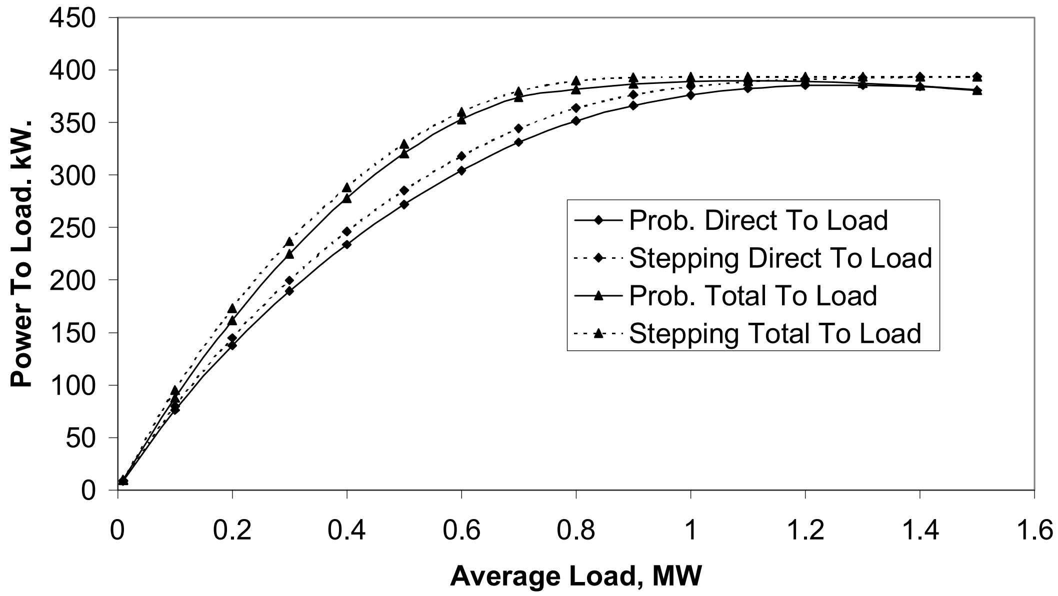

Figure 3.38 Power Supplied to the Loads With and Without the Store

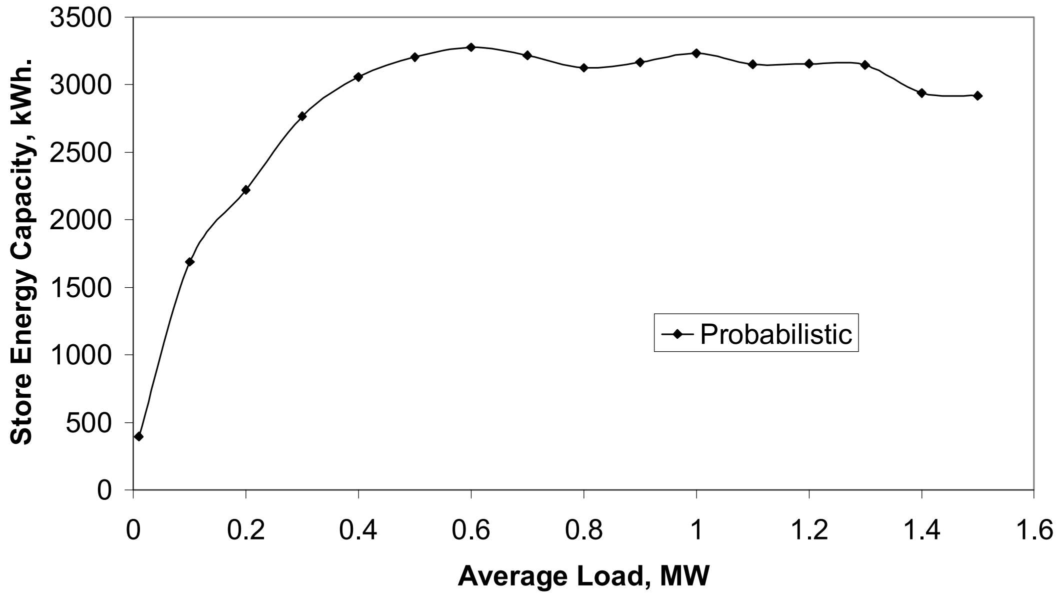

Figure 3.39 Required Store Size for a Cycle Time of 24 Hours

At low levels of demand, the store spends most of the time full, fig. 3.35, the unsatisfied demand is very low, fig. 3.33, and the curtailed power is very high, fig. 3.32. Then at high levels of demand, the store spends most of its time empty, fig. 3.34, the unsatisfied demand is high, fig. 3.33 and the curtailed power drops to zero, fig. 3.32.

The store is most useful when supply and demand are balanced, around a demand of 400kW, as indicated by the large amount of time spent emptying and filling, fig. 3.36 and fig. 3.37 and by the increase in power supplied to the demand as a result of the store, fig. 3.38.

The required store energy capacity, fig. 3.39 shows that at low levels of demand, the store is small. Here it is limited by the small and occasional power deficits created by the small demand when wind speeds are low. At high levels of demand, the store size increases to a plateau where it is determined by the size of the energy surpluses on windy days.

3.2.2 Varying Turbine Power Capacity

Figs. 3.40 to 3.47 show the effect of varying the turbine power rating while all other parameters are held constant. In many ways, these results are the inverse of varying the demand at constant turbine power. However, the following graphs show more clearly the effect of an average demand that is much larger than average supply, when the turbine power rating is very low.

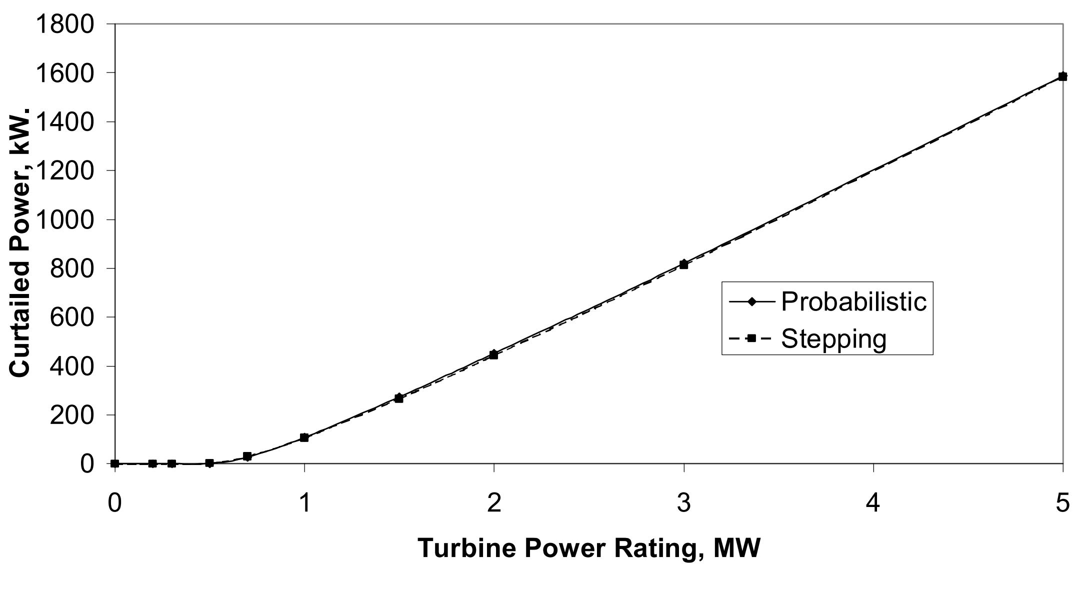

Figure 3.40 Power Curtailed Due To Finite Store Being Full

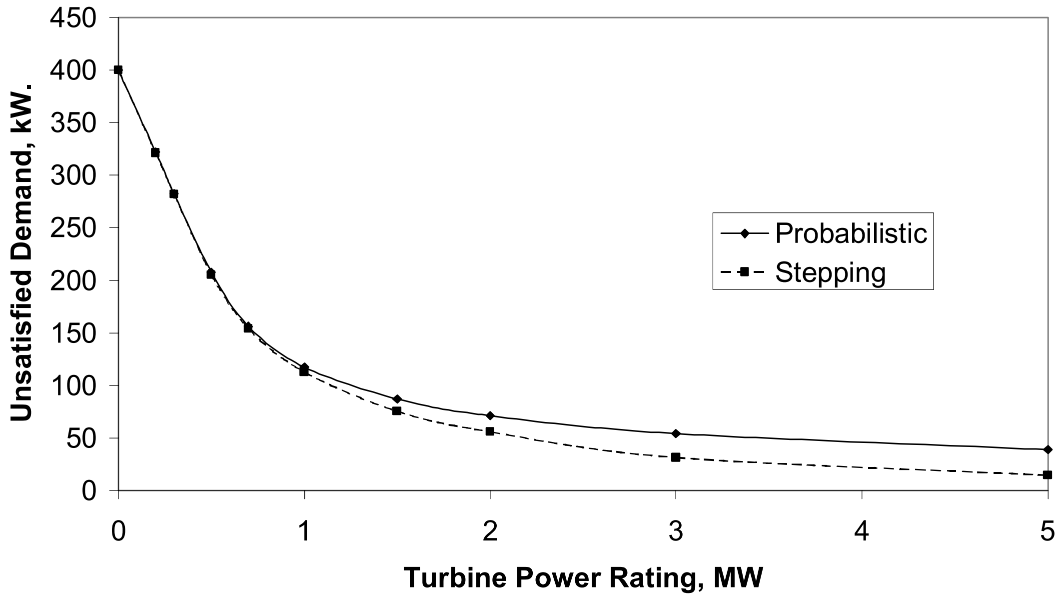

Figure 3.41 Unsatisfied Demand Due to Store Being Empty

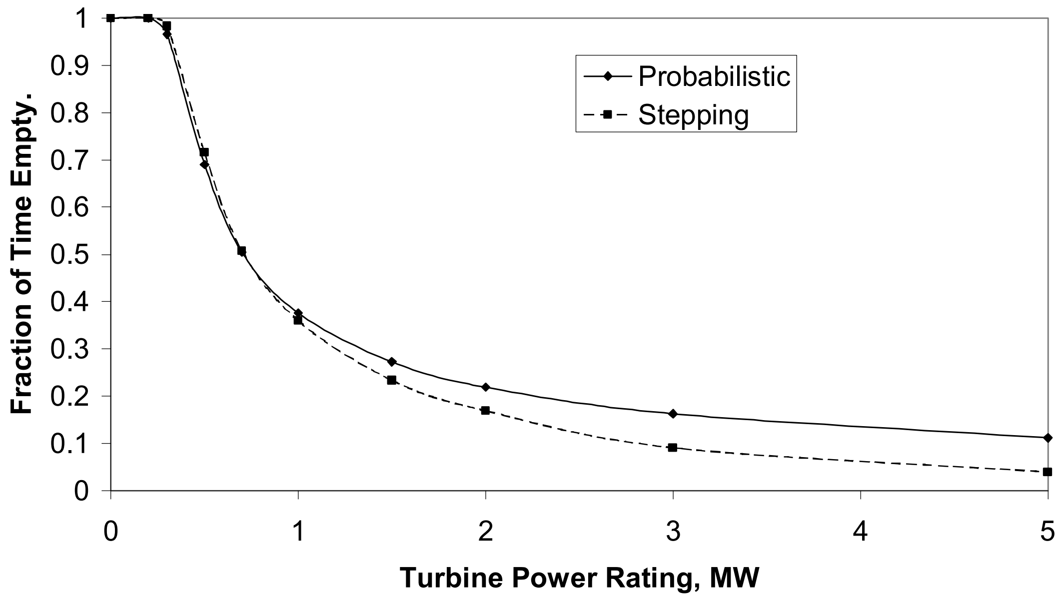

Figure 3.42 Fraction of Time that the Store is Empty

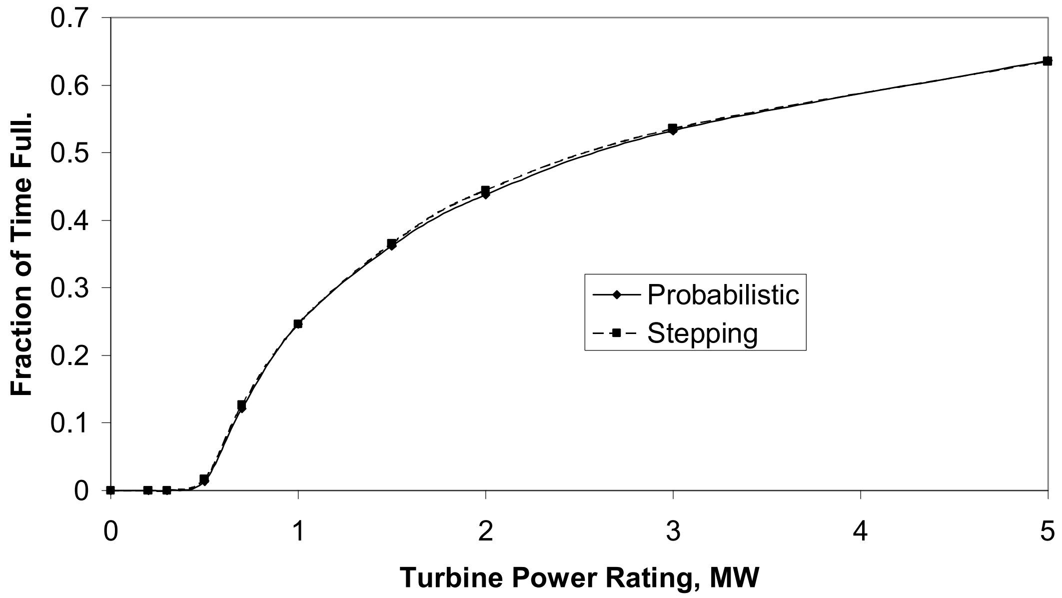

Figure 3.43 Fraction of Time that the Store is Full

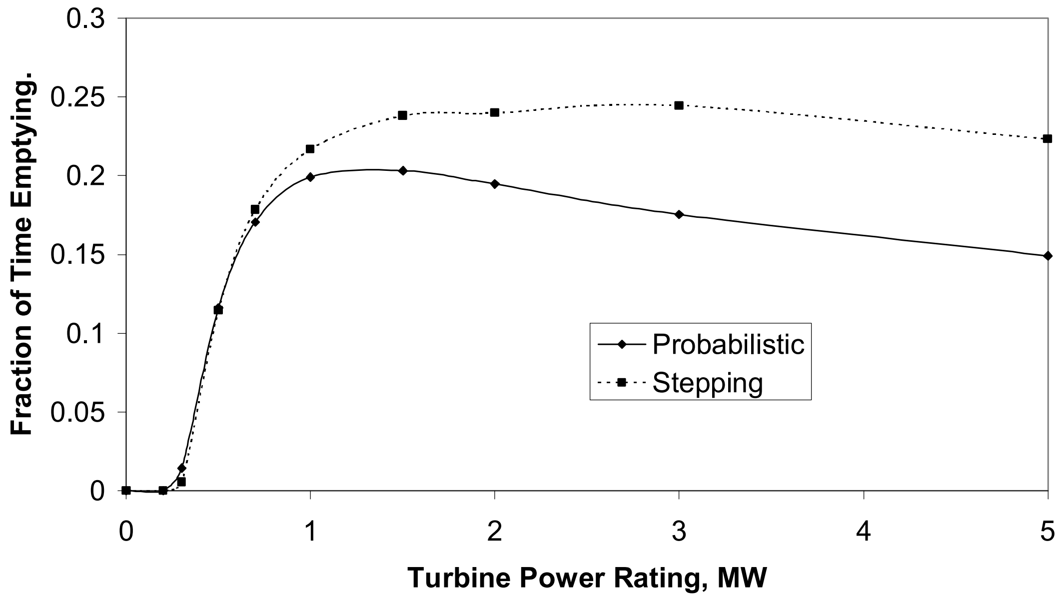

Figure 3.44 Fraction of Time that the Store is Emptying

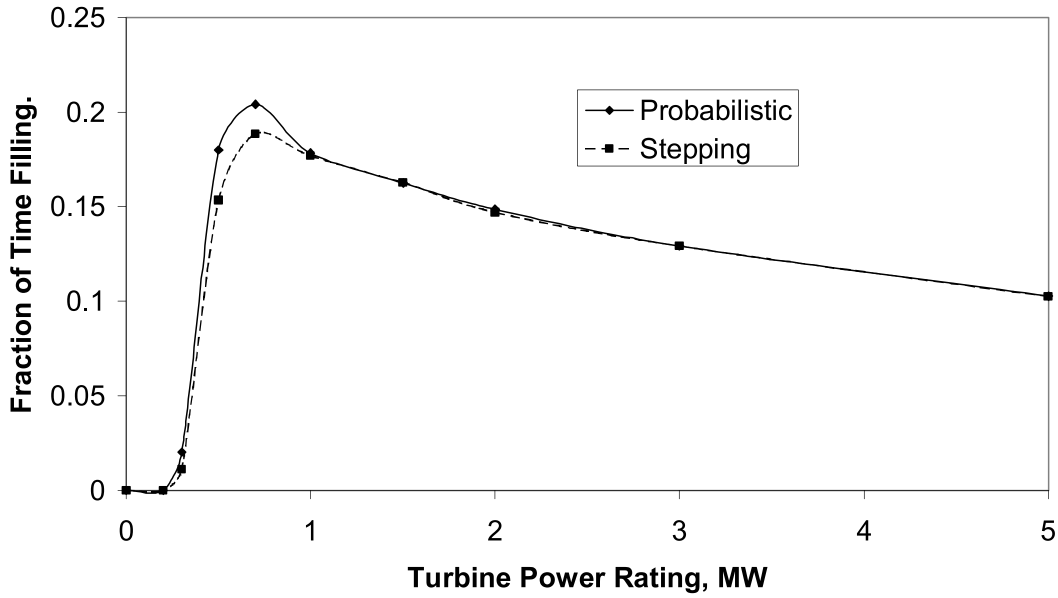

Figure 3.45 Fraction of Time that the Store is Filling

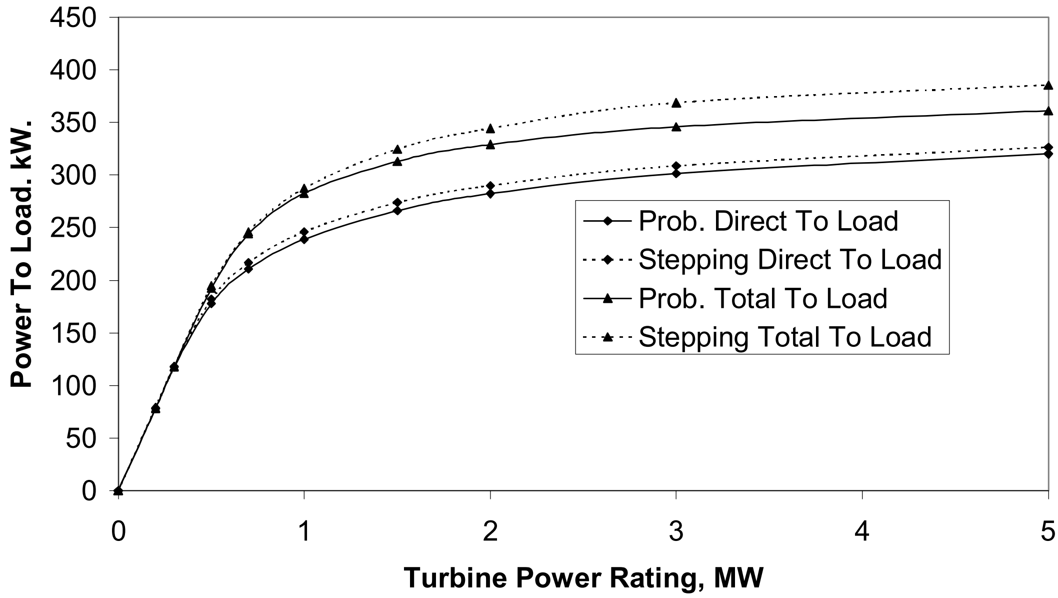

Figure 3.46 Power Supplied to the Loads With and Without the Store

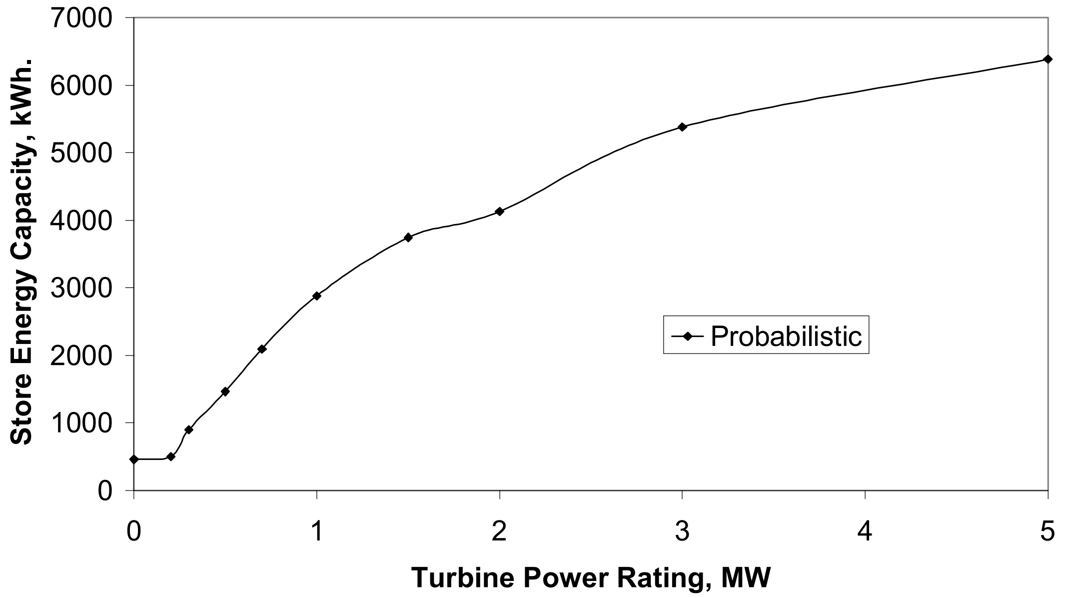

Figure 3.47 Required Store Size for a Cycle Time of 24 Hours

At low levels of turbine capacity, the store spends most of the time empty, fig. 3.42, the unsatisfied demand is very high, fig. 3.41 and the curtailed power is zero, fig. 3.40. Then at high levels of turbine capacity, the store spends most of its time full, fig. 3.43, the unsatisfied demand is low, fig. 3.41 and the curtailed power becomes large, fig. 3.40.

The store is most useful when supply and demand are balanced, at a turbine power of 1MW or larger, as indicated by the large amount of time spent emptying, fig. 3.44 and filling, fig. 3.45 and by the average increase in satisfied demand as a result of the store, fig. 3.46. The store continues to be equally useful at much larger turbine capacities than 1MW, even though the average wind power is well above 400kW (the average demand). The store continues to be useful on days when wind speeds are low.

The required store energy capacity, fig. 3.47 shows that at small turbine capacities the store is small. Here it is determined by the small and occasional power surpluses created when wind speeds are high. At high levels of turbine capacity, the store size increases but more slowly, determined more by the size of the average demand.

3.2.3 Varying Charge and Discharge Rates

Figs. 3.48 to 3.55 show the effect of varying the charge and discharge rates of the store. For simplicity, charge and discharge rates are made equal and are both varied at the same time. The probabilistic and time stepping programs are both capable of modelling charge and discharge rates independently, but real energy storage systems, especially electrochemical systems, often have maximum charge and discharge rates that are similar in magnitude.

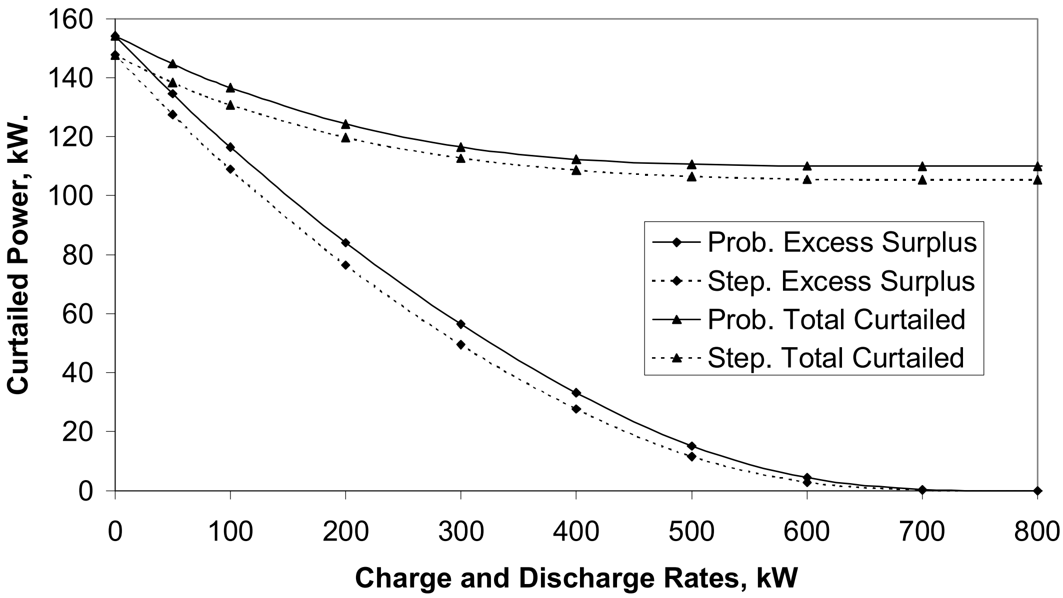

Figure 3.48 Power Curtailed Due To Surplus Power Exceeding Maximum Charging Rate (Excess Surplus) and Due to Store Being Full (Difference between Excess and Total)

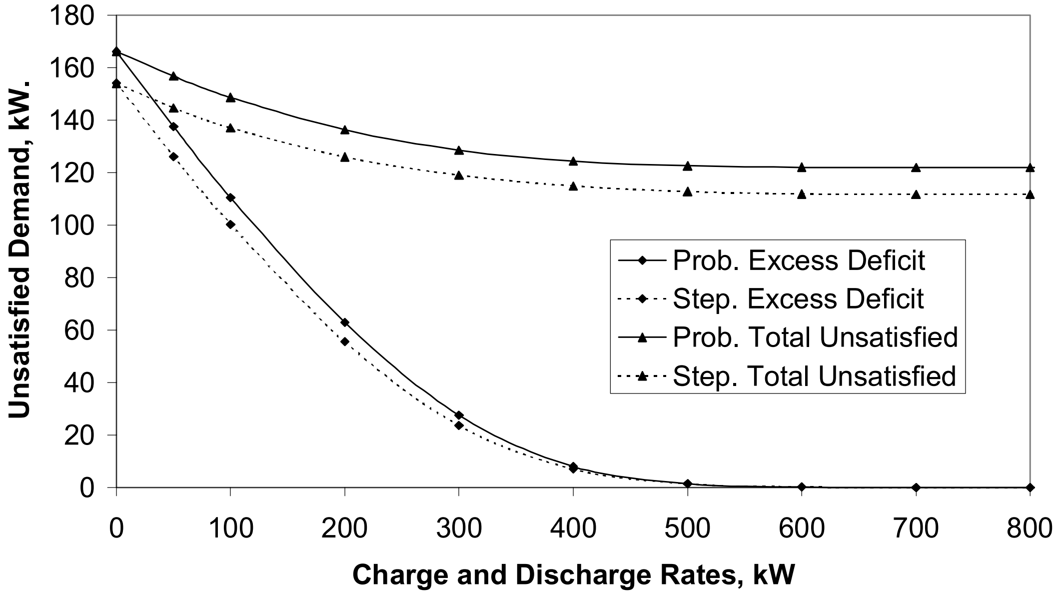

Figure 3.49 Unsatisfied Demand Due to Deficit Exceeding Discharge Capacity (Excess Deficit) and Due to Store Being Empty (Difference between Excess and Total)

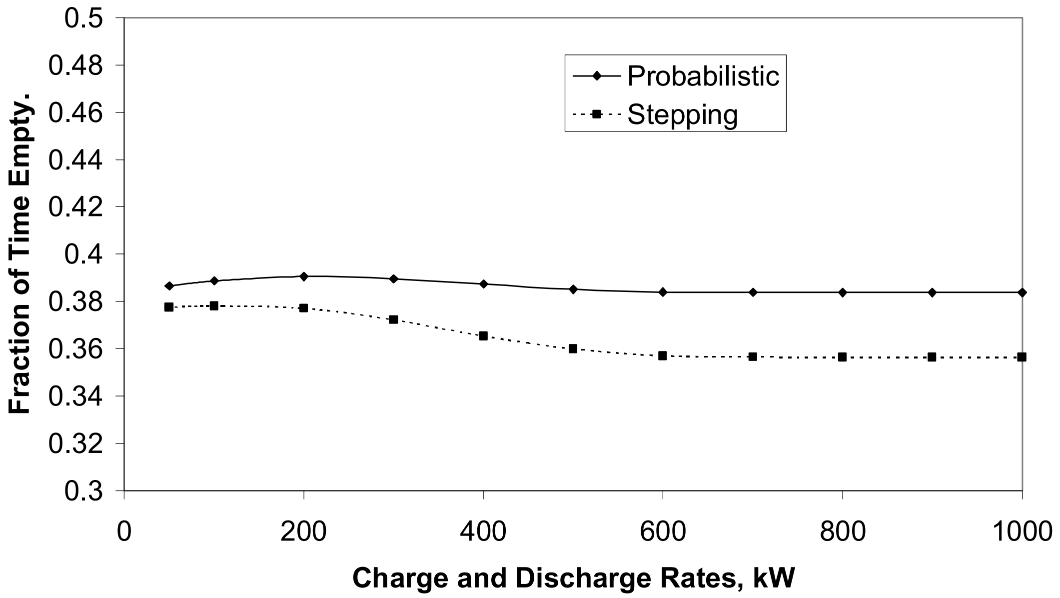

Figure 3.50 Fraction of Time that the Store is Empty

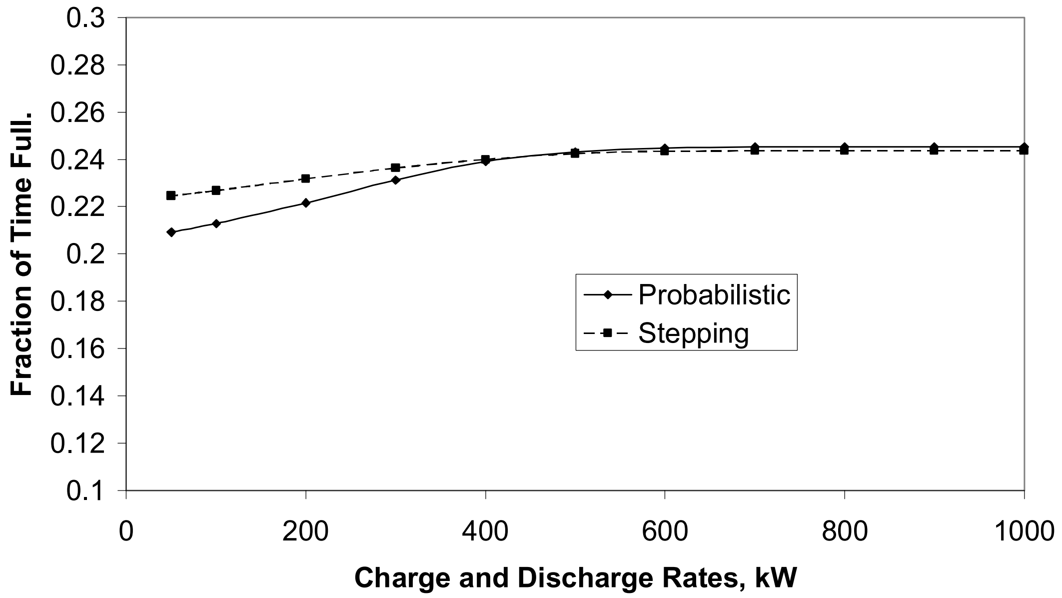

Figure 3.51 Fraction of Time that the Store is Full

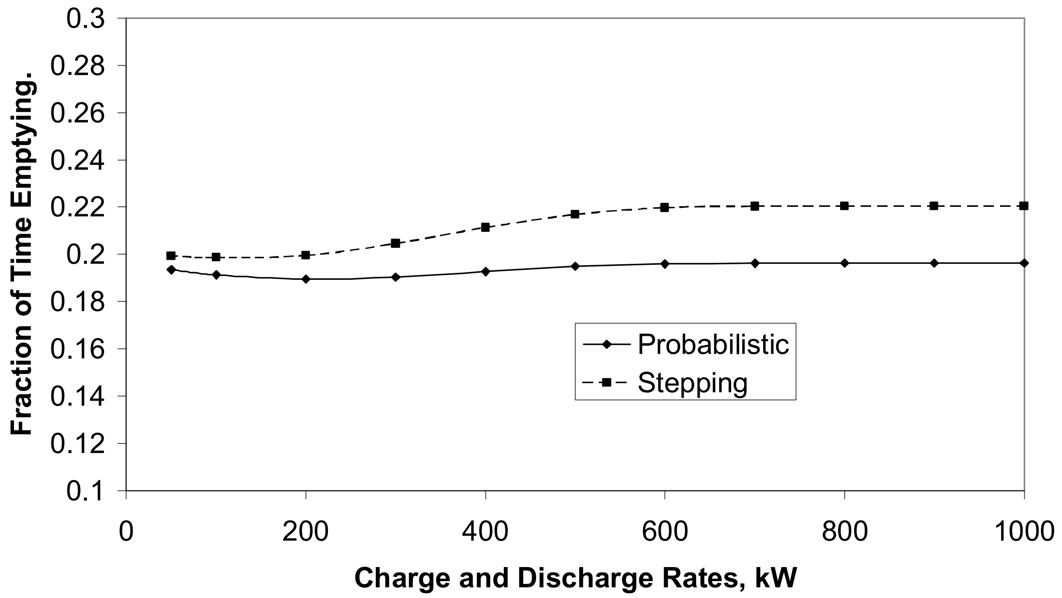

Figure 3.52 Fraction of Time that the Store is Emptying

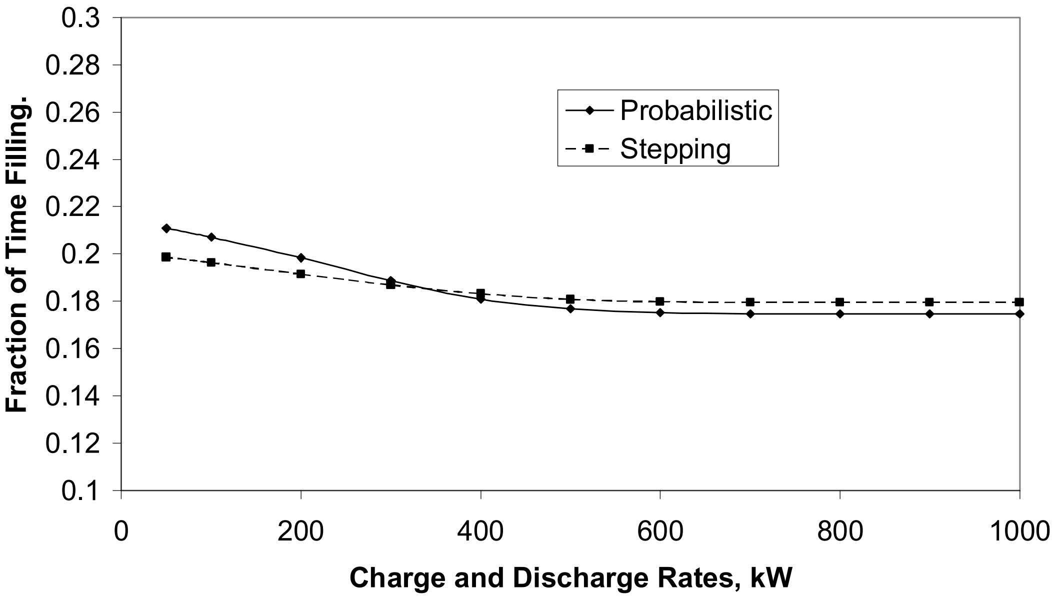

Figure 3.53 Fraction of Time that the Store is Filling

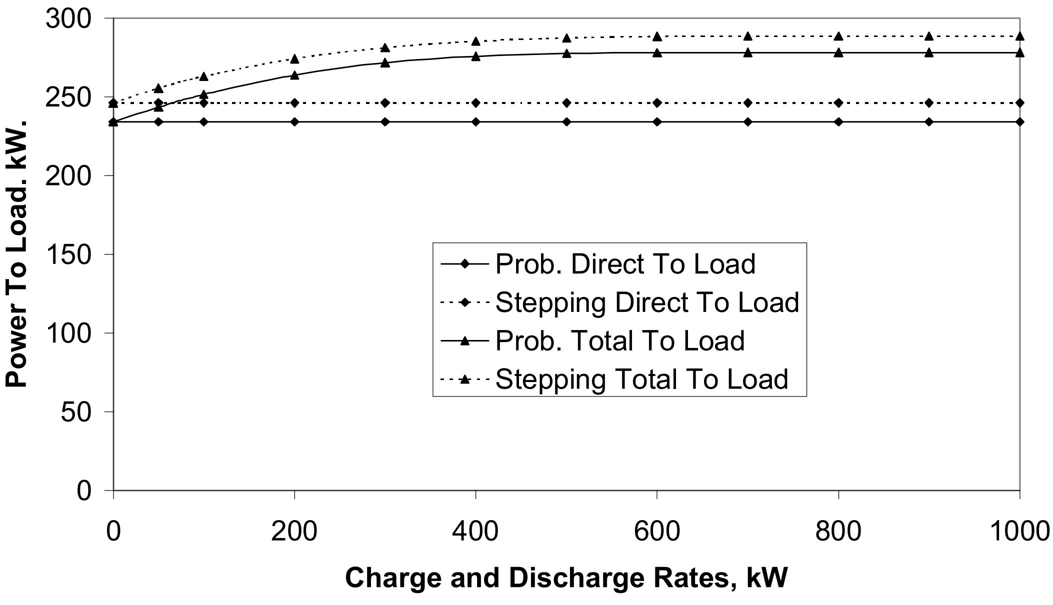

Figure 3.54 Power Supplied to the Loads With and Without the Store

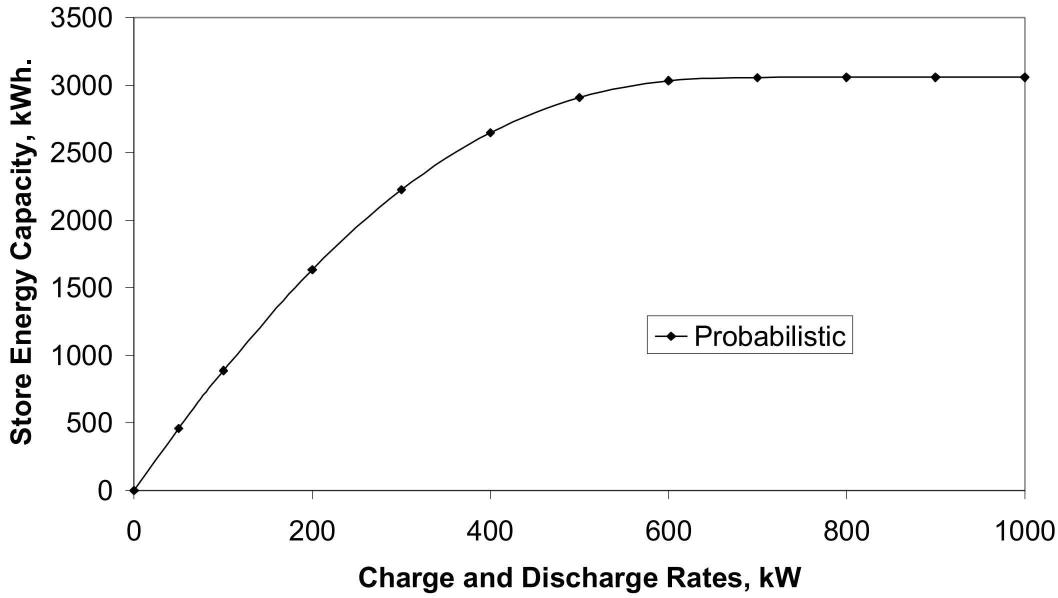

Figure 3.55 Required Store Size for a Cycle Time of 24 Hours

Restricting the maximum charge and discharge rates of the store effectively reduces the power entering and leaving the store. In times of surplus power, some power is curtailed even though the store may not be full. In fig. 3.48, the lower lines represent power curtailed due to the limited charge rate of the store. The upper lines represent the total average curtailed power, and the difference between them is the power curtailed due to the store being full. When the maximum charging rate is small, the full-curtailed power is reduced, but the total curtailed power is increased.

Similarly, during periods of power deficit, some demand is not satisfied even though the store may not be empty. In fig. 3.49, the lower lines represent the unsatisfied demand due to the limited discharge rate. The upper lines represent the total average unsatisfied demand and the difference between them is the demand left unsatisfied due to the store being empty. When the maximum discharging rate is small, the empty-unsatisfied demand is reduced, but the total unsatisfied demand is increased.

The fractions of time that the store is full, empty, filling and emptying are almost unchanged by the charge and discharge rate limitations, if both are changed together, figs. 3.50 to 3.53. The case of zero charge rate and discharge rate is a special case that has been excluded from these graphs; fractions of time spent full and empty etc. are meaningless because zero charge and discharge rates mean the store is not connected to the system!

As the maximum charge and discharge rates are increased, the store achieves maximum usefulness at about 600kW, when all possible surpluses and deficits can be accommodated by the store, fig. 3.54. Beyond this point, the store cannot further increase the average power supplied to the loads without increasing the cycle time of the store.

The required energy capacity of the store increases with charge and discharge rates until all possible surpluses and deficits within 24 hours can be accommodated, fig. 3.55. Again, this occurs at charge and discharge rates of 600kW.

3.2.4 Varying Store Efficiency

Figs. 3.56 to 3.63 show the effect of varying the round-trip efficiency of the store. Parasitic/standing losses are still set to zero. The efficiency is then simply the ratio of energy out to power in. Efficiency is varied between 0% and 100%.

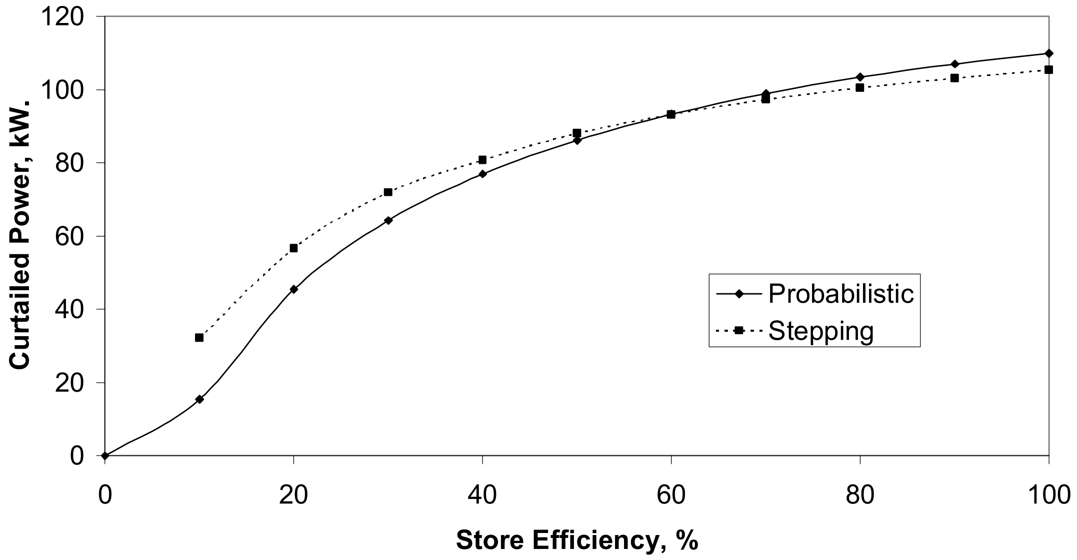

Figure 3.56 Power Curtailed Due To Store Being Full

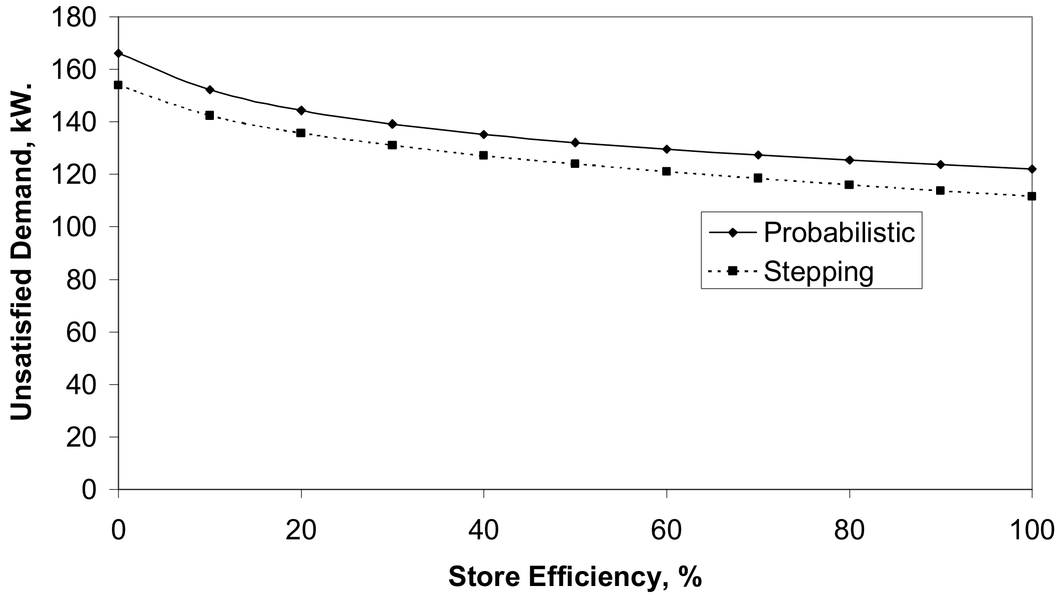

Figure 3.57 Unsatisfied Demand Due to Store Being Empty

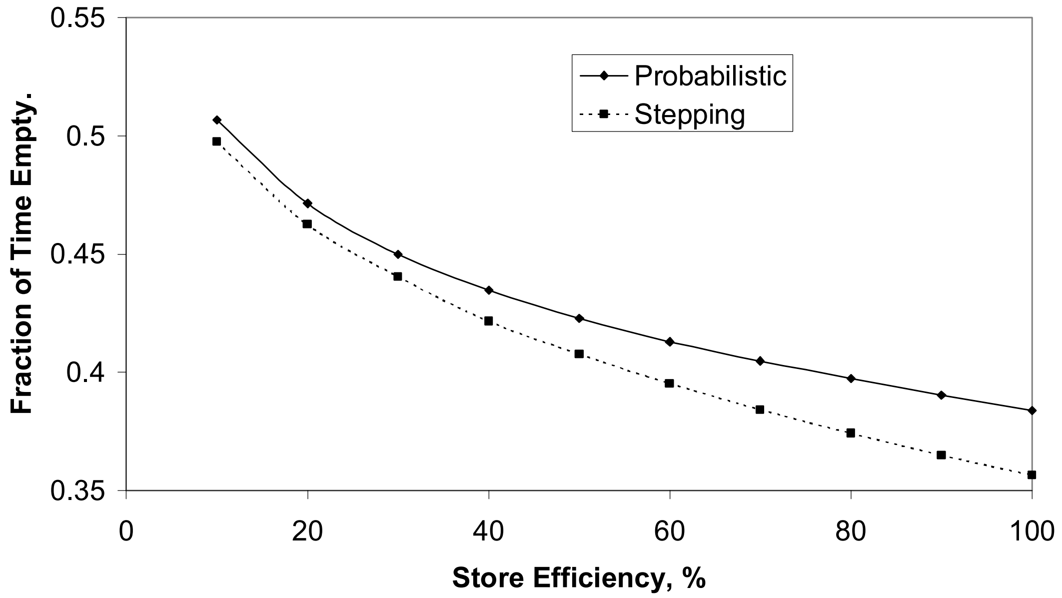

Figure 3.58 Fraction of Time that the Store is Empty

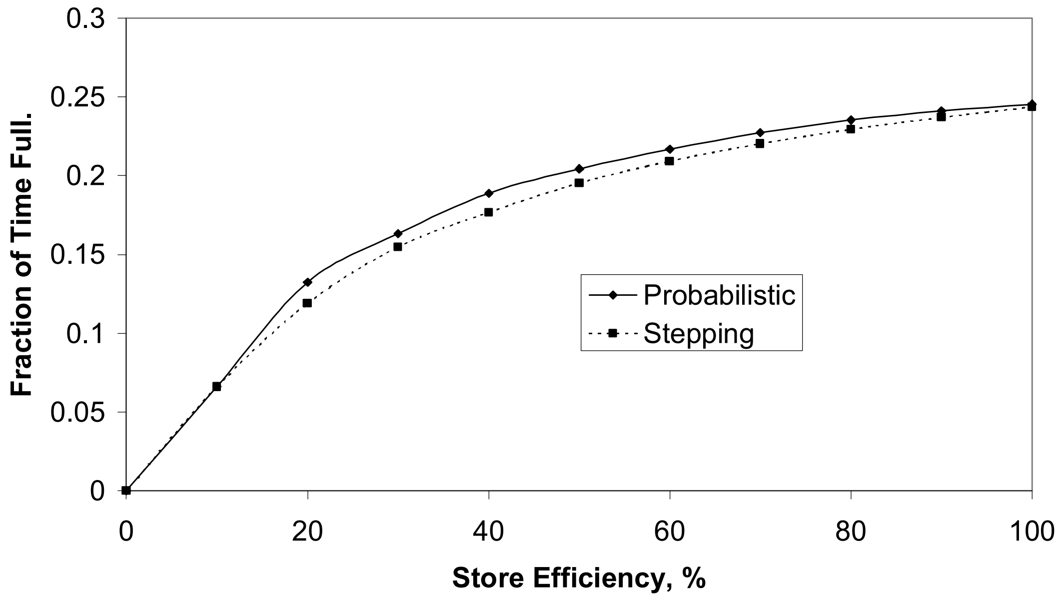

Figure 3.59 Fraction of Time that the Store is Full

Figure 3.60 Fraction of Time that the Store is Emptying

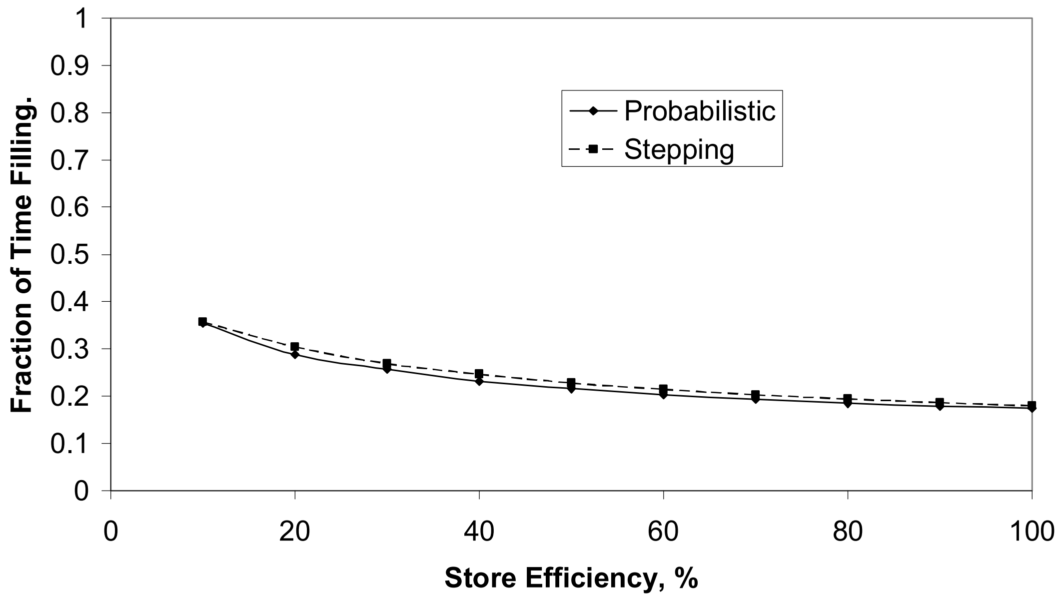

Figure 3.61 Fraction of Time that the Store is Filling

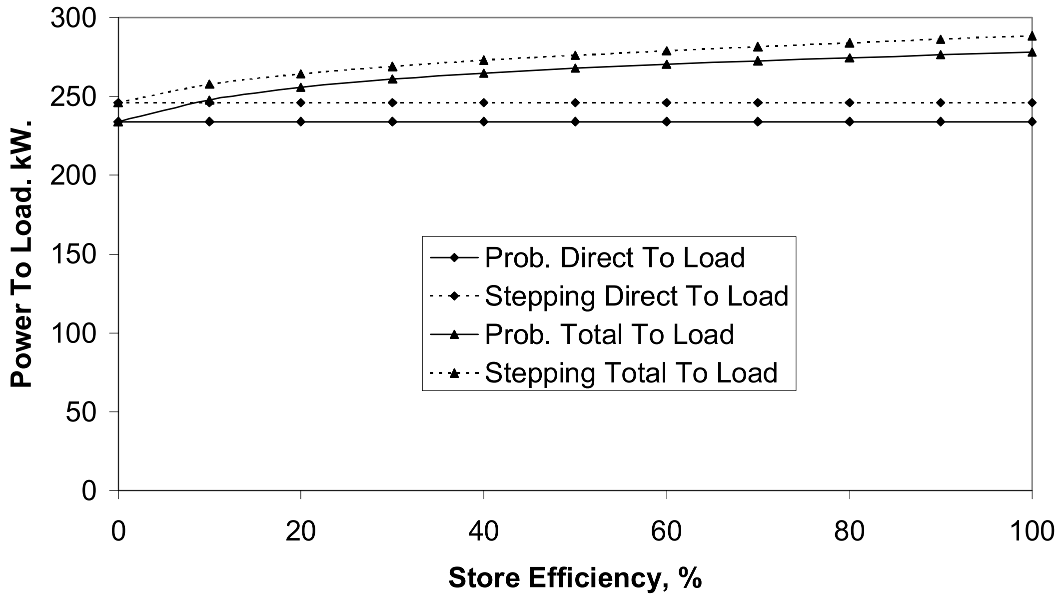

Figure 3.62 Power Supplied to the Loads With and Without the Store

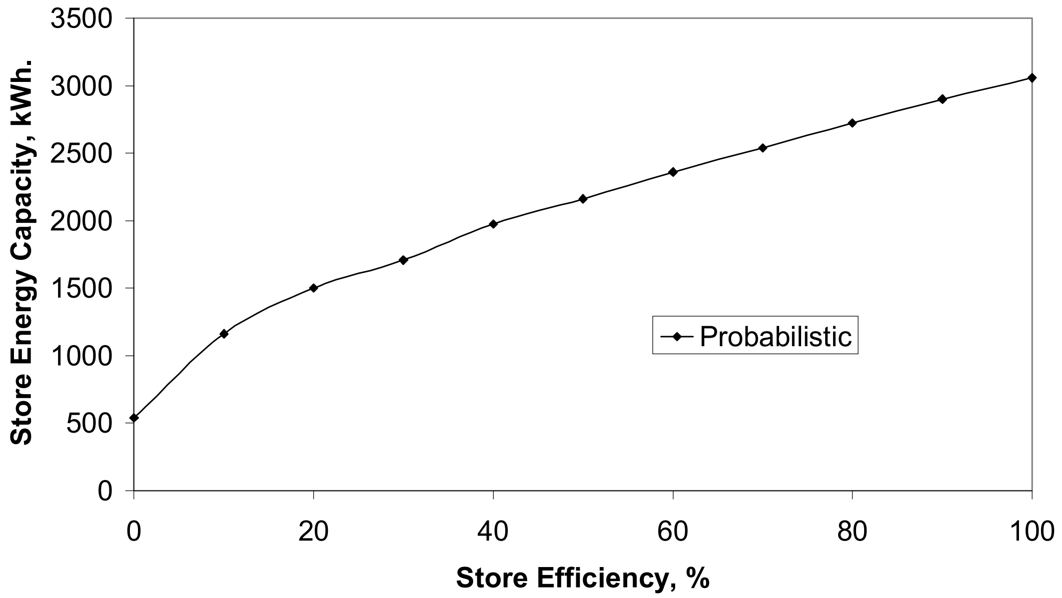

Figure 3.63 Required Store Size for a Cycle Time of 24 Hours

Reducing the efficiency of the store means that a proportion (1-efficiency) of the energy directed to the store is lost. This reduces the energy entering the store and causes the store to be less full on average than if the efficiency were 100%. Therefore, at low store efficiencies the full-curtailed power is reduced, fig. 3.56 and the empty-unsatisfied demand is increased, fig. 3.57. As expected, the store spends less time full, fig. 3.59 and more time empty, fig. 3.58. It also spends more time filling but at a slower rate, fig. 3.61. The store is on average less full and so empties more quickly in times of power deficit and spends less time emptying, fig. 3.60. The 0% efficiency case is another special case in which the store is effectively not connected to the system.

The reduced efficiency means a reduced benefit from the store in terms of total power delivered to loads, fig. 3.62.

The required energy capacity of the store declines with reduced efficiency, fig. 3.63 because a reduced amount of energy enters and leaves the store. Again, the losses are accounted as energy enters the store, not as it leaves.

3.2.5 Varying Store Parasitic Loss

Figs. 3.64 to 3.71 show the effect of varying the parasitic loss (or standing loss) of the store between 0 and 700kW. The 24-hour modelling programs treat parasitic loss in a simplistic way, in that it always draws power, even when the store is empty. In extreme cases, the power delivered to loads can effectively be negative.

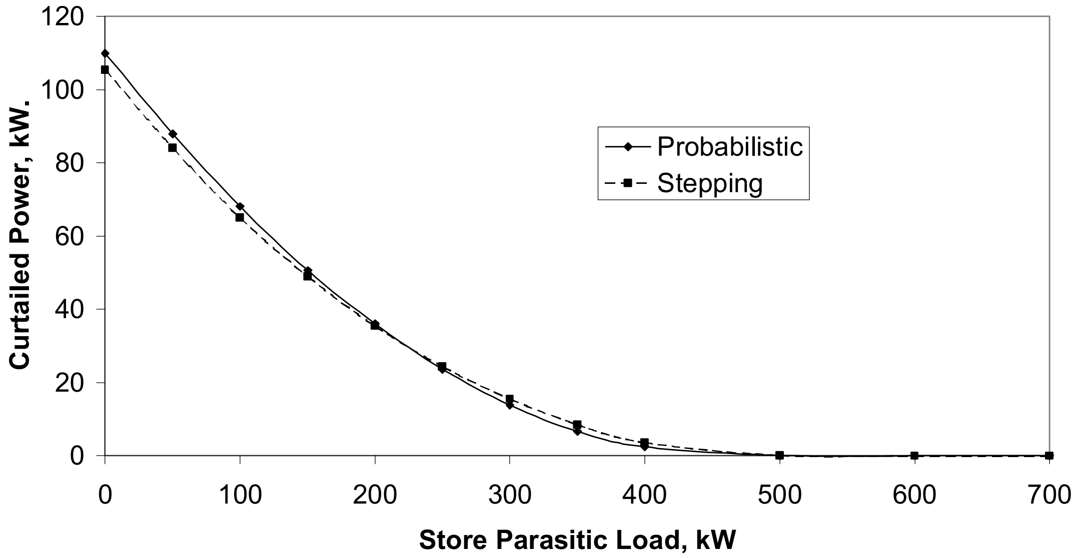

Figure 3.64 Power Curtailed Due To Store Being Full

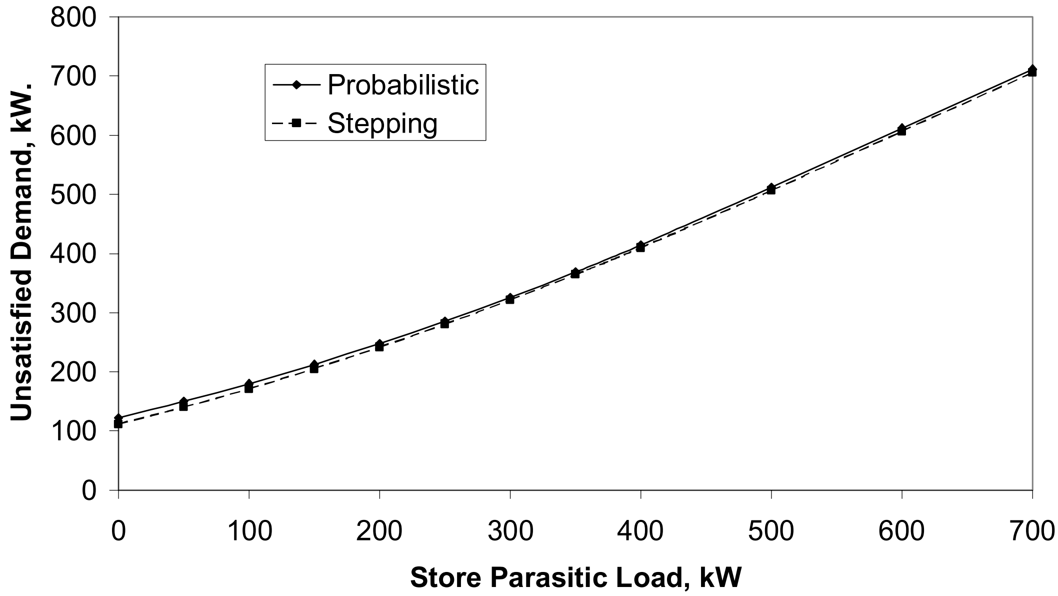

Figure 3.65 Unsatisfied Demand Due to Store Being Empty

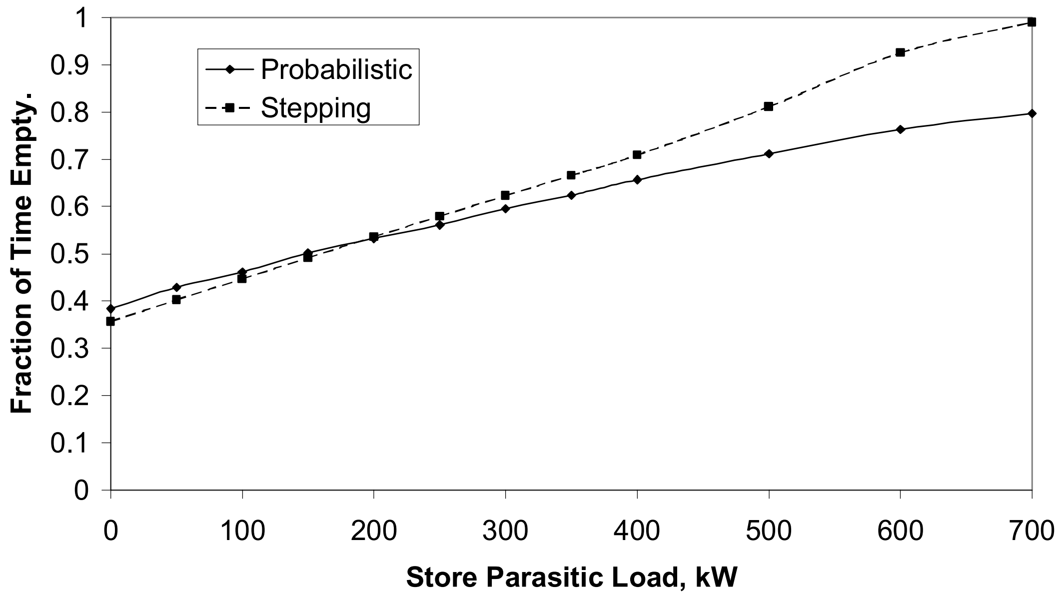

Figure 3.66 Fraction of Time that the Store is Empty

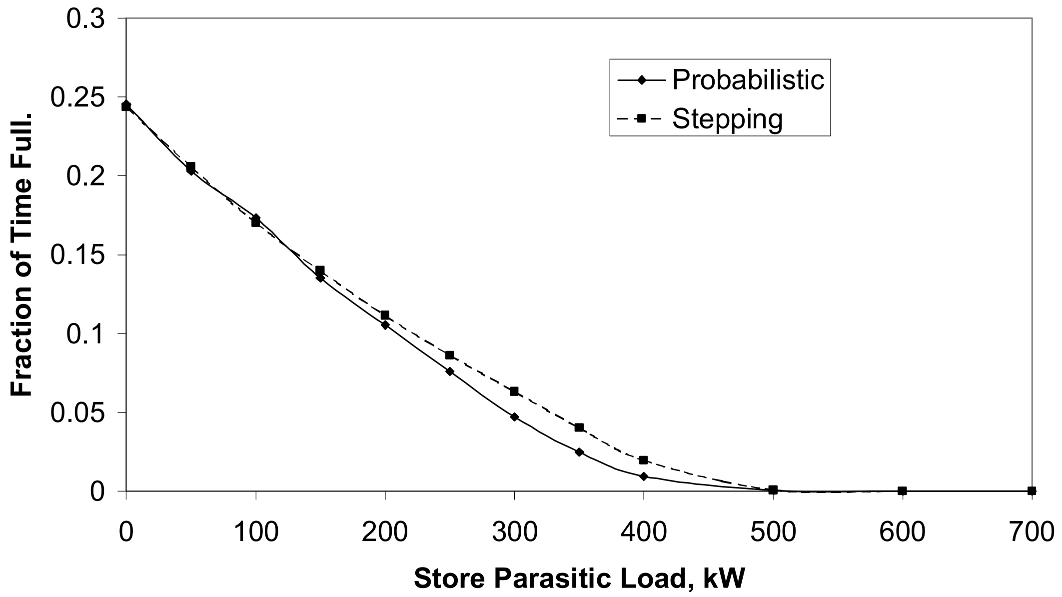

Figure 3.67 Fraction of Time that the Store is Full

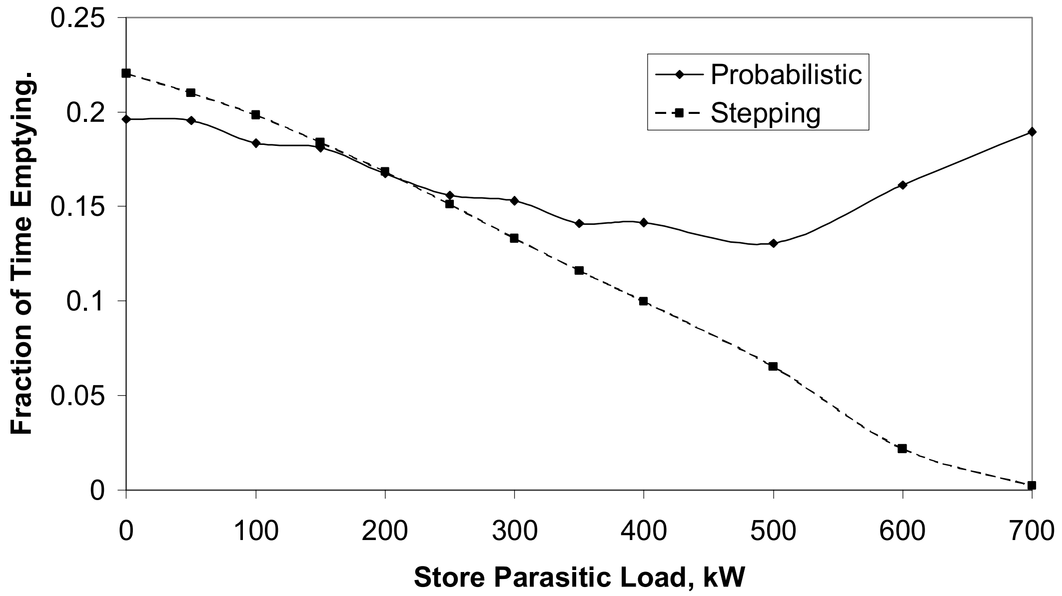

Figure 3.68 Fraction of Time that the Store is Emptying

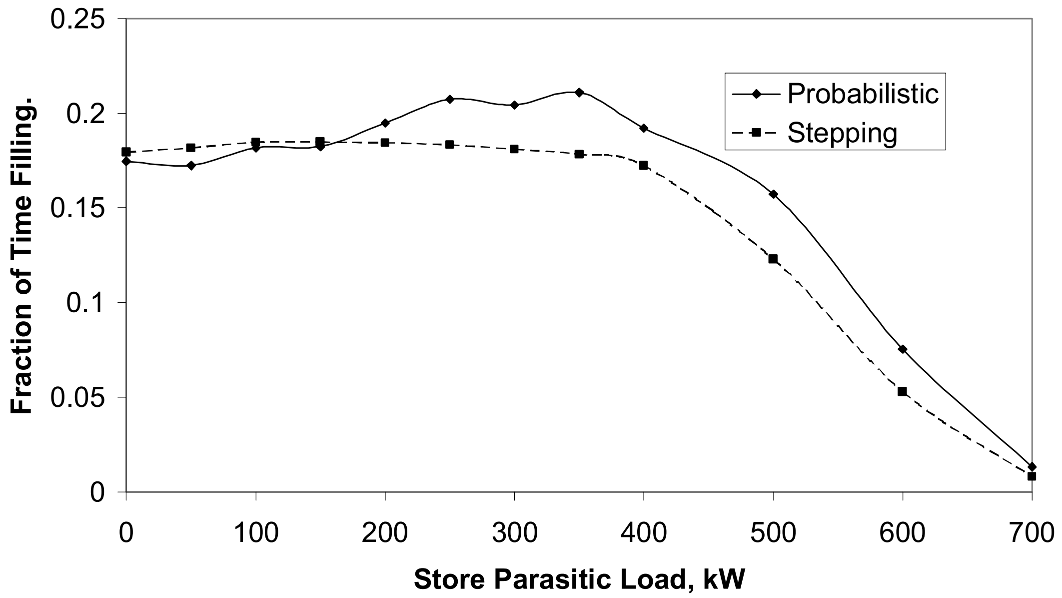

Figure 3.69 Fraction of Time that the Store is Filling

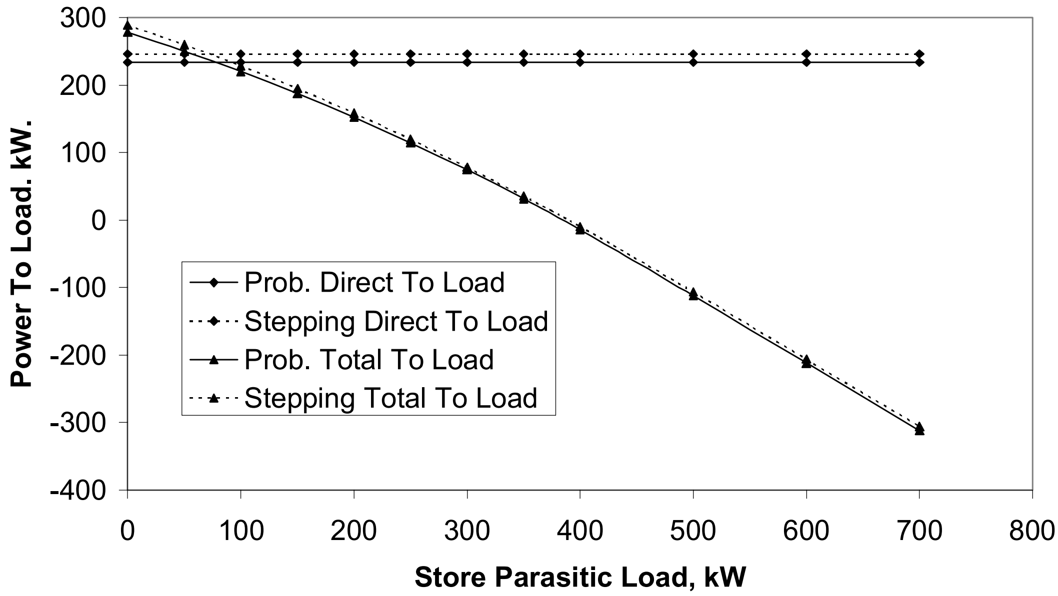

Figure 3.70 Power Supplied to the Loads With and Without the Store

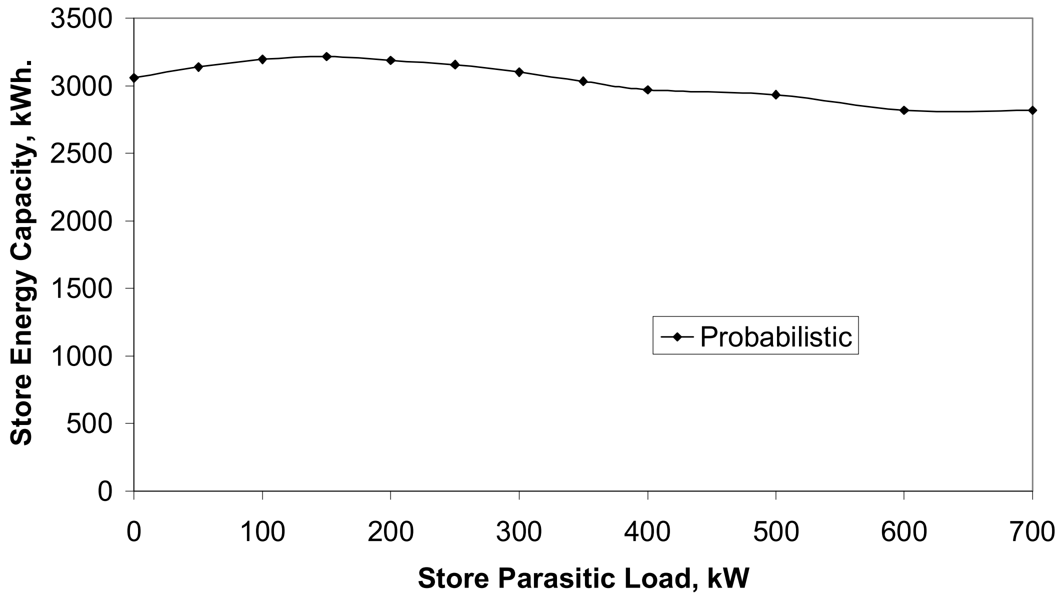

Figure 3.71 Required Store Size for a Cycle Time of 24 Hours

Parasitic or standing loss reduces the average energy in the store but in a different way from a reduced round-trip efficiency. Parasitic loss reduces the curtailed power, fig. 3.64, both by reducing the fraction of time that the store spends full, fig. 3.67 but also by directly absorbing surplus power as it enters the store. It is equivalent to an increase in the load.

In the worst case shown here, a parasitic loss of 700kW, the store loses even the largest possible power surplus and draws power from the loads! The parasitic loss increases the unsatisfied demand, fig. 3.65, until at 700kW of parasitic loss the unsatisfied demand is larger than the average demand of 400kW.

As expected, the parasitic loss reduces the time that the store is full, fig. 3.67 and increases the time that the store is empty, fig. 3.66. Parasitic loss reduces the time spent emptying, fig. 3.68, presumably because in times of power deficit, the store is often already empty, fig. 3.66. However, the probabilistic calculation predicts a strange increase in emptying time when parasitic loss is greater than about 400kW. This could be due to a software bug in the probabilistic 24-hour program, but it only appears when parasitic loss is unrealistically big.

The time spent filling, fig. 3.69 remains fairly constant up to about 400kW of parasitic loss, then reduces rapidly, perhaps as parasitic losses absorb all surpluses.

The benefit of the store becomes negative if the parasitic loss is greater than about 75kW, fig. 3.70, but the predicted store energy capacity required remains almost constant, fig. 3.71. The probabilistic method is calculating the variability of net power, but not necessarily relating the store size to the actual accumulated energy (which becomes zero with large parasitic loss).

3.2.6 Varying Solar Power Fraction

Until now, in all the presented calculations, the only power input to the electricity system has been wind power. Figs. 3.72 to 3.79 show the effect of progressively reducing the fraction of wind power and replacing it with an equal average amount of solar PV power. Since the wind capacity factor is about 0.390 and the solar capacity factor is 0.119 (for the weather conditions at RAL), the exchange rate of solar PV capacity to wind turbine power rating is 3.28.

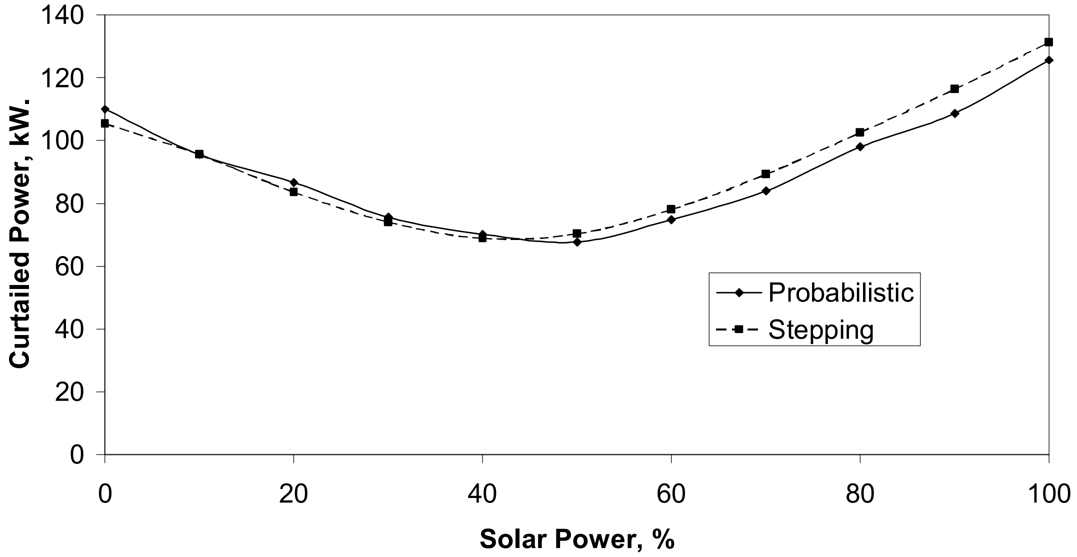

Figure 3.72 Power Curtailed Due To Store Being Full

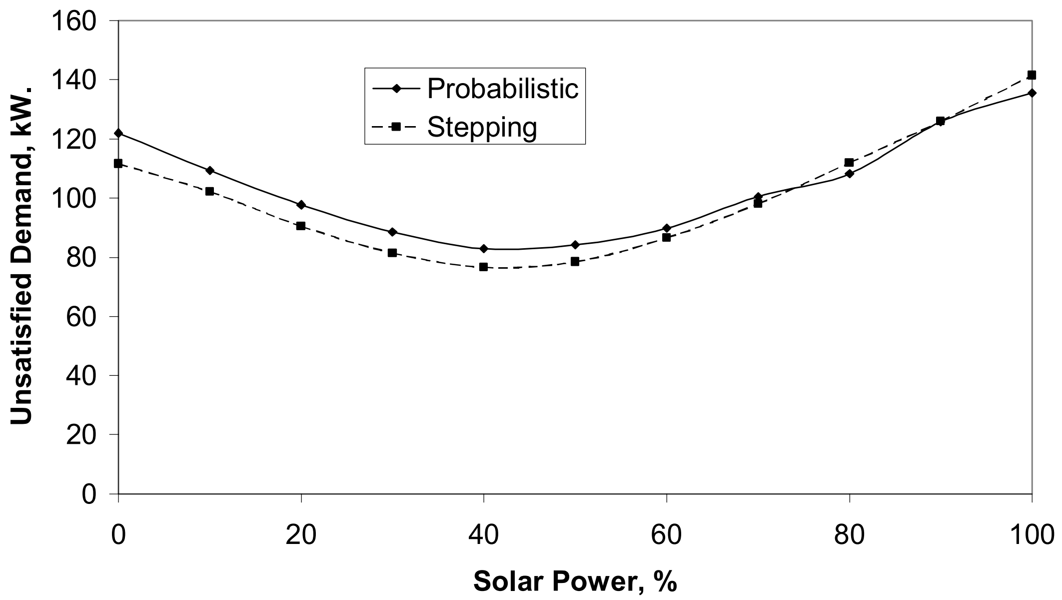

Figure 3.73 Unsatisfied Demand Due to Store Being Empty

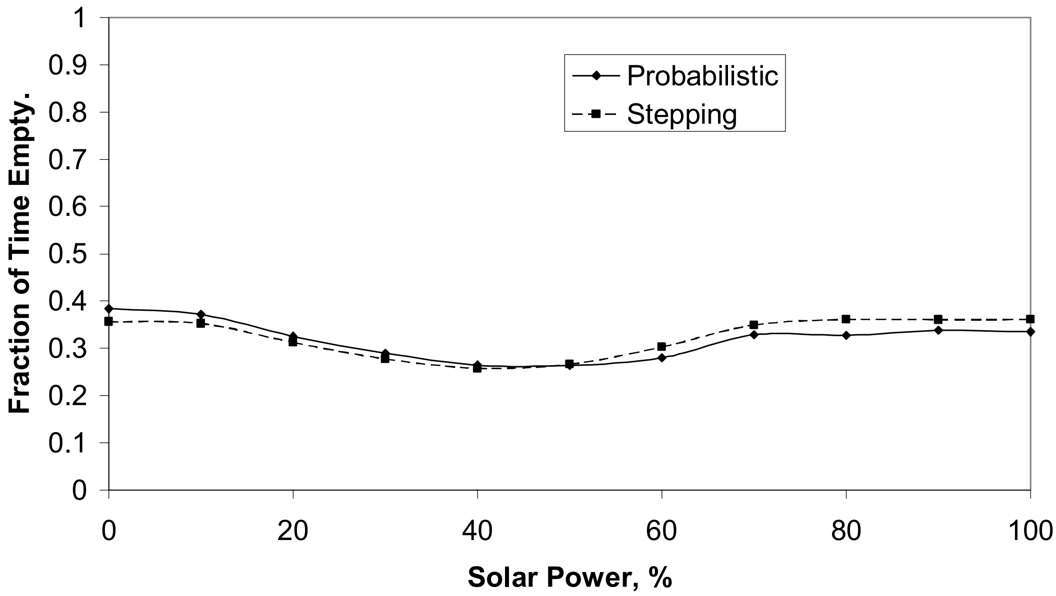

Figure 3.74 Fraction of Time that the Store is Empty

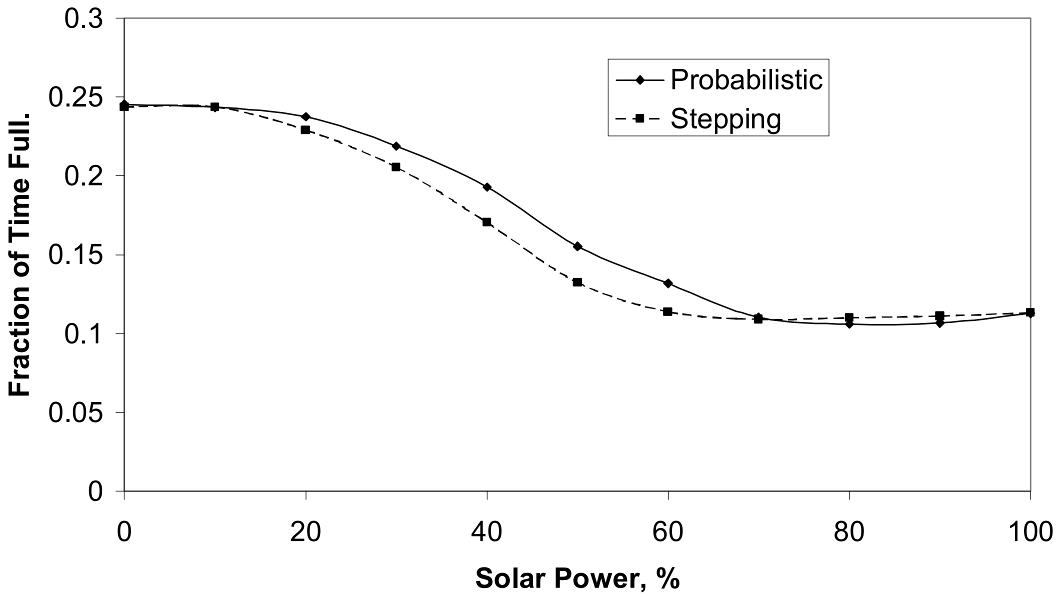

Figure 3.75 Fraction of Time that the Store is Full

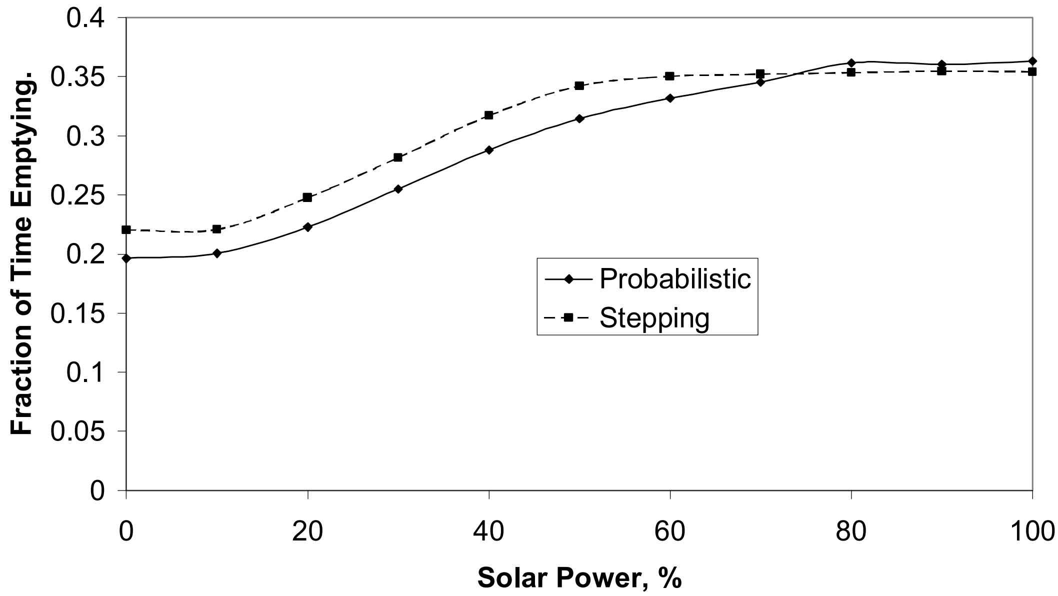

Figure 3.76 Fraction of Time that the Store is Emptying

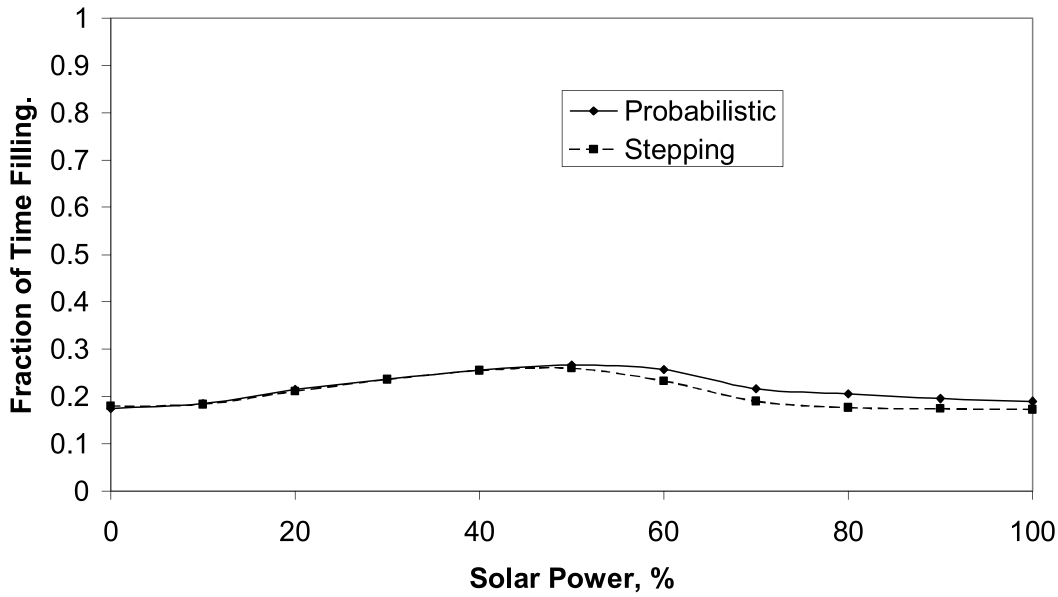

Figure 3.77 Fraction of Time that the Store is Filling

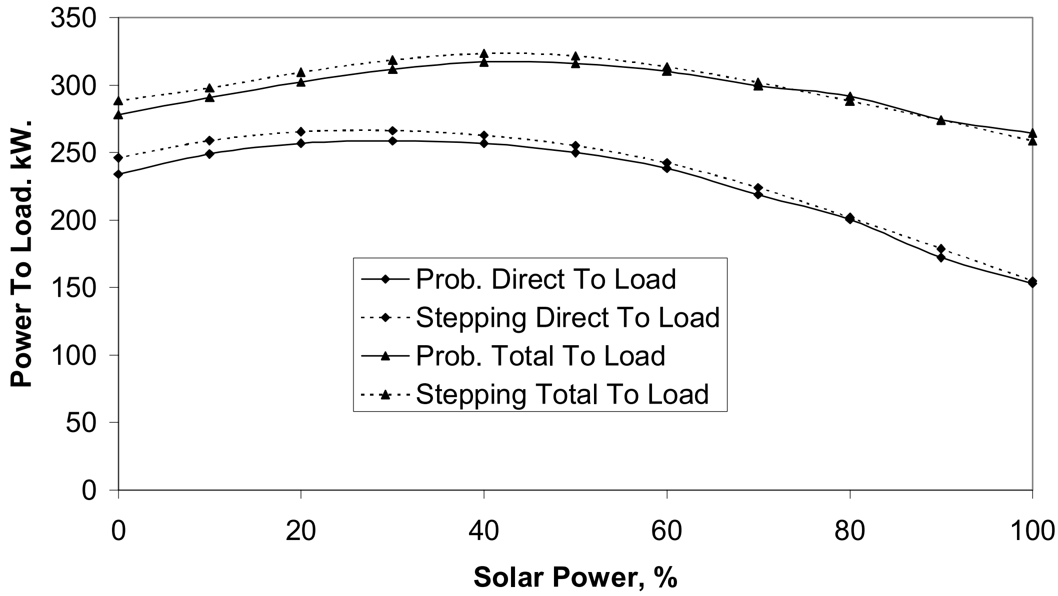

Figure 3.78 Power Supplied to the Loads With and Without the Store

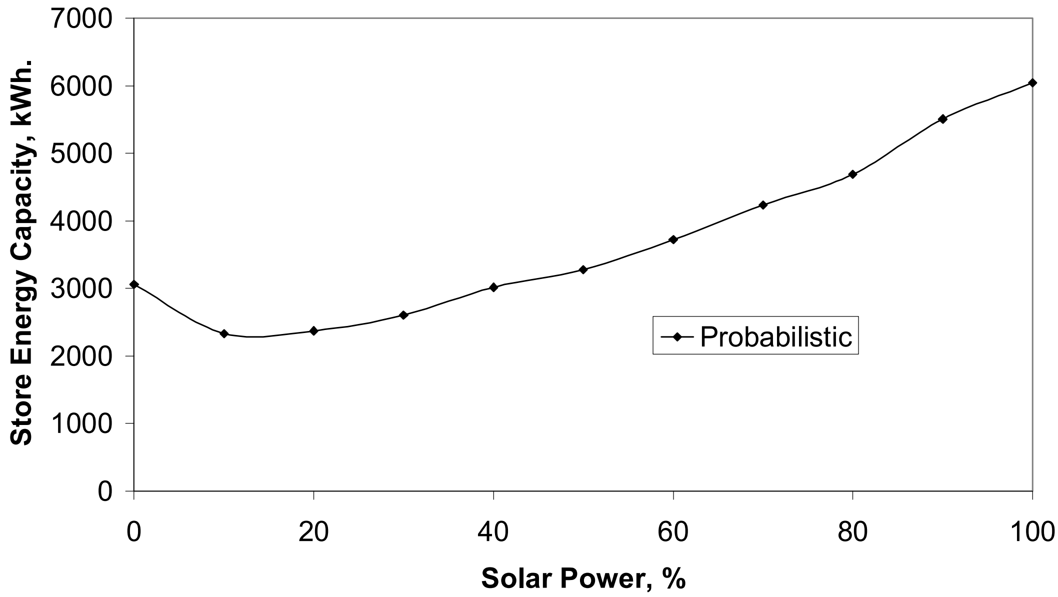

Figure 3.79 Required Store Size for a Cycle Time of 24 Hours

As the solar power fraction increases and wind power fraction decreases, there is a technical optimum apparent as a minimum of average curtailed power, fig. 3.72 and a minimum of average unsatisfied demand, fig. 3.73. This optimum occurs at between 40% and 50% solar PV power. This is the point at which the average total of wind power and solar power is most reliable (least variable), both because wind power and solar power are seasonally complementary, section 2.4.9 and because the sum of any two random, uncorrelated variables has a lower percentage standard deviation than each variable on its own, according to the central limit theorem. This optimum of 40% to 50% solar is also apparent as a minimum in the fraction of time that the store spends empty, fig. 3.74, and a maximum of renewable power supplied to the loads, fig. 3.78. This optimum solar fraction even maximises the time that the store spends filling, fig. 3.77.

Comparing the ends of graphs 3.74 to 3.77 (100% wind vs. 100% solar), solar power on its own causes the store to spend less time full, fig. 3.75 and more time emptying, fig. 3.76, perhaps because solar power has a highly skewed PDF, with high probability of zero or low power, and a low probability of high power, fig. 2.1 in section 2.4.2.

The store increases the average power delivered to the load, more so for solar power than for wind power, fig. 3.78. This is because solar power has a greater within-day variability than wind power. The required 24-hour store energy capacity is also greater with 100% solar power than with 100% wind power, fig. 3.79, again because of greater within-day variability of solar power. The store energy capacity shows a minimum at between 10% and 20% solar power, corresponding to a minimum of within-day variability of the net solar and wind power.

3.2.7 Varying Solar Power With No Wind Power

Figs. 3.80 to 3.87 show the effect of varying solar power capacity in a system with time-varying demand but zero wind power. Like the varying wind turbine capacity case, section 3.2.2, average renewable power supply is at first much smaller than demand, then much larger. The solar PV capacity is taken to a ridiculously large size of 64MW, just to see how the computer programs cope.

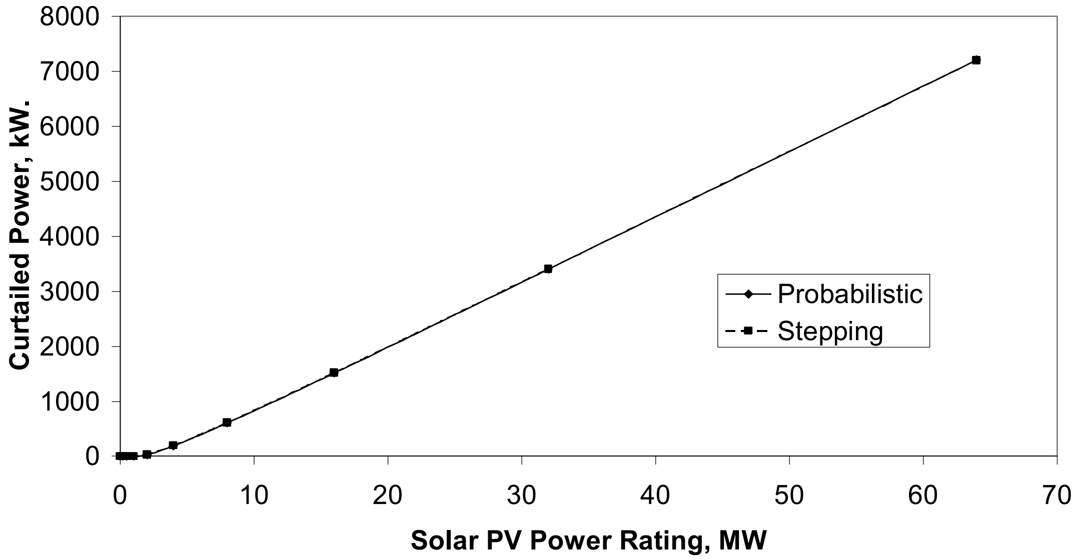

Figure 3.80 Power Curtailed Due To Finite Store Being Full

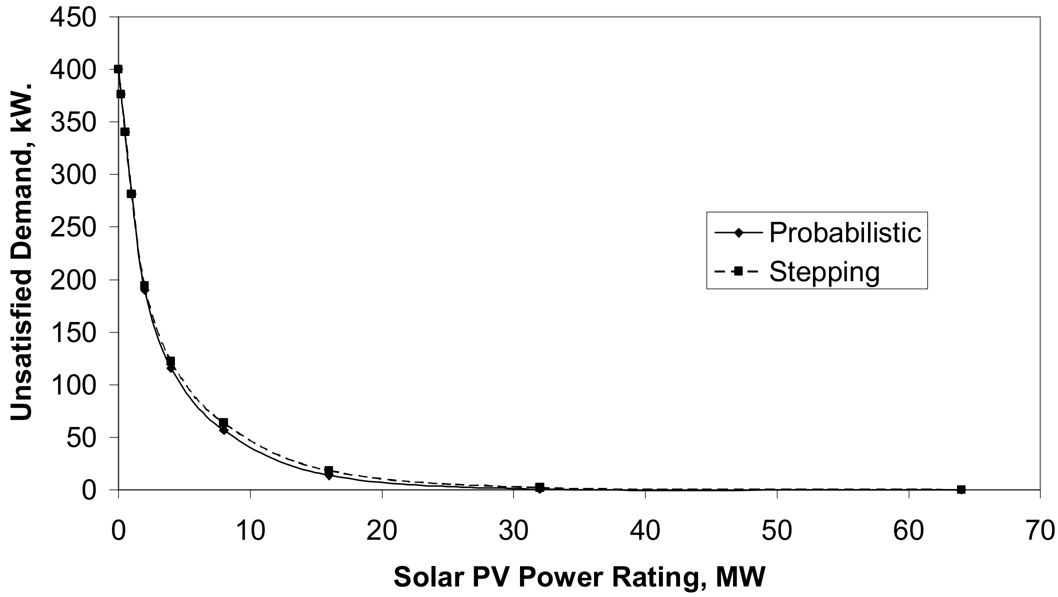

Figure 3.81 Unsatisfied Demand Due to Store Being Empty

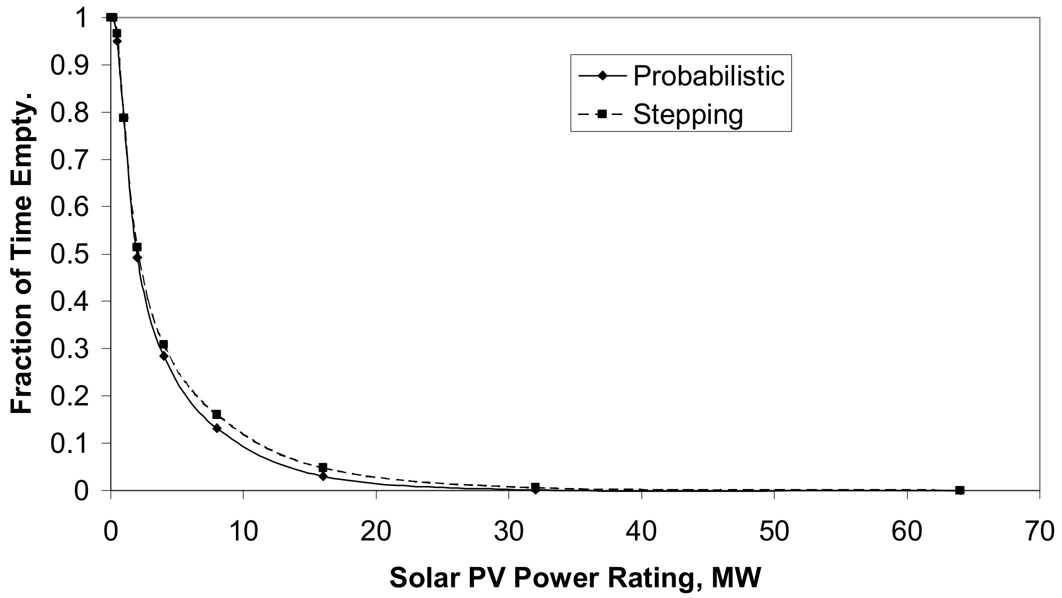

Figure 3.82 Fraction of Time that the Store is Empty

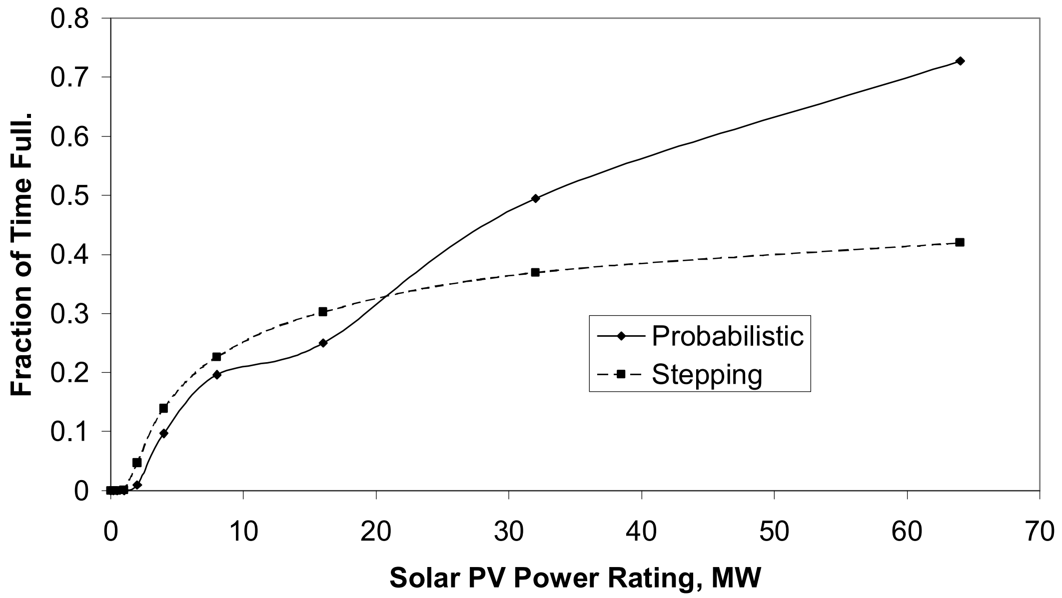

Figure 3.83 Fraction of Time that the Store is Full

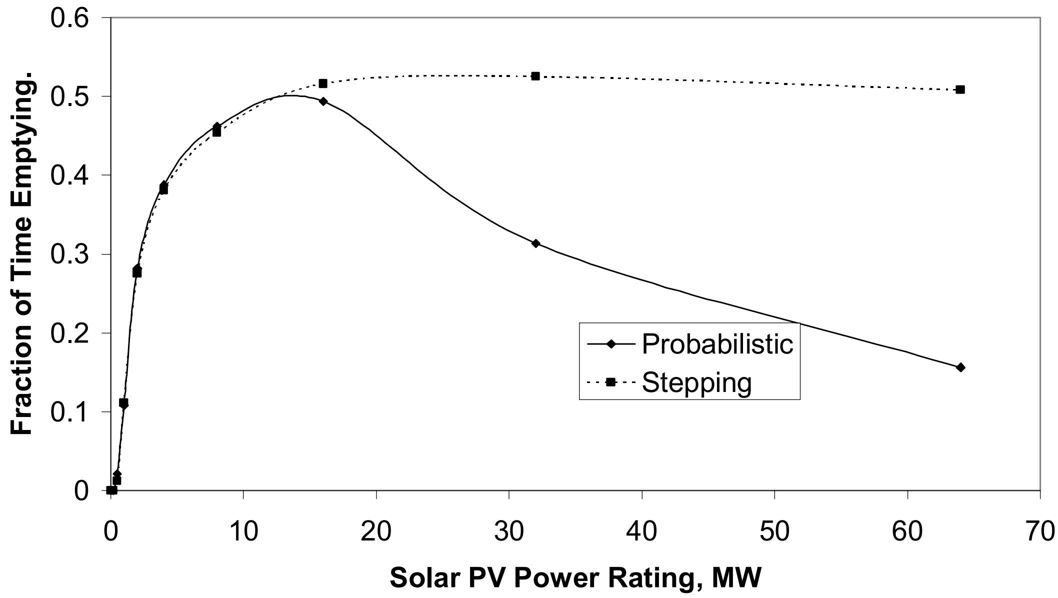

Figure 3.84 Fraction of Time that the Store is Emptying

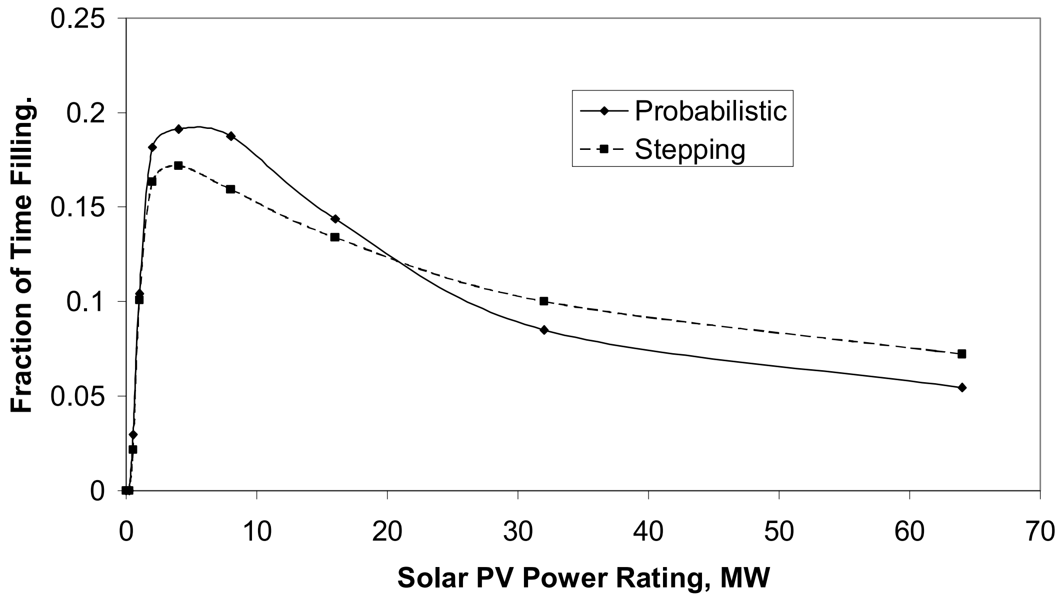

Figure 3.85 Fraction of Time that the Store is Filling

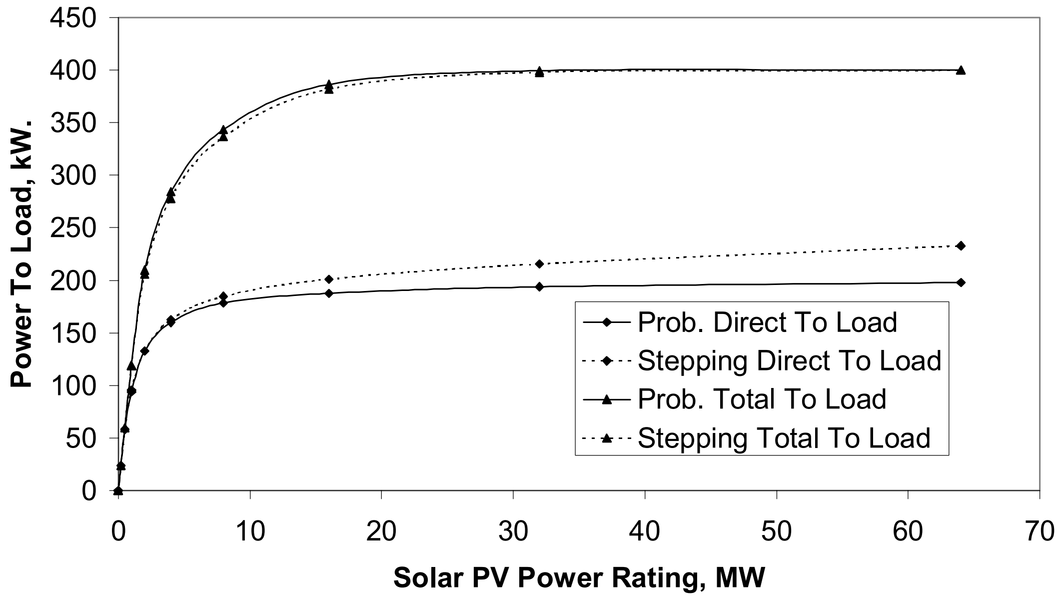

Figure 3.86 Power Supplied to the Loads With and Without the Store

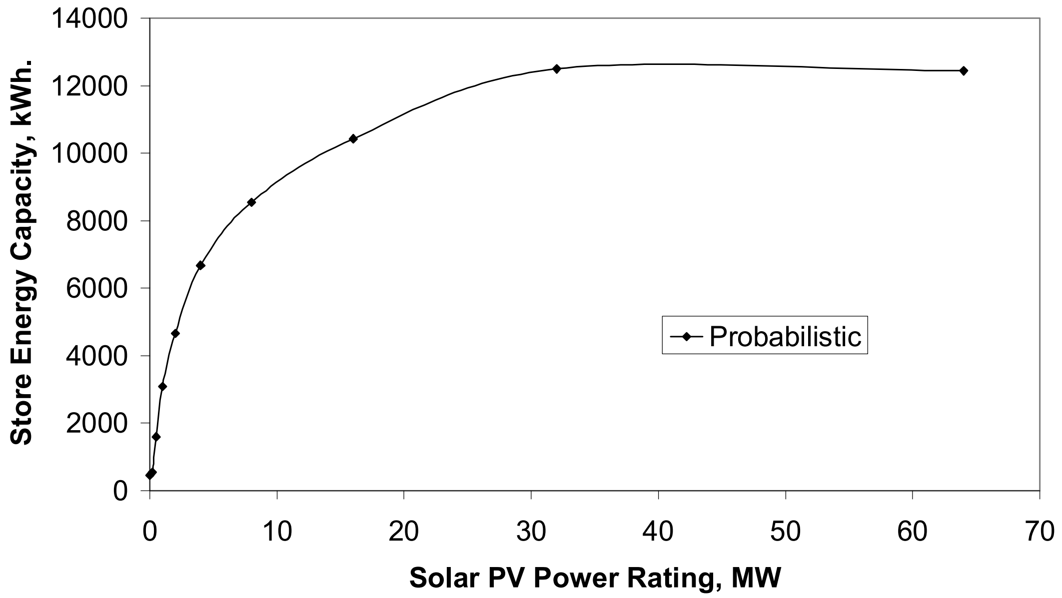

Figure 3.87 Required Store Size for a Cycle Time of 24 Hours

At low levels of solar PV capacity, the store spends most of the time empty, fig. 3.82; the unsatisfied demand is very high, fig. 3.81 and the curtailed power is zero, fig. 3.80. Then at high levels of solar capacity, the store spends most of its time full, fig. 3.83, the unsatisfied demand is low or even zero, fig. 3.81, and the curtailed power becomes large, fig. 3.80. Unlike wind power, even the worst (dullest) days of the year have some solar power, so if the solar capacity is large enough, a 24-hour store will be sufficient to supply the entire demand. But this would require a solar capacity 80 times the average load, or 32MW!

Increasing solar capacity causes the store to spend less time empty, fig. 3.82 and more time full, fig. 3.83. With increasing solar power capacity, the fraction of time spent filling, fig. 3.85 initially increases then reduces, to be replaced by time spent full.

However, the time spent emptying, fig. 3.84 is predicted differently by the probabilistic and time step models. It either increases then levels out at almost 0.5 (time stepping) or increases then drops again (probabilistic). Here we have to believe the time-step model, since we know that on average, the sun is below the horizon for 50% of the time. The demand never drops to zero, so the store must be empty or emptying for at least 50% of the time. The probabilistic calculations of times spent filling, full, emptying and empty, appendix B, do not have access to the real time sequence of power flows. Instead they make assumptions based on the discretised PDFs of net power within each 24-hour period. As the solar power capacity increases, the power interval of this discretisation gets larger and the accuracy of the method may decline. This is the likely explanation for the under-estimate of time spent emptying in the probabilistic 24-hour program.

The store does more ‘work’ with increasing solar capacity, fig. 3.86 up to about 16MW of solar capacity. Above this level, the store does not increase the average power delivered to loads. The demand is almost entirely satisfied above this level of solar power, fig. 3.81.

The required store energy capacity increases with solar power capacity, but then levels out, determined by the size of variations in demand, fig. 3.87.

3.2.8 Varying Grid Connection Capacity

Figs. 3.88 to 3.95 show the effect of a grid connection of varying capacity. For simplicity, the maximum grid import power is set equal to the maximum grid export power and both are changed at the same time. Three store control options are shown, section 2.12.4. The first option keeps the store as full as possible, to minimise the probability of the store being empty and so minimise the unsatisfied demand. The second option uses the grid to balance surpluses and deficits as much as possible, and as far as possible minimising the risk of either energy curtailment or unsatisfied demand. The third option keeps the store as empty as possible, so minimising the probability of curtailed energy. The option numbers are show in the legends of the figures.

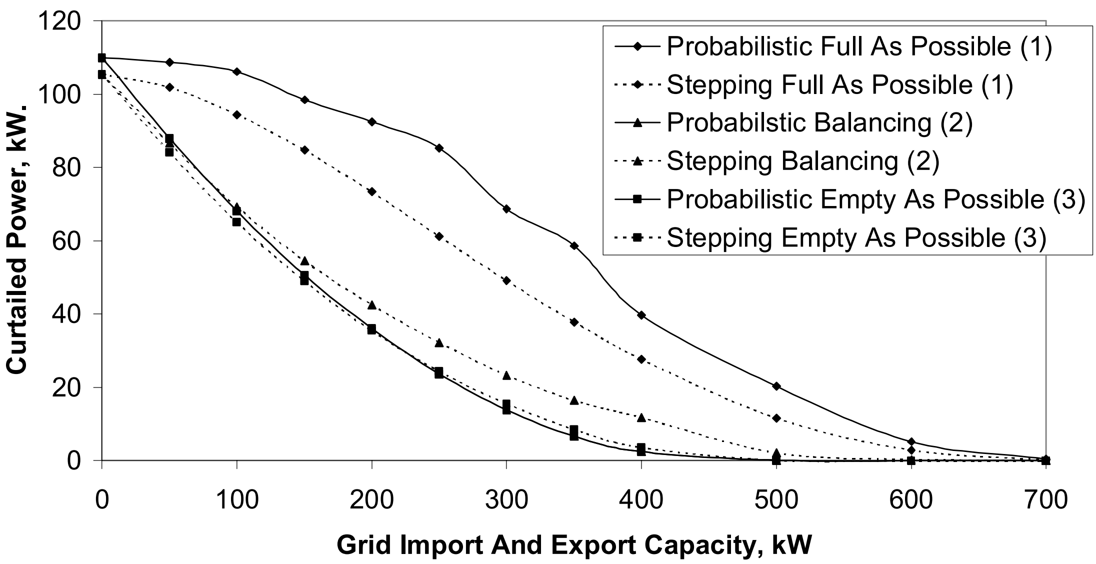

Figure 3.88 Power Curtailed Due To Finite Store Being Full

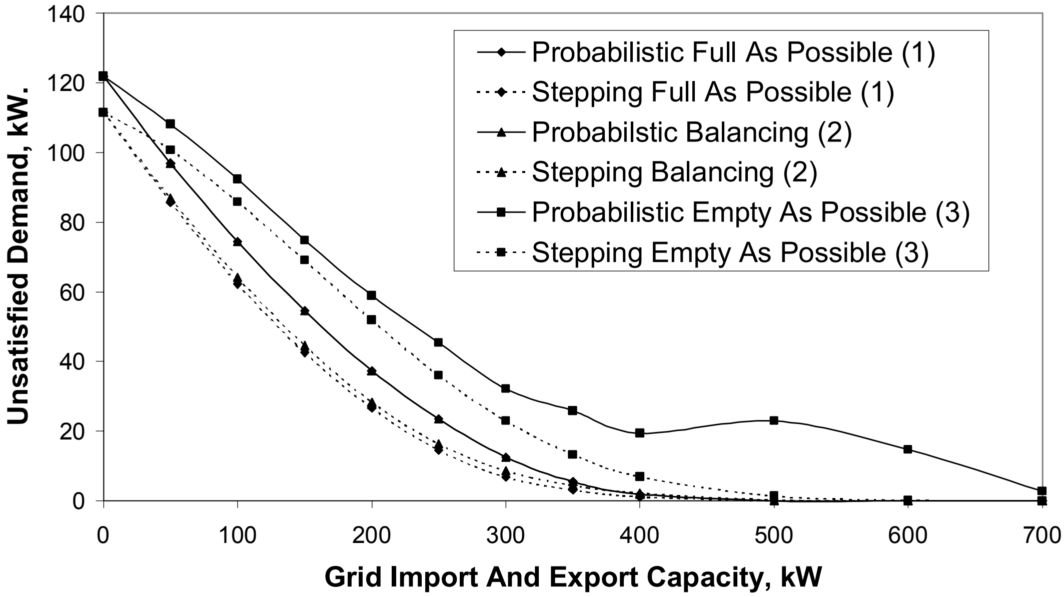

Figure 3.89 Unsatisfied Demand Due to Store Being Empty

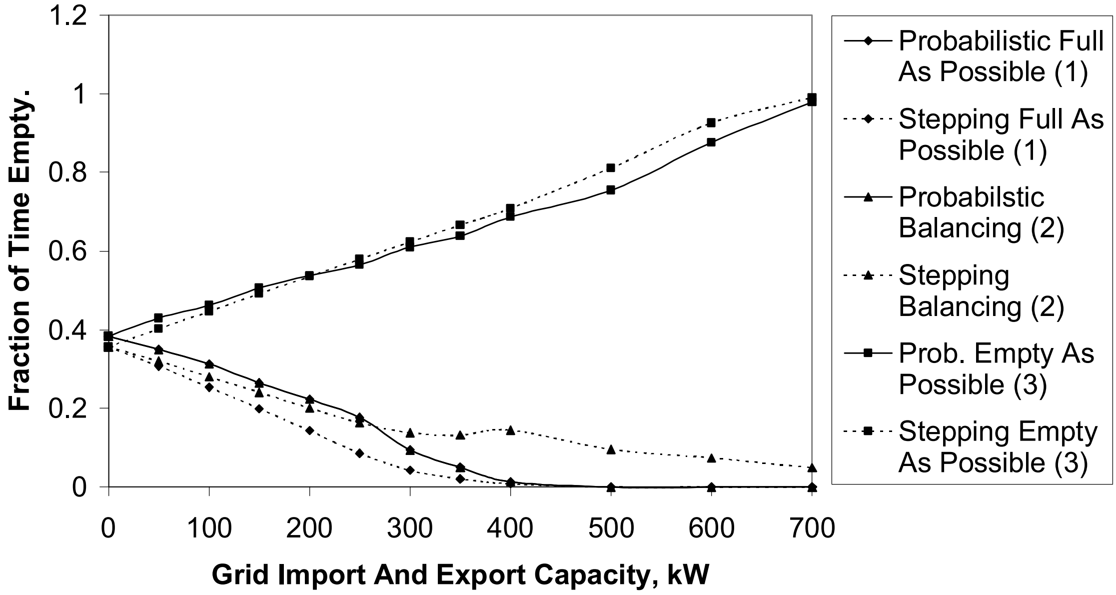

Figure 3.90 Fraction of Time that the Store is Empty

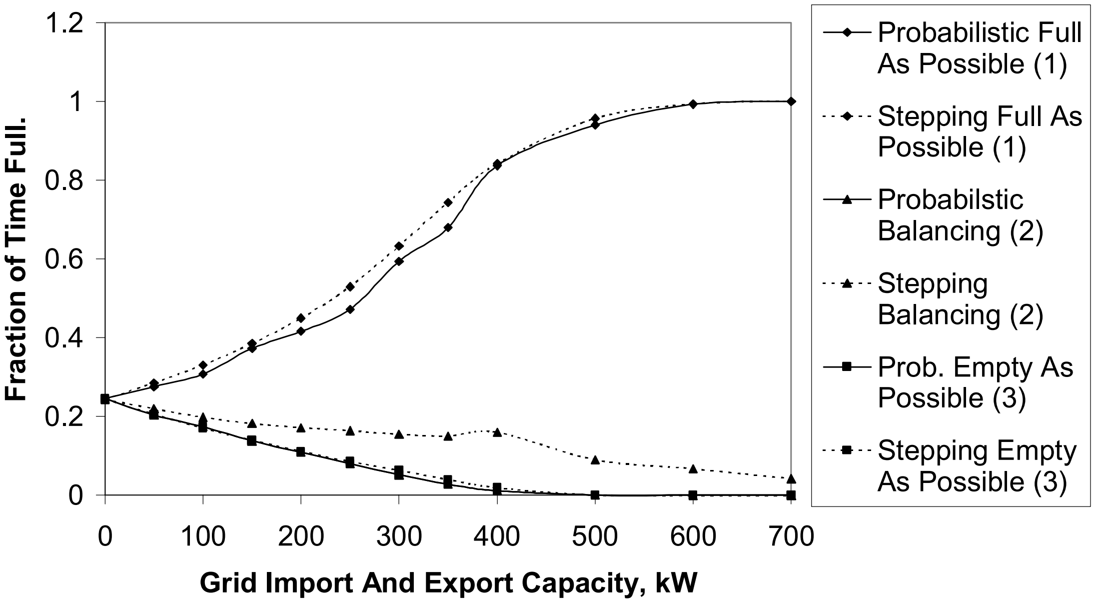

Figure 3.91 Fraction of Time that the Store is Full

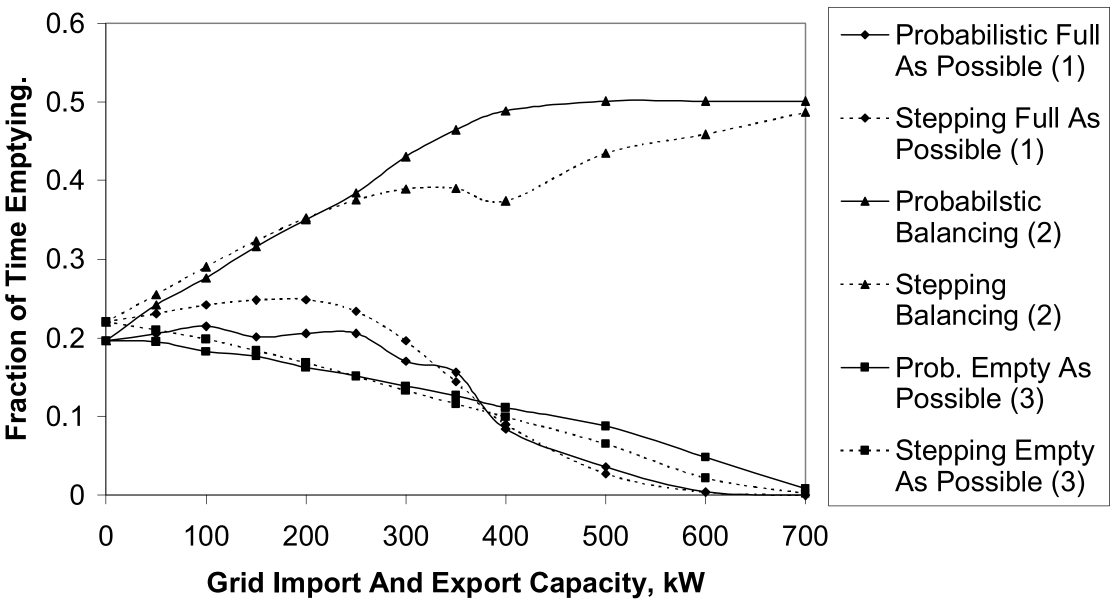

Figure 3.92 Fraction of Time that the Store is Emptying

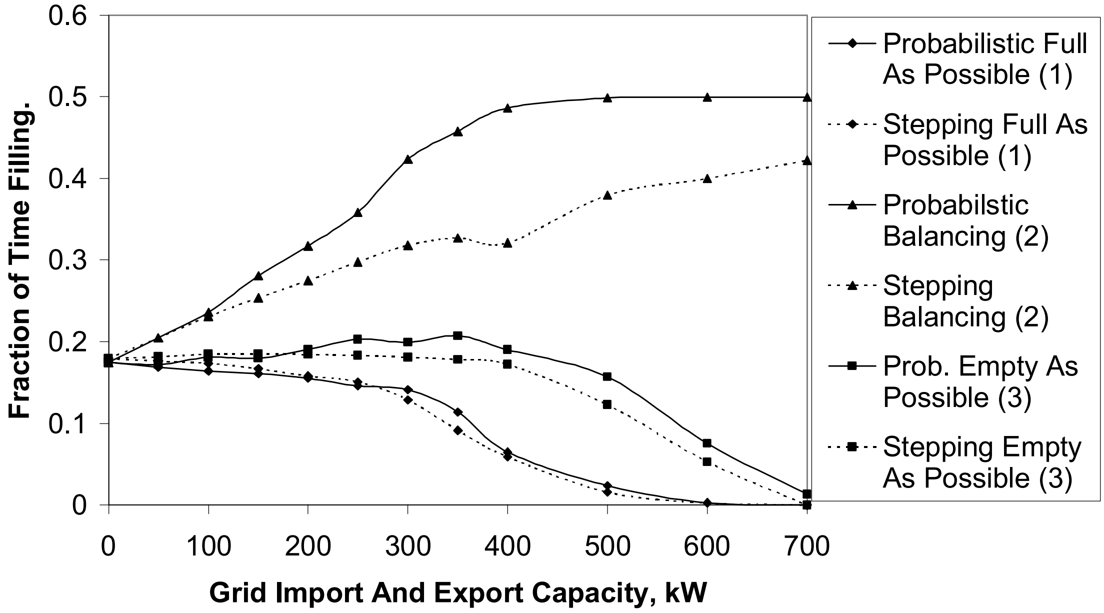

Figure 3.93 Fraction of Time that the Store is Filling

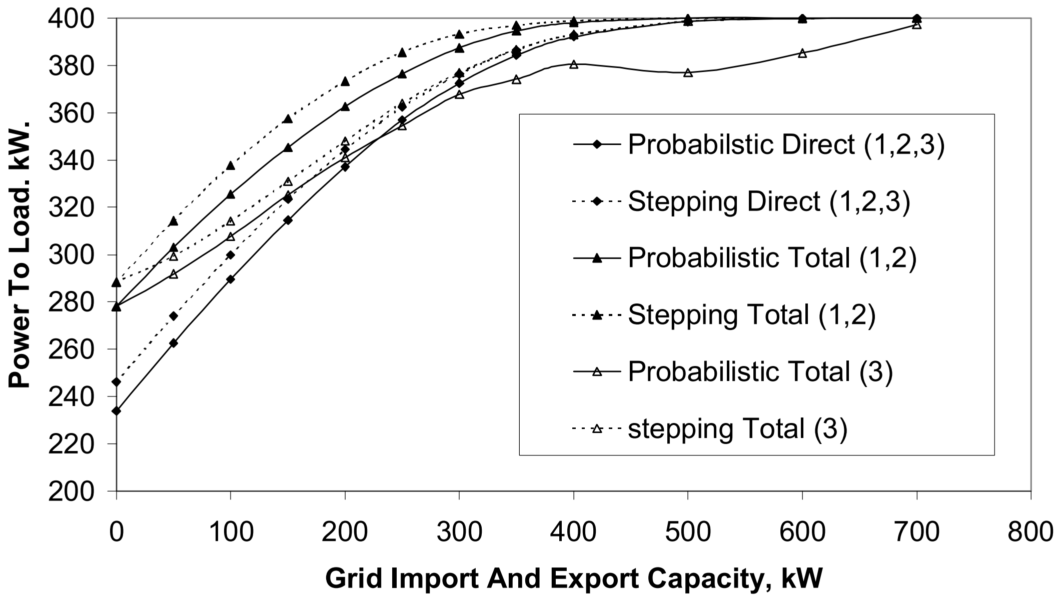

Figure 3.94 Power Supplied to the Loads With and Without the Store

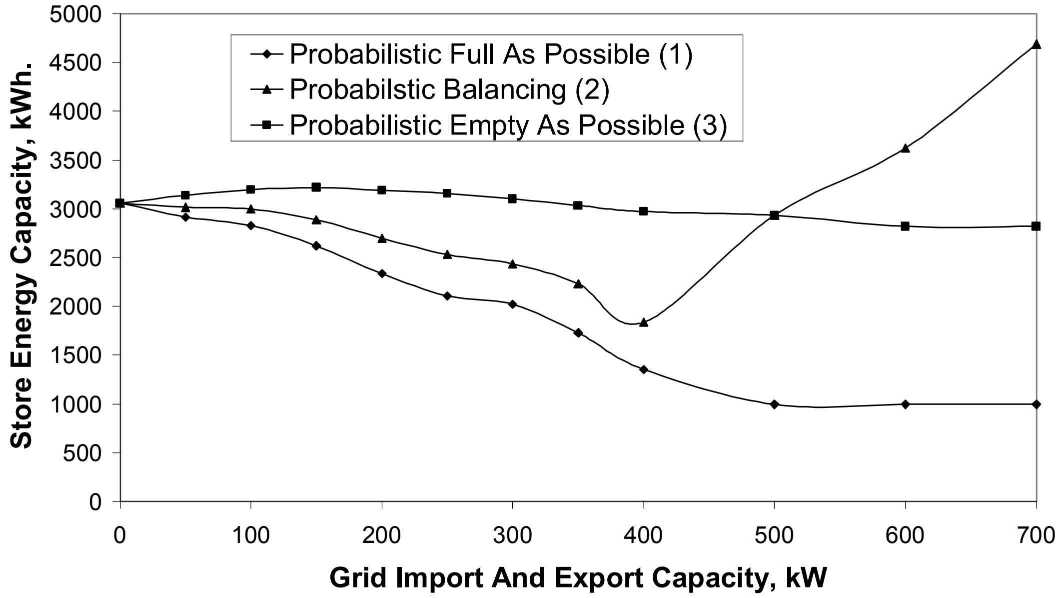

Figure 3.95 Required Store Size for a Cycle Time of 24 Hours

The curtailed power, fig. 3.88 is minimised in control options 2 and 3. In the probabilistic program, options 2 and 3 perform equally well, since the model assumes perfect weather forecasting, but in the time step program option 3 performs slightly better than option 2 due to large short-term surpluses of power affecting the ability of option 2 to absorb energy. Option 1 uses the grid to recharge the store even when there is a small power surplus, but then curtails power later in the day if the surplus is too large. Option 1 therefore causes a much larger curtailment of power, both in the probabilistic program and the time step program.

The unsatisfied demand, fig. 3.89 is minimised in control options 1 and 2. In the probabilistic program, options 1 and 2 perform equally well, since the model assumes perfect weather forecasting, but in the time step program option 1 performs slightly better than option 2 due to large short-term deficits of power affecting the ability of option 2 to supply energy. Option 3 uses the grid to export power even when there is a small deficit, but then leaves some demand unsatisfied later in the day if the deficit is too large. Option 3 therefore causes a much larger unsatisfied demand, both in the probabilistic program and the time step program.

Option 1 causes the store to spend more time full than options 2 or 3, fig. 3.91. Option 3 causes the store to spend more time empty than options 1 or 2, fig. 3.90 and option 2 causes the store to spend more time filling, fig. 3.93 and emptying, fig. 3.92, and the minimum of time full or empty. Option 1 causes the store to spend a little less time filling than option 3, fig. 3.93 probably because the store fills more rapidly when importing from the grid at the same time as absorbing surplus wind power.

Fig. 3.94 shows the increase in satisfied demand resulting from the presence of the energy store. The ‘direct’ power, without the benefit of the store, increases with the strength of the grid connection, but does not change with the store control option. As the grid import and export capacity increase to about 600kW, all possible power deficits can be met by the grid, and the average delivered power approaches the average demand of 400kW. The store in control options 1 or 2 increases the delivered power, with greatest benefit when the grid import and export are small. However, the store in option 3 increases the delivered power by a smaller margin. The probabilistic program may have a bug in its calculation at this point; at large grid capacities it actually predicts that the store reduces the delivered power compared to the no-store case when in control option 3.

The required store energy capacity, fig. 3.95 depends on the control option. A store in option 1 is sized to supply power through periods of deficit when the grid and renewables together cannot meet the demand. In control option 1, as the grid import capacity increases, the days when the store does most work change to days of very low wind speed. When the daily wind speed is low, its variability is also low, and the accumulated and discharged energy is small.

A store in control option 3 is sized to accumulate energy on days when the grid cannot export the entire surplus. As the grid export capacity increases, the days when the store does most work change to days of higher wind speed. Initially, this means more variable wind speeds and greater store size, but then at very large grid export capacity, the net surpluses become smaller and the store size reduces again, but the reduction in energy capacity is small.

A store in control option 2 behaves sometimes like option 1 and sometimes like option 3, so at small grid strengths, its energy capacity is an average of options 1 and 3. However, a store in option 2 also attempts to make the grid import and exported power as constant as possible. This results in the store absorbing lots of variability when the grid import and export capacities are large. The required store size apparently increases as a result, fig. 3.95, even though the store is actually needed much less when the grid capacity is large.

3.2.9 Varying Backup Generation

Figs. 3.96 to 3.103 show the effect of an increasing backup generation capacity. This effectively behaves like a grid connection that can import power but not export surplus power. Again, control options 1, 2 and 3 are compared and the option numbers are show in the legends of the figures.

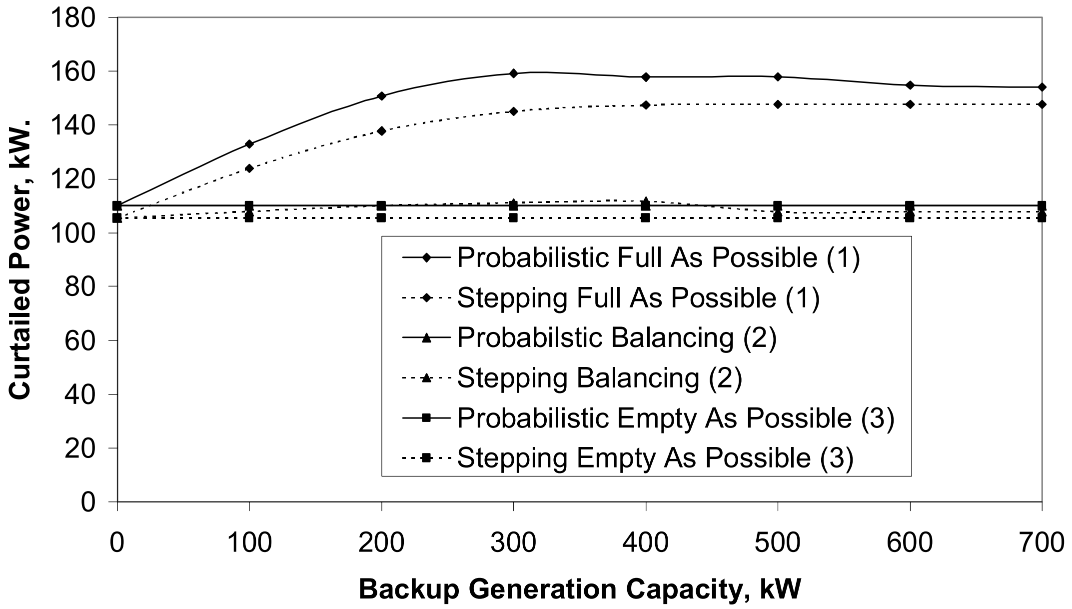

Figure 3.96 Power Curtailed Due To Finite Store Being Full

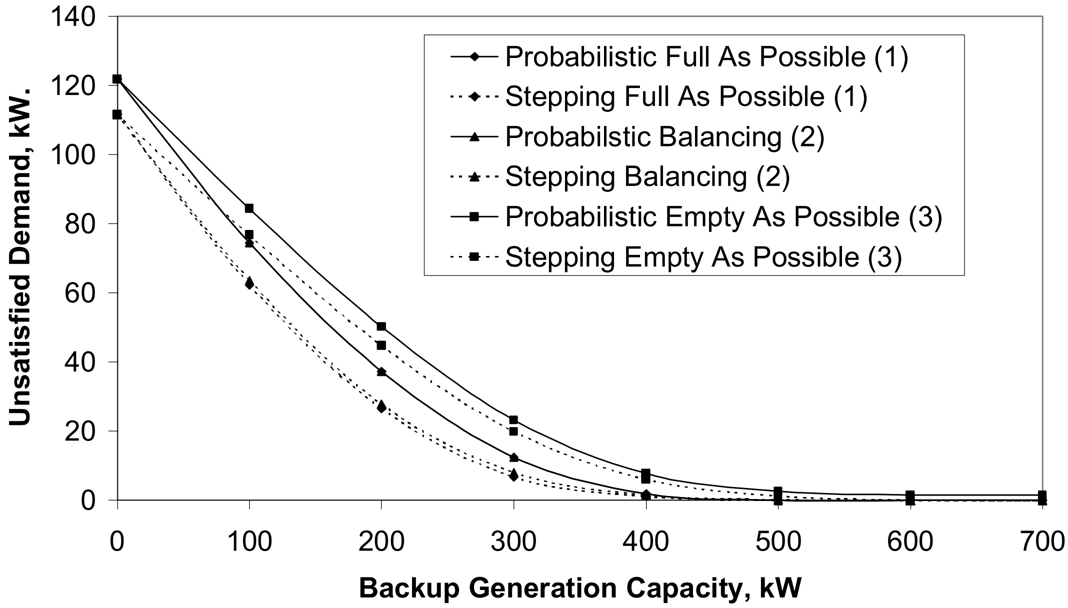

Figure 3.97 Unsatisfied Demand Due to Store Being Empty

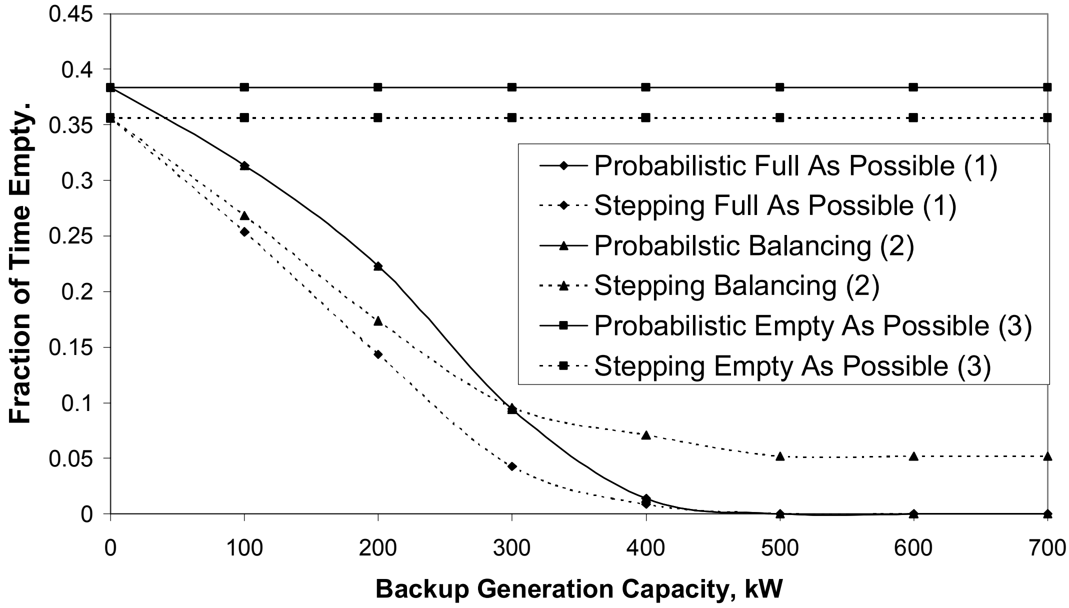

Figure 3.98 Fraction of Time that the Store is Empty

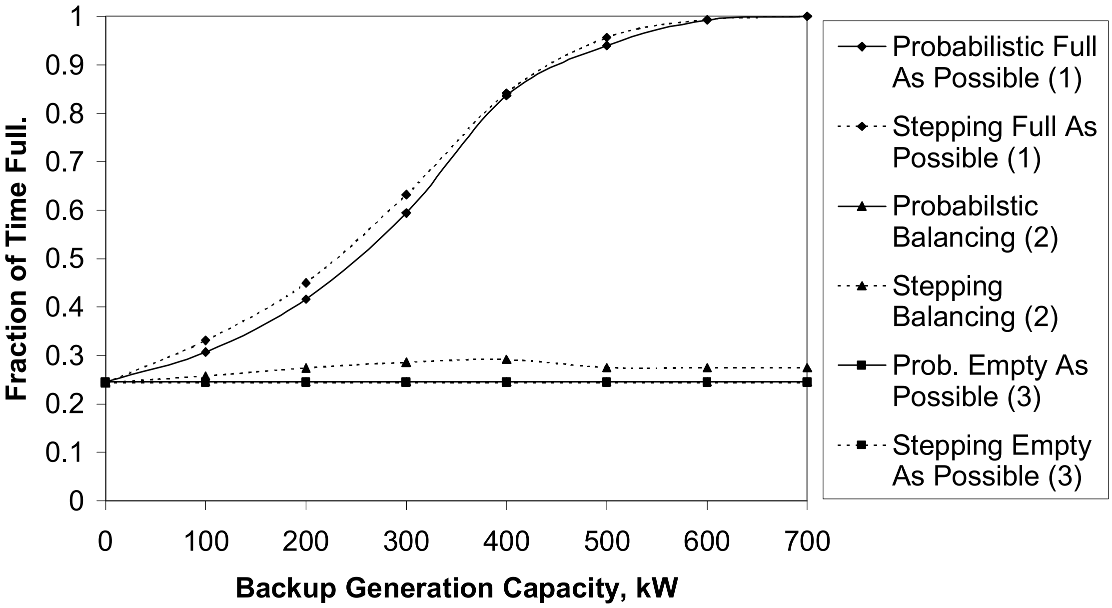

Figure 3.99 Fraction of Time that the Store is Full

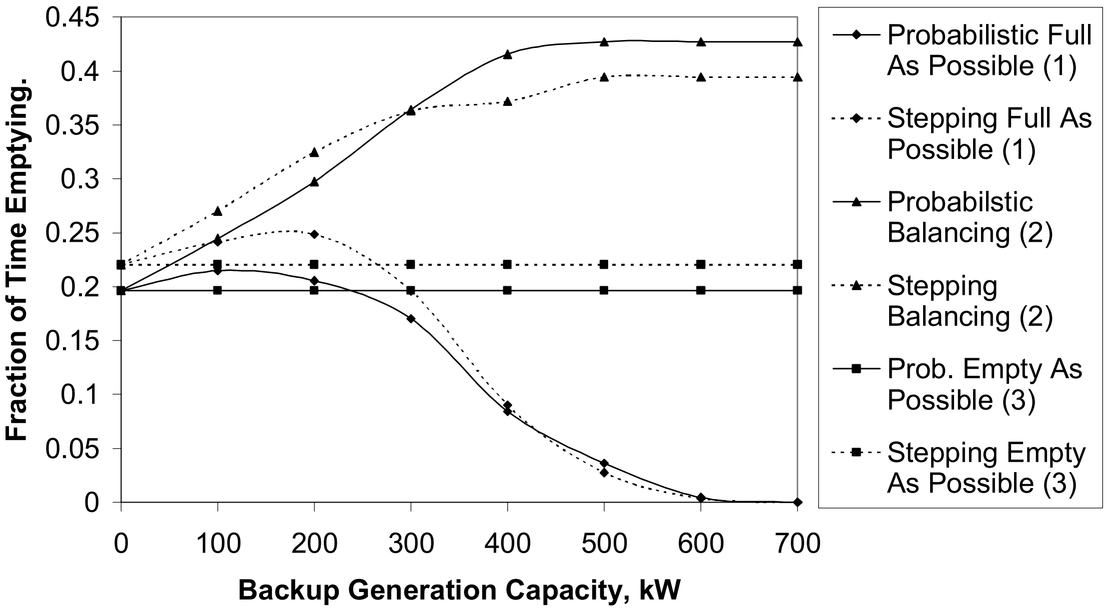

Figure 3.100 Fraction of Time that the Store is Emptying

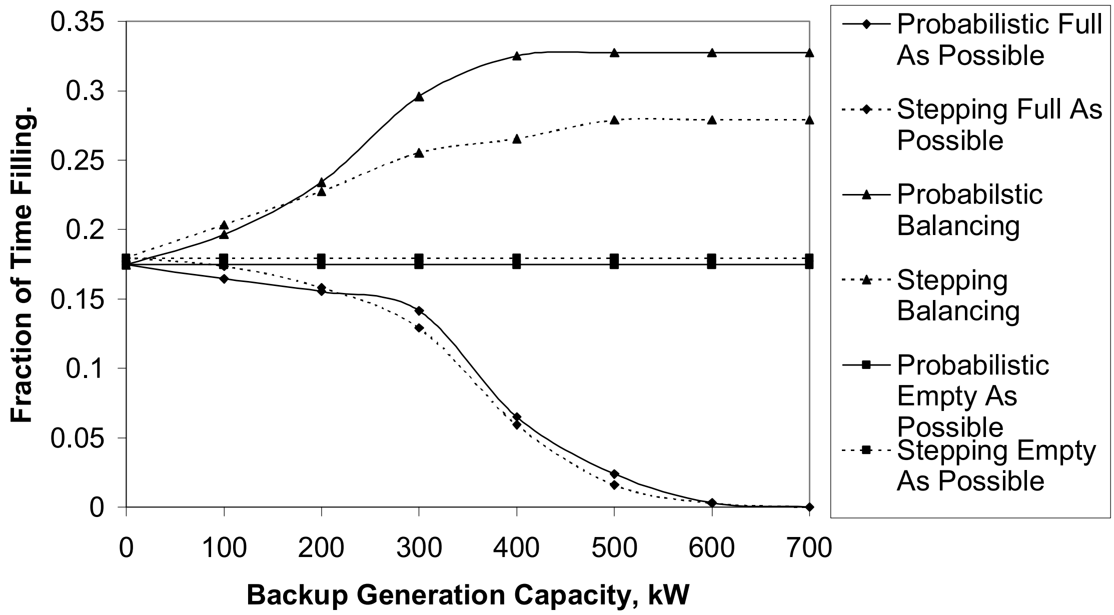

Figure 3.101 Fraction of Time that the Store is Filling

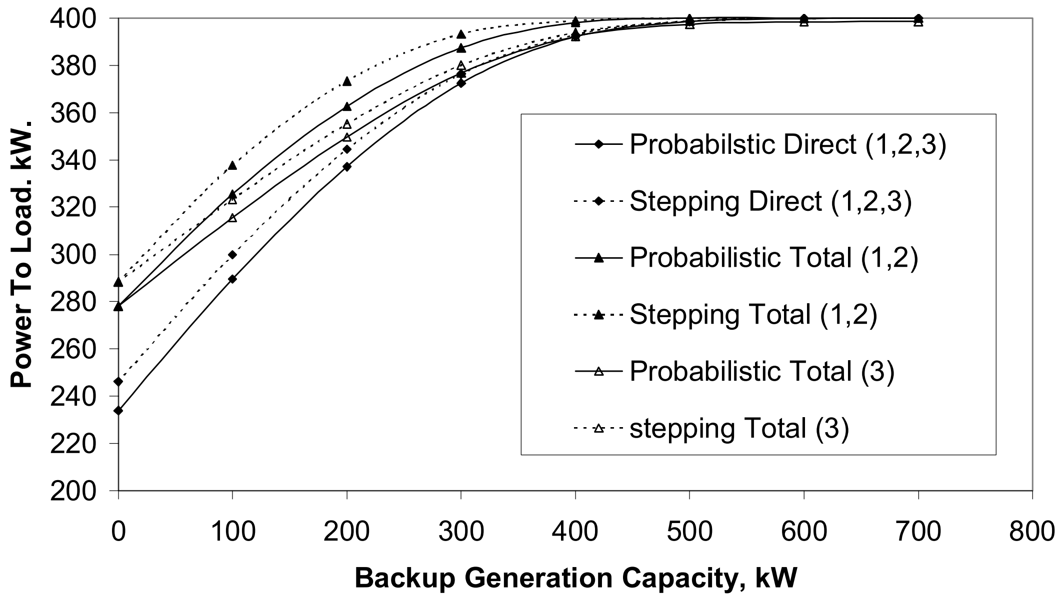

Figure 3.102 Power Supplied to the Loads With and Without the Store

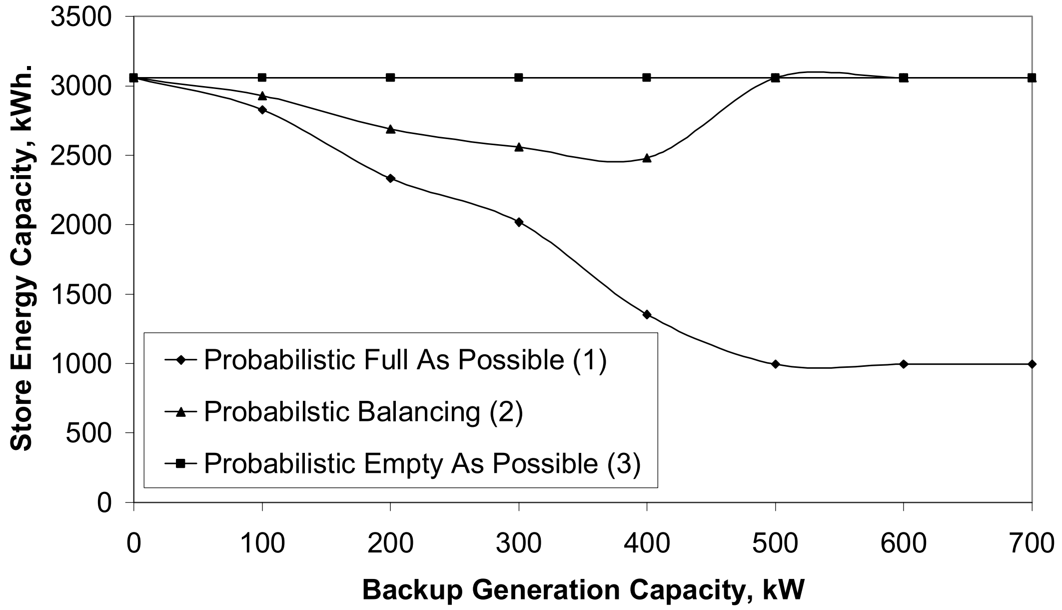

Figure 3.103 Required Store Size for a Cycle Time of 24 Hours

The curtailed power, fig. 3.96 is minimised in control options 2 and 3. Option 1 uses the backup generation to recharge the store even when there is a small power surplus, but then curtails power later in the day if the surplus continues and the store is full. Option 1 therefore causes a much larger curtailment of power, both in the probabilistic program and the time step program.

The unsatisfied demand, fig. 3.97 is minimised in control options 1 and 2. Option 3 does not operate the backup generation until the store is empty, but then leaves some demand unsatisfied later in the day if the deficit is too large. Option 3 therefore causes a much larger unsatisfied demand, both in the probabilistic program and the time step program.

Option 1 causes the store to spend more time full than the other control options, fig. 3.99. Option 3 causes the store to spend more time empty than other control options, fig. 3.98. In fact option 3 has no effect on the fraction of time that the store spends empty, full, emptying or filling compared to the zero backup case, because it only uses backup generation when the store is empty. Control option 2 causes the store to spend more time filling, fig. 3.101 and emptying, fig. 3.100, and the minimum of time full or empty.

Fig. 3.102 shows the increase in satisfied demand resulting from the presence of the energy store. The ‘direct’ power, without the benefit of the store, increases with the backup generation capacity, until about 600kW at which point all power deficits are satisfied by the backup generator. The store in control options 1 and 2 increase the delivered power more than option 3, with greatest benefit when the backup generation capacity is small. This is consistent with fig. 3.97 where control options 1 and 2 reduce the unsatisfied demand more than option 3 does.

The required store energy capacity, fig. 3.103 is again least for control option 1. An option 1 store is sized to supply power through periods when renewable energy together with backup generation cannot meet the demand. As the backup generation increases, the store does most work on days of very low wind speed. Wind power variability is then also low, and the accumulated and discharged energy is small.

The operation of a store in control option 3 is unaffected by the backup generation. The store size is therefore unchanged by the backup generation capacity.

A store in option 2 behaves sometimes like option 1 and sometimes like option 3, so at small backup generation capacities, the store size is an average of options 1 and 3. At large backup generation capacities, the store size reverts to being the same as in option 3 because the risk of unsatisfied demand is removed.

3.2.10 Comparison of Probabilistic with Time Step Methods

All the figures of section 3.2 show good agreement between the probabilistic and time-step models. As in the wind-only modelling programs, the probabilistic method often agrees better with the time step method when predicting power flows, e.g. curtailed power and unsatisfied demand, e.g. fig.

3.32 and fig. 3.33, than when predicting times spent full, empty, filling and emptying, e.g. figs. 3.34 to 3.37. This may be because the probabilistic method first calculates average power flows from renewable sources, to the load, to and from the store and curtailed power etc. It then uses these as inputs to its estimates of time fractions spent full and empty etc. Thus the calculations of time fractions are one step more removed from input data than the power flow calculations.

The 24-hour programs often give better agreement between probabilistic and time step methods, e.g. figs. 3.32 and 3.33 than the wind-only programs, e.g. figs. 3.9 and 3.10 when modelling power flows. A principle difference is the presence of seasonal variation in the24-hour probabilistic program, and this could be responsible for the improvement.

The 24-hour modelling programs show less good agreement between probabilistic and time-step methods when run in solar-only mode, figs. 3.80 to 3.87. This suggests that the solar probabilistic model needs refinement, or perhaps less should be expected of a probabilistic solar model. The solar model has certainly required much more complex calculation than the wind model.

The probabilistic method always shows similar trends to the time step method, even if it does not always agree in absolute level.

3.3 Computer Run Time and Program Complexity

Table 3.1 compares the size and performance of the four computer programs written for this thesis.

Table 3.1 Computer Program Lengths and Run Times

Program name	Description	Lines of code	Preparation and loading data, seconds	Time per case, seconds
finitestore MSc13.m	Wind-only probabilistic	655	Very small	1.6
finitestore stepping MSc7.m	Wind-only stepping	273	17	104
Probabilistic11.m	24-Hour wind, solar and load probabilistic	2036	Very small	42
Stepping3.m	24-Hour wind, solar and load stepping	607	34	98

The probabilistic program codes are 2 to 3 times longer than the time step codes that model the same electricity system. This increase is due to the added complexity of creating the probabilistic models. The mathematical complexity represents the physical complexity of random wind speed variations and random solar irradiance variations on top of daily and seasonal variations. The number of lines in each program includes the comments, but even the comments represent some of the thought required to write the code.

The time step programs take a measurable time to load the time series of weather data. The wind-only program takes approximately 17 seconds to load 4 years of 1-minute wind speed data. The 24-hour program takes approximately 34 seconds to load 4 years of 1-minute wind speed data, 4 years of 1-minute solar irradiance data and 1 year of half-hourly demand data. The demand data probably takes an insignificant length of time to load compared to the weather data. In contrast, the probabilistic programs appear to start immediately. The time required to load the demand profiles, the wind speed power spectrum and the solar irradiance power spectrum is negligible.

Both probabilistic programs do save a significant amount of computer calculation time compared to the time step programs. In the wind-only programs, the calculation time is cut by a factor of over 60. The wind-only probabilistic program only requires 2 of the nested loops: the period-average wind speed loop and the within-period wind speed loop, fig. 2.74.

In the 24-hour programs, the calculation time is cut by a factor of only 2.33. The 24-hour probabilistic program requires many nested loops: the month, daily average wind speed, daily average cloudiness factor, week or weekend, within-day wind speed, within-day cloudiness factor and hour-of-day, as described in section 2.10.3 and fig. 2.7.4. The effect of the nested loops is offset by efficient coding and by a reduction in the number of possible wind speed and wind power levels. However, one of the major benefits of a probabilistic method, the reduction in computer run time, has almost evaporated.

The 24-hour time step program takes slightly less time per case than the wind-only time step program, probably as a result of more efficient coding. This is despite having to perform calculations on solar irradiance and electrical demand as well as wind speed data.

4 Conclusions

A probabilistic method has been developed for modelling energy storage used with wind power and solar power. This method is based on a spectral analysis of variations in wind speed and solar irradiance variations and is an alternative to the standard time step method. It has the advantages that the required input data is less and that it is faster to run. Key features and results are listed below.

1. The probabilistic method gives promising results that are generally accurate when compared to the time step method.
2. The generic wind speed spectrum derived from various sites around the British Isles gives results as good as those derived from the RAL spectrum when used in the probabilistic method and when compared with time step results obtained using RAL wind speed data. This is a good indication that the generic spectrum is generally applicable and useful for modelling wind speed at any site in the British Isles, and possibly in any maritime climate.
3. The probabilistic predictions of system power flows and losses are mostly very accurate when compared with time step predictions.
4. The probabilistic predictions of the fractions of time that a store spends full, empty, filling or emptying are usually good, but not as close as the power flow predictions when compared with time step predictions.
5. The wind-only probabilistic program was relatively simple to code and easily models stores that operate on all cycle times from minutes to years. However, the 24-hour probabilistic program, modelling wind, solar and load variations, is more complex in terms of both the underlying science and the number of lines of computer code.
6. The wind-only probabilistic program runs at approximately sixty times the speed of the time step wind-only program. The 24-hour probabilistic program runs at approximately twice the speed of its time step equivalent and retains the advantage of requiring far less input data.

7. The 24-hour probabilistic program is less accurate when run in solar-only mode than when run in wind-only mode or wind-and-solar model.
8. In the UK, aggregated electrical demand standard deviation is 24% of the mean value; typical wind power standard deviation is 90% of its mean value; typical solar power standard deviation is 170% of its mean value. Wind power is therefore much more variable than electrical demand. Solar power is even more variable than wind power, if rather more predictable than wind power, and accounts for the increased difficulty in determining an accurate probabilistic representation.
9. The current versions of the probabilistic programs are not optimised for accuracy and computer run time. Some parts of the code make gross modelling simplifications while other areas may be overly detailed. The computer codes need refinement, especially in the modelling of solar power. This should be the focus of further work.

4.1 Advantages and Disadvantages of the Probabilistic Method Confirmed by This Thesis:

Advantages

1. The probabilistic method requires less input data
2. The computer run time is less than the standard time step method, especially for the wind-only probabilistic method.
3. Its results are generally accurate when compared to the time step method, especially when predicting total average energy flows.

Disadvantages

1. It is computationally more complex, especially when modelling solar power variations
2. The probabilistic method is not quite as accurate as the time step method when predicting the fractions of time that a store spends full, empty, filling or emptying.
3. It cannot easily model increased complexities, such as complex control strategies or other intermittent sources of renewable generation, e.g. tidal power or wave power.

4.2 Further Work and Potential Applications

The probabilistic programs are not yet in a form that could be used commercially or by anyone not familiar with the method. No user interface has been written; inputs are changed by directly modifying the Matlab code. A graphical user interface would make the programs more accessible.

More functionality could and should be added to the programs. A future program should be capable of modelling wind power, solar power and variable demand over any working period (store cycle time). It may also be desirable to model other sources of intermittent renewable energy and an electricity demand that depends on weather conditions. It would definitely be useful to model two or more types of storage working together, e.g. flywheels for the short term, batteries for medium term and hydrogen for long term storage.

At the same time, the program could be simplified, perhaps by making simpler assumptions about the variations in solar power and 'lumping' all stochastic variation into one composite spectrum. The probabilistic method would then use just two spectra: a periodic one in which seasonal and diurnal correlations can be accounted, and a stochastic one in which wind power, solar power and any other random variables are assumed to be independent. This method would considerably reduce the number of nested loops, figure 2.74 and further reduce computer run time.

Returning to the value of energy storage, section 1.4, the program already includes a very simple method of accounting costs of each component of an electricity system (although not presented in this thesis). This method could be developed to perform an economic comparison between systems with and without energy storage. If the probabilistic program can be made to run fast enough, it may even be possible to make the program perform an economic optimisation with regard to the size of components in the system.

It is anticipated that the probabilistic method would be useful for first stage feasibility studies of renewable energy systems and approximate sizing of components. Programs already exist for more detailed studies, e.g. HOMER

(Lambert, Lilienthal 2003) and HYBRID2 (Manwell et al. 1998) but these demand more input data and specific size options of system components.

Appendix A

List of Publications

Some of the results presented in this thesis have been published as contributions to conferences or articles in journals; they are listed below.

Journal: (Barton, Infield 2004) This journal paper describes the potential benefit of energy storage in increasing the penetration of wind powered generation onto a weak grid in which the level of embedded generation is limited by voltage rise. This paper also describes the operation of a probability matrix and compares various different energy storage technologies.

Conference: (Barton, Infield 2005a) This invited, refereed paper, was presented at the IEEE Power Engineering Society General Meeting conference, 2005, and is included in the conference proceedings. It describes the operation of a probability matrix and includes calculations of fractions of time that a store spends full, empty, filling and emptying using the probabilistic method. The probabilistic method is used on all time scales from minutes to months.

Journal: (Barton, Infield 2005b) This journal paper and oral presentation shows results of a probabilistic calculation performed on an electricity system including wind power, solar PV power and an energy store suitable for a cycle time of 24 hours and supplying a time-varying load of variable size.

Copies of the publications are attached below.

Appendix B

Calculation of Fractions of Time Spent Full or Empty

B.1 Operating States of an Energy Store

The probabilistic method assumes that at any instant in time, the energy store may be in one of four operating states:

Full: Energy supply exceeds demand and the store is full
Filling: Energy supply exceeds demand but the store is not yet full
Empty: Energy demand exceeds supply and the store is empty
Emptying: Energy demand exceeds supply but the store is not yet empty

The model makes an important simplifying assumption that the periods of time (store operating periods) used in the analysis may be divided into ones where supply exceeds demand, and other periods in which demand exceeds supply. Corrections must be made for store losses, store charge rates and discharge rates, to calculate actual net supply and demand, and from these, the power flows into and out of the store.

The 24-hour probabilistic modelling program allows different options for backup generation / grid import or grid export. These options are discussed in section 2.12.4 and may result in store operating periods in which supply and demand are balanced. The remainder of this appendix is devoted to periods when the backup generation and/or grid connection is absent or is insufficient to balance the net surplus or deficit of energy to or from the store.

During some time periods, more power flows into the store than out, on average. In these periods, the store is filling, full or emptying, but is assumed to be never empty, fig. B.1. The net surplus of energy is balanced by curtailed power when the store is full.

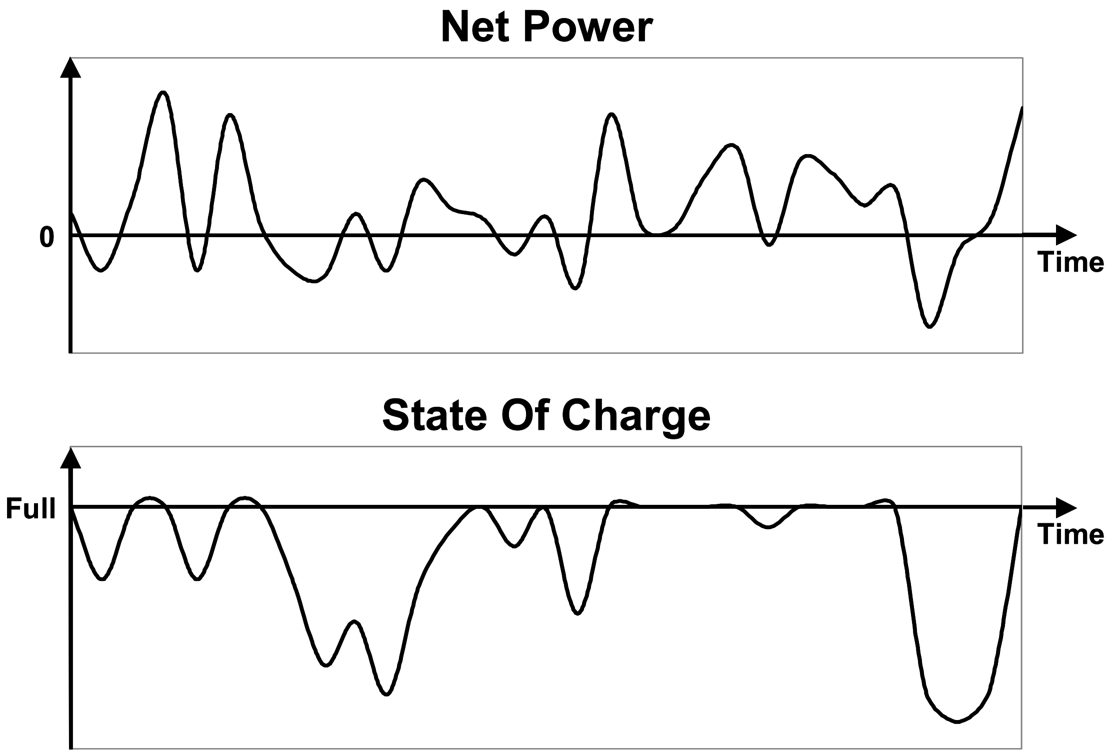

Figure B.1. Operation of a store where energy supply exceeds demand

For other periods, more power flows out of the store on average. Then the model assumes that the store is either empty, emptying or filling but is never full, fig. B.2. The net shortfall of energy is balanced by the unsatisfied demand when the store is empty.

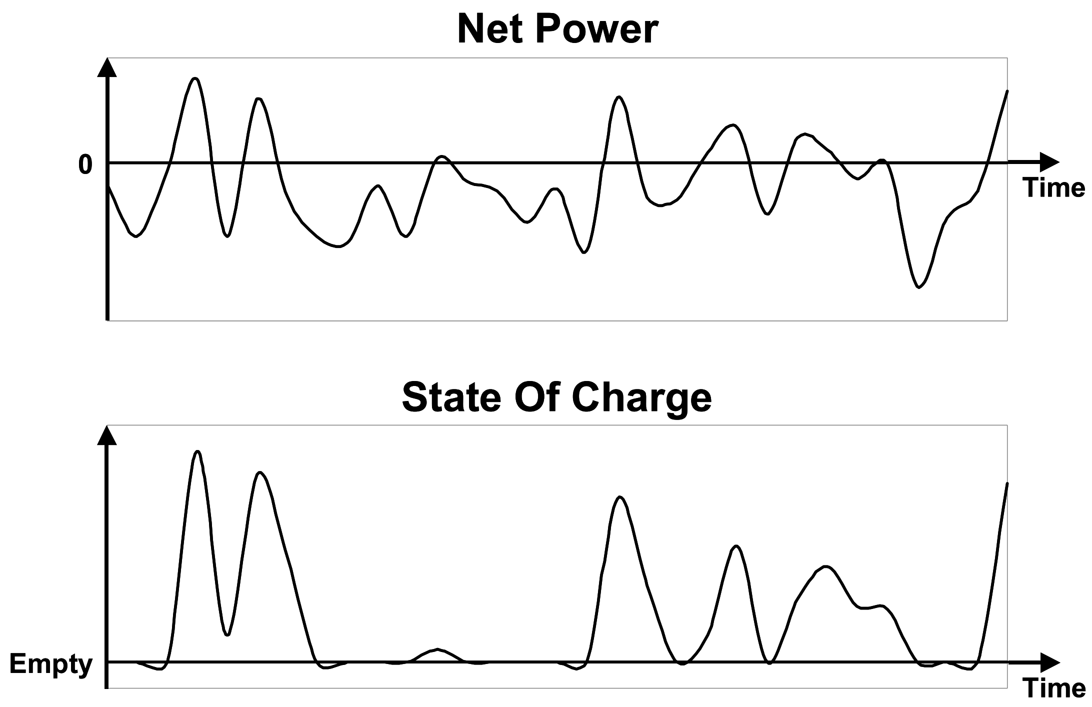

Figure B.2. Operation of a store where energy demand exceeds supply

The following paragraphs describe how the probabilistic model calculates the fractions of time spent full, empty, filling and emptying.

B.2 When Energy Into and Out of the Store is Almost Balanced

The algorithm uses a previously calculated probability distribution of net power into or out of the store, shown diagrammatically in figs. B.3 and B.4 below. The probability distribution has already been corrected for store efficiency, parasitic store losses and maximum store charging and discharging rates, section 2.10.7.

B.2.1 When The Store is More Empty than Full

Let us consider a period in which average demand exceeds average supply, fig. B.3.

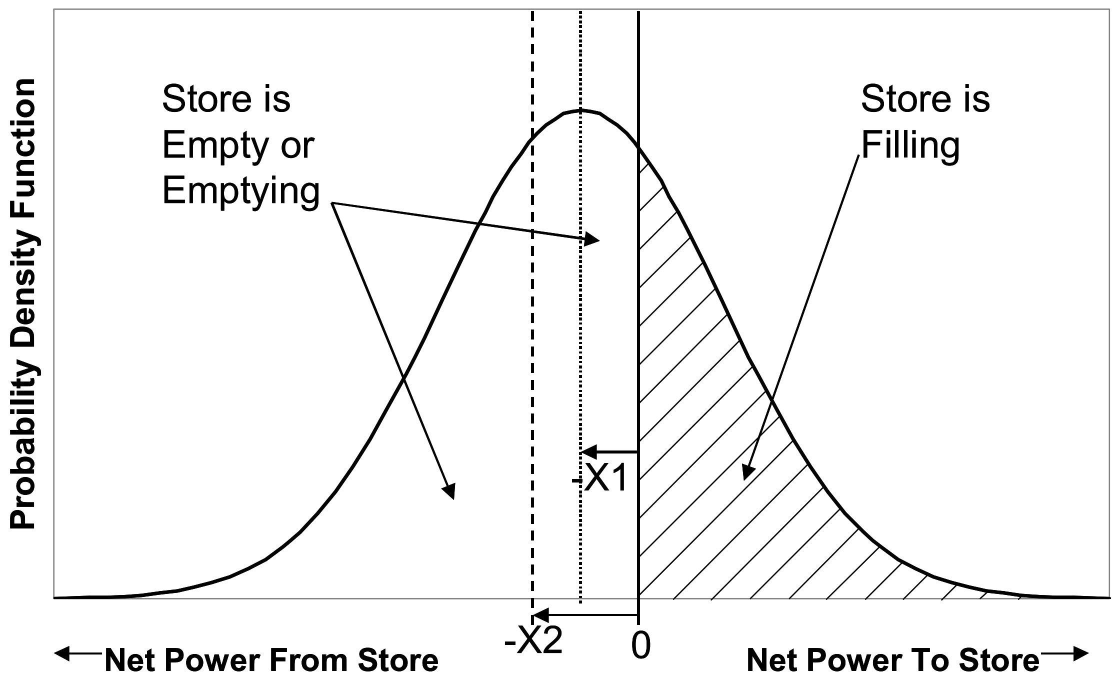

Figure B.3. Probability distribution of net power in a period when average demand exceeds supply.

The shaded area represents the probability that at any instant the supply exceeds demand. During this time, the store cannot be empty or emptying, and the model assumes that the store must be filling. Thus the fraction of time that the store is filling can be directly calculated from the probability distribution. The unshaded area represents the probability of empty and emptying. More calculation is required to split this probability according to these two states. The model assumes that the net power rapidly reverses, and that the distribution of net power is the same whether the store is empty or emptying. There is no assumed time-dependence of net power on net state-of charge of the store. We do know the average power deficit over the whole period, X_1. This is the demand that is not satisfied due to the store being empty. We also know that when the store is empty, the instantaneous net power must be negative (otherwise the store would start to refill). The model assumes that the average power deficit when empty or emptying is given by the centre of gravity, X_2 of the unshaded portion of the probability distribution, to the left of zero.

Then: $Empty\_Time \times X_2 = Total\_Time \times X_1$

And: $Fraction\_of\_time\_that\_store\_is\_empty = \dfrac{X_1}{X_2}$

The fraction of time that the store is emptying is the remaining time in the period:

$$Emptying\_fraction = 1 - Empty\_fraction - Filling\_fraction$$

Note that when the store is empty, some of the demand is still satisfied, and only a portion of it is unsatisfied. The 'hours of power cut' calculated later is not the same as the fraction of time that the store is empty.

B.2.2 When The Store is More Full than Empty

Now let us consider a period in which average supply exceeds average demand, fig. B.4.

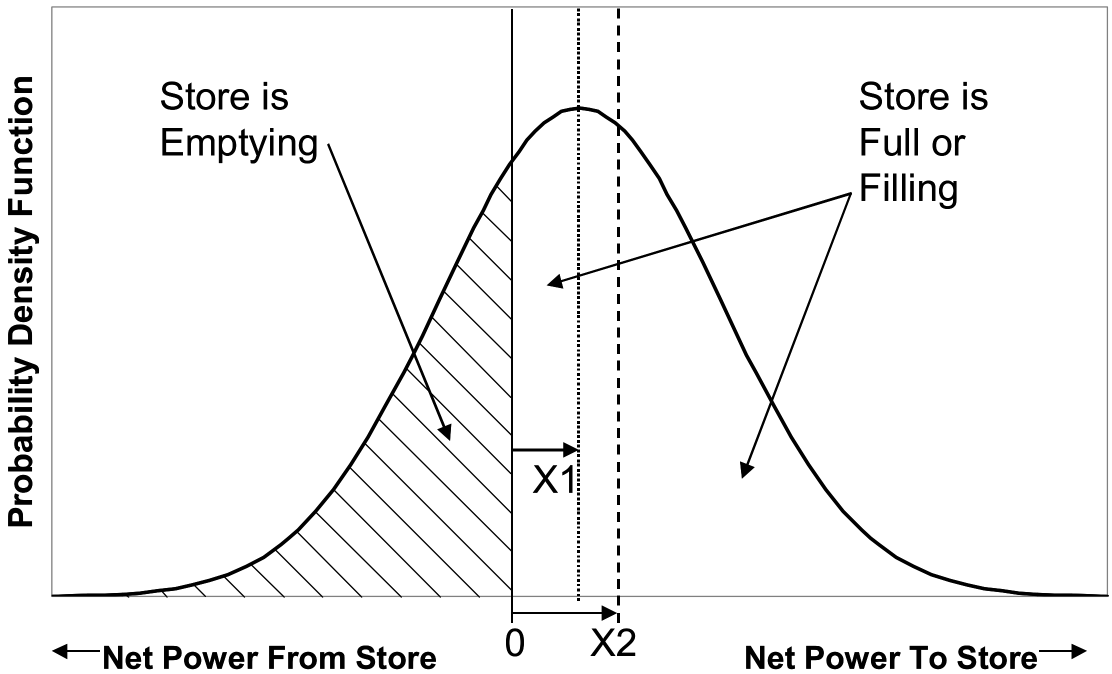

Figure B.4. Probability distribution of net power in a period when average supply exceeds demand.

The shaded area in fig. B.4 represents the probability that at any instant the demand exceeds supply. During this time, the store cannot be full or filling, and the model assumes that the store must be emptying. Thus the fraction of time that the store is emptying can be directly calculated from the probability distribution. The unshaded area now represents the probability of full and filling. As before, this probability must be split into the two time fractions. The model assumes that the distribution of net power is the same whether the store is full or filling. We know the average power surplus over the whole period is X_1. We also know that when the store is full, the instantaneous net power must be positive (otherwise the store would start to empty). The model assumes that the average power surplus when full is given by the centre of gravity, X_2 of the unshaded portion of the probability distribution, to the right of zero.

Then: $Full\_Time \times X_2 = Total\_Time \times X_1$

And: $Fraction\_of\_time\_that\_store\_is\_full = \dfrac{X_1}{X_2}$

The fraction of time that the store is filling is the remaining time in the period:

$$Filling\_fraction = 1 - Emptying\_fraction - Full\_fraction$$

B.3 When Supply And Demand Are Very Different – The Triangular Wave Method

The above method works very well over a wide range of operating conditions. However, due to autocorrelation of the power time series, when the net power is above average, it stays above average for some time, and when the net power is below average, it stays below for some time. It is concluded that the net power and state-of-charge are not truly independent.

When average demand is much greater than supply, the store spends most of its time empty. When the store does occasionally accumulate a little energy at instants of above-average power, it therefore takes a relatively long time to empty again compared to the rate if the instantaneous power were closer to the mean for the period.

Conversely, when average supply is much greater than demand, the store spends most of its time full. When the store occasionally empties a little at instants of below-average power, it therefore takes a relatively long time to refill compared to the rate if the power were closer to the mean for the period.

The times that the store spends empty, full, emptying and filling cannot be directly calculated because all the temporal information has been lost; all we have is a spectrum of wind speed variations and a probability distribution of net power. To deal with this, the probabilistic model uses an approximation to the time variations. It assumes that the net power time variation follows a triangular waveform. Based on this, the required calculations can be made as described below. Let us first consider a period in which average demand exceeds average supply, fig. B.5.

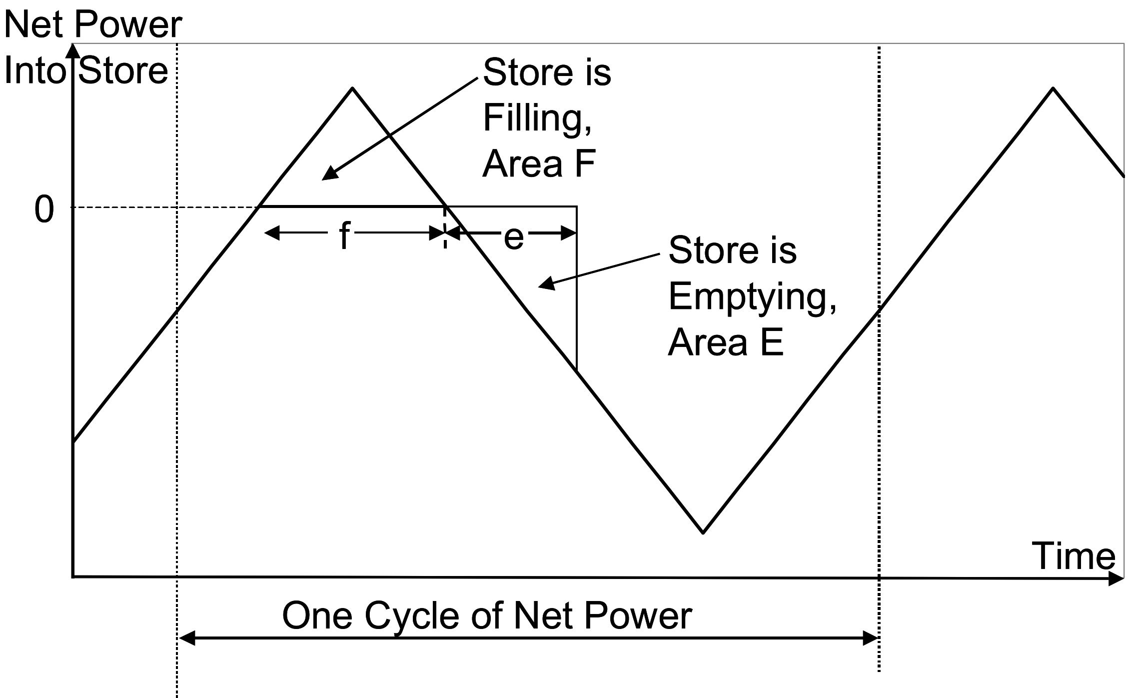

Figure B.5. Modelling a store that spends most of its time empty

The fraction of time that the store spends filling, f is calculated from the probability distribution, as before, fig. B.3. We do not have to know the length of a typical cycle of power variation (the period of the triangular wave). All time fractions are normalised by this period. The time taken to empty again, e is calculated such that the areas of the filling triangle, F and emptying triangle, E are equal in fig. B.5, after adjustment for store efficiency. If the efficiency, η is constant, defined as:

$$\eta = \frac{EnergyOut}{EnergyIn}$$

And assuming that the state of charge at the end of any time period is the same as at the beginning.

Then: $e = f\sqrt{\frac{\eta}{2}}$

Now let us consider a period in which average supply exceeds average demand, fig. B.6.

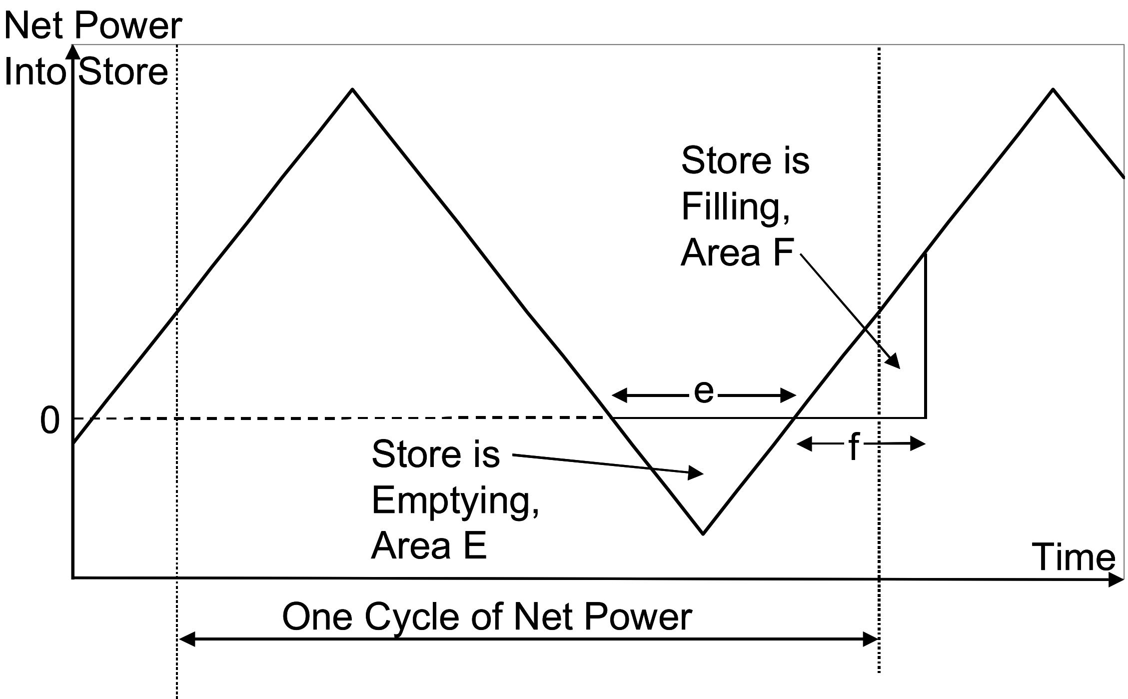

Figure B.6. Modelling a store that spends most of its time full

The fraction of time that the store spends emptying, e is calculated from the probability distribution, as before, fig. B.4. The time taken to refill, f is calculated such that the areas of the filling triangle, F and emptying triangle, E are equal in fig. B.6, after adjustment for store efficiency. If the efficiency, η is constant, defined as:

$$\eta = \frac{EnergyOut}{EnergyIn}$$

And assuming that the state of charge at the end of any time period is the same as at the beginning.

Then: $f = \frac{e}{\sqrt{2\eta}}$

This triangular wave method was found to give such good results that it was considered as an option for all conditions, even when energy into and out of the store is almost balanced. This required a more complex formula, using one root of a quadratic equation. In a store experiencing slightly more energy out than in, the emptying after filling would take not only the ‘down’ slope but also part of the following ‘up’ slope, fig. B.7.

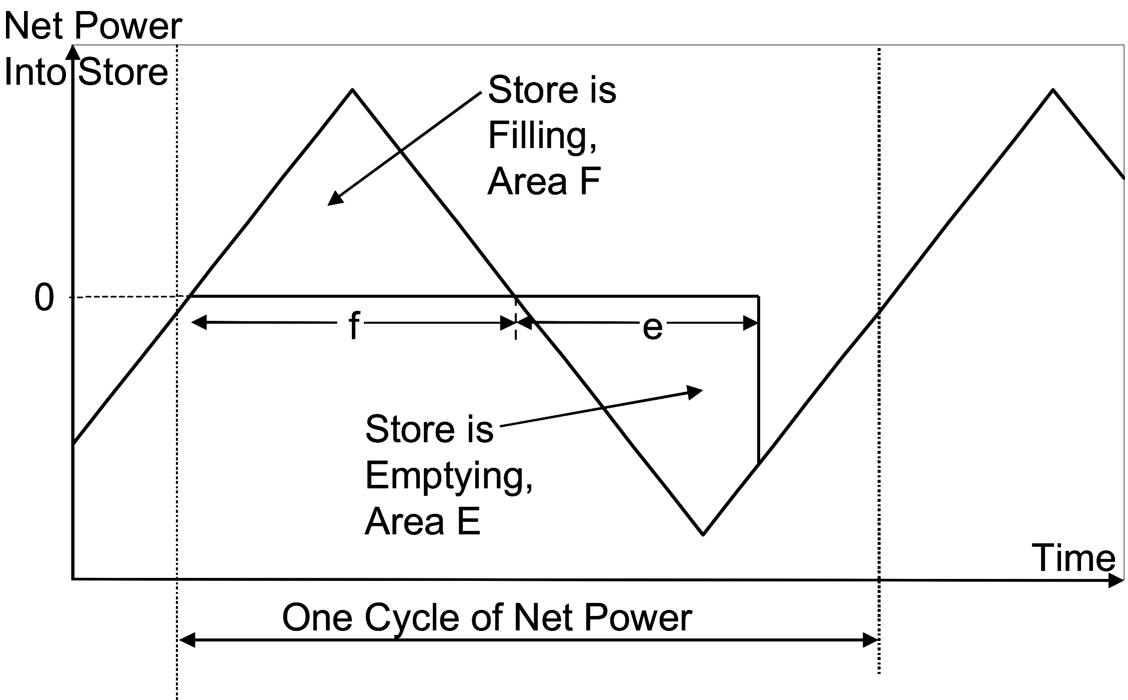

Figure B.7. Modelling a store that spends a small fraction of time empty

Similarly, in a store with slightly more energy in than out, the refilling after emptying would take not only the 'up' slope but also part of the following 'down' slope. This approach was tried, but the results were not as good as the first method shown in section B.2, when compared with the time-stepping results.

The probabilistic method now achieves the best of both worlds. When energy into and out of the store is almost balanced, the first method is used. When the energy is far from balanced, the triangular wave method is used. When energy out of the store exceeds energy into the store, the triangular wave method always gives the longer emptying time. When energy into the store exceeds energy out of the store, the triangular wave method always gives the longer refilling time. Thus the probabilistic method always evaluates both methods and takes the longer of the two emptying times, or the longer of the two refilling times:

So when energy out of the store exceeds energy into the store:

Emptying time, $e = \max\left[1 - f - \frac{X1}{X2}, f\sqrt{\frac{\eta}{2}}\right]$

Then empty time $= 1 - e - f$

And when energy into the store exceeds energy out of the store:

Refilling time, $f = \max\left[1 - e - \frac{X1}{X2}, \frac{e}{\sqrt{2\eta}}\right]$

Then full time $= 1 - e - f$

B.4 modification of the Triangular Wave Method for Truncated Probability Distributions

The formulation described above works well for many cases, but not when the probability distribution ends abruptly, for example due to the shape of the turbine power curve, or due to the store having a limited charge or discharge rate. For example, in fig. B.8, the left-hand tail of the probability distribution function has been shortened and concentrated into a second peak. The peak may represent a single operating state of the system with a significant probability, e.g. when the wind speed is too low to generate any electricity.

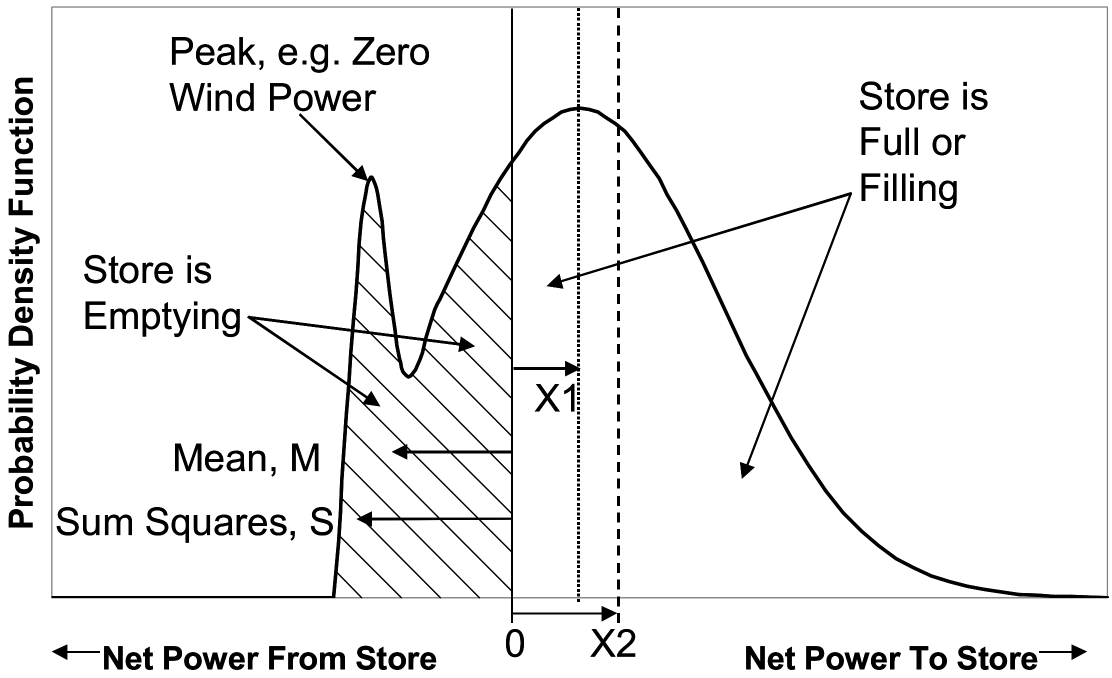

Figure B.8. Truncated probability distribution of net power, in a period when energy to the store exceeds energy from the store.

This situation has again been modelled using a triangular wave, but now with a truncated triangular wave, fig. B.9.

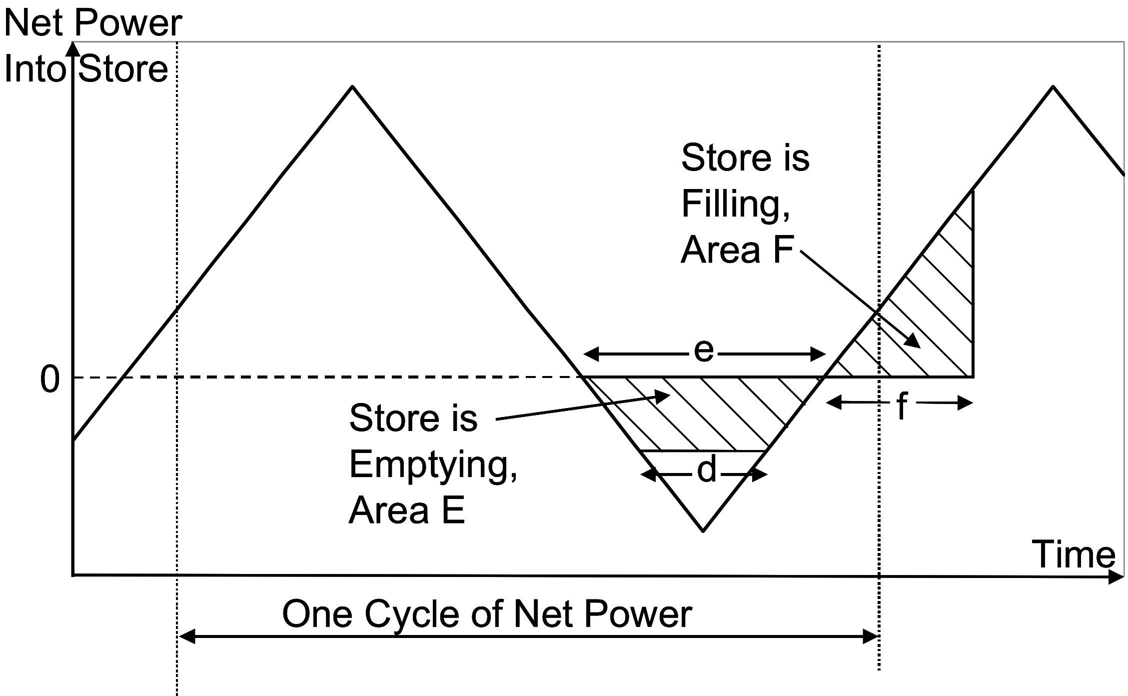

Figure B.9. Modelling a store that spends most of its time full, using a truncated triangular wave

The maximum power out of the store, represented by the truncation of the triangle of time d, may not be the same as the actual maximum power coming from the store as described by the probability distribution function. The actual distribution may include a small tail of low probability e.g. due to variations in electrical demand. To accommodate this, and other possible variations in the shape of the probability distribution, the truncation of the triangular wave, d is calculated from the first and second moments of the actual probability density function, M and S, fig. B.8.

Mean, $M = \dfrac{\sum_{x=-\infty}^{x=0} -p_i x_i}{\sum_{x=-\infty}^{x=0} p_i}$

Sum of Squares, $S = \dfrac{\sum_{x=-\infty}^{x=0} p_i x_i^2}{\sum_{x=-\infty}^{x=0} p_i}$

Then the fraction of time, d that the triangular wave spends at the limit is calculated to give the same values of M and S as the actual probability distribution

$$d = e\left\{\frac{4M^2}{3S} - 1 + \frac{M}{3S}\sqrt{\left[16M^2 - 12S\right]}\right\}$$

When $4M^2 > 3S$, d is real and positive and is used to calculate the refilling time. However, when $4M^2 < 3S$, d is complex and its real part is negative. Then d is set to zero.

The full formula for the fractional refilling time is: $f = \sqrt{\frac{e^2 - d^2}{2\eta}}$

Thus the formula reverts to the simple triangular wave formula when d=0

In the general case: Refilling time, $f = \max\left[1 - e - \frac{X1}{X2}, \sqrt{\frac{e^2 - d^2}{2\eta}}\right]$

And full time is still $= 1 - e - f$

Similar modifications apply to the case of a period in which power out of the store exceeds power into the store. Now M and S refer to the positive portion of the net power probability distribution:

Mean, $M = \frac{\sum_{x=0}^{x=+\infty} p_i x_i}{\sum_{x=0}^{x=+\infty} p_i}$

Sum of Squares, $S = \frac{\sum_{x=0}^{x=+\infty} p_i x_i^2}{\sum_{x=0}^{x=+\infty} p_i}$

Let c be the fraction of time that the power into a store is limited. The total fraction of time that the store is filling is f. The formula for the fractional time of limited power is otherwise unchanged:

$$c = f\left\{\frac{4M^2}{3S} - 1 + \frac{M}{3S}\sqrt{\left[16M^2 - 12S\right]}\right\}$$

The formula for the fractional emptying time is: $e = \sqrt{\frac{\eta\left(f^2 - c^2\right)}{2}}$

This also reverts to the original triangular wave formula when c=0.

In the general case: Emptying time, $e = \max\left[1 - f - \frac{X1}{X2}, \sqrt{\frac{\eta\left(f^2 - c^2\right)}{2}}\right]$

Appendix C

The Search for an Improved Wind Speed Distribution

Weibull distributions tend to fit wind speed data well for relatively large standard deviations (small Weibull shape factor). However, as seen in section 2.6.3, when the standard deviation is small compared to the mean, and the shape factor of the equivalent Weibull distribution (with the same mean and standard deviation) becomes large, then the skewness of the Weibull distribution becomes low or even negative. Real wind speed distributions almost always have highly positive skewness, even when the standard deviation is small. However, neither Weibull nor normal distributions possess this characteristic. This has led to a search for a more suitable family of distributions to model wind speed probabilities. Several distributions are promising: the log-normal, Chi, Chi squared and gamma distributions for example. As will be seen below, the non-central chi squared distribution with 2 degrees of freedom offers an alternative that always has a positive level of skewness.

The Rayleigh distribution (a Weibull distribution with a shape factor of exactly 2) is very limited but has some basis in theory (Hassan, Sykes 1990). This derivation of the Rayleigh distribution from a bi-variate normal distribution assumes that the wind is isotropic; that it is uniformly distributed with no prevailing wind direction. This raises an interesting question: What would be the distribution if the wind did have a prevailing direction? What if the north-south and east-west wind components were still each normally distributed, but the centre of those distributions were not at zero wind speed? Here we refer to this distribution as an Offset Circle - Circular Normal distribution.

The following derivation is an extension of the theory given in (Hassan, Sykes 1990). Consider wind speed components in the east-west and north-south directions of x and y respectively. The resultant wind speed is:

$$U = \sqrt{x^2 + y^2}$$

Assuming that the wind is isotropic (uniformly distributed with no prevailing wind), then each wind speed component will be normally distributed about zero:

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[\frac{-(x-\mu_x)^2}{2\sigma^2}\right]$$ Where μ_x is the mean in the x-direction

$$f(y) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[\frac{-(y-\mu_y)^2}{2\sigma^2}\right]$$ Where μ_y is the mean in the y-direction

Then $f(x,y) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[\frac{-(x-\mu_x)^2}{2\sigma^2}\right] \times \frac{1}{\sigma\sqrt{2\pi}} \exp\left[\frac{-(y-\mu_y)^2}{2\sigma^2}\right]$

$$= \frac{1}{2\pi\sigma^2} \exp\left[\frac{-(x-\mu_x)^2-(y-\mu_y)^2}{2\sigma^2}\right]$$

$= \frac{1}{2\pi\sigma^2} \exp\left[\frac{-r_\mu{}^2}{2\sigma^2}\right]$ Where r_μ is the wind speed vector from point (μ_x,μ_y) to point (*x*,*y*).

For the isotropic case, $\mu_x = \mu_y = 0$, and $r_\mu = U$. The probability density is then:

$$f(x,y) = \frac{1}{2\pi\sigma^2} \exp\left[\frac{-U^2}{2\sigma^2}\right]$$

To obtain the probability density as a function of wind speed, U this equation must be integrated in polar coordinates with respect to angle (wind direction), θ:

$$f(U) = \int_{\theta=0}^{2\pi} \frac{U}{2\pi\sigma^2} \exp\left[\frac{-U^2}{2\sigma^2}\right] d\theta = \frac{U}{\sigma^2} \exp\left[\frac{-U^2}{2\sigma^2}\right]$$

This is the standard Rayleigh distribution, as expected.

When the distribution is non-central, $\mu_x \neq 0$ and/or $\mu_y \neq 0$, and the integration is more difficult. Let us consider a wind speed interval, δU at wind speed, U, with wind speed direction, θ and sector angle interval, $\delta\theta$.

The ‘area’ of this interval is $U \times \delta U \times \delta\theta$.

And the probability of the wind speed occupying that area is:

$$p(\delta U, \delta\theta) = U\delta U\delta\theta \frac{1}{2\pi\sigma^2} \exp\left[\frac{-(x-\mu_x)^2 - (y-\mu_y)^2}{2\sigma^2}\right]$$

$$= U\delta U\delta\theta \frac{1}{2\pi\sigma^2} \exp\left[\frac{-(U\cos\theta-\mu_x)^2 - (U\sin\theta-\mu_y)^2}{2\sigma^2}\right]$$

In the limit, as $\delta\theta$ tends to zero, the total probability within the ring δU is:

$$p(\delta U) = \frac{U\delta U}{2\pi\sigma^2} \int_{\theta=0}^{2\pi} \exp\left[\frac{-(U\cos\theta-\mu_x)^2 - (U\sin\theta-\mu_y)^2}{2\sigma^2}\right] d\theta$$

And the wind speed probability density is:

$$f(U) = \frac{U}{2\pi\sigma^2} \int_{\theta=0}^{2\pi} \exp\left[\frac{-(U\cos\theta-\mu_x)^2 - (U\sin\theta-\mu_y)^2}{2\sigma^2}\right] d\theta$$

This probability density function has previously been evaluated and is identified as an ‘integral of the circular normal distribution over an offset circle’, (Abramowitz, Stegun 1964).

If R is the radius of the circle and $r = \sqrt{\mu_x^2 + \mu_y^2}$

Then the cumulative probability of the wind being less than R is $P\left(\frac{R^2}{2}, r^2\right)$

Where P is the cumulative form of the non-central chi-square distribution with two degrees of freedom, (Abramowitz, Stegun 1964). This function is also known as the generalised Rayleigh, Rayleigh-Rice or Rice distribution.

The above integral has been evaluated numerically for a range of offsets and results have also been compared with Weibull distributions with the same means and standard deviations in fig. C.1.

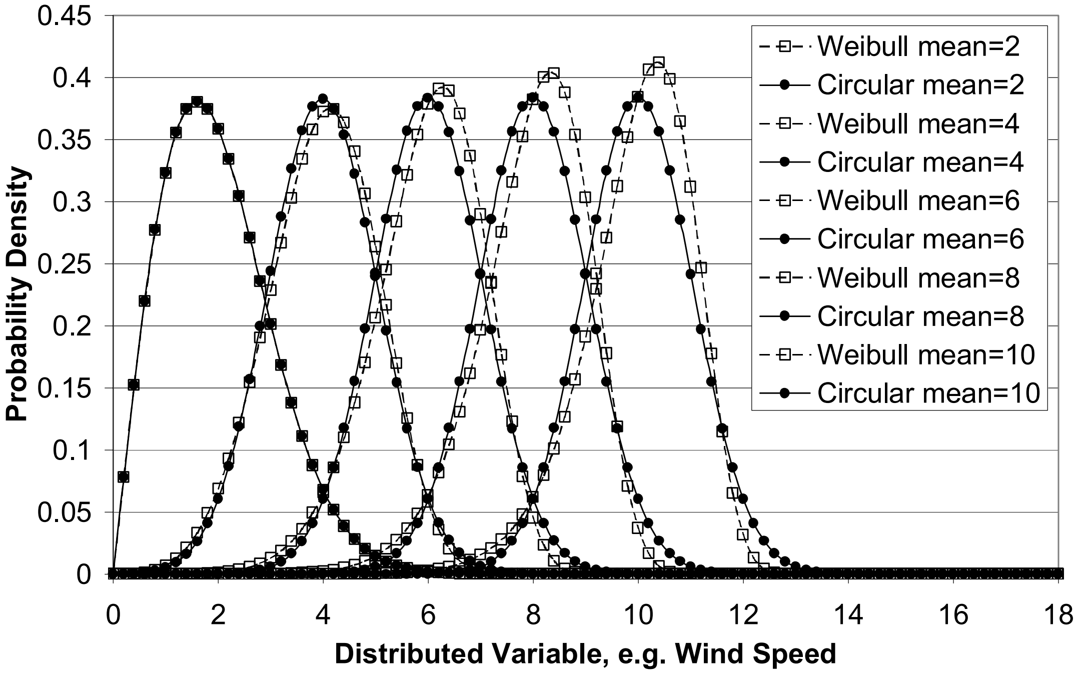

Figure C.1 Weibull distributions compared with offset-circle - circular normal distributions with the same means and standard deviations

In fig. C.1 all the distributions have standard deviations of 1.04, and various different mean values, in pairs. Each pair consists of a Weibull distribution and an offset circle - circular normal distribution with the same mean.

The first pair has a mean of 2, and the first Weibull distribution has a shape factor of exactly 2. Thus it can be seen that a Rayleigh distribution is exactly the same as a circular normal distribution with the same mean value.

Subsequent pairs have means of 4, 6, 8 and 10. In these distributions, the offset circle - circular normal PDFs all have positive skewness, decreasing in skewness as the mean increases. The Weibull distributions all have negative skewness, and their shapes get progressively more different from the corresponding offset circular normal distributions.

The above offset circle - circular normal distribution, or non-central chi-squared distribution, offers the possibility of a single, mathematically elegant family of wind speed distribution functions with a wide range of standard deviation to mean ratios, all with positive skewness. However, this is left for future work. The process of constructing a suitable distribution from a given mean and standard deviation requires iteration and would be too time consuming for the probabilistic method. The probabilistic method therefore uses the tried and tested Weibull and normal distributions.

5 References

Abramowitz, M. & Stegun, I.A. 1964, *Handbook of mathematical functions with formulas, graphs, and mathematical tables,* 1st edn, National Bureau of Standards, Washington, D.C.

Aguiar, R. & Collares-Pereira, M. 1992a, "Statistical properties of hourly global radiation", *Solar Energy,* vol. 48, no. 3, pp. 157-167.

Aguiar, R. & Collares-Pereira, M. 1992b, "TAG: a time-dependent, autoregressive, Gaussian model for generating synthetic hourly radiation", *Solar Energy,* vol. 49, no. 3, pp. 167-174.

Akuffo, F.O. & Brew-Hammond, A. 1993, "The frequency distribution of daily global irradiation at Kumasi", *Solar Energy,* vol. 50, no. 2, pp. 145-154.

Allen, P. & Todd, B. 1995, *Off the grid: managing electricity from renewable sources,* 1st edn, CAT Publications, National Centre for Alternative Technology.

Amato, U., Andretta, A., Bartoli, B., Coluzzi, B., Cuomo, V. & Serio, C. 1985, "Stochastic modelling of solar-radiation data", *Nuovo Cimento C,* vol. 8C, no. 3, pp. 248-258.

ASPO 2005, 4th August 2005-last update*, Association for the Study of Peak Oil and Gas* [Homepage of Association for the Study of Peak Oil and Gas], [Online]. Available: http://www.peakoil.net/ [2005, August 16] .

Babatunde, E.B. & Aro, T.O. 1995, "Relationship between "clearness index" and "cloudiness index" at a tropical station (Ilorin, Nigeria)", *Renewable Energy,* vol. 6, no. 7, pp. 801-805.

Barley, C.D. & Winn, C.B. 1996, "Optimal dispatch strategy in remote hybrid power systems", *Solar Energy,* vol. 58, no. 4-6, pp. 165-179.

Barton, J.P. & Infield, D.G. 2005a, "Energy storage and its use with wind power", IEEE Power Engineering Society, .

Barton, J.P. & Infield, D.G. 2005b, "A probabilistic method for calculating the usefulness of a store with finite energy capacity for smoothing electricity generation from wind and solar power", *Journal of Power Sources,* vol. To be published.

Barton, J.P. & Infield, D.G. 2004, "Energy storage and its use with intermittent renewable energy", *IEEE Transactions on Energy Conversion,* vol. 19, no. 2, pp. 441-448.

Baumgaertner, J. 1995, *Statistical representation of a hybrid photovoltaic - wind system for controller design*, MPhil edn, Loughborough University of Technology Department of Electronic and Electrical Engineering, Loughborough.

Bendat, J.S. & Piersol, A.G. 1993, *Engineering applications of correlation and spectral analysis. Second edition,* John Wiley, New York, NY, USA.

Biga, A.J. & Rosa, R. 1981, "Statistical behaviour of solar irradiation over consecutive days", *Solar Energy,* vol. 27, no. 2, pp. 149-157.

Bossanyi, E.A. & Anderson, M.B. 1984, "COMPARISON OF TECHNIQUES FOR EVALUATING THE FREQUENCY OF WIND TURBINE SHUT-DOWNS", *Wind Energy Conversion 1984, Proceedings of the Sixth BWEA Wind Energy Conference.* Cambridge Univ Press, Cambridge, Engl, Reading, Engl, pp. 228.

Bryden, I.G. & Macfarlane, D.M. 2000, "The utilisation of short term energy storage with tidal current generation systems", *Energy,* vol. 25, no. 9, pp. 893-907.

Bundesverband WindEnergie e.V. 2000, *Windenergie 2000 Wind Turbine Catalogue,* 1st edn, Bundesverband, Germany.

Burdon, I.P. 1998, "Options for mid-merit power generation in the UK electricity market", *Power Engineering Journal,* vol. 12, no. 3, pp. 115-122.

Castro, M.A., Carpio, J., Peire, J. & Rodriguez, J.A. 1996, "Renewable-energy integration assessment through a dedicated computer program", *Solar Energy,* vol. 57, no. 6, pp. 471-484.

Celik, A. & Marshall, R. 1998, *The system performance of autonomous photovoltaic and wind hybrid energy systems using synthetically generated weather data*, Professional Eng. Publishing, Cardiff, UK.

Celik, A.N. 2002a, "A simplified model based on clearness index for estimating yearly performance of hybrid PV energy systems", *Progress in Photovoltaics: Research and Applications,* vol. 10, no. 8, pp. 545-554.

Celik, A.N. 2002b, "The system performance of autonomous photovoltaic-wind hybrid energy systems using synthetically generated weather data", *Renewable Energy,* vol. 27, no. 1, pp. 107-121.

Chadee, J.C. & Sharma, C. 2001, "Wind speed distributions: a new catalogue of defined models", *Wind Engineering,* vol. 25, no. 6, pp. 319-337.

Chun Che Fung, Rattanongphisat, W. & Nayar, C. 2002, *A simulation study on the economic aspects of hybrid energy systems for remote islands in Thailand*, IEEE, Beijing, China.

Duffie, J.A. & Beckman, W.A. 1974, *Solar energy thermal processes,* Wiley-Interscience, Chichester, Sussex, UK.

Durisch, W. & Bulgheroni, W. 1999, *Climatological investigation for solar-power stations in the Swiss Alps*, Elsevier; Elsevier, UK; Bahrain.

Frean, P.B. 1983, *Optimal characteristics of components for small wind energy conversion systems*, MSc edn, University of Strathclyde Department of Operational Research, Glasgow.

Garcia, A., Torres, J.L., Prieto, E. & De Francisco, A. 1998, "Fitting wind speed distributions: a case study [for wind farms]", *Solar Energy,* vol. 62, no. 2, pp. 139-144.

Giraud, F. & Salameh, Z.M. 2001, "Steady-state performance of a grid-connected rooftop hybrid wind-photovoltaic power system with battery storage", *IEEE Transactions on Energy Conversion,* vol. 16, no. 1, pp. 1-7.

Gonzalez, J.-. & Calbo, J. 1999, "Influence of the global radiation variability on the hourly diffuse fraction correlations", *Solar Energy,* vol. 65, no. 2, pp. 119-131.

Gordon, J.M. & Reddy, T.A. 1989, "Stationary statistics and sequential properties of normal beam and global solar radiation on tilted surfaces", *Solar Energy,* vol. 42, no. 1, pp. 35-44.

Graham, V.A. & Hollands, K.G.T. 1990, "A method to generate synthetic hourly solar radiation globally", *Solar Energy,* vol. 44, no. 6, pp. 333-341.

Habibi, M.S. 2001, *Model for impact of storage on spinning reserve requirements and distributed generation*, IEEE, Athens, OH, USA.

Halliday, J.A. 1988, *Wind meteorology and the integration of electricity generated by wind turbines*, University of Strathclyde Department of Applied Physics and Energy Studies unit.

Hassan, U. & Sykes, D.M. 1990, *"Wind Structure and Statistics" in Wind Energy Conversion Systems chapter 2 pp. 11-32,* 1st edn, Prentice Hall, UK.

Hastings, N. A. J. & Peacock, J.B. 1975, *Statistical distributions. A handbook for students and practitioners,* 1st edn, John Wiley & Sons, New York-Toronto, Ont.

Hollands, K.G.T. & Huget, R.G. 1983, "A probability density function for the clearness index, with applications", *Solar Energy,* vol. 30, no. 3, pp. 195-209.

Hove, T. 2000, "A method for predicting long-term average performance of photovoltaic systems", *Renewable Energy,* vol. 21, no. 2, pp. 207-229.

Ibanez, M., Beckman, W.A. & Klein, S.A. 2002, "Frequency distributions for hourly and daily clearness indices", *Transactions of the ASME. Journal of Solar Energy Engineering,* vol. 124, no. 1, pp. 28-33.

Ibanez, M., Rosell, J.I. & Beckman, W.A. 2003, "A bi-variable probability density function for the daily clearness index", *Solar Energy,* vol. 75, no. 1, pp. 73-80.

Ideriah, F.J.K. 1992, "On the characteristic distribution of hourly diffuse and global solar radiation at Ibadan", *Renewable Energy,* vol. 2, no. 4-5, pp. 461-468.

Infield, D.G. 1990, *"Storage and Logistics" in Wind Energy Conversion Systems chapter 16 pp. 341-356,* 1st edn, Prentice Hall, UK.

Infield, D.G. 1984, *A study of electricity storage and central electricity generation*, Rutherford Appleton Laboratory Energy Research Unit, UK.

Infield, D., Scotney, A., Lundsager, P., Bindner, H., Pierik, J., Uhlen, K., Falchetta, M., Manninem, L. & Van Dijk, V. 1994, *Engineering Design Tools for Wind Diesel Systems - Final report for CEC contract JOUR-0078. Volume 8 - Logistical Package: Program documentation*, Energy Research Unit, Rutherford Appleton Laboratory, UK.

International Energy Agency 2002, November 2002-last update*, Renewables in Global Energy Supply - An IEA Fact Sheet* [Homepage of International Energy Agency], [Online]. Available: http://www.iea.org/dbtw-wpd/textbase/papers/2002/leaflet.pdf [2004, November 23] .

Iqbal, M.T. 2004, "A feasibility study of a zero energy home in Newfoundland", *Renewable Energy,* vol. 29, no. 2, pp. 277-289.

Iqbal, M.T. 2003, "Prefeasibility study of a wind-diesel system for St. Brendan's, Newfoundland", *Wind Engineering,* vol. 27, no. 1, pp. 39-51.

Jurado, M., Caridad, J.M. & Ruiz, V. 1995, "Statistical distribution of the clearness index with radiation data integrated over five minute intervals", *Solar Energy,* vol. 55, no. 6, pp. 469-473.

Khallat, M.A. & Rahman, S. 1986, "A probabilistic approach to photovoltaic generator performance prediction", *IEEE Transactions on Energy Conversion,* vol. EC-1, no. 3, pp. 34-40.

Kudish, A.I. & Ianetz, A. 1996, "Analysis of daily clearness index, global and beam radiation for Beer Sheva, Israel: partition according to day type and statistical analysis", *Energy Conversion and Management,* vol. 37, no. 4, pp. 405-416.

Lachs, W.R. & Sutanto, D. 1995, *Application of battery energy storage in power systems*, IEEE, Singapore.

Laird, J.L. & Harshvardhan 1997, "Analysis of cumulus solar irradiance reflectance (CSIR) events", *Atmospheric Research,* vol. 44, no. 3-4, pp. 317-332.

Lam, J.C. & Li, D.H.W. 1996, "Study of solar radiation data for Hong Kong", *Energy Conversion and Management,* vol. 37, no. 3, pp. 343-351.

Lambert, T. & Lilienthal, P. 2003, May 2003-last update*, HOMER - The Optimization Model for Distributed Power. Getting Started Guide Version 2.0 DRAFT* [Homepage of National Renewable Energy Laboratory], [Online]. Available: http://www.nrel.gov/homer/downloads/HOMERGettingStarted.pdf [2004, November 26] .

Lopez, G., Rubio, M.A. & Batles, F.J. 2000, "Estimation of hourly direct normal from measured global solar irradiance in Spain", *Renewable Energy,* vol. 21, no. 2, pp. 175-186.

Loutfi, H. & Khtira, A. 1992, "Stochastic analysis and generation of synthetic sequences of daily global solar irradiation: Rabat site (Morocco)", *Renewable Energy,* vol. 2, no. 2, pp. 129-138.

Luna, R.E. & Church, H.W. 1974, "Estimation of long-term concentrations using a `universal' wind speed distribution", *Journal of Applied Meteorology,* vol. 13, no. 8, pp. 910-916.

MacDonald, I.L. & Zucchini, W. 1997, *Hidden Markov and other models for discrete-valued time series,* 1st edn, Chapman & Hall, London.

Mandal, P., Senjyu, T., Uezato, K. & Funabashi, T. 2004, *Forecasting several-hours-ahead electricity demand using neural network*, IEEE, Hong Kong, China.

Manwell, J.F., Rogers, A., Hayman, G., Avelar, C.T. & McGowan, J.G. 1998, 2 November 1998-last update*, HYBRID2 - A Hybrid System Simulation Model. Theory Manual* [Homepage of Renewable Energy Research Laboratory, University of Massachusetts], [Online]. Available: http://www.ecs.umass.edu/mie/labs/rerl/hy2/theory/pdf/complete.pdf [2004, November 25] .

Masters, C.L., Mutale, J., Strbac, G., Curcic, S. & Jenkins, N. 2000, "Statistical evaluation of voltages in distribution systems with embedded wind generation", *IEE Proceedings-Generation, Transmission and Distribution,* vol. 147, no. 4, pp. 207-212.

McSharry, P.E., Bouwman, S. & Bloemhof, G. 2005, "Probabilistic forecasts of the magnitude and timing of peak electricity demand", *IEEE Transactions on Power Systems,* vol. 20, no. 2, pp. 1166-1172.

Mefti, A., Bouroubi, M.Y. & Adane, A. 2003, "Generation of hourly solar radiation for inclined surfaces using monthly mean sunshine duration in Algeria", *Energy Conversion and Management,* vol. 44, no. 19, pp. 3125-3141.

Milborrow, D. 2000, "Revolutionary Potential", *Wind Power Monthly,* vol. 16, no. 10.

Milborrow, D. 2002, "Assimilating wind [wind energy contribution to UK energy needs]", *IEE Review,* vol. 48, no. 1, pp. 9-13.

Mills, A. & Al-Hallaj, S. 2004, "Simulation, of hydrogen-based hybrid systems using Hybrid2", *International Journal of Hydrogen Energy,* vol. 29, no. 10, pp. 991-999.

Morgan, T.R., Marshall, R.H. & Brinkworth, B.J. 1997, "`ARES'-a refined simulation program for the sizing and optimisation of autonomous hybrid energy systems", *Solar Energy,* vol. 59, no. 4-6, pp. 205-215.

Muselli, M., Poggi, P., Notton, G. & Louche, A. 2001, "First order Markov chain model for generating synthetic "typical days" series of global irradiation in order to design photovoltaic stand alone systems", *Energy Conversion and Management,* vol. 42, no. 6, pp. 675-687.

National Environmental Research Council 2004, 30 November 2004-last update*, British Atmospheric Data Centre* [Homepage of NERC Centres for Atmospheric Science], [Online]. Available: http://badc.nerc.ac.uk/home/ [2004, December 3rd] .

National Grid Company plc 2003, March 2003-last update*, Seven Year StatementMarch 2003* [Homepage of The National Grid Company plc], [Online]. Available: http://www.nationalgrid.com/uk/library/documents/sys_03/ [2004, November 22nd] .

Norris, J.R. 1997, *Markov chains,* 1st edn, Cambridge University Press, Cambridge.

Olmo, F.J., Vida, J., Foyo, I., Castro-Diez, Y. & Alados-Arboledas, L. 1999, "Prediction of global irradiance on inclined surfaces from horizontal global irradiance", *Energy,* vol. 24, no. 8, pp. 689-704.

Palutikof, J.P., Bass, J.H., Halliday, J.A., Davies, T.D. & Barthelmie, R.J. 1989, *Estimation of the wind resource at proposed wind turbine sites: the problems of spatial and temporal variability*, Peter Peregrinus, Glasgow, UK.

Panickar, P.S., Islam, S.M. & Bleijs, J.A.M. 1999, *Potential of load management to improve the economics of wind-diesel hybrid systems*, Univ. Leicester, Leicester, UK.

Panickar, P.S., Islam, S.M. & Nayar, C.V. 1998, "A new quasi-optimal control algorithm for a wind-diesel hybrid system", *Wind Engineering,* vol. 22, no. 3, pp. 159-170.

Patel, J.K. 1976, *Handbook of statistical distributions,* 1st edn, Dekker, New York.

Poggi, P., Notton, G., Muselli, M. & Louche, A. 2000, "Stochastic study of hourly total solar radiation in Corsica using a Markov model", *International Journal of Climatology,* vol. 20, no. 14, pp. 1843-1860.

Rahman, S., Khallat, M.A. & Salameh, Z.M. 1988, "Characterization of insolation data for use in photovoltaic system analysis models", *Energy,* vol. 13, no. 1, pp. 63-72.

RCEP 2004, 27th January 2004-last update*, THE ROYAL COMMISSION ON ENVIRONMENTAL POLLUTION'S 22nd REPORT* [Homepage of Royal Commission on Envirnonmental Pollution], [Online]. Available: http://www.rcep.org.uk/pdf/summary.pdf [2004, November 23] .

Rifkin, J. 2002, *The Hydrogen Economy,* 1st edn, Penguin Putnam Inc., New York.

Sahin, A.D. & Sen, Z. 1998, "Statistical analysis of the Angstrom formula coefficients and application for Turkey", *Solar Energy,* vol. 62, no. 1, pp. 29-38.

Schlez, W. 2000, *Voltage Fluctuations Caused by Groups of Wind Turbines*, Loughborough University.

Shaw, R.H. & McCartney, H.A. 1985, "Gust penetration into plant canopies", *Atmospheric Environment,* vol. 19, no. 5, pp. 827-830.

Skartveit, A., Olseth, J.A. & Tuft, M.E. 1998, "An hourly diffuse fraction model with correction for variability and surface albedo", *Solar Energy,* vol. 63, no. 3, pp. 173-183.

Spera, D.A. 1994, *Wind turbine technology :fundamental concepts of wind turbine engineering,* 1st edn, ASME Press, New York.

Stevens, M.J.M. & Smulders, P.T. 1979, "The estimation of the parameters of the Weibull wind speed distribution for wind energy utilization purposes", *Wind Engineering,* vol. 3, no. 2, pp. 132-145.

Suehrcke, H. & McCormick, P.G. 1989, "The distribution of average instantaneous terrestrial solar radiation over the day", *Solar Energy,* vol. 42, no. 4, pp. 303-309.

Sulaiman, M.Y., Oo, W.M.H., Wahab, M.A. & Zakaria, A. 1999, "Application of beta distribution model to Malaysian sunshine data", *Renewable Energy,* vol. 18, no. 4, pp. 573-579.

Swift-Hook, D.T. & Ter-Gazarian, A.G. 1994, "The Value of storage on power systems with intermittent energy sources", *Renewable Energy,* vol. 5, no. 5-8, pp. 1479-1482.

Ter-Gazarian, A.G. & Kagan, N. 1992, "Design model for electrical distribution systems considering renewable, conventional and energy storage units", *IEE Proceedings C (Generation, Transmission and Distribution),* vol. 139, no. 6, pp. 499-504.

Tiris, M. & Tiris, C. 1998, "Optimum collector slope and model evaluation: case study for Gebze, Turkey", *Energy Conversion and Management,* vol. 39, no. 3-4, pp. 167-172.

Torre, M.C., Poggi, P. & Louche, A. 2001, "Markovian model for studying wind speed time series in Corsica", *International Journal of Renewable Energy Engineering,* vol. 3, no. 2, pp. 311-312-319.

Tovar, J., Olmo, F.J. & Alados-Arboledas, L. 1998, "One-minute global irradiance probability density distributions conditioned to the optical air mass", *Solar Energy,* vol. 62, no. 6, pp. 387-393.

Tovar, J., Olmo, F.J., Batlles, F.J. & Alados-Arboledas, L. 2001, "Dependence of one-minute global irradiance probability density distributions on hourly irradiation", *Energy,* vol. 26, no. 7, pp. 659-668.

Tovar, J., Olmo, F.J., Batlles, F.J. & Alados-Arboledas, L. 1999, "One minute k/sub b/ and k/sub d/ probability density distributions conditioned to the optical air mass", *Solar Energy,* vol. 65, no. 5, pp. 297-304.

Townsend, A.A. 1979, "Flow patterns of large eddies in a wake and in a boundary layer", *Journal of Fluid Mechanics,* vol. 95, pp. 515-537.

Trabea, A.A. 2000, "Analysis of solar radiation measurements at Al-Arish area, North Sinai, Egypt", *Renewable Energy,* vol. 20, no. 1, pp. 109-125.

Tuller, S.E. & Brett, A.C. 1985, "The goodness of fit of the Weibull and Rayleigh distributions to the distributions of observed wind speeds in a topographically diverse area", *Journal of Climatology,* vol. 5, no. 1, pp. 79-94.

Tuller, S.E. & Brett, A.C. 1984, "The characteristics of wind velocity that favor the fitting of a Weibull distribution in wind speed analysis", *Journal of Climate and Applied Meteorology,* vol. 23, no. 1, pp. 124-134.

Ulgen, K. & Hepbasli, A. 2002, "Prediction of solar radiation parameters through clearness index for Izmir, Turkey", *Energy Sources,* vol. 24, no. 8, pp. 773-785.

UNDP, UNDESA & WEC 2000, September 2000-last update*, World Energy Assessment - Key Messages* [Homepage of United Nations Development Program], [Online]. Available: http://www.undp.org/energy/activities/wea/brief-frame.html [2005, August 15] .

Unozawa, H., Otani, K. & Kurokawa, K. 2001, *A simplified estimating model for in-plane irradiation using minute horizontal irradiation*, Elsevier; Elsevier, Netherlands; Royton Sapporo, Japan.

Van der Auwera, L., de Meyer, F. & Malet, L.M. 1980, "The use of the Weibull three-parameter model for estimating mean wind power densities", *Journal of Applied Meteorology,* vol. 19, no. 7, pp. 819-825.

Van Der Hoven, I. 1957, "Power spectrum of horizontal wind speed in the frequency range from 0.0007 to 900 cycles per hour", *Journal of Meteorology,* vol. 14, no. 1, pp. 160-164.

Wan-Kai Pang, Forster, J.J. & Troutt, M.D. 2001, "Estimation of wind speed distribution using Markov chain Monte Carlo techniques", *Journal of Applied Meteorology,* vol. 40, no. 8, pp. 1476-1484.

Watson, S.J., Landberg, L. & Halliday, J.A. 1994, "Application of wind speed forecasting to the integration of wind energy into a large scale power system", *IEE Proceedings-Generation, Transmission and Distribution,* vol. 141, no. 4, pp. 357-362.

Watson, S.J. & Ter-Gazarian, A.G. 1996, *The optimisation of renewable energy sources in an electrical power system by use of simulation and deterministic planning models*, Elsevier; Elsevier, UK; London, UK.

Yakin, M.Z. 1984, *A stochastic programming model for determining optimal electricity generation*, Pergamon, San Francisco, CA, USA.

Youcef Ettoumi, F., Mefti, A., Adane, A. & Bouroubi, M.Y. 2002, "Statistical analysis of solar measurements in Algeria using beta distributions", *Renewable Energy,* vol. 26, no. 1, pp. 47-67.