

The timing and societal synchronisation of energy demand

Giulio Mattioli, Elizabeth Shove and Jacopo Torriti

This working paper summarises a presentation given on 10th December 2013 to the Department of Energy and Climate Change. All the authors are part of the DEMAND research centre. WWW.DEMAND.AC.UK

It is increasingly important to know about when energy is used in the home, at work and on the move. Issues of time and timing have not featured strongly in energy policy analysis and in modelling, much of which has focused on estimating and reducing total average annual demand per capita. If smarter ways of balancing supply and demand are to take hold, and if we are to make better use of decarbonised forms of supply, it is essential to understand and intervene in patterns of societal synchronisation. This calls for detailed knowledge of when, and on what occasions many people engage in the same activities at the same time, of how such patterns are changing, and of how might they be shaped. In addition, the impact of smart meters and controls partly depends on whether there is, in fact scope for shifting the timing of what people do, and for changing the rhythm of the day. Is the scheduling of daily life an arena that policy can influence, and if so how? The DEMAND Centre has been linking time use, energy consumption and travel diary data as a means of addressing these questions and in this working paper we present some of the issues and results arising from that exercise.

Introduction

Looking ahead, a decarbonised energy system, a system that depends on more renewable sources of power, and that is more reliant on decarbonised electricity than at present, calls for ‘smarter’ ways of balancing supply and demand. Many of the themes addressed in this document relate to this concern, and to the challenge of shifting ‘peaks’ in energy demand. In taking a social science approach to the time and timing of practices and hence consumption we contribute to an understanding of the social and temporal patterns into which current programmes – for instance - of smart metering fit.

More broadly, we are interested in how changing social practices generate new patterns of demand and in the impact of what people do on future demand (overall). As part of this we want to develop methods of capturing the relationship between mobility (where people go, and when) and energy demand, thereby linking what are normally separate domains of energy and transport studies.

With this ambition in mind, our research addresses a number of specific questions: first, what combinations of practices make up morning and evening peaks, and how do peak demands and practices vary across the week, and over the year? Second, when and how do patterns of societal synchronisation occur (i.e. moments when many people are doing the same thing at the same time) and how do peaks in practice relate to peaks in energy demand? Third, and this is especially relevant for the potential to shift the timing of demand, how do sequences of practice fit together to constitute ‘blocks’ that hang together for one reason or another. Fourth, how does the range of practices enacted in society itself change, - over decades or epochs – and how do these longer term trends play out with respect to mobility and energy demand? Finally, we ask where and how policy might influence any of these trends and patterns.

Our approach to these questions is shaped by the DEMAND centre’s theoretical focus and orientation. In brief, this is informed by three core propositions: that energy is used not for its own sake, but in the course of accomplishing social practices; that social practices and energy demand

are both shaped by infrastructures and institutions, and that these systems reproduce interpretations of need and entitlement, and of normal and acceptable ways of life. These propositions inform the DEMAND Centre's overall research programme, with specific themes focusing on related questions: for instance – how and why do end use practices vary? How and why do such practices change over time? How do infrastructures of supply and demand shape end use practices? What are the implications for concepts of normality, need and entitlement? And, finally, how is energy demand constituted, transformed and steered? The work presented here relates to Theme 1: Trends and Patterns in Energy Demand, one aim of which is to use existing data on time use, energy consumption and transport/travel to better understand the temporal order, sequencing and synchronisation of daily life.

This ambition is complicated by the fact that existing data collection in each of these areas has a different history and purpose. For example, in the field of time use studies, respondents are often asked to keep diaries of what they do, at ten minute intervals. In energy research, the home or the appliances are more commonly the focus of attention, with energy consumption being metered at time intervals of seconds. Within transport studies, there is a tradition of collecting travel diaries in which many details are established but only about journeys or travel events across a week or so. Each of these styles of data collection provides more or less insight into the patterning of infrequent events versus routines, and into issues of timing, duration and sequence. More problematic, from our point of view, is the fact that the 2005 UK time use survey was not very detailed, meaning that the most comprehensive data we have is from 2000. In what follows, we consider possible analytic strategies and methods of proceeding with our cross-cutting enquiries, despite these limitations.

What are peak demands made of?

In order to understand what people do at times of peak energy demand we started by looking into available UK time use data. Time use data record what people are doing at ten minute intervals during the day. Unfortunately in the UK there is no recent nationally representative time use study. The most recent large time use study in this country is the 2000 Office for National Statistics Time Use Survey (the next survey is scheduled to take place in 2014-2015). The more recent Household Electricity Survey has some diaries with information regarding the timing when appliances were on, but less on people's activities. Although the 'Trajectory' dataset is limited it includes time use data from 500 respondents equipped with GPS devices collecting 10 minute interval data on location over 3 days. In addition, questionnaire information revealed what people were doing, and with whom, at any given time of the day. Basic demographic information about respondent, including age, gender, individual income and household income are included in the dataset and can be weighted.

We used a sub-sample of this dataset (50 people) to examine the timing and duration of activities during morning (7.00 - 10.00 am) and afternoon (4.30 – 7.30 pm) residential peak electricity demand periods. So far the results are not so surprising. They show that morning and evening peaks are consist of different activities. For instance, on Monday morning, as sleeping phases out, preparations to get ready for the day (i.e. washing, dressing etc.) ensue, before people start working. In the evening, as work phases out, watching TV picks up along with food preparation. In our subsample on Monday there is more moving around by car in the evening than in the morning. Thanks to the GPS information, we can follow people around. For instance, we see that between 4.00 and 10.00 pm three respondents who moved around by car were engaged in the following sequences of activities:

Person 8 – church > driving > cleaning, tidying house
Person 11 – resting > driving > shopping > driving > work for job
Person 13 – driving > resting > preparing food & drink

We have return to the issue of sequencing later in the discussion.

More interestingly, Monday and Friday evening peaks are different. For instance, on Friday work ends earlier and there are some people washing and getting dressed, presumably preparing to go out. This is relevant in that understanding differences in exactly which activities take place at times of peak energy demand might be the starting point for a more subtle analysis aimed at understanding what (lack of) flexibility is inherent in what people do at specific times of the day. This could have implications for energy systems purposes (e.g. building algorithms for demand side controllers), but it is not the only type of variation we are interested in.

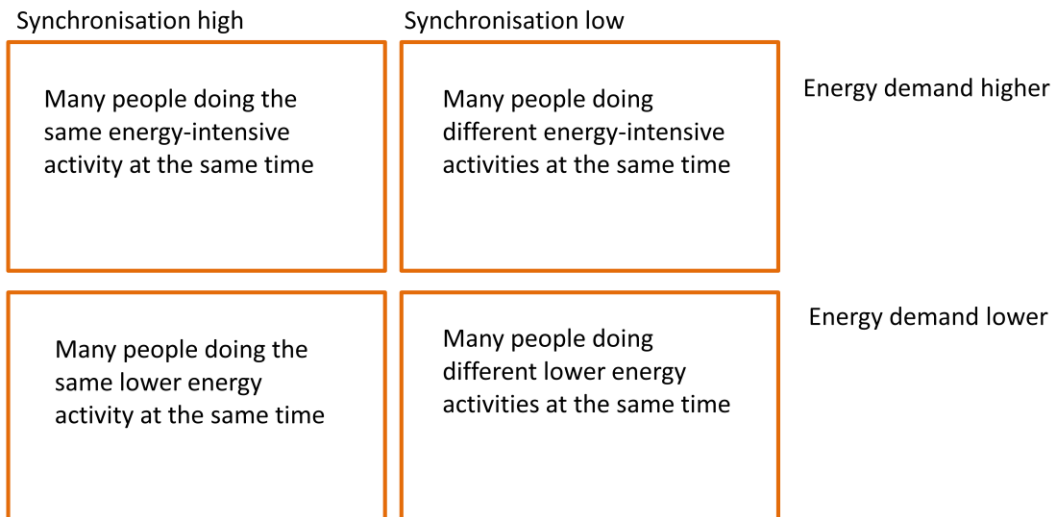
By following the timing and scheduling of certain practices (e.g. preparing food), we are able to say more about variation across the days of the week. For instance, data from the 2000 Office for National Statistics Time Use Survey shows that much more time is dedicated to food preparation on Sundays than on Saturdays and especially weekdays. Changing scale, we can also observe seasonal differences in practices like cycling (which could be for leisure or for commuting), which takes place less frequently in winter and autumn compared with spring and summer. In general, the 2000 Office for National Statistics Time Use Survey shows that many practices are conducted –or were conducted in the year 2000- repeatedly throughout the year regardless of weather, outdoor temperatures, etc.¹ Another analytic method, in theory made possible by large time use datasets (e.g. the ONS Time use survey) would be to track what happens at the same time of the day throughout the year. For instance, we could ask what happens at 11.00 AM on Saturdays all year around? This would provide a different method of revealing variations in people’s activities at the same time of the day, and in capturing some of the energy related implications of the temporal ordering of social practice.

In analysing time use data and relating this to peaks in demand we recognise that people do not live in isolation. This simple statement has significant implications for micro-level synchronicity (e.g. between occupants within the same household) and energy demand –depending on the extent to which there is a shared use of appliances, lighting for co-inhabited rooms, etc. Levels of multi-occupancy are typically addressed by energy modellers via stochastic approaches predicting the probability that any additional tenant/occupant might enter or leave the household for a specific time period. The Trajectory dataset allows us to identify who respondents were with at different times of the day. Perhaps not surprisingly, during weekdays, early mornings, evenings and nights are often spent with partner/spouse and children. The rest of the day is predominantly spent with work colleagues and/or by oneself. Weekends follow different patterns as the time spent with colleagues is much more limited and more time is spent with people in the family and/or friends. On its own, this is not especially revealing. But it becomes important when we think about the types of activities that constitute peak demand, and the extent to which these do or do not depend on the co-presence of others.

Peak demand and societal synchronisation

One of the goals of Theme 1 is to develop methods of representing societal synchronisation (that is the extent to which members of society are engaged in the same practices at the same time) and use these methods to analyse existing datasets in order to draw conclusions about the relationship between societal synchronisation and energy demand. From a theoretical point of view, there is no necessary or direct relationship between societal synchronisation and energy demand. Several combinations are possible, as illustrated in the matrix below.

¹ Though we do need to do more work on this, for instance focusing on ‘daylight needy’ practices and relating these to changing patterns of daylight through the year.

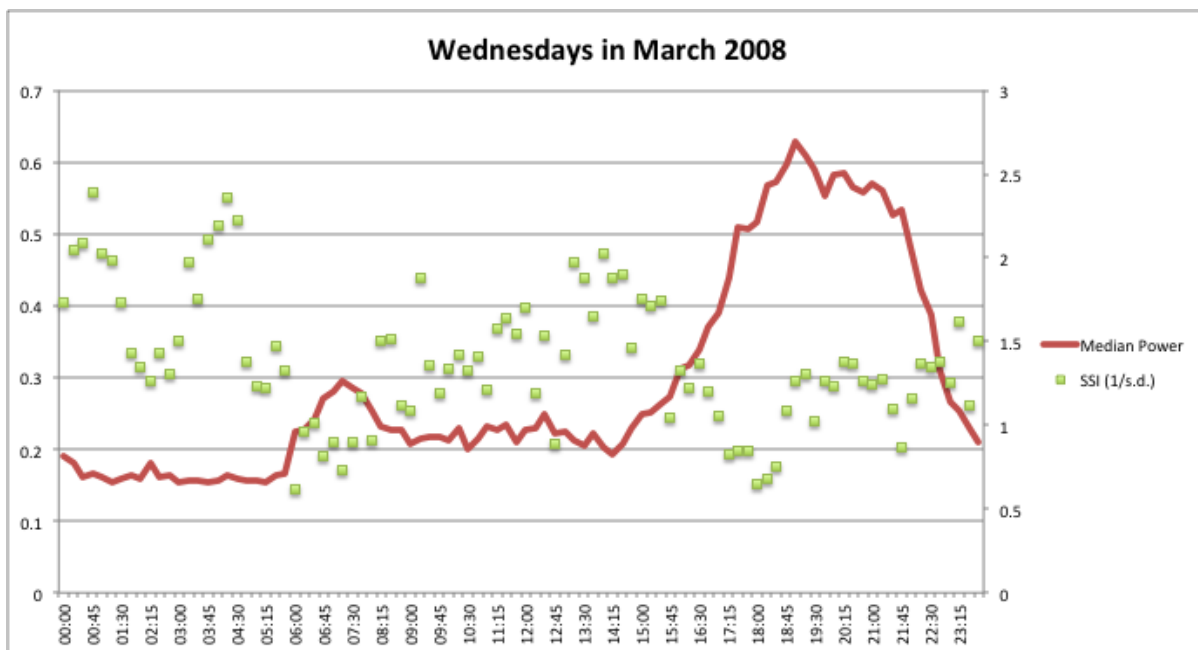


To elaborate, when many people are doing the same relatively energy-intensive activity at the same time, such as TV watching, or preparing an evening meal, we observe a peak in energy demand and high levels of synchronisation. Conversely, when many people are doing different energy-intensive activities at the same time, we observe a peak in energy consumption even though there is a low levels of societal synchronisation. An example of this might be the travel peak that is observed on Saturday around noon: unlike other travel peaks, this is the result of a great variety of journey purposes, and is thus associated with low levels of societal synchronisation.

Similarly, low levels of energy demand might be the result of many people doing different lower energy activities at the same time, or it might be a consequence of many people doing the same lower energy activity at the same time (an obvious example of this is sleeping).

If we are to take these ideas further we need a method of measuring the synchronisation of social practice. One simple solution is to suggest that for a given moment in time, synchronisation can be defined as the inverse of variation. Therefore, when using measured data (such as electricity imported from the grid) we define our “societal synchronisation index” as the inverse of the standard deviation. When using category data, such as the activity codes used in time use surveys, we define it as the difference between 100 and the standardised Shannon’s H – an established measure of entropy in time use studies.

With these ideas in place, we now consider examples of the four combinations of synchronisation and energy demand, using both measured and category data, and both household energy and travel data. The goal is to exemplify the kind of findings that can be generated by using this approach.



Original dataset: Richardson, I. and Thomson, M., *One-Minute Resolution Domestic Electricity Use Data, 2008-2009* [computer file]. Colchester, Essex: UK Data Archive [distributor], October 2010. SN: 6583, <http://dx.doi.org/10.5255/UKDA-SN-6583-1>.

Analysis by Ben Anderson.

This graph shows the results of a study based on one-minute resolution domestic electricity data for 22 British households on Wednesdays in March 2008. The red line corresponds to the median kw imported from the grid per 15 minutes, while the green dots show our social synchronisation index. The graph for Wednesdays in March shows that in general peaks in demand match low values of the societal synchronisation index. It is noticeable that whilst overall demand is lower in the mornings than the evenings, there is low synchronisation (i.e. a lot of variation) in the mornings (e.g. between 6 and 9 am). We also see high levels of synchronisation and low levels of electricity use during the night when most people are sleeping.

Similar analysis of Sundays in March 2008 shows that when compared with Wednesdays, Sunday stands out as not having a morning peak in demand and also having much more variation during the day.

Another exercise in characterising synchronisation involved data drawn from the British Time Use Survey for the year 2000, focusing on transport on a winter Tuesday. Data on the percentage of respondents engaged in travel activities throughout the day reveals two peaks, one in the morning and one in the evening. However, our societal synchronisation index suggests that there is a lot more “societal synchronisation” in terms of “why people travel” in the morning peak (more than 60 out of 100) as compared to the evening peak. In other words, people travel for a wider variety of reasons in the evening peak. What we might think of as ‘varied’ as opposed to ‘synchronised’ peaks are even more important as regards travel during the week end. If we take a winter Saturday, we can see a single, large peak around noon, and we can also see that this corresponds to very low levels of synchronisation.

When talking about synchronisation, we are looking ‘vertically’ at the data: for a given moment in time, we are basically asking about the variety of practices in which people are engaged: is it very

varied, or is it not? Another approach is to look “horizontally” at the data, focusing on how practices follow one another in the course of a day.

Sequences of practice and temporal flexibility

Not surprisingly, practices are linked over time, and as described below, some are very *tightly locked* together. Tightly coupled sequences make ‘blocks’ that structure the rhythm of the day. These arrangements have distinctive features both of timing and duration. Some practices have to come before or after others. Some practices often come before or after others. Some practices rarely come before or after others. Others are more *flexible*, and can be done at different times. Practices of short duration can be slotted in between longer ‘blocks’. These patterns are not fixed and as we know, ‘blocking’ and the coupling and de-coupling of practices changes historically. This is illustrated by the decreasing significance of seasonality and what seems to be an increasing potential for multi-tasking. If we are to understand the timing and scheduling of practices (and the implications for energy demand), we need to develop techniques for identifying and analysing sequential patterns.

The method we have used up to this point (identifying which specific practices constitute peaks in demand) is clearly limited if the goal is to represent series and sequences of practices, and to understand when and how these are reproduced. It is possible to use the 2000 Time Use Survey data to show the frequency of occurrence of three practices that we assumed to be linked – food preparation, eating and washing up – and to do so for different days of the week: for example, for Wednesdays and Sundays. This exercise suggests that preparing and eating food and washing up are related in that the peaks for the three activities tend to be synchronised. It also shows that they are slightly lagged, suggesting that – as expected, food preparation tends to precede eating, which is in turn followed by dish washing. However, graphs of this kind are not very informative: they tell us little about the predictability of such ‘blocks’, about how rigidly such blocks are defined, or about how they are in turn related to other tightly or loosely coupled sequences. To address these issues we need to explore other methods of data analysis.

The Visual-TimePACTS software has been developed at Linköping University (in Sweden) in order to visualize and compare sequences of activity. It provides a method for representing sequences in time use data that allows us to say more about how different activities/practices connect, one to another. One visual representation consists of vertical bars, showing the time at which different activities are undertaken, each activity being coded with a different colour. Data for around 100 individuals clearly shows the synchronization of the working day: and it shows that a few individuals follow other rhythms (such as working night shifts, etc.). What makes this software interesting for us is the inbuilt tool for sequential pattern mining. Based on the apriori algorithm, it allows the researcher to ‘fish for patterns’, based on several user-specified criteria. For example, we can look for occurrences of the sequence “cook dinner, eat dinner, wash dishes”. Interestingly, the tool makes it possible to detect this pattern even when other activities (such as “having a smoke”) interrupt the preparing-eating-washing sequence – these intervening activities are considered to be ‘gaps’. Such blocks and sequences can be displayed in different ways. They can be represented in the context of one respondent’s day and also mapped for all respondents, potentially allowing us to detect differences and similarities in the synchronisation and sequencing of activities between social groups. For example, we can see that the sequence “breakfast – read newspaper – drive car – work” is more frequent among men than among women. The pattern can also be taken as a unit of analysis, such that the whole sequence, rather than a single activity becomes the object of study.

Another tool used to find patterns in time use data is called “Activitree”. This is based on ‘page-rank’ algorithms similar to those used by web-search services like Google. The idea is to measure the degree of ‘connectivity’ between all different activities recorded in the diary, and then to assess

which are more likely to be associated with an activity of choice. Interestingly for those who study transport, the single most well-connected activity is “travel by car”. The “Activitree” tool makes it possible to iteratively add activities to the candidate sequence. At each step, the connectivity scores are computed again, allowing the researcher to explore the data further. One can, for example, examine scores for the activity sequence “travel by car – work”. One possible use of this tool is to identify sequence patterns that are ‘car dependent’, i.e. that they are much more likely to be preceded or followed by ‘travel by car’ than by ‘travel by other modes’. Overall, this software provides a range of tools for identifying sequences of activities that are interesting from an energy use perspective.

This far we have considered methods of analysis that allow us to look at how peaks of demand are constituted – and at ways of identifying the practices of which such peaks are made on a daily, weekly and seasonal basis. We have also discussed methods of distinguishing between different types of societal synchronisation, and of representing these different forms and their impact on energy demand. Our third step was to consider relations between practices, and to explore techniques for identifying loosely and tightly coupled ‘bundles’ or complexes of practice. All this is to take the range of practices enacted in society today pretty much for granted. In reality, what people do and when and how they do it is anything but static. In bringing this discussion to a close we draw attention to the historical dynamics of practice and to the implications of such longer term change for energy and mobility demand.

How practices change over time

Two simple examples serve to illustrate the types of transformation that we need to consider. The first relates to the practice of eating meals at home. In the UK, this practice has changed significantly over the last 50 years. In the 1960s many more meals were eaten at home. These meals were mostly consumed at 8am, 1pm (many still had lunch at home) and 6pm, with a later, smaller spike of ‘supper’ at 10 pm. In 2001 breakfast and lunch are much less commonly eaten at home. Whilst some 12% do have dinner at home in 2001, they do so a bit later – more like 7pm. Even so, this figure is much less than the 23% or so that pertained in 1961.

The second example has to do with women’s employment. Graphs showing when men and women are at work (over the same period – 1961-2001, also from Gershuny 2011) provide a tangible reminder of the rapid increase in women’s participation in paid employment. Sweeping socio-economic changes of this nature are important for where energy demand occurs, and also for the rhythm of the day. Amongst much else, such trends are important for eating habits, for where lunch is prepared and consumed, and for how breakfast and dinner are, on average, organised.

Other developments, for instance in technologies (TVs, internet, central heating systems, timers, cooking appliances etc.) also have a bearing on how people spend their and on how different practices are, and are not connected to energy demand. Technologies play into the timing of demand in different ways. To give another simple example, a controller on a central heating system makes it possible to heat the home even when everyone is out (e.g. when all are at work). Other devices such as batteries and portable power systems allow people to enact energy-demanding practices in new locations (working on the train, etc.). Both modify the relation between space, time and energy demand.

Perhaps more important, technologies are not ‘innocent’ in the sense that they are themselves integral to the conduct of specific practices. New devices (and opportunities to use energy) are part and parcel of the ongoing dynamics of daily life. There is much more to say about the recursive relation between technologies/infrastructures and social practice, but for now it is enough to make

the point that such relations are crucial for longer term trends in overall energy demand, and in the detail of when and where that demand occurs.

Opportunities for policy intervention

Much current discussion about smart meters – and their role in balancing supply and demand in a more renewable/decarbonised system – supposes that energy demand can be shifted somewhat at will. Our work provides a means of assessing the social/sociological plausibility of that assumption. It also draws attention to some of the many ways in which non-energy policy (e.g. on women's employment) impacts on daily and weekly rhythms of mobility and energy demand. Whether aware of it or not, regulations about opening hours (Sunday trading, GP services, pubs, school holidays); working time (shift work); daylight savings time, and more widely many forms of urban planning have a collective impact on shared socio-temporal rhythms. Other institutions also play a key role – for instance, many businesses are variously involved in setting on- and off-peak charges (for energy, for parking, for train travel etc.); promoting 'early bird' offers, extended opening hours and the like.

More work is required to establish the fixity and flexibility of contemporary temporal rhythms, and of how blocks of practices are sequenced and scheduled. But it is already obvious that those involved in energy policy would do well to take note of the many ways in which governments have a hand in orchestrating the timing and synchronisation of daily life and hence energy demand – and thus in configuring (enabling and also limiting) the potential for shifting peak loads of one kind or another. In addition, and as hinted at here but developed more fully in the rest of the DEMAND programme, a longer term challenge is to recognise how social practices change over time, and how these dynamic processes matter for energy demand.