

# Energy Consumption in The Office

Anton Stefanek

Uli Harder

Jeremy T. Bradley

Department of Computing, Imperial College London  
Huxley Building, 180 Queen's Gate, London, SW7 2RH  
{as1005,uh,jb}@doc.ic.ac.uk

**Abstract.** In this paper we present measurements of energy usage of standard office computing equipment. Using a data trace lasting for all of March 2012 we analyse the energy use of office equipment such as desktop computers, a printer and a fridge. The interest in a more detailed knowledge of the energy usage patterns of these appliances is driven by the desire to manage, and if possible reduce, the energy consumption of computing equipment in a university department. The reason behind this can be financial to reduce electricity costs and/or environmental to reduce the carbon foot print of an office environment. We analyse the data and show simple autoregressive time series models to predict the energy usage of appliances. We also show that it's feasible to accurately approximate the power consumption of a desktop computer using the CPU utilisation information. We describe a future set-up where we plan to monitor the energy usage of a student lab.

**Keywords:** Energy consumption, Measurement, Forecasting

## 1 Introduction

In this paper we present data collected with several energy monitors in an office environment. The data comprises a number of different desktop computers, a fridge, a printer, a mobile phone / laptop charger and a digital radio. In addition we also monitor the office temperature. The data we present here can be downloaded from [3].

The motivation behind this measurement exercise is to understand what the environmental impact of computing is. In order to make computing greener we first need to know how much energy is used by computers. This measurement exercise is a warm-up for a larger experiment, where we aim to monitor the energy usage of an entire computer lab used by undergraduate students. Given that conservative estimates of peak usage indicate that departments and companies pay thousands of GBP for the electricity usage of their computing equipment, being able to quantify how money could be saved is very useful in the current economic climate. In the long-term we would also like to be able to investigate how different scheduling policies for CPU scavenging applications like Condor [7] could be improved to use less energy.

In the following sections we first describe the experimental setup. We perform a survey of the data and show a few insightful visualisations. In the next section we describe simple statistical models forecasting future demand. Finally we discuss the implications of our observations and sketch a future larger measurement exercise

## 2 Experimental setup

For this experiment we used ten *SmartPlug* meters from *AlertMe*[1]. Each Smart-Plug was placed between the device and the mains network and measured the actual power consumption with an accuracy of around 1W. The plugs reported power measurements via the Zigbee wireless protocol to a *SmartHub* controller, approximately every 10 seconds. The hub is a low power ARM based computer running a version of the Linux operating system. It is connected to the Internet and reports the measurements from the plugs to the AlertMe website. In addition, we used a *SmartDisplay* device to report the current temperature of the office to the hub.

The data from AlertMe can either be viewed on a dashboard with limited capabilities or obtained via an API. However, the API does not allow detailed history requests. Therefore, we ran a simple script that retrieved the current measurements from the AlertMe API at regular intervals of 12 seconds (a rate just below the limits imposed by AlertMe) and uploaded all the data to *Cosm* [2] website, which provides an online storage of time series data.

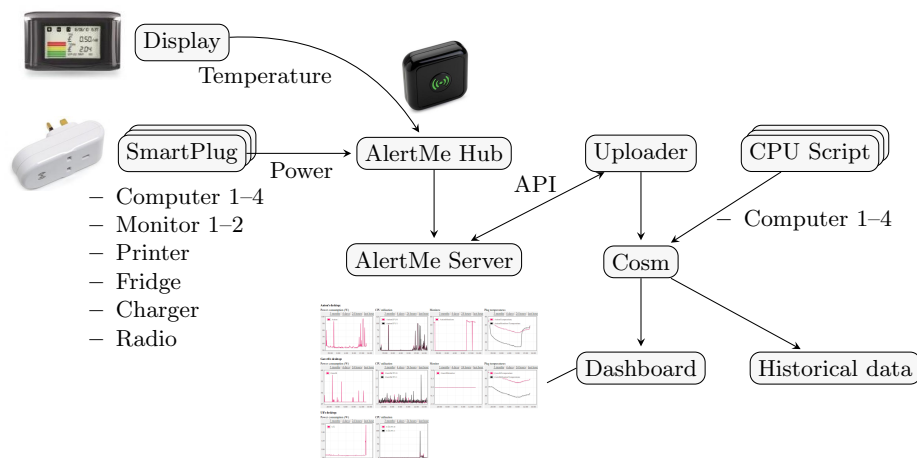
On each monitored desktop computer we ran a simple script that measured the CPU utilisation and deposited the data onto Cosm every 12 seconds.

For convenience, we wrote a simple HTML dashboard that displays an overview of all the current and historical measurements via calls to Cosm. We retrieved detailed historical data from Cosm via the Cosm API and used various data visualisation and statistical analysis tools to gain insights from the experiment. Figure 1 shows an overview of the whole setup.

### 2.1 Monitored devices

We monitored the following appliances:

- Four desktop computers. Computer 1 is an HP model with 3GHz Intel Core 2 Duo E8400 processor. Computer 2 is a Dell model with 2.8GHz Intel Core 2 Duo E7400 processor. Computer 3 has a quad core 3.4 GHz Intel Core i5-670 processor. Computer 4 has a single core 3.6 GHz Intel Pentium 4 processor. All the computers are running a custom version of the Ubuntu Linux operating system and are used for common office tasks such as web browsing, programming and scientific computing.



**Fig. 1.** Overview of the experimental setup. Images of the SmartPlug, Display and Hub taken from [www.alertme.com](http://www.alertme.com)

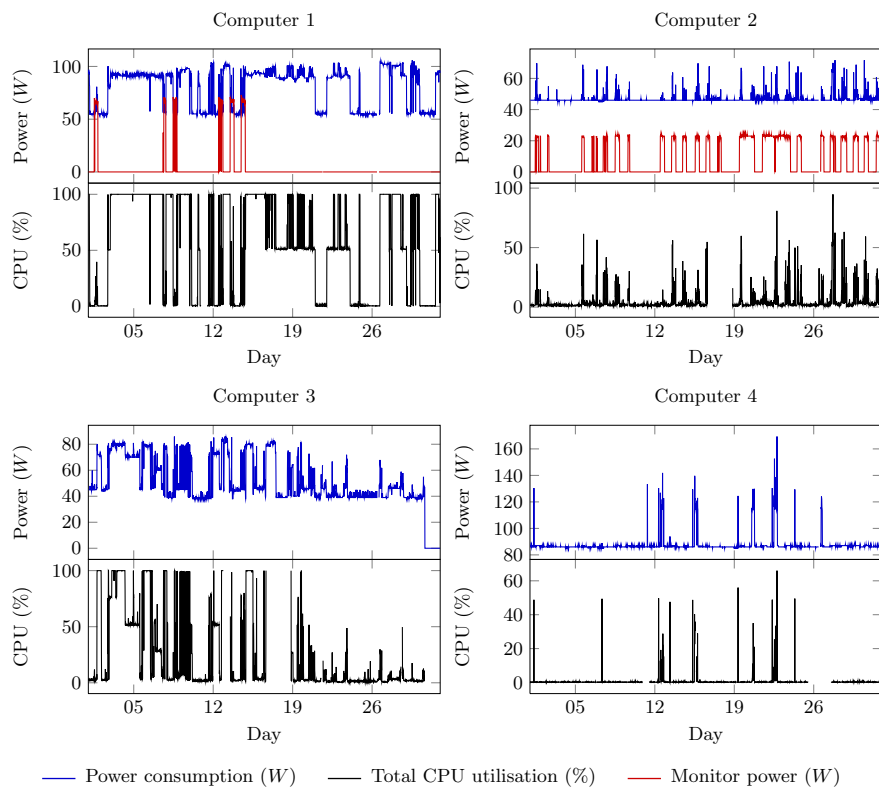
- Monitors. We monitored two screens (21 inch NEC LCD monitor and 20 inch Dell LCD monitor) attached to Computer 1 and one 22 inch Sun monitor attached to Computer 2.
- A large shared HP printer. The printer is part of a college-wide printing system. Users can submit jobs to a central queue and then start printing after physically logging into the printer.
- A fridge.
- A mains socket used for occasional charging of a mobile phone and a laptop.
- A digital radio.

### 3 Survey of the data

We ran the experiment described above throughout March 2012. Overall, there were only a few hours when we lost data due to a crashed uploader script and outages of the AlertMe and Cosm servers. In total, we collected around 200k power measurement data points for each device and 230k CPU measurement data points for each CPU core of each computer.

Due to various delays in the system, the resulting time series data points do not come from regular time intervals. We performed a simple re-sampling of the data, producing a regular time series where each data point corresponds to a time-average of all the original points in its interval. Figure 2 shows the measured data for the four computers, Figure 3 for the other devices, when re-sampled into 20 minute intervals.

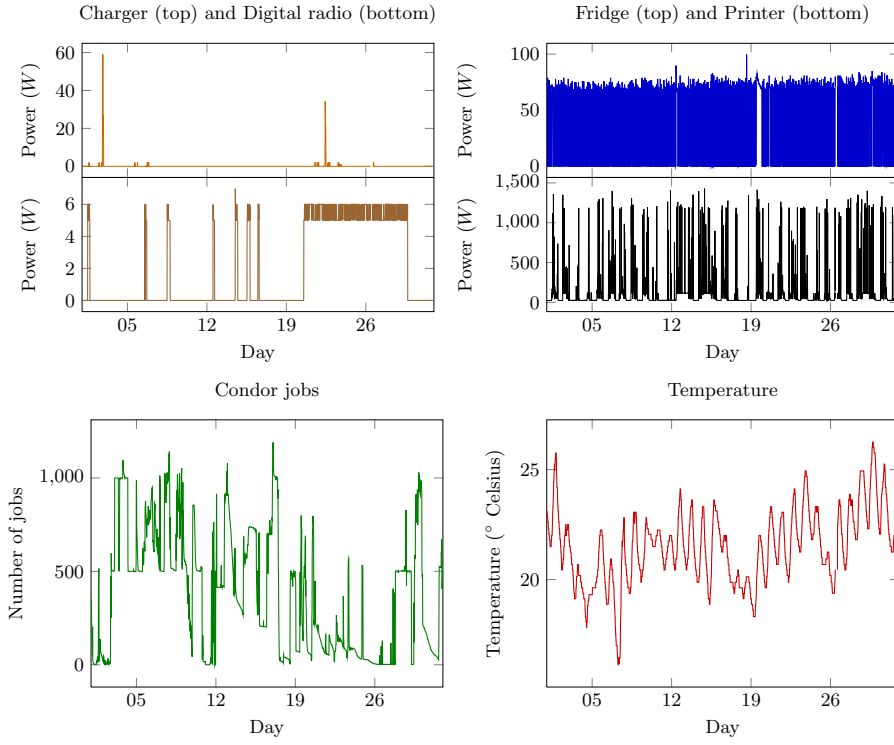
Figure 3 also shows the corresponding room temperature and number of active jobs submitted to the Condor system obtained via the `condor_stats` command.



**Fig. 2.** Energy use and CPU utilisation of the monitored computers. The shown dates correspond to Mondays.

Condor monitors the activity of a large number of computers. If the computers are idle for a period of time, Condor schedules any pending jobs for processing on the available computer. In our experiments, Computers 1 and 3 are assigned to a Condor pool. This is manifested in some of the high CPU utilisation during periods of time when the monitors are switched off.

The plots show quite clearly that the usage patterns are very different for all four computers. It is clear that the power consumption is highly correlated with the total CPU utilisation. We will look at this relationship in Section 3.3. Moreover, similar load patterns for Computer 1 and Computer 3 suggest that there is a correlation between the number of active Condor jobs and the CPU utilisation of the computers. We will look at this relationship in Section 3.4. Table 1 summarises the total energy consumption of all the devices.



**Fig. 3.** Energy used by the non-computing devices and the number of active jobs in the Condor pool and the temperature in the office.

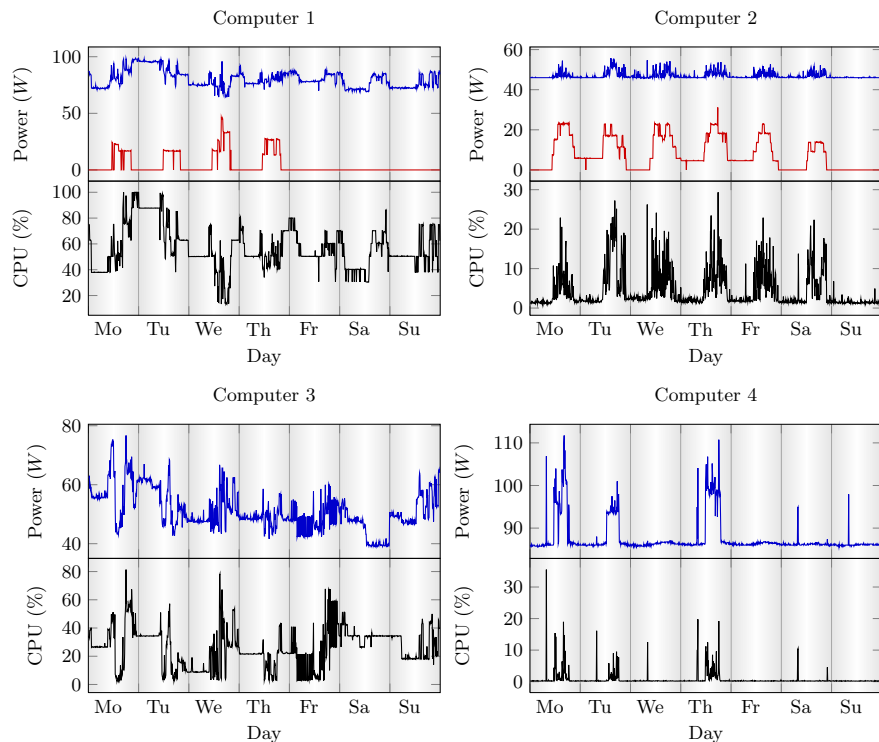
**Table 1.** Total energy consumption of all the devices in March 2012.

Device	Total (kWh)
Computer 1	59.2
Computer 2	34.6
Computer 3	37.4
Computer 4	64.5
Monitor 1	2.9
Monitor 2	5.5

	Total (kWh)
Charger	0.16
Radio	1.3
Fridge	19.9
Printer	108.4

### 3.1 Day of week averages

We average the above data for each week day. Figure 4 shows the averaged data for the 4 computers. We can clearly distinguish the weekend and standard working hours. Figure 5 shows the averaged data for the non-computer devices and for the number of condor jobs and the temperature.

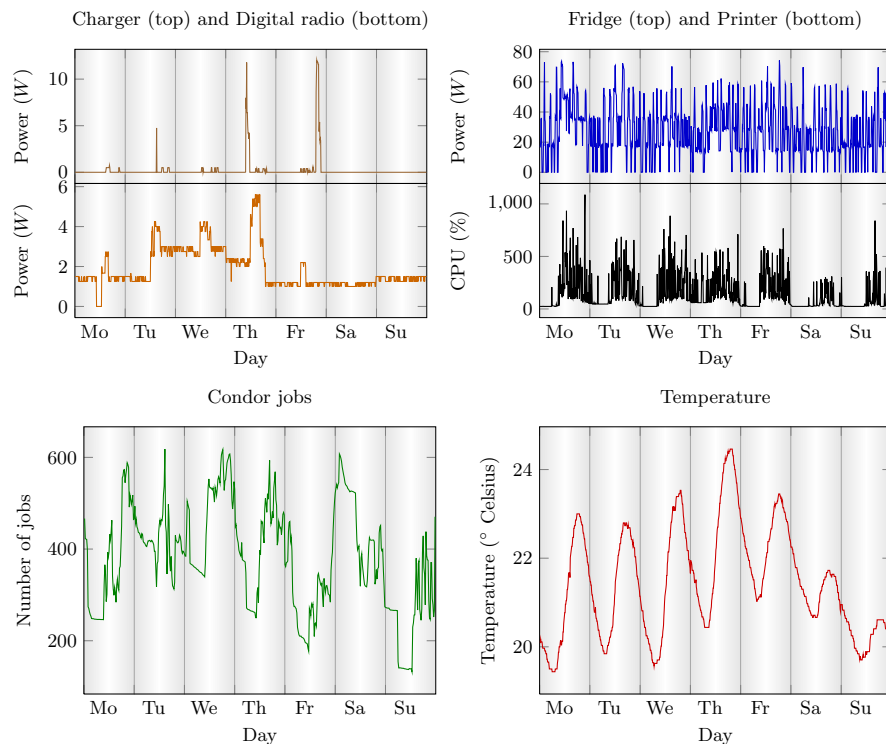


**Fig. 4.** Week day average of the CPU utilisation and power consumption of the 4 computers.

### 3.2 Duration histograms

We look at the proportion of time the devices stay at particular power levels. For each time series, we divide the range of its values into 25 intervals. We keep track of the duration of how long the value stays within a single interval. At each time of a change to a different interval, we record the total duration and the power level interval. We then plot a histogram of the total time the device stayed within an interval. Figure 6 shows these histograms for the 4 computers.

For example, the upper left figure shows the histogram for the power consumption of Computer 1. It can be seen that there are roughly 3 distinct power levels at which the computer stays most of the time. The lowest power level of around 56 W, corresponding to the machine being idle, is mostly maintained for periods of time around 1000 seconds long. The plot on the right hand side shows the CPU utilisation histogram for the same machine. Here, the 3 different levels are even more clear and show that most of the long running tasks are at the highest utilisation. The tall peak showing this is not present in the power histogram, because the power values had show larger fluctuations and do not stay within



**Fig. 5.** Week day average of the power consumption of the non-computer devices and the week day average of the number of active Condor jobs and the temperature.

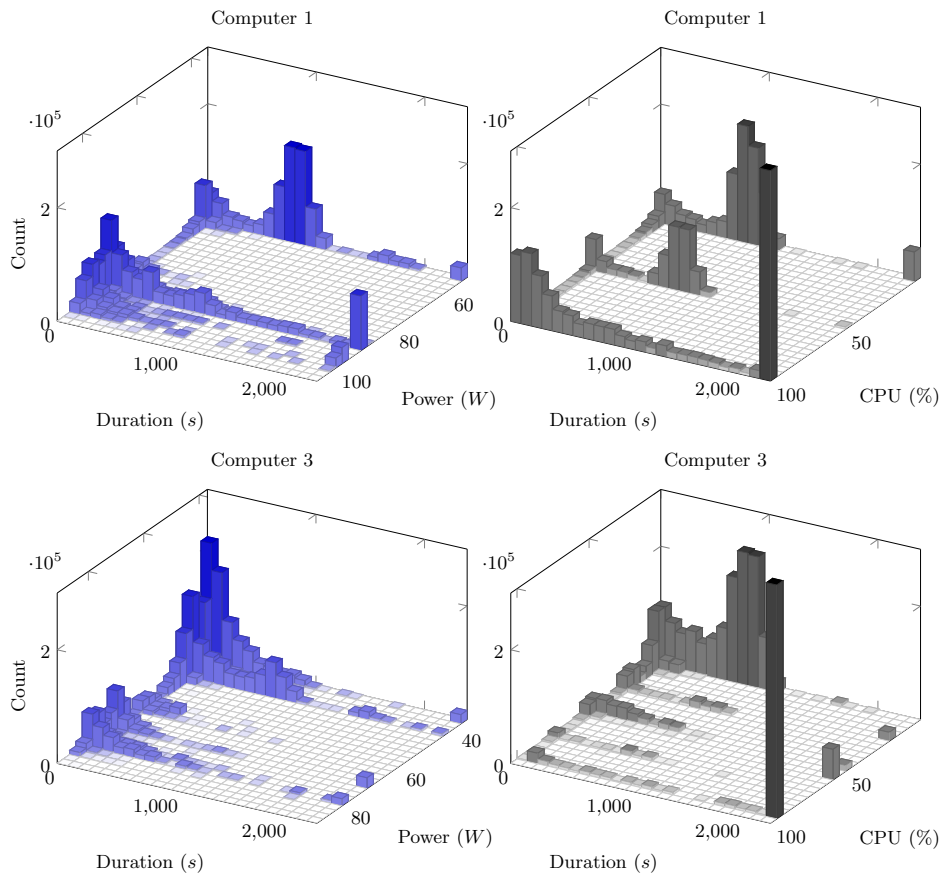
the same tight interval for a long time. This is probably a manifestation of the fact that the power consumption also depends on various other factors, such as the usage of the disk drive, graphics card or the room temperature.

Figure 8 in the following section shows the histograms for the Computers 1 and 3 where the values are summed over the range of possible durations.

Figure 7 shows the histogram for the Fridge and the Printer. The histogram for the Fridge clearly shows the two different states of the fridge. In case of the Printer, it can be seen that it is mostly idle, except for very short spikes of very high power consumption corresponding to the print jobs.

### 3.3 Power vs. CPU relationship

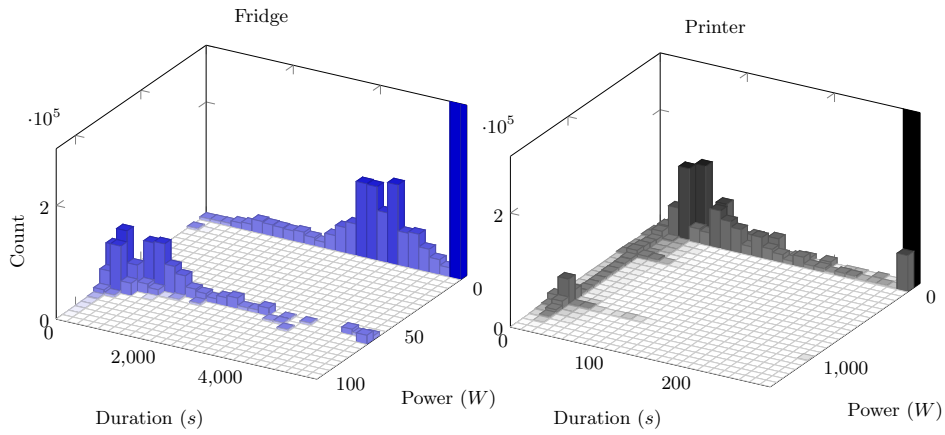
In this section, we look at the relationship between the power consumption and CPU utilisation for the monitored computers. We will try to come up with a simple model that will estimate the power from just the CPU utilisation information.




**Fig. 6.** Plot of the total time the computers stay at different power/CPU utilisation levels. For each power/utilisation interval, the plots show all possible durations in seconds for which the device stayed within the interval. For each level/duration pair, the total time in such state is plot. The time spent in states where the duration is larger than shown is included in the total value for the largest duration.

Figure 8 shows the distribution of the power consumption and CPU utilisation for Computers 1 and 3, which are part of a Condor pool and therefore showed a larger variety of utilisation and power levels during the experiment. There are clearly several distinct power and CPU levels at which the computers stay for most of the time. For example, the utilisation of the Computer 1 is around 0%, 50% or 100% most of the time and similarly 0%, 25%, 50%, 7% and 100% for Computer 3. This corresponds to the number of CPU cores of the computers. Usually, each Condor job fully utilises one core so this is expected. Figure 9 further shows the histogram of the relative changes of the power consumption





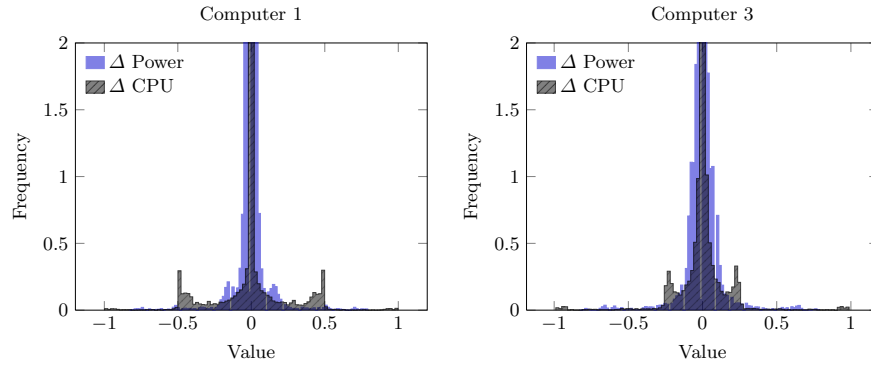
**Fig. 7.** The total time the fridge and the printer spent at each power level, divided according to the length of time they consecutively spend at the level. The bars in the upper right corners aggregate the count of idle periods longer than the shown maximum durations.



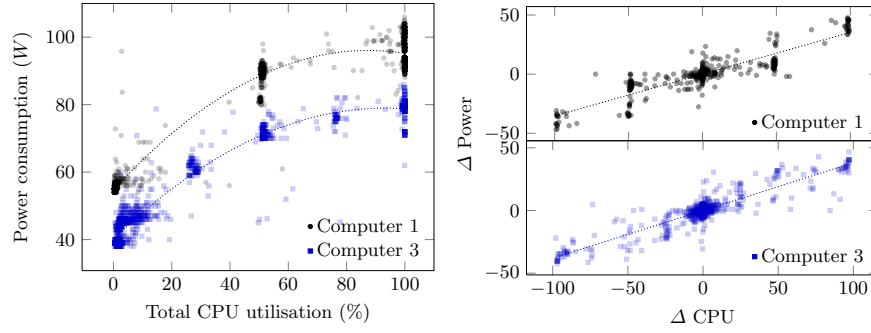
**Fig. 8.** Histogram of the power consumption and CPU utilisation levels for Computer 1 and 3.

and CPU utilisation for Computers 1 and 3. As expected, the values for the computer with more cores change in smaller relative steps.

Figure 10 plots the power consumption against CPU utilisation. The left hand side plot shows that the power consumption can be approximated by a quadratic function of the CPU utilisation. This confirms expectations from literature. It is suggested that adding a higher order term to a linear function can often accurately capture the relationship between the power consumption and CPU speed [6].



**Fig. 9.** Histogram of the normalised changes in power consumption and CPU utilisation.



**Fig. 10.** Plots of power consumption (change in) against CPU utilisation (change in, respectively). The left hand side plot suggests that the power consumption can be characterised as a quadratic function in the CPU utilisation.


The solid lines in Figure 10 are the fitted quadratic polynomials for Computer 1 and 3. Their parameters are

$$\text{Power}_1(x) = 54.7W + x \cdot 0.94W/\%CPU - x^2 \cdot 5.3 \cdot 10^{-3}W/\%CPU^2$$

$$\text{Power}_3(x) = 40.0W + x \cdot 0.83W/\%CPU - x^2 \cdot 4.4 \cdot 10^{-3}W/\%CPU^2$$

respectively, where  $x$  is the CPU utilisation.

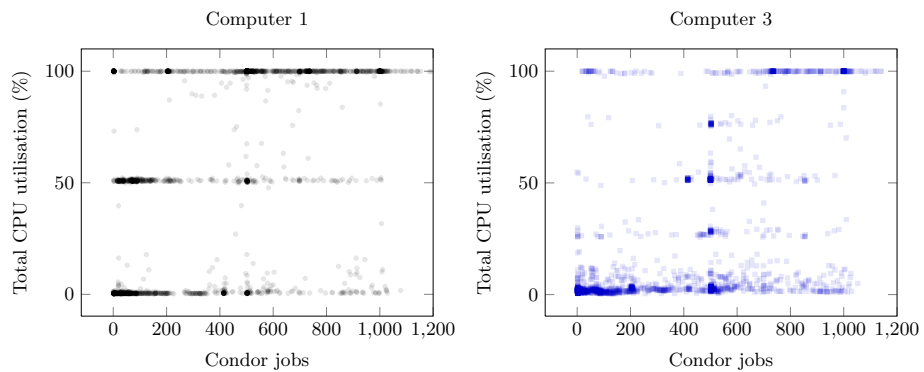
Figure 11 shows the power consumption of Computers 1 and 3 obtained from CPU utilisation and the above best fit functions. The sum of the absolute error at each time point for both computers is around 10% of the total power consumption.



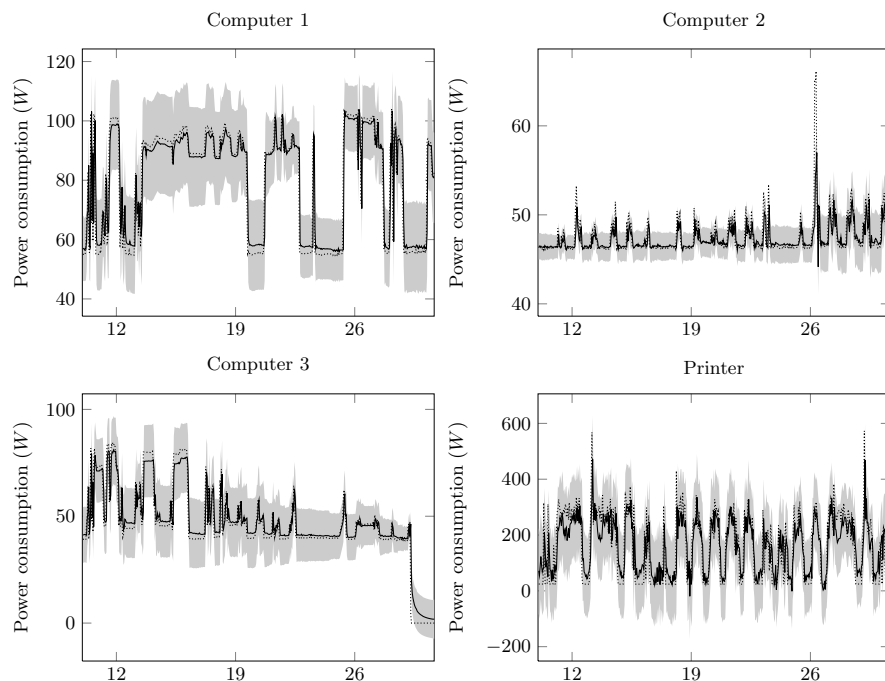
**Fig. 11.** Power consumption as obtained from the CPU utilisation and the quadratic polynomial fit function from Figure 10.

### 3.4 CPU vs. Condor relationship

Much of the load on Computers 1 and 3 is caused by jobs executed by the Condor system. Figure 12 looks at the relationship between the CPU utilisation and the number of active jobs submitted to Condor. Apart from some load for Computer 1 when there are a few Condor jobs, it looks like the load increases almost as a step function as soon as the number of jobs is over a certain threshold (around 500 jobs).



**Fig. 12.** Plots of CPU utilisation against the number of active Condor jobs.



**Fig. 13.** Autoregression model for Computers 1-3 and the Printer. At each hour shown, the data from the last ten days was used to predict the consumption for the next hour. The shaded area shows the 95% confidence interval of the prediction. The dotted lines are the actual values.

## 4 Forecast model of energy consumption

Predicting the CPU utilisation and power consumption for an immediate future time interval can be useful in many applications. For example, if a scheme is employed that puts inactive computers to sleep, a forecast model could detect possible candidates by looking for computers with 0% CPU utilisation. We demonstrate that such a prediction could be feasible. Figure 13 shows the results from a simple auto-regression forecast model, similar to a model previously used for predicting arrivals in a health-care system [5]. We take the data from the first 10 days as an initial training set. At each hour after the 10-th day, we forecast the value for the next hour and plot the corresponding 95% confidence interval. We form a new training set, replacing the data for the first hour with the real data for the predicted hour. We continue until the end of the data set.

The predictions look fairly accurate. However, they do not immediately react to sudden changes in the time series value. This can be crucial in many applications such as the sleep algorithm mentioned above. Therefore we plan to explore alternative, more accurate models.

## 5 Conclusion

In this paper we have shown how the energy usage of standard office equipment can be measured using common power monitors. We presented some insightful visualisations of the data and observed a simple quadratic relationship between the power consumption and CPU utilisation. This relationship could be used to reduce the number of needed power monitors. After a suitable calibration, the power consumption could be estimated from the easily accessible CPU utilisation information.

We demonstrated that it is possible to use the measured data to predict future use. This can lead to many potential applications. For example, the Condor scheduling algorithms could use the predictions to choose computers suitable for running submitted jobs and ones suitable for hibernation. Another possibility is to provide guidance to the users who could be encouraged to adjust their behaviour accordingly. In any case, more sophisticated time series models would be required to capture the different feedback mechanisms.

We are planning to run a larger experiment using the experience described here. We plan to improve the reliability of the data recording software and run the measurements on a larger scale. We also plan to include other sources of energy measurements, such as clamp meters. We hope to develop a technique that will help to discover possible improvements to the current infrastructure and energy management. The data and analysis scripts in R [4] have been made available on our website [3].

## Acknowledgment

Anton Stefanek, Uli Harder and Jeremy T. Bradley are funded by EPSRC on the Analysis of Massively Parallel Stochastic Systems (AMPS) project (reference EP/G011737/1).

## References

1. *AlertMe*. Available at <http://www.alertme.com/>.
2. *Cosm*. Available at <https://cosm.com/>.
3. *Energy usage data*. Available at <http://aesop.doc.ic.ac.uk/tools/energy/>.
4. *The R Project for Statistical Computing*. Available at <http://www.r-project.org>.
5. Susanna Wau Men Au-Yeung (2008): *Response Times in Healthcare Systems*. Ph.D. thesis, Imperial College London. Available at <http://pubs.doc.ic.ac.uk/response-times-in-healthcare/>.
6. Suzanne Rivoire, Parthasarathy Ranganathan & Christos Kozyrakis (2008): *A comparison of high-level full-system power models*. In: *HotPower'08 Proceedings of the 2008 conference on Power aware computing and systems*, p. 3.
7. Douglas Thain, Todd Tannenbaum & Miron Livny (2005): *Distributed computing in practice: the Condor experience*. *Concurrency - Practice and Experience* 17(2-4), pp. 323–356.