CHLORIDE

Secure Power Always

Future power trends in dense single-user IT environments

Data Centre power density (W/m^2) is growing at an exponential rate. What is the impact upon UPS systems?

Building Services & Environmental Engineer, October 07

From the first microprocessor introduction in 1971 to the introduction of the Pentium IV in 2001 the heat output rose from 2W to 70W. The drivers for the losses were all related to the amount of work done by the chip itself in terms of how much current was switched and how often:

- The number of transistors increased x18,000 to 42m
- The switching speed increased x13,000 to 1.4GHz
- The track spacing decreased by x83 to 0.012 Microns

It is interesting to note that the track spacing is down to such a small dimension that the DC voltage used for the chip has to be 0.4VDC - so that a remarkable 350A flows through the tiny silicon chip at the heart of the processor package.

Since 2001 the trend has continued and the typical power consumed has risen to 140W. At the same time co-processors and piggy-back topologies have been introduced to increase the computing capacity (MIPS, millions of instructions per second) capacity in a given volume ($MIPS/cm^3$). The result has been the almost exponential growth in data centre power density (W/m^2) and the advent of specialised cooling methods and systems.

The rise of the specialist cooled cabinets and the resultant need for UPS power for their mechanical heat extraction systems are current issues for data centre designers.

Against this background some remarkable news has been (certainly deliberately) leaking out of the microprocessor OEM's: At a recent conference a representative of Intel announced that Watts/microprocessor was just about to fall from the recent industry norm of 140W to 60W or even less over the next 1-2 years. This will, of course, overlay the installed base of servers (with 3-4 years installed cycle time) and, therefore, probably work through to 90% of the population only over the next 4-5 years. It was clear from the review of the technological trends that co-processors (duo and quad etc) and piggy-back topology were already accelerating as the power ceiling per chip had been reached at around 140W.

At first hearing I wondered what effect that would have on our UPS business. After a little historical reading it was possible to plot and extrapolate the Watts/Chip trend from 1996-2011. **See Fig. 1.**

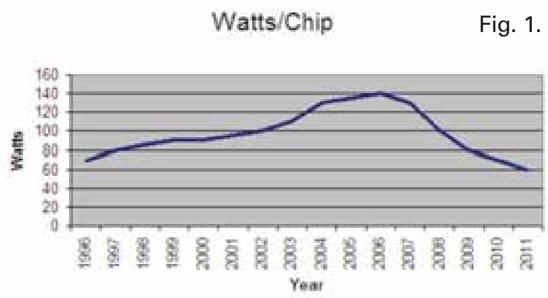

Fig. 1.

Looking back into the historical dimensions (height in U) of servers compared with the numbers of chips installed it can be seen that the heat rejection and mechanical packages has shaped the development and that the density rose from around 0.15/U in 1996 and can be expected to reach 5/U by 2011.

The second factor that has affected the historical power consumption is that of the partial load efficiency of the microprocessors themselves. They have traditionally been very poor in that when processing very little information compared to their capacity the power consumed has been above 80%. This has been improving and it is expected to reach 50% by 2011. To counter this trend the software technique of Virtualisation (where processor loads are maximised by the supervisory software) is being adopted very quickly. This is predicted to be fully adopted (processors running at full load) by 2011.

Taking the exploration of trends to the demand side it would appear from the historical data that the demand for computing capacity per 'box' (MIPS) has actually outstripped the MIPS/Chip curve. See Fig. 2.This is possibly a reaction to the marketing of the enterprise software but could also be a negative result of poor (processor hungry) software compilation. Taking the demand for MIPS at 50% growth per year but the available capacity per chip at only 30% growth per year produces the trend lines for MIPS (relative to 1996 being 100, over x40,000 by 2011) and the trend for the number of 42U high IT cabinets to house the servers/microprocessors.

If in the 'cabinet business' you will have to rely on the specialised cooling solutions because clearly the number of cabinets is going to reduce even with a huge growth in MIPS.

Three interdependent trend results are now sought from the model.

- The total UPS power
- The data-centre Watts/m^2
- The raised floor area, m^2

In **Fig. 3** the left hand scale is 'Number of Cabinets' in units and 'Raised Floor Area' in m^2. This model starts in 1996 with a notional 100 cabinet data-centre. It can clearly be seen that the W/m^2 will continue to rise and, therefore, drive the demand for cooling plant UPS CAPEX.

For the 'power demand' the model **(Fig. 4)** produces a very interesting result. It predicts that the IT load will level out (per installation but the number of installations will increase) but that an increasing number of installations will apply UPS to the mechanical load and so drive the power demand upwards.

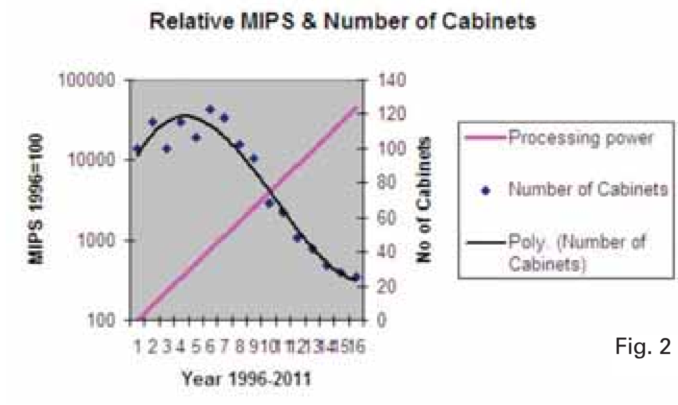

Fig. 2

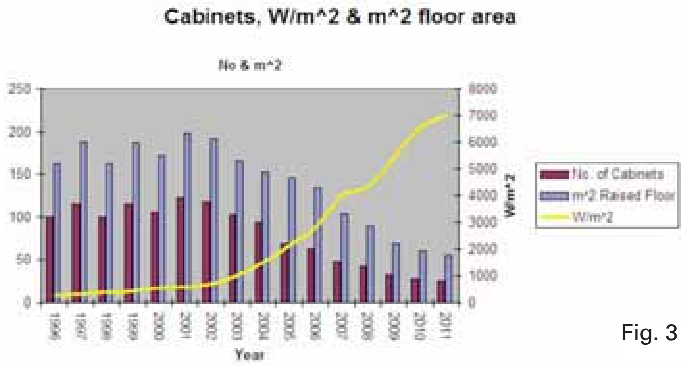

Fig. 3

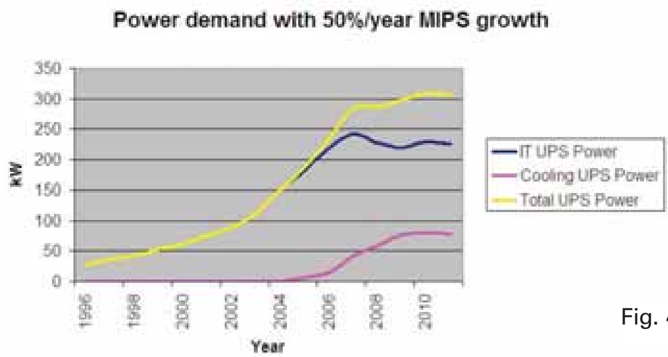

Fig. 4

About the author

Ian F Bitterlin
BSc(Hons) DipDesInn MCIBSE MIET

With over 36 years experience of writing technical articles for leading companies and institutions, Ian is a world renowned author and speaker and an expert in all aspects of critical power and building services.