

## Voltage, throughput, power, data reliability and multi-core scaling

PRiME is a five year EPSRC funded research programme (2013-2018), in which four UK universities address the challenges of **power consumption** and **reliability** of future high-performance embedded systems utilising many-core processors.

### The four PRiME parameters

In digital CMOS systems, a higher supply voltage ( $V_{DD}$ ) usually allows a higher operating (clock) frequency and hence a higher data throughput, but at the cost of higher power consumption. When power is limited, it is possible to obtain an increase in system throughput by scaling to multiple on-chip processor cores - if the computation can be parallelized and spread out amongst the cores.

These general knowledge points are usually obtained by considering the reliability (or integrity) of the data as a separate issue, i.e. by assuming that the data reliability is 100%. PRiME has instead studied the inter-relationship of all these factors together - voltage, throughput, power, and reliability in the context of multi-core scaling.

power consumption. In this context, the system must operate above a minimum frequency/ $V_{DD}$ , and below some constant power constraint. In addition, the frequency- $V_{DD}$  relationship must ensure that the logic can complete its operation between two clock pulses. **This is the timing reliability curve.** PRiME has carried out experiments with a number of real systems to characterize their behaviours in relation to Figure 1. Figure 2 shows data from one set of experiments, in this case with a Xilinx Zynq FPGA device.




Figure 1. Voltage, performance, reliability and power.

Figure 1 illustrates the usual design limitations of CMOS systems. Reliability can be defined as error-free or bounded-error operations satisfying some minimum throughput requirement and maximum




Figure 2. FPGA power, reliability,  $V_{DD}$  and frequency maps.

The panel on the left shows the power map with warmer colours indicating higher power consumption. The panel on the right shows the data reliability map. Reliability in this case is indicated by the signal to noise ratio (SNR) when the device is used to process a video sequence. White is infinity, i.e. full reliability, and colours moving from warm to cool indicate gradually decreasing SNR. When considering only the clock/timing reliability curve, the white area is the reliable operation region - where reliability is defined as  $SNR = \infty$ . As can be seen in figure 2, with non-infinity SNR values deemed as “reliable”, the timing

reliability curve would move up allowing higher frequency to be traded for relative reliability at any  $V_{DD}$ .

### The effect of scaling

PRiME has studied the effect on the relationships between the four PRiME parameters when multi-/many-core scaling happens. A system was first characterized for its single-core behaviour through its legal  $V_{DD}$  range, then scaling under fixed power budgets was studied with regard to the parameter interplays.

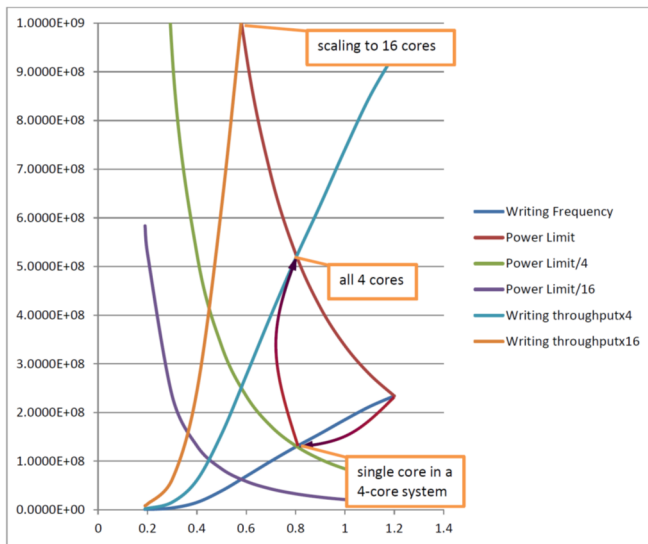


Figure 3. Scaling to multiple cores.

For a fixed power budget, enough to operate a single core at the nominal  $V_{DD}$  and maximum frequency without losing timing correctness, the system is scaled up to 4 and 16 cores. With 4 cores, each core would operate under  $\frac{1}{4}$  of the total power budget, this roughly corresponds with 0.8V and 129MHz. All 4 cores together give 0.8V and a throughput of 516MHz, a vast improvement over the single-core's capability of 234MHz at 1.2V. Scaling up to 16 cores will see the system working at around 0.6V and over 1GHz.

In the near-, trans- and sub-threshold voltage regions, the constant power curve no longer fits the shape shown in Figure 3. PRiME has found that below about 0.6V the benefits of scaling decrease drastically. This is true across the multiple systems that PRiME has studied.

### Systems with different timing reliability curves

Timing reliability curves collected across a number of systems have been normalized over  $V_{DD}$  and frequency (with the maximum values of each considered to be 1) and plotted together:




Figure 4. Steepness/shalowness of timing reliability curves.

Figure 4 shows how different systems may have differences in the steepness of their timing reliability curves. This impacts on how effective parallelization scaling (increasing the number of cores) can be for the purpose of increasing throughput, saving power, or both. The larger  $df/dV_{DD}$  is, the less effective scaling would be. In other words, the curve should ideally be shallow with a small  $df/dV_{DD}$ . Commercial systems tend to have higher  $df/dV_{DD}$  than a theoretical "perfect" CMOS system and this leads to lower effectiveness for parallelization scaling on the PRiME parameters.

### Future work

Parametric modelling will be incorporated into the more qualitative modelling methods within PRiME.

### More information

Visit the PRiME programme web site, including the opportunity to sign-up for programme updates, or contact the PRiME Collaboration Manager (including industry liaison):

**Gerry Scott**  
Email: gerry.scott@soton.ac.uk  
Tel: +44 (0)23 8059 2749



[www.prime-project.org](http://www.prime-project.org)