(When) Will CMPs hit the Power Wall?

Cor Meenderinck and Ben Juurlink
Computer Engineering Department
Delft University of Technology, the Netherlands
Email: {cor,benj}@ce.et.tudelft.nl

Abstract— **Currently the power wall is one of the major obstacles chip industry is facing. At the same time processor architecture shifts towards chip multiprocessors (CMPs), which are believed to alleviate the power problem. In this paper we analyze the impact of the power wall on CMP design. As a case study we model a CMP consisting of Alpha 21264 cores, scaled to future technology nodes according to the ITRS roadmap. In 2020 such a CMP would contain 625 cores, and when running on the maximum possible frequency of 73 GHz would consume 7 kW, while the power budget is predicted to be 198 W. From these figures it is clear that power will be a major bottleneck for performance increase. However, we also calculated the power constrained performance increase which shows that technology improvements enables a doubling of performance increase every three years for CMPs. Overhead and limited amounts of available thread level parallelism eventually cause this increase to flatten out. At the architectural level, performance can be increased by designing for power efficiency. We conclude from this study that for the next decade CMPs can provide significant performance increases without hitting the power wall.**

I. Introduction

It is said that we are reaching or have reached the power wall, meaning that performance increase comes to an end due to power constraints. Considering the fact that industry has shifted towards developing chip multiprocessors (CMPs), we can conclude that monocores have hit, or at least are very close to the power wall indeed. At the same time it is believed that multicore architectures alleviate the problem (see Figure 1). First of all, this is because multicores allow to exploit Thread Level Parallelism (TLP), which increases performance. Second, and more importantly, CMPs allow a power efficient way of exploiting parallelism. In contrast, many aggressive techniques to exploit Instruction Level Parallelism (ILP), e.g., superscalar, out-of-order, and hyper-pipelining, are power inefficient [1].

When designing future CMPs, power-performance trade-offs and optimizations have to be made. The question arises how serious the power problem is, and thus how big an architectural paradigm shift is required. Power efficiency of CMPs can be improved by designing asymmetric multicores. For example, several domain specific accelerators could be employed which are turned on and shut off according to the actual workload. But, is the power saving it provides worth the area cost? In this paper we try to answer those questions.

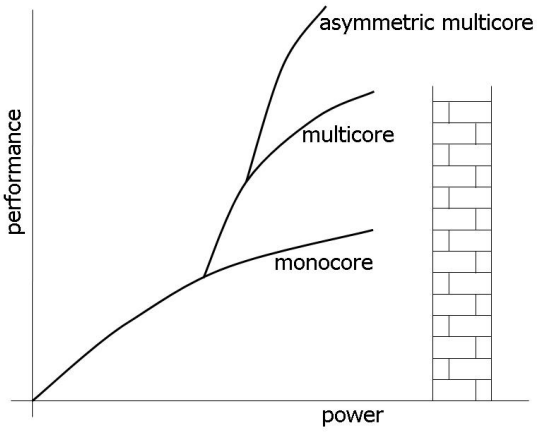

Fig. 1. The power wall problem.

Specifically, in this paper we focus on technology improvements as they have been one of the main drivers of performance increase in the near past. According to the ITRS roadmap [2], technology improvements are expected to remain. It predicts clock frequencies of 73 GHz for the year 2020 and an astonishing amount of transistors will be available. Because of power constraints, those technology improvements, however, might possibly not be able to be exploited completely. In the next section we analyze the limits of performance increase due to technology improvements with respect to power constraints.

II. The alpha experiment

To analyze the effect of technology improvements on the performance of future CMPs, and to investigate the power consumption trend, the following experiment was performed. Take an Alpha 21264 core,

TABLE I
DETAILS OF THE ALPHA 21264 CORE.

year	1998
technology node	$350nm$
supply voltage	$2.2V$
die area	$314mm^2$
dyn. power ($400MHz$)	$48W$
dyn. power ($600MHz$)	$70W$

scale it to future technology nodes according to the ITRS roadmap, create a hypothetical CMP consisting of the scaled cores, and calculate the power numbers. Specifically, we calculate the power consumption of a CMP for full blown operation, i.e., all cores are active and run at the maximum possible frequency. Furthermore, we analyze the performance increase over time if the power consumption is restricted to the power budget allowed by packaging.

The Alpha 21264 [3] core was chosen as subject of this experiment for two reasons. First, the Alpha cores have been described in detail in scientific articles providing the required data for the experiment. Second, the 21264 is a moderate sized core lacking the aggressively exploiting ILP techniques of current high performance cores. Thus, it is a good representation of what is generally expected to be normal in future many-core CMPs. Table I provides an overview of the key details of the 21264 relevant for this analysis.

Scaling the 21264 is done according to data in the 2005 edition of the International Technology Roadmap for Semiconductors (ITRS) [2]. The relevant parameters are stated in Table II. The time frame considered is the year of writing to the end of the roadmap: 2007-2020. The values of the technology node and the on-chip frequency were taken from the executive summary page 79. The on-chip frequency is based on the fundamental transistor delay, and an assumed maximum number of 12 inverter delays. The die area values were taken from the table on page 81 of the executive summary. Finally, the values of the supply voltage and the gate capacitance (per micron device) were taken from the table starting at page 11 of the 'process integration, devices, and structures' section of the roadmap.

To model the experimental CMP for future technology nodes, we scale all required parameters of the 21264 core. The values that are available in the ITRS, we use as such. The others we scale using the available parameters by taking the ratio between the original 21264 parameter values and the predictions of the roadmap. The gate capacitance of the 21264 was not found in literature, thus we extrapolated the values of the roadmap and found a value of $1.1 * 10^{-15} F/\mu m$ for 1998.

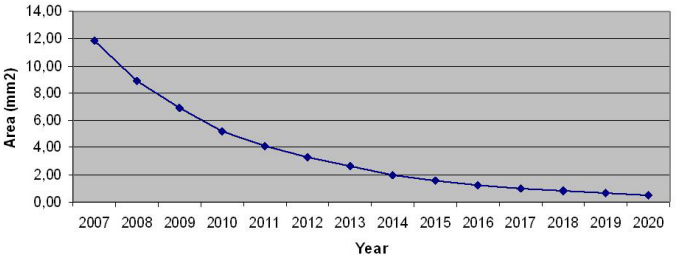

Fig. 2. Area of one core.

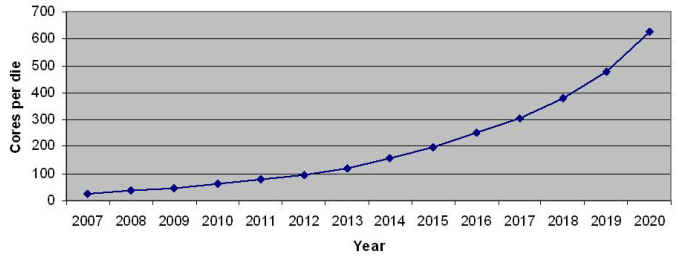

Fig. 3. Number of cores per die.

First, the area of one core was scaled. Let $L(t)$ be the process technology size for year t and let L_{orig} be the process technology size of the original core. The area of one 21264 core in year t will be $A_1(t) = A_{orig} * \left(\frac{L(t)}{L_{orig}}\right)^2$. Figure 2 depicts the results for the time frame considered. The area of one core decreases quadratically over time, and will be about half a square millimeter in 2020.

Second, using the scaled area of one core, the number of cores that fit on a die was calculated. The ITRS roadmap assumes a die area of $310mm^2$ for the entire time frame. Thus, the total number of cores per die in year t is $\#_{cores}(t) = \frac{310}{A_1(t)}$, and is depicted in Figure 3. For 2007 it was calculated that a 26-core CMP would be possible. This is reasonable as two examples of current state-of-the art CMPs are the Tilera Tile64 [4] and the Clearspeed CSX600 [5]. The first has 64 cores, each having their own L1 and L2 cache and can run a full operating system autonomously. These cores are very similar in size to the Alpha 21264. The latter has 96 cores which are simpler and intended to exploit data level parallelism. In 2020 our calculations predict a CMP with 625 cores. This might be considered an extremely large number, but it is in line with the expected doubling of cores every three years [6].

TABLE II
TECHNOLOGY PARAMETERS OF THE ITRS ROADMAP.

	2007	2008	2009	2010	2011	2012	2013
technology (nm)	68	59	52	45	40	36	32
frequency (MHz)	9285	10972	12369	15079	17658	20065	22980
die area (mm^2)	310	310	310	310	310	310	310
supply voltage (V)	1,1	1	1	1	1	0,9	0,9
$C_{g,total}$ ($F/\mu m$)	6,99E-16	9,04E-16	7,55E-16	7,35E-16	6,5E-16	6,29E-16	6,28E-16
	2014	2015	2016	2017	2018	2019	2020
technology (nm)	28	25	22	20	18	16	14
frequency (MHz)	28356	33403	39683	45535	53207	62443	73122
die area (mm^2)	310	310	310	310	310	310	310
supply voltage (V)	0,9	0,8	0,8	0,7	0,7	0,7	0,7
$C_{g,total}$ ($F/\mu m$)	5,59E-16	5,25E-16	5,25E-16	4,78E-16	4,48E-16	4,1E-16	3,62E-16

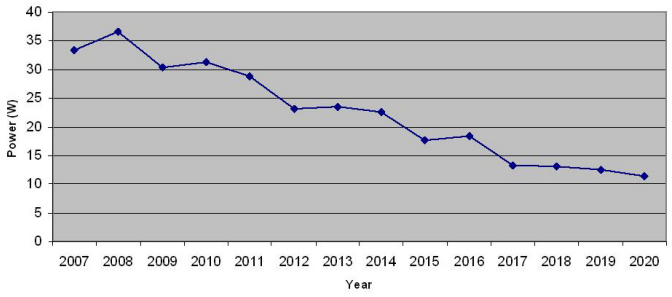

Fig. 4. Power of one core.

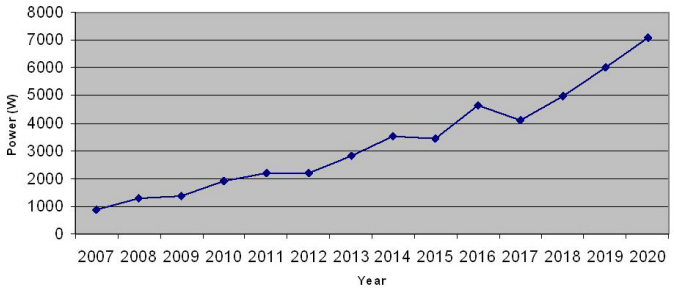

Fig. 5. Total power of the case study CMP.

Finally, the power of one core was scaled. Power consumption consists of a dynamic and a static part, of which the latter is dominated by leakage. The data, required to scale the static power, is not available to us and thus we restrict this power analysis to dynamic power. It is expected that leakage remains a problem and thus our calculations are on the safe side.

The equation to calculate the dynamic power is $P_{dyn} = \alpha C f V^2$, where α is the transistor activity factor, C is the gate capacitance, f is the clock frequency, and V is the power supply voltage. The activity factor α of the 21264 processor is unknown, but as the architecture does not change with scaling, it is assumed to be constant. The capacitance C (F) in the equation is different from capacitance $C_{g,total}$ $(F/\mu m)$ in Table II, but they relate to each other according to $C \propto C_{g,total} * L$. Thus, the dynamic power at year (t) is calculated as:

$$P(t) = P_{orig} * \frac{C_{g,total}(t) * L(t)}{C_{g,total,orig} * L_{orig}} * \frac{f(t)}{f_{orig}} * \left(\frac{V(t)}{V_{orig}}\right)^2. \quad (1)$$

The reader should notice that this analysis assumes that cores run at the maximum possible frequency. Figure 4 depicts the power of one core over time. As the curve shows it roughly halves every seven years, resulting in 11 W in 2020.

Now all parameters have been scaled, it is possible to calculate the power consumption of the total CMP. It is assumed that each core is active and thus $P_{total}(t) = \#_{cores}(t) * P(t)$. Figure 5 depicts the total power over time and shows that for the assumptions of this analysis the total power consumption gradually increases and reaches 7 kW in 2020. It is clear that in practice this large power consumption is not possible. The roadmap predicts that the power budget allowed due to packaging constraints is 198 W. This is why power has become one of the main design constraints nowadays.

The figure also shows that the difference between the power budget and the power consumption of the full blown hypothetical CMP is increasing over time. That means that a large part of the technology improvement can not be put into effect for performance increase. For example, between 2011 and 2015 technology allows doubling the on-chip frequency, however the power consumption would increase with a factor

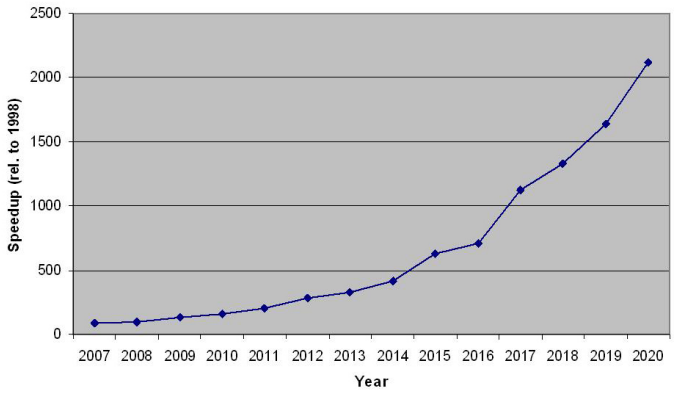

Fig. 6. Power constrained performance increase.

1.5. Thus, for equal power only a small frequency improvement would be possible.

This poses the question to what extend technology improvements can contribute to the performance increase of CMPs. To analyze this a performance measure is required. Actual performance is depending on many factors and can only be measured on real systems. Thus we express performance in terms of hardware capabilities. For the case study of this paper, the parameters that influence performance are frequency and the number of cores. We define a speedup S, relative to the original 21264 core, as $S = \frac{f(t)}{f_{orig}} * \frac{\#_{cores}(t)}{1}$.

We are interested in the speedup of CMPs that meet the power budget of 198 W. As the results show using the maximum possible frequency and using all cores concurrently, the power budget is exceeded. To meet the power budget, either the frequency could be scaled down or a number of cores could be shut down. Both measures are linear to the speedup, and thus the power constrained speedup can be defined as:

$$S_{p.constr.} = \frac{f(t)}{f_{orig}} * \frac{\#_{cores}(t)}{1} * \frac{P_{budget}}{P(t)}. \qquad (2)$$

Figure 6 depicts the power constrained performance of the case study CMP over time. The curve shows a doubling of performance every three years. Although power on becomes a bottleneck for performance, using the CMP paradigm performance increase is possible.

III. Conclusions

In this paper we analyzed the impact of the power wall on CMP design. Specifically, we investigated the limits to performance increase of CMPs due to technology improvements with respect to power constraints. As a case study we modelled a CMP consisting of Alpha 21264 cores, scaled to future technology nodes according to the ITRS roadmap. We calculated that in 2020 such a CMP would contain 625 cores, each consuming 11 W when running on the maximum possible frequency of 73 GHz. The total CMP, at full blown operation, would consume 7 kW while the power budget is predicted by the ITRS to be 198 W.

From these figures it is clear that power has become a major design constraint, and will be a major bottleneck for performance increase. However, it does not mean that the power wall has been hit for CMPs. We calculated the power constrained performance increase and showed that technology improvements enables a doubling of performance every three years for CMPs. Eventually, this performance increase will flatten out due to increasing overhead and limits to the amount of TLP that can be exploited.

At the architectural level performance improvements can be achieved by using power efficient techniques such as asymmetric CMPs and domain specific accelerators. Our analysis showed that the transistor budget will not be the bottleneck. Thus, putting several extremely power efficient domain specific accelerators on a CMP improves overall performance, although they might be shut off a significant amount of time.

From this study we conclude that for the next decade CMPs can provide significant performance increases without hitting the power wall. Technology improvements will provide the means, however, to be exploited, at the architectural level power efficiency should be the main design criterion.

Acknowledgment

The authors would like to thank Stefanos Kaxiras for his input on the methodology used in this paper.

References

[1] H. Hofstee, "Power efficient processor architecture and the cell processor," in *Int. Symp. on High-Performance Computer Architecture*, 2005, pp. 258–262.

[2] "International Technology Roadmap for Semiconductors, 2005 Edition," 2005. [Online]. Available: http://www.itrs.net

[3] R. Kessler, "The Alpha 21264 microprocessor," *Micro, IEEE*, vol. 19, no. 2, pp. 24–36, 1999.

[4] Tilera, "TILE64(TM) Processor Family." [Online]. Available: http://www.tilera.com

[5] ClearSpeed, "The CSX600 Processor." [Online]. Available: http://www.clearspeed.com

[6] P. Stenström, "Chip-multiprocessing and Beyond," in *Proc. Twelfth Int. Symp. on High-Performance Computer Architecture*, 2006, pp. 109–109.