An Exploration of Tiled Architectures for Space Applications

Peter Kogge, *Fellow, IEEE*, **Megan Vance**, *Student Member, IEEE*

***Abstract*—For many reasons, the chip technology for spacecraft computing has lagged commercial systems by decades. Equally disconcerting, however, has been a similar lag in the computer architectures used. This paper will look at an emerging class of multi-core processor architectures, called "tiled," and extrapolate what they might look like in the future, and how they might be adapted to space applications. This extrapolation will include physical characteristics (speed, power and area), system architectures, and programming models.**

I. INTRODUCTION

Multi-core architectures have become the standard for modern commercial microprocessor chips, with several different classes system architectures [6]. One such is "tiled," where a self-contained processor is replicated in an array-like pattern over a die. Two such chips with different characteristics are EXECUBE[9,10] and Tile64 [3]. The first was built on a DRAM technology base, the second on a logic base. The first is a scalable single part type design without need for external memory; the second requires external DRAM but has a rich I/O suite. The first supports a SIMD and SPMD message-passing; the second shared memory. While both are fixed point, the first has a simple pipeline where each core is tied directly to on-chip memory; the second supports multiple instruction issue with a two level coherent cache design.

II. SCALING

The full paper will place each design point into a common technology base, and extrapolate to the end of the silicon roadmap[4], as was done in [2] and [5]. Three different scenarios will be explored:

- Simple die shrink: where new technology will be used to reduce power and area, but not the characteristics.
- Constant die size, same core: die area is kept constant, and filled with as many cores as possible.
- Constant die size & core count: the die area is kept constant, as is the number of cores, with the additional space used for more on-chip memory.

Power, speed, and performance will be compared, as will points in time where other constraints such as pins or bandwidth will limit the chips' potential.

III. 3. SPACE SYSTEM ARCHITECTURE NEEDS

Each of the baseline chips has different characteristics that are valuable for embedded space applications, but neither is "optimal." Floating point support, fault tolerance in both hardware and software, inherently lower power and enhanced power management, robust programming models, and the ability to seamlessly scale to different design points of both memory, performance, and multiple generations of technology could all be enhanced. In the full paper, we will develop each more fully, and propose a potential tiled architecture that may be more suitable. Many of the observations made in a recent projection to exascale computing[2] will be used in this extrapolation, as will prior investigations into deep space processing needs[8] and power management approaches[7].

IV. 4. EXPLORING A SPACE-OPTIMIZED TILED CHIP

Using the above tiled architecture, the full paper will use a new tool[1] to explore the range of feasible chip implementations through multiple generations of technology, and taking into account constraints such as limited contacts, the need to keep per chip power dissipation under control, and the potential for advanced packaging, such as 3D stacks.

V. 5. FUTURE WORK

The full paper will also discuss what else needs to be done to make such tiled chips useful for space applications. Particular attention will be made to programming models that support both robust scaling and significant fault tolerance.

REFERENCES

1. M. Vance & P. Kogge, "SASIC Design Space & Performance Exploration," Notre Dame CSE TR-2008-17.
2. P. Kogge (editor), "ExaScale Computing Study: Technology Challenges in Achieving Exascale Systems," Notre Dame CSE TR-2008-13.
3. S. Bell, et al., "TILE64 Processor: A 64-Core SoC with Mesh Interconnect," ISSCC, Feb. 2008.
4. -, Int. Technology Roadmap for Semiconductors, 2007 ed.
5. P. Kogge, "Some Initial Explorations into the Hierarchical Multi-Core Chip Design Space for HPC Systems" IWIA, Jan. 2007.
6. P. Kogge, "An Exploration of the Technology Space for Multi-Core Memory/Logic Chips for Highly Scalable Parallel Systems," IWIA, Jan. 18-19, 2005.
7. P. Kogge, et al, "A Comparative Analysis of Power and Energy Management Techniques in Real Embedded Applications," IWIA, Jan. 2003.
8. P. Kogge, et al, "Characterization of Future Deep Space Computing Loads," SMC-IT, July, 2003.
9. T. Sunaga, P. Kogge, et al, "A Processor In Memory Chip for Massively Parallel Embedded Applications," IEEE J. Solid State Circuits, Oct. 1996.

10. P. Kogge, "The EXECUBE Approach to Massively Parallel Processing," ICPP, August, 1994.