10x10 Heterogeneous Architecture with OpenDwarfs on FPGA

Anshuman Verma

Power Wall and Multi core generation

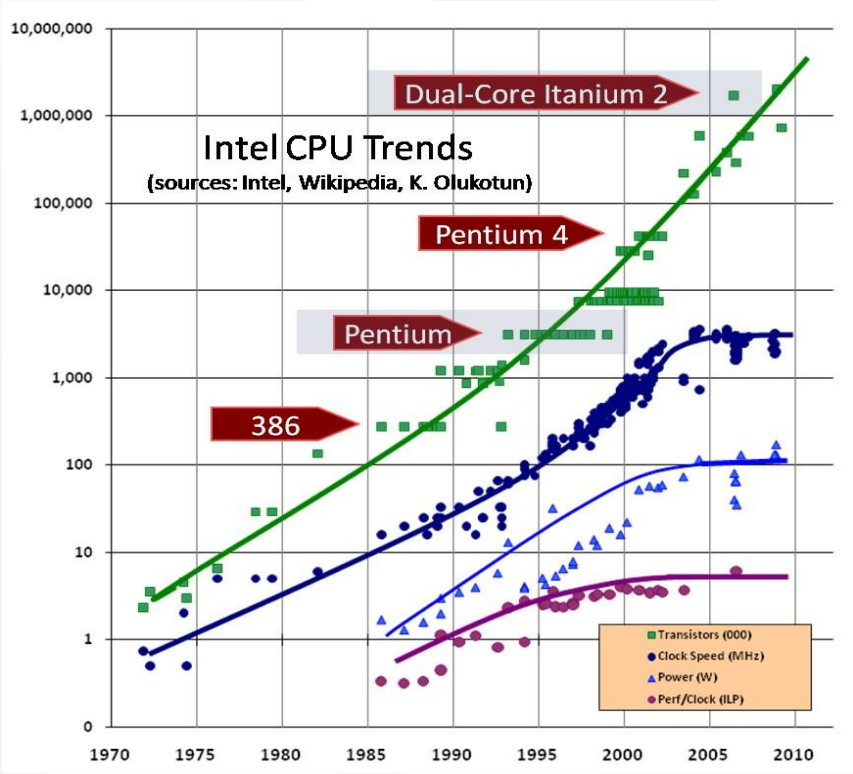

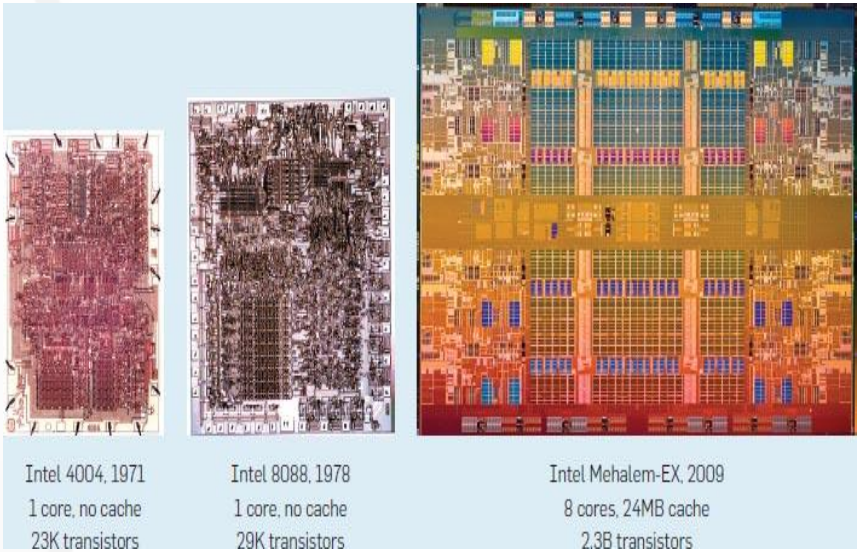

Dark Silicon

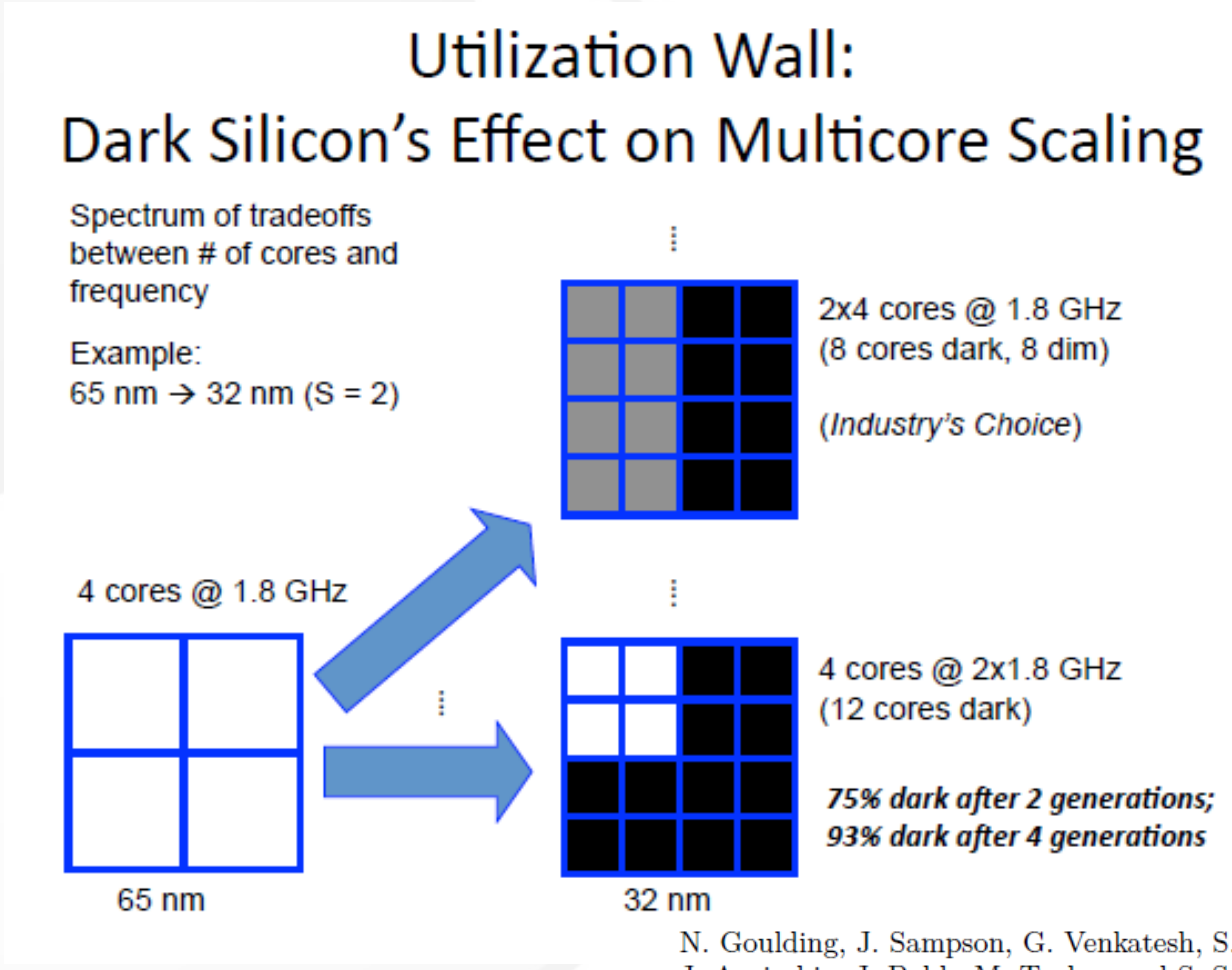

N. Goulding, J. Sampson, G. Venkatesh, S. Garcia, J. Auricchio, J. Babb, M. Taylor, and S. Swanson. "GreenDroid: A mobile application processor for a future of dark silicon." In *HOTCHIPS*, 2010.

Traditional Optimization: 90/10 Paradigm

Workloads

"ILP"
"reuse locality"
"linear access"
"bit-field opns"
"branch patterns"

Abstracted
"common" cases

"pipelining"
"superscalar"
"caches & blocks"
"mmedia"
"branch pred"

Optimizations

- Workloads: analyze and derive common cases(90%)
- Invent arch features, implementation optimizations with broad impact (90%)
- Improve performance by adding optimizations
- Aggregation and Efficiency: 8080 80 insts => SB 500+ instructions

10x10 Optimization Paradigm

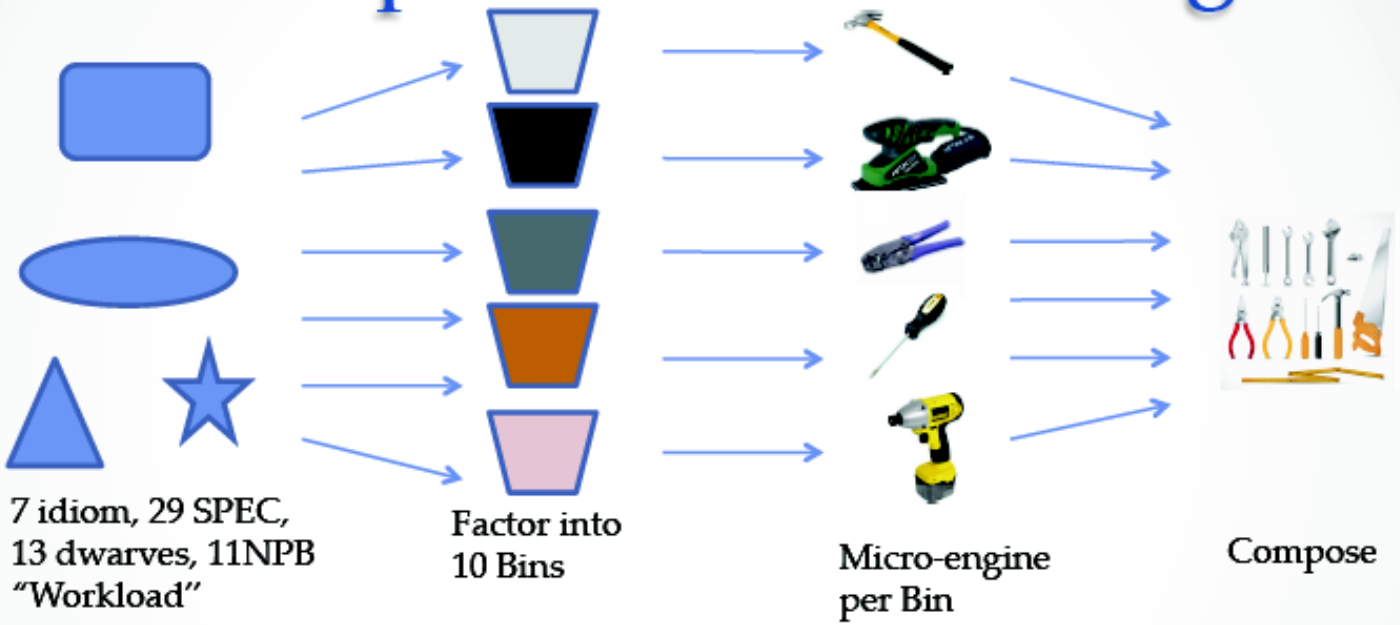

- Identify 10 application clusters; compute structures; datatypes (focus on 10 distinct bins)
- Optimize architecture, Optimize implementation of each separately (improve energy-delay product by 10-100x)
- Compose together sharing memory hierarchy and interconnect (preserve the benefits of customization)

10x10 Architecture

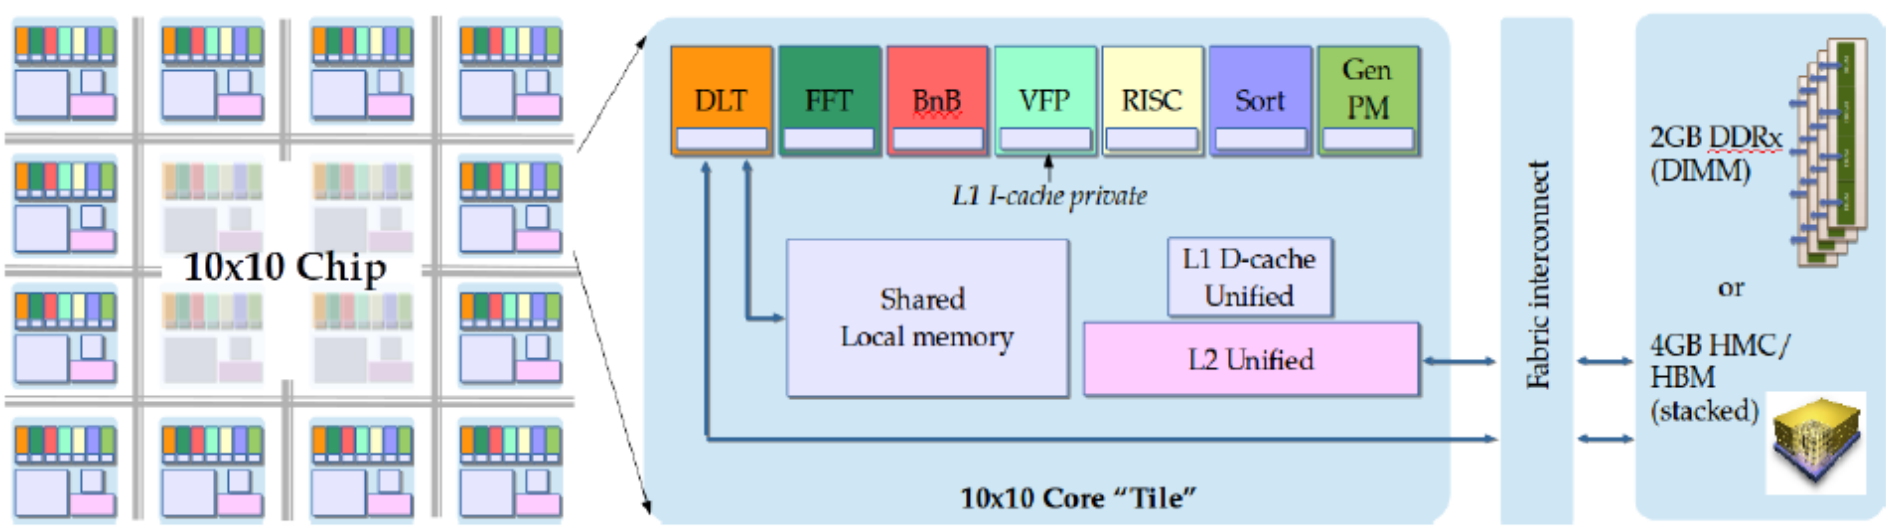

. 10x10 Chip includes tiled 10x10 cores; each has several Micro-engines.

What is OpenDwarf ?

- A (computational) dwarf is an algorithmic method that *captures a pattern of computation and communication*
  - Inspired by Phil Colella, who identified *seven* numerical methods important for science and engineering

- Benchmark Suite of 13 Computational Dwarfs & Apps

Sparse Linear Algebra
Dense Linear Algebra
Spectral Methods
Unstructured Grids
N-Body Methods
Structured Grids
Monte Carlo → MapReduce

+

Combinational Logic
Graph Traversal
Dynamic Programming
Backtrack & Branch+Bound
Graphical Models
Finite State Machine

OpenCL synthesis on FPGA

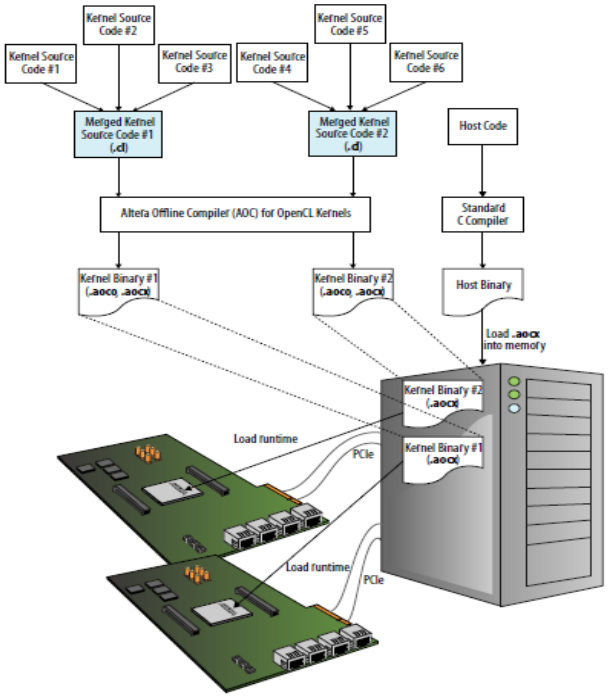

Table 1. Monte Carlo Black-Scholes Results

Platform	Power Watts (W)	Performance Simulations per second (Bsims/s)	Efficiency Simulations per second per Watt (Msims/s/W)
W3690 Xeon Processor	130	0.032	0.0025
nVidia Kepler20	212	10.1	48
BittWare S5-PCIe-HQ	45	12.0	266

- Pros
  - Customized Hardware
  - Better performance/Watt
- Cons
  - Large synthesis time

FPGA Advantage

- Pros
  - Customized Hardware
  - Better performance/Watt
- Cons
  - Large synthesis time

Table 1. Monte Carlo Black-Scholes Results

Platform	Power Watts (W)	Performance Simulations per second (Bsims/s)	Efficiency Simulations per second per Watt (Msims/s/W)
W3690 Xeon Processor	130	0.032	0.0025
nVidia Kepler20	212	10.1	48
BittWare S5-PCIe-HQ	45	12.0	266

https://www.altera.com/en_US/pdfs/literature/wp/wp-01173-opencl.pdf

Refactoring OpenDwarfs

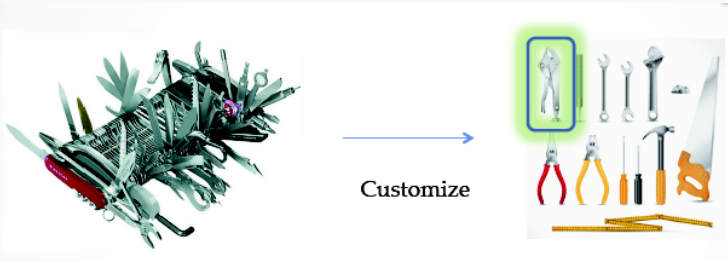

Contribution

- Lower synthesis time for Kernels
- Improvement in performance/Watt
- Proving the concept on FPGA

Thank You