Benchmarking for SuperComputer Ranking

Veljko Milutinovic

Agenda

- Supercomputing
- HPC Rankings
- The Maxeler Data flow
- Suggestion: A new metric
- Suggestion: A different workload
- Conclusion

Supercomputing

- A supercomputer is a computer at the frontline of the current computational capacity
- Supercomputers were introduced in the 1960s
- Very large systems

Applications

- Molecular Dynamics Simulation (Tianhe-A1)
- Artificial Neuron Simulation (BlueGene/P)
- Weather Forecasting
- Oil and Gas
- Market Analyses
- Etc...

The List (Top500)

- Established in 1993
- Gives the ranking of the top 500 World's most powerful Supercomputers
- Focus on the performance of general purpose systems
- Updated every 6 months

1. K computer, 2011 Fujitsu, Japan
2. NUDT YH MPP, 2010, NUDT, China
3. Cray XT5-HE 2009, Cray Inc, USA
4. Dawning TC3600 Blade 2010, Dawning, China

The Benchmark

- The Top 500 list is built upon the performance on the LINPACK benchmark
- Solving a dense system of linear equations.
- Matrix vector multiplication, a common problem in scientific computations
- Gives an estimate of peak performance

The Metric (FLOPS)

- Floating Point Operations per second
- Gives an estimate of how fast
  a computer solves
  floating point-intensive problems
- TOP500 is based on the peak FLOPS
- We are currently in the PetaFLOPS range

The Issues

- Floating point operations do not dominate the execution time in modern systems!
  (they used to, though)
- With the growing complexity of HPC systems, it is becoming more difficult to make full use of peak FLOPS
- Novel approaches exist, for which FLOPS are even less relevant
  (Example: Maxeler systems)

Example: The Maxeler Data Flow

- A custom chip for a specific application
- No instructions no instruction decode logic
- No branches no branch prediction
- Explicit parallelism: No out-of-order scheduling
- Data streamed onto-chip: No multi-level caches

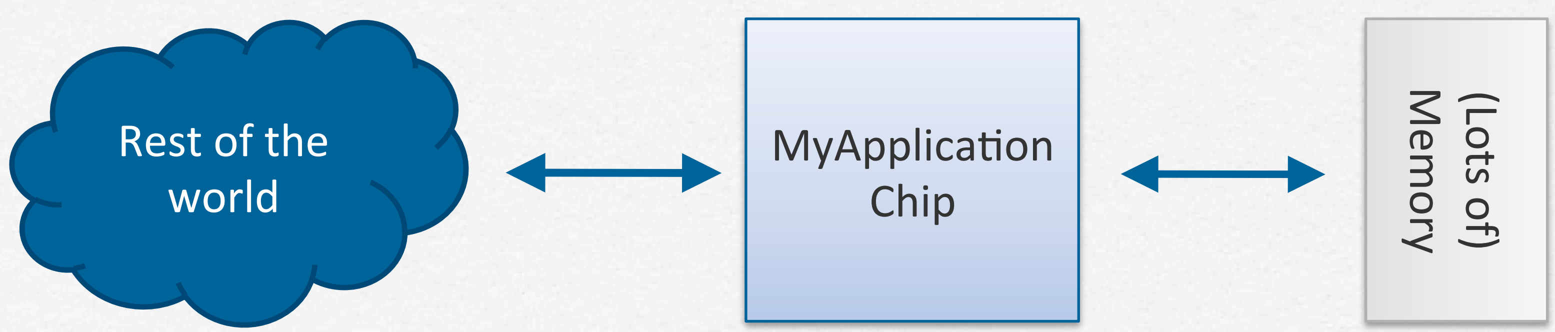

Example: The Maxeler Data flow (2)

- But we have more than one application
- Generally impractical to have machines that are completely optimized for only one code
- Need to run many applications on a typical cluster

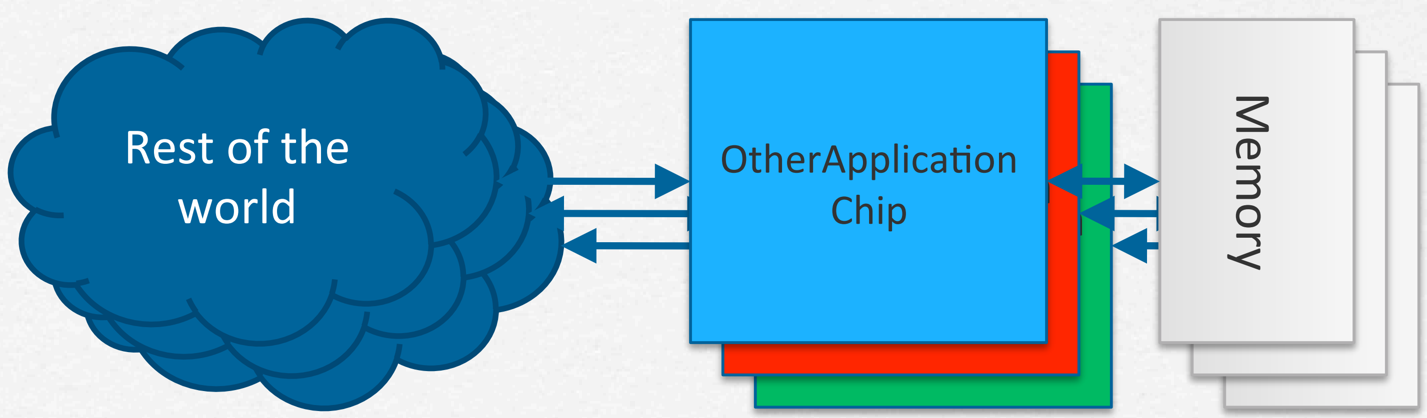

Example:
The Maxeler Data flow (3)

- Use a reconfigurable chip
  that can be reprogrammed at runtime
  to implement:
  - Different applications
  - Or different versions of the same application

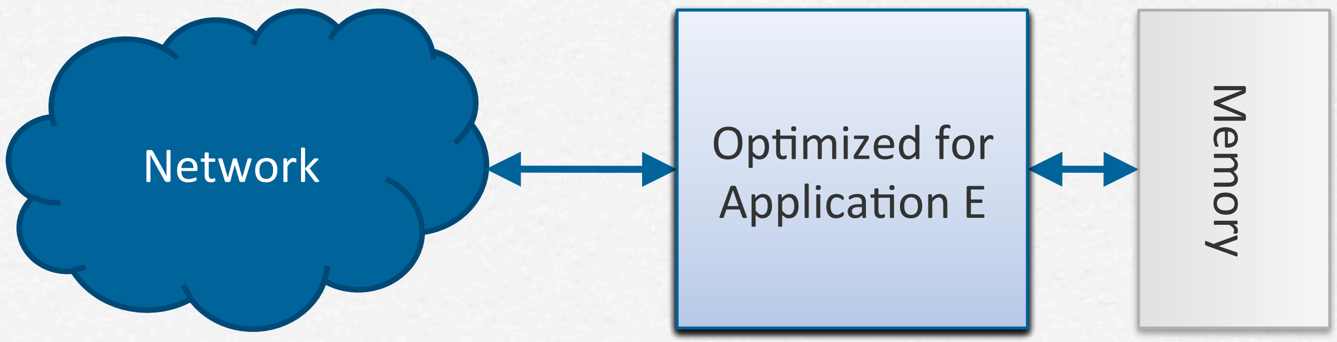

Example:
The Maxeler Data flow (4)

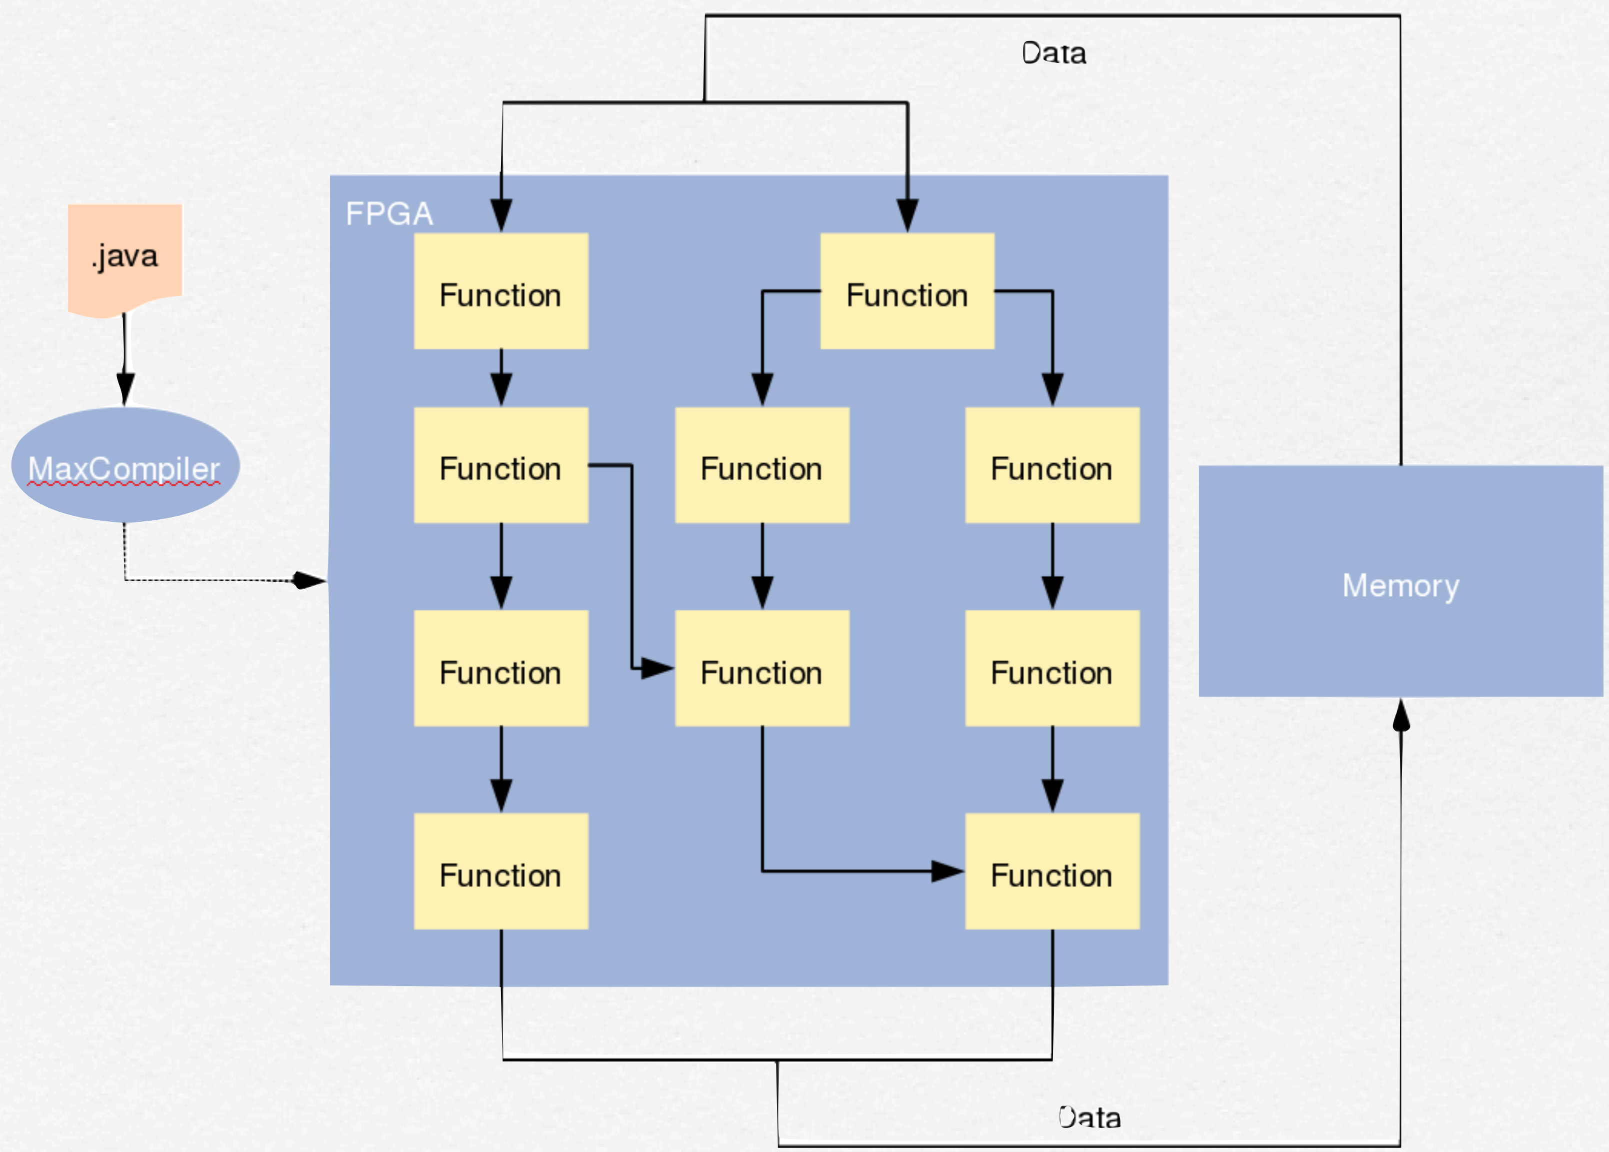

Example:
The Maxeler Data flow (4)

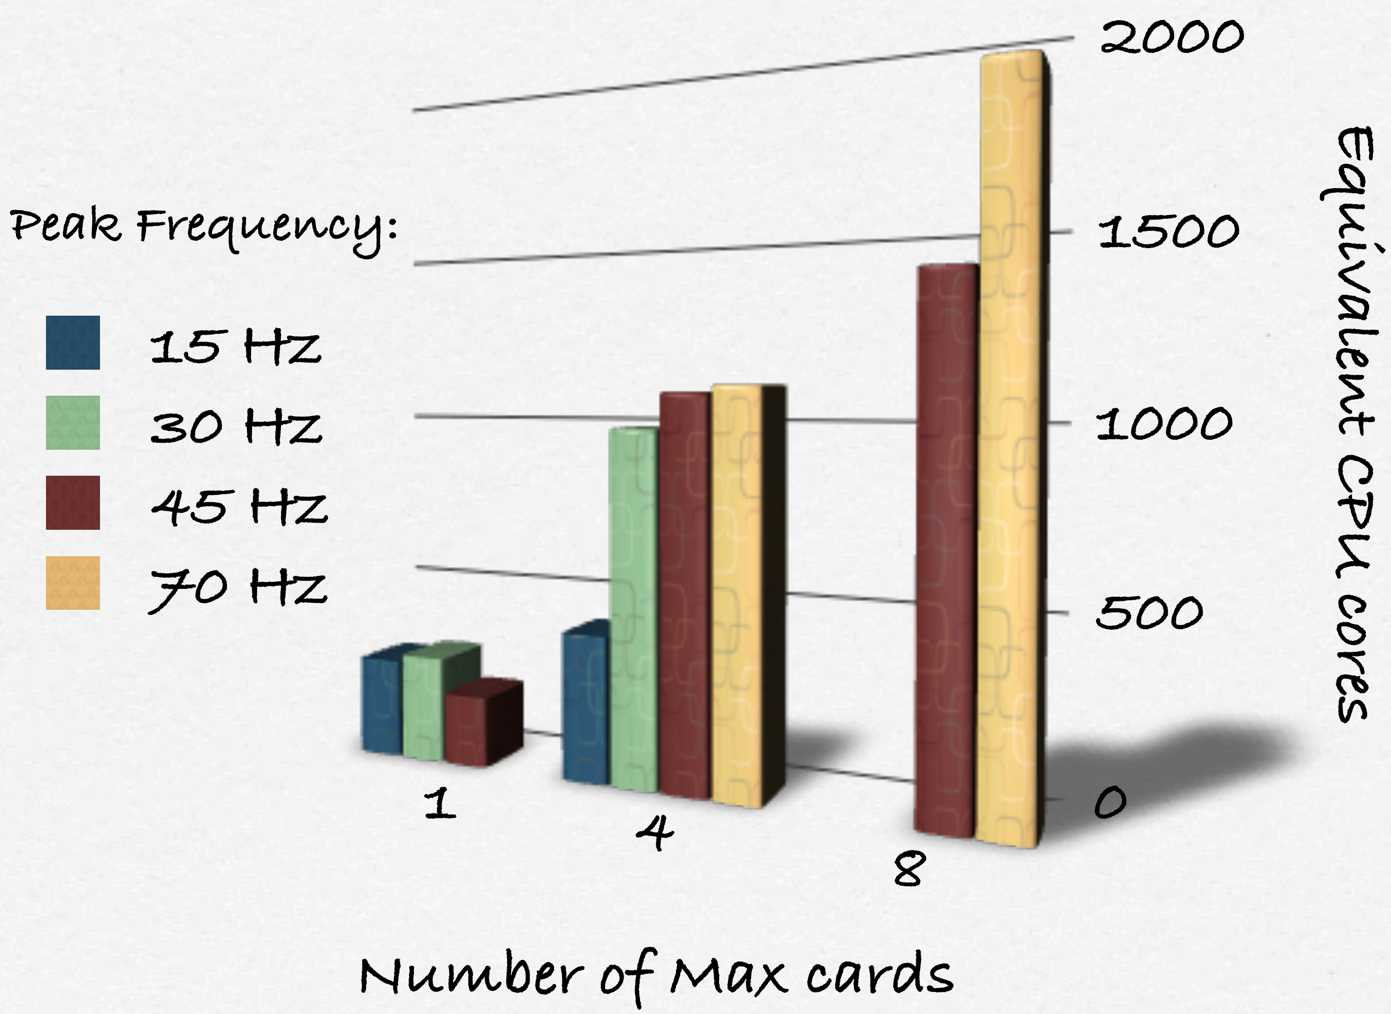

Performance of Maxeler-accelerated Finite Difference Modeling

Platform	Idle	Load
Dual Xeon 2.66GHz	185W	255W
with Max 2 cards	210W	240W

Power usage published by J. P. Morgan, the credit derivatives risk calculation

Platform	Speedup
Full precision	31x
Reduced precision	37x

Speedup versus 8 core Xeon server published by J. P. Morgan, the credit derivatives risk calculation

Why FLOPS do not tell the whole story?

- Memory access rate goes down, due to a large on-chip memory
- In a given computation, each result has many operations behind it, but that does not matter in dataflow computing since the computation is a side effect of the data flowing through the chip
- No load, store or brach instructions

Why LINPACK does not tell the whole story?

- Focus on FLOPS
- A highly regular generic workload
- Expansions have been suggested:
  - Graph500
  - HPC Challenge
- But these still focus on control flow systems

Suggestion 1: Datarate

- The rate of result production
  (e.g. Petabyte per second)
- A data-centric approach,
  how much results a system can produce
- per watt
- per cubic foot

Suggestion 1: Datarate (2)

- Financial considerations
  play a major role in computing!
- Unreasonable to include non-transparent and ever negotiated pricing information
- But, cost of computer systems is dictated
  by the cost of the chips:
  - regularity of the design, VLSI process, chip area, volume...
- This issue remains a challenge!

Suggestion 2: Real applications

- Besides generic benchmarks
  rank on results when solving real problems
  - Shows the real and full potential
    of an approach
  - Idea present in popular rankings GPU
- Useful to customers, choose a system
  that's good on a problem similar to their own

Suggestion 2: Real applications (2)

- But, how to choose the applications?
- Idea: examine the Top500, Graph500, or the list of HPC systems with most investment
- Select a number of applications top ranking systems were most used for
- Evolves with science!

Conclusion

- Whenever a paradigm shift happens
  in computer technology,
  computer architecture,
  or computer applications,
  a new approach has to be introduced
- The Exascale era requires this major shift!

Goran Rakocevic, Oskar Mencer, Michael J. Flynn, Mateo Valero, Roman Trobec, Per Stenstrom, Veljko Milutinovic

Benchmarking for SuperComputer Ranking

grakocevic@gmail.com