Maximum Performance Computing at Exascale

Running Exascale applications in 3MW by 2018

Maxeler Technologies is a leader in the market for High Performance Computing solutions.

Applications accelerated on Maxeler's Dataflow Compute Engines have shown speed ups in processing time of 20 to 50x. Thus, a 1U MaxNode enables over 90% reduction in energy usage and can provide the same compute performance as a 20-50x larger conventional system. We provide complete and scalable solutions including hardware platforms, software programming tools and application consulting.

Simply scaling conventional multicore processors will not deliver exascale performance in less than 100MW. Maxeler's answer is Maximum Performance Computing, which achieves exascale performance for specific applications with over 20x lower power consumption.

The Challenges of Exascale

We identify five major challenges that are impeding progress towards exascale computing:

- Power
- Programmability
- Communication
- Space
- Reliability

We address these challenges by having 20 to 50 times fewer accelerated compute nodes, reducing space and power consumption, while increasing reliability.

We define exascale performance as delivering application performance equivalent to a conventional supercomputer with 1 exaflop peak performance.

To illustrate this, we can compare programming a exascale supercomputer built from conventional multi-core processors to one built using Maxeler accelerated compute nodes, to run a large 3D convolution.

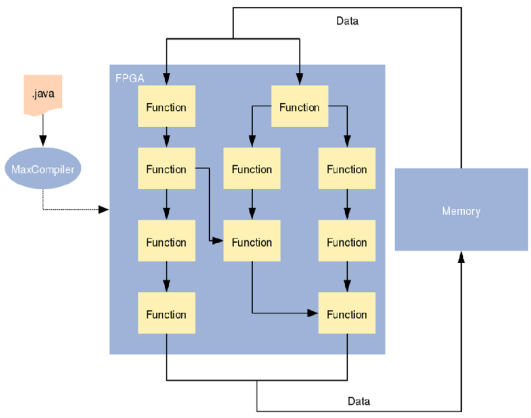

In a Maxeler dataflow compute engine, large volumes of data stream from memory through multiple function elements on chip before the final results are written back to memory

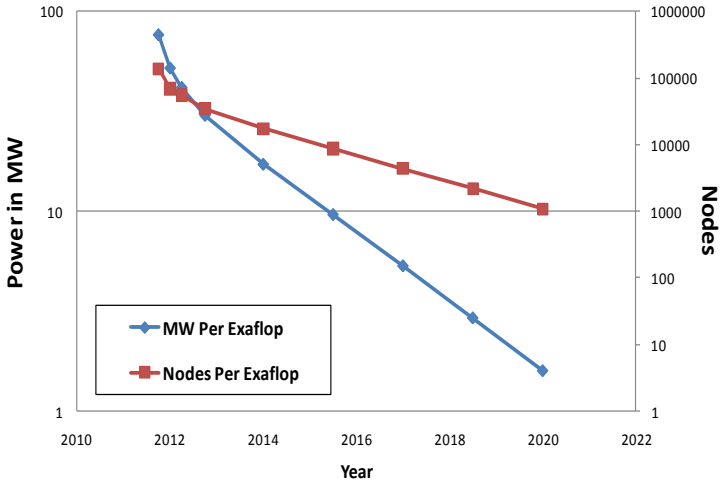

Graph showing how the number of Maxeler nodes and power consumption required to reach Exascale application performance decreases over time.

To achieve 1 exaflop using conventional CPU cores, with 8FLOPS/clock at 2.5GHz, requires 50 million cores, which would currently require over 1000MW of power, and is likely to require at least 100MW in the 2018-19 timeframe.

We define our convolution problem as follows:

- Regular 3D grid, 10,000 points in each dimension = **1 terapoint**
- 13x13x13 convolution stencil

With 50M cores, our 1 terapoint grid gives us only 20k points in each core's subdomain, which means that for every 1 output point calculated, the core must access 2 halo points from other cores. The convolution calculation is therefore no longer localized, and becomes completely limited by the performance of the interconnect.

By 2018 we project that we will be capable of building an exascale supercomputer within 3MW with around 17,500 Dataflow Compute Engines. This means 60 million points per compute engine, a simpler programming environment and a drastically reduced burden of halo communication, which is further mitigated by the use of our low-latency high-bandwidth interconnect between processing elements, MaxRing.

Reliability is significantly improved by having many fewer compute nodes, since the mean-time-to-failure of a single node remains the same. Our custom compute fabric allows us to further increase reliability by implementing application-specific error checking and rollback mechanisms.

For more information visit our website:

http://www.maxeler.com/

MAXELER Technologies
MAXIMUM PERFORMANCE COMPUTING