# Moving from Petaflops to Petadata

MAXELER Technologies

*This is a summary of a paper published in Communications of the ACM, May 2013.*

## About the Authors

Veljko Milutinovic is a professor at the University of Belgrade. Goran Rakocevic is a research assistant at the Mathematical Institute of the Serbian Academy of Sciences and Arts. Per Stenstrom is a professor at Chalmers University. Roman Trobec is an associate professor at the Jozef Stefan Institute, Slovenia. Mateo Valero is the director of the Barcelona Supercomputing Centre. Oskar Mencer is the CEO and Michael J. Flynn is Chairman at Maxeler Technologies and Professor Emeritus at Stanford University.

## Summary

As the supercomputing community sets its sights on breaking the exascale barrier, the very definition of what constitutes our most powerful supercomputers has come into question. This call for reevaluation is a result of observations that for various sophisticated, real world applications some computers outperform the supercomputers on the TOP500 list by orders of magnitude, but are not included on this list. These selected computers do not make the cut because they do not measure up based on FLOPS (Floating-point Operations Per Second). The legitimacy of this measure is being questioned even by a creator of the TOP500 list, calling for a more comprehensive performance measure for ranking with a much wider scope.

The dataflow approach, developed by Maxeler Technologies, is one of the unconventional computer systems which necessitate this reevaluation, such as the MPC-X Node shown in Figure 1. Alternative measures proposed to include unconventional systems include performance per watt, performance per cubic foot and performance per monetary unit.

These alternative implementations, such as the dataflow approach, are becoming considerably more attractive across the supercomputing community as a growing consensus believes breaking the exascale barrier will require a shift towards a hardware/software co-design. The dataflow approach entails implementation of kernels compiled into a dataflow engine, resulting in an array structure that can include hundreds to thousands of pipelined stages. While the clock cycle on dataflow engines (DFEs) is relatively slow, this drawback is overcome through increased efficiency and magnitude of parallelism employed on the DFE. Also, the lower clock speed results in less consumed power, less space, and less money compared to systems using a faster clock.

Examination of the TOP500 list reveals the top is dominated by traditional, control flow systems. It is only natural to expect these systems to deliver the highest performance, but that is not always true. In the case of a data-intensive petroleum application, a dataflow system offers a speedup close to 200 times. Yet, a custom hardware implementation of LINPACK only offers a speedup of 2-6 times.

Considering these findings one can make the conclusion that FLOPS count does not sufficiently cover all aspects of HPC systems. Other proposed ideas include data generation rate per second, per watt, or per cubic foot for a particular algorithm and data size. Another compelling argument for the justification of expanding the FLOPS count is that floating-point arithmetic no longer dominates the execution time. Additionally, LINPACK, the workload used to create the TOP500 supercomputers list, is not a sufficient predictor of performance. LINPACK is primarily focused on control flow computing rather than a more data-centric view relevant to real world applications. This paper proposes a benchmark based on the performance of a system solving real-life problems that the top ranking systems have been most commonly used for. Overall, we need a measure that most effectively translates into usable performance.

*Figure 1: The MPC-X Node which serves as a building block for the Maxeler MaxRacks.*