

High Performance Communication

Tom Vander Aa, imec ExaScience Life Lab



Communication is key

In the late 90s, I was teaching parallel programming in C using MPI to students. The most important lesson I wanted them to remember is that communication is much more important than computation. The form of the benchmark couldn't be more common: a set of convolutional filters applied to an image, one filter after the other in a pipelined fashion.

The task of the students was to make a parallel version of the program run as fast as possible on a gorgeous 4 node Parsytec x'plorer cluster with 16 Transputer processors in total. The most obvious solution was to chop to image up in pieces and process each piece on a different processor. The trick was to choose your input images large enough or otherwise the parallel version would be communication-dominated and you would not get any speedup

Exascale computers bring exascale problems

After 15 years in low-energy embedded systems, I am now back in HPC, and some things haven't changed. HPC is still mostly about programming in Fortran/C/C++ using MPI, with OpenMP bolted on for node-level parallelism. Avoiding communication is still the aim of the game, now more than ever, and that's because other things did change. I've gone from 16 processors with limited internal parallelism to clusters with thousands of nodes where each node contains tens of cores with wide vector parallelism on each core. More compute power and more parallelism means more bits are computed further apart. Getting all these bits to talk to each other using C+MPI+OpenMP is for sure a recipe for disaster; similar to what my students discovered 15 years ago, only much worse.

These trends will continue towards Exascale.

No silver bullet

There isn't a single-point solution to solve these problems. Instead, we need a many-point solution, one that covers all abstraction levels and makes sure there is parallelism extracted and efficiently used at each abstraction level. We need algorithms that do more computation and less communication, or at least communication that is less blocking, to accomplish a certain task. We need flexible programming models that make it easy to map our algorithms on the wide variety of architectures and to efficiently use the wide variety of types of parallelism. That's where EXA2CT steps in.

EXA2CT-ly what we need :-)

The EXA2CT project aims to integrate development of algorithms and programming models tuned to future Exascale supercomputer architectures.

We will show how an optimized combination of both aspects is needed to benefit from the massive amount of compute power these exascale computers will provide. This will be demonstrated by providing open source proto-applications, which will help to boot-strap the creation of genuine exascale codes.

The EXA2CT project (www.exa2ct.eu [1]), comprising 10 partners from six different countries and funded by the European commission, started in September 2013. We are eager to chat about our first results with you at ISC'14 in Leipzig.

EXA2CT is part of a network about research on Exascale in Europe. Meet EXA2CT and the other partners of the network at our joint BoF session "[Exascale Research: The European Approach](#) [2]" on Tuesday, June 24, 2014, 2.15pm - 3.15pm in Hall 5. We are also happy to welcome you at our booth 833.

Tom Vander Aa is a researcher/project coordinator in the ExaScience Life Lab at imec. This lab creates new supercomputer solutions to generate breakthroughs in life sciences and biotechnology. Before joining the Exascience Lab he was at Target Compiler Technologies and at the architecture and compiler group in imec working on low energy high performance architectures and compilation techniques. In 2005, Tom obtained a PhD in electrical engineering from KULeuven, Leuven, Belgium on energy optimization for instruction memory of embedded processors. He has a master in computer science degree, also from KULeuven and is a member of the IEEE.

This blog was originally published on [ISC Events](#) [3] and is reproduced here with permission.

Source URL (retrieved on 01/29/2015 - 11:27am):

<http://www.scientificcomputing.com/blogs/2014/06/high-performance->

[communication](#)

Links:

[1] <http://www.exa2ct.eu/>

[2] http://www.isc-events.com/isc14_ap/sessiondetails.htm?t=session&o=108&a=select&ra=search

[3] <http://www.isc-events.com/isc14/blogs.html>