

Rapid Performance Assessment



You might have heard about the major speed-ups GPUs and FPGAs have promised, but also about the fact that this speed-up will depend a lot on the type of software/algorithm. Investing in OpenCL or CUDA can therefore feel risky, since going in costs time and money, while keeping out can potentially give too much space to the competition. But if you want your customers to get the best experience without paying an unnecessary high price, you'll need to know what the return of your investment could be. With this quick assessment we will help you determine exactly that.

What we've done before

Most assessments were on answering the question "How much speed-up can I get using GPUs?". Other questions were:

- Does this algorithm work on this specific mobile processor?
- Can we better use CUDA, OpenCL or OpenGL shaders for this algorithm?
- Does the HPC code run best on a Tesla K40 or FirePro S9150?
- How many weeks/months would it take to port all code?
- How many GPUs do I need for under 1 second responses?
- Does this code port to an FPGA?
 - Which OpenCL device best suites by algorithm: CPU, GPU, APU, DSP, FPGA or something else?

Is your question in the list?

Program

Within a week we can fully analyse your code, or two weeks if the codebase is large or complex. During the assessment we write/port/optimize code, to be able to support our conclusions with numbers.

After the assessment you get an overview of the hotspots, an indication of total speed-up when using OpenCL (or comparable technology), and the answers to your questions.

Preparations

Send a mail to contact@streamhpc.com for more information, and we'll call you back to talk about your requirements. Please provide times when you want to be called back.