Subject: Re: Experimental CUDA support
Posted by mirek on Fri, 14 Feb 2025 19:02:36 GMT
View Forum Message <> Reply to Message

Tom1 wrote on Fri, 14 February 2025 15:16Hi,

Very interesting development!

Maybe I should get some NVIDIA hardware to test this.

I wonder if Intel oneAPI could be supported too in a similar way...? That would offer cross-platform compatibility for Windows, Linux (and gradually MacOS) for hardware from NVIDIA, AMD and Intel. Therefore, oneAPI looks like a good fit for u++, but of course I don't know the internals.

Best regards,

Tom

Uhm, too soon to tell. Quick glance at oneAPI and the SYCL authors does not bring too much confidence.

Originally I tried with Vulkan / GLSL, but that is really unusable - I am not bothered with 500+ lines to run simple kernel, but GLSL is unusable language for anything complex. It is like C without pointers... Meanwhile, CUDA is for years full blown C++ with hardly any limitations. Impossible for oneAPI / SYCL to catch up with something incompatible IMO.

Amd HIP on the other hand - that one looks promising.

Right now I am a bit unhappy that CUDA requires Visual C++ compiler, but maybe a bit of effort I might be able to persuade it to use mingw-clang instead...