inverted CERN School of Computing 2016

Contribution ID: **8**

Type: **not specified**

III - Template Metaprogramming for massively parallel scientific computing - Templates for Iteration; Thread-level Parallelism

Wednesday 2 March 2016 11:30 (1 hour)

Large scale scientific computing raises questions on different levels ranging from the fomulation of the problems to the choice of the best algorithms and their implementation for a specific platform. There are similarities in these different topics that can be exploited by modern-style C++ template metaprogramming techniques to produce readable, maintainable and generic code. Traditional low-level code tend to be fast but platform-dependent, and it obfuscates the meaning of the algorithm. On the other hand, object-oriented approach is nice to read, but may come with an inherent performance penalty.

These lectures aim to present he basics of the Expression Template (ET) idiom which allows us to keep the object-oriented approach without sacrificing performance. We will in particular show to to enhance ET to include SIMD vectorization. We will then introduce techniques for abstracting iteration, and introduce thread-level parallelism for use in heavy data-centric loads. We will show to to apply these methods in a way which keeps the "front end" code very readable.

LECTURE 3
In this lecture, we will look into a specific technique to parallelize a large data-centric workload iterating over a multi-dimensional array. We will show how to separate iteration and computation and how the "front-end" algorithm can then be made independent on the dimensionality, coordinate system, or order of numerical approximation. We will show how this separation further helps to implement thread-level parallelism into the "back-end" and explore some common cases of data dependency. We will finally take a look at an example code combining the ideas of all three lectures.

Presenter: VYSKOČIL, Jiří (Czech Technical University in Prague)