



Contribution ID: 34

Type: not specified

## Parallelized Kalman-Filter-Based Reconstruction of Particle Tracks on Many-Core Processors and GPUs

*Wednesday, March 8, 2017 5:15 PM (30 minutes)*

For over a decade now, physical and energy constraints have limited clock speed improvements in commodity microprocessors. Instead, chipmakers have been pushed into producing lower-power, multi-core processors such as GPGPU, ARM and Intel MIC. Broad-based efforts from manufacturers and developers have been devoted to making these processors user-friendly enough to perform general computations. However, extracting performance from a larger number of cores, as well as specialized vector or SIMD units, requires special care in algorithm design and code optimization.

One of the most computationally challenging problems in high-energy particle experiments is finding and fitting the charged-particle tracks during event reconstruction. This is expected to become by far the dominant problem in the High-Luminosity Large Hadron Collider (HL-LHC), for example. Today the most common track finding methods are those based on the Kalman filter. Experience with Kalman techniques on real tracking detector systems has shown that they are robust and provide high physics performance. This is why they are currently in use at the LHC, both in the trigger and offline.

Previously we reported on the significant parallel speedups that resulted from our investigations to adapt Kalman filters to track fitting and track building on Intel Xeon and Xeon Phi. We continue to make progress toward the understanding of these processors while progressively introducing more realistic physics. These processors, in particular Xeon Phi, provided a good foundation for porting these algorithms to NVIDIA GPUs, for which parallelization and vectorization is of utmost importance. The challenge lies mostly in the ability to feed these graphical devices with enough data to keep them busy. We also discuss strategies for minimizing code duplication while still being able to keep the previously cited algorithms as close to the hardware as possible.

**Primary authors:** Dr CERATI, Giuseppe (Fermilab (US)); Dr ELMER, Peter (Princeton University (US)); Dr KRUTELYOV, Slava (Univ. of California San Diego (US)); Dr LANTZ, Steven R (Cornell University (US)); LEFEBVRE, Matthieu (Princeton University (US)); Mr MCDERMOTT, Kevin (Cornell University (US)); Dr RILEY, Daniel Sherman (Cornell University (US)); Dr TADEL, Matevz (Univ. of California San Diego (US)); Prof. WITTICH, Peter (Cornell University (US)); Prof. WURTHWEIN, Frank (Univ. of California San Diego (US)); Prof. YAGIL, Avi (Univ. of California San Diego (US))

**Presenter:** LEFEBVRE, Matthieu (Princeton University (US))

**Track Classification:** 1: Parallel and discrete pattern recognition