

Kalman Filter Tracking on Parallel Architectures

Giuseppe **Cerati**^{1,a}, Peter **Elmer**^{2,b}, Slava **Krutelyov**^{1,c}, Steven **Lantz**^{3,d}, Matthieu **Lefebvre**^{2,e}, Kevin **McDermott**^{3,f}, Daniel **Riley**^{3,g}, Matevž **Tadel**^{1,h}, Peter **Wittich**^{3,i}, Frank **Würthwein**^{1,j}, and Avi **Yagil**^{1,k}

¹UC San Diego, La Jolla, CA, USA

²Princeton University, Princeton, NJ, USA

³Cornell University, Ithaca, NY, USA

Abstract. Power density constraints are limiting the performance improvements of modern CPUs. To address this we have seen the introduction of lower-power, multi-core processors such as GPGPU, ARM and Intel MIC. In order to achieve the theoretical performance gains of these processors, it will be necessary to parallelize algorithms to exploit larger numbers of lightweight cores and specialized functions like large vector units. Track finding and fitting is one of the most computationally challenging problems for event reconstruction in particle physics. At the High-Luminosity Large Hadron Collider (HL-LHC), for example, this will be by far the dominant problem. The need for greater parallelism has driven investigations of very different track finding techniques such as Cellular Automata or Hough Transforms. The most common track finding techniques in use today, however, are those based on a Kalman filter approach. Significant experience has been accumulated with these techniques on real tracking detector systems, both in the trigger and offline. They are known to provide high physics performance, are robust, and are in use today at the LHC. Given the utility of the Kalman filter in track finding, we have begun to port these algorithms to parallel architectures, namely Intel Xeon and Xeon Phi. We report here on our progress towards an end-to-end track reconstruction algorithm fully exploiting vectorization and parallelization techniques in a simplified experimental environment.

^ae-mail: giuseppe.cerati@cern.ch

^be-mail: peter.elmer@cern.ch

^ce-mail: vyacheslav.krutelyov@cern.ch

^de-mail: steve.lantz@cornell.edu

^ee-mail: ml15@princeton.edu

^fe-mail: kevin.mcdermott@cern.ch

^ge-mail: daniel.riley@cornell.edu

^he-mail: matevz.tadel@cern.ch

ⁱe-mail: wittich@cornell.edu

^je-mail: fkw@ucsd.edu

^ke-mail: ayagil@physics.ucsd.edu

1 Introduction

The Large Hadron Collider (LHC) at CERN is the highest energy collider ever constructed. It consists of two counter-circulating proton beams made to collide in four locations around a 27 kilometer ring straddling the border between Switzerland and France. By several means it is the largest man-made scientific device on the planet. The goal of the LHC is to probe the basic building blocks of matter and their interactions. In 2012, the Higgs boson was discovered by the CMS and ATLAS collaborations [1, 2]. By measuring the energy and momentum of particles escaping the collision, we can infer the existence of massive particles that were created and measure those massive particles' properties. This process is known as event reconstruction and consists of integrating information from different detector components. Track reconstruction, also known as tracking, is one step in event reconstruction and determines the trajectories of charged particles ("tracks") from a set of positions of energy deposits within the various layers of our detector ("hits"). Tracking, in comparison to other aspects of event reconstruction, is the most computationally time consuming, the most sensitive to increased activity in the detector, and traditionally, the least amenable to parallelized processing. The speed of online reconstruction has a direct impact on how much data can be stored from the 40 MHz collisions rate, while the speed on the offline reconstruction limits how much data can be processed for physics analyses. This research is aimed at vastly speeding up tracking.

The large time spent in tracking will become even more important in the HL-LHC era of the Large Hadron Collider. The increase in event rate will lead to an increase in detector occupancy ("pile-up", PU), leading to an exponential gain in time taken to perform track reconstruction, as can be seen in Figure 1 [3]. In the Figure, PU25 corresponds to the data taken during 2012, and PU140 corresponds to the low end of estimates for the HL-LHC era. Clearly this research on tracking performance will become increasingly important during this era.

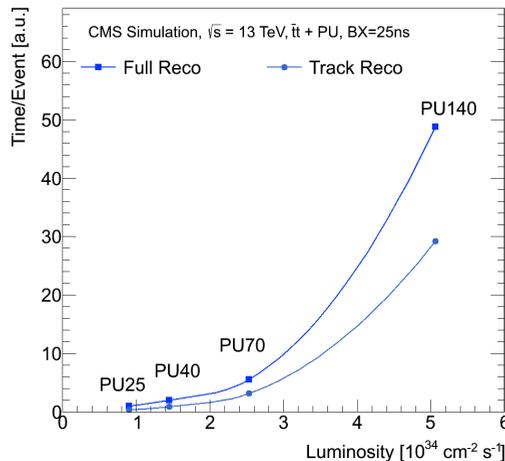


Figure 1. CPU time per event versus instantaneous luminosity, for both full reconstruction and the dominant tracking portion. PU25 corresponds to the data taken during 2012, and PU140 corresponds to the HL-LHC era. The time of the reconstruction is dominated by track reconstruction [3].

A correlated issue to speeding up tracking is the change in computing architectures in the last decade. Around 2005, the computing processor market reached a turning point: power density lim-

itations in chips ended the long-time trend of ever-increasing clock speeds, and our applications no longer immediately run exponentially faster on subsequent generations of processors. This is true even though the underlying transistor count continues to increase per Moore's law. Increases in processing speed such as those required by the time increases in Figure 1 will no longer come 'for free' from faster processors. New, parallel processors instead are aggregates of 'small cores' that in total still show large performance gains from generation to generation, but their usage requires a re-work of our mostly serial software to exploit these gains. The processors in question include ARM, GPGPU and the Intel Xeon and Xeon Phi; in this work we target the Xeon and Xeon Phi architectures.

2 Kalman filter tracking

The algorithm we are targeting for parallelized processing is a Kalman filter (KF) based algorithm [4]. KF-based tracking algorithms are widely used since they are fast and incorporate estimates of multiple scattering off massive detectors directly into the trajectory of the particle. In addition, a combinatorial KF provides an elegant way to resolve ambiguities when considering multiple hits to add to a track [5]. Other track finding algorithms, more naturally suited to parallelization and coming from the image processing community, have been explored by others. These include Hough Transforms and Cellular Automata, among others (see, for instance, [6]). However, these are not the main algorithms in use at the LHC today, whereas there is extensive understanding on how KF algorithms perform. KF algorithms have proven to be robust and perform well in the difficult experimental environment of the LHC [3]. Rather than abandon the collective knowledge gained from KF algorithms, we wish to extend this robust tool by porting KF tracking to parallel architectures.

KF tracking proceeds in three main stages: seeding, building, and fitting. Seeding provides the initial estimate of the track parameters based on a few hits in a subset of layers. Realistic seeding is currently under development and will not be reported here. Building then collects additional hits in other detector layers to form a complete track, using the combinatorial KF as a basis for predicting which hits to consider and keep. Track building is by far the most time consuming step of tracking, as it requires branching to explore potentially more than one candidate track per seed after finding compatible hits on a given layer. After hits have been assigned to tracks, a final fit with a KF can be performed over each track to provide the best estimate of each tracks' parameters.

To realize performance gains, we need to exploit two types of parallelism: vectorization and parallelization. Vectorization aims to perform a single instruction on multiple data at the same time, i.e. performing the same operation across different data in lockstep. KF tracking cannot be ported in a straightforward way to run in parallel on many-core processors for the following reasons. In building, as explained above, we can branch to explore multiple candidates per seed, thus destroying the lock-step synchronization required for vectorization performance. Parallelization aims to perform different instructions at the same time on different data. The challenge to tracking then is workload balancing across different threads, as track occupancy in a detector is not uniformly distributed on a per event basis. Past work by our group has shown progress in porting sub-stages of KF tracking to support parallelism in simplified detectors (see, e.g. our presentations at ACAT2014 [7], CHEP2015 [8], and NSS-MIC2015 [9]). As the hit collection is completely determined after track building, track fitting can repeatedly apply the KF algorithm without branching, making this the ideal place to start in porting KF tracking to Xeon and Xeon Phi, with our first results shown at ACAT2014 [7].

2.1 Optimized matrix library MATRIPLEX

The computational problem of KF-based tracking consists of a sequence of matrix operations. Unlike [4], we parameterize the measurement state and the track state in global coordinates with matrices of

size $N \times N = 3 \times 3$ and $N \times N = 6 \times 6$, respectively. To allow maximum flexibility for exploring SIMD operations on small-dimensional matrices, and to decouple the specific computations from the high level algorithm, we have developed a new matrix library, **MATRIPLEX**. The **MATRIPLEX** memory layout is optimized for the loading of vector units for SIMD operations on a set of matrices. **MATRIPLEX** includes a code generator for defining optimized matrix operations, including support for symmetric matrices and on-the-fly matrix transposition. Patterns of elements which are known by construction to be zero or one can be specified, and the resulting generated code will be optimized accordingly to reduce unnecessary register loads and arithmetic operations. The generated code can be either standard C++ or simple intrinsic macros that can be easily mapped to architecture-specific intrinsic functions.

3 Track building

Combinatorial KF track building shares the same core KF calculations as track fitting, but has two major complications for utilizing vectorizable operations compared to track fitting. The first such problem is the “nHits” problem: when performing the KF operations, the hit set for a given candidate track in building is undefined and one potentially has to choose between $O(10^4)$ hits per layer. The second problem is the “combinatorial” problem: when more than one hit on a given layer is compatible with a given input candidate, branching occurs to explore all possible candidate tracks up to some cutoff number of candidates per seed.

The key to reducing the nHits problem is to partition the tracks and hits spatially in order to reduce the search space for a given track. We partition space without reference to the detector structure, placing hits in each layer into directionally projective bins. We first define η bins, where $\eta = -\log[\tan^{-1}(\theta/2)]$, θ being the polar angle with respect to the beam line. The η bins are self-consistent; because in our detector, tracks do not bend in η . The bins of tracks are staggered with respect to bins of hits, so that tracks never search for hits outside two overlapping hit bins when using the combinatorial KF. In this way, the number of η bins for tracks will always be an odd number. These bins provide simple boundaries for thread associations using the OpenMP parallelization library [10]. We begin with 21 η bins and distribute threads equally across the bins, such that the number of threads per η bin is either a multiple or a factor of the number of bins. We also define ϕ bins, where ϕ is the azimuthal angle, perpendicular to the beam line. As the hits are sorted in ϕ for every η bin and layer, the ϕ bins simply provide a fast look-up of compatible hits in a layer for a given track given the estimate from the KF.

With regards to the combinatorial problem, we first developed track building to add only the best hit out of all compatible hits on a layer for each track candidate. By definition, then, each seed only produces one track per layer and does not require copying of track candidates to explore multiple hits per layer, as in the full combinatorial KF approach. The best hit is defined as the hit that produced the lowest χ^2 increment to the track candidate. The vectorization and parallelization performance of this “best hit” approach were presented at CHEP2015 [8]. After demonstrating the feasibility of track building in the best hit approach, we then allowed for more than one track candidate per seed per layer to be produced, as in the full combinatorial KF approach.

3.1 Memory management

With the full combinatorial approach in place, we performed extensive studies of the performance of our software, in terms of both the physics performance and the code performance. For the latter, we used the Intel VTune 2016 [11] suite of tools to identify bottlenecks and understand the impacts of

our optimization attempts. In particular, as can be expected, we determined that memory management is of critical importance. For this effect we describe below several studies to optimize memory performance, and discuss the results of these studies.

The first such impact on performance arose out significant time being spent loading data into our local caches. By using vectorizable intrinsic functions to perform copying, we saw 20% speedup. We also noticed a significant amount of time was being spent resizing the C++ vector object that stored the indices of the hits associated to a given track, and simply reserving the memory did not help. By moving to a statically sized array, the resizing was avoided, yielding a 45% speedup. The size of the data structures used in our algorithm has a crucial impact on the timing performance. In particular, the data structures for the hits and tracks must be optimized in size to fit into the fastest cache memory. We were able to reduce the size of the track objects by 20% and the hit object by 40%, leading to a 30% speedup. There were additional objects used in the algorithm that were being instantiated after each event that could instead be reset and recycled, which led to a 25% speedup. It is important to note that each of these improvements lead to individual, multiplicative speedups.

Lastly, to mitigate the impact from serial work in the copying of track candidates after finding compatible hits on each in the full combinatorial approach, we moved copying outside of all vectorizable operations in what we term the “clone engine”. The clone engine approach only copies the best candidate tracks up to N candidates per seed after reading and sorting a bookkeeping list of all compatible hits. This is in contrast to a first attempt at the combinatorial KF which copied a candidate after each time a hit was deemed compatible, and sorting and keeping only the best N candidates per seed after all the possible hits on a given layer for all input candidates were explored. A more detailed discussion of this work on memory management and impacts on performance was presented at NSS-MIC2015 [9].

3.2 Results

We present here the latest vectorization and parallelization benchmarks in track building since NSS-MIC2015 given all of the developments described previously. Figure 2 contains two plots displaying the vectorization performance of three different track building approaches on Xeon, shown as functions of the number of vector units enabled. The left plot displays the absolute time in seconds for building tracks for events with 20,000 tracks as a function of the code parameter that sets the effective width of the vector unit in floats. The plot on the right shows the speedup in time compared to the time from the first data point as a function of this parameter. Only one thread is enabled in these tests.

The blue curve is the “best hit” approach described previously. Naturally, this approach will have the lowest absolute time in comparison to the two fully combinatorial approaches, the black and pink curves. The black curve is the original approach to combinatorial track finding where the copying of the track candidates is inside the vectorizable KF operations. The pink curve shows the clone engine approach, which moves the serial work of copying outside of the KF vectorizable operations. As expected, the clone engine approach has a lower absolute time and higher speedup than the original. It is important to note that while the best hit approach is the fastest, the physics performance has inefficiencies in hit finding and track finding from not being fully able to explore multiple track candidates per seed. Even in our simplified model, this behavior is already apparent, and the best hit approach is expected to become even more inefficient as we add in more realistic detector effects.

Figure 3 shows the same vectorization plots as Figure 2, now with Xeon Phi, which has AVX-512 vector registers that are twice as wide as those in Xeon. The original combinatorial approach is not shown, as it was apparent from Figure 2 that it did not perform as well as the clone engine approach. With one thread enabled, Xeon Phi sees the same $1.5\times$ to $2\times$ speedup as Xeon. There is a noticeable

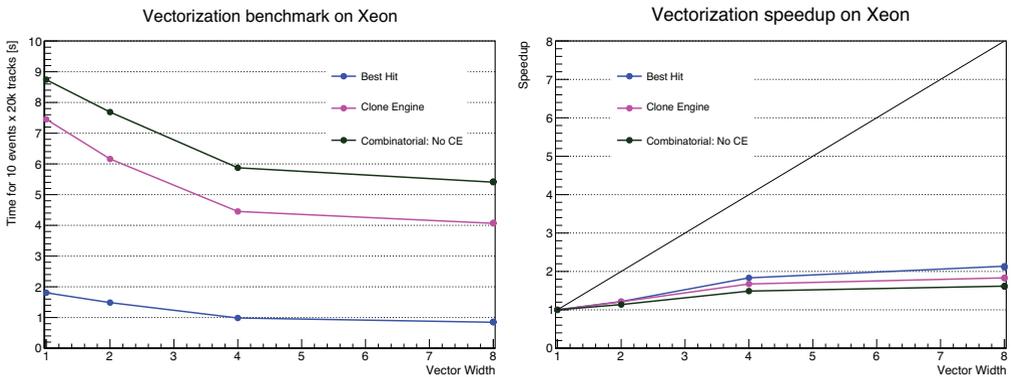


Figure 2. Left: Absolute time as a function of the code parameter that sets the effective vector width. Right: Speedup as a function of this parameter on Xeon. Both plots show results for three different track building approaches. The blue curve is the no combinatorial tracking approach. The black and pink curves use the full combinatorial approach in track building. The pink curve moves the copying of tracks outside of the vectorizable operations. Only one thread is enabled.

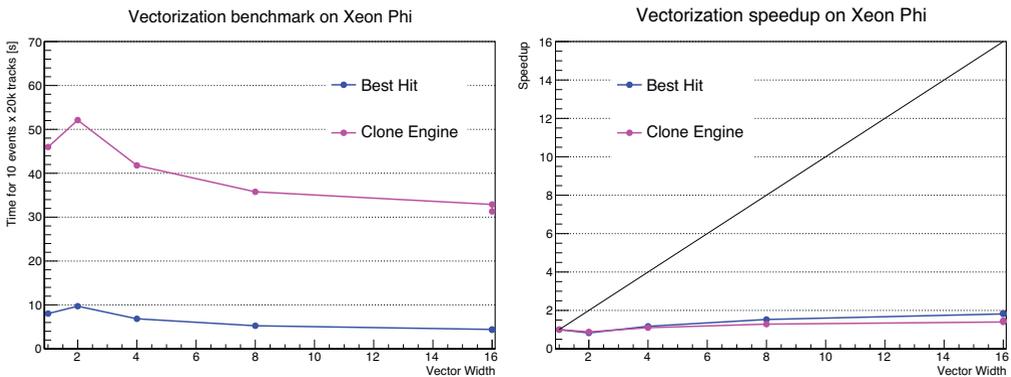


Figure 3. Left: Absolute time as a function of the code parameter that sets the effective vector width on Xeon Phi. Right: Speedup as a function of this parameter on Xeon Phi. Both plots show results for two different track building approaches. The blue curve is the no combinatorial tracking approach. The pink curve uses the full combinatorial approach in track building and moves the copying of tracks outside of the vectorizable operations. Only one thread is enabled.

increase in time after the vector unit size is first doubled. This is likely overhead from enabling the vector units. As seen on Xeon, the best hit approach still has the lowest absolute timing.

Figure 4 displays the parallelization performance in terms of the absolute time and speedup of the clone engine approach on Xeon as a function of the number of threads enabled using OpenMP v. 4.0. The best hit approach is removed from these plots, because in reality, the full combinatorial approach is one we will ultimately use. The Xeon machine we are using has 12 physical cores which appear to be 24 logical cores to the OS due to hyperthreading being enabled. As described previously in Section 3, we enable threads as multiples or factors of the number of η bins. Since we start with 21 η

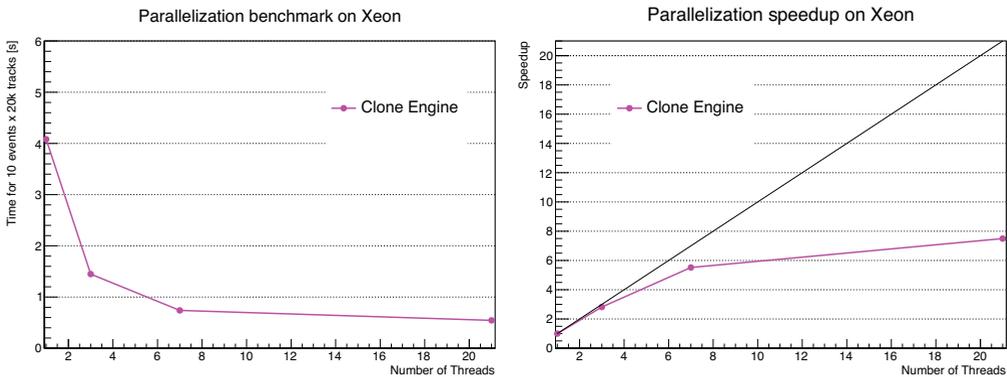


Figure 4. Left: Absolute time as a function of the number of threads enabled on Xeon. Right: Speedup as a function of the number of threads enabled on Xeon. Both plots show results for the full combinatorial approach in track building that moves the copying of tracks outside of the vectorizable operations. The full Xeon vector width (8 floats) is assumed by the code.

bins, we cannot explore the critical boundary of 12 threads with OpenMP. The slight deviation away from ideal scaling before hyperthreading arises from the fact that some threads are unlucky and end up with more work per event than the rest, stalling the other threads, which was shown with VTune. The large deviation away from ideal scaling after enabling hyperthreading is due to resource limitations: the two threads per core are contending for the same instruction pipelines and data caches. Even so, nearly 8 \times speedup is seen for 21 threads.

Figure 5 displays the same parallelization performance plots as Figure 4, now in terms of the Xeon Phi, with the full vector width of 16 floats enabled using OpenMP v. 4.0. The Xeon Phi we are using has 61 physical cores requiring 122 independent instruction streams for full utilization, due to the fact that the Xeon Phi issues instructions for a given thread every other clock cycle. Therefore, to keep a given physical core busy every clock cycle, one has to schedule two threads per core alternating in instruction execution. A form of hyperthreading is also present on Xeon Phi, yielding a total of 244 hardware threads (logical cores). Thus we notice two features in the speedup curve for Xeon Phi, occurring just past the physical core count (61) and twice that number (122). However, this loss in speedup is eventually recovered, and a factor of 30 \times speedup is observed with 210 threads enabled.

4 Conclusion and outlook

We have made significant progress in parallelized and vectorized Kalman filter-based end-to-end tracking R&D on Xeon and Xeon Phi architectures. Through the use of a wide toolkit which includes Intel VTune, we have developed a good understanding of bottlenecks and limitations of our implementation which has led to further improvements. We are beginning to explore Intel Threaded Building Blocks (TBB), which allows for more dynamic scheduling to prevent stalls in resource usage. In this way, threads can be spawned in numbers more naturally suited to the number of cores, unlike OpenMP. The project has seen promising initial results; however, much work remains.

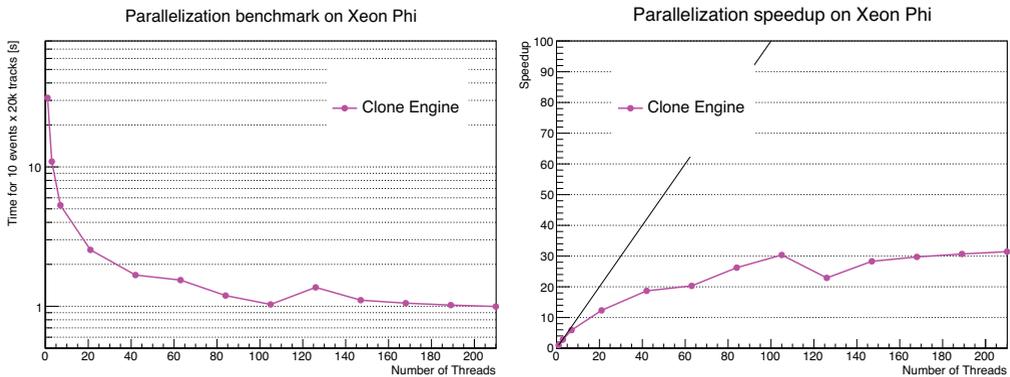


Figure 5. Left: Absolute time as a function of the number of threads enabled on Xeon Phi. Right: Speedup as a function of the number of threads enabled on Xeon Phi. Both plots show results for the full combinatorial approach in track building that moves the copying of tracks outside of the vectorizable operations. The full Xeon Phi vector width (16 floats) is assumed by the code.

Acknowledgment

This work is supported by the U.S. National Science Foundation, under the grants PHY-1520969, PHY-1521042, PHY-1520942 and PHY-1120138.

References

- [1] CMS Collaboration, Phys. Lett. **B716**, 30 (2012)
- [2] ATLAS Collaboration, Phys. Lett. **B716**, 1 (2012)
- [3] G. Cerati, Proceedings of Science **VERTEX2014**, 037 (2015), http://pos.sissa.it/archive/conferences/227/037/Vertex2014_037.pdf
- [4] R. Frühwirth, Nucl. Inst. & Meth. A **262**, 444 (1987)
- [5] P. Billoir, Comp. Phys. Comm. **225**, 352 (1984)
- [6] V. Haylo, P. LeGresley, P. Lujan, V. Karpusenkov, and A. Vladimirov, JINST **9**, P04005 (2014)
- [7] G. Cerati et al., J. Phys. Conf. Ser. **608**, 012057 (2014)
- [8] G. Cerati et al., J. Phys. Conf. Ser. **664**, 072008 (2015)
- [9] G. Cerati et al., *Kalman-filter-based particle tracking on parallel architectures at hadron colliders* (2015), [arXiv:1601.08245](https://arxiv.org/abs/1601.08245)
- [10] OpenMP Architecture Review Board, *Openmp application program interface version 4.0* (2013), <http://www.openmp.org/mp-documents/OpenMP4.0.0.pdf>
- [11] See <https://software.intel.com/en-us/intel-vtune-amplifier-xe>, accessed 23 March 2016