

Segment Linking: A Highly Parallelizable Track Reconstruction Algorithm for HL-LHC

P Chang¹, P Elmer², Y Gu¹, V Krutelyov¹, G Niendorf³, M Reid³,
 B V Sathia Narayanan¹, M Tadel¹, E Vourliotis^{1a}, B Wang², P
 Wittich³, A Yagil¹

¹ University of California San Diego, CA, US

² Princeton University, NJ, US

³ Cornell University, NY, US

E-mail: ^a emmanouil.vourliotis@cern.ch

Abstract. The High Luminosity upgrade of the Large Hadron Collider (HL-LHC) will produce particle collisions with up to 200 simultaneous proton-proton interactions. These unprecedented conditions will create a combinatorial complexity for charged-particle track reconstruction that demands a computational cost that is expected to surpass the projected computing budget using conventional CPUs. Motivated by this and taking into account the prevalence of heterogeneous computing in cutting-edge High Performance Computing centers, we propose an efficient, fast and highly parallelizable bottom-up approach to track reconstruction for the HL-LHC, along with an associated implementation on GPUs, in the context of the Phase 2 CMS outer tracker. Our algorithm, called Segment Linking (or Line Segment Tracking), takes advantage of localized track stub creation, combining individual stubs to progressively form higher level objects that are subject to kinematical and geometrical requirements compatible with genuine physics tracks. The local nature of the algorithm makes it ideal for parallelization under the Single Instruction, Multiple Data paradigm, as hundreds of objects can be built simultaneously. The computing and physics performance of the algorithm has been tested on an NVIDIA Tesla V100 GPU, already yielding efficiency and timing measurements that are on par with the latest, multi-CPU versions of existing CMS tracking algorithms.

1. Introduction and Motivation

The Large Hadron Collider (LHC) of CERN is currently the largest accelerator in the world. In order to take advantage of the extensive infrastructure for its operation for as long as possible, a major upgrade to the LHC is planned for 2027 (Phase 2). This upgrade, called the High Luminosity LHC (HL-LHC), aims at collecting more than 3000 fb^{-1} of integrated luminosity of proton-proton (p-p) collision data, increasing the corresponding nominal value of the LHC by a factor of ten [1]. To achieve such high values of integrated luminosity, the instantaneous luminosity, i.e. rate of collisions, is pushed to its limits. This leads to an increased pileup (PU), i.e. number of simultaneous p-p interactions per bunch crossing. The average PU during the full LHC operation (including Run 3) is expected to stay under 70, while the same number for the HL-LHC will rise to around 200.

The much harsher PU conditions anticipated in the HL-LHC lead to an escalation of the complexity of event reconstruction. Especially for the reconstruction of charged-particle tracks, which relies on the combination of multiple hits and is an inherently combinatorial



Content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

problem, several times more PU results in unprecedented tracker occupancies and computational complexity. As shown in reference [2], this translates to an exponential growth of time needed for the track reconstruction. To mitigate this effect, increased computational resources would be required, leading to highly elevated costs of operation.

Despite the projected increase in the processing power of central processing units (CPUs), which dominated the computational resources of Run 1 and 2 of the LHC, the computational complexity of event reconstruction at the HL-LHC is expected to demand a computational cost that exceeds the computing budget using conventional single-thread CPUs programming. As a result, alternative approaches need to be explored. A solution to this problem is the efficient parallelization of existing reconstruction algorithms, where this is feasible, or the creation of new algorithms that can benefit from parallelism. The concept of parallelization can be combined with the modern paradigm of heterogeneous computing, which aims at utilizing a large variety of processing units, such as graphics processing units (GPUs) or field-programmable gate arrays (FPGAs).

This paradigm shift is already implemented in Run 3 related applications in experiments of the LHC. Taking the example of the Compact Muon Solenoid (CMS) experiment [3], the pixel track and vertex reconstruction (Patatrack [4]), the outer tracker strip local reconstruction, and the electromagnetic calorimeter and hadronic calorimeter reconstruction algorithms have been ported to GPUs. The combination of CPU and GPU usage leads to $\sim 25\%$ timing reductions for the total Run 3 CMS High Level Trigger (HLT) reconstruction. The HL-LHC cost projections show that offloading 50% (80%) of the computational work to GPUs decreases the HLT processing farm cost by 35% (75%) [5].

Moving in this direction, this work presents a highly parallelizable algorithm for track reconstruction at the HL-LHC, called *Segment Linking* or *Line Segment Tracking (LST)*. The algorithm has been developed to benefit from heterogeneous computing architectures, more specifically from GPU performance, and a relevant implementation of it on an NVIDIA Tesla V100 GPU is described.

2. The Segment Linking Algorithm

The track reconstruction at the LHC is usually based on Kalman filter methods [6], which are inherently sequential in nature. This has repercussions in terms of computational time for their application. For example, for the CMS experiment, track reconstruction used to require almost 60% of the time needed for the whole event reconstruction procedure at 50 PU. There has been a recent effort towards the parallelization of Kalman filter approaches (*mkFit* project [7]) that improves the timing performance with comparable physics results. Nevertheless, there is a clear need for track reconstruction algorithms with architectures specifically designed for parallelization.

LST is an algorithm developed with the concept of parallelization built into it. The algorithm relies on local hits in the tracker to start producing short tracks. Given the local nature of these short tracks, the algorithm utilizes the Single Instruction, Multiple Data (SIMD) paradigm, where all of the short tracks can be created concurrently with the same set of instructions, since they are independent from one another. Progressively, the short tracks are linked together to form longer tracks, leading to collections of objects with different characteristics. Finally, objects from these collections are combined to create a Track Candidate collection with a high efficiency and a low fake rate. LST is inspired by the XFT algorithm in the Collider Detector at Fermilab (CDF) at the Tevatron [8] and its prototype was presented at the International Conference of High Energy Physics (ICHEP) in 2016 [9].

Before delving into the details of the track building logic of the algorithm, it is worth noting a real example for its application. The CMS Phase 2 Outer Tracker, shown in figure 1 (blue and red lines) is a suitable detector for the LST algorithm. This comes from the fact that

each layer is composed by multiple “ p_T modules” (lines in figure 1), each with two closely-spaced sensors which can provide a very localized combination of hits, called *stub*, as shown in figure 2. In the LST algorithm presented here, a stub is accepted only if the hits are within a selection window (marked in green in figure 2) which corresponds to a p_T threshold of 0.8 GeV for the traversing particle. The usage of stubs instead of individual hits can significantly help in reducing the combinatorics in the initial steps of the LST algorithm up to 7.5 times. As an example, the number of hits in the first layer of the outer track in one $t\bar{t}$ event at 200 PU event is approximately 36000, while this number goes down to 6000 for stubs in the same layer.

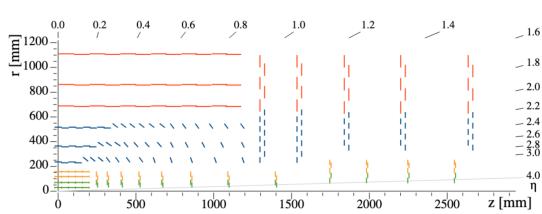


Figure 1. $r - z$ perspective of one quarter of the Phase-2 CMS tracker. The green and orange lines represent the inner tracker modules, while the blue and red lines represent the outer tracker modules [10].

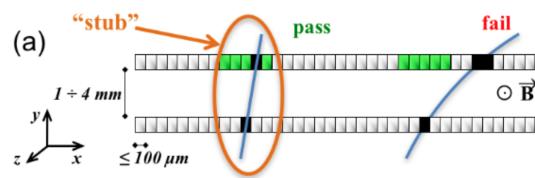


Figure 2. Sketch of a tracker p_T module. The series of gray squares represent the closely-spaced sensors of a module, while the green squares indicate the channels compatible with a particle which has passed through reference channel (black square). The lower the particle p_T , the lower the probability to produce an accepted stub [10].

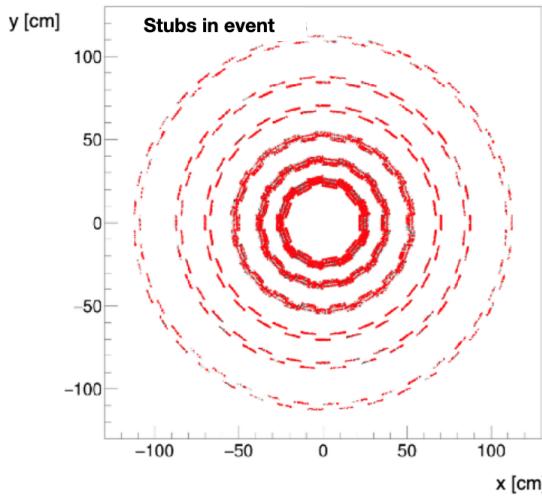


Figure 3. Occupancy distribution for stubs in a single $t\bar{t}$ event at 200 PU in the barrel part of the Phase-2 CMS tracker as a function of the x and y coordinates of the CMS detector.

2.1. The Logic of the Algorithm

The track building procedure starts with the creation of stubs in each of the p_T module of the outer tracker. The occupancy distribution of stubs in the $x - y$ two-dimensional plane of the CMS coordinate system is shown in figure 3 for a typical $t\bar{t}$ event at 200 PU. The next step is the linking of stubs in neighboring layers (figure 6) together to form a *segment* (figure 7). With around 6000 stubs in a single layer, the combinatorics for all the stub connections is immense. To decrease the number of combinations, physics requirements are taken into account. More specifically, helices of particles originating from the center of the detector, with both positive

and negative sign hypotheses and with p_T at the previously mentioned threshold of 0.8 GeV, are considered. Tracks from simulation are considered as well, in order to include cases that may have been missed by the idealistic physics scenario of helices. This leads to restrictions in the number of modules that can be linked to a single (reference) module, as it can be seen in figure 4.

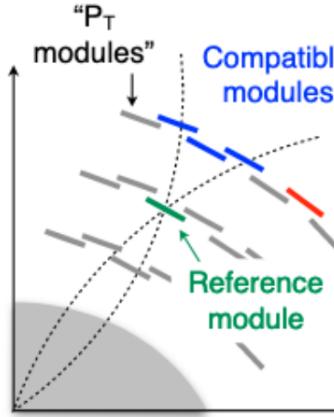


Figure 4. Demonstration of the logic behind the construction of a map of modules compatible for linking algorithm objects. The dashed line represents helices of positively and negatively particles with $p_T = 0.8$ GeV.

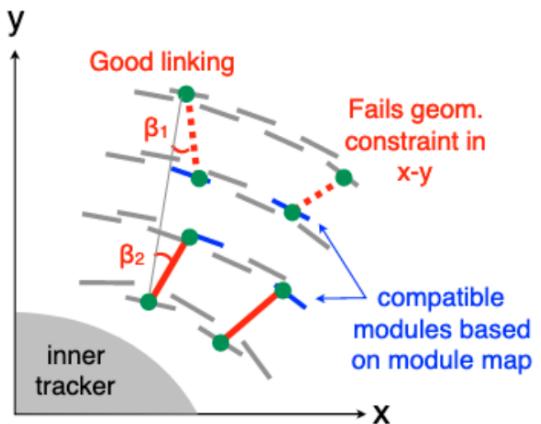


Figure 5. Two examples showcasing how geometrical requirements, in combination with the list of compatible modules, lead to accepting or rejecting the linking of algorithm objects.

The creation of segments is followed by the creation of *triplets* ($T3$) and *quintuplets* ($T5$). Triplets are formed from the linking of two segments that share a common stub (figure 8), while quintuplets involves the linking of two triplets with a common stub (figure 9). As objects get longer, more handles can be utilized for determining whether their linking is consistent with a track hypothesis. As an example, figure 5 shows that the relative angle between segments can be used to reject their combination if that leads to an unphysical track pattern, such as the right set of segments in the figure.

In addition to the outer-tracker-only objects, exploiting information from the inner tracker can assist in creating collections with higher purity of physical tracks. Triplets and quintuplets are linked to seeds from the pixel detector, called *pixel line segments* (pLS), hence producing even longer objects, *pixel triplets* ($pT3$) (figure 10) and *pixel quintuplets* ($pT5$) (figure 11), respectively.

After all of the objects have been created, a “cross-cleaning” procedure is applied. This ensures that duplicates from different object collections are removed from the final output *track candidate* (TC) collection and is implemented as follows: First, any pixel triplets that are close to any pixel quintuplet in the $\eta - \phi$ plane are marked as duplicates. The same procedure then takes place for quintuplets versus both pixel quintuplets and cleaned pixel triplets. Finally, remaining pixel line segments, unused in the previous linking steps, are cross-cleaned versus all the other objects. In the above cross-cleaning steps, it is also required that objects do not share any hits. As the last step of the algorithm, the track candidate collection is created by adding to it all of the following, non-duplicate objects:

- pixel quintuplets,

- pixel triplets,
- quintuplets, and
- unlinked pixel line segments.



Figure 6. Stubs.



Figure 7. Segment.



Figure 8. Triplet.



Figure 9. Quintuplet.

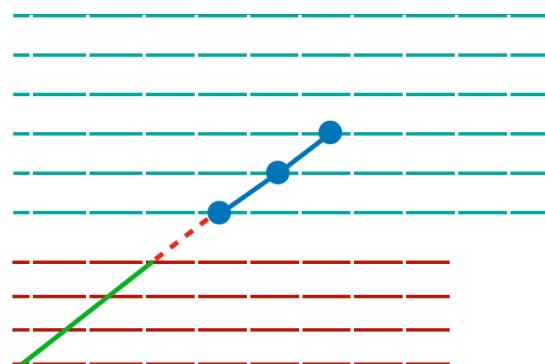


Figure 10. Pixel triplet.

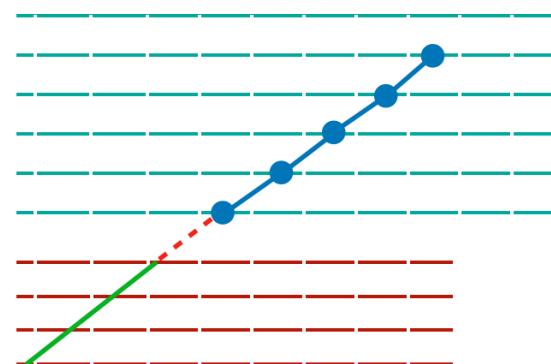


Figure 11. Pixel quintuplet.

The different objects created in the LST algorithm. The cyan lines represent the outer tracker layers, while the dark red lines represent the inner tracker layers.

2.2. Physics Performance

The selection of the objects that comprise the final track candidate collection is based on maximizing the track finding efficiency with the lowest possible fake rate. Each different object category aims at recovering efficiency of tracks with different characteristics. This is showcased in figures 12-14, which show the breakdown of the track candidate efficiency (TC, in black) in different objects as a function of different variables.

Starting from figure 12, which presents the efficiency as a function of p_T , it is evident that the pixel quintuplets (pT5, in red), which are the longest objects, covering almost the whole

range of the detector layers, are the driver of the performance. The same figure indicates that the pixel triplets (pT3, in green) and unlinked pixel line segments (pLS, in magenta) recover some efficiency at lower p_T . This is because low p_T tracks have a lower probability of leaving hits in a lot of tracker layers. The total efficiency for track candidates plateaus at its maximum value of 90% already at 5 GeV. The real benefit of including the unlinked pixel line segments in the track candidate collection reveals itself in the efficiency plot as a function of η , where they noticeably are the drivers of the performance at the high $|\eta|$ range. Finally, figure 14 shows that the inclusion of quintuplets (T5, in blue) is the unique handle of reconstructing displaced tracks.

In terms of fake rate, figures 15 and 16 present it as a function of p_T and η respectively. Above the lower p_T threshold of 0.8 GeV, the fake rate remains less than 10% for the bulk of the events, up till ~ 5 GeV. For higher p_T values, some requirements on the objects are loosened to improve efficiency, which leads to an increase in the fake rate. The fake rate tends to be higher in the central region, where less accurate p_T modules are employed with respect to the region of $|\eta| > 2.2$. Due to the extensive cross-cleaning, the duplicate rate is less than 5%, flat as a function of p_T , with the main contribution coming from unused pixel segments at high $|\eta|$.

As a reference for the performance of the LST algorithm, the current CMS algorithms, adapted for HL-LHC, can be used. Looking at the “Patatrack Trimmed” configuration of reference [11], the efficiency of the LST algorithm is comparable with that, consistently around 90% over the whole $|\eta|$ range. At the same time, the LST fake rate is higher, reaching $\sim 15\%$ versus 4% in the barrel and $\sim 20\%$ versus 14% in the endcap. It is worth mentioning that the comparison can only be very rough, since important components of track reconstruction, such as the track final fitting, have not been implemented yet for the LST algorithm.

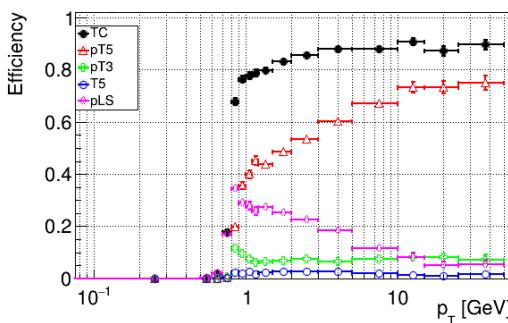


Figure 12. Efficiency for reconstructed tracks in a $t\bar{t}$ sample at 200 PU, as a function of the generator track p_T . The color code is explained in the text.

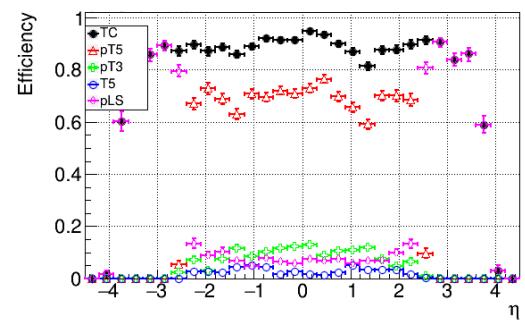


Figure 13. Efficiency for reconstructed tracks in a $t\bar{t}$ sample at 200 PU, as a function of the generator track η . The color code is explained in the text.

3. GPU Implementation

GPUs have both advantages and disadvantages when compared to CPUs, making each one of these processing units suitable for different applications. GPUs host a much larger number of cores, $\mathcal{O}(10^3)$, than CPUs do, which have at most a few tens of cores. Due to this, GPUs can achieve greater numbers of floating point operations per second, hence being able to maintain high throughput. On the other hand, GPUs have limited memory resources, which can lead to high latency, when host-to-device data transfers are needed, while CPUs have lower memory latency, owing to more extensive use of caches. Given these characteristics, CPUs are more suitable for serial processing, while GPUs excel in parallel computations.

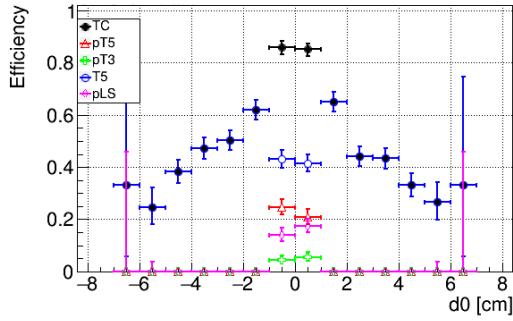


Figure 14. Efficiency for reconstructed tracks in a muon gun sample at 0 PU with the muon production vertex uniformly distributed in a cube with edge of 5 cm, as a function of the generator track d_0 , i.e. its distance from the center of the detector. The color code is explained in the text.

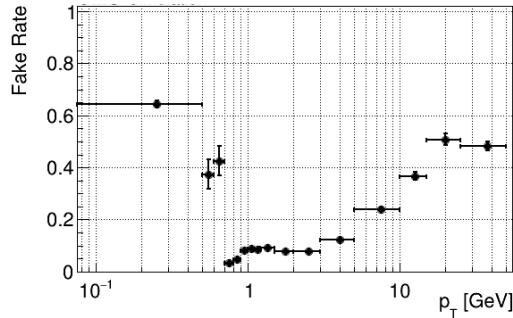


Figure 15. Fake rate for reconstructed tracks in a $t\bar{t}$ sample at 200 PU, as a function of the reconstructed track p_T .

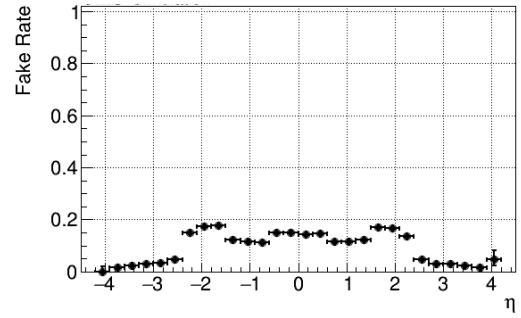


Figure 16. Fake rate for reconstructed tracks in a $t\bar{t}$ sample at 200 PU, as a function of the reconstructed track η .

From the considerations above, it is clear that GPU programming requires both a parallelizable algorithm and an appropriate implementation to exploit the full capabilities of GPUs. As already mentioned, LST is an algorithm that is fit for parallelization according to the SIMD principle and our implementation, described below, aims to take advantage of the multitude of GPU cores while efficiently handling memory assignments and transfers. The LST algorithm implementation on GPU has been developed with CUDA, a programming framework for general-purpose computing on GPUs, developed by NVIDIA.

One cornerstone of the LST algorithm implementation is the usage of Structure of Arrays format, commonly referred to as “SoA”, with explicit memory management. Instead of using the common practice of object-oriented programming to create classes that hold the properties of a single object, separate arrays are created for separate object properties. These arrays hold each relevant quantity in continuous memory addresses, hence allowing for successive memory reads of the same quantity to be contiguous (memory coalescing). In this way, the time needed for accessing memory is minimized. Additionally, a custom cache memory allocation is utilized for small and commonly used objects. The cache memory is of very limited size but it allows for extremely fast memory transfers, helping further to decrease the total execution time of the algorithm. Finally, extensive studies on the final occupancy of all algorithm objects have led to truncations of their output numbers at the level of 99.9% or 99.99%, resulting in a 5-fold reduction of the total memory footprint for the algorithm, while sacrificing only 1–2% of efficiency.

In terms of organization of its algorithmic steps, the LST code is split into different kernels for the creation of each object. A kernel is a set of operations that can run in parallel on independent GPU threads. Due to the fact that inputs for each kernel are highly localized, a large number of

independent object reconstruction threads can run in the same kernel concurrently. The order in which kernels are launched roughly follows the description of the algorithm in subsection 2.1 and is shown in figure 17.

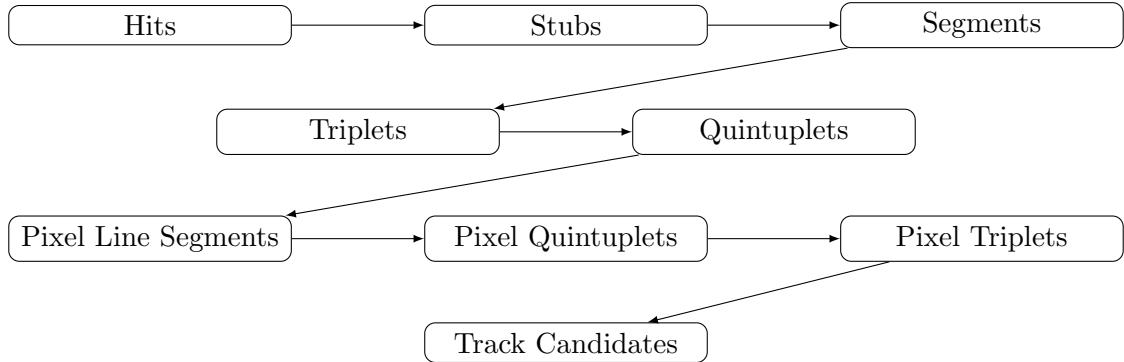


Figure 17. Flowchart of the order of kernel execution.

There is another layer of parallelization that can be exploited: event-level parallelization. Instead of running only one event in one GPU stream, multiple streams can be employed, each one of which runs a different event. This technique, called “multi-streaming”, allows for better utilization of the GPU resources and provides more opportunities for the kernel scheduler to insert more work. Moreover, it can hide the high GPU latency by assigning work to processing elements which already have data available, while others wait for data to be transferred. This can be seen in figure 18, where light blue indicates the total work done by the GPU and dark blue shows the work done by kernels in different streams, both as a function of time. In the 1-stream case (upper schematic), the total GPU work is interrupted while the single kernel running waits for new data. On the other hand, in the 8-stream case (lower schematic), even though kernels in separate stream have similar waiting times, the total work done by the GPU is continuous in time, as data loading is hidden by work in other streams.



Figure 18. Comparison of algorithm execution on 1-stream (upper) versus 8-streams (lower), using the NVIDIA Visual Profiler. Multi-streaming allows for running parts of multiple kernels concurrently.

The benefits of multi-streaming are evident when measuring the time needed for the execution of the LST algorithm. This timing measurement is performed in a $t\bar{t}$ sample at 200 PU and excludes the initial data transfer host-to-device for the hits, as well as any final fitting procedure

on the output tracks, which has not been implemented in the current version of the algorithm. The average time needed for the execution of the algorithm on a single event is 32 ms for the 1-stream case, while it decreases to 26 ms for the 8-stream case, resulting in $\sim 20\%$ improvement with respect to the 1-stream case. These timing measurements are in the same range as the latest CMS tracking timing measurements on CPUs [11]. It is worth noting that the LST algorithm is on par also price-wise with multi-CPU efforts, since two 32-core Intel Skylake Gold Xeon processors, commonly used in such efforts, cost approximately the same as one NVIDIA Tesla V100 GPU that has been used in our case.

4. Summary and Outlook

In this work, Segment Linking (or Line Segment Tracking), a highly parallelizable track finding algorithm, aimed at alleviating the challenges that the HL-LHC will pose for tracking, has been presented. The algorithm has been successfully implemented on an NVIDIA Tesla V100 GPU and has been shown to have both physically and computationally comparable performance with cutting edge tracking efforts on CPUs.

More improvements of the algorithm are planned for the future. On the physics side, a re-evaluation of the selection criteria for objects is underway to improve the efficiency and lower the fake rate, especially for displaced tracks. Integration with other, similar GPU developments, such as Patatrack, is also foreseen. Computationally, mathematical optimizations for the calculation of physics parameters, as well as the extension of the usage of alternative data types, e.g. half-precision floats, are being explored. Plans to refine memory coalescing will potentially help reducing timing even further. Ultimately, the grand design for the algorithm involves its integration to the central CMS software for HLT and offline usage during the HL-LHC operation.

5. Acknowledgements

This work was supported by the U.S. National Science Foundation under Cooperative Agreements OAC-1836650 and PHY-2121686 and grant NSF-PHY-1912813.

References

- [1] Aberle O et al 2020 High-Luminosity Large Hadron Collider (HL-LHC): Technical design report CERN-2020-010 URL <https://cds.cern.ch/record/2749422>
- [2] Cerati G B 2014 Tracking and vertexing algorithms at high pileup CMS-CR-2014-345 URL <https://cds.cern.ch/record/1966040>
- [3] The CMS Collaboration 2008 *Journal of Instrumentation* **3** S08004
- [4] Bocci A, Innocente V, Kortelainen M, Pantaleo F and Rovere M 2020 *Frontiers in Big Data* **3**
- [5] The CMS Collaboration 2021 The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger CERN-LHCC-2021-007, CMS-TDR-022 URL <https://cds.cern.ch/record/2759072>
- [6] Fröhwirth R 1987 *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **262** 444
- [7] Cerati G, Elmer P, Gravelle B, Kortelainen M, Krutelyov V, Lantz S, Masciovecchio M, McDermott K, Norris B, Reid M, Hall A R, Riley D, Tadel M, Wittich P, Würthwein F, Yagil A 2019 Speeding up Particle Track Reconstruction in the CMS Detector using a Vectorized and Parallelized Kalman Filter Algorithm URL <https://arxiv.org/abs/1906.11744>
- [8] Ciobanu C, Gerstenlager J, Hoftiezer J, Hughes R, Johnson M, Koehn P, Neu C, Sanchez C, Winer B L, Freeman J, Holm S, Lewis J, Shaw T, Wesson T, Bloom K, Gerdes D, Dawson, J W, Haberichter W N 1999 *IEEE Transactions on Nuclear Science* **46** 933
- [9] Krutelyov V, Cerati G, Tadel M, Wuerthwein F, Yagil A 2017 *Journal of Physics: Conference Series* **898** 042023
- [10] The CMS Collaboration 2017 The Phase-2 Upgrade of the CMS Tracker CERN-LHCC-2017-009, CMS-TDR-014 URL <https://cds.cern.ch/record/2272264>
- [11] The CMS Collaboration 2021 Performance of Phase-2 HLT Reconstruction and GPU offloading benchmarks CMS-DP-2021-013 URL <https://cds.cern.ch/record/2792313>