January 6, 2017

Report on the L1 Tracking Review held on December 8-9, 2016

D. Abbaneo, C. Foudas, A. Marchioro, E. Perez, R. Van Berg, F. Vasey

1. Scope of the review

Over the past years three teams within the tracker collaboration have developed demonstration systems to establish the feasibility of a system which will provide tracking information to the CMS Level-1 trigger during the High-Luminosity operation of the LHC, using data from the future upgraded outer tracker of CMS. The three groups have developed hardware demonstrators and the associated firmware and emulator software, and have evaluated their designs.

The results were presented at CERN[1] during the 8th and the 9th of December, and further discussed during Q&A sessions that took place on December 12th. In this report the main findings from this exercise are summarized, along with the outlook on the development of a possible final system. The three designs are evaluated based on performance, robustness, flexibility, scalability, development effort and risks, expected cost and size of the system, according to criteria defined and agreed upfront[2,3].

The extrapolation from the existing demonstrators to the final system involves assumptions that were discussed during the review; all three teams proposed development plans that were precise enough to assess scalability.

This report is intended to provide useful input to the collaboration to optimize and rationalize the design of the final system, considering the importance of this project for the CMS upgrade, and to define an appropriate development plan.

2. Performance from simulation

Results from demonstrators were validated against simulations (although not all three demonstrators have been fully completed). Although some samples came late, the performance has been explored in a large phase space, up to 200 Pile-Up (PU), extending the acceptance down to 2 GeV and including effects of detector inefficiencies.

The performance is fully satisfactory in terms of track finding efficiency, resolution and output rate. It is worth noting that the results were obtained using the flat TBPS geometry (inner three layers of the outer tracker, the most crowded region of the detector), which gives a factor of two higher rate of stubs than the final geometry, which is a huge overhead. The small differences that can be seen in the performance plots shown during the review were discussed in the Q&A sessions and are largely understood. It is expected that additional optimisation shall reduce them even further. The R&D carried out on the demonstrators has generated new ideas for further improvements, that at this point have only been partially developed, since the agreed priority was to complete the implementation in hardware of the baseline options, in time for the review. Such new ideas as well as the details of the extrapolations to the final systems have been discussed in the Q&A sessions.

---

[1] Agenda of the meeting

[2] Charge of the Review

[3] Specifications of the demonstrator systems

All three methods are able to extend the tracking acceptance down to 2 GeV. For detector inefficiencies, the Tracklet implementation is intrinsically insensitive to the loss of efficiency in one layer, while AM and TMTT in the current implementation can recover performance through a change of configuration parameters, which appears to be adequate (and probably can be further improved).

Overall, although substantial work will be needed to develop an optimal implementation, as concepts all three methods are proven to be fully appropriate in terms of achievable tracking performance (efficiency, output rates and track parameter resolution), meeting all requirements in the most challenging conditions and with ample margin.

3. Demonstrators

The three demonstrators have reached different levels of completeness, but are sufficient to validate the performance results from simulations, and to give a solid indication of the achievable latency, which was one of the main purposes of the exercise.

3.1 Brief description of the demonstrators

The TMTT demonstrator is complete, based on five micro-TCA MP7 cards (excluding data source and sink): it covers one time-slice (1/36), and one full sector (one octant [$2\pi/8$], plus+minus sides). The MP7 cards implement Virtex-7 690 FPGAs. The algorithm firmware uses a Hough transform for pattern recognition followed by a Kalman filter.

The Tracklet demonstrator has been completed for two subsector configurations separately, based on three micro-TCA CTP7 cards (excluding data source and sink). One configuration covers one time-slice (1/6) and a half barrel sector ([$2\pi/28$], barrel plus side)], the other configuration covers a half hybrid sector ([$2\pi/28$], [disks and part of barrel] plus side). An integrated half sector ([$2\pi/28$], plus side) was reported as almost complete, except for the duplicate removal step, but no results were available, yet. The CTP7 cards also use Virtex-7 690 FPGAs. The algorithm firmware uses the “tracklet” road search algorithm, followed by a χ^2 fit to extract the track parameters.

The AM demonstrator is partly complete, based on ten ATCA Pulsar IIb cards (excluding data sources): all 10 time-slices and a full trigger tower (1/48) are in-shelf, but the Pattern Recognition Mezzanine (PRM) is incomplete: the AM is emulated in a Kintex Ultrascale FPGA (1’000 patterns instead of 250’000), the communication link from Pattern Recognition Board (PRB) to PRM is missing, and the duplicate removal step is not yet implemented in hardware. An alternative PRM based on up to 12 AM06 ASICs (128’000 patterns per ASIC) was described as functional and being integrated with the PRB; such board represents a step towards system integration, but it is not relevant to the latency demonstration since it is not designed to meet CMS timing requirements.

3.2 Measured latency and extrapolation to final system

Latency estimates were given with different levels of uncertainty, depending on the level of completeness of the demonstrators, and with slightly different assumptions on the functionality of the DTC boards. The quoted numbers are adequate to give a first quantitative assessment of the required latency for L1 tracking, keeping in mind that several options can still be explored to optimize the implementation of the final system, hence speculating on small differences at this stage is not meaningful.

TMTT Proposal: The full trigger chain is demonstrated and the measured latency is 3.763 µsec. Using the FW emulator for a Kintex Ultrascale (KY115) FPGA, a predicted latency of 2.852 µsec is obtained for the final system (18 time slices).

Tracklet Proposal: The trigger chain is demonstrated separately in the barrel and the hybrid barrel/endcap regions. The measured latency is 3.48 µsec. A considerable part of the latency budget (≈ 1 µsec) is spent in the sector-to-sector communication, which may be avoided by re-optimizing the architecture. Further gains will come from the use of next-generation FPGAs.

AM Proposal: the full trigger chain is not demonstrated; the latency is measured in two sections of the chain. The resulting measured latency is 2.53 µsec, which increases to 3.025 µsec when including a further ×2 TMUX factor in the track fitting step, to improve performance in high-density environment. Such value is still not including the latency of the duplicate track removal algorithm, which is not yet implemented. The latency of that step is estimated to be around 400 ns. Despite the incomplete chain, it is reasonable to expect that with further optimization, for the final system an overall latency below 3 µsec should be achievable also for this method.

In summary, based on the measurements and simulations performed, we can conclude that a latency below 4 µsec is demonstrated, and a latency below 3 µsec is expected to be achievable for the final system, for all three methods.

4. Developments towards final system

Some of the main characteristics of the final system, as proposed for each of the three methods and presented in the review, are listed in the table below.

	AM (24 crates scenario)	**TMTT** (baseline scenario)	**Tracklets**
Cost (MCHF,MUSD)	**6.6**	**4.3**	**4.3**
Power/blade (W)	**≈ 250** (incl. 150 for 2xPRMs)	**< 200**	**< 200** (est. to be similar to TMTT)
System footprint	**240 boards** **24 shelves**	**144 boards** **16 shelves**	**144 boards** **12 shelves**

The costs and system footprints reported in the table are the figures presented by the three teams. The estimates are not perfectly coherent in terms of assumptions and level of details considered in the extrapolation exercise, but they give a correct view of the scalability of the systems at the current level of knowledge, when there are still several options to explore to optimize the systems.

The hardware developments planned for the final systems involve uncertainties that translate to risks in budget and schedule. Development of firmware and software can also generate important schedule risks, if the project is not structured in an efficient way. Some key issues are listed below, applicable to one or more of the proposed methods.

- Multiple hardware and/or firmware development paths based on different frameworks and hardware platforms lead to inefficient use of the available resources, resulting in schedule risks.

- The Track Finding system has a deep interplay with the DTC systems. The two systems need to be developed coherently to avoid inefficient efforts and schedule risks.
- The implementation of high-speed links will require substantial effort. While 10 Gb/s links are already in use in the HEP community, the next-generation of 16 Gb/s links appears to be a conservative choice on the time scale of the project. Next-to-next generation 25 Gb/s links may be more difficult to implement: a risk-to-benefit analysis based on link demonstrators in hardware and firmware will have to be carried out before committing to such a choice.
- Full size ATCA shelves when used to their full capability are not simple to operate, especially if heavy use is made of either the backplane mesh or RTM I/O resources. Adequate time for development of tools and methods has to be accounted for in the project planning. The successful development of the Pulsar II-b can be considered as an important stepping stone towards the final system.
- Development of an integrated circuit in a new technology is not trivial, even if the circuit is considered relatively simple. The 28 nm technology point is much more complex than the 65 nm and 130 nm technologies, that are by now relatively common in HEP. Complete and verified design kits for 28 nm do not yet exist in the HEP domain.
- 3-D or two-tier technologies are becoming somewhat more common in commercial areas, but despite a long development effort in HEP, no HEP project has yet been able to take advantage of those technologies.

Based on the above considerations, the following specific comments apply for the proposed methods.

AM Proposal

- Two projects were presented, a planar 28 nm version and a 3D stacked version of the associative memory, carried out by separate engineering teams with parallel development plans. Some observations are given below for the two options.
    - 3D stacked. 3D chip stacking combined with advanced TSV technologies is the likely and very exciting future for many applications in particle physics, mainly in the detector to read-out electronics integration. The remarkable progress shown by industry in the last few years is the result of heavy investments, far beyond what HEP can produce in the short or medium timescale. The two-tiers approach offers more routing resources, however there is no evidence that the design of the chip would be limited by routing, (given the large area of the single associative memory cell), when using a 40 or 28 nm planar technology with up to 9 metal layers. Hence there is no compelling reason to link the CMS project to such R&D.
    - 28 nm planar. The development could be feasible in the required timescale, provided that sufficient engineering resources are made available. The design presents significant challenges, such as: (i) the speed of the associative memory does not automatically scale with the technology, substantial architectural changes are required to operate a large (i.e. many patterns) chip at or above 200 MHz; (ii) massive use of parallel links between chips to reduce latency, instead of a serial link, may introduce serious timing issues; (iii) the projected power consumption of up to 10 W calls for complex and expensive packages, that should be designed in a proper chip-package co-design environment from the early stages of design. The complexity of the chip package to some extent transfers to added complexity in the board design.

- For both chip technology options, it is optimistic to assume that a project of such complexity will succeed in producing the final design with a single mask set. Contingency should be included, both in budget and schedule, for fixing problems in a further iteration that may cost an additional mask set or a fraction of it, and one additional year in the development schedule.
- It is evident that any schedule risk in the chip development propagates to the development of the pattern recognition mezzanine boards, and to the production and commissioning of the whole system.

Tracklet & TMTT Proposals

- For what concerns hardware developments, the two FPGA-based proposals rely on commercially available technologies. The projects can be considered low-risk provided that the transition to high-speed links is planned and implemented cautiously, and the capabilities of the ATCA boards are not exploited to their limits (which is the case for both proposed implementations).
- It is desirable to continue exploring different algorithmic options for track reconstruction, but such R&D should be based on a single well-defined hardware platform and a common development framework, to avoid generating problems in the project schedule.

5. Conclusions

Thanks to the remarkable efforts of the participating groups over the past months, the results achieved demonstrate that track reconstruction at Level-1 is a realistic option for CMS at HL-LHC. Although substantial optimization work remains to be done, the demonstration systems that have been realized have shown that the three proposed methods can meet, and even exceed, the demanding requirements of the CMS trigger, with the necessary margin and robustness.

While the AM-based method requires a challenging chip development in a novel (for HEP) technology, carrying many unknowns that may significantly impact cost and schedule, a method based only on commercial FPGAs can be considered as low-risk, provided that the required developments are properly planned and managed.

The committee members wish to congratulate the groups for the remarkable efforts and the exciting results obtained.

Special thanks to the Working Group conveners K. Hahn and P. Wittich for the excellent organization of the review. Having common specifications for the demonstrators and standardized slides for the performance plots has been extremely helpful, enabling an objective comparison of the three proposed methods.