Feasibility of Hardware Acceleration in the LHC Orbit Feedback Controller (OFC)

L. Grech, G. Valentino
D. Alves, S. Jackson, J.Wenninger

L-Università ta' Malta

Abstract: Orbit correction in accelerators typically make use of a linear model of the machine, called the Response Matrix (RM), that relates local beam deflections to position changes.The RM is used to obtain a Pseudo-Inverse (PI), which is used in a feedback configuration, where positional errors from the reference orbit as measured by Beam Position Monitors (BPMs) are used to calculate the required change in the current flowing through the Closed Orbit Dipoles (CODs). The calculation of the PIs from the RMs is a crucial part in the LHC's Orbit Feedback Controller (OFC), however in the present implementation of the OFC this calculation is omitted as it takes too much time to calculate and thus is unsuitable in a real-time system. As a temporary solution the LHC operators pre-calculate the new PIs outside the OFC, and then manually upload them to the OFC in advance. In this paper we aim to find a solution to this computational bottleneck through hardware acceleration in order to act automatically and as quickly as possible to COD and/or BPM failures by recalculating the PIs within the OFC. These results will eventually be used in the renovation of the OFC for the LHC's Run 3.

OFC and SVD

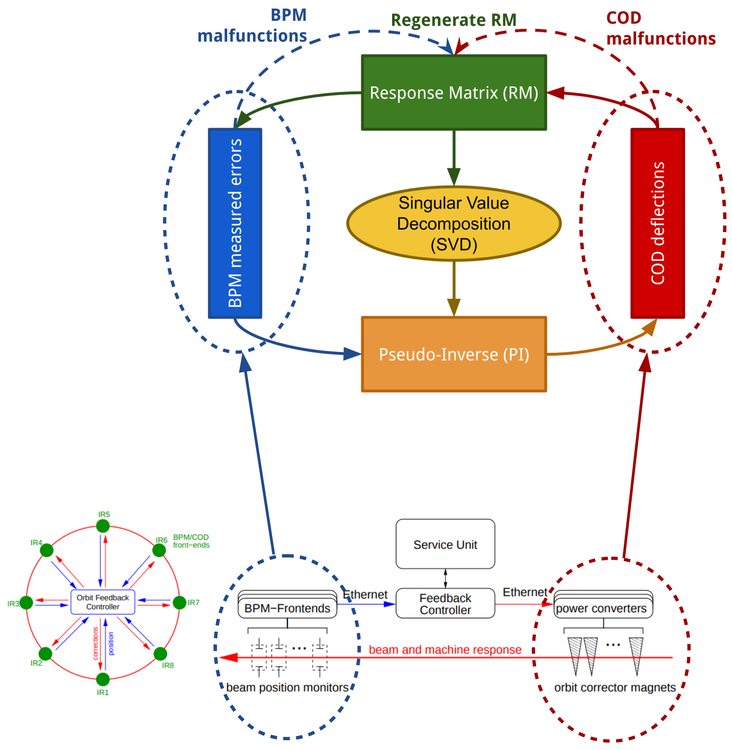

- RM is generated from the LHC optics
- RM relates COD deflections to positional vector
- Malfunctions in BPMs and CODs
  - Regenerate RM
  - Recalculate PI

Problem Definition

- During normal LHC operation **several RMs** are used .
  - Different optics, e.g. during RAMP 13-14 RMs are used
- **BPM and COD malfunctions** prompt regeneration of RM and consecutive PI re-calculation
- Re-calculation of PI for each RM is a **computational bottleneck** in the OFC

Proposed Solution

- **Hardware acceleration** to speed-up PI re-computation through SVD

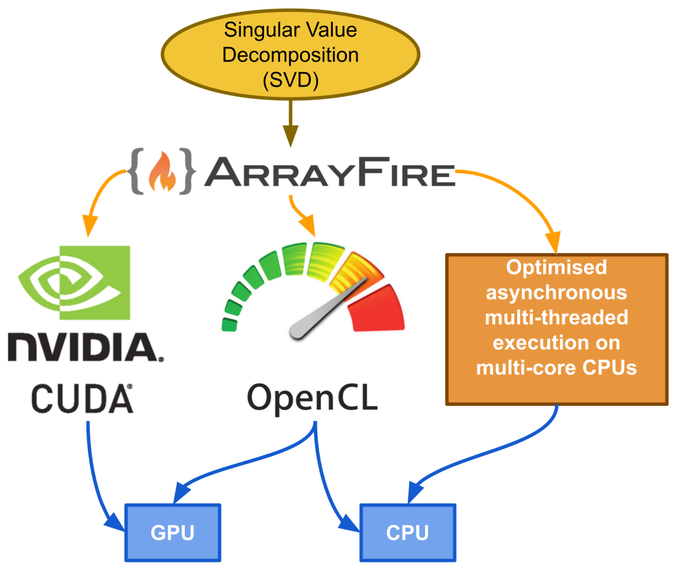

- **ArrayFire** → Hardware acceleration library
- CUDA, OpenCL or optimised CPU implementations of SVD

Results

- Benchmarking test shows that ArrayFire using CPU obtained the best performance

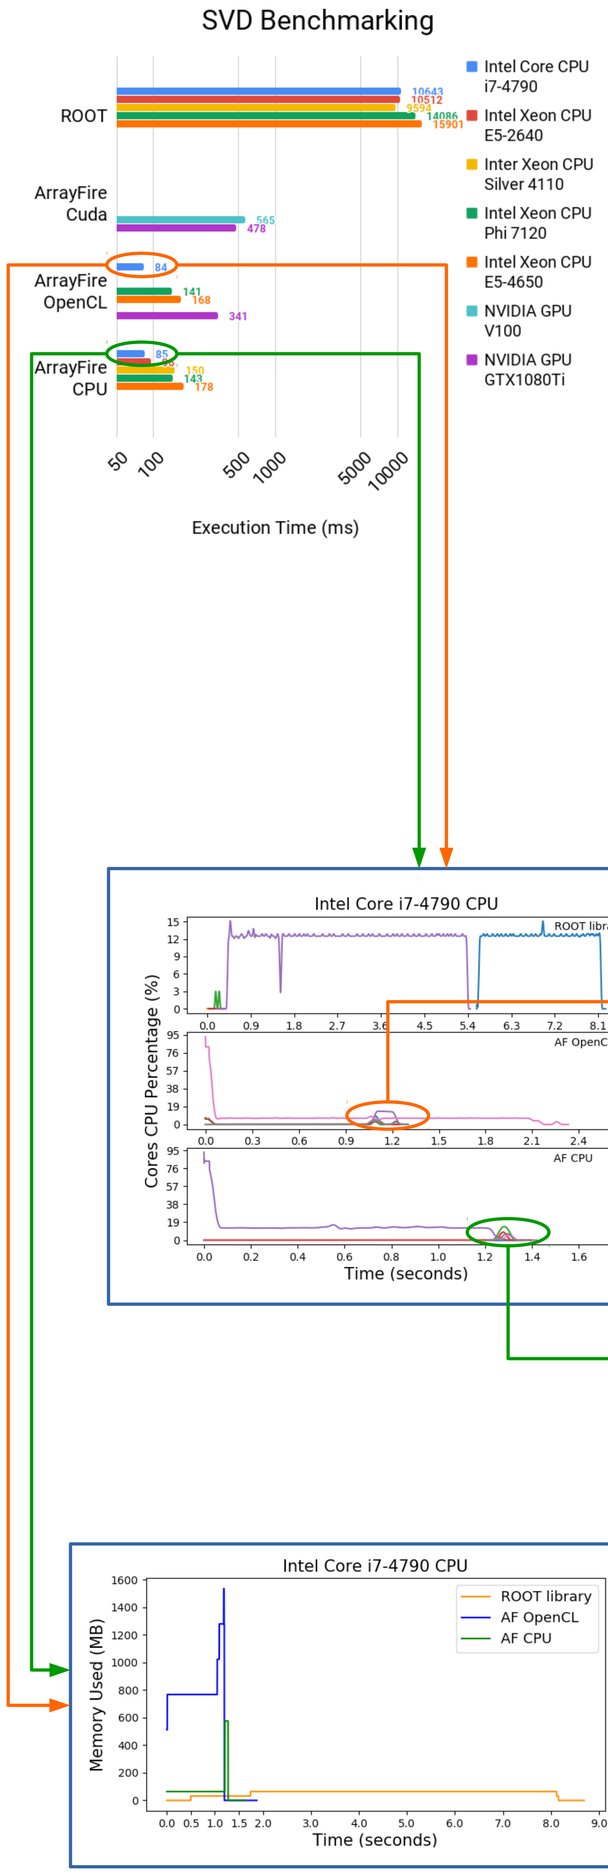

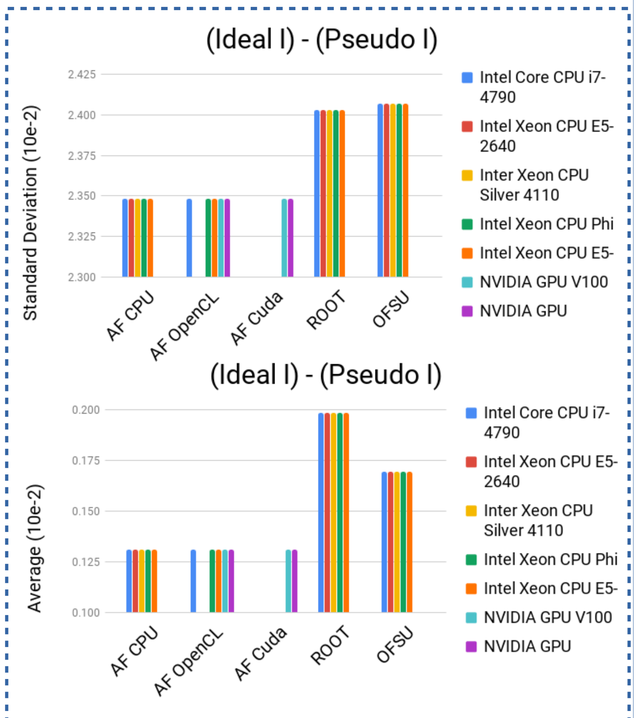

- Accuracy test shows that libraries obtain similar accuracy in the PI calculation

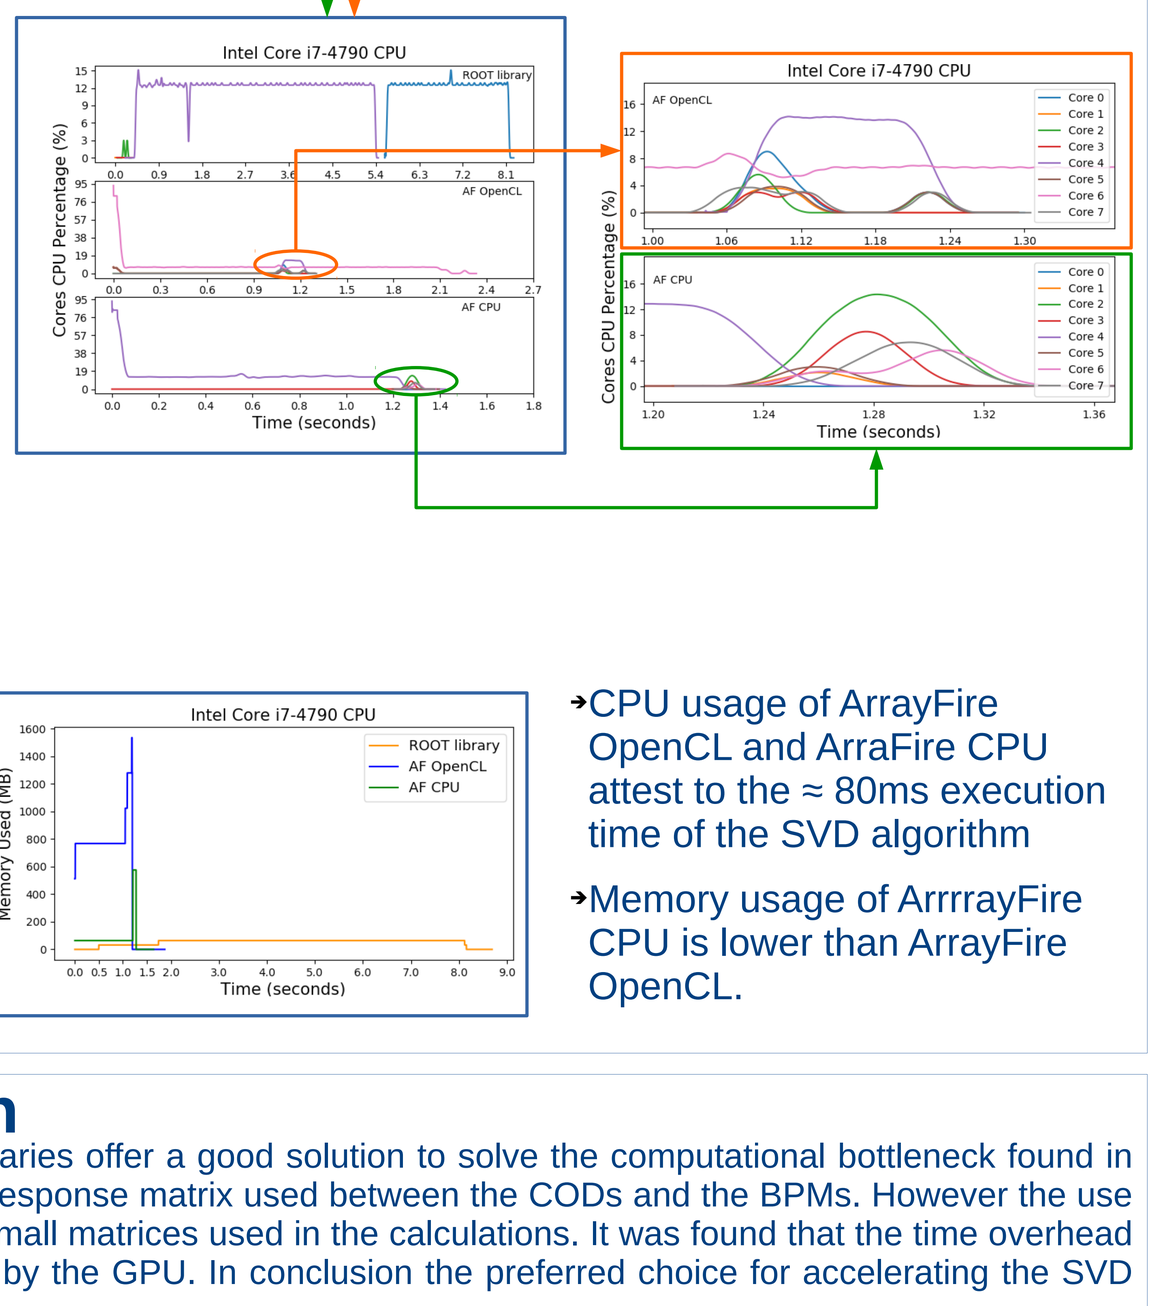

- CPU usage of ArrayFire OpenCL and ArraFire CPU attest to the ≈ 80ms execution time of the SVD algorithm
- Memory usage of ArrrrayFire CPU is lower than ArrayFire OpenCL.

Conclusion

From the work presented in this paper, it was concluded that hardware acceleration libraries offer a good solution to solve the computational bottleneck found in the current implementation of the OFC when it is calculating the pseudo-inverse for the response matrix used between the CODs and the BPMs. However the use of GPUs in future designs of the OFC would not be recommended due to the relatively small matrices used in the calculations. It was found that the time overhead of copying the input matrix to the GPU would balance out the accelerated computation by the GPU. In conclusion the preferred choice for accelerating the SVD computation within the OFC is ArrayFire using CPUs.

References

- R J Steinhagen. LHC Beam Stability and Feedback Control – Orbit and Energy. PhD thesis, RWTH Aachen U., 2007.
- CERN. ROOT a data analysis framework | ROOT a data analysis framework. https://root.cern.ch/. Accessed: 2018-8-15.
- Michel Arruat, Leandro Fernandez, Stephen Jackson, Frank Locci, Jean-Luc Nougaret, Maciej Peryt, Anastasiya Radeva, Maciej Sobczak, and Marc Vanden Eynden. Front-end software architecture. In ICALEPCS07, volume 7, pages 310–312, 2007.
- J Wenninger and R Steinhagen. LHC orbit feedback control requirements. Technical report, CERN AB-OP, March 2007.
- CUDA (compute unified device architecture) definition. https://techterms.com/definition/cuda, July 2015. Accessed: 2018-12-3.I
- OpenCL - the open standard for parallel programming of heterogeneous systems. https://www.khronos.org/opencl/, July 2013. Accessed: 2018-12-3.
- Techlab. https://techlab.web.cern.ch/. Accessed: 2018-12-3.
- James G Malcolm, Pavan Yalamanchili, Chris McClanahan, Vishwanath Venugopalakrishnan, John Melonakos, and others. ArrayFire: a GPU acceleration platform. Proceedings of SPIE - The International Society for Optical Engineering, May 2012.
- David Levinthal. Performance analysis guide for intel® core i7 processor and intel® xeon 5500 processors. Technical report, Intel Corporation, 2009.

Code

MOPHA151

leander.grech@cern.ch

ICALEPCS 2019