A GPU VARIANT OF MBTRACK AND ITS APPLICATION IN SLS-2

U. Locans[1,2,*], H.S. Xu[1], A. Adelmann[1], L. Stingelin[1]
Paul Scherrer Institut, Villigen, Switzerland
University of Latvia, Riga, Latvia

Abstract

Mbtrack is a widely used multi-bunch tracking code for modeling collective instabilities in electron storage rings. It has been applied to the Swiss Light Source upgrade proposal (SLS-2) for the study of single bunch instabilities. However, an n-bunch simulation using mbtrack requires to run n+1 MPI processes. Therefore, a large scale computing cluster may be necessary to perform the simulation. In order to reduce the demands of computing resources for multi-bunch simulations, a CUDA version of mbtrack was developed, in which the computations of mbtrack are offloaded to a graphics processing unit (GPU). With the mbtrack-cuda variant, multi-bunch simulations can now run in a standalone workstation equipped with an Nvidia graphics card for scientific computing. The implementation and benchmark of the mbtrack-cuda code together with the applications in the study of longitudinal instabilities for SLS-2 will be presented.

INTRODUCTION

Mbtrack is a multi-bunch tracking code for the study of collective instabilities in synchrotrons and electron storage rings. The limiting factor of mbtrack is the computational resources required by large scale multi-bunch simulations. A the CUDA version of mbtrack was developed in order to allow computations to be offload to the graphical processing unit (GPU). This allows to perform multi-bunch simulations on systems which lack the necessary CPU computing power.

The Mbtrack GPU version is using CUDA to target Nvidia GPUs. This version offers the same simulation capabilities as the original mbtrack code. Unlike the CPU version, which parallelizes the simulation over the bunches, the CUDA version takes advantage of massively parallel architecture of the GPU to parallelize the tracking over macro-particles in the bunch. The resulting version was tested on a system equipped with Nvidia Tesla K40c GPU - the results and computing times were compared with the CPU version of mbtrack.

GPU ACCELERATION OF MBTRACK

CUDA Version of mbtrack

The mbtrack-cuda version of multi-bunch tracking software uses CUDA to offload all the heavy computations to the GPU cores, as shown in the Fig. 1, while the host CPU is only used for the simulation setup, saving the statistics and writing the results to output files. The mbtrack-cuda version is capable to simulate the same physics as the original mbtrack code. For full description of mbtrack and its capabilities see [1].

All the algorithms in mbtrack-cuda were rewritten using CUDA to allow the application to offload calculations to Nvidia GPUs. The calculations in the GPU are parallelized over macro particles in the bunch, which means that every thread executed on the GPU performs the calculations for a different particle. In order to calculate statistics and wake functions it is necessary to accumulate data over all the particles in the bunch. For these operations the Thrust library parallel reductions and CUDAs atomic operations are used. During the simulation, all the particle data are kept in the GPU memory. To minimize the data transfer between CPU and GPU the host side only receives data that is needed to write the output files. The flow diagram of transformations performed on the GPU and data transfer between host and the device is shown in Fig. 2.

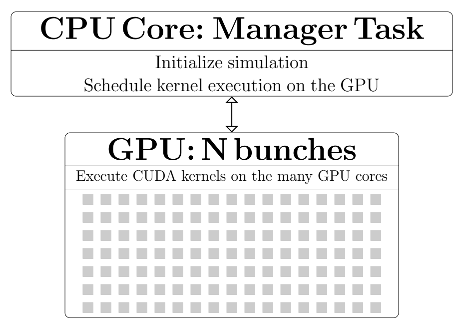

Figure 1: Mbtrack-cuda computation model.

Verification of the Results

To verify the mbtrack-cuda results a comparison with the original code was performed. Since mbtrack-cuda provides the same features as the original code the results from both simulations should match. To test the code lattice 'dc12c', which is an option of the SLS-2 storage ring, was used for the simulations. The main parameters of the 'dc12c' lattice are shown in the Table 1.

Table 1: The Main Parameters of the 'dc12c' Lattice

Parameters	Values
Circumference C_{ring}	290.4 m
Beam Energy E_0	2.4 GeV
Radiation Energy Loss per Turn U_0	569.604 keV
Momentum Compaction Factor α_c	-1.356×10^{-4}
Betatron Tune $\nu_x \setminus \nu_y$	37.221 \ 10.323
Natural Chromaticity $\xi_x \setminus \xi_y$	-66.591 \ -40.445
Transverse Emittance $\epsilon_x \setminus \epsilon_y$	138.55 pm \ 0.12 pm
Harmonic Number h	484
Peak RF Voltage V_{RF}	1.4 MV
Synchrotron Tune ν_s	2.48×10^{-3}

* locans.uldis@psi.ch

ISBN 978-3-95450-182-3

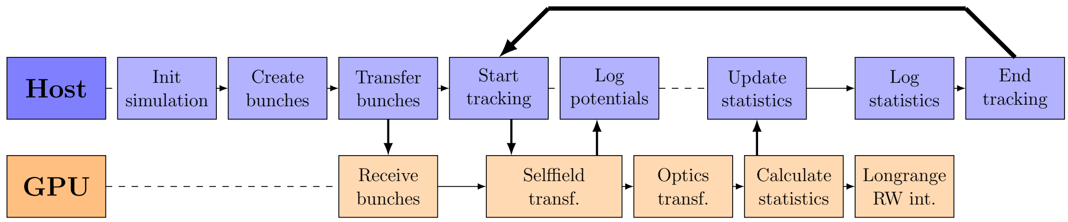

Figure 2: Flow diagram of mbtrack-cuda - tasks executed by CPU and GPU.

The simulations to compare the results from the two versions are performed with identical initial bunch distributions. The synchrotron radiation effects were disabled in these simulations, since these effects require random numbers and that would introduce differences in the results since different random number generators are used by both codes.

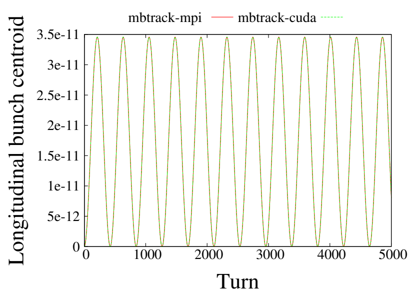

Figure 3: Comparison of mbtrack-mpi and mbtrack-cuda results - both versions produce identical results.

The results of these comparisons are shown in Fig 3. The simulation includes 10 bunches and the results show the mean longitudinal bunch centroid position at different turns. Based on the results from Fig. 3 we conclude that both codes give the same results, which indicates that tracking in mbtrack-cuda works as expected.

BENCHMARKS OF THE GPU VERSION

To test the performance of the mbtrack-cuda version benchmarks were performed on two systems. The first system was equipped with 2x Intel E5-2609 v2 CPUs and the second system was equipped with 2x Intel E5-2697 v4 CPUs. Additionally both systems had a Nvidia Tesla K40c GPU. The software was compiled using gcc 4.8.4 compiler with openmpi 1.8.2. The GPU code was compiled using CUDA 8. The benchmarks aimed to test the ability of mbtrack-cuda to perform large multi bunch simulations as well as compare the performance of CPU and GPU versions for simulations with smaller number of bunches.

In the benchmark simulation each bunch consists of 100,000 macro particles and the simulation was carried out for 20,000 turns. Two separate benchmark tests were performed with synchrotron radiation effects enabled and disabled. The reported results show full execution time of the simulation, which includes all the simulation setup, generation of initial bunch distribution and result output to files. The results from these simulations are shown in Table 2.

Table 2: Benchmarks of mbtrack-mpi and mbtrack-cuda on 8 core CPU, 36 core CPU and a GPU

Bunches	No SR effects			With SR effects		
	CPU1	CPU2	GPU	CPU1	CPU2	GPU
1	260s	155s	93s	810s	549s	93s
2	266s	157s	157s	814s	558s	157s
3	287s	166s	278s	832s	562s	221s
10	2400s	166s	544s	3596s	577s	546s
20	-	396s	1004s	-	757s	1006s
121	-	3260s	5855s	-	2883s	5860s
363	-	6707s	18693s	-	10931s	18754s

The results from the benchmark tests show, that for simulation with small number of bunches the GPU version offers a significant speedup in computing time over the CPU version. This speedup is achieved because the simulation in GPU version is parallelized over macro-particles in the bunch, which allows to use the available parallel resources efficiently even if the number of bunches in the simulation is low. The results also show that the GPU version is able to handle multi-bunch simulations in reasonable time, which would not be possible without a many core CPU or a CPU cluster.

SIMULATIONS FOR THE SLS-2 STORAGE RING

We implement the mbtrack-cuda code in the studies of the microwave instability and the longitudinal multi-bunch instability in SLS-2. The above mentioned lattice 'dc12c' has been used in the simulations.

In the study of microwave instability, we've found that the simulations still can be accelerated by the mbtrack-cuda code since it parallelizes over macro-particles. Both the Resistive-Wall (RW) wake field and the short-range wake field of the first design of the beam position monitor (BPM) are included in the simulations of the microwave instability. We assume a round copper chamber with 10 mm inner radius along the whole ring. In the ring 150 identical BPMs are proposed. The short-range wake function used in the simulations is shown in Fig. 4, while the dependence of the bunch lengths and the energy spread on the single-bunch currents are shown in Fig. 5.

The study shows that an ideal third harmonic cavity (3HC) can help increase the microwave instability threshold. The threshold of the microwave instability is about 3 mA even though there is no third harmonic cavity and both the above

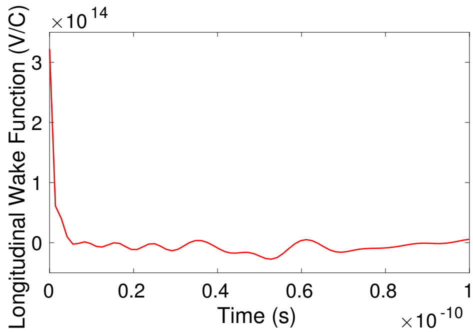

Figure 4: The total longitudinal wake function used in the simulations.

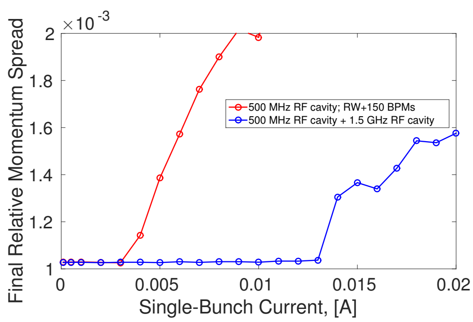

Figure 5: The single-bunch energy spread in the cases with and without the third harmonic cavity.

mentioned RW wake field and the BPM's wake field are included in the simulations (as shown by the red circle marks). Even in this case, the threshold is higher than the required single-bunch current in the bunch train (about 0.83 mA in the 390 bunches + 1 camshaft bunch mode). However, the margin is not enough in this case. In the case with the ideal bunch lengthening cavity, which is shown by the blue circles, the microwave instability threshold is about 14 mA, which is still much higher than the required current.

Since we might reuse the existing 500MHz ELETTRA-type cavities in SLS-2, we could use the information of the longitudinal HOMs (higher order modes) of the existing cavity in our study. The operation experience of this kind of cavity in SLS shows that the cooling water temperature is the key parameter for tuning away the HOMs. Therefore, we need to find the stable temperature window by calculating the growth rates of the longitudinal coupled-bunch instability (LCBI). We use both analytic method and tracking method to calculate the temperature window.

In the SLS cavity, there are nine dangerous longitudinal modes, called L1 to L9. We use the L5 mode as an example in our calculation. We first calculate the growth rate of the LCBI driven by the L5 mode in the temperature range of 25 $^\circ C$ and 45 $^\circ C$, which is shown by the green curve in Fig 6. Then, we use the mbtrack-cuda code to simulate the case using uniform filling (484 bunches) and the 3/4 filling pattern (363 bunches in the continuous buckets). The growth rate in the uniform filling pattern is shown by the red curve in the figure. Furthermore, the growth rate of the first bunch, the one in the middle, and the last bunch in the bunch train in the 3/4 filling pattern are also shown in the figure.

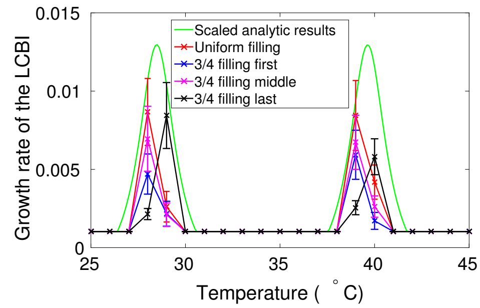

Figure 6: The comparison of the calculation of the longitudinal coupled-bunch instability driven by the longitudinal Higher-Order Mode L5 of the Cavity #3.

The results show that the analytic formula can give a quite reliable estimation of the stable temperature window. The difference along the bunch train doesn't influence the stable temperature window too much.

CONCLUSIONS

The development of GPU version of mbtrack code is presented in this paper. The GPU version uses CUDA to allow mbtrack to utilize the computing resources of GPUs for the most time consuming parts of the calculations. The CUDA version of the code has been verified against the original mbtrack code and the benchmark tests were performed to measure the performance of both codes. The mbtrack-cuda code offers a significant speedup over the CPU version for single bunch simulations and provides the possibility to run multi-bunch simulations in a stand alone CPU + GPU system in reasonable time.

The GPU version was used at PSI for the study of the longitudinal collective instabilities of the SLS-2 storage ring. The studies include both single bunch and multi bunch simulations.

ACKNOWLEDGEMENTS

The authors would like to thank Ryutaro Nagaoka and Francis Cullinan from Synchrotron SOLEIL with there help and support with mbtrack.

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n.o290605 (PSI-FELLOW / COFUND).

REFERENCES

[1] G.Skripka, R.Nagaoka, M.Klein, F.Cullinan and P.F.Tavares, Simultaneous computation of intrabunch and interbunch collective beam motions in storage rings, in *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, Volume 806, 2016

ISBN 978-3-95450-182-3