# PSB timing and scaling with Synergia

Eric G. Stern
Accelerator Simulation group/Scientific Computing Division
Fermilab, Batavia, IL USA

## Abstract

I compare the running time per turn to calculate the beam dynamics in the PS Booster with Synergia on several computational platforms. I use nodes with the two architectures of the Fermilab Wilson cluster, as well as the ALCF BlueGene/Q supercomputer. I vary the total number of nodes and cores per node to investigate the scaling behavior. I also look at how the performance changes varying the problem specification in terms of numbers of macro particles and space charge computation grid size. For comparison, I also show the the run speed of the same problem run on the CERN 48 core Spacecharge cluster, but with a different program (PTC-ORBIT)[1].

## PSB Case1 conditions:

The PSB example Case1 benchmark was run with Synergia. This simulation conditions are:

- Single RF
- 201 Space charge kicks/turn
- 2.5D Open-Hockney SC solver with grid size 128x128x128
- Input particles read from the sample input file

## Computing resources considered

A number of clusters with different hardware and architectures were considered in this study.

| Resource | Sockets/node | Cores/socket | Manufacturer | Arch | Clock GHz |
|---|---|---|---|---|---|
| Wilson Intel12 | 2 | 6 | Intel “Westmere” | X5650 | 2.67 |
| Wilson AMD32 | 4 | 8 | AMD Opteron | 6128 HE | 2.0 |
| CERN Spacecharge | 4 | 12 | AMD Opteron | 6164 HE | 1.7 |

1 Provided by F. Schmidt, CERN Space Charge Working Group

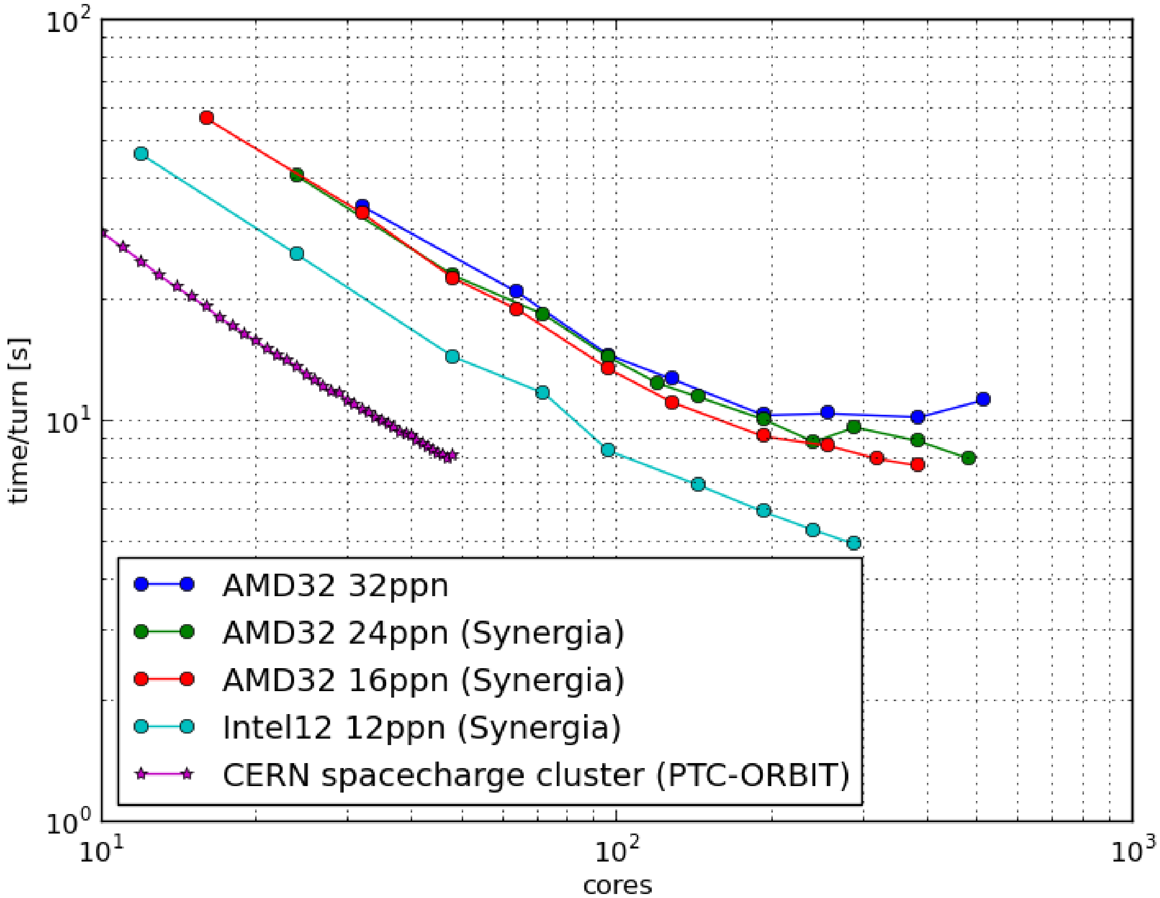


*Illustration 1: Time per turn for running the PSB Case1 example with Synergia on various node configurations on the Wilson cluster and PTC ORBIT on the CERN Spacecharge cluster.*

## PSB Case1 scaling studies

The cluster has nodes with two different architectures. The Intel12 nodes have 2.67 GHz Intel CPUs with 12 cores/node. The AMD32 nodes have 2.0 GHz processors with 32 cores/node. All the nodes have Double-Data-Rate Infiniband connections. In the legend on the plot, PPN refers to processors per node.

The Intel nodes scale very well and with the higher clock speed do in fact perform better per core. The AMD32 nodes besides being slower, clearly have some kind of network bandwidth contention going on because running with fewer active cores/node improves the scaling and the raw speed. The Intel12 nodes with only 12 cores/node don't experience the network congestion with this number of cores (and this is all the nodes we have of this type.) Both kinds of nodes scale to 200 cores with this problem. The best time of slightly more than 5 sec/turn is with 288 cores on the Intel12 nodes.

Looking at the CERN calculation, the slope of the time/turn with the number of computational cores is the same, but since they are running a different program (PTC-ORBIT), we can't directly compare the absolute run times.

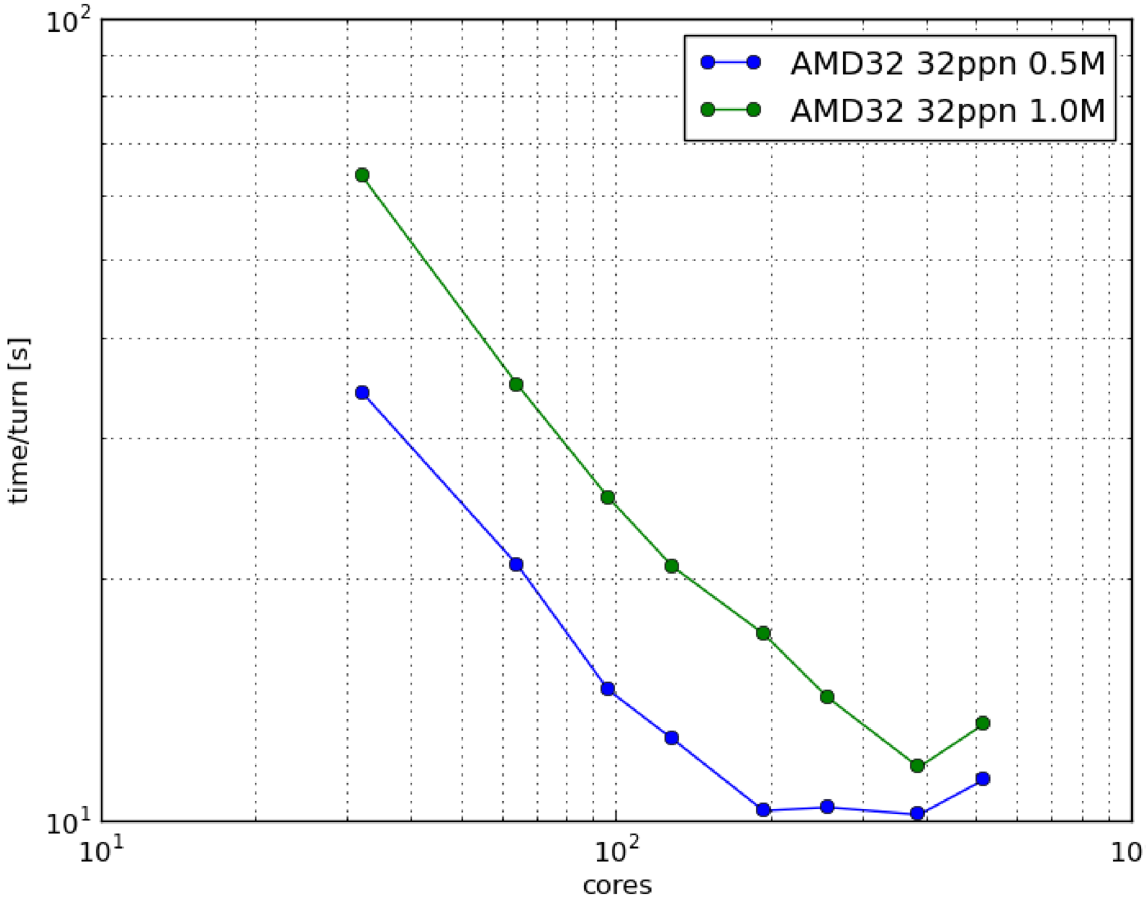


*Illustration 2: Time per turn for running the PSB Case1 example with different numbers of macroparticles.*

Changing the problem parameters changes the scaling behavior. The above plot is the scaling on AMD32 nodes for 0.5M and 1M macroparticles. The scaling with 1.0M particles is better (out to about 400 cores instead of 200 cores) because the calculation spends a smaller proportion of time in the section that experiences network congestion.

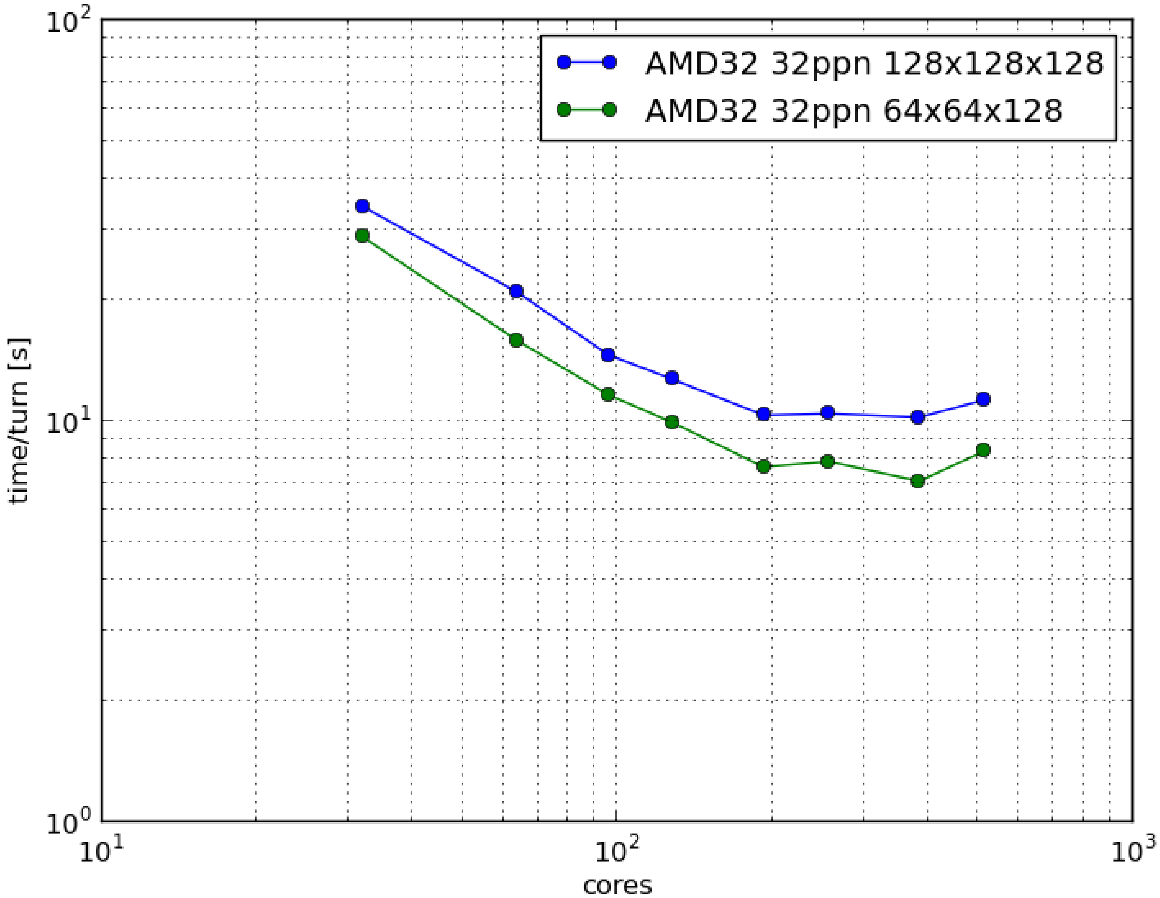


*Illustration 3: Time per turn running the PSB Case1 example varying the grid size used in the solver.*

On the AMD32 nodes, changing the SC grid size from 128x128x128 to 64x64x128 does not change the scaling limit appreciably, but the overall performance improves with the smaller grid. This is consistent with previous results showing that the majority of the time in the SC calculation is communication of the charge density and the potential over the network.

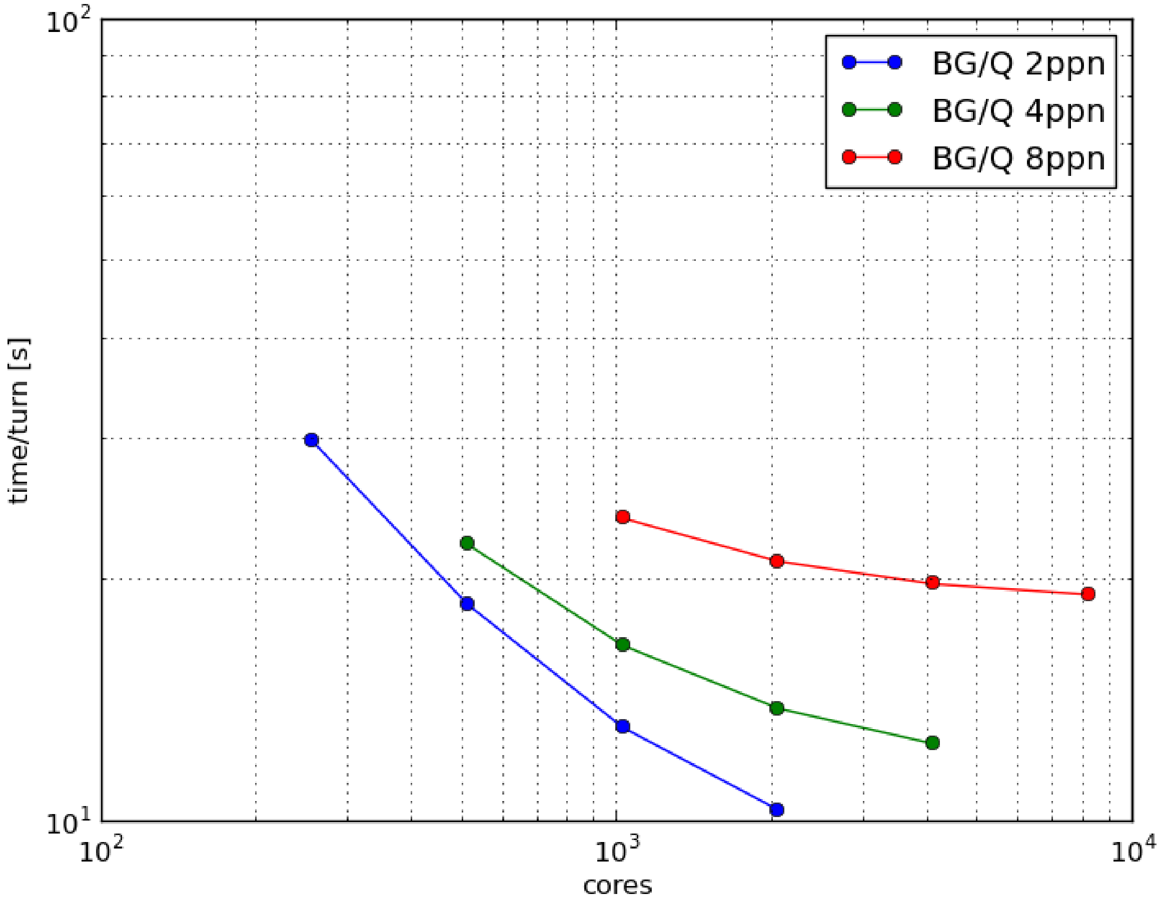


*Illustration 4: Time per turn running the PSB Case1 example on the BlueGene/Q nodes at ALCF with different numbers of cores/node.*

Scaling on the BlueGene/Q nodes for the PSB example is not good. Given the size of the computation, networking issues seem to dominate the performance. We also see significant network congestion getting data in and out of the nodes. The BlueGene/Q nodes are individually very slow. With Synergia we find that the BlueGene/Q nodes work best for problems that can be partitioned into weakly interacting pieces such as bunch-bunch coupling by wakefields. Individual bunch computations don't scale past about 4096 cores.

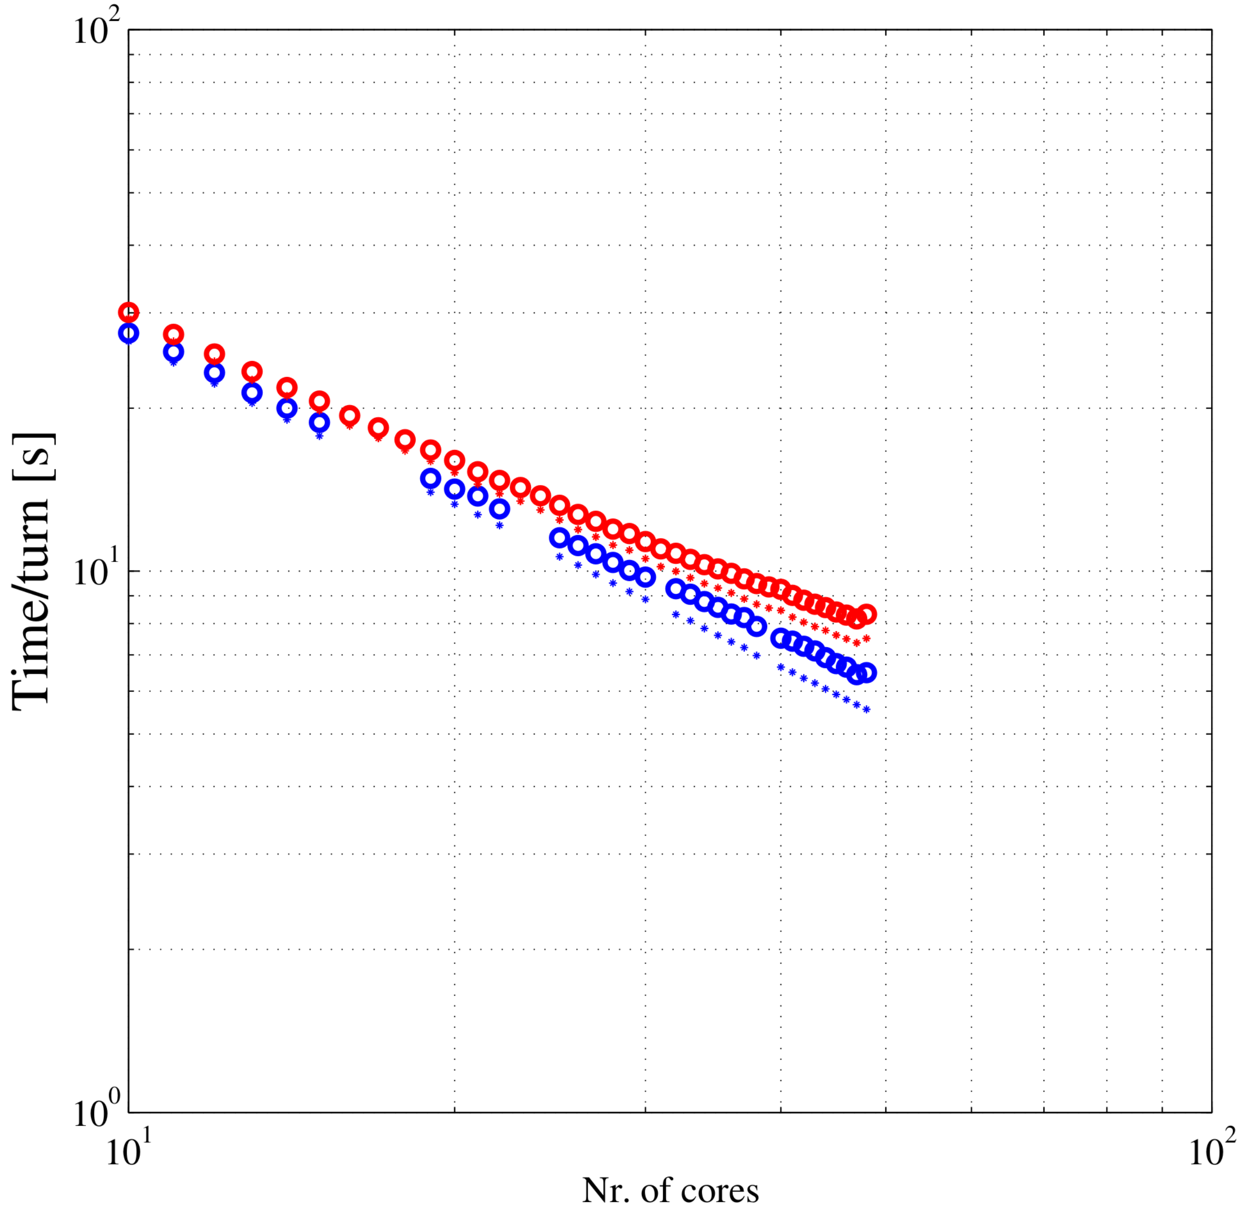


*Illustration 5: The time per turn running the PSB Case1 example on the CERN Spacechargecluster. The legend: blue is no space charge, red has space charge. The asterisks are CPUtime; circles are system time.*