Geospace simulations on the Cell BE processor

Kai Germaschewski

Space Science Center
Department of Physics
University of New Hampshire

December 15, 2008

Collaborators: Douglas Larson, Jimmy Raeder, Hartmut Ruhl

Outline

The OpenGGCM code

OpenGGCM: Global Magnetosphere Modeling

The Open Geospace General Circulation Model:

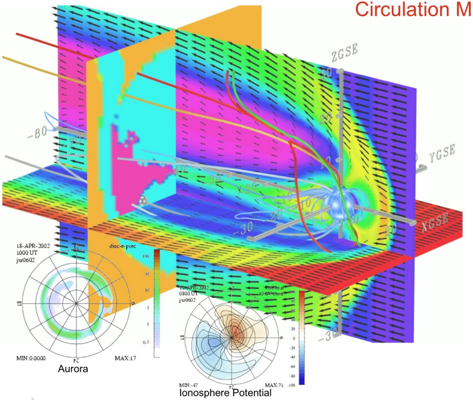

- Coupled global magnetosphere - ionosphere - thermosphere model.
- 3d Magnetohydrodynamic magnetosphere model.
- Coupled with NOAA/SEC 3d dynamic/chemistry ionosphere - thermosphere model (CTIM).
- Model runs on demand provided at the Community Coordinated Modeling Center (CCMC at NASA/GSFC
  http://ccmc.gsfc.nasa.gov/
- Will be coupled with ring current models (RCM, Fok/Jordanova models) in the near future.
- Fully parallelized code, real-time capable.
- Used for basic research, data analysis support, mission planning, space weather studies, and Space Weather Forecasting in the future.

Personnel: J. Raeder, J. Dorelli, K. Germaschewski, D. Larson, E. Kaghashvili, T. Fuller-Rowell (NOAA/SEC), F. Toffoletto (Rice U.),M.-C. Fok (GSFC), W. Li, B. Loring, T. Fogal, B. Barry, R. Vega

The OpenGGCM code

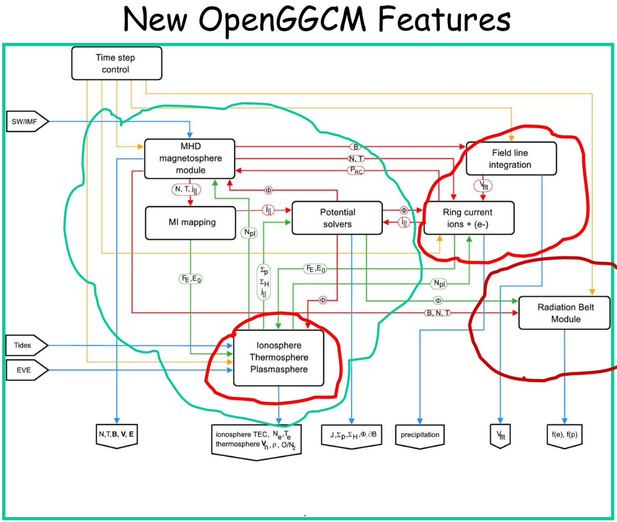

Cell Broadband Engine

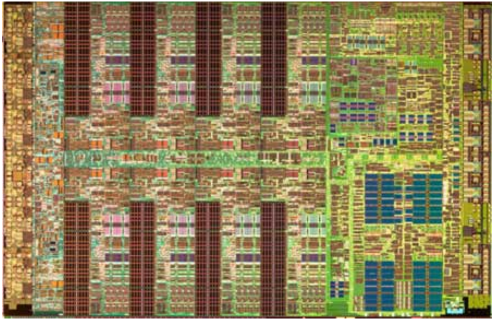

Cell BE processor

- 1 PPE (Power Processing Element)
- 8 SPEs (Synergistic Processing Elements)
- 218 GFlops

Sony PlayStation 3 with 1 Cell BE chip

PS3

- 1 Cell chip
- only 6 usable SPEs
- Cost: ~~$ 599~~ $399

Porting OpenGGCM to the Cell

Similarities to conventional architectures

- PS3/Cell Blades run Linux OS
- GNU toolchain: gcc, gfortran, gdb
- IBM provides tools / SDK for porting applications

Challenges

- Need to decompose / load-balance problem to run on the 6 SPEs
- Only 256KB local store available per SPE, data needs to be moved to/from main memory by DMA
- SIMD (single instruction, multiple data)

Problem decomposition

Numerically integrate the equations slice-by-slice.

- For an explicit scheme, the r.h.s. at (i, j, k) depends on a stencil $[i - B : i + B] \times [j - B : j + B] \times [k - B : k + B]$.
- Calculate the r.h.s. for one k-slice at a time, which needs slices $k - B ... k + B$ to be loaded into memory.

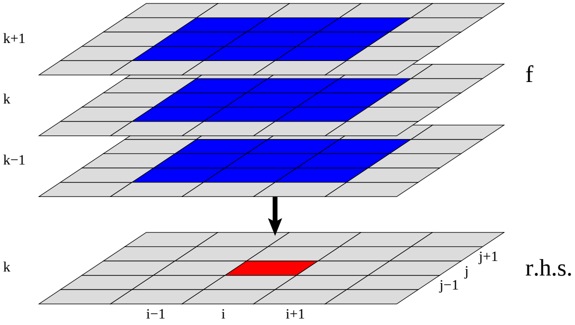

Problem decomposition

- Once the r.h.s. on a slice k is done, proceed to the next slice $k+1$, storing the just calculated slice k, expiring the current $k-B$ fields and loading the next $(k+B+1)$ slice of field data.

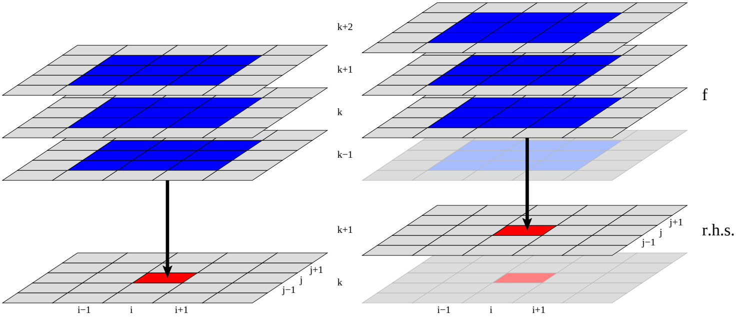

Problem decomposition

For MHD, need to load 8 fields $\times$ 5 slices of field data, calculating 8 fields of output data, plus double-buffering and temporary storage.

- It is not possible to fit whole *xy*-slices into local store.
- Need to decompose into smaller units in order to utilize all SPUs.

Solution: Domain-decompose in the *xy* domain!

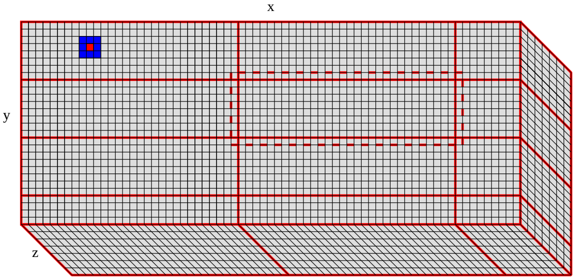

Double buffering

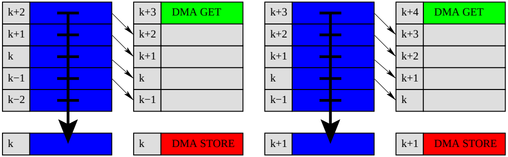

Hide DMA latencies by overlapping DMA and computation

- Preload $k + B + 1$-slice while calculating on $k - B \dots k + B$.
- Continue calculation in different buffer while storing k-slice.

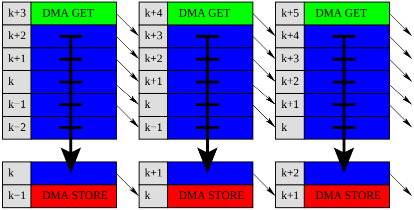

Implementation

Objectives

- Maintability and Readability (KISS).
- Adapts to different architectures to provide high performance.

One codebase can be compiled into SIMD or scalar code, as well as *xy*-domain decomposed or stepping through the entire domain at once.

```
slice_for_each(fi, 0, 0) {
  F3_(rr1  ,0,0,0) += dt * F3(rrr  ,0,0,0);
  F3_(uu1  ,0,0,0) += dt * F3(uur  ,0,0,0);
  F3_(rv1.x,0,0,0) += dt * F3(rvr.x,0,0,0);
  F3_(rv1.y,0,0,0) += dt * F3(rvr.y,0,0,0);
  F3_(rv1.z,0,0,0) += dt * F3(rvr.z,0,0,0);
} slice_for_each_end;
```

Implementation

curr()

```
// ----------------------------------------------------------------
// curr
//
// compute current xj from magnetic field b

static inline void
curr(vfld_t xj, vfld_t b, crd_t bdx4, int ix, int iy, int iz)
{
  F3_(xj.x,0,0,0) =
    (F3(b.z,0,1,0) - F3(b.z,0,0,0)) * CRDy(bdx4,0) -
    (F3(b.y,0,0,1) - F3(b.y,0,0,0)) * CRDz(bdx4,0);
  F3_(xj.y,0,0,0) =
    (F3(b.x,0,0,1) - F3(b.x,0,0,0)) * CRDz(bdx4,0) -
    (F3(b.z,1,0,0) - F3(b.z,0,0,0)) * CRDx(bdx4,0);
  F3_(xj.z,0,0,0) =
    (F3(b.y,1,0,0) - F3(b.y,0,0,0)) * CRDx(bdx4,0) -
    (F3(b.x,0,1,0) - F3(b.x,0,0,0)) * CRDy(bdx4,0);
}
```

OpenGGCM timing results

OpenGGCM test case: 128x64x64 stretched grid

- Timestep Δt in the simulation: $\Delta t = 0.599$
- Wallclock time for advancing by Δt on PPE: 4.350
- Wallclock time for advancing by Δt using 6 SPEs: 0.176

Using 6 SPEs, OpenGGCM is running 3x faster than realtime!

Scaling with number of SPEs used

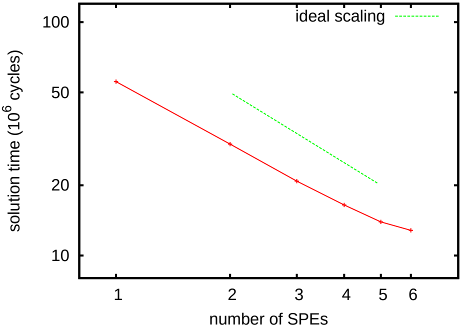

job	percent	cycles	calls	cycles/call
PRIMVAR	2.23	442717485	3840	115291
PREP1	5.70	1131278425	4096	276191
PRED	14.42	2859987371	3840	744788
CORR	8.61	1707791789	3840	444737
PUSH_AXPY	2.88	571812076	3840	148909
read_mbox	66.15	13117982581	26758	490245

page 1UL0201A0623 CMYK

Page A2 • NEW HAMPSHIRE UNION LEADER, Monday, June 23, 2008

UNH's supercomputer to predict 'space weather'

◆**Advanced math:** By combining 40 PS3 gaming consoles, the institute created a computer that, in theory, can perform 8 trillion calculations per second.

From left, University of New Hampshire researchers Kai Germaschewski, Andrew Foulks, Joachim Raeder and Doug Larson show off the 40 Playstation 3 game consoles they linked to form a supercomputer.

COURTESY

By CLYNTON NAMUO
Union Leader Correspondent

DURHAM — Video-game nerds rejoice: Researchers at the University of New Hampshire have bundled 40 Playstation3 game consoles to form a supercomputer they will use to help predict "space weather."

Associate physics professor Jimmy Raeder said he and other researchers at UNH's Institute for the Study of Earth, Oceans and Space have combined the gaming systems into a supercomputer that, theoretically, can perform 8 trillion calculations a second.

"The gaming industry is insatiable," Raeder said. "The more computational power they have, the more realistic they can make the games."

Raeder and his associates will use the supercomputer's vast power to study what's commonly referred to as space weather, or the sun's interaction with the Earth's magnetosphere. The sun periodically produces solar flares and solar wind that interacts with earth in many ways and can interfere with satellites, ground communication and even cause power outages, he said.

"Satellites can die because of space weather effects," he said.

The supercomputer will also be used to study Aurora, known also as the Northern or Southern Lights depending on where they occur. Raeder said researchers eventually hope to be able to predict space weather much in the same way meteorologists predict regular weather.

And like regular weather, the study of space weather requires increasingly complex calculations. This is where the PS3-powered supercomputer comes in handy, Raeder said.

The console, introduced during 2006 holiday shopping season, is well-known in gaming circles for its cutting-edge graphics, which is made possible by an advance computer chip designed specifically for the system called the cell broadband engine.

Raeder said the chip itself is the key to the system's performance. By combining 40 PS3s, the UNH researchers have created a supercomputer that, in theory, can perform 8 trillion, or 8 thousand billion, calculations per second. Take that H&R Block.

That computing power pales in comparison to a recently announced supercomputer called "Roadrunner," created by scientists at the Los Alamos National Laboratory in New Mexico. That computer combined hundreds of PS3 chips to form the fastest supercomputer on earth.

Raeder said Roadrunner is the first supercomputer to achieve a petaflop, or one quadrillion calculations per second. That's 10 to the 15th power or one thousand trillion calculations per second.

"It's just absolutely mind blowing," Raeder said.

Roadrunner cost about $133 million, but Raeder and his associates spent only $24,000, including the cost of the systems and the parts necessary to combine them, to construct their scaled-down version. The UNH endeavor is being funded with a four-year $1.5 million National Science Foundation grant.

UNH's PS3-driven supercomputer has enough computing power to match the UNH institute's other supercomputer, which weighs 8,000 pounds and cost $750,000. Plus, even with the modifications, the PS3 machines can still play video games, Raeder said.

At left, Republican presidential candidate Sen. John McCain, R-Ariz., walks to his waiting SUV after worship services at the North Phoenix Baptist Church in Phoenix, Ariz., yesterday.

At right, Democratic presidential candidate Sen.

Tilton man dies after car accident

LACONIA — A Tilton man who was involved in a single-car accident Saturday has died, police said yesterday.

Howard Cilley, 53, of 7 Pine Hill Road was hurt when his silver Nissan

Scaling with number of PS3's

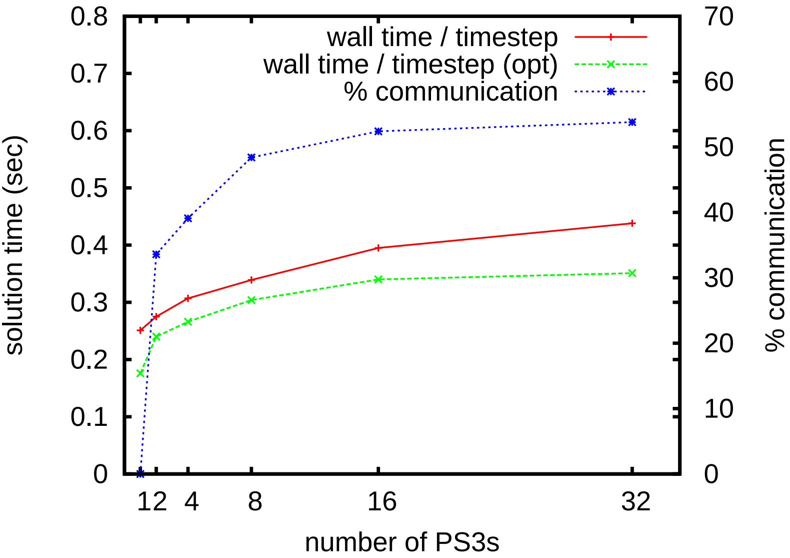

Particle-in-Cell

- Governing equations: Vlasov-Boltzmann + Maxwell
  Important for magnetic reconnection as well as laser-plasma interaction
- PIC method: solve by introducing metaparticles, collisions are challenging (but done and working).

$$\begin{aligned} \dot{\vec{x}} &= \vec{v} \\ \dot{\vec{v}} &= \vec{F}/m \end{aligned}$$

- Performance: Advance field and push metaparticles (1000 particles per cell)
- Particle pusher on Cell: Using SPUs we obtained a speed-up of a factor of 60 over using main processor (PPU) only.

Summary

- The Cell processor has enormous potential to accelerate PDE codes.
- Code generation is a key technology in putting things together while keeping codes manageable

Outlook

- Continue bringing technologies together:
  - Adaptive mesh refinement
  - 3D XMHD implicit time integration
  - Getting the most out of modern HPC architectures
- Particle-in-Cell on Cell