

Accelerating Climate Models with the IBM Cell Processor

Shujia Zhou

NASA GSFC SIVO

Northrop Grumman Corporation/TASC

Collaborators: Tom Clune, Dan Duffy, Samuel
Williams, Max Suarez, Milton Halem

Shujia Zhou, 2008


Outline

- Climate (e.g., NASA GEOS5) models
 - Column physics
 - Port the solar radiation code
 - Dynamics
 - Analyze the finite volume dynamics core
 - 4D Var data assimilation
 - Be very compute-intensive

Background


- NASA is interested in the potential performance and cost benefits of adapting some science applications to emerging nontraditional processors such as the IBM Cell
 - Motivation
 - Increase performance by one to two orders of magnitude over traditional processors.
 - Challenges
 - SPE's small local memory (256 KB)
 - A low-level communication mechanism
 - Direct memory address management

The IBM Cell Processor



- 205 single-precision GFLOPS
- High-speed data ring (EIB) with a sustained bandwidth of 205GB/s
- 25.6 GB/s processor-to-memory bandwidth
- 256KB local store at SPE


Processes in a Climate Model



<http://www.ucar.edu/communications/CCSM/overview.html>

Shujia Zhou, 2008

NASA GEOS 5 Code Structure




Climate and Weather Models

- Constraints
 - A few hundred thousand lines of code written in Fortran over 20+ years
 - Some modularity in F90/95
 - Still evolving
 - Production requirement
 - Cannot rewrite completely
 - Minimal intrusion
 - Good ratio of performance to cost
- Solutions
 - Select the computationally intensive model components
 - I/O is smaller compared to computation
 - The number of lines of code is manageable

Porting Strategies for A Cell Processor


- To simplify porting, put the calculations involving dependency into one SPE
 - Take in extra data to make it self-contained
 - The communication cost for extra data should be smaller than its calculation cost
- Minimize the intrusion to the original code as much as possible

Data Decomposition for Cell Streaming Model: Climate and Weather Applications



Port the Solar Radiation Code to IBM Cell Processor


Global Circulation Model



Main Memory

Data Transfer via DMA

Local Store



The solar radiation component of NASA's Goddard Earth Observing System Model, Version 5 (GEOS-5) was chosen to evaluate the Cell's programming paradigm due to these factors:


- One of the most computationally intensive parts of GEOS-5, at least 20% (including infrared radiation)
- The time required for I/O is much smaller than for numeric computation
- The independent vertical column calculation greatly simplifies parallel programming

Flow Diagram of Data Transfer via DMA (4 SPEs)

Main memory (512MB)

Multiple arrays





- array1D
- array2D
- array3D




Global Array Index

PPE

Thread




cycles




LS (256KB)

1	5	9	13
2	6	10	14
3	7	11	15
4	8	12	16


SPE0



SPE1



SPE2




SPE3

Local Array Index

Shujia Zhou, 2008

Detailed Memory Analysis




Profile of The Cell-Version Solar Radiation Code

Function call	DMA get		soluv	solir	cldscale	deledd	cldflx
Time (usec)	40		25186	143308	77	99	4551
iteration	1	soluv	1		1	20 = 5x4	5
		solir		1	3	120= 3x10x4	30= 3x10
		total	1	1	4	140	35

8 SPEs, 1024 columns, single precision of deledd(), no vectorization

Performance

IBM BladeCenter QS20	Itanium 2	Xeon	
single precision, vectorization and unrolling of deled() and cldfx()	single precision	single precision	
1 PPE + 8 SPEs PPE: 64-bit PowerPC, 3.2GHz SPE: 128-bit SIMD, 3.2 GHz	1 core of Intel Itanium2, 1.5 GHz, with 4MB Cache	1 core of Intel Xeon (dempsey), 3.2 GHz, with 4 MB cache (2MB per core)	1 core of Intel Xeon (woodcrest), 2.66 GHz, with 4 MB cache (2MB per core)



Shujia Zhou, 2008


Performance

- For 1024 columns, **latest results** show that the Cell ran the C-version code **8.8, 11.6, 12.8** times faster than a core on Intel's Woodcrest, Dempsey, and Itanium2 processors for the baseline code (Fortran, single-precision), respectively, with
 - Vectorization
 - Unrolling (~11% improvement)
 - IBM XL SPU C compiler (~20% improvement over gcc)

Accelerate

The Key Climate and Weather Community Component Cubed-Sphere Finite Volume Dynamics (CSFVD)

- CSFVD becomes next-generation climate models and weather forecast systems of NASA GMAO, NOAA GFDL, and NSF NCAR
- CSFVD takes at least 25% of total computing time
- Our analysis indicates that CSFVD can use Cell's streaming programming model to accelerate



Shujia Zhou, 2008

Computing Requirement for 4D-Var

- The 4D Var data assimilation system better utilize satellite observations and consequently improve forecasting skill
 - Typically request **hundreds** of iterations of the linearized forecast model and its adjoint for **each** forecast run.
 - Require significantly shorten the execution time of the model to satisfy the **operational** requirement.

Summary

- We have demonstrated that IBM Cell technology can dramatically accelerate the physics components of climate and weather applications
- Our analysis shows that the dynamics component should be amenable to accelerate
- Data assimilation is extremely compute-intensive, which desires extraordinary acceleration

Acknowledgements

- We would like to thank
 - Tsengdar Lee (NASA High End Computing Program) for funding support
 - Mike Seablom for inspiration and helpful discussion
 - Bruce Van Aartsen and Carlos Cruz (NASA-GSFC/Northrop Grumman) for assisting with code translation from Fortran to C
 - NASA NCCS for installing the IBM Cell Simulator for code development
 - The Dice Project for training support
 - The UMBC Multicore Computational Center for providing access to the IBM Bluegrit system as a hardware test environment for our research