27th CESM workshop June 15, 2022

Enabling the execution of PUMAS on GPUs

Jian Sun, John Dennis, Sheri Mickelson, Brian Vanderwende, Andrew Gettelman, Katherine Thayer-Calder





This material is based upon work supported by the National Center for Atmospheric Research, which is a major facility sponsor ed by the National Science Foundation under Cooperative Agreement No. 1852977.

Outline

- ✤ What is cloud microphysics?
- Code overview of cloud microphysics scheme in CAM (PUMAS)
- Methodology
- Preliminary results & discussion
- Summary & future work



What is cloud microphysics?

"...small-scale (from sub-micron to cm) processes driving the formation and evolution of cloud and precipitation particles..."





(Morrison et al., 2020)



✤ PUMAS:

- micro_pumas_v1.F90: ~2900 lines, 7 subroutines/functions
- micro_pumas_utils.F90: ~2000 lines, 54 subroutines/functions
- ✤ CAM:
 - wv_sat_methods.F90: ~400 lines, 29 subroutines/functions
 - wv_saturation.F90: ~700 lines, 30 subroutines/functions
 - 33 additional CAM codes related to wv_* F90
- CAM physics source codes: ~410,000 lines
 - PUMAS represents ~1.2% of CAM physics codes
 - PUMAS contributes to ~8% of computational time of CAM physics

Offload to GPU

- Start from a KGen kernel (easy for code development and debug)
- Use OpenACC and OpenMP offload to port CPU codes to GPU
 - Directives-based parallel programming model \rightarrow Single source code \checkmark
 - Convert CPU codes to GPU codes by adding pragmas (Readability √)
 - Intel's auto-migration tool from OpenACC to OpenMP offload (https://github.com/intel/intel-application-migration-tool-for-openacc-to-openmp)

Code size of OpenACC version (similar for OpenMP offload) **◆** PUMAS:

- micro_pumas_v1.F90: ~2900 lines \rightarrow ~3200 lines
- micro_pumas_utils.F90: ~2000 lines \rightarrow ~2200 lines
- ✤ CAM:

NCAR

- wv_sat_methods.F90: ~400 lines → ~460 lines
- wv_saturation.F90: ~700 lines \rightarrow ~790 lines

Increase the code size by ~10%



- Does the GPU version of PUMAS/CAM codes return the bit-for-bit results compared with the CPU version of codes?
 - ➢ If yes, that is great!
 - If no, we need to ask ourselves "Is this difference expected?"
 - If yes, run a verification test (e.g., ECT, AMWG diagnostics package, etc)
 - If no, it could be something we do not understand fully (likely) or a code bug (more likely)
- ✤ Always check the correctness before looking at the performance.



Performance of KGen kernel

- One CPU node (CPU run) vs. one MPI rank + one GPU (GPU run)
- Log scale for X- and Y-axis
- GPU performs worse than CPU consistently



Number of columns per node or GPU

Profiling KGen kernel

• Example of computational hotspots in the GPU code: loops with dependency

Original implementation



• Solution: reverse the loop order



Profiling KGen kernel (cont'd)

- Example of computational hotspots in the GPU code: sedimentation subroutines run in serial
- Solution: run them asynchronously

Original implementation

New implementation

Lice	l ice
call Sedimentation(moncol.nlev.do cldice.deltat.fi.fni.pdel inv. &	call Sedimentation(moncol.nlev.do cldice.deltat.fi.fni.pdel inv. &
gitend.nitend.gisedten.dumi.dumni.prect.iflx. &	aitend.nitend.aisedten.dumi.dumni.prect i.iflx IQUEUE &
xxlx=xxls.gxsevap=gisevap.tlat=tlat.gvlat=gvlat. &	xxlx=xxls.gxsevap=gisevap.tlat=tlat i.gvlat=gvlat i. &
xcldm=icldm.preci=preci)	xcldm=icldm.preci=preci i)
! lig	! lig
call Sedimentation(moncol.nlev.,TRUE.,deltat.fc.fnc.pdel inv. &	call Sedimentation(moncol.nlev.TRUE.deltat.fc.fnc.pdel inv. &
actend.nctend.acsedten.dumc.brect.lflx. &	actend.nctend.acsedten.dumc.prect 1.1flx LQUEUE &
xxlx=xxlv, gxsevap=gcsevap,tlat=tlat,gvlat=gvlat,xcldm=lcldm)	xxlx=xxlv,qxsevap=qcsevap,tlat=tlat l,qvlat=qvlat l,xcldm=lcldm)
! rain	! rain
call Sedimentation(moncol,nlev,.TRUE,,deltat,fr,fnr,pdel inv, &	call Sedimentation(mgncol,nlev,.TRUE,,deltat,fr,fnr,pdel inv, &
artend, nrtend, arsedten, dumr, dumn, prect, rflx)	artend, nrtend, arsedten, dumr, dumnr, prect r, rflx RQUEUE
! snow	! snow
call Sedimentation(moncol,nlev,.TRUE,,deltat,fs,fns,pdel inv, &	call Sedimentation(moncol,nlev,.TRUE,,deltat,fs,fns,pdel inv, &
astend, nstend, assedten, dums, dumns, prect, sflx, preci=preci)	astend, nstend, assedten, dums, dumns, prect s, sflx SQUEUE preci=preci s)
! graupel	! graupel
call Sedimentation(mgncol,nlev,.TRUE.,deltat,fg,fng,pdel inv, &	call Sedimentation(mgncol,nlev,.TRUE.,deltat,fg,fng,pdel inv, &
<pre>qgtend,ngtend,qgsedten,dumg,dumng,prect,gflx,preci=preci)</pre>	qgtend, ngtend, qgsedten, dumg, dumng, prect_g, gflx GQUEUE preci=preci_g)



Performance of KGen kernel (new implementations)



Enabling the execution of PUMAS on GPUs

NCAR

After some optimizations, we find:

- The performance of OpenACC and OpenMP offload is largely improved, which the CPU performance is marginally affected
- The performance of OpenACC and OpenMP offload is competitive, but OpenACC is better in general
- GPU outperforms CPU for large problem size

CAM configuration

- Test configurations:
 - Compset: F2000dev
 - Cloud microphysics: MG3 configuration
 - Dycore: FV
 - Resolution: 1 degree for FV
 - Simulation length: 1 day
 - Machine: Cheyenne (CPU), Casper (CPU and GPU)
 - Resource: one compute node with 36 CPU cores, one V100 GPU per node
 - PCOLS: different data sizes offloaded to GPU per time

Performance of PUMAS in CAM



• GPU-enabled PUMAS can outperform its CPU version in a practical CAM simulation

Computation vs. data movement in a GPU run



• Data transfer could be more time-consuming than computation



Summary

- What we have done/learned:
 - Port PUMAS to GPU by OpenACC and OpenMP offload
 - Evaluate the performance on CPU and GPU
 - > Even one GPU per node is showing promising results
 - > OpenMP offload is relatively immature compared to OpenACC \rightarrow more evaluations needed
- What is next?

NCAR

- Multiple GPUs per node
- High-resolution CAM simulation
- Sensitivity tests with different PUMAS configuration
- > One manuscript is under preparation

Acknowledge

- This work is funded by the following project:
 - ➢ NSF CSSI EarthWorks (Award number: 2005137)
 - NSF NCAR-base funding
- Many thanks to contributions/helps from different labs/organizations:
 - CISL: Supreeth Madapur Suresh, Cena Miller, Allison Baker, Daniel Howard
 - CSEG: Jim Edwards
 - CGD: Cheryl Craig, Steve Goldhaber
 - Areanddee: Rich Loft
 - NVIDIA: Raghu Raj Prasanna Kumar
 - INTEL: Harald Servat

NCAR

UCAR

ORNL: Youngsung Kim