

Visuals on the House: Optimizing HPC Workflows with No-Cost CPU Visualization

Victor A. Mateevits^{†¶}, Andres Sewell^{*†}, Mathis Bode[‡], Paul Fischer^{##†}, Jens Henrik Göbbert[‡], Joseph A. Insley^{†¶}, Ioannis Kavroulakis[§], Damaskinos Konioris[§], Yu-Hsiang Lan[#], Misun Min[†], Dimitrios Papageorgiou[§], Michael E. Papka^{†¶}, Steve Petruzza^{*}, Silvio Rizzi[†], Ananias Tomboulides[§]

[†]Argonne National Laboratory, ^{*}Utah State University, [‡]Forschungszentrum Jülich, [§]Aristotle University of Thessaloniki, [¶]University of Illinois Chicago, ^{||}Northern Illinois University, [#]University of Illinois Urbana-Champaign

Abstract

The rise of heterogeneous resources in modern HPC systems has driven the scientific community beyond the exascale threshold. However, relying on GPUs for simulations often leaves CPUs underutilized. *In situ* techniques reduce data movement by operating in-memory but still involve blocking operations. We propose copying GPU-based timestep data to host memory, allowing CPUs to **concurrently handle visualization and analysis**, thus enabling simulations to continue without interruption.

Implementation

- Inline Visualization and Analysis (GPU Blocking)
- Concurrent Visualization and Analysis (sim on the GPU, vis on the CPU)

```
// Thread pool to manage non-blocking threads
ThreadPool::ThreadPool(size_t numThreads) : stop(false), activeTasks(0) {
    for (size_t i = 0; i < numThreads; ++i) {
        workers.emplace_back(6*ThreadPool::workerThread, this);
    }
}

ThreadPool::~ThreadPool() {
    std::unique_lock<std::mutex> lock(queueMutex);
    stop = true;
    condition.notify_all();
    for (std::thread &worker : workers) {
        worker.join();
    }
}

void ThreadPool::enqueue(std::function<void*> task) {
    int taskID = taskIDCounter++;
    std::unique_lock<std::mutex> lock(queueMutex);
    tasks.push(std::move(task));
    ++activeTasks;
    condition.notify_one();
}

void ThreadPool::waitForCompletion() {
    std::unique_lock<std::mutex> lock(queueMutex);
    completionCondition.wait(lock, [this]() { return tasks.empty() && (activeTasks == 0); });
}

void ThreadPool::workerThread() {
    while (true) {
        std::function<void*> task;
        {
            std::unique_lock<std::mutex> lock(queueMutex);
            condition.wait(lock, [this]() { return stop || !tasks.empty(); });
            if (stop && tasks.empty()) {
                return;
            }
            task = std::move(tasks.front());
            tasks.pop();
        }
        task();
        --activeTasks;
        completionCondition.notify_all();
    }
}

int ThreadPool::getTaskIDCounter() const {
    return taskIDCounter.load();
}
```

Use Case


The GABLS1 benchmark simulates an atmospheric boundary layer (ABL) driven by a uniform geostrophic wind and a prescribed surface, and is used for weather forecasts, climate models, and our understanding of atmospheric phenomena.

Next Steps


- Preliminary results are promising
- Run at scale on ALCF and JSC resources: 70 (280 GPUs), 140 (560 GPUs), 280 (1120 GPUs) nodes planned.

Acknowledgements

This research used resources of the Argonne Leadership Computing Facility, a U.S. Department of Energy (DOE) Office of Science user facility at Argonne National Laboratory and is based on research supported by the U.S. DOE Office of Science-Advanced Scientific Computing Research Program, under Contract No. DE-AC02-06CH11357. This work was supported by Northern Illinois University. The authors from JSC acknowledge computing time grants for the project TurbulenceSL by the JARA-HPC Vergabegremium provided on the JARA-HPC Partition part of the supercomputer JURECA at Jülich Supercomputing Centre, Forschungszentrum Jülich, the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time on the GCS Supercomputer JUWELS at Jülich Supercomputing Centre (JSC), and funding from Innovative Algorithms for Applications on European Exascale Supercomputers (Inno4Scale) program as well as the Center of Excellence for Exascale CFD (CEEC) funded by the EU. Support by the Joint Laboratory for Extreme Scale Computing (JLESC, <https://jlesc.github.io/>) for traveling is acknowledged.



The GABLS case (MFEV/SMG model) simulated with NekRS, showcasing the turbulence fine structures in the X-Z plane at $y = 100$ m.



Slices of potential temperature in the GABLS case, highlighting the stratification of the Atmospheric Boundary Layer

References

- [1] A. C. Bauer, J. Abbasi, J. Ahrens, H. Childs, B. Geveci, S. Klasky, K. Moreland, P. O'Leary, V. Vishwanath, B. Whitlock, et al. In situ methods, infrastructures, and applications on high performance computing platforms. In *Computer Graphics Forum*, vol. 35, pp. 577–597. Wiley Online Library, 2016. 1
- [2] R. J. Beare, M. K. MacVean, A. A. M. Holtslag, J. Cuxart, I. Esau, J. C. Golaz, M. A. Jimenez, M. Khairoutdinov, B. Kosovic, D. Lewellen, T. S. Lund, J. K. Lundquist, A. McCabe, A. F. Moene, Y. Noh, S. Raasch, and P. Sullivan. An Intercomparison of Large-Eddy Simulations of the Stable Boundary Layer. *Boundary-Layer Meteorology*, 118(2):247–272, Feb. 2006. doi: 10.1007/s10546-004-2820-6 2
- [3] N. El-Sayed and B. Schroeder. To checkpoint or not to checkpoint: Understanding energy-performance-i/o tradeoffs in hpc checkpointing. In 2014 IEEE International Conference on Cluster Computing (CLUSTER), pp. 93–102. IEEE, 2014. 1
- [4] M. Isakov, E. d. Rosario, S. Madireddy, P. Balaprakash, P. Carns, R. B. Ross, and M. A. Kinsky. Hpc i/o throughput bottleneck analysis with explainable local models. In SC20: International Conference for High Performance Computing, Networking, Storage and Analysis, pp. 1–13, 2020. doi: 10.1109/SC41405.2020.00037 1
- [5] B. Kosovic and J. A. Curry. A large eddy simulation study of a quasi steady, stably stratified atmospheric boundary layer. *Journal of the Atmospheric Sciences*, 57(8):1052 – 1068, 2000. doi: 10.1175/1520-0469(2000)057<1052:ALESSO>2.0.CO;2 2