

For computer time, this research used the resources of the Supercomputing Laboratory at KAUST.

MOTIVATIONS

Extreme thread concurrency for parallel performance and data locality for traffic reduction associated with dynamic core frequencies, portend the end of the uniform bulk synchronous programming model. It is paramount to design new algorithms, programming models, and amenable software tools in order to relax the synchrony of existing algorithms and their implementations. Many standard means of expressing algorithms, including loops and function calls, introduce artifactual ordering and synchronization from which future software must be freed.

The industrial project codenamed **ExaWave** addresses these challenges for seismic modeling applications by introducing:

1. Fine-grained computations to exploit the thread concurrency and maximize strong scalability on a single shared-memory node
2. Asynchronous execution to cope with the dynamic core frequencies and bus bandwidth saturation while mitigating the load imbalance overhead
3. Data motion reduction techniques for the stencil time integration.


Beside the traditional spatial blocking (SB) technique for stencil computations, we integrate the Multicore-optimized Wavefront Diamond (MWD) tiling approach [1, 2], which takes into account the three aforementioned ingredients, into the **ExaWave** software framework.

However, the fine-grained asynchronous handling of the Convolutional Perfectly Matched Layer (CPML) [3] for absorbing boundary conditions turns out to be a challenging open research problem, due to severe inherent data dependencies constraints, which impedes MWD performance impact.

We combine techniques of loop fusion to reduce memory traffic and sliding windows to cut down the engendered extra flops, in order to consolidate CPML integration with the overall asynchronous MWD technique.

THE EXAWAVE PROJECT


- Reverse Time Migration
- Finite Difference Time Domain on Staggered Grid [4]
- Stencil computations with SB and MWD [1, 2]
- CPML Absorbing Boundary Conditions [3]
- Programming models: Hybrid MPI+OpenMP




- Solving the 2nd order formulation of the 3D acoustic wave equation with constant density: $\frac{1}{c^2} \frac{\partial^2 p}{\partial t^2} - \Delta p = s$
- Discretizing with 8th order in space and 2nd order in time

STENCIL COMPUTATIONS

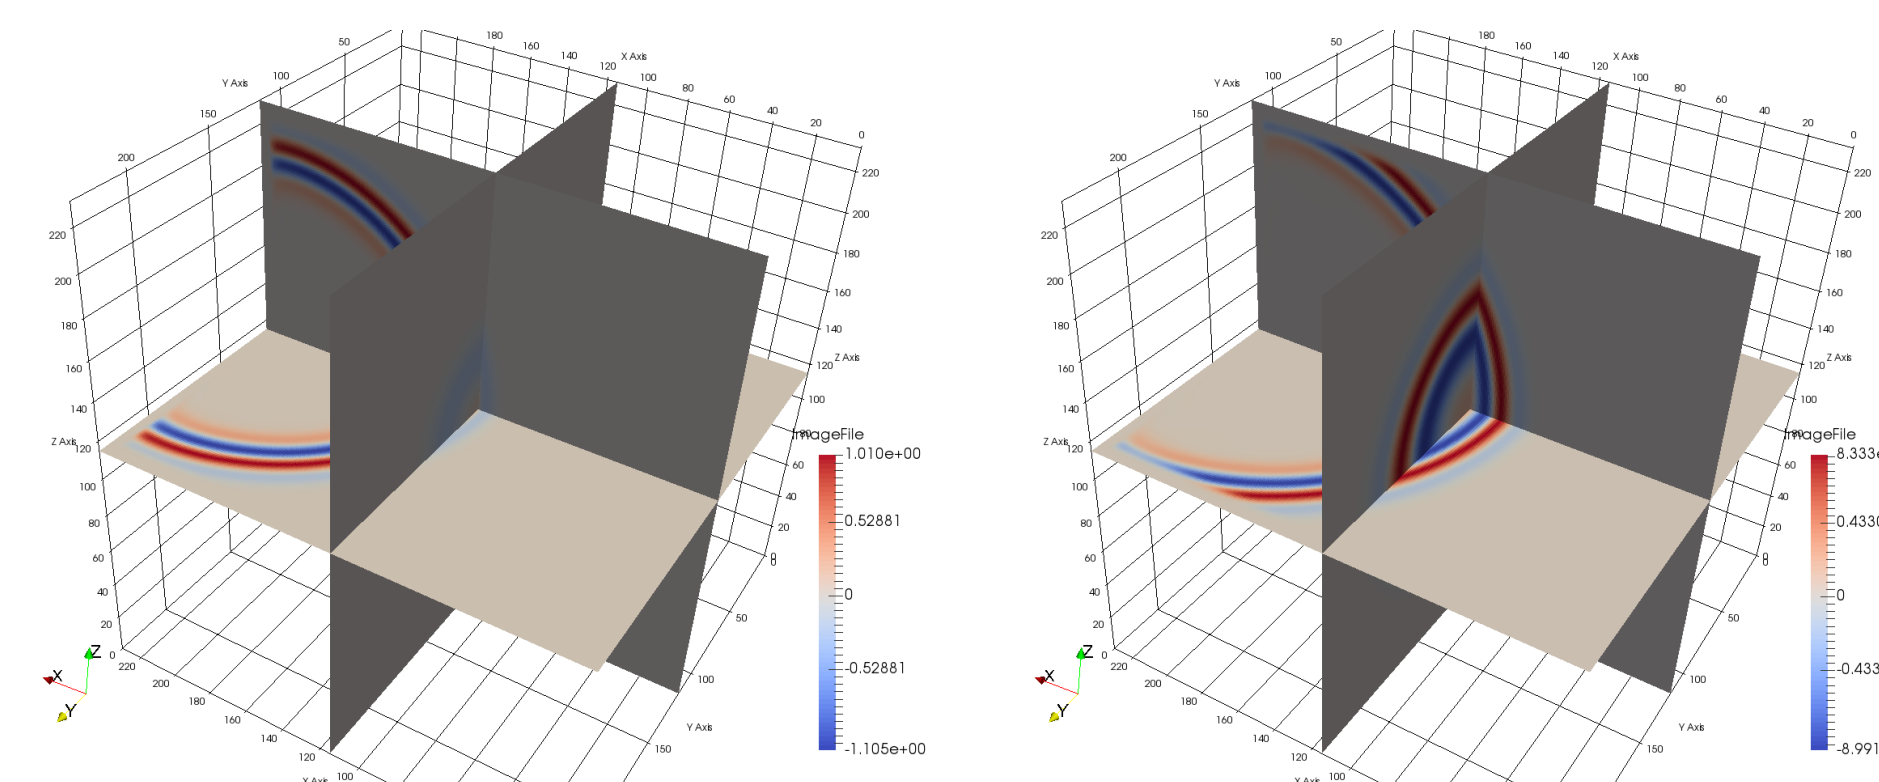
- Z dimension as fastest index w/ SIMD vectorizations
- Reference SB: block spatial points in X and Y dimensions



- MWD [1, 2] with asynchronous fine-grained computations

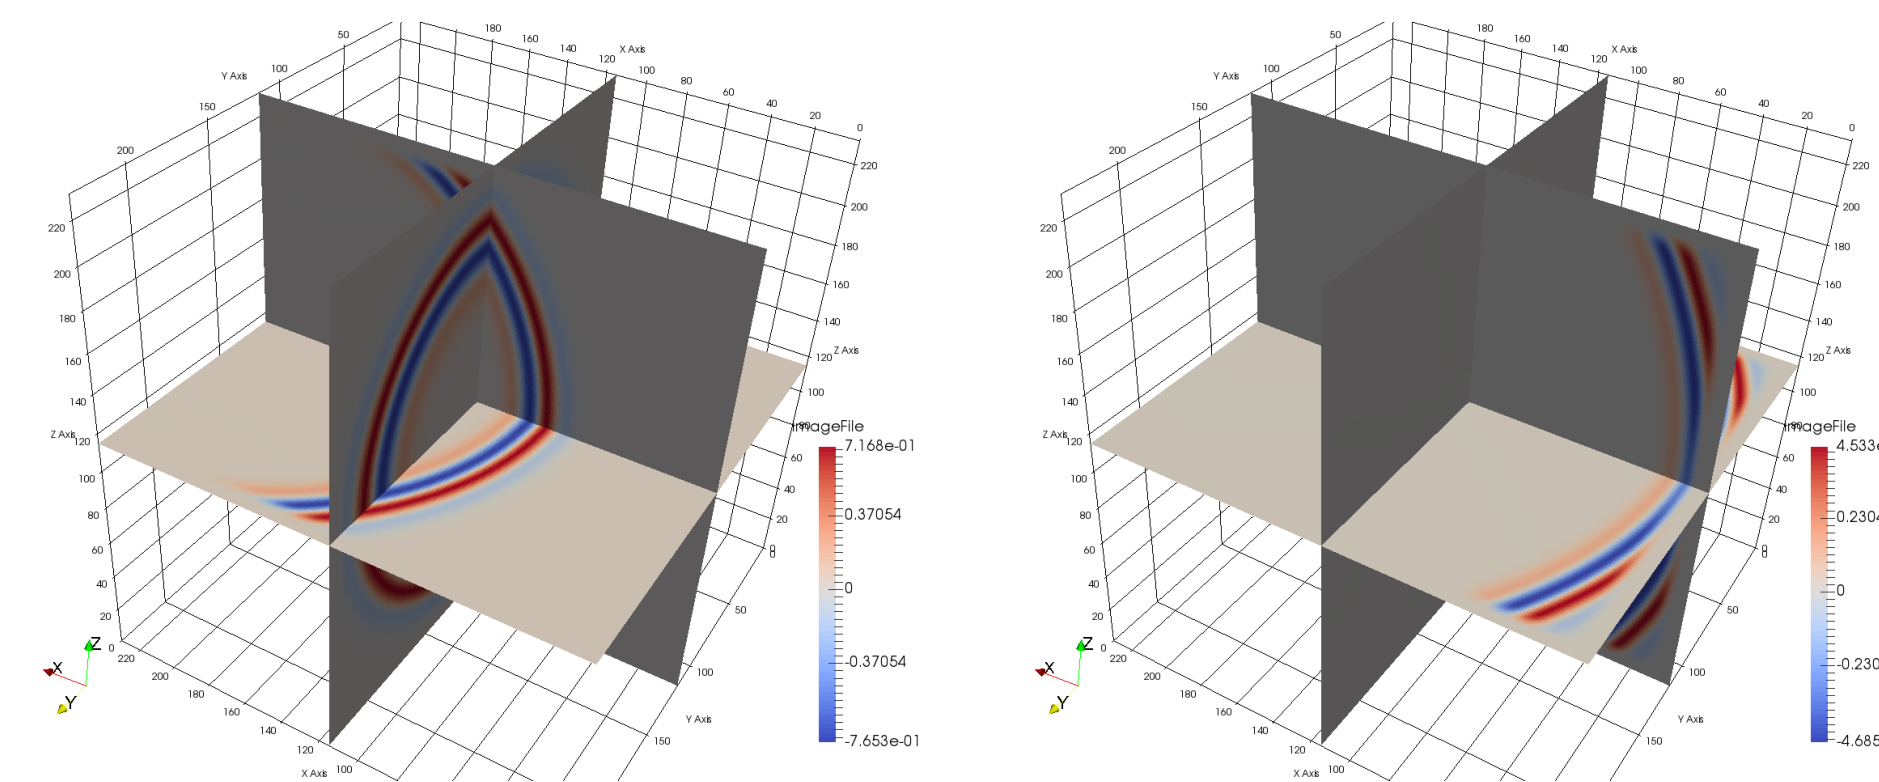


NUMERICAL ACCURACY



Time step 608


Time step 712



Time step 808

Time step 1208

- Maximum difference : 1.16434e-05 • Mean Absolute Error (MAE) : 1.10e-06
- Root Mean Square (RMSE) : 1.70e-05



CPML IMPLEMENTATION DETAILS


- CPML Formulation

$$\frac{\partial^2 p}{\partial t^2} = c^2 (P_x + P_y + P_z)$$

$$P_z = \frac{\partial A_z}{\partial z} + \Psi_z \quad \Psi_z^n = b_z \Psi_z^{n-1} + a_z \partial_z^{n-\frac{1}{2}} A_z$$

$$A_z = \frac{\partial p}{\partial z} + \Phi_z \quad \Phi_z^n = b_z \Phi_z^{n-1} + a_z \partial_z^{n-\frac{1}{2}} P_z$$

- CPML Redesign for the Z dimension (contiguous memory accesses)




- Extension to the X and Y dimensions


- The 1D sliding window mechanism breaks vectorization due to non-contiguous memory accesses
→ we implement a sliding surface mechanism

PERFORMANCE RESULTS #TIME STEPS=2000, CPML=10 POINTS KMP_AFFINITY=COMPACT, GRANULARITY=FINE, NUMACTL -INTERLEAVE=ALL

Intel Haswell (HSW) • E5-2698v3 @ 2.30 GHz • 40 MB L3 cache • 2 sockets x 16 cores • Sustained BW: 111 GB/s • AVX256 • 128 GB RAM • icc 18.0.1



Intel Skylake (SKL) • Platinum 8176 @ 2.10 GHz • 38.5 MB L3 cache • 2 sockets x 28 cores • Sustained BW: 182 GB/s • AVX512 • 192 GB RAM • icc 18.0.1



SUMMARY AND FUTURE WORK

- Integration of asynchronous MWD with CPML into ExaWave
- Performance improvements of MWD over SB on large grid sizes
- Application on large scale seismic surveys
- Opportunities for asynchronous I/O operations
- Support for elastic and VTI/TTI approximations
- Extension to random boundaries and various stencil orders

REFERENCES

- [1] T. Malas, G. Hager, H. Ltaief, H. Stengel, G. Wellein, and D. Keyes. Multicore Optimized Wavefront Diamond Blocking for Optimizing Stencil Updates. *SIAM Journal on Scientific Computing*, 37(4):439–464, 2015.
- [2] T. Malas, G. Hager, H. Ltaief, and D. E. Keyes. Multidimensional Intratile Parallelization for Memory-Starved Stencil Computations. *ACM Trans. Parallel Comput.*, 4(3):12:1–12:32, December 2017.
- [3] D. Komatitsch and R. Martin. An Unsplit Convolutional Perfectly Matched Layer Improved at Grazing Incidence for the Seismic Wave Equation. *GEOPHYSICS*, 72(5):SM155–SM167, 2007.
- [4] V. Etienne, T. Tonellot, P. Thierry, V. Berthoumieux, and C. Andreoli. Optimization of the Seismic Modeling with the FDTD Method. *SEG Annual Meeting, Expanded Abstracts*, pages 3536–3540, 2014.