GE Global Research

Tri-hybrid Computational Fluid Dynamics

on DoE's Cray XK7, Titan.

Aaron Vose, Brian Mitchell, and John Levesque

CUG - 5/8/2014

Hybrid CFD - Overview

- Introduction to GE's TACOMA.
- OpenMP porting effort.
  - Lessons learned.
- OpenACC porting effort.
  - Lessons learned.
- Performance results.

Expressing Underlying Algorithmic Parallelism

Turbomachinery CFD @ GE

- Essential to the design process.
- Jet engines, gas turbines, steam turbines, and more.
- 10's of cores → >10K cores.
- US National Lab supercomputers prove computational technology.

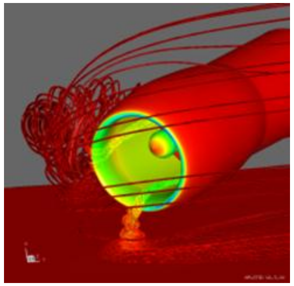

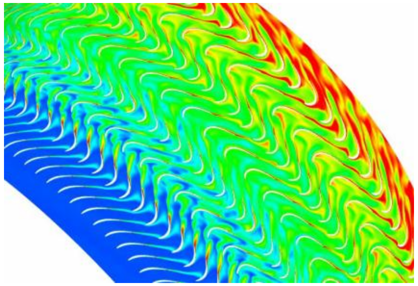

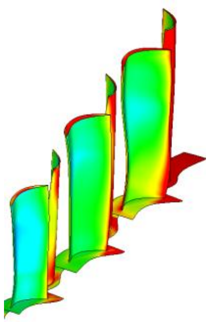

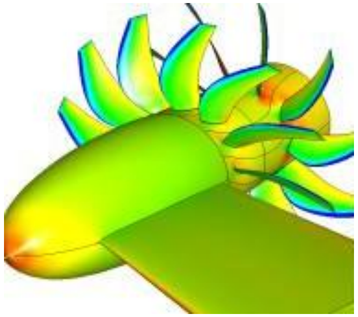

http://www.olcf.ornl.gov/wp-content/themes/olcf/titan/images/gallery/titan1.jpg

Hybrid CFD - TACOMA

- Computational Fluid Dynamics (CFD) code.
  - Used heavily for turbomachinary airfoil design
  - 2nd order accurate in time and space.
  - Finite-volume, block-structured, compressible flow solver.
  - Stability achieved via JST [1].
  - Convergence accelerated with pseudo-time marching and multi-grid.

[1] JAMES, A., SCHMIDT, W., AND TURKEL, E. "Numerical solution of the euler equations by finite volume methods using runge-kutta time-stepping schemes." AIAA (1981).

Hybrid CFD - OpenMP Porting

- Express existing loop-level parallelism with OMP directives.
- Port to OpenMP before OpenACC.
  - OpenACC can reuse most OMP scoping.
  - OMP is easier; data motion can be ignored.
- Modify loops to expose more of underlying algorithms' parallelism.

Hybrid CFD - Algo. Parallelism

```
do k=1,n3
  do j=1,n2
    do i=1,n1
      df(1:3) = dflux(i,j,k)
      R(i,j,k)    += df(1) +
                     df(2) +
                     df(3)
      R(i-1,j,k) -= df(1)
      R(i,j-1,k) -= df(2)
      R(i,j,k-1) -= df(3)
    end do
  end do
end do
```

Hybrid CFD - Algo. Parallelism

```
do k=1,n3
  do j=1,n2
    do i=1,n1
      df(1:3) = dflux(i,j,k)
      R(i,j,k)    += df(1) +
                     df(2) +
                     df(3)
      R(i-1,j,k)  -= df(1)
      R(i,j-1,k)  -= df(2)
      R(i,j,k-1)  -= df(3)
    end do
  end do
end do
```

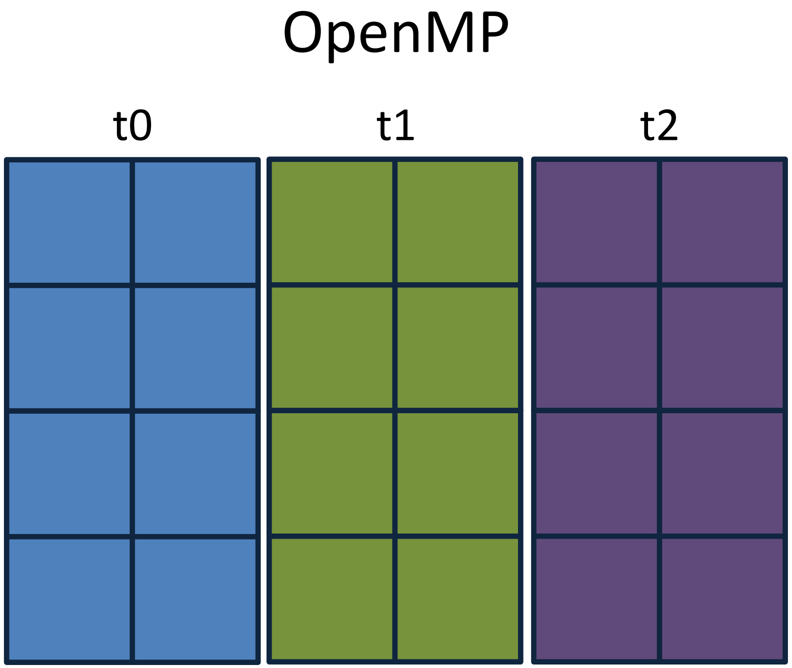

Hybrid CFD - Algo. Parallelism

```
do k=1,n3
  do j=1,n2
    do i=1,n1
      df(1:3) = dflux(i,j,k)
      R(i,j,k)   += df(1) +
                    df(2) +
                    df(3)
      R(i-1,j,k) -= df(1)
      R(i,j-1,k) -= df(2)
      R(i,j,k-1) -= df(3)
    end do
  end do
end do
```

```
do k=1,n3
  do j=1,n2
    do i=1,n1
      df(i,j,k,1:3) = dflux(i,j,k)
    end do
  end do
end do
```

OpenACC

```
do k=1,n3
  do j=1,n2
    do i=1,n1
      R(i,j,k) += df(i,j,k,1) +
                  df(i,j,k,2) +
                  df(i,j,k,3)
      R(i,j,k) -= df(i+1,j,k,1) +
                  df(i,j+1,k,2) +
                  df(i,j,k+1,3)
    end do
  end do
end do
```

Hybrid CFD - OpenACC Porting

- Identify candidate loops:
  - Check loops' trip/iteration count (CrayPAT).
- Add OpenACC directives / Optimize Kernels:
  - Check compiler listing for proper vectorization.
  - Ignore data motion (best performed once kernels are done and have known data requirements).
- Finally, optimize device <-> host data motion.

Hybrid CFD - OpenACC Data

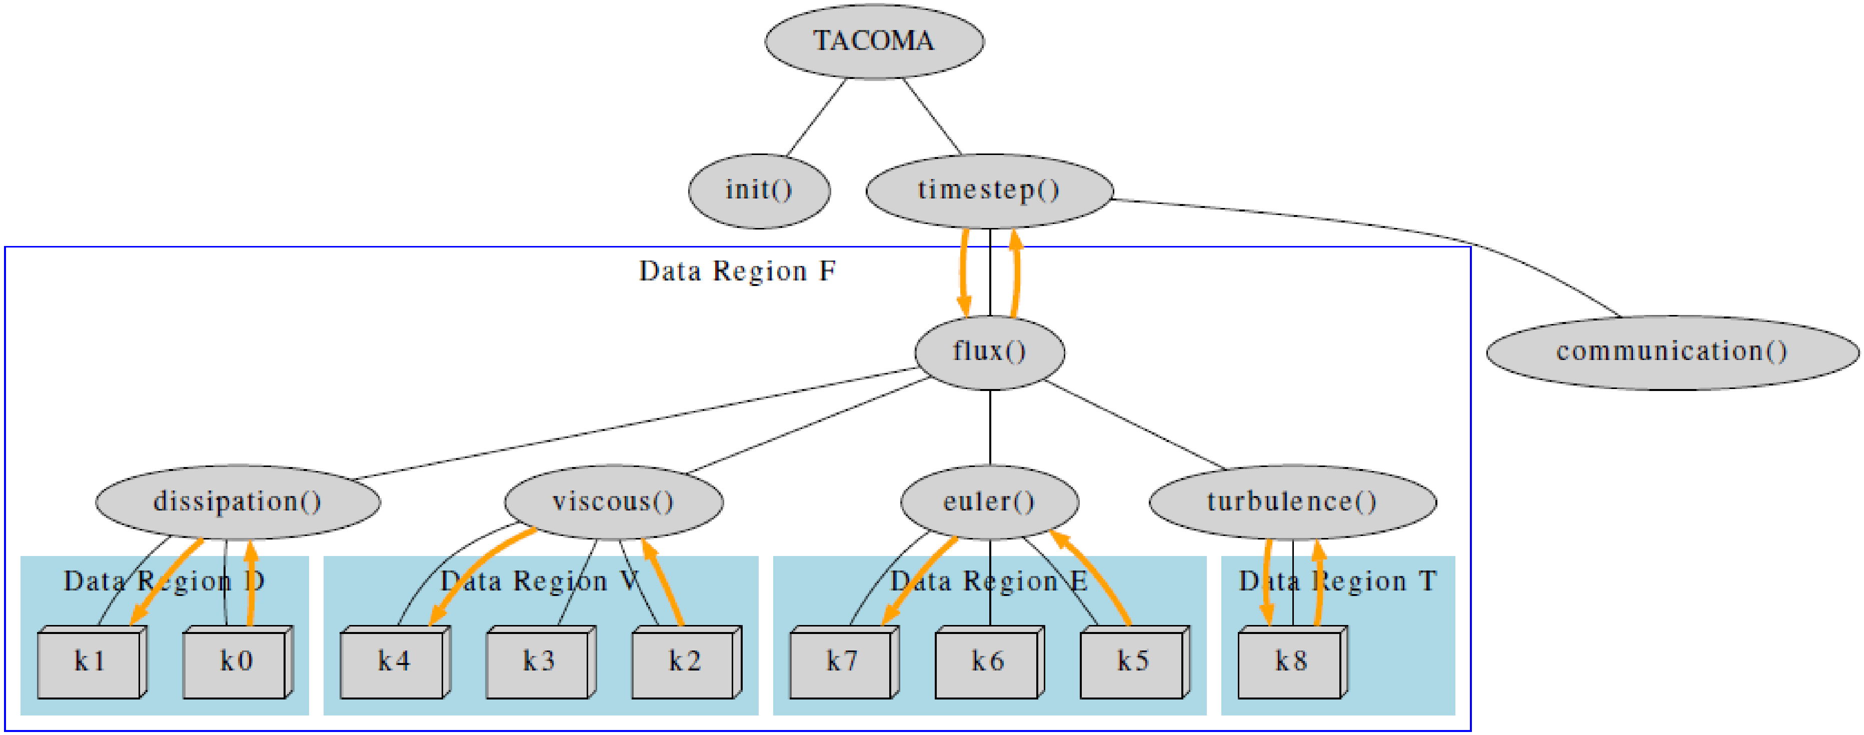

- Create OpenACC data regions:
  - Keep data on the GPU device as long as possible.

Hybrid CFD - Performance Results

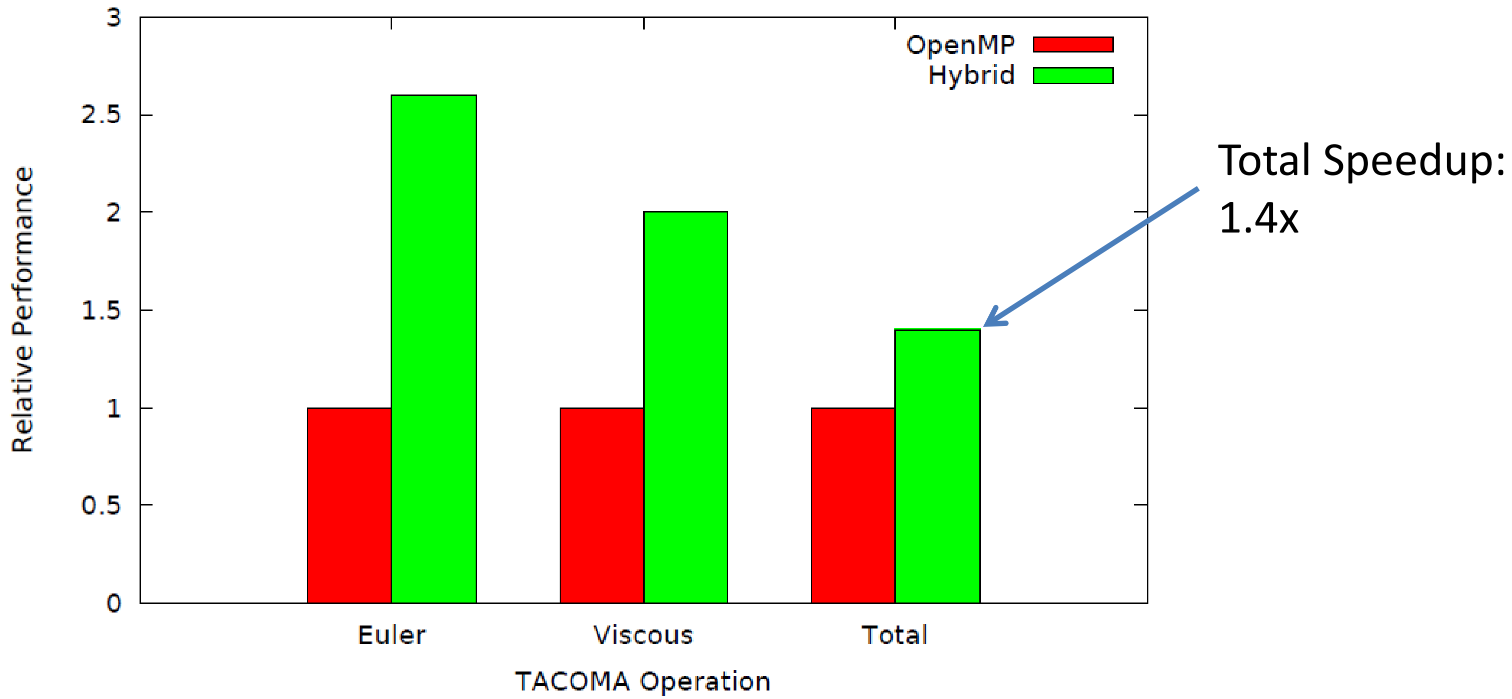

- MPI+OpenMP vs. MPI+OpenMP+OpenACC.

Hybrid CFD - Conclusion

- Lessons learned:
  - Port to OpenMP before OpenACC.
    - Reuse scoping work.
  - Optimize OpenACC data motion last.
    - Perform bottom-up, hierarchical data optimization.
  - Express underlying algorithm's parallelism.
    - Don't limit parallelism by existing implementation.

- TACOMA's OpenACC achieves 1.4x speedup.

GE Global Research

Hybrid CFD - Legal

- **Acknowledgement:** This research used resources of the OLCF at ORNL, which is supported by the Office of Science of the U.S. DoE under Contract No. DE-AC05-00OR22725.

- **Safe Harbor Statement:** This presentation may contain forward-looking statements that are based on our current expectations. Forward looking statements may include statements about our financial guidance and expected operating results, our opportunities and future potential, our product development and new product introduction plans, our ability to expand and penetrate our addressable markets and other statements that are not historical facts. These statements are only predictions and actual results may materially vary from those projected. Please refer to Cray's documents filed with the SEC from time to time concerning factors that could affect the Company and these forward-looking statements.

-

