SciDAC/LCF's/ADIOS

SciDAC 1 SciDAC 2 SciDAC 3 SciDAC 4

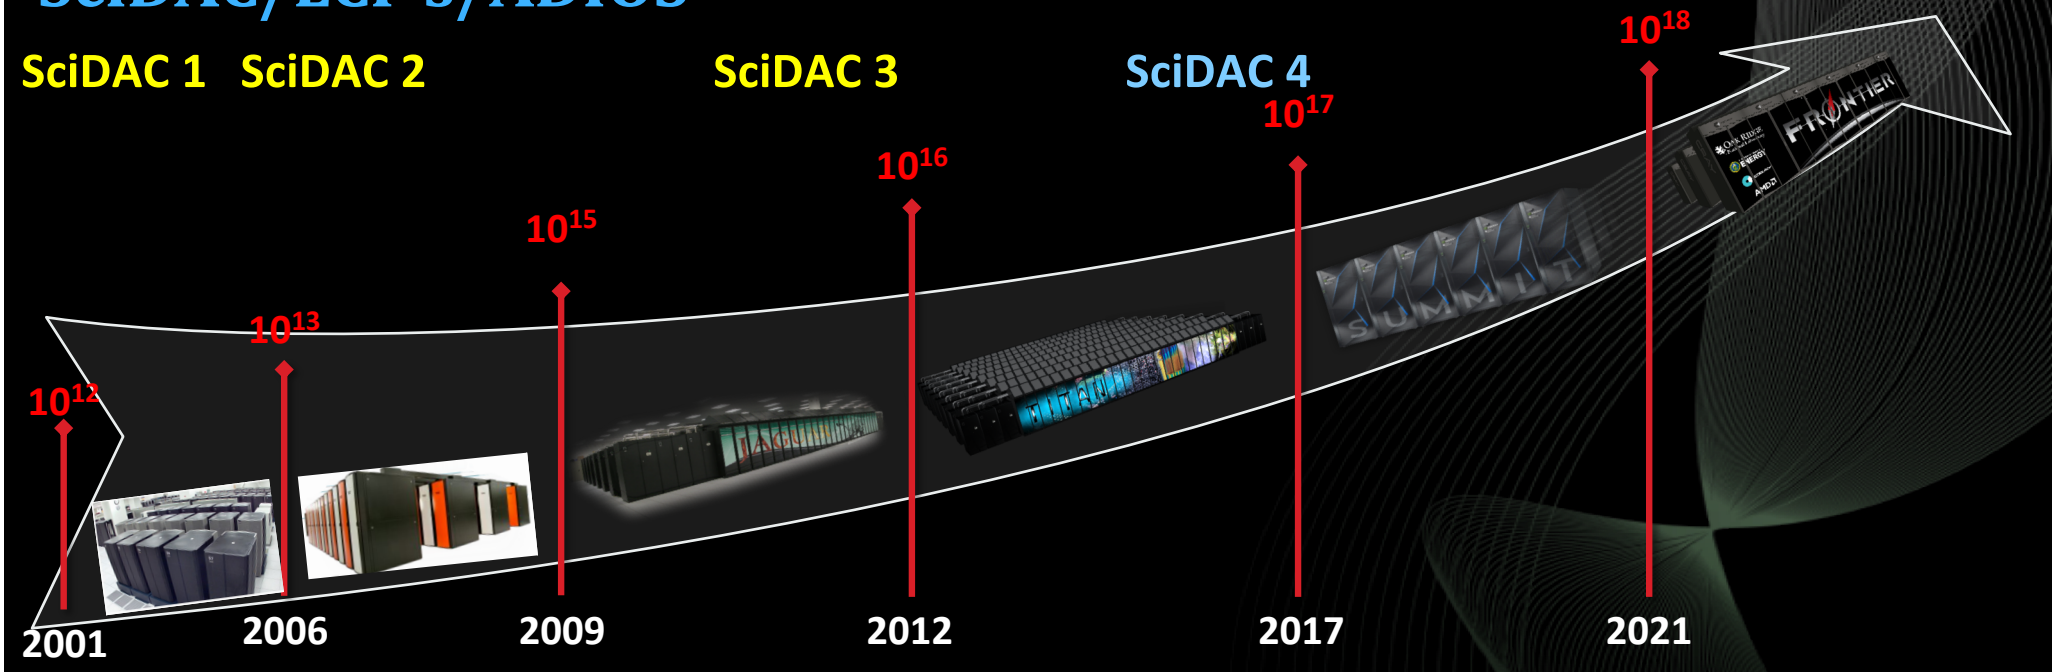

ADIOS timeline

BP self-describing file format ADIOS 0.01

Data Staging ADIOS 0.1

Research writing performance

Research reading performance

1.2 Data Staging

ADIOS 1.0 released

1.4 weak code coupling

1.6 index, compression

1.8 Query /Indexing

1.10 Recovery/Ease of use

1.11 compression

1.12 Burst buffer

2.2 New framework in C++

2.3 Summit optimizations

SciDAC funding in purple boxes

SciDAC1– Fusion microturbulence, CPES-1 W. W. Lee, W. Wang, Z. Lin, C. S. Chang, W. Tang, S. Ethier - 2 TB of data written & viz

- Science GOAL: Run ITER-size simulations to understand the behavior of Turbulence Spreading
- Computer Science GOAL: Enable the fast output of billions of particles on full-scale runs of the GTC and GTS code, to post-process over 2 TB of data from 2006 runs, and visualize the turbulence spreading
  - Initial creation of ADIOS 0.1 to enable the efficient output of self-describing data on the new Jaguar OLCF computer, on multiple full-machine runs, 100X faster than previous solutions
  - The visualization created from all of the data eventually made the FSP cover, and was shown in the Hollywood movie "Wall Street: Money Never Sleeps"

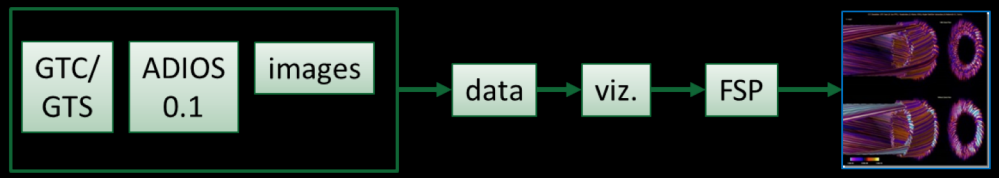

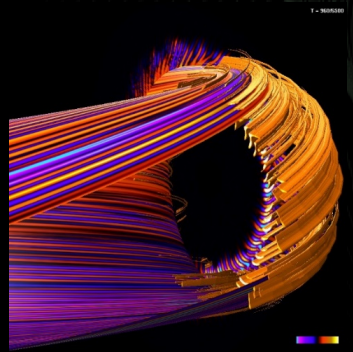

Klasky, S., Ethier, S., Lin, Z., Martins, K., McCune, D. and Samtaney, R., 2003, November. Grid-based parallel data streaming implemented for the gyrokinetic toroidal code. *Proceedings of the 2003 ACM/IEEE conference on Supercomputing* (p. 24). ACM.

GTC simulations show that turbulence spreading causes Bohm scaling

- GTC: Turbulence spreading is responsible for transport Bohm scaling in small devices
  [Lin, Hahm, PoP2004]
- ITG simulations using GTC for shaped plasmas have confirmed the nonlocal property of turbulent transport
  [Wang et al., PoP2006]
- Using an early version of ADIOS 0.1, GTC was able to write out 2 TB of data, in <10% of total runtime

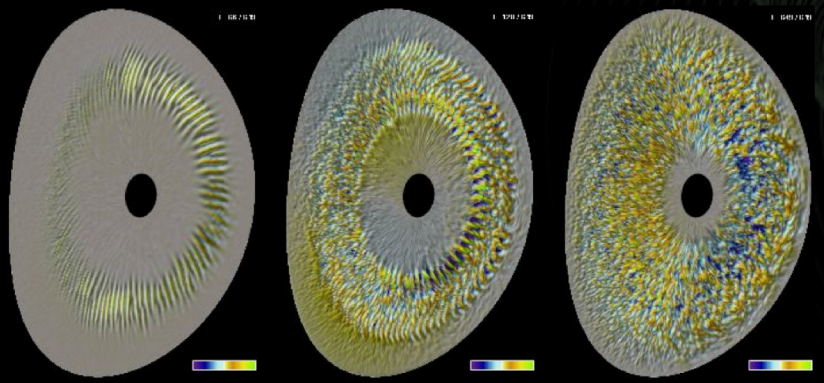

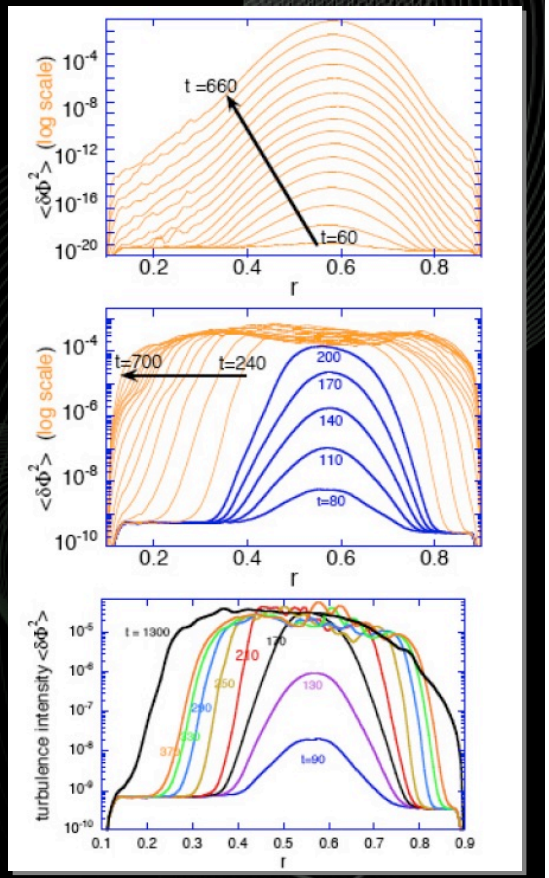

Computer Science goals – SciDAC 1

Create an I/O abstraction (ADIOS) to allow I/O to move data from a producer (application) to a consumer (file system or another application)

Create a new self-describing file format based on a log file format

Reduce the time simulations write data to the storage system (i.e. Lustre)

Results showed 10X – 100X faster than other I/O solutions

SciDAC2 - (SDM, CPES-1, CPES-2) – C. S. Chang, S. Ku, et al.

- Science GOAL: Run Whole-volume integrated gyrokinetic simulation of plasma turbulence in realistic diverted-tokamak geometry
- Computer Science GOAL: Enable efficient Checkpoint/Restart output, Analysis output, and in situ code coupling on Jaguar
  - ADIOS 1.0 fully implemented and coupled to the M3D and Elite codes using EFFIS (Kepler, ADIOS, eSiMon dashboard) to the I/O from 2 hours to 10 minutes, improving both the I/O of other solutions and enable (file-based) in situ coupling and monitoring on Jaguar
- Impact
  - Enabled in situ data processing and code coupling
  - Enabled efficient self-describing output of all the XGC1 data

Chang, C. S., et al. "Whole-volume integrated gyrokinetic simulation of plasma turbulence in realistic diverted-tokamak geometry." *Journal of Physics: Conference Series*. Vol. 180. No. 1. IOP Publishing, 2009.

Center for Plasma Edge Simulation: PI, Chang

- Whole device simulation of turbulence in XGC1 shows that core transport is connected to edge pedestal property
- ITER is based on experimental observation that a good edge pedestal yields a good core plasma
- We succeeded, for the first time, in gyrokinetic simulation from wall to magnetic axis around the whole torus in DIII-D diverted geometry
- Multiscale background and turbulence dynamics is simulated together in full-f 5D gyrokinetic code
- The edge generated turbulence (top figure) propagates inward (bottom figure), connecting nonlocally the edge plasma property to the core plasma transport
- Another discovery: Central heating generates isolated turbulence in the central core. The isolation is by a self-generated strong ExB flow
- Jaguar is used at full scale

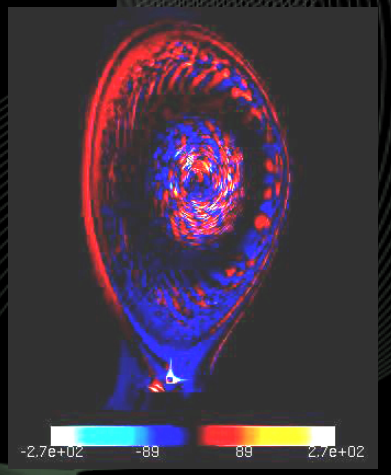

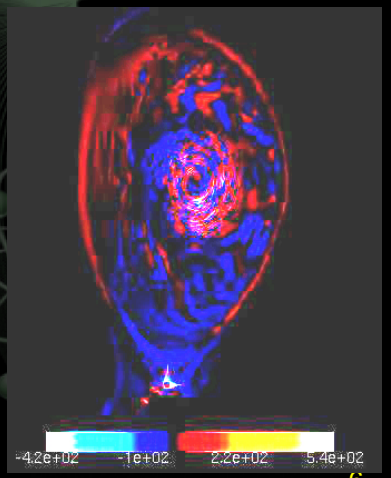

Computer Science goals – SciDAC 2

Enhance the framework for larger scales by creating new I/O algorithms for efficient I/O

New mechanisms for fast reading of the voluminous data

Allow loose/weak code coupling

SciDAC3 - (SDAV, EPSI) – C. S. Chang, S. Ku, et al.

- Science GOAL: Gyrokinetic-neoclassical study of the bootstrap current in H-mode pedestals
- Computer Science GOAL: Enable efficient Checkpoint/Restart output, Analysis output, and in situ code coupling on Titan
  - ADIOS enabled efficient output and code coupling on Titan, enabling over 150 GB/s of over a PB of output
- Impact
  - Enabled memory-based in situ code, coupling, and efficient output
  - Enabled efficient self-describing output of all the XGC1 data

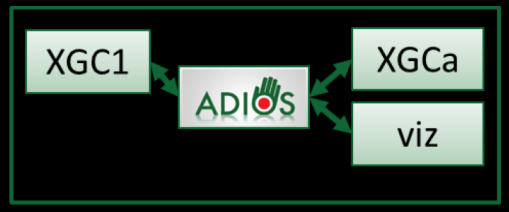

Klasky, S., Abbasi, H., Logan, J., Parashar, M., Schwan, K., Shoshani, A., Wolf, M., Ahern, S., Altintas, I., Bethel, W. and Chacon, L., In situ data processing for extreme-scale computing. *Proceedings of SciDAC*. 2011.

EPSI: PI C. S. Chang

EPSi Edge Physics Simulation

- High-fidelity simulation of the spontaneous low-to-high mode turbulence bifurcation, for the first time in 35 years after the discovery of the high-mode plasma operation in tokamak plasma edge
- ADIOS enabled the large scale run by reducing the I/O time by 10X from our previous I/O solution (100X total from our starting I/O in CPES)
- Frequent machine problems made resiliency a top concern for our project
- https://www.energy.gov/science/fes/articles/extreme-scale-code-models-extremely-hot-plasma-explain-spontaneous-transition

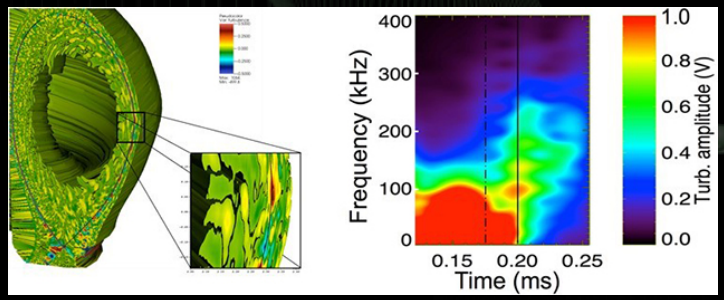

Computer Science Goals – SciDAC 3

Efficient I/O on Titan (Lustre) with over 1K OSTs

New staging mechanisms to support stronger coupling of codes

Coupling of Visualization with simulations using the ADIOS APIS

Schema incorporated into ADIOS to support complex data models

Broader impact from SciDAC

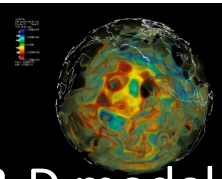

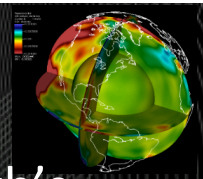

- Seismic Tomography Workflow - Enabled the most detailed 3-D model of Earth's interior showing the entire globe from the surface to the core–mantle boundary, a depth of 1,800 miles with (SPECFEM3D_GLOBE)
  - http://www.spacewar.com/reports/A_seismic_mapping_milestone_999.html
  - Generated over 1 PB of output on Titan in a 6 hour run, at speeds over 150 GBs
  - Created the Adaptable Seismic Data Format (ASDF) was developed to allow for recording, reproducing, and analyzing data on large-scale supercomputers
- Used in Accelerator Modeling (PIConGPU) enabling the fast output of over 3 PB of data generated from their Titan simulations for the understanding of miniature (table-top) accelerators to target cancer cells
  - 2013 Gordon Bell finalist
  - https://www.ornl.gov/news/titan-targets-tumors

Broader impact from SciDAC

- Used in the SKA project to enable efficient storage of their streaming data (over 1 PB/s, 24x7)
  - Currently being incorporated into the CASACORE framework for Radio Astronomy data
- Used by FM Global Inc (OpenFOAM) to enable the understanding of the spreading of fires in warehouses, which made it into the 2018 DOE budget request.
  - https://www.olcf.ornl.gov/2016/01/05/fighting-fire-with-firefoam/
- S3D - perform the first direct numerical simulations of reacting jets in cross flow. These transverse jets are a class of flows used in practical applications in which high mixing rates are desirable — for example, in fuel injection nozzles in stationary gas turbines for power generation or in aero-gas turbines.
  - https://www.hpcwire.com/2009/10/29/adios_ignites_combustion_simulations/