

A Maya Use Case: Adaptable Scientific Workflows with ADIOS for General Relativistic Astrophysics

Magdalena Slawinska¹ Tanja Bode² Jeremy Logan³
Michael Clark² Hongbo Zou¹ Matthew Kinsey²
Matthew Wolf¹ Pablo Laguna² Scott Klasky⁴

¹College of Computing
Georgia Institute of Technology, Atlanta, GA

²School of Physics
Georgia Institute of Technology, Atlanta, GA

³Remote Data Analysis and Visualization Center
University of Tennessee, Oak Ridge, TN

⁴Scientific Data Group
Oak Ridge National Laboratory, Oak Ridge, TN

Outline

Maya-RDAV Introduction

Numerical Relativistic Simulations with Maya

Maya Code

Data Processing and Analysis

Our Approach to Maya Workflow Improvement

Incorporating ADIOS into Maya as Fast I/O System

Incorporating ADIOS into Maya


Magneto-hydrodynamics

Particles

How to Make Maya Simulations Interactive and Steerable at Runtime




Exploiting Flexpath/ADIOS in Maya

Summary

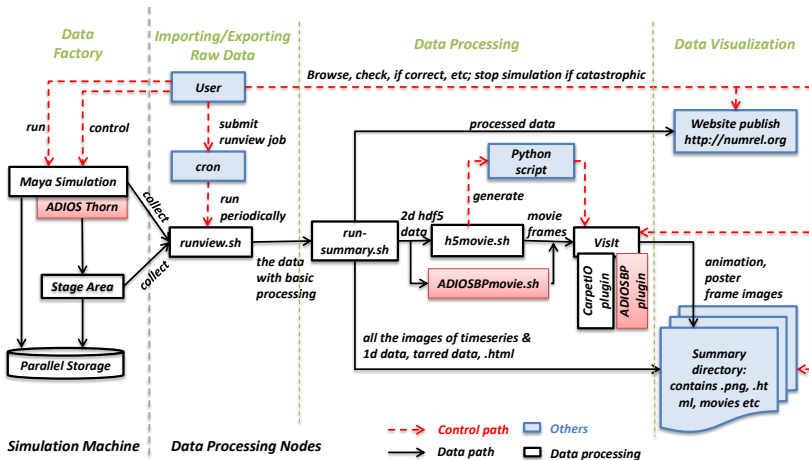


Maya Code

- ▶ 3D numerical relativistic calculations simulating creation and collision of black holes
- ▶ Solves Einstein's equations for black holes, stars, etc
- ▶ Based on Einstein Toolkit, Cactus, and Kranc; uses Carpet AMR module
- ▶ Uses Georgia Tech's custom spacetime and hydrodynamics evolution code



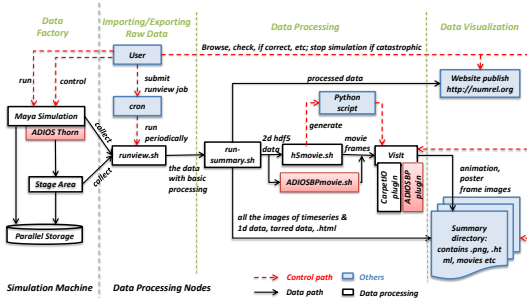
Maya Workflow



Data Processing and Analysis in Maya


Data

- Physical quantities of interest, e.g., matter density, magnetic field components, gravitational radiation intensity
- Periodic checkpoints



Processing and analysis

- Some physical quantities scrutinized during in-simulation analysis (analysis thorns)
- Post-processing including interactive analysis, scripted plots, and movies



Maya I/O Characteristics and Challenges

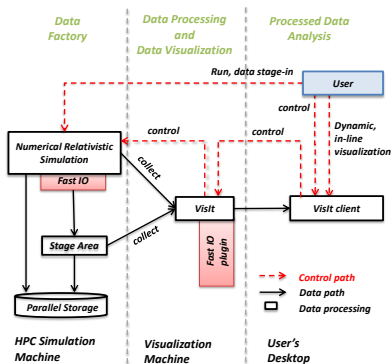
- ▶ Expensive I/O: assumption of keeping I/O overhead at 40% of walltime → write out a small number of variables with a lower dimensionality
- ▶ Mostly static, shell-based, post-processing-based data analysis and visualization
- ▶ A typical simulation:
 - ▶ A couple of weeks on a local cluster
 - ▶ Recording 15 2D physical variables of interest
 - ▶ 21GB output data (15 3D variables would generate 4.5TB)
 - ▶ Utilize a combination of HDF5 and ASCII for writes

As a consequence

- ▶ Not all variables get written out. Experimenters pre-select a small subset for each run
- ▶ Requires reruns of the simulations if correct variables were not output
- ▶ Simulation data sets prone to corruptions with large files

Our Approach to Maya Workflow Improvement

- ▶ Incorporating new staging I/O technologies to reduce I/O overhead
 - ▶ Enable outputting more physical variables of interest
 - ▶ Enable outputting higher dimensionality of variables
- ▶ Expanding the existing Maya workflow to address inline visualization and steering



ADIOS – Adaptable I/O System as Our Fast I/O

- ▶ A componentization of the I/O layer with a selection of various I/O methods
- ▶ Designed to provide a scalable, portable, and efficient I/O solution for various computer infrastructures
- ▶ In certain cases demonstrated I/O performance improvements of factor 1000 over established parallel file formats
- ▶ Successfully incorporated into several scientific codes including GTC fusion code, GTS, Chimera, XGC0, XGC1, Flash, S3D

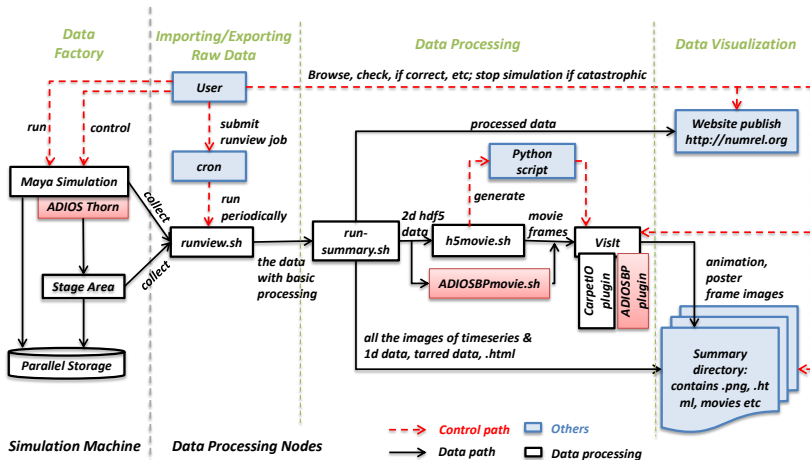
Advantages of Using ADIOS

- ▶ Ease switching to other transport method (MPI, MPI_AMR, Flexpath, etc)
- ▶ Offers memory-to-memory coupling solutions (Flexpath)
- ▶ Self-describing BP file format
 - ▶ Ability to avoid runtime costs related to file consistency validation
 - ▶ Footer index—does not necessitate moving it whenever data length change occurs
 - ▶ Resilience support

- ▶ More info

www.olcf.ornl.gov/center-projects/adios/




How to Improve Maya Workflow with ADIOS



Incorporating ADIOS into Maya: CarpetLOADIOS Thorn

Magneto-hydrodynamics simulations



- ▶ Early prototype demonstrated an improvement in output performance from 3-5 minutes for a full I/O checkpoint using HDF5 to 6 seconds using ADIOS on Kraken
- ▶ After a number of tuning operations on NICS machines, the equivalent HDF5 runs take 15 secs on Kraken
- ▶ This is still a 60% improvement in output performance
- ▶ Hopefully, the CarpetLOADIOS thorn will be released in a couple of months



Incorporating ADIOS into Maya: Particles


Particle-based simulations

- ▶ Early prototype shows 35% improvement over standard I/O methods in the simulation of 1 million particles running on 192 threads with 8 OpenMP threads per MPI process
- ▶ Ease of adding ADIOS support to particle based simulations (one afternoon + accelerators ;))



Incorporating ADIOS into Maya: Current Progress

- ▶ Due to recent change in ADIOS API, we had to modify our ADIOS-based checkpoint file format
- ▶ Validating the new file format
 - ▶ Debugging and performance evaluation of the newest CarpetLOADIOS thorn on Kraken and Nautilus



Exploiting Flexpath – ADIOS Transport Method – in Maya

- ▶ Included in ADIOS spec and distribution
- ▶ One of standard ADIOS transport methods – switching to Flexpath is easy:

```
// config.xml
<method group="temperature" method="FLEXPATH">QUEUE_SIZE=4;</method>

// writer
adios_init ("config.xml", comm);

// reader
adios_read_init_method(ADIOS_READ_METHOD_FLEXPATH, comm, "");
ADIOS_FILE* afile = adios_read_open("config.xml",
    ADIOS_READ_METHOD_FLEXPATH, comm, ADIOS_LOCKMODE_NONE, 0.0);
....
adios_read_finalize_method(ADIOS_READ_METHOD_FLEXPATH);
```

Advantages of using Flexpath

- ▶ Will allow to gain on-demand insights into simulations at runtime
- ▶ Will offer dynamic exploration of data
- ▶ Fast I/O and binary filtering will enable the scientist to formulate very specific queries related to NR systems being evolved

Flexpath Technical Details

- ▶ Provides memory-to-memory coupling
- ▶ Allows for memory-to-disk and disk-to-memory operations using efficient buffered I/O in the staging area (as part of ADIOS)
- ▶ An event-based method (EVPath) with inline transformation support
- ▶ Uses in-system JIT compiler to deploy binary filters at runtime
- ▶ Offers multi-protocol overlay networking, including multi-machine, advanced shared memory support, networking over RDMA (Remote Direct Memory Access)

Summary


- ▶ Maya-RDAV project
 - ▶ Enrich runtime for numerical relativistic simulations and allow physicists to interactively monitor and steer computations at runtime
- ▶ Approach
 - ▶ Improve Maya workflow by incorporating ADIOS
- ▶ Implementation
 - ▶ Cataloged existing Maya workflow
 - ▶ ADIOS integration: Carpet I/O thorn + VisIt ADIOS plugins for (magneto-)hydrodynamics and particle-based NR simulations
 - ▶ Flexpath: providing memory-to-memory coupling + advanced binary filtering to enable dynamic exploration of data

Acknowledgments

- ▶ NSF grant:
ARRA-NSF-OCI-0906324

Future Work

- ▶ Performing more aggressive analytics with large datasets
- ▶ Reslicing data across all space-time
- ▶ Adding HTML5 extensions to the existing workflow



THANK YOU
? || #