

ORNL visit

Wednesday, January 25, 2017

4:54 PM

Prep: add synchronous SOS aggregation call to TAU.

Process:

- configure / build SOS
- configure / build TAU (configured with SOS)
- configure / build ADIOS (all deps)
- configure / build ADIOS TAU wrapper
- configure / build SOS examples
- configure / build skel examples

Meeting

Jeremy overview of Skel

Jeremy overview of Diesel

- Model -> DSL -> outputs
- Templates -> DSL -> outputs
- Model is defined in YAML
 - o model.dsl
 - o define all subcommands (like skel)
 - source
 - makefile
 - xml
 - etc.
- Templates written with variable substitution

Post-lunch

Paper plan

3 past paper examples for the need to avoid introducing network jitter.

- FlexIO - Fang Zheng
- DataTap - Hasan Abbasi
- Manish Parashar (something related to jitter)

Target:

- SC? - not an I/O track. March 20
 - o Clouds and Distributed computing
 - o Data analytics, vis and storage
- Cluster? - April 24 / May 1
- SC Workshops?
 - o ISAV
- Middleware - could be a hard sell.
- EScience - 19 May

Story:

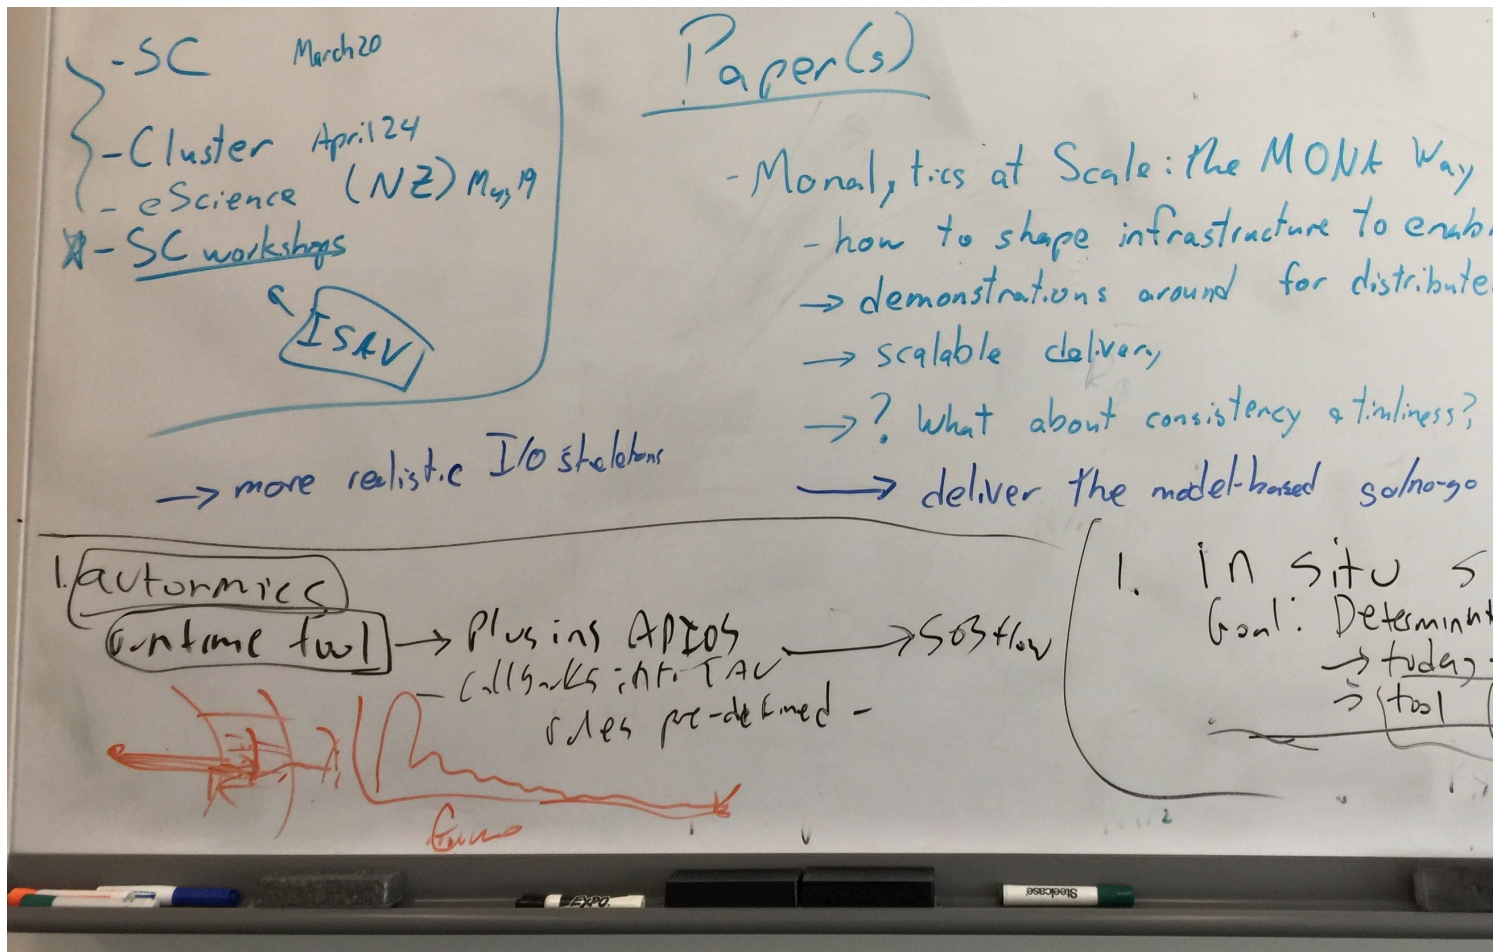
- Monalytics at scale: The MONA way
 - o Infrastructure focus
 - o how to shape infrastructure to enable requirements
 - o NOT a framework paper
 - o demonstrations around the different resource choices for distributed store of performance data
 - o scalable delivery
 - What about consistency?
 - Timeliness
 - o Actually delivering the model-based go / no-go analysis
 - preventing jitter
 - clock synchronization in the database
 - need to implement the scheduling pull
 - o *More realistic I/O skeletons*

Discussion with Scott Klasky:

- What are the metrics of success for the MONA project
 - o Demonstrable effect on application performance








Generation of quadchart

DAY 2

Discussion of
synthetic workflow integration (see photo)



to requirements
d store of perf. data

analysis WDM

service appra.
on of performance
- currently difficult

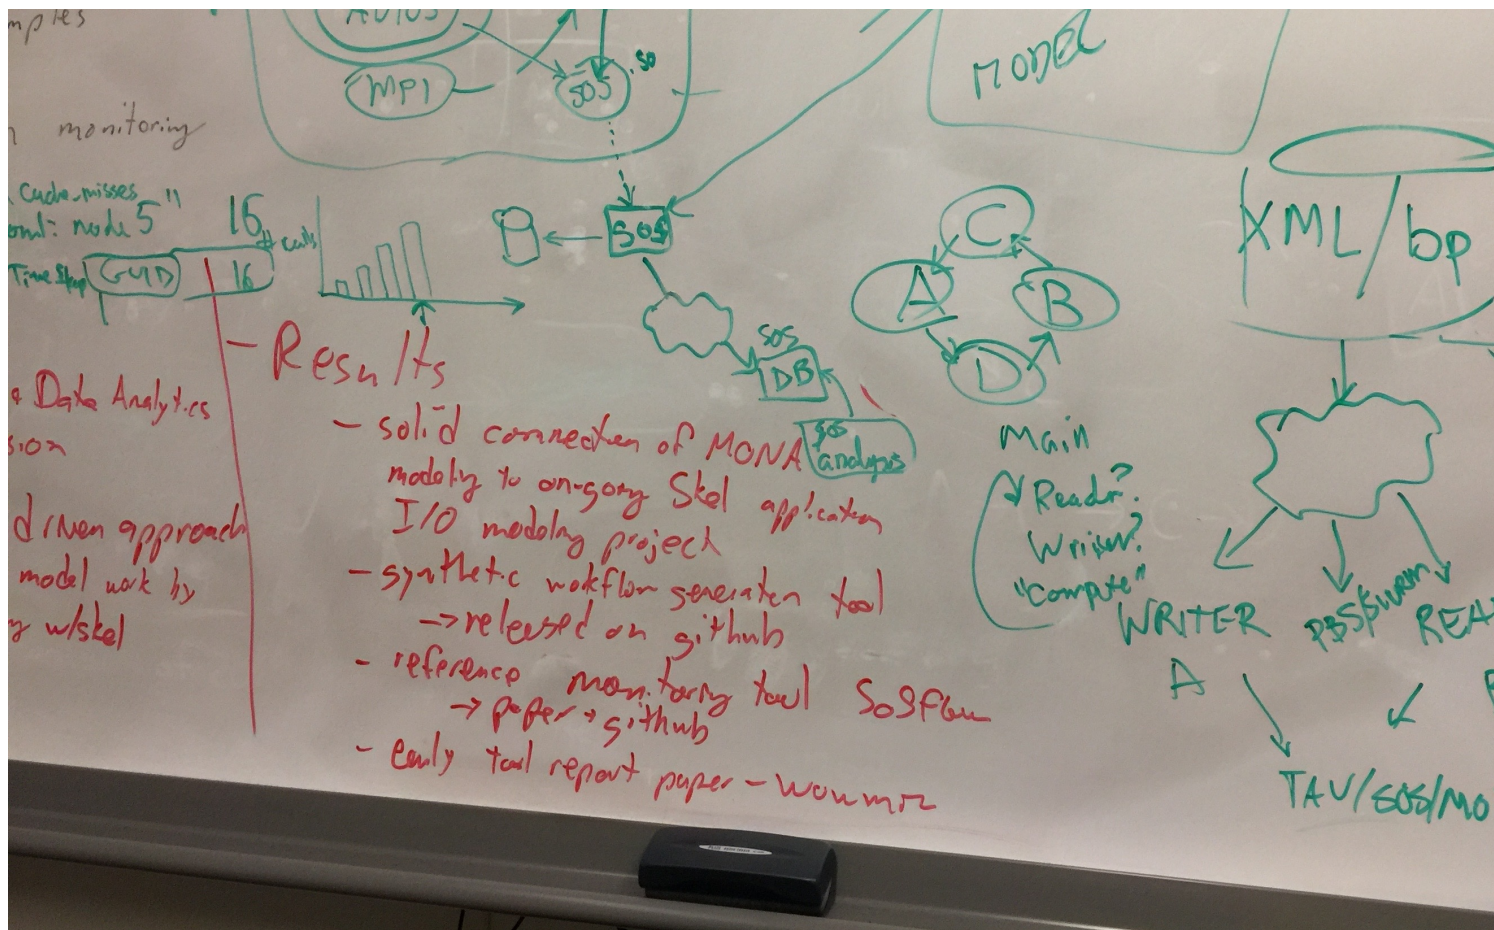
slow → web page
view

ureboard

ms


rr



B




Integration plan

- reduced jitter / timed pulls / a la datatap
- "generic" test harness for doing analysis development and testing
- datatap code is in ADIOS
- AMR, load balancing codes are tricky, because of irregular behavior
- Gyrokinetic codes are tricky because of multiple, interacting periodicity.
- I/O time is not taken into account - need to consider multiple periodicities (writing vectors, writing particles)
- Parameters: loop length, offset within the loop, length of good window, realtime offset, I/O frequency
- GTS, LAMMPS examples







Captureboard

Network Contention avoidance
Memory Pressure

Falcon

Falcon



tan
nland/Kids 80
DDR IB +40
mmutation
tr61
n

enland
path for scheduling

Integration
apture of MPI
computation Collectives
computation

SS
P WORKFLOW

Titan
Keenland/Kids 80
+40
QDR IB

VO/GT Flexpath instrumentation

VO/GT Feedback & control

VO SOS → E Vpath

VO/ORNL MPMD on Titan

GT/VO Stand up on Keenland

GT Modify IB in Flexpath for scheduling

ORNL Analytics

PPPL Use case data

VO/UTK Skel/Synthetic integration

VO SOS/TAU event capture of MPI
Collectives
adios - begin - computation
End - computation

All Test experiments
Run with Real APP WORKFLOW

Paper plan

- planning postponed until phone calls

Call with Stephane

- Current holdup - waiting on jobs on Edison to validate ADIOS integration in GTS

current holdup - waiting on jobs on Edison to validate ADIOS integration in GTS update.


- Matt suggesting Kids/Keenland/Falcon as backup plan?
- GTS is running, needs Petsc
- After GTS is running, need analysis component
- Then...could add Visit as a third component

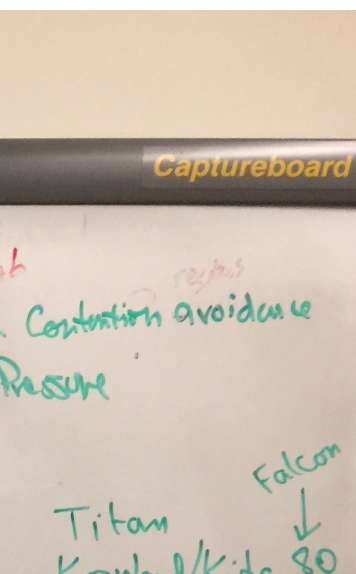
Lunch

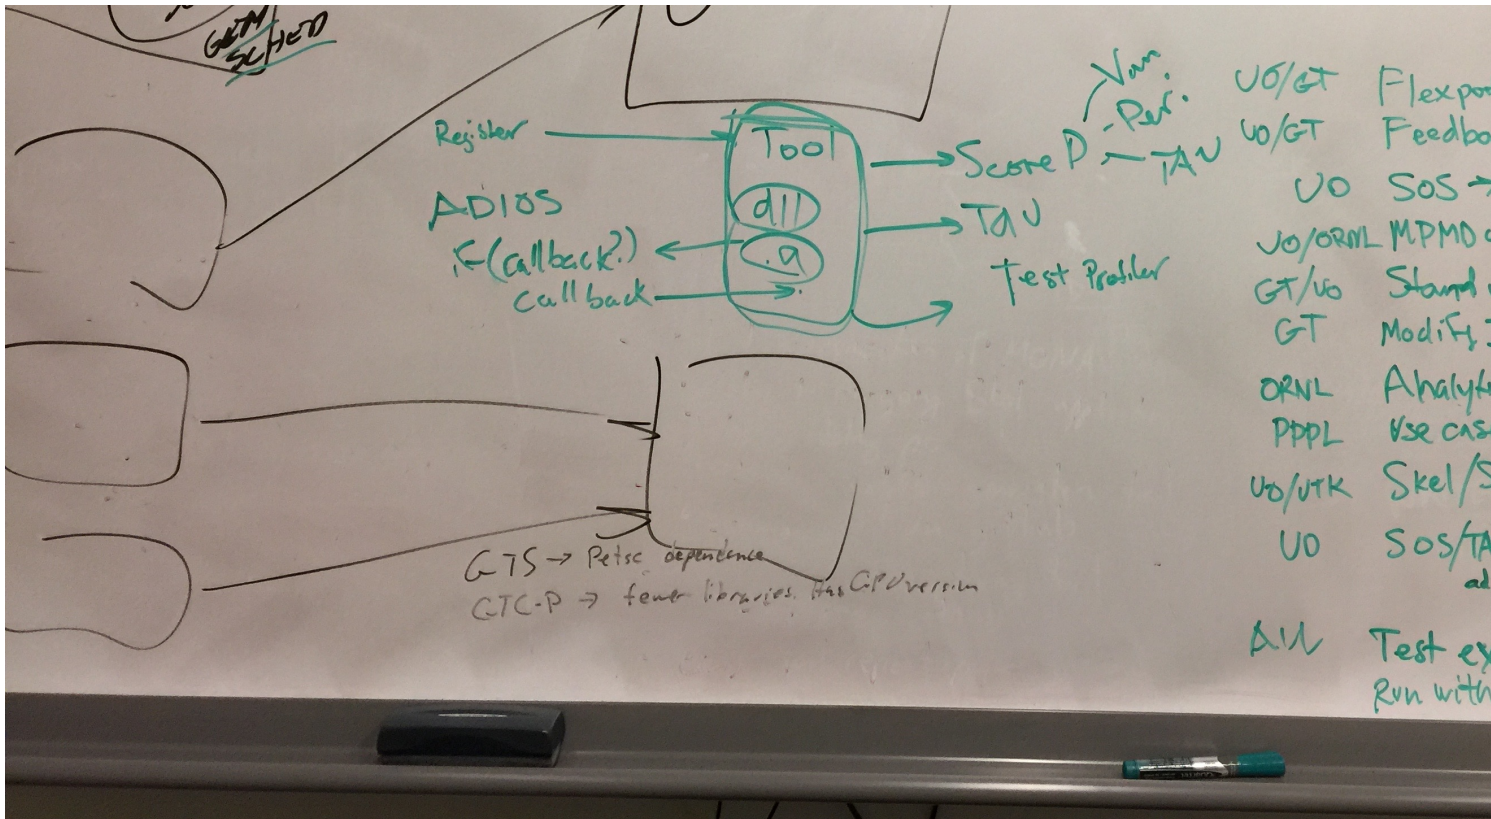
Data science group meeting

Meeting with Norbert

- What in ADIOS should be measured by TAU - if we were going to have a more "permanent" instrumentation
- Persistent capability to observe ADIOS behavior (like Vampir tracing)
- is External API sufficient?
- What about internal transport information?
- Bytes written, metadata size, bytes in memory.
- callback interface - register a tool, the tool does the timing/measurement
- We care about:
 - o interval timers
 - o memory footprint (especially max size)
 - o bytes transferred
 - o synchronization delays
- Work with Norbert on creating callback API in ADIOS
- Would be useful for CODAR "always on" data collection, Interaction with Score-P / Vampir







CODAR overview

- Center for Online Data Analytics and Reduction
- one of 4 codesign centers for ECP - all others are math/methods
- EVERYONE needs to deal with the I/O bottleneck on exascale machines
- Divided into 3 teams
 - Analysis
 - Broader scope than reduction
 - Not graph/topological analytics
 - Focused on another type of reduction (ML, statistical inference, summarization)
 - Reduction (compression, some lossless, mostly lossy)
 - Guidance on which compression methods to use for which kinds of data
 - Infrastructure
 - How do we deliver this?
 - If things are good, what is the platform for delivery?
 - How do users evaluate the solution for themselves?
 - Testing harness that drives initial tables and evaluation processes of the raw data and compression methods and algorithms.

Keenland/K145 +40
DDR IB
th instrumentation
ck & control
E Upath
on Titan
up on Keenland
IB in Flashpath for scheduling
is
e data
Synthetic integration
W event capture of MPI
Collectives
ios - begin - computation
End - computation
periments
Real App workflow

- Something that looks a lot like Skel, ADIOS running on SwiftT
 - How to extract performance, throughput, memory footprint, etc
- Complications - temporal compression, buffering techniques
- Not just single timestep, single processor. The complex examples are the hard things.
- ADIOS people, Justin Wosniak (SwiftT / ANL), Kerstin's group (provenance, Brookhaven NL)
- Expected to have crossover with other application, software project groups

