Track 2 : NRT workflows

Scott Klasky, Norbert Podhorszki, Fred SUTER

REDWOOD Status upate

10/02/2024

Near Real-Time Workflow and Resilience

Why Near real time is needed

- Avoiding **disasters** (e.g. distruption in fusion)
- Saving **computation/energy** (e.g. Numerical instability)
- Saving **post processing time** for knowledge discovery
- Saving **time** between the next experiment/simulation/etc.

Need for resilience

- **Failures** or prohibitive **delays** may happen

Example of NRT **Workflow in Fusion Energy Science**

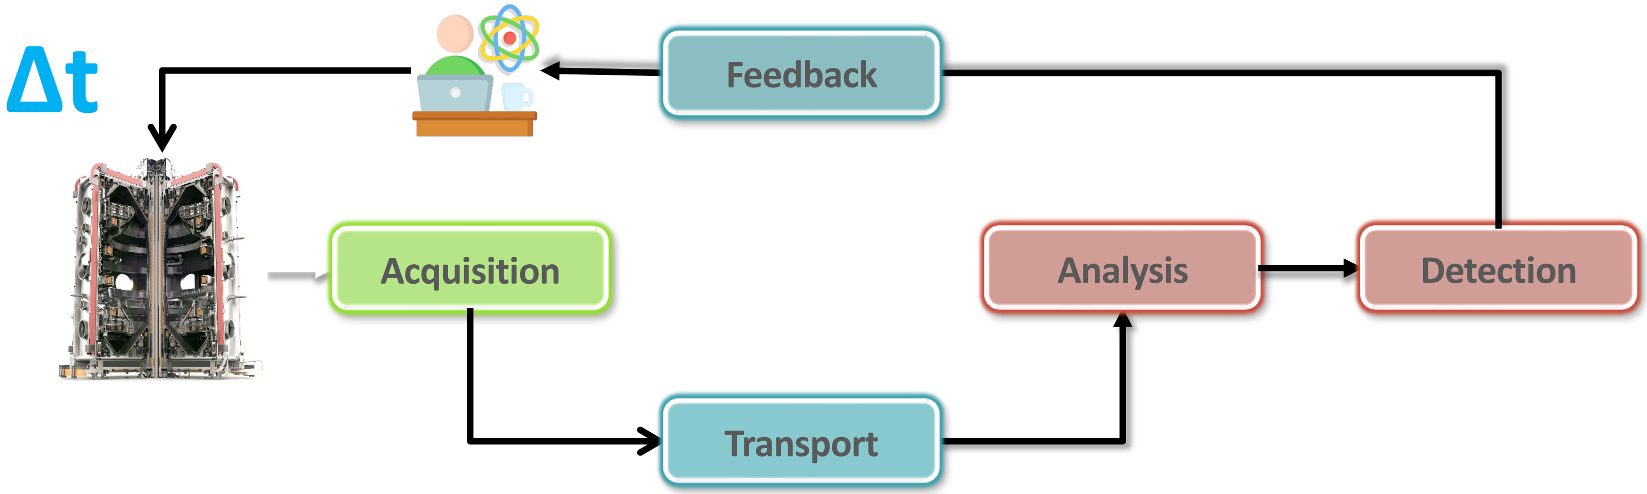

Scenario 1: **Δt** = about **1 hour**

- **Reconfiguration** between shots of the fusion reactor

Scenario 2: **Δt** = only a **few seconds**

- Early **detection of precursors** of a catastrophic event

Traditional Approaches

Near Real-Time

- **Filter** out data as soon as produced by the instrument
- Rely on **simpler** and **faster** analysis methods

→Important information may be completely lost

Resilience

- **Replicate** data
- **Replicate** computations
- **Distribute** over multiple distinct computing systems

→ This has a cost!

Roadmap Towards Resilient Near Real-Time Workflows

1. Enabling Data Streaming and Data Reduction
2. Generating Multiple Reduced Data Streams
3. Reducing Workflow Execution Time
4. Designing Multi-Objective Optimization Algorithms
5. Evaluating the Performance of Resilient Resource Allocation Strategies

1) Enabling Data Streaming and Data Reduction

Approach

- **Decouple** data acquisition, **reduction**, and **transport** from the workflow components that **produce** or **consume** data

Software

- **ADIOS:** High-performance I/O framework
- **MGARD:** MultiGrid Adaptive Reduction of Data

ADIOS

https://adios-io.org

Declarative **publish/subscribe** API

- Decoupled from the I/O strategy

Self-describing data

Multiple implementations (engines)

- **File-based** and **data streaming**

Data reduction operators

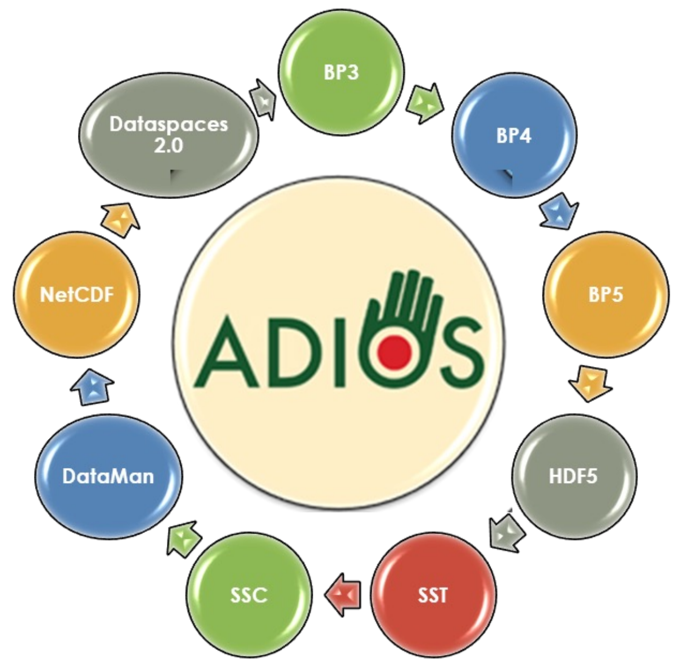

Transform-based compressor (multi-resolution, multi-precision)

- Data transformation → Quantization → Encoding
- Mathematically control errors in reconstructed data

⟶ **Large compression ratios + guaranteed errors**

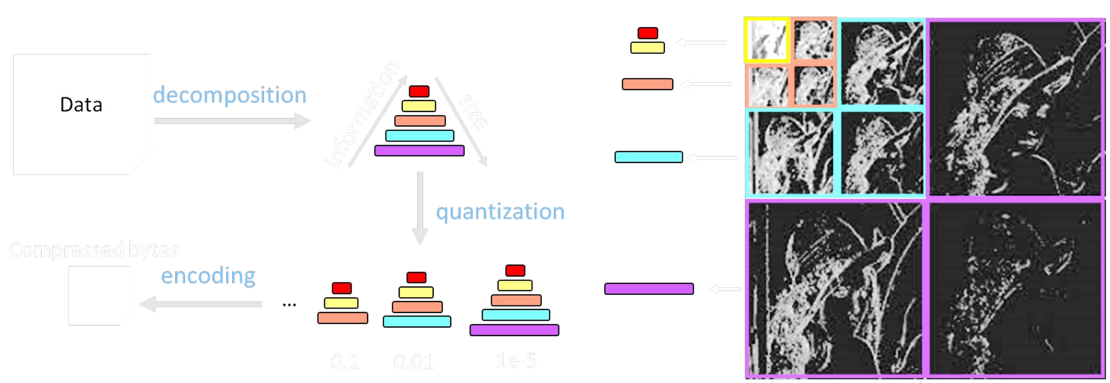

Trust

Speed

Adaptive compression

Progressive retrieval

2) Generating Multiple Reduced Data Streams

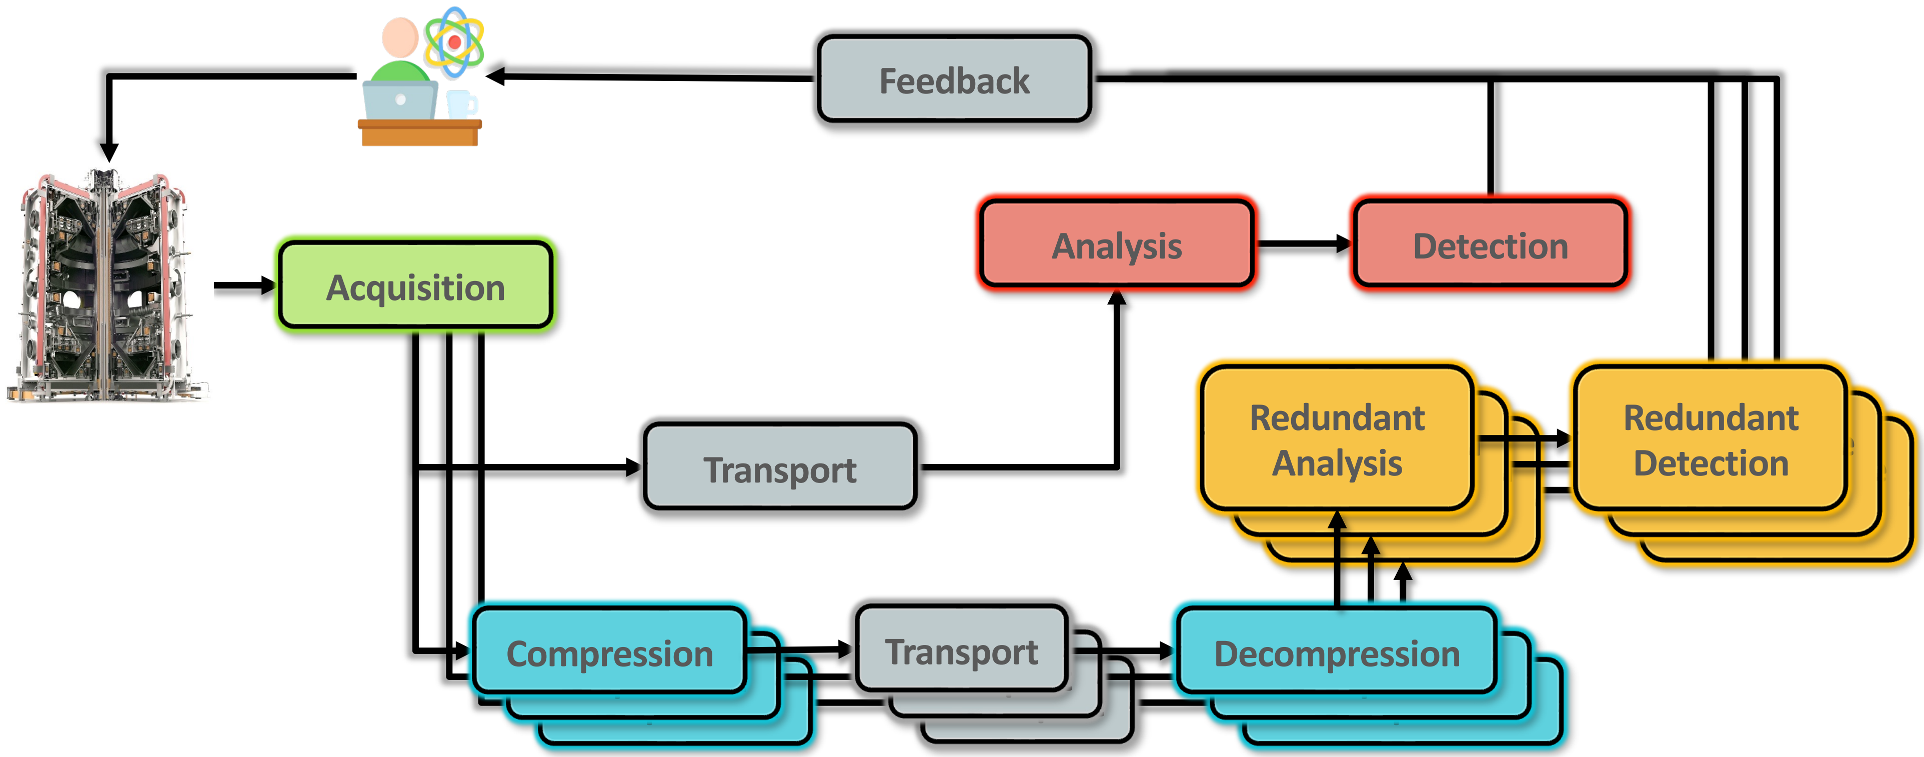

3a) Reducing Data Transfer Time

Original
67.5 MB

stride {10,10,10}
690KB

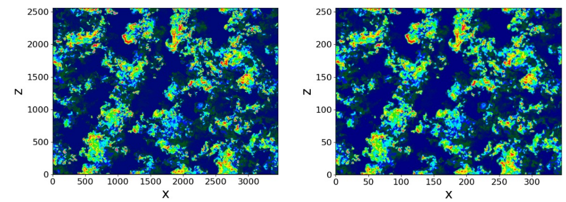

Read by **resolution** (using **striding**)

- Set a stride on N-dimensional (uniform grid) data

  ⟶ **Decreased size, but no guarantee on error , but we can determine the error and quantize to this error**

Read by **accuracy** (using **data refactoring**)

- Assume data is refactored at writing time
  - Allows for reconstruction to a user-specified error (MGARD)

    → **Same selection size, but less bits needed to reconstruct**

On-going: Read with resolution + accuracy

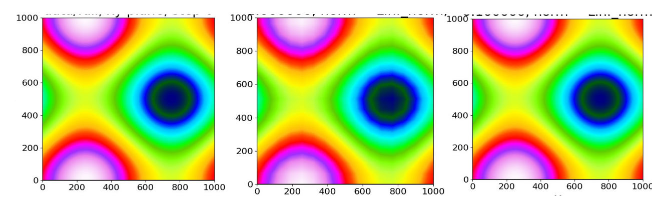

Original data

error 1.0

error 0.1

3b) Reducing Analysis Compute Time

Leverage Refactoring and Progressive Retrieval

– MGARD **multilevel decomposition** and bitplane encoding
  - Prioritize information
  - Adapt to resource constraints

– **Progressively retrieve** data **to desired accuracy**

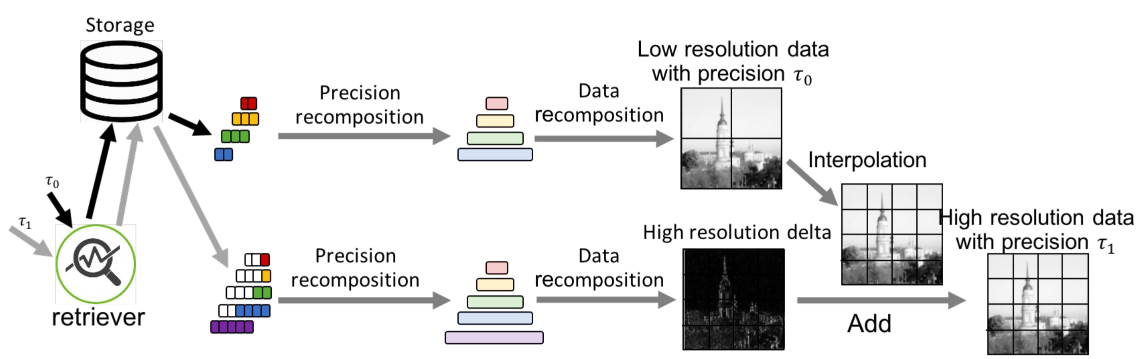

Smaller data after recomposition → shorter analysis time

4) Designing Multi-Objective Optimization Algorithms

Objective: Given a workflow **W** processing a dataset **D**

- Determine a **resilient execution scheme** that satisfies
  - **Time** constraint (**Δt**)
  - Minimum **resilience factor** (**R**)
  - Maximum **resource budget** (**B**) for resilience

Approach

- Generate **R reduced data streams** $D_{e_i}^{j}$ (e_i = error after reduction)
- Allocate **R reduced workflow replicas** to C_k resources
- Find ***i*** such that:

$$max(T(W, D_{e_i}^{j}, C_k)) \leq \Delta t \text{ and } \sum_k cost_k \leq B$$

→ $T(W, D_{e_i}^{j}, C_k)$ $cost_k$

5) Evaluating Performance

Determine **how an algorithm would perform** in various scenarios

- Optimization goals, application configuration, target infrastructure, failure injection pattern, …

Need application **execution metrics**

- Benchmarking is limited in scope or infeasible (i.e., failure injection)
- Analytical modeling becomes highly combinatorial
  - Leads to simplifying assumptions

Resort to **discrete-event simulation** instead

- Build on the **SimGrid** toolkit
  - **Open-source** project since 1998 – 2,300+ citations and 640+ usages

SIMGRID
https://simgrid.org

Versatile Simulated Data Transport Layer

Objective

- A simulated **ADIOS**
- Pub/Sub API – Self-describing data – Multiples engines
- Change the transport method, not the code

Implementation

- **SimGrid-based**
- Easy prototyping
- Accuracy and scalability
- Delay and failure injection

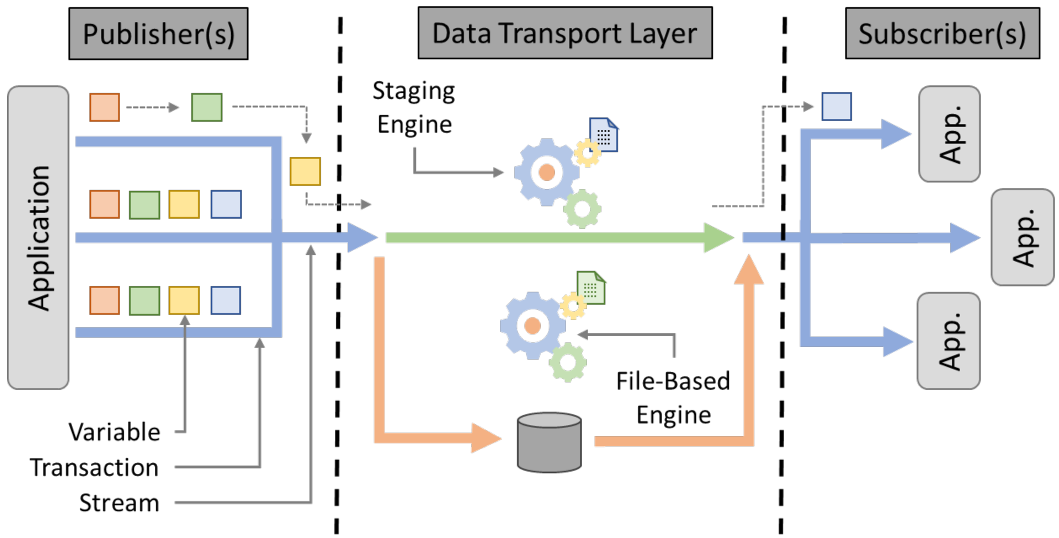

GE Aerospace - Background

Redesigned the I/O module in hpMusic/GENESIS using **ADIOS as backend**

On Summit: 1,000 node run, 1.4 TB/step

- Cut I/O cost from previous **900% to 2% runtime overhead**
- Allow enough data to be saved and for **resilient CR**
- Large datasets available on OLCF storage

Resilient NRT workflows

- Data **streamed** from OLCF to GE HQ
- Analysis workflow run during the simulation
  - Compute and visualize **derived quantities** (Total Velocity, Temperature, Speed of sound, Mach number, Total pressure, Loss)

Example: VKI case, 1000 nodes, 6000 GPUs on Summit, 5B DOF

- 1.4 TB data per output step
- Iteration speed 3 seconds wallclock time
- Desired output frequency:
  - Every 100 iterations
  - Every 300 wallclock seconds
- Runtime: 10 hours, total 168 TB data

I/O performance (write time)

- CGNS: 2700s per step
- ADIOS: 6s per step on average

Data Staging with reduced accuracy: need to work with security for a GE/ORNL DMZ

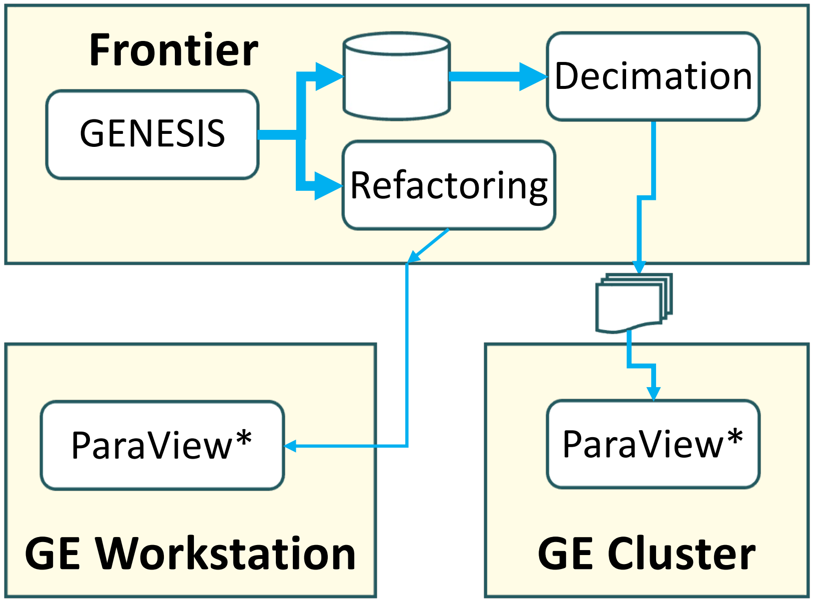

NRT decisions

Verification at full fidelity

ParaView*: Interactive visualization + analysis IDE

VKI use case

- Each step is reduced (Reduction Ratio) by 214X
- Each step contains ~10 variables, and only need to look at QoIs with ~5 variables
- Streamed to GE to meet the NRT constraint of **300 seconds/step moving 685 MB/variable**
- **We need to get GE approval to allow data to be streamed from OLCF to GE without a firewall**

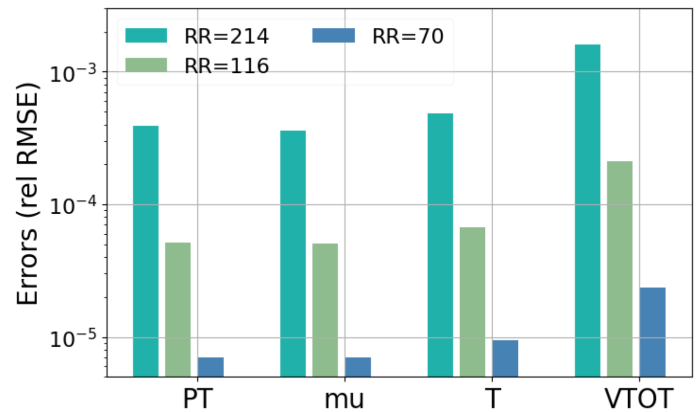

UK work

- Visiting UKAEA and Oxford Paul Watry, Shaun de Witt, Rob Akers
- Working on MAST data integration workflows

Other activities from ORNL

OAK RIDGE National Laboratory

Towards Resilient Near Real-time Workflows

Scott Klasky, Norbert Podhorszki, Fred SUTER

Questions?

U.S. DEPARTMENT OF ENERGY