NITRD Networking and Information Technology Research and Development

The government seeks individual input; attendees/participants may provide individual advice only.

Middleware and Grid Interagency Coordination (MAGIC) Meeting Minutes[1]

April 3, 2019, 12-2 pm
NCO, 490 L'Enfant Plaza, Ste. 8001
Washington, D.C. 20024

Participants (*In-Person Participants)

Ilya Baldin (RENCI)
Jessica Breet (ORNL)
Richard Carlson (DOE/SC)*
Vipin Chaudhary (NSF)
Yong Chen (TTU)
Kaushik De (UTA)
Margaret Johnson (NCSA)
Padma Krishnaswamy (FCC)
Joyce Lee (NCO)*
Deep Medhi (NSF)
Gilberto Pastorello (LBL)
Don Petravick (NCSA)
Hakizumwami Birali Runesah (UChicago)
Alan Sill (TTU)
Suhas Somnath (ORNL)
Dale Stansberry (ORNL)
Nathan Tallent (PNNL)
Wei Yang (Stanford/SLAC)

Proceedings

This meeting was chaired by Richard Carlson (DOE/SC) and Vipin Chaudhary (NSF). The March 2019 meeting minutes were approved.

I. **Speaker Series: Data Life Cycle: Margaret Johnson, Assistant Director, National Center of Supercomputing Applications (NCSA) and Don Petravick, Senior Project Manager, NCSA, *Continuous Learning About Data: Experience from the Dark Energy Survey and NCSA***
   a. Dark Energy Survey (DES) program is designed to characterize and understand the nature of dark energy by measuring cosmic expansion with high precision. It consistently observes a large fraction of the southern sky to build a statistical data set.

   b. DES: A Statistical Sky Survey (Slide 3) To build up a sky survey and support this precision cosmology, you need to build a large, uniform high quality data set, with systematics very well characterized, to analyze cosmological objects:
      i. Observing: program completed Jan 2019
      ii. Data processing, archiving: NCSA managed and served data products for downstream analysis
      iii. Data retention and management: framework supports this process

[1] Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program.

c. DESDM: Data Processing (Slide 4) Run multiple pipelines to produce data sets needed for experiment:
    i. Collect about 18k images (1 TB) /night: data processed in 2 cadences within 24 hours to enable supernova (SNe) follow-up observations and provide feedback of quality of exposures
    ii. Review data annually to reprocess uniformly with calibrations to generate deep catalogs - basis for most of scientific analysis, particularly cosmology.

d. Nightly Cadence (Slide 5) Nightly production:
    i. Run SNe Difference imaging pipeline: compare SNe exposure to templates to generate candidates for follow-up observations on different instruments
    ii. FIRSTCUT processing: perform quality assessment of all observations; evaluate individual exposures for survey quality and send quickly to observatory to observe images and ensure uniform coverage (Uniform survey)

e. Annual Cadence (Slide 6) Annual processing is more rigorous and includes:
    i. Review all well-calibrated images, producing master calibrations
    ii. Filter raw scientific images based on quality assessment of FIRSTCUT process
    iii. Criteria for inclusion: whether involved in co-addition (higher level, deeper images)
    iv. Other criteria: asymmetric, photometric calibrations of entire data set
    v. Re-perform single epoch processing: run each image through pipeline to produce clean-up images and single-epoch catalogs of objects detected in images
    vi. Global calibration for photometry
    vii. Add images to co-adds (re-catalogued; basis for cosmological analysis)
    viii. Other downstream processes fed in higher level data products
    ix. Curate and finalize data products from annual release processing campaign
    x. Summer activities: re-process SN images to drive improvements of algorithms, as needed

f. Ad-Hoc Campaigns (Slide 7) Support processes of the survey within same processing framework and with same rigor; examples:
    i. Parameter sweeps: identify configurations to be used in next annual release processing campaign
    ii. Calibrations: processed in ad hoc manner
    iii. One-off campaigns requested by collaborations

g. DES Data Management System- Provenance and Metadata Aspects (Slide 8) Processing supported by unified framework used for all campaigns:
    i. Enables the generation of diverse data products needed on cadences needed by survey
    ii. Framework built to run any scientific code from community and maintains single metadata and single provenance framework

h. How DESDM Developed (Slide 9) Originally, all files should be accompanied by rich metadata produced when data was produced, however:
    i. Representation of provenance is varied

ii. While information may have been present, it was not stored in an efficient way for usability
iii. Generated files were mutable – a process could overwrite them and eliminate the history
iv. Metadata was only collected at time of generation; no mechanization to further annotate data.

i. Revised Approach to Provenance: Adopted aspects of Open Provenance Model (OPM) (Slides 10-12) OPM implementation - changed workflow control framework:
   i. Records state of provenance alongside scientific data file as pipeline gets executed
   ii. Provenance Files returned to central archive alongside scientific data files and get ingested as provenance tuples into NCSA's operational data base; tuples then available to be queried through provenance framework
   iii. Centralized information about provenance made QA processes more efficient (example Slide 12)

j. Need for additional information about data (examples Slide 13) Important to feed information learned about data quality into further processing campaigns, creating higher level data products.

k. Tags Implemented in the operational database to capture this additional information (Slide 14):
   i. Effect on processing campaign: higher level data products can be created off of data products; generated in ad hoc campaign
   ii. NCSA approach: Metadata and provenance systems should:
      1. Present all the knowledge gained about the quality and packaging of data products to any pipeline making increasingly tailored data products, and
      2. Preserve and make accessible all relevant data **vs.** Embedding information so decisions made at run-time; use organization of the files/file system to understand provenance. Difficult to scale as higher level products created; decision making not well exposed.

l. DESDM Result (Slide 15) Implementation of DESDM provenance system.

m. NSF DIBBS: Data information building blocks (Slide 16, Tableau includes Clowder) Represents incorporation of advanced thinking; approached provenance issue more systematically.

n. Clowder (Slide 17, link): Integrated into 100 different, mostly smaller systems. Reaches into analysis tool chains where knowledge is derived vs. DESDM (only produces data products on demand).

o. Supporting Scientific Research Data (Slide 18) Supports sharing between independent researchers including curation (social and auto-curation, custom previews), publication, and touches on reproducibility, reuse (enables collaboration in building code and sharing it).

p. Clowder "Tags" (Slide 19) Applies label to data element/entity, analogous to DESDM tags, but no ontology agreement; DES has agreement regarding ontology within small collaborations.

q. Clowder supports use of ontologies (Slide 20) Clowder puts RDF-based ontology toolkit into stack, which has all necessary machinery to produce ontological type relationship; can apply to "birth" metadata and metadata added by production; instead of casual tags, has full ontological toolkit.

r. Clowder supports gradual formalization of a provenance model (Slide 21)
    i. Can build graphs based on the tools; amount of data that can be represented as tuples and related to each other in ontologically based query languages is much greater.
    ii. Supports model that DES didn't reach; can approach an ontology as scientists learn about data by using labeling; as consensus emerges, can promote it to tuple and ontological world; important as scale of application of these systems grows.

s. Incorporating basic information science as a tool (Slide 22-23) can address a large number of users.

t. Summary (Slide 24) DES shows what can be done incrementally in standards-aware silo and what can be reused from program of agency supported work.

II. **Speaker Series: Data Life Cycle: Yong Chen, Associate Professor, Computer Science, Director, Data-Intensive Scalable Computing Lab Texas Tech University, Site Director, *NSF Cloud and Autonomic Computing, Empowering Data-driven Discovery with a Lightweight Provenance Service for High Performance Computing***

a. Data Challenges (Slide1-2) Scientific discovery is highly data intensive: Data groups come from experimental data (generated from simulation models and calculations) and observational data (collected from sensors and instruments).

b. Reasons behind Data Revolution (Slide 3):
    i. More powerful computers and a proliferation of computing devices generate data faster
    ii. More data is required and produced by high-resolution, multi-model scientific discovery
    iii. Fundamental paradigm shift for data-driven science vs. model driven computational science; data-driven science does not have a model of formula or equation to describe correlation
    iv. More scientific breakthroughs are powered by advanced HPC and data understanding capabilities

c. Our Vision (Slide 4) Create lightweight and transparent provenance service for HPC.

d. What is Provenance (Slide 5 -6, diagram)

    i. Artwork - prove originality; documented history of an object

ii. Computer science - lineage of data from all system entities (include users, machines, processes, programs, threads, data files) and relationships among all these entities.

e. How to Represent Provenance: Graph-based Model for HPC Provenance Data (Slide 7-8) property graph model: vertex, edge, properties/attributes (can describe vertices in detail):
   i. Map HPC provenance onto property graph model to describe data: entities (users, files), and relationships among entities
   ii. Map entity to vertex: data objects/files, executions/simulation runs
   iii. Map relationships to edge (example Slide 8)
   iv. Attributes – properties: allows us to define and customize

f. Requirements on Managing Provenance in HPC (Slide 9) HPC system has specific requirements:
   i. Performance - users are very performance sensitive – want to manage with less than 1% slowdown; more bold goal than some existing provenance (10-20% overhead) and 1MB memory footprint per core
   ii. Coverage - provenance generated from multiple physical locations (e.g., login node, management node); want various granularities of support, depending on overhead we are monitoring
   iii. Transparency - users should not need to change/recompile their codes for provenance; should allow provenance to be collected and managed (not disabled)

g. How to Collect and Manage Provenance (Slide 10-15) LPS Overall Architecture in HPC: three major components:
   i. LPS Tracer (Slide 11): leverage kernel instrument to collect detailed runtime events to build provenance; adjust to define granularity
   ii. LSP Aggregator (Slide 12-14): monitor overhead, direct granularity change, retain critical executions and eliminate unimportant processes to increase system efficiency (Slide 14, Figure)
   iii. LPS Builder (Slide 15) View provenance globally, reference to any subsystem or provenance events, build provenance with versioning; track different versions

h. What Can We Do with Provenance (Slides 16 -17) Use cases:
   i. Evaluate new system: identify benchmarks, configurations, parameters that lead to particular scenarios (Slide 16, Diagram)
   ii. User/Project/Job level Audit: using transversal of provenance graphs across fields of operations to identify specific user/project/job at different levels.
   iii. Organization of Data Space (currently, traditional POSIX): once captured, provenance information of comprehensive relationships, the view can dynamically shift from standard fixed, hierarchical view of data sets
   iv. Data sharing, publishing
   v. Reproducibility, workflow management

i. Summary (Slide 19)
   i. Data driven discovery is new driving force for sciences (4^{th} paradigm), in addition to 3^{rd} paradigm (model-driven, computational paradigm)

ii. Holistic approach towards provenance data can help with understanding data and mining insights
iii. Working on LPS for HPC systems
iv. Call to community for more R&D efforts in the space

III. Discussion

a. Provenance – how to prevent changes being made to what has already been done? Depends on value of data or value of doing it?
   i. DES- database ingested in production system; then goes to read-only database table to users.
   ii. Secure campus: off-campus user must use VPN or secure gateway to access system. Always issue of whether users on HPC systems privacy protections vs. users have running jobs and information should be accessible. Consider privacy from job angle.

b. NCSA collect data at certain time of day or 24 hours? Is part of analysis how data is changing from 1 night to the next? Reduce amount of data stored? Also, in context of scientific analysis.
   i. Cosmological survey run by optical telescope is only run at night; further constrained by 2 factors: 1) need to observe when object is directly overhead (avoid galactic cap); 2) moon, weather
   ii. Took data in northern hemisphere fall/winter to avoid galactic cap; irregular processing/workload
   iii. Astronomy has efficient domain-specific data compression to reduce data; 1TB or less/night was compressed data
   iv. Primarily cosmology study - so distant galaxies static; supernovas change (standard candles) – special processing supports SN survey; only domain specific compression algorithm for images
   v. Get data from instruments during the day, but these are calibrations

c. External researchers can access NCSA data?
   i. Data proprietary to DES collaboration (international and domestic researchers)
   ii. DES role stops at certain level of data refinement, then goes to working groups

d. NCSA's current needs? What is driving the research agenda? Clowder and DES informed each other; DES closes down in 1.5 year

e. Next 2 things needed for your system?
   i. Clowder able to work on missing features; notes work with geospatial communities
   ii. Yong Chen:
      1. Working on the prototype; representing provenance in property graph model; assembling and processing pieces of code and get feedback from community
      2. Provenance survey – from HPC system and system administrator's perspective – how to leverage provenance collected here and how to leverage and optimize valuable resources

IV. **<u>Large Scale Networking IWG (LSN) Deliverables:</u>**

a. <u>Containerization and DevOps Reports</u>: Brief reports to be derived from last year's speaker series and will be delivered to LSN. Dhruva Chakravorty (TAM) is doing the containerization report. Contact Joyce Lee if interested in volunteering.

V. **<u>Data Life Cycle Series Planning:</u>**

a. <u>Data Triage:</u> Can be drawn out into 2 months (May and June). Will check with Fran Berman (RPI); check with Dr. Greenberg for June. How did you decide which data to discard and determine what needs to be processed quickly? Processing to be delayed. Other aspects of data life cycle. Different set of priorities based on certain factors? Looking for community insights and speaker suggestions. As science communities grow, how to store voluminous data that will be generated? (LSST, LHC)

b. <u>Confirmed speakers:</u>

 i. May

  1. Ilkay Altintas, Chief Data Science Officer, San Diego Supercomputer Center
  2. Ben Blaiszik (Computation institute; UChicago); Material Data Facility
  3. Glenn Lockwood (NERSC) – large scale data storage

 ii. June

  1. Fran Berman (Rensselaer Polytechnic Institute and co-founder of RDA (RDA – developed working groups, interested in output) (to invite)
  2. Dr. Jane Greenberg (Metadata Research Center, Drexel University) (will invite)

c. <u>Future Topic</u>:

 i. Science meets business thinking (Don Petravick)
 ii. Standards of data retention and keep data after project ends? May have insufficient amount in grant.
 iii. Relying on community to define major topics to discuss during monthly calls. Send Rich Carlson, Vipin Chaudhury and Joyce Lee.

VI. **<u>Roundtable/Events:</u>** Deep Medhi (NSF): Making progress on mid-scale infrastructure activities ($6M-20M range). Three proposals have been reviewed and invitations will go out soon.

VII. **<u>Next meeting:</u>** May 1, 2019 (12 pm ET)