

# Job Provenance — Insight into very large provenance datasets

Aleš Křenek and Zdeněk Šustr

*on behalf of CESNET JP team*

# Overview and Focus of the Demo

- Job Provenance service architecture
  - brief overview
- Demo application
  - based on Provenance Challenges workflow (Head et al.)
  - simple graphical user interface
  - quite artificial (this is not so important)
- **Quantity of processed data**
  - records on approx. 1 million of jobs
  - all of them really executed
  - JP allows repeated queries with interactive response

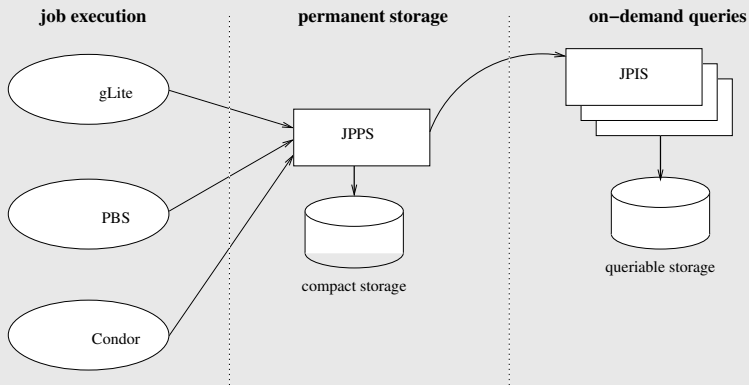
# Overview and Focus of the Demo

- Job Provenance service architecture
  - brief overview
- Demo application
  - based on Provenance Challenges workflow (Head et al.)
  - simple graphical user interface
  - quite artificial (this is not so important)
- **Quantity of processed data**
  - records on approx. 1 million of jobs
  - all of them really executed
  - JP allows repeated queries with interactive response

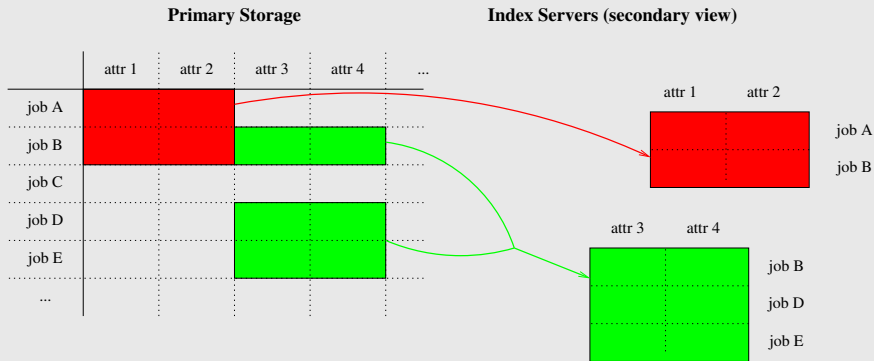
# Overview and Focus of the Demo

- Job Provenance service architecture
  - brief overview
- Demo application
  - based on Provenance Challenges workflow (Head et al.)
  - simple graphical user interface
  - quite artificial (this is not so important)
- **Quantity of processed data**
  - records on approx. 1 million of jobs
  - all of them really executed
  - JP allows repeated queries with interactive response


# Job Provenance Architecture



# Provision of Specific Views




- Workflow used in both Provenance Challenges




- Message of the motivation paper (Head et al., 2005)
  - Hippocampus volume decreases with dementia
  - Volume measurement done manually by expert operator
- Assumed development of automated measurement
  - Need of calibration wrt. clinical results
  - Look for optimal parameter settings

- Workflow used in both Provenance Challenges



- Message of the motivation paper (Head et al., 2005)
  - Hippocampus volume decreases with dementia
  - Volume measurement done manually by expert operator
- Assumed development of automated measurement
  - Need of calibration wrt. clinical results
  - Look for optimal parameter settings

- Workflow used in both Provenance Challenges



- Message of the motivation paper (Head et al., 2005)
  - Hippocampus volume decreases with dementia
  - Volume measurement done manually by expert operator
- Assumed development of automated measurement
  - Need of calibration wrt. clinical results
  - Look for optimal parameter settings

## Demo Application (2)

- One computed job
  - whole workflow, including hippocampus volume computation
  - one subject, specific setting of parameters
- Parameters to calibrate
  - **Intensity threshold**
    - ▶ principal to distinguish hippocampus
    - ▶ values 100, 200, ..., 4000
  - **Warping order**
    - ▶ impact on computational complexity
    - ▶ varying sensitivity
    - ▶ all four independently, values 3, 6, 9, 12
  - $4^4 \times 40 = 1\,024\,000$  combinations

- Job occurrence diagram
  - number of jobs – color saturation
  - subject age –  $x$  axis
  - hippocampus volume –  $y$  axis
  - clinical dementia rating – red/green
  - **good volume measurement yields vertical separation**
- Job selection controls
  - restrict range of all parameters
  - only matching jobs shown in occurrence diagram
- Parameter occurrence histograms
  - show distribution of parameter values
  - counted selection in occurrence diagram

# The Demo: Threshold Working Range

- random color distribution
  - too many jobs fail to calculate hippocampus volume correctly
- middle area should be empty
  - what are properties of jobs there?
  - look at histograms
  - higher occurrence of low and high values of threshold
- threshold restriction to 1000–2500 yields visible separation

# The Demo: Warping Parameters

- visible red/green overlaps remain
- select these areas, examine warping parameters distribution
- higher occurrence of low values of `warp1` and `warp2`
- some images more sensitive to these warpings
- restriction to 12 only improves separation

# The Demo: Defective Input

- switch to logarithmic view
- visible anomaly at age 82–83
  - vertical orientation suggests one subject
  - selection shows only few involved
  - examination of inputs is feasible, wrong one easily detected
- invalidation of affected jobs
  - shortcut – only secondary (JPIS) data
  - immediate effect

- implementation
  - no image processing due to performance limitations
  - simple calculation of hippocampus volume from input parameters and subject age+cdr
- execution
  - 16-core machine (AMD Opteron 885), 64GB RAM
  - hosting 128 virtual machines, managed by PBS
  - 1 M jobs run within 1 week (in advance)
- job records
  - Logging and Bookkeeping service (EGEE)
  - application-specific data as *user tags*

## Behind the Scene – JP setup

- JPPS – primary data
  - two compressed files for each job
    - ▶ job execution record
    - ▶ user tags
  - necessary format to scale to billions of jobs
  - approx. 1 million of jobs (several GB of uncompressed data)
  - cannot be queried interactively
- JPIS – secondary data
  - populated in advance (several hours)
  - selection relevant for the experiment
    - ▶ only jobs of interest (still 1 M)
    - ▶ only relevant job attributes ( ~ 10)
  - further processing (e.g. round hippocampus volume to 0.3 bins)
- JPIS queries
  - SQL database – one of public interfaces
  - single table “group by” queries to populate GUI

## Behind the Scene – JP setup

- JPPS – primary data
  - two compressed files for each job
    - ▶ job execution record
    - ▶ user tags
  - necessary format to scale to billions of jobs
  - approx. 1 million of jobs (several GB of uncompressed data)
  - cannot be queried interactively
- JPIS – secondary data
  - populated in advance (several hours)
  - selection relevant for the experiment
    - ▶ only jobs of interest (still 1 M)
    - ▶ only relevant job attributes ( ~ 10)
  - further processing (e.g. round hippocampus volume to 0.3 bins)
- JPIS queries
  - SQL database – one of public interfaces
  - single table “group by” queries to populate GUI

## Behind the Scene – JP setup

- JPPS – primary data
  - two compressed files for each job
    - ▶ job execution record
    - ▶ user tags
  - necessary format to scale to billions of jobs
  - approx. 1 million of jobs (several GB of uncompressed data)
  - cannot be queried interactively
- JPIS – secondary data
  - populated in advance (several hours)
  - selection relevant for the experiment
    - ▶ only jobs of interest (still 1 M)
    - ▶ only relevant job attributes (  $\sim 10$  )
  - further processing (e.g. round hippocampus volume to 0.3 bins)
- JPIS queries
  - SQL database – one of public interfaces
  - single table “group by” queries to populate GUI

- switch to JPIS subscribed to receive updates
- invalidate defective jobs via primary data
  - record an “invalid” tag to 10,000 jobs
- submit replacement jobs
  - using restricted (discovered working) range of parameters
  - 256 jobs in few minutes
- monitor effect of both actions

- Job Provenance
  - fairly simple concepts
  - focus on large number of provenance records
  - high customizability
- The demo
  - approx. 1 million of jobs, all really executed
  - application-specific transformation achieved with JP config
  - data queryable via SQL, interactive response