Data Pipelines: Real Life Fully Automated Fault-tolerant Data Movement and Processing

George Kola
Computer Sciences Department
University of Wisconsin-Madison
kola@cs.wisc.edu
http://www.cs.wisc.edu/condor

Outline

- What users want ?
- Data pipeline overview
- Real life Data pipelines
  - NCSA and WCER pipelines
- Conclusions

What users want ?

- Make data available at different sites
- Process data and make results available at different sites
- Use distributed computing resources for processing
- Full automation and fault-tolerance

What users want ?

- Can we press a button and expect it to complete ?
- Can we not bother about failures ?
- Can we get acceptable throughput ?

- Yes… Data pipeline is the solution!

Data Pipeline Overview

- Fully automated framework for data movement and processing
- Fault tolerant & resilient to failures
  - Understands failures and handles them
- Self-tuning
- Rich statistics
- Dynamic visualization of system state

Data Pipelines Design

- View data placement and computation as full fledged jobs
- Data placement handled by Stork
- Computation handled by Condor/Condor-G
- Dependencies between jobs handled by DAGMan
- Tunable statistics generation/collection tool
- Visualization handled by DEVise

Fault Tolerance

- Failure makes automation difficult
- Variety of failures happen in real life
  - Network, software, hardware
- System designed taking failure into account
- Hierarchical fault tolerance
  - Stork/Condor, DAGMan
- Understands failures
  - Stork switches protocols
- Persistent logging. Recovers from machine crashes

Self Tuning

- Users are domain experts and not necessarily computer experts
- Data movement tuned using
  - Storage system characteristics
  - Dynamic network characteristics
- Computation scheduled on data availability

Statistics/Visualization

- Network statistics
- Job run-times, data transfer times
- Tunable statistics collection
- Statistics entered into Postgres database
- Interesting facts can be derived from the data
- Dynamic system visualization using DEVise

Real life Data Pipelines

NCSA

ASTRONOMY
UNIVERSITY OF ILLINOIS AT URBANA–CHAMPAIGN

- Astronomy data processing pipeline
  - ~3 TB (2611 x 1.1 GB files)
  - Joint work with Robert Brunner, Michael Remijan et al. at NCSA

WCER

Wisconsin Center for Education Research
at the School of Education, University of Wisconsin-Madison

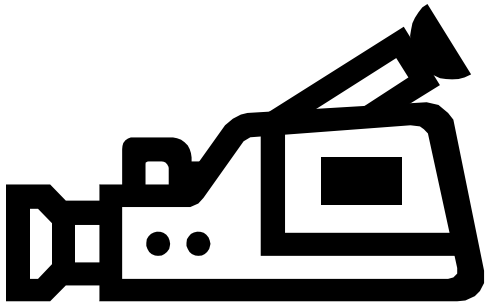

- WCER educational video pipeline
  - ~6TB (13 GB files)
  - Joint work with Chris Thorn et al at WCER

Digital Insight
WCER

DPOSS Data

NCSA

ASTRONOMY
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

- Palomar-Oschin photographic plates used to map one half of celestial sphere
- Each photographic plate digitized into a single image
- Calibration done by software pipeline at Caltech
- Want to run SExtractor on the images

NCSA Pipeline

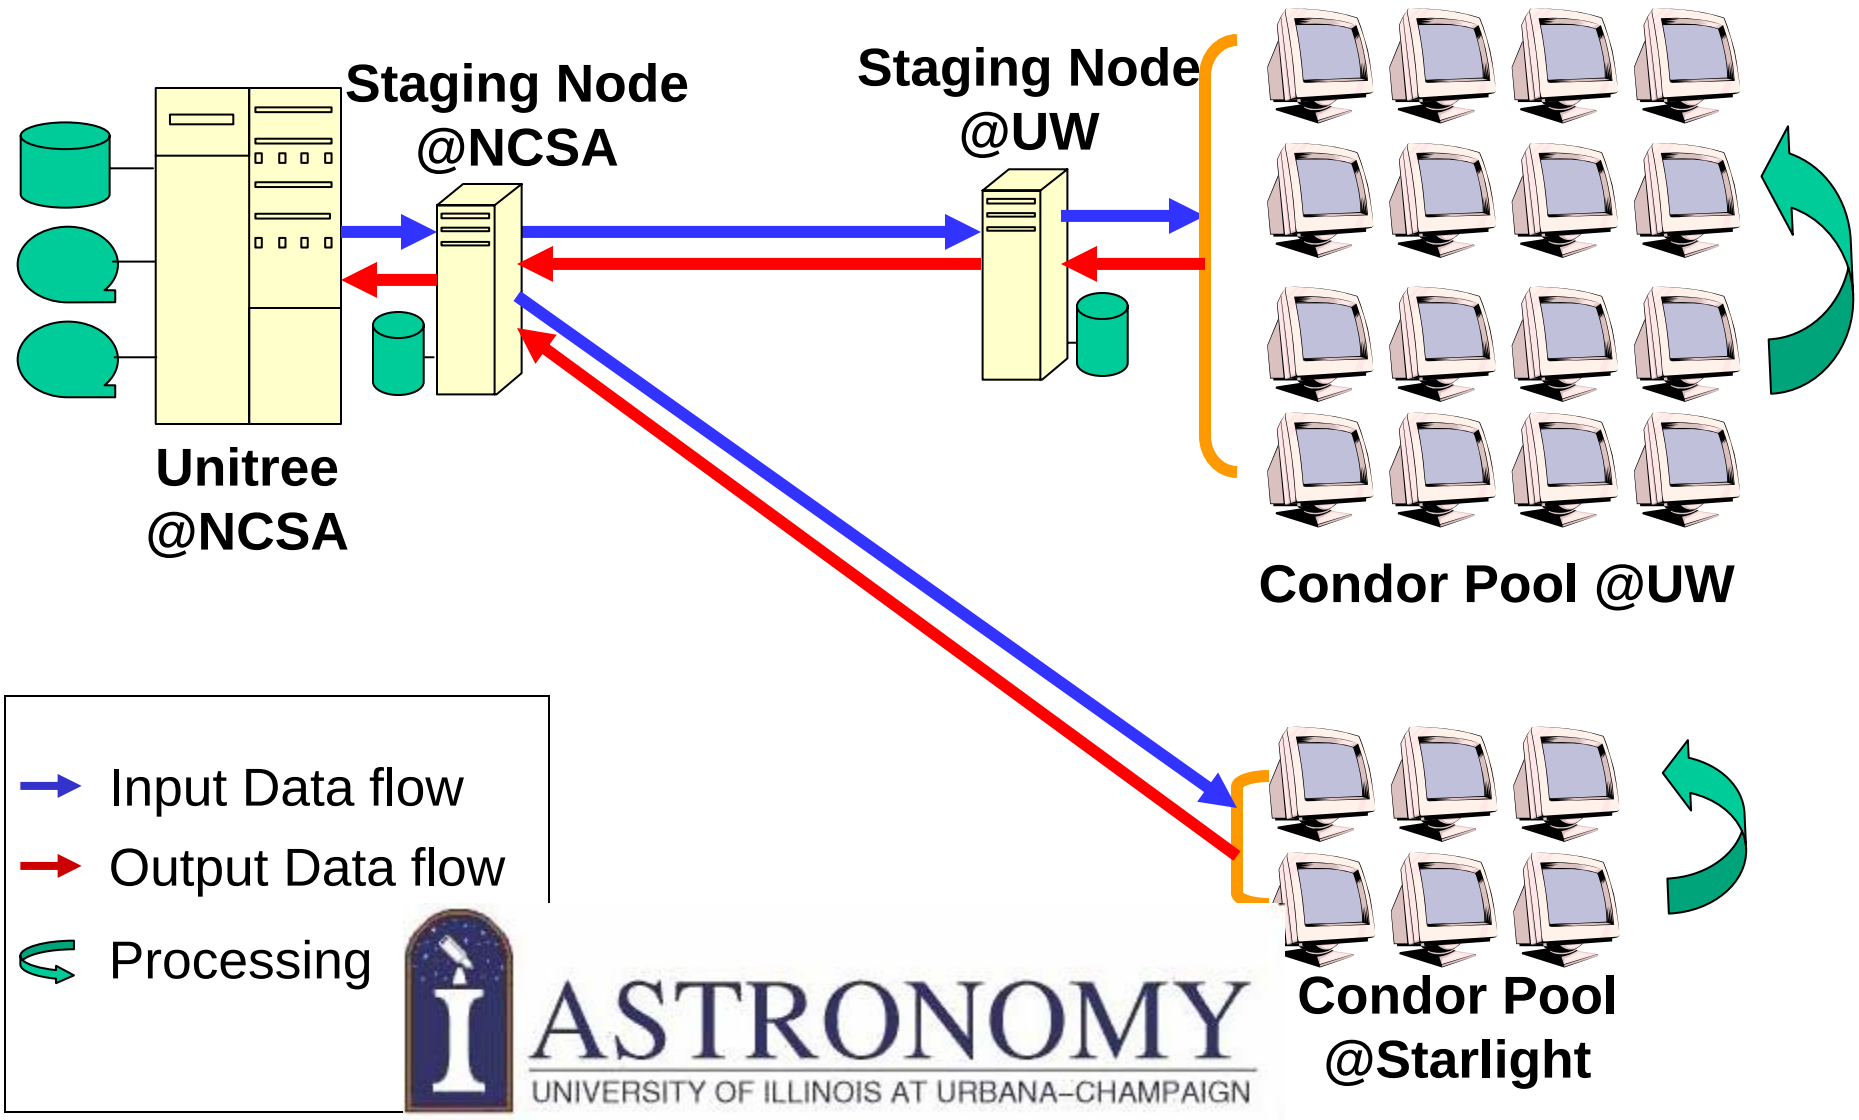

ASTRONOMY
UNIVERSITY OF ILLINOIS AT URBANA–CHAMPAIGN

NCSA Pipeline

- Moved & Processed 3 TB of DPOSS image data in under 6 days
  - Most powerful astronomy data processing facility!
- Adapt for other datasets (Petabytes): Quest2, CARMA, NOAO, NRAO, LSST
- Key component in future Astronomy Cyber infrastructure

ASTRONOMY
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

WCER Pipeline

- Need to convert DV videos to MPEG-1, MPEG-2 and MPEG-4
- Each 1 hour video is 13 GB

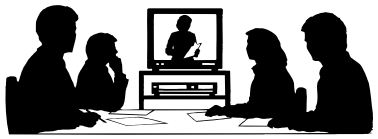

- Videos accessible through 'transana' software
- Need to stage the original and processed videos to SDSC

WCER Pipeline

- First attempt at such large scale distributed video processing
- Decoder problems with large 13 GB files
- Uses bleeding edge technology

Encoding	Resolution	File Size	Average Time
MPEG-1	Half (320 x 240)	600 MB	2 hours
MPEG-2	Full (720x480)	2 GB	8 hours
MPEG-4	Half (320 x 480)	250 MB	4 hours

WCER Pipeline

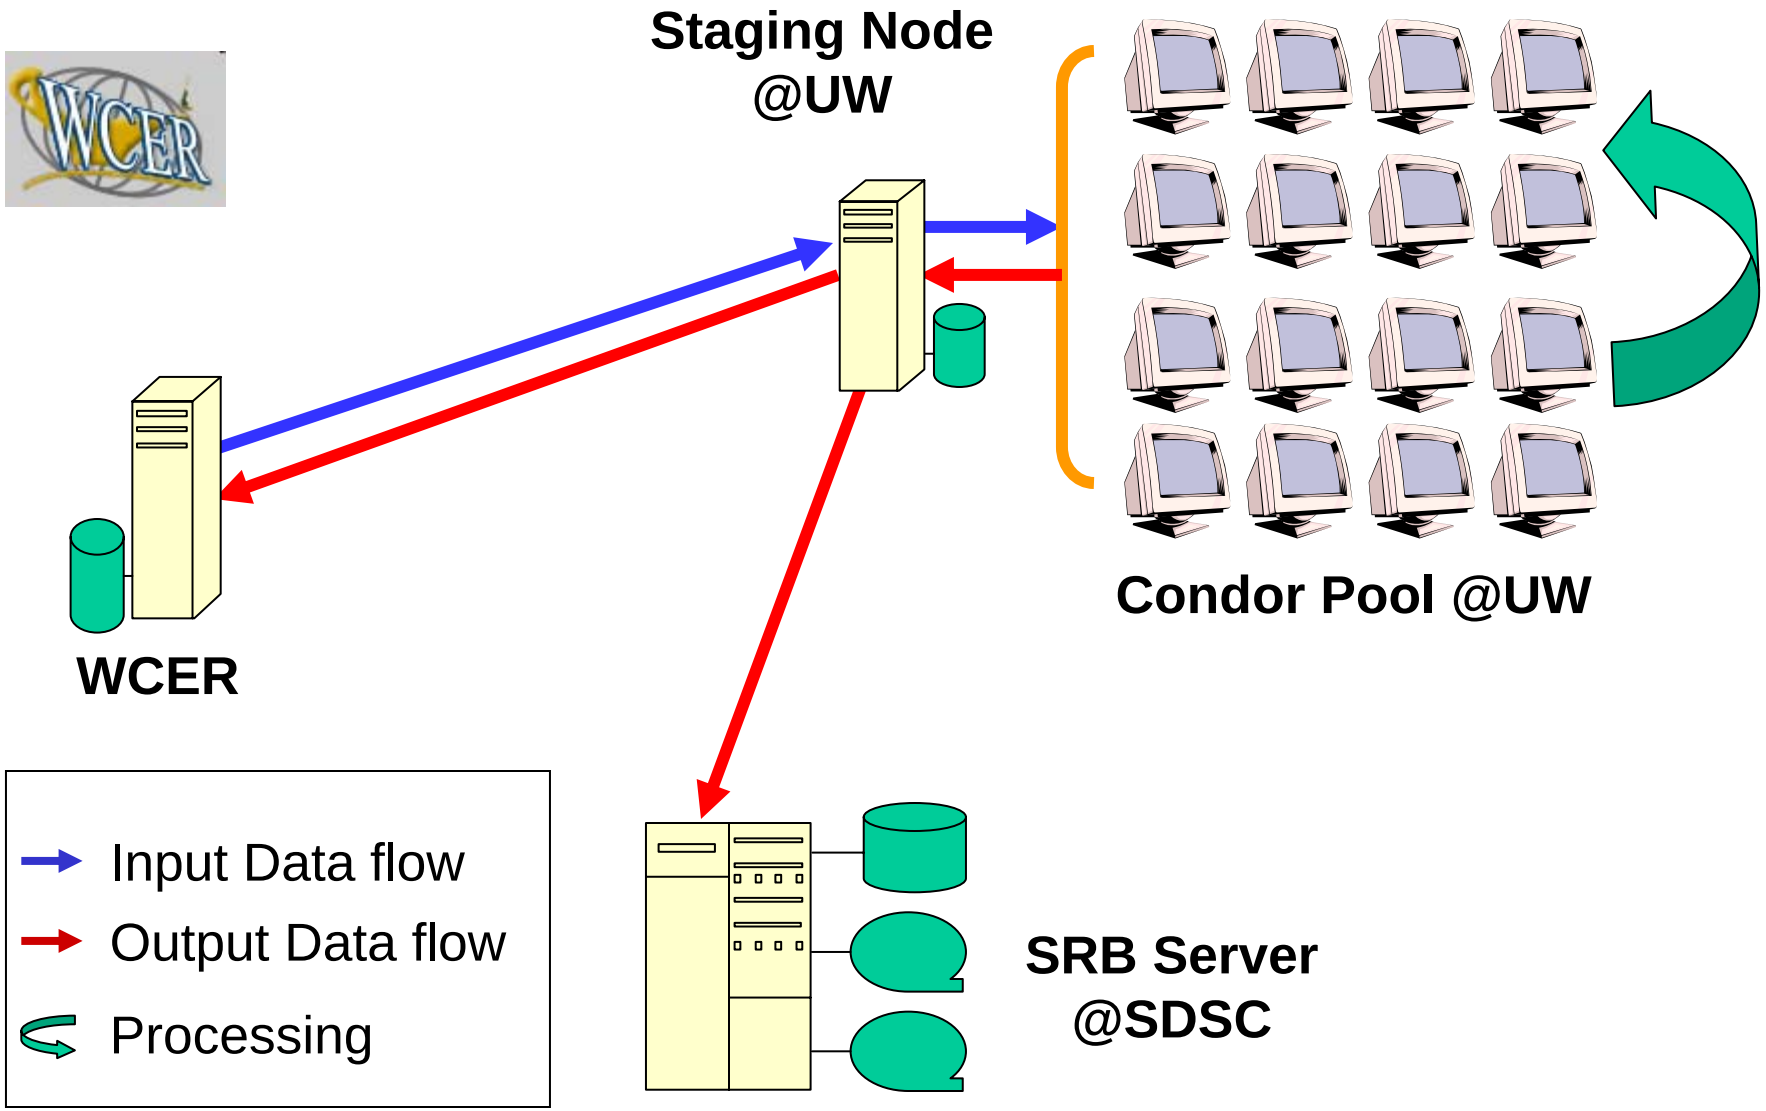

Conclusion

- Large scale data movement & processing can be fully automated!
- Successfully processed terabytes of data
- Data pipelines are useful for diverse fields
- We have shown two working case studies in astronomy and educational research
- We are working with our collaborators to make this production quality

Questions

- Thanks for listening

- Contact Information
  George Kola kola@cs.wisc.edu
  Tevfik Kosar kosart@cs.wisc.edu
  Office: 3361 Computer Science

- Collaborators
  NCSA Robert Brunner (rb@astro.uiuc.edu)
  NCSA Michael Remijan (remijan@ncsa.uiuc.edu)
  WCER Chris Thorn (cathorn@wisc.edu)