A First Look at Reproducibility and Non-Determinism in CMS Software and ROOT Data

Peter Ivie, Charles Zheng, and Douglas Thain
University of Notre Dame

CHEP 2016

What we usually work on:

HTCondor

Lobster Data Analysis

Global Filesystems

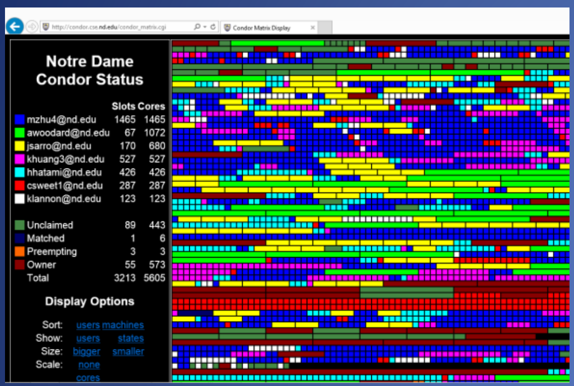

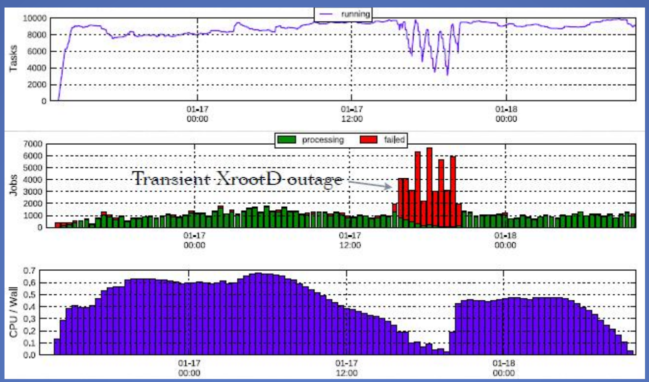

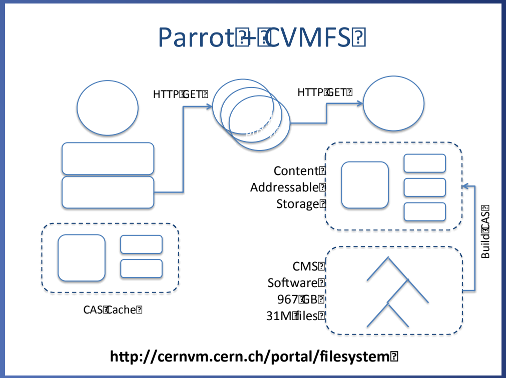

Today, a different question:

What happens if we attempt to run the same thing twice?

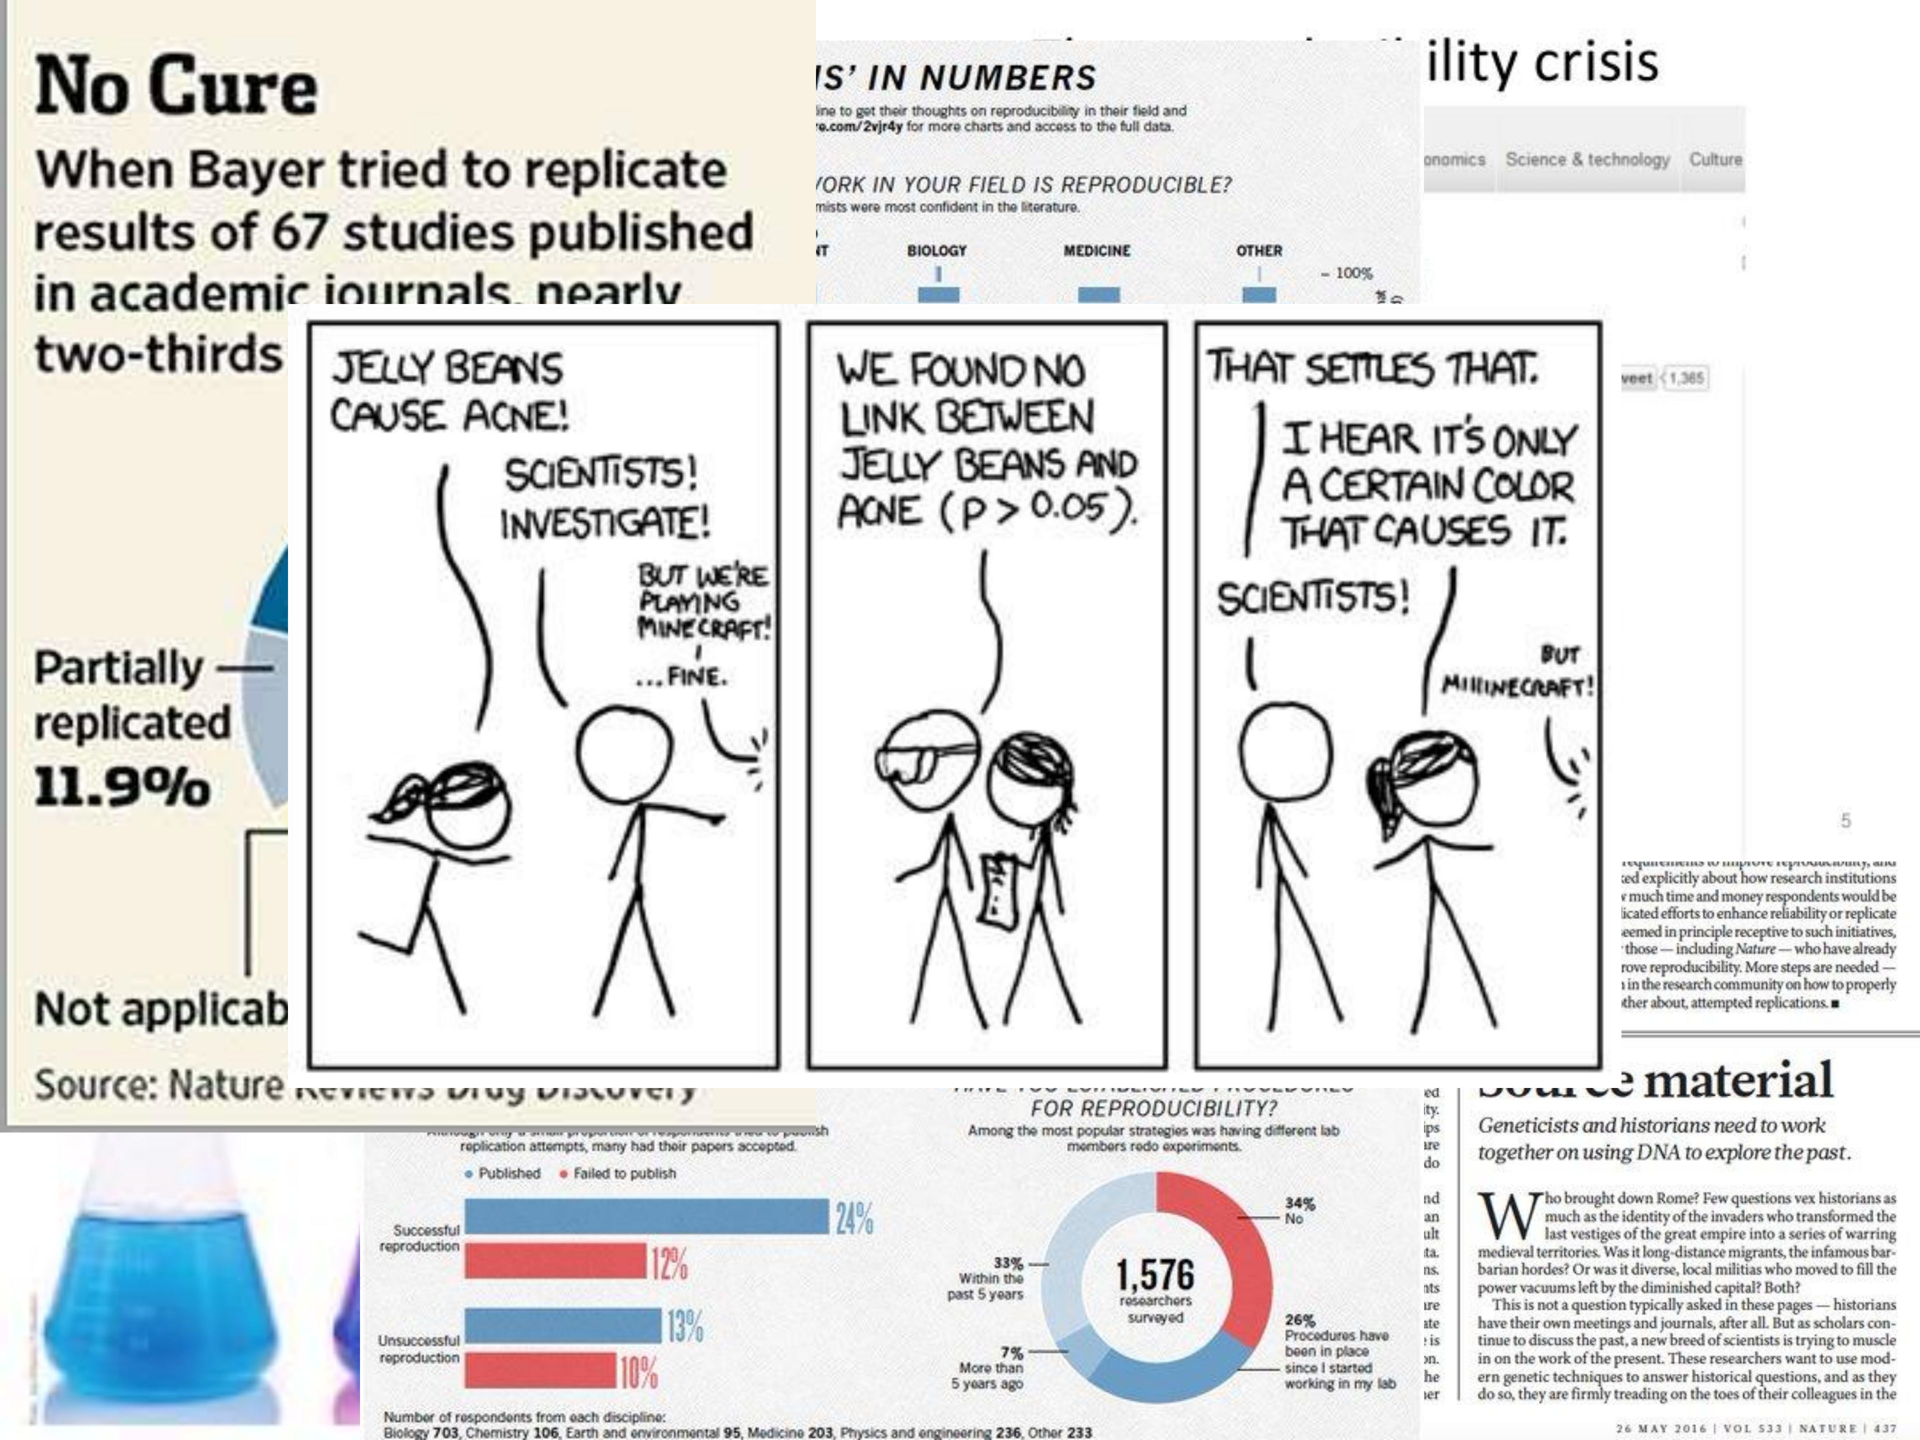

ility crisis
No Cure
When Bayer tried to replicate results of 67 studies published in academic journals, nearly two-thirds
Partially replicated 11.9%
Not applicab
Source: Nature
'S' IN NUMBERS
ine to get their thoughts on reproducibility in their field and
re.com/2vjr4y for more charts and access to the full data.
ORK IN YOUR FIELD IS REPRODUCIBLE?
mists were most confident in the literature.
BIOLOGY
MEDICINE
OTHER
100%
onomics
Science & technology
Culture
JELLY BEANS CAUSE ACNE!
SCIENTISTS! INVESTIGATE!
BUT WE'RE PLAYING MINECRAFT!
... FINE.
WE FOUND NO LINK BETWEEN JELLY BEANS AND ACNE (P > 0.05).
THAT SETTLES THAT.
I HEAR IT'S ONLY A CERTAIN COLOR THAT CAUSES IT.
SCIENTISTS!
BUT MIIIINECRAFT!
5
FOR REPRODUCIBILITY?
Among the most popular strategies was having different lab members redo experiments.
replication attempts, many had their papers accepted.
Published
Failed to publish
Successful reproduction
24%
12%
Unsuccessful reproduction
13%
10%
34% No
33% Within the past 5 years
1,576 researchers surveyed
26% Procedures have been in place since I started working in my lab
7% More than 5 years ago
Number of respondents from each discipline:
Biology 703, Chemistry 106, Earth and environmental 95, Medicine 203, Physics and engineering 236, Other 233
ed explicitly about how research institutions
much time and money respondents would be
licated efforts to enhance reliability or replicate
eemed in principle receptive to such initiatives,
those — including Nature — who have already
rove reproducibility. More steps are needed —
in the research community on how to properly
ther about, attempted replications.
material
Geneticists and historians need to work together on using DNA to explore the past.
Who brought down Rome? Few questions vex historians as much as the identity of the invaders who transformed the last vestiges of the great empire into a series of warring medieval territories. Was it long-distance migrants, the infamous barbarian hordes? Or was it diverse, local militias who moved to fill the power vacuums left by the diminished capital? Both?
This is not a question typically asked in these pages — historians have their own meetings and journals, after all. But as scholars continue to discuss the past, a new breed of scientists is trying to muscle in on the work of the present. These researchers want to use modern genetic techniques to answer historical questions, and as they do so, they are firmly treading on the toes of their colleagues in the
26 MAY 2016 | VOL 533 | NATURE | 437

DASPOS Data and Software Preservation for Open Science

www.daspos.org

ABOUT | PEOPLE | WORKSHOPS | RESEARCH | REPORTS

The massive data sets accumulated by High Energy Physics (HEP) experiments represent the most direct result of the often decades-long process of construction, commissioning and data aquisition that characterize this science. Many of these data are unique and represent an irreplaceable resource for potential future studies. Forward-thinking efforts for preservation are necessary now in order to achieve the relevant parameters, analysis paths and software to preserve the usefulness of these rich and varied data sets.

"Ten or 20 years ago we might have been able to repeat an experiment. They were simpler, cheaper and on a smaller scale. Today that is not the case. So if we need to re-evaluate the data we collect to test a new theory, or adjust it to a new development, we are going to have to be able to resuse it. That means we are going to need to save it as open data..."

Rolf-Dieter Heur 2008
Director General, CERN

First Workshop Scheduled

The first DASPOS Workshop has been scheduled for Thursday - Friday, March 21-22, 2013, at CERN. More information

Data and Software Preservation for Open Science, DASPOS, represents an initial exploration of the key technical problems that must be solved to provide appropriate data, software and algorithmic preservation for HEP, including the contexts necessary to understand, trust and reuse the data. While the archiving of HEP data may require some HEP-specific technical solutions, DASPOS will create a template for preservation that will be useful across many different disciplines, leading to a broad, coordinated effort.

Workshop 1
2012-12-17 19:11:04

WORKSHOP 1 Establishment of Use Cases for Archived Data and Software in HEP Date: Thursday-Friday...

Workshop 2
2012-12-17 19:11:04

WORKSHOP 2 Survey of Commonality with other Disciplines Attendees: Broad participation from many...

Discovery and Coordination

Series of highly-structured public workshops to define, discuss and document the details of data and software preservation

Prototyping and Experimentation

Key areas of reserach: data and query models and software sustainability models

The DASPOS Team

Computer science experts, experienced digital librarians, and experts in data-intensive fields, such as physics, astrophysics and bioinformatics

Different Layers of Reproducibility

(1) Can someone else follow simple instructions and get the same software environment?

– Requires careful dependency management.

(2) Can I run the exact same code on the same computer twice and get the same result?

?

– Requires deterministic computing techniques.

. . .

(99) Can someone else construct equivalent expts and observe the same physical laws?

ATLAS
vs
CMS

– Difficult to do this without steps 1-98 working.

Some Reasons for Non-Determinism

- Provenance data intermixed with output.
- Unique identifiers embedded in objects.
- Pseudo-random numbers seeded by environmental data (pid,uid,time)
- Entropic randomness (/dev/urandom)
- Parallelism results in re-ordering of data.
- I/O interactions with OS not consistent.
- OS/Compiler details affect ordering/precision of floating point operations.

A very preliminary look…

An initial experiment.

CVMFS + SCRAM for consistent environment

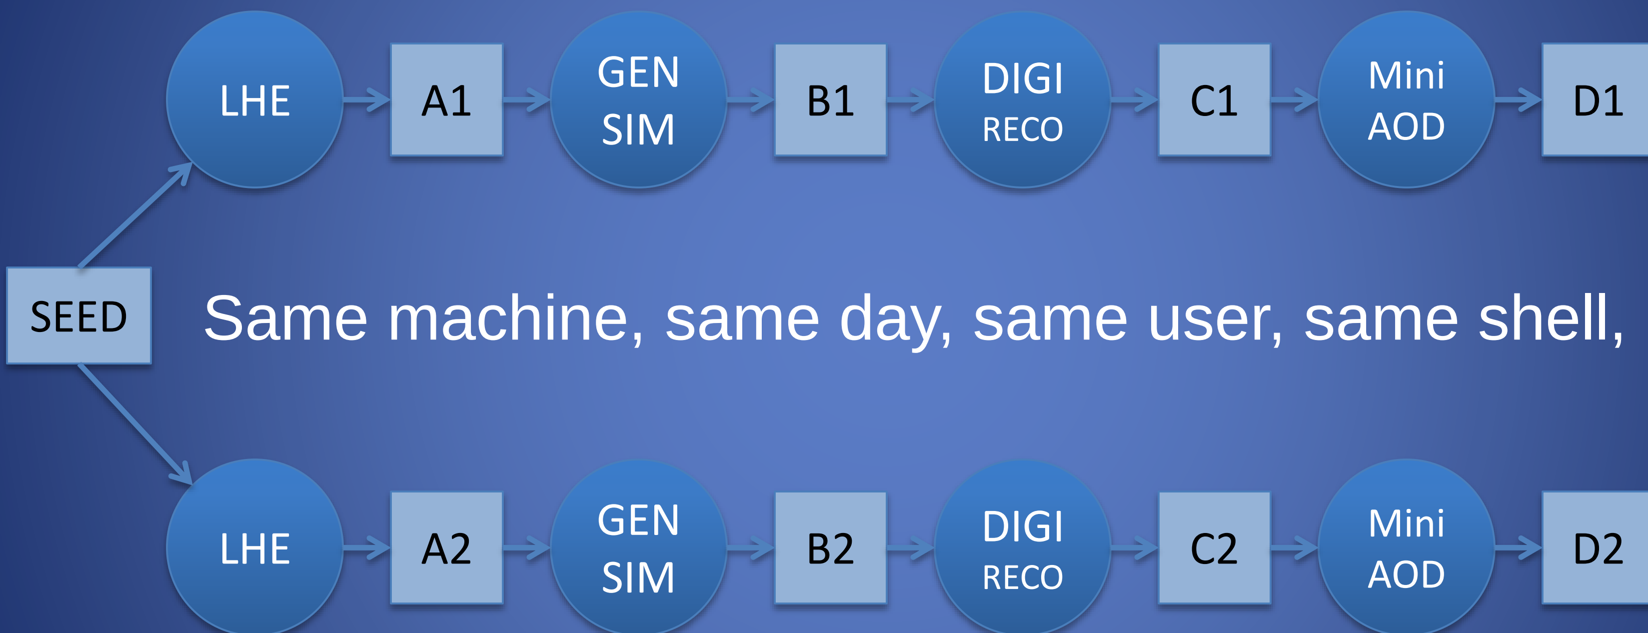

Same machine, same day, same user, same shell, . . .

Not one of the output files was the same!

What constitutes ***the same output***?

The same bytes?
(take a checksum)

```
0000000 5089 474e 0a0d 0a1a 0000
0000020 0000 8002 0000 e001 0308
0000040 00d6 0100 5032 544c ff45
0000060 a0a0 00ff 0000 00c0 8000
0000100 c0ee 0040 c8c8 4100 e169
0000120 80c0 30ff 8060 008b 4000
0000140 d4ff 2aa5 ff2a 00ff e040
0000160 331a 3333 4d4d 664d 6666
0000200 b3b3 c0b3 c0c0 cccc e5cc
```

```
0000000 5089 474e 0a0d 0a1a 0000
0000020 0000 8002 0000 e001 0308
0000040 00d6 0100 5032 544c ff45
0000060 a0a0 00ff 0000 00c0 8000
0000100 c0ee 0040 c8c8 4100 e169
0000120 80c0 30ff 8060 008b 4000
0000140 d4ff 2aa5 ff2a 00ff e040
0000160 331a 3333 4d4d 664d 6666
0000200 b3b3 c0b3 c0c0 cccc e5cc
```

The same data?
(ignore metadata)

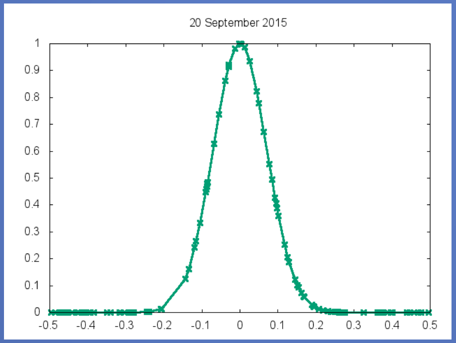

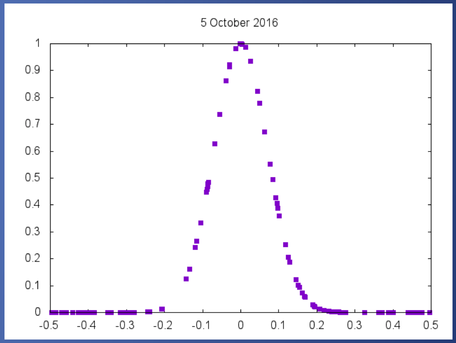

The same physics?
(fit to a model)

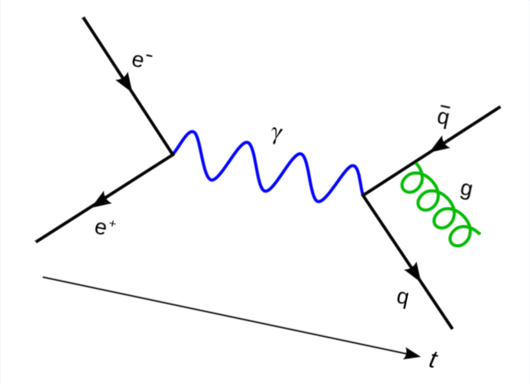

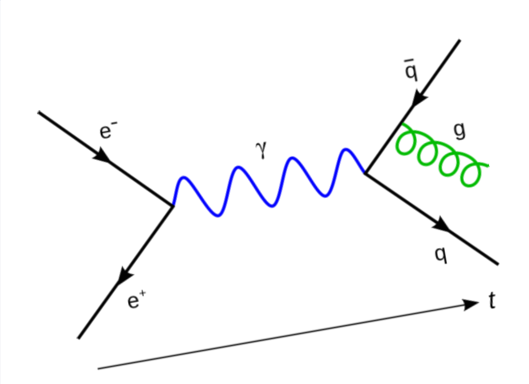

Structure of a Typical ROOT File

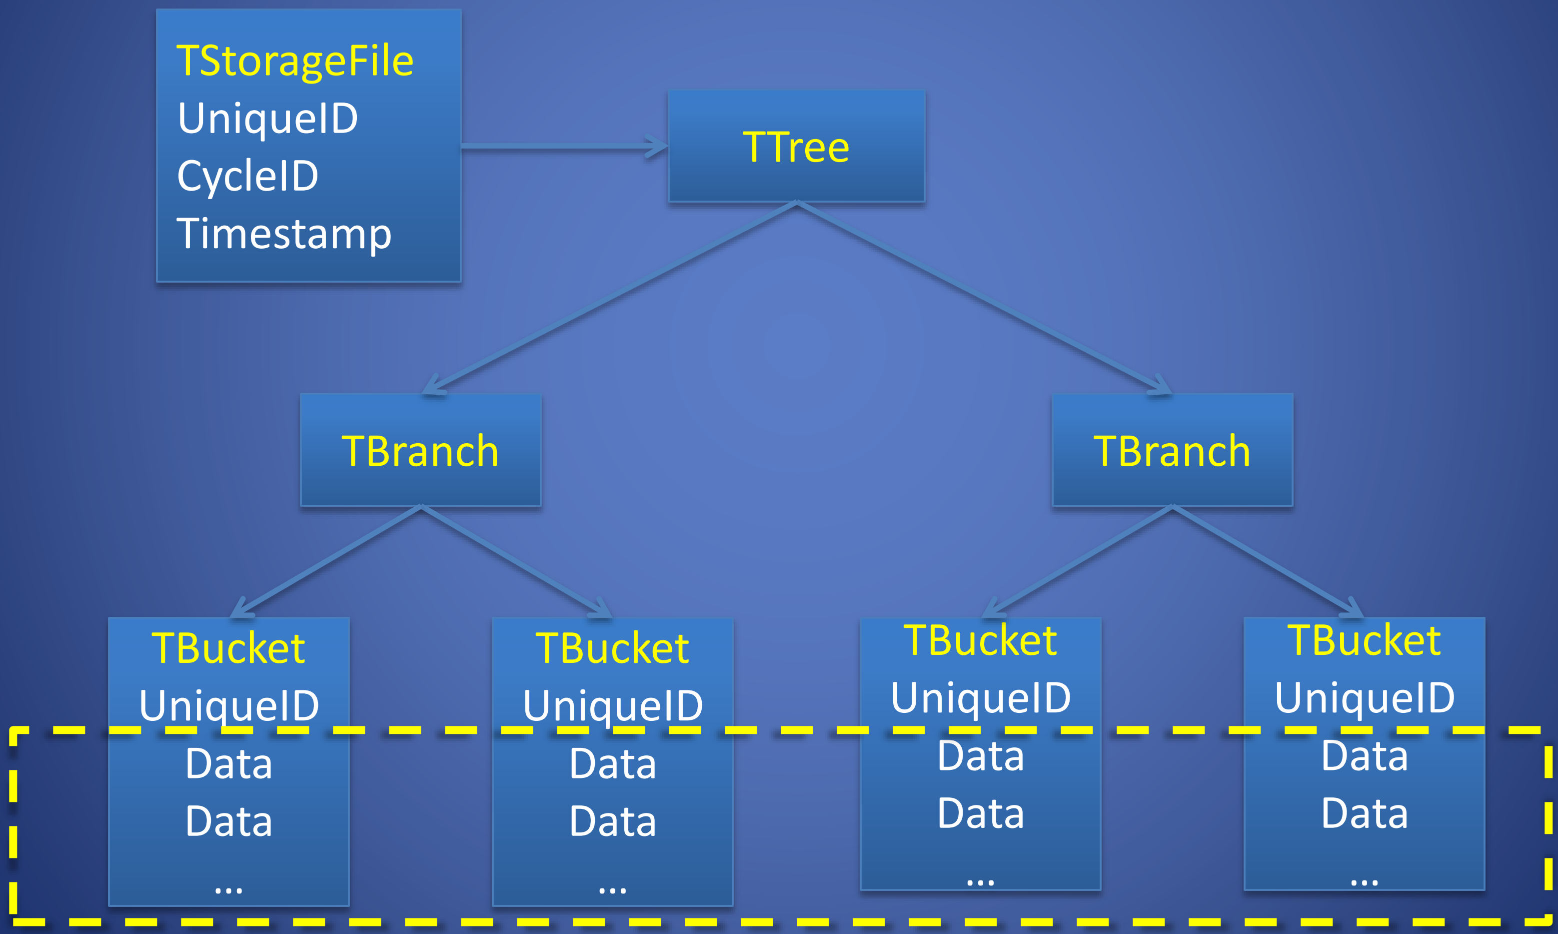

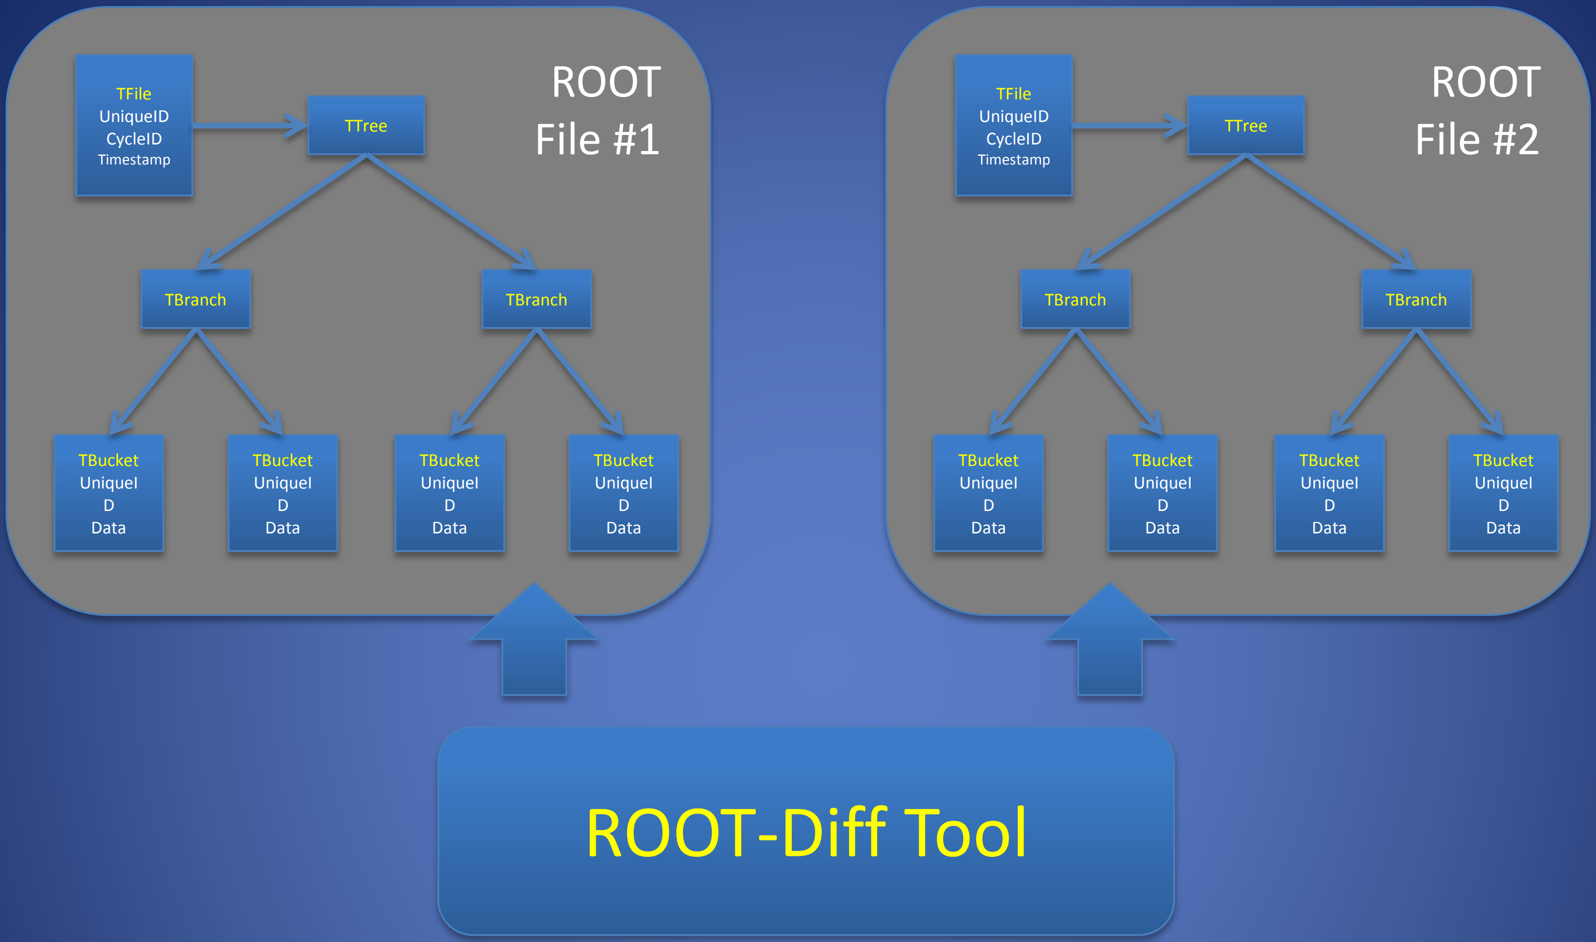

1. Different structure.
2. Same structure, different content.
3. Same structure, same content, different metadata.
4. Same bytes everywhere.

Now, what’s different?

Stage	# Objects	Ignored	Not Equal	Structural Equal	Content Equal	Bitwise Equal
LHE	605	9	4	592	580	0
GENSIM	717	9	1	707	694	0
DIGI – 1	1171	9	1	1161	1158	0
DIGI – 2	6563	9	21	6533	6526	0
DIGI – 3	204	21	7	176	175	0
MINIAOD	2723	9	4	2710	2671	0

Physics data is not the same.

What is the cost of root-diff?

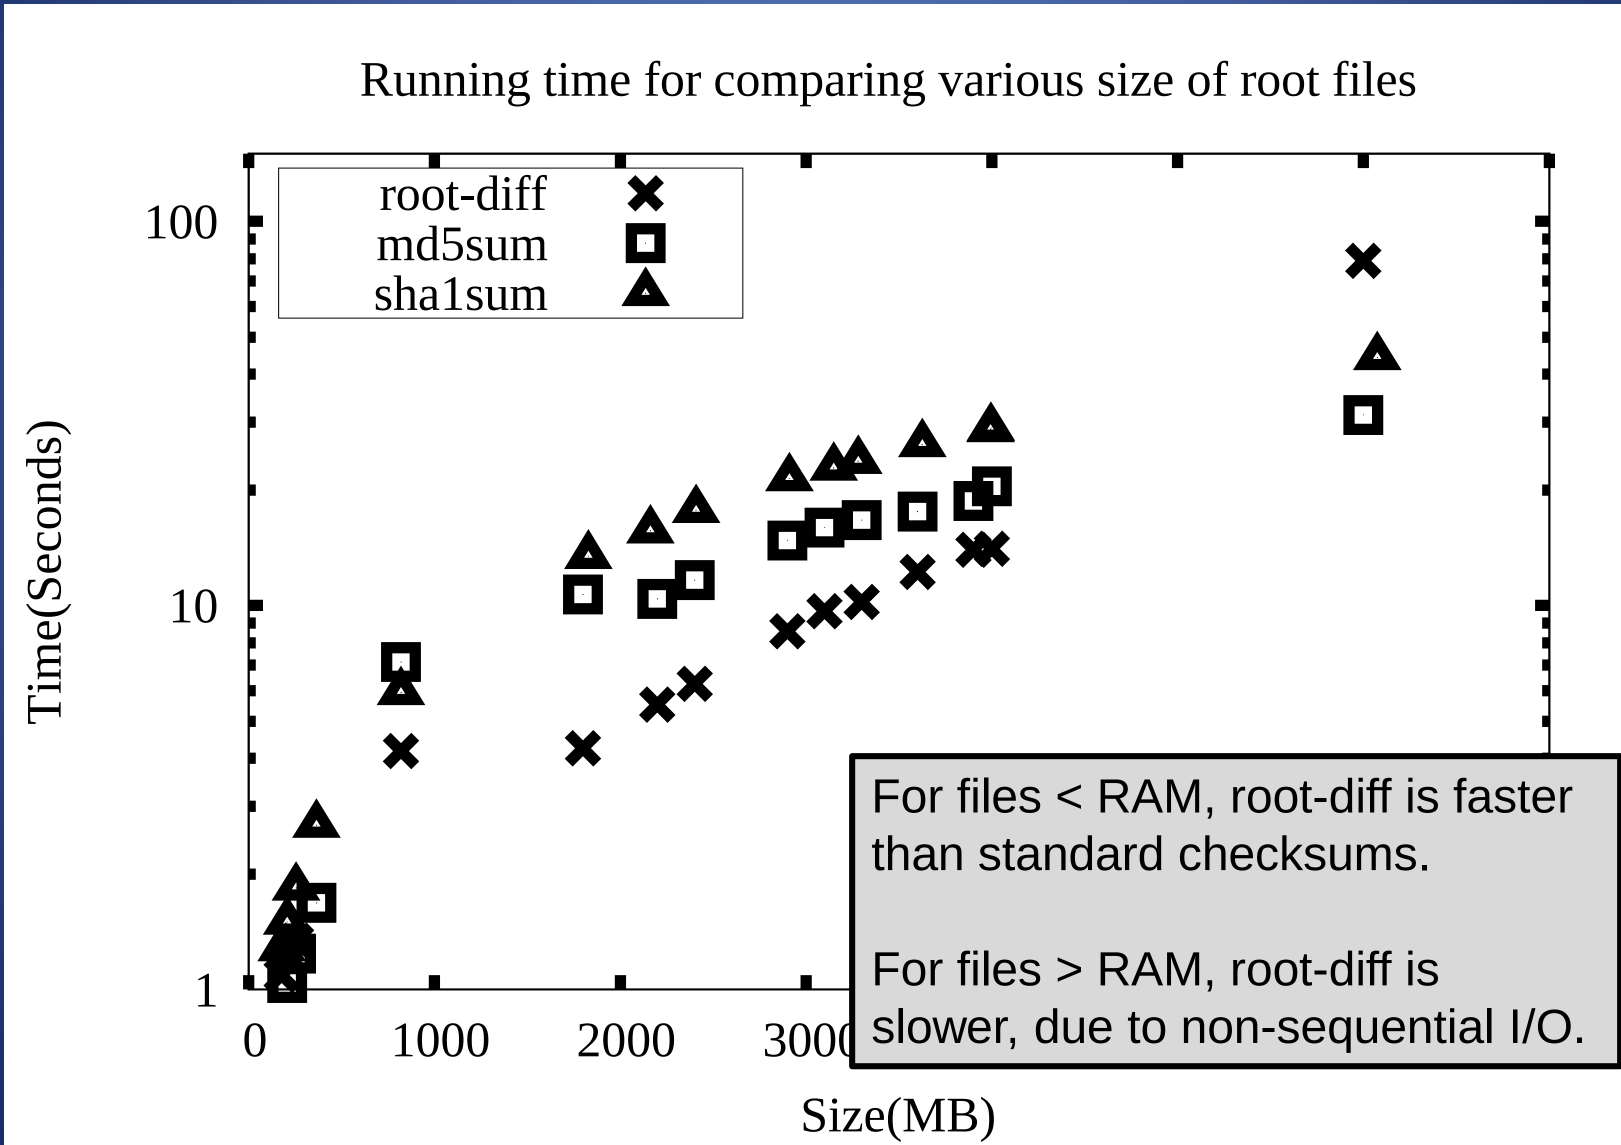

Can we eliminate or prevent some kinds of non-determinism?

Parrot Modifies Syscalls

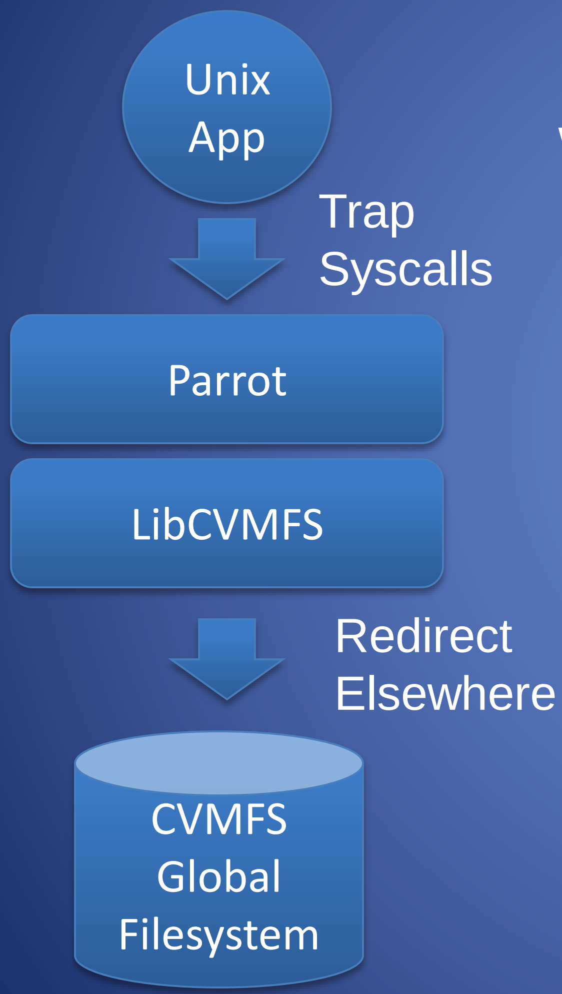

Use Parrot to mount remote filesystems without requiring root access.

```
% parrot_run bash

% cd /cvmfs/cms.cern.ch
% source cmsset_default.sh
% scram

. . .
```

http://ccl.cse.nd.edu/software/parrot

Use Parrot to Eliminate Some Sources of Non-Determinism

```
parrot_run --time-warp
           --uid 0
           --hostname thematrix
           --mount /dev/urandom /some/fixed/data
```

```
% whoami
root
% uname
thematrix
% date –utc
Mon Jan  1  00:00:05 UTC 2001
```

http://ccl.cse.nd.edu/software/parrot

Does it work?   Almost!

“Time Warp” makes LHE produce bitwise-equal results.

(But only for <121 events.)

Results of a First Look

- ROOT-DIFF is a general algorithm that can determine equivalence of ROOT files, ignoring the incidental metadata inherent to ROOT.
- CMS codes (at first look) *do not* produce identical physics data on consecutive runs, when compared with ROOT-DIFF.
- "Time Warp" mode causes LHE to give *bitwise equal* results for a small number of events, therefore the code is sensitive to OS effects, even with a constant seed.
- We have not fully isolated the factors that result in non-deterministic behavior!
- Going forward:
  - Push towards one good deterministic example to be shared.
  - Clearly identify all conventional sources of non-determinism.
  - Evaluate effect of increasing concurrency on determinism.

Example code here:

https://github.com/cooperative-computing-lab
/root-diff
/repro-example-chep-2016

Questions and Discussion Welcome

Peter Ivie
pivie@nd.edu

Charles Zheng
czheng2@nd.edu

Douglas Thain
dthain@nd.edu

@ProfThain