Summary
pre-GDB / Morning Session

Jeff Templon

2015.03.11

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Sessions

- Morning session co-organized with NLeSC
- Introduce NLeSC to HEP and HEP to NLeSC; investigate project synergies with Astronomy

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Tuesday, 10 March 2015

09:30 - 12:30

HEP and Other Sciences

Discussion on computing challenges with other sciences, co-organized with Netherlands e-Science Center

Convener: Jeff Templon (NIKHEF (NL))

09:30 **Welcome** *10'*

Speaker: Jeff Templon (NIKHEF (NL))

Material: **Slides**

Link to agenda

09:40 **NLeSC and Challenges in Astronomy** *30'*

Speaker: Dr. Rob van Nieuwpoort (Netherlands e-Science Center)

Material: **Slides**

10:10 **Performant Scientific Software: Theory vs Reality** *30'*

Speaker: Gerhard Raven (NIKHEF (NL))

Material: **Slides**

10:40 **Coffee** *20'*

11:00 **Dealing with High-Energy Physics Data in 2020 and beyond** *30'*

Speaker: Massimo Lamanna (CERN)

Material: **Slides** **Video**

Core eScience Technologies

- Optimized Data Handling
- Big Data Analytics
- Efficient Computing

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Imaging pipeline

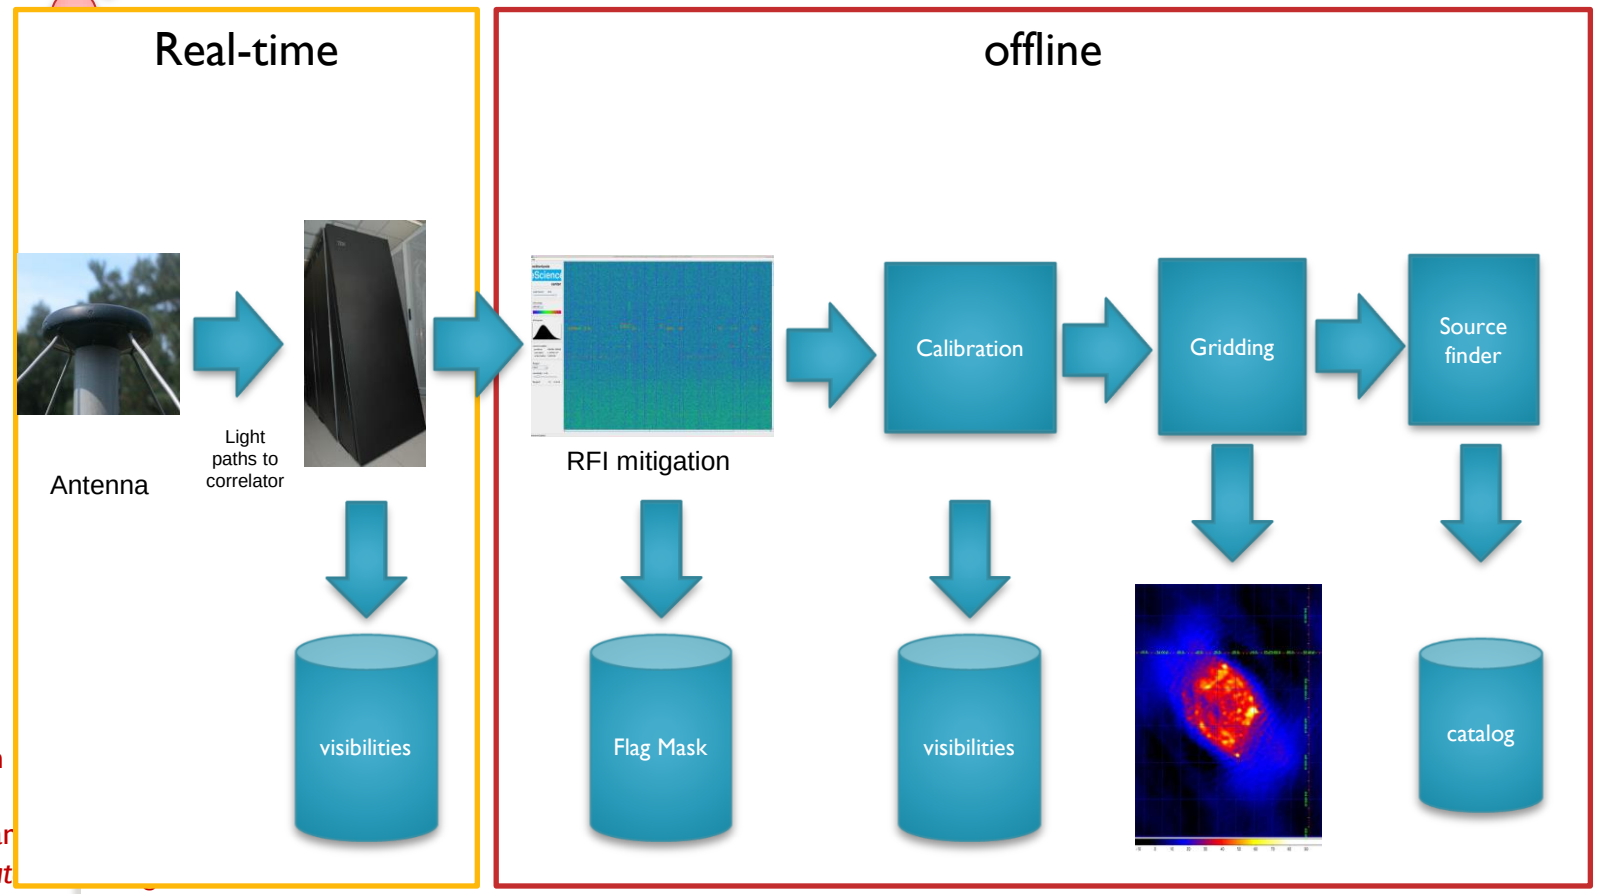

J. Templon
Nikhef
Amsterdar
Physics Dat
Group

Imaging pipeline: scaling up

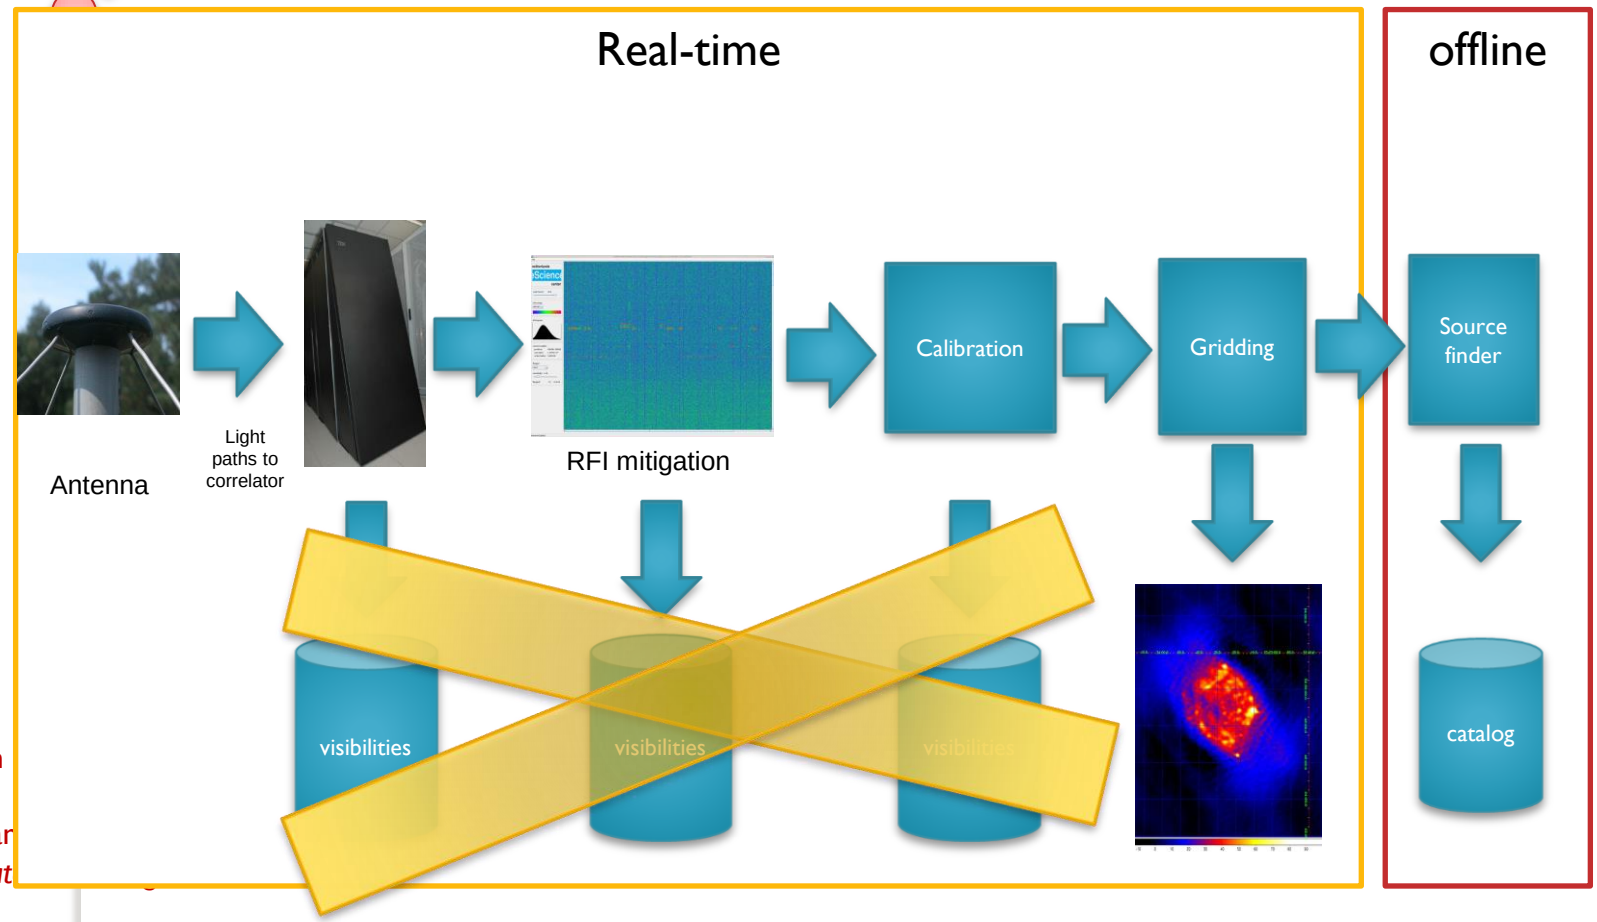

J. Templon
Nikhef
Amsterdar
Physics Dat
Group

HEP & Astro

- NLeSC – scaling up means revisiting 70's algorithms back when computers were not fast enough
- LHCb – "inside of the computer becoming important again"

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

The Challenge of Modern Computing

Exploit all 7 "parallel" dimensions of modern computing architecture for HPC

–Inside a core (climb the ILP wall)

1. Superscalar: Fill the ports (maximize instruction per cycle)
2. Pipelined: Fill the stages (avoid stalls)
3. SIMD (vector): Fill the register width (exploit SSE, AVX)

–Inside a Box (climb the memory wall)

4. HW threads: Fill up a core (share core & caches)
5. Processor cores: Fill up a processor (share of low level resources)
6. Sockets: Fill up a box (share high level resources)

–LAN & WAN (climb the network wall)

7. Optimize scheduling and resource sharing on the Grid

HEP has been traditionally good (only) in the latter

Vincenzo Innocente, CMS, 2012

J. Templon
Nikhef
Amsterdam
Physics Data P
Group

Every Cycle Counts!

- Pattern Recognition
  - Hough transform: vectorize Hough function
- Track Reconstruction in magnetic field
  - Reduce precision of magnetic field (
    to single precis
  - Vectorize interp
  - Reduce precisio
    computation (en
  - But keep double
    integration!!!
  - Vectorization of ma
- Everywhere:
  - Reduce copying (p
    std::experimental::st
  - Don't dynamic_cast
    have to, revisit the de
  - Separate & inline fast path from slow path
  - Use contiguous containers (std::vector, std::array, …)
    — and if not, rewrite so you can…
  - Fight (heap) memory abuse — 'reserve' properly,
    move 'hot' data together, bitfields instead of lots of
    bools
  - Write vectorized APIs — avoid 'raw loops'

1. Measure & Benchmark!
2. Don't do more work than strictly needed
3. Improve memory usage
4. Vectorization — utilize SIMD

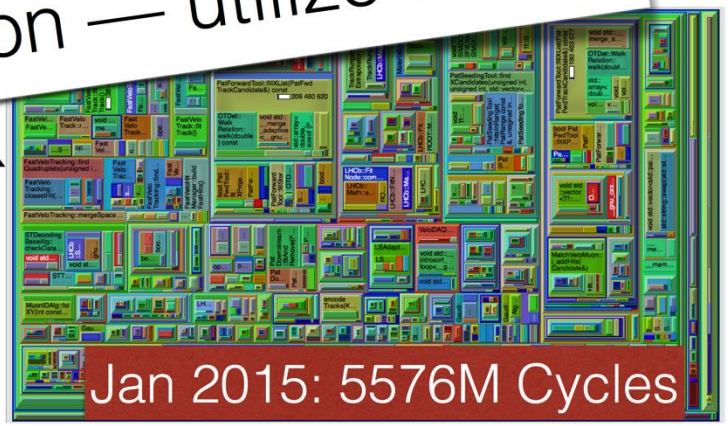

J. Templon
Nikhef
Amsterdam
Physics Data P
Group

Take-home messages

- Prepared for the next steps towards 2020 (and beyond)
  - LHC and High-Luminosity LHC
  - Open to collaboration with sister sciences
- Multiple federated data centres
  - Xroot/HTTP
  - Large-farm operations understood
  - Federation → Business continuity
  - EOS as heavy-duty system for data analysis

- "Unified" storage from LHC data analysis farms to desktop computing
  - EOS + CERNBOX
- Experience in archiving and preserving scientific data
  - 100 PB repository / 1 PB/week (run1) with CASTOR
    - Substantially growing (LHC run 2)

My Conclusions I

- NLeSC interested in further discussions
- WLCG audience seemed interested in / impressed with NLeSC
- Number of synergies with Astro
- More info please visit NLeSC website
- Please contact them if you're interested in collaboration (and let us (Nikhef) know!)
- Thanks to speakers and participants for the excellent presentations and discussions

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Lunch Conclusions II

- NLeSC staff involved in SoftwareX
  - Has anybody published root or Geant yet?
- Some overlap with HSF
  - Sustainability, career paths sw people

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group