GLAST Large Area Telescope:

Collaboration Compute Needs

Richard Dubois
Stanford Linear Accelerator Center
richard@slac.stanford.edu

Outline

- Introduction to the Service Challenges: some context for need
- Computing Resource Projections
- Major Milestones and Manpower Needs

From Steve's DC2 closeout talk

Work to be Done: Responsibilities

Systematic & Sensitivity Studies

Analysis Coordinator and Science groups

- pt sources, extended sources, transients; upper limits
- diffuse analyses
- variability (incl. pulsars)
- neighboring sources
- flaring & diffuse effects
- **focus on 1st papers analyses**

Other Studies

C&A group and ISOC jointly

- PSR ("handoff review") performance
- analysis tuning (signal/bkgd, quality knobs by topic)
- update simulation (s/c model, tune from beam test and IA data…)
- first light observations (simulate point, then scan); early ops analyses
- effects of burst repoints
- sky survey strategy checks
- background fluxes evaluation early ops

Readiness

ISOC

- digital data problems
- instrument problems (bad channels, wrong rates, recognizing a few wrong constants, …)
- ASP (aka quicklook running and burst handling)
- receiving data dumps, running the pipeline, benchmarking resources and times, reliability
- idiosyncrasies vs. problems
- day(s) in the life
- performance monitoring
- documentation

Collaboration participation needed for each of these!

Simplified Diagram for ISOC Data Flow

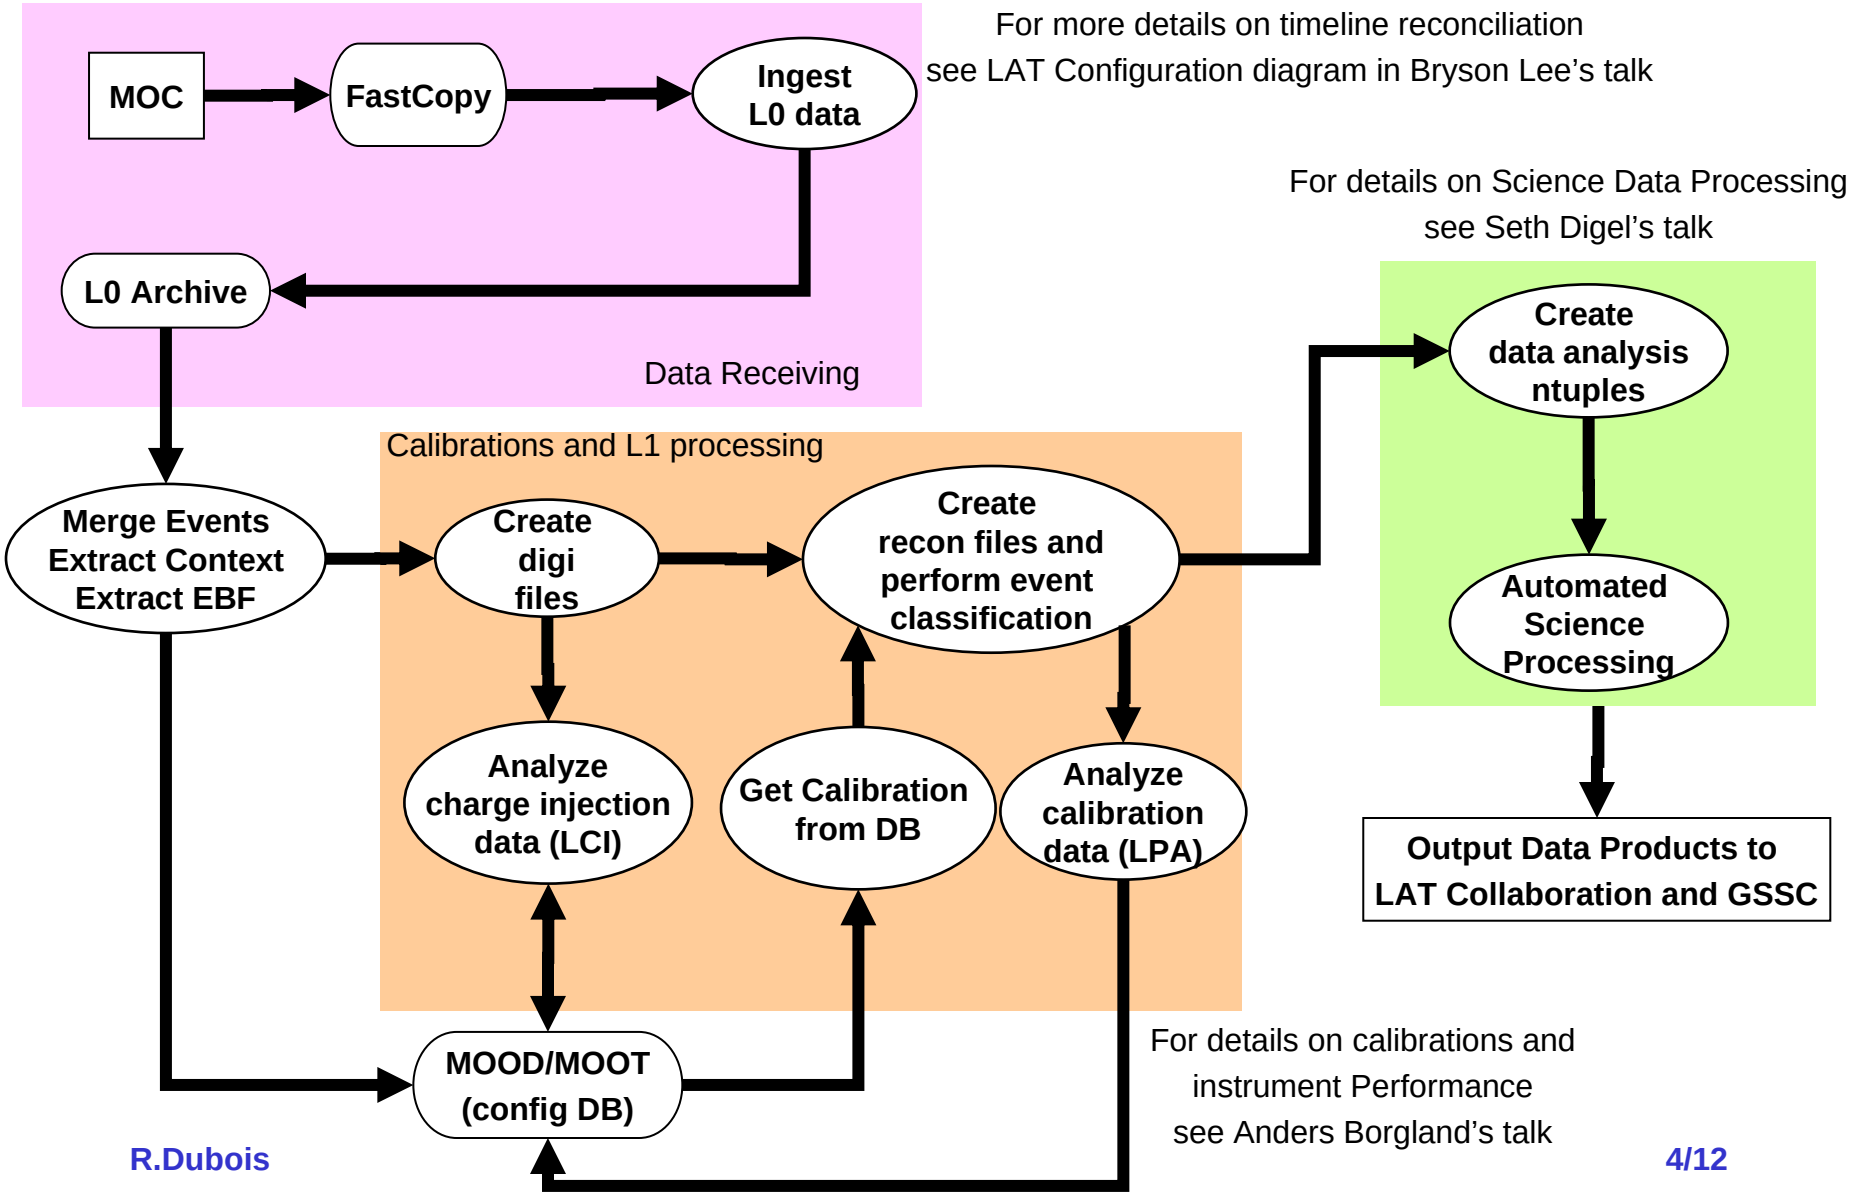

In a NutShell

QuickTime™ and a
TIFF (LZW) decompressor
are needed to see this picture.

To be demonstrated by ~Feb 2007

Integrate with Science Group efforts and GRTs

We are identifying the full set of functions to test with SCs:
https://confluence.slac.stanford.edu/display/ISOC/Test+List

Connection to Science Groups

- There are moves afoot for several datasets to be created for the Science Groups:
  - 1 yr obssim survey run
  - 55 day obssim run, with pointed observations:
    - Earth occultation not currently in exposure calculations
    - ARR
    - New orbit
    - LAT/SC misalignment
    - Final DC2 IRFs
  - 55 day Gleam run – ibid
    - Background interleave for pointed observations
  - 1 year Gleam run
    - Potentially huge backgrounds run needed
    - Use of Lyon, Italian farms?

- Possible timescales are October, December and February for these 3 datasets

Rough Timeline

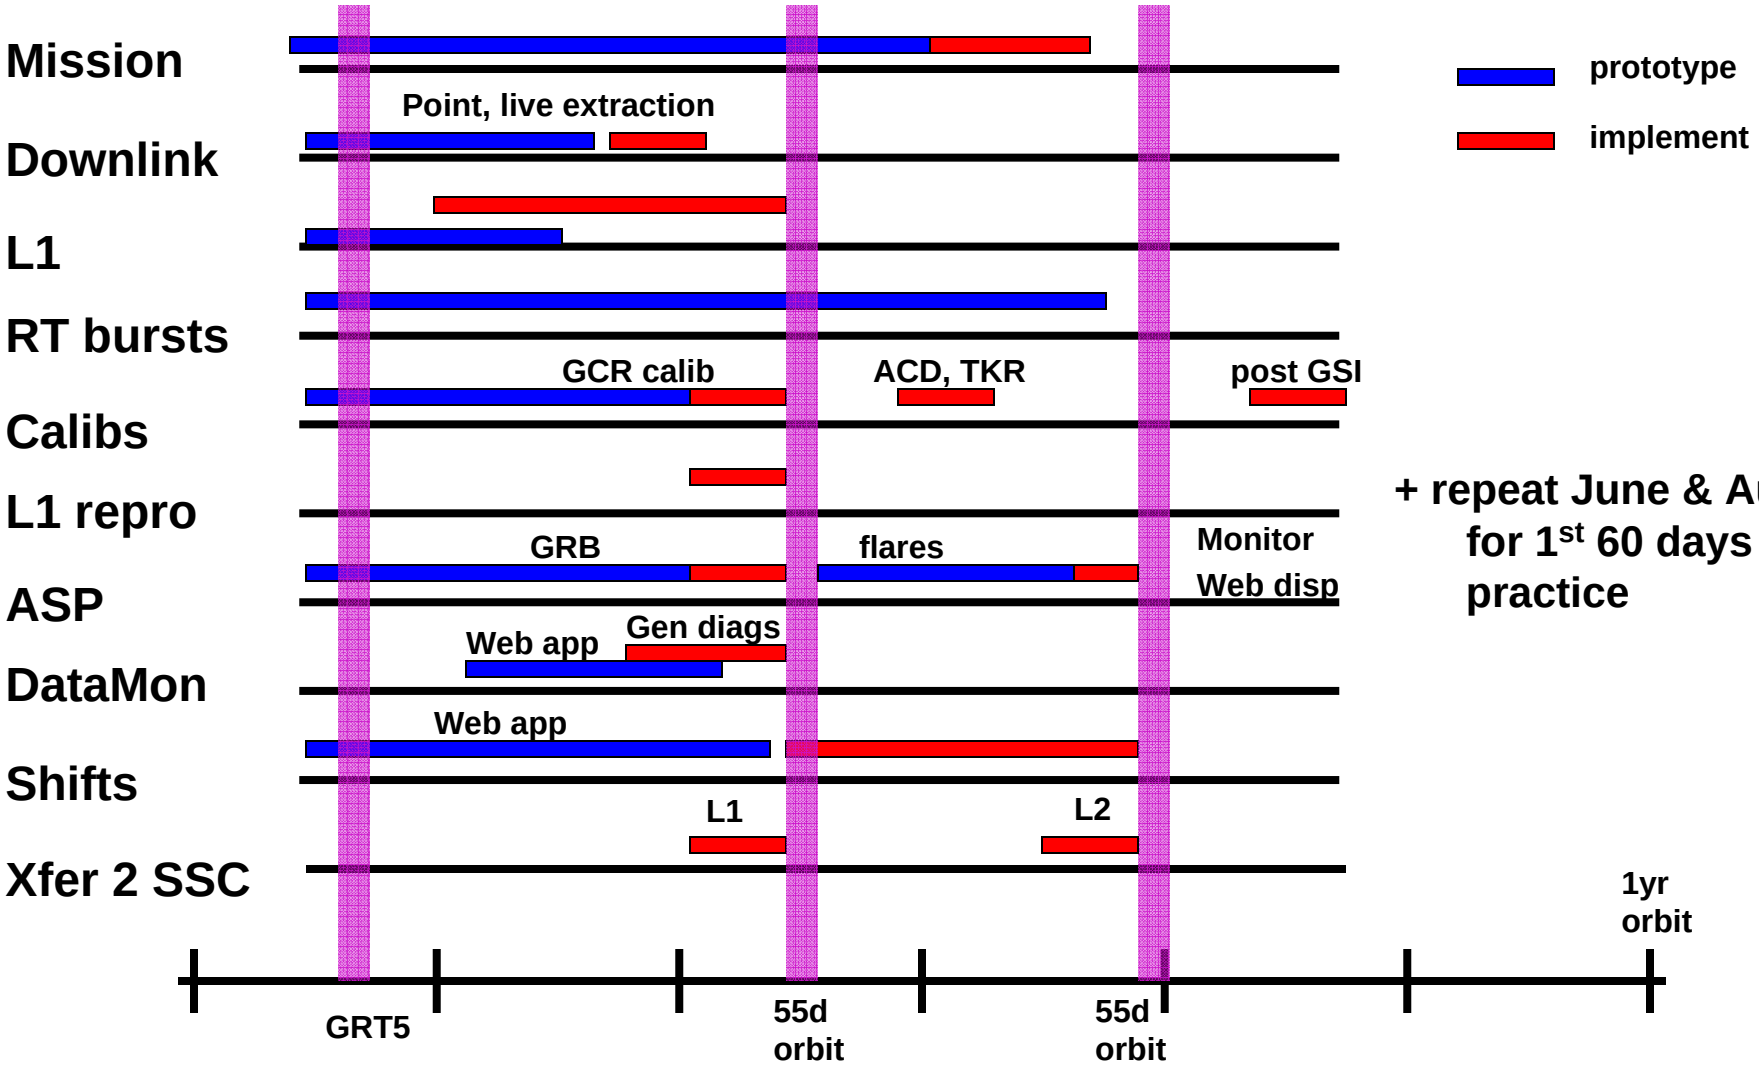

+ repeat June & Aug for 1st 60 days practice

Aiming for initial tests of everything by December with Feb as fallback

Current Compute Resources

In 2nd year of projected annual $300k Capital Equipment Projects
- Supplying, batch farm disk & CPU, as well as dedicated servers
- Optimize purchases based on best deals SCS can come up with

- 100 TB disk (33 remaining) –
  - LAT Commissioning (~50)
  - DC2 (10)
  - Beamtest (5)
  - Pre-ship review (5)
  - Infrastructure needs (code builds; system tests; user disk)
  - Hoping IFC to fund 50 TB more in CY06 - current price is $4k/TB

- Tremendous use of SLAC Batch farm!
  - 20 AMD Opteron dual core dual CPU boxes added to SLAC batch farm in GLAST's name; 20 being installed → 160 effective CPUs - aiming for 400 by launch
  - Have leveraged these first 80 into routine use of 400 CPUs
  - SCS wants us to use this model
    - Contribute boxes to the farm – they will 'guarantee' turnaround as if we had dedicated machines
    - The more we contribute the more assured we'll be of on-demand service
  - The SLAC farm is available to all GLAST collaborators!

- Asking France and Italy about farm capacity there for collaboration science support: minimum 100 CPUs, 50 TB disk

Pipeline Throughput

Pipeline Statistics

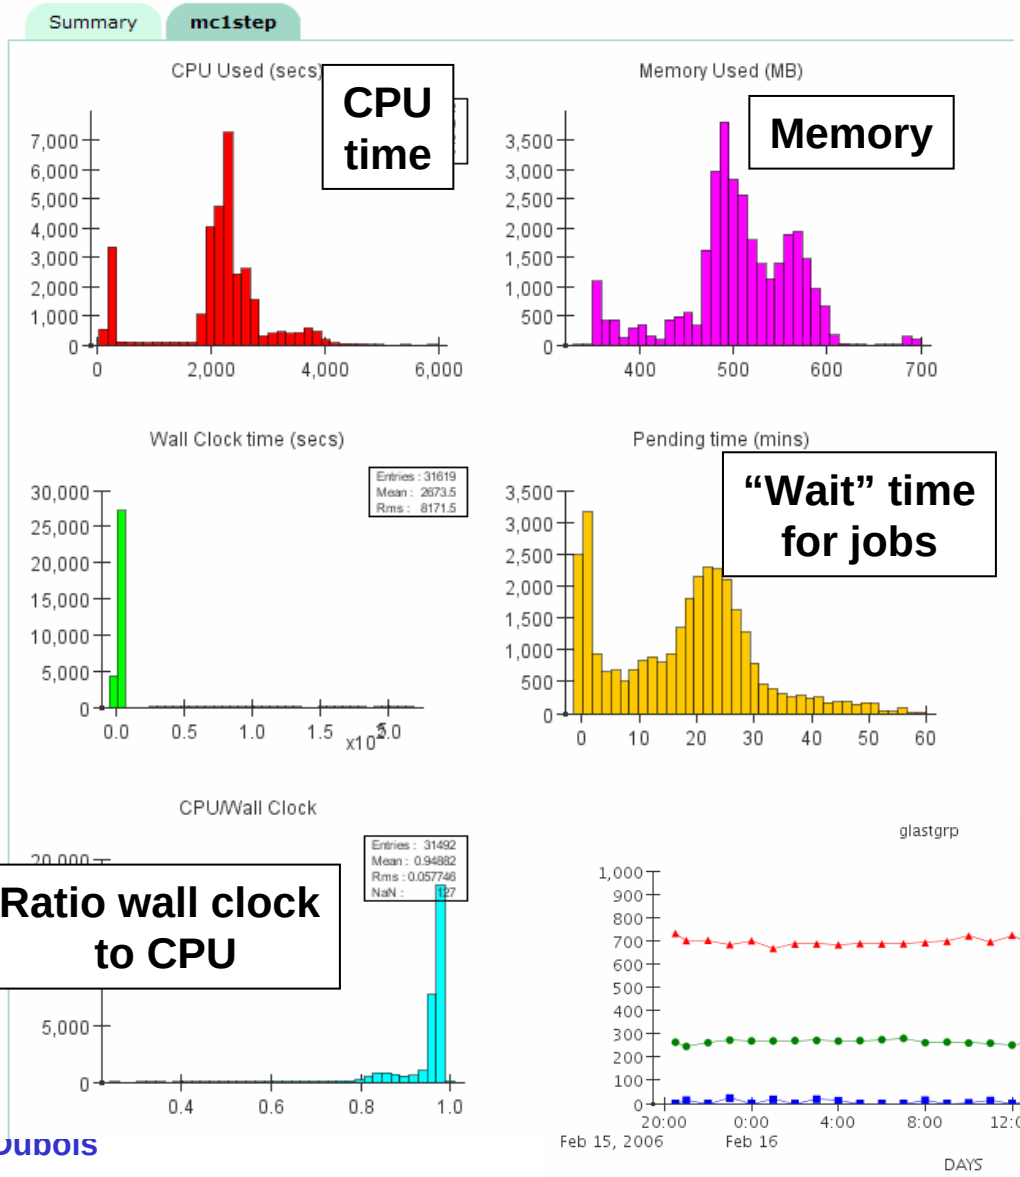

Resources Used

~28 CPU years in pipeline so far overall

~200,000 batch jobs for DC2 + ~5TB disk

~150k batch jobs for PSR + similar amount of disk

1 yr orbit could be 1.5M jobs + 30-50 TB

Better not to depend on SCS' largess

Future Storage Options: Orbit Data

- 10% Solution:
  - 60 TB/yr @ ~$240k for disk; ~$50k tape (or $120k if all recon kept on tape)
  - Assumes we can safely cut 90% of the background - we don't have the cuts derterined yet
- Full disk in '07; 10% in '08 and onwards (add 25 TB contingency to 150)
  - ~175 TB overall @ $700k for disk; $140k for tape
  - Free up 110 TB in '08 which covers '08 and ½ of '09

- No recon on disk
  - ~30-40 TB/yr @ 120-160k$ disk; 25-35 k$ tape
  - suffer latency in retrieval when needed
    - 1 TB retrieved in 1 day currently
    - Affects reprocessing merit; event displays; calibrations
    - Is there some need we'll find for fast turnaround that would make us regret not having things right there? Don't know yet.
- Wrinkles
  - Price of disk will continue to drop. How soon and how much? 20% last year
  - SCS has told us not to worry about tape costs (yet)

SAS Manpower Needs

- Already working closely with the collaboration – 20 of our ~30 FTEs come from outside SLAC!
- Could not be done without this group-wide effort
- Quite a bit of ebb and flow in who does what
- Contributions outside SAS/SLAC:
  - Italy:
    - Core tools, G4 support – 2,5 FTE
    - TKR alg development – 1 FTE
    - SciTools development – 2,5 FTE
  - France:
    - Core tools: 0,5 FTE
    - CAL sim/recon/calibs – ~5 FTE
    - Source Catalogue – 3 FTE
  - US:
    - Core tools – 3 FTE (UW, Goddard)
    - CAL sim/calibs – 2 FTE (NRL)
    - ACD – 1 FTE (Goddard, SVAC)
    - TKR – 1 FTE (UCSC)
    - Sundry contributions
- Responsibilities outside SAS/SLAC:
  - Instrument code release manager - UW
  - SciTools release manager - GSSC
  - Builds manager – Goddard
  - System tests manager – Goddard
  - Documentation - Goddard
  - Code build tool (CMT) – LLR/France
  - Build tool gui & Event display – Udine/Italy
  - Source Catalogue – CEA/France
- Responsibilities inside SAS/SLAC:
  - Overall management
  - Data handling
  - Calibrations infrastructure
  - Overall Reconstruction; TKR
- Shortfall:
  - ~ 4 FTE for infrastructure work
  - Hardest to get assistance here

What to Take Away

- **Prudent approach is to:**
  - Acquire ~400 cores available to GLAST
    - Lesson learned from 5-ring circus of DC2, BT, EM
  - Keep full event details on disk in '08 : ~175 TB
    - Hard to do both these without IFC contribution of 50 TB in CY06, 80 TB in CY07
    - Would shortchange CPU if push comes to shove
  - GLAST will do better science the more compute power it has access to
    - Have not hit the plateau yet!
    - Extend Pipeline 2 to France and Italy

- **SAS manpower is barebones in core software**
  - Expect to get 1 new FTE at GSFC via Steve
  - Short 3 with no prospects of getting them from the rest of the collaboration