NATIONAL ENERGY RESEARCH
SCIENTIFIC COMPUTING CENTER

The Big Data Revolution in Astrophysics: A Case Study of the Palomar Transient Factory

Peter Nugent (LBNL/UCB)

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Future Concerns for HPC

- Power
- Data
- Programming
- All of them hit your wallet - $’s

Which Swim-Lane to Choose

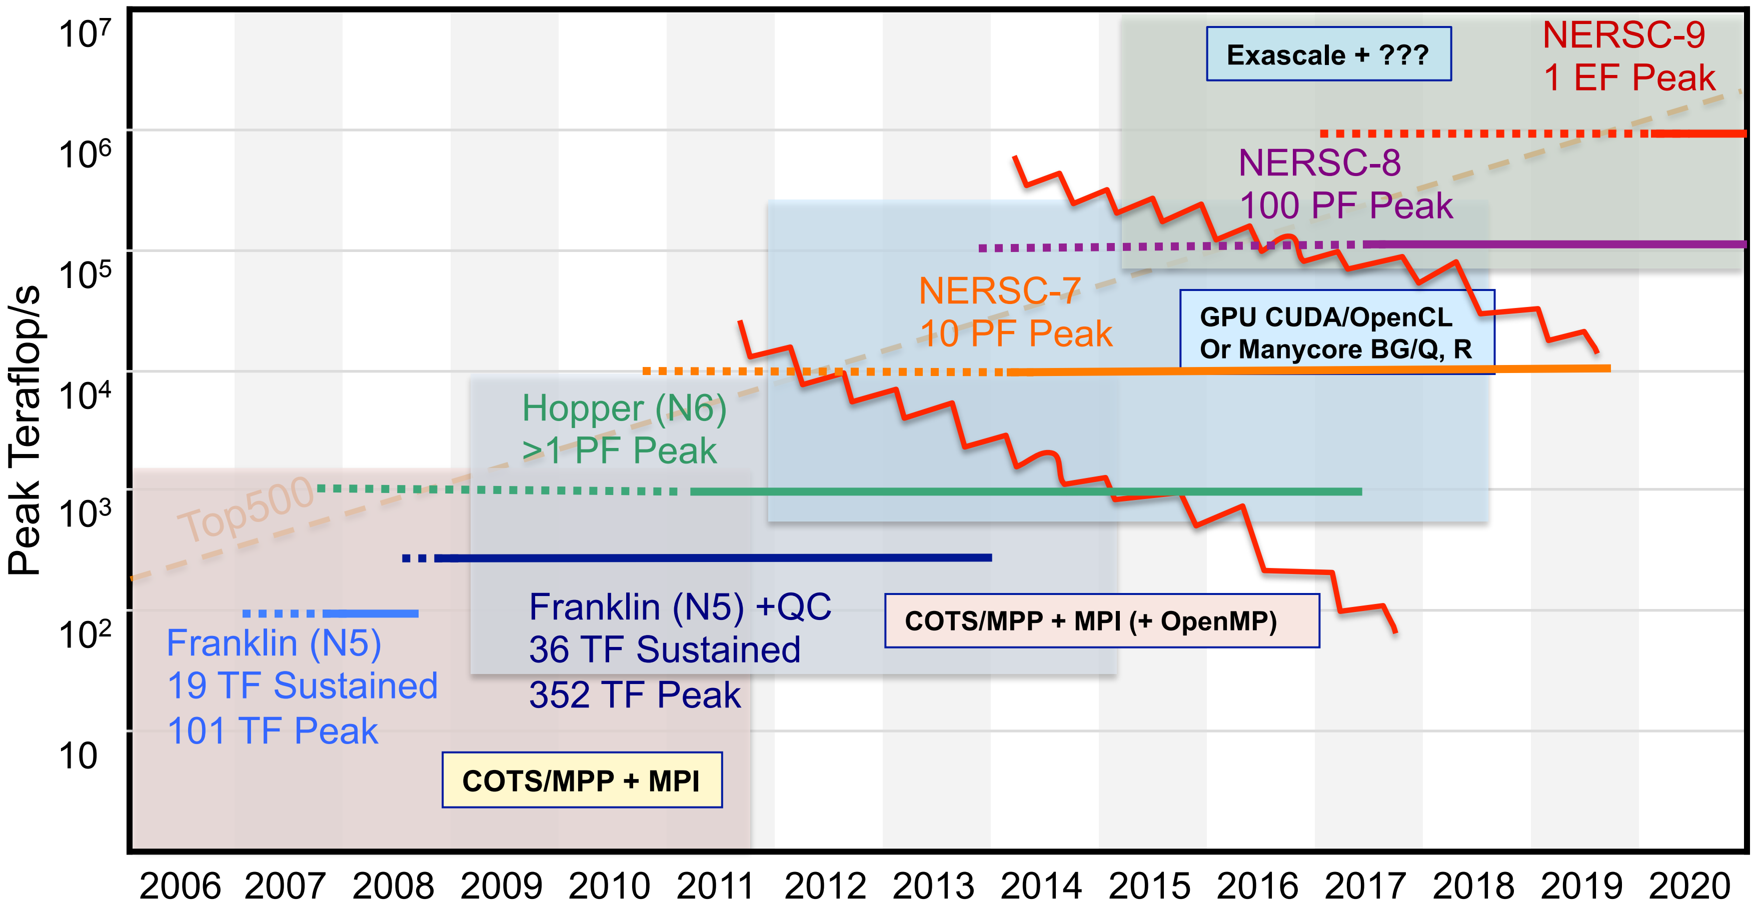

Key Questions for GPU Testbed

- What parts of the NERSC workload will benefit from GPU acceleration
- What portions of workload see no benefit
  - GPU costs as much as the CPU host
  - GPU consumes same power as host (~200-250W)
- Is GPU "swim lane" going to supplant the CPU swim lane, or just expand into "different" space
  - Are GPU's the future, or just a feature?
  - Should NERSC expand GPU-based systems to serve specific user needs

Dirac GPU Testbed Configuration

Hardware

- 44 nodes w/ 1 GPU per node
  - integrated into carver cluster
- Host Node
  - dual-socket Xeon E5530 (Nehalem)
  - 76.5GF Peak DP (153 GF SP)
  - QDR Infiniband
  - 24GB GB DDR-1066 memory
  - 51 GB/s peak mem BW
- GPU
  - Nvidia Tesla C2050 (Fermi)
  - 515GF peak DP (1030GF SP) : 6x more than host
  - 3 GB memory
  - 144 GB/s peak mem BW (3x)

Software

- CUDA 3.1
- PGI Compilers
- GPU Direct
  - OpenMPI
  - MVAPICH
- *Matlab Parallel Computing Toolbox coming soon*

A Likely Trajectory - Collision or Convergence?

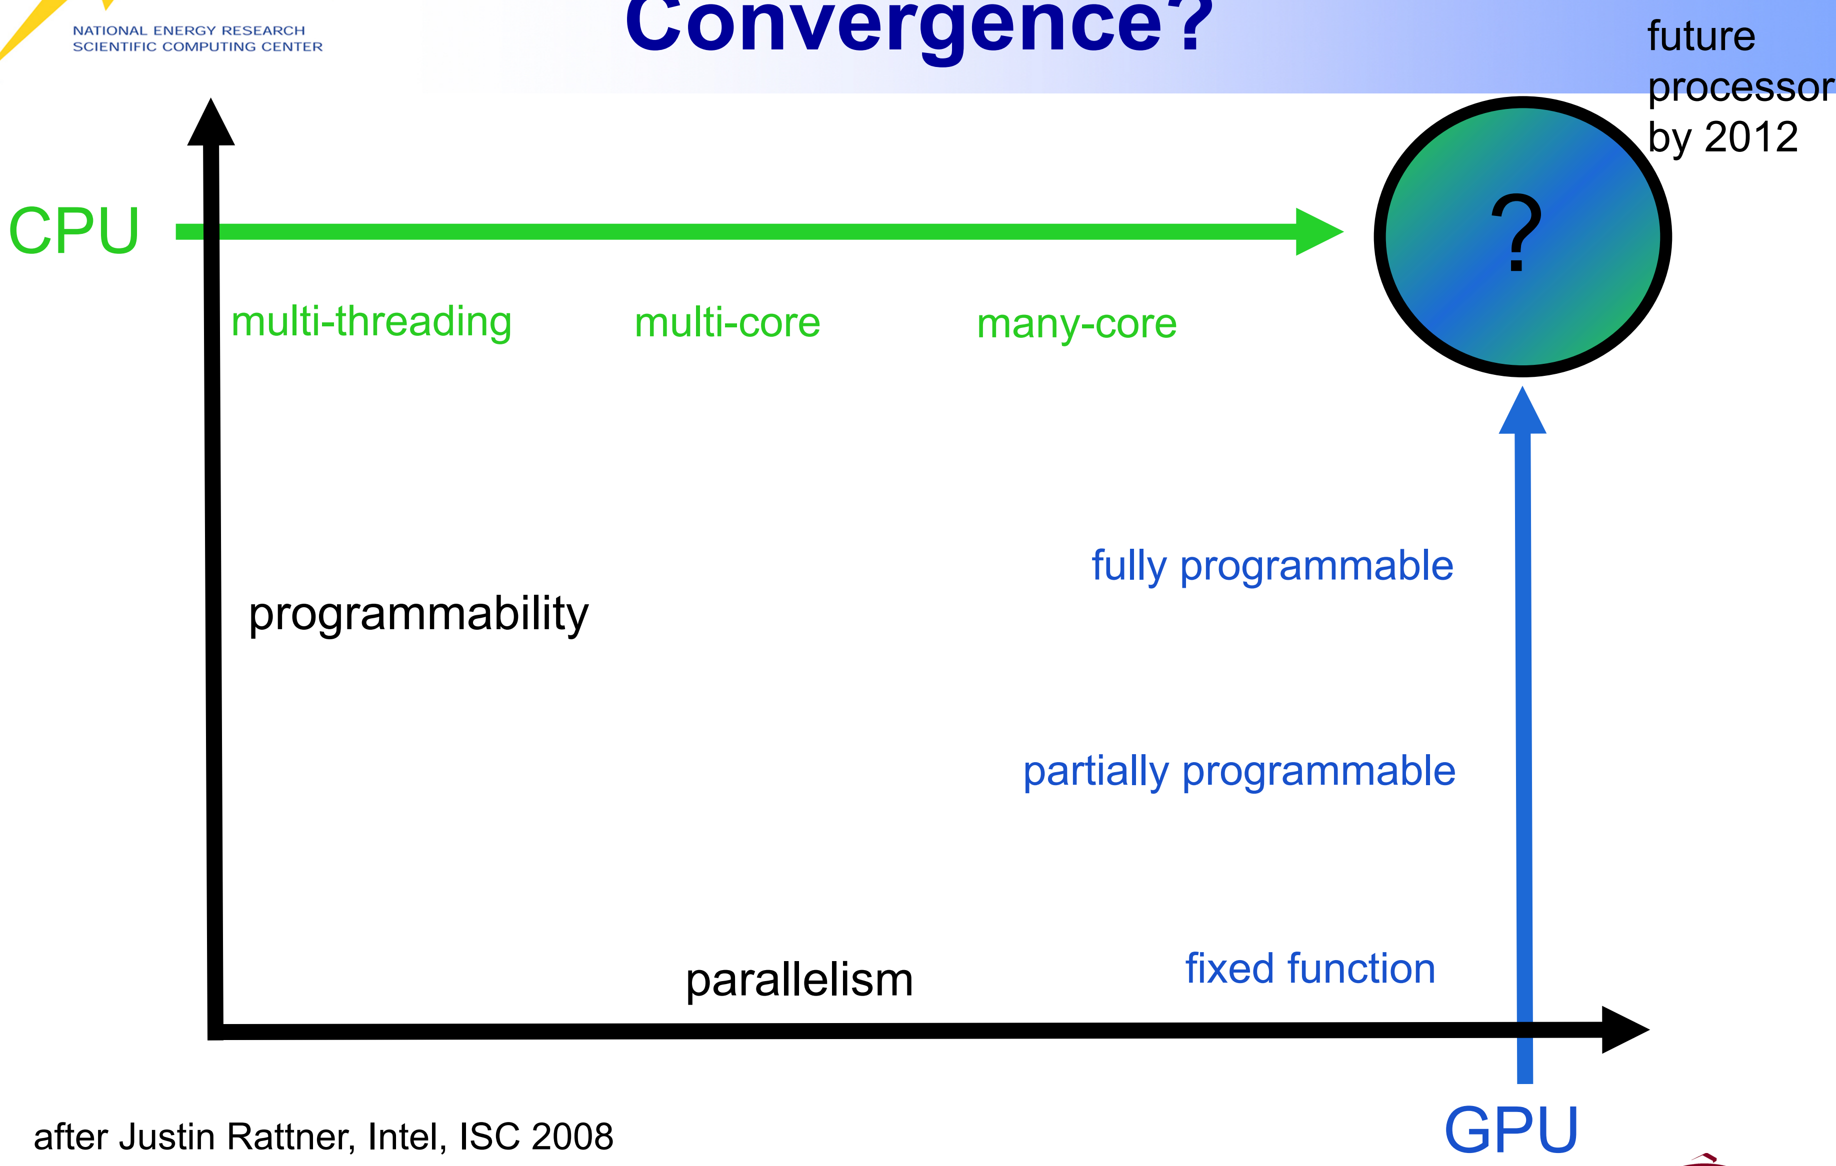

after Justin Rattner, Intel, ISC 2008

Cost of Data Movement

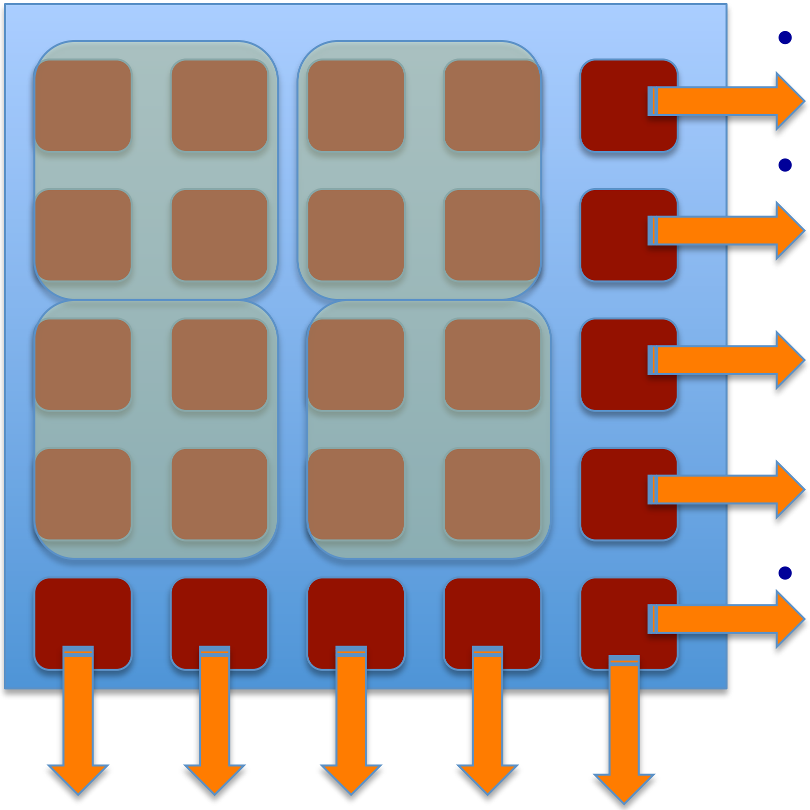

- ~1000-10k simple cores
  4-8 wide SIMD or VLIW bundles
- Cost of moving long-distances on chip motivates clustering on-chip
  - 1mm costs ~6pj (today & 2018)
  - 20mm costs ~120 pj (today & 2018)
  - FLOP costs ~100pj today
  - FLOP costs ~25pj in 2018
- Different Architectural Directions
  - GPU: WARPs of hardware threads clustered around shared register file
  - Limited area cache-coherence
  - Hardware multithreading clusters

Performance Summary

Domain	Algorithm	Performance Summary cf. 8 core Nehalem
QCD	Sparse Matrix Conjugate Gradient	~10x
Molecular Dynamics (HOOMD)	N-body	~6-7x
Lattice Boltzmann CFD	Lattice Boltzmann	~1.17x
Geophysical Modeling	quasi-minimum-residual (QMR) solver	~3.33x
QCD	Krylov space solvers to compute intensive matrix inversion	~3.0x (matrix multiply) ~2.5 multi-shifted bi-conjugate gradient algorithm
Astrophysics	AMR	~5x

- **Comparisons of optimized CPU version to optimized GPU version are rare**
  - **3x memory and 6x flop rate advantage for GPU**
  - **Speedups in excess of 50x-100x usually indicate methodological errors**
- **No tests performed on analysis of large data sets.**

Early User Observations

- "Domain Scientist input is essential. Debugging can be hard"
- "Programming on the GPU can be a bit challenging, as many of the resources to learn from are geared towards programmers rather than scientists"
- "For advanced programmer GPU and CUDA is good team"
- "Most CUDA can be learned from downloading example code. First get it working and then optimizations like cache, memory but application specific. NVIDIA doc is very dense. Mixing MPI and CUDA is straightforward but challenge is in using all cores on node efficiently"

…Not Everyone Was Happy…

- "disappointing – the 200GF for DGEMM was not worth the trade off of re-writing code in CUDA"
- "CUDA is inadequate for most scientific computation since i) most algorithms are not trivial data-parallel ii) NVIDIA documentation on performance tuning is not always correct"

Results: Astrophysics

- **Astrophysics: H. Y. Schive**
  - **Code:** *GPU-accelerated Adaptive MEsh Refinment (GAMER)*
  - **Method***: 2^{nd} order PDEs using adaptive meshes*
  - **Performance:** *1GPU 5x faster than 8 CPUs*
  - **Notes:** *Would be faster, but AMR datastructures must be manipulated on CPU*

- **Astrophysics: Jose Fiestas**
  - **Code:** *PhiGPU*
  - **Method:** *N-body application with 4^{th}-8^{th} order Hermite integrators*
  - **Performance:** *performance improves 60% for C1060 vs. C2050 (don't know speedup vs. CPU)*
  - **Notes:** *mixed precision, running on 32 nodes*

“Current” Optical Surveys

Photometric:

Palomar Transient Factory
La Silla Supernova Search
SkyMapper
PanSTARRS

Spectroscopic:
SDSS III

All of these surveys span astrophysics from planets to cosmology, from the static to the transient universe.

Competition

The competition were two wide-field multi-color surveys with cadences that were either unpredictable (SkyMapper) or from days to weeks (PanSTARRS) in a given filter.

How could we do something better/different?

- Start quickly - P48" coupled with the CFHT12k camera
- Don't do multiple colors
- Explore the temporal domains in unique ways
- Take full advantage of the big-iron at Super-Computing Centers
- Get all the science we possibly can out of this program

Thus we need the capability of providing immediate follow-up of ***unique*** transients, using 4 to 10-m class telescopes.

Phase-Space

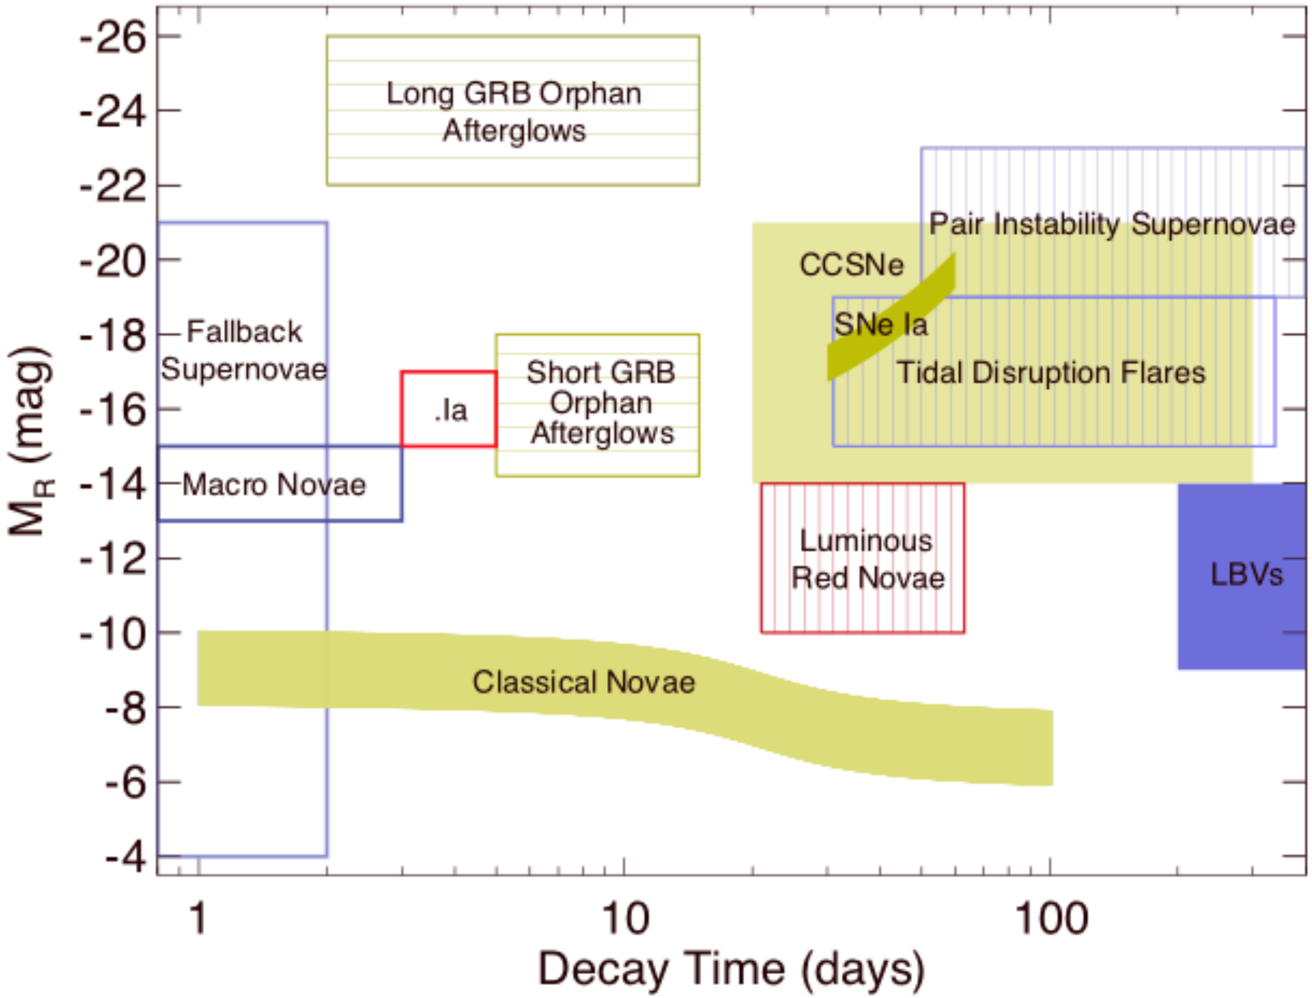

Long GRB Orphan Afterglows
Pair Instability Supernovae
CCSNe
SNe Ia
Tidal Disruption Flares
Fallback Supernovae
.Ia
Short GRB Orphan Afterglows
Macro Novae
Luminous Red Novae
LBVs
Classical Novae
M_R (mag)
-26
-24
-22
-20
-18
-16
-14
-12
-10
-8
-6
-4
1
10
100
Decay Time (days)

PTF (2009-2013)

- CFH12k camera on the Palomar Oschin Schmidt telescope
  - 7.8 sq deg field of view, 1” pixels
  - 60s exposures with 15-20s readout in r, g and H-alpha
  - First light Nov. 24, 2008.
  - First useful science images on Jan 13th, 2009.

- 2 Cadences (Mar. - Nov.)
  - Nightly (35% of time) on nearby galaxies and clusters (g/r)
  - Every 3 nights (65% of time) on mostly SDSS fields with minimum coverage of 2500 sq deg. (r) to 20th mag 10-sigma
  - H-alpha during bright time (full +/-2 days)

Nov-Feb, minute cadences on select fields.

PTF Science

PTF Key Projects	
Various SNe	Dwarf novae
Transients in nearby galaxies	Core collapse SNe
RR Lyrae	Solar system objects
CVs	AGN
AM CVn	Blazars
Galactic dynamics	LIGO & Neutrino transients
Flare stars	Hostless transients
Nearby star kinematics	Orphan GRB afterglows
Rotation in clusters	Eclipsing stars and planets
Tidal events	H-alpha sky-survey

The power of PTF resides in its diverse science goals and follow-up.

PTF Science

- Detected transients will be followed up using a wide variety of optical and IR, photometric and spectroscopic followup facilities.

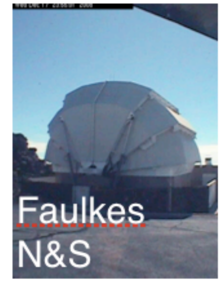

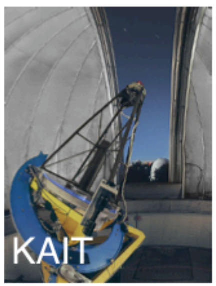

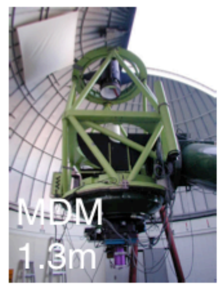

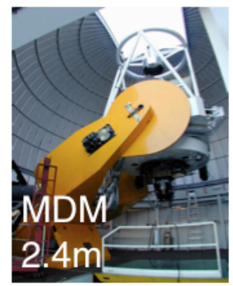

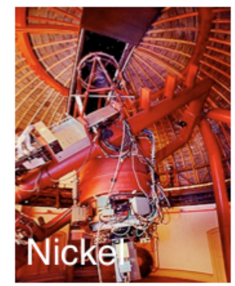

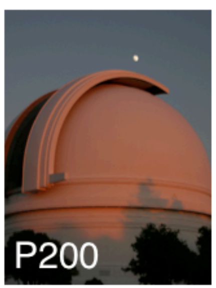

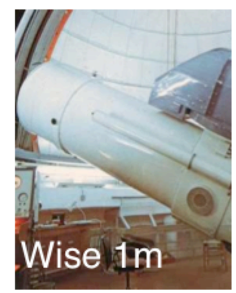

The power of PTF resides in its diverse science goals and follow-up.

PTF Pipeline

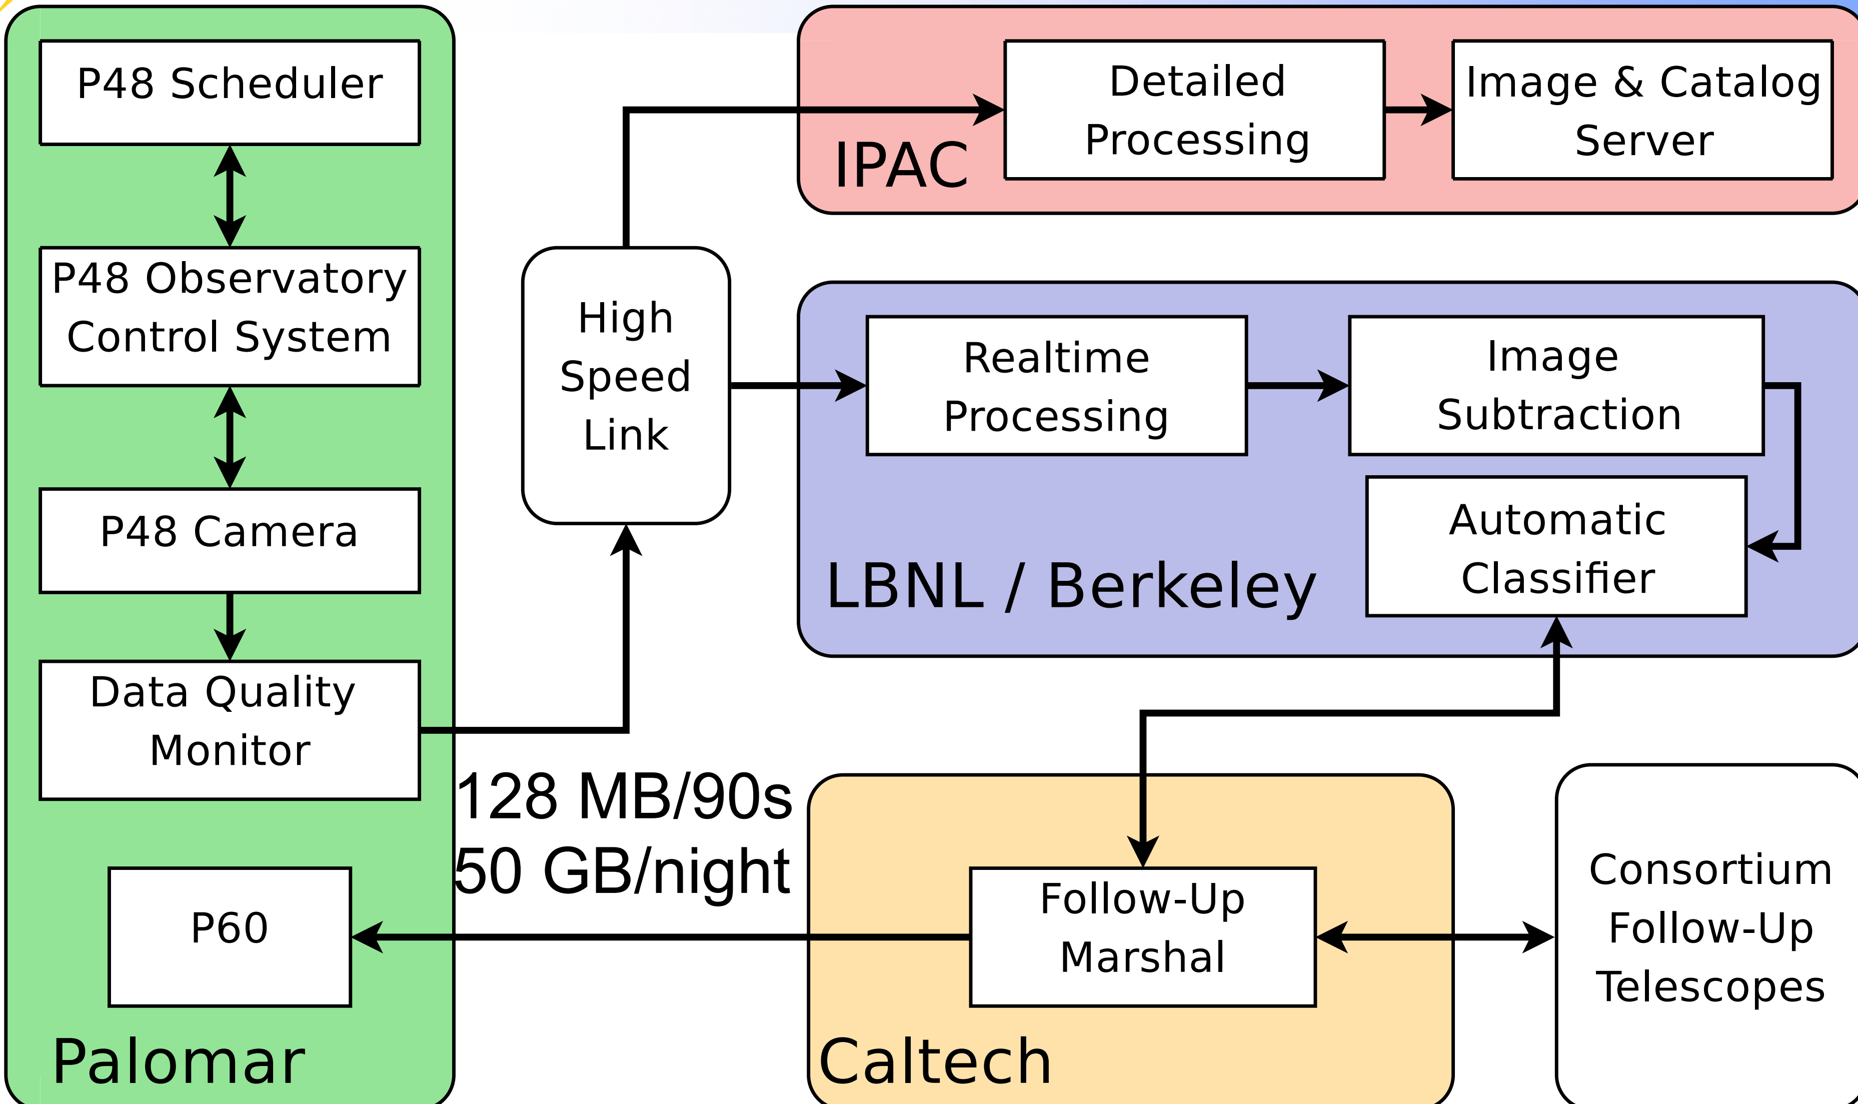

Pipeline

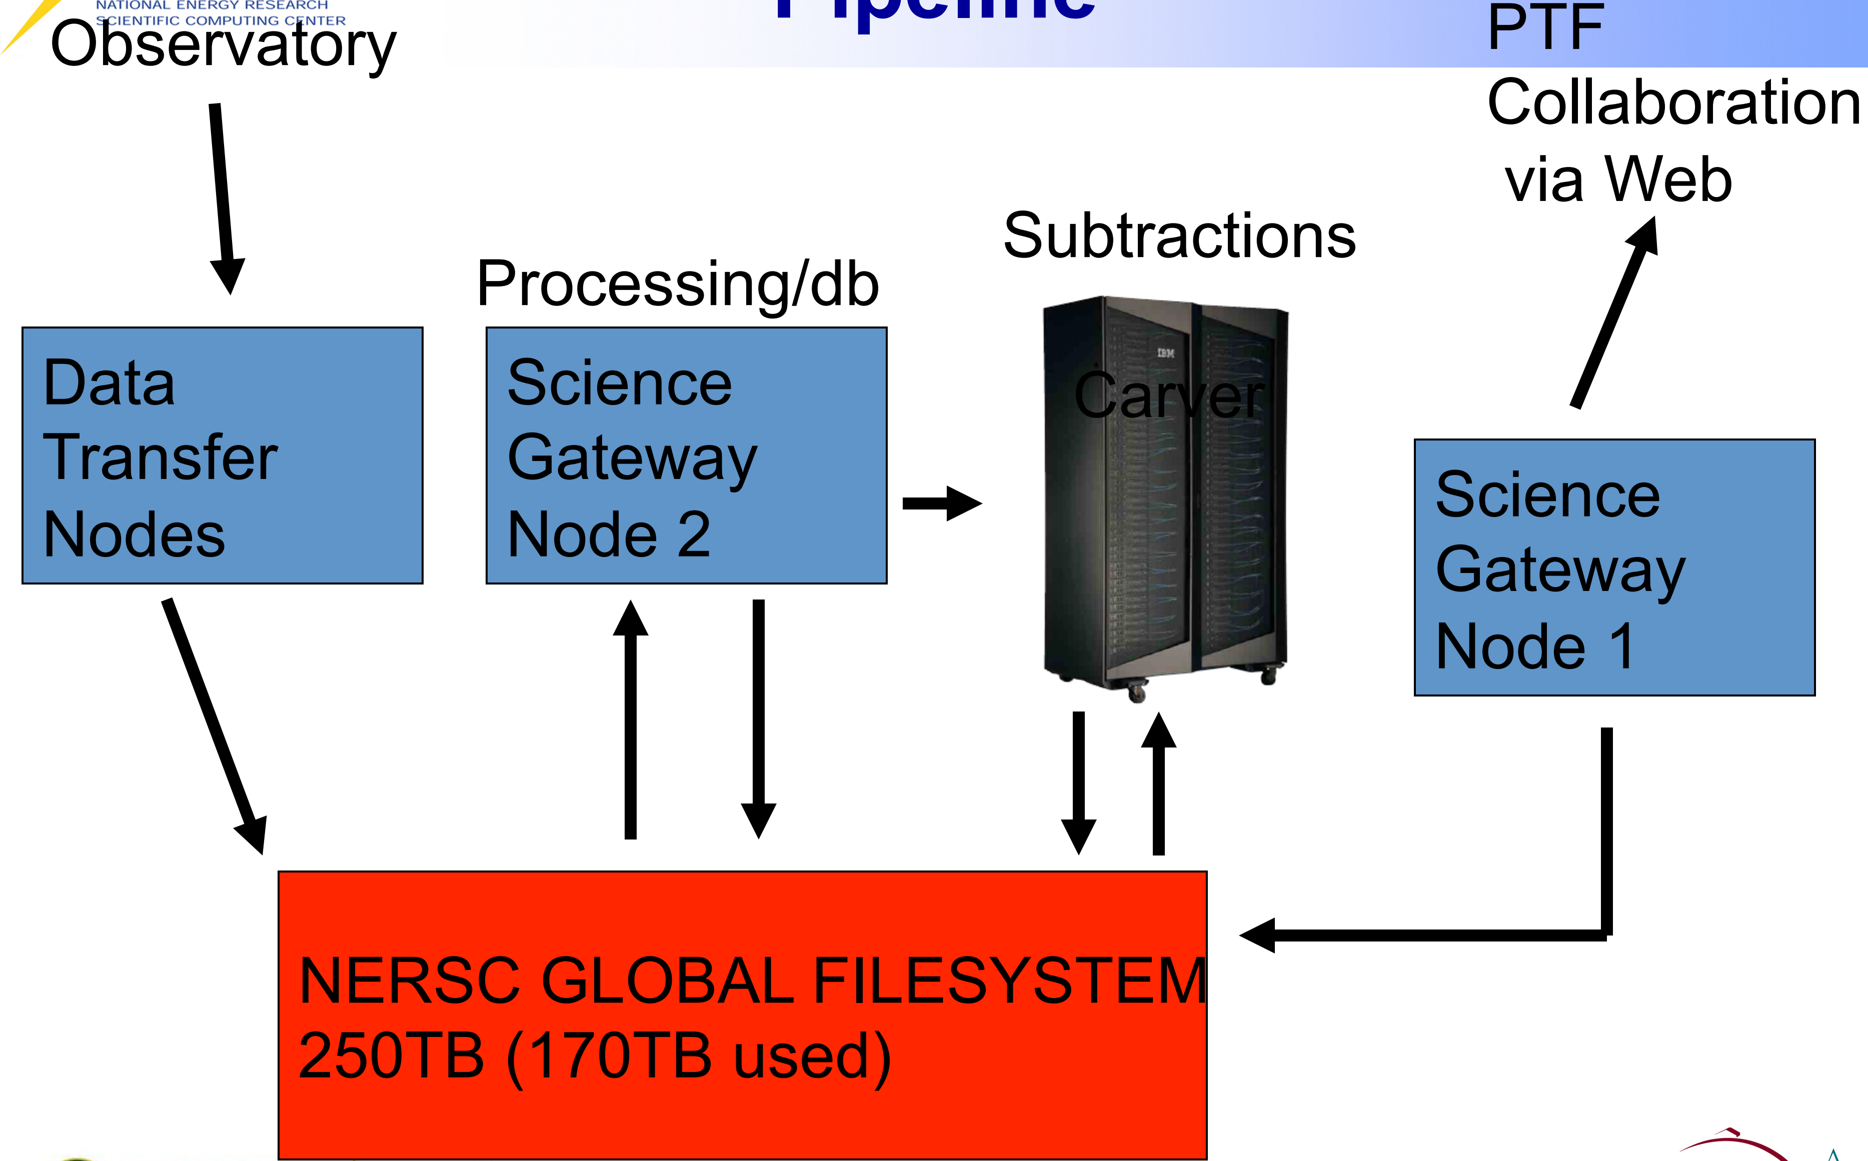
Observatory
PTF Collaboration via Web
Subtractions
Processing/db
Data Transfer Nodes
Science Gateway Node 2
Carver
IBM
Science Gateway Node 1
NERSC GLOBAL FILESYSTEM
250TB (170TB used)

PTF Database

- 1.66M images
- 29k references
- 1.18M subtractions
- 765M candidates
- 40k saved transients

(many more unsaved)

All in just ~700 nights.

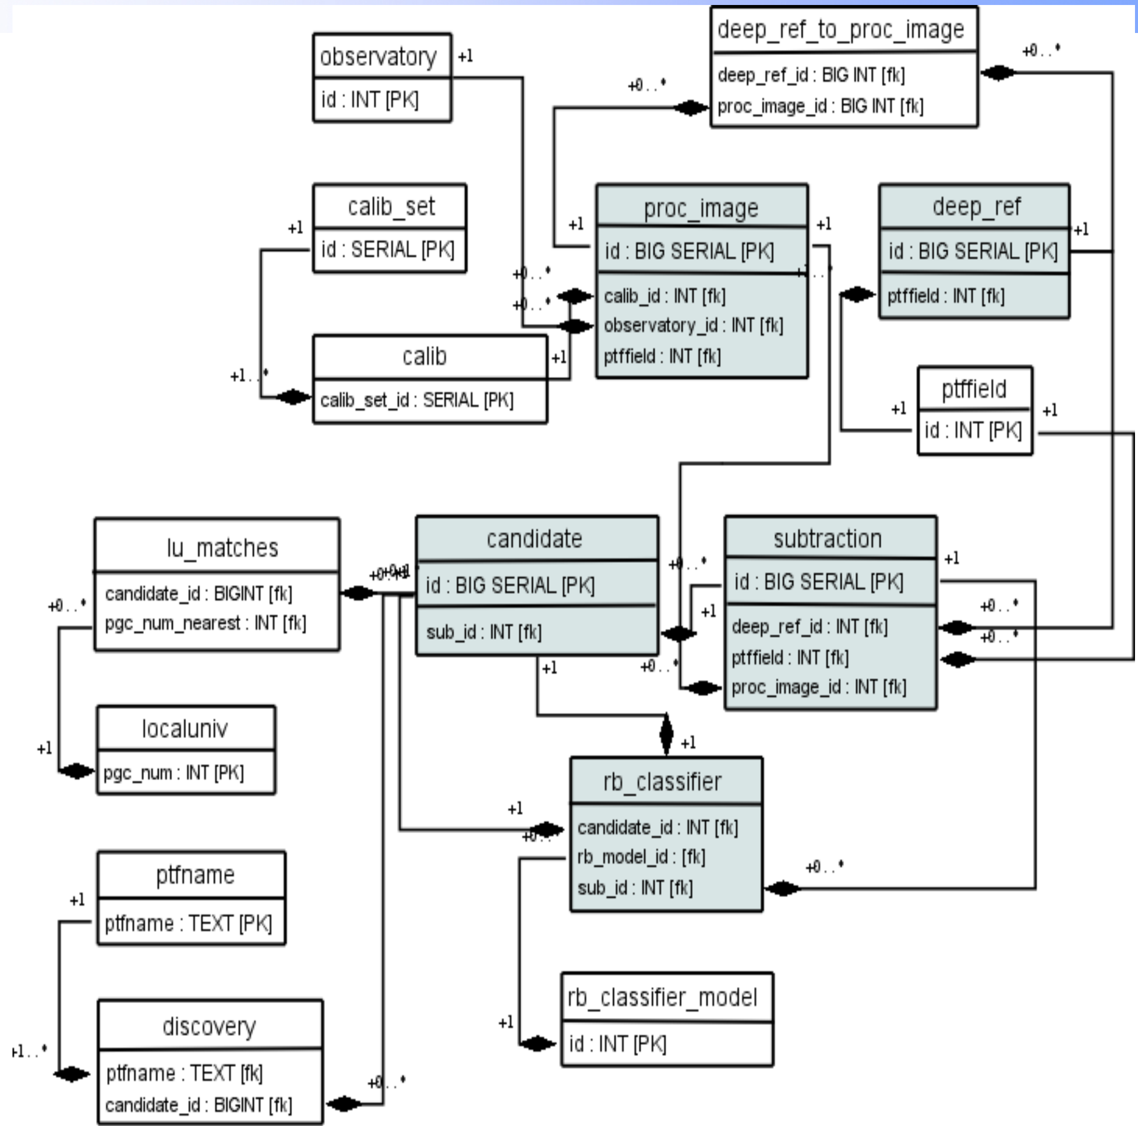

PTF Sky Coverage

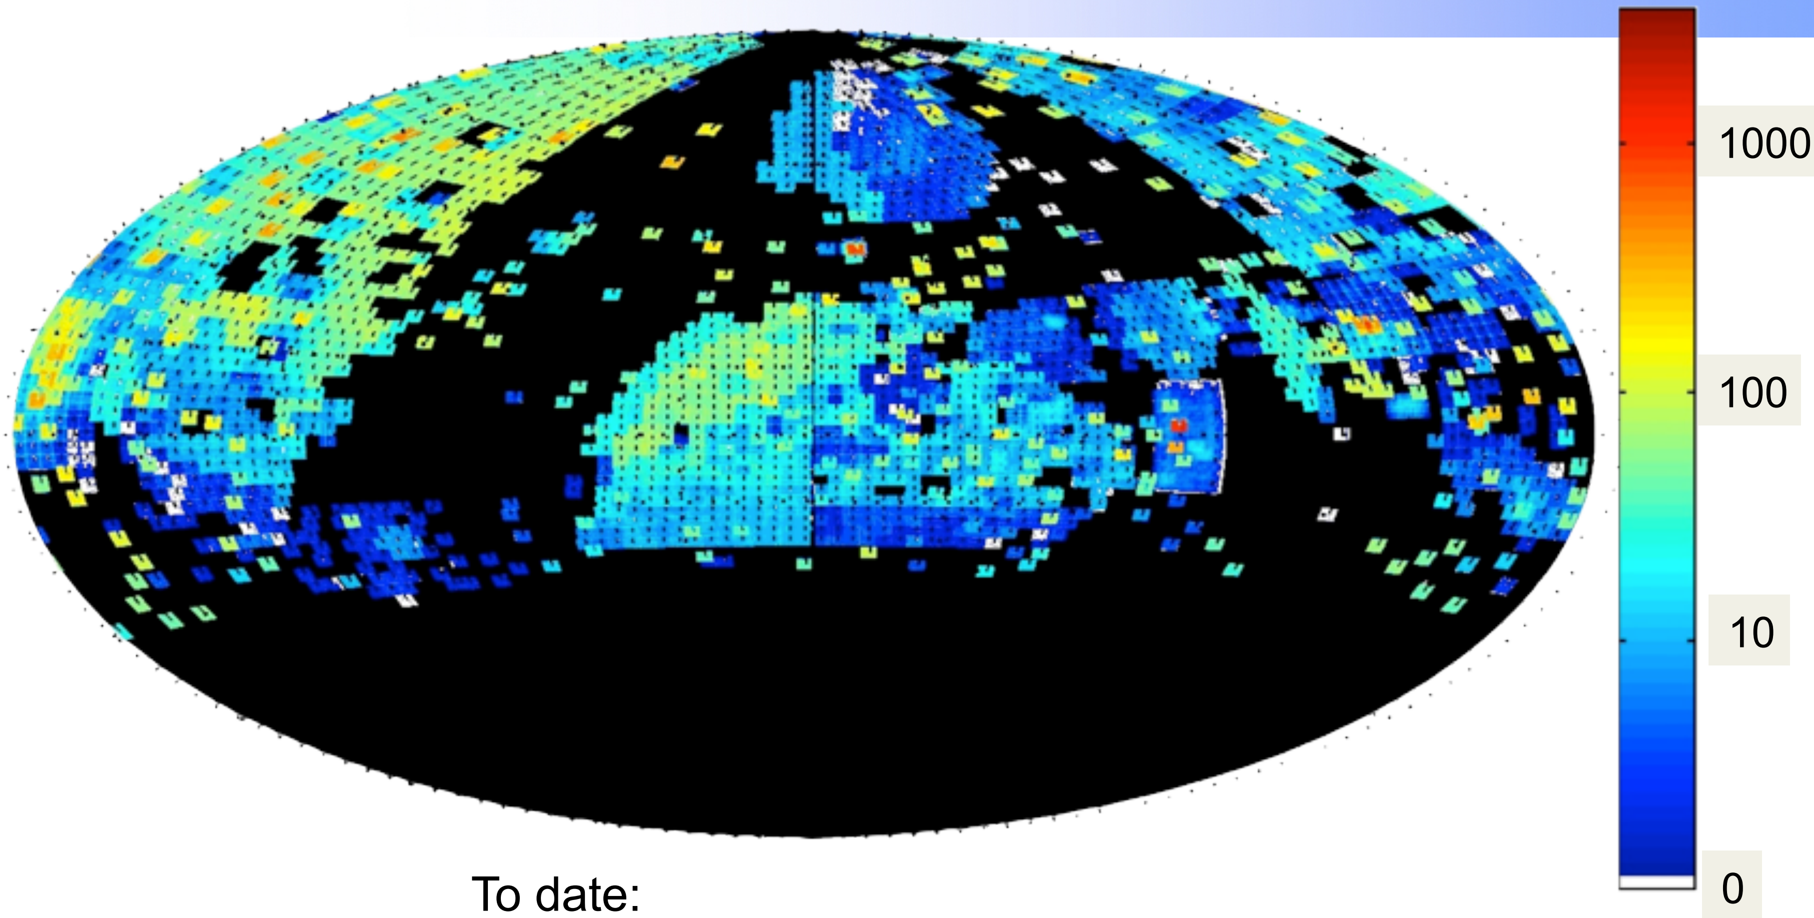

To date:

- 1500 Spectroscopically typed supernovae
- 10^5 Galactic Transients
- 10^4 Transients in M31
- 22nd/23rd/24th magnitude total depth (blue/green/orange)

PTF: Real or Bogus

PTF produces 1 million candidates during a typical night:

- Most of these are not real
  - Image Artifacts
  - Misalignment of images due to poor sky conditions
  - Image saturation from bright stars

- 50k are asteroids
- 1-2k are variable stars
- 100 supernovae
- 3-4 ***new, young*** supernovae or other explosions

Real or Bogus

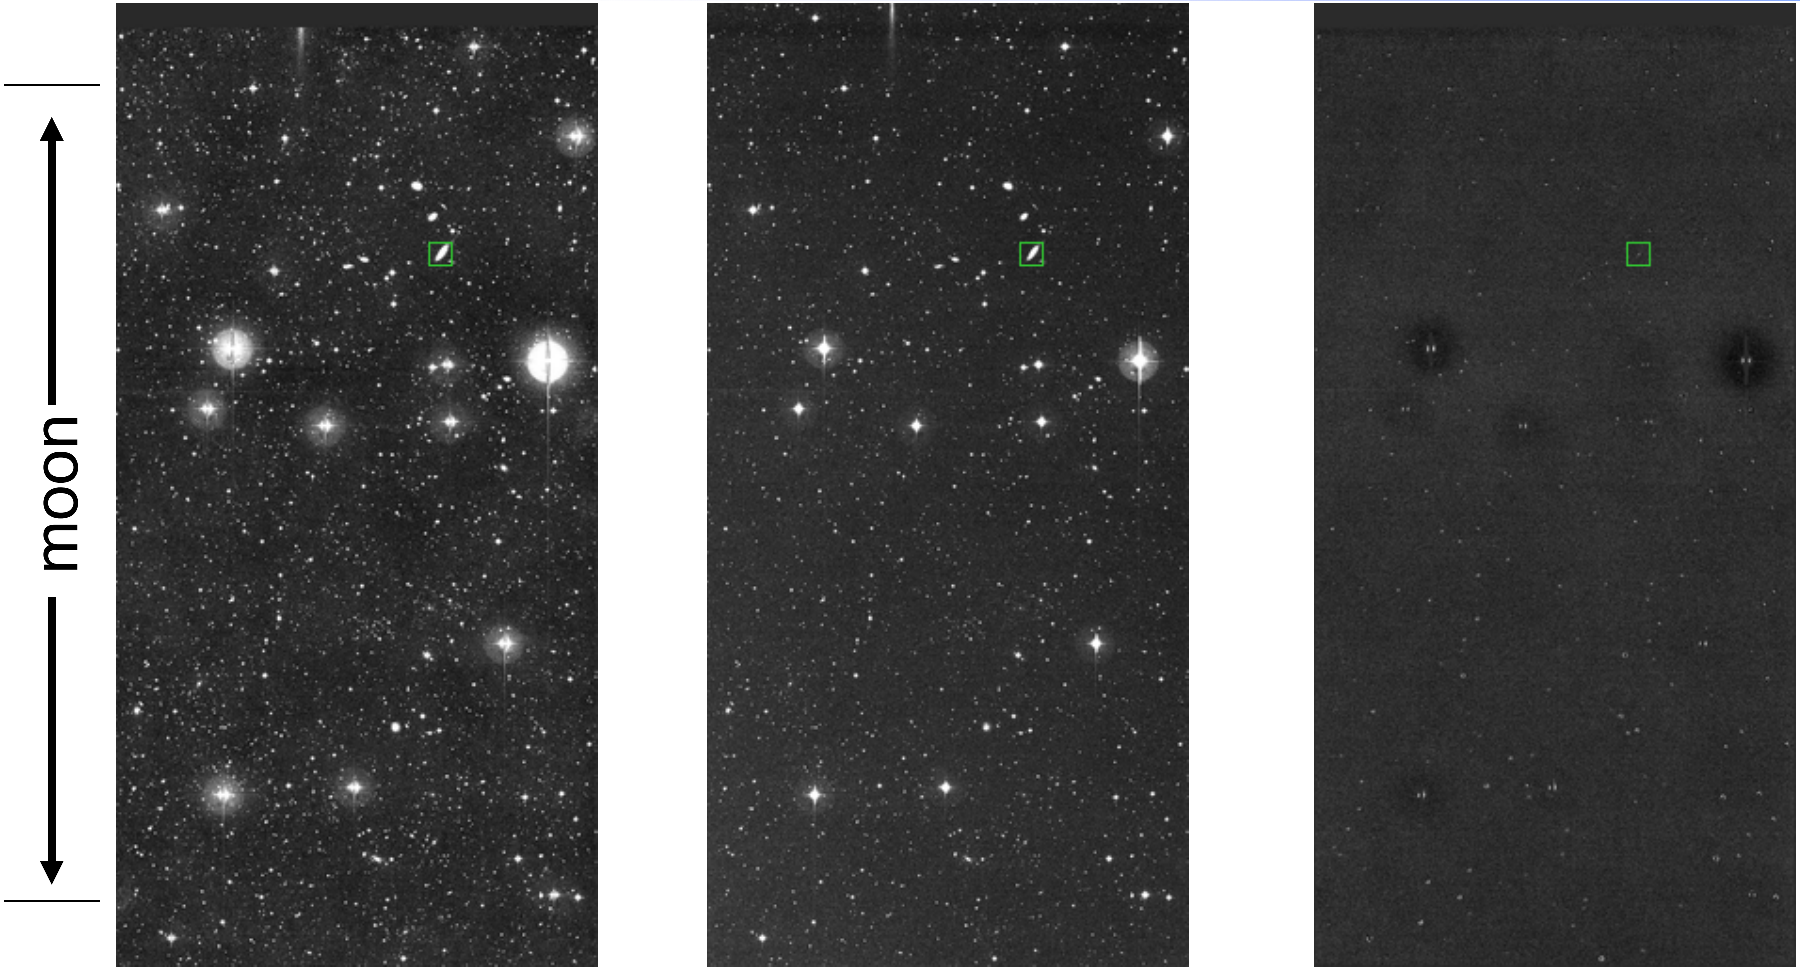

4096 X 2048 CCD images - over 3000 per night

Real or Bogus

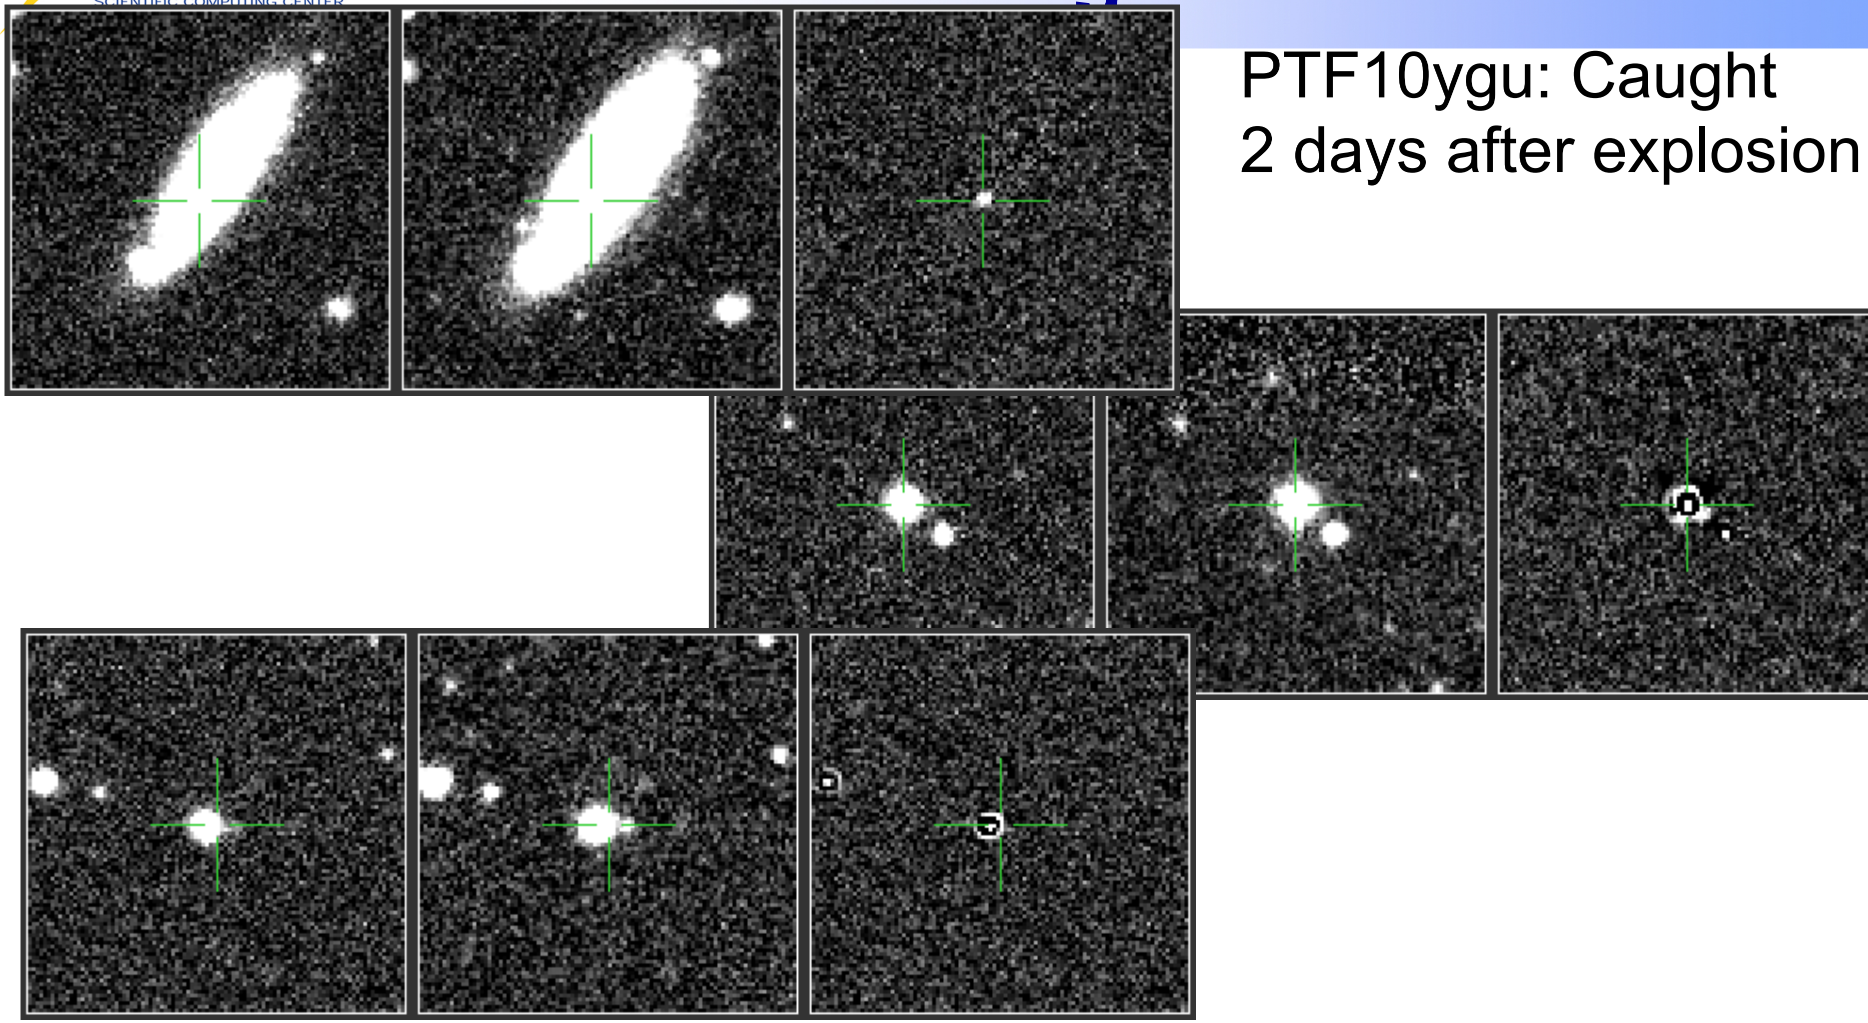

PTF10ygu: Caught
2 days after explosion

230 bogus candidates, 2 variable stars, 4 asteroids and the ~~youngest~~ Type Ia supernovae observed to date.

SN in M101

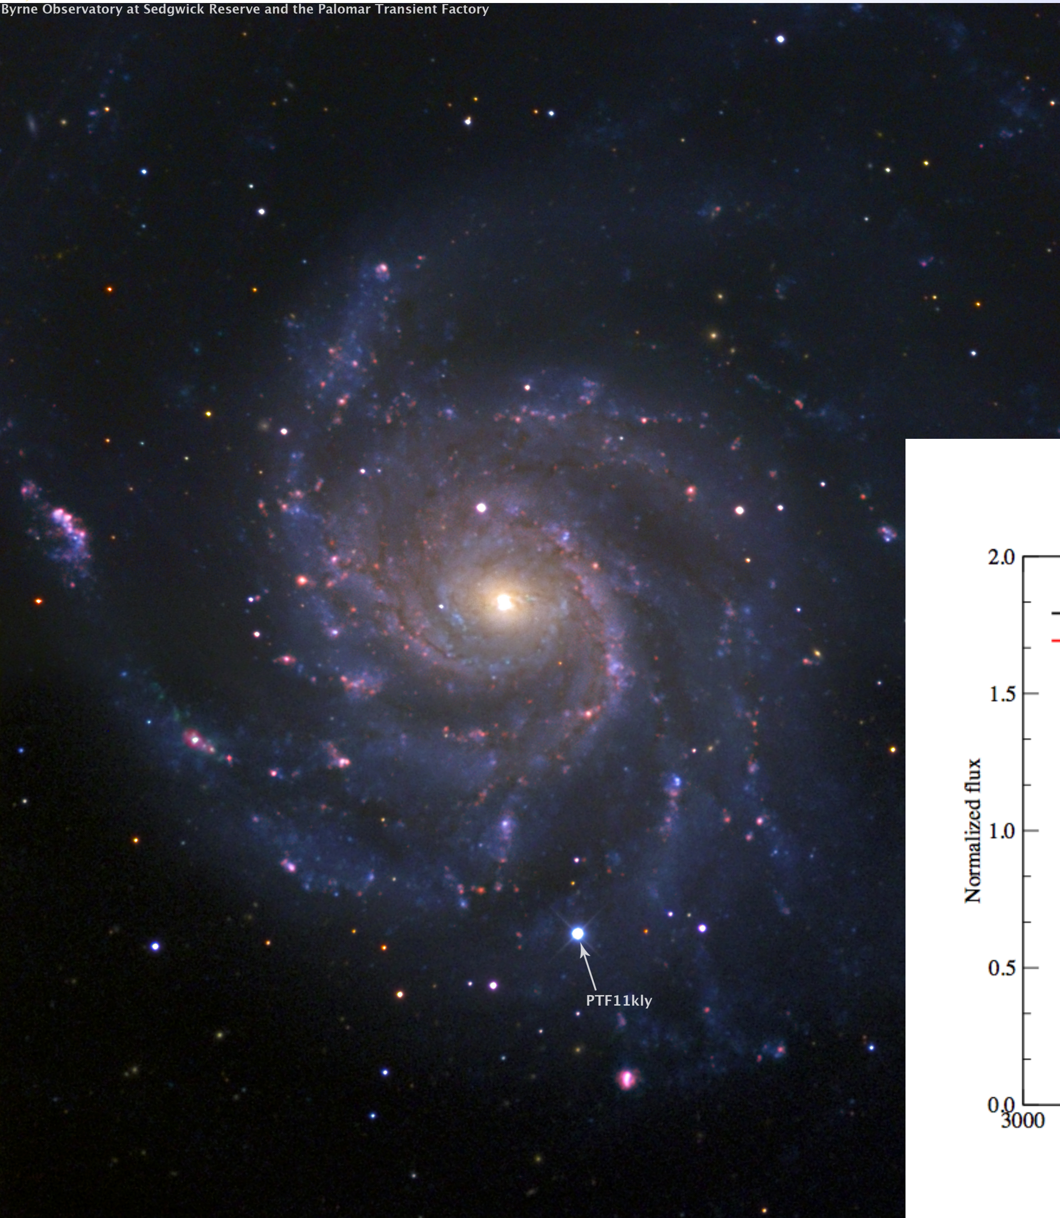

PTF11kly: Caught 11 hours after explosion

Quick query – what's the best candidate from the previous night…

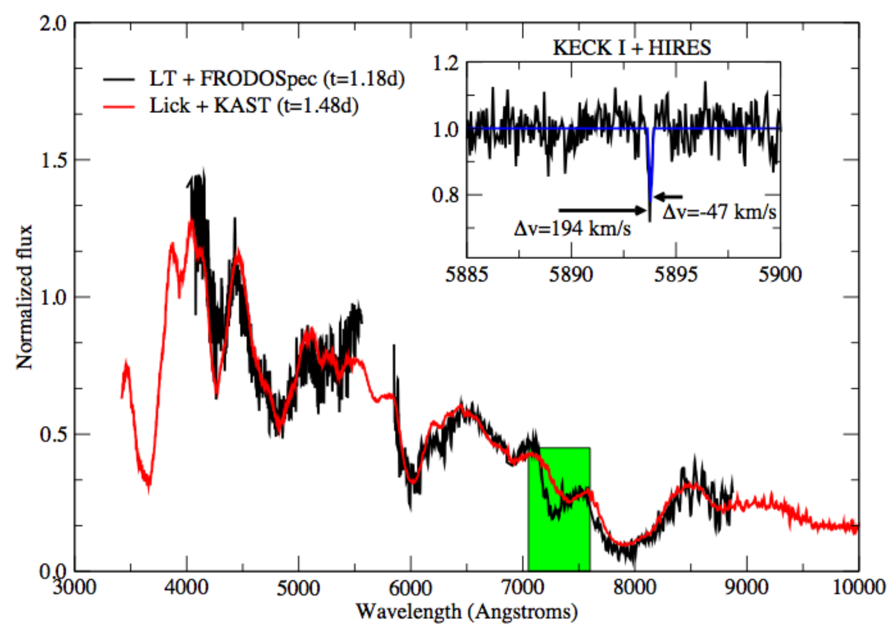

Users…

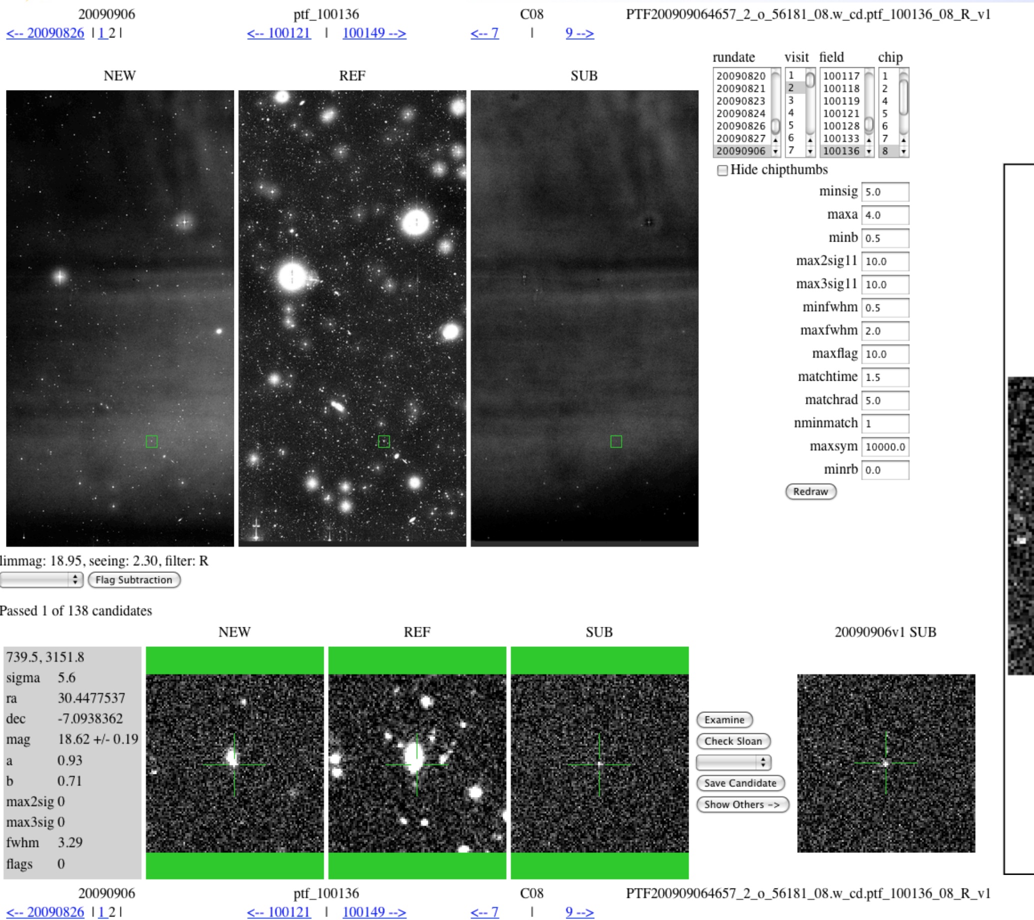

20090906 ptf_100136 C08 PTF200909064657_2_o_56181_08.w_cd.ptf_100136_08_R_v1
<-- 20090826 | 1 2 | <-- 100121 | 100149 --> <-- 7 | 9 -->
NEW REF SUB
rundate visit field chip
Hide chipthumbs
minsig 5.0
maxa 4.0
minb 0.5
max2sig11 10.0
max3sig11 10.0
minfwhm 0.5
maxfwhm 2.0
maxflag 10.0
matchtime 1.5
matchrad 5.0
nminmatch 1
maxsym 10000.0
minrb 0.0
Redraw
limmag: 18.95, seeing: 2.30, filter: R
Flag Subtraction
Passed 1 of 138 candidates
NEW REF SUB 20090906v1 SUB
739.5, 3151.8
sigma 5.6
ra 30.4477537
dec -7.0938362
mag 18.62 +/- 0.19
a 0.93
b 0.71
max2sig 0
max3sig 0
fwhm 3.29
flags 0
Examine
Check Sloan
Save Candidate
Show Others -->
20090906 ptf_100136 C08 PTF200909064657_2_o_56181_08.w_cd.ptf_100136_08_R_v1
<-- 20090826 | 1 2 | <-- 100121 | 100149 --> <-- 7 | 9 -->

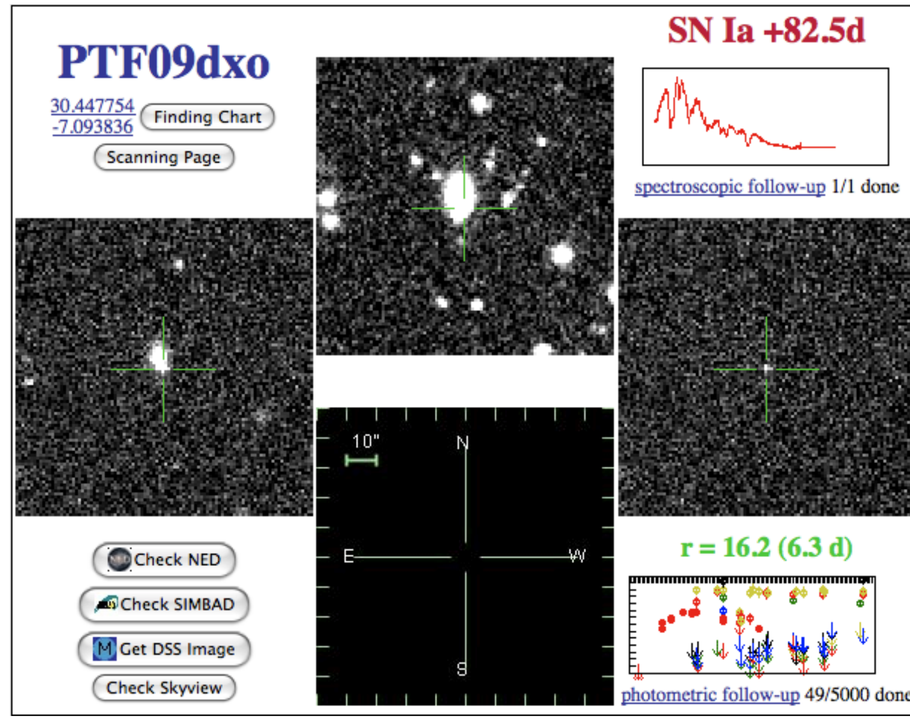

PTF09dxo
30.447754
-7.093836
Finding Chart
Scanning Page
SN Ia +82.5d
spectroscopic follow-up 1/1 done
10"
N
E
W
S
r = 16.2 (6.3 d)
Check NED
Check SIMBAD
Get DSS Image
Check Skyview
photometric follow-up 49/5000 done

Citizen Scientists…

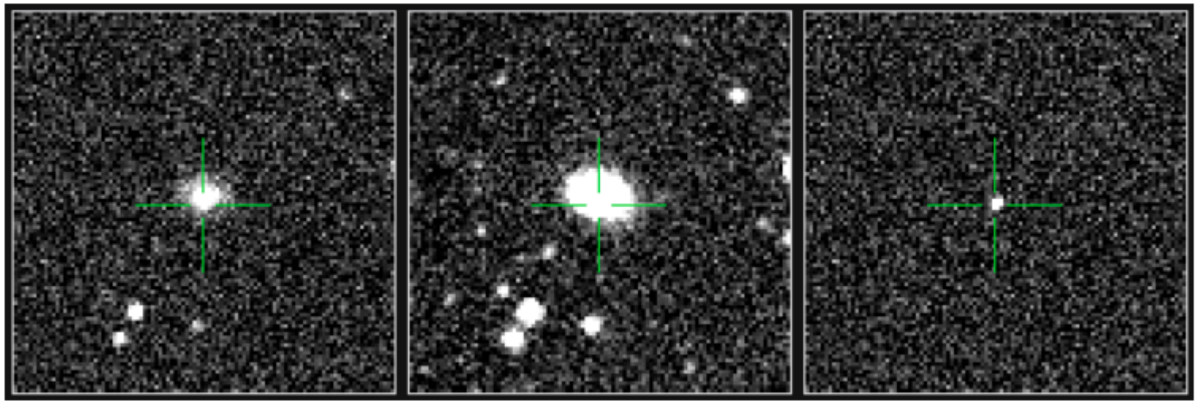

http:// supernova.galaxyzoo.org is now up and running!

A beta version appeared last year to support the SN Ia program in PTF and a WHT spectroscopy run. I spent a week with the folks at Oxford setting up the db and giving them training sets of good and bad candidates. They did the rest… 1200 members of galaxy zoo screened all the candidates between Aug 1 and Aug 12 in 3 hrs. The top 50 hits were all SNe/variable stars and they found 3 before we did. They scanned ~25,000 objects - 3 objects/min. They now do ~200 nightly and we have 15,000 users.

NERSC
NATIONAL ENERGY RESEARCH
SCIENTIFIC COMPUTING CENTER

Robot

Transient/VarStar Candidates

Name	ID	Viz	RB	ieg1	ieg2	irock	igal	best class	oarical class (origin)	discovery score	medscore	mag	mag_ref	number of matches	LBL ID matches
PTF10ghq	225381638 [jsb = 6608] Oarical...		0.330	0.254	1.684	-1.284	-1.491	circumnuclear event*	qso (simbad)	0.425	0.190	19.93	18.66	68	223563196 217748241 214828278 214352929 212441675 210402097 208825939 208620780 206781298 206405176 and 58 more...
PTF10hin	225447619 [jsb = 1714] Oarical...	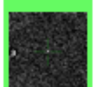	0.360	0.284	1.883	-1.436	-1.668	circumnuclear event*	qso (simbad)	0.430	0.393	19.79	18.39	18	225260857 205025445 204836371 196484063 189731647 183038011 173817454 168951242 162602090 162438316 and 8 more...
PTF10mwu	225440151 [jsb = 5340] Oarical...	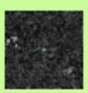	0.318	-1.246	-1.445	-1.654	1.520	varstar/galactic event*	varstar (sdss)	0.390	0.144	20.17	18.81	30	217613094 216468866 216294050 216241139 205491181 205235136 204075266 204009765 203830442 203711898 and 20 more...

A robot (built by Josh Bloom at UCB) queries the db every 20 min and compares new transients with archival information to ascertain its likely nature and publishes them to the collaboration - ***classification***.

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Robot

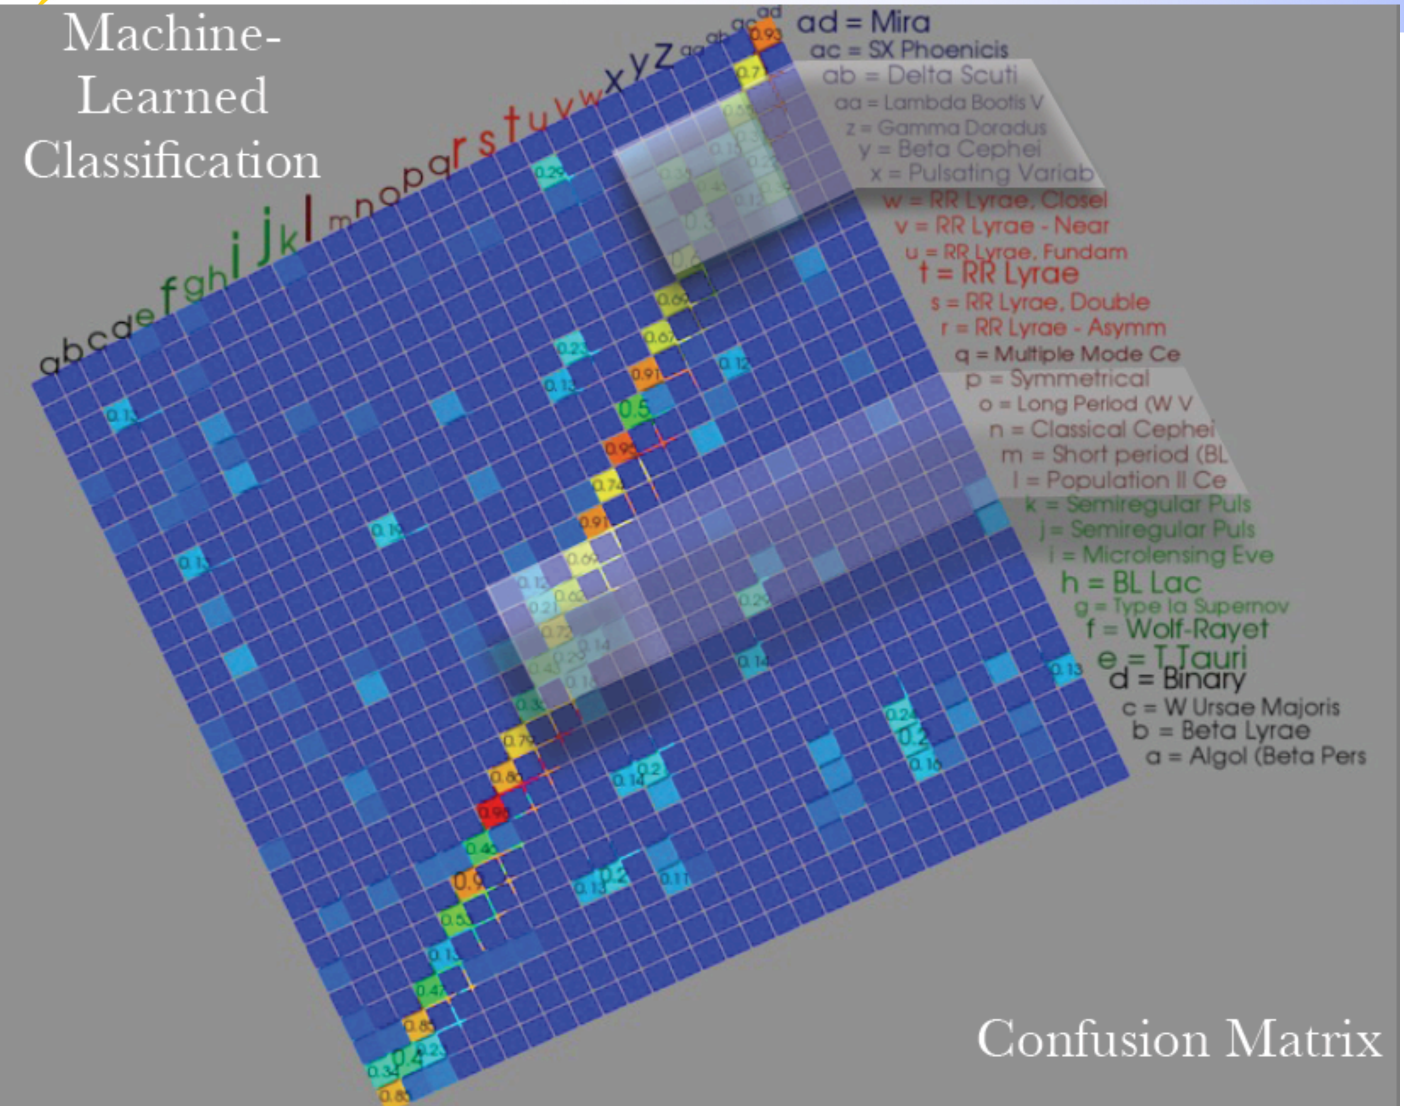

Complications to traditional methods include varying uncertainties in data, non-structured temporal sequence (bad weather, etc.), differing levels of historical information (in SDSS or not, known host in NED, etc.)

And this is just for stars…we also have ones for SNe, AGN…

Turn-around

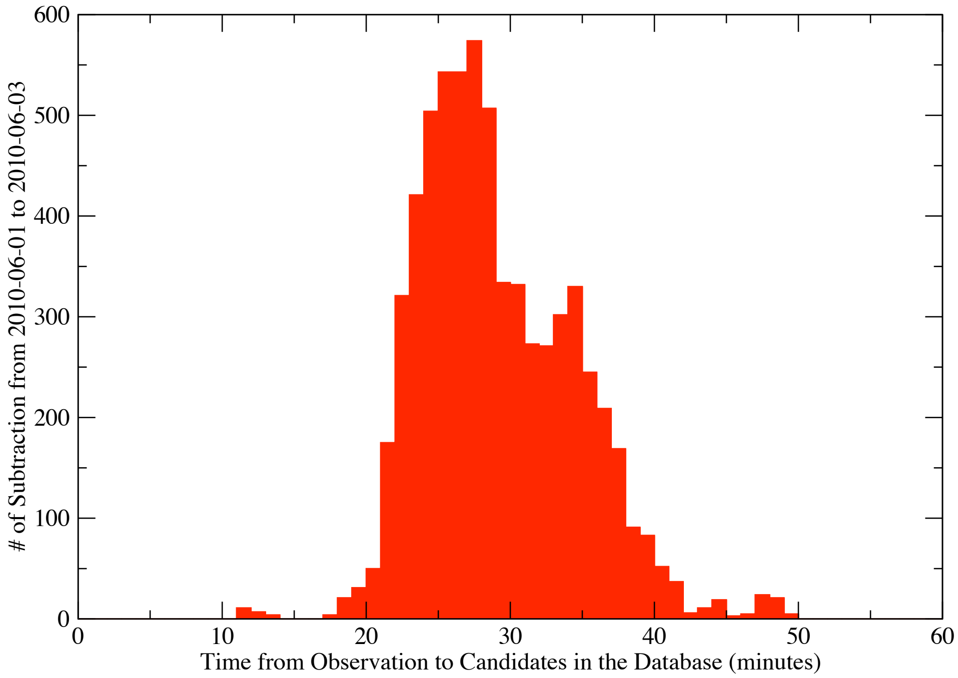

The scanning is handled in three ways:

(1) Individuals can look through anything they want and save things to the PTF database
(2) SN Zoo
(3) UCB machine learning algorithm is applied to all candidates and reports are generated on the best targets and what they are likely to be (SN, AGN, varstar) by comparison to extant catalogs as well as the PTF reference catalog. These come out ~15 min after a group of subtractions are loaded into the database.

On June 3, 2010 we were able to photometrically screen 4 SN candidates with the Palomar 60" telescope in *g*, *r* and *i*-band (50% of the time on P60 is devoted to this) within 2.5 hrs of discovery on the Palomar Schmidt and take spectra of them at Keck the same night. Now a nightly occurrence.

NERSC NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

U.S. DEPARTMENT OF ENERGY Office of Science

BERKELEY LAB

Robot -10vdl

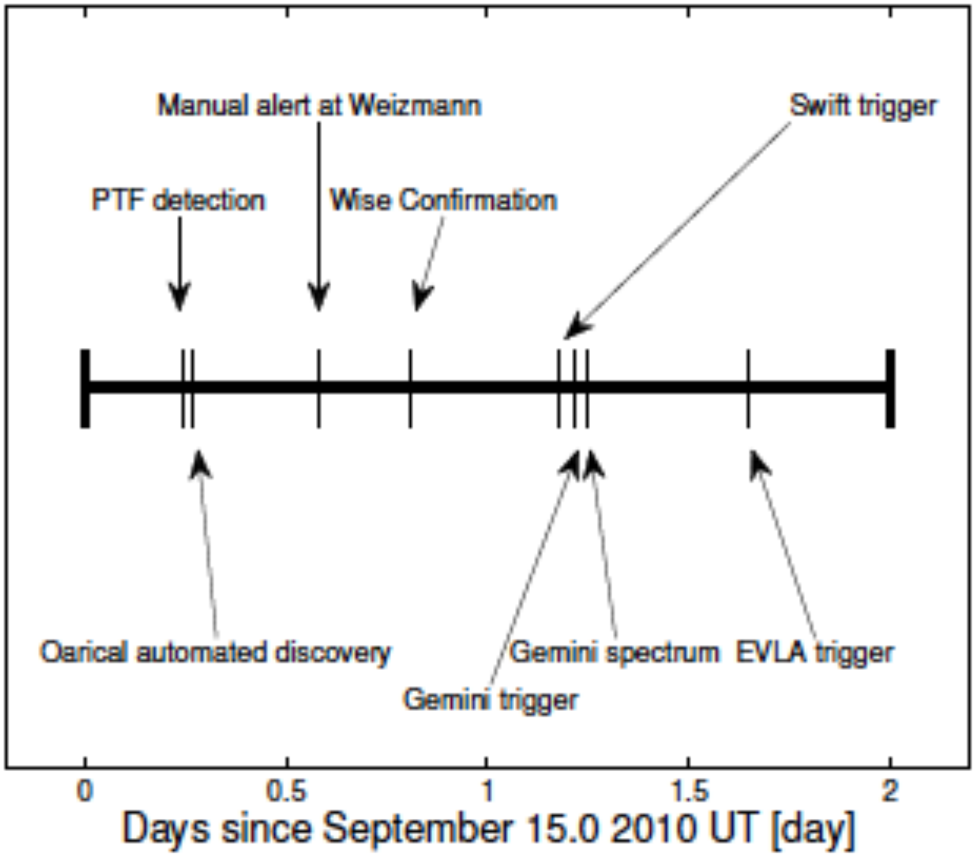

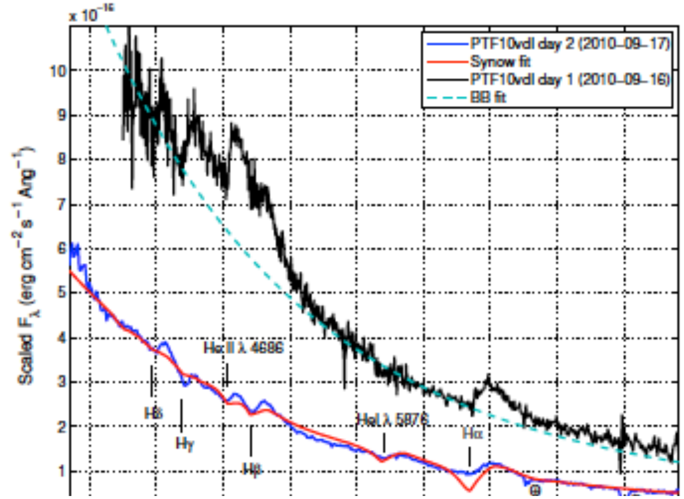

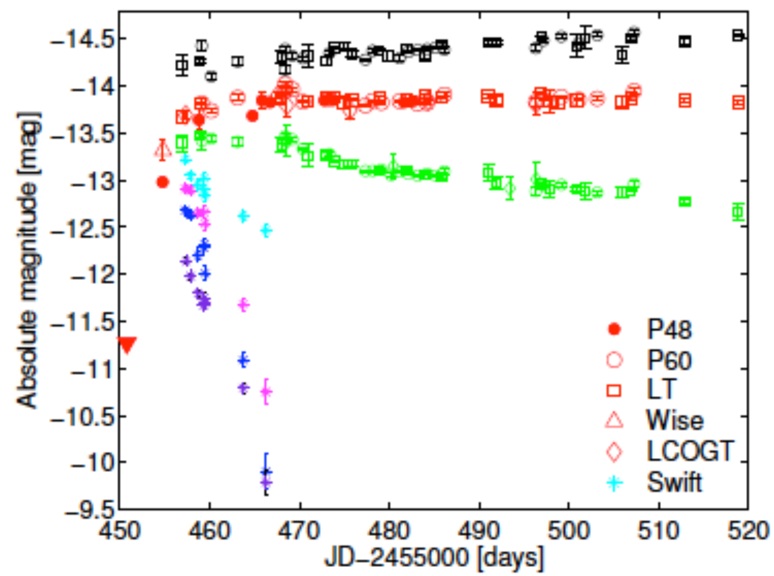

Discovery and follow-up of PTF 10vdl a SN II.

PTF Totals

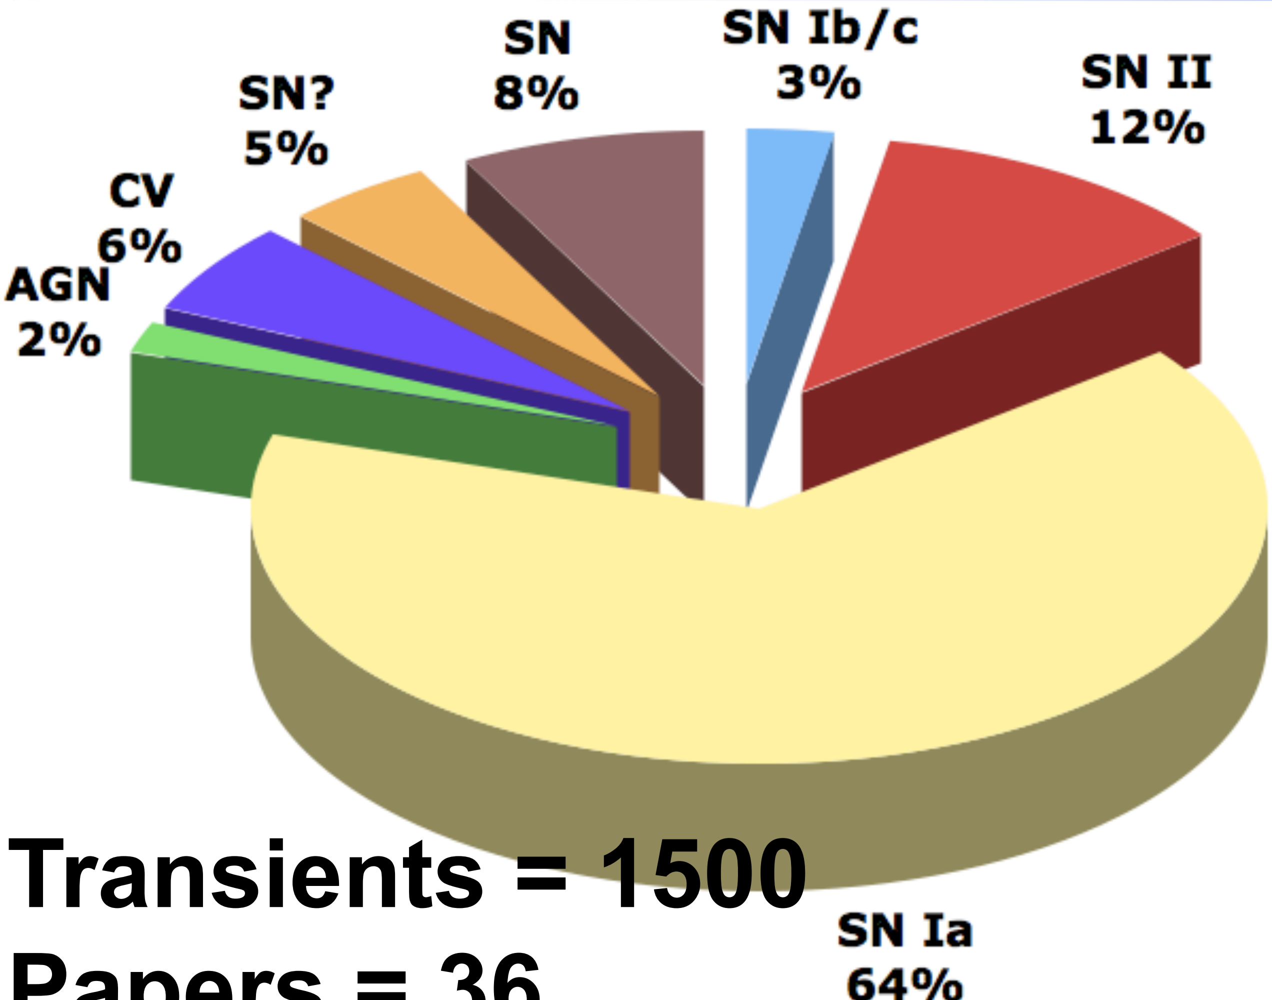

Transients = 1500
Papers = 36

In addition to these we have followed 2 triggers from IceCube and one from LIGO.

We estimate that at the end of the survey we will have 40B detections in the individual images and 40B detections in the deep co-additions.

PTF Totals

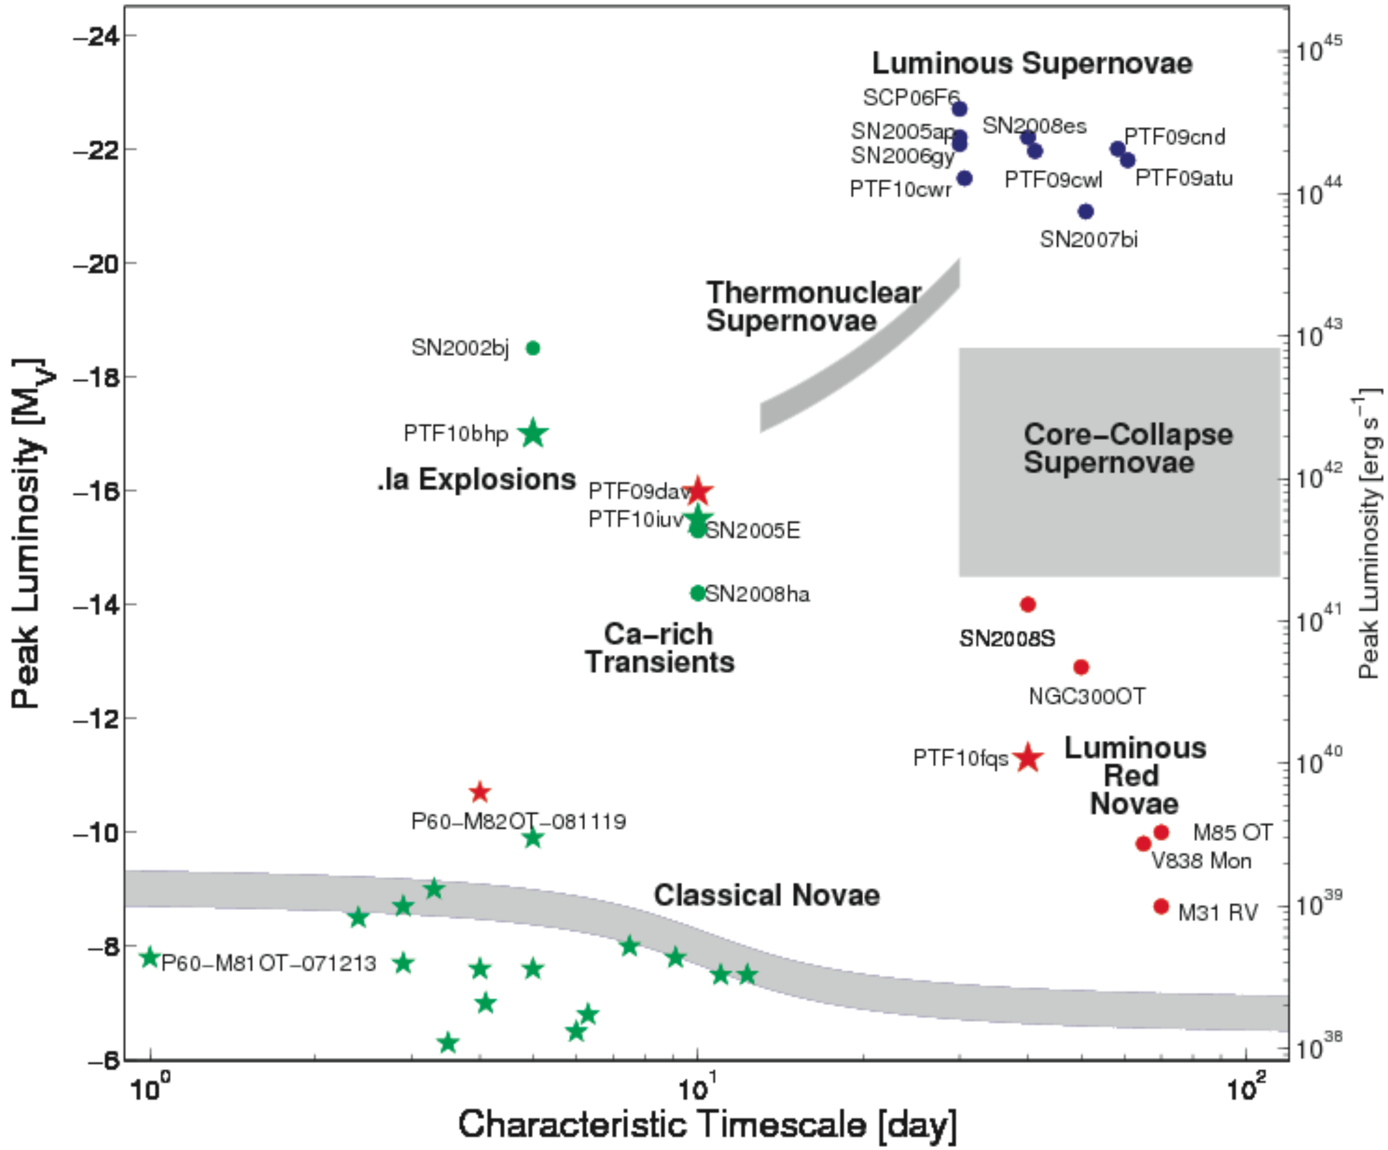

Luminous Supernovae
SCP06F6
SN2005ap
SN2006gy
PTF10cwr
SN2008es
PTF09cwl
PTF09cnd
PTF09atu
SN2007bi
Thermonuclear Supernovae
Core–Collapse Supernovae
SN2002bj
PTF10bhp
.Ia Explosions
PTF09dav
PTF10iuv
SN2005E
SN2008ha
Ca–rich Transients
SN2008S
NGC300OT
PTF10fqs
Luminous Red Novae
M85 OT
V838 Mon
M31 RV
P60–M82OT–081119
Classical Novae
P60–M81OT–071213
Peak Luminosity [M_V]
Peak Luminosity [erg s^{-1}]
Characteristic Timescale [day]
–24
–22
–20
–18
–16
–14
–12
–10
–8
–6
10^{0}
10^{1}
10^{2}
10^{45}
10^{44}
10^{43}
10^{42}
10^{41}
10^{40}
10^{39}
10^{38}

Near Future

Next Generation Transient Survey (aka PTF-II)

- Upgrade to 5X PTF: 36 sq. deg. (~ 1 billion pixels)
- Would like to explore the sky on 100s timescales
- Turnaround in 10-20 minutes with list of new candidates
- Ingest SDSS, BOSS, NED, etc. catalogs to refine our understanding of these candidates in real-time
- Able to handle Advanced LIGO, neutrino detectors, etc.

NERSC
NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Bottlenecks…

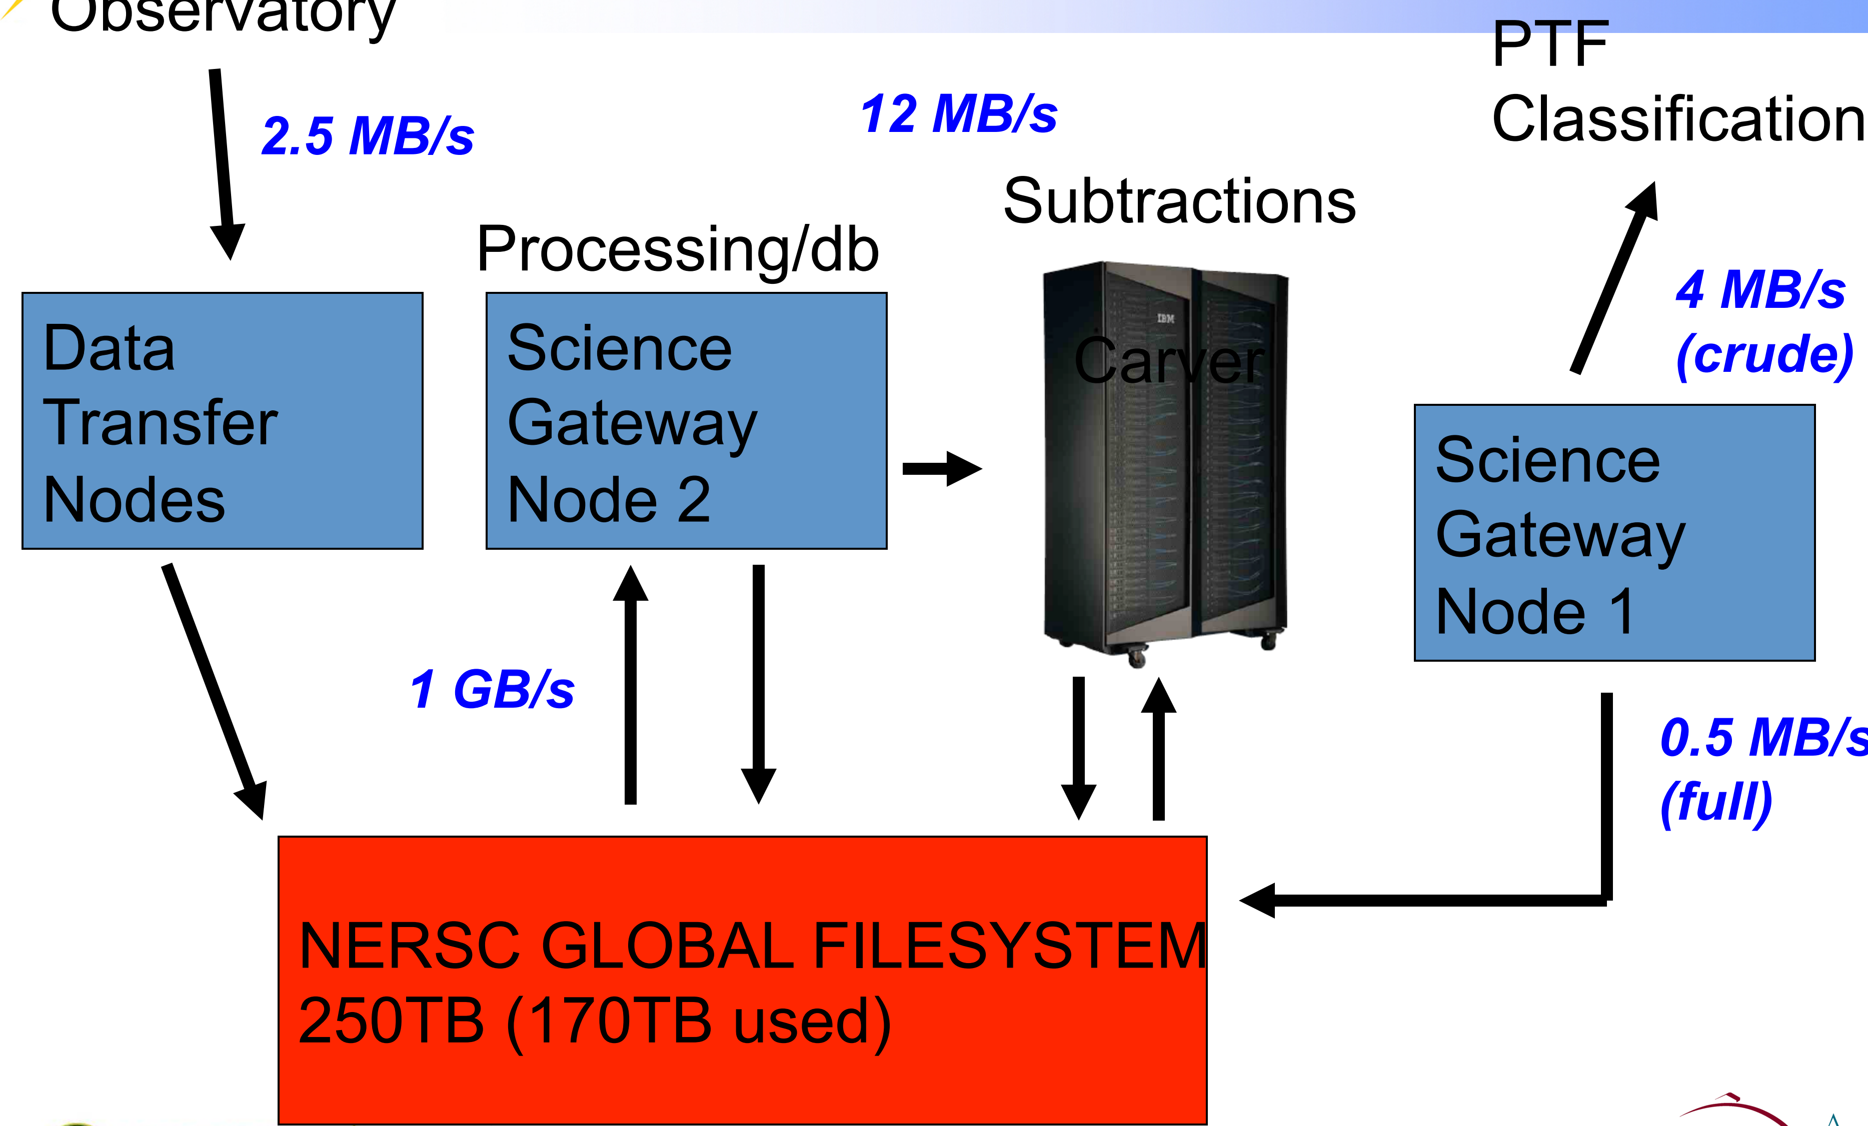

Bottlenecks…crude *vs.* real

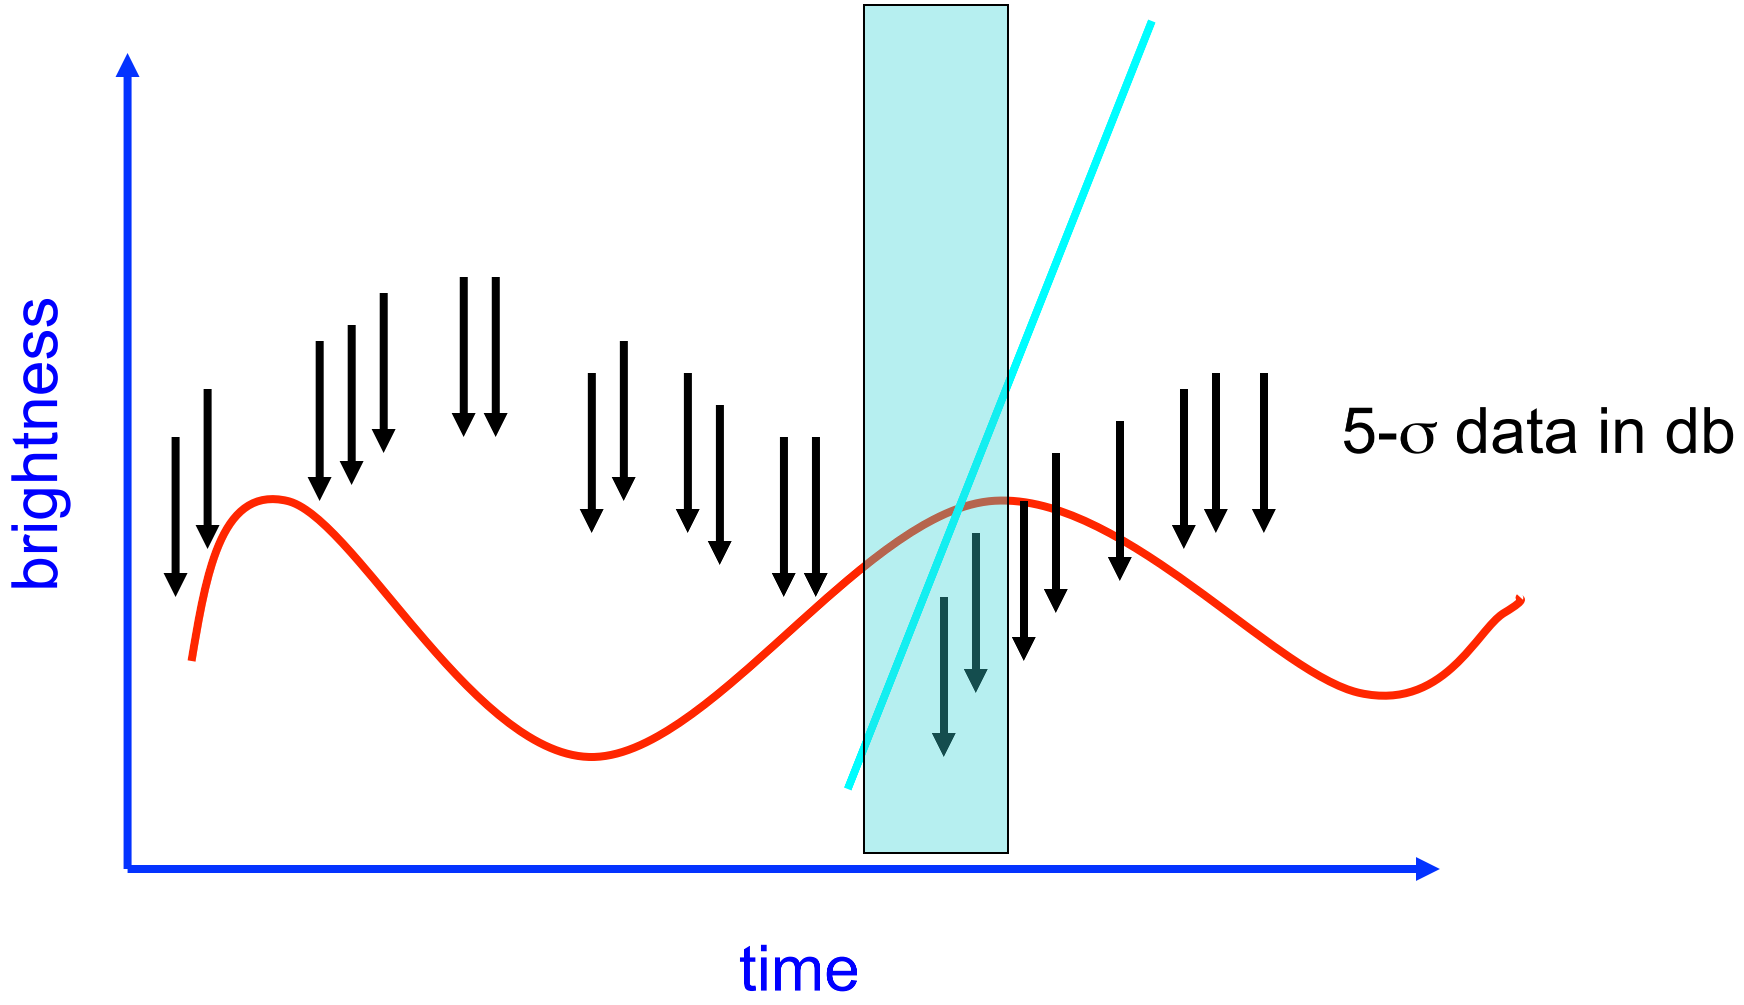

Heavy *Random* I/O

11.20.09

SDSC, UC San Diego, LBNL Team Wins SC09 'Storage Challenge' Award

Team Highlights Flash-Memory of SDSC's New "Dash" and "Gordon" Systems

SDSC Storage Challenge team members (L to R) Jiahua He, Michael Norman, Arun Jagatheesan, and Allan Snavely. SDSC, along with LBNL and UC San Diego researchers, won the Storage Challenge competition, announced this week at SC09 in Portland, Oregon.

A research team from the San Diego Supercomputer Center (SDSC) at UC San Diego and the University of California's Lawrence Berkeley National Laboratory has won the Storage Challenge competition at SC09, the leading international conference on high-performance computing, networking, storage and analysis being held in this week in Portland, Oregon.

The research team based its Storage Challenge submission for the annual conference on the architecture of SDSC's recently

	Forward Q1	Backward Q1
DASH-IO-SSD	11s (145x)	100s (24x)
Existing DB	1600s	2400s

SC09 Storage challenge allowed us to couple both the SDSS db and the PTF candidate db to ask the question, which objects that we think are qso in the static SDSS data vary like one in the PTF data. PTF db is now 165GB and growing nightly!

Heavy *Random* I/O + analytics

Aster Data: Analytics Application Platform

Aster's 'Data-Application Server'

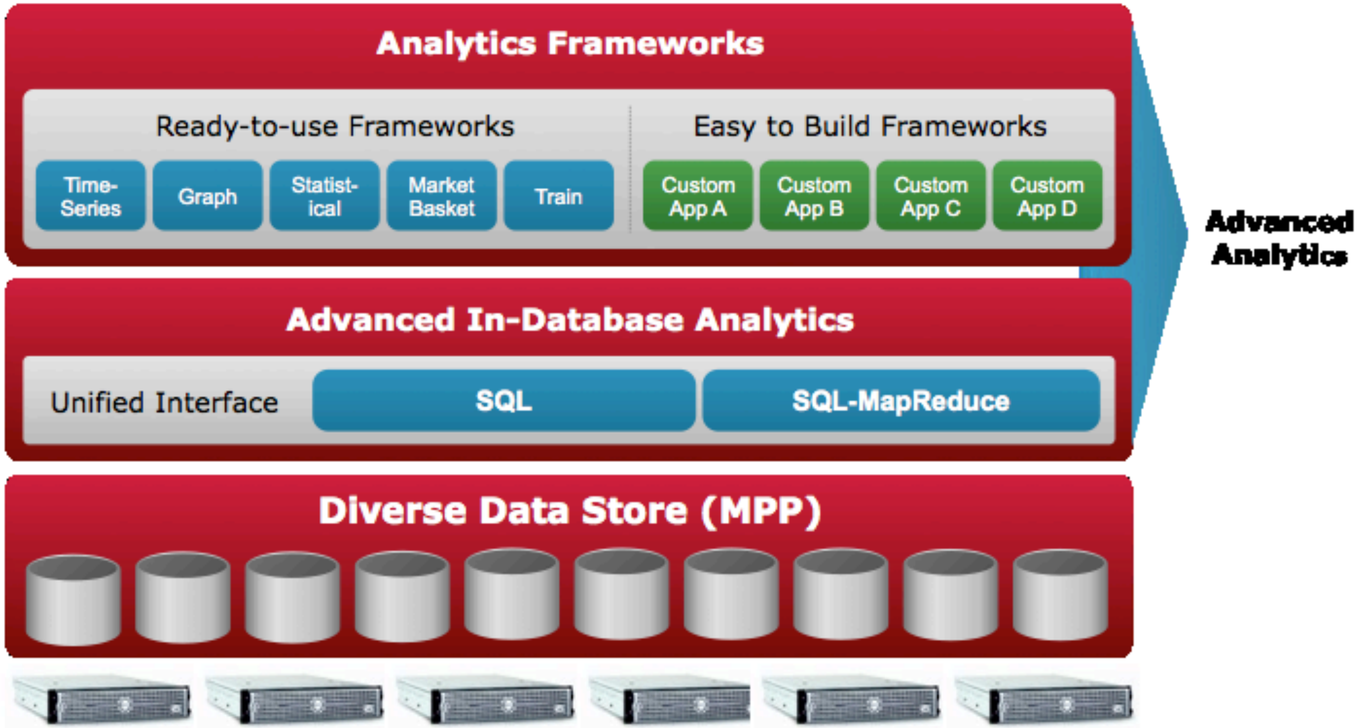

17

Aster Data provides a parallel db solution that also allows us to embed many of our machine learning algorithms. Already handle PB datasets.

Likely will couple both solutions (Aster + SSD).

Conclusions - Future

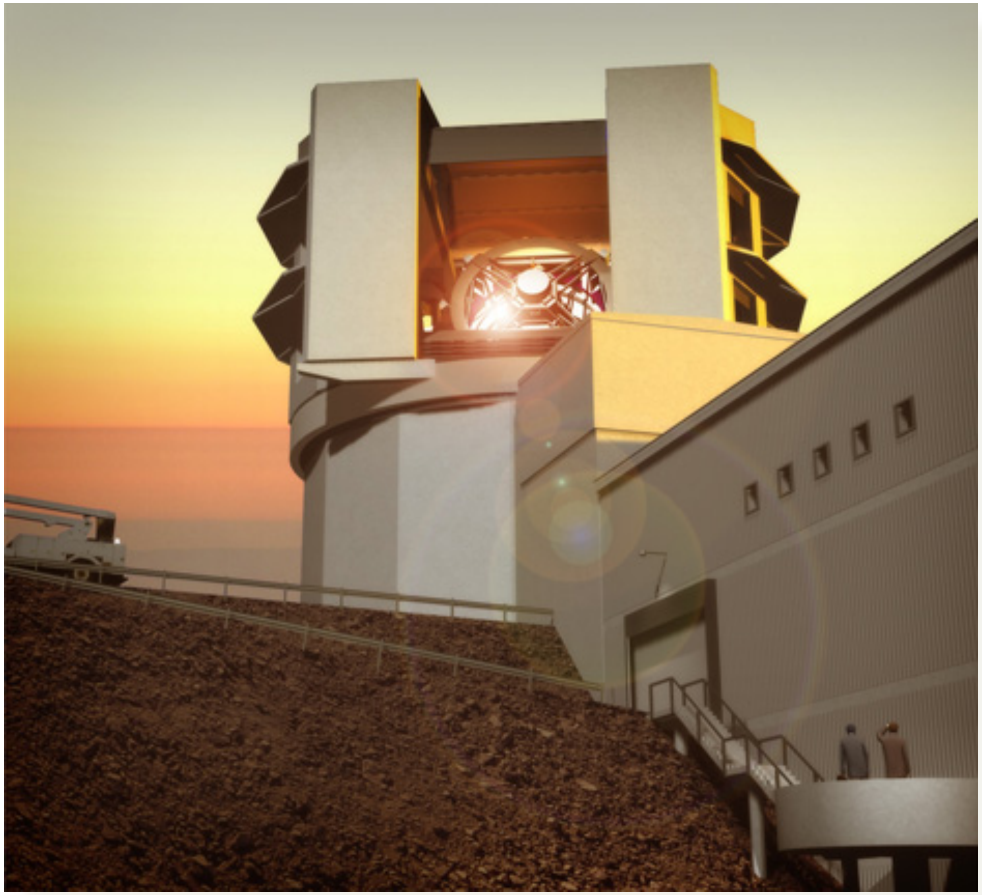

LSST - 15TB data/night
Only one 30-m telescope

Future Concerns

- Power – 20MW per facility
- Data – PB’s per simulation
- Programming – Investment to rewrite?