Low Latency Transient Searches

Jameson Rollins

Department of Physics
Columbia University
New York, NY

LIGO Seminar
California Institute of Technology
March 30, 2010

LIGO-G1000405

Outline

1. **Motivation**
   - Initiating electromagnetic followup observations
   - Detector and data characterization
2. **The Low-latency Analysis Pipeline**
   - Low-latency data calibration and distribution
   - Event trigger generation
   - Event followup
3. **Looking Ahead**

Motivation

Relative times of gravitational and other emissions

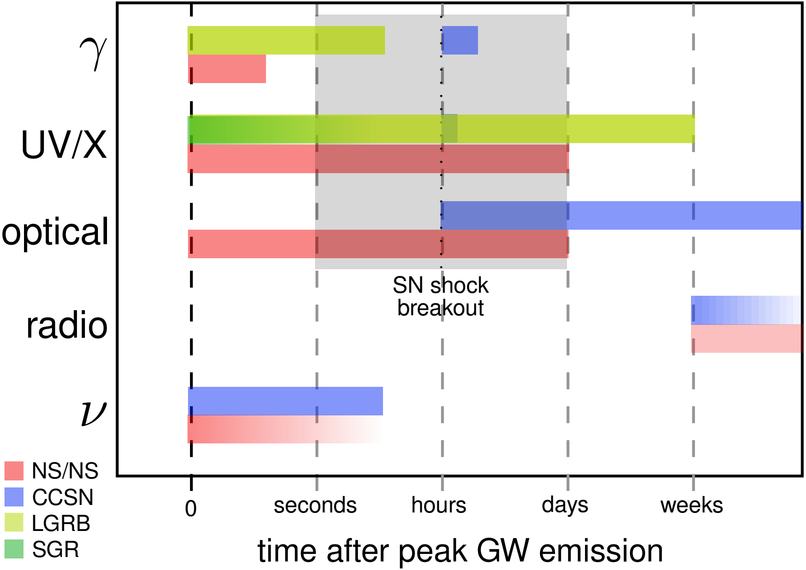

Thank you C. Ott, A. Corsi, I. Bartos, S. Marka!

Relative times of gravitational and other emissions

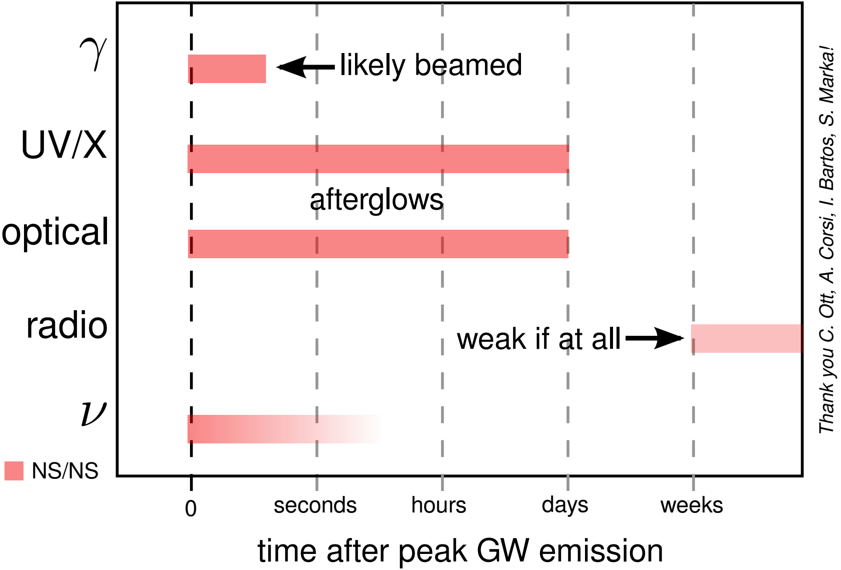

Relative times of gravitational and other emissions

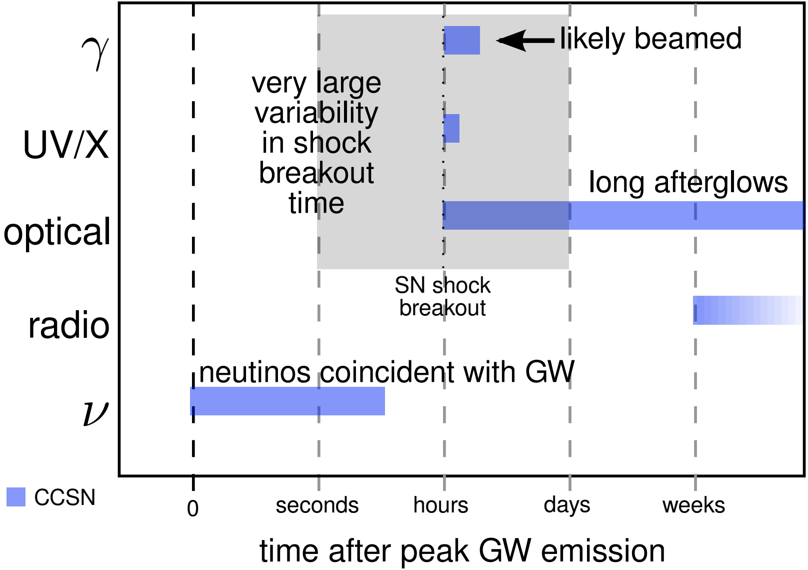

Thank you C. Ott, A. Corsi, I. Bartos, S. Marka!

Relative times of gravitational and other emissions

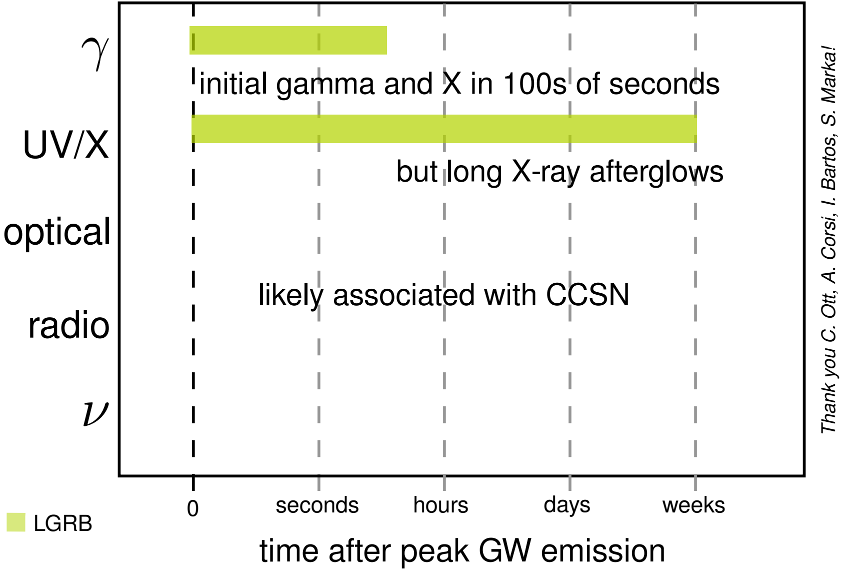

Relative times of gravitational and other emissions

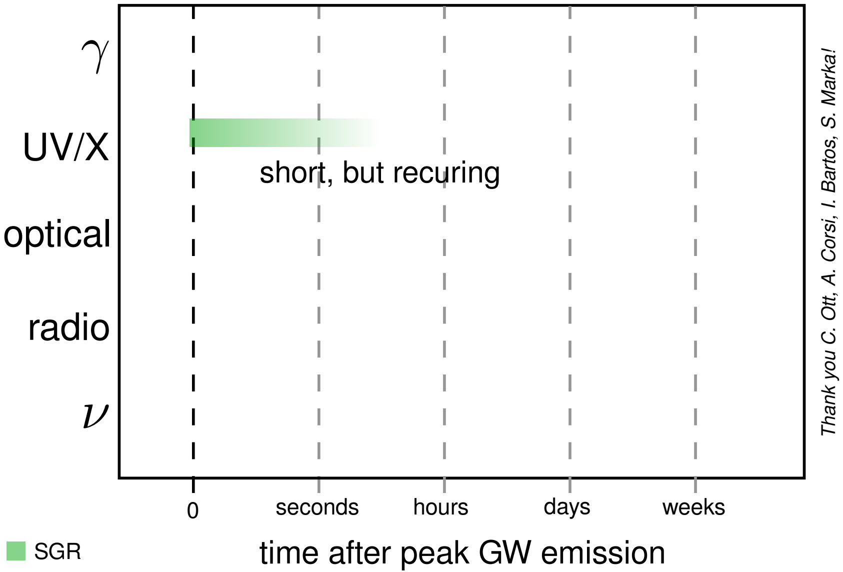

Relative times of gravitational and other emissions

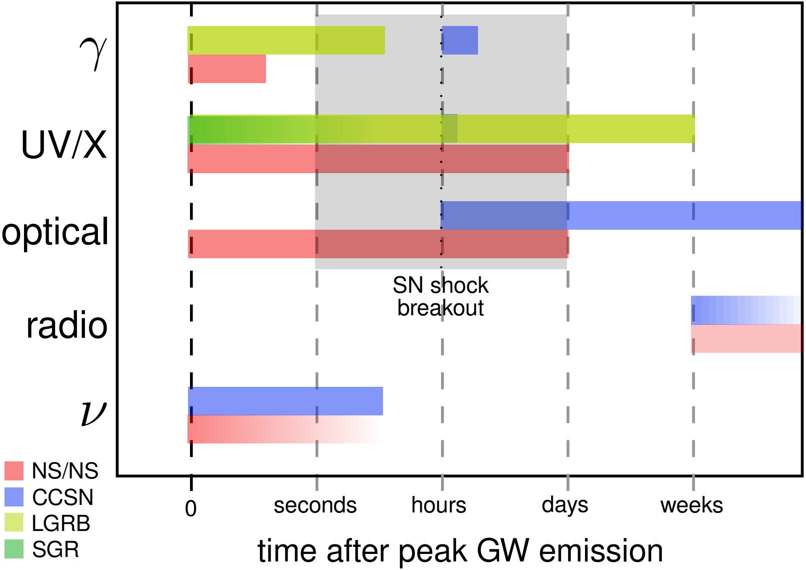

Thank you C. Ott, A. Corsi, I. Bartos, S. Marka!

Initiating electromagnetic followup observations

Very strong scientific motivation for promptly identifying gravitational wave event candidates:

observation of a non-GW counterpart to a GW event

- greatly increase detection confidence, effectively increasing the sensitivity of a search
- dramatically increase information about the source (position, distance, composition, etc.)
- catch what other non-GW observatories might miss (trigger observations of GRB afterglows for off-axis GRBs)

See the original LOOC UP paper (CQG/25/184034).

Detector and data characterization

The S5 analyses clearly indicated that excessive non-Gaussian noise and glitches adversely affect the searches, compromising sensitivity.

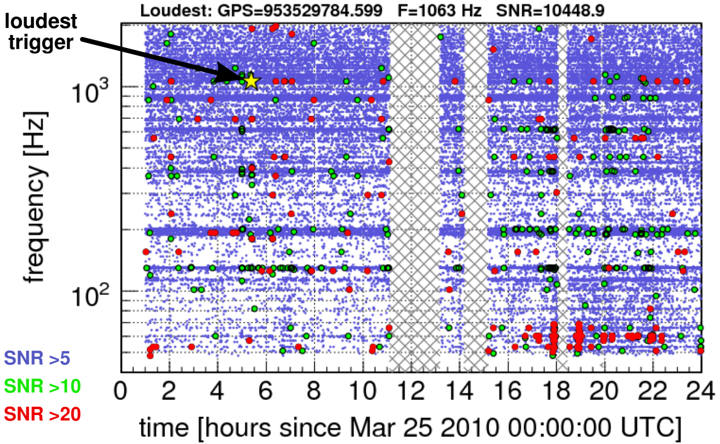

Low-latency analysis = faster feedback to commissioners.

The Low-latency Analysis Pipeline

Development history

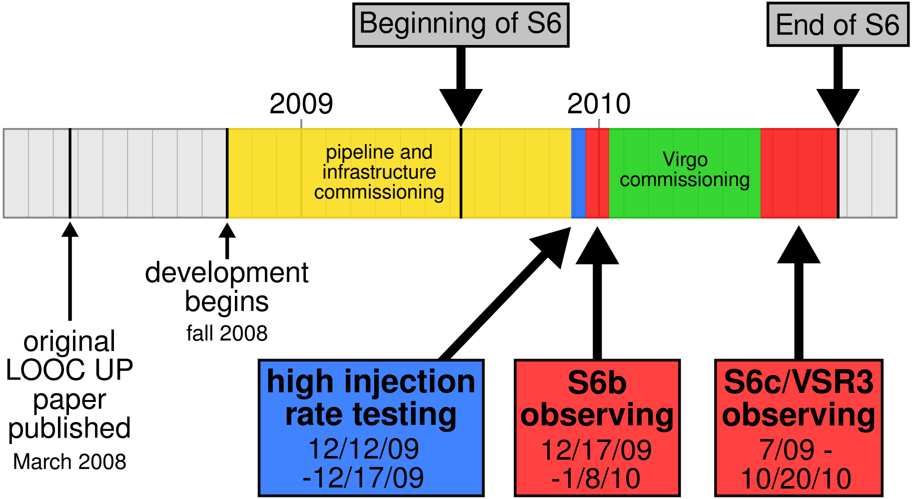

Development of low-latency infrastructure

Much new infrastructure needed to be built:

- low-latency data calibration
- low-latency segment generation (IE. DQ/vetos)
- distribution of above to analysis sites
- notification distribution system

Massive kudos to DASWG for helping make all of this happen:

Jordi Burguet-Castell, John Zweizig, Xavier Siemens, Ping Wei, Duncan Brown, Larne Pekowsky, Greg Mendell, Igor Yakushin, Dan Kozak, Adam Mercer, Scott Koranda, Patrick Brady, Xavier Amador, Josh Smith, Carsten Aulbert, Stuart Anderson, Brian Moe, Larry Price, et. al.

Modification of the event generators

In addition, the existing event generators (Omega Pipeline and Coherent WaveBurst) had to be modified as well:

- Position reconstruction code had to be developed and integrated.

Accurate and precise position reconstruction is critical for doing any followup science.

- Pipelines had to be modified/expanded to run in a low-latency mode.

Not to mention **review**!

The S6 low-latency pipeline

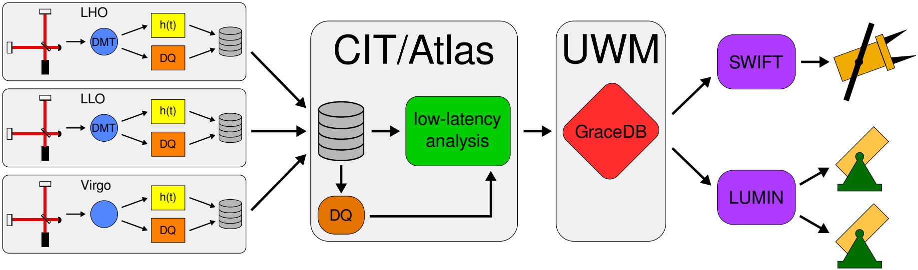

Low-latency data calibration and distribution

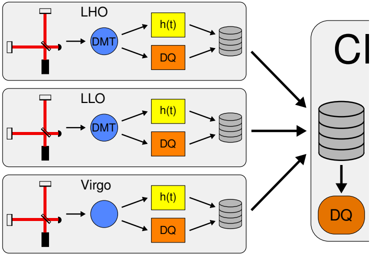

Low-latency data calibration and distribution

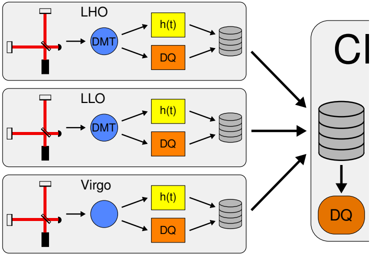

- new LIGO DMT/LAL low-latency calibrated frame generation
  - 16 second frames
  - ~16 second latency
  - include state vector
- new low-latency DQ/veto system
  - stored in XML files for faster processing
- frames and DQ XML both distributed to analysis sites via `rsync`

Next stage of pipeline...

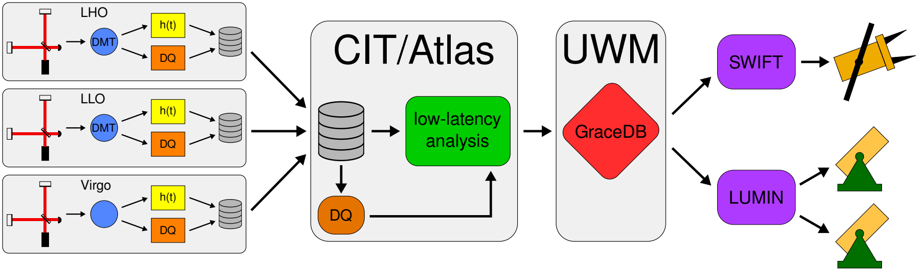

Event trigger generation

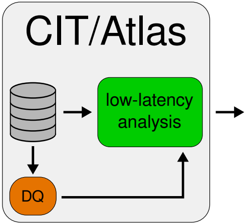

Omega Rapid Online Analysis

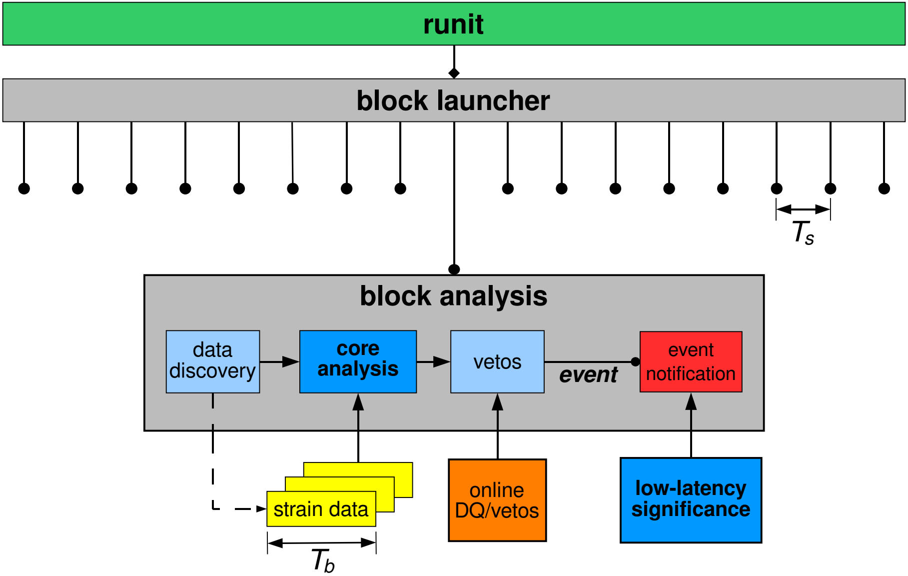

OROA: block launcher

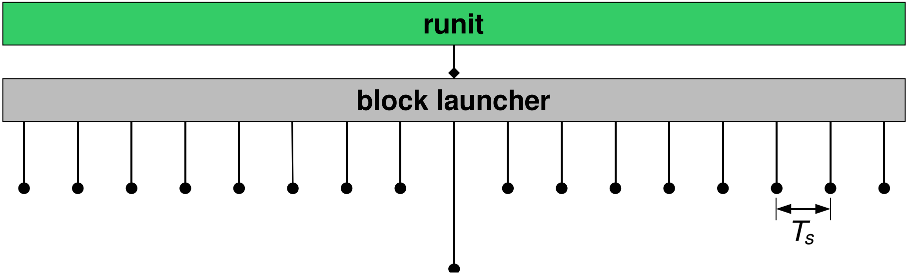

OROA: block launcher

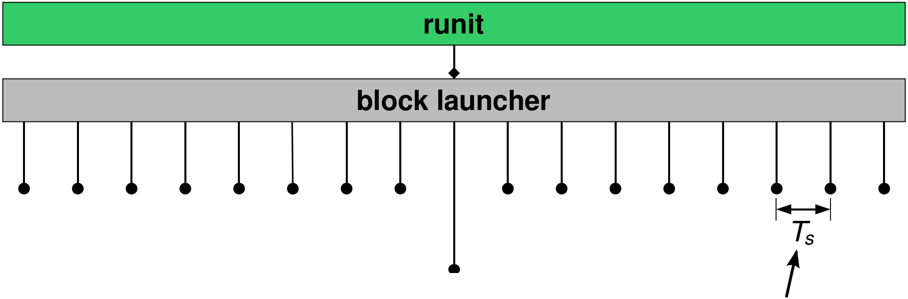

block launcher runs continuously, launching block analyses in the background

- each block completely independent
- each block analysed as fast as data for that block becomes available
- blocks monitored for problems

launcher is monitored by `runit` process supervision

Omega Rapid Online Analysis

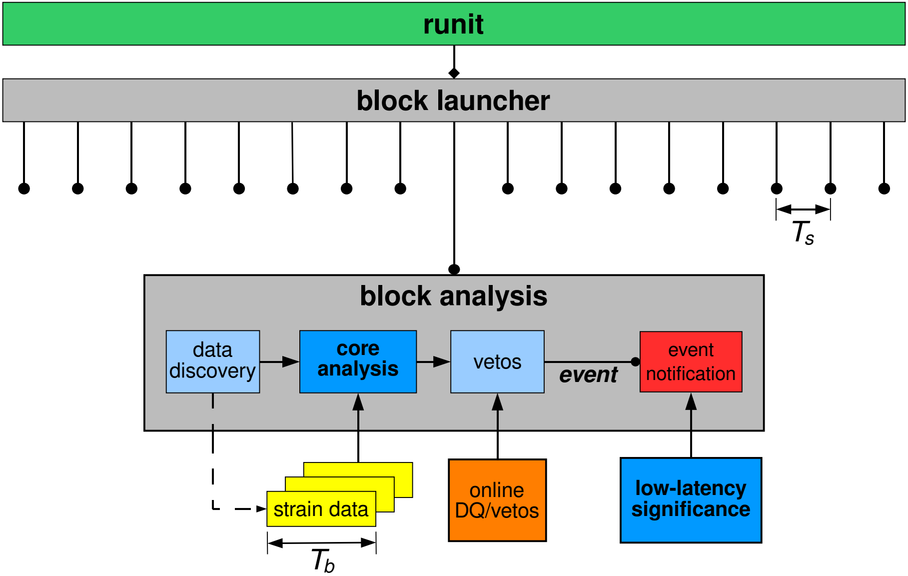

OROA: block analysis

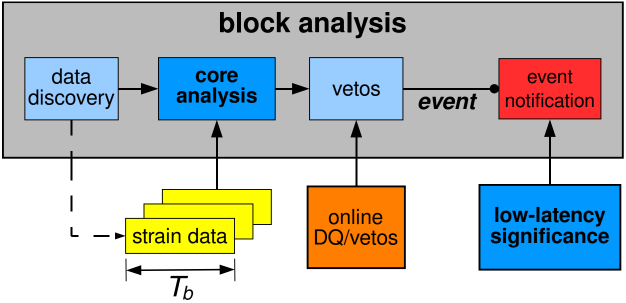

OROA: block analysis

block analysis handles all processing of single block of data

- data discovery: waits for data from all IFOs, or timeout
- core analysis: Omega Pipeline identifies most significant event in block
- apply vetos: full IFO analyzed
  minimal fixed threshold
  Cat2 DQ vetos
- event notification: passed to GraceDB, with significance

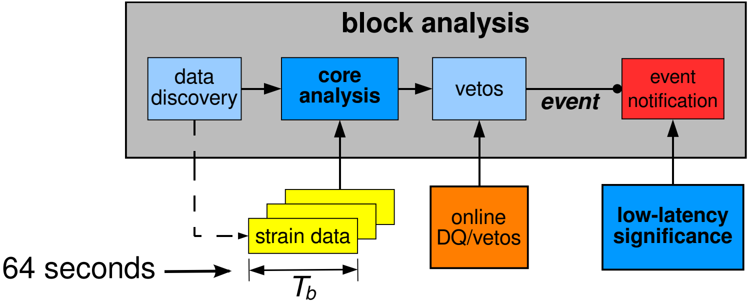

OROA core

- core analysis: Omega Pipeline identifies most significant event in block

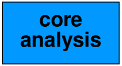

OROA core: Omega Pipeline

Omega Pipeline: hierarchical, coherent, transient search pipeline, built on Q-Pipeline, written mostly in MATLAB and Python.

- initially, each data stream (IFO) is analyzed independently:
  - each block is convolved with multiple sine-gaussian templates of different Qs
  - significant time/frequency/Q tiles are clustered
- clusters from each IFO are tested for time and frequency coincidence
- the loudest coincident cluster is passed to internal followup analyses:
  - **Bayesian**: calculates direction posterior
  - **xCoherent**: computes X-Pipeline-like coherent statistics based on sky location determined from Bayesian analysis

Bayesian position reconstruction

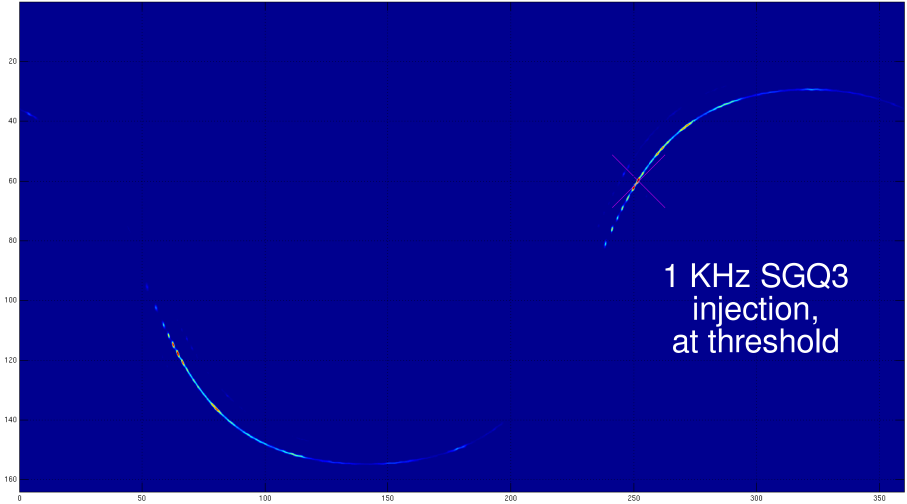

Bayesian position reconstruction

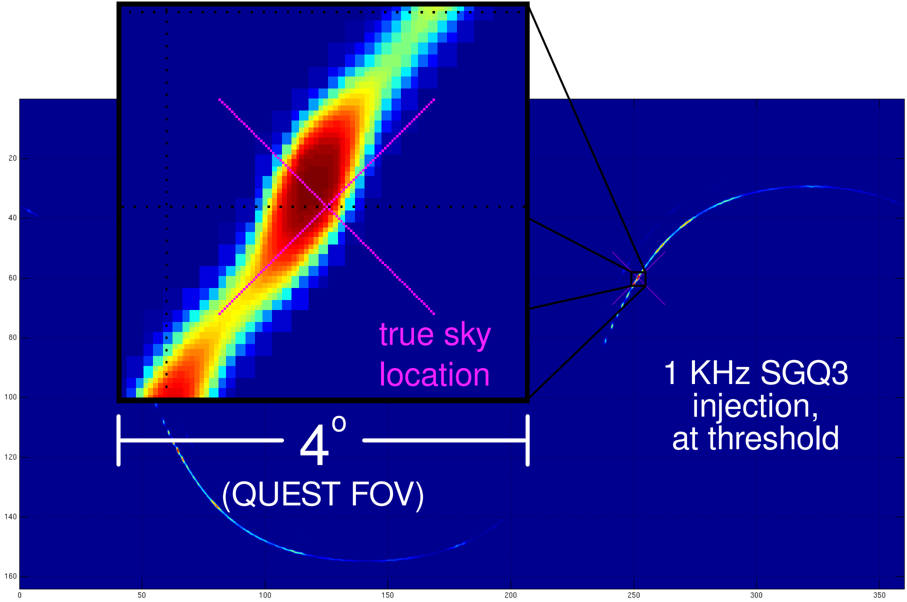

OROA: resource requirements

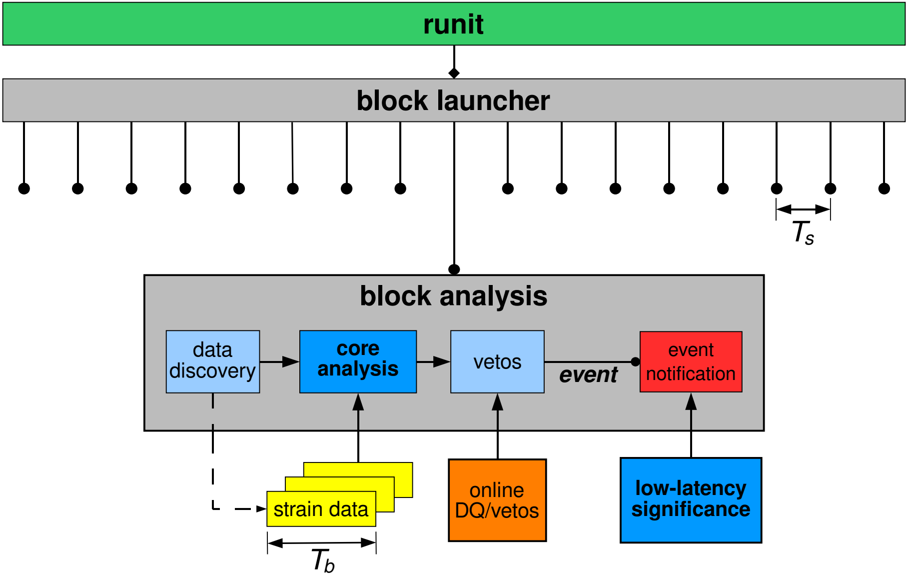

OROA: resource requirements

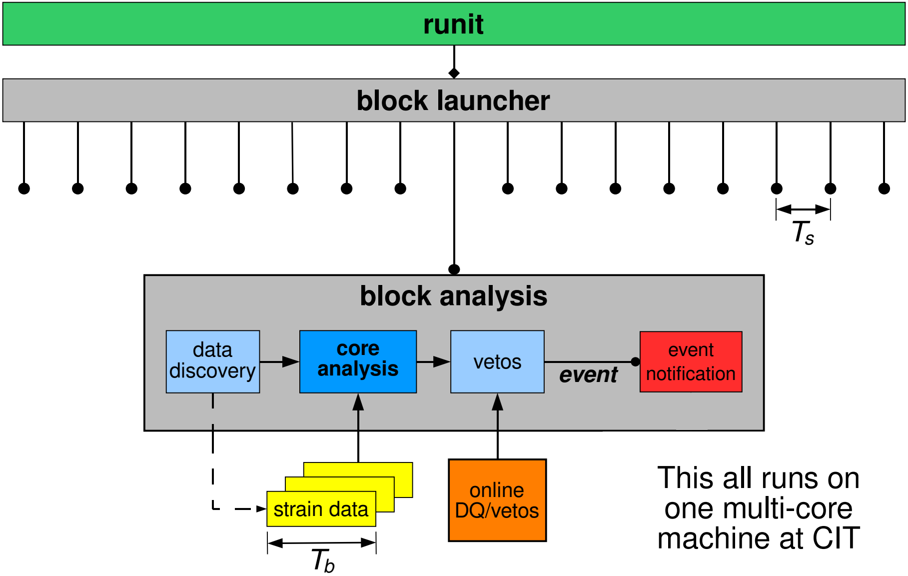

This all runs on one multi-core machine at CIT

OROA: but...

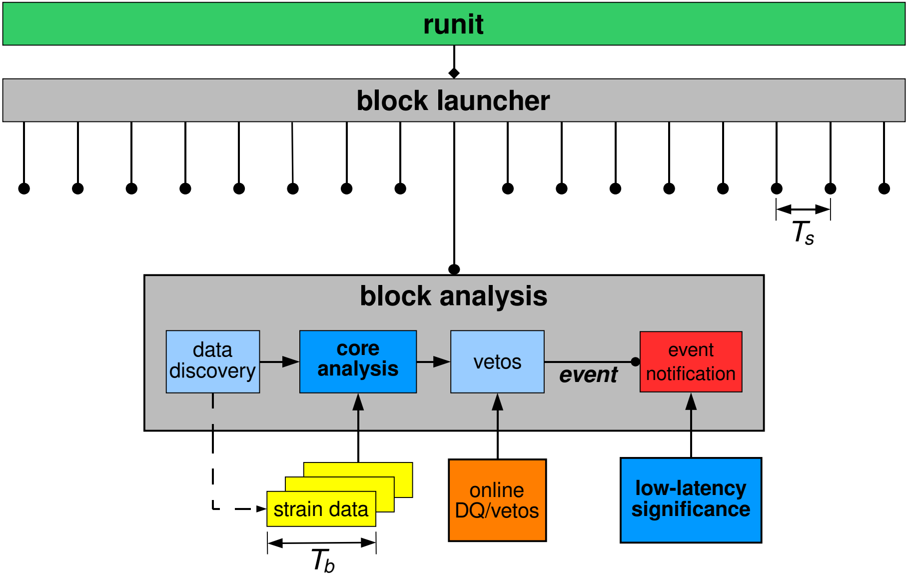
runit
block launcher
T_s
block analysis
data discovery
core analysis
vetos
event
event notification
strain data
T_b
online DQ/vetos
low-latency significance

OROA: low-latency significance

Low-latency significance most computationally intensive part of the pipeline

- 100 time-slide analyses of last 30 minutes every 30 minutes
- significance of every event is measured against background from:
  - most recent run
  - last day of background
  - background from entire run "epoch"

Requires ~100 cores from CIT cluster continuously

low-latency significance

OROA: results

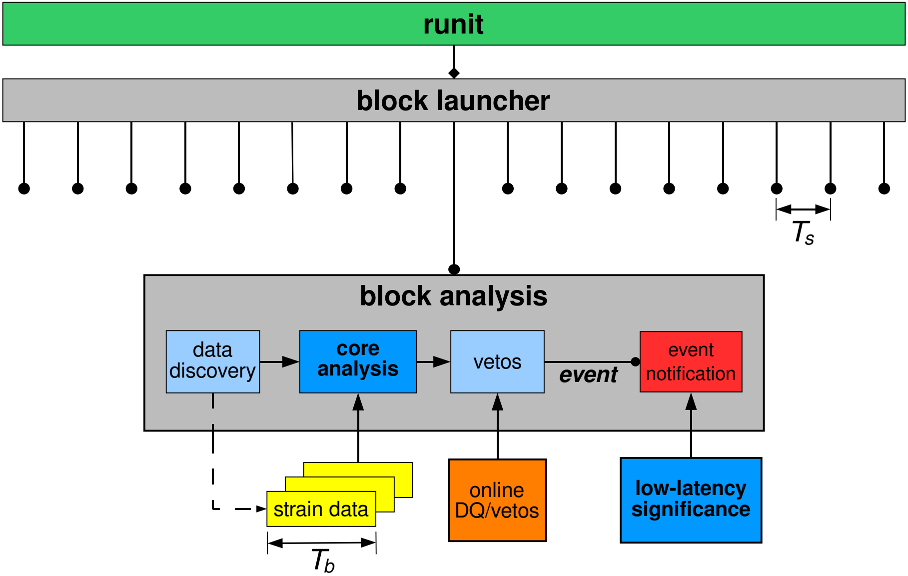

OROA: results

OROA main H1L1V1 analysis web page

Event: 946135007

⇑

first ever triggered observation! (QUEST)

OROA: results

0946135007.718750000 - Tue Dec 29 15:16:33 2009 UTC

event: 0946135007.718750000
discovery: 946135155
latency: 147 seconds

- wevent
- event.info
- scans:
  - LHO
  - LLO
  - Virgo
- segment / log
- permalink
- gracedb: G3821

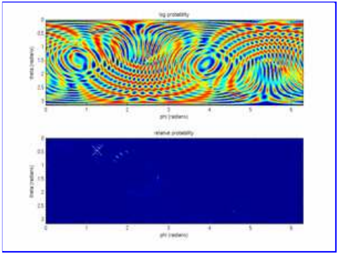

logOdds: 6.244
probSignal: 0.998

duration: 0.068 s
frequency: 429.000 Hz
bandwidth: 14.700 Hz

modeTheta: 0.459 rad
modePhi: 1.254 rad

significance	rank	livetime (years)	events/year
background, all:	323 / 656314	1.16465079727	277.336349021
background, last day:	24 / 56211	0.0997482698304	240.605677079
background, last run:	0 / 0		
zero lag (includes injections):	58 / 11340	0.0201232032854	2882.24489796

Latencies

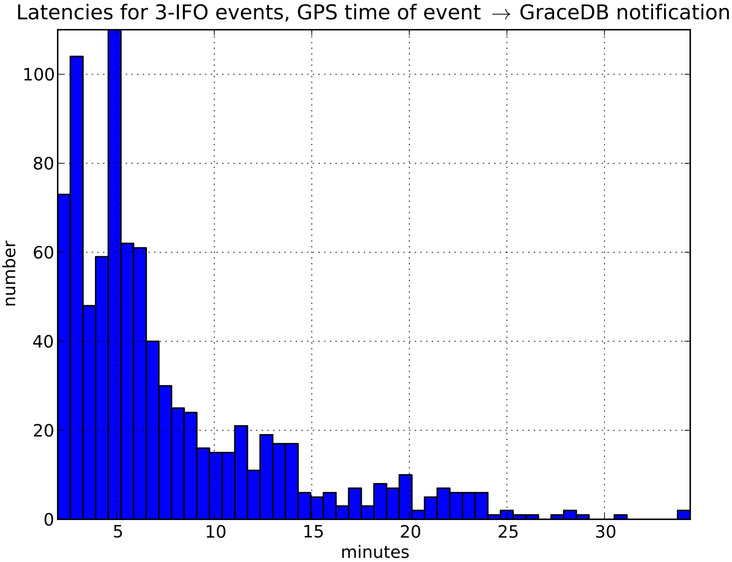
Latencies for 3-IFO events, GPS time of event → GraceDB notification
number
minutes
0
20
40
60
80
100
5
10
15
20
25
30

The "OmegaGrams" (shown in slide 6) are created by a single interferometer version of the OROA (first two stages of the full pipeline):

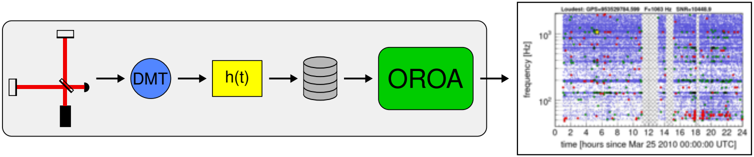

These are running independently on dedicated nodes at the LIGO sites and Virgo.

Last stage of pipeline...

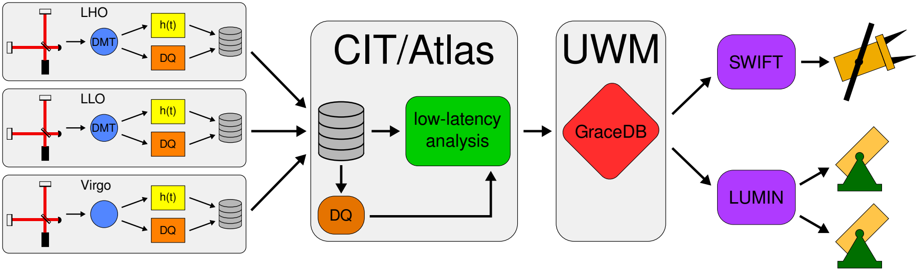

Event followups

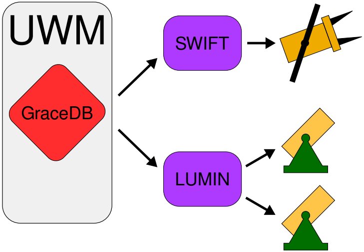

Event followups

- new **GraceDB** event database and notification/distribution system
  - events reported via XMPP
  - event info stored in searchable database
  - listeners can subscribe to event notifications (XMPP, email, RSS)
- followup analyses receive event notification via GraceDB

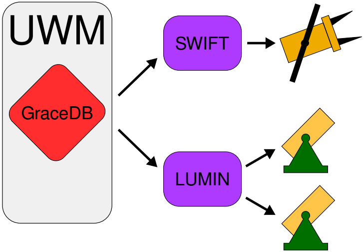

LOOC UP

The LOOC-UP followup analysis uses triggers to point wide-field optical telescopes (QUEST and TAROT).

- events collected from GraceDB
- event position reconstruction information and catalog of local galaxies are used to determine telescope pointings

LOOC UP observations

ETG	GPS	Lag	Observed
Omega	945320899	37 minutes	none
Omega	945895177	30 minutes	none
Omega	946135007	20 minutes	QUEST
cWB	946465347	25 minutes	none
cWB	946786257	25 minutes	none
Omega	946795800	20 minutes	QUEST
cWB	946586256	re-visit	TAROT
cWB	946889212	95 minutes	QUEST

Looking Ahead

What have we learned?

A lot.

- **monitoring is important** - calibration, distribution, ETGs, NFS, etc. all occasionally go down
- **redundancy is needed** - for the ETGs themselves, and for data distribution
- **significance estimation can be tricky** - computing requirements were under-appreciated

Resource requirements

There will likely (hopefully) be a *lot* more low-latency analyses running in Advanced LIGO

- for searches.
- for detector/data characterization.

Hopefully the other searches (CBC, CW?, stochastic?) can get in the game as well.

- MBTAOnline low-latency CBC search is working and will hopefully be sending triggers to GraceDB this summer.

For searches, it's particularly important to pay attention to requirements for significance estimation.

- 10 searches x 1000 time-slides = 10,000 cores **continuously**

Need to get more clever about how to determine event significance.

Streaming data

The use of data streams, as opposed to data read off disk, will likely become increasingly important in aLIGO.

$\Longrightarrow$ **NDS2**

Will probably need things like stream proxies at the analysis sites to keep from overburdening the network.

Conclusion

Low-latency burst analysis in S6 has been a great success.

- 4 observations where made with two observatories based triggers passed from the low-latency analyses via LOOC UP.
- Hopefully many more will be made this summer.

Low-latency analysis is critical moving into Advanced LIGO.

- We have learned a lot, but there is much still to do.

Conclusion

Thank you.