Computational challenges of gravitational-wave searches

Michał Bejger (CAMK PAN)

GPU Days / 23.06.17

Gravitation: Newton vs Einstein

Eternal two body systems:

- ⋆ Absolute time and space,
- ⋆ deterministic solutions.

- ⋆ Stable two-body solution does not exist.
- ⋆ Evolution due to the existence of a third "body", the spacetime.

Evolution of a binary system

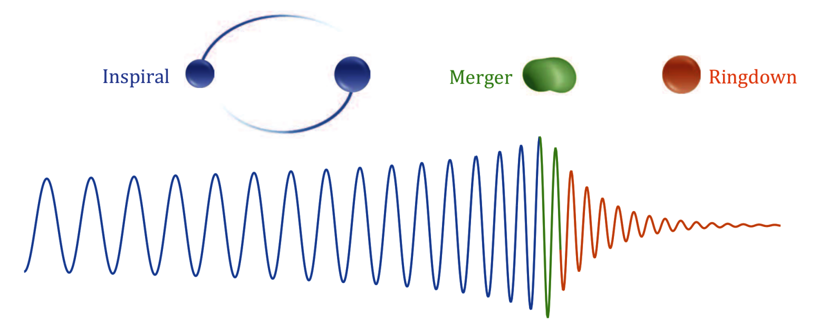

Gravitational waves

Einstein (1916) - in linear regime there are wave solutions to GR equations *(time-varying distortions of the curvature propagating with the speed of light)*:

- ⋆ In realistic astrophysical situations, length-scale of the wave λ is much smaller than other important curvatures $\mathcal{L}$,
- ⋆ Split of the Riemann curvature tensor

$$R_{\alpha\beta\gamma\delta} = R^{GW}_{\alpha\beta\gamma\delta} + R^{B}_{\alpha\beta\gamma\delta}$$

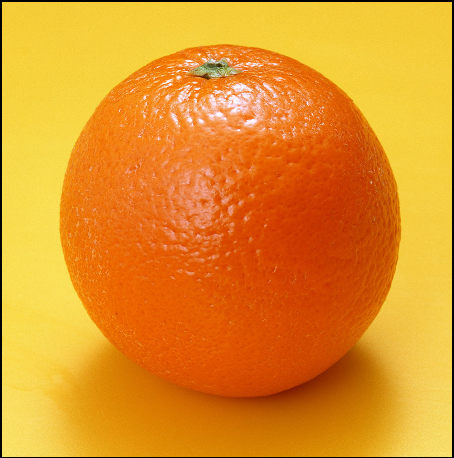

"Kip Thorne's orange": B - large-scale background ($\mathcal{L} \simeq 10$ cm), GW - fine-scale distortions/waves ($\lambda \simeq$ few mm).

LIGO O1 and O2: 3 (”and a half”) events

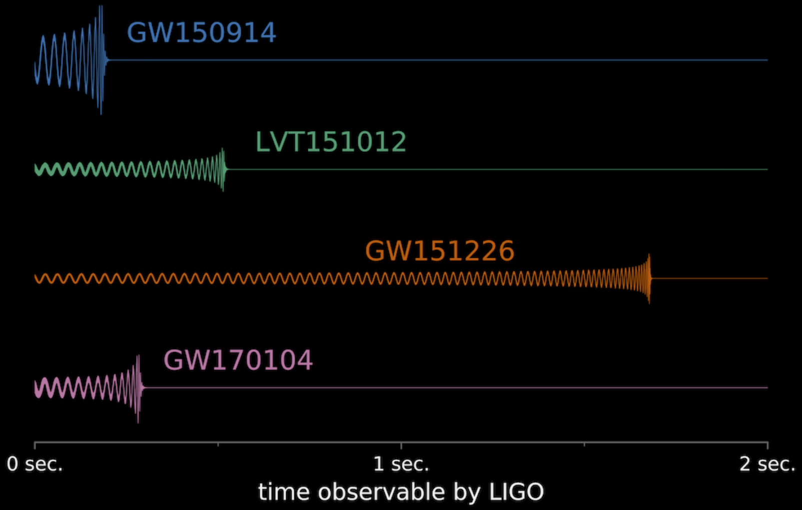

Cataclismic vs continuous types of GW sources

(Hokusai "The Great Wave off Kanagawa")

One-time cataclysmic events, e.g. last moments of binary systems of

- ⋆ black holes (GW150914, GW151226, LVT151012, GW170104),
- ⋆ neutron stars,
- ⋆ black hole and neutron star.

(Shoson "Cranes landing")

Periodic phenomena, e.g.

- ⋆ rotating non-axisymmetric neutron stars ("*gravitational pulsars*"),
- ⋆ wide binary systems.

Detection principle: spacetime distance measurement

Gravitational waves are tiny distortions of spacetime:

$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}$.

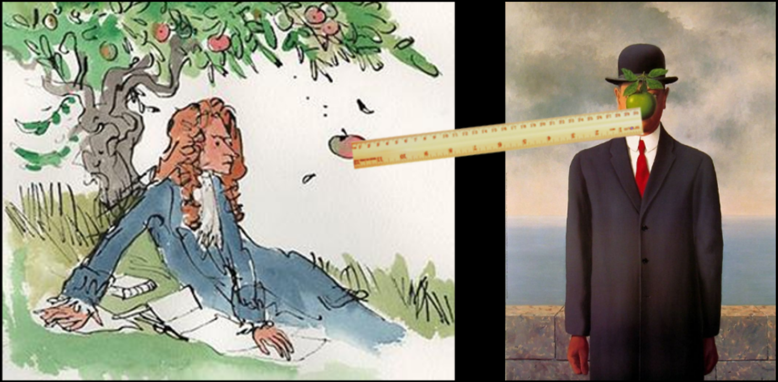

(Quentin Blake "Izaak Newton")

(Rene Magritte "The Son of Man")

"How to measure distance when the ruler also changes length???"

Detection principle: laser interferometry

Light as a ruler. Changes in length $\Longrightarrow$ variation in light travel time:

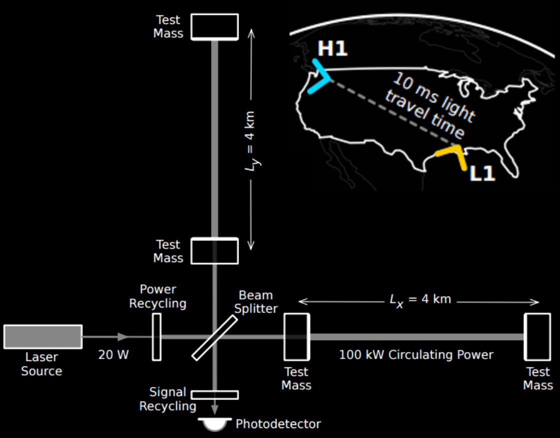

- ⋆ Round-trip time in the x-arm: $t_x = hL/c$, y-arm: $t_y = -hL/c$,
- ⋆ Round-trip times difference:

$$\boxed{\Delta\tau = 2hL/c}.$$

- ⋆ **Phase difference** (dividing $\Delta\tau$ by the radian period of light $2\pi/\lambda$):

$$\boxed{\Delta\phi = \frac{4\pi}{\lambda} hL = \frac{2\pi c}{\lambda} h\tau}.$$

Changes in arms length are **very** small: $\delta L_x - \delta L_y = \Delta L < 10^{-18}$ m (smaller than the size of the proton). Wave amplitude $h = \Delta L/L \leq 10^{-21}$.

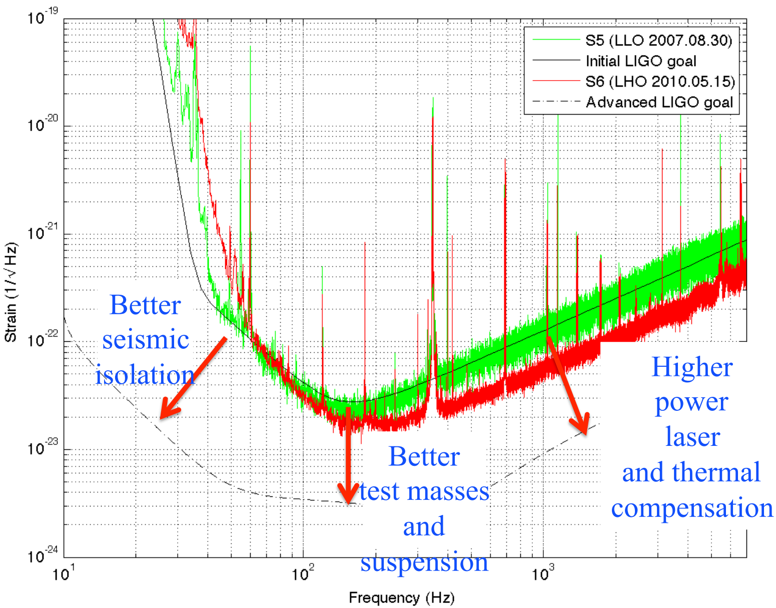

S5 (LLO 2007.08.30)
Initial LIGO goal
S6 (LHO 2010.05.15)
Advanced LIGO goal
Strain (1/√Hz)
Frequency (Hz)
Better seismic isolation
Better test masses and suspension
Higher power laser and thermal compensation

LIGO O1 and O2 runs

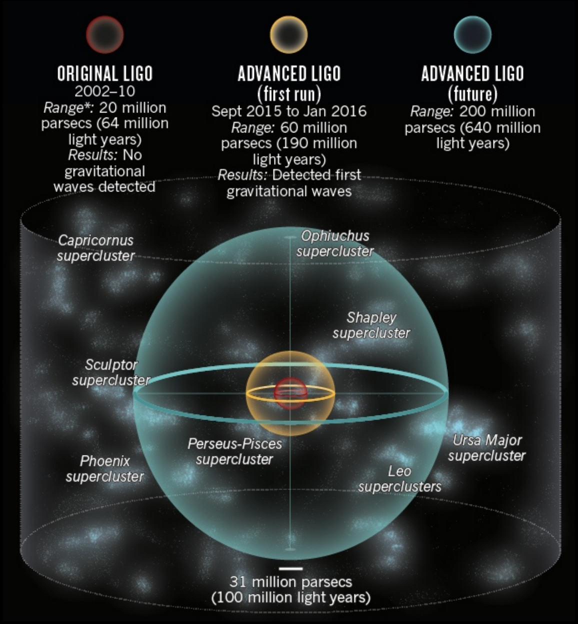

- ⋆ LIGO O1: September 12th 2015 - January 19th 2016,
- ⋆ LIGO O2 starting November 30th 2016, divided into two parts,
- ⋆ Virgo joining the O2b in summer 2017.

The Gravitational Wave Signal Tableau

Short duration ⟷ Long duration

Waveform known ⟷ Waveform unknown

Cosmic string cusp / kink

NS / BH ringdown

Low-mass inspiral

Asymmetric spinning NS

LISA binary

High-mass inspiral

Newborn NS Rotation-driven instability

Binary merger

Cosmological stochastic background

Stellar core collapse

Astrophysical stochastic background

???

???

???

courtesy of Peter Shawhan

Summary of Data Analysis Methods

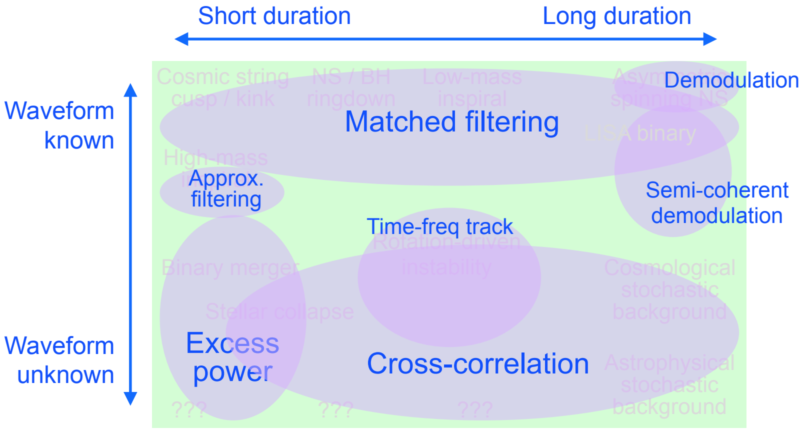

courtesy of Peter Shawhan

Neutron stars = very dense, magnetized stars

The most relativistic **material** objects in the Universe: compactness $M/R \simeq 0.5$, observed in all EM spectrum as pulsars, magnetars, in supernovæ remnants, in accreting systems, in double neutron star binaries...

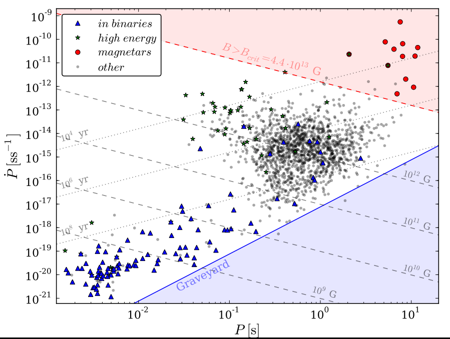

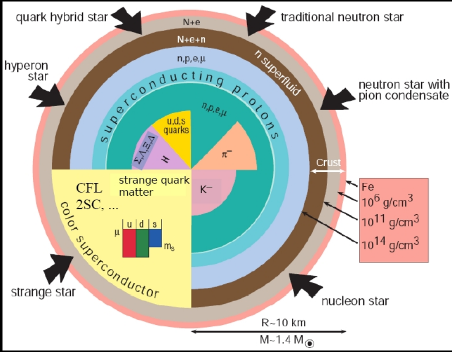

About 2500 NS observed to date, $10^8 - 10^9$ in the Galaxy.

Continuous GWs from spinning neutron stars

Characteristics:

- ⋆ Long-lived: $T > T_{obs}$,
- ⋆ Nearly periodic: $f_{GW} \propto \ f_{rot}$

Mechanisms that can create time-varying quadrupole moment:

- ⋆ "Mountains" (elastic and/or magnetic stresses, $f_{GW} = 2f_{rot}$),
- ⋆ Oscillations (r-modes, $f_{GW} = 4/3 f_{rot}$),
- ⋆ Free precession ($f_{GW} \propto f_{rot} + f_{prec}$)
- ⋆ Accretion (drives deformations from r-modes, thermal gradients, magnetic fields, $f_{GW} \simeq f_{rot}$)

(see PASA **32**, 34, 2015 for a review)

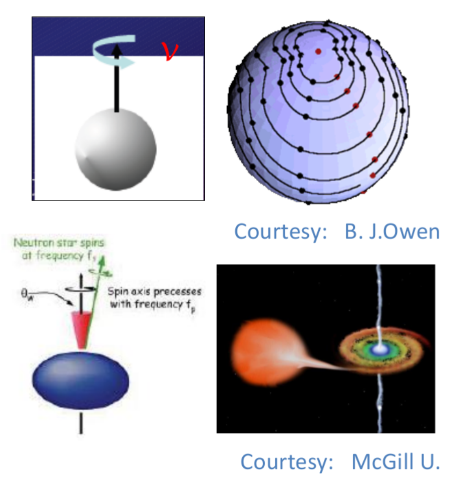

Example: a monochromatic signal

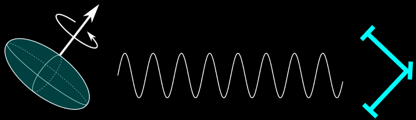

In this case a Fourier transform is sufficient to detect the signal (simplest matched filter method):

$$F = \left| \int_0^{T_0} x(t) \exp(-i\omega t) dt \right|^2$$

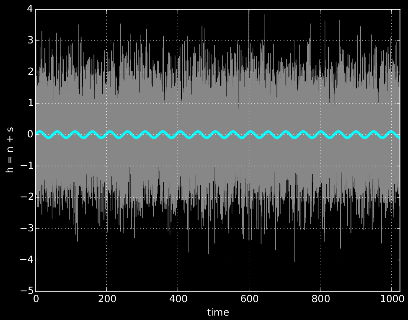

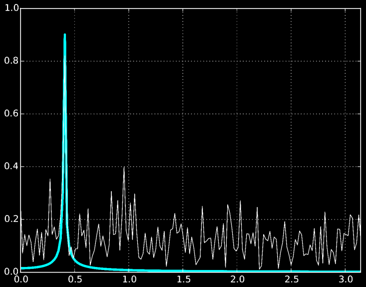

T_0 - time series duration, S_0 - spectral density of the data.

Signal-to-noise $SNR = h_0 \sqrt{\frac{T_0}{S_0}}$

In reality: signal is modulated

Since the detector is on Earth, planets and Earth's rotation influences signal's amplitude and phase.

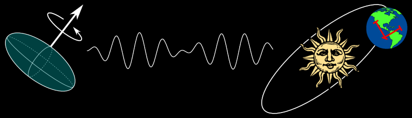

- ⋆ Signal is almost monochromatic: sources may slow down/spin up,
- ⋆ it has to demodulated (detector is moving),

→ precise ephemerides of the Solar System needed.

Detector movement distinguishes a real signal from detector's spectral artifacts ("lines").

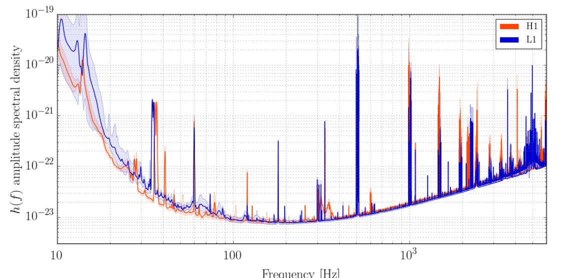

Example: the $\mathcal{F}$-statistic

$\mathcal{F}$-statistic estimates how well the amplitude and phase modulated model matches the data $x(t)$

$$\mathcal{F} = \frac{2}{S_0 T_0}\left(\frac{|F_a|^2}{\langle a^2\rangle} + \frac{|F_b|^2}{\langle b^2\rangle}\right)$$

where S_0 is the spectral density, T_0 is the observation time, and

$$F_a = \int_0^{T_0} x(t) a(t) \exp(-i\phi(t)) dt, \qquad F_b = \ldots$$

$a(t)$, $b(t)$ - amplitude modulation functions that depend on the sources' sky position $(\alpha,\ \delta)$,

$\phi(t)$ - phase modulation function that depends on $(f, \dot{f}, \alpha, \delta)$

(PRD **58**, 063001, 1998)

Taxonomy of continuous GW search methods

- Targeted searches
  - based on matched filtering (data of length T_0 correlated with signal templates). Position, f and $\dot{f}$, sometimes source orientation, are known.

$$h_0 \propto \sqrt{S(f)/T_0}$$

- Directed searches
  - Cases when some parameters are known, e.g. the position:

  $\rightarrow$ Supernovæ remnants, Sco X-1, the Galactic center, globular clusters etc.

- All-sky searches
  - Source parameters and position not known $\rightarrow$ parameter space is large $\rightarrow$ problem becomes computationally bound

  $\rightarrow$ *Hierarchical approach:* analysis of N data segments of length T_s coherently, combining the results incoherently (most sophisticated example: the Einstein@Home project)

$$h_0 \propto \sqrt{S(f)/T_s}/N^{1/4}$$

Example: computational cost for an all-sky search

In order to optimally cover a range of $(f, \dot{f}, \alpha, \delta)$ parameters,

$$\text{computing power} \propto \underbrace{T_0^2}_{\dot{f}} \times \underbrace{T_0^{[0-3]}}_{\alpha,\delta} \times \underbrace{T_0 \log(T_0)}_{f \text{ by FFT}} = T_0^{[3-6]} \log(T_0).$$

(see PRD **90**, 122010, 2014). Coherent search of $T_0 \simeq 1\,yr$ of data would require zettaFLOPS (10^{21} FLOPS) scale computers $\rightarrow$ currently impossible $\ddot{\frown}$

Solution: divide data into shorter length time frames ($T_s \simeq$ days)

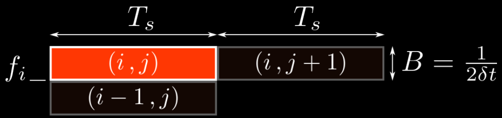

- ⋆ Perform a search in narrow frequency bands: sampling time $\delta t = 1/2B$, number of data points $N_p = T_s/\delta t = 2T_sB$

$\rightarrow$ feasible on a petaFLOP computer.

Second stage: look for coincidences between different T_s segments.

Third stage: Analyze interesting outliers ("targeted search").

POLGRAW all-sky search pipeline implementation

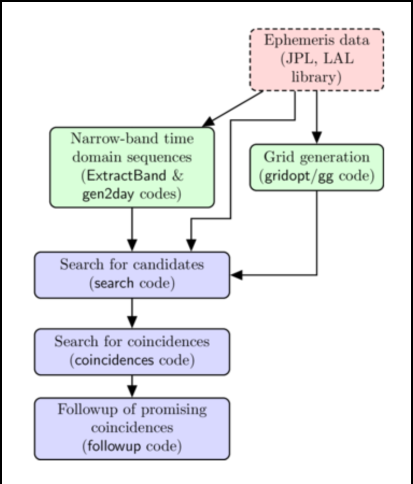

- ⋆ Stage 1. Search in segments of length T_s - a *lot* of FFTs, trigonometry.
- ⋆ Stage 2. Concidences between different segments - sorting large arrays of integers,
- ⋆ Stage 3. Followup of outliers - mostly additions and some trigonometry.

`OpenCL` version in progress, ported by Máté:
`https://github.com/mbejger/polgraw-allsky/tree/xplat`

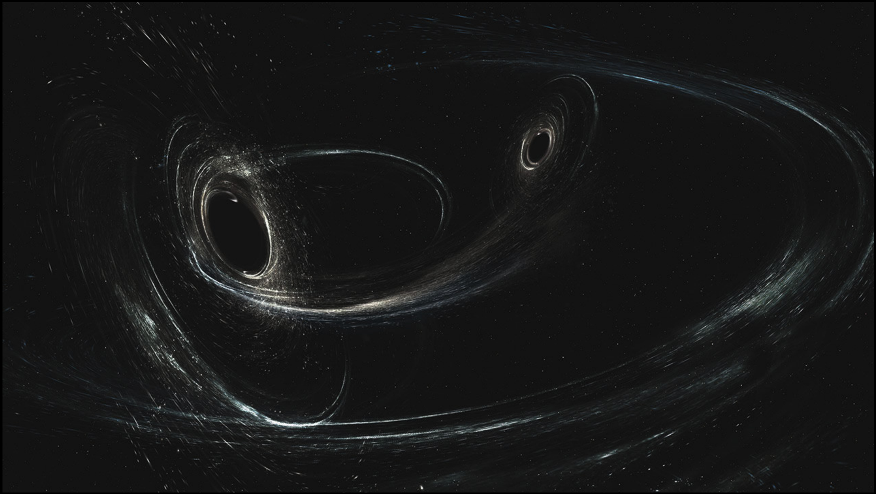

Exploiting the phenomenology of inspiral, merger, ringdown

- Post-Newtonian description of inspiral
  - Expansion of e.g. gravitational wave phase in powers of (v/c)
  - Do the coefficients depend on masses, spins as predicted by GR?
- Tidal effects during inspiral
  - “Black hole mimickers”: boson stars, dark matter stars, gravastars, ...
  - If less compact than neutron stars, can have large tidal effects
- Plunge and merger
  - Most dynamical regime
- Consistency between inspiral and post-inspiral regimes
- Ringdown
  - From the quasi-normal mode spectrum: (indirect) test of no-hair theorem
- Gravitational wave echoes
  - Quantum-modified black holes, exotic objects:
repeated bursts of GWs after ringdown
- Anomalous propagation of gravitational waves over large distances
  - Massive graviton, violations of local Lorentz invariance

(courtesy Chris Van Den Broek)

Source Parameters vs. Signal Parameters

- **Inspiral source parameters (19!)**
  - » Masses (m_1, m_2) → masses and spins are "intrinsic" parameters
  - » Spins $\mathbf{S}_1$, $\mathbf{S}_2$ → Negligible for neutron stars, at least
  - » Orbital phase at coalescence, φ → Maximize analytically when filtering
  - » Inclination of orbital plane, ι
  - » Sky location (α, δ)
    → Simply multiplicative for a given detector (long-wavelength limit)
  - » Distance d → Simply multiplicative
  - » Coalescence time t_c

Filter with orthogonal templates, take quadrature sum

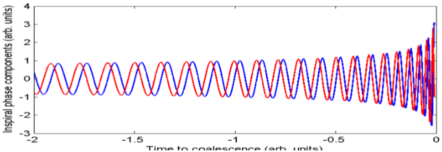

➔ Only have to explicitly search over masses and coalescence time ("intrinsic parameters")

(courtesy Alan Weinstein)

Matched-filtering

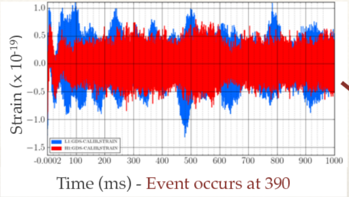

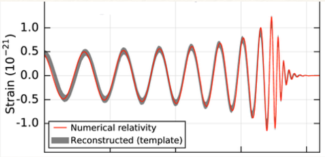

$$\langle s|h\rangle = \int_0^\infty \frac{\tilde{s}(f)\tilde{h}^*(f)}{S_n(f))} e^{-2i\pi t f}$$

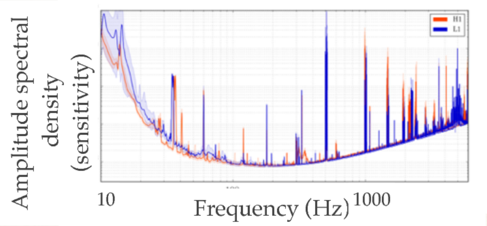

Allen et al. PRD 85 (2012) 122006

SB, …, IH, SP et al. PRD 87 (2013) 024003

(courtesy Ian Harry)

A large parameter space

- Non-eccentric BBH mergers are described by 15 parameters
  - We wish to cover the full space

Bayes Theorem

$$P(h(\hat{\lambda})|s) = \frac{P(h(\hat{\lambda}))P(s|h(\hat{\lambda}))}{P(s)}$$

$$\downarrow$$

$$P(h(\hat{\lambda})|s) = \frac{P(h(\hat{\lambda}))}{A}\left(\left\langle s|h(\hat{\lambda})\right\rangle - 0.5\left\langle h(\hat{\lambda})|h(\hat{\lambda})\right\rangle\right)$$

- For detection searches analytically maximise over phase, amplitude and time. Then loop over masses and spins.
- For parameter inference calculate the posterior explicitly.

(courtesy Ian Harry)

Template bank placement

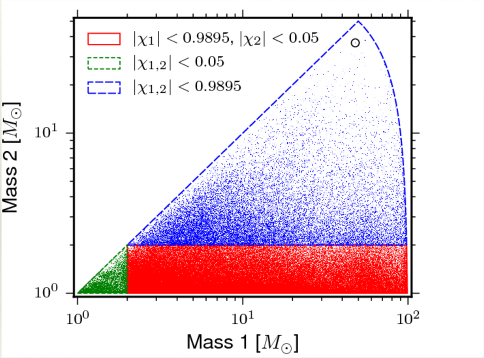

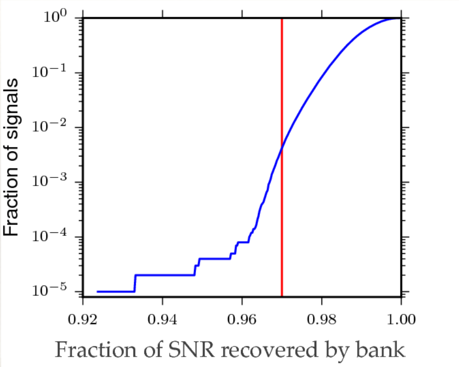

IH et al. PRD 80 (2009) 104014
Brown, IH et al. PRD 86 (2012) 084017
IH et al. PRD 89 (2014) 024010
Ajith, SP, et al. PRD 89 (2014) 084041
SP et al. PRD 89 (2014) 024003
Capano, IH, SP and AB 1602.03509

(courtesy Ian Harry)

Detection Statistic

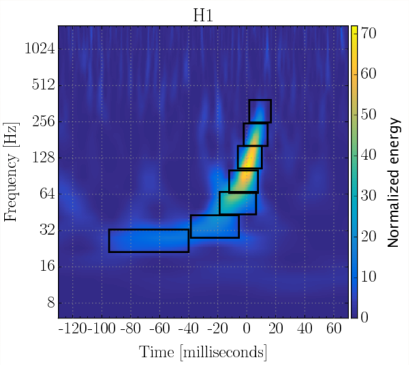

Real signal

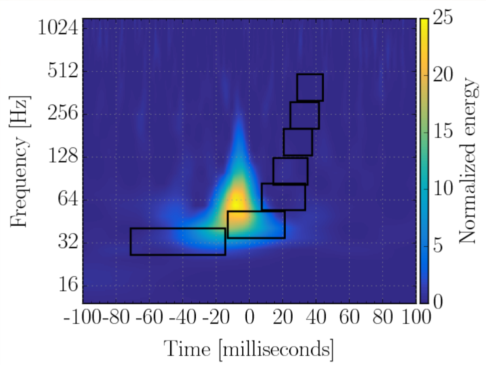

Instrumental artifact

Allen PRD 71 (2005) 062001
SB, …, IH, SP et al. PRD 87 (2013) 024003

(courtesy Ian Harry)

Sources of transient noise

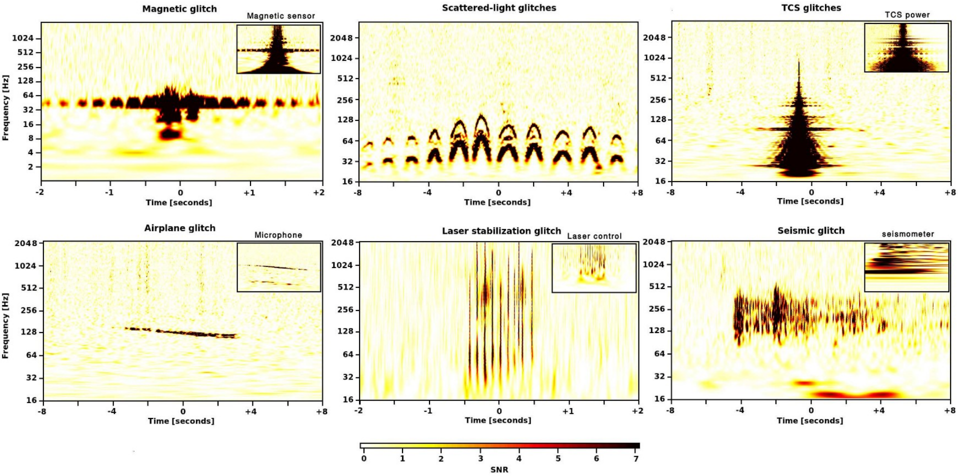

The sensitivity of transient searches is limited by transients:"glitches". (courtesy Florent Robinet)

Interferometer Layout: Beams and Benches

Several beams are extracted from the detector to measure the position and orientation of the mirrors, the lengths and alignments of the optical cavities, etc.

The photodiodes that measure the laser power in these beams are kept on various in-vacuum optical benches.

B1: from the antisymmetric port of the detector, measured on Suspended Detection Bench 2 (SDB2). This is where the GW signal is measured.

B2: from the symmetric port of the detector (also called the reflected port), the beam reflected by the PR. from Suspended Injection Bench 2 (SIB2).

B4: a small pickoff (180ppm) of the light circulating inside the Power Recycling cavity (PRC). From SPRB.

B7, B8: light transmitted by the north & west arms, on SNEB/SWEB.

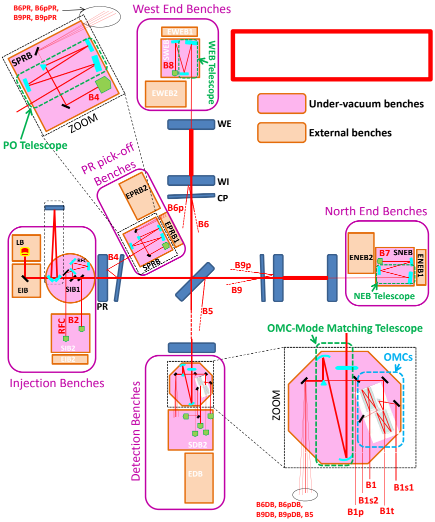

(courtesy Daniel Hoak)

Q transform (Omicron)

Q transform = SFT with a Gaussian window (window size ~ 1/f)

$$X(\tau, \phi, Q) = \int_{-\infty}^{+\infty} x(t) w(t - \tau, \phi, Q) e^{-2i\pi\phi t} dt.$$

→ Tiling: Q planes with frequency rows (log) subdivided in time tiles (linear)

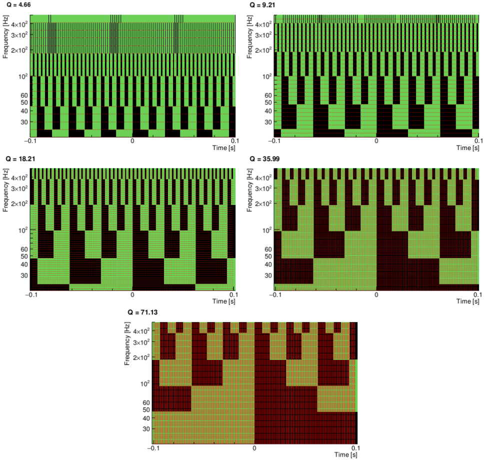

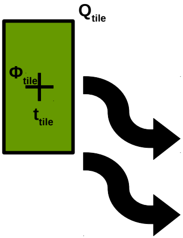

Q transform coefficient
X_{tile}

Use of whitened data
→ SNR
SNR_{tile}

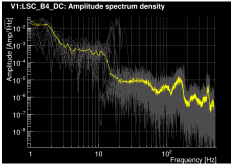

Noise PSD is estimated
→ whitening

Omicron maps

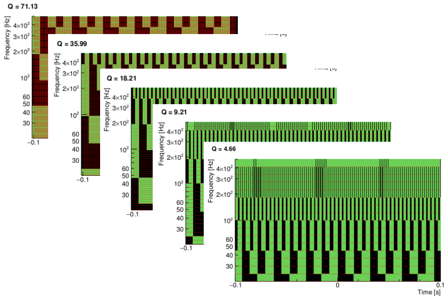

The Q transform is computed for every tiles → SNR → maps

Maps are stacked up, tiles with highest SNR values are displayed on top

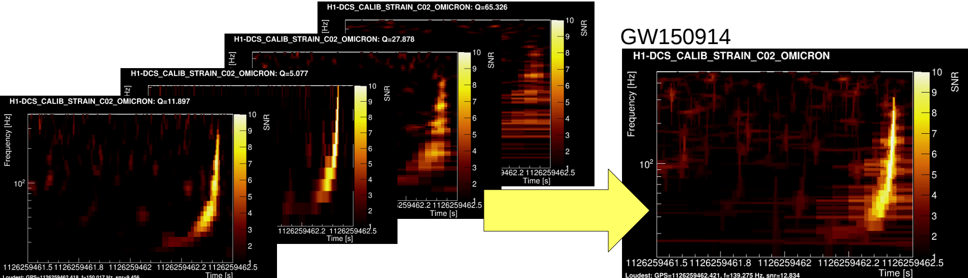

Noise characterisation - knowledge of the clusters

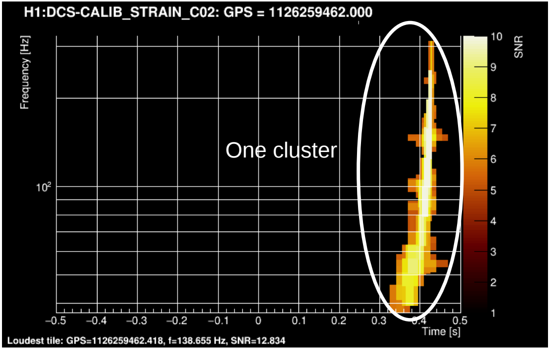

Triggers are clustered in time

One cluster:
– time
– frequency
– Q value
– SNR
→ given by the tile with the highest SNR

+ duration, bandwidth...

In addition to speeding up the production of spectrograms, clasification of various types glitches seems like a good job for GPUs. (Omicron slides: courtesy Florent Robinet)

Calculating a significance (how many sigmas?)

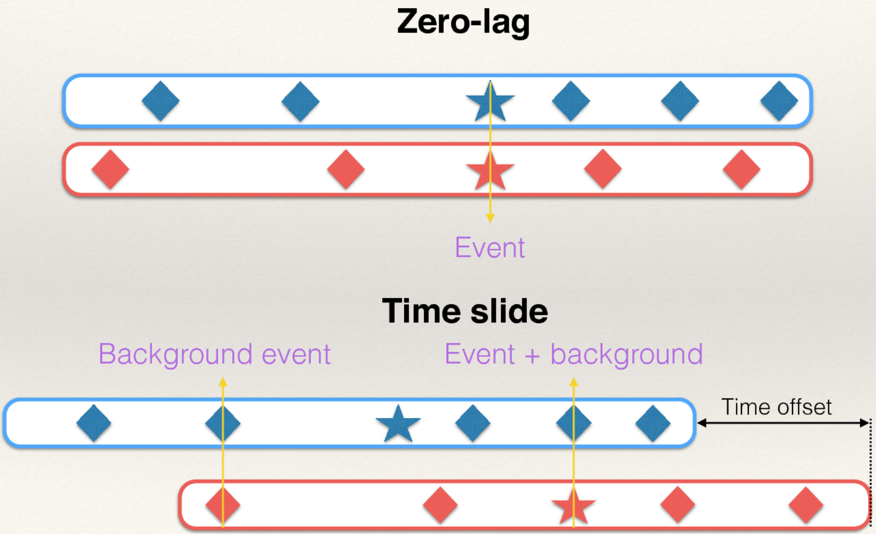

(courtesy Ian Harry)

Computational cost of O1 CBC search

- 249,077 template waveforms
- 79.04 days (Hanford) and 67.35 days (Livingston) "analysis ready" data in 130 calendar day window (x2 observatories). 51.5 days coincident.
- 1 core (hyper-threaded) can filter 5,500 templates (256s of data in 256s of time)
- 110,000 CPU hours to

  O(11M) CPU hours

  - x 2 search pipelines
  - x O(10) "simulation campaign" runs for rate statements and testing
  - x Reruns
- Parameter estimation on GW150914 and LVT151012 (O(1M) CPU hours))

(courtesy Ian Harry)

Current usage of LDG since O1 start

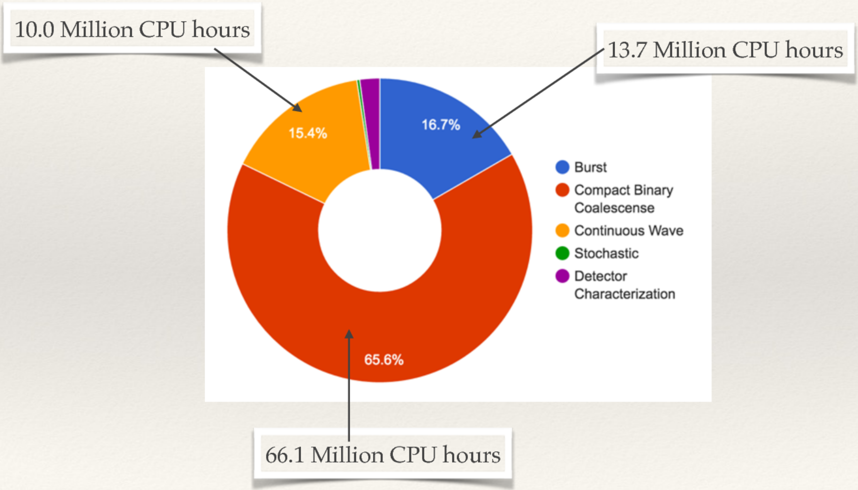

(courtesy Ian Harry)