Loosely coherent algorithms - robust and computationally efficient search of large parameter spaces.

Vladimir Dergachev
(Caltech)

Caltech, December 12th 2011

DCC: LIGO-G1101278-v2

Observation of a single target

- Detection of signals from a single target is relatively straightforward:
  - Record data
  - Compute correlation with known waveform
  - Analyze significance
- This procedure is not computationally limited, so we can rely on well-known statistical techniques.

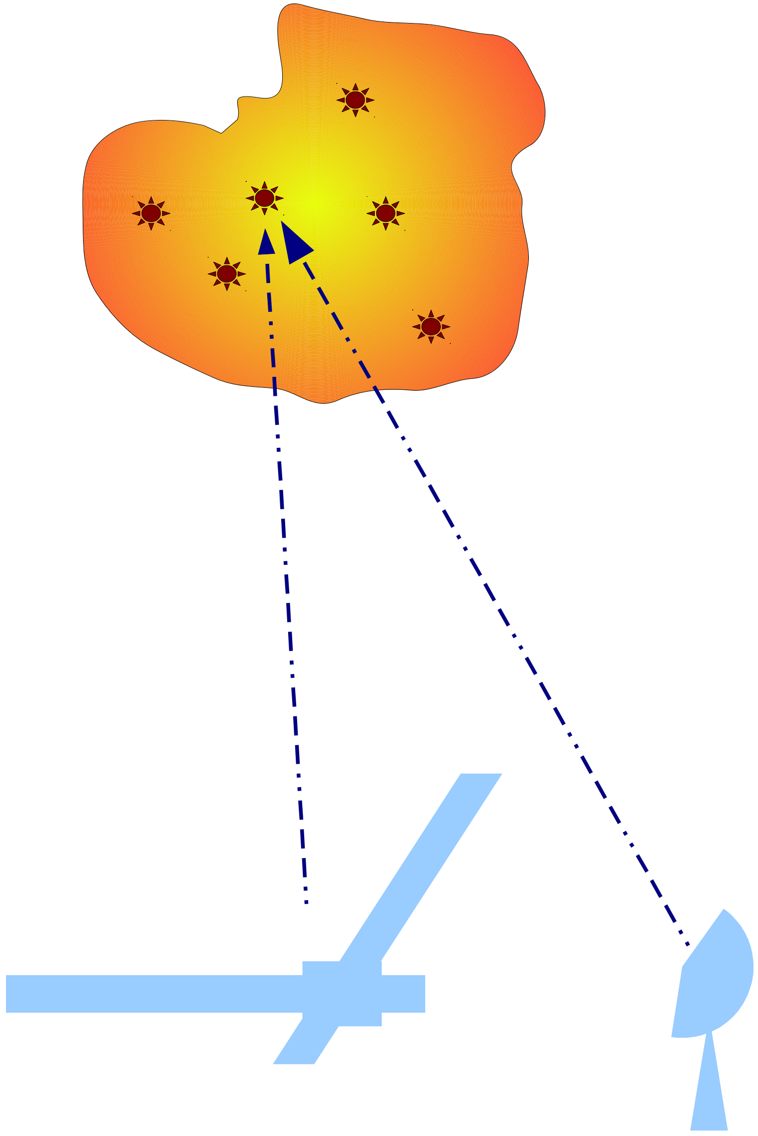

Observation of multiple targets

- Multiple targets are slightly trickier:
  - Record data
  - Compute correlation with known waveform
  - Analyze significance
  - Repeat for all targets of interest
- Computing power starts to scale, need to pay attention to trials factor

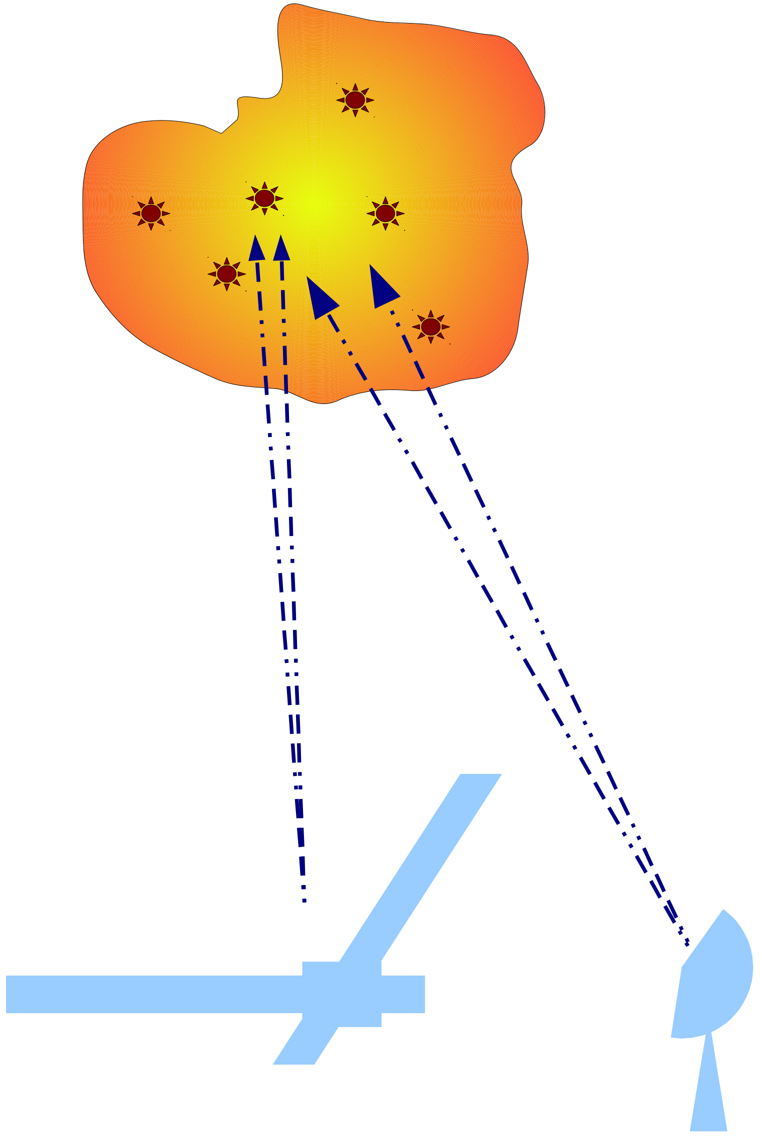

Blind search

- Blind search requires large computing power:
  - Record data
  - Compute correlation with known waveform
  - Analyze significance
  - Sweep large area of parameter space
- Pure noise triggers events at 5σ level, need to pay close attention to computational efficiency

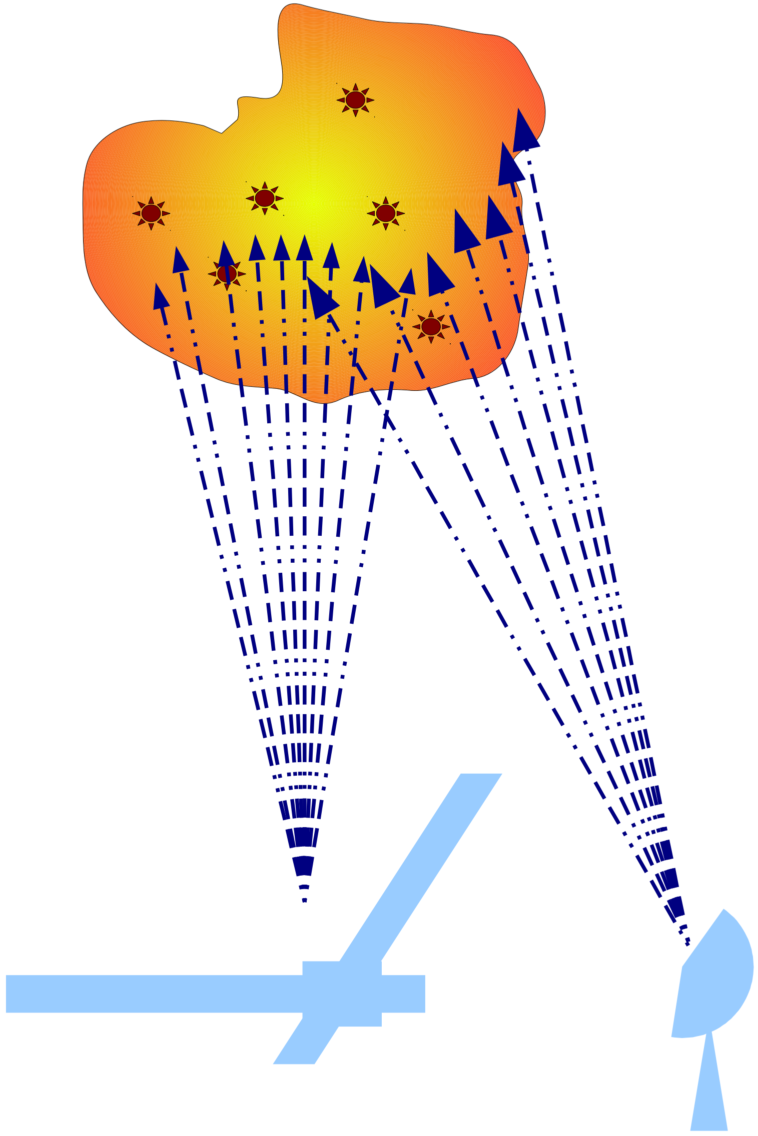

Blind search of large datasets

- Blind search of large datasets:
  - Record data
  - Use carefully chosen algorithm for signal detection
  - Analyze significance
  - Sweep large area of parameter space
- Pure Gaussian noise triggers events at 6σ level
- Real data produces 7σ artifacts
- Signal template computation becomes challenging
- Computation efficiency becomes as important as statistical efficiency.

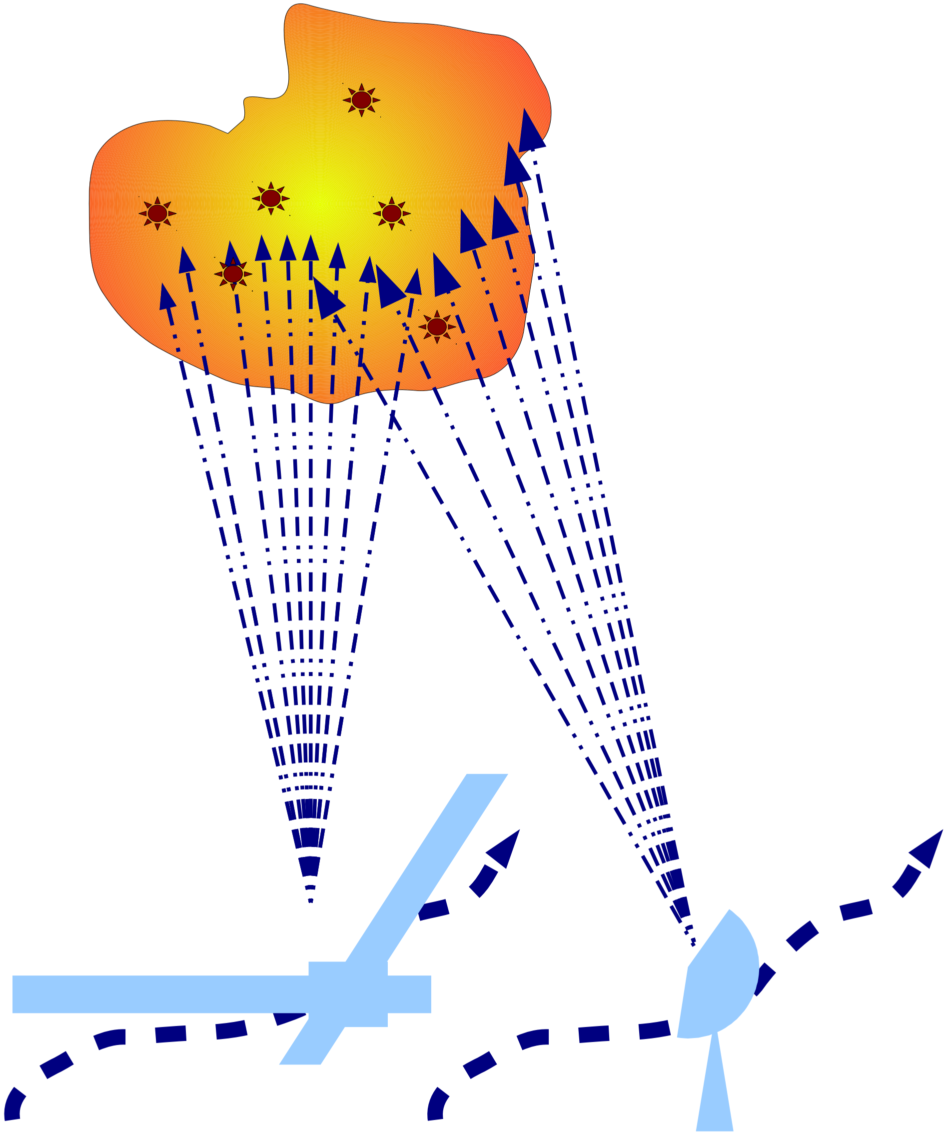

Computational power

- In the near future we can expect access to clusters of ~1000 CPUs delivering ~1 TFlop each
- One year of run time provides $3 \cdot 10^{22}$=0.05 mol of floating point operations.
- How can we use this computational power wisely ?
  - For weak signals we really have to look at each template. Figure of merit: cycles per template.
  - Storing (even in RAM) is expensive. Aggregate data for multiple templates.

Loosely coherent search

- Analyze data for a set of nearby templates at once.
- Report cumulative statistics for the entire set:
  - Is there a signal ?
  - Data quality ?
- Interesting mathematical problem. Best algorithm depends on
  - set of signal templates
  - collected statistics
  - computing hardware

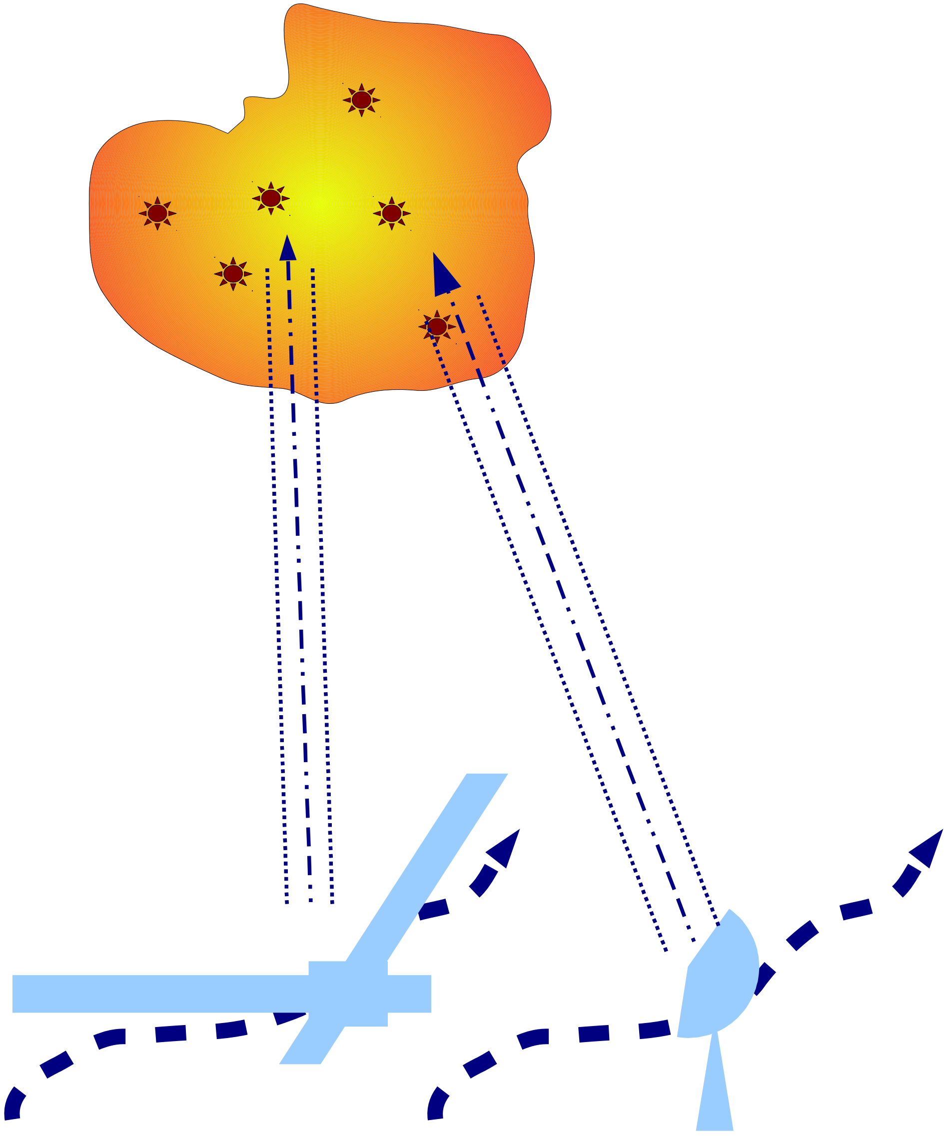

V Dergachev 2010 Class. Quantum Grav. 27 205017

Example: search for CW gravitational signals

• Searches for continuous gravitational wave signals typically use 6-24 months of data (2 TB)
• We search millions of sky templates, millions of frequency templates and thousands of auxiliary parameters. Computationally limited.
• A modified version can be used in radio astronomy, radar.

Semi-coherent version runs on Einstein@Home →

	Cycles per template
ComputeFStatistic (plain matched filter)	~ 1000000
Resampling	~ 20000
Loosely coherent $\delta=0$	<1400

Loosely coherent searches for sets of well-modeled signals
arXiv:1110.3297

Loosely coherent algorithm can cover $2\cdot 10^{19}$ templates using 1 year of petaflop cluster CPU time. Efficiency increases with dimension of parameter space.

Some *Loosely Coherent* ideas

- Consider a subspace containing a single set of nearby templates.
- The subspace is usually hard to describe exactly, but one can construct an over-determined basis that covers it.
- Evaluate groups of templates as a single linear operator in this basis (for example a convolution)
- Speed of operator evaluation determines efficiency of the search.

Loosely coherent searches for sets of well-modeled signals
arXiv:1110.3297

Summary

- Computationally limited searches of large parameter spaces require new algorithms that balance statistical efficiency with computational efficiency.
- Loosely coherent algorithms iterate over *sets* of nearby templates, gaining efficiency as they scale.
- The more parameters, the better !

End of talk

(supporting slides for questions follow)

LIGO detectors

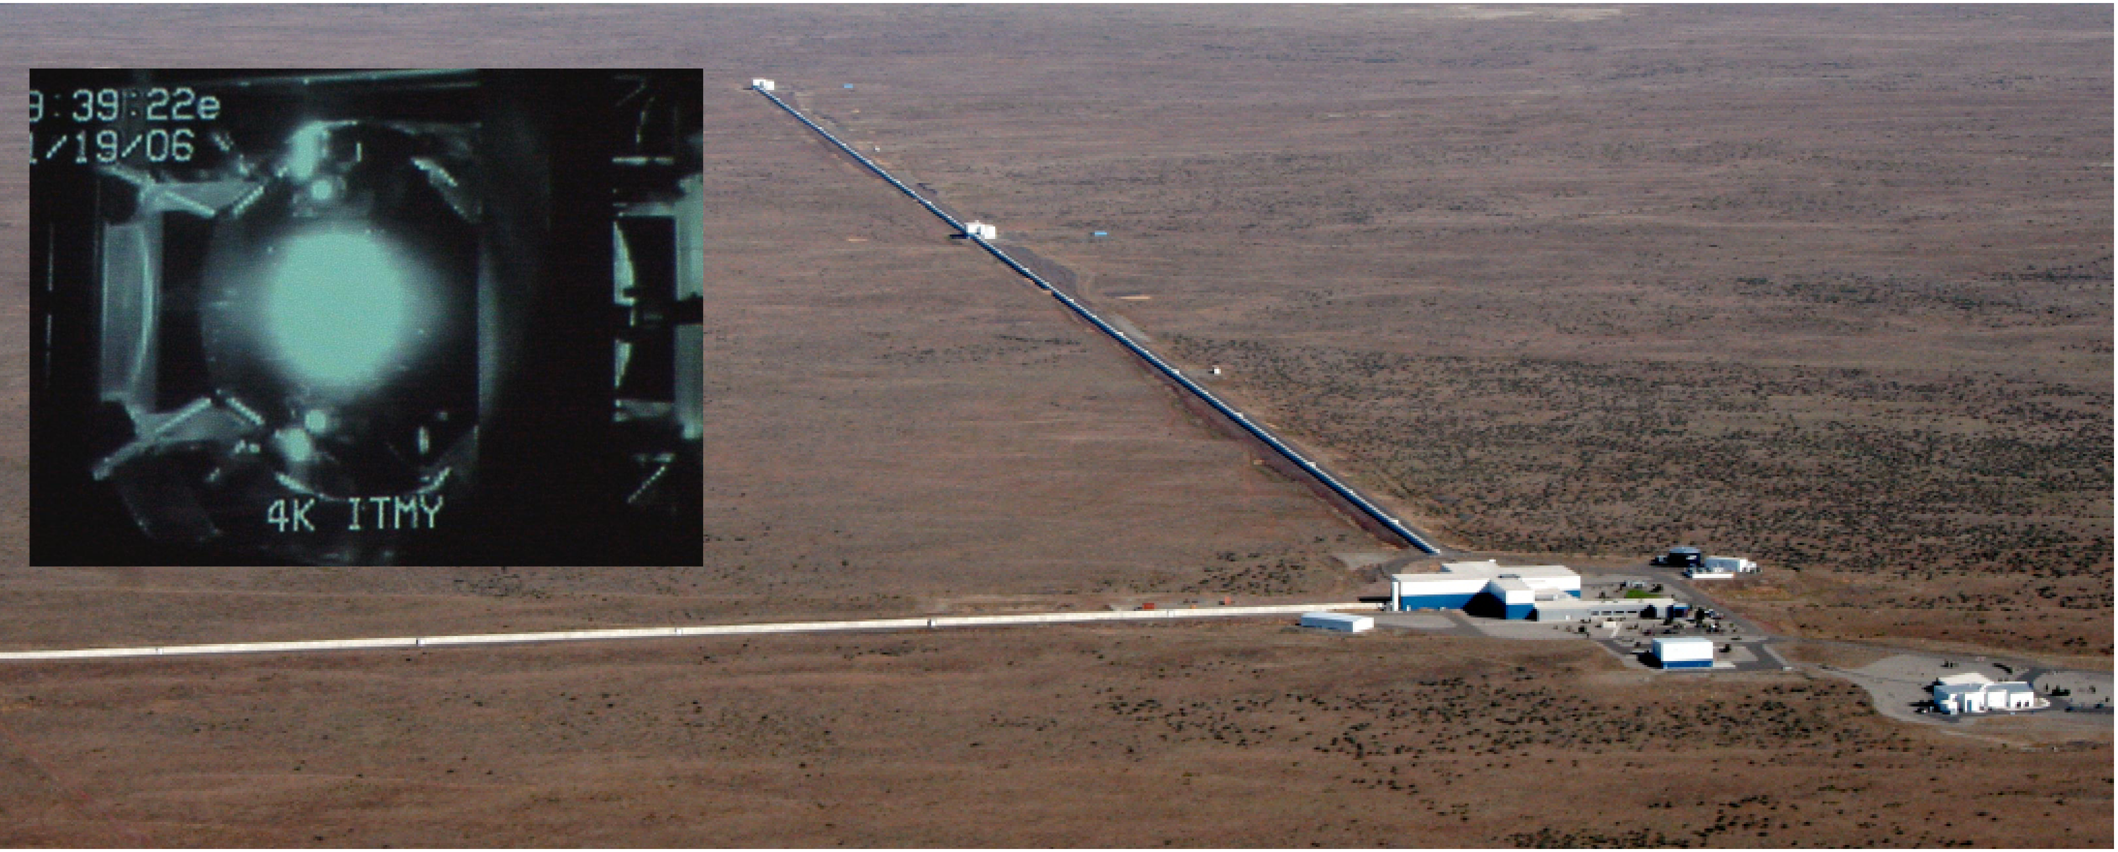

LIGO Hanford observatory

- 4km long vacuum tubes with ~100kW laser beams inside (2010)
- The detector can measure relative displacement between mirrors at the end of the arms with precision of 10^{-19} m/sqrt(Hz)

Continuous gravitational waves

Rotating neutron star

- We know of many rotating neutron stars with frequencies from below 1 Hz to more than 700 Hz
- Gravitational radiation is expected to be emitted at twice the frequency
- Not all rotating neutron stars have to emit radio waves or X rays
- Are any convenient sources nearby ?

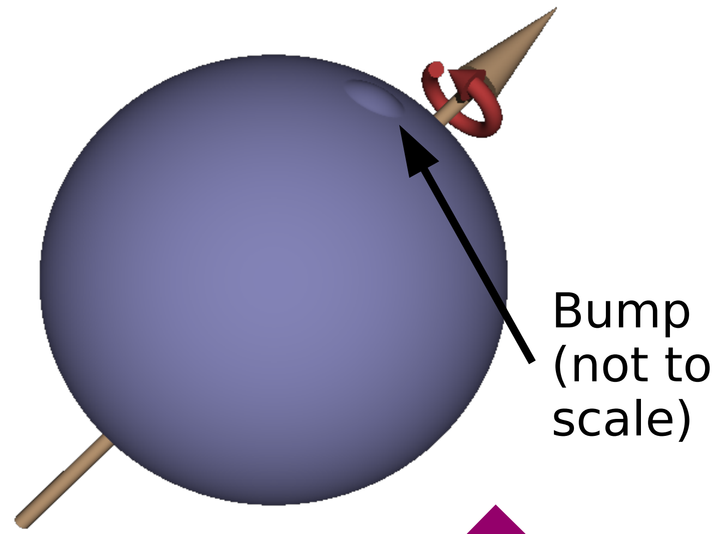

Linearly polarized gravitational waves

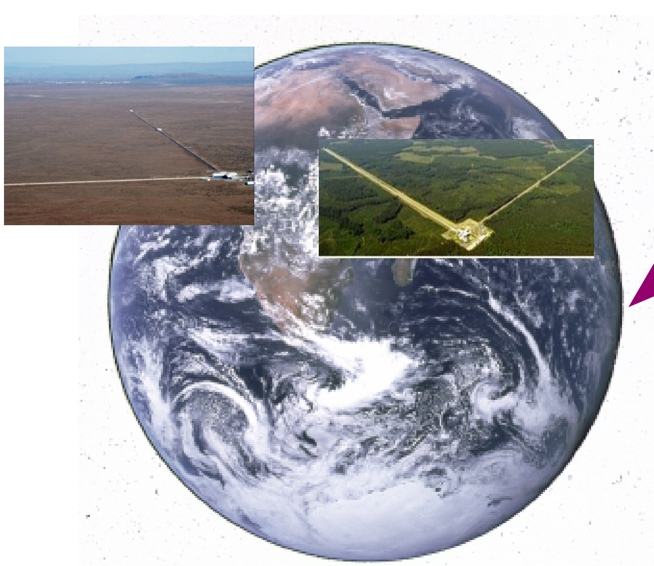

Circularly polarized gravitational waves

Detection methods

Semi-coherent

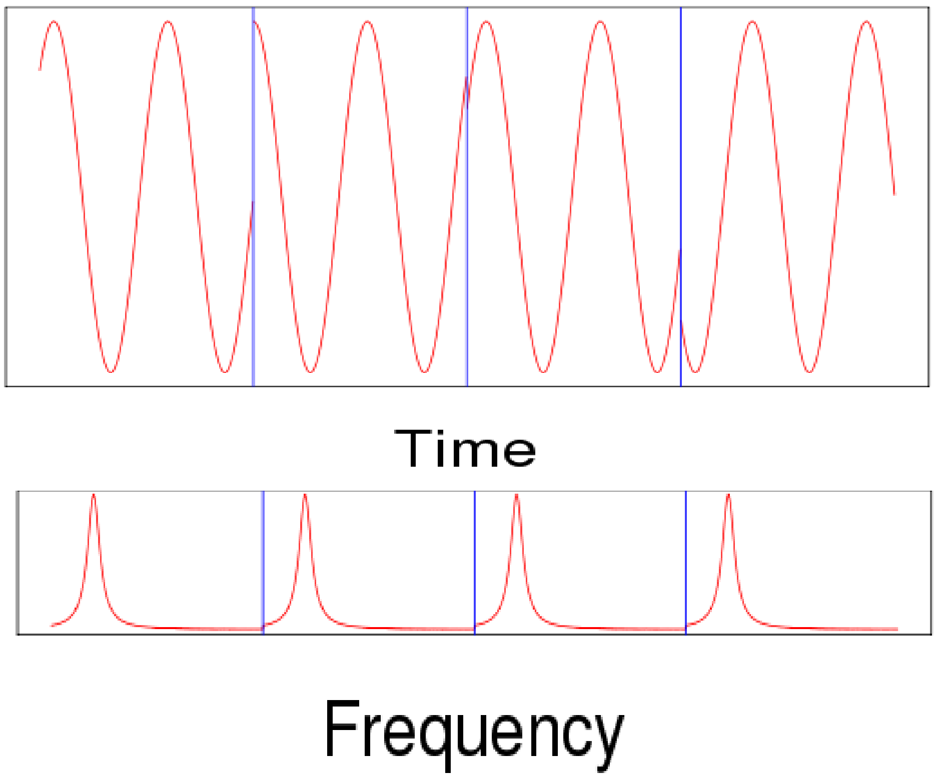

- Chop data into equal size (30 min) chunks
- Sum powers from all chunks
- Ignores phase information between chunks
- Sensitivity scales $1/T^{0.25}$
- CPU cycles scale as T^4

Coherent

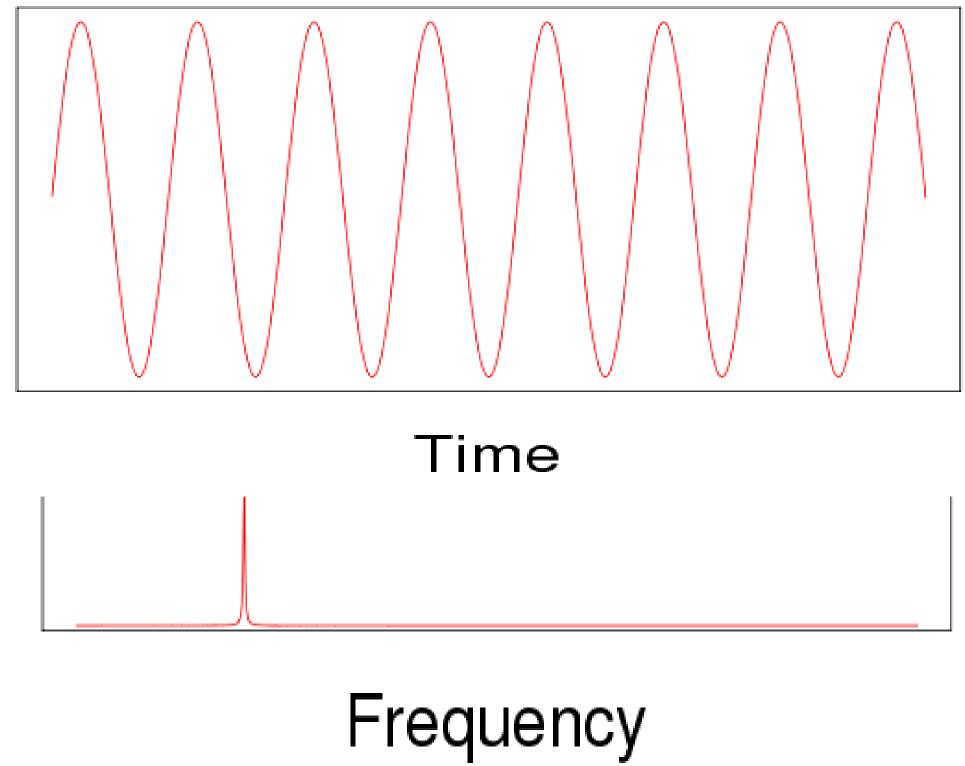

- Use matched filter to achieve high sensitivity
- Need to know exact signal form
- Sensitivity scales as $1/T^{0.5}$
- CPU cycles scale as T^6 or faster

Loosely coherent search

- A loosely coherent code computes a statistic indicating whether a signal from a given family S is present in the data.
- This is in contrast to traditional codes that focus on detecting a particular signal corresponding to a fixed template.
- Ideally, the statistic is easy to compute and falls off sharply outside S.
- A prototype large δ (and large S) implementation used in Full S5 analysis.

Loosely coherent search

Semi-coherent

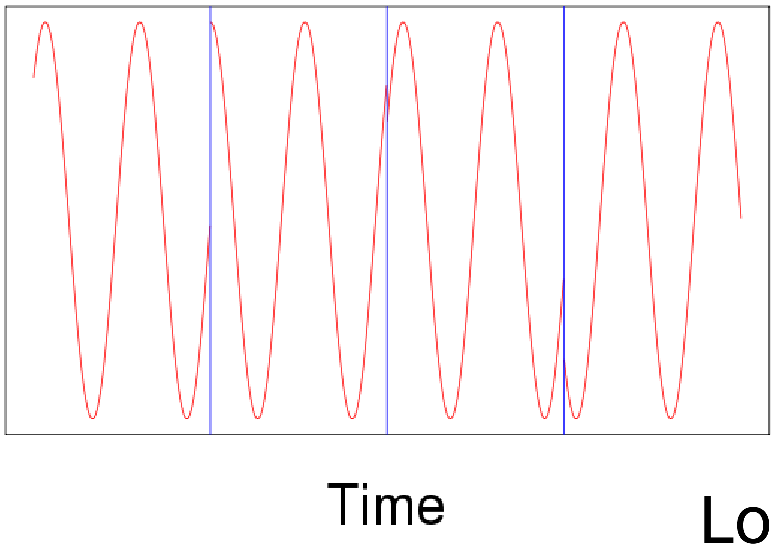

Coherent

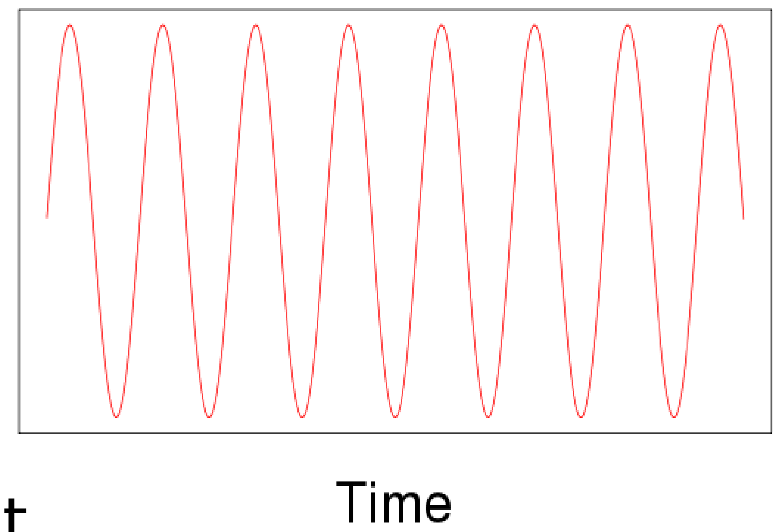

Loosely coherent

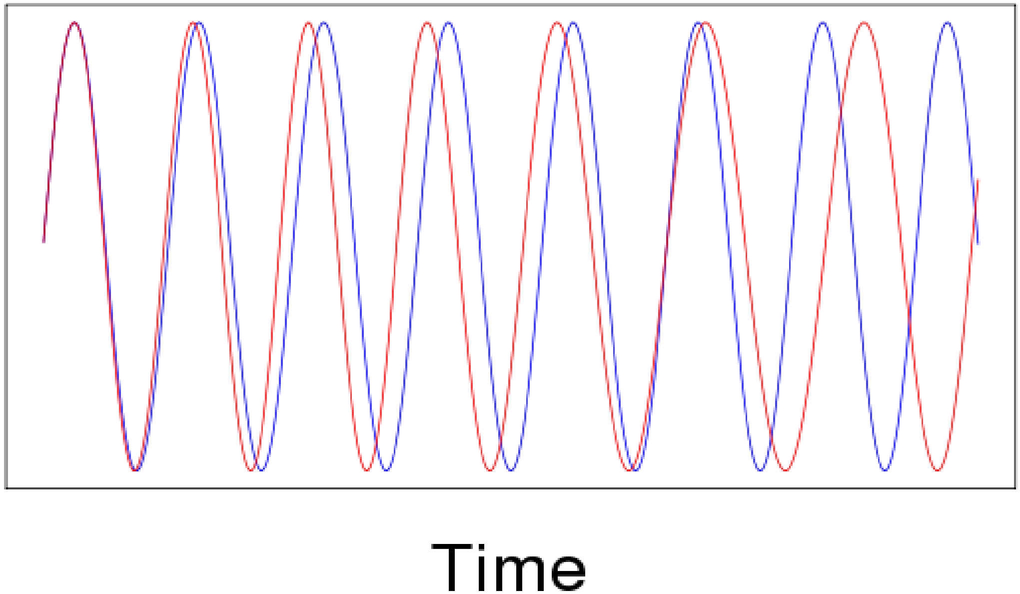

How do we implement a long coherence length loosely coherent search ?
(a simplified description)

- We will have to compute statistics similar to matched filter
- We will need to do it fast
- We will need to do it for a lot of waveforms

Signal model

Suppose we are looking for a signal $s(t,a)$ of unknown amplitude A and frequency ν:

$$s(t,a) = Ae^{2\pi i\nu(t+\Xi(t,a))}$$

Assume that phase modulation term $\Xi(t,a)$ varies slowly with t and a.

$$\Xi(t,a) = \text{Earth motion} + \text{Source evolution}$$

$$\frac{\partial \Xi(t,a)}{\partial t} \sim \text{Spindown} + \text{Doppler shift}$$

Conventional matched filter

Our data $f(t)$ is a sum of some unknown signal plus uncorrelated noise $\xi(t)$:

$$s(t, a_0) = Ae^{2\pi i\nu(t+\Xi(t,a_0))}$$

$$f(t) = s(t, a_0) + \xi(t)$$

$$\langle \xi(t)\xi(t') \rangle = \sigma^2(t)\delta(t - t')$$

Then the matched filter for $a=a_0$ is an optimal way to detect this signal:

$$A(\nu, a) = \frac{1}{\mathfrak{W}} \int_{t_0}^{t_1} f(t) e^{-2\pi i\nu(t+\Xi(t,a))} \frac{dt}{\sigma^2(t)}$$

If a_0 is unknown we simply iterate over a set of parameters a . This is not quite optimal, but close enough.

$$\mathfrak{W} = \int_{t_0}^{t_1} \frac{dt}{\sigma^2(t)}$$

Computational issues.

$$A(\nu, a) = \frac{1}{\mathfrak{W}} \int_{t_0}^{t_1} f(t) e^{-2\pi i \nu (t + \Xi(t,a))} \frac{dt}{\sigma^2(t)}$$

Computing $A(\nu,a)$ directly is rather expensive – not only one needs to compute the integral, but $\Xi(t,a)$ is a complicated function. This can be done by limiting the timebase $t_1 - t_0$ and by applying sufficient computational power, such as available to Einstein@Home.

Resampling

A large speedup can be achieved by noticing that the expression for $A(\nu, a)$ is almost a Fourier transform:

$$A(\nu, a) = \frac{1}{\mathfrak{W}} \int_{t_0}^{t_1} f(t) e^{-2\pi i \nu (t + \Xi(t,a))} \frac{dt}{\sigma^2(t)}$$

We can rewrite it to look exactly as a Fourier transform by redefining time to “straighten it out”:

$$t' = t + \Xi(t, a)$$

$$A(\nu, a) = \frac{1}{\mathfrak{W}} \int_{t_0}^{t_1} \frac{f(t')}{\sigma^2(t')(1 + \Xi_t(t', a))} e^{-2\pi i \nu t'} dt'$$

For practical computation with a DFT, this requires *resampling* $f(t)$ to a regular grid in t' variable.

Stepping in the sky

$$A(\nu, a) = \frac{1}{\mathfrak{W}} \int_{t_0}^{t_1} \frac{f(t')}{\sigma^2(t')(1 + \Xi_t(t', a))} e^{-2\pi i \nu t'} dt'$$

Resampling computes $A(\nu, a)$ for a series of frequency bins. It is reasonable to expect that nearby values of parameter a can be computed by convolving with a suitable kernel:

$$A(\nu, a) = \int A(\mu, a_0) K(\nu, \mu, a) d\mu$$

This was mentioned as a stepping in the sky method by B.F.Schutz in “The detection of gravitational waves” - edited by D.G.Blair. As far as we know no implementation was ever made.

Loosely coherent search

A *loosely coherent* search deals with families of signals, rather than a single waveform. So we consider a patch of parameter space near a signal with parameters ν_0, a_0.

It is too difficult to compute $A(\nu, a)$ from raw $f(t)$ for all points in this patch, so we need to transform our $f(t)$ into a more convenient basis.

As our signals are almost sine waves, a Fourier transform will do:

$$F(\lambda; \nu_0, a_0) = \frac{1}{\mathfrak{W}} \int_{t_0}^{t_1} \frac{f(t)}{\sigma^2(t)} e^{-2\pi i \nu_0 (t + \Xi(t, a_0))} e^{-2\pi i \lambda t} dt$$

$F(\nu_0; \nu_0, a_0)$ is exactly the matched filter $A(\nu_0, a_0)$, but for $\lambda \neq \nu_0$ it returns a distorted version of it.

Loosely coherent search

The distortion can be removed with a (pseudo) convolution:

$$A(\nu, a) = \frac{1}{\mathfrak{W}} \int_{t_0}^{t_1} f(t) e^{-2\pi i \nu (t + \Xi(t,a))} \frac{1}{\sigma^2(t)} dt =$$

$$= \frac{1}{\mathfrak{W}} \int_{t_0}^{t_1} f(t) e^{-2\pi i \nu_0 (t + \Xi(t,a_0))} \frac{1}{\sigma^2(t)} \cdot$$

$$\cdot e^{-2\pi i (\nu - \nu_0) t - 2\pi i (\nu - \nu_0) \Xi(t,a_0) - 2\pi i \nu (\Xi(t,a) - \Xi(t,a_0))} dt =$$

$$= \int F(\nu - \nu_0 - \mu; \nu_0, a_0) \cdot$$

$$\cdot \int_{t_0}^{t_1} e^{-2\pi i \mu t - 2\pi i (\nu - \nu_0) \Xi(t,a_0) - 2\pi i \nu (\Xi(t,a) - \Xi(t,a_0))} dt d\mu$$

Practical implementation

- Actual signals are more complicated - need to keep track of polarization.
- Need to control the number of terms in the convolutions – inner-most loop uses 11.
- Need to compute statistics – it is hard to compute robust statistics with a few dozen operations to be comparable with the cost of computing $A(\nu, a)$

	Cycles per template
ComputeFStatistic (plain matched filter)	~ 1000000
Resampling	~ 20000
Loosely coherent δ=0	<1400

Advantages

- An efficient code for analysis of “blob” regions on the sky.
- Easy to optimize with vector arithmetic, GPUs, FPGAs, etc.
- Good sensitivity with only 1.5 months of data – opportunity for early results when Advanced LIGO comes online.

Code performance on simulated data.

Monte-Carlo run on 1.5 months of Gaussian data

- Correct reconstruction of injections
- Sensitivity of 1.5-month coherent code
- Injections have uniformly distributed polarizations and sky locations

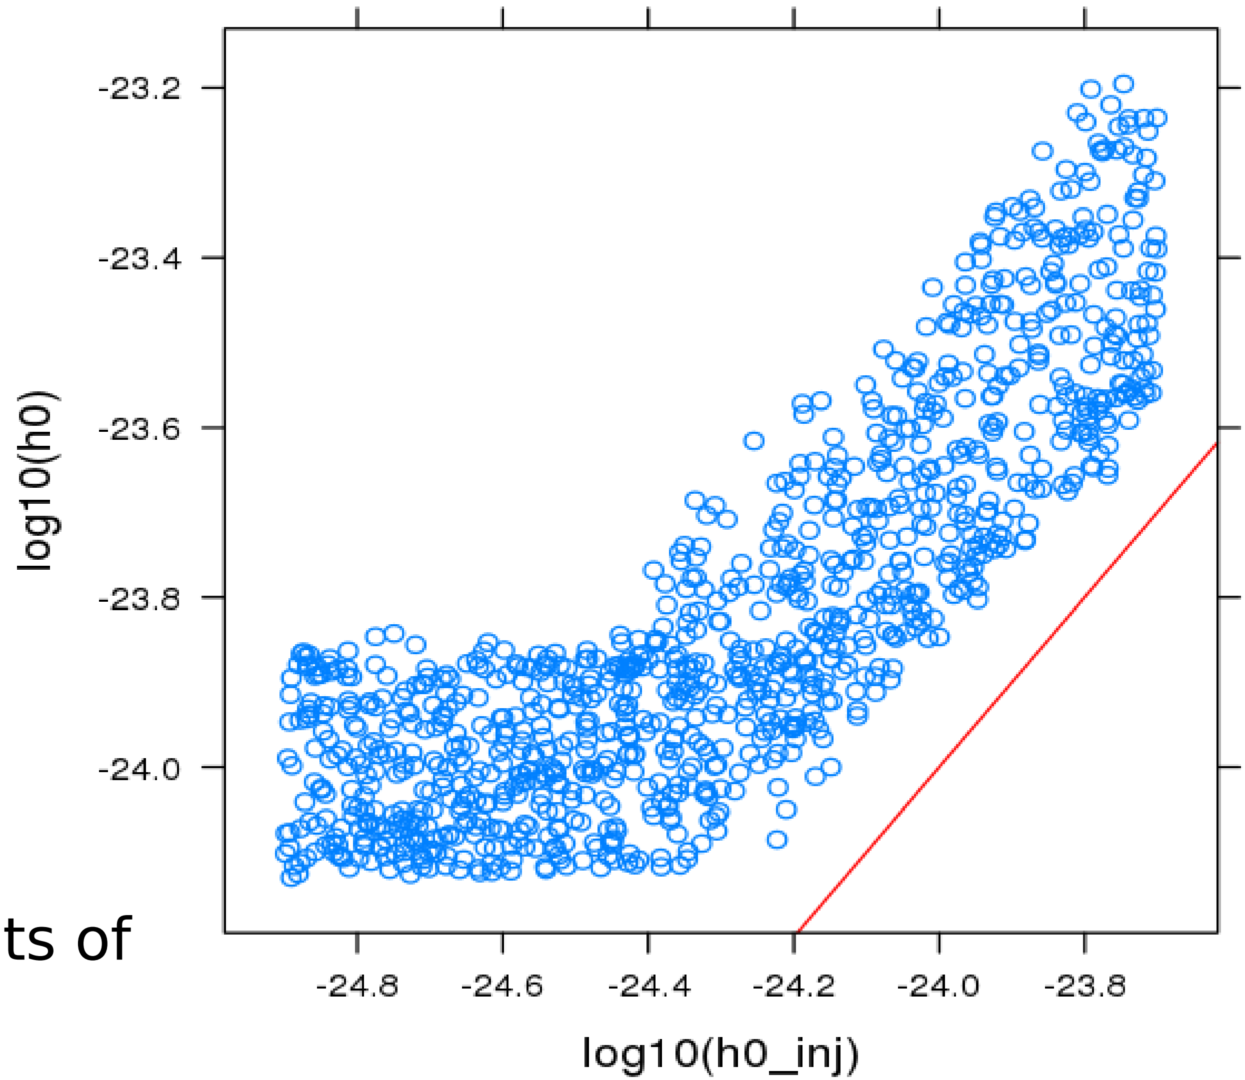

Paper in DCC: Loosely coherent searches for sets of well-modelled signals.

Frequency reconstruction

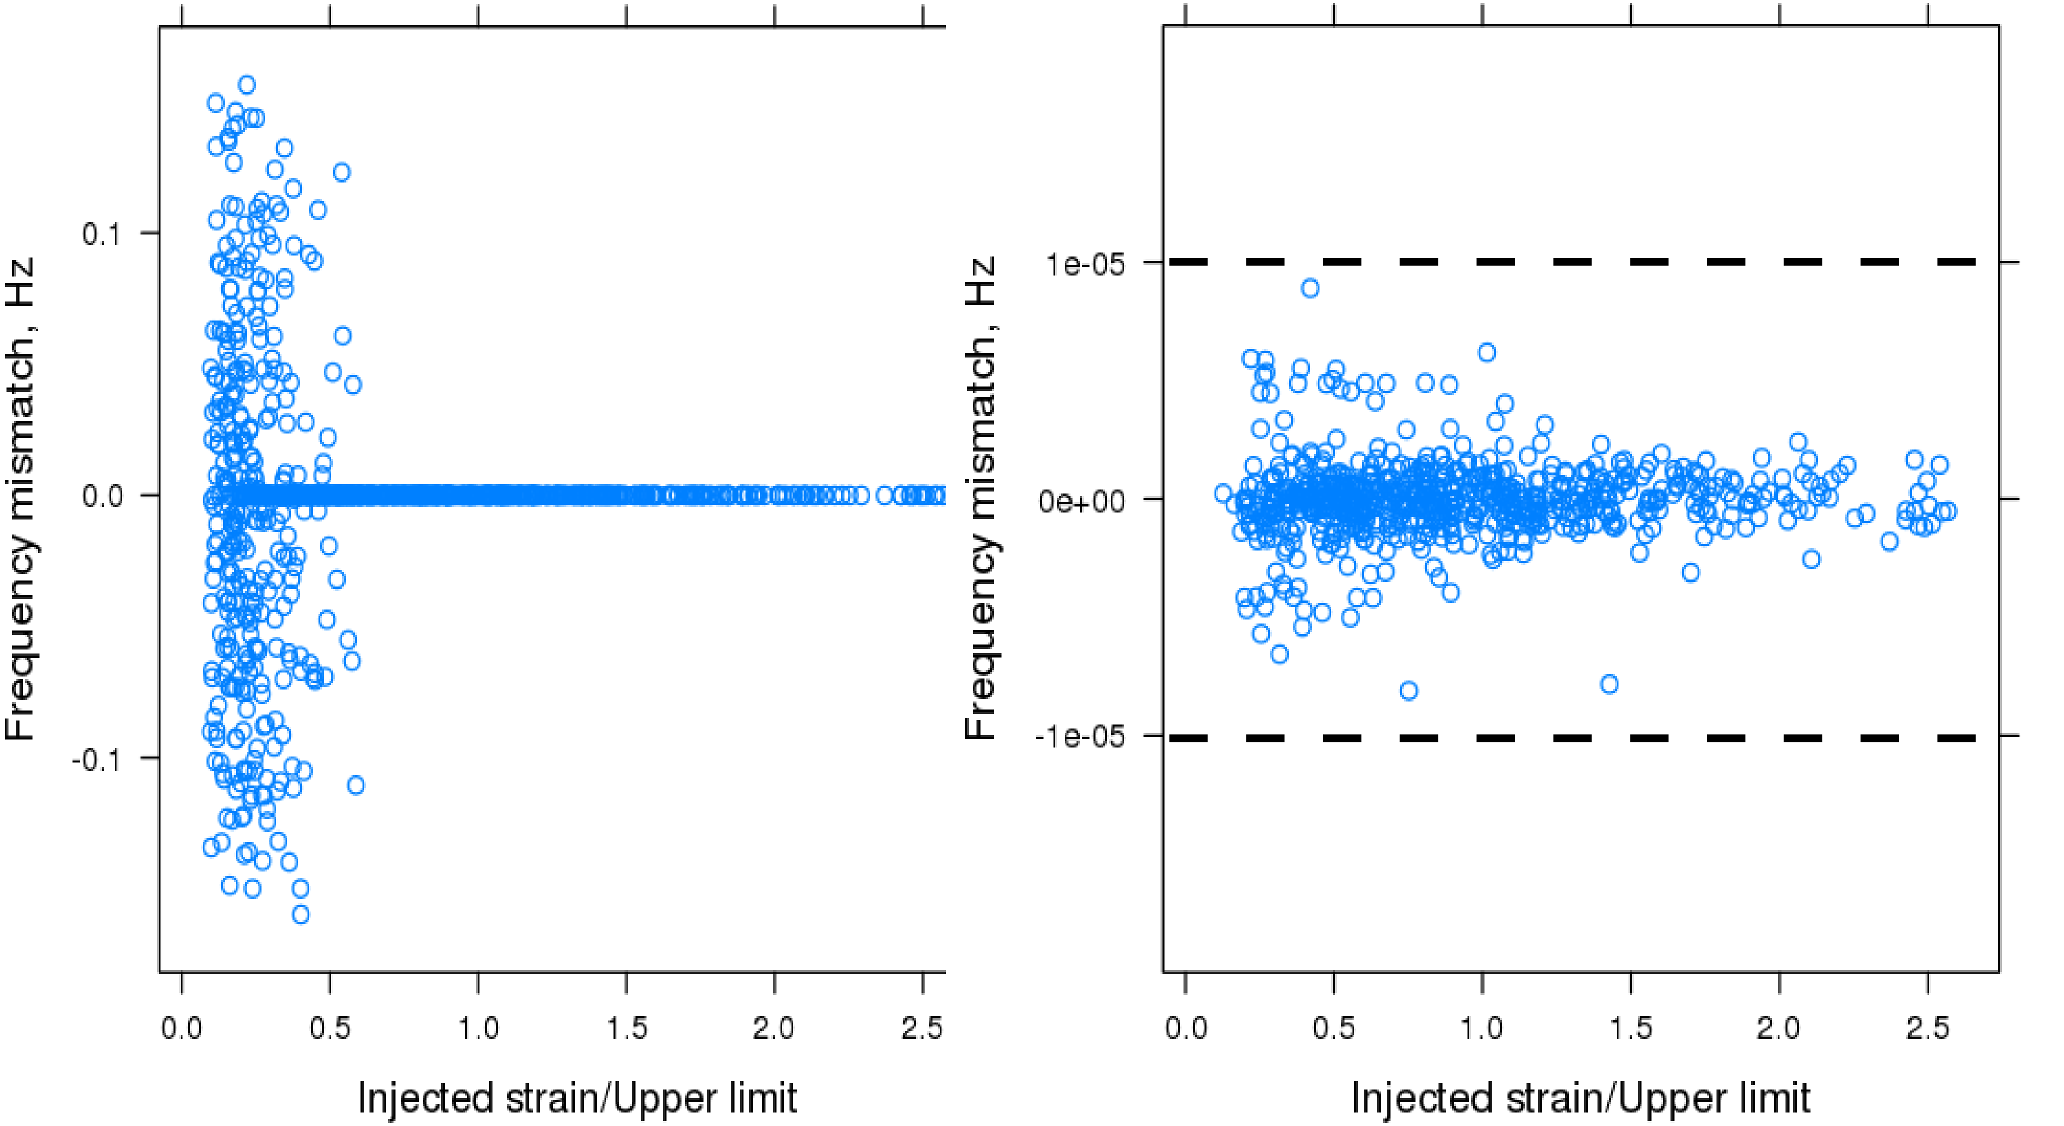

Detection efficiency

- Injections uniformly distributed on the sky and polarizations
- Injection was considered found if the frequency of largest outlier matched injection frequency with 10^{-5} Hz tolerance.

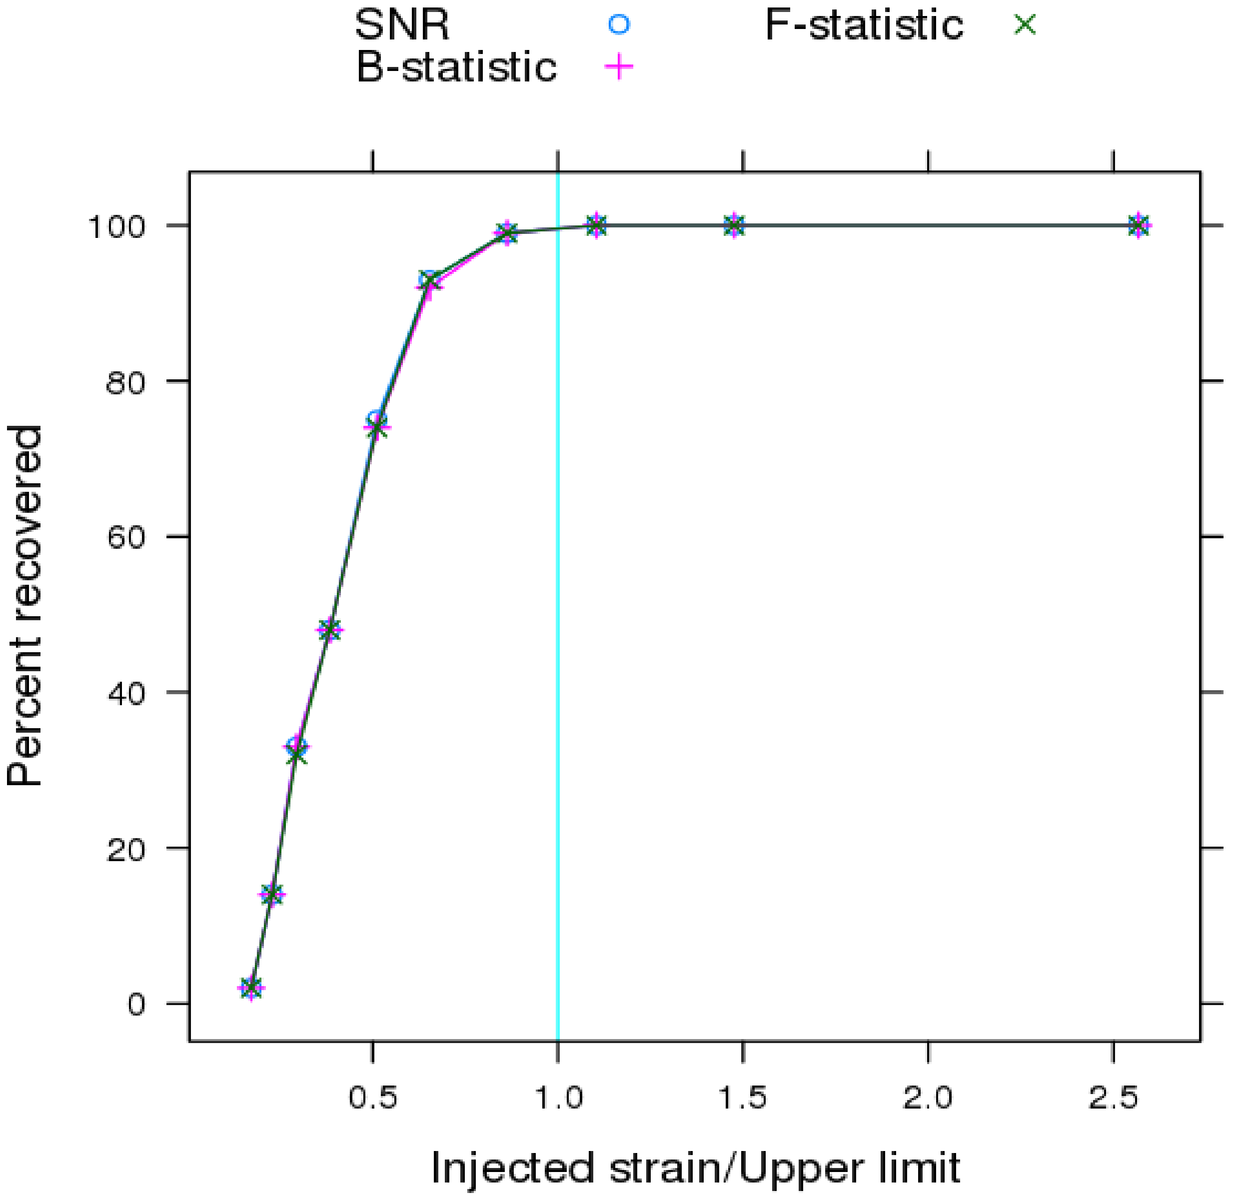

Code performance

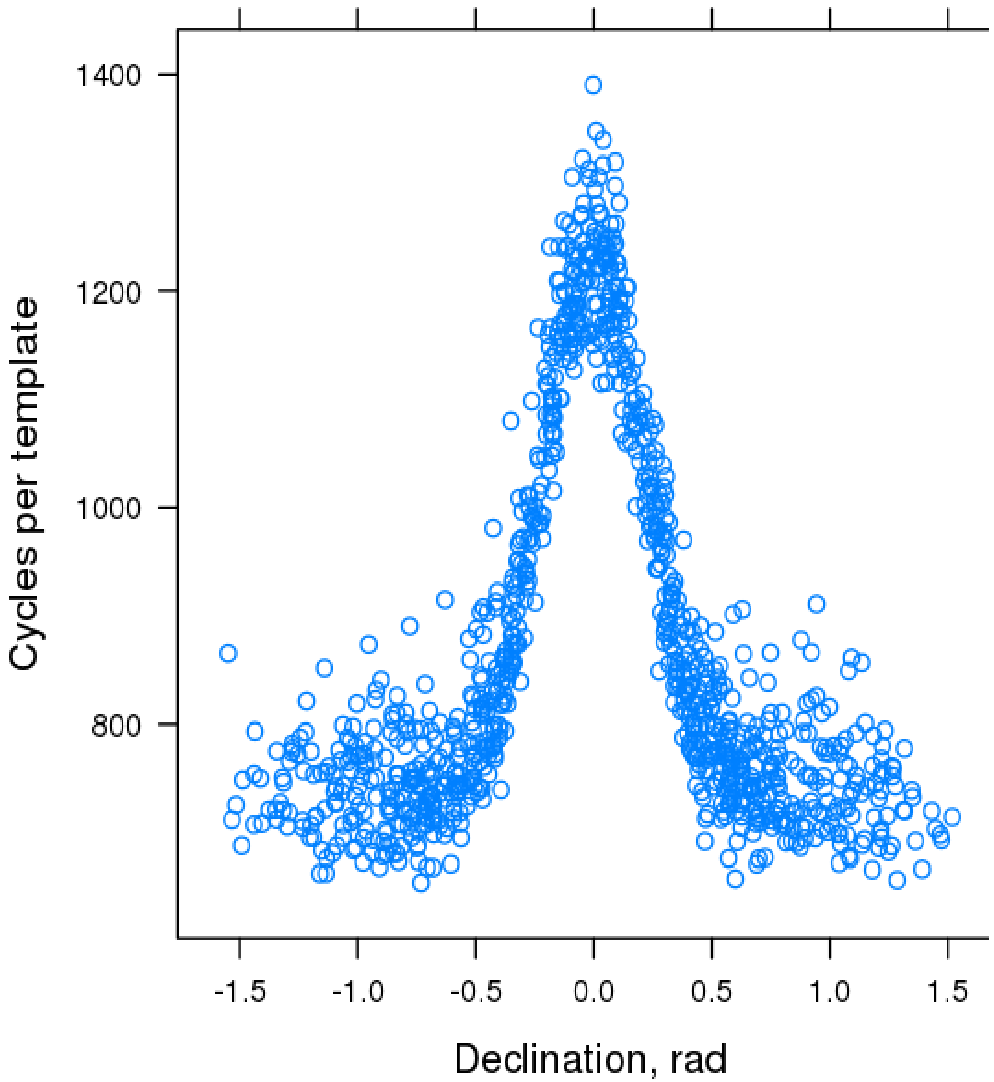

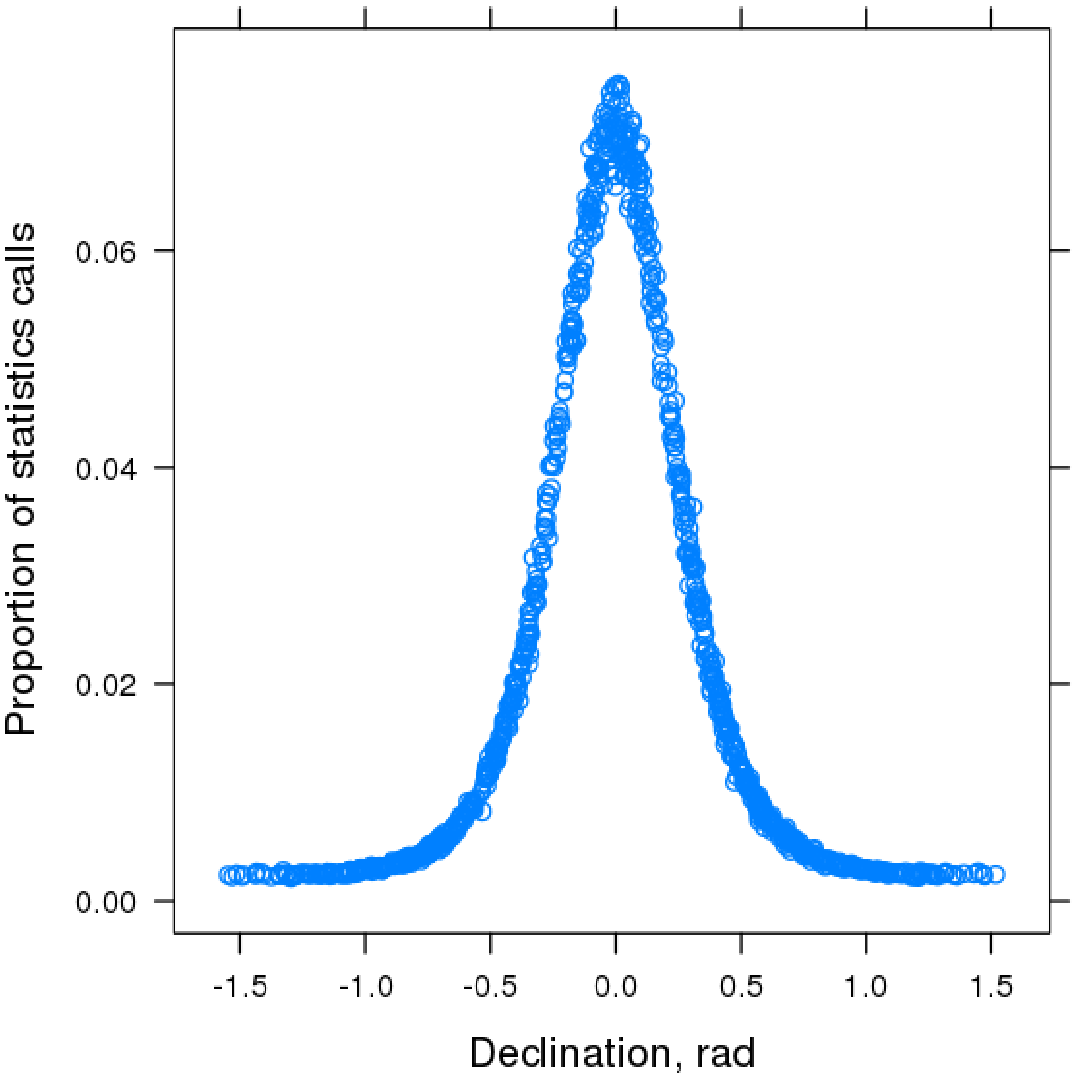

Summary/TODO

- We have designed and implemented $\delta=0$ loosely coherent search that achieves a factor of 10 improvement over previously available coherent codes.
- Second iteration of the code will focus on robustness to signal waveform and additional frequency evolution parameters
- Try analysis on a suitable “blob” region

Directed searches

• Directed searches focus on a particular spot on sky and search an accessible frequency range
• Due to larger number of templates strain sensitivity is lower
• Plot on the right: results from Cas A search (Astrophys. J. 722 (2010) 1504)

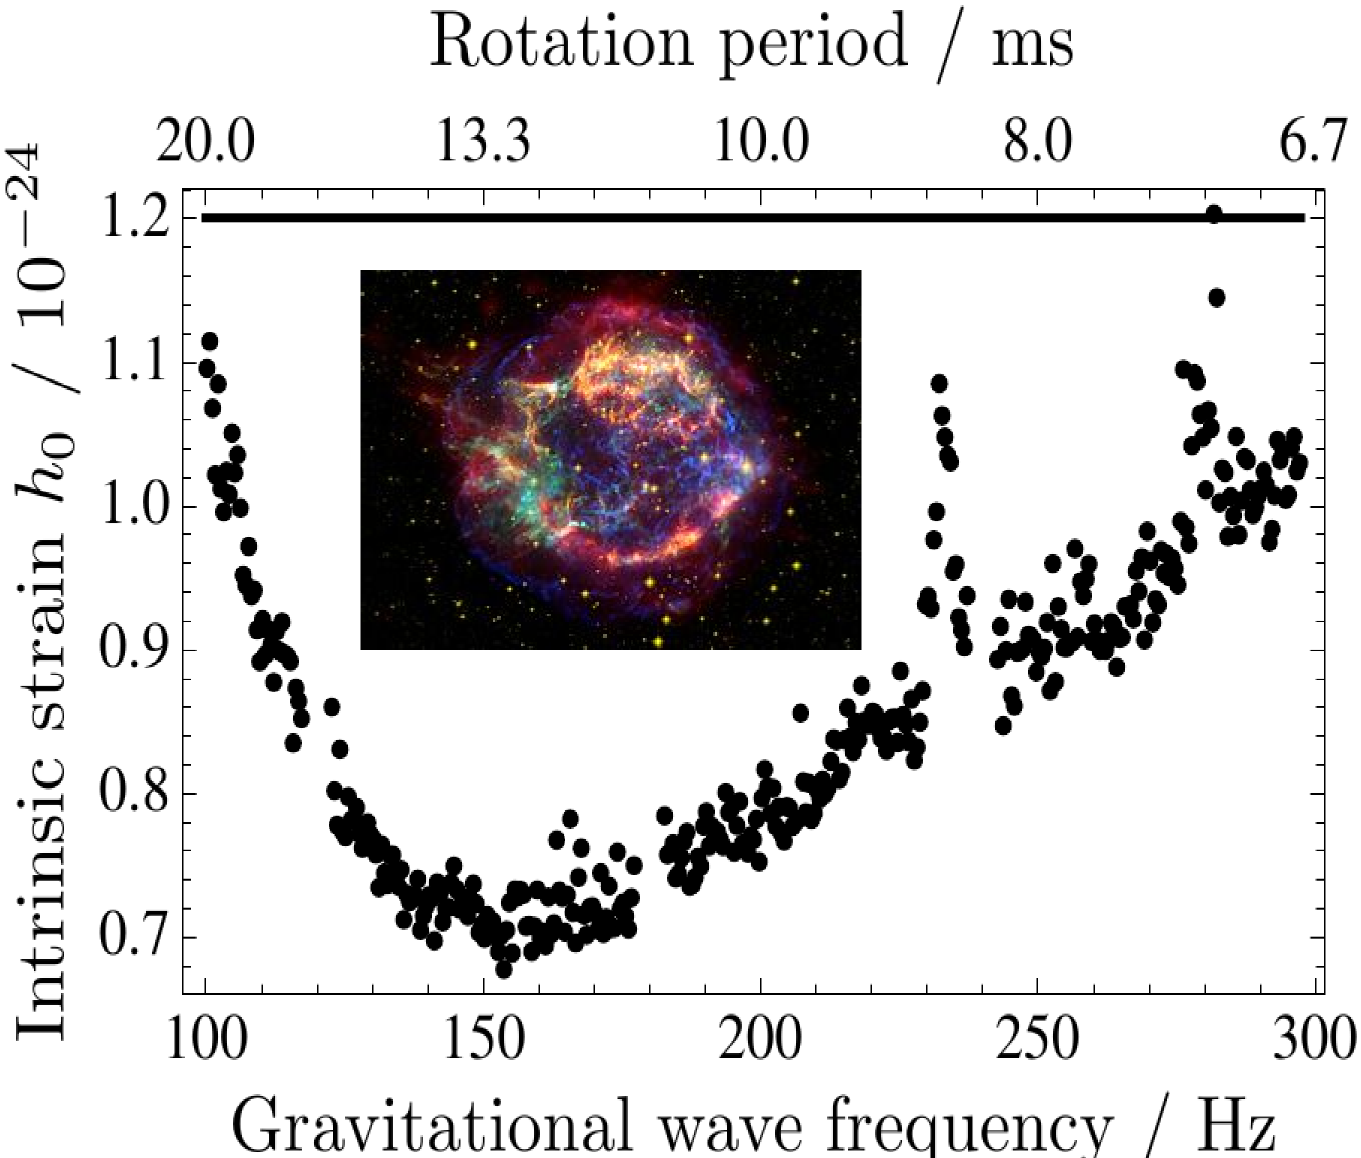

All-sky searches

- Modern all-sky searches are computationally limited.
- They are based on semi-coherent approach where results of coherent integration in several separate time stretches are incoherently combined to yield the final statistic
  - Einstein@Home - distributed search with computer time contributed by numerous volunteers
  - Hough search – semi-coherent search based on Hough algorithm
  - PowerFlux – semi-coherent search that estimates power coming from a particular direction on the sky
  - StackSlide – estimates power from particular direction on the sky using different weighting scheme

Sensitivity of CW searches

A semi-coherent search operating on data of length T and coherence length αT has a sensitivity scaling law of

$$\sim T^{-0.5} \alpha^{-0.25}$$

- For example, for a search using 30-min coherent segments over period of 1 year

  $T^{-0.5} \alpha^{-0.25}$ =0.002/sqrt(s)

- A purely coherent search of just 70 hours has a similar factor $T^{-0.5}$ =0.002/sqrt(s), but the necessary frequency resolution is 140 times finer.
- Actual sensitivity depends strongly on implementation particulars and response to detector artifacts.

Sensitivity of CW searches

The formula $T^{-0.5} * \alpha^{-0.25}$ has a slow dependence on α so a search with 1 year coherence time gains a factor of ~10 over semi-coherent search using 30-minute SFTs that uses the same data.

- This improvement is seen in broadband searches targeted at a specific sky location.
- It is highly desirable to do long-duration coherent searches over an extended area.

Intriguing known neutron stars

- There are over 2000 known neutron stars, with estimates of 10^8 to 10^9 in our galaxy.
- PSR J0108-1431 – 130 parsecs away.
- PSR B1508+55 – 1100 km/s velocity.
- PSR J1748-2446ad – rotation frequency of 716 Hz (if gravitational waves are emitted they will likely be at 1432 Hz).
- What else is out there ?

Neutron star ellipticity

- Simulations show that neutron star crust has breaking strain of ~0.1
- This corresponds to maximal ellipticity of $4*10^{-6}$ for a neutron star of 1.4 solar mass and 10 km radius.

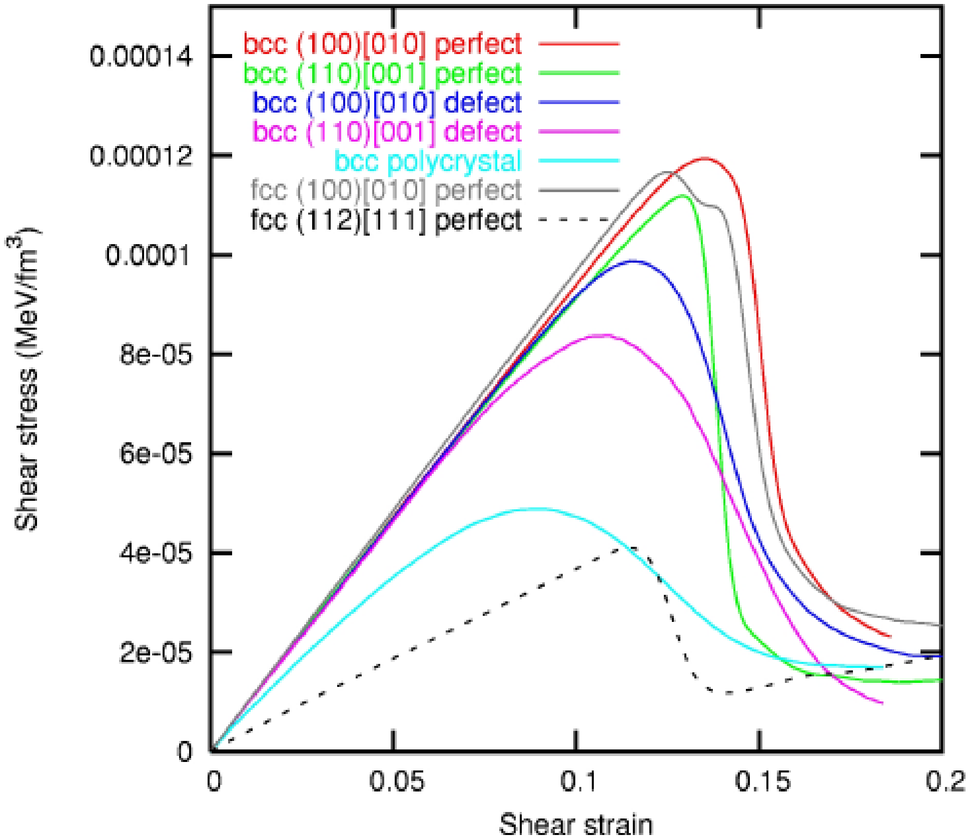

FIG. 1: Shear stress versus strain for perfect and defective bcc and fcc single crystals containig about 2 million ions for different shear systems —as given by the crystallographic orientation of the shear plane and the direction of the applied stress [17] — are shown. In addition a polycrystalline sample containing 12.8 million ions and 8 randomly oriented grains with an average grain diameter of 3962 fm is shown. Results were obtained at a strain rate of 4×10^{-7} c/fm.

Breaking Strain of Neutron Star Crust and Gravitational Waves, C. J. Horowitz and Kai Kadau, Phys. Rev. Lett. 102, 191102 (2009)

Intriguing known neutron stars

- There are over 2000 known neutron stars, with estimates of 10^8 to 10^9 in our galaxy.
- PSR J0108-1431 – 130 parsecs away.
- PSR B1508+55 – 1100 km/s velocity.
- PSR J1748-2446ad – rotation frequency of 716 Hz (if gravitational waves are emitted they will likely be at 1432 Hz).
- PSR B1257+12 has three planets, closest at 0.19 AU.
- PSR B1620-26 has a 2.5 Jupiter mass planet that orbits both it and a white dwarf companion.
- What else is out there ?

Neutron star facts

- A small mass m on the surface of the neutron star of 1.4 solar masses and 20km diameter has a gravitational potential ~ $GMm/r= GM/(rc^2)\ (mc^2) \sim 0.2\ mc^2$
- This is just Schwarzschild radius divided by the radius of the neutron star.
- For each proton/neutron of mass 1GeV this energy is ~200 MeV
- For comparison, typical chemical binding energy is 1eV, hydrogen fusion energy is 14 MeV and full energy from U^{235} fission is 20 MeV
- When material falls onto the surface of the neutron star this energy gets released. Accreting neutron stars glow in X-rays !

Searches for binary systems

- There are several searches for binary systems:
  - Sideband search – based on coherent F-statistic with a comb template
  - Radiometer search – measures directional crosscorrelation between detectors
  - TwoSpect – uses secondary Fourier transform to find binary periodicity
  - Polynomial search – uses polynomial approximation to signal model

Loosely coherent search

Semi-coherent

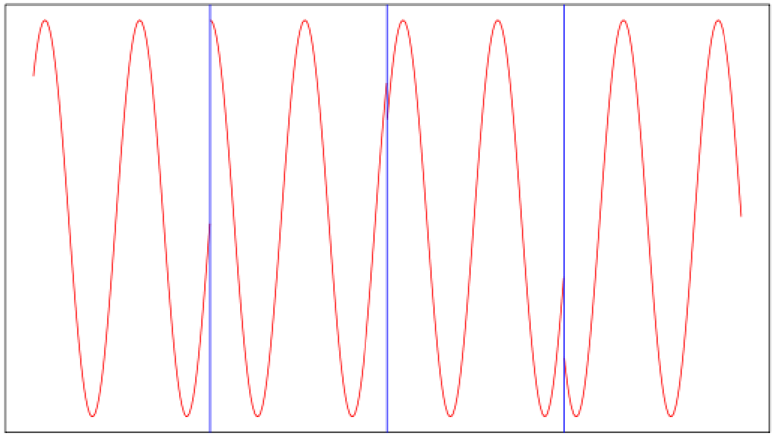

Time

Coherent

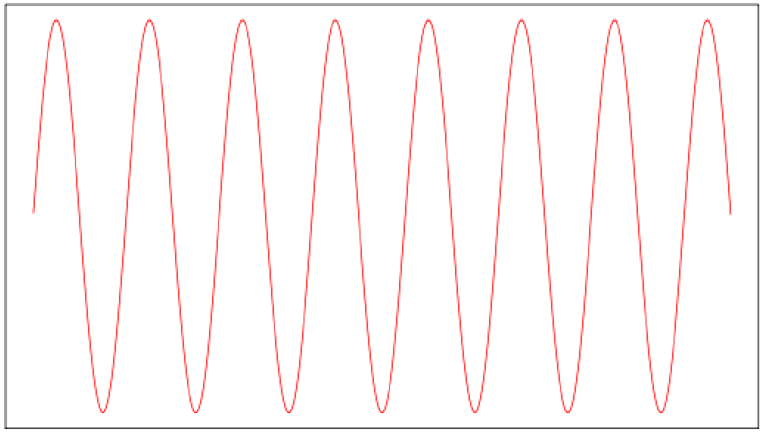

Time

Loosely coherent

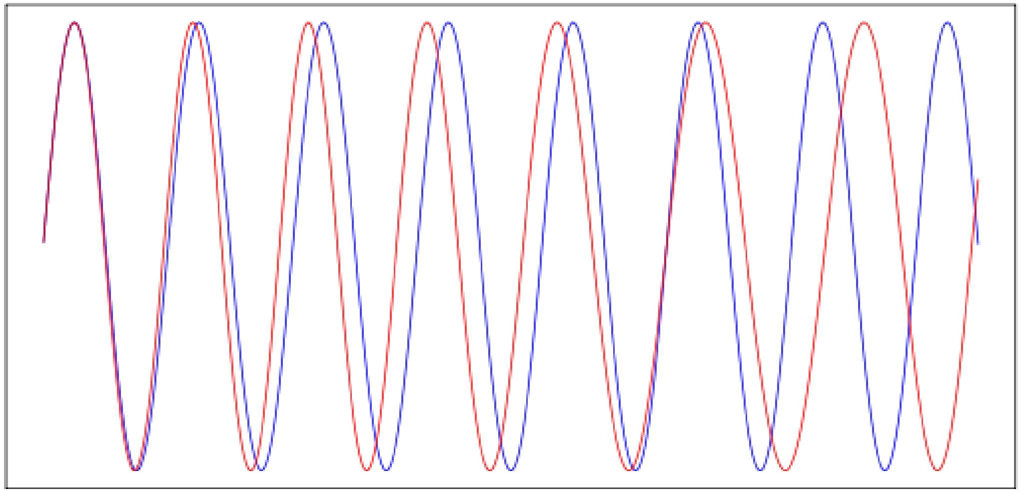

Time

Example

- Start with a fully coherent sum of a series of SFT bins (one bin from each SFT) $P = |\sum a_k e^{i\phi_k}|^2$
- Average over all phases such that $|\phi_k - \phi_{k+1}| \leqslant \delta$
- Obtain statistic with reduced sensitivity to phase drift:

$$P' = \sum \bar{a}_k a_l \left(\frac{\sin(\delta)}{\delta}\right)^{|k-l|}$$

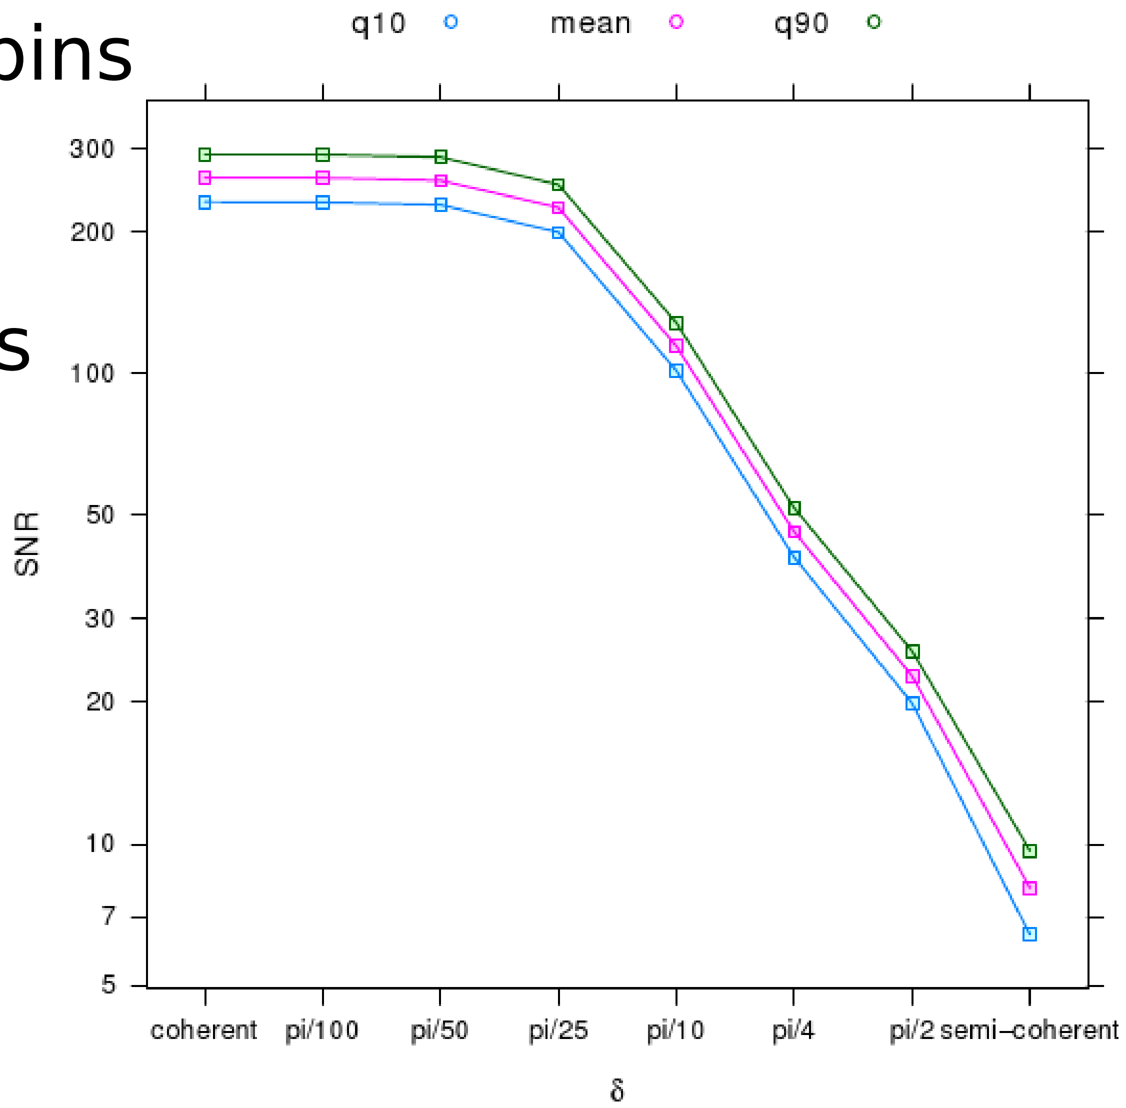

Another point of view

Assume that SFT bins have already been corrected for Doppler shift (resampled):

Data

a_1 a_2 a_3 a_4 ....

$h*\exp(i*\varphi_1)$ $h*\exp(i*\varphi_2)$ $h*\exp(i*\varphi_3)$ $h*\exp(i*\varphi_4)$ ....

Expected signal

Since phases vary slowly $|\varphi_k - \varphi_{k+1}|<\delta$ we can use a low-pass filter to reject rapidly varying noise and then compute the power sum as usual.

This method is optimal if all slowly varying signals are physically valid – which is the case for very small δ.

PowerFlux parameter reconstruction

- Find highest SNR coincidences using method close to one that will be used in full S5 analysis (still fine tuning the constants)
- Plots show difference between coincidence and injection

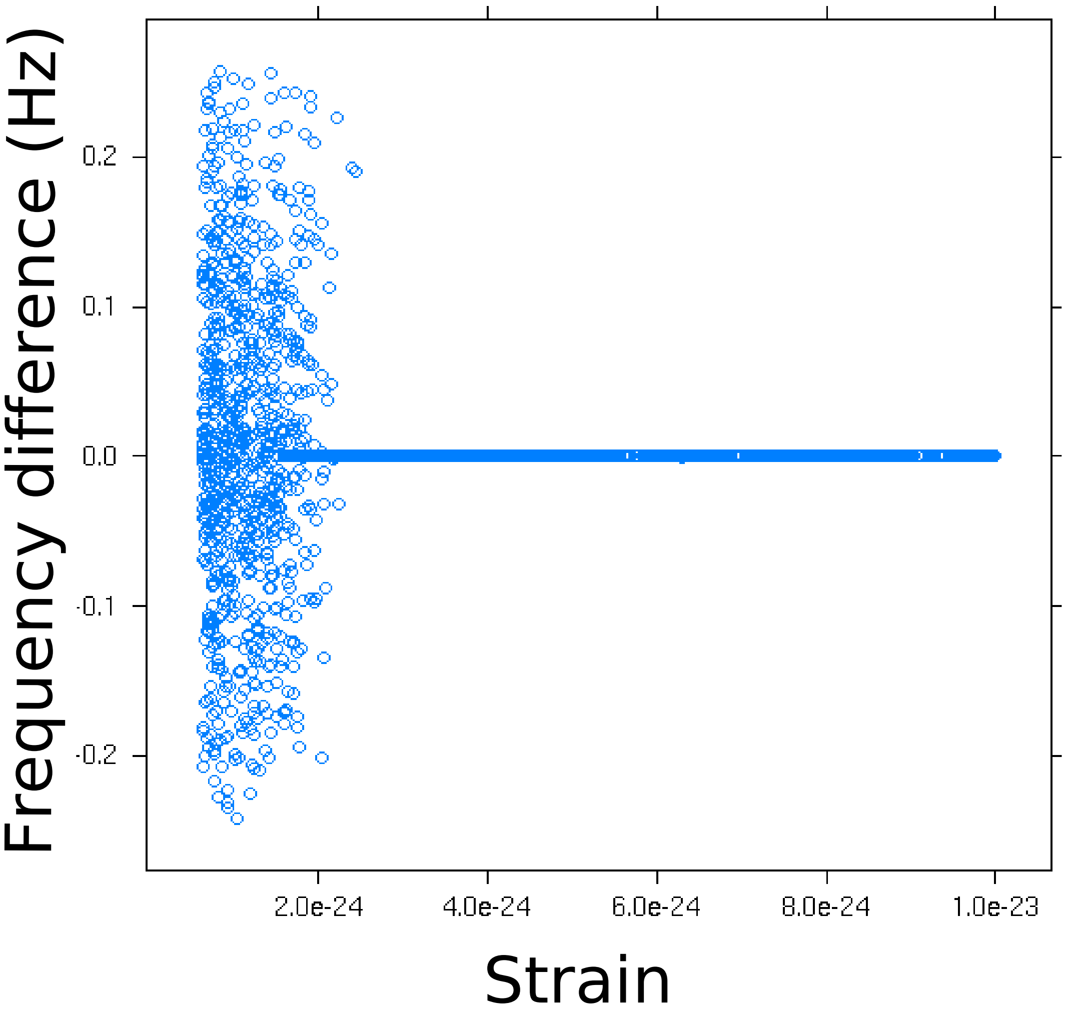

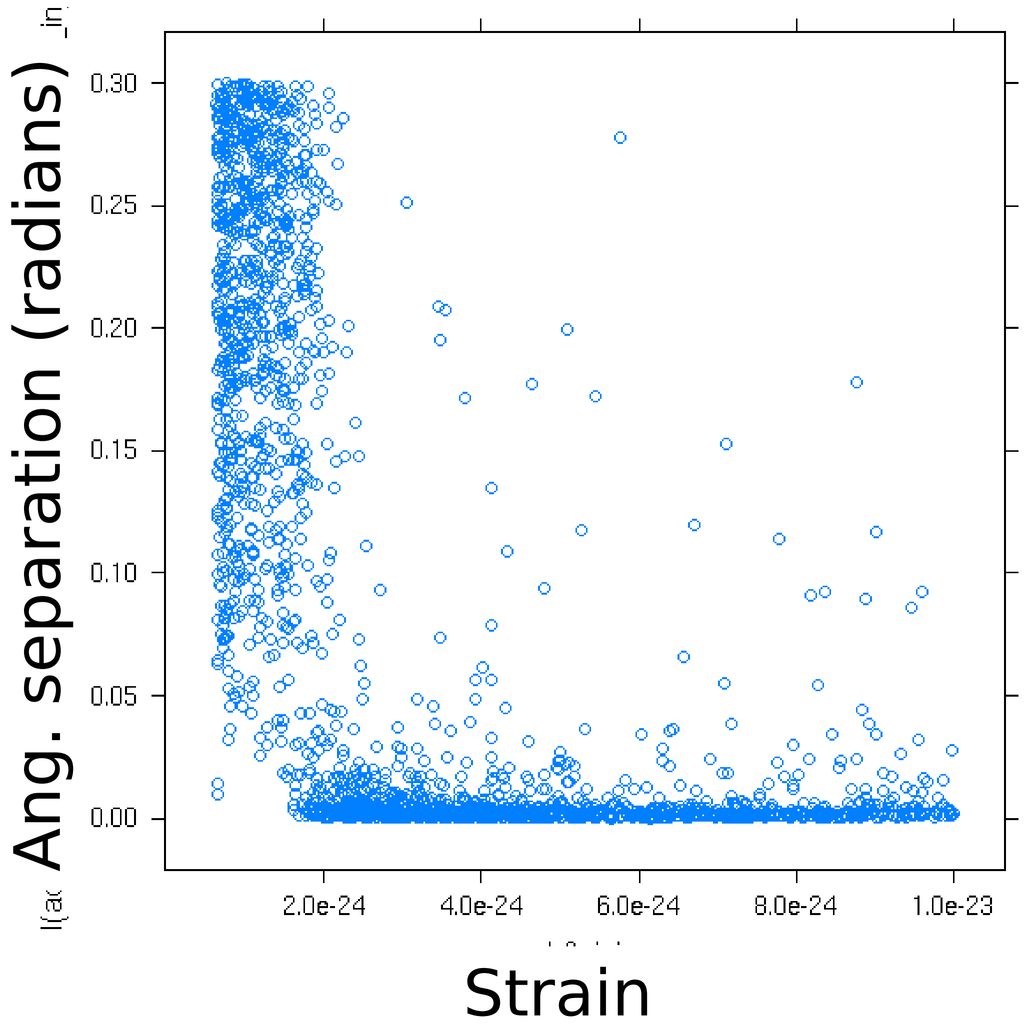

PowerFlux parameter reconstruction

- Find highest SNR coincidences using method close to one that will be used in full S5 analysis (still fine tuning the constants)
- Plots show difference between coincidence and injection

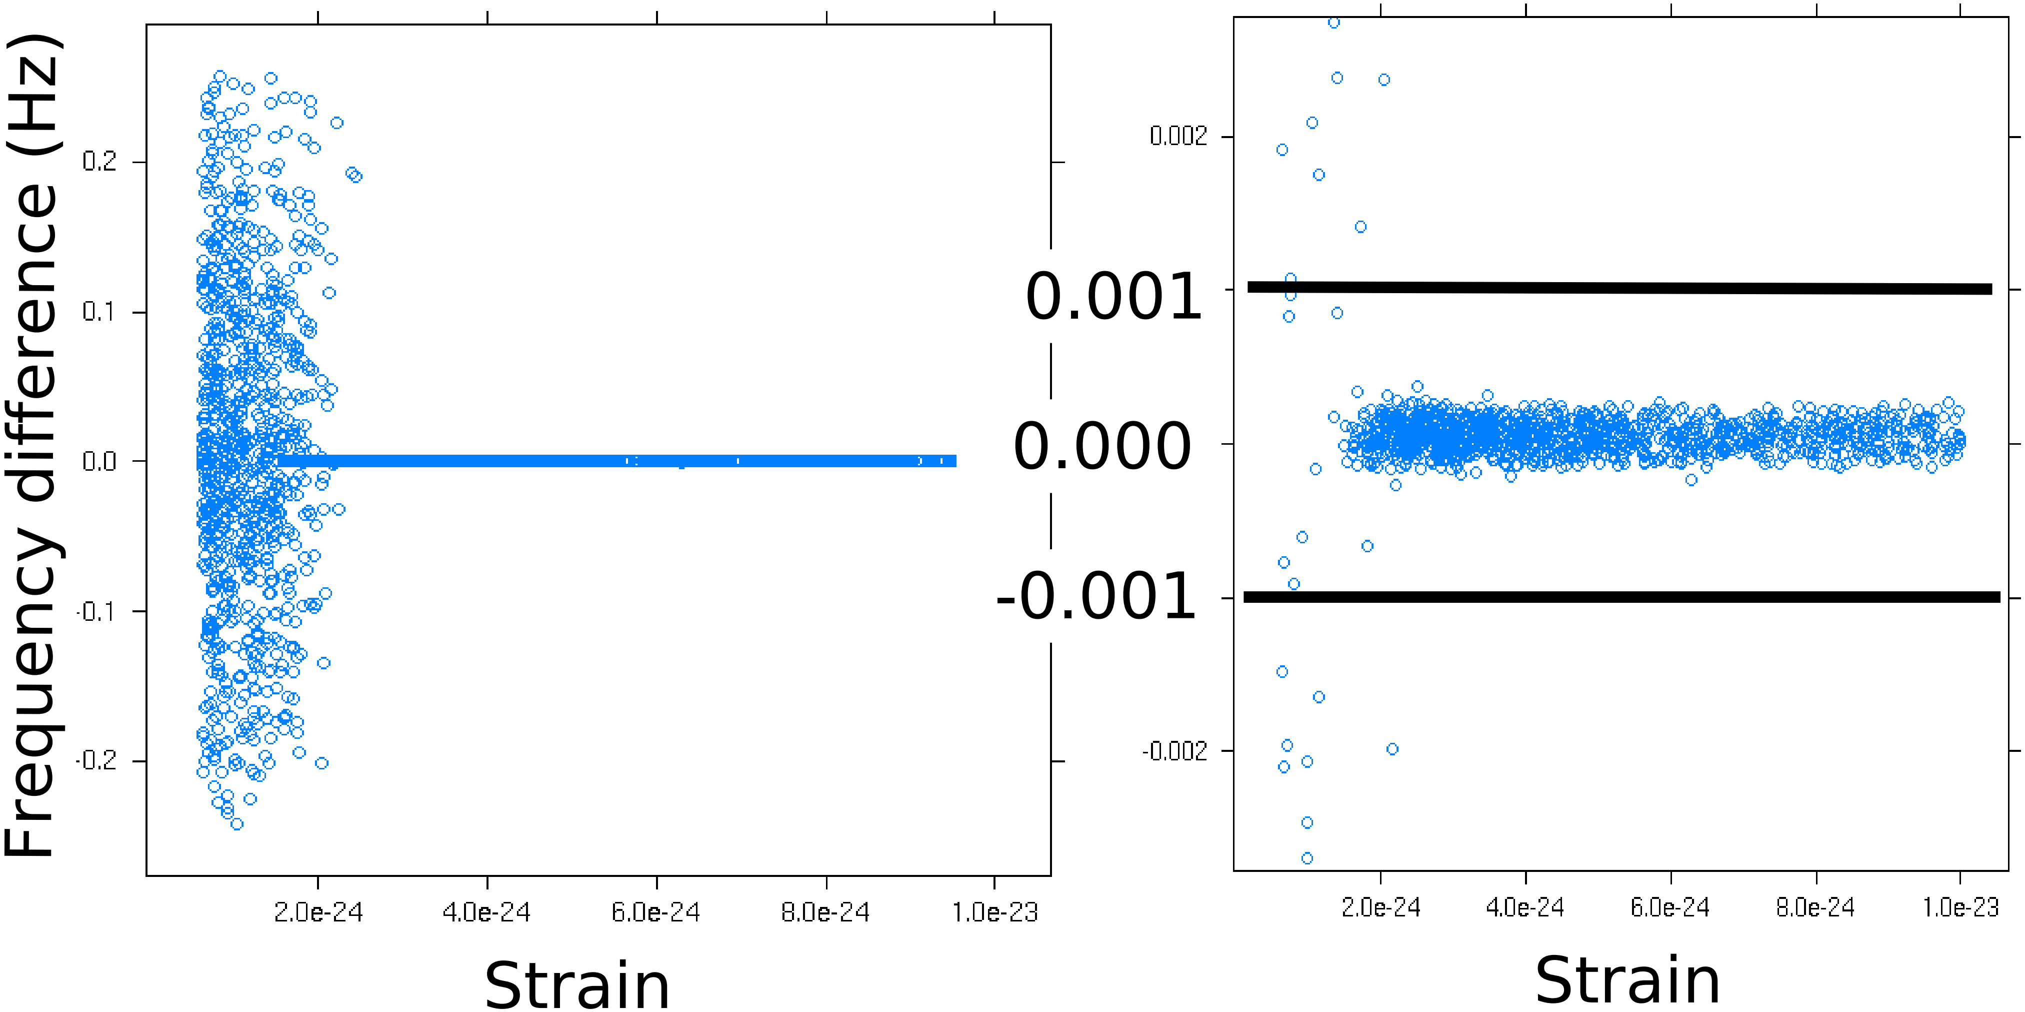

Size of frequency shift in 1/1800 Hz bins due to different causes

	Relative	50 Hz	500 Hz	1500 Hz
Earth rotation	1e-6	0.1	1	3
Earth orbital motion	1e-4	9	90	270

	3 months	1 year	2 years
1e-11 Hz/s spindown	0.14	0.56	1.1
1e-10 Hz/s spindown	1.4	5.6	11
1e-9 Hz/s spindown	14	56	112

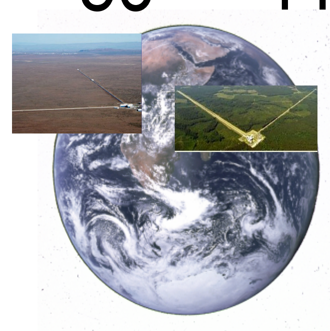

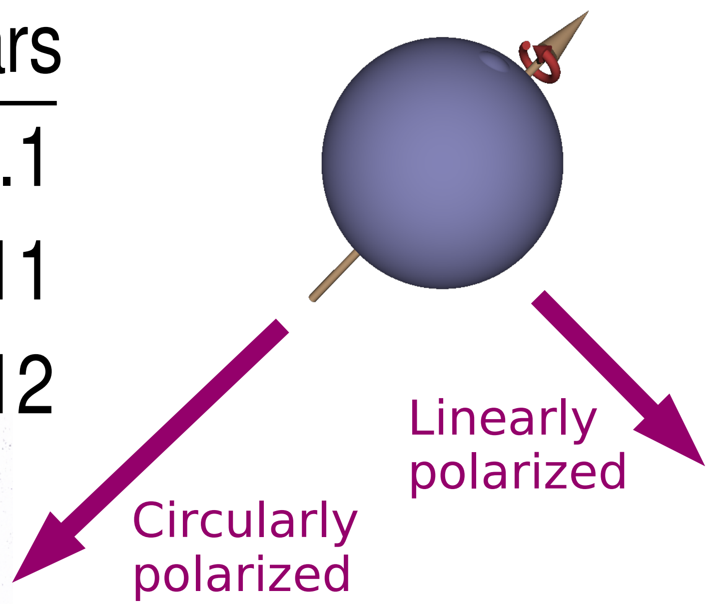

PowerFlux 2.x analysis pipeline

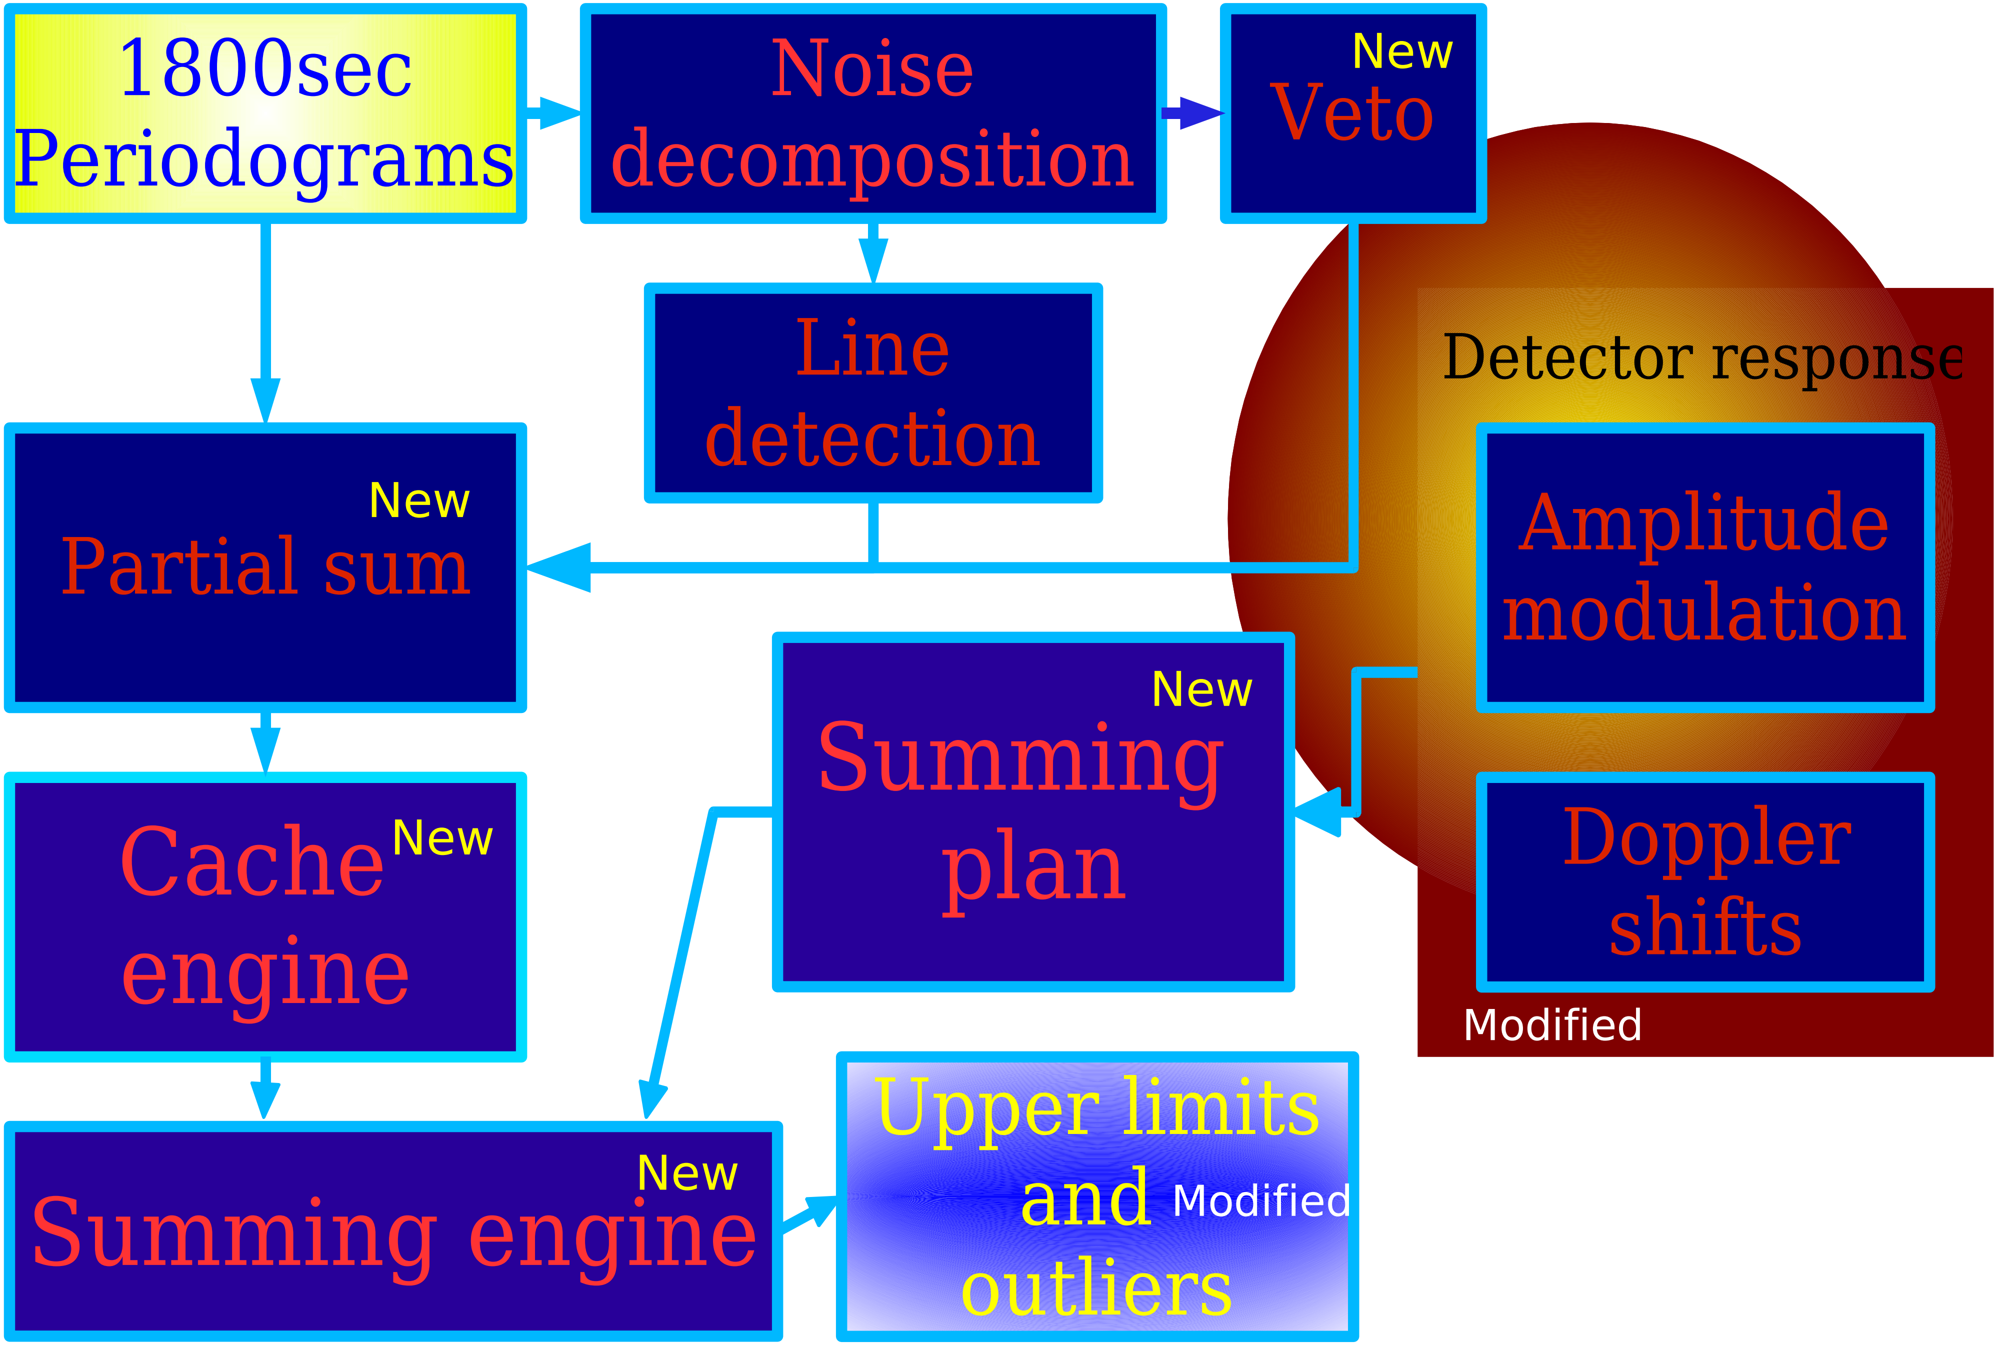

Monte-Carlo injection runs

• Paper shows sample plots from 400 Hz Monte-Carlo injection run
• Review covered Monte-Carlo runs in 50, 100, 400 and 800 Hz bands.
• Right: reconstructed strain versus injected value

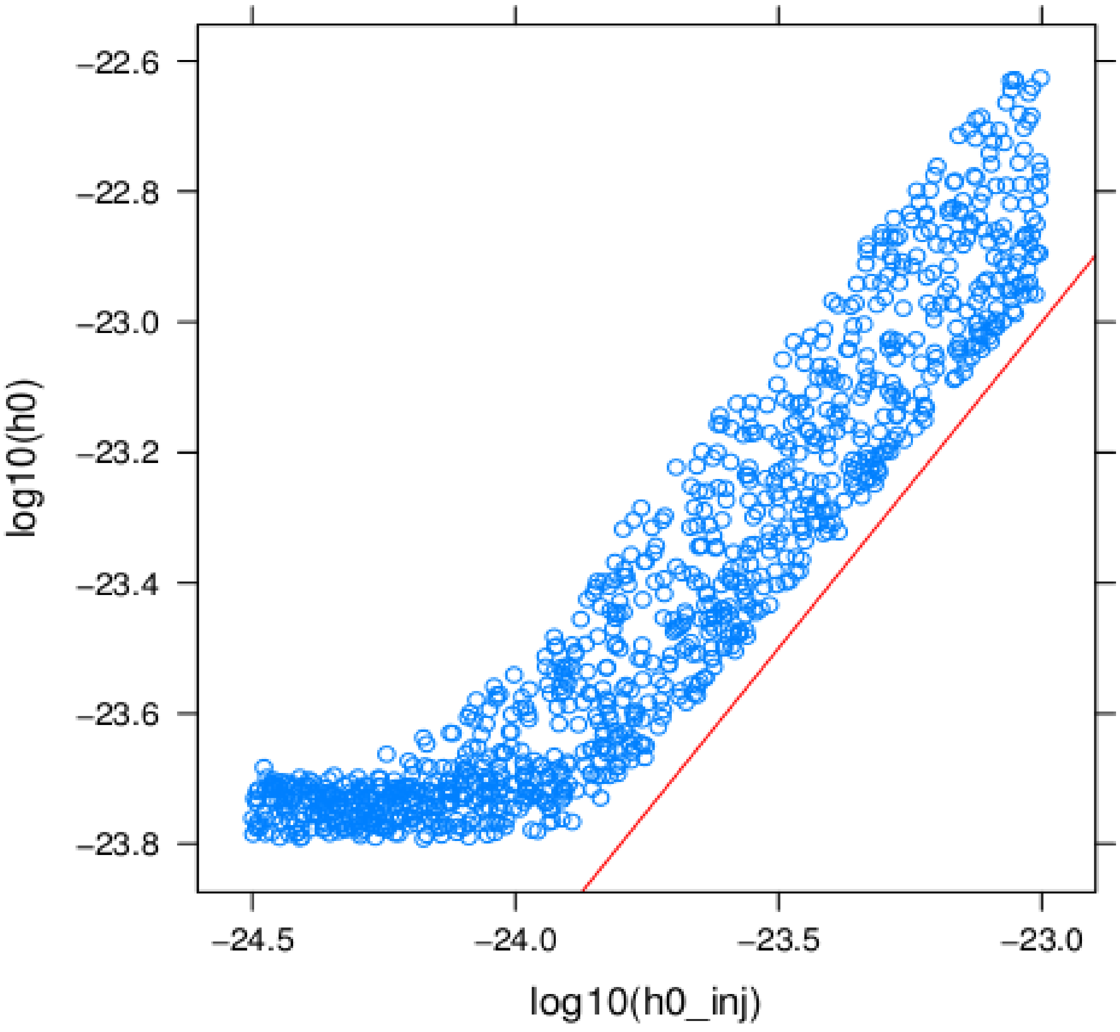

Loosely coherent search

It is likely that the brightest CW object is extreme in some way and has a special reason (like a companion star or gas giant) for large quadrupole moment.

This can cause phase evolution different from perfect monochromatic emitter model.

Semi-coherent searches are insensitive to large variations in phase, but are limited in sensitivity for targeted search applications.

Fully coherent searches assume very close adherence to monochromatic emitter model.

Need *Loosely coherent search* that is sensitive to signals with slow phase evolution.

Zooming in onto software injection

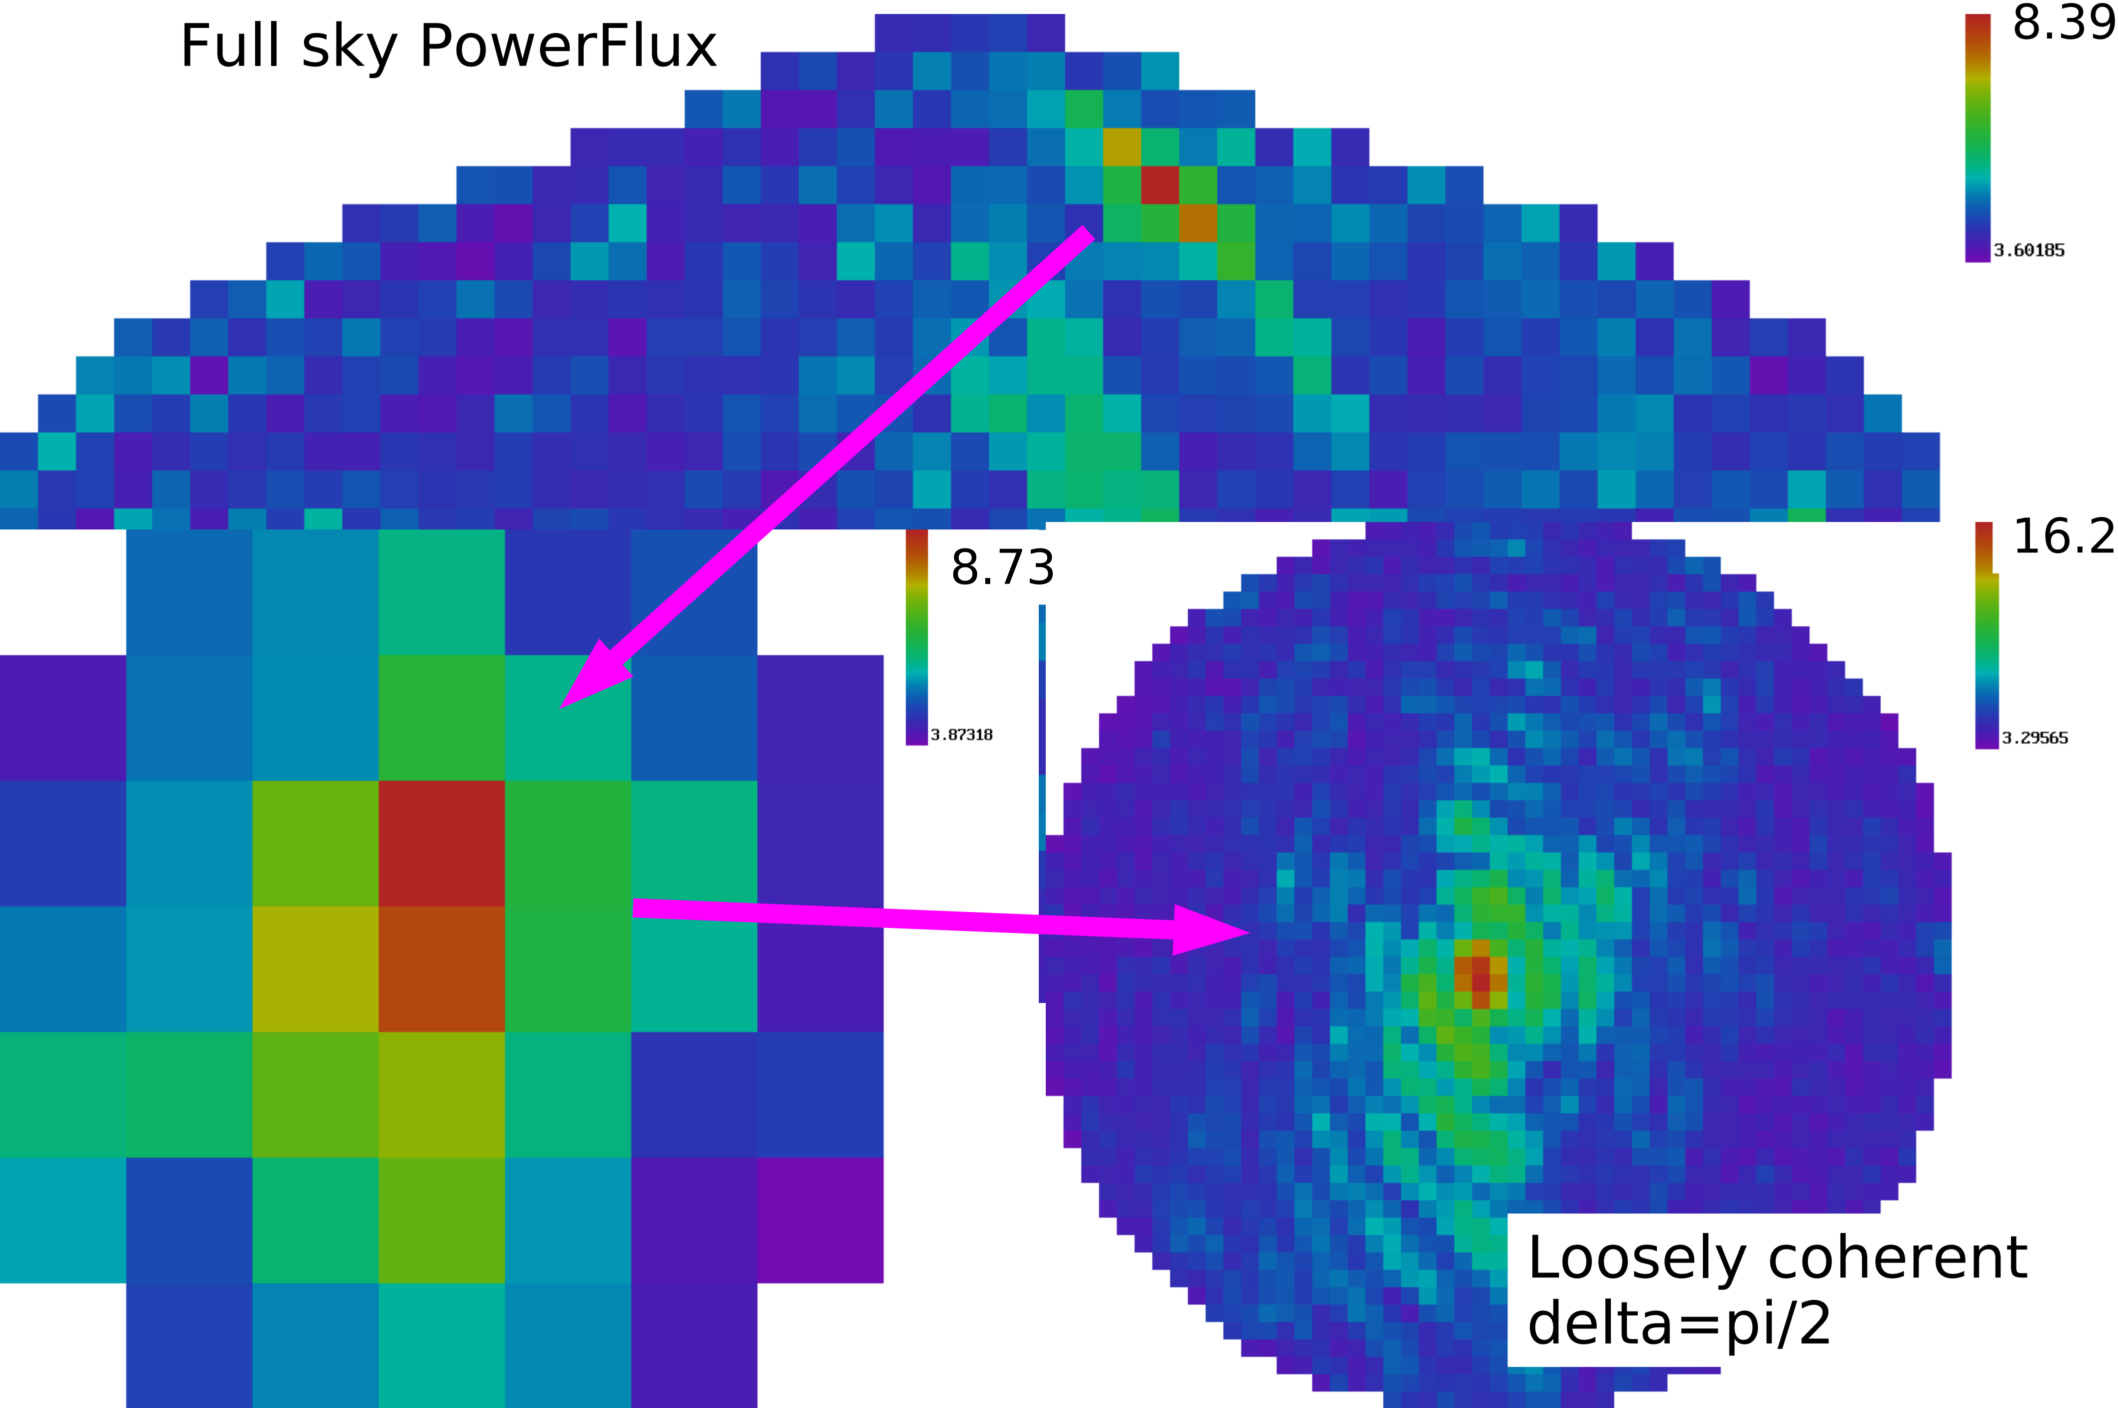

Injection recovery in 400 Hz band

- First stage is a semi-coherent PowerFlux run
- After coincidence test the outliers are passed to pi/2 loosely coherent search
- After cuts based on expected SNR increase the outliers are passed to pi/2 loosely coherent search with coherent combination of data between different interferometers

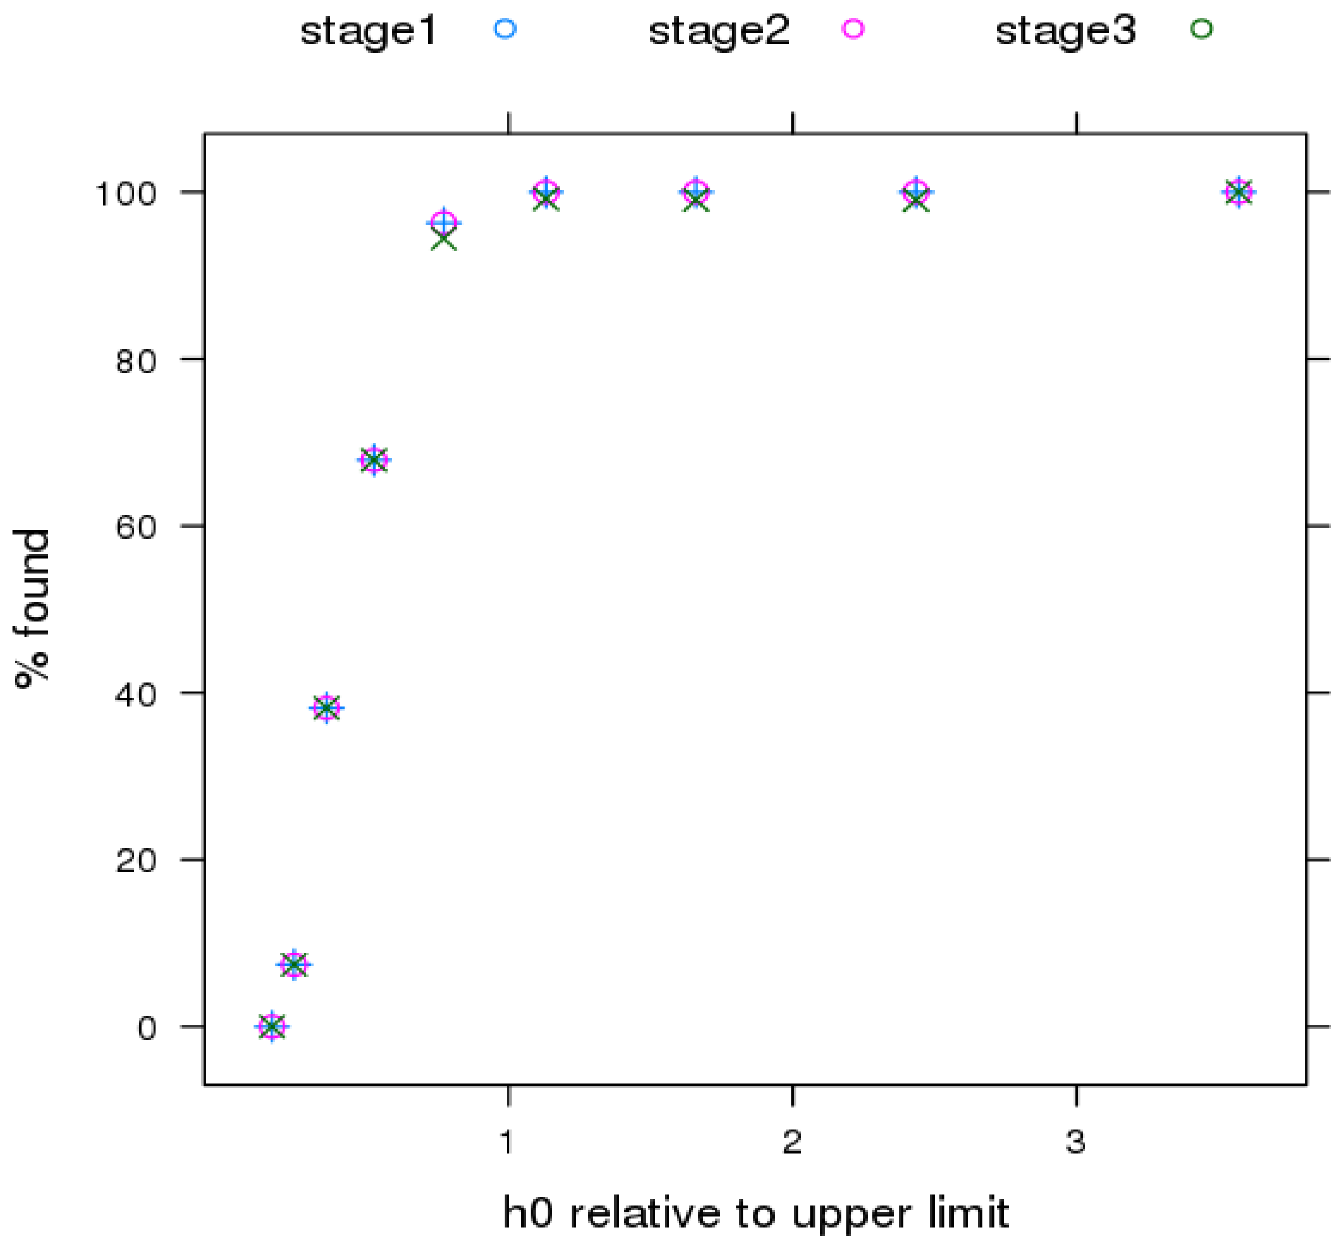

Large delta loosely coherent code

Implemented loosely coherent code on top of PowerFlux code base.

http://arxiv.org/abs/1003.2178

Tested for delta values of pi, pi/2 and pi/5 using simulated Gaussian data and actual interferometer noise.

Not as fast as we would like, but (almost ?) good enough for doing followups.

Very tempting to make a pass at optimizing the code to speed up simulations.

Dedicated loosely coherent code

- Purpose: wide- and narrow-band searches of small sky regions with long coherence times and control of phase evolution.
- Explore δ resulting in phase wrap around of 10-100 times over S5/S6 run.
- Produce upper limits and perform detection search similar to PowerFlux (and reusing PowerFlux infrastructure).

Code architecture

- Start with 4-second long SFTs. 1 Month of data = 1.2 million SFTs. Need new program to produce concatenated v2 SFTs right away.
- 4-second timebase = 0.25 Hz bins. No need to shift bins around, just apply phases.
- After applying phases, take a Fourier transform (similar to resampling).
- Compute power and sum it up using a sliding window – i.e. power after applying low pass filter.

Development status

- Ported a lot of PowerFlux code for use by dedicated loosely coherent codes
- Playground sets of simulated and real 4 second SFTs
- New injection engine based on LALBarycenter
- Have easy to modify “toy” code that implements resampling-like pipeline on a small grid
- Work will speed up once Full S5 analysis is complete

Monte-Carlo injection run

- 1000 injections into 400 Hz band.
- Search done within 0.03 radians of injection point (similar to coherent pipeline).
- Wide SNR band – not all injections get a high SNR due to the offset problem.

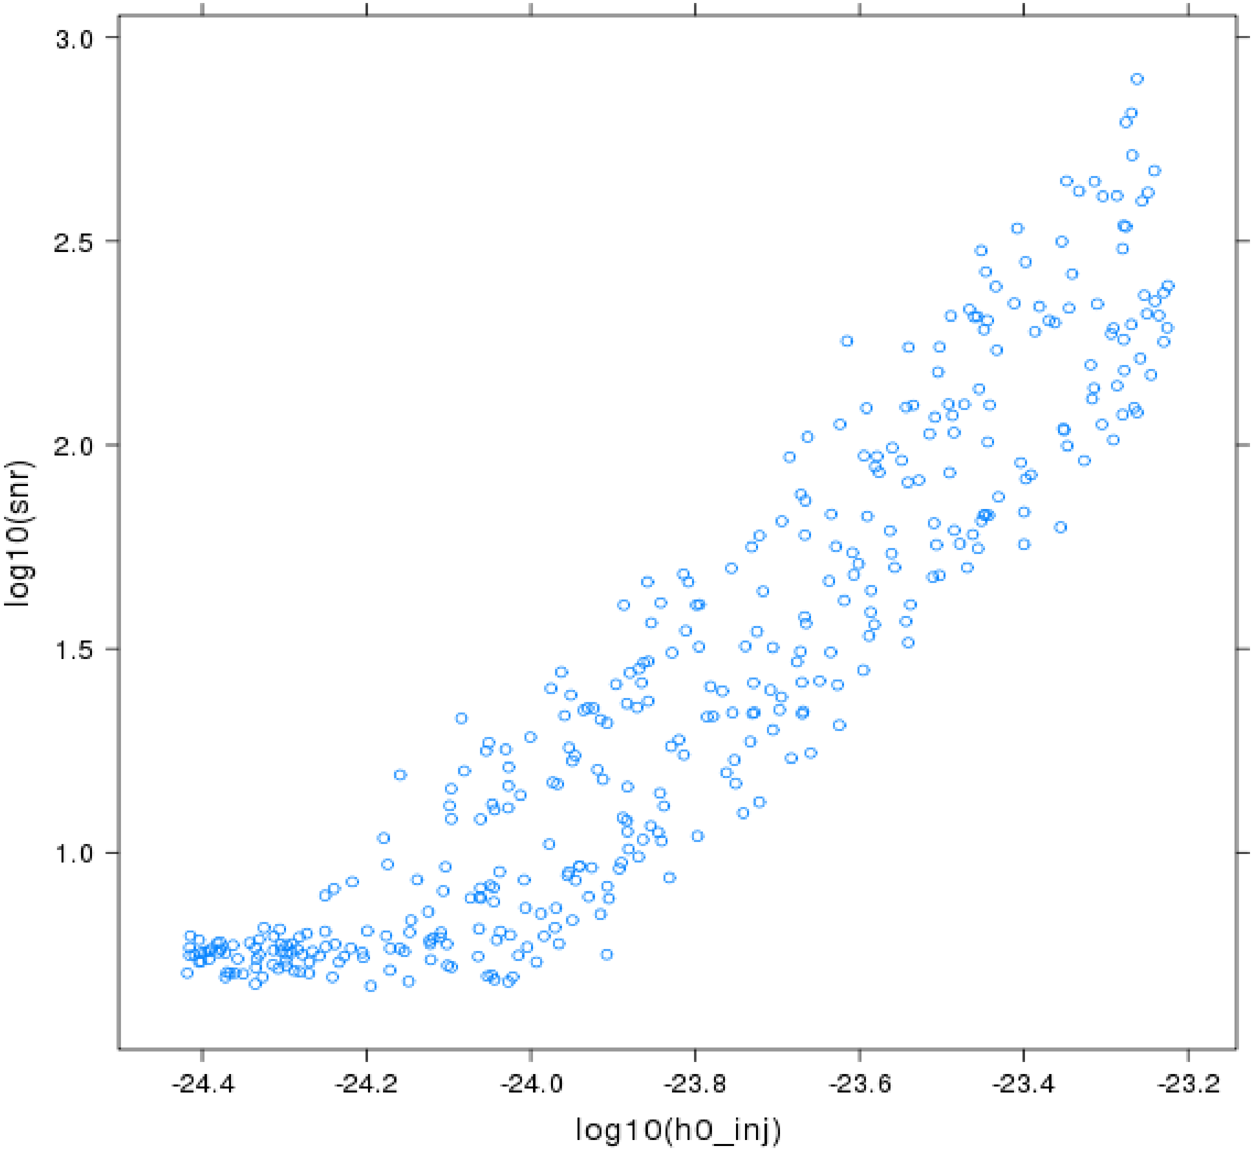

We do see finer parameter reconstruction

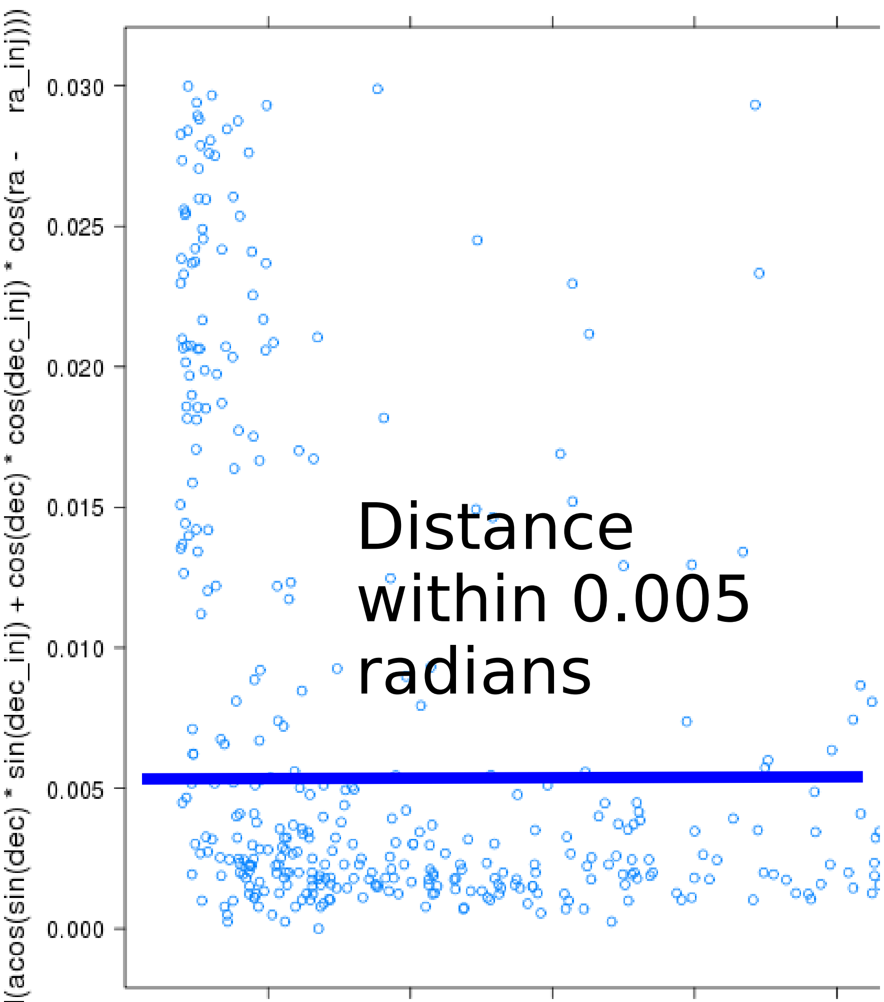

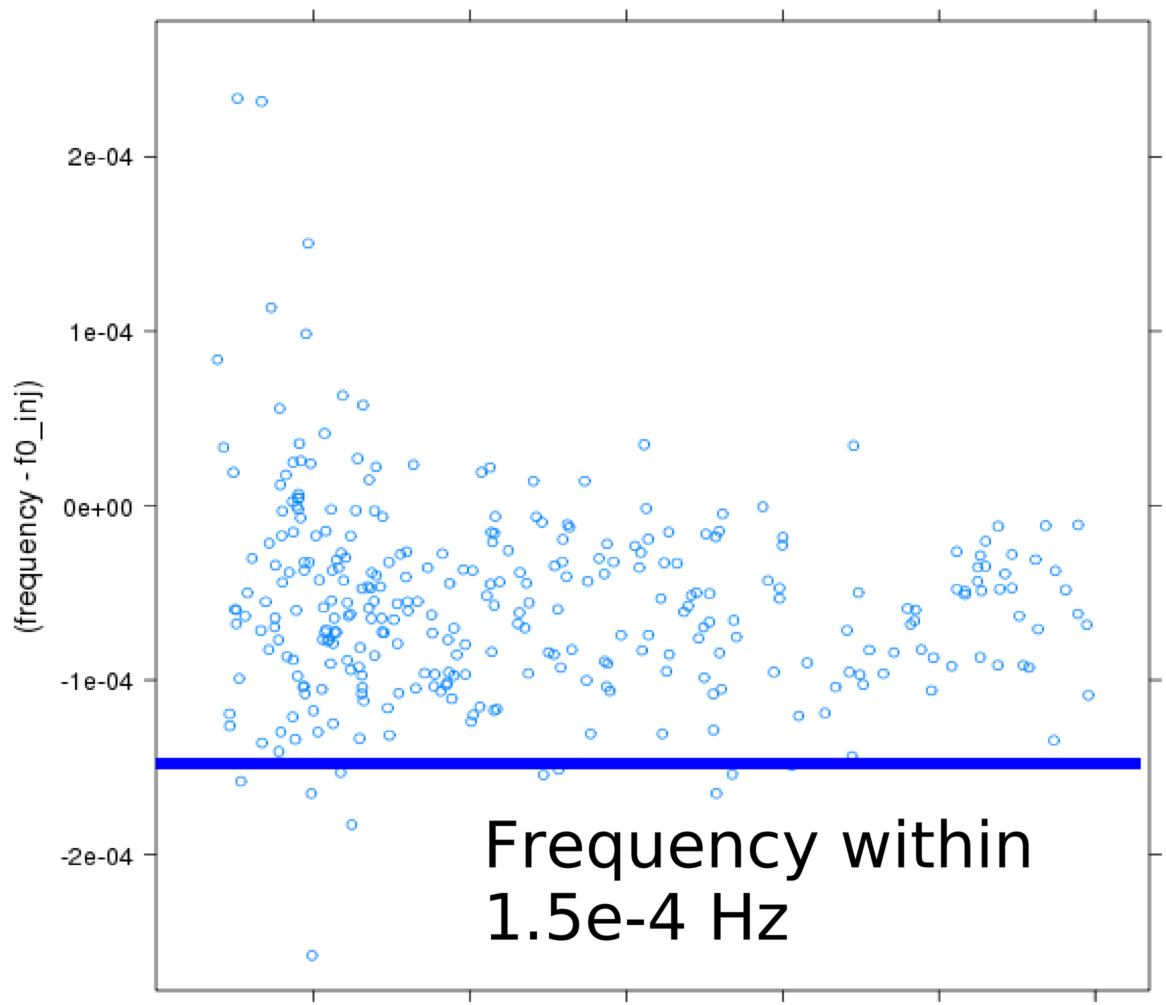

Lessons from initial PowerFlux implementation

Present code is about 400 times slower than single-bin PowerFlux sum – good enough for full S5 followups but more optimization can be done.

Two major factors

need much finer parameter steppings to correctly treat phase

inefficient sum over two indices

Also need to use proper barycentric code instead of approximation from Doppler shifts alone (not possible to do loosely coherent multi-IFO sum without it).

Not surprising as we are building on top of infrastructure developed for a different algorithm.

It would be best to have dedicated loosely coherent code

Response to frequency mismatch

- Linearly polarized injections with h0=1.8e-23 into Gaussian data
- Fixed sky location
- Frequencies within 400-410 Hz
- Loosely coherent search with Lanczos kernel

$$\frac{\sin(\delta\,\Delta\,t/30\text{min})\sin(0.333\cdot\delta\,\Delta\,t/30\text{min})}{0.333\cdot\delta^2\cdot(t/30\text{min})^2}$$

(zero when $\delta\Delta t/30\text{min}$ exceeds 3π)

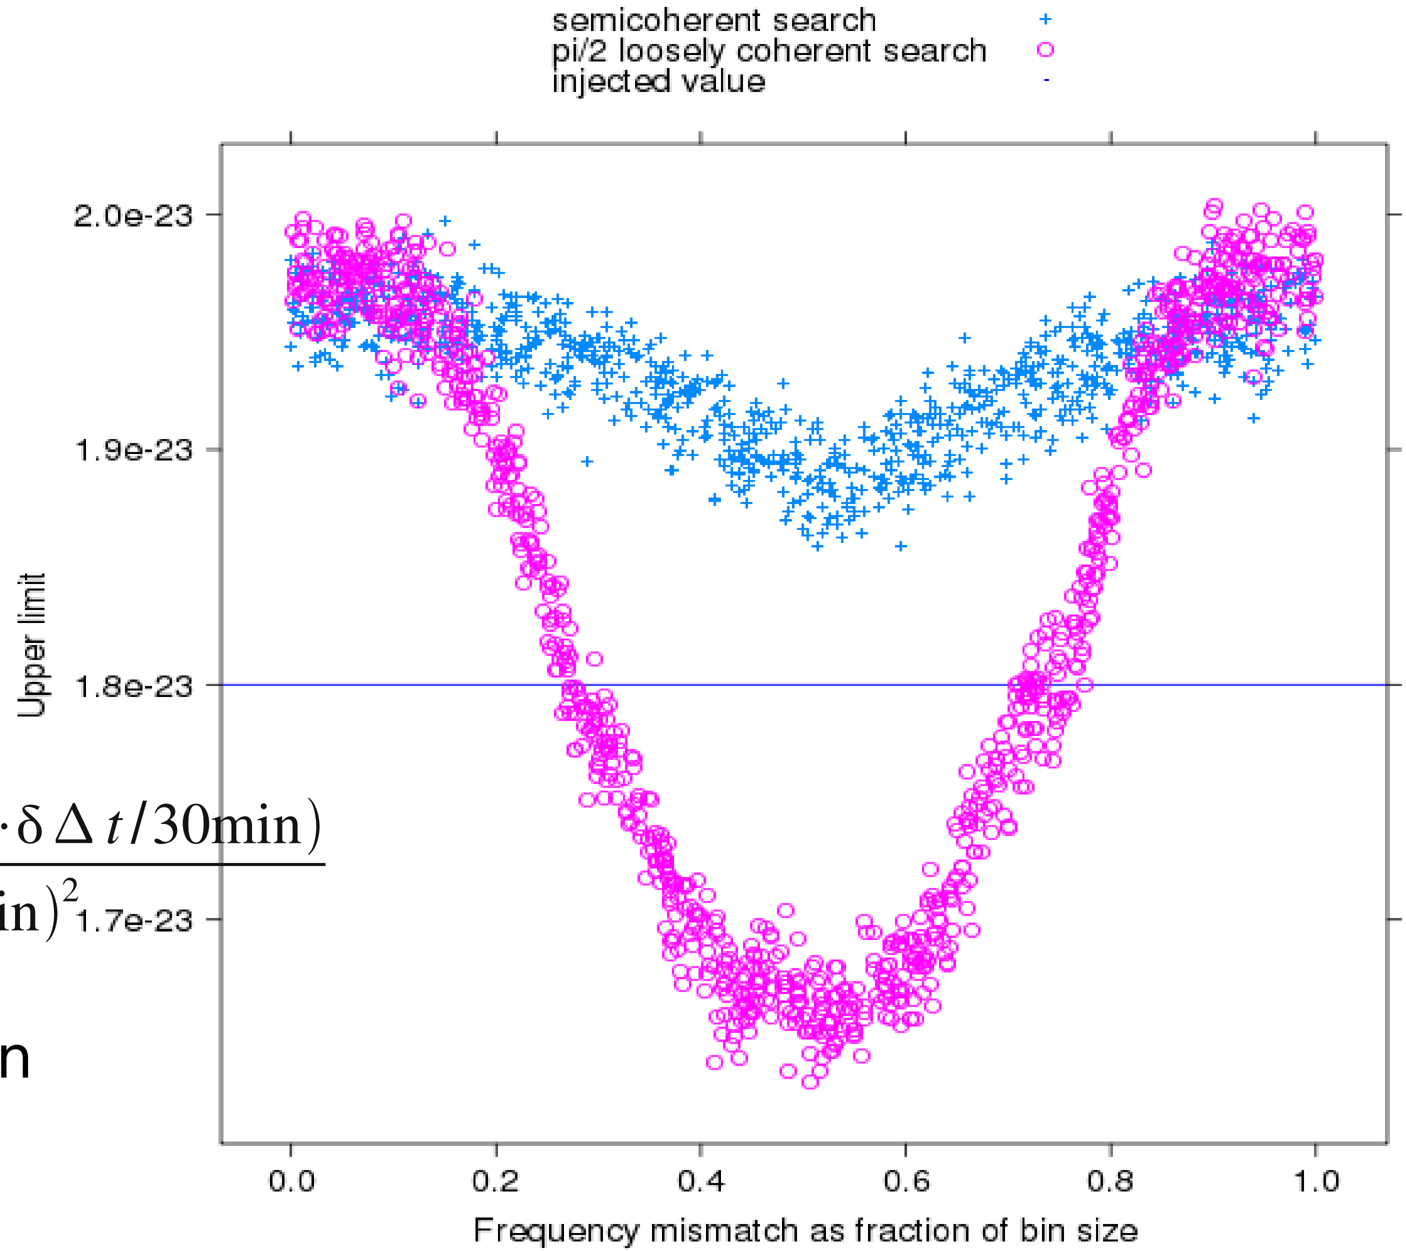

SNR improvement

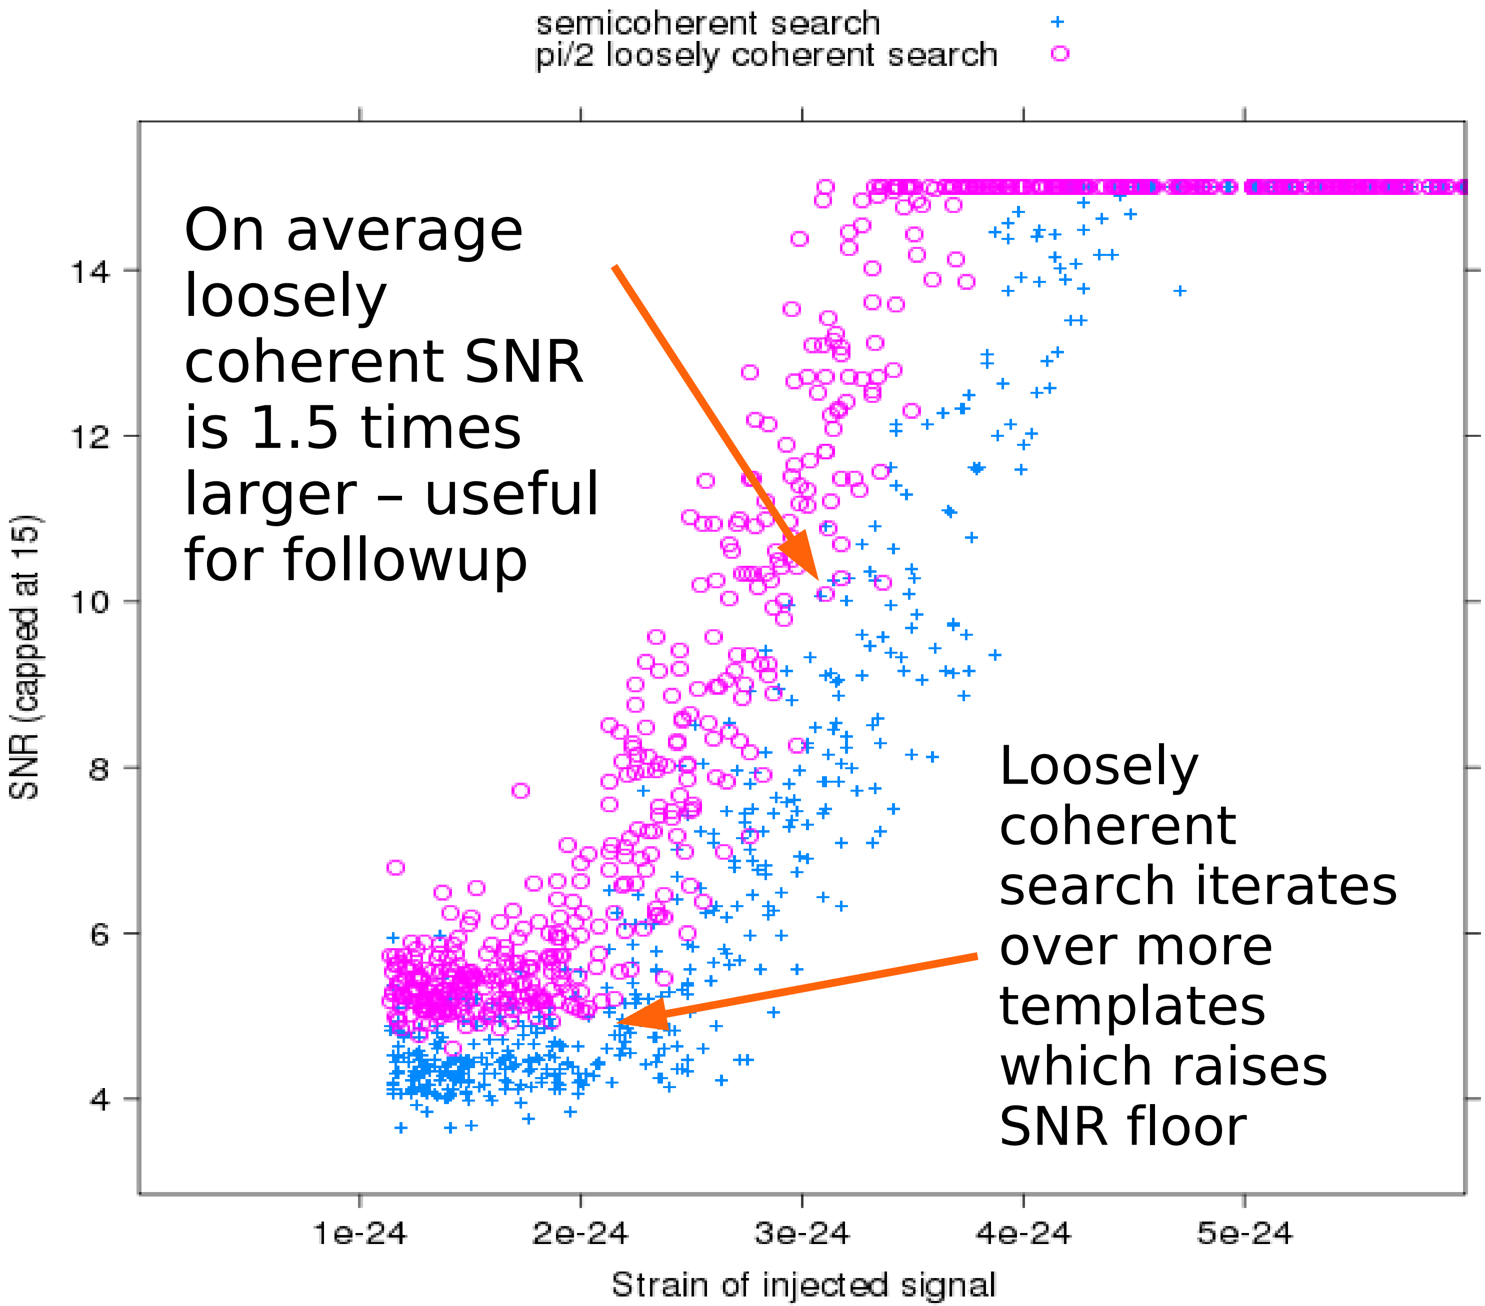

FPGA implementation ?

Modern FPGAs are very powerful, but cannot compete in floating point computations with CPUs.

However, we are looking for noise dominated signals – no reason to use 24 bits precision.

One 0.25 Hz bin can be efficiently encoded as 8 bits, 1 Month of data – 600 KB.

Similar to GPUs there are I/O issues.

But more flexibility with logic.

Stable development environment – Verilog or VHDL.

FPGA board example ($225)

32 MB
100 MHz
memory

1.6 million
gate FPGA

100 Mbit
Ethernet