MMCW workshop, NIKHEF, 2023-07-11

DCC: LIGO-G2300970-v3

Continuous Waves – a review so far

David Keitel (IAC3 / UIB)
for the LIGO-Virgo-KAGRA collaboration

Universitat de les Illes Balears | IAC3 Institute of Applied Computing & Community Code. | LIGO VIRGO KAGRA

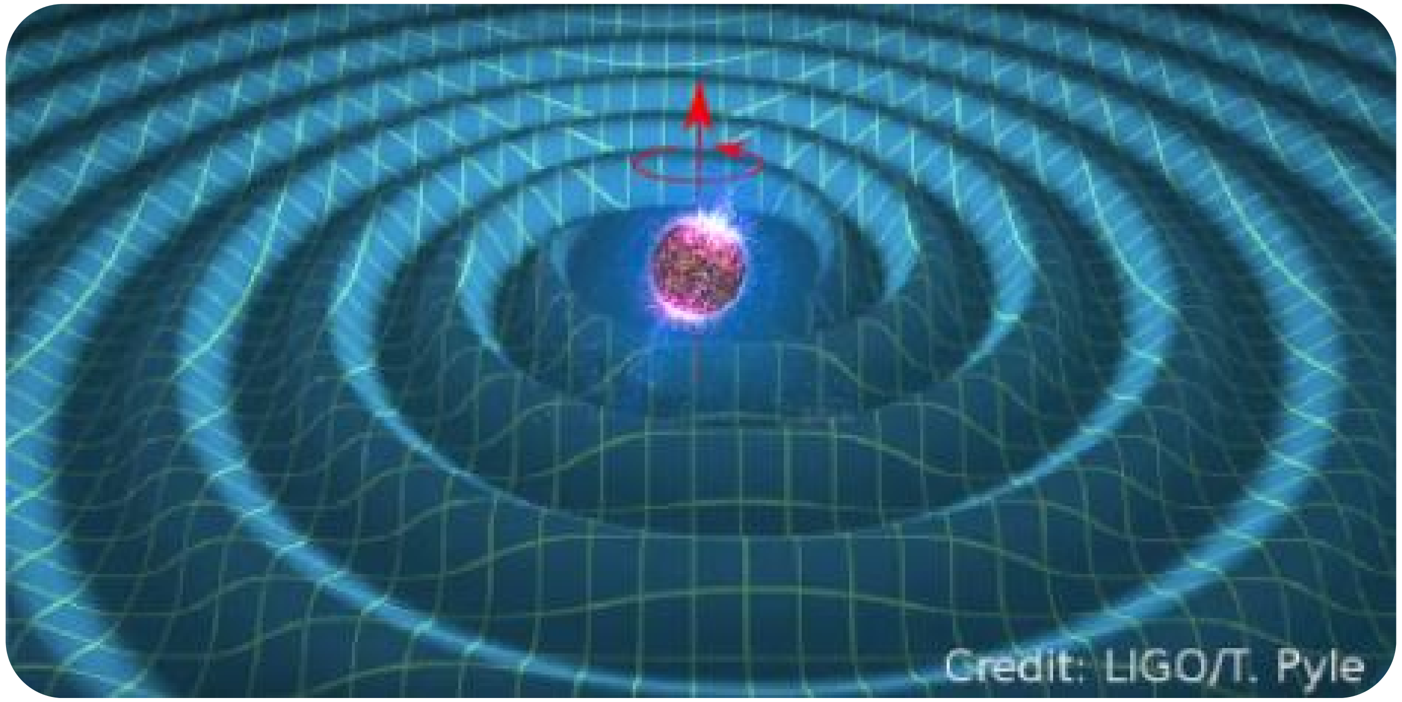

beyond CBCs: GW signal types

(for LVK network)

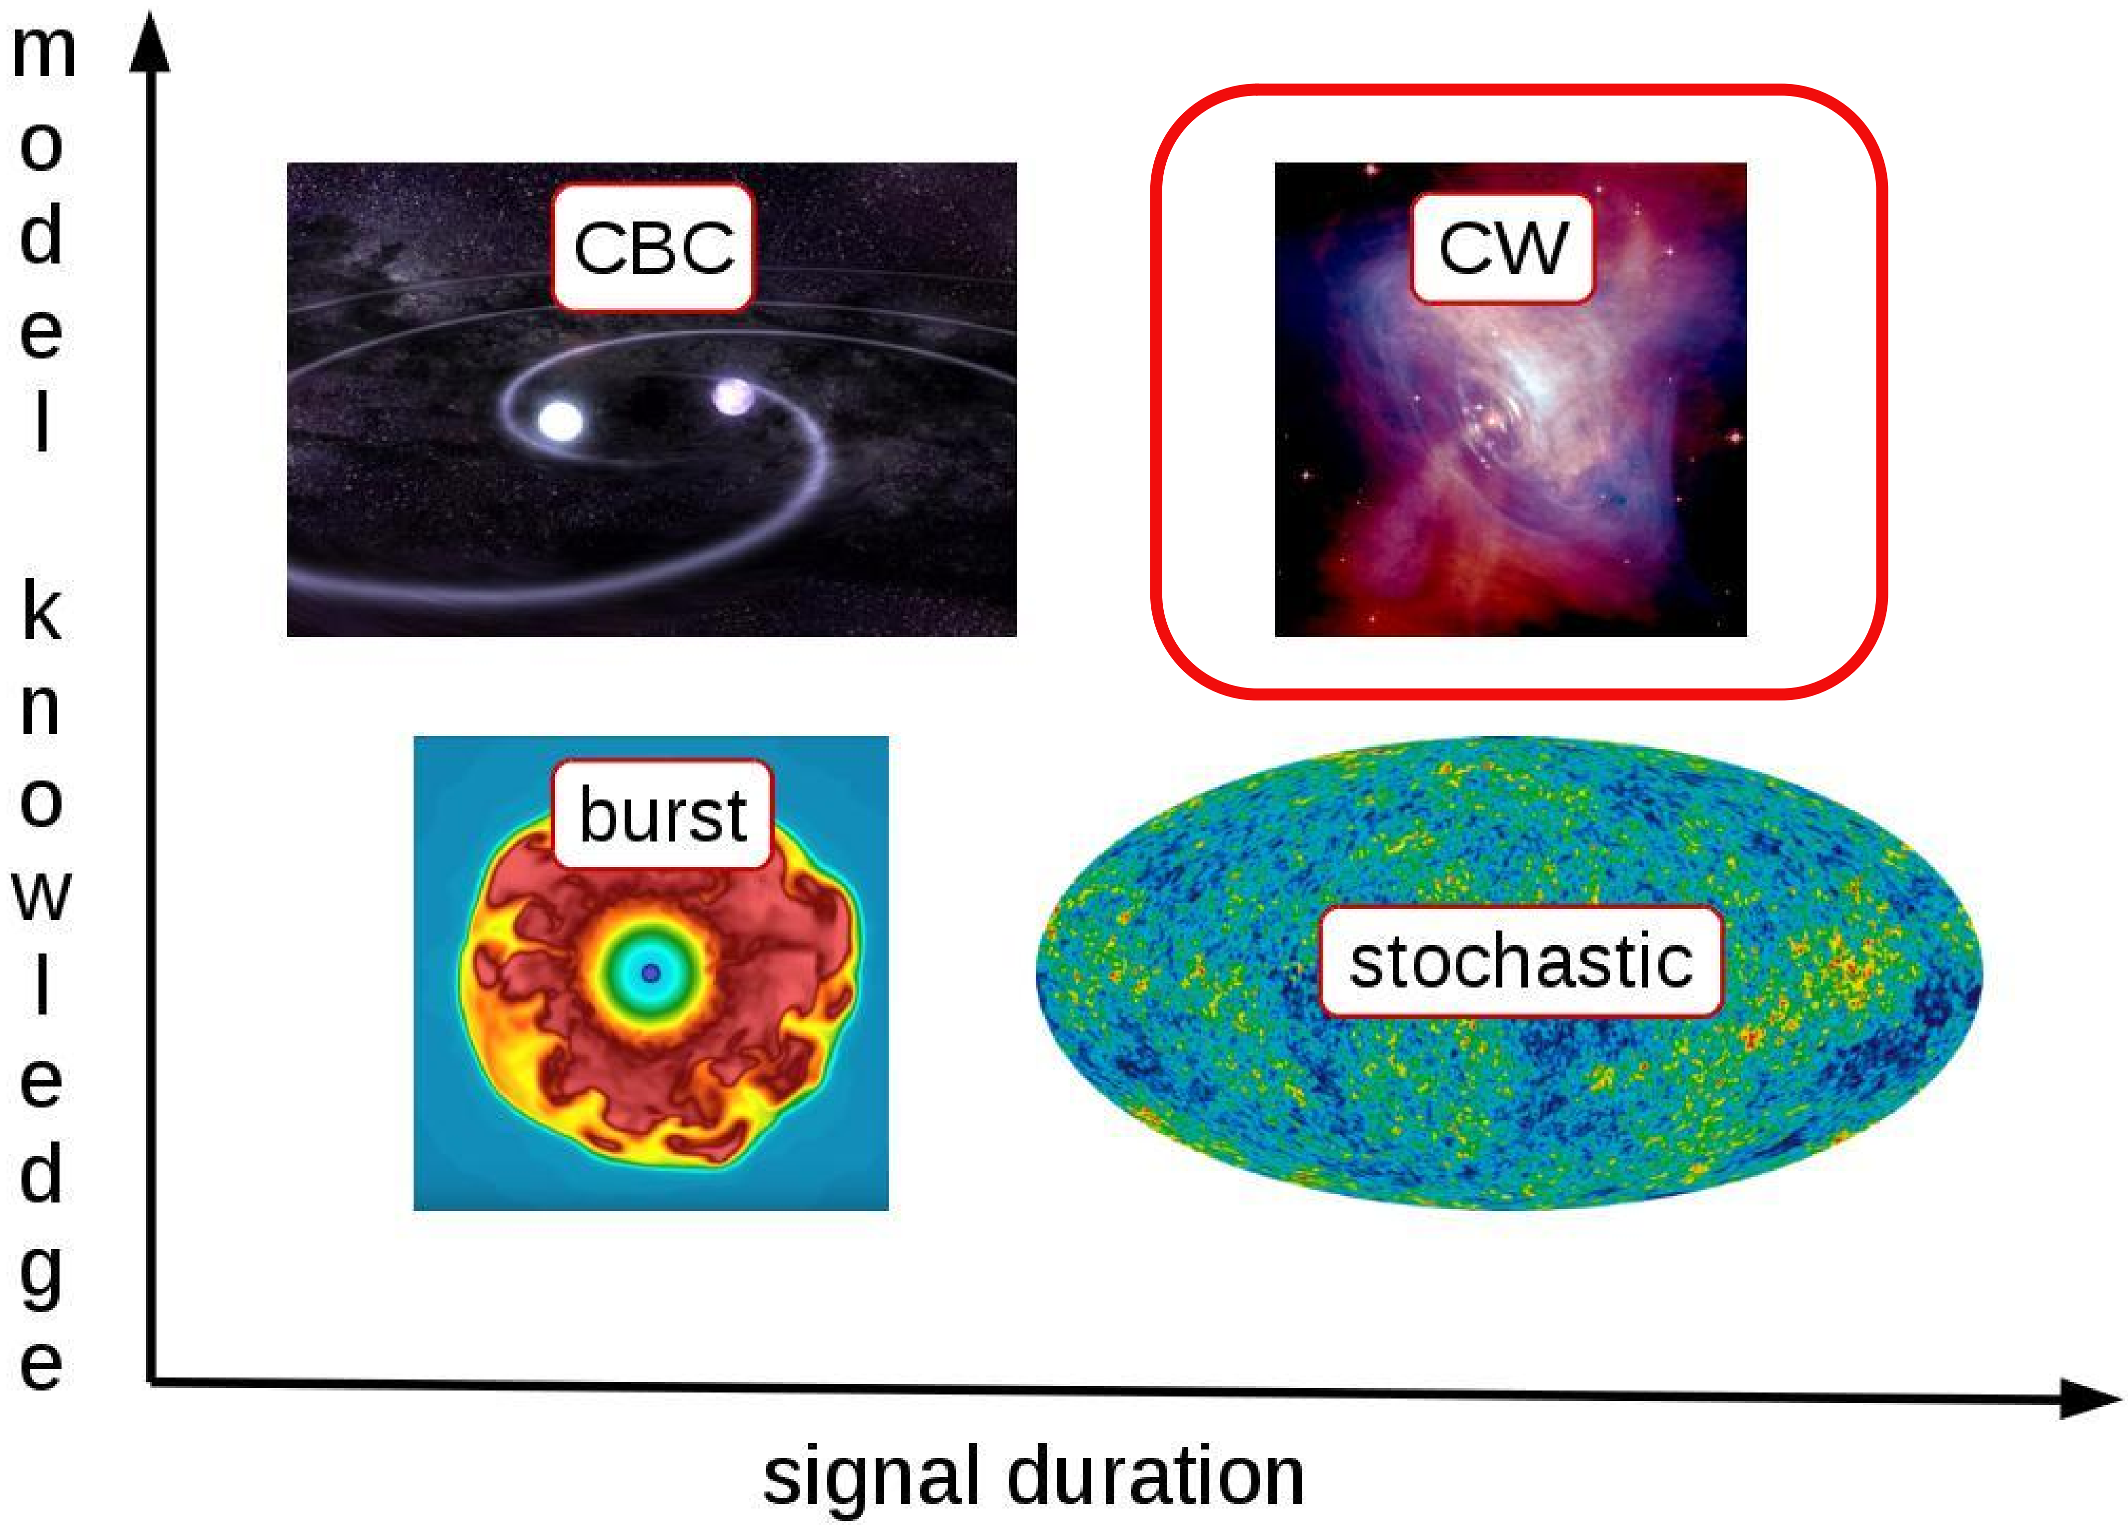

CW sources

- Long-duration signals with steady frequency evolution
  → non-catastrophic emitters with stable quadrupolar deformations,
  but still need to be rotating/moving fast enough to emit in LVK detector band
- Prime candidates: spinning neutron stars with non-axisymmetric deformations ("mountains" of cm or smaller size).

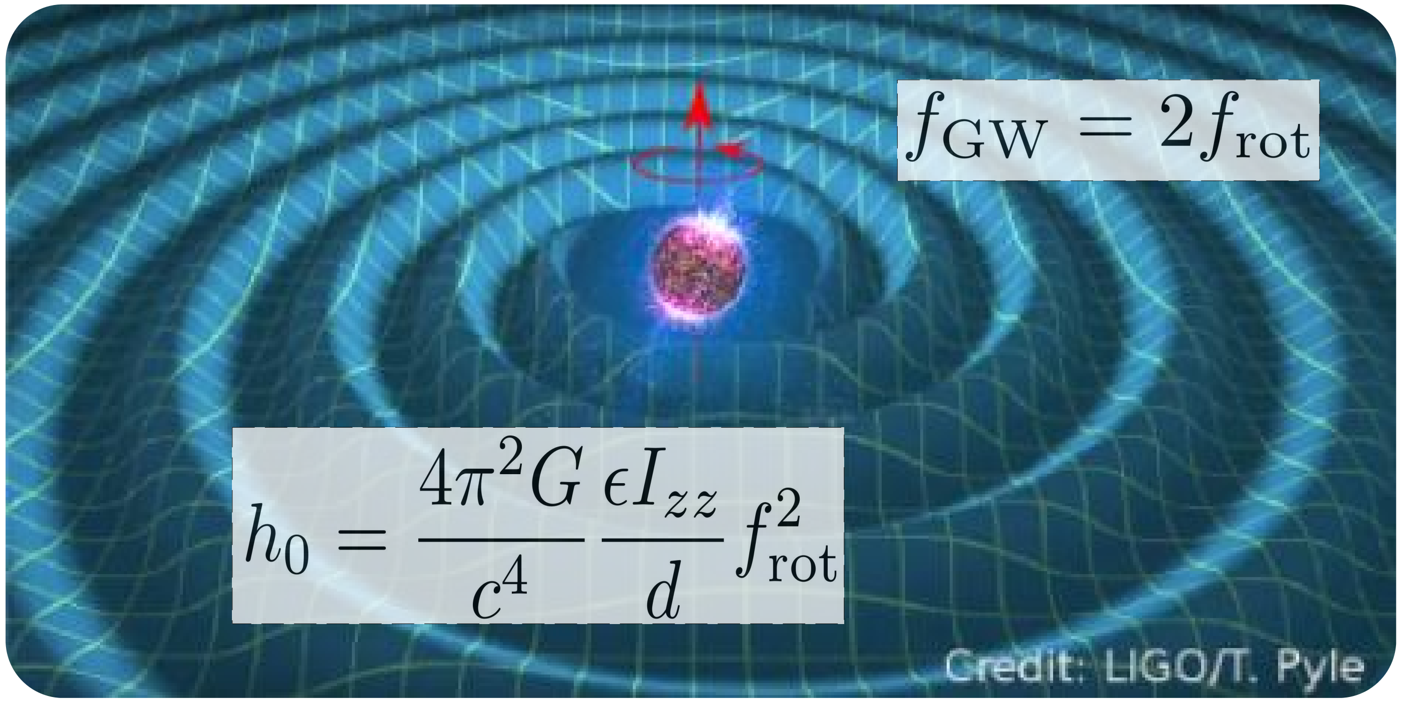

- This makes CWs a promising novel probe of these astrophysical laboratories of nuclear physics at extreme densities.
- Other NS emission channels: global oscillations (e.g. r-modes), free precession (different f_{rot}-to-f_{GW} scaling).
- Long CW-like transients from newborn neutron stars or pulsar glitches.
- Other sources: exotic physics such as boson clouds around spinning black holes, or early inspiral of low-mass compact binaries (e.g. primordial black holes).

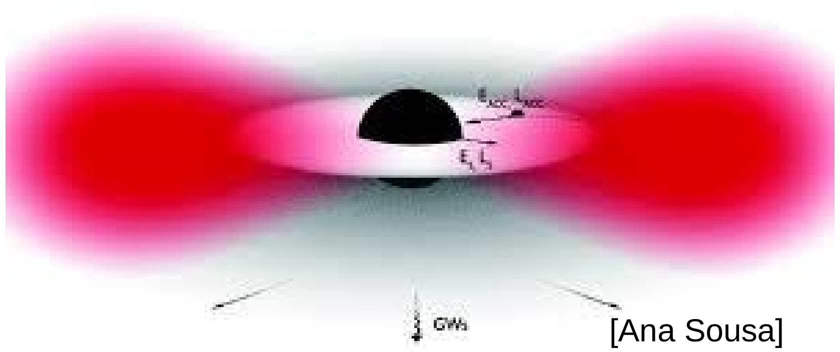

[Ana Sousa]

- → listen to contributed talks this week for the actual source modelling and astrophysics/ nuclear physics inference potential!

a brief review of (recent) reviews

universe [2019] MDPI

Review

Continuous Gravitational Waves from Neutron Stars: Current Status and Prospects

Magdalena Sieniawska * and Michał Bejger

Chapter 1

Handbook of Gravitational Wave Astronomy: Isolated Neutron Stars [2020]

Brynmor Haskell and Kai Schwenzer

universe [2021] MDPI

Review

Search Methods for Continuous Gravitational-Wave Signals from Unknown Sources in the Advanced-Detector Era

Rodrigo Tenorio *, David Keitel and Alicia M. Sintes

galaxies [2022]

Review

Status and Perspectives of Continuous Gravitational Wave Searches

Ornella Juliana Piccinni [1,2]

Living Reviews in Relativity (2023)26:3 [2023]
https://doi.org/10.1007/s41114-023-00044-3

REVIEW ARTICLE

Searches for continuous-wave gravitational radiation

Keith Riles[1]

clear signs of a mature
but vibrant field

Searches for continuous gravitational waves from neutron stars:
A twenty-year retrospective [2023]

Karl Wette[a,b]

CW signals

- over short times, signal looks like perfect single-frequency sinusoid (“pure tone”) : “quasi-stationary” and “monochromatic”.
- Slow frequency & amplitude modulation:
  - intrinsic pulsar spin-down (energy loss)
  - daily rotation of Earth
  - yearly orbit of Earth around Sun
  - optional: source binary orbit
- Still much simpler than CBC waveforms! We usually do not need to care about simulations-informed waveform models.
- But the longer the signal, the more sensitive a matched filter becomes to tiny offsets in template parameters.
- → searches (for unknown sources) very computationally expensive: have to cover the parameter space (frequency, frequency derivatives, sky location) extremely densely with templates, up to 10^{17} in all-sky searches!

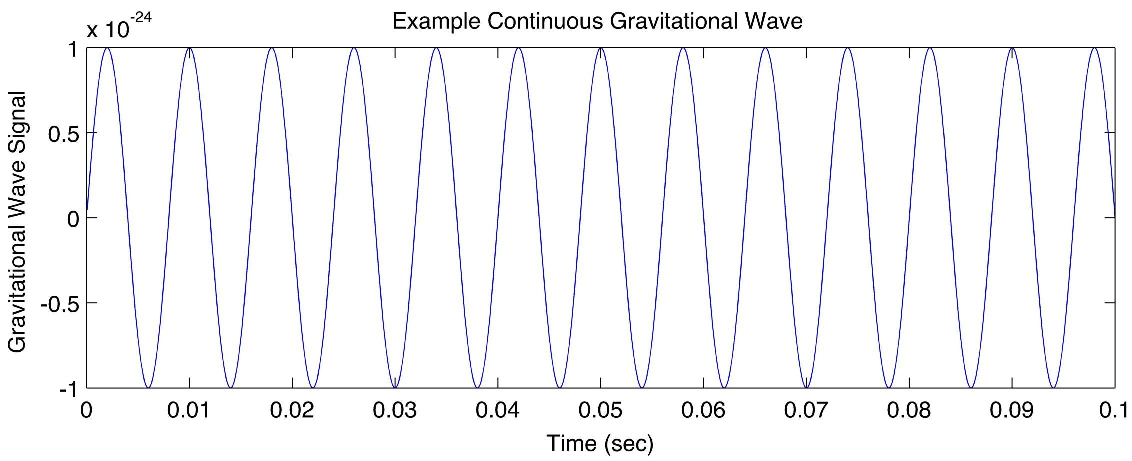

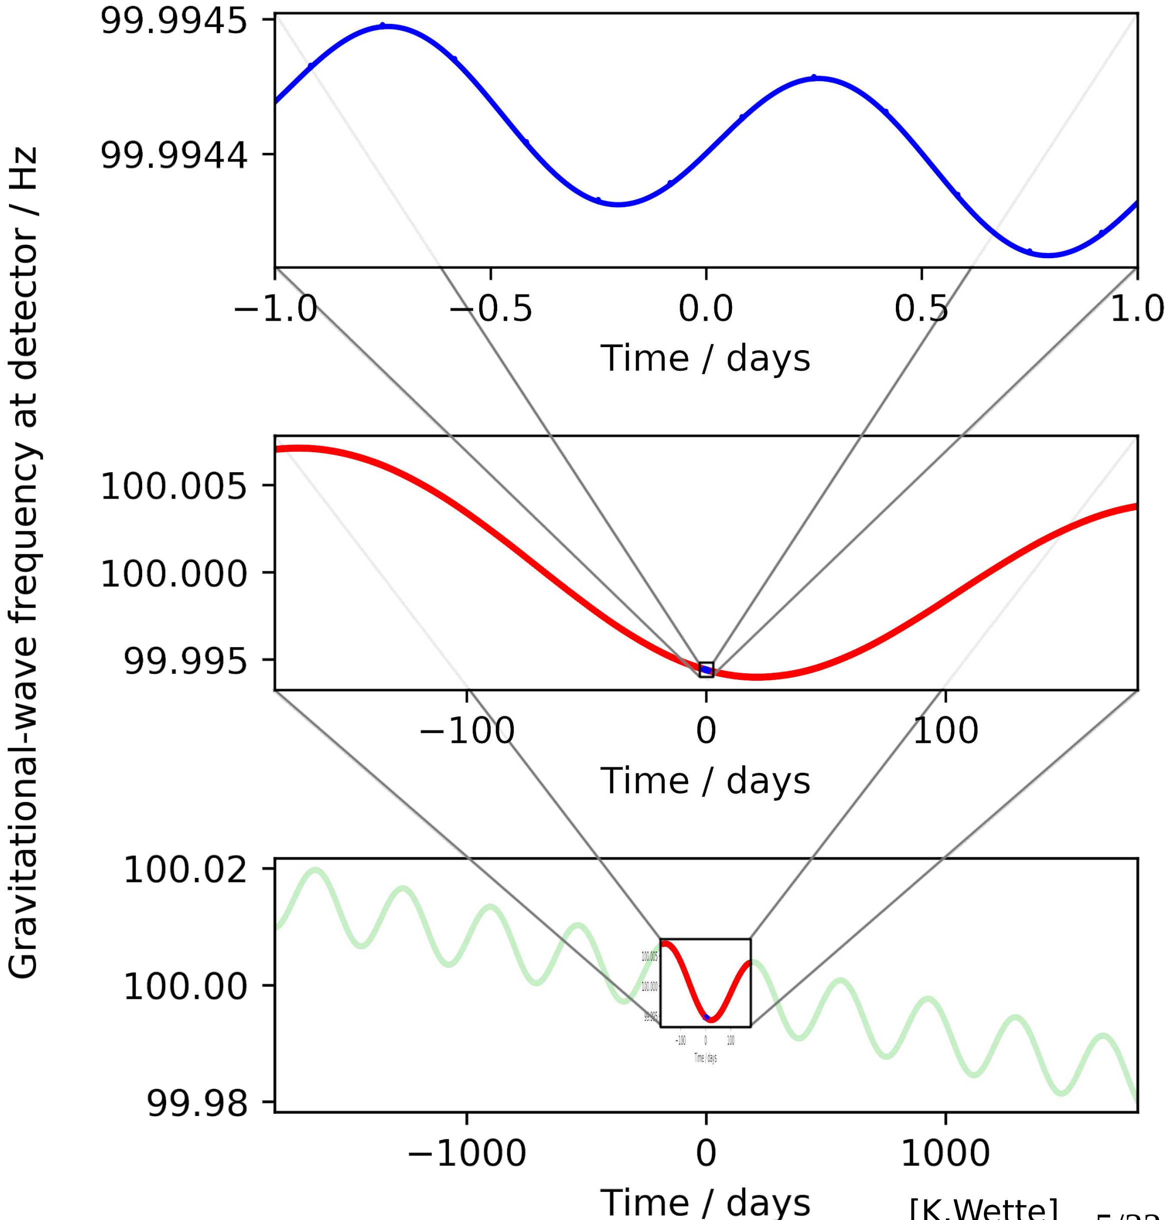

[K.Wette]

CW signals: Taylor series spindown model

CW signal emitted by a spinning deformed NS:

- the usual two polarization components:

$$h_+(\tau) \equiv A_+ \cos\Phi(\tau)\,, \quad h_\times(\tau) \equiv A_\times \sin\Phi(\tau)$$

- phase evolution:

$$\Phi(\tau) = \phi_0 + 2\pi \sum_{s=0}^{s_{\max}} \frac{f^{(s)}(\tau_{\mathrm{ref}})}{(s+1)!} (\tau - \tau_{\mathrm{ref}})^{s+1}$$

with $f^{(s)}(\tau_{\mathrm{ref}}) \equiv \left.\frac{\mathrm{d}^s f(\tau)}{\mathrm{d}\tau^s}\right|_{\tau_{\mathrm{ref}}}$

- Same as in radio timing, just at GW frequencies, e.g. $f_{\mathrm{GW}} = 2f_{\mathrm{rot}}$ for “mountains”

  (→ Andrea Possenti’s talk tomorrow)

- If the pulsar has negligible proper motion *and* we had an ideal omnidirectional detector at the solar system barycenter, this would be all there is to it!

- *BUT* need to take into account actual detector response, and timing corrections between SSB and detector frame.

CW signals: Doppler modulation and detector response

- real GW detectors on Earth: Doppler modulation from daily&yearly motion
- can be expressed as timing relation between wavefront arrivals in detector frame and in SSB:

$$\tau(t;\boldsymbol{n},\mathbf{b}) = t + \frac{\boldsymbol{r}(t)\cdot\boldsymbol{n}}{c}$$

- modulated signal waveform at detector:

$$h(t;\mathcal{A},\boldsymbol{\lambda}) = F_+(t;\boldsymbol{n},\psi)\,A_+\cos\left[\phi_0+\phi(t;\boldsymbol{\lambda})\right] + F_\times(t;\boldsymbol{n},\psi)\,A_\times\sin\left[\phi_0+\phi(t;\boldsymbol{\lambda})\right]$$

(including detector response / antenna pattern)

- CW signal frequency evolution parameters ("Doppler parameters", λ):

  intrinsic spindown terms,
  sky position (alpha,delta)

- correction for this effect in data analysis also called "barycentring" and is a main cost factor

→ Andrzej Krolak's talk tomorrow

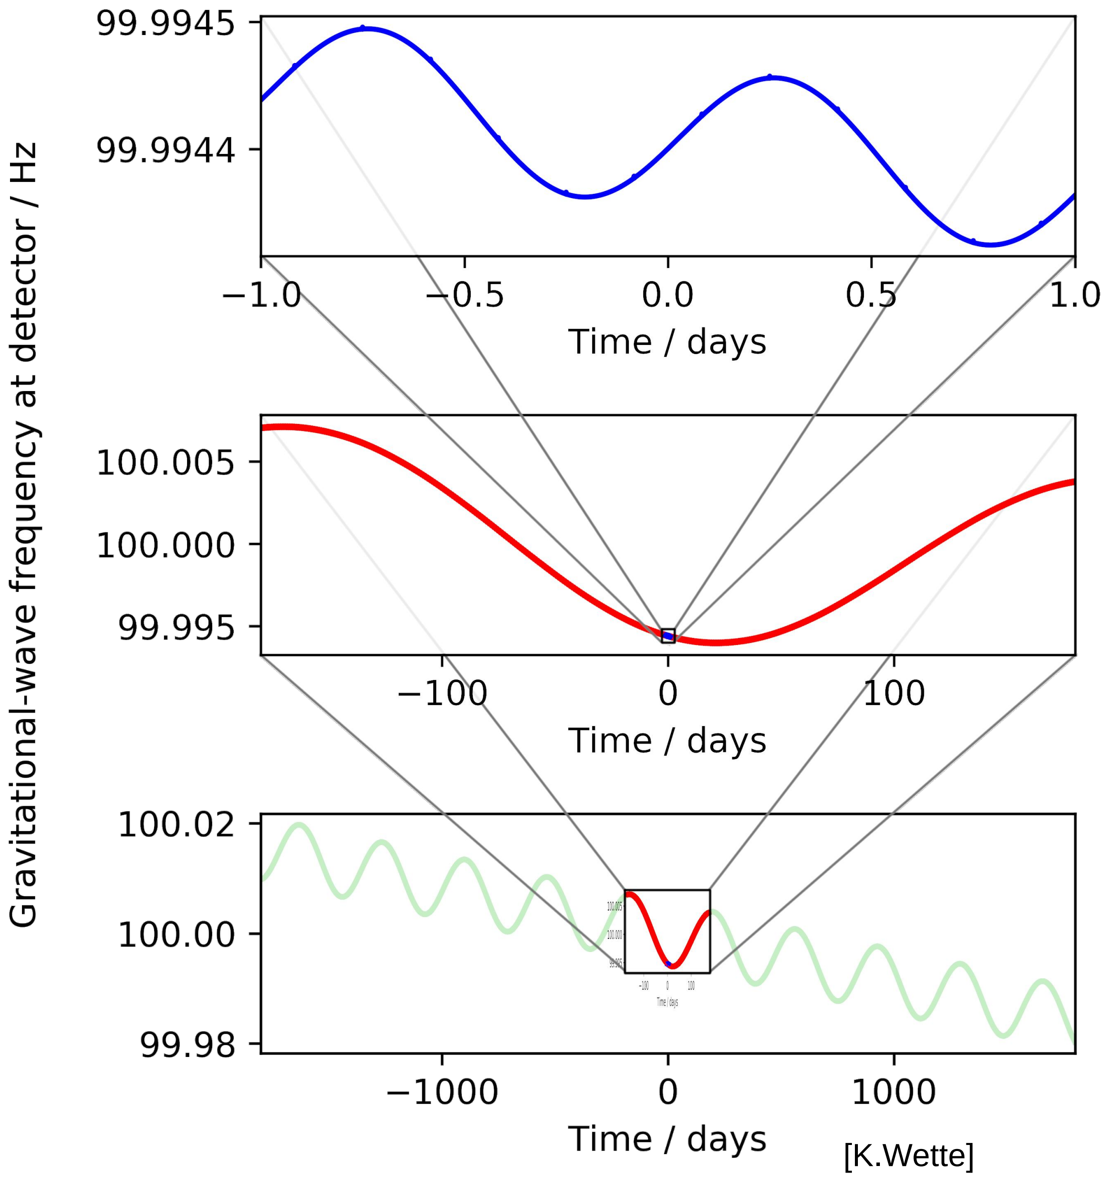

[K.Wette]

detector noise

- To first approximation, GW detector noise is ***Gaussian*** (especially when averaging over long durations).
- Fully described by Power Spectral Density (PSD).
- Frequency-dependent PSD (“coloured noise”):

$$S_{\mathrm{n}}(f) \approx \frac{2}{T_{\mathrm{obs}}} E\left[|\widetilde{n}(f)|^2 \right]$$

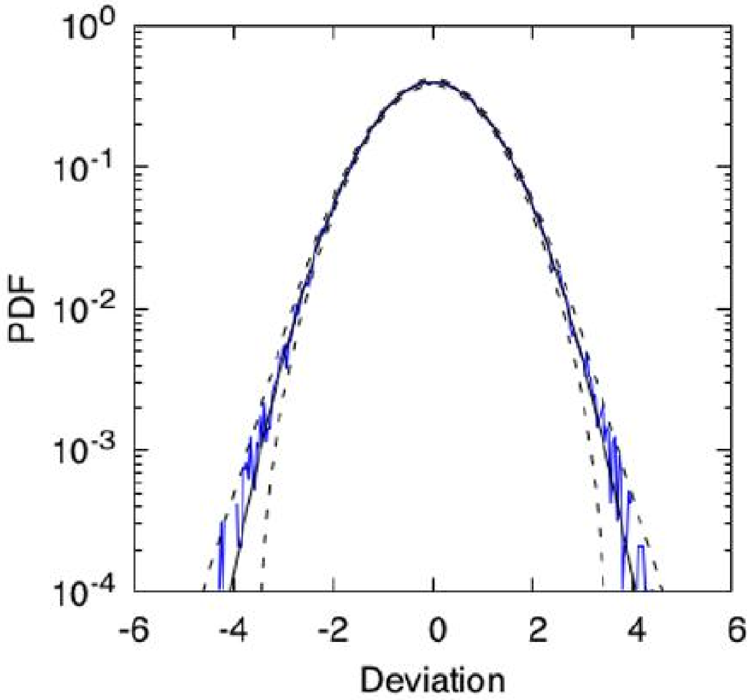

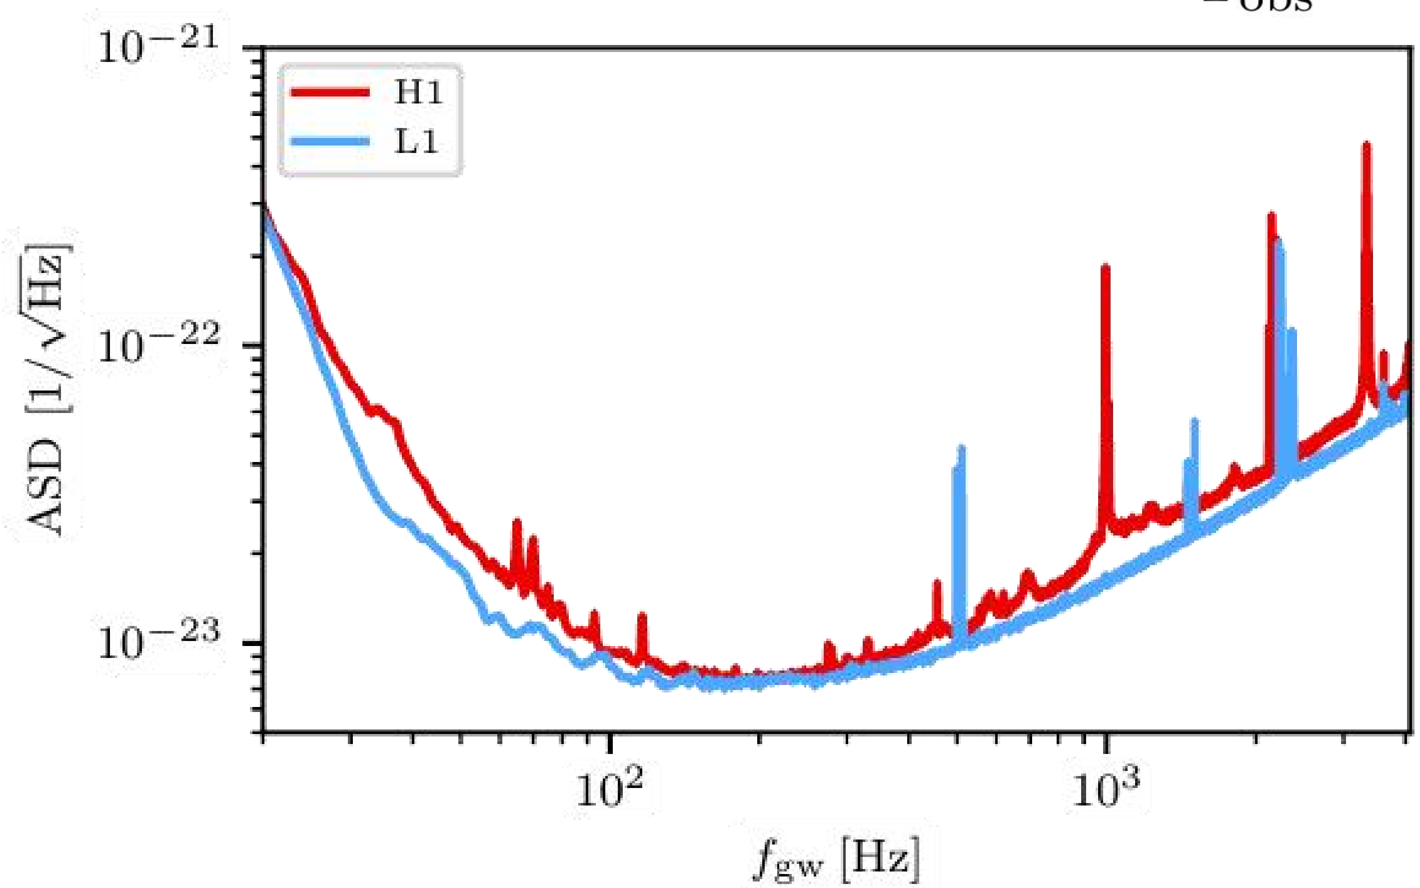

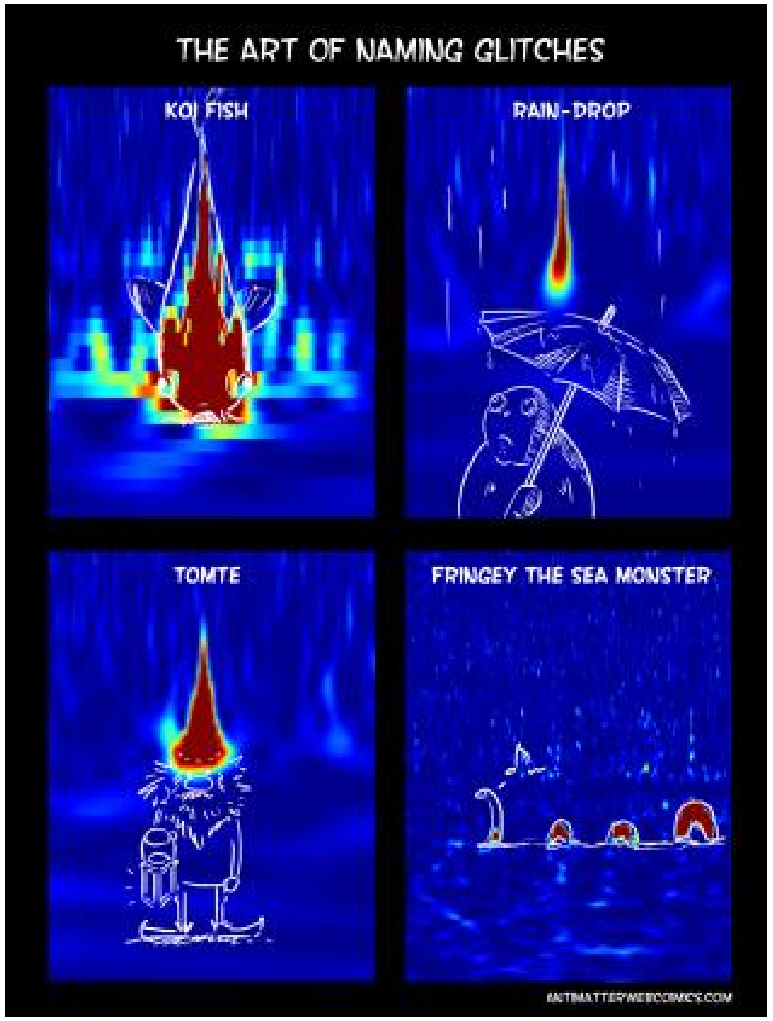

Real noise ***not*** perfectly Gaussian, contains artifacts like

– *glitches* (short duration, complex shapes)

– *lines* (fixed frequency, can be persistent, main CW headache)

→ Ansel Neunzert’s talk Thursday

detector noise & CW data analysis

- For quasimonochromatic CW signals, we usually work in the Fourier domain.
- Allows us to extract "narrowband" data sets, and assume the noise PSD is almost constant over the range of interest:

$$(\boldsymbol{x}|\boldsymbol{y}) \approx 2 \sum_{X}^{N_{\mathrm{Det}}} S_X^{-1}(f_{\mathrm{s}}) \int_0^T x^X(t)\, y^X(t)\, dt$$

→ timeseries inner product, related to matched filter
→ Andrzej Krolak's talk tomorrow

- Data is usually split up into Short Fourier Transforms (SFTs), typically of T_{SFT}=1800s or similar. Then we only have to assume the PSD is constant over each SFT:

$$\langle \boldsymbol{x} \,|\boldsymbol{y}\,\rangle \approx 2 \sum_{X=1}^{N_{\mathrm{det}}} \sum_{\alpha=1}^{N_{\mathrm{SFT}}^X} \frac{1}{S_\alpha^X} \int_0^{T_{\mathrm{SFT}}} x_\alpha^X(t')\, y_\alpha^X(t')\, \mathrm{d}t'$$

(Virgo groups use different, but conceptionally similar formats)

- PSD can be estimated from periodogram of the per-SFT data, and averaged over longer durations:

$$\widehat{S}^X(f') \equiv \frac{1}{N_{\mathrm{SFT}}^X} \sum_{\alpha=1}^{N_{\mathrm{SFT}}^X} \frac{2\left|\tilde{x}_\alpha^X(f')\right|^2}{T_{\mathrm{SFT}}}$$

CWs: the key points

- Quasi-monochromatic: the templates are simple.
- Incredibly weak.
- Long-duration:
  - We can gain SNR by integrating longer.
  - Data becomes very close to Gaussian (except near narrow disturbances).
  - Very precise frequency resolution from long-term phase coherence requirement.
  - Very precise sky localisation, even with a single detector, because the Earth moves during the observation.
  - Computational cost for unknown targets grows steeply with observing time (or, more precisely, with *coherence time* – more later).
    - This is not mainly because of the cost of a single long matched filter ($\sim T$).
    - Mainly because the template bank to cover a certain parameter space at an acceptable mismatch becomes extremely dense over long periods.
    - This is the logical flip-side of getting the great resolution benefit.

CW searches

- And now on to practical applications and results!
- Categorisation by amount of prior information:
  - targeted searches
  - narrowband searches
  - directed searches
  - spotlight searches
  - allsky searches

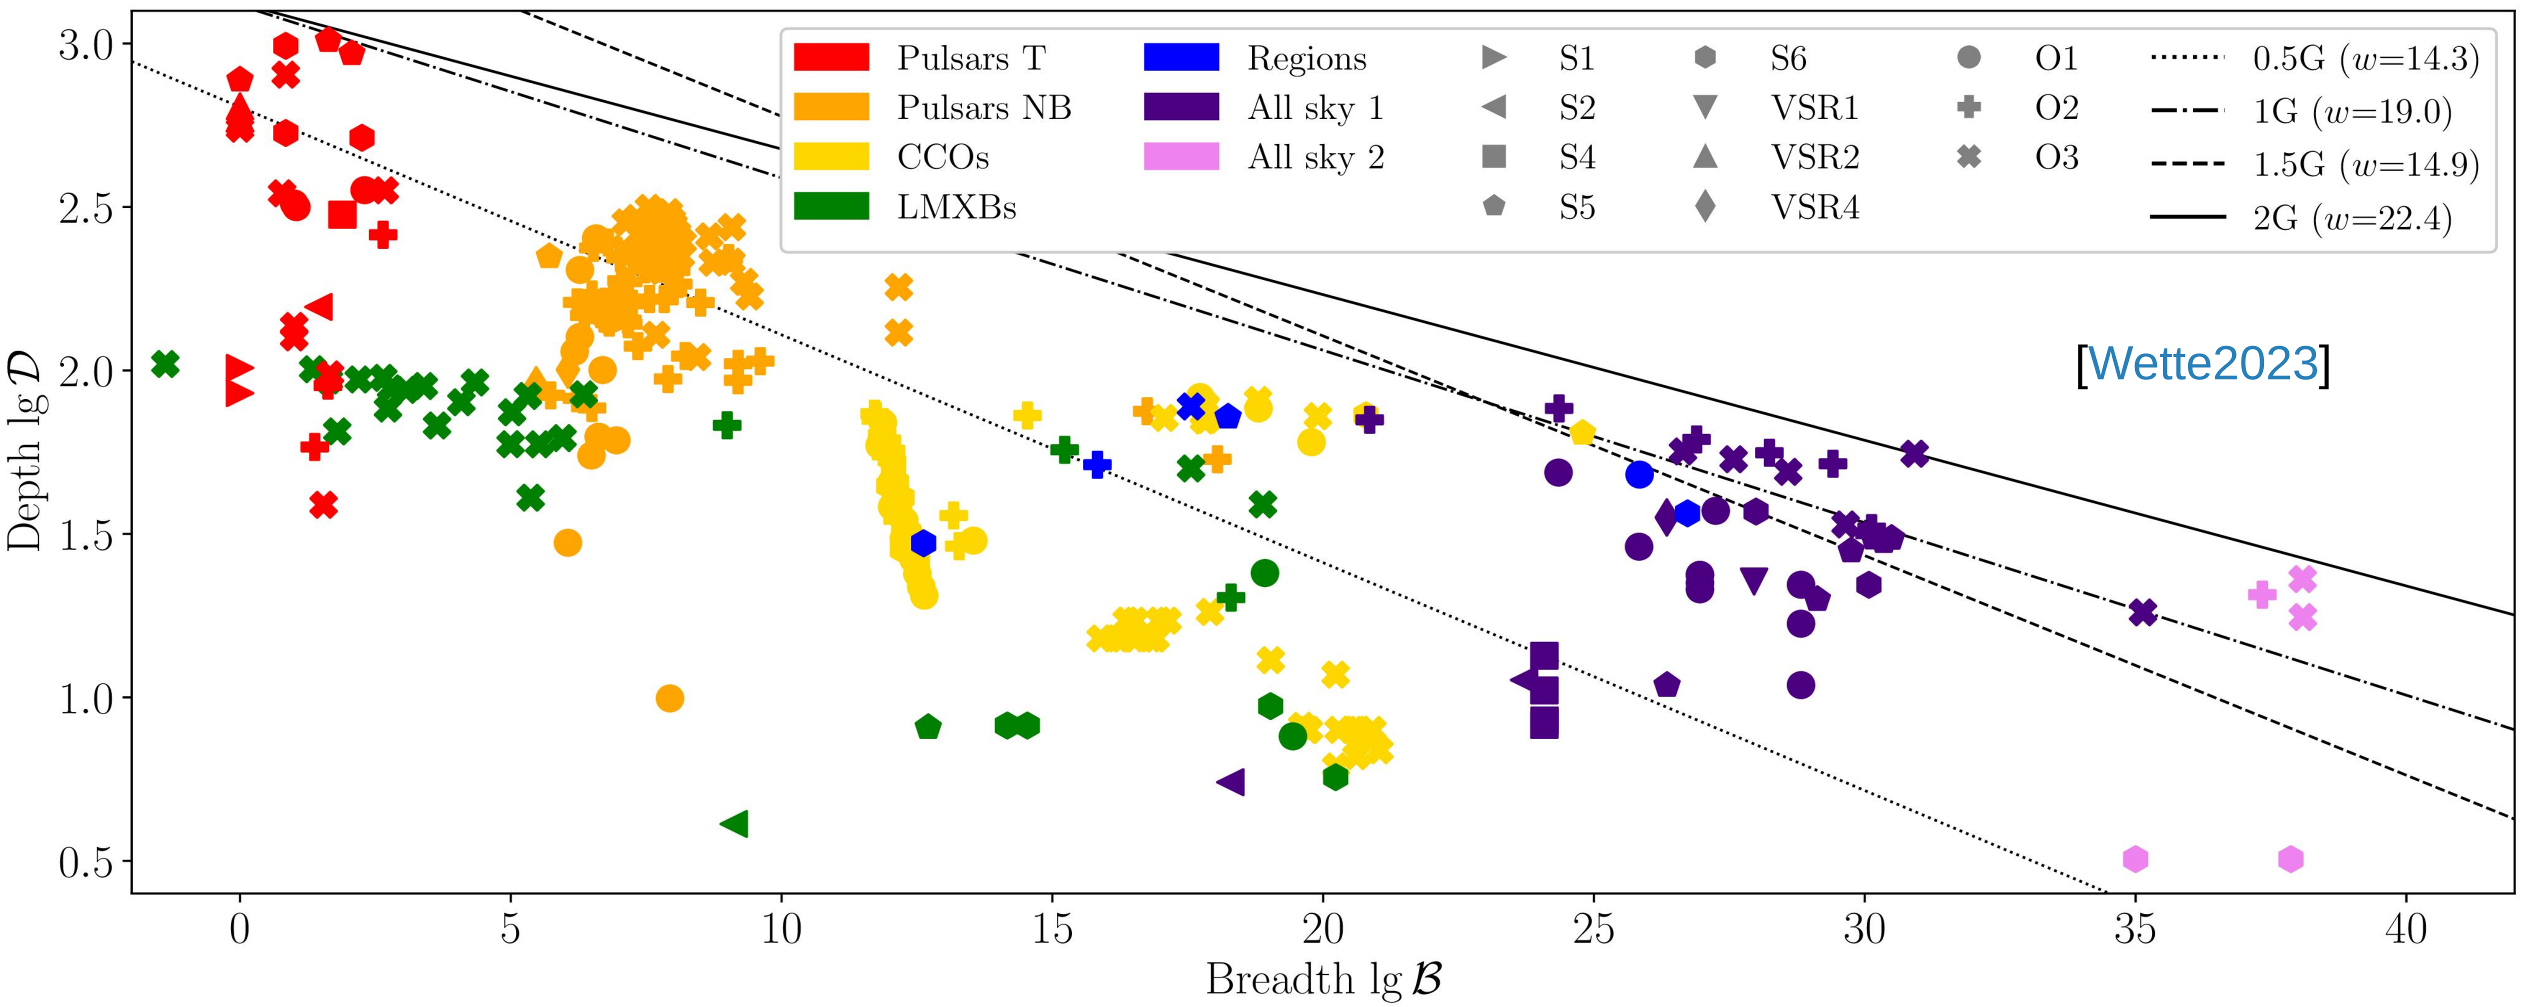

targeted searches

- Detailed ephemerides of known pulsars from radio, X-ray, gamma-ray observations ((→ Andrea Possenti's talk) allow cheap and very sensitive ***fully-coherent analysis***: assume that the Taylor expansion signal model holds for the full GW data set, without deviations between rotation and GW frequencies, and without any phase jumps.
- Crucial milestone for each target: indirect ***spindown upper limit*** assumes all energy loss into GWs:

$$h_0 \leq h_{\rm sd} = \frac{1}{d}\sqrt{\frac{5G\,I_{zz}}{2c^3}\,\frac{|\dot{\nu}|}{\nu}}$$

  In the absence of statistically significant detections, GW searches produce ***observational upper limits*** on the actual strain at the detector:
  if the source were emitting h_0 above this threshold, we would have detected a louder outlier in our analysis with e.g. 90% confidence
- For Crab and Vela pulsars, the spindown limit was already beaten with initial LIGO/Virgo in the 2000s.

THE ASTROPHYSICAL JOURNAL, 683: L45–L49, 2008 August 10
© 2008. The American Astronomical Society. All rights reserved. Printed in U.S.A.

BEATING THE SPIN-DOWN LIMIT ON GRAVITATIONAL WAVE EMISSION FROM THE CRAB PULSAR

targeted searches: results

- Upper limits on 236 targets, at both $f_{\rm gw} = 2f_{\rm spin}$, $f_{\rm gw} = f_{\rm spin}$
  Abbott+2021, "Searches for Gravitational Waves from Known Pulsars at Two Harmonics in the Second and Third LIGO-Virgo Observing Runs"

- Glasgow time-domain Bayesian method [Dupuis&Woan2005, Pitkin+2017]
- Warsaw time-domain *F*-statistic [Jaranowski-Królak-Schutz1998, Jaranowski&Królak2010]
- Rome 5n-vector method [Astone+2014, Mastrogiovanni+2017]
- Searches with non-GR templates can constrain non-standard polarisation content (scalar, vector modes: Isi+2015, Isi+2017, Verma2021)

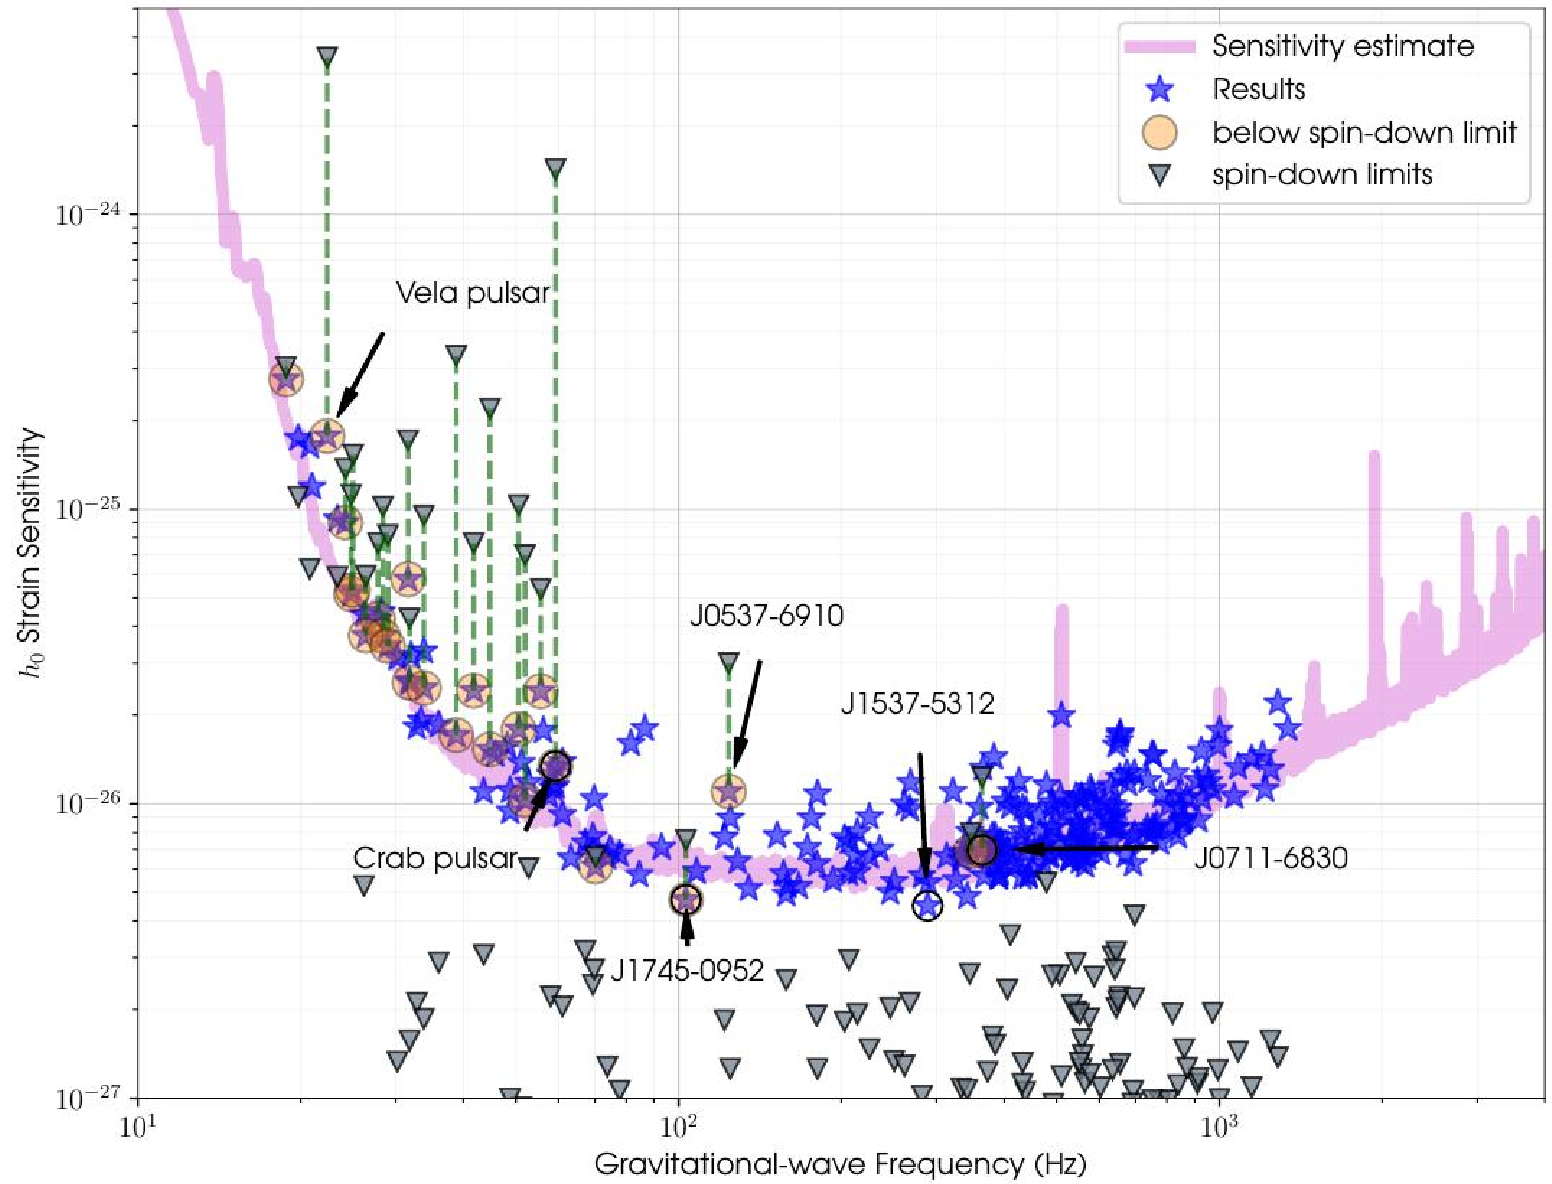

- Non-LVK searches: e.g. Nieder+2021 on Einstein@Home gamma-ray pulsars, and → see talk by Anjana Ashok tomorrow

narrow-band searches

- Still targeting known pulsars with ephemerides provided by EM observers.
- But relax EM-GW frequency equality assumption, covering the EM uncertainties and allowing for a physical EM-GW mismatch (e.g. because the EM pulsar signals come from the magnetosphere and measure the rotation of the NS surface, while GWs come from the interior which may rotate differently).
- Small template banks, typically O(10^6).
- O3 results (Abbott+2022): 18 pulsars (5n-vec method and freq-domain $\mathcal{F}$-stat), plus transient $\mathcal{F}$-stat search on 9 glitches from 6 pulsars
- J0537 O3 search (Abbott+2021): r-modes, covering f_{GW} in 86–97Hz for f_{rot} = 62 Hz, 5n-vec and time-domain $\mathcal{F}$-stat.

(earlier O1+O2 search: Fesik&Papa2021)

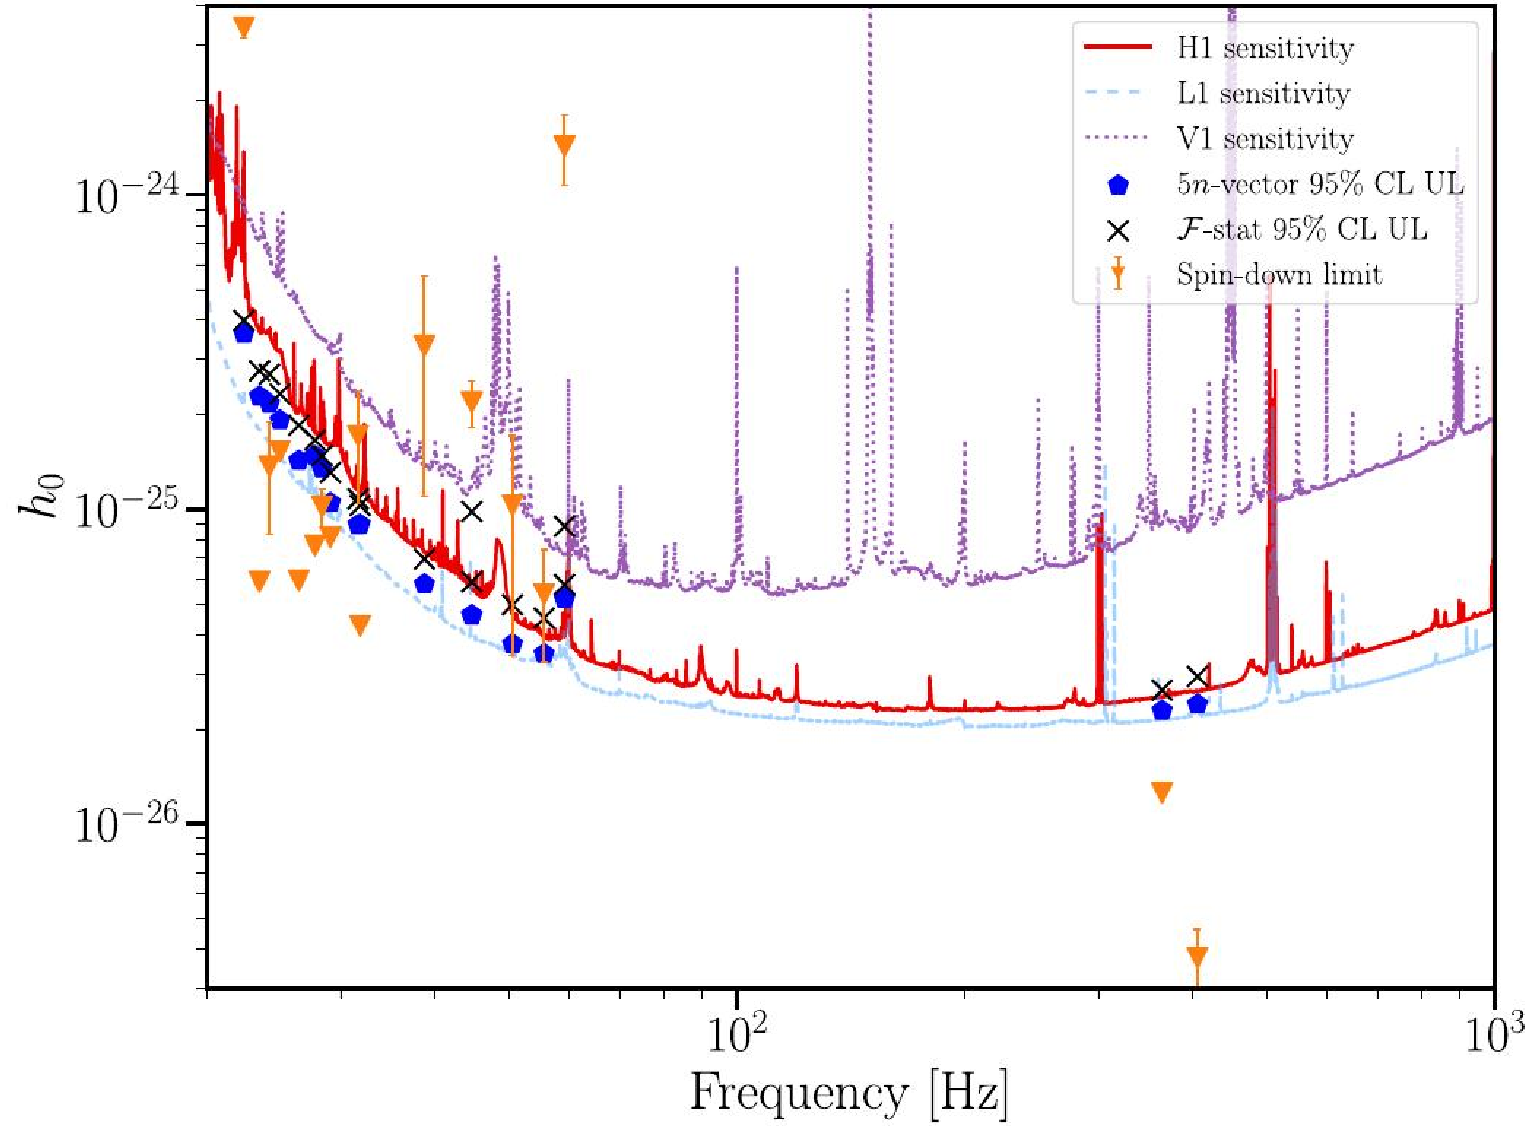

- Crab r-modes (O1+O2): Rajbhandari+2021

directed & all-sky searches

- Cover a large parameter space (frequency and spindowns, possibly sky) at affordable computational cost. Our signal model is simple and dimensionality is not too high (e.g. 4 if only one spin-down term included), but parameter space is curved and highly structured.
- metric can be computed (approximately) as time-average of phase model derivatives:

$$g_{ij} \sim \langle \partial_i\phi\, \partial_j\phi \rangle - \langle \partial_i\phi \rangle \langle \partial_j\phi \rangle \longrightarrow g_{\theta\theta} \propto f^2\, T_{\rm obs}^2\, (V/c)^2\,, \qquad g_{ff} \propto T_{\rm obs}^2\,, \qquad g_{\dot f \dot f} \propto T_{\rm obs}^4$$

  where $V/c \sim 10^{-4}$ is the max Doppler shift from Earth's motion [Prix2007]

- For a 4D parameter space $d^4\boldsymbol{\lambda} = d\Omega \times df \times d\dot{f}$, number of templates:

$$dN_p \propto \sqrt{|\det g_{ij}|}\; d^4\boldsymbol{\lambda} \propto T_{\rm obs}^5\, f^2\, d^4\boldsymbol{\lambda}$$ *(or even steeper!)*

- Computational cost $\propto T_{\rm obs}^6\, f^2\, d^4\boldsymbol{\lambda}$ (per-template MF cost is $\propto T_{\rm obs}$)
- Intuitive reason for steep scaling of N_p:
  with growing T_{obs}, tiny differences in parameters lead to total dephasing of the signal
  → growing mismatch → shrinking ellipses covered by each template point
- Larger template banks also have higher *trials factor* (chance of spurious noise outliers)
  → less significance for the same signal!

semi-coherent searches

- Fully-coherent searches are always a bad choice for broad searches with T_{obs} > months.
- Basic idea for better sensitivity at fixed cost:

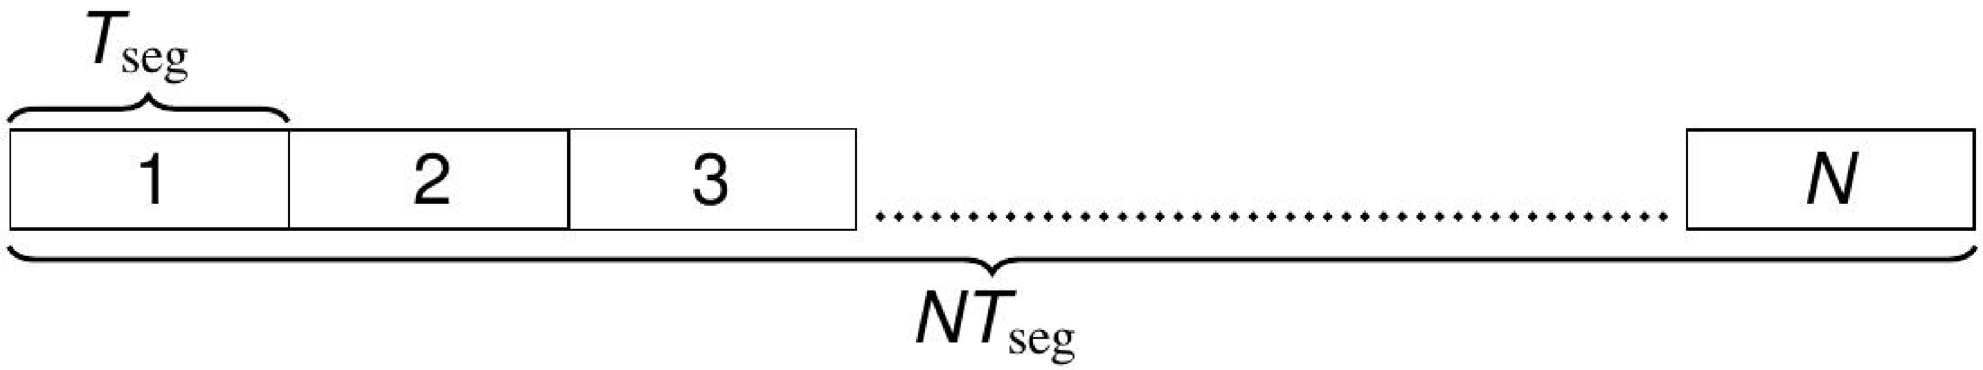

1. split data $\boldsymbol{x}$ into N shorter segments $\boldsymbol{x}_k$
2. compute coherent $\widetilde{\mathcal{F}}_k$ in each segment
3. incoherent combination, e.g. $\widehat{\mathcal{F}}(\boldsymbol{x}, \lambda) \equiv \sum_{k=1}^{N} \widetilde{\mathcal{F}}_k(\boldsymbol{x}_k, \lambda)$

- Required template density only scales with T_{seg} instead of T_{obs}.
- No longer require phase-coherence across the whole T_{obs}:
  → overall sensitivity is reduced
  → more susceptible to spurious instrumental artifacts that look like a CW in individual segments even if not across the whole run
  → but also more robust to astrophysical variations in the source (e.g. NS glitches)
- We gain so much computational efficiency → higher depth at fixed budget!
- Spurious candidates can be taken out with hierarchical follow-ups.

semi-coherent searches

- T_{seg} can be as short as a single SFT (e.g. 1800s for most "Hough" type searches) and as long as several days (for the most expensive Einstein@Home distributed computing searches)

- Instead of simple semi-coherent sum $\widehat{\mathcal{F}}(\boldsymbol{x}, \lambda) \equiv \sum_{k=1}^{N} \widetilde{\mathcal{F}}_k(\boldsymbol{x}_k, \lambda)$

  more sophisticated methods exist, e.g. using ***refinement***:

  – use "coarse grid" template banks $\{\lambda_k\}$ in each segment with resolution given by T_{seg}

  – evaluate final detection statistic on a "fine grid" $\{\lambda\}$ with resolution given by T_{obs}

  – Get that final $F(x,\lambda)$ from summing up $F_k(x_k,\lambda_k)$ along the λ time-frequency track

  – Optimal fine-grid construction and coarse/fine computational cost balancing is tricky and requires detailed understanding of parameter space structure (correlations/degeneracies between parameters).

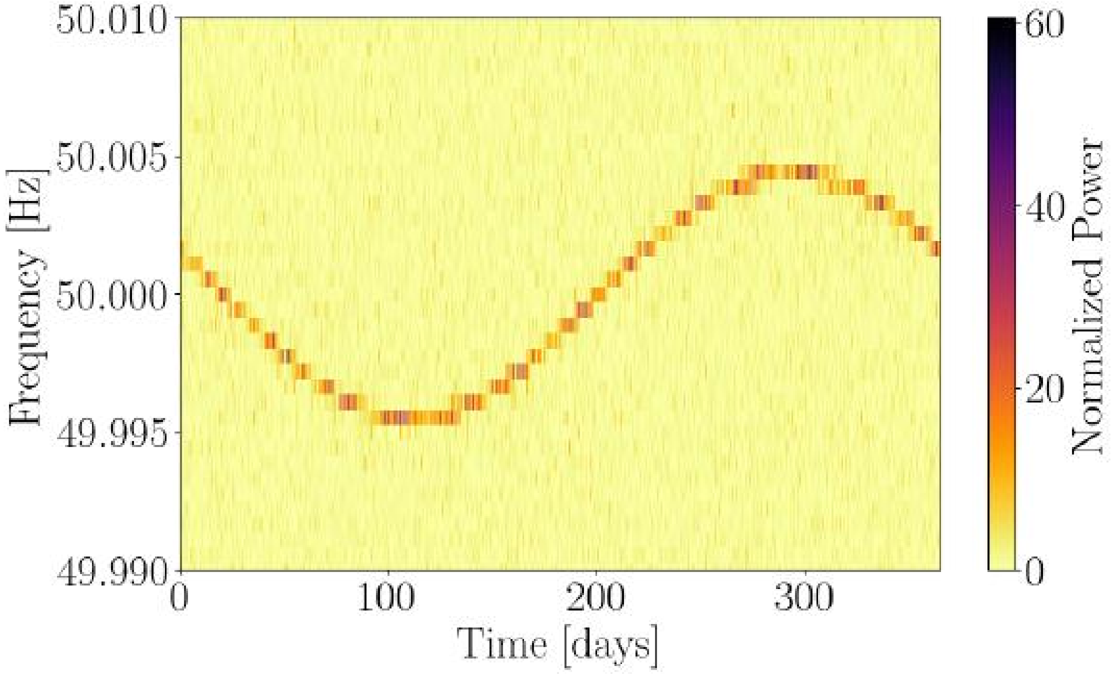

candidate post-processing

- Wide-parameter space searches (directed, all-sky) typically produce lots of statistical outliers that need to be post-processed.
- Typical steps (in variable order):
  - vetos: use simple characteristics of noise or expected signals to “kill” candidates en masse
  - clustering: reduce number of candidates by identifying small volumes in parameter space with multiple outliers that could come from the same physical source (instrumental disturbance or real CW signal)
  - follow-up: run a new search around interesting candidates, using different methods or settings: switching to matched filter if not used in first stage, increasing the coherence times, etc
  - upper limits: if no detection ☹ : software injections of simulated signals to estimate the h_0 at which we’d detect 95% of signals (averaged over other parameters)
    - can then be astrophysically interpreted as max allowed ellipticity for a NS at a certain distance
    - or equivalently exclusion distance for NSs at given max ellipticity
    - or e.g. under r-mode model

hierarchical MCMC follow-up

- Just one example of post-processing steps → tutorial Thursday
- Basic idea: multiple stages of follow-up chained in a *hierarchical search* [Brady&Creighton1998, Papa+2016, Ashton&Prix2018, Tenorio+2021]
- increasing T_{coh} in each stage:
  1) parameter space resolution becomes finer
  2) better background-signal separation
  3) suppress spurious non-CW artifacts (e.g. phase jumps at segment boundaries)

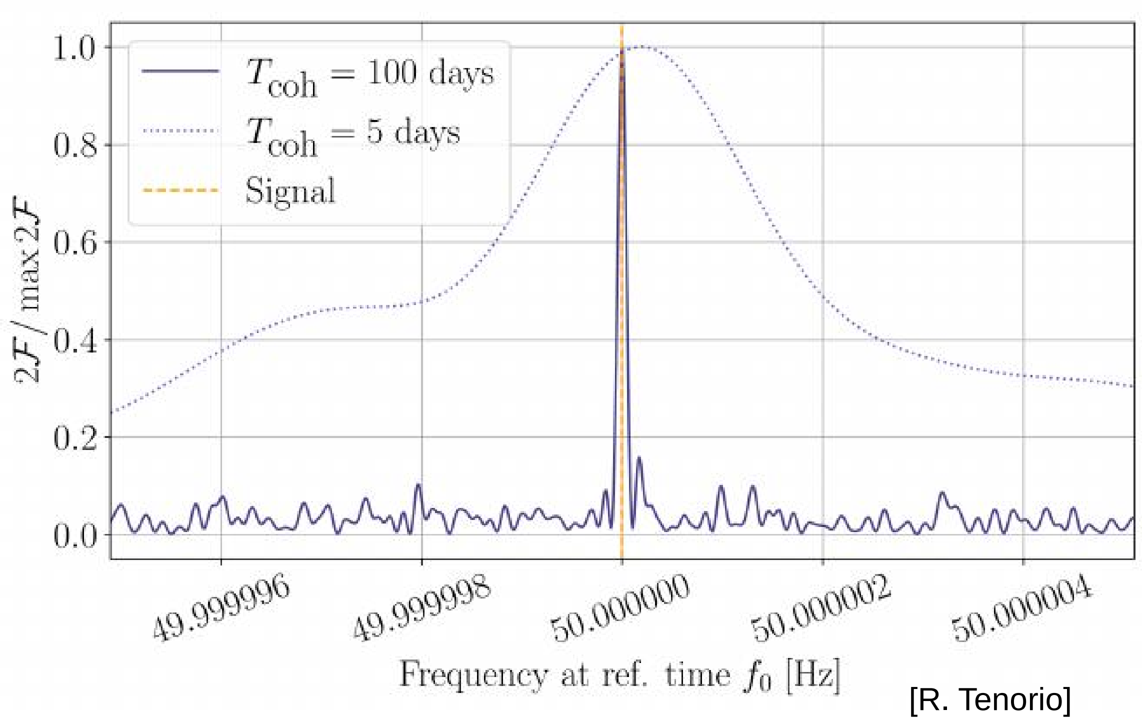

[R. Tenorio]

- PyFstat multi-stage MCMC approach (Ashton&Prix2018, Keitel+2021, Tenorio+2021): "naturally" zoom in from stage to stage, without having to fine-tune grids, following a "ladder" of increasing T_{coh}.

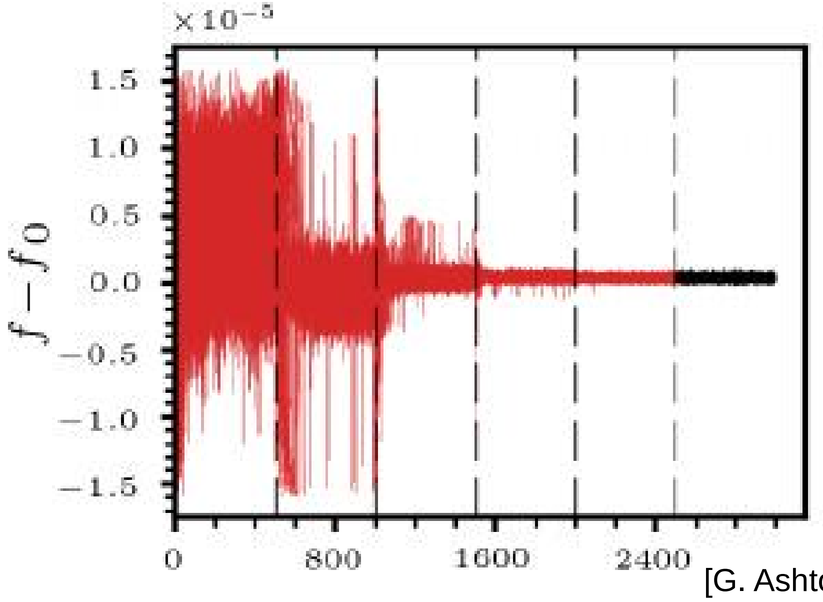

[G. Ashton]

- Used e.g. in Tenorio+2021 to rule out various O2 candidates, and in six O3 papers by the LVK.

Viterbi methods

- Another relatively recent innovation in CW searches, based on a well established signal processing method:
- Consider CW signal as a "hidden Markov model" and the GW model as the observable derived from it.
- points along a time-frequency track are the "states" of that model
- "Viterbi algorithm" is an efficient way to find the best track across [t,f] data range.
- → extremely cheap CW search

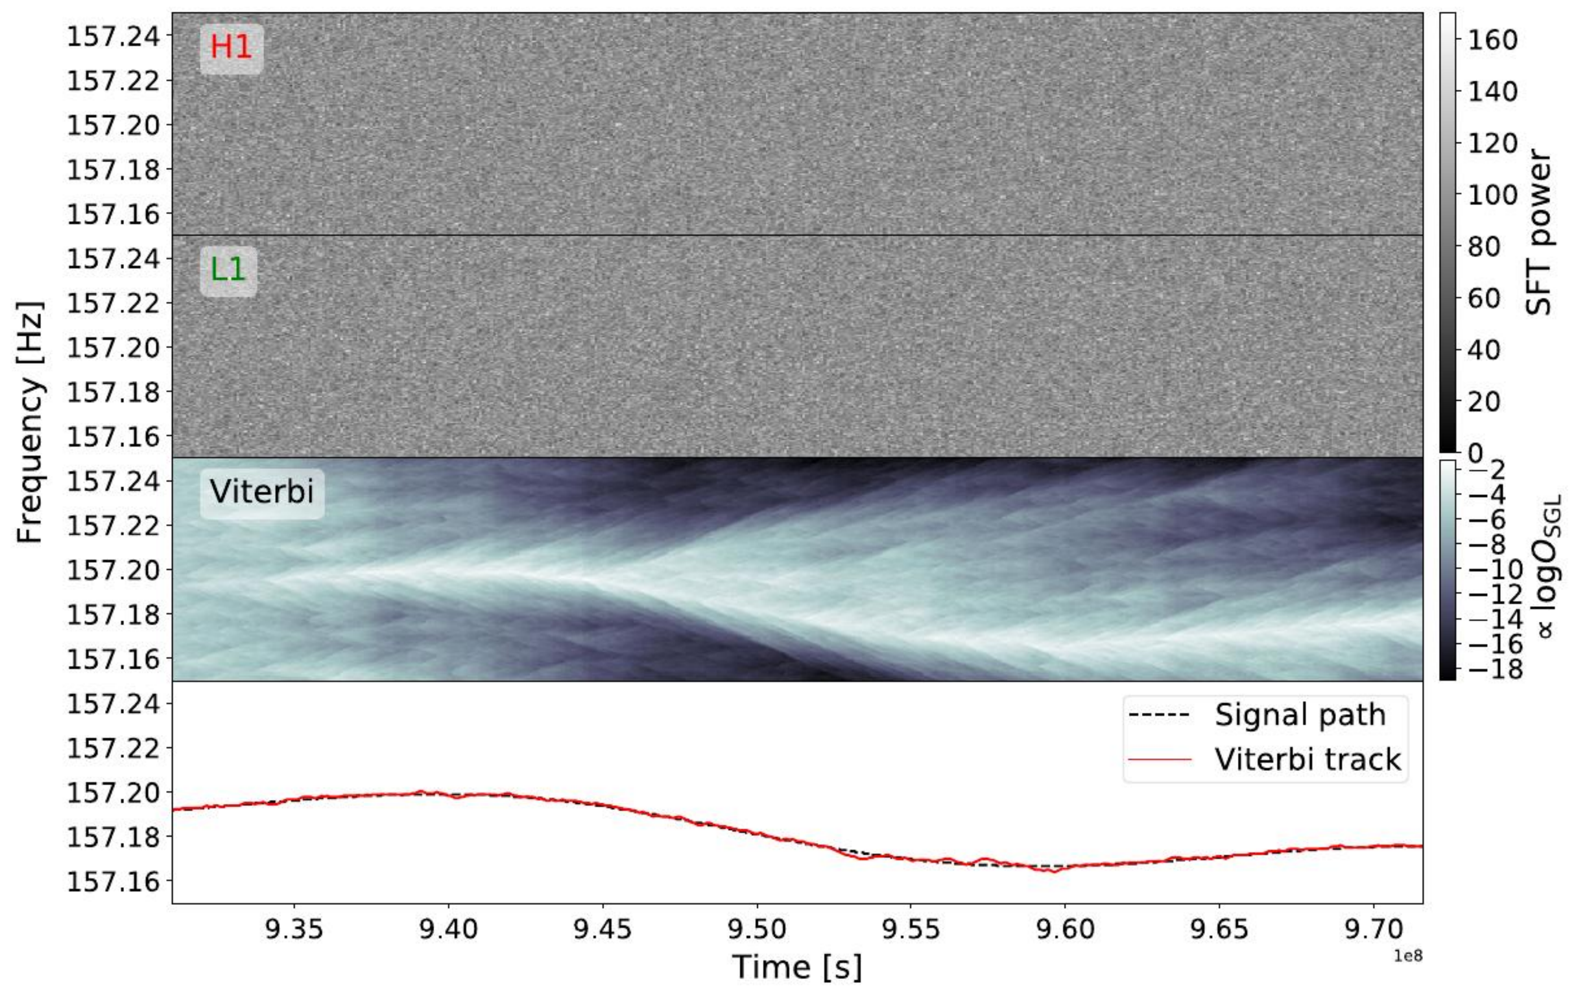

[Bayley+2019]

- robust against non-ideal signal evolution, e.g. NS glitchs, timing noise, spin wandering due to choppy accretion, …
- Suvorova+2016, Suvorova+2017, Sun+2017, Sun+2019, Bayley+2019
- Used in various directed and all-sky searches (see following slides).
- Also see → Joe Bayley's and Hannah Middleton's talks today, Andrés Vargas' tomorrow

directed search results (SNRs)

- We know where to look on the sky, but nothing about frequency and spindown.
- E.g. searches for supernova remnants with central compact object not yet found, or only through thermal emission, without pulsar signal telling us its rotation frequency.

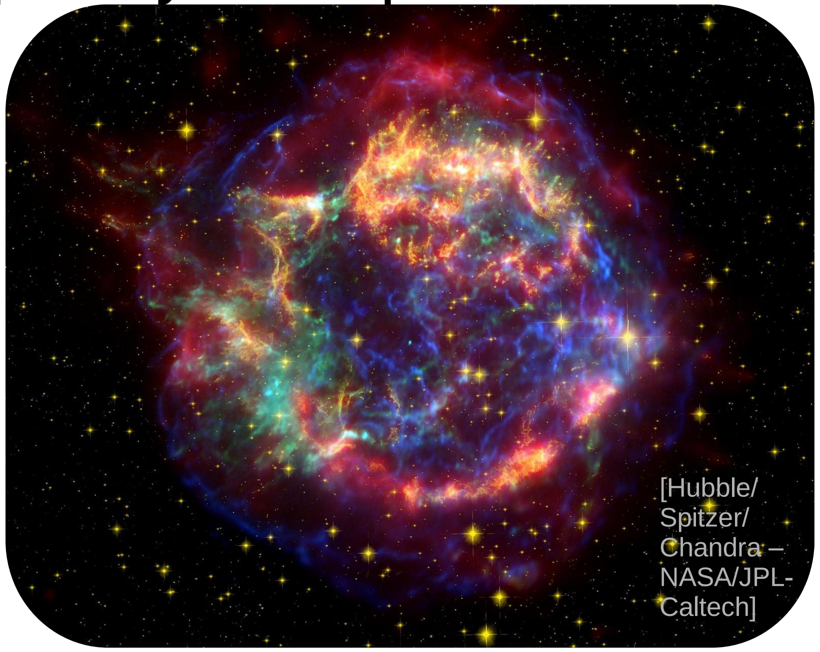

- Latest LVK supernova remnant results from O3:
  Abbott+2021 for 15 young targets,
  Abbott+2022 for Cassiopeia A and Vela Jr.
- Viterbi pipelines (standard and 1*f*+2*f*) and BSD ("Band-Sampled Data", version of FrequencyHough semi-coherent search with an efficient data format)
- As wide-freq searches, results can be inter-preted for both mass quadrupoles (ellipticity) and r-modes.

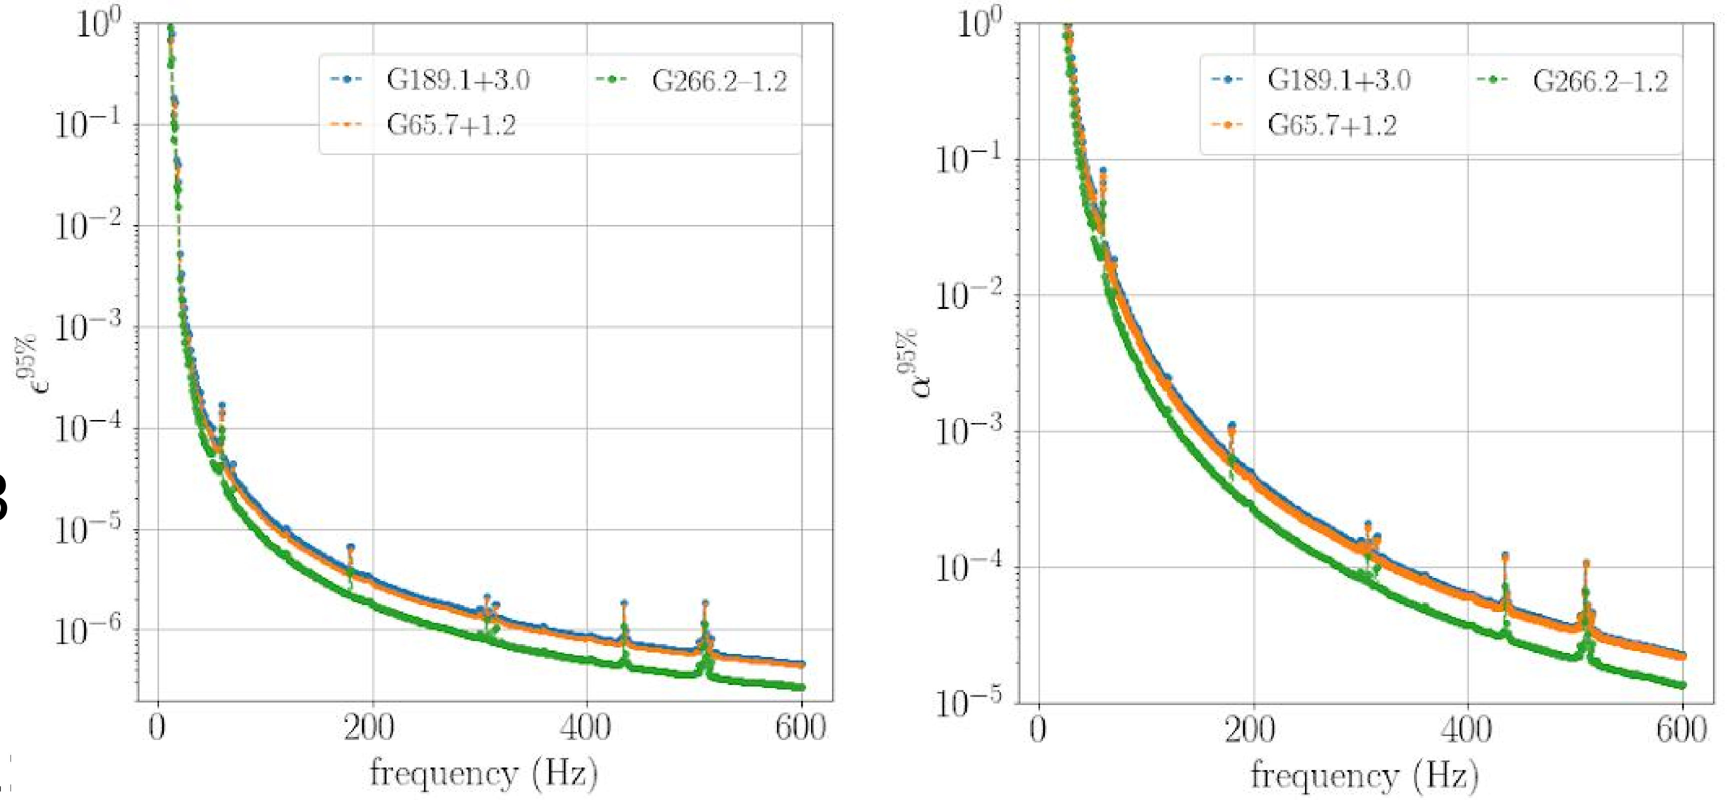

- Deep O2 search for G347.3 with Einstein@Home: Ming+2022
- Kilonova remnant G4.8+6.2 Liu&Zou2022 (O3 data)

directed search results (GC)

- galactic centre: 8kpc away from Earth, effectively a single sky pixel at our resolution
- Abbott+2022 O3 search with BSD pipeline
- Again, UL results can be interpreted in different ways:

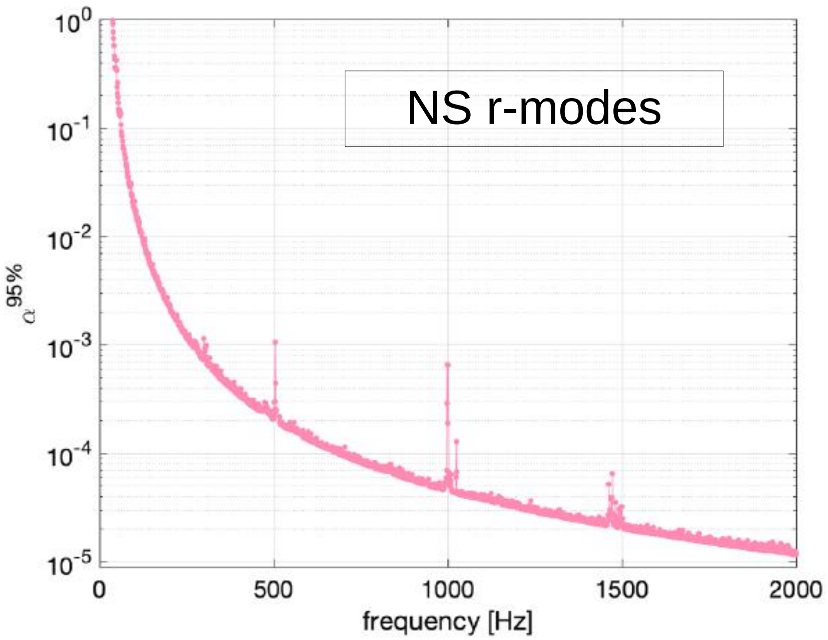

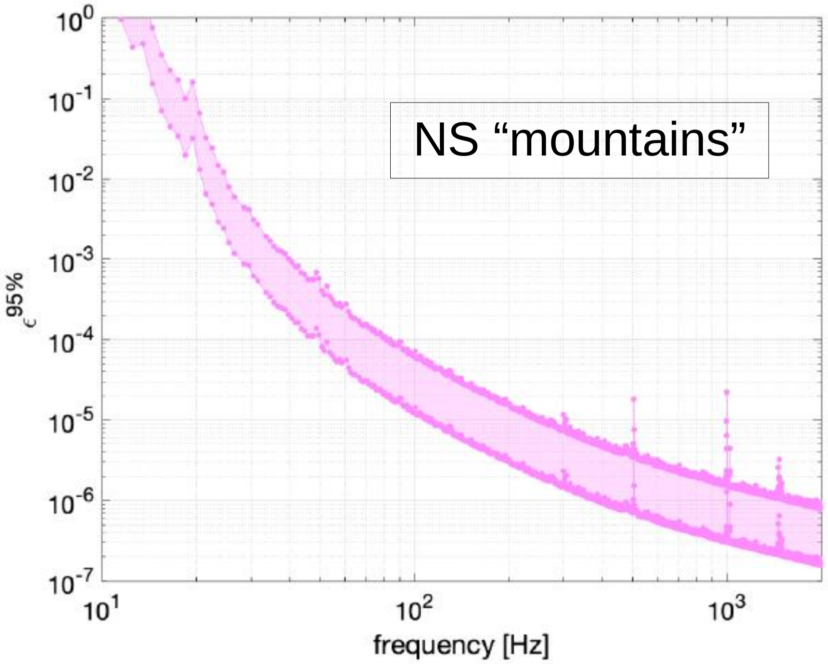

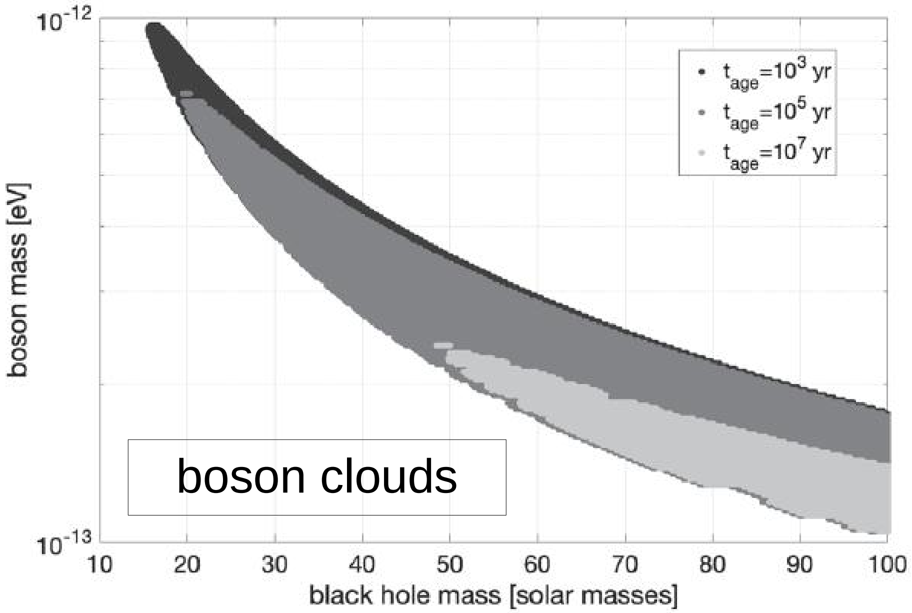

- Other quasi-point-like targets in the past: globular clusters, e.g. Abbott+2017 for NGC6544 (iLIGO S6 data)

directed search results (binaries)

- Most promising target for *directed* binary searches: Scorpius X-1
- Accreting binary, rotation frequency of NS unknown, possible that it is *torque-balanced*: GW emission exactly cancels out accretion spinup

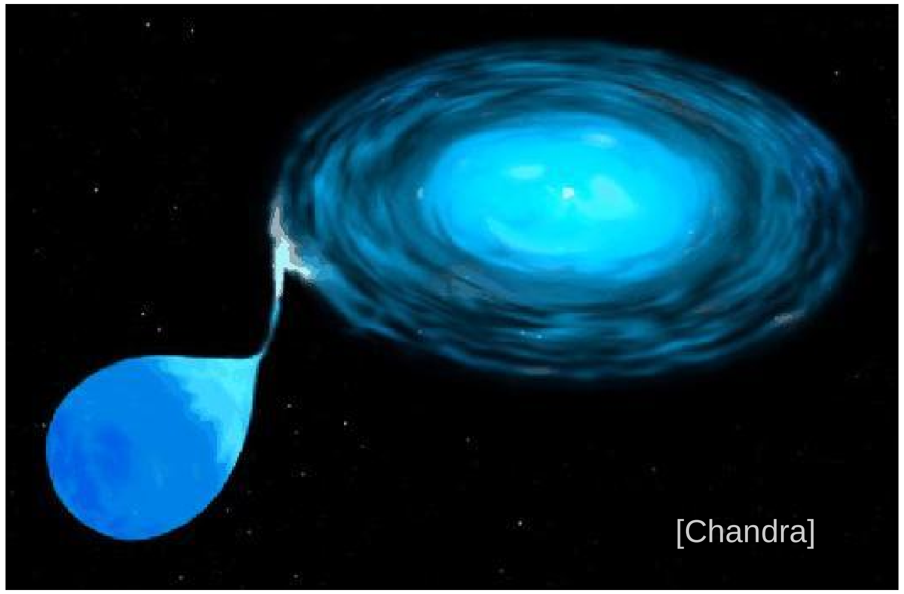

- Challenging due to large parameter space (frequency + orbital uncertainties)
- repeatedly analysed with different methods, in O3:
  Abbott+2022a fast&robust Viterbi search
  (→ Andrés Vargas' talk tomorrow)
  Abbott+2022b deeper cross-correlation search
  Whelan+2023 cross-corr rerun with corrected ephemeris
  (→ Thomas Killestein' talk Thursday)

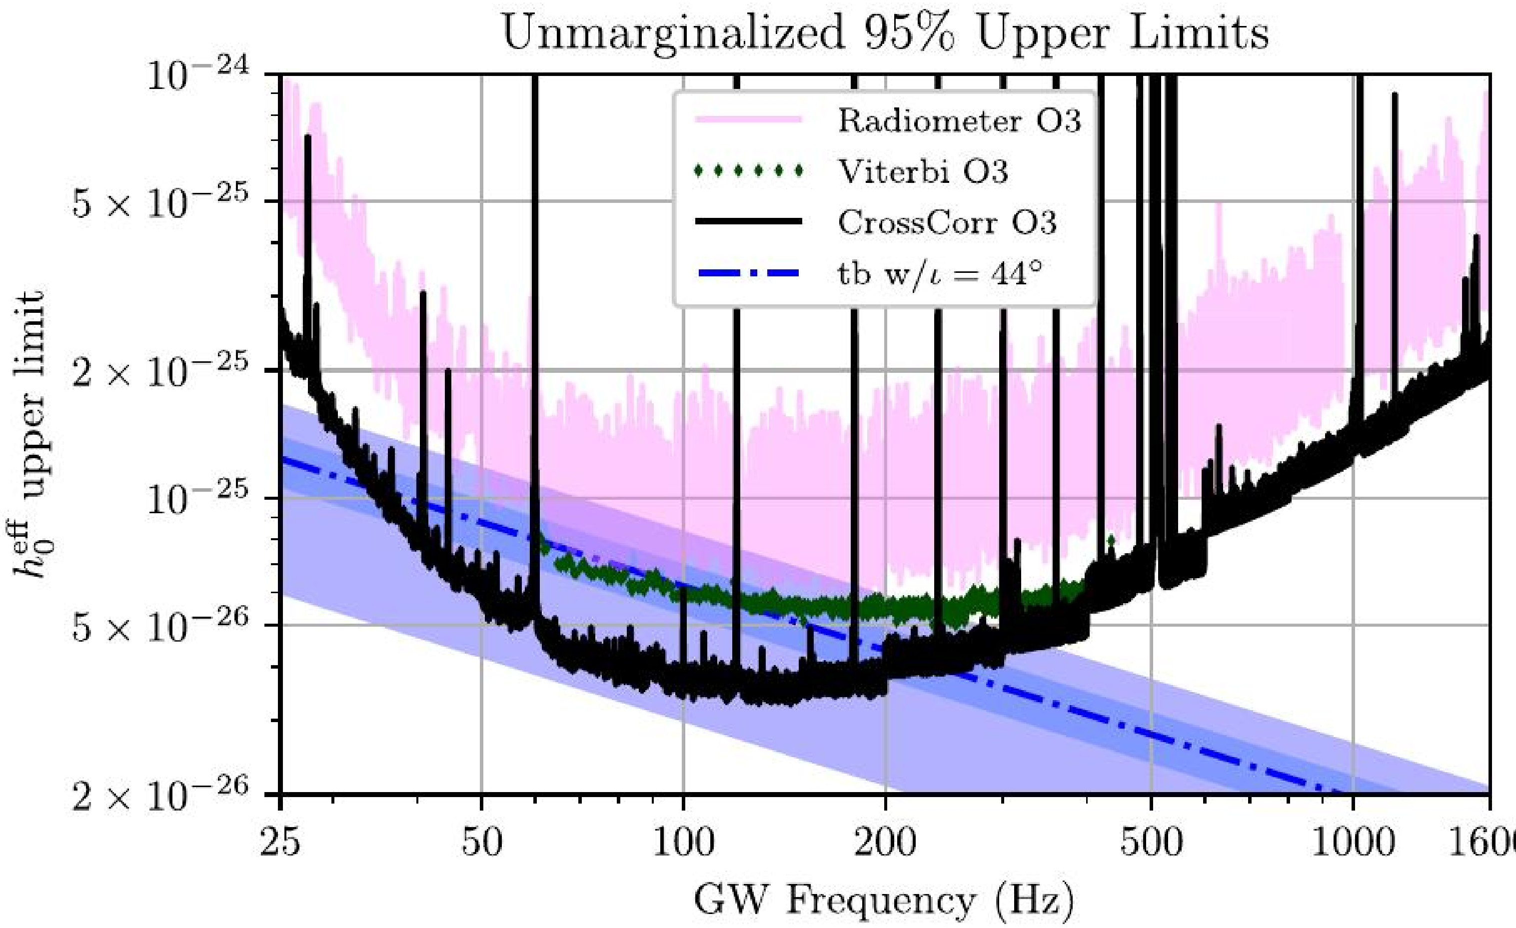

- Abbott+2021: searches for 20 AMXPs with Viterbi
  (→ Andrés Vargas' talk tomorrow)

“spotlight” or guided searches

- Some extended regions on the sky (e.g. galactic spiral arms) are expected to be overdense in NSs. Or if we focus on the most nearby NSs, structure of solar neighbourhood should be taken into account.
- Can also choose to guide frequency/spindown ranges by EM-observed pulsar population, or by population modelling.
- Historic example: Aasi+2016 search of the Orion spur using iLIGO S6 data
- See talks by → Rodrigo Tenorio later today

  and → Gianluca Pagliaro tomorrow

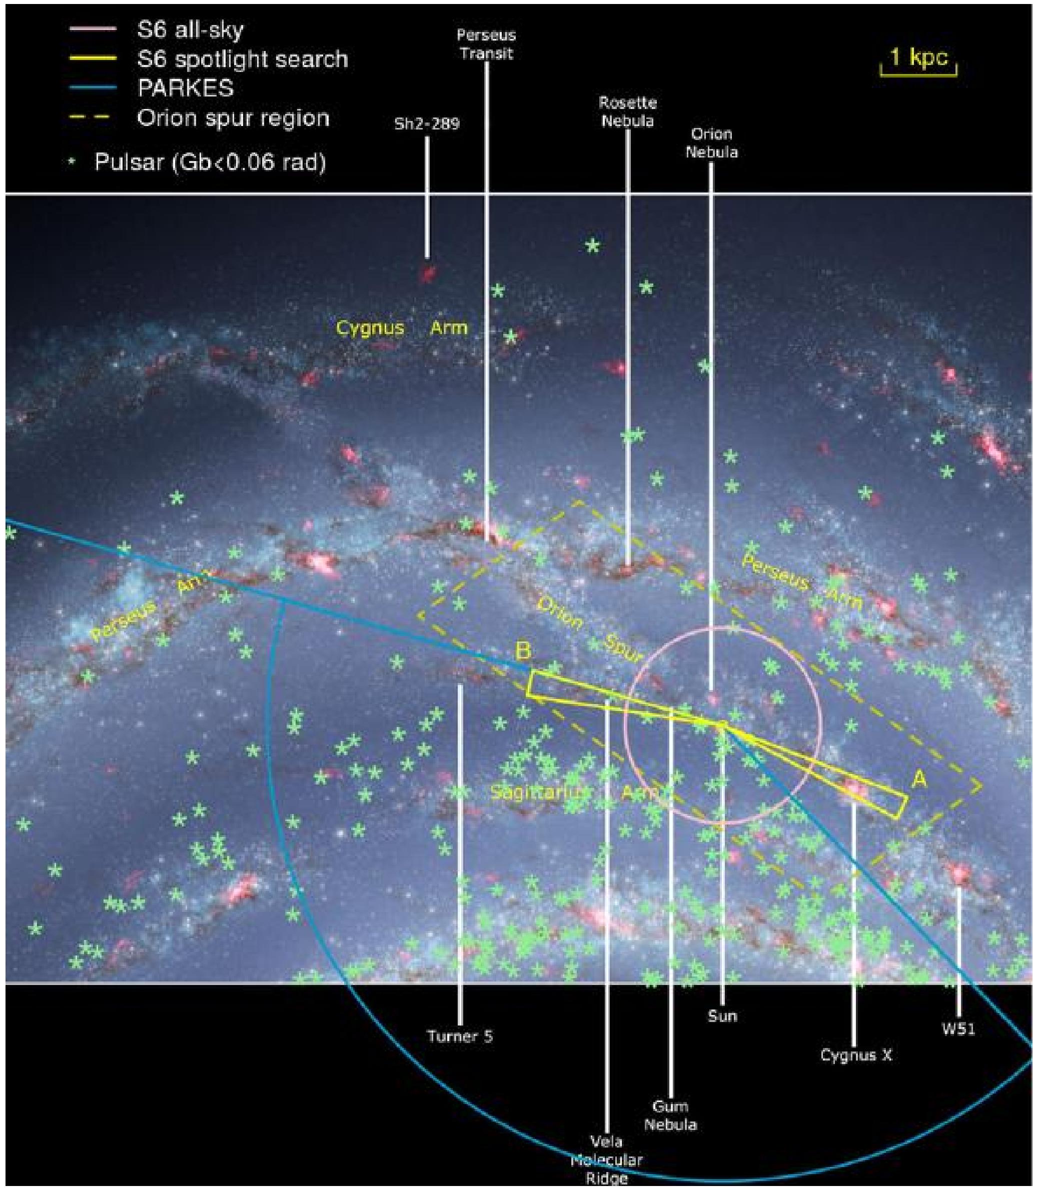

all-sky search results

- LVC results [Abbott+2021] from O3a data on unknown isolated NSs, anywhere in the sky, at any frequency and over broad spindown range
- PowerFlux pipeline (semi-coherent weighted sums of SFT power, going back to Abbott+2007) with “loosely coherent” follow-up stages [Dergachev2010, Dergachev2011]

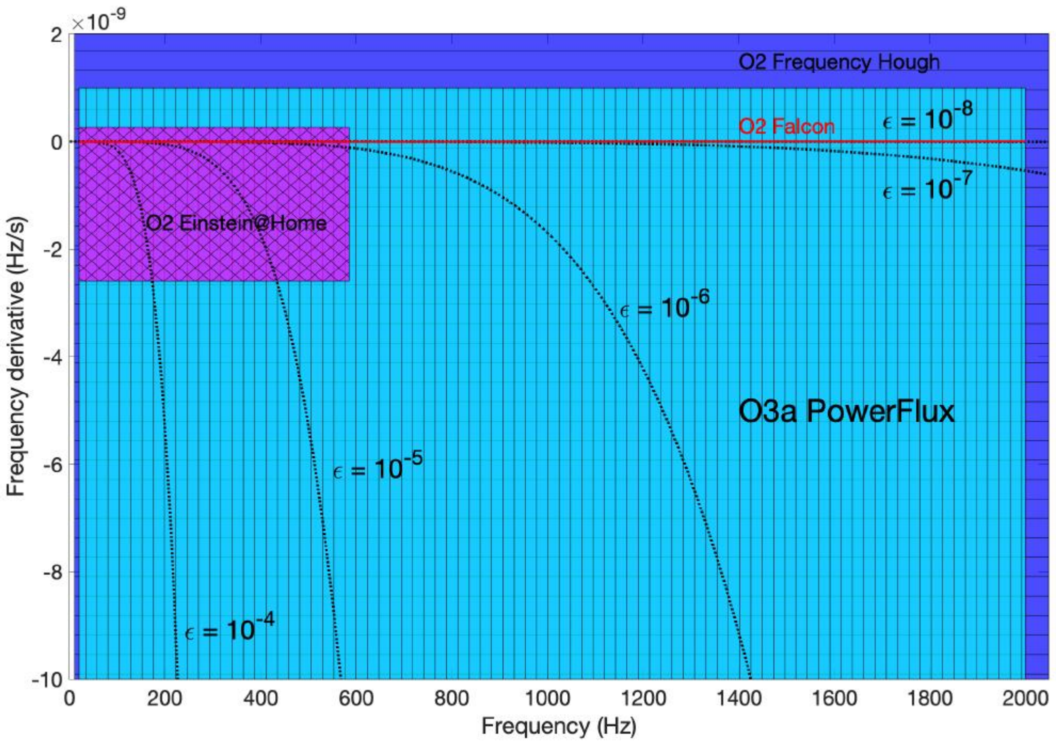

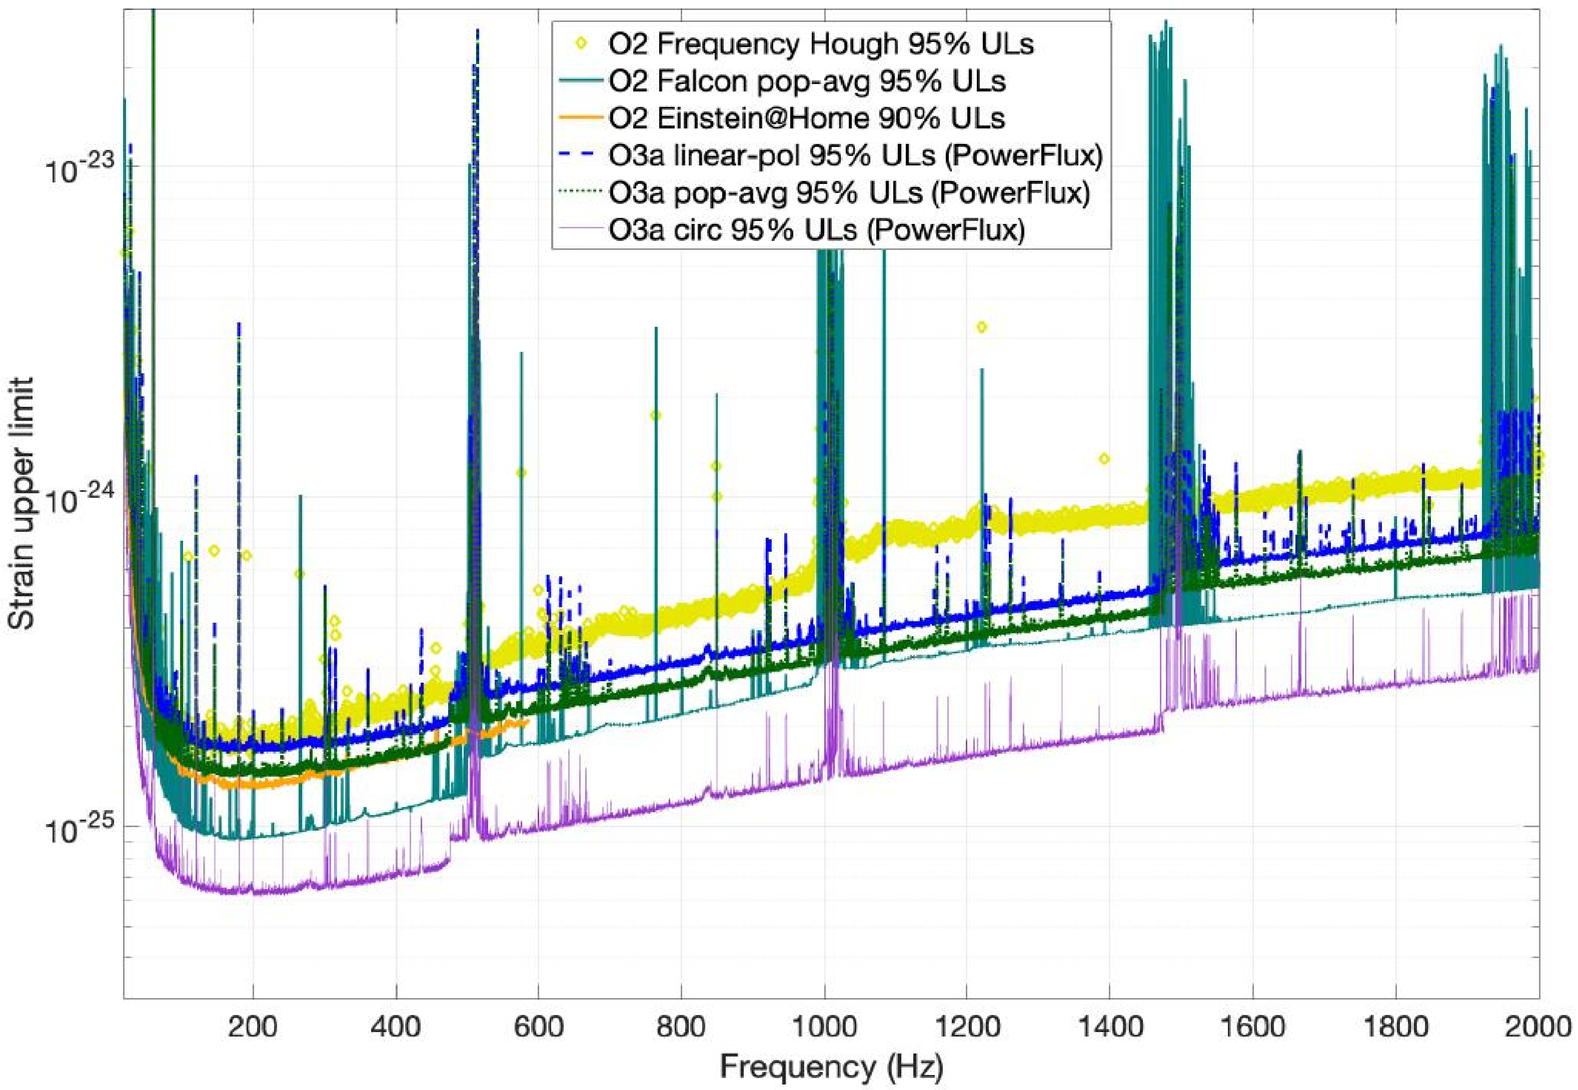

all-sky search results

- latest LVK results [Abbott+2022] from full-O3 data on CWs from unknown isolated sources:

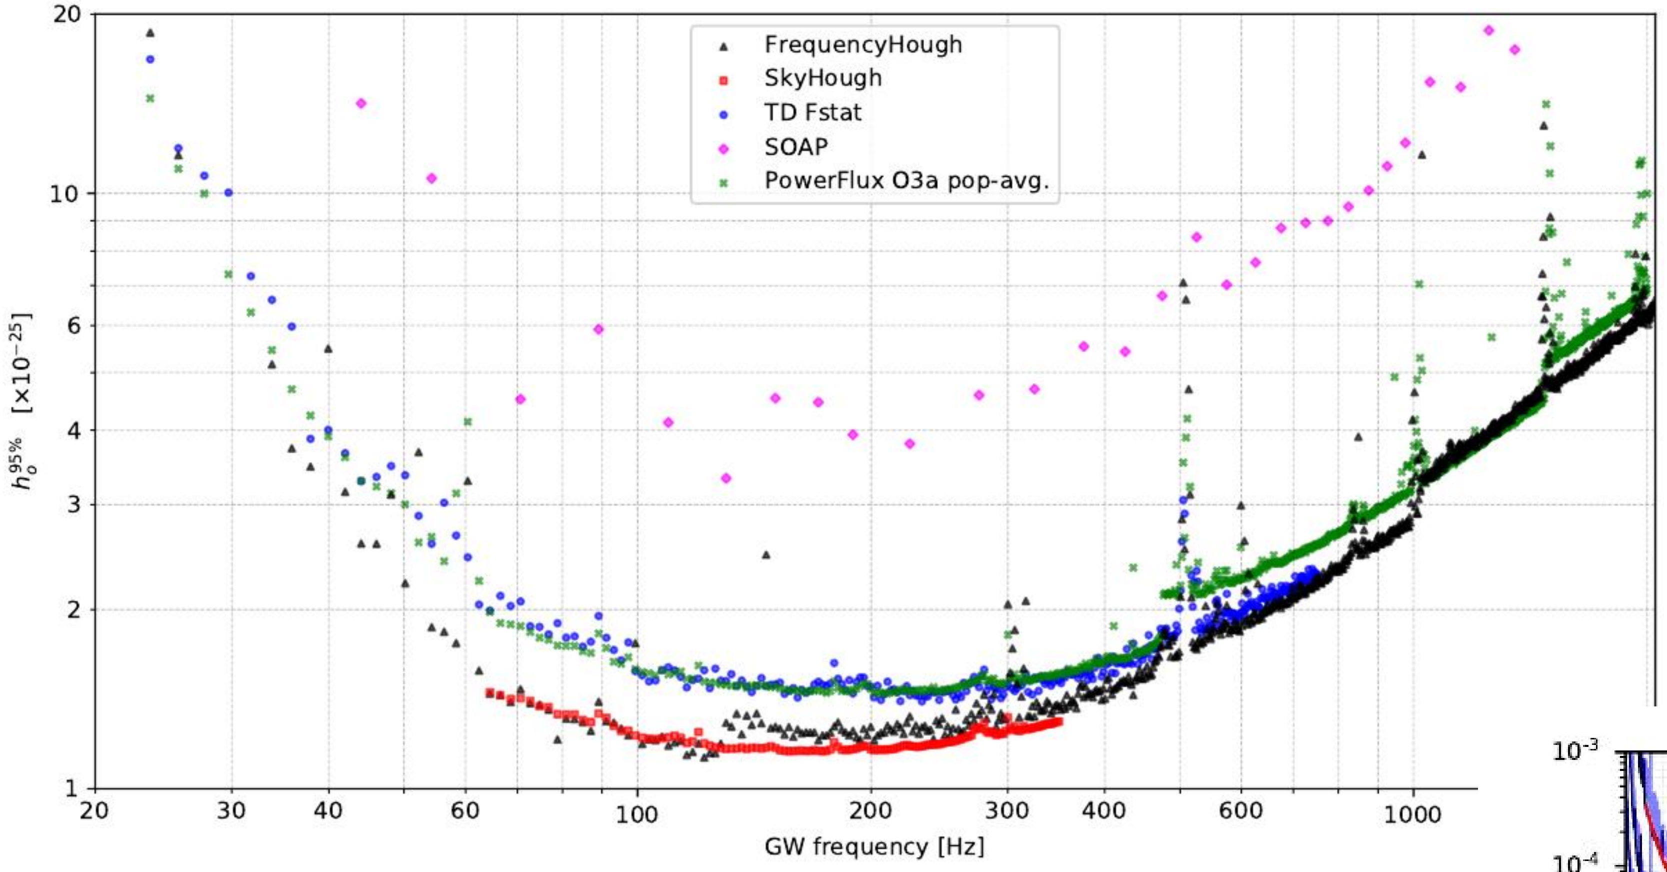

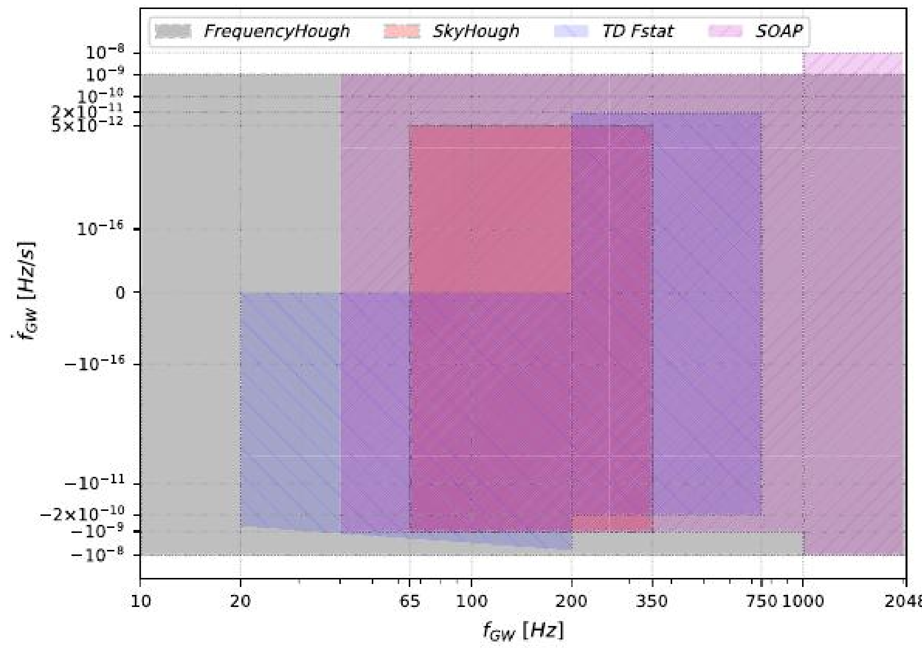

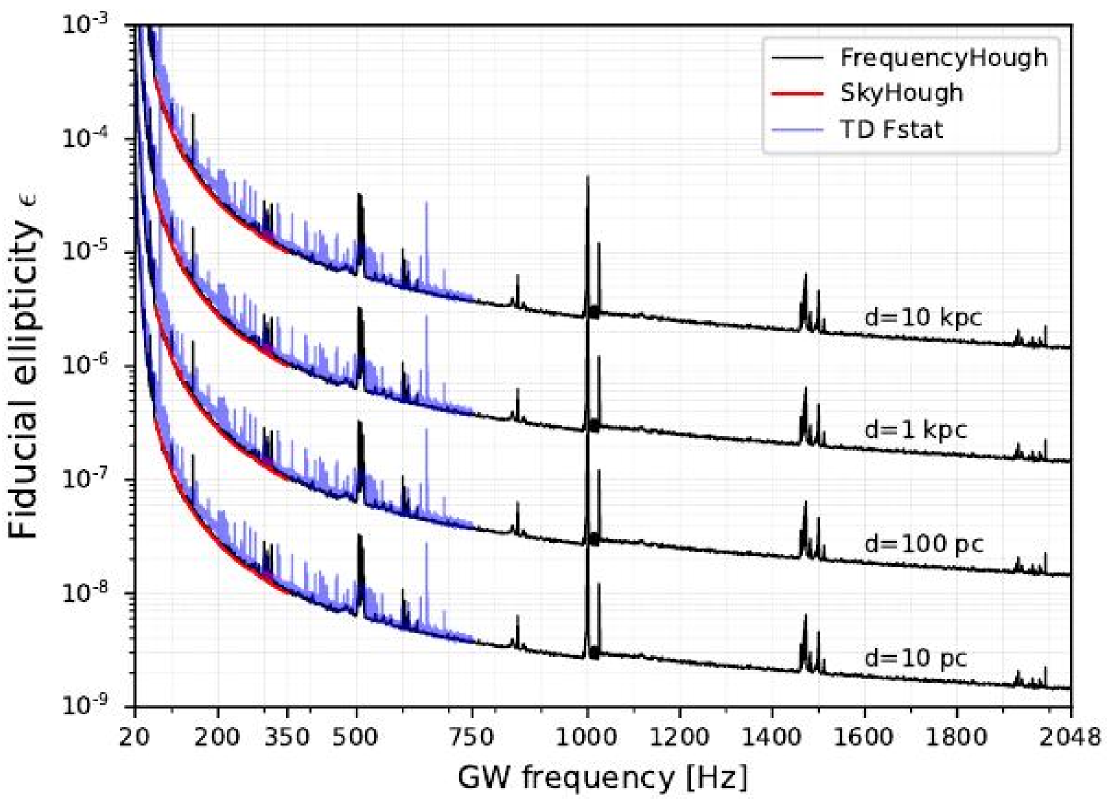

- Four pipelines:
  - SOAP (Viterbi, → Joe Bayley's talk today)
  - time-domain *F*-stat
  - SkyHough [Krishnan+2004, Sintes+2006]
  - FrequencyHough [Astone+2014]

all-sky search results

- Recent non-LVK results on O3 public data: Dergachev&Papa2023
- deep ULs for narrow parameter space inspired by pulsar population
- “Frequency-Resolved Atlas”:
  freq-and-sky-resolved data release,
  including ULs
  and deep candidate lists

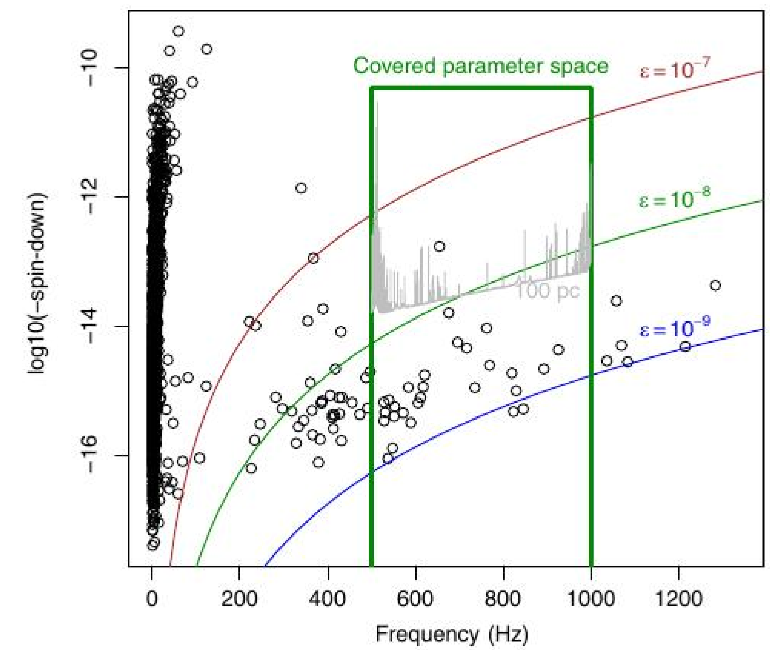

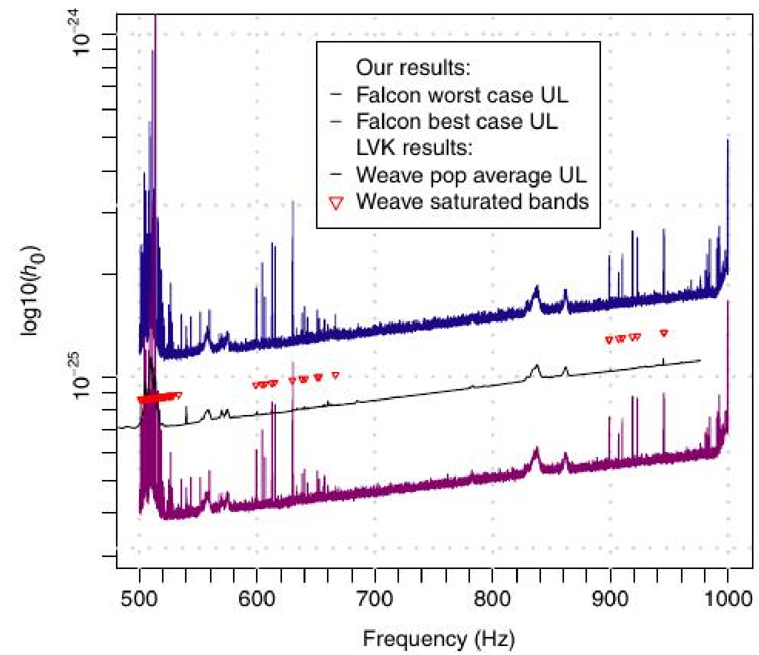

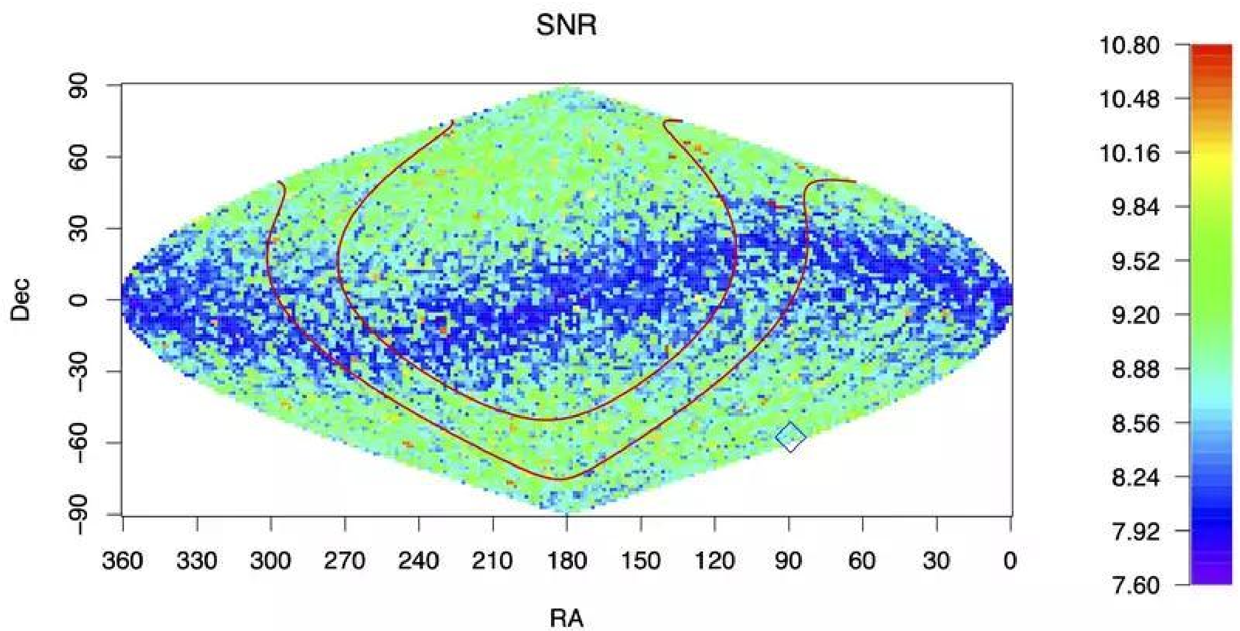

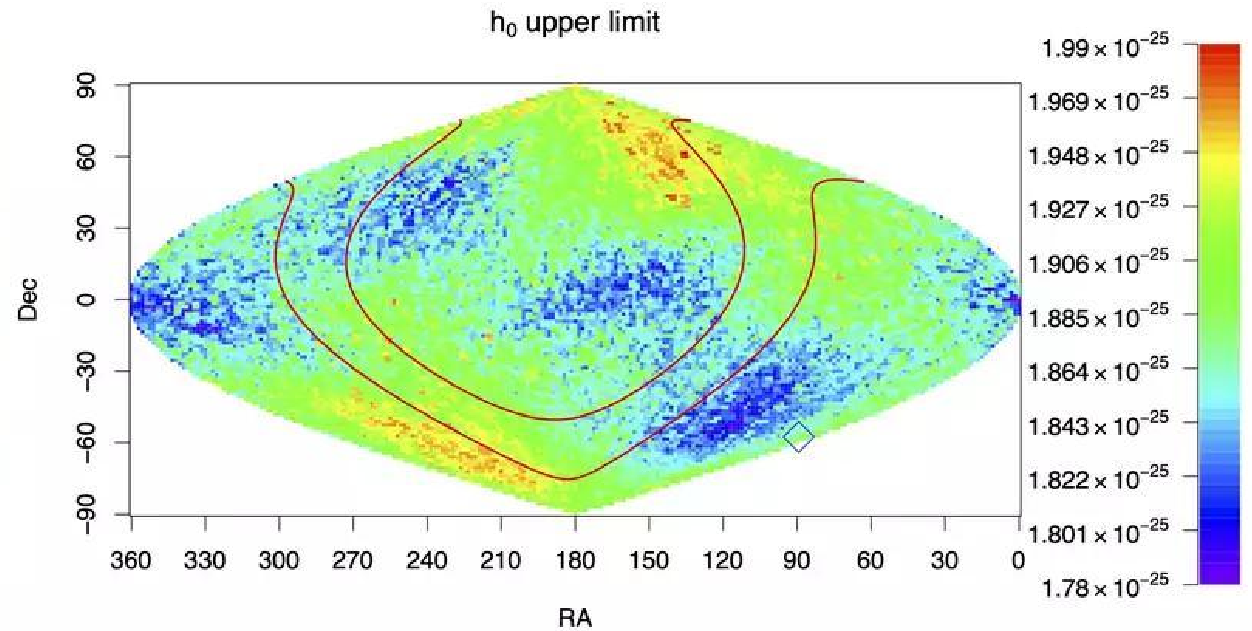

all-sky search results

- Recent non-LVK results on O3 public data: Steltner+2023 Einstein@Home results

→
see talk by
Benjamin Steltner
later today

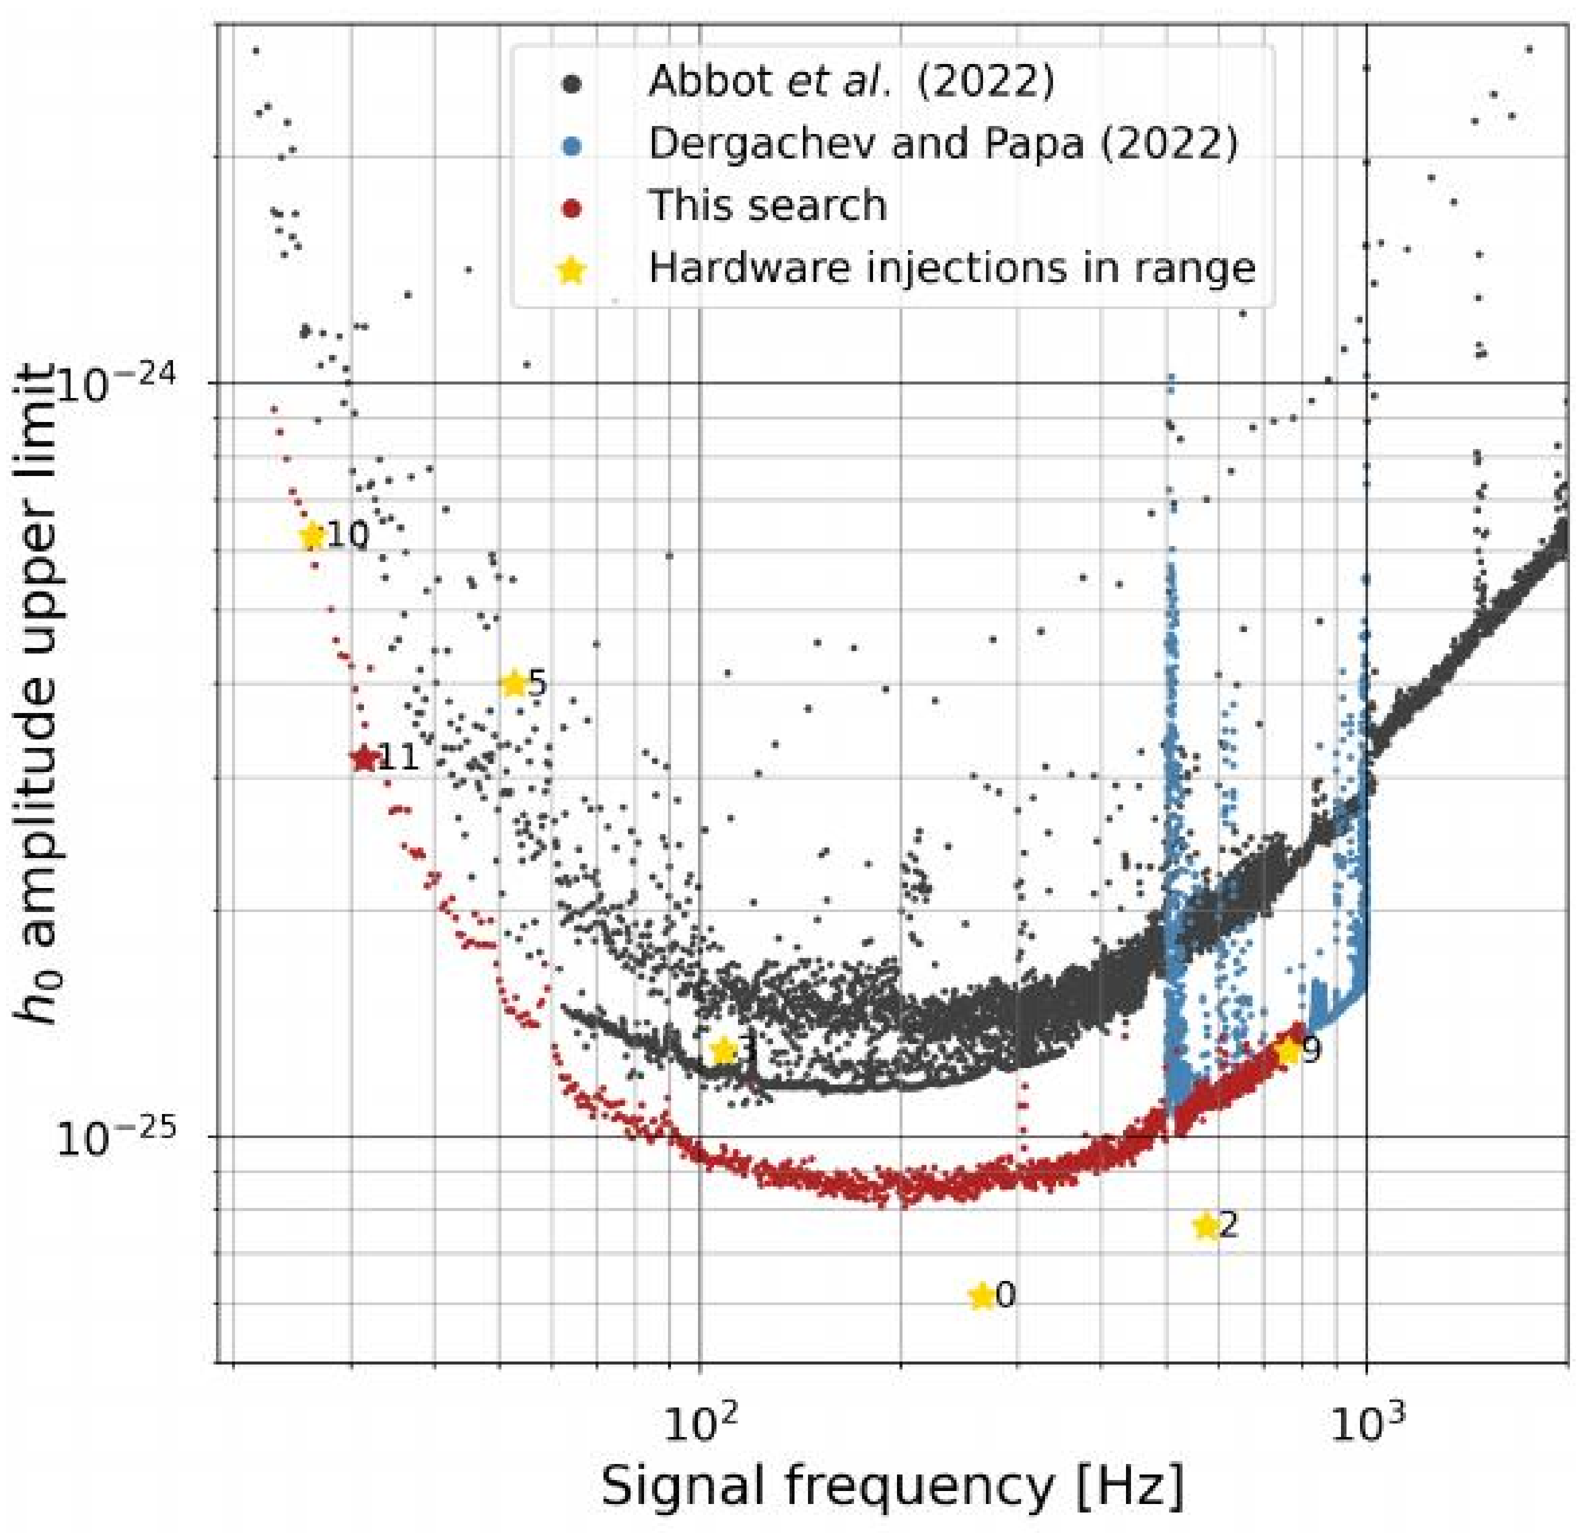

all-sky search results (binaries)

all-sky binary CW search on O3a data [Abbott+2020, Tenorio+2021] with BinarySkyHough pipeline

[Covas&Sintes2019]

~10^{14} templates

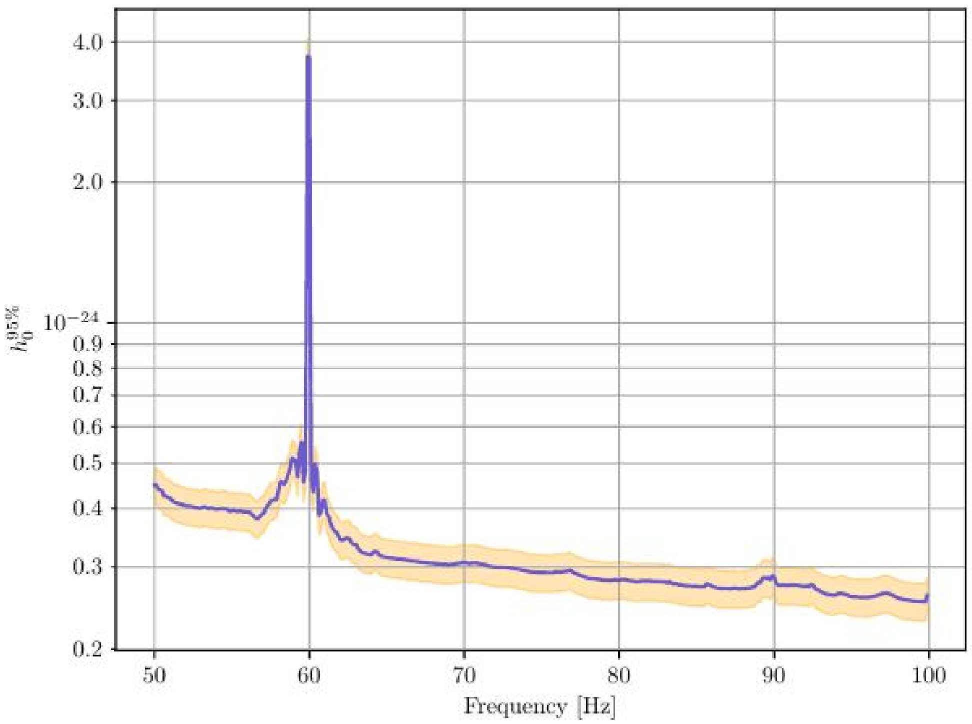

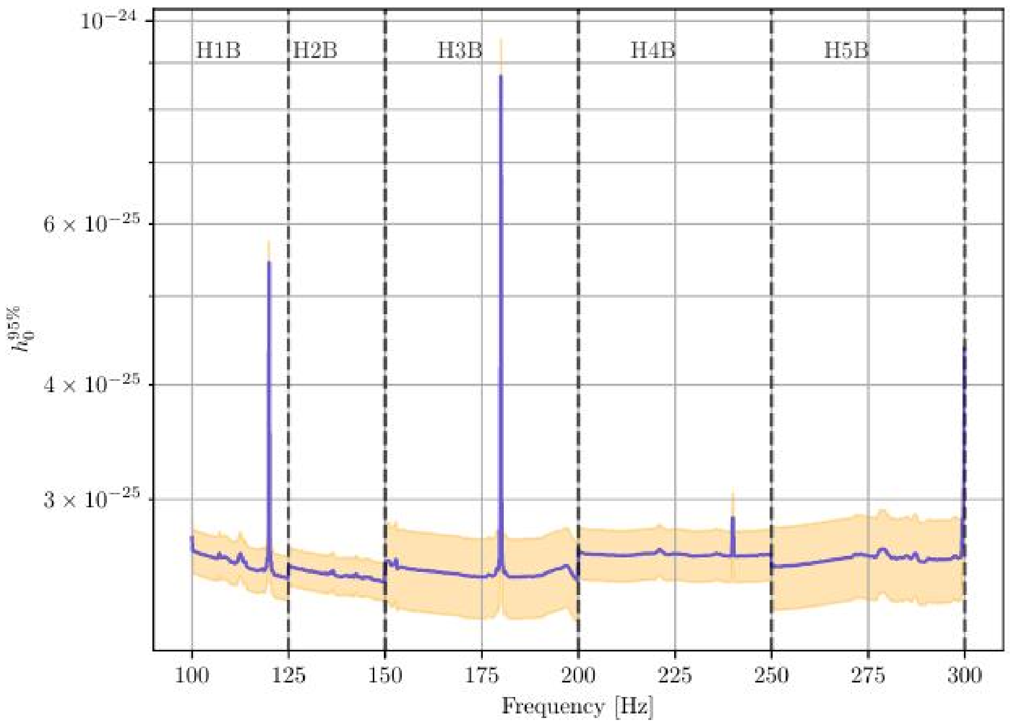

h_0 upper limits translated into *astrophysical reach* [kpc] at fixed ellipticity or into ellipticity constraints at fixed distance:

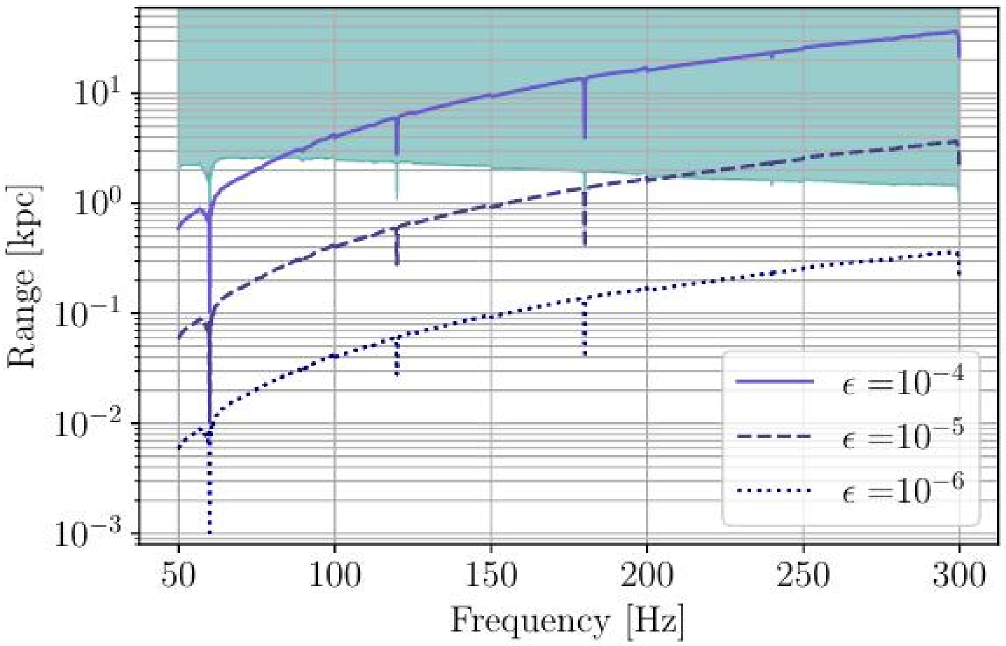

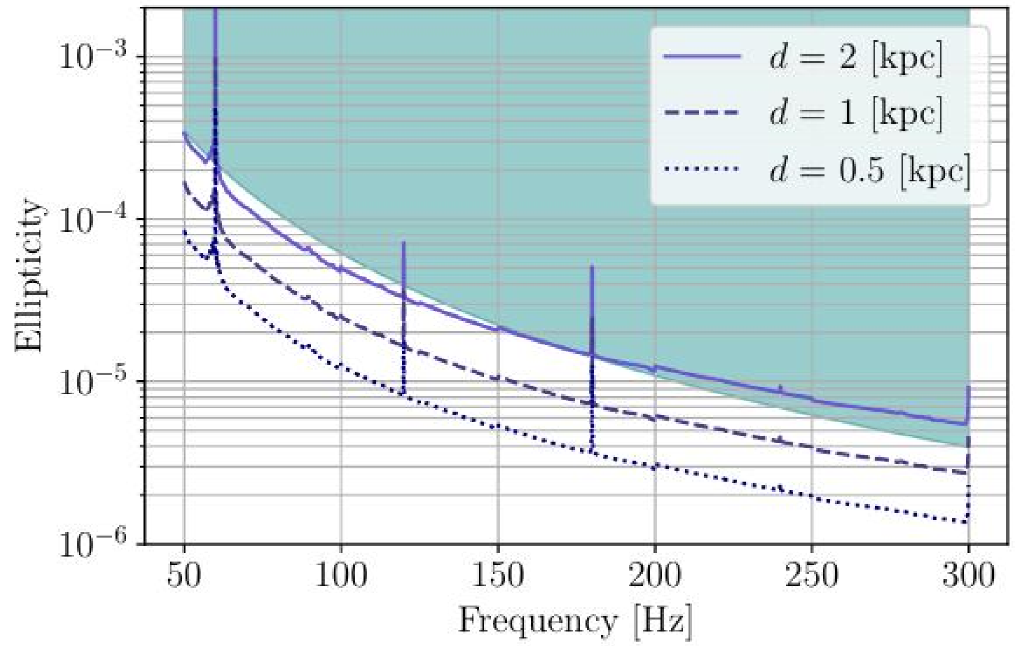

[https://arxiv.org/abs/2012.12128]

all-sky search results (binaries)

- Recent non-LVK results on O3 public data: Covas+2022

→ see talk by Pep Covas later today

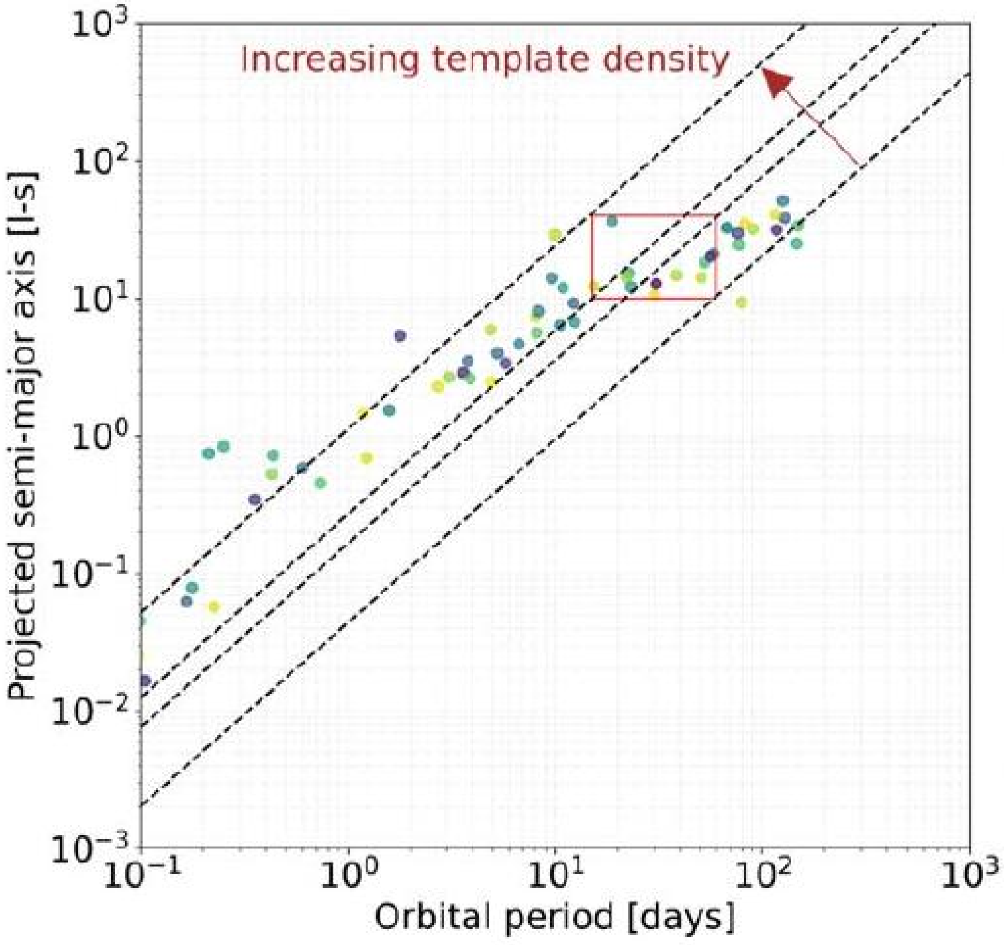

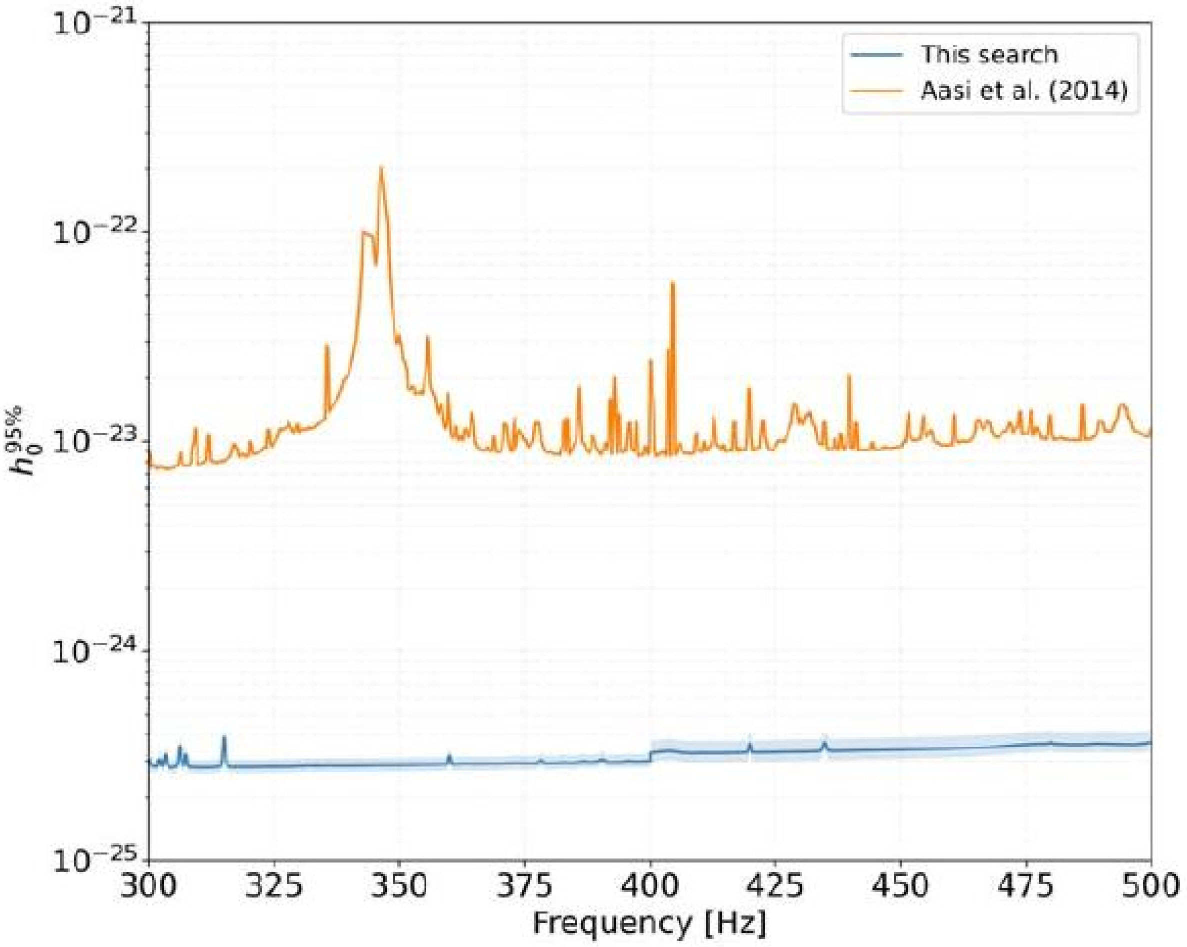

- ”Abbott et al. (2021d) also searched O3a data (but only up to 300 Hz) and attained a sensitivity comparable to ours”

outlook: CWs beyond LVK and beyond NSs

- Future detectors (Einstein Telescope, Cosmic Explorer): → see Ben Owen’s talk tomorrow
- CW-like signals possible from other more exotic sources. Many search methods can be transferred.

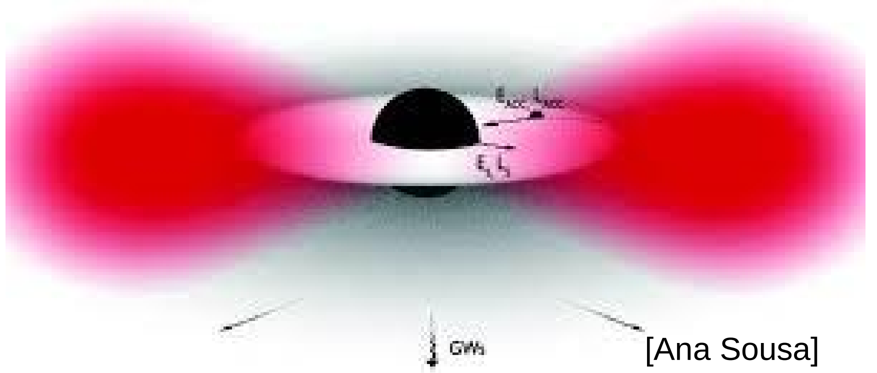

[Ana Sousa]

- New physics: light bosons (e.g. “axions”) could form “clouds” around spinning black holes and extract energy.
  → CW-like emission with frequency related to particle mass
  - LVK search on O3 data: Abbott+2022

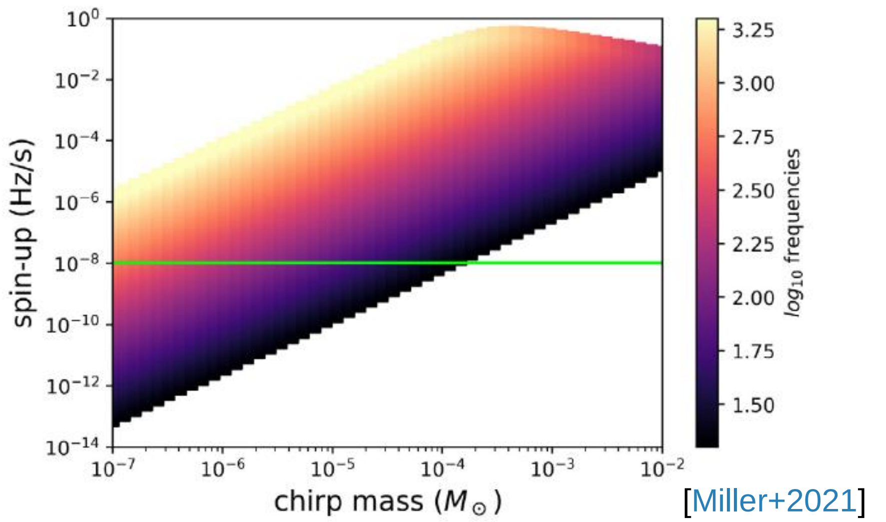

[Miller+2021]

- Low-mass compact binaries are CW-like in early inspiral, e.g. primordial black holes (Miller+2021, Miller+2022).
- Direct dark matter interaction with GW detectors, e.g. “dark photon” search in O3 (Abbott+2022).

$$\mathcal{L} = -\frac{1}{4\mu_0} F^{\mu\nu} F_{\mu\nu} + \frac{1}{2\mu_0}\left(\frac{m_A c}{\hbar}\right)^2 A^\mu A_\mu - \epsilon e J^\mu A_\mu$$

- Less speculative, but further in the future (2034+): LISA space-based detector will see CW-like early inspiral of white dwarf – white dwarf binaries.
  - We already know such systems are out there in the Milky Way and ***will*** be detectable

→ “verification binaries”!

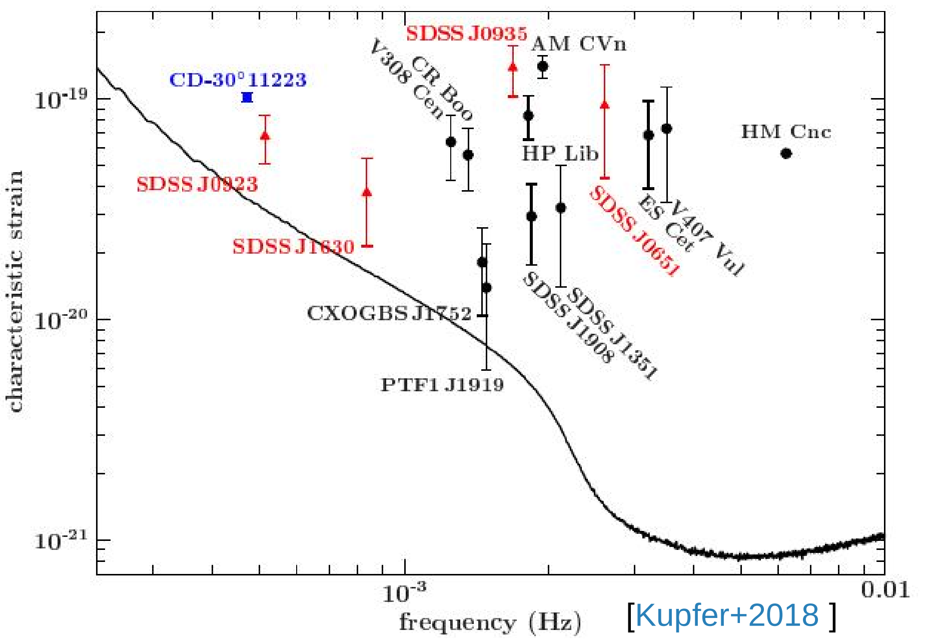

[Kupfer+2018]

tutorials on Thursday

- Enjoy the rest of the workshop!
- Scheduled for later: a practical data analysis tutorial
- If you don’t have a working environment with LALSuite and PyFstat yet, follow instructions at https://github.com/PyFstat/PyFstat/wiki/conda-environments or try Google Colab to run the tutorials online.
- Tutorials will be based on https://pyfstat.readthedocs.io/en/stable/examples.html
- PyFstat project home: https://github.com/PyFstat/PyFstat
- PyFstat reference: https://doi.org/10.21105/joss.03000
- Note: this is a relatively small project. It builds on top of LALSuite (thanks to Karl Wette’s SWIG bindings which allow python to call C libraries, Wette2020) but it is not as deeply tested as LALSuite itself, and only LVK-reviewed for a few specific applications (mainly MCMC candidate follow-up).

  → Only the second workshop where we run tutorials, and it’s quite possible that you run into some bugs in corner cases, or missing features.

  → Issue reports or pull requests via github are very welcome!

acknowledgments

David Keitel is supported by the Spanish Ministerio de Ciencia, Innovación y Universidades (ref. BEAGAL 18/00148) and cofinanced by the Universitat de les Illes Balears. Rodrigo Tenorio is supported by the Spanish Ministerio de Universidades (ref. FPU 18/00694). This work was supported by the Universitat de les Illes Balears (UIB); the Spanish Ministry of Science and Innovation (MCIN) and the Spanish Agencia Estatal de Investigación (AEI) grants PID2019-106416GB-I00/MCIN/AEI/10.13039/501100011033, RED2022-134204-E, RED2022-134411-T; the MCIN with funding from the European Union NextGenerationEU (PRTR-C17.I1); the FEDER Operational Program 2021-2027 of the Balearic Islands; the Comunitat Autònoma de les Illes Balears through the Direcció General de Política Universitaria i Recerca with funds from the Tourist Stay Tax Law ITS 2017-006 (PRD2018/23, PDR2020/11); the Conselleria de Fons Europeus, Universitat i Cultura del Govern de les Illes Balears; and EU COST Action CA18108. This material is based upon work supported by NSF's LIGO Laboratory which is a major facility fully funded by the National Science Foundation.