Title: Extreme-mass-ratio inspirals: Waveform models and signal space

Speakers: Alvin Chua

Series: Strong Gravity

Date: October 22, 2020 - 1:00 PM

URL: http://pirsa.org/20100007

Abstract: Extreme-mass-ratio inspirals (EMRIs) are the only gravitational-wave sources for the future LISA detector that combine the issue of strong-field complexity with that of long-lived signals. The result is a profoundly difficult inverse problem, with many theoretical and computational challenges presented both by the forward modeling of the predicted EMRI waveform, and by the recovery of an inverse solution for the presence and properties of actual EMRI signals in LISA data. I outline recent progress and ongoing work on both fronts. Specifically, I describe: i) the next generation of accurate, efficient and extensive models for (the response of LISA to) an EMRI signal; and ii) a heuristic study of degeneracy in the space of LISA-observable EMRIs, with a view to informing analysis strategies for EMRI search and characterization.

Extreme-mass-ratio inspirals:
Waveform models and signal space

Alvin Chua
JPL-Caltech

Perimeter Institute (virtual)
22 October 2020

Jet Propulsion Laboratory
California Institute of Technology

The inverse problem in GW astronomy

- Waveform modeling provides the basic forward models of GW science:

$$\text{Parameters} \mapsto \text{observables}$$

  - Other forward models include those for detector noise, or for source populations
- Data analysis uses forward models to find inverse solutions:

$$\text{Data} \mapsto \text{parameters}$$

  - “Parameters” here include presence of signal itself (traditional LVK name: detection)
- Both parts are solved in very different ways, often by separate communities
- But it is useful to keep in mind that they are flip sides of the same problem

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vl

LISA: An overview

- Laser Interferometer Space Antenna
  - 3 spacecraft, 2.5×10^9 m separation
  - Earth-trailing “cartwheel” orbit
  - Mission lifetime: 4-10 years

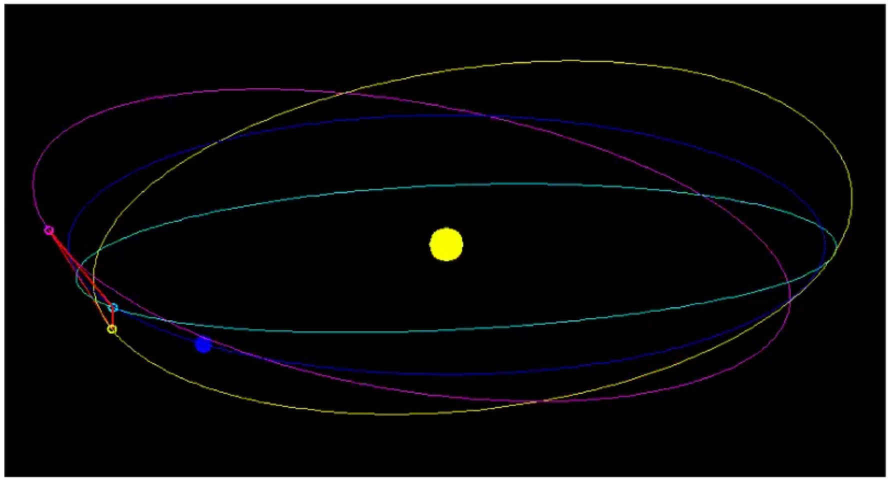

N. Douillet

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vl

LISA: An overview

- Laser Interferometer Space Antenna
  - 3 spacecraft, 2.5×10^9 m separation
  - Earth-trailing "cartwheel" orbit
  - Mission lifetime: 4-10 years
- GW sensitivity band: 10^{-4}-1 Hz
- Effectively 3 interferometers
  - 2 Michelson: Measure GW polarizations
  - 1 Sagnac: Measure non-GW noise

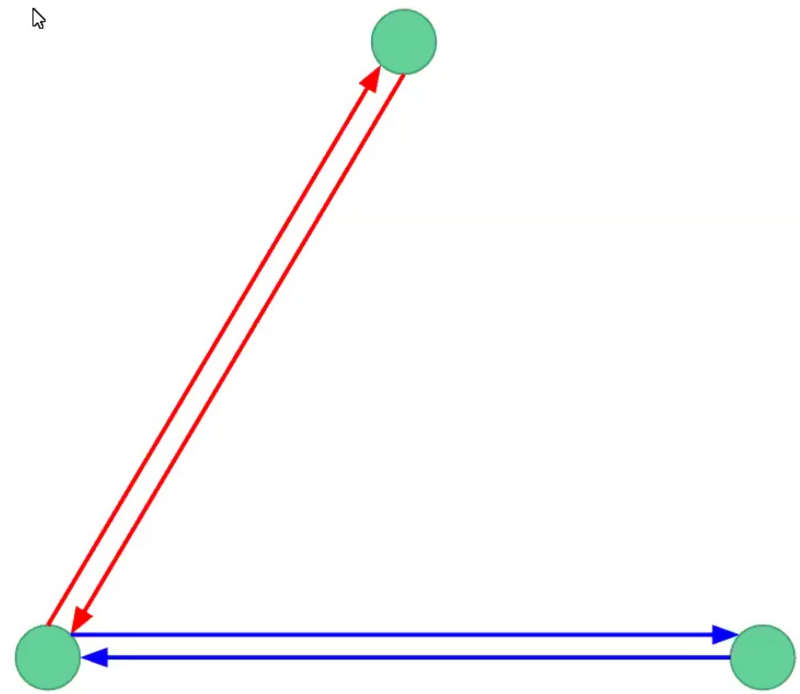

LISA: An overview

- Laser Interferometer Space Antenna
  - 3 spacecraft, 2.5×10^9 m separation
  - Earth-trailing "cartwheel" orbit
  - Mission lifetime: 4-10 years
- GW sensitivity band: 10^{-4}-1 Hz
- Effectively 3 interferometers
  - 2 Michelson: Measure GW polarizations
  - 1 Sagnac: Measure non-GW noise

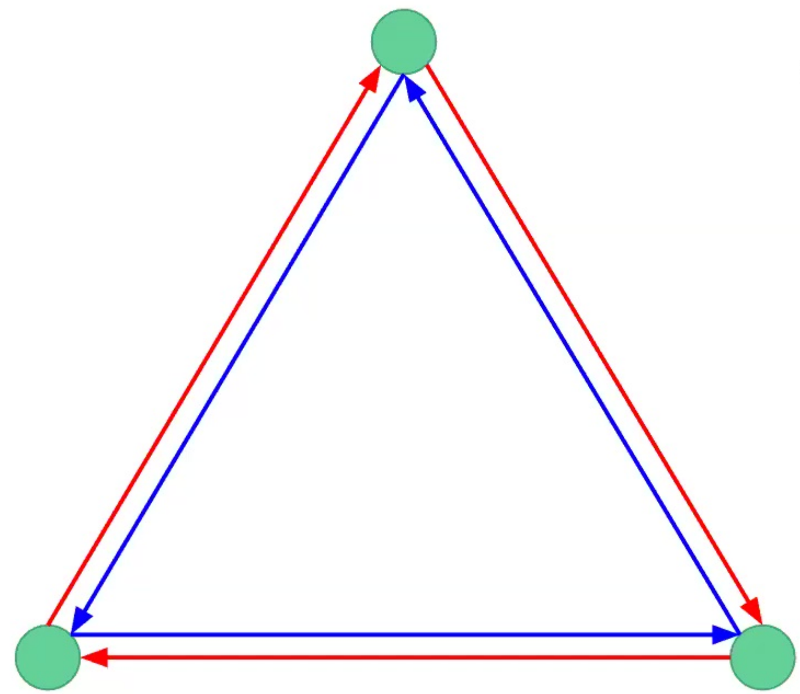

LISA: An overview

- Laser Interferometer Space Antenna
  - 3 spacecraft, 2.5×10^9 m separation
  - Earth-trailing "cartwheel" orbit
  - Mission lifetime: 4-10 years
- GW sensitivity band: 10^{-4}-1 Hz
- Effectively 3 interferometers
  - 2 Michelson: Measure GW polarizations
  - 1 Sagnac: Measure non-GW noise
- Selected as ESA L3, launch 2034
  - Mission adoption within 5 years

ESA

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

LISA: An overview

- LISA Pathfinder
  - 2 inertial masses, 40 cm separation
  - Orbit around Sun-Earth L1
  - Mission lifetime: Dec 2015-Jun 2017
- Not sensitive to GWs

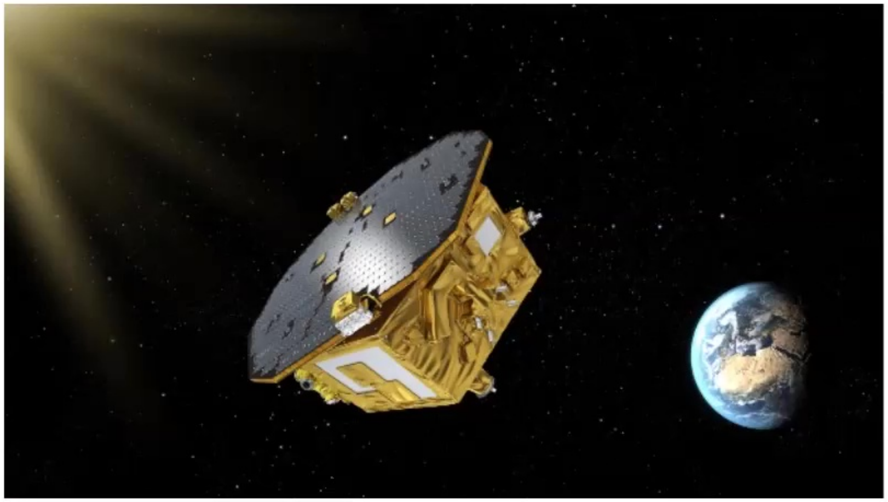

ESA

LISA: An overview

- LISA Pathfinder
  - 2 inertial masses, 40 cm separation
  - Orbit around Sun-Earth L1
  - Mission lifetime: Dec 2015-Jun 2017
- Not sensitive to GWs
- Demonstrated drag-free technology
  - Measured displacement to < 10^{-11} m
  - Achieved LISA noise requirement!

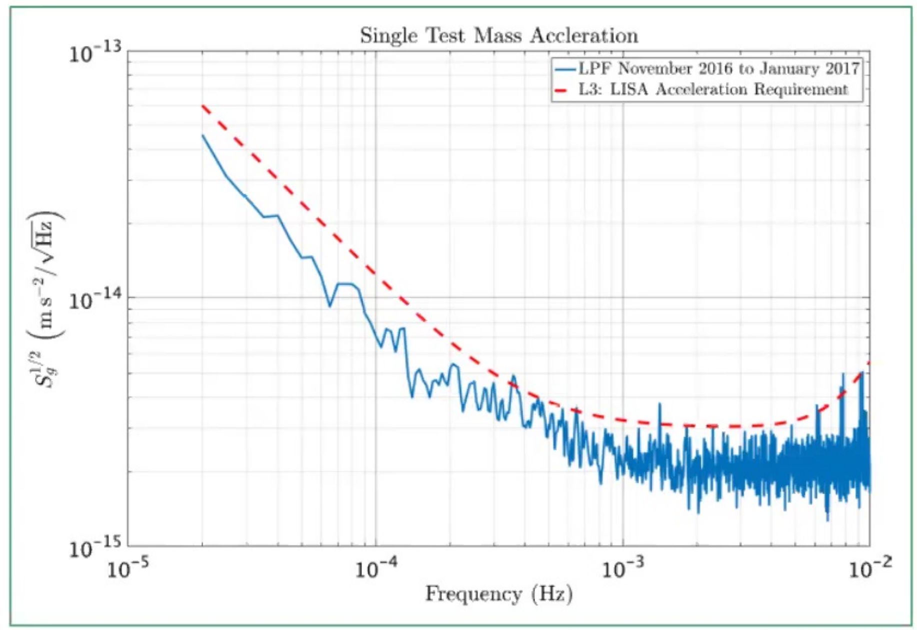

Danzmann et al. (2017)

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

The LISA inverse problem: Key issues

Issue 0: The LISA instrument

- Complicated instrument response
  - Orbits of 3 spacecraft & 6 test masses
  - Time-delay interferometry
- Data gaps & glitches
  - 7-hour gaps every 2 weeks
  - Optical-path & acceleration glitches
- Non-stationary noise
  - Time-evolving noise PSD

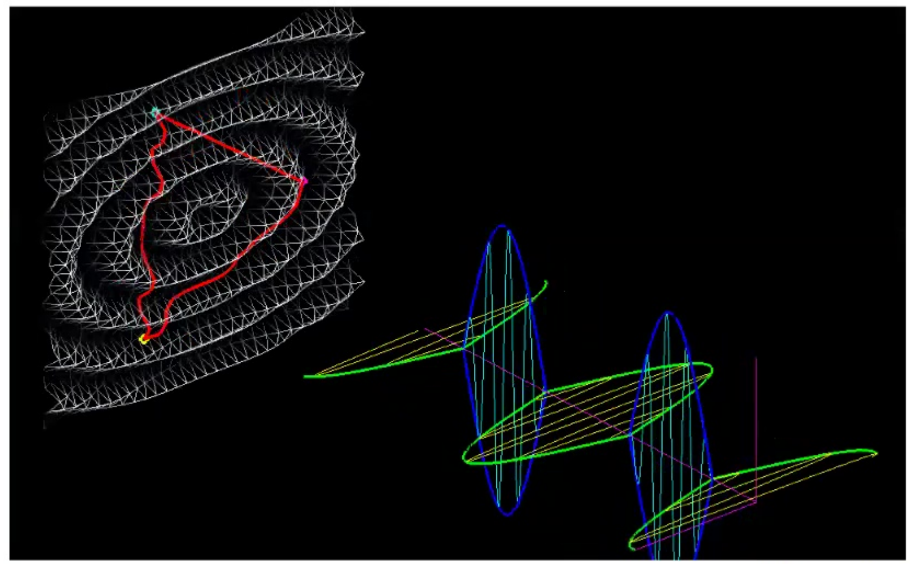

N. Douillet

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

The LISA inverse problem: Key issues

Issue 1: Signal confusion

- Many signals overlap in time/frequency
  - $10\text{-}10^3$ MBH mergers
  - $1\text{-}10^4$ EMRIs
  - $10^4\text{-}10^5$ resolvable Galactic WD binaries
  - Et cetera
- Cannot just subtract then move on

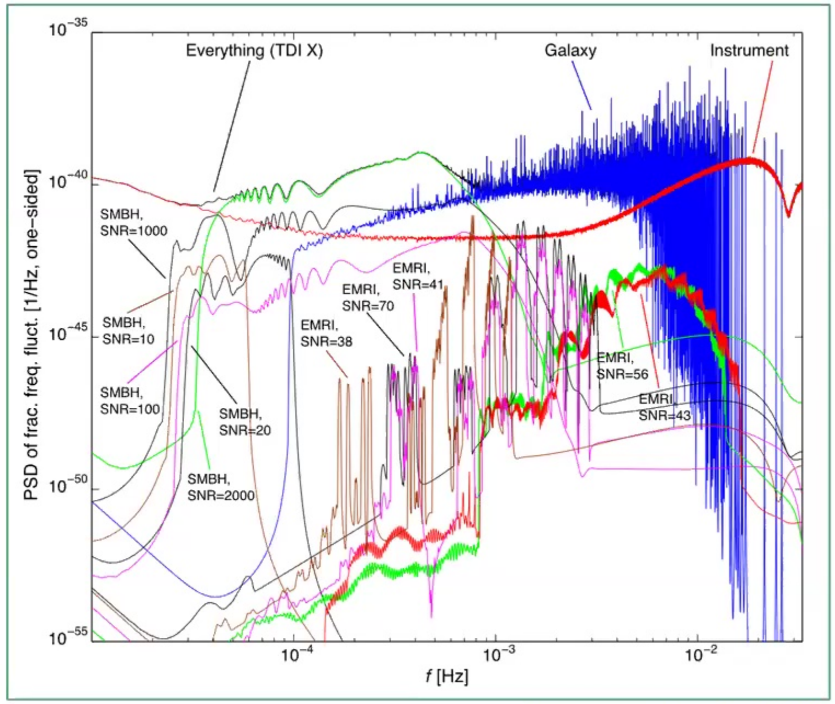

Arnaud et al. (2007)

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

The LISA inverse problem: Key issues

Issue 1: Signal confusion

- Many signals overlap in time/frequency
  - 10-10^3 MBH mergers
  - 1-10^4 EMRIs
  - 10^4-10^5 resolvable Galactic WD binaries
  - Et cetera
- Cannot just subtract then move on
- Global-solution algorithm
  - Not practical to do fully simultaneous fit
  - Separate pipelines for each source type
  - Repeated catalog update & lookup

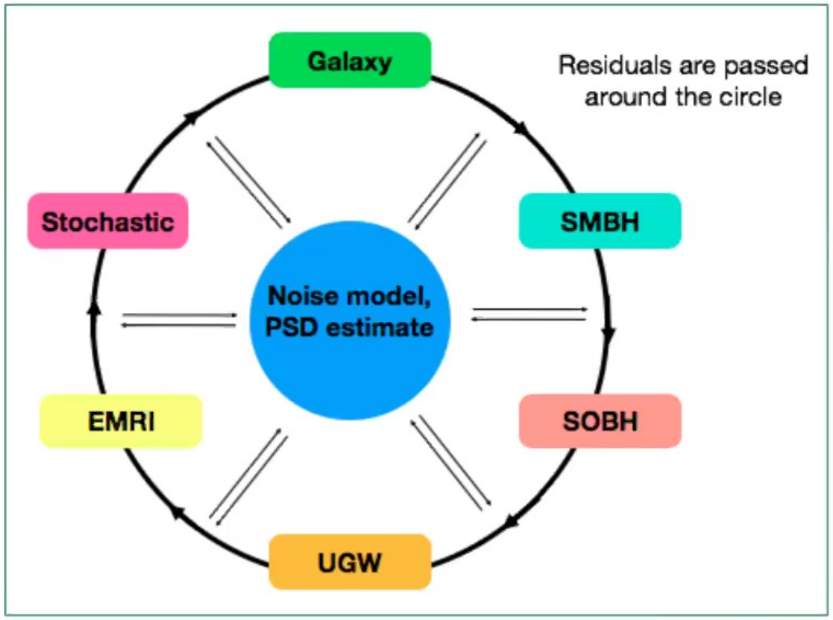

N. Cornish

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vl

The LISA inverse problem: Key issues

Issue 2: Search algorithms

- Parameter space can be huge
  - Dimensionality: Up to 20 per source type
  - Volume measure: Fisher information metric
  - Credible regions can be very localized
- Probability surface can be multimodal
  - Multiple signals per source type
  - Possibly severe degeneracies (more later)

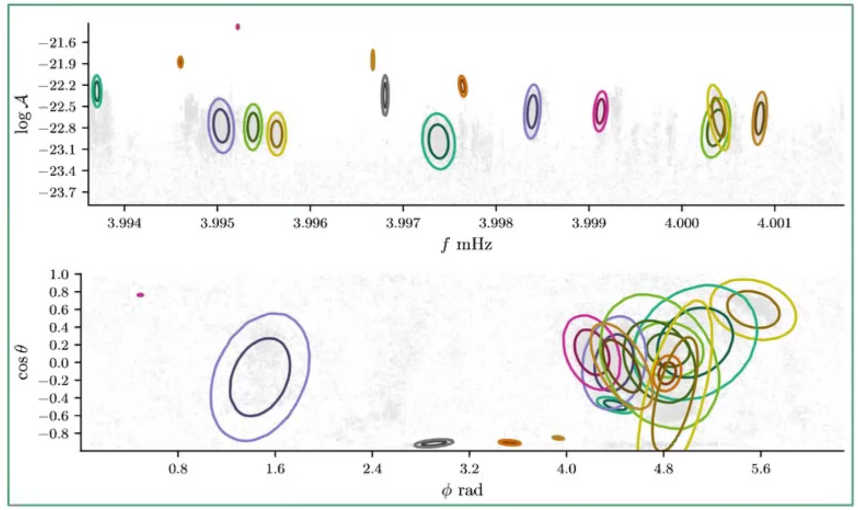

Littenberg et al. (2020)

The LISA inverse problem: Key issues

Issue 2: Search algorithms

- Parameter space can be huge
  - Dimensionality: Up to 20 per source type
  - Volume measure: Fisher information metric
  - Credible regions can be very localized
- Probability surface can be multimodal
  - Multiple signals per source type
  - Possibly severe degeneracies (more later)
- Main option is stochastic search
  - May be hierarchical: Coarse- to fine-grained
  - Many waveform evaluations are required

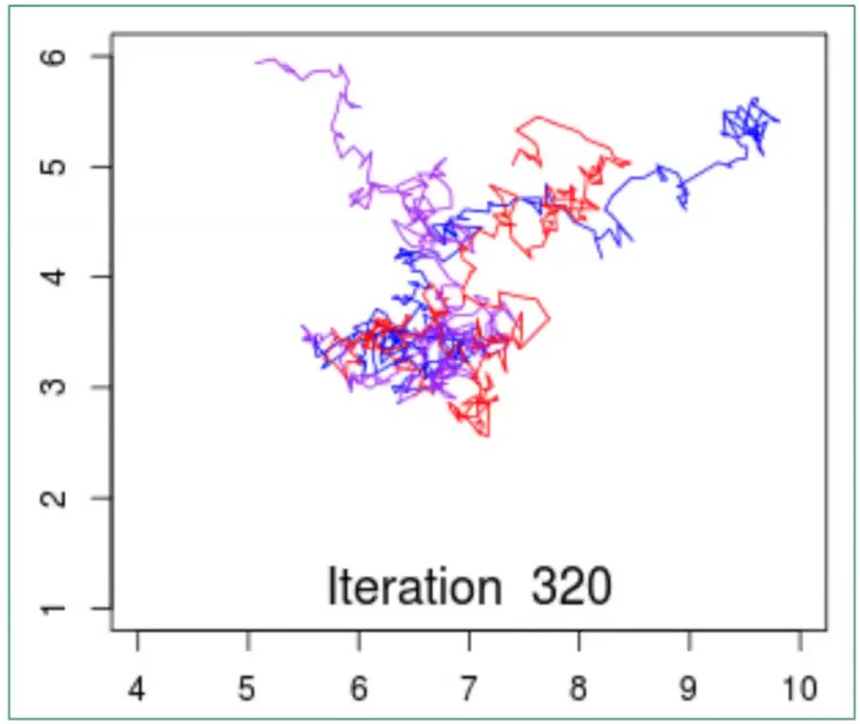

M. Joseph

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

The LISA inverse problem: Key issues

Issue 2: Search algorithms

- Parameter space can be huge
  - Dimensionality: Up to 20 per source type
  - Volume measure: Fisher information metric
  - Credible regions can be very localized
- Probability surface can be multimodal
  - Multiple signals per source type
  - Possibly severe degeneracies (more later)
- Main option is stochastic search
  - May be hierarchical: Coarse- to fine-grained
  - Many waveform evaluations are required

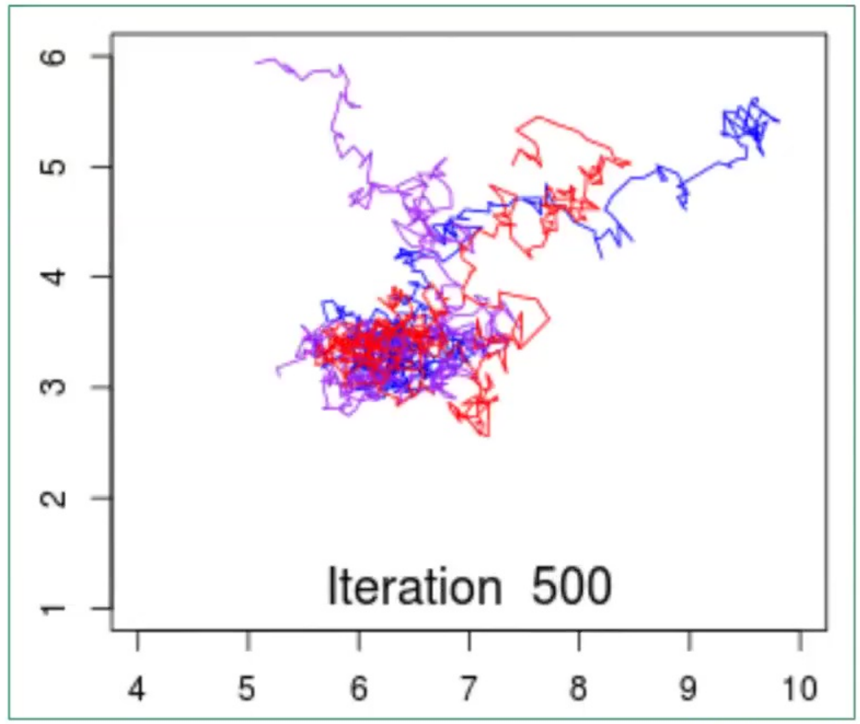

M. Joseph

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

The LISA inverse problem: Key issues

Issue 2: Search algorithms

- Parameter space can be huge
  - Dimensionality: Up to 20 per source type
  - Volume measure: Fisher information metric
  - Credible regions can be very localized
- Probability surface can be multimodal
  - Multiple signals per source type
  - Possibly severe degeneracies (more later)
- Main option is stochastic search
  - May be hierarchical: Coarse- to fine-grained
  - Many waveform evaluations are required

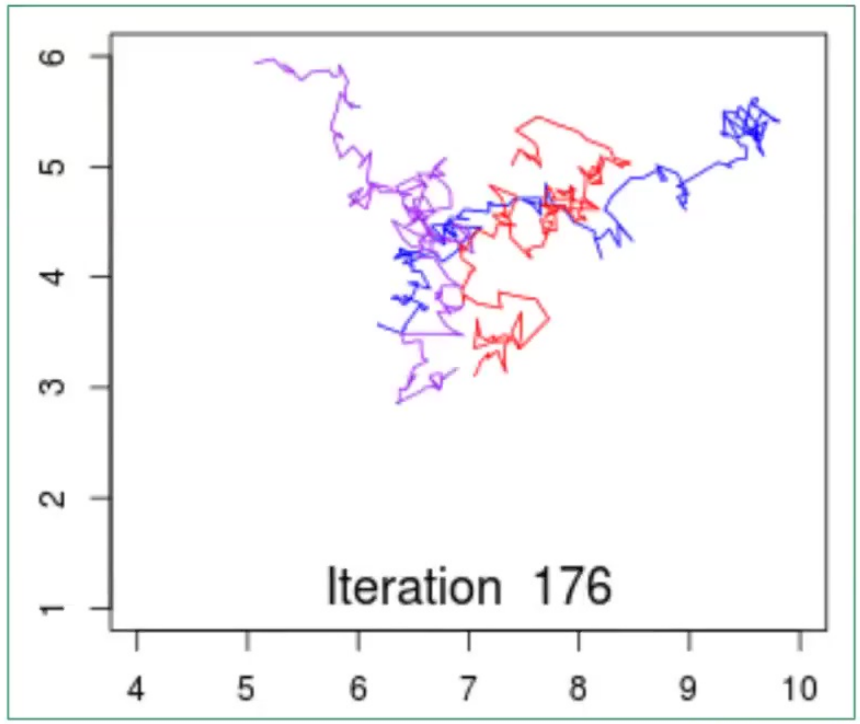

M. Joseph

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

The LISA inverse problem: Key issues

Issue 2: Search algorithms

- Parameter space can be huge
  - Dimensionality: Up to 20 per source type
  - Volume measure: Fisher information metric
  - Credible regions can be very localized
- Probability surface can be multimodal
  - Multiple signals per source type
  - Possibly severe degeneracies (more later)
- Main option is stochastic search
  - May be hierarchical: Coarse- to fine-grained
  - Many waveform evaluations are required

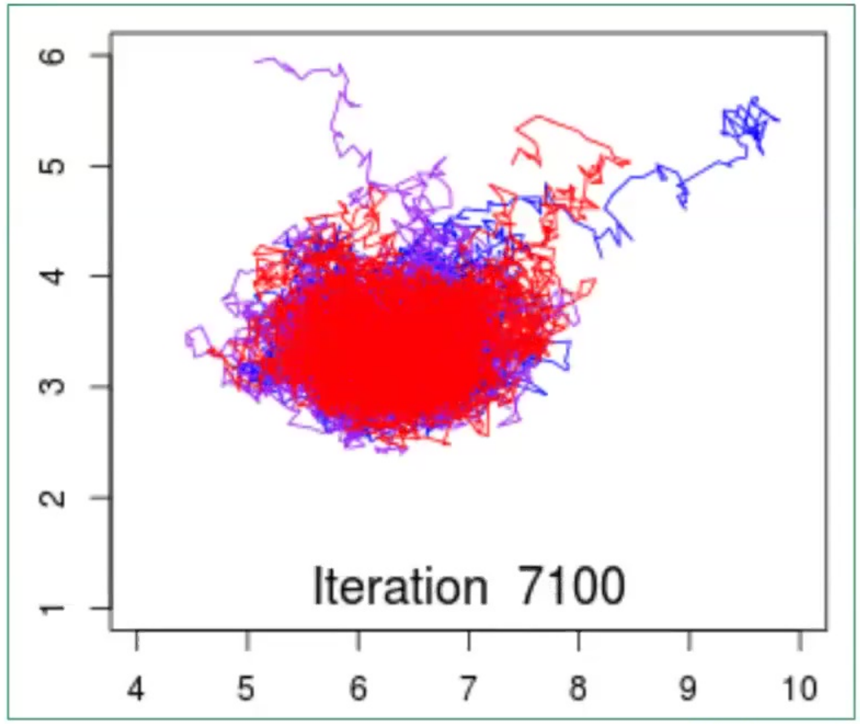

M. Joseph

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

The LISA inverse problem: Key issues

Issue 3: High-accuracy modeling

- Some strong-field, high-SNR sources
  - Complexity + strength ⇒ variability
  - Most interesting, but most difficult to model
- Need theoretical error < statistical error
  - Waveform mismatches of < 10^{-4}

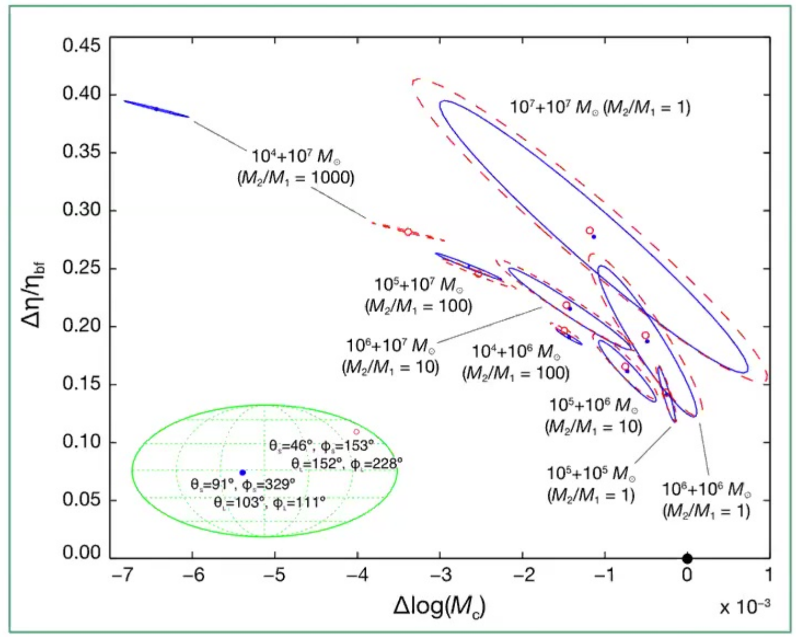

Cutler & Vallisneri (2007)

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

The LISA inverse problem: Key issues

Issue 3: High-accuracy modeling

- Some strong-field, high-SNR sources
  - Complexity + strength ⇒ variability
  - Most interesting, but most difficult to model
- Need theoretical error < statistical error
  - Waveform mismatches of < 10^{-4}
- Achieving accuracy is only the first step
  - Use accurate models to inform efficient ones
  - Efficient models must still be accurate enough
  - Standard LVK techniques may not scale

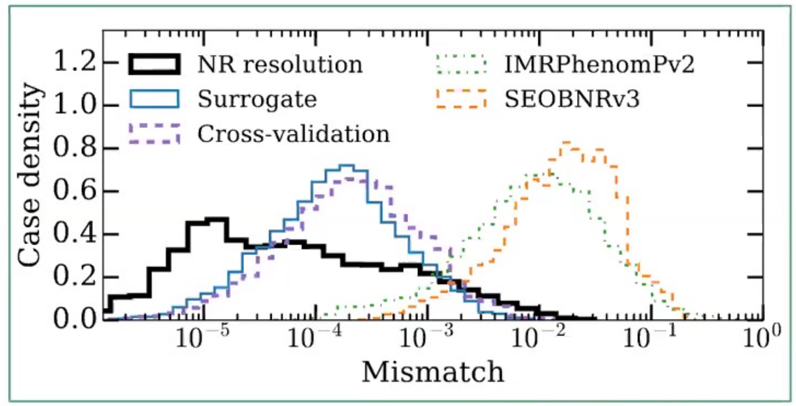

Blackman et al. (2017)

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vh

EMRIs: An overview

- Extreme-mass-ratio inspirals
  - Capture of stellar-mass compact object by MBH
  - Long-lived in LISA band: 10^4-10^5 strong-field orbits
  - Extreme precession, can be eccentric up to plunge

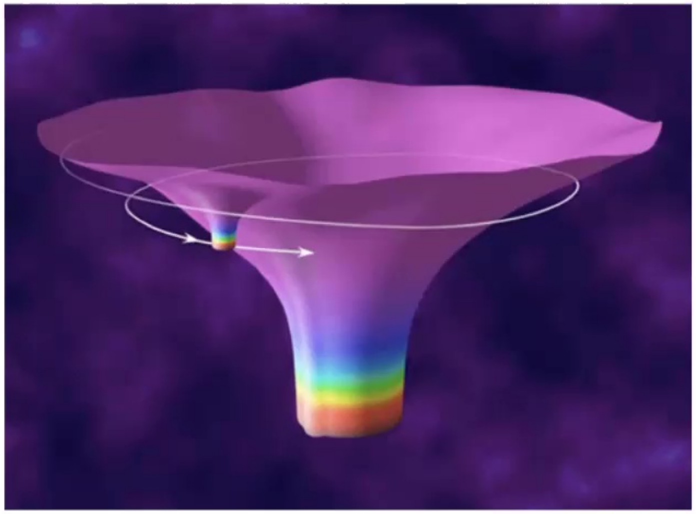

NASA

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vl

EMRIs: An overview

- Extreme-mass-ratio inspirals
  - Capture of stellar-mass compact object by MBH
  - Long-lived in LISA band: 10^4-10^5 strong-field orbits
  - Extreme precession, can be eccentric up to plunge
- The physics of EMRIs
  - Cannot use NR
  - BH perturbation theory with small mass ratio ($< 10^{-4}$)
  - Calculate effective SF on Kerr orbits
  - Need SF up to 2nd-order dissipative
  - Recent breakthrough at 2nd-order [Pound et al., 2020]

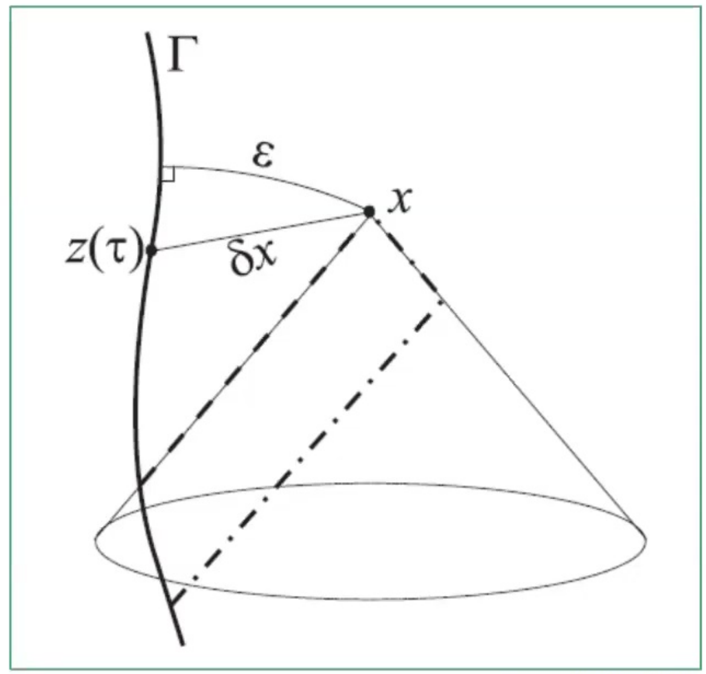

Barack (2009)

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

EMRIs: An overview

- The astrophysics of EMRIs
  - Event-rate estimates: 1-10^4 (per LISA) [Babak et al., 2017]
  - Brown-dwarf “problem” [Gourgoulhon et al., 2019; Amaro-Seoane, 2019; Amaro-Seoane, private comm.]
  - Possibly many other environmental effects

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

EMRIs: An overview

- The astrophysics of EMRIs
  - Event-rate estimates: 1-10^4 (per LISA) [Babak et al., 2017]
  - Brown-dwarf "problem" [Gourgoulhon et al., 2019; Amaro-Seoane, 2019; Amaro-Seoane, private comm.]
  - Possibly many other environmental effects
- Why bother? Environment may mess up modeling/analysis, or even existence
- There are several good reasons
  - High-precision science: MBH & galaxy astrophysics, tests of fundamental physics
  - Global fit: Even if LISA data contains just 1 EMRI signal, it will have to be accurately subtracted
  - Challenge: Everybody likes one

The EMRI inverse problem: Difficulties

- Everything is hard for EMRIs (+ astrophysics, but let's not go into that here)
- Difficulty 1: Accuracy
  - EMRIs are strong-field, high-SNR sources that require accurate modeling to find & characterize
  - Phasing accurate to post-1-adiabatic order will be enough, but we are not there yet
- Difficulty 2: Efficiency
  - EMRI signals are long-lived with rich harmonic content; they are costly to model & analyze
  - Stochastic algorithms in data analysis require bulk generation of waveforms (at least billions)
- Difficulty 3: Extensiveness
  - Even the "leading-order" space of EMRI orbits is gargantuan in terms of information volume
  - Waveforms are only half the battle: detector response encodes important extrinsic effects

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

Outline

- The inverse problem in GW astronomy
- LISA: An overview
  - The LISA inverse problem
- EMRIs: An overview
  - The EMRI inverse problem
- Next-generation EMRI waveform models
  - Model 1: Eccentric Schwarzschild, adiabatic
  - Model 2: Generic Kerr, PN adiabatic
- The uncharted space of EMRI signals
  - Approximative techniques
  - Here there be degeneracies

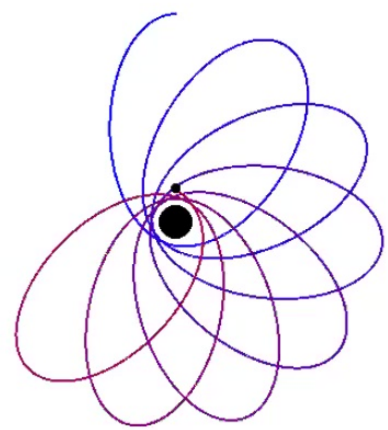

Next-generation models: Definitions

- Waveform model:
  - Accurate (fully relativistic)
  - Not necessarily efficient & not necessarily extensive
  - Only examples are adiabatic (Teukolsky-based) models
  - Previous adiabatic models are very slow & limited in scope

Next-generation models: Definitions

- Waveform model:
    - Accurate (fully relativistic)
    - Not necessarily efficient & not necessarily extensive
    - Only examples are adiabatic (Teukolsky-based) models
    - Previous adiabatic models are very slow & limited in scope
- Template model:
    - Not necessarily accurate
    - Efficient & extensive
    - Only examples are semi-relativistic (kludge) models
    - Semi-relativistic models have an inherent cap on accuracy

Next-generation models: Definitions

- Waveform model:
  - Accurate (fully relativistic)
  - Not necessarily efficient & not necessarily extensive
  - Only examples are adiabatic (Teukolsky-based) models
  - Previous adiabatic models are very slow & limited in scope
- Template model:
  - Not necessarily accurate
  - Efficient & extensive
  - Only examples are semi-relativistic (kludge) models
  - Semi-relativistic models have an inherent cap on accuracy
- Aim is for next generation of models to be accurate, efficient & extensive
- The 2 new models are accurate & efficient, but not yet extensive

Next-generation models: Framework

- Standard modular description
  - Angular & frequency-based decomposition
  - Osculating geodesics
- Generic Kerr orbits
  - Need schemes to evolve through resonances
  - Need secondary spin, mass/spin evolution, etc.
- Angular dependence
  - Spheroidal harmonics with spin weight -2
- Inspiral trajectory (+ mode phasing)
  - Post-1-adiabatic order
- Mode amplitudes
  - Adiabatic order

$$h(t) = \frac{1}{r}\sum_{lmkn} A_{lmkn}(t) e^{-i\Phi_{mkn}(t)} V_{lmkn}(\theta,\phi)$$

$$G(t) \equiv (p(t), e(t), \iota(t))$$

$$V_{lmkn}(\theta,\phi) = {}_{-2}S_{lmkn}(\theta) e^{im\phi}$$

$$\Phi_{mkn}(t) = \text{init.} + \int_{t_0}^{t} dt'\, \omega_{mkn}(G(t')) + \text{osc.}$$

$$A_{lmkn}(t) = -\frac{2Z^{\infty}_{lmkn}(G(t))}{\omega^2_{mkn}(G(t))}$$

Next-generation models: Inspiral trajectory

- Trajectory models
  - PN flux-based (partial coverage, only for search)
  - Teukolsky flux-based (only for search)
  - NIT [van de Meent & Warburton, 2018]
  - Two-timescale framework

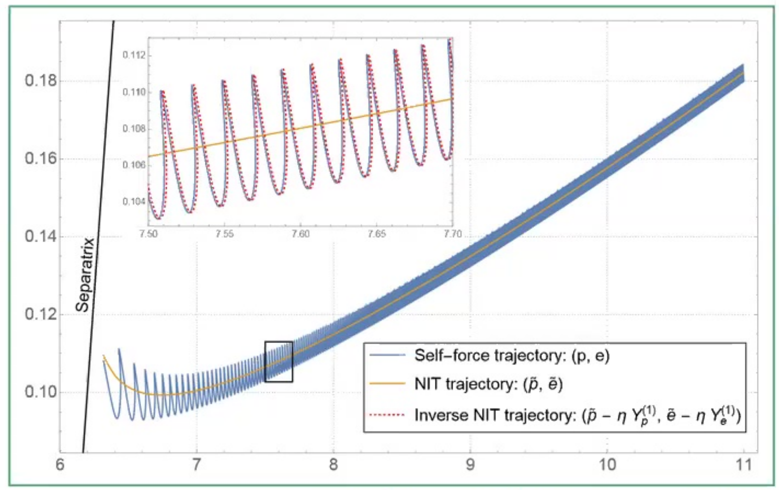

van de Meent & Warburton (2018)

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

Next-generation models: Inspiral trajectory

- Trajectory models
  - PN flux-based (partial coverage, only for search)
  - Teukolsky flux-based (only for search)
  - NIT [van de Meent & Warburton, 2018]
  - Two-timescale framework
- Resonance schemes
  - Fit jump vectors over resonance surfaces
  - Switch between fast & slow evolution

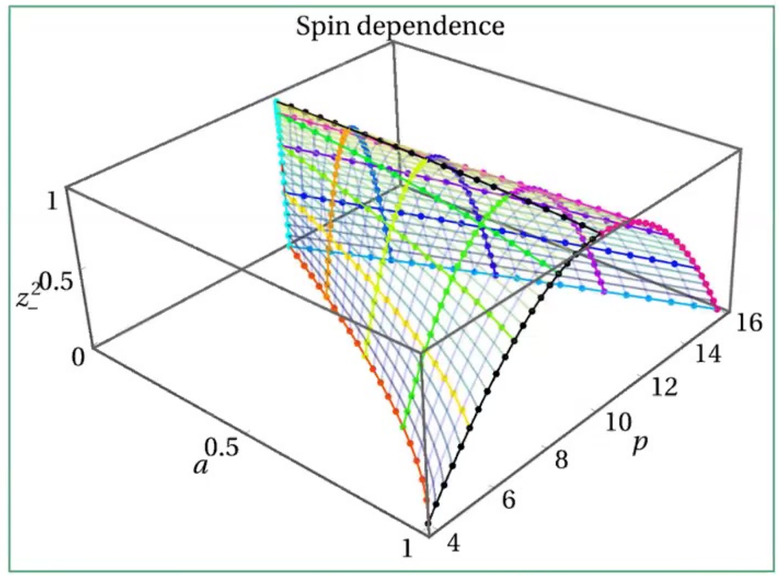

Brink, Geyer & Hinderer (2013)

Next-generation models: Mode amplitudes

- Amplitude models
  - PN amplitudes (partial coverage)
  - Interpolate Teukolsky amplitudes (local)
  - Fit Teukolsky amplitudes (global)
- Are adiabatic amplitudes OK?

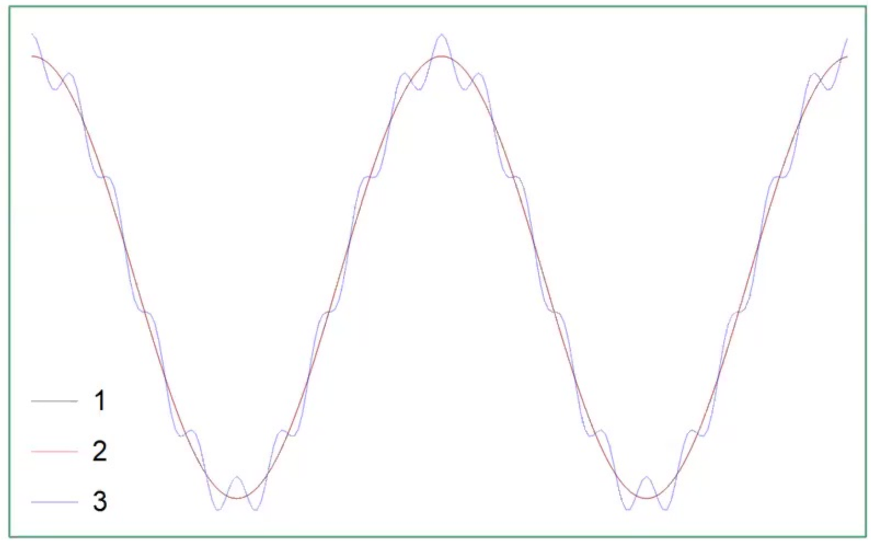

$1 : \cos(\omega t),\ T = 10^3$ cycles

$2 : \cos((1 + 10^{-4})\omega t)$

$3 : \cos(\omega t) + 10^{-1}\cos(100\omega t)$

Next-generation models: Mode amplitudes

- Amplitude models
  - PN amplitudes (partial coverage)
  - Interpolate Teukolsky amplitudes (local)
  - Fit Teukolsky amplitudes (global)
- Are adiabatic amplitudes OK?
  - Yes, probably
- Interpolation/fitting is difficult though
  - High-dimensional mode space: 10^3-10^5
  - Kerr: > 2-dimensional geodesic space
  - Many evaluations along trajectory

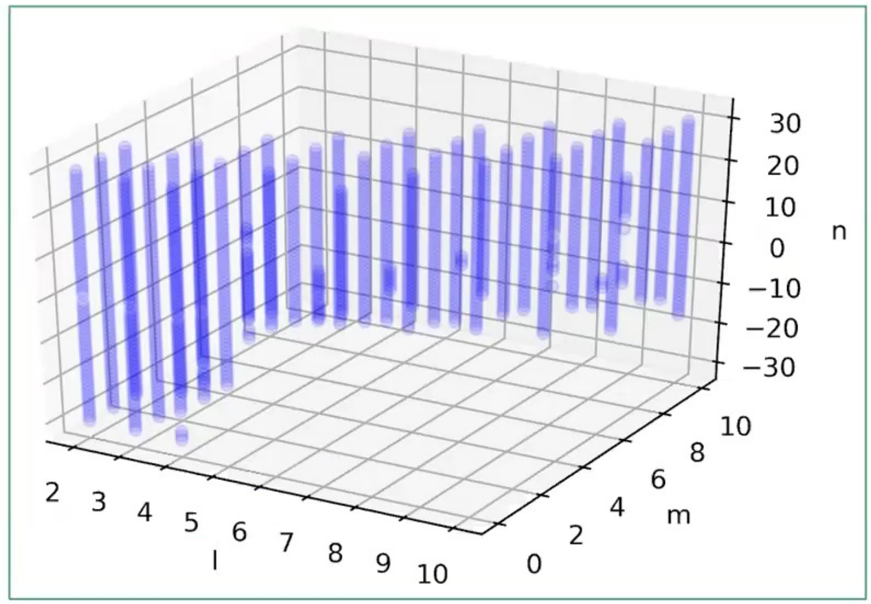

Model 1 [AC, Katz, Warburton & Hughes, in rev.]

- Standard modular description
  - Angular & frequency-based decomposition
  - Osculating geodesics
- Eccentric Schwarzschild orbits
  - Neglect resonances
  - Neglect secondary spin, mass/spin evolution, etc.
- Angular dependence
  - Spherical harmonics with spin weight -2
- Inspiral trajectory (+ mode phasing)
  - Adiabatic order
- Mode amplitudes
  - Adiabatic order

$$h(t) = \frac{1}{r}\sum_{lmn} A_{lmn}(t) e^{-i\Phi_{mn}(t)} V_{lm}(\theta,\phi)$$

$$G(t) \equiv (p(t), e(t))$$

$$V_{lm}(\theta,\phi) = {}_{-2}Y_{lm}(\theta) e^{im\phi}$$

$$\Phi_{mn}(t) = \text{init.} + \int_{t_0}^{t} dt'\, \omega_{mn}(G(t'))$$

$$A_{lmn}(t) = -\frac{2Z^{\infty}_{lmn}(G(t))}{\omega^2_{mn}(G(t))}$$

Model 1 [AC, Katz, Warburton & Hughes, in rev.]

- Trajectory model: Teukolsky flux-based
  - Numerical data: 10^{-12} fractional error on amplitudes, 1640 points in geodesic space
  - Bicubic flux interpolation, 8th-order Runge-Kutta trajectory integration

$$(p_0, e_0, \eta) \mapsto (G(t), \Phi_{mn}(t))$$

- Amplitude model: Neural-network fit to Teukolsky amplitudes
  - Numerical data: Same as trajectory model

$$2 \leq l \leq 10,\ m \leq l,\ |n| \leq 30$$

$$(p, e) \mapsto \mathrm{vec}(A_{lmn}) \in \mathbb{C}^{3843} \cong \mathbb{R}^{7686}$$

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

Model 1: Mode amplitudes

- ROMAN [AC, Galley & Vallisneri, 2019]
  - Combination of ROM & deep learning
  - Alternative to surrogate + ROQ framework
- Construct reduced basis for mode set

$$\mathrm{vec}(A_{lmn})(p,e) = \sum_i \alpha_i(p,e)\mathbf{e}_i \equiv \alpha(p,e)$$

$$(p,e) \mapsto \alpha \in \mathbb{C}^{99} \cong \mathbb{R}^{198}$$

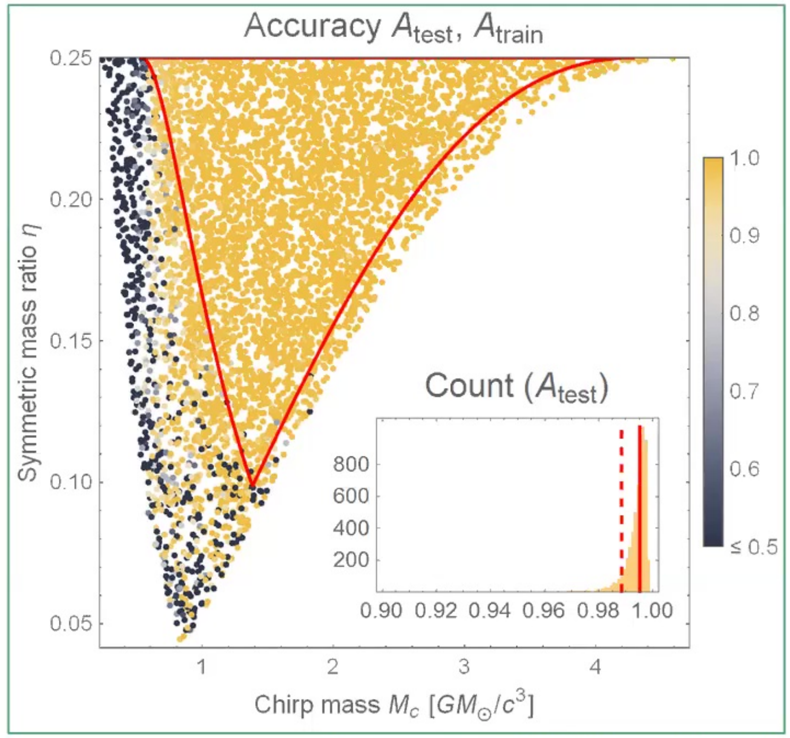

Chua, Galley & Vallisneri (2019)

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XlSCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vl

Model 1: Mode amplitudes

- ROMAN [AC, Galley & Vallisneri, 2019]
  - Combination of ROM & deep learning
  - Alternative to surrogate + ROQ framework
- Construct reduced basis for mode set

$$\mathrm{vec}(A_{lmn})(p,e) = \sum_i \alpha_i(p,e)\mathbf{e}_i \equiv \alpha(p,e)$$

$$(p,e) \mapsto \alpha \in \mathbb{C}^{99} \cong \mathbb{R}^{198}$$

- Train neural network on reduced map
  - Network: Multilayer perceptron, 20 layers
  - Initial eccentricities up to 0.7
  - Plunge eccentricities up to 0.5
  - Separations from LSO + 0.2M to LSO + 10M

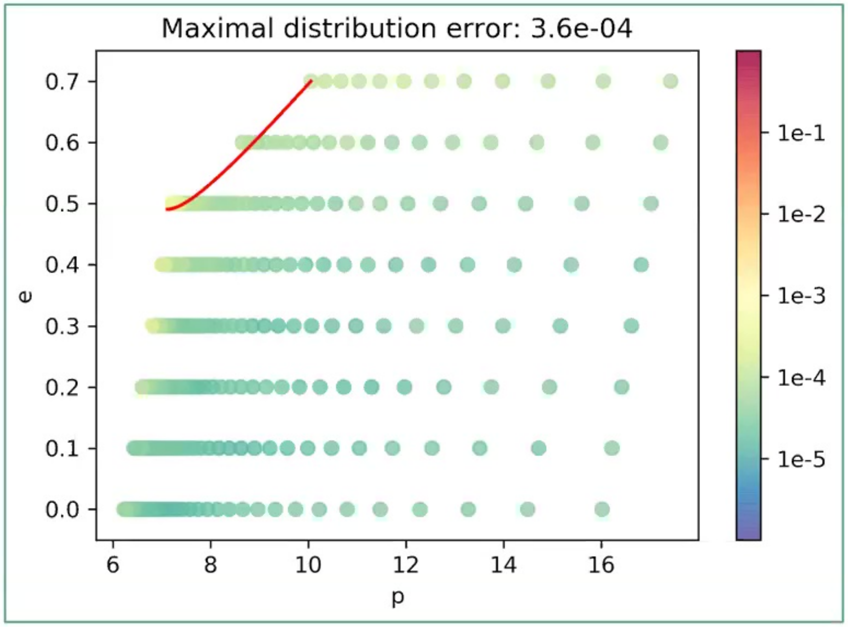

$$\mathrm{error} = 1 - \frac{\Re(\alpha^\dagger \alpha_{\mathrm{num}})}{\sqrt{(\alpha^\dagger\alpha)(\alpha^\dagger_{\mathrm{num}}\alpha_{\mathrm{num}})}}$$

Model 1: Mode amplitudes (an aside)

- Why not apply ROM directly to waveforms?
  - Circular Schwarzschild IMRI: 1 parameter, < 200 cycles, 22 modes [Rifat et al., 2020]
  - A usable EMRI surrogate is unlikely to cover more than a miniscule patch of parameter space
  - Accuracy is also an issue: Even best NR surrogates have maximal mismatches > 10^{-3}

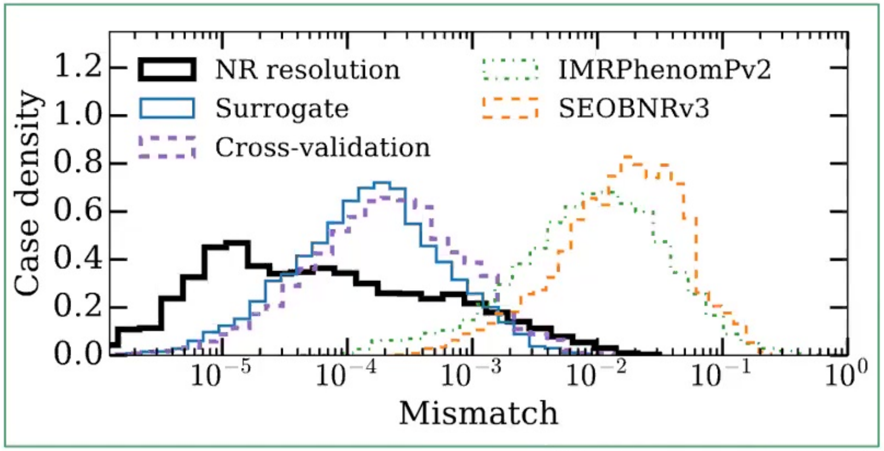

Blackman et al. (2017)

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

Model 1: Implementation & benchmarking

- GPU acceleration in GW modeling
  - MBH waveforms (PhenomHM) [Katz et al. (+ AC), 2020]
  - Kludge waveforms (AAK) [AC with Katz]
- 3 separate implementations
  - Slow: Bicubic amplitude interpolation, dense trajectory integration, no approximations
  - Fast (CPU): Neural-network amplitude fit, sparse trajectory integration, mode selection
  - Fast (GPU): Same as above, but native to GPU

bhptoolkit.org/FastEMRIWaveforms

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vl

Model 1: Implementation & benchmarking

- Speed
  - Analysis-length waveforms: 1 year at 0.1 Hz
  - Full harmonic content: All relevant modes
  - Maximal wall time (for GPU): < 200 ms

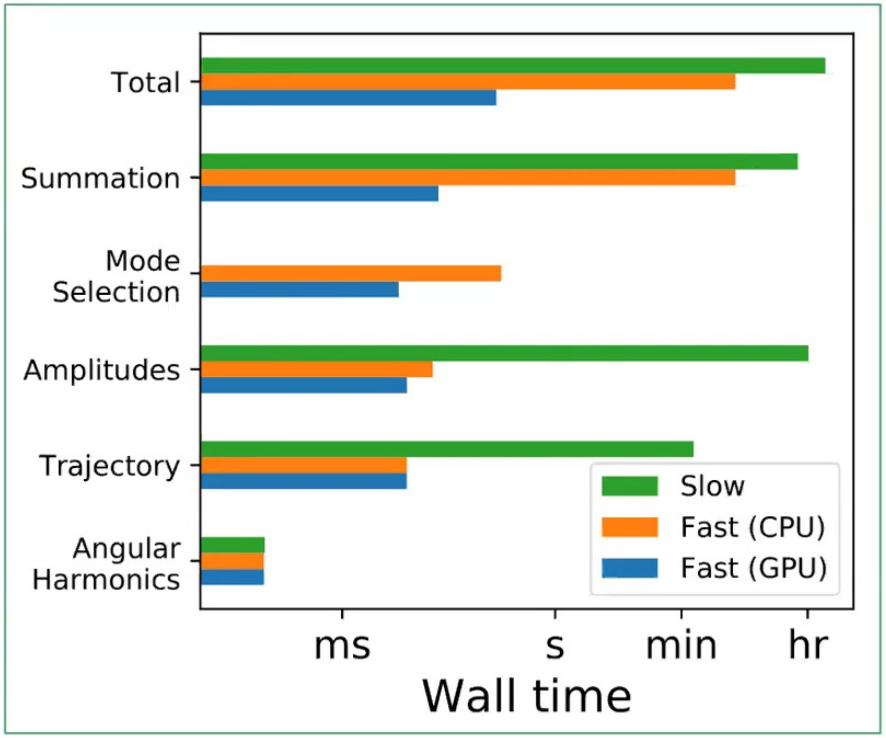

Model 1: Implementation & benchmarking

- Speed
  - Analysis-length waveforms: 1 year at 0.1 Hz
  - Full harmonic content: All relevant modes
  - Maximal wall time (for GPU): < 200 ms
- Accuracy
  - Slow waveform is taken as fiducial
  - Maximal mismatch: $< 4 \times 10^{-4}$
  - No theoretical-error bias up to SNR 100

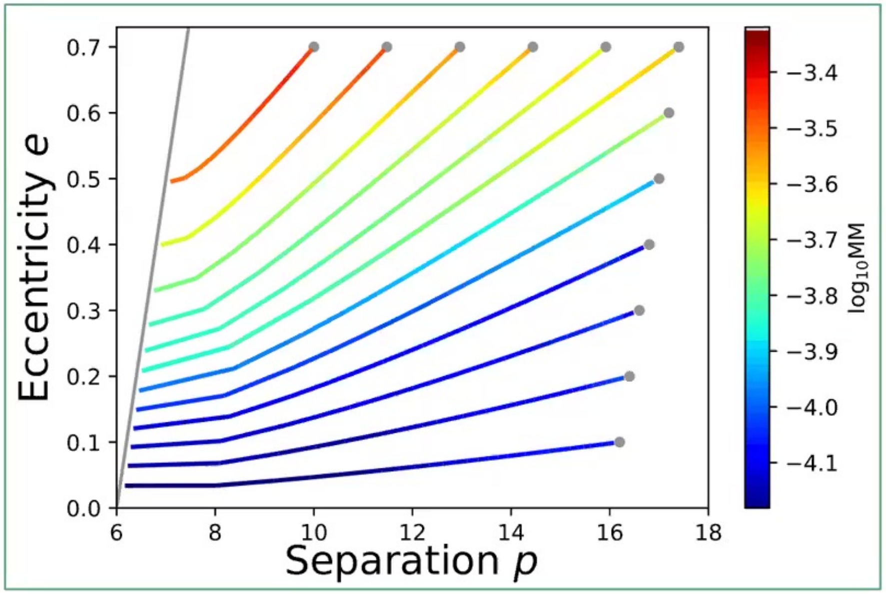

Model 1: Implementation & benchmarking

- Speed
  - Analysis-length waveforms: 1 year at 0.1 Hz
  - Full harmonic content: All relevant modes
  - Maximal wall time (for GPU): < 200 ms
- Accuracy
  - Slow waveform is taken as fiducial
  - Maximal mismatch: $< 4 \times 10^{-4}$
  - No theoretical-error bias up to SNR 100

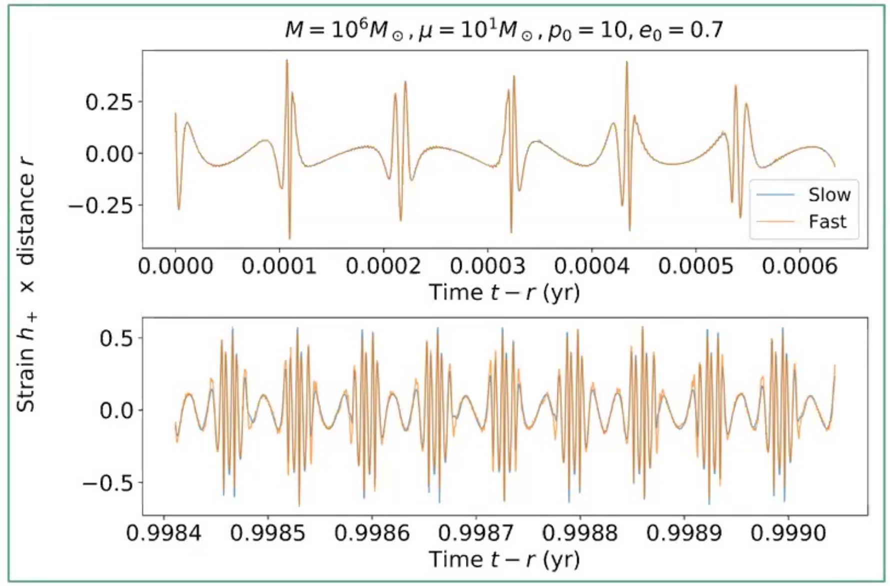

Model 2 [Isoyama et al. (+ AC), in prep.]

- Standard modular description
  - Angular & frequency-based decomposition
  - Osculating geodesics
- Generic Kerr orbits
  - Approximate resonances
  - Neglect secondary spin, mass/spin evolution, etc.
- Angular dependence
  - Spheroidal harmonics with spin weight -2
- Inspiral trajectory (+ mode phasing)
  - PN approximation to adiabatic order
- Mode amplitudes
  - PN approximation to adiabatic order

$$h(t) = \frac{1}{r} \sum_{lmkn} A_{lmkn}(t) e^{-i\Phi_{mkn}(t)} V_{lmkn}(\theta, \phi)$$

$$G(t) \equiv (p(t), e(t), \iota(t))$$

$$V_{lmkn}(\theta, \phi) = {}_{-2}S_{lmkn}(\theta) e^{im\phi}$$

$$\Phi_{mkn}(t) = \text{init.} + \int_{t_0}^{t} dt' \, \omega_{mkn}(G(t'))$$

$$A_{lmkn}(t) = -\frac{2Z^{\infty}_{lmkn}(G(t))}{\omega^2_{mkn}(G(t))}$$

Model 2 [Isoyama et al. (+ AC), in prep.]

- Trajectory model: 5PN $O(e^{10})$ flux-based
    - Generally more accurate (even at < 10*M*) than Teukolsky-fitted fluxes [Gair & Glampedakis, 2006]

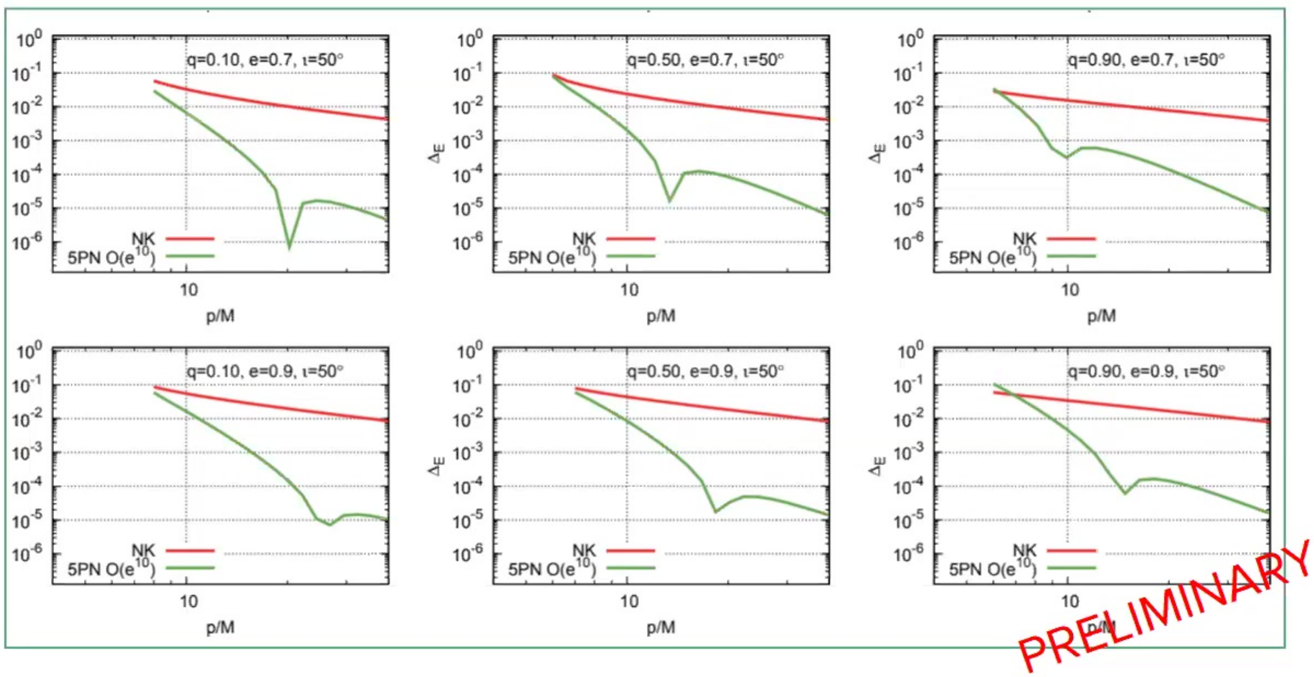

Model 2 [Isoyama et al. (+ AC), in prep.]

- Trajectory model: 5PN $O(e^{10})$ flux-based
  - Generally more accurate (even at < 10M) than Teukolsky-fitted fluxes [Gair & Glampedakis, 2006]
- Amplitude model: 5PN $O(e^{10})$ amplitudes
  - Domain of validity: TBD
- CPU implementation: < 10 times slower than CPU implementation of Model 1
- Can include approximate resonances
  - Resonant fluxes from mode amplitudes (+ horizon amplitudes) [Flanagan, Hughes & Ruangsri, 2014]
  - Directly compute leading-order jumps at resonance surfaces

Next-generation models: Future work

- Improve source-end extensiveness (Model 1)
  - Eccentric equatorial Kerr: Add primary spin
  - Partial post-adiabatic trajectory: Add 1st-order SF
  - Incorporate resonance schemes
- Improve efficiency (Model 2): Probably GPU implementation
- Improve detector-end efficiency & extensiveness
  - Integrate with GPU-accelerated full LISA response models [AC with Katz, Bayle, Vallisneri]
- Other representations
  - Frequency domain: Higher-order SPA [AC with Hughes, Warburton, Katz]
  - Time-frequency domain: STFT

The uncharted space of EMRI signals

- Consider a single LISA data channel of 2 years at 0.1 Hz
  - Data space D is a 6×10^6-dimensional Hilbert space
  - Inner product is the usual noise-weighted one
- Consider an EMRI waveform model with *n* parameters
  - Model is an embedding of parameter space P in D
  - Signal space S is an *n*-submanifold of D
  - Fisher information metric on P is pullback of metric on D

$$\langle a|b\rangle = 4\Re\sum_{f>0}^{\mathrm{Nyq}} df\,\frac{\tilde{a}^*(f)\tilde{b}(f)}{S_n(f)}$$

$$\theta \mapsto h(\theta), \quad \Gamma_{ij} = \langle \partial_{\theta_i} h | \partial_{\theta_j} h \rangle$$

The uncharted space of EMRI signals

- Consider a single LISA data channel of 2 years at 0.1 Hz
  - Data space D is a 6 × 10[6]-dimensional Hilbert space
  - Inner product is the usual noise-weighted one

$$\langle a|b\rangle = 4\Re\sum_{f>0}^{\mathrm{Nyq}} df\,\frac{\tilde{a}^*(f)\tilde{b}(f)}{S_n(f)}$$

- Consider an EMRI waveform model with *n* parameters
  - Model is an embedding of parameter space P in D
  - Signal space S is an *n*-submanifold of D
  - Fisher information metric on P is pullback of metric on D

$$\theta \mapsto h(\theta), \quad \Gamma_{ij} = \langle \partial_{\theta_i} h | \partial_{\theta_j} h \rangle$$

- Consider negative log-likelihood with noise-free injection as a function on S
  - Essentially distance (in D) from some fixed point on S

$$-2\ln L = \langle h - h_\star | h - h_\star \rangle$$

The uncharted space of EMRI signals

- What do we know about the space S?
    - It is very big (integral of Fisher volume form) [Gair et al., 2004]
    - Local likelihood correlations are understood (Fisher correlations)
    - Likelihood has nonlocal secondary modes (degeneracies) [MLDCs]
    - That's all: Nothing is known about prevalence, distribution & properties of secondaries

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

The uncharted space of EMRI signals

- What do we know about the space S?
  - It is very big (integral of Fisher volume form) [Gair et al., 2004]
  - Local likelihood correlations are understood (Fisher correlations)
  - Likelihood has nonlocal secondary modes (degeneracies) [MLDCs]
  - That's all: Nothing is known about prevalence, distribution & properties of secondaries
- Heuristic study of EMRI signal space [AC & Cutler, in prep.]
  - Pick a relatively central point in parameter space as injection
  - Map out the likelihood or some related surface at high resolution
  - Target coverage is full parameter space, or as close as we can get
  - Surprisingly nontrivial task: All standard waveform, sampling & clustering tools are ill-suited

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

Approximative techniques

- Stripped-down kludge waveform
  - AAK + Cutler response [AC, Moore & Gair, 2017]
  - Only 4 modes (reasonable at low eccentricity)
  - Noise-weighted amplitude & phase trajectories

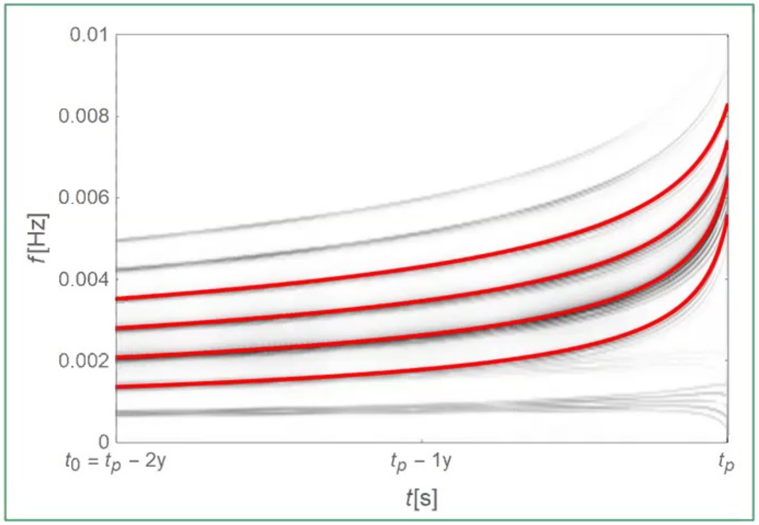

Approximative techniques

- Stripped-down kludge waveform
  - AAK + Cutler response [AC, Moore & Gair, 2017]
  - Only 4 modes (reasonable at low eccentricity)
  - Noise-weighted amplitude & phase trajectories
- Phase-proximity inner product
  - Amplitude product where phase difference < 1 rad
  - Approximate envelope of full inner product
  - Overestimates full overlaps by factor of 2-3
- Combination is faster than full AAK on GPU

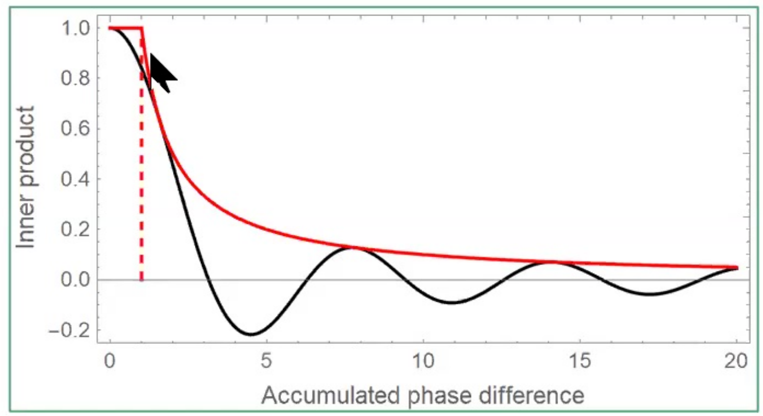

$$\langle h_1 | h_2 \rangle \approx \sum_{t \in T} dt \, A_1(t) A_2(t)$$

$$T = \{t : |\Phi_1(t) - \Phi_2(t)| < 1\}$$

Approximative techniques

- Traditional likelihood has large-scale SNR gradients
  - We really want to look at the overlap surface instead
- Exploratory likelihood
  - Suppresses volume in probability tails, so as to better locate & explore secondaries
  - Secondaries still cannot be resolved directly from posterior samples
  - Used to obtain large set of high-overlap points for clustering analysis

$$\mathrm{O}(h, h_\star) = \frac{\langle h | h_\star \rangle}{\sqrt{\langle h | h \rangle \langle h_\star | h_\star \rangle}}$$

$$\ln L_{\mathrm{expl}} = \mathrm{O}(h, h_\star) + 100 \min\left\{\mathrm{O}(h, h_\star) - 0.5, 0\right\} - 1$$

Approximative techniques

$$O(h, h_\star) = \frac{\langle h | h_\star \rangle}{\sqrt{\langle h|h\rangle \langle h_\star | h_\star \rangle}}$$

- Traditional likelihood has large-scale SNR gradients
  - We really want to look at the overlap surface instead
- Exploratory likelihood
  - Suppresses volume in probability tails, so as to better locate & explore secondaries
  - Secondaries still cannot be resolved directly from posterior samples
  - Used to obtain large set of high-overlap points for clustering analysis

$$\ln L_{\text{expl}} = O(h, h_\star) + 100 \min \{O(h, h_\star) - 0.5, 0\} - 1$$

- Bespoke clustering algorithm
  - Centroid-based algorithms do not count clusters
  - Density-based algorithms rely on perfect sampling & define cluster membership arbitrarily
  - Our algorithm uses overlaps along connecting lines to define notions of connectedness

Degeneracies: Evidence

- Threshold-SNR (20) injection, 6 intrinsic parameters, posterior bounds × 10
  - 30 secondaries: Full overlaps from 0.45 to 0.72

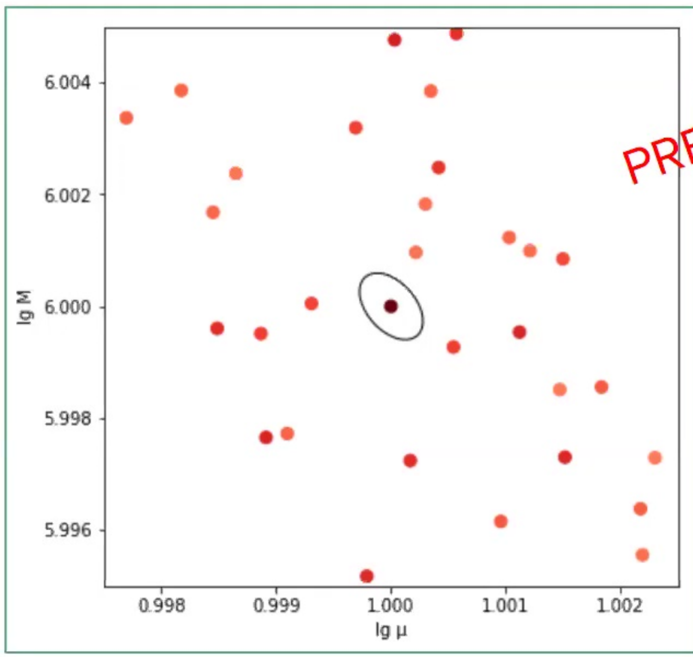

PRELIMINARY

Degeneracies: Evidence

- Threshold-SNR (20) injection, 6 intrinsic parameters, posterior bounds × 10
  - 30 secondaries: Full overlaps from 0.45 to 0.72

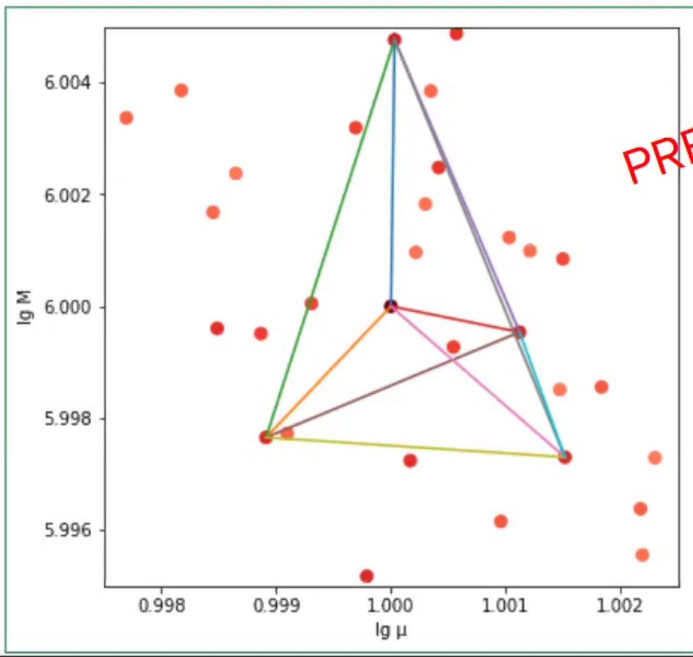

PRELIMINARY

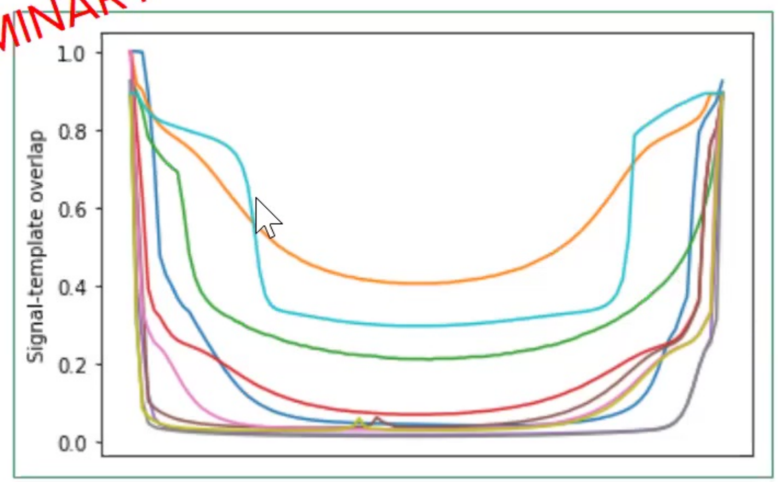

Degeneracies: Evidence

- Threshold-SNR (20) injection, 6 intrinsic parameters, posterior bounds × 100
  - 675 additional secondaries: Full overlaps from 0.23 to 0.76

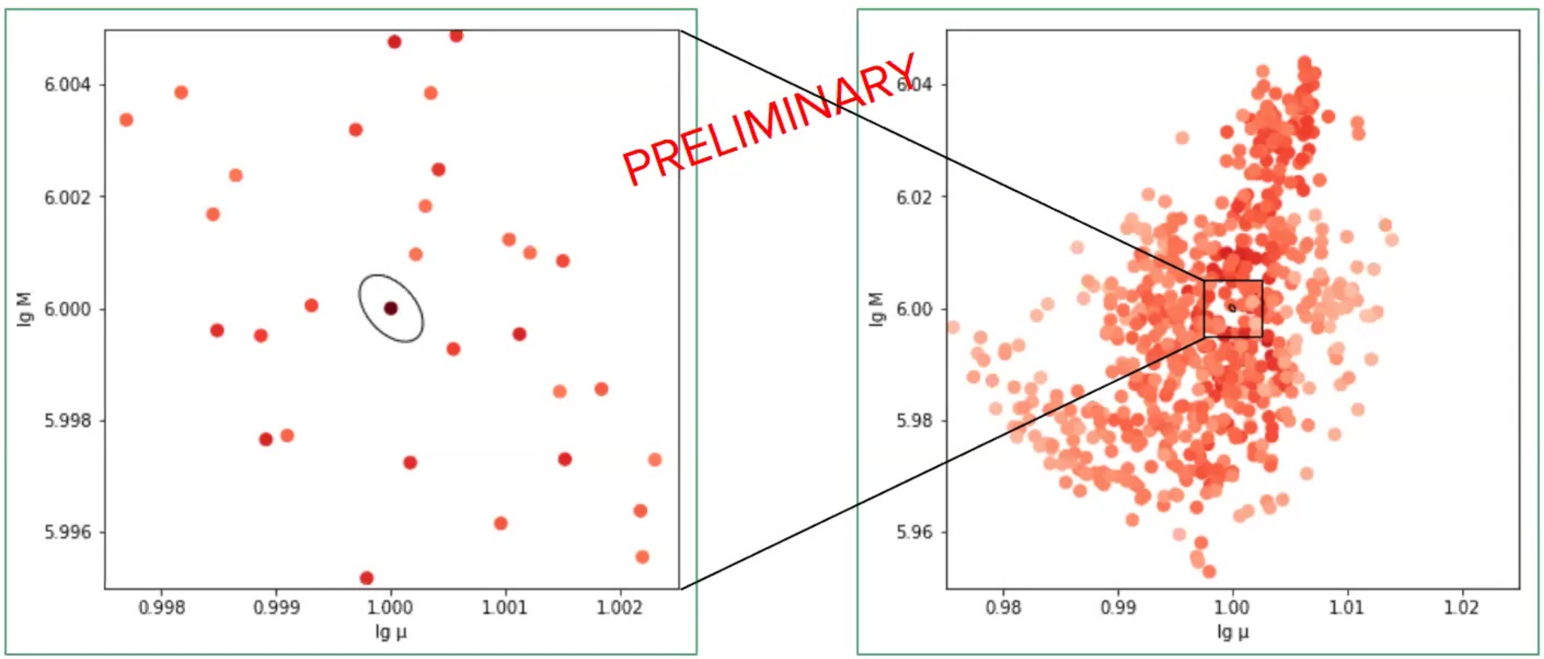

Degeneracies: Implications

- Sum of 2 secondaries from different signals > either primary?
  - Unlikely to be an issue: Primary overlap will be high also
  - In 10^4 pairs of signals drawn from astrophysical priors, none have overlap > 1%
  - More detailed analysis TBD
- Secondary + noise > primary?
  - Unlikely to be an issue: At threshold SNR, probability is < 1% if no secondary overlap > 0.78
- Sampling difficulties
  - Secondaries are suppressed, but primary is highly localized & hard to find
  - Degeneracies will not be addressed by “mode-hopping” MCMC proposals [Cornish, 2011]
  - Parallel tempering & nested sampling may help in principle, but will need high resolution
  - Gradient-based sampling (e.g. HMC) will not help

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vt

Model 1: Implementation & benchmarking

- Speed
  - Analysis-length waveforms: 1 year at 0.1 Hz
  - Full harmonic content: All relevant modes
  - Maximal wall time (for GPU): < 200 ms

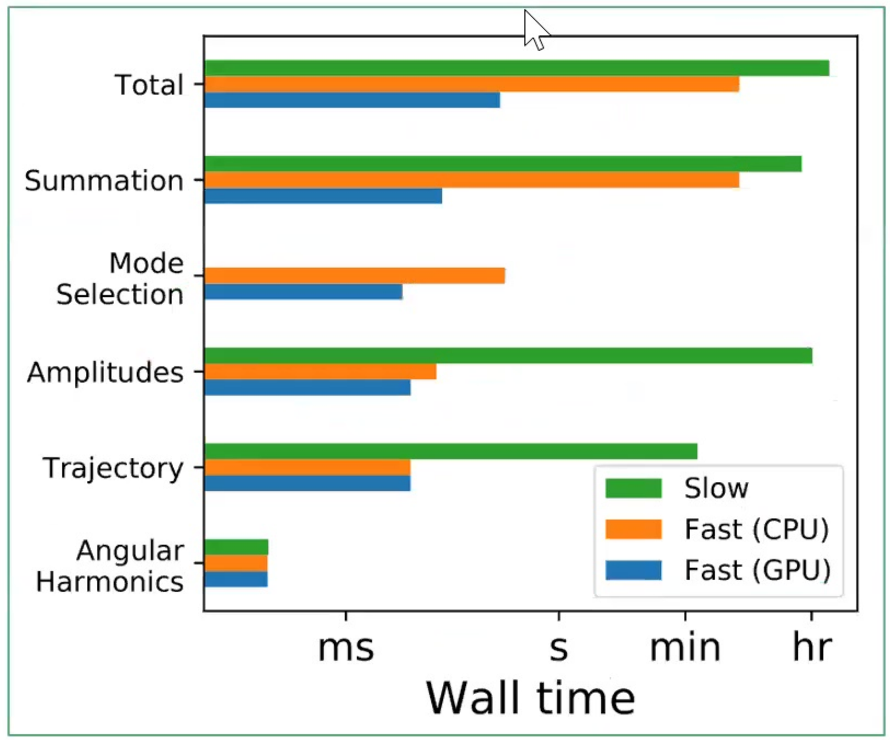

Next-generation models: Framework

- Standard modular description
  - Angular & frequency-based decomposition
  - Osculating geodesics
- Generic Kerr orbits
  - Need schemes to evolve through resonances
  - Need secondary spin, mass/spin evolution, etc.
- Angular dependence
  - Spheroidal harmonics with spin weight -2
- Inspiral trajectory (+ mode phasing)
  - Post-1-adiabatic order
- Mode amplitudes
  - Adiabatic order

$$h(t) = \frac{1}{r} \sum_{lmkn} A_{lmkn}(t) e^{-i\Phi_{mkn}(t)} V_{lmkn}(\theta, \phi)$$

$$G(t) \equiv (p(t), e(t), \iota(t))$$

$$V_{lmkn}(\theta, \phi) = {}_{-2}S_{lmkn}(\theta) e^{im\phi}$$

$$\Phi_{mkn}(t) = \text{init.} + \int_{t_0}^{t} dt' \, \omega_{mkn}(G(t')) + \text{osc.}$$

$$A_{lmkn}(t) = -\frac{2Z^{\infty}_{lmkn}(G(t))}{\omega^2_{mkn}(G(t))}$$

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

EMRIs: An overview

- The astrophysics of EMRIs
  - Event-rate estimates: 1-10^4 (per LISA) [Babak et al., 2017]
  - Brown-dwarf “problem” [Gourgoulhon et al., 2019; Amaro-Seoane, 2019; Amaro-Seoane, private comm.]
  - Possibly many other environmental effects

The EMRI inverse problem: Difficulties

- Everything is hard for EMRIs (+ astrophysics, but let's not go into that here)
- Difficulty 1: Accuracy
  - EMRIs are strong-field, high-SNR sources that require accurate modeling to find & characterize
  - Phasing accurate to post-1-adiabatic order will be enough, but we are not there yet
- Difficulty 2: Efficiency
  - EMRI signals are long-lived with rich harmonic content; they are costly to model & analyze
  - Stochastic algorithms in data analysis require bulk generation of waveforms (at least billions)
- Difficulty 3: Extensiveness
  - Even the "leading-order" space of EMRI orbits is gargantuan in terms of information volume
  - Waveforms are only half the battle: detector response encodes important extrinsic effects

Next-generation models: Definitions

- Waveform model:
  - Accurate (fully relativistic)
  - Not necessarily efficient & not necessarily extensive
  - Only examples are adiabatic (Teukolsky-based) models
  - Previous adiabatic models are very slow & limited in scope
- Template model:
  - Not necessarily accurate
  - Efficient & extensive
  - Only examples are semi-relativistic (kludge) models
  - Semi-relativistic models have an inherent cap on accuracy
- Aim is for next generation of models to be accurate, efficient & extensive
- The 2 new models are accurate & efficient, but not yet extensive

Next-generation models: Inspiral trajectory

- Trajectory models
  - PN flux-based (partial coverage, only for search)
  - Teukolsky flux-based (only for search)
  - NIT [van de Meent & Warburton, 2018]
  - Two-timescale framework

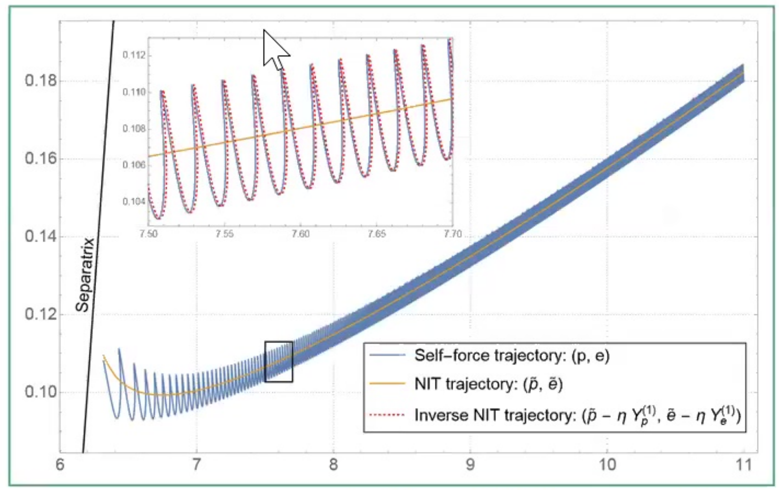

van de Meent & Warburton (2018)

201022 PI - Google Slides - Mozilla Firefox

https://docs.google.com/presentation/d/1q1XISCncWV8Xqsb4NNJ1gC8Ut3CrlRLeSk2Sp9c8HJ8/present?token=AC4w5Vł

Next-generation models: Future work

- Improve source-end extensiveness (Model 1)
  - Eccentric equatorial Kerr: Add primary spin
  - Partial post-adiabatic trajectory: Add 1st-order SF
  - Incorporate resonance schemes
- Improve efficiency (Model 2): Probably GPU implementation
- Improve detector-end efficiency & extensiveness
  - Integrate with GPU-accelerated full LISA response models [AC with Katz, Bayle, Vallisneri]
- Other representations
  - Frequency domain: Higher-order SPA [AC with Hughes, Warburton, Katz]
  - Time-frequency domain: STFT