Listening to black holes with supercomputers

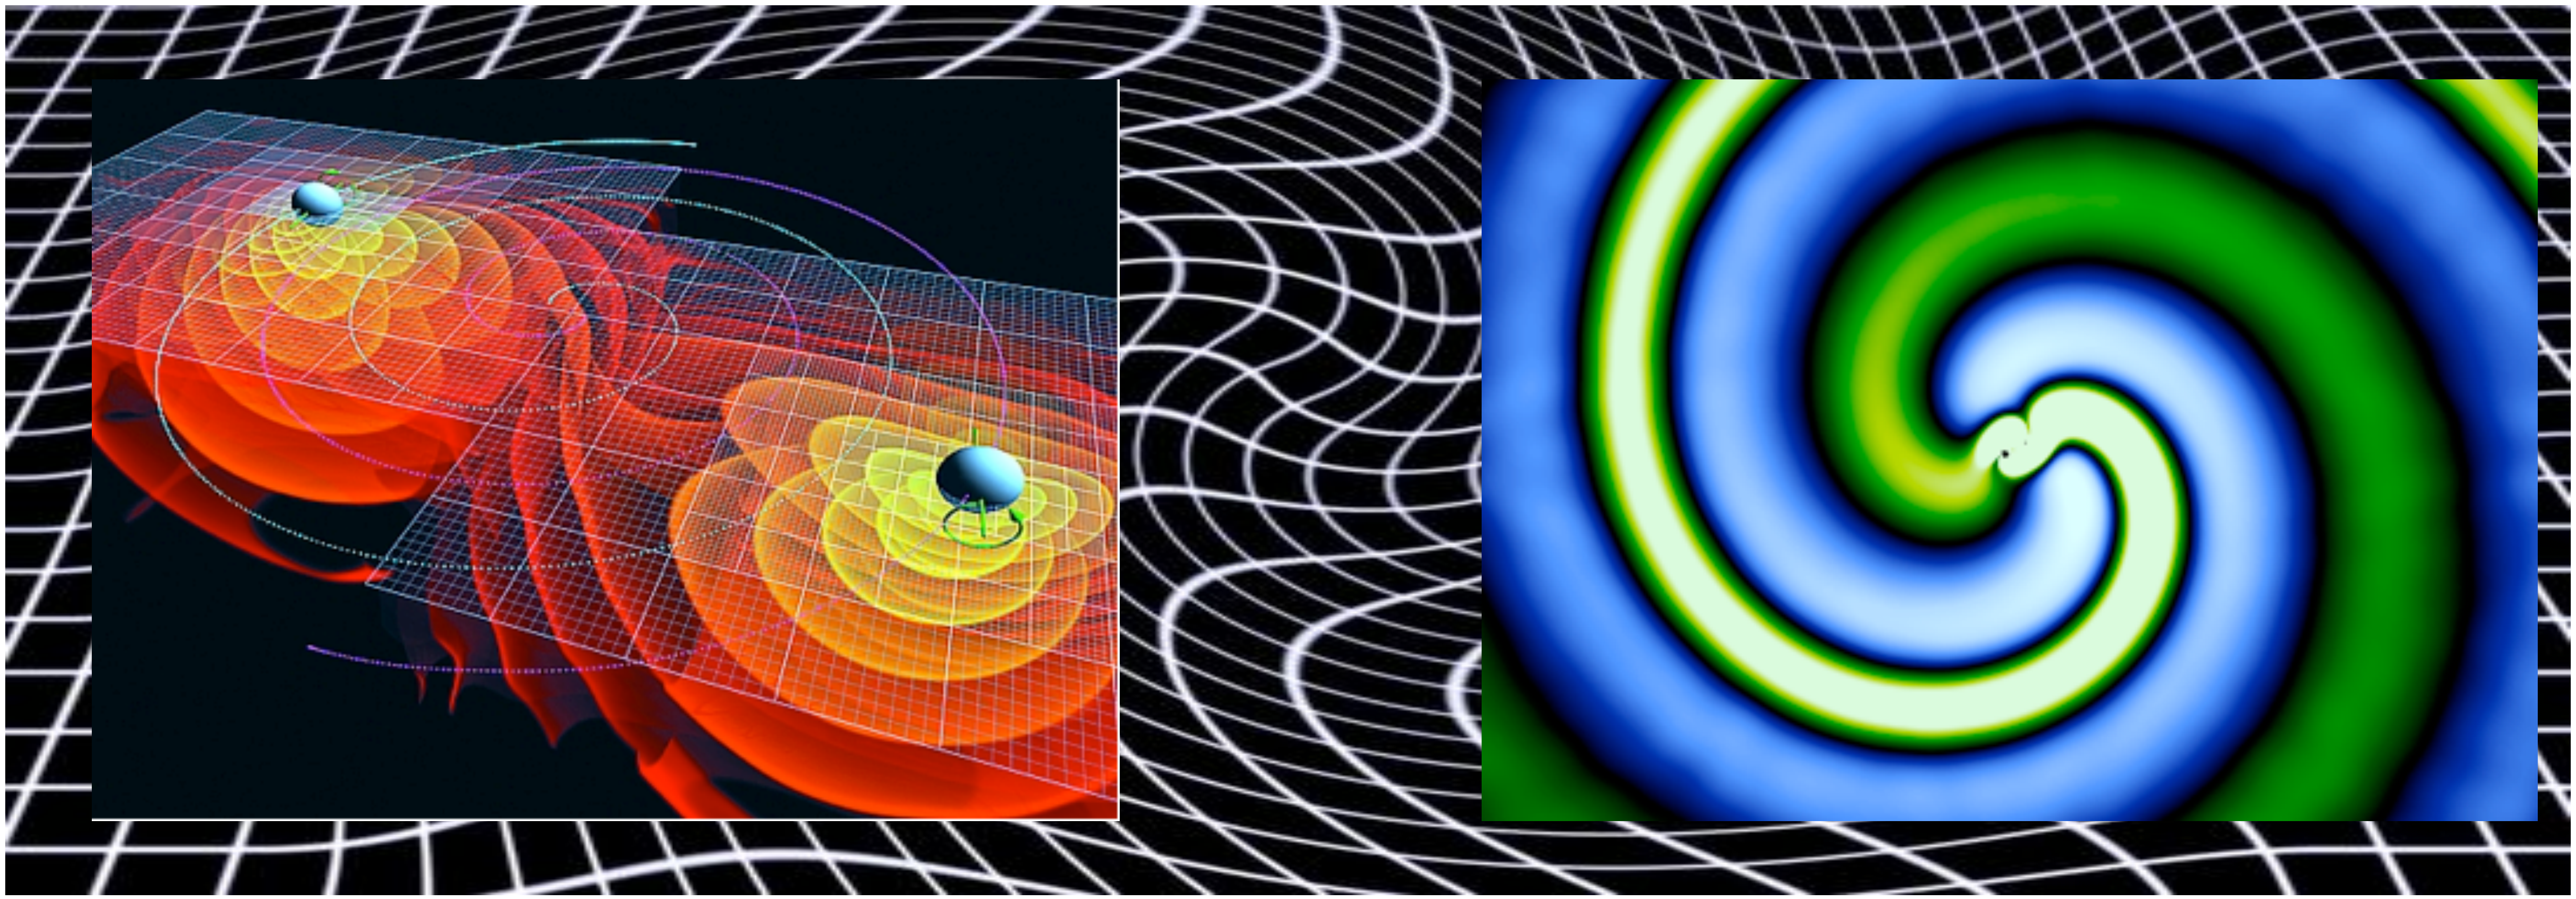

General Relativity: when heavy stars run out of fuel, nothing can halt their gravitational collapse to a black hole.

Our aim: understand observational signatures of merging BH binaries, required to identify such events in gravitational wave observations.

S. Husa, Universitat de les Illes Balears
PRACEdays15, 26.5. 2015

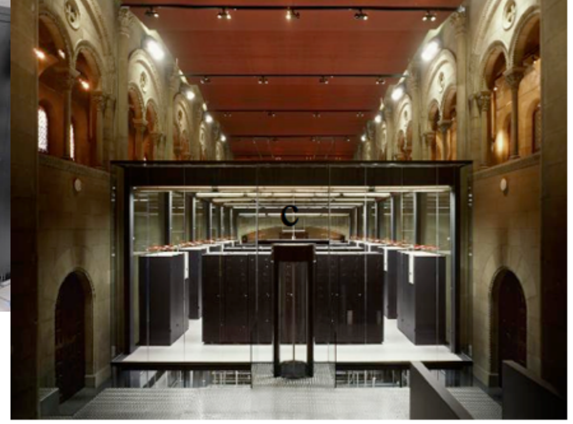

Thanks to collaborators and funding!

Contributors: Michael Pürrer, Mark Hannam, Sebastian Khan, Frank Ohme, Alejandro Bohé, Francisco Jimenez, Juan Calderon, Bernd Brügmann, Nathan Johnson-McDaniel, Denis Pollney, Christian Reisswig, Milton Ruiz, Patricia Schmidt, Marcus Thierfelder, Vijay Varma, Parameswaran Ajith.

The Dark Universe

- Black holes have taken center stage in astrophysics & fundamental physics.
  - Extraordinarily clean systems, described by their mass M and spin $\chi = \frac{|\vec{J}|}{M^2}$
  - Allow precision astrophysics and fundamental physics:
    - life cycle of stars
    - supermassive black holes in galaxy cores
    - testing general relativity - find new physics?
  - direct observation?

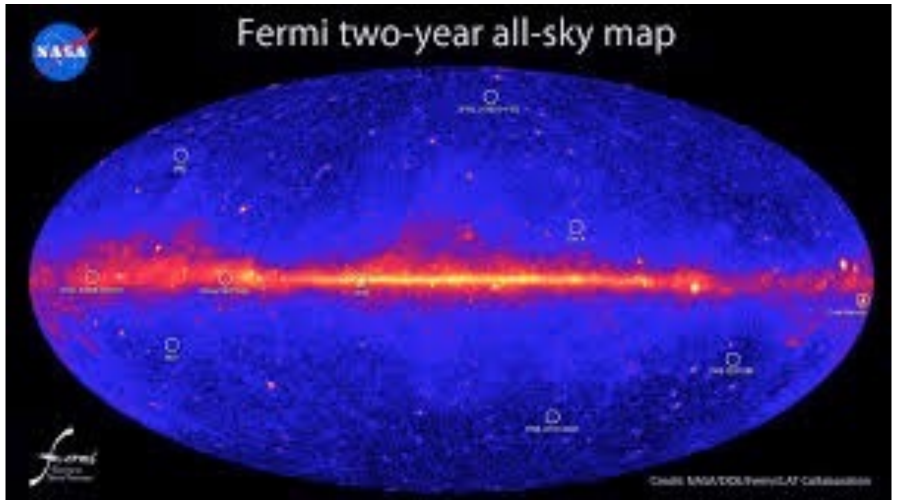

- Electromagnetic waves taught us what we know about the universe. Superposition of waves from many particles
  - —> image of the source.
- Electromagnetic spectrum is blind to some of the most violent and exotic objects in the universe.

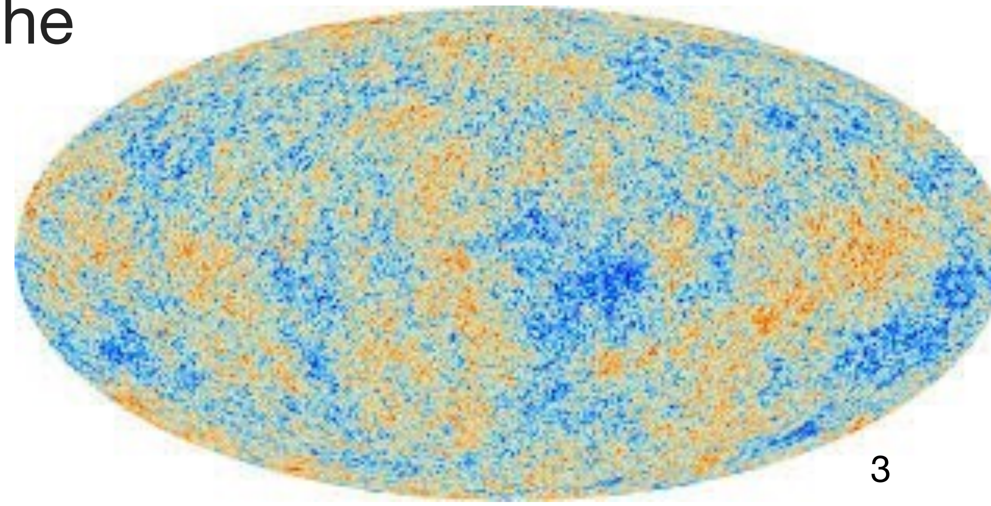

Gravitational Waves

- Spacetime in general relativity is a deformable entity, ripples in spacetime travel at the speed of light and carry with the information on their source.
- Close binary systems of BH/NS are most efficient sources of gravitational waves.

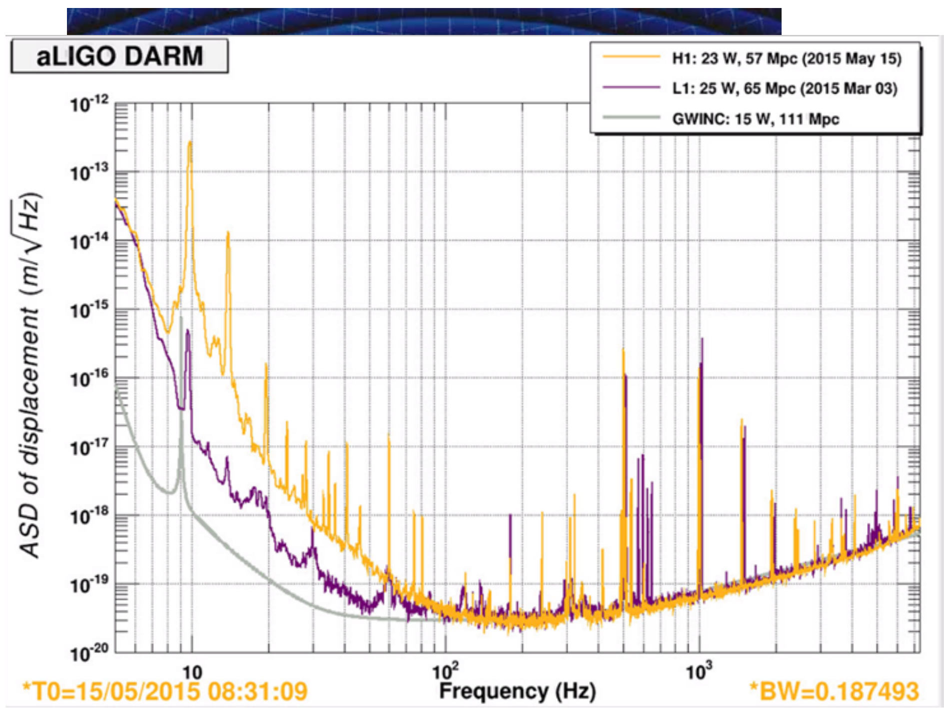

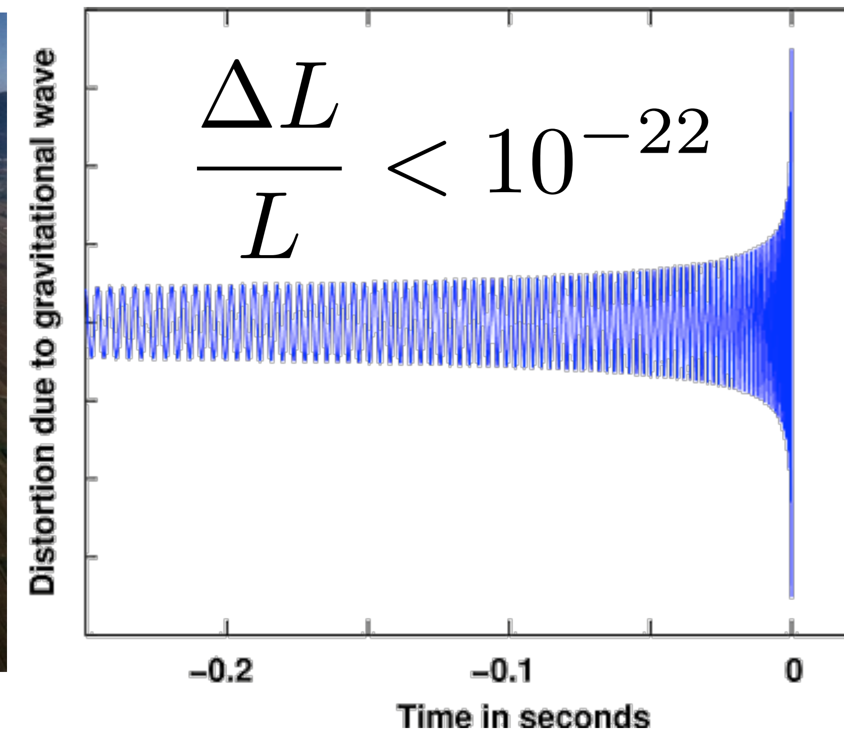

- GW signal carries information about the bulk motion of objects: analogous to hearing sound.
  - “soundtrack of the universe”
- GW signal encodes masses, spins, eccentricity of binary & possibly new physics - needs to be decoded.

Gravitational wave detectors and data analysis

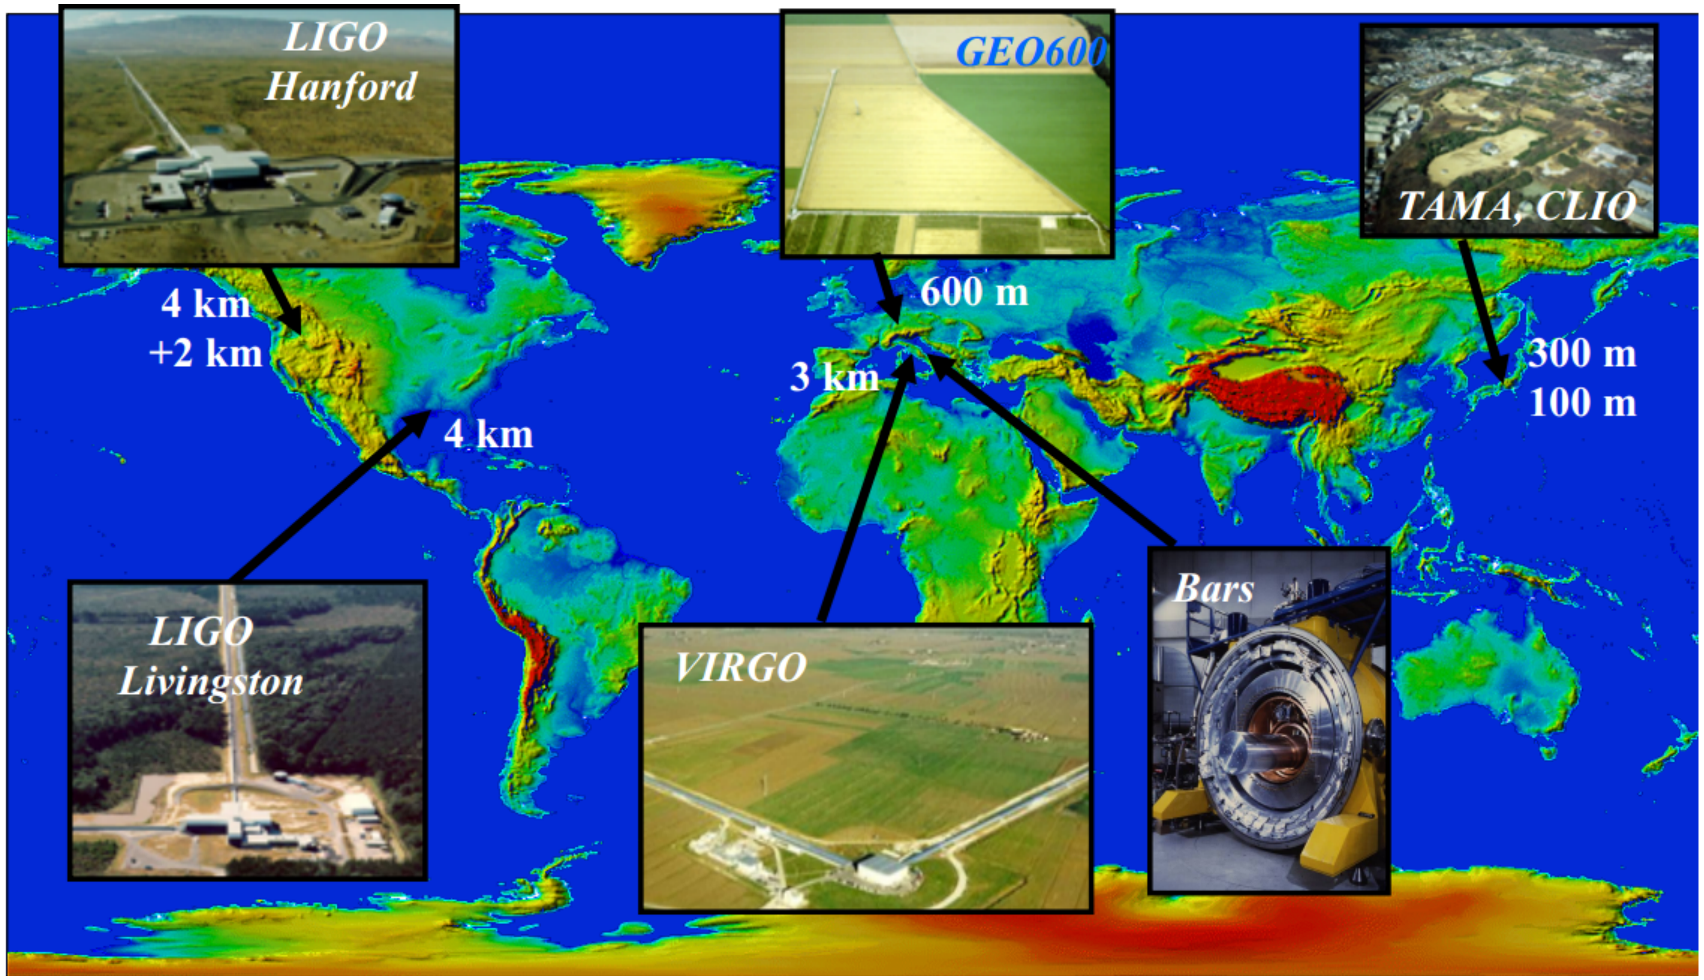

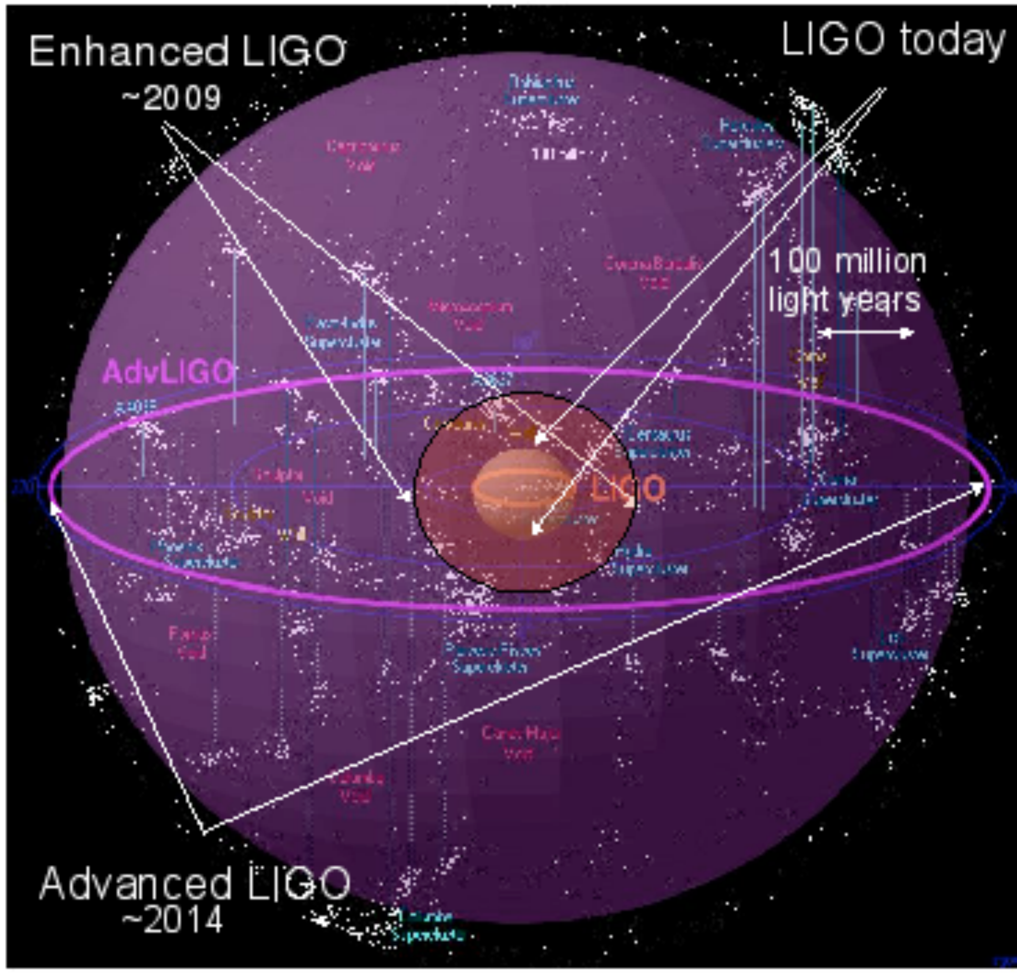

- Since first LIGO science run in 2002 upper limits have been set, but no direct detection, LIGO-Virgo Scientific collaboration has grown to ~ 1000 scientists.
- Computational challenge 1: Searches for BH merger events are based on “matched filtering” with template waveforms.
- Computational challenge 2: “template banks” need to be computed in general relativity - Inaccurate templates: lose events & incorrectly identify them (masses & spins of a binary, identification as BH or NS).

Gravitational wave detectors and data analysis

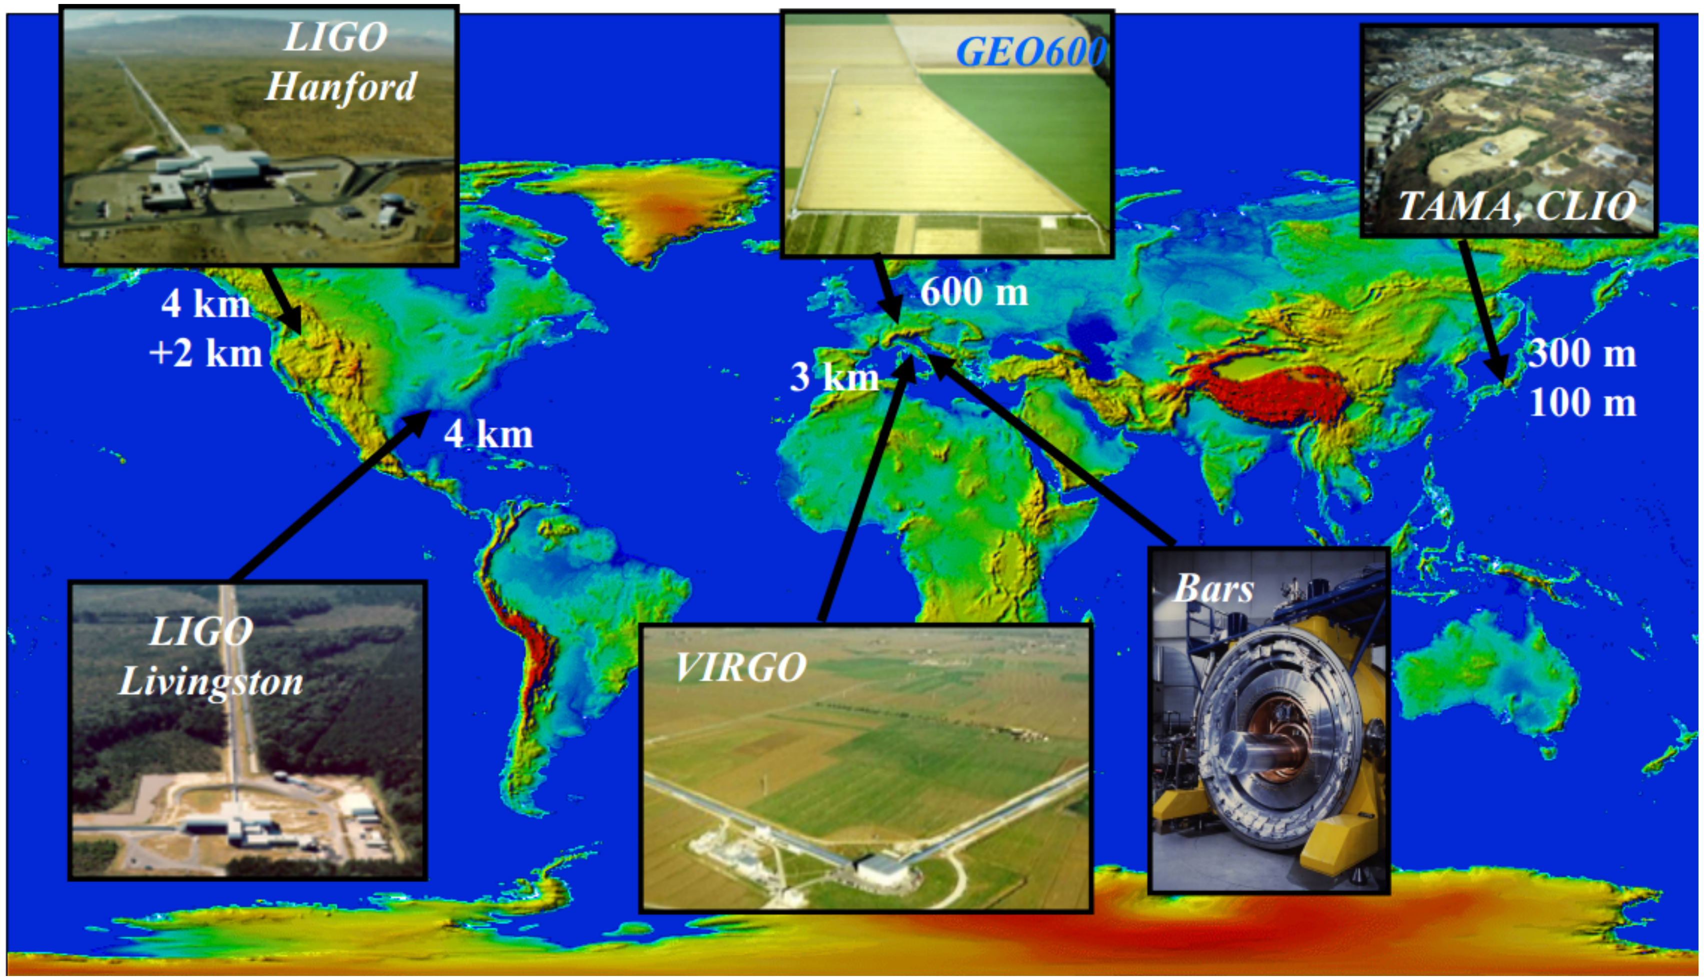

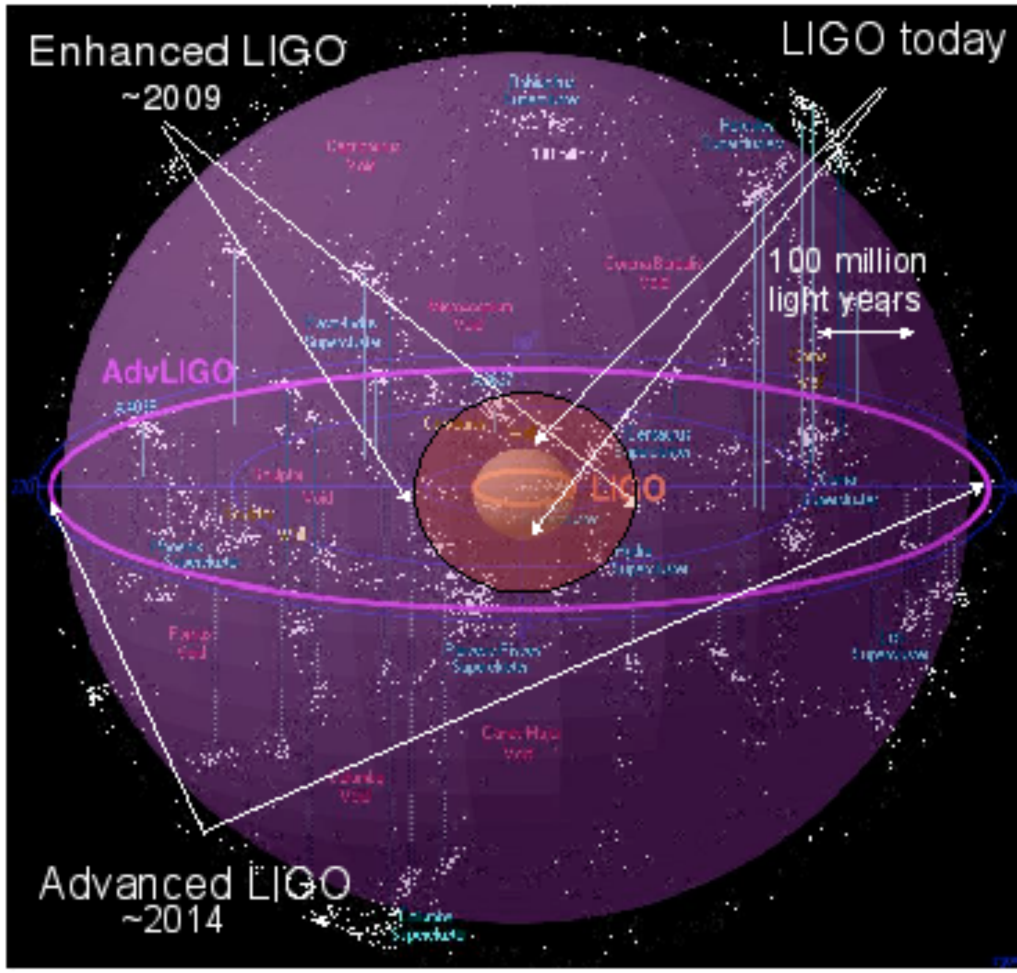

- Since first LIGO science run in 2002 upper limits have been set, but no direct detection, LIGO-Virgo Scientific collaboration has grown to ~ 1000 scientists.
- Computational challenge 1: Searches for BH merger events are based on “matched filtering” with template waveforms.
- Computational challenge 2: “template banks” need to be computed in general relativity - Inaccurate templates: lose events & incorrectly identify them (masses & spins of a binary, identification as BH or NS).

Gravitational wave detectors and data analysis

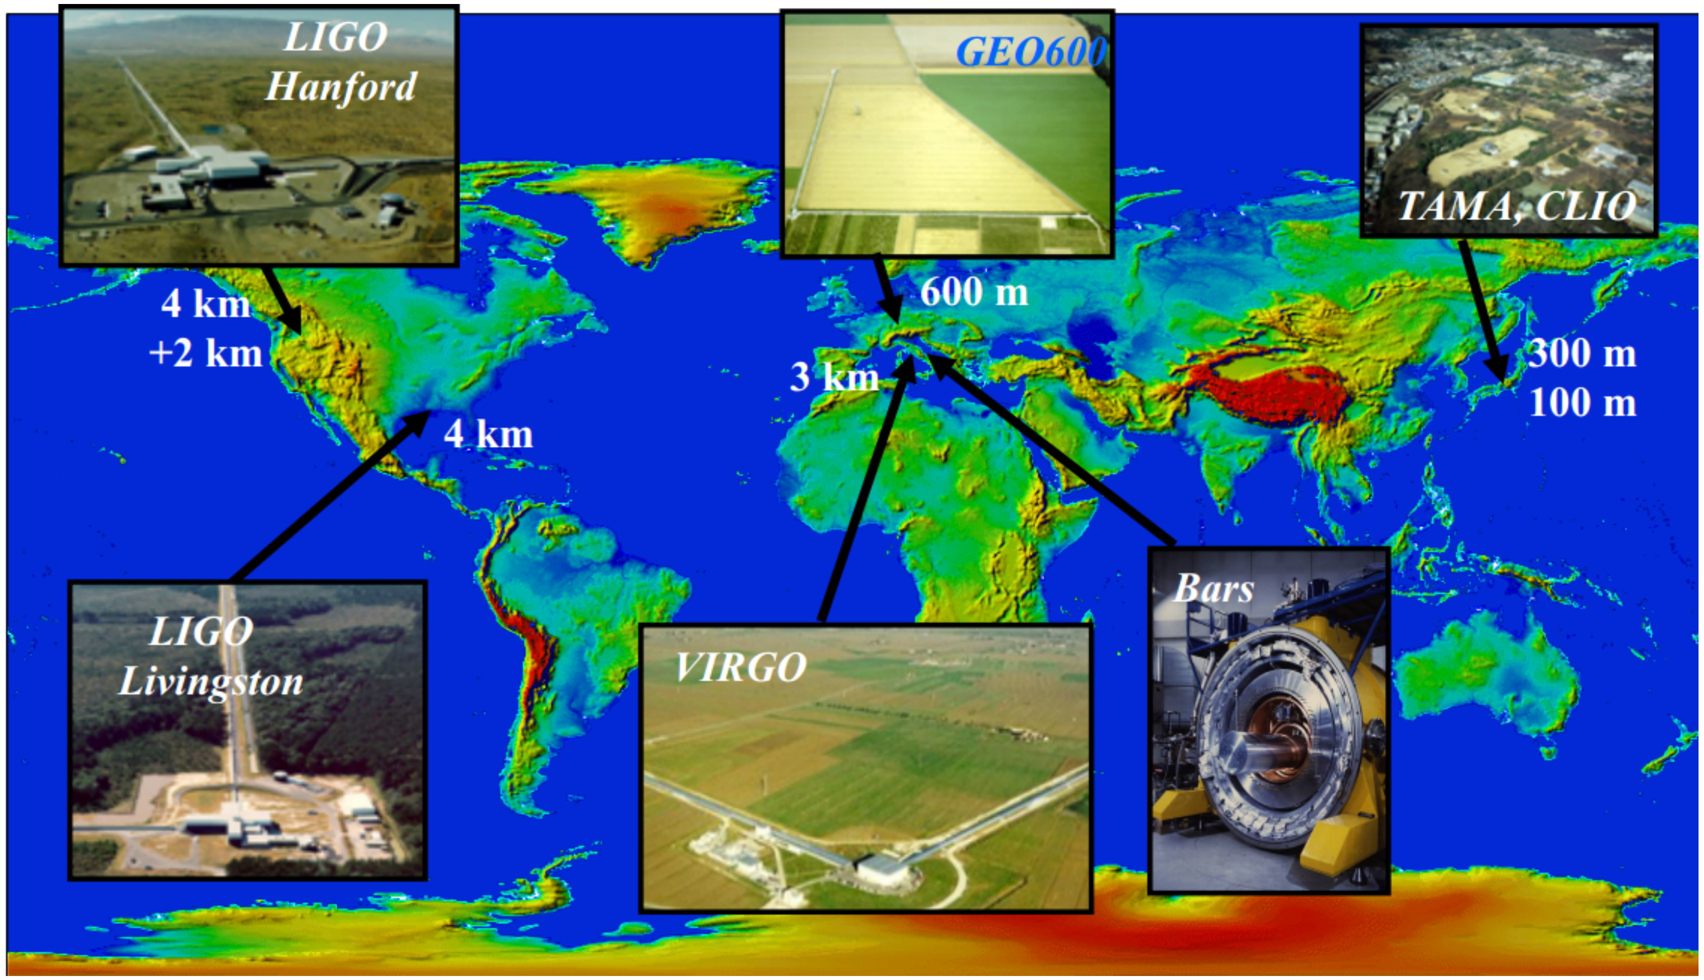

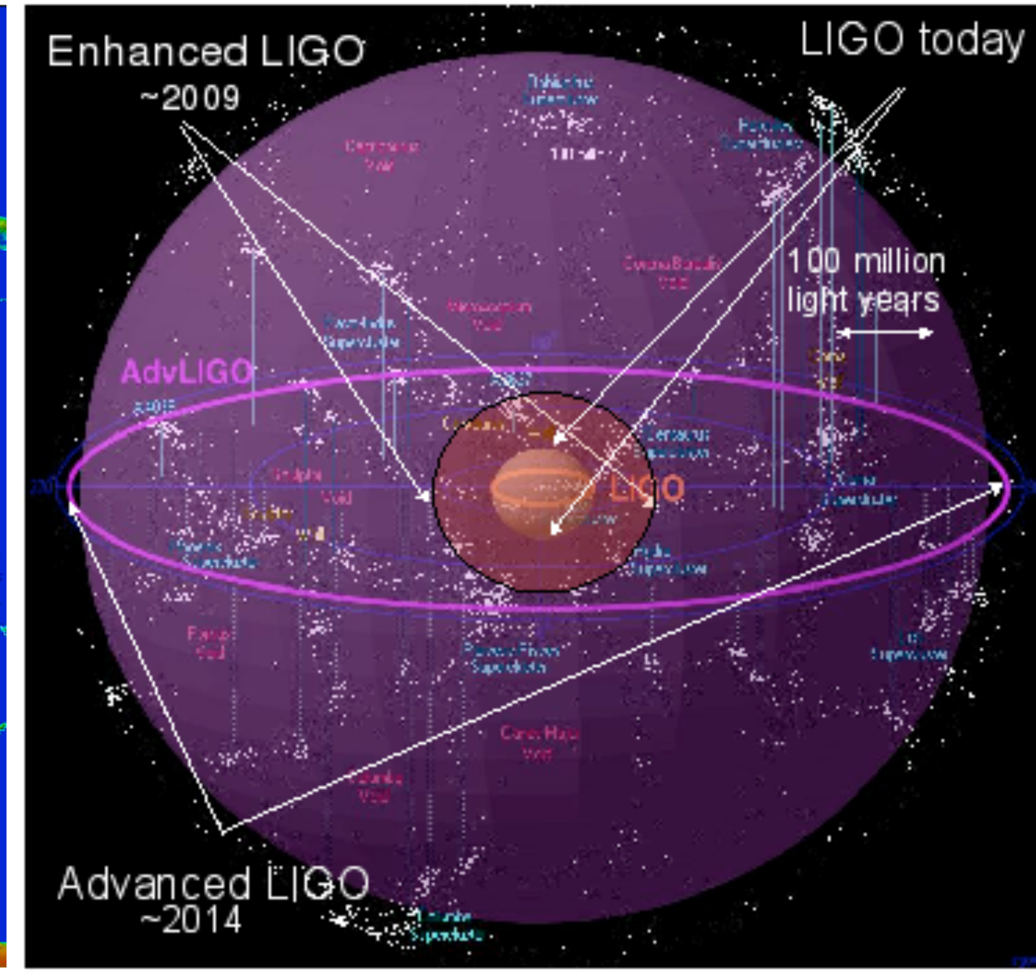

- Since first LIGO science run in 2002 upper limits have been set, but no direct detection, LIGO-Virgo Scientific collaboration has grown to ~ 1000 scientists.
- Computational challenge 1: Searches for BH merger events are based on “matched filtering” with template waveforms.
- Computational challenge 2: “template banks” need to be computed in general relativity - Inaccurate templates: lose events & incorrectly identify them (masses & spins of a binary, identification as BH or NS). **7D parameter space**

Anatomy of BH mergers

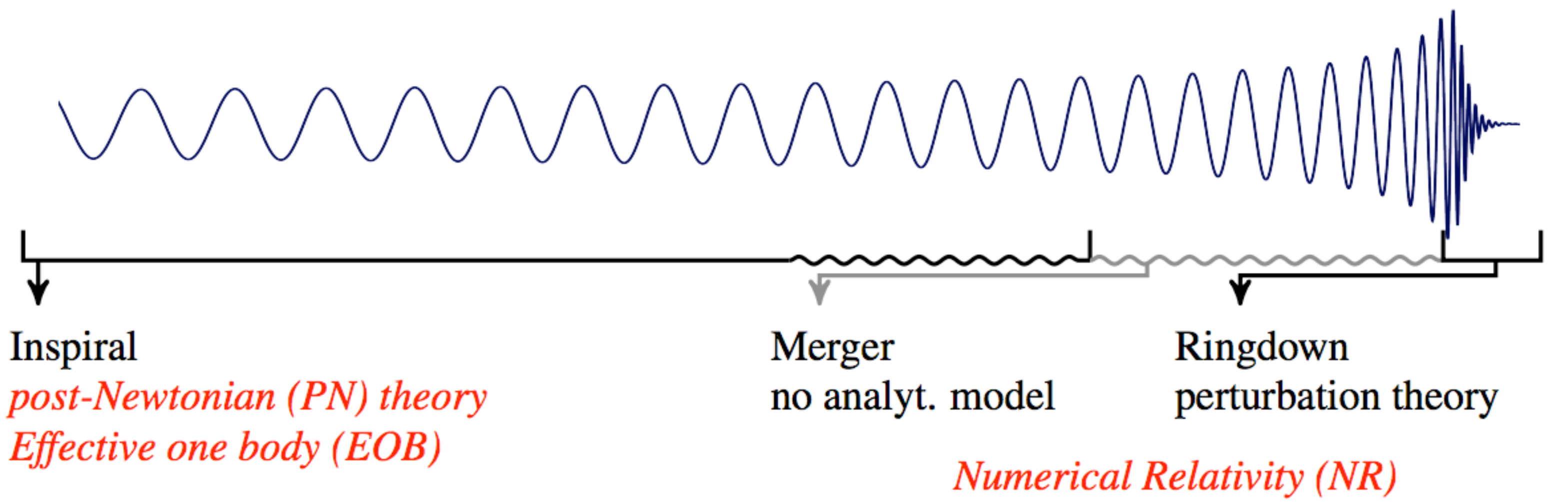

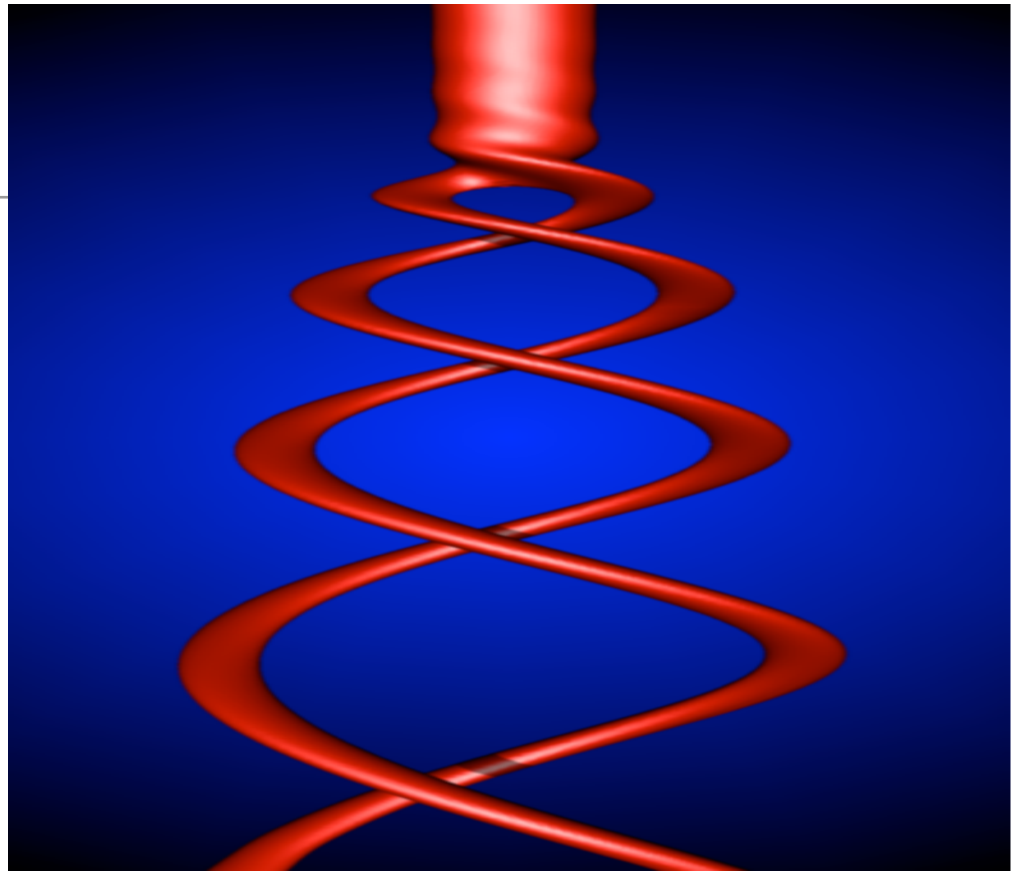

- Inspiral: energy loss to GWs leads to adiabatic inspiral, well described by post-Newtonian perturbation theory.

- Late inspiral & merger: post-Newtonian expansion breaks
  - solve full Einstein equations numerically as PDEs, "match" to post-Newtonian inspiral.
  - Most of the energy released (< 12 % of the mass).
- Ringdown: superposition of damped harmonics, frequencies known from perturbation theory.

Anatomy of BH mergers

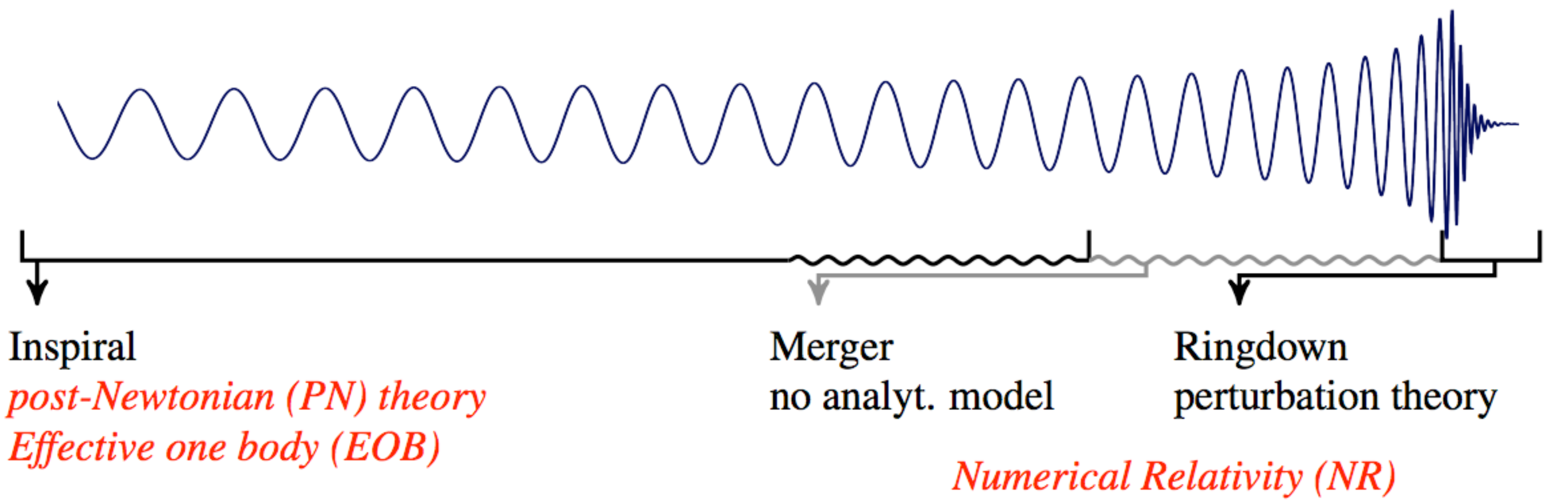

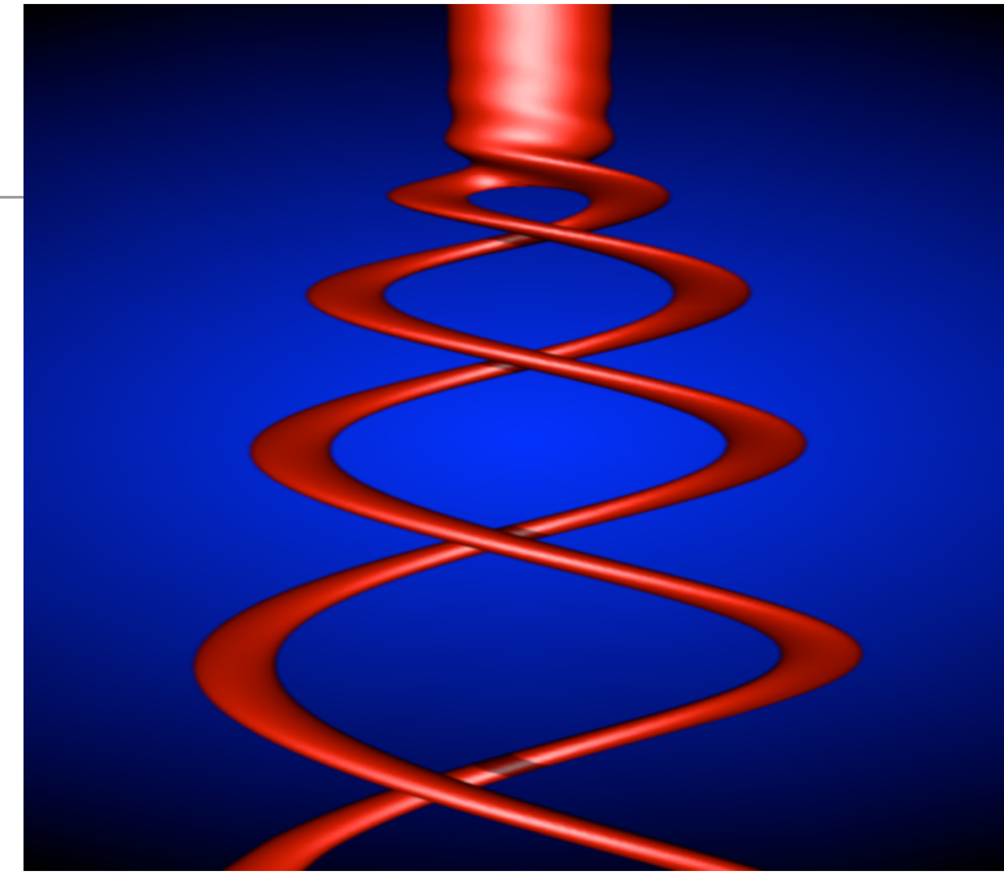

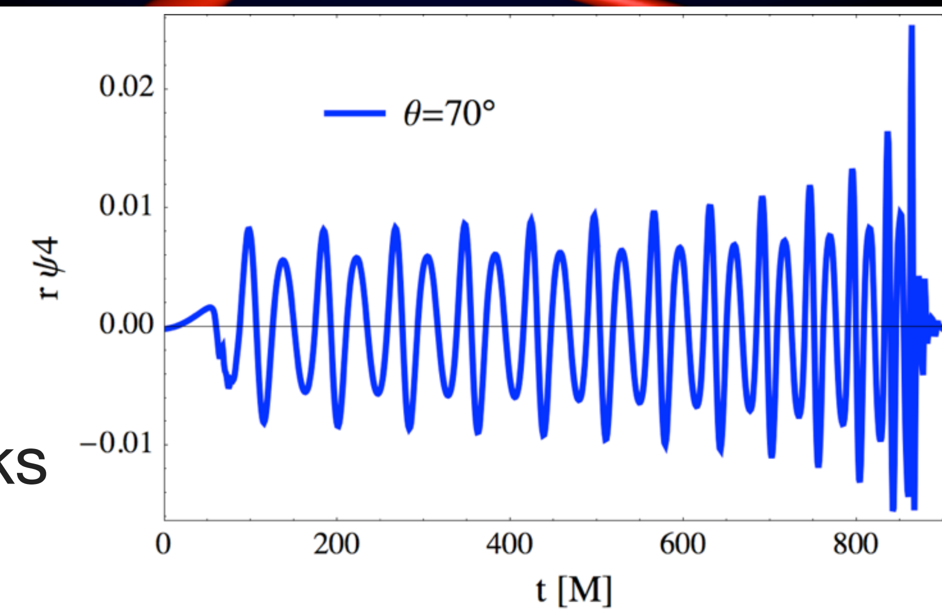

- Inspiral: energy loss to GWs leads to adiabatic inspiral, well described by post-Newtonian perturbation theory.
- Late inspiral & merger: post-Newtonian expansion breaks
  - solve full Einstein equations numerically as PDEs, “match” to post-Newtonian inspiral.
  - Most of the energy released (< 12 % of the mass).
- Ringdown: superposition of damped harmonics, frequencies known from perturbation theory.

Anatomy of BH mergers

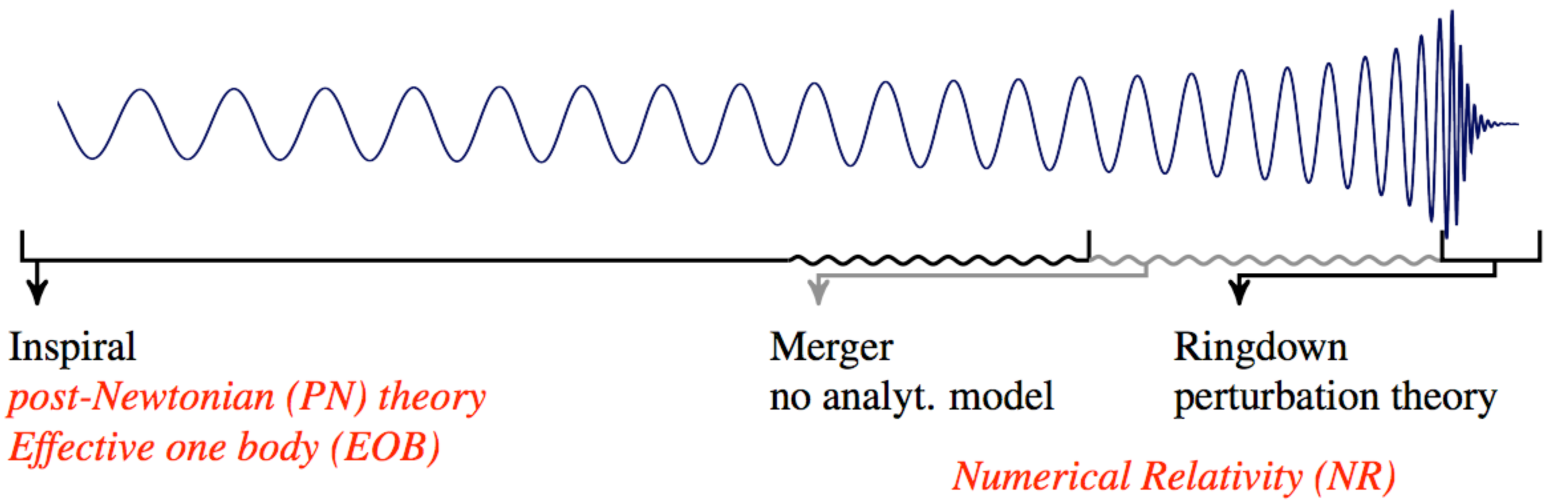

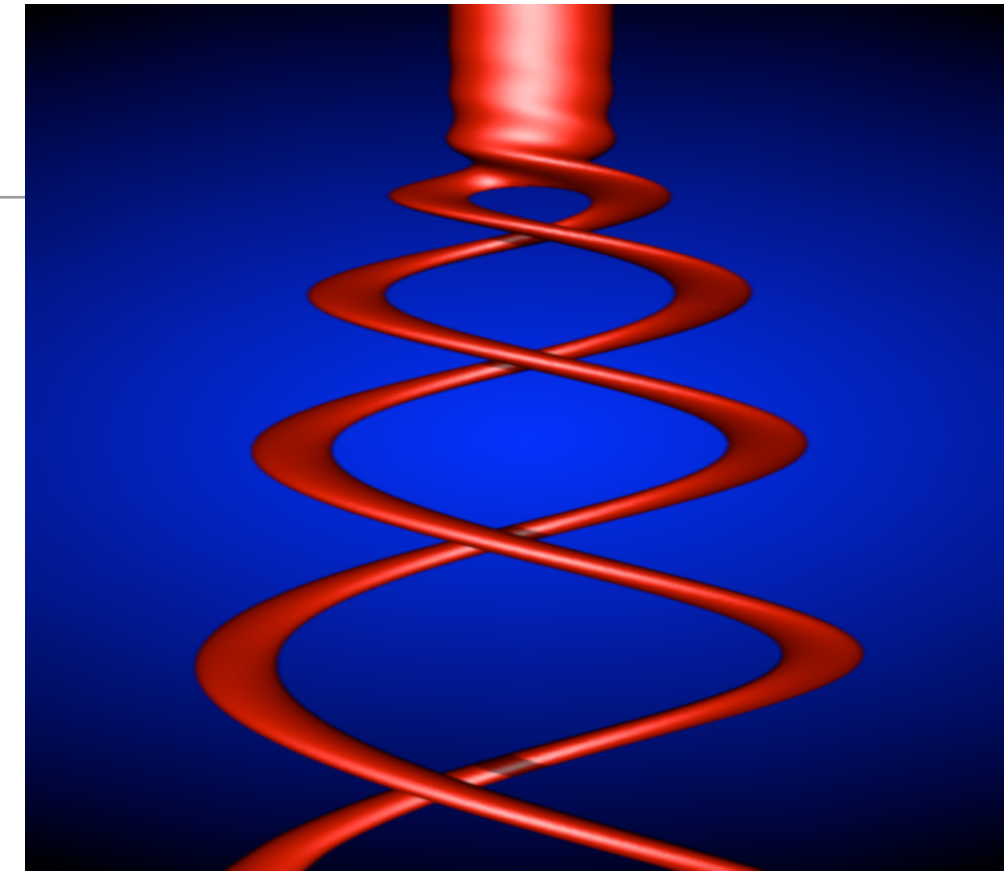

- Inspiral: energy loss to GWs leads to adiabatic inspiral, well described by post-Newtonian perturbation theory.
- Late inspiral & merger: post-Newtonian expansion breaks
  - solve full Einstein equations numerically as PDEs, “match” to post-Newtonian inspiral.
  - Most of the energy released (< 12 % of the mass).
- Ringdown: superposition of damped harmonics, frequencies known from perturbation theory.

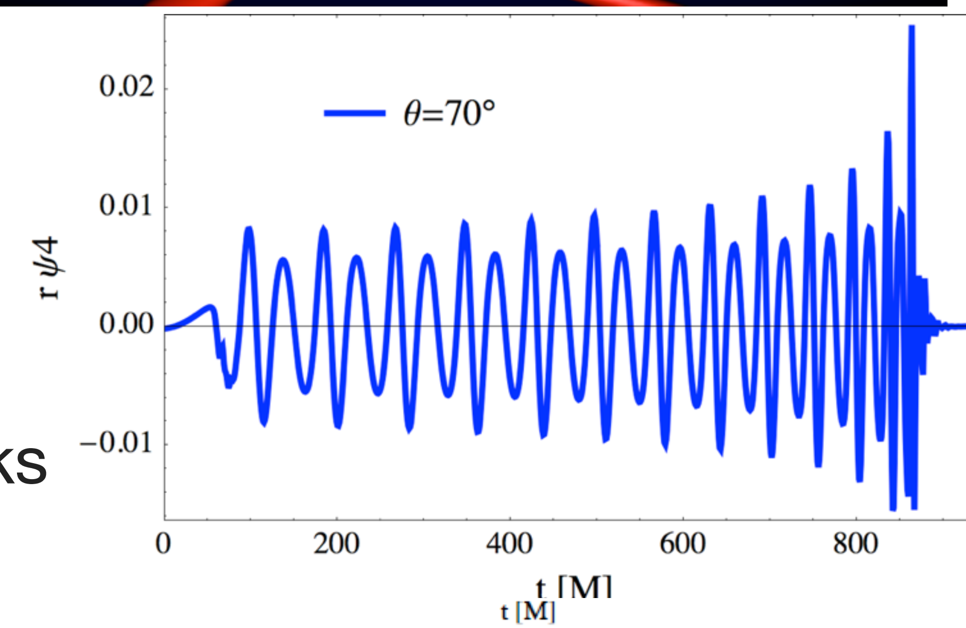

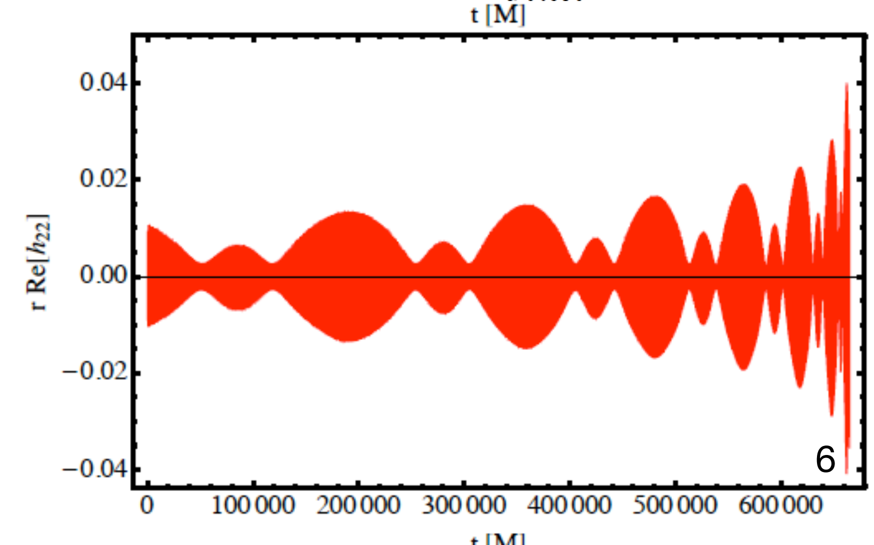

Anatomy of BH mergers

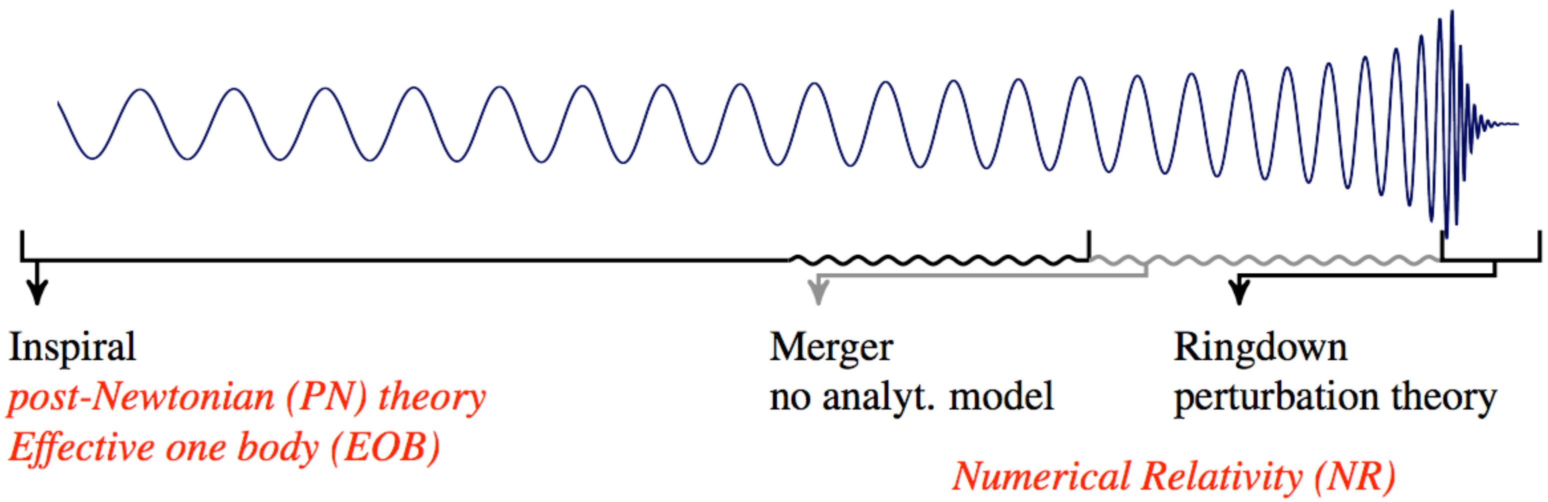

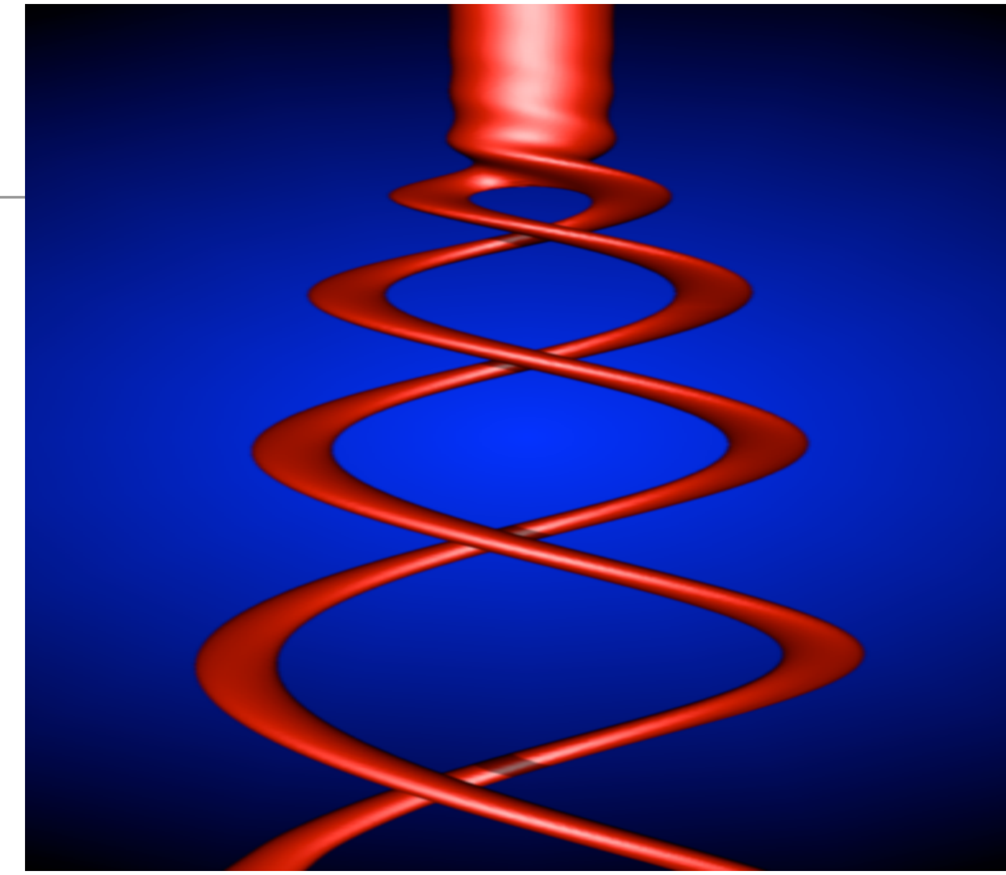

- Inspiral: energy loss to GWs leads to adiabatic inspiral, well described by post-Newtonian perturbation theory.

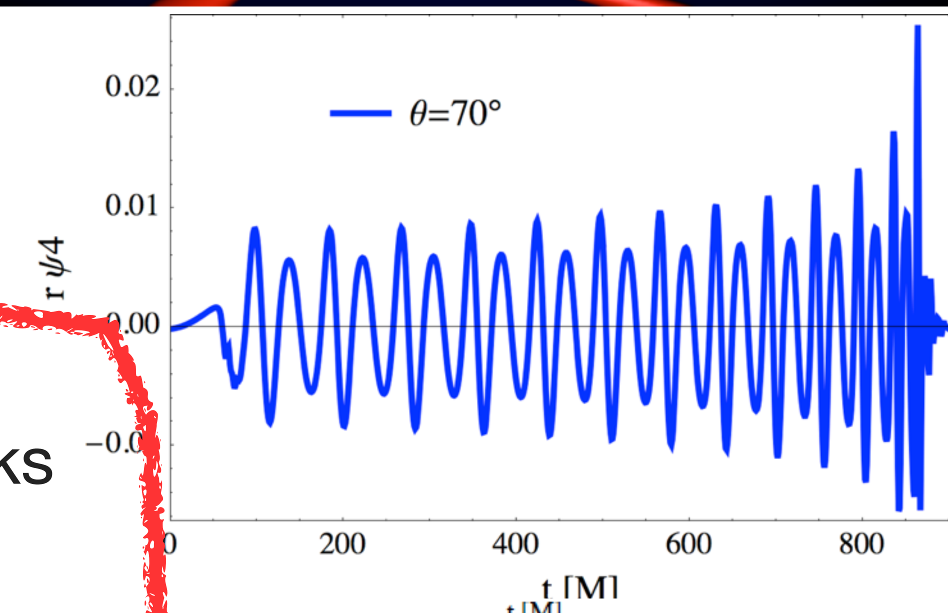

- Late inspiral & merger: post-Newtonian expansion breaks
  - solve full Einstein equations numerically as PDEs, “match” to post-Newtonian inspiral.
  - Most of the energy released (< 12 % of the mass).
- Ringdown: superposition of damped harmonics, frequencies known from perturbation theory.

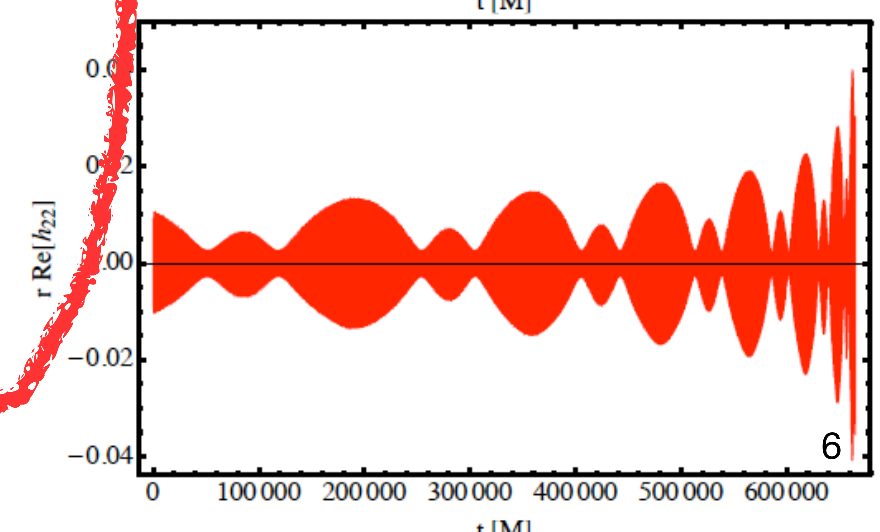

Anatomy of BH mergers

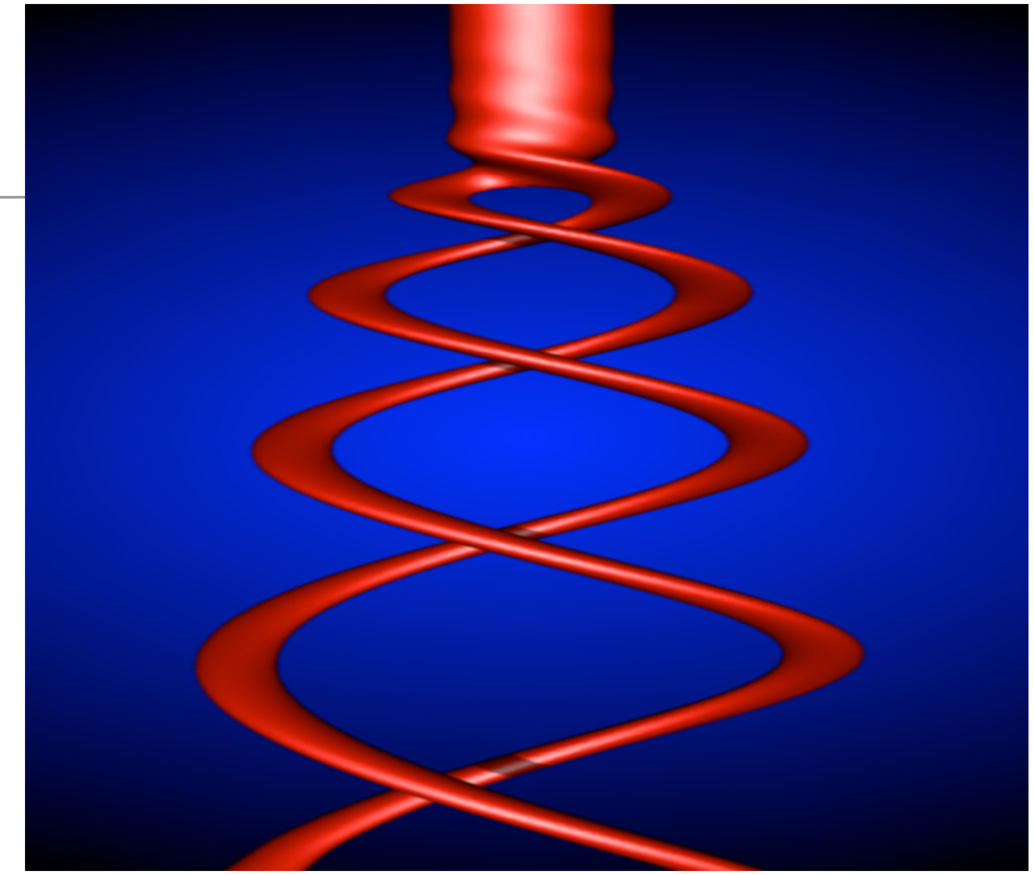

well described by post-Newtonian perturbation theory.

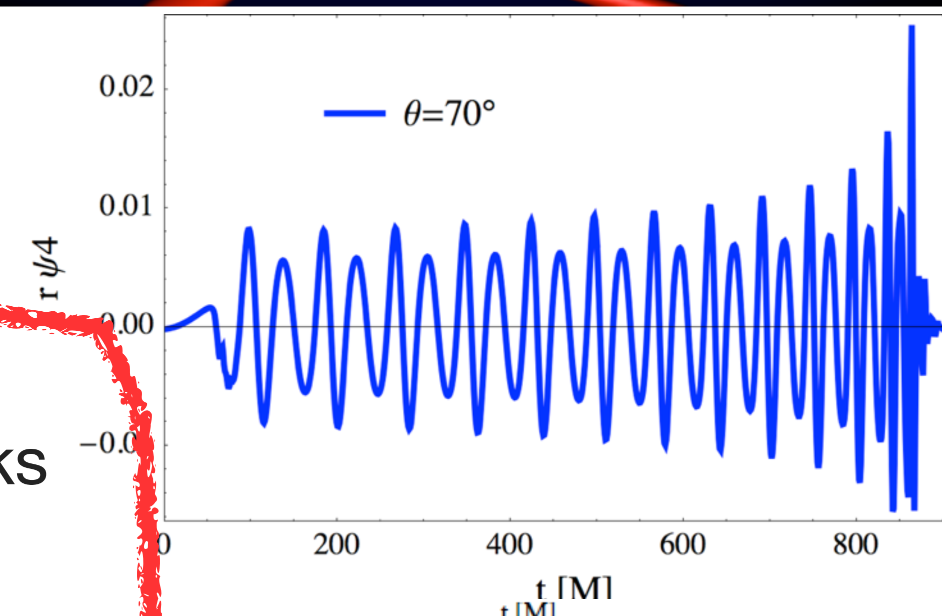

- Late inspiral & merger: post-Newtonian expansion breaks
  - solve full Einstein equations numerically as PDEs, “match” to post-Newtonian inspiral.
  - Most of the energy released (< 12 % of the mass).
- Ringdown: superposition of damped harmonics, frequencies known from perturbation theory.

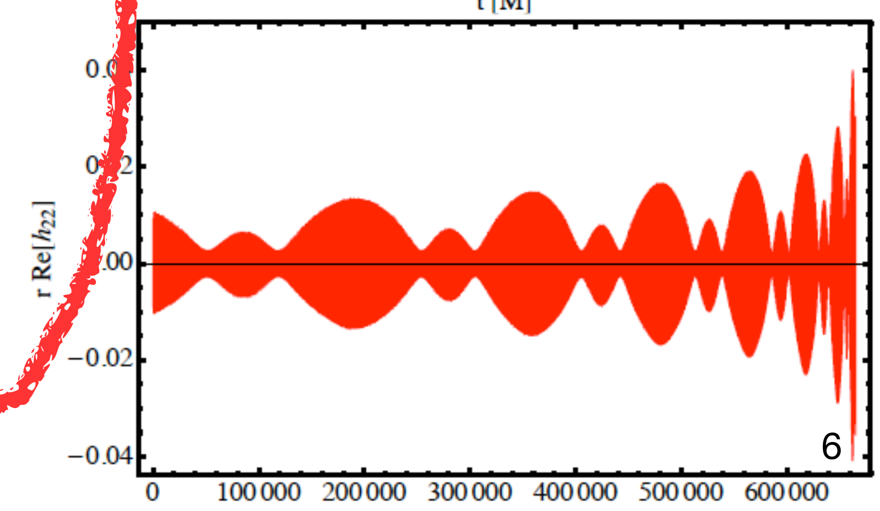

Solving the Einstein equations

- Einstein equations describe geometry of spacetime, not fields on a fixed spacetime -> technical problems took 4 decades to solve for 2-body problem.
- EE can be viewed as 10 coupled nonlinear wave equations (hyperbolic) plus elliptic constraint equations (solved initially, then just monitored).
- We evolve BSSNOK version: 24 evolution equations + monitor 9 constraints.
  - 1000s of terms, hard to optimize for compiler.

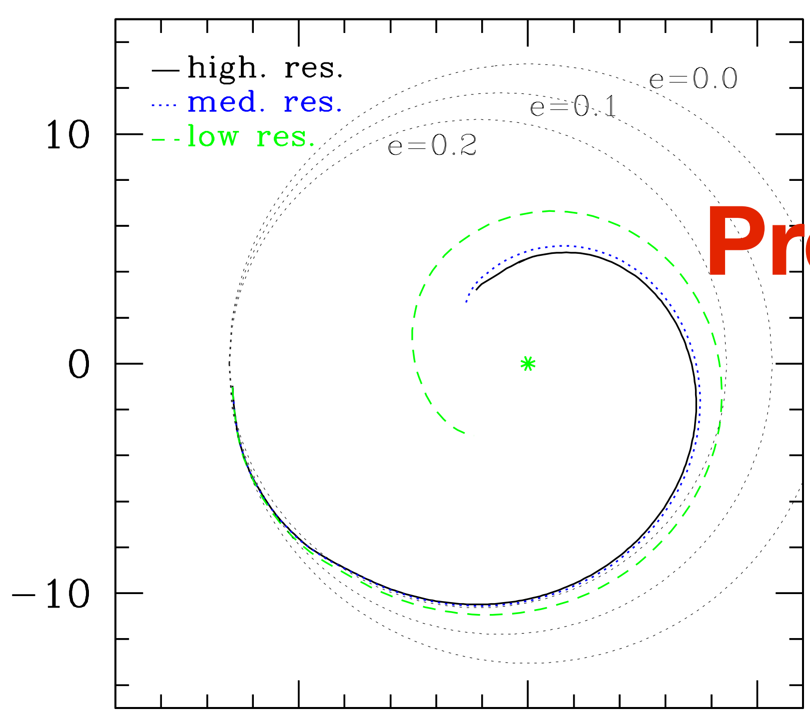

Pretorius 2005 PRL

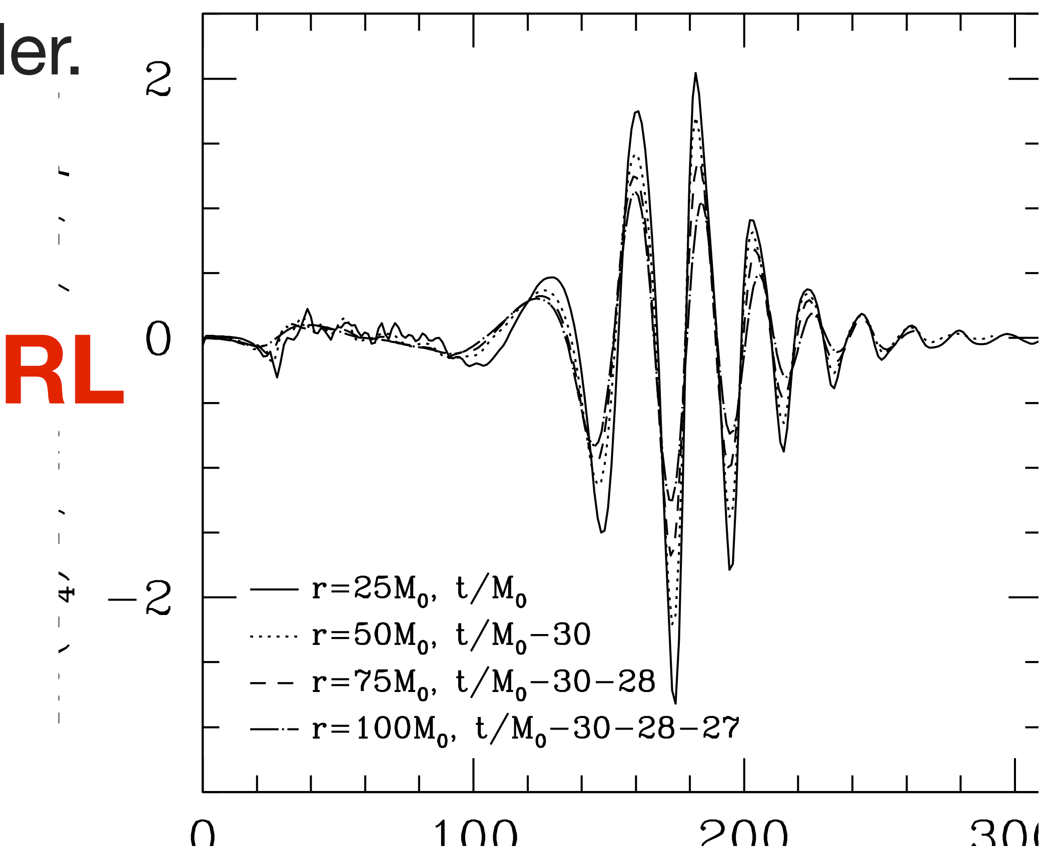

- Einstein equations form singularities hidden inside BHs, not shocks -> use 6-8th oder finite differencing. More efficient, but less robust: spectral.

Scales and mesh refinement

- BH binaries have several length & time scales:
  - individual BHs (most compact objects)
    - resolution around BHs determines accuracy of tracking orbital phase
    - “recipe” to configure box sizes
  - orbital scale: typically start at separations > 15 km M/Msun
- wave scale frequency increases ~ factor of 10
- 1/distance^n background falloff
- ambiguity in boundary conditions:
  - causally isolate boundaries
    - -> 1000s of km M/Msun

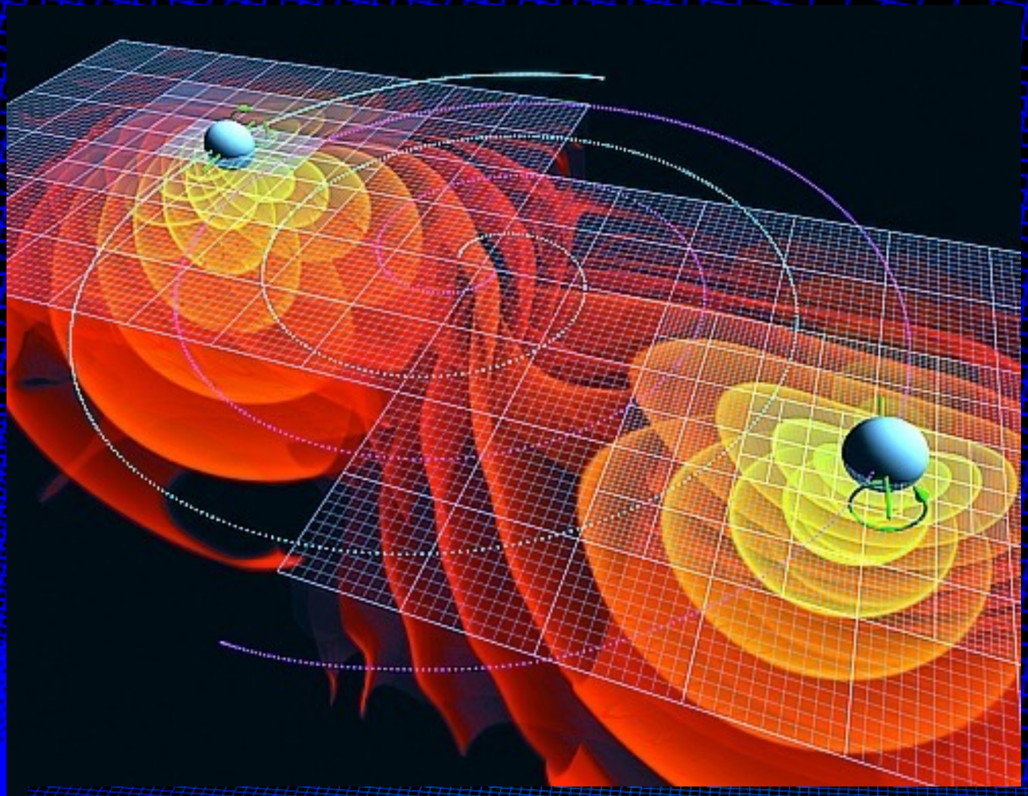

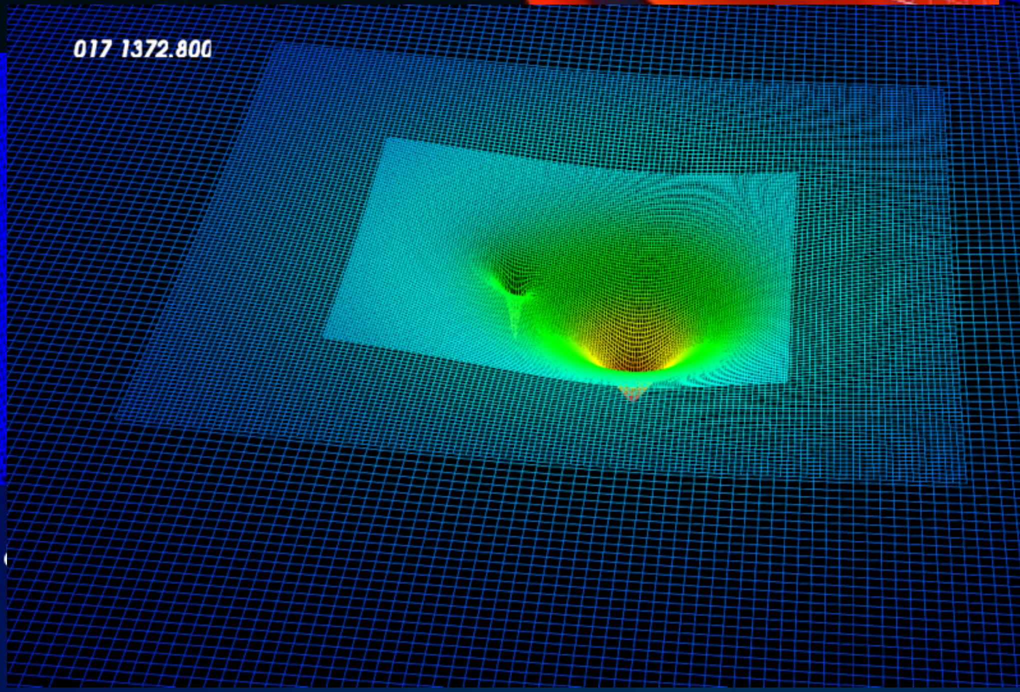

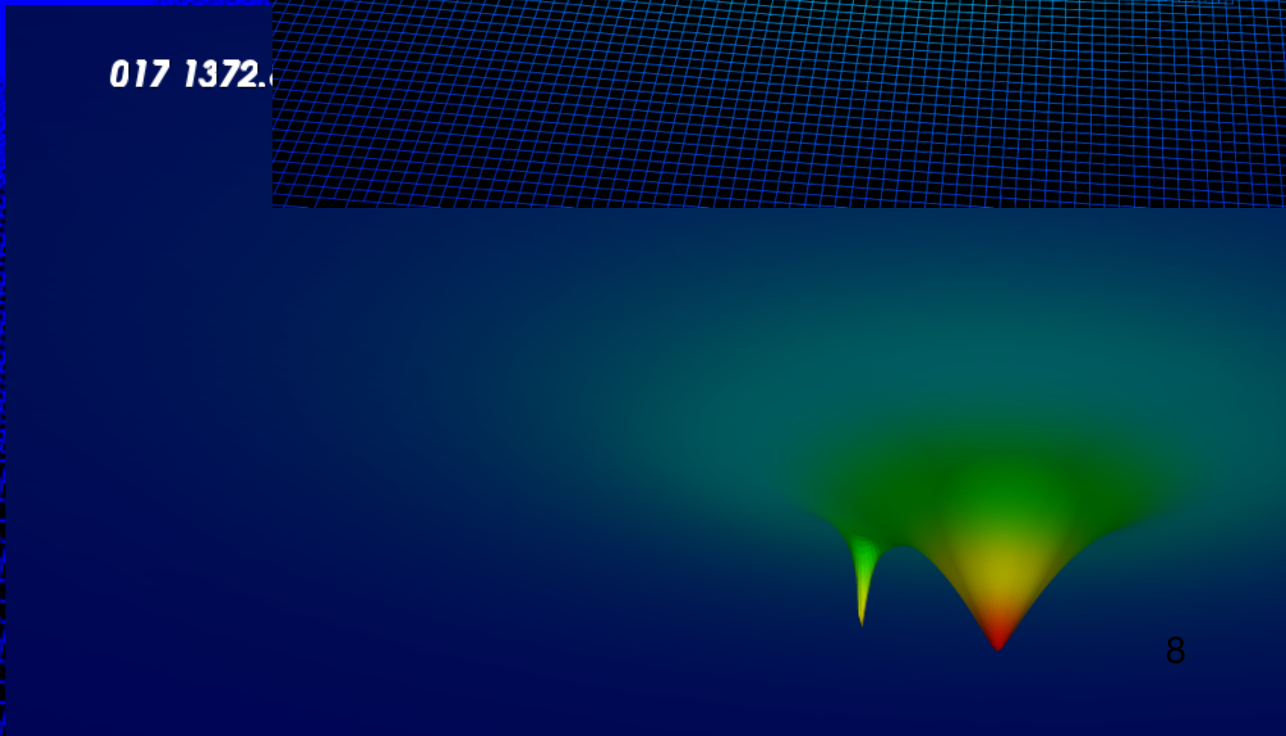

Project History & Strategy

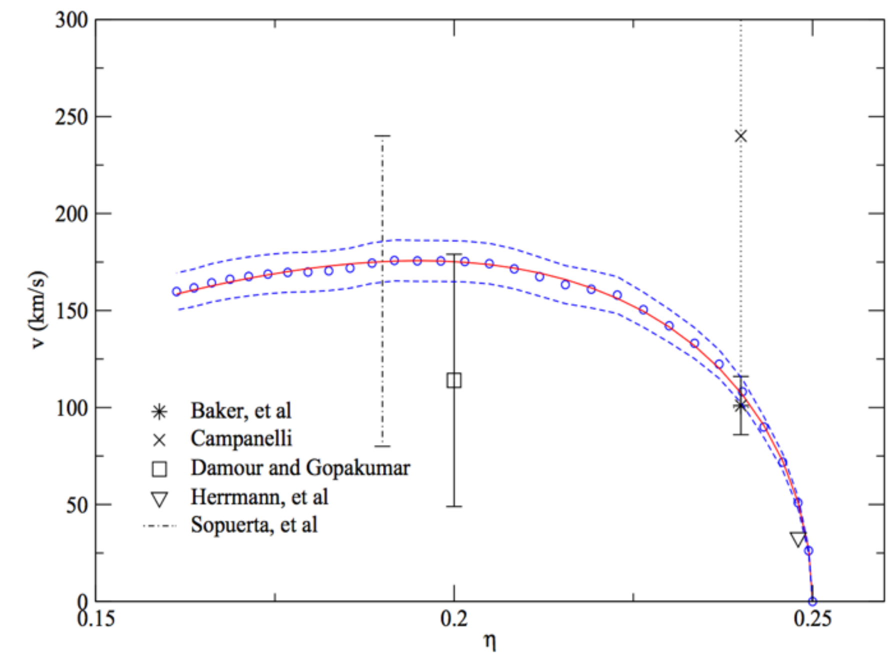

- 2006: Uni Jena + 1 million hours from DECI LRZ/HLRB-II + MPI Hannover/Potsdam
  - **phenomenological waveform program: analytical waveform models in the frequency domain**
    - make use of degeneracies and hierarchy in 7D parameter space
    - robust finite difference code to explore parameter space
  - 43 publications
- 16.7 + 37 million hours in PRACE 3+5, waiting for current evaluation
  - need hundreds of cases @ 10^4 - 10^6 CPU hours to ensure detection and parameter estimation ~100s of million of CPU hours.
  - High throughput for many independent simulations at hundreds of cores.

Project History & Strategy

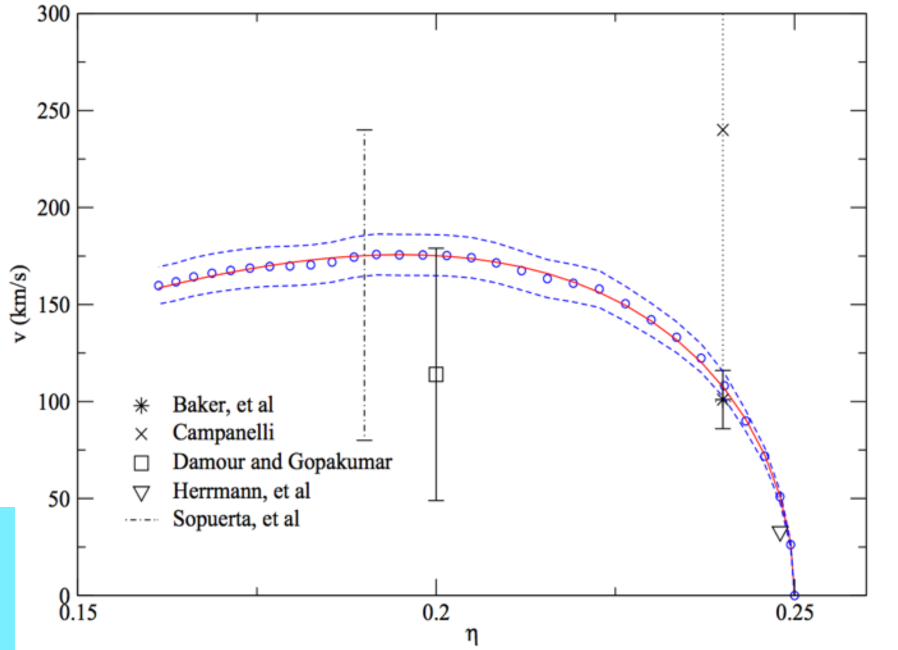

- 2006: Uni Jena + 1 million hours from DECI LRZ/HLRB-II + MPI Hannover/Potsdam

Thanks to DECI optimization support: Iris Christadler (LRZ)

  - **phenomenological waveform program: analytical waveform models in the frequency domain**
    - make use of degeneracies and hierarchy in 7D parameter space
    - robust finite difference code to explore parameter space
  - 43 publications
- 16.7 + 37 million hours in PRACE 3+5, waiting for current evaluation
  - need hundreds of cases @ 10^4 - 10^6 CPU hours to ensure detection and parameter estimation ~100s of million of CPU hours.
  - High throughput for many independent simulations at hundreds of cores.

Project History & Strategy

- 2006: Uni Jena + 1 million hours from DECI
  LRZ/HLRB-II + MPI Hannover/Potsdam

Thanks to DECI optimization support: Iris Christadler (LRZ)

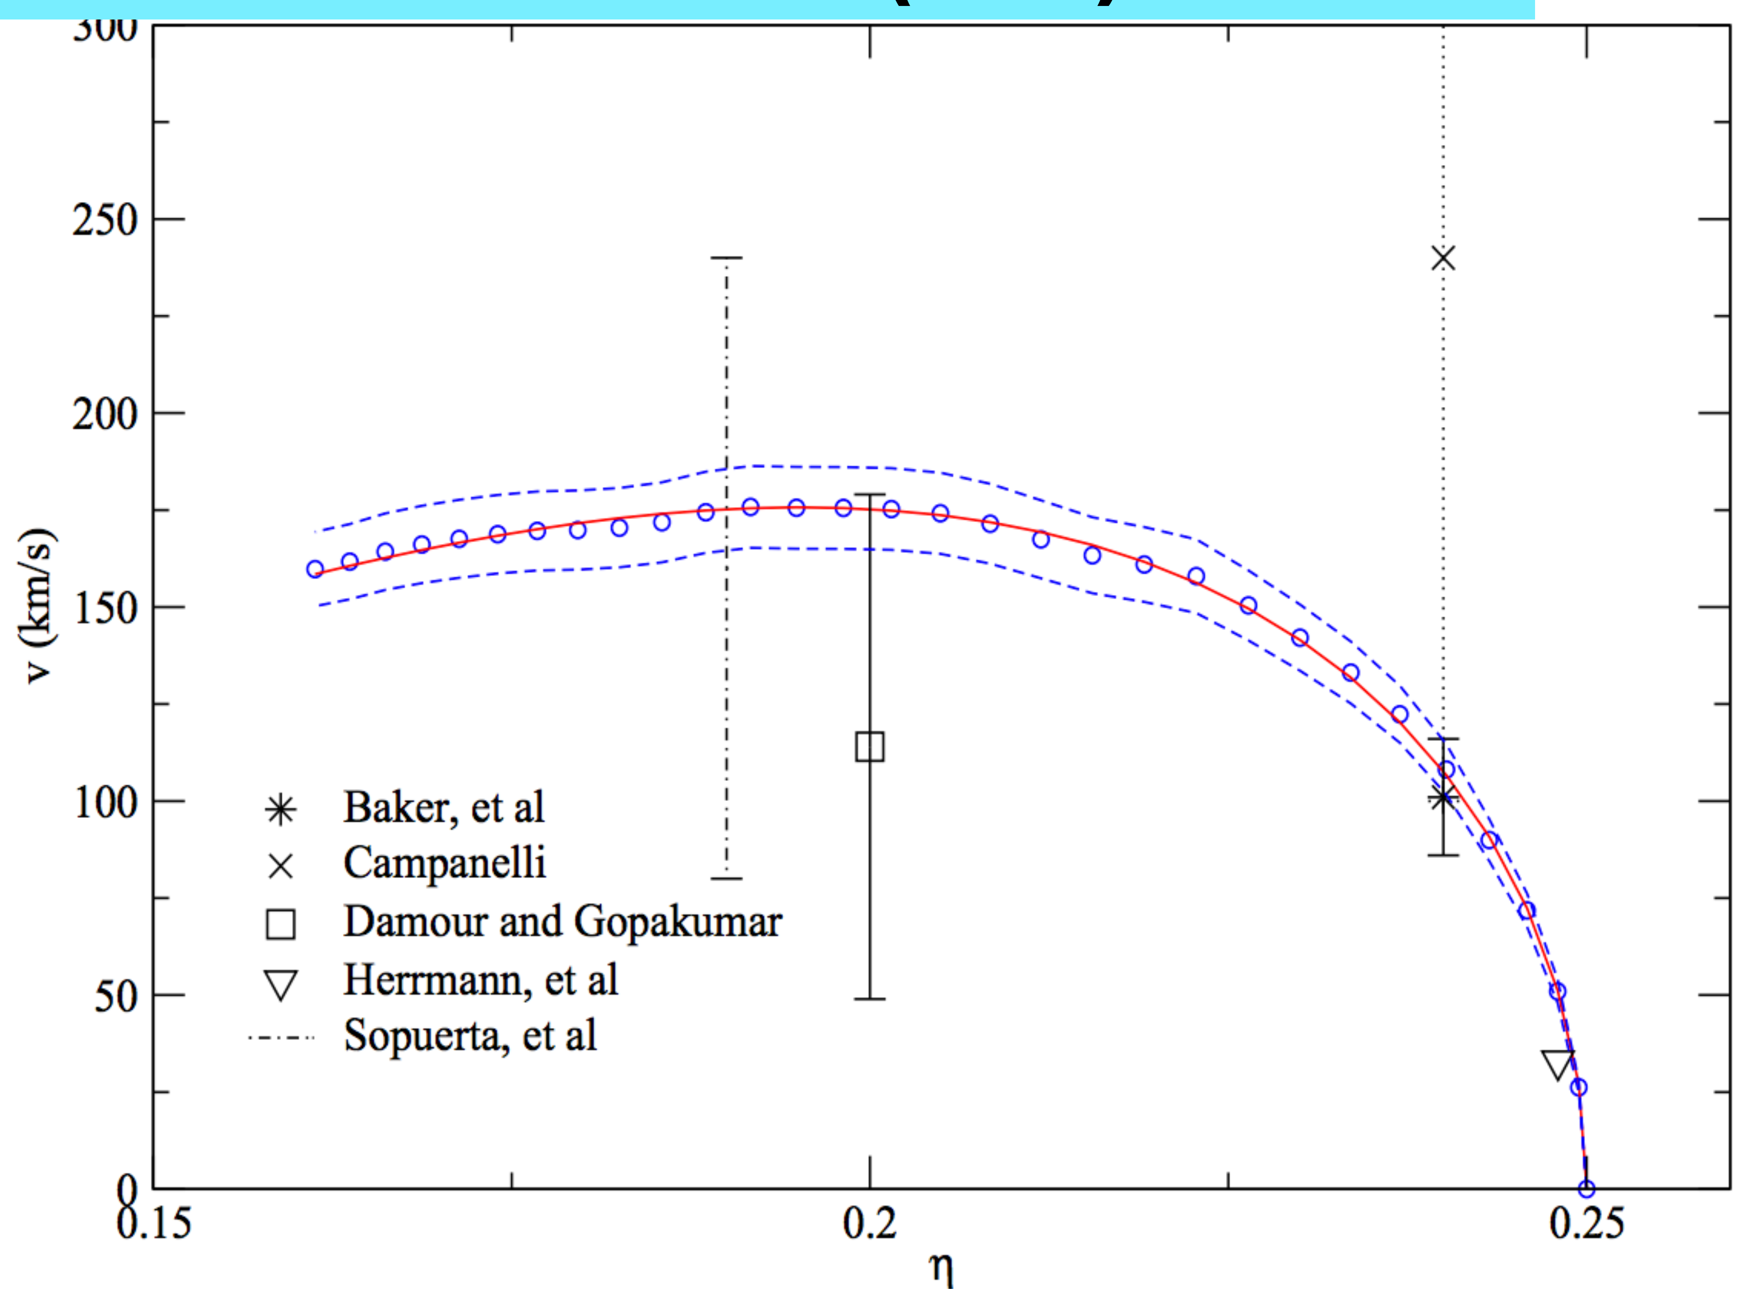

-
- r space
-
- 16 ... uation
  - etection and
  - eds of cores.

Project History & Strategy

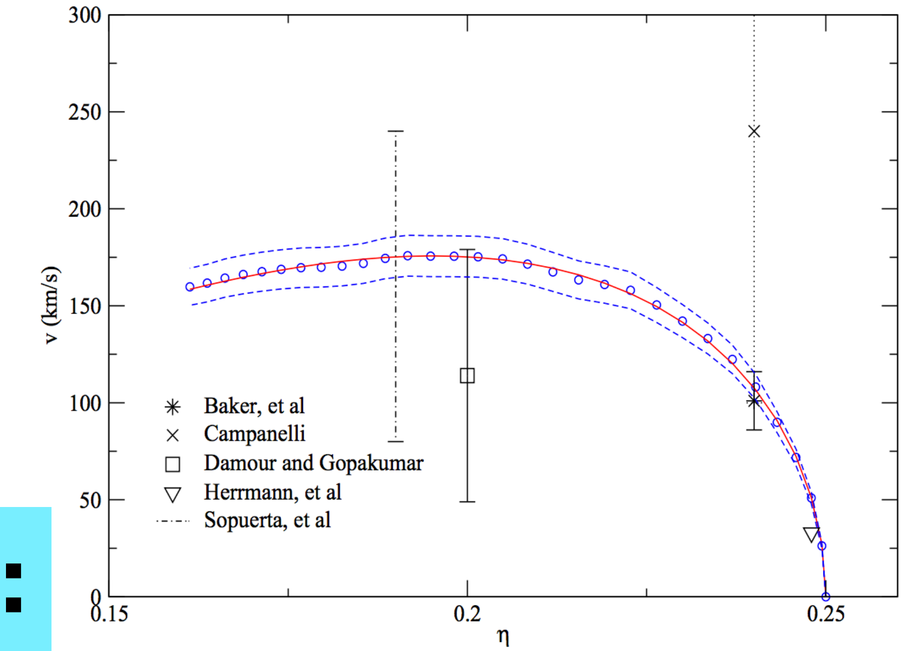

- 2006: Uni Jena + 1 million hours from DECI LRZ/HLRB-II + MPI Hannover/Potsdam

Thanks to DECI optimization support: Iris Christadler (LRZ)

- **phenomenological waveform program: analytical waveform models in the frequency domain**
  - make use of degeneracies and hierarchy in 7D parameter space
  - robust finite difference code to explore parameter space
- 43 publications
- 16.7 + 37 million hours in PRACE 3+5, waiting for current evaluation
  - need hundreds of cases @ 10^4 - 10^6 CPU hours to ensure detection and parameter estimation ~100s of million of CPU hours.
  - High throughput for many independent simulations at hundreds of cores.

Project History & Strategy

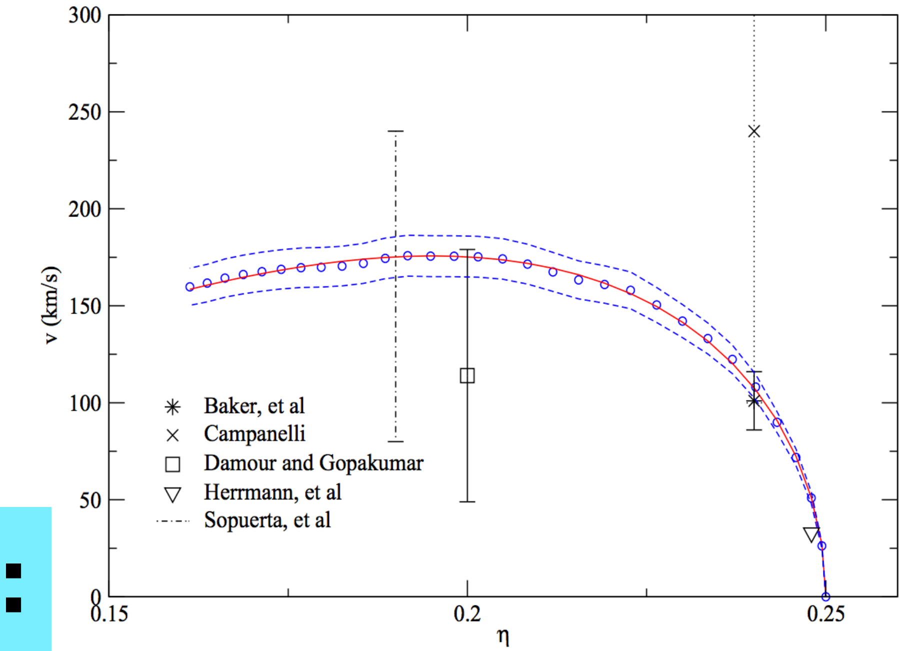

- 2006: Uni Jena + 1 million hours from DECI LRZ/HLRB-II + MPI Hannover/Potsdam

Thanks to DECI optimization support: Iris Christadler (LRZ)

- **phenomenological waveform program: analytical waveform models in the frequency domain**
  - make use of degeneracies and hierarchy in 7D parameter space
  - robust finite difference code to explore parameter space
- 43 publications

Competing project: SXS collaboration
Calibration of Effective-One-Body models in time domain with long simulations, use spectral code.

- High throughput for many independent simulations at hundreds of cores.

Project History & Strategy

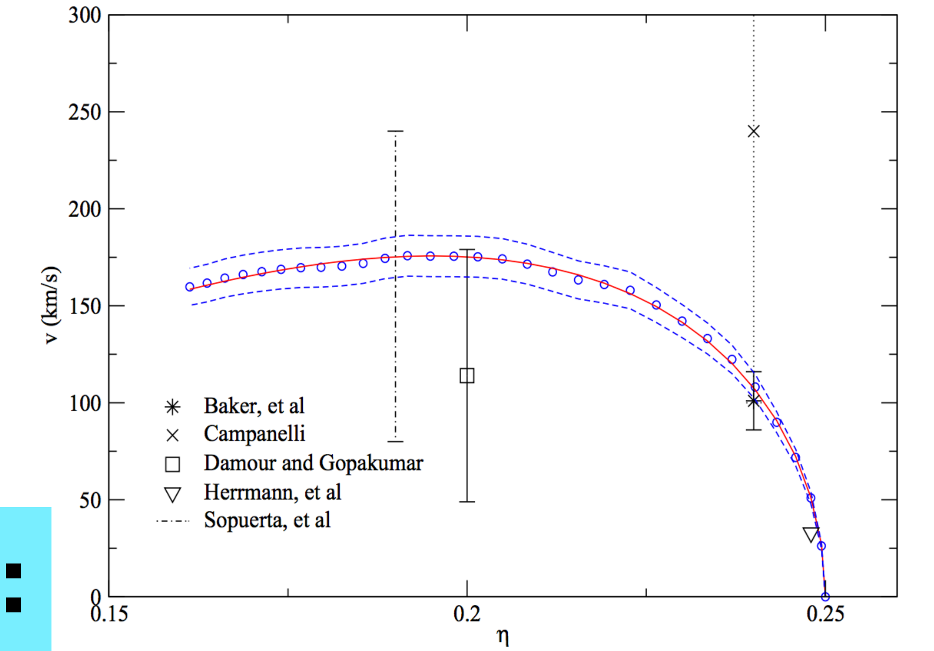

- 2006: Uni Jena + 1 million hours from DECI LRZ/HLRB-II + MPI Hannover/Potsdam

Thanks to DECI optimization support: Iris Christadler (LRZ)

  - **phenomenological waveform program: analytical waveform models in the frequency domain**
    - make use of degeneracies and hierarchy in 7D parameter space
    - robust finite difference code to explore parameter space
  - 43 publications
- 16.7 + 37 million hours in PRACE 3+5, waiting for current evaluation
  - need hundreds of cases @ 10^4 - 10^6 CPU hours to ensure detection and parameter estimation ~100s of million of CPU hours.
  - High throughput for many independent simulations at hundreds of cores.

Code infrastructure

- Use 2 codes, use MPI and OpenMP domain decomposition parallelization:
    - BAM (developed originally at Uni Jena, C) - used for most production runs
    - Einstein toolkit (open source, C, C++, F90) - very active development
    - performance & scaling very similar
- Use explicit Runge-Kutta time-stepping: time step limited by Courant condition
    - Use ghost-point based variant of Berger-Oliger to refine temporal and spatial resolution.

- Outer grids dominate memory requirements, innermost grids speed.
- Checkpointing: longest simulation ran ~ 4 months.
- Run on the minimal number of cores for the problem, use available memory/core.

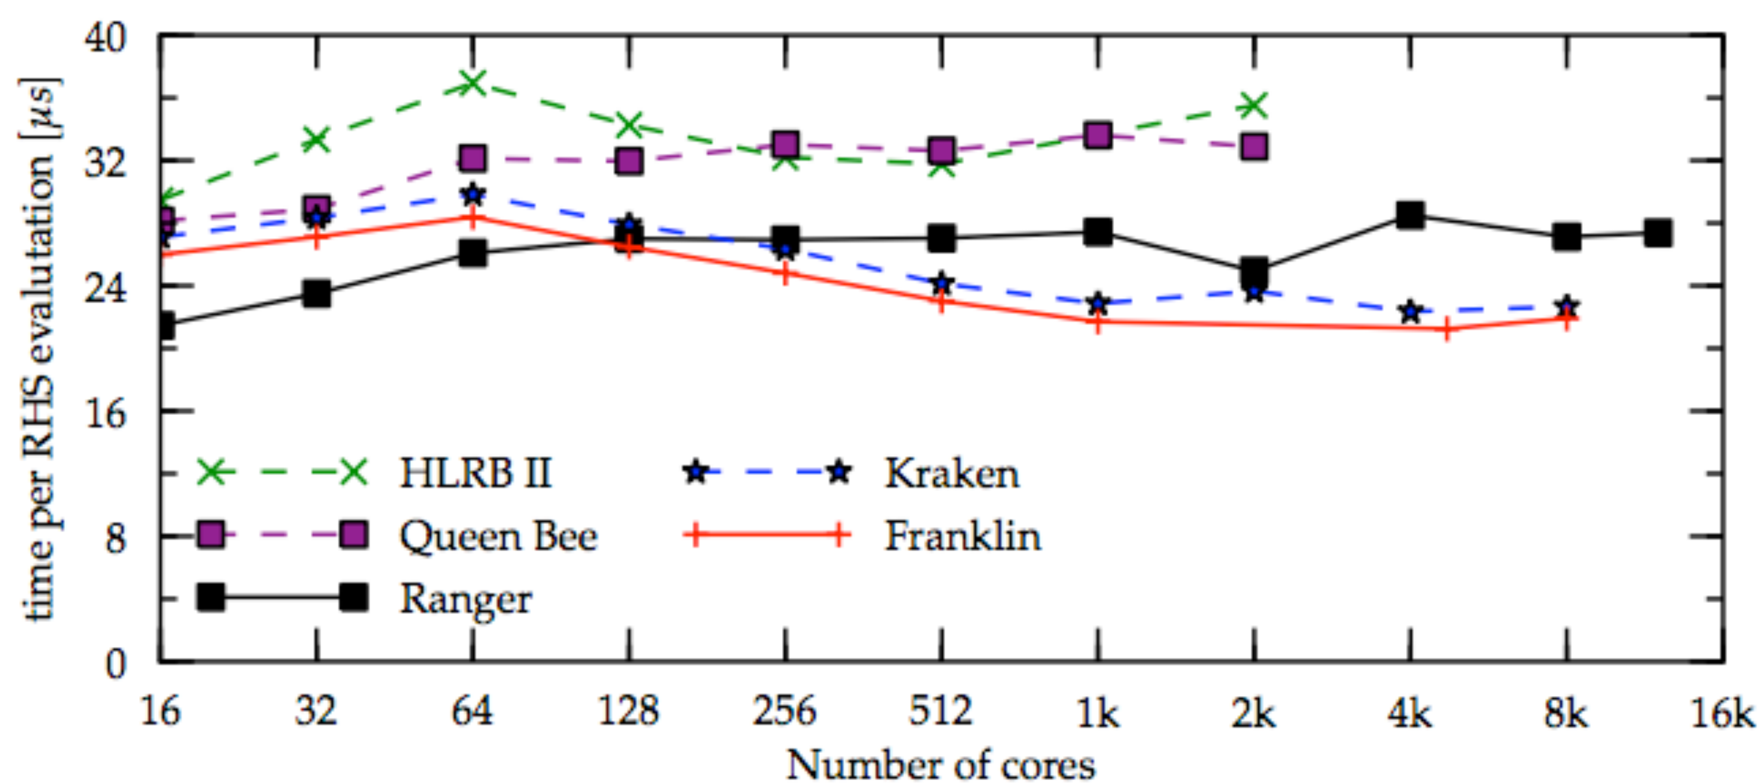

Figure 2. Results from weak scaling tests evolving the Einstein equations on
refinement grid structure with nine levels. This shows the time required per grid poi

Code infrastructure

- Use 2 codes, use MPI and OpenMP domain decomposition parallelization:
    - BAM (developed originally at Uni Jena, C) - used for most production runs
    - Einstein toolkit (open source, C, C++, F90) - very active development
    - performance & scaling very similar
- Use explicit Runge-Kutta time-stepping: time step limited by Courant condition
    - Use ghost-point based variant of Berger-Oliger to refine temporal and spatial resolution.

- Outer grids dominate memory requirements, innermost grids speed.
- Checkpointing: longest simulation ran ~ 4 months.
- Run on the minimal number of cores for the problem, use available memory/core.

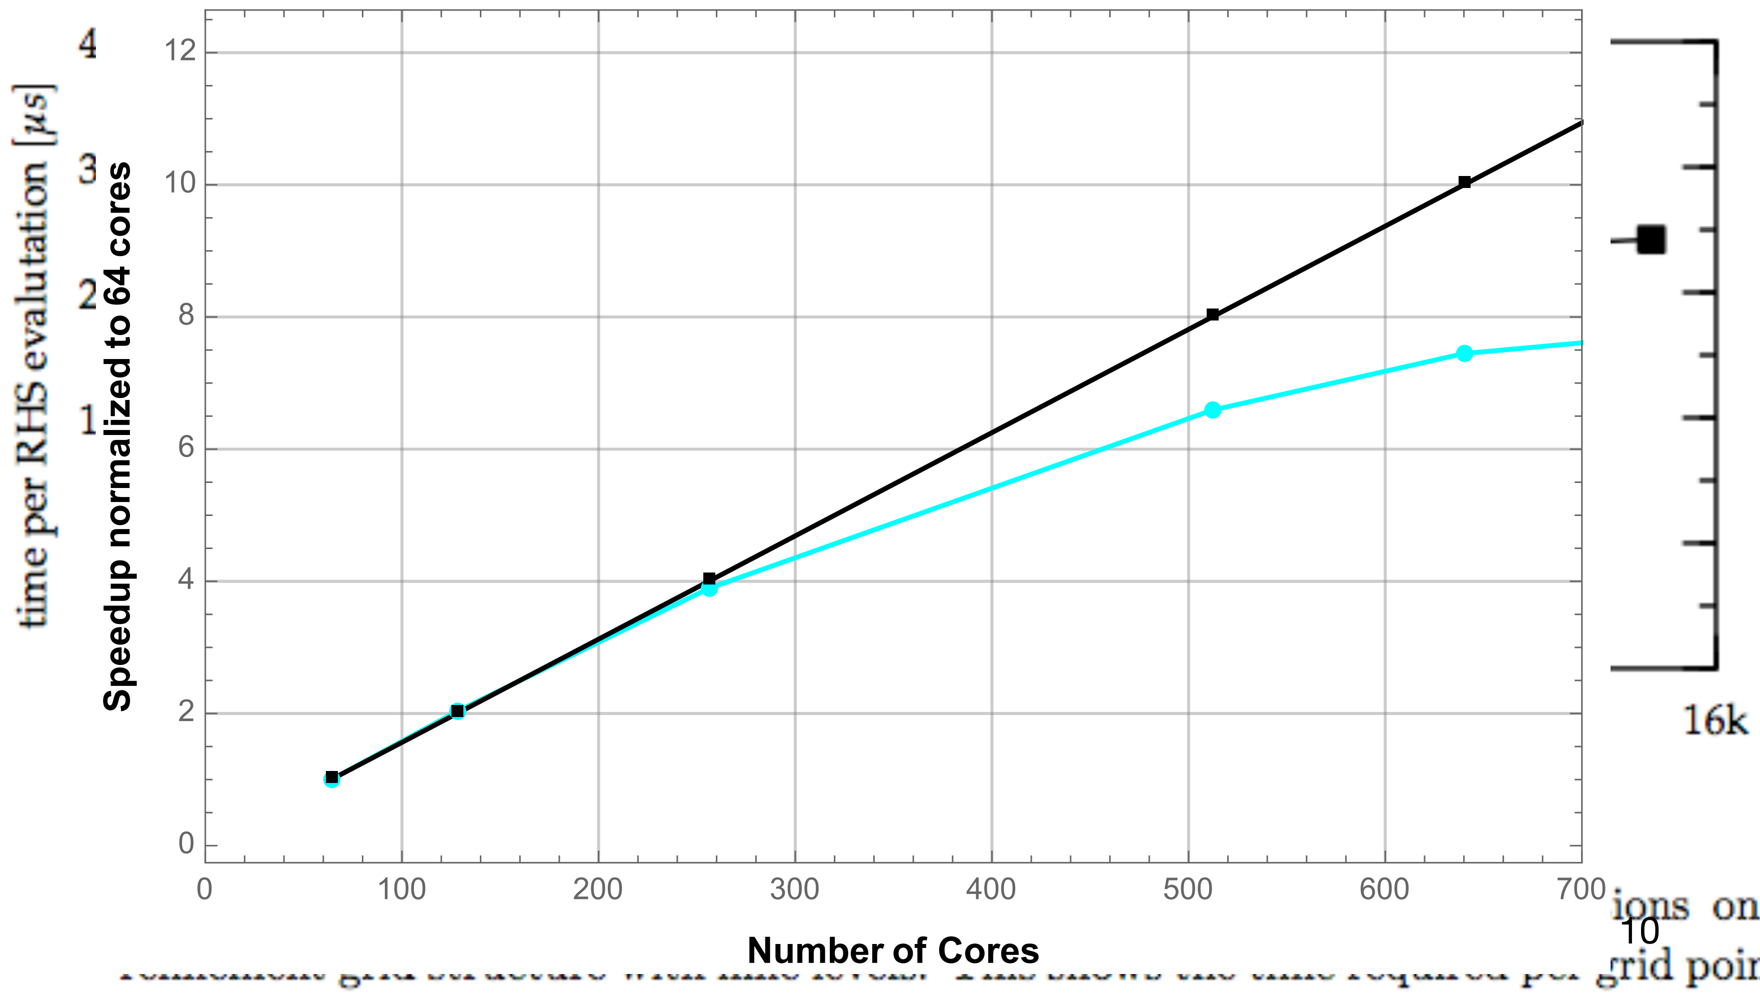

Simulations performed & Job-bundling strategy

- PRACE-5: calibrate non-precessing model up to mass ratio 18, 37 million hours
- 12 high mass ratio cases: most expensive BH-simulations we are aware of > 10^6 hours.
- Prepare large scale precessing study.
- Bundle several cases into bigger jobs, possibly reconfigure after each checkpoint.
- Monitor throughput, queue times, manage job chaining etc. with cron.
- LRZ provided workaround for bug in IBM parallel environment:
    - https://www.lrz.de/services/compute/supermuc/loadleveler/special/index.html#subjobs-intel
- Used > 20 million hours during last ~4 months of allocation.

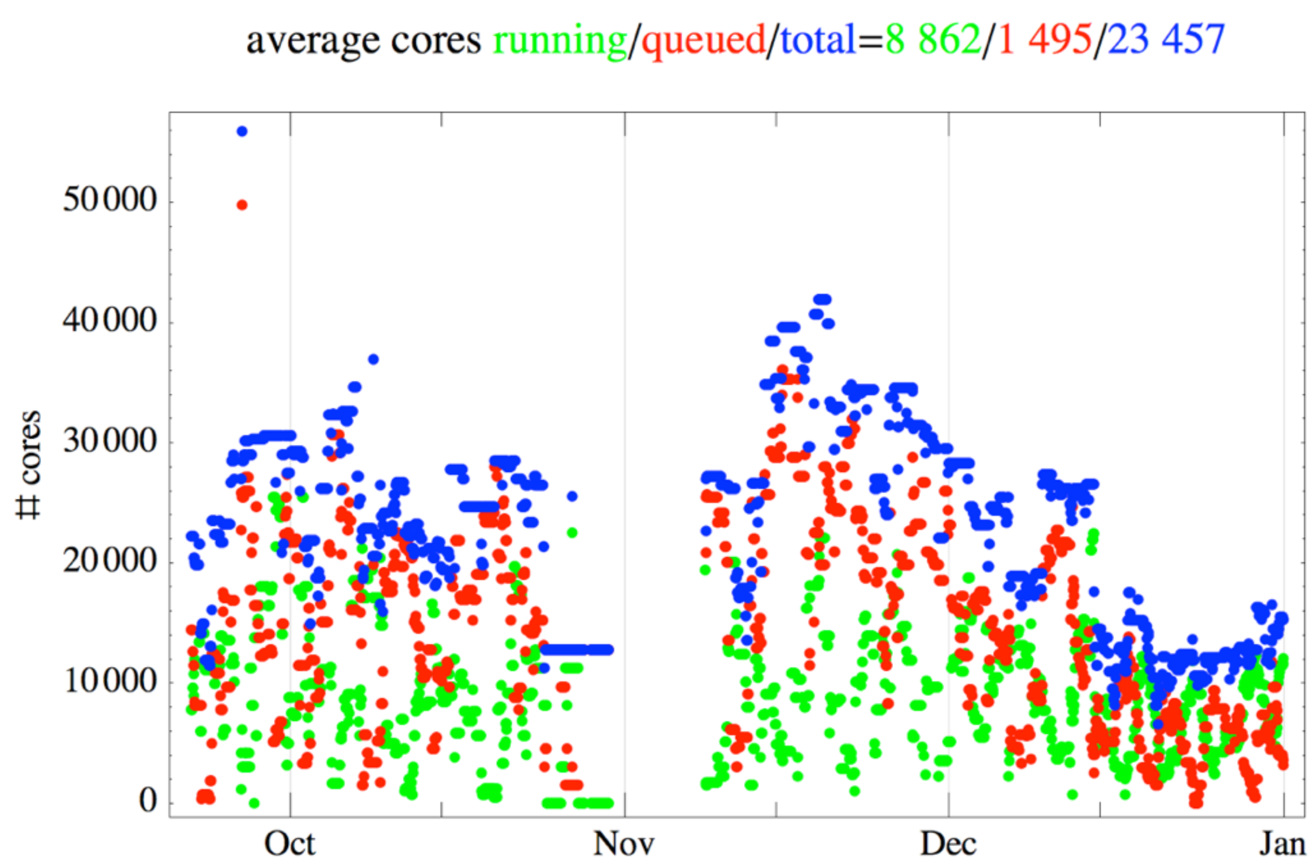

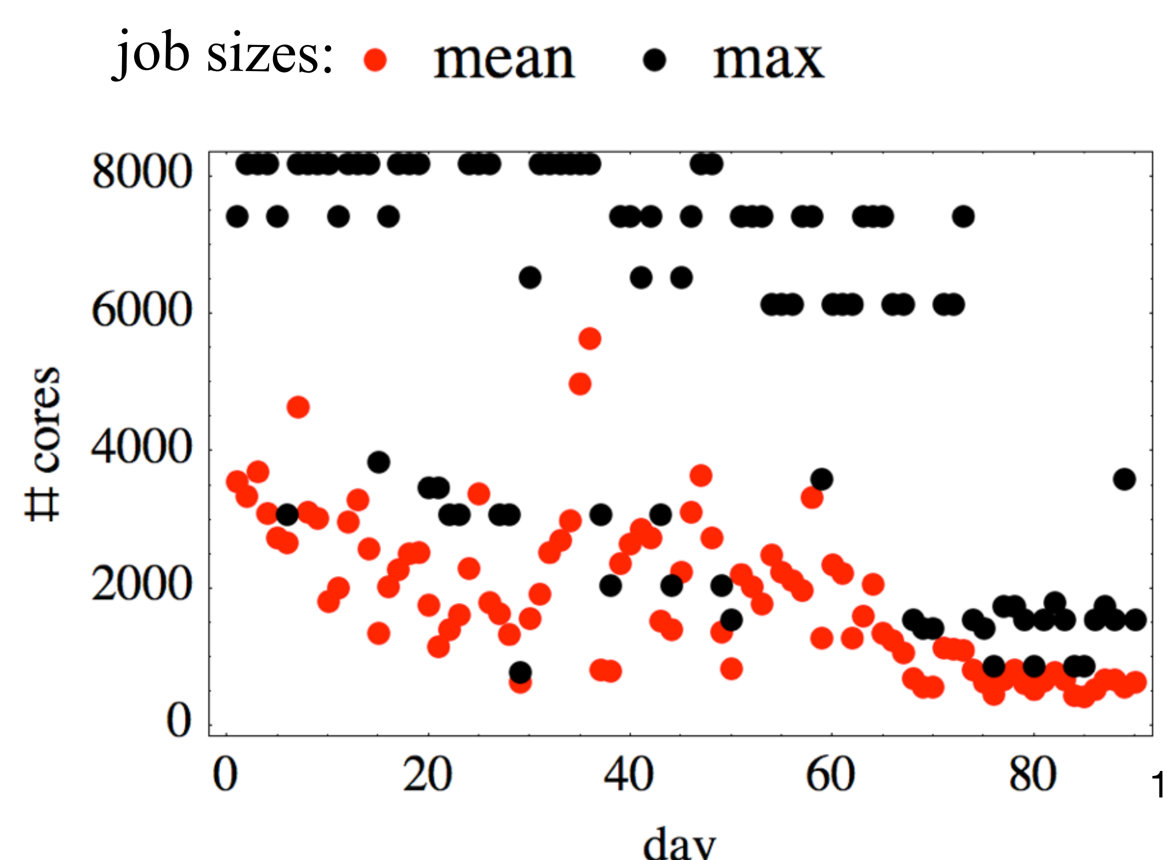

The Challenge of Precession

- Spins parallel to orbital angular momentum: no precession, orbital plane preserved.

- Orbital angular momentum and individual spins slowly precess around total angular momentum.

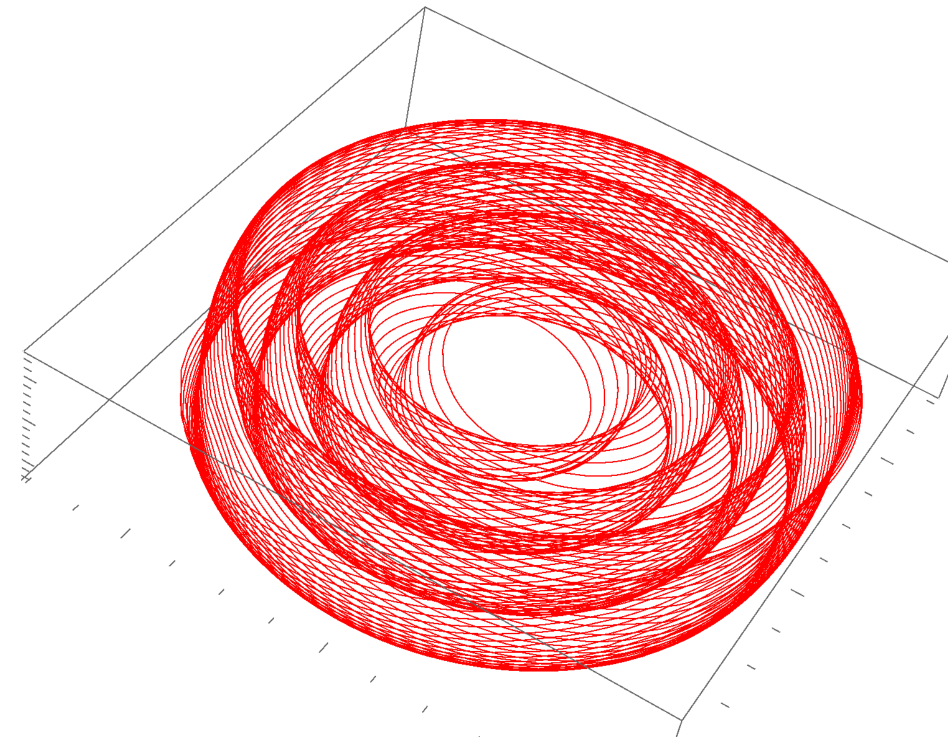

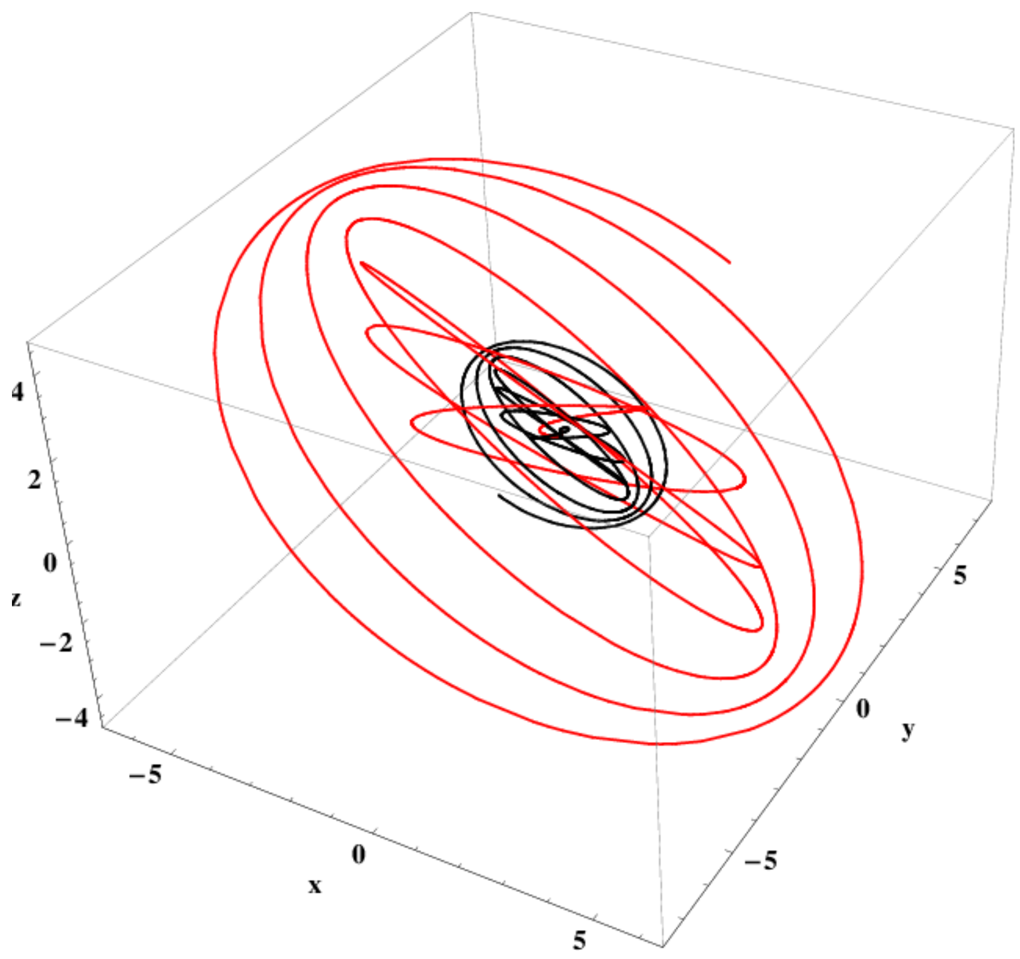

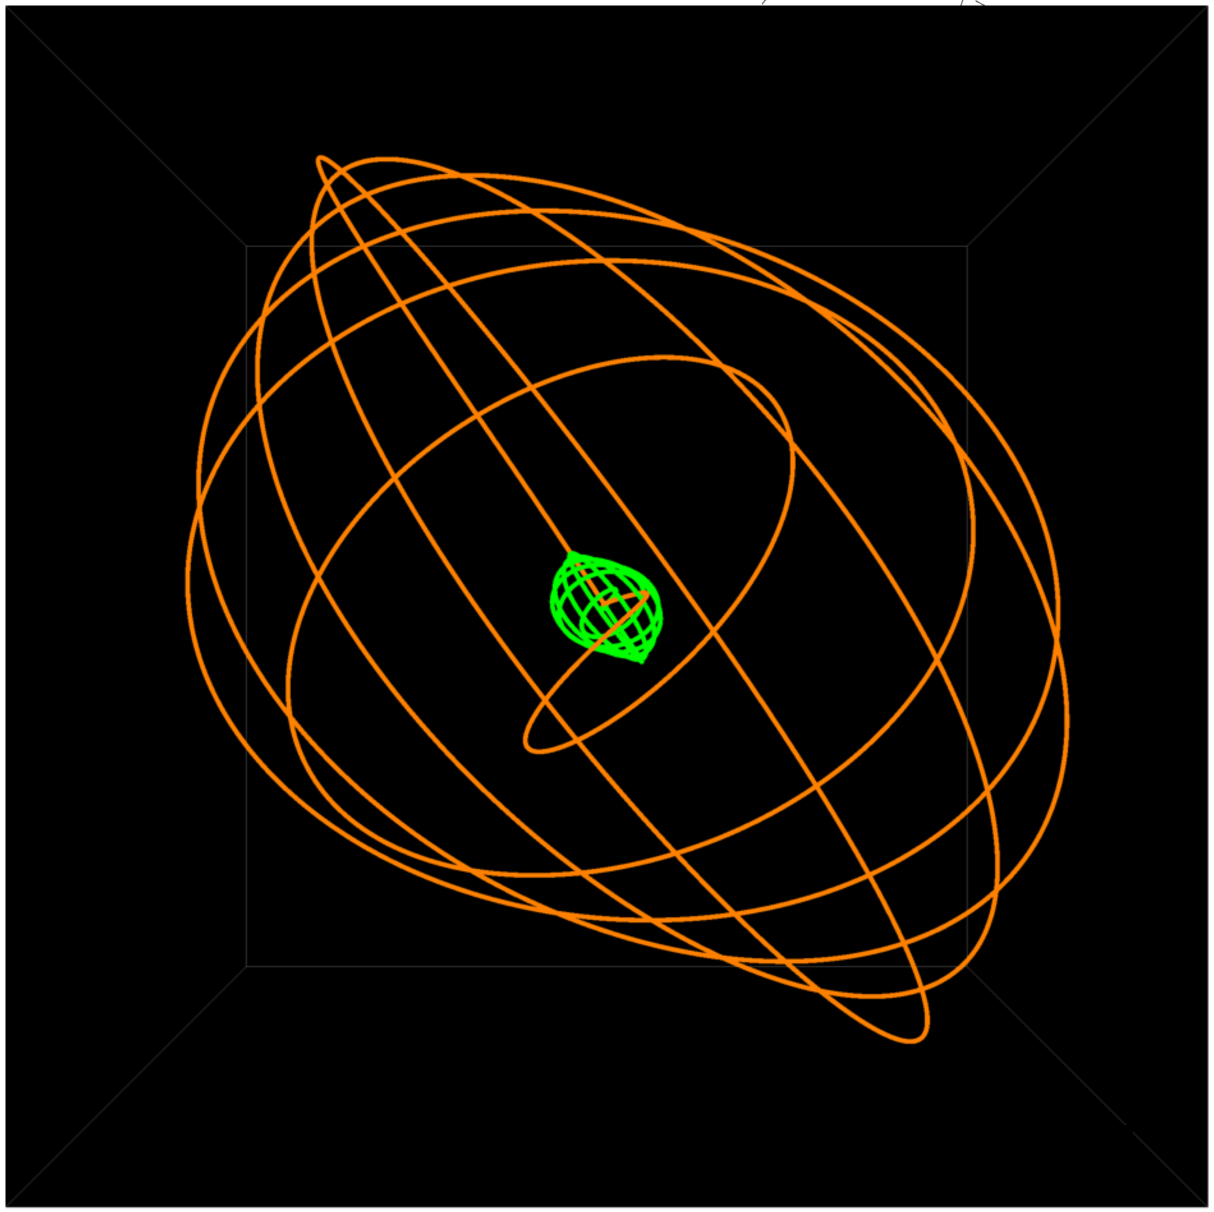

A path toward understanding precession

- In a co-rotating frame the phasing and radiated angular momentum are essentially unaffected by precession - “simple standard form” of a precessing WF: align z-axis with principal axis of the radiation quadrupole moment [Schmidt+ PRD 2011]
- Spherical harmonic mode structure in standard frame corresponds to non-precessing case -> “twisting up” accurate aligned spin model with “post-Newtonian” Euler angles works well [Schmidt+ PRD 2012, Hannam+ PRL 2013]

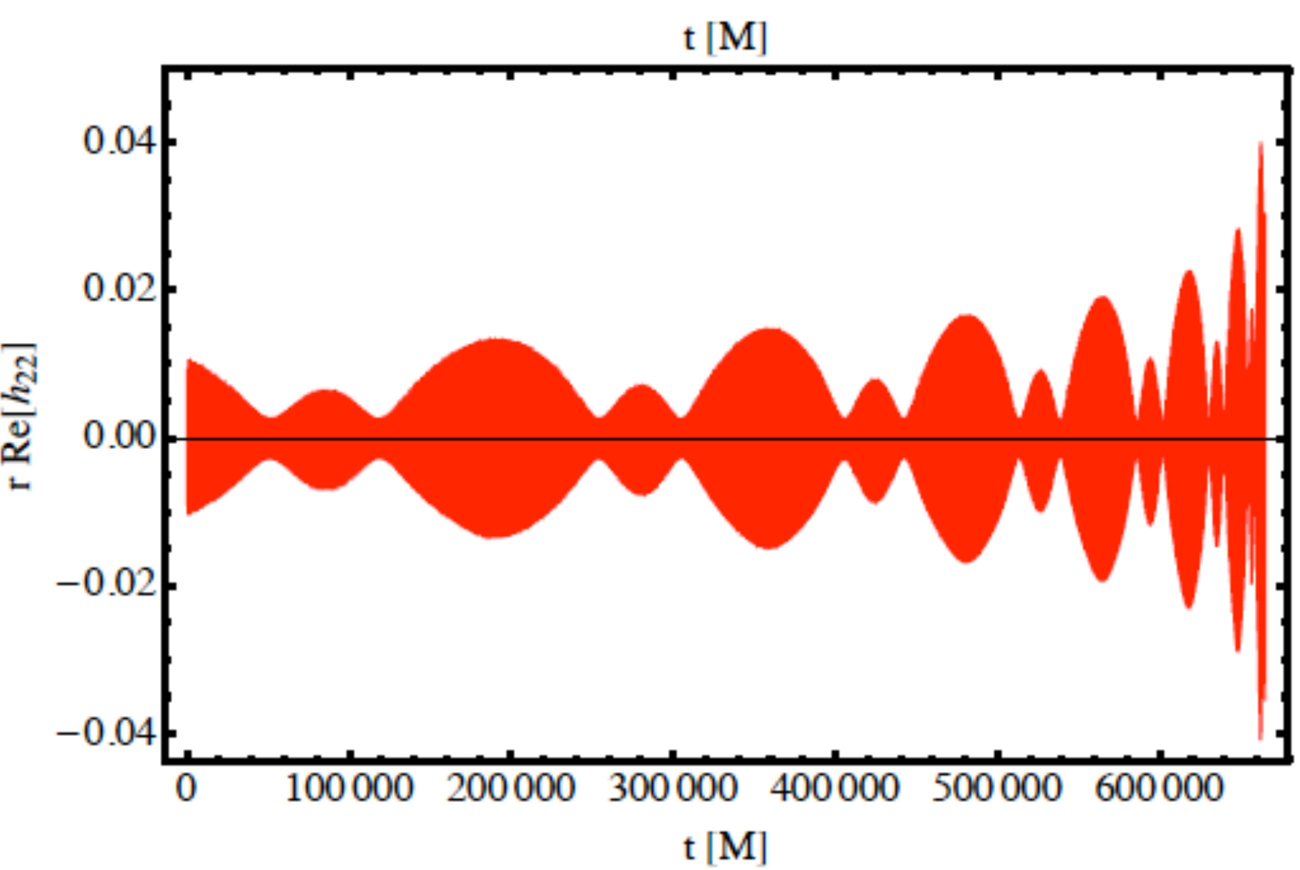

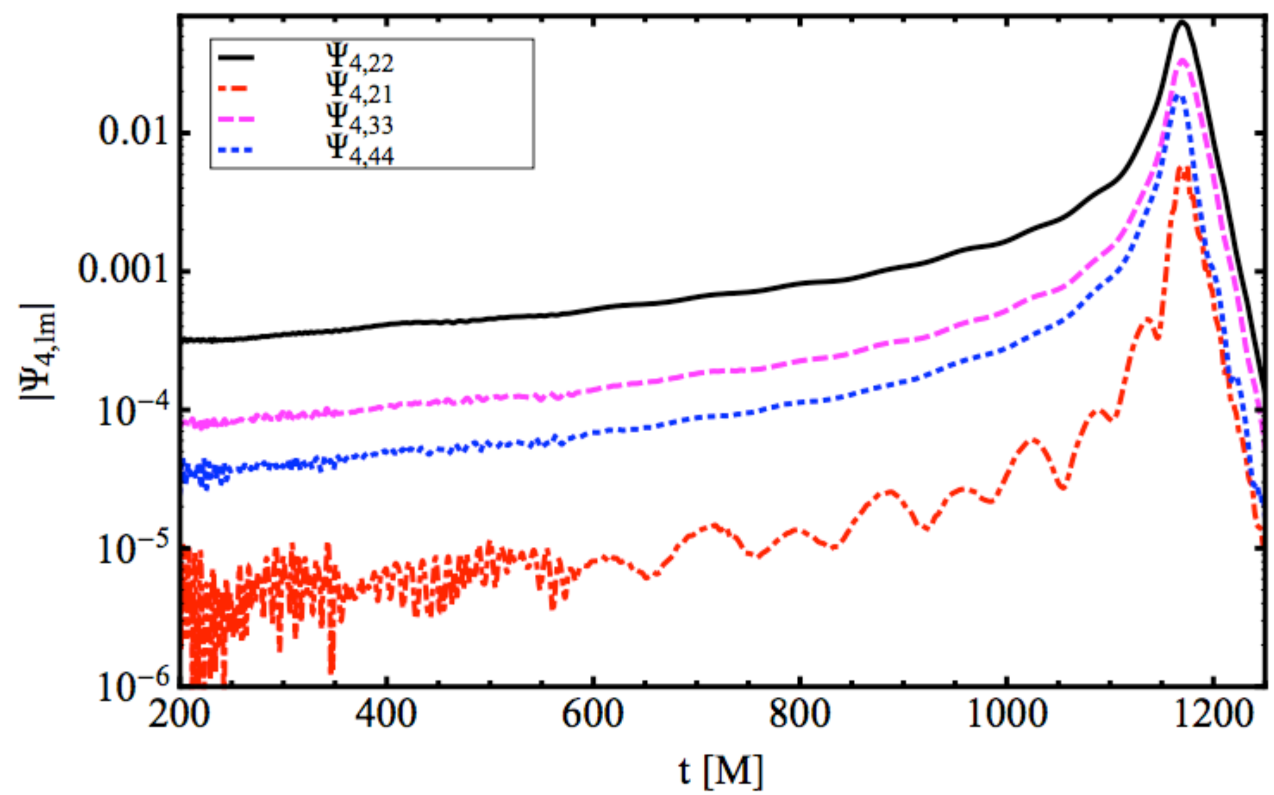

Fitting factors: Models vs. PN-NR-hybrids

FF ~ detection efficiency: 0.97 standard detection limit

$$\langle h_1, h_2 \rangle = \max_{\phi_0, t_0} 4\Re \int_{f_1}^{f_2} \frac{\tilde{h}_1(f)\, \tilde{h}_2^*(f)}{S_n(f)}\, df$$

- PhenomC: nonprecessing model
- PhenomP: PhenomC twisted up with PN

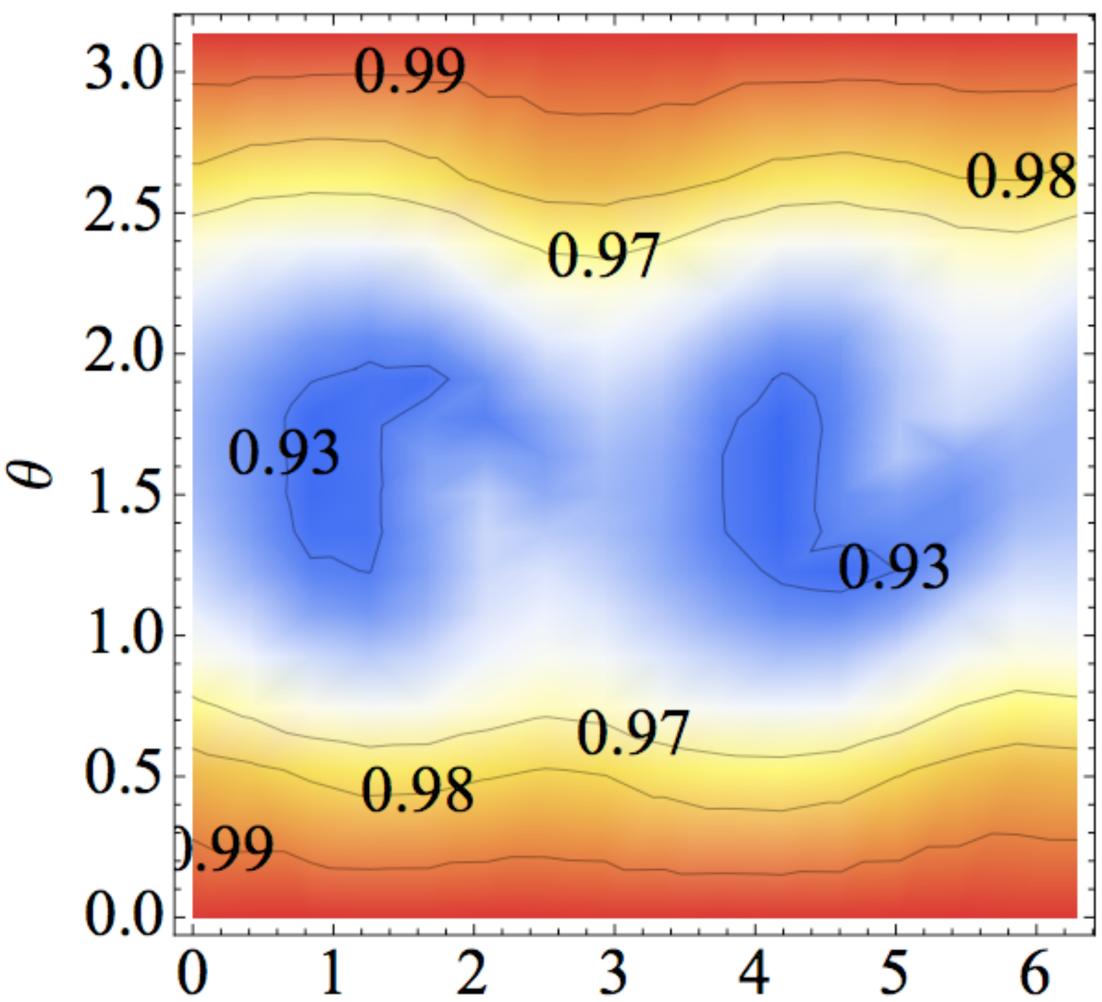

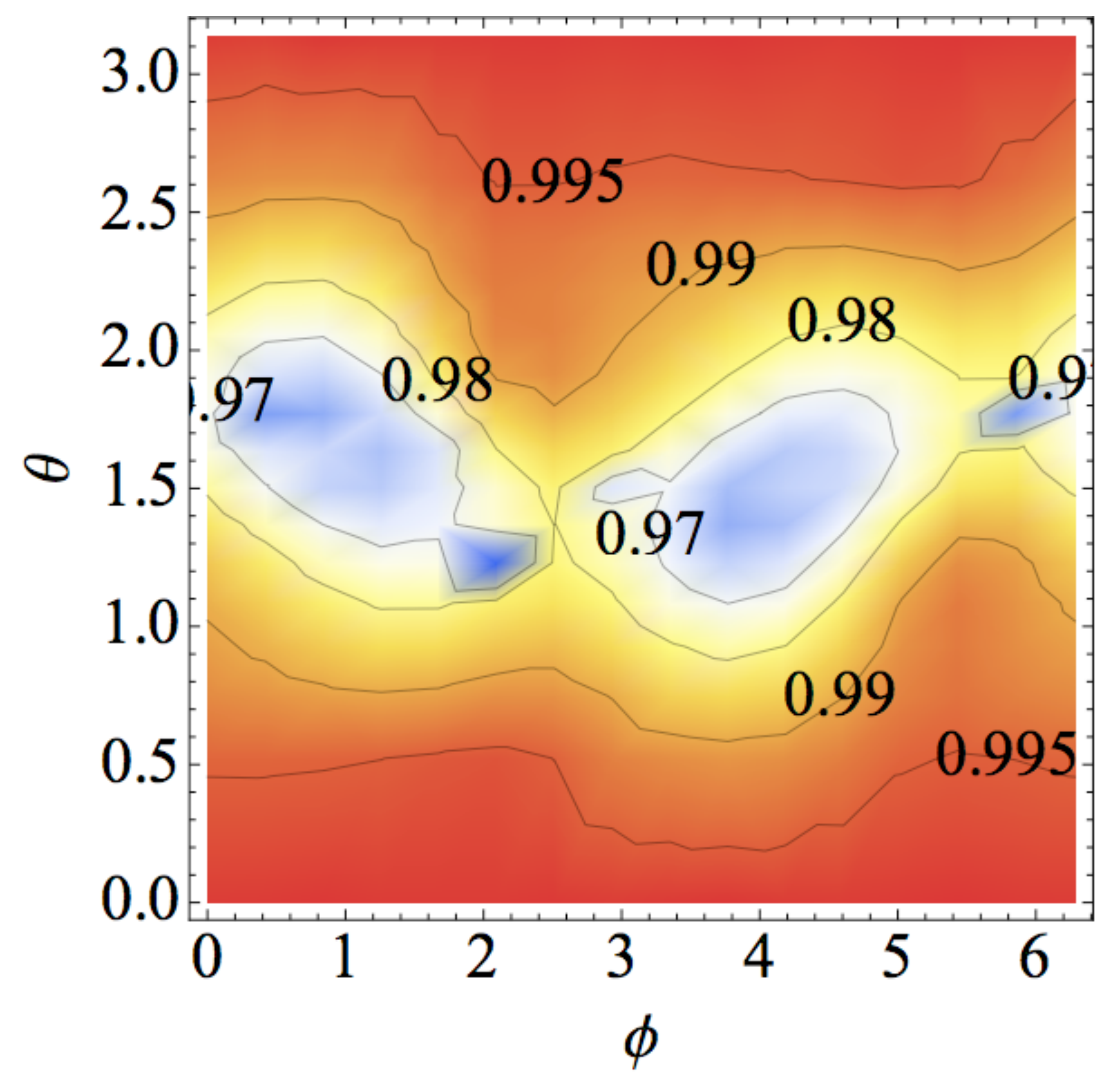

- Future: need to calibrate merger/ringdown to actual precessing NR waveforms.

Understanding the significance of subdominant modes

- Learned how to systematically glue post-Newtonian and numerical relativity data for general spherical harmonic modes.
- Understand where in parameter space higher modes are important when neglecting spin, starting to analyze general spinning case.

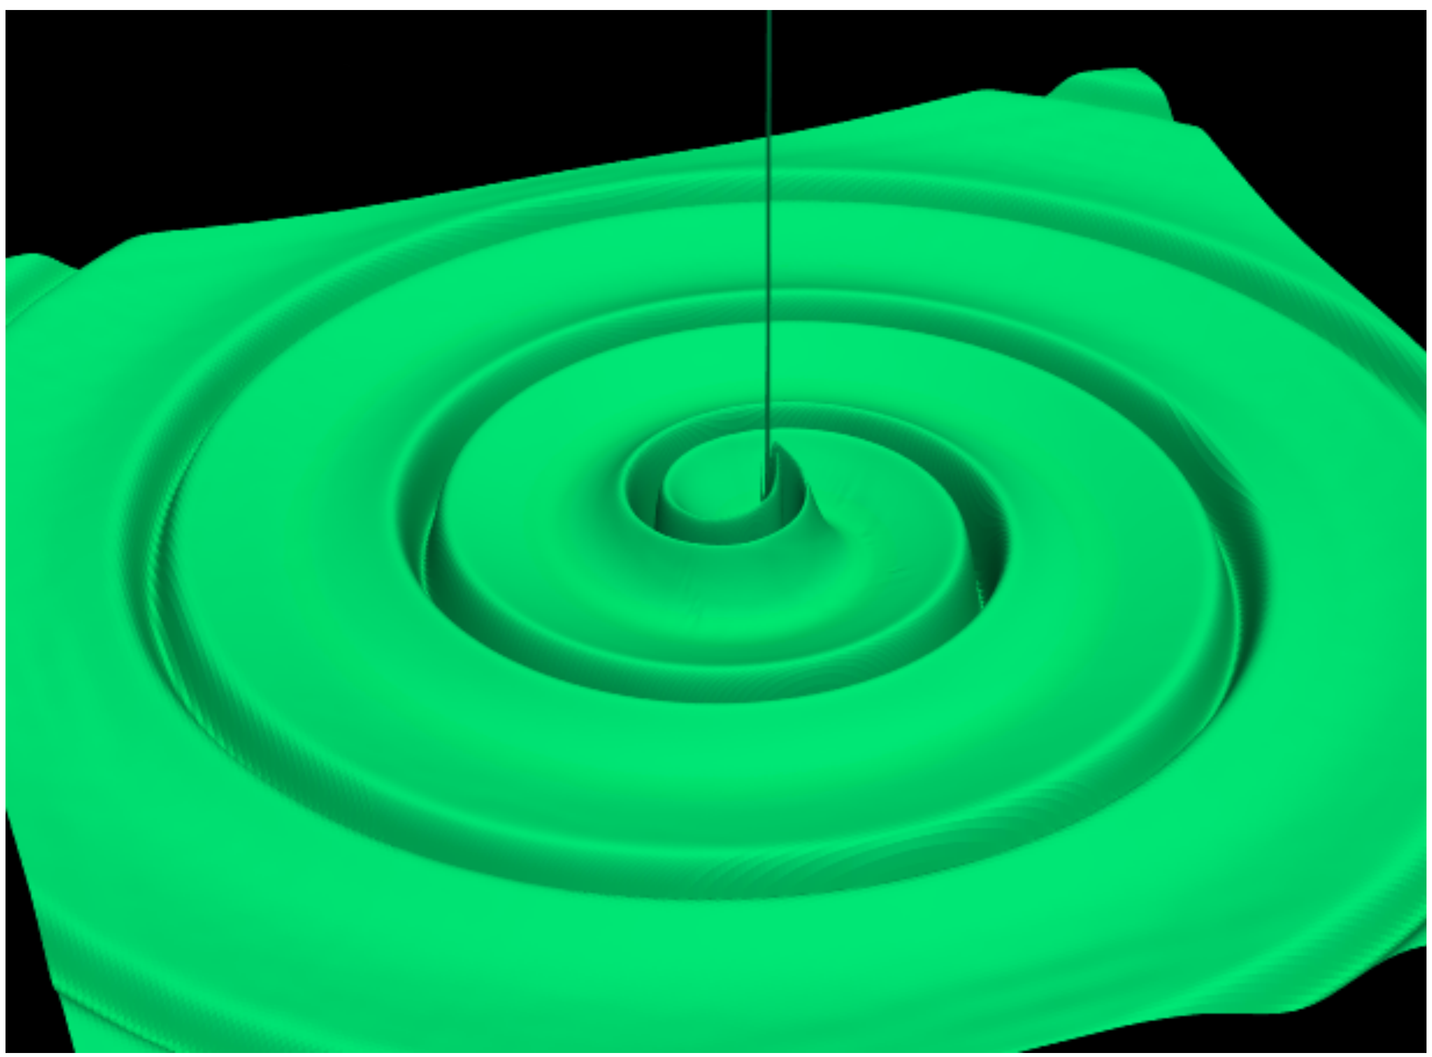

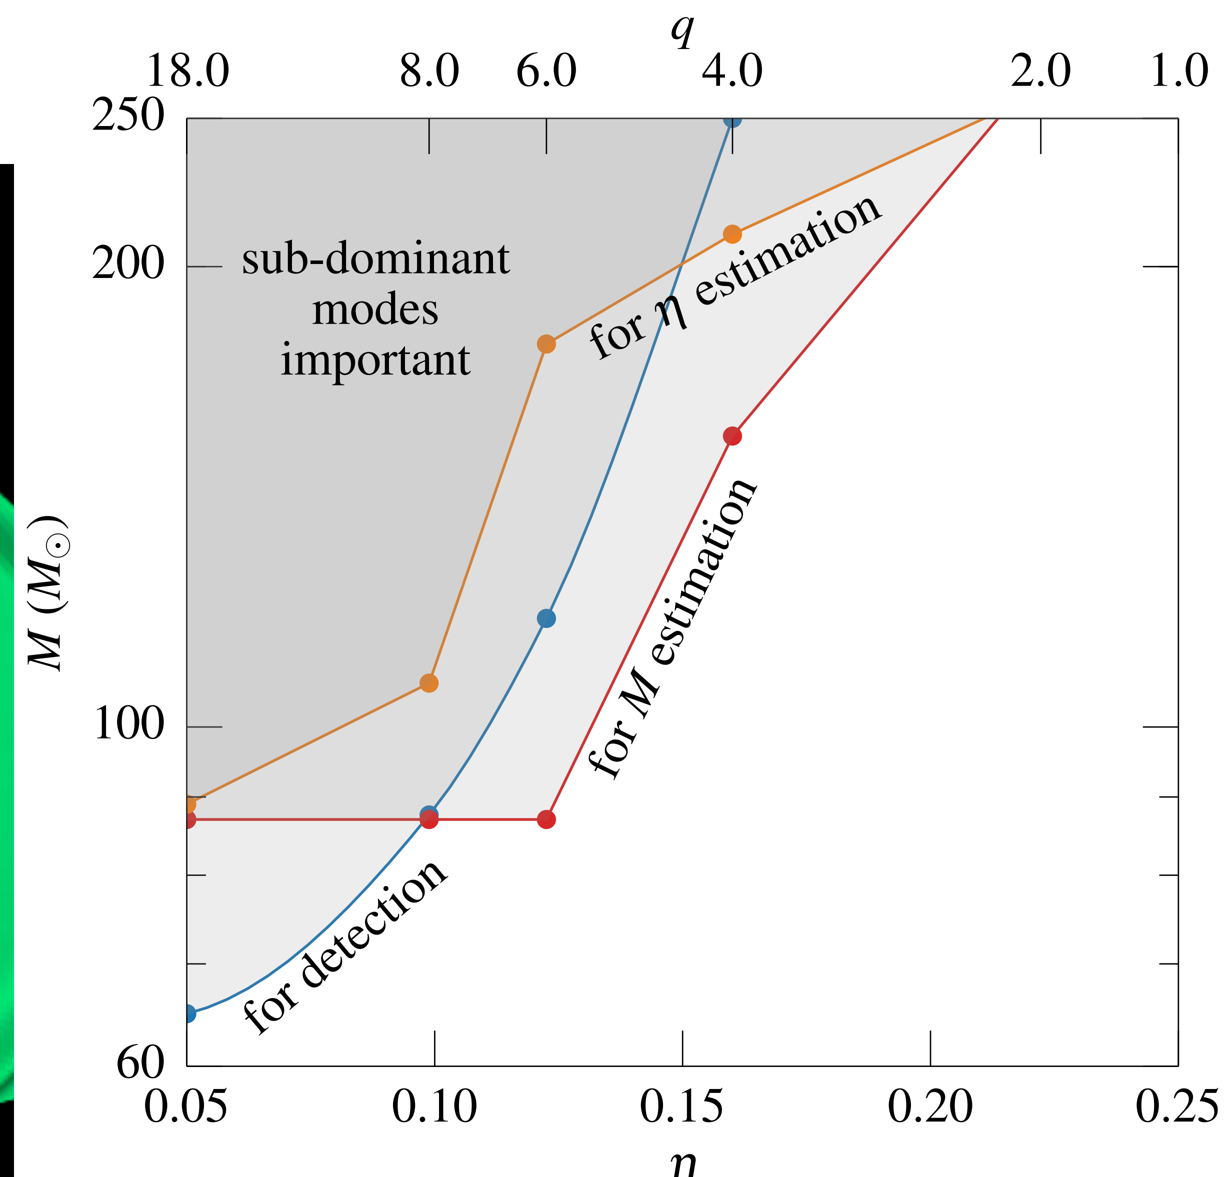

Understanding the significance of subdominant modes

- Learned how to systematically glue post-Newtonian and numerical relativity data for general spherical harmonic modes.
- Understand where in parameter space higher modes are important when neglecting spin, starting to analyze general spinning case.

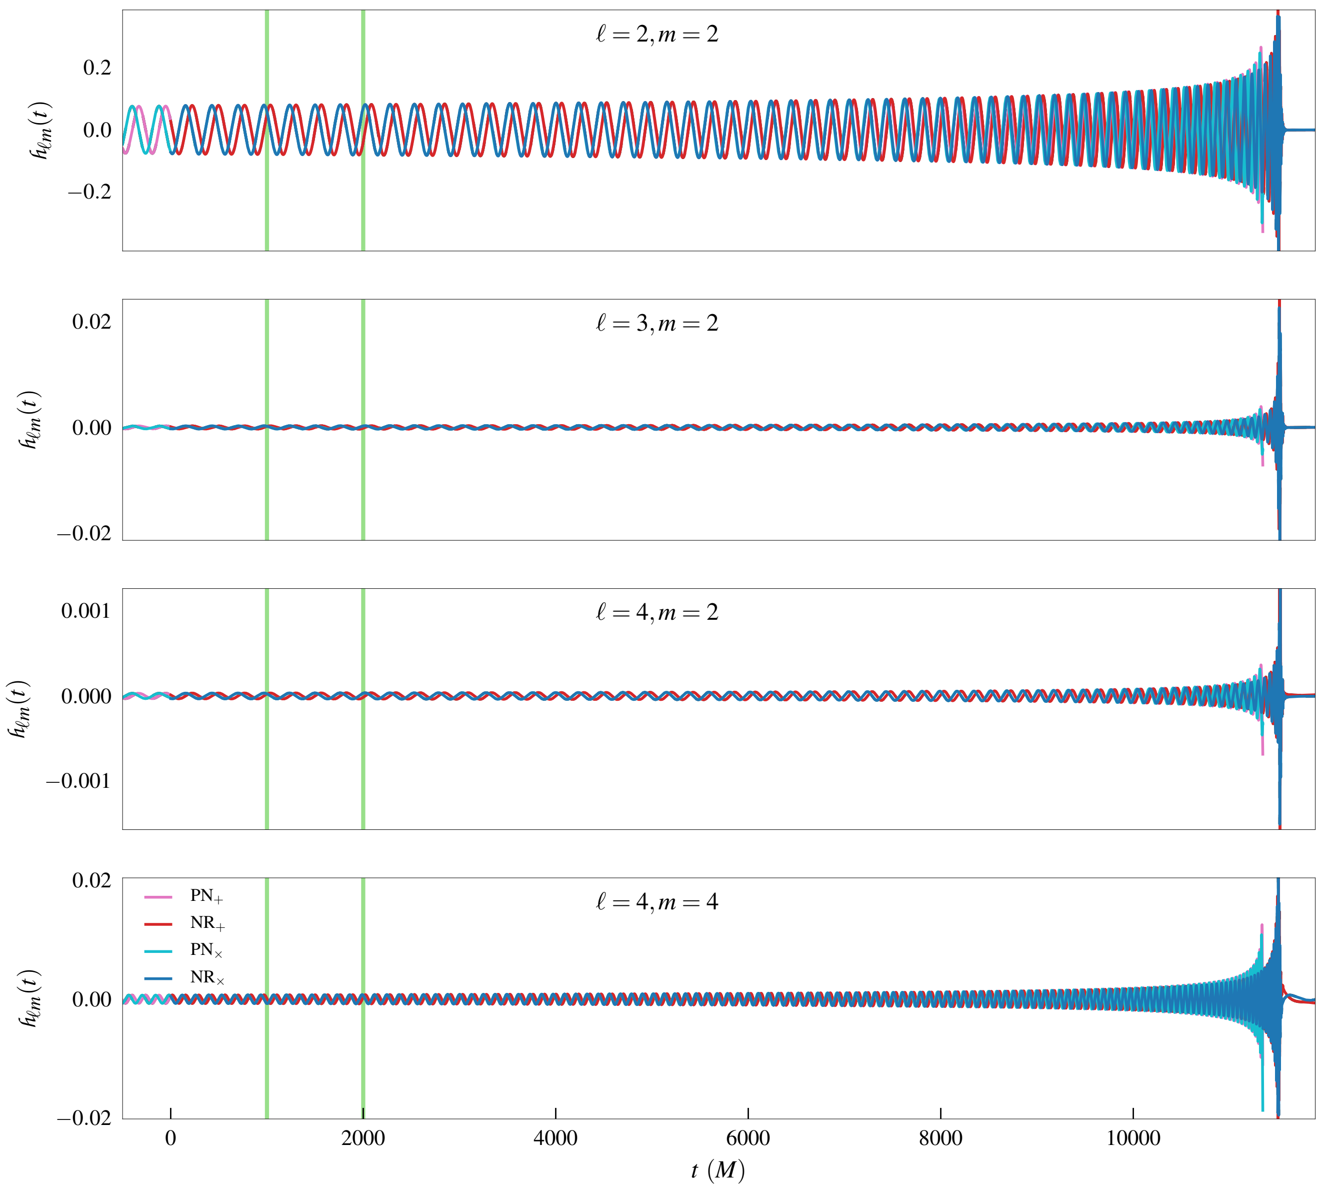

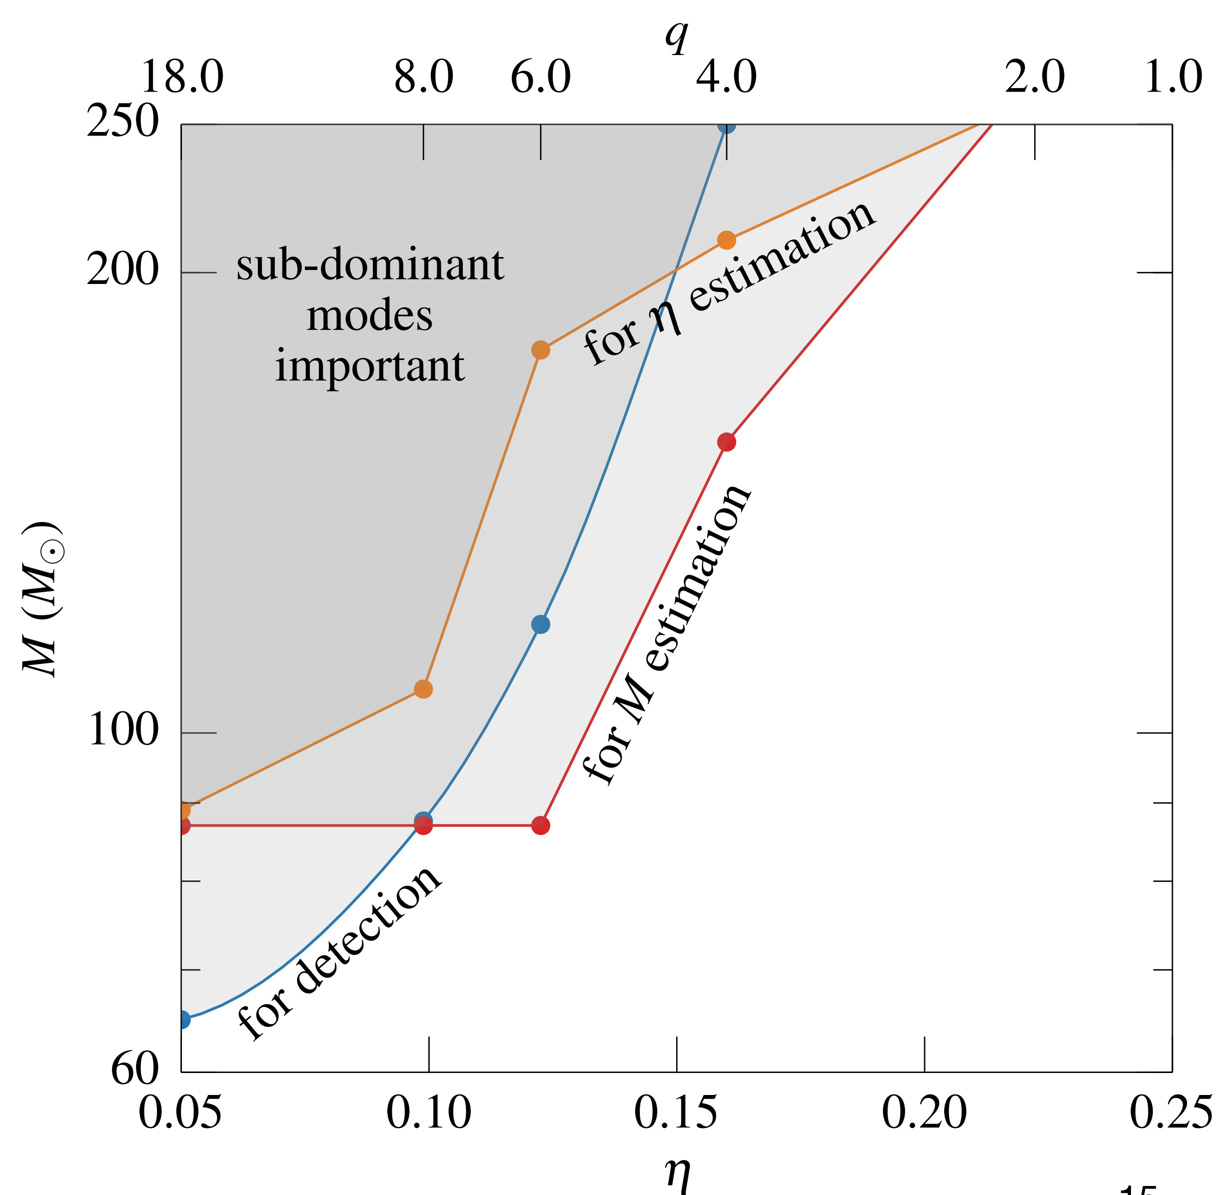

PhenomD - Phenom re-imagined

- Currently used in GW data analysis: PhenomA/B/C/P
  - calibrated to m1/m2=4, moderate spins
    - want |spins| ~ 1, m1/m2 ~ 100
  - good for detection, parameter estimation “toy models”

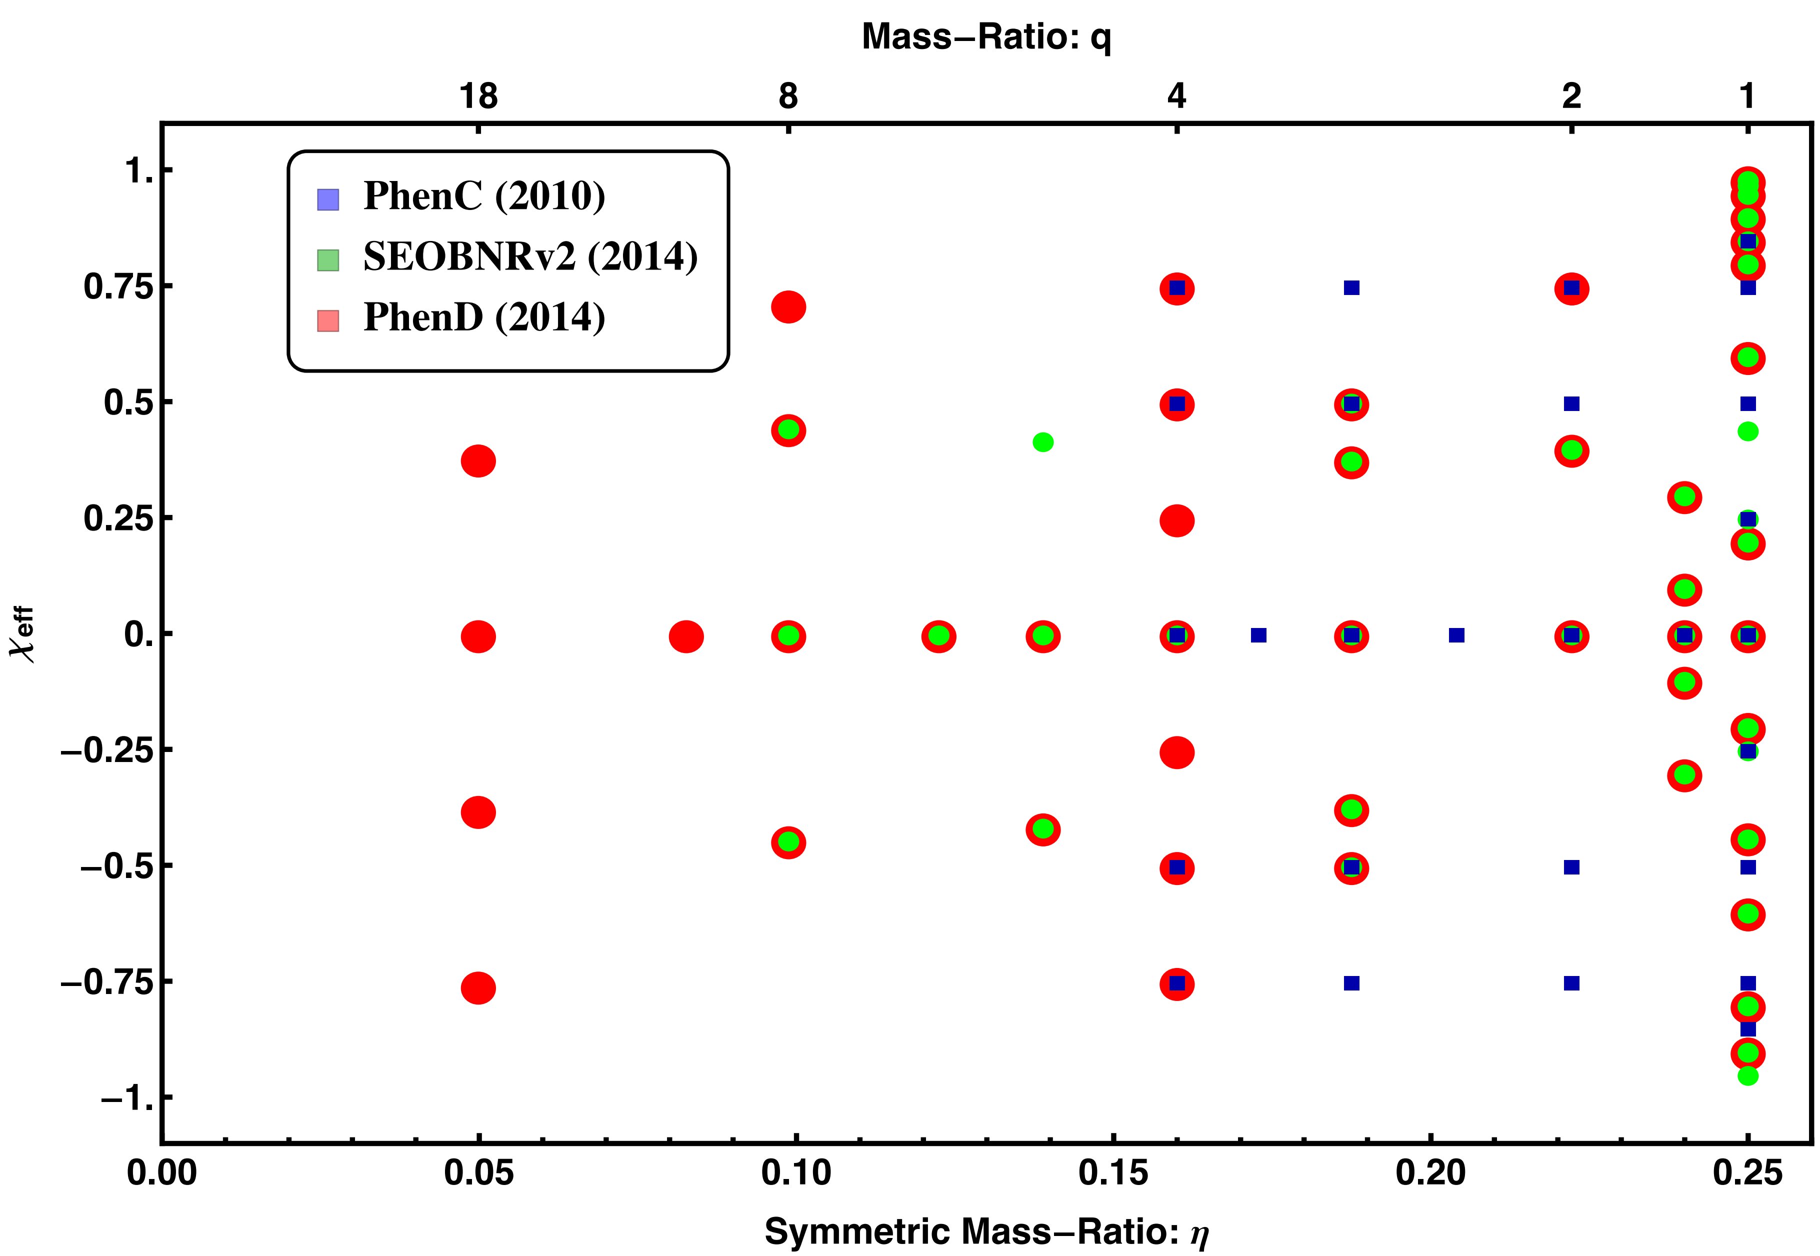

Based on non-precessing equal spin SXS & BAM simulations up to q=18.

Raw data for modelling

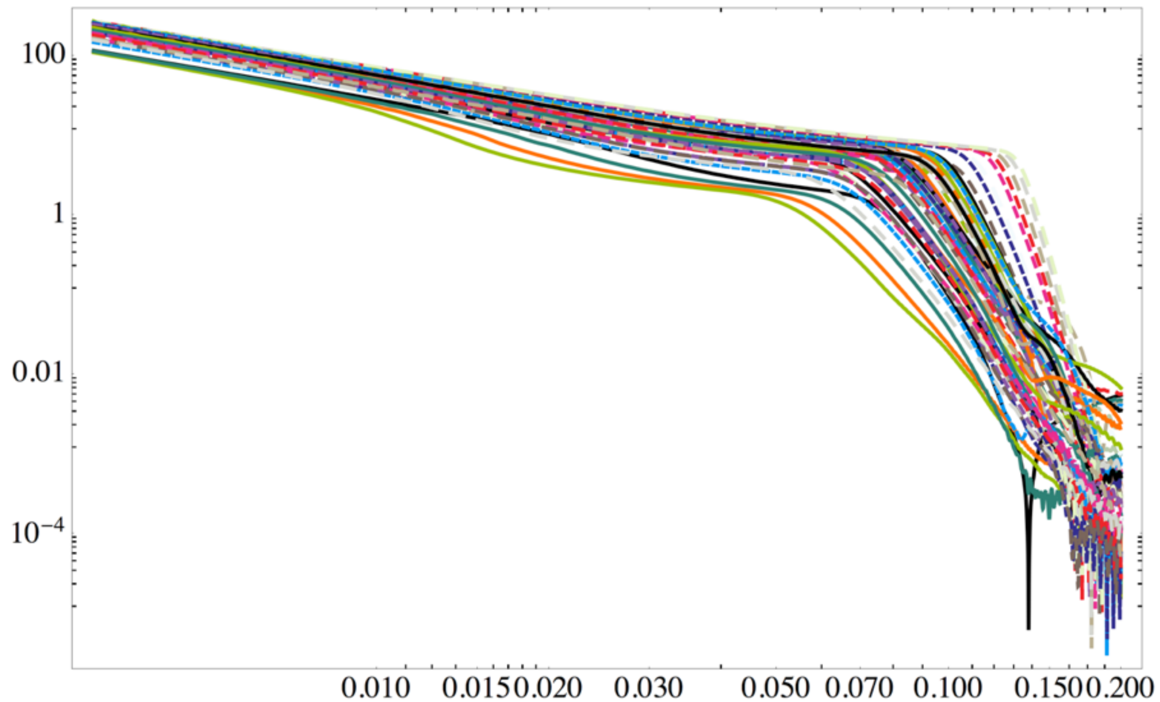

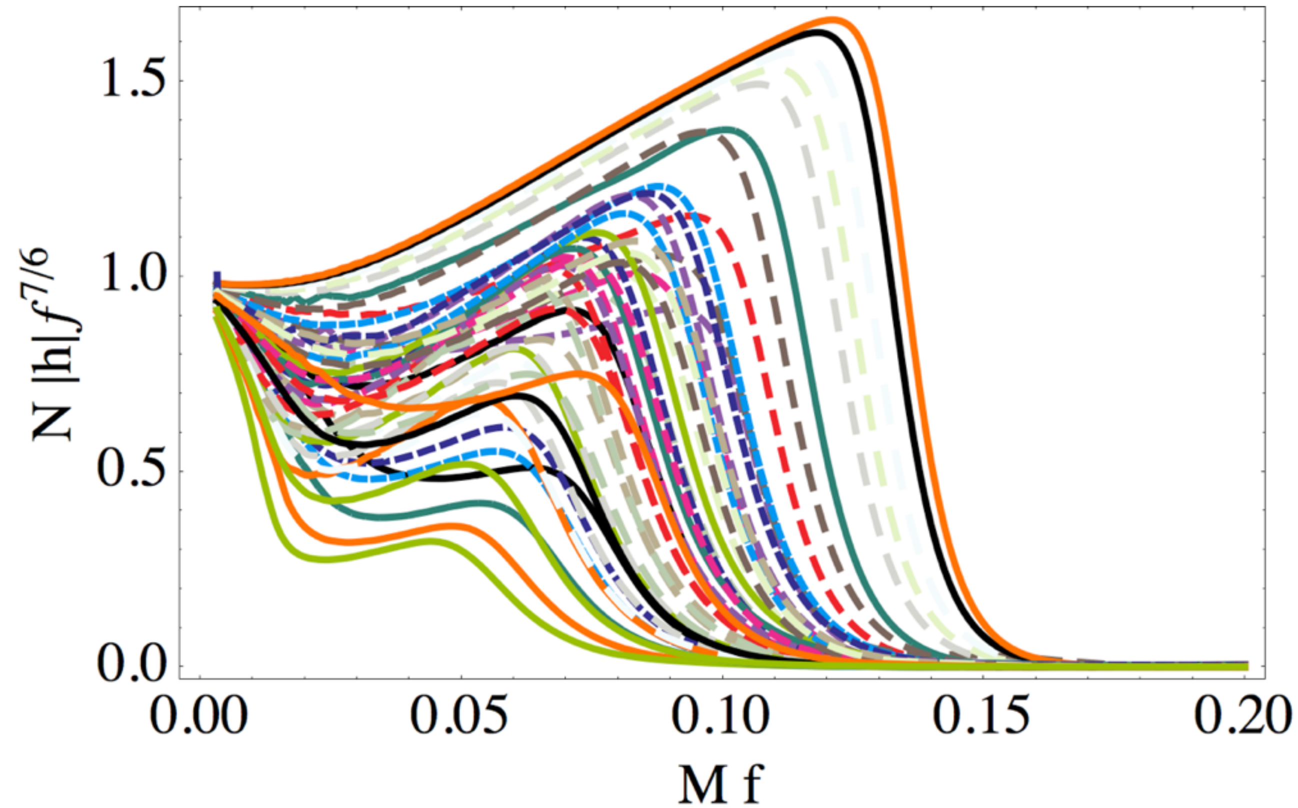

- Dominant spherical harmonic only:
  l=2, |m|=2
- Frequency domain
  amplitude & phase
- Model rescaled amplitude
  & phase derivative

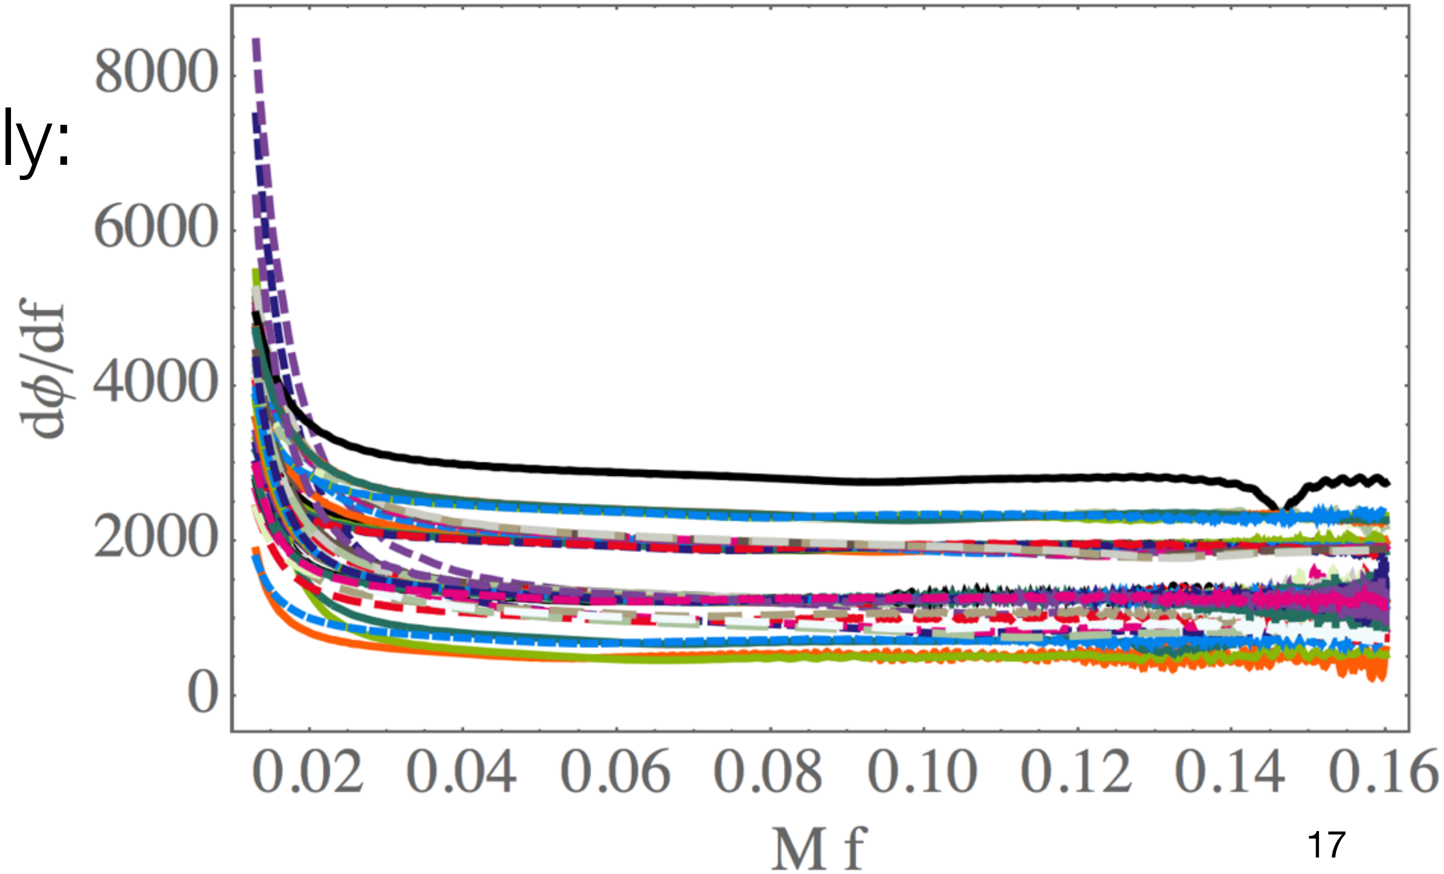

Phenomenological parameter fits

Example: merger/ringdown:

$$h_{MR} = \frac{a e^{-\lambda(f - f_{\rm ring})}}{f_{\rm damp}^2 + (f - f_{\rm ring})^2} \qquad \Phi'_{\rm MR} = \alpha_1 + \alpha_2 f^{-2} + \alpha_3 f^{-1/4} + \frac{\alpha_4 f_{\rm damp}}{f_{\rm damp}^2 + (f - f_{\rm ring})^2}$$

Total: 2 x 8-10 parameters as functions of (η, χ_{eff})

$$\eta = \frac{m1 m2}{(m1 + m2)^2}, \qquad \chi_{eff} = \frac{m_1\chi_1 + m_2\chi_2}{m_1 + m_2} - \frac{76}{113}\frac{1}{2}(\chi_1 + \chi_2)\eta$$

P. Ajith, Phys. Rev. D 84, 084037 (2011)

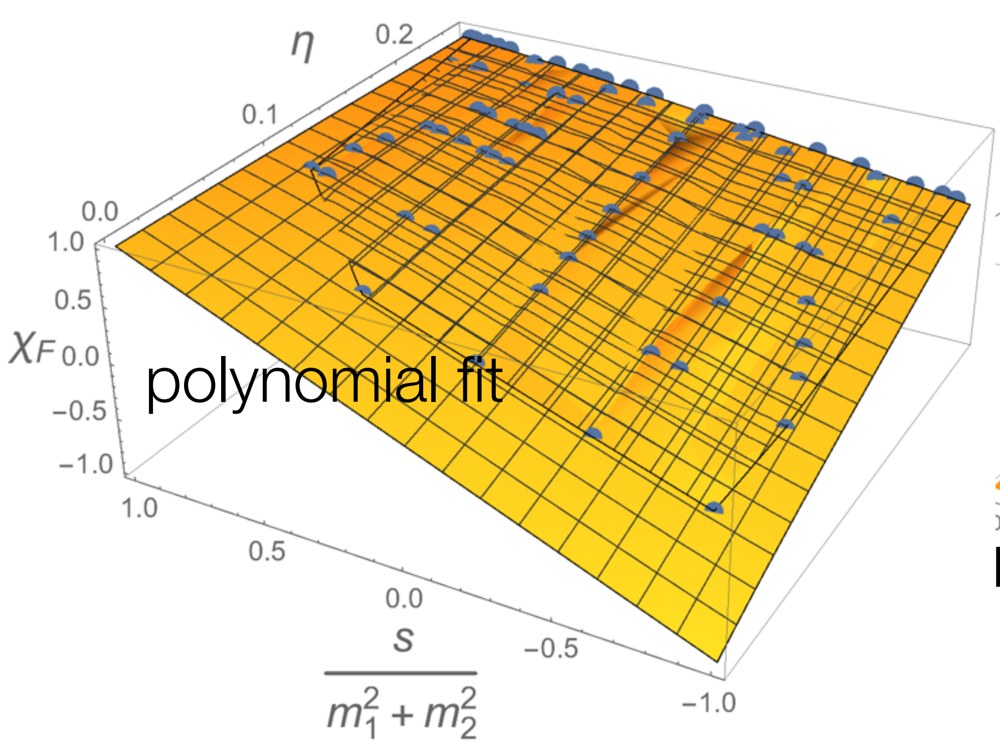

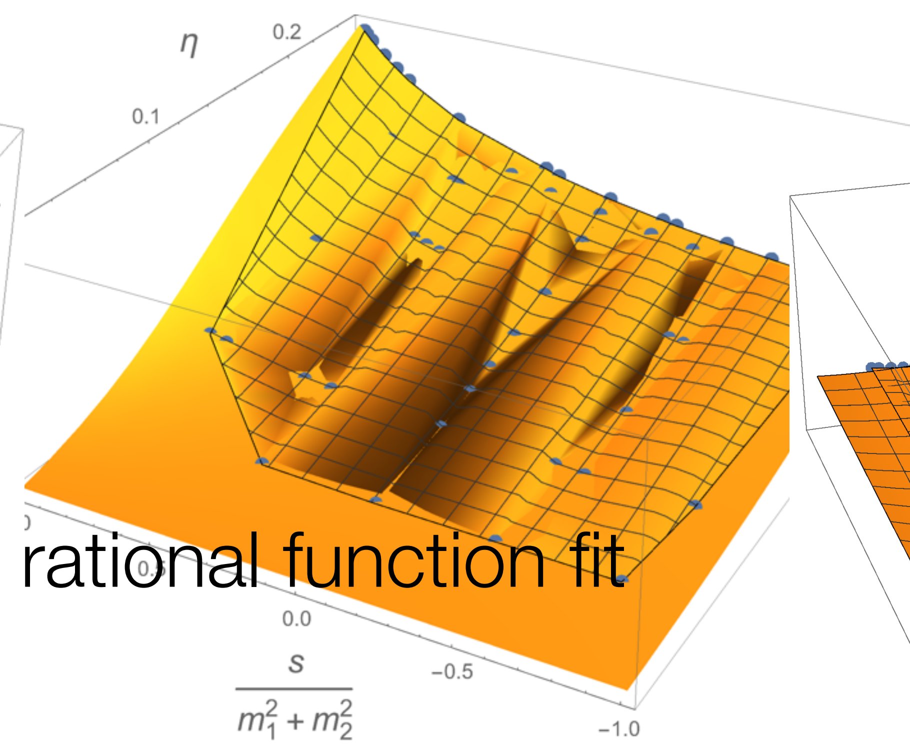

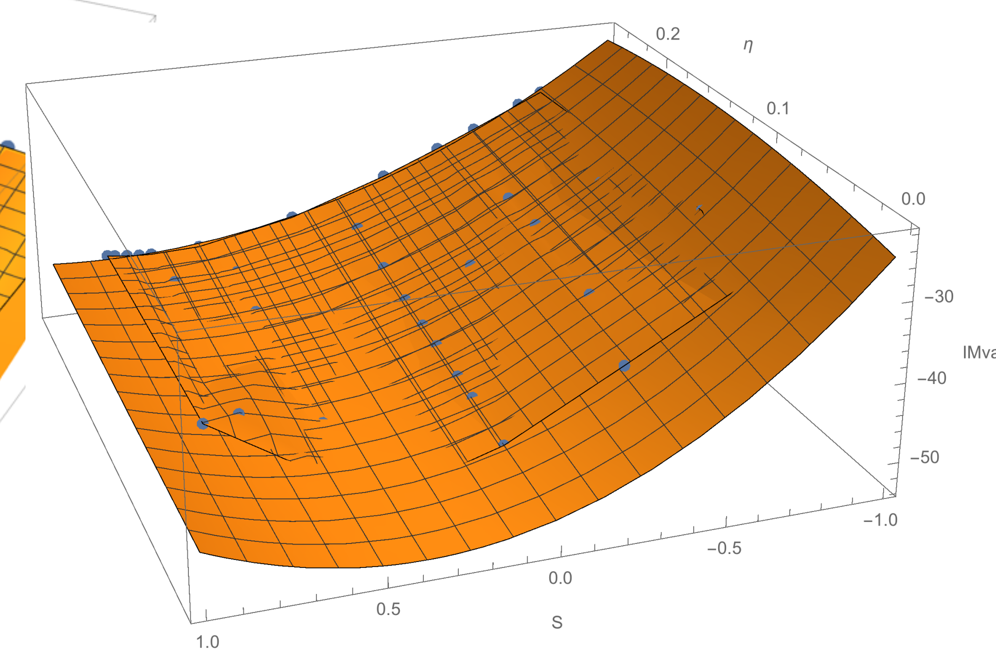

Phenomenological parameter fits

Example: merger/ringdown:

$$h_{MR} = \frac{a e^{-\lambda(f - f_{\rm ring})}}{f_{\rm damp}^2 + (f - f_{\rm ring})^2} \qquad \Phi'_{\rm MR} = \alpha_1 + \alpha_2 f^{-2} + \alpha_3 f^{-1/4} + \frac{\alpha_4 f_{\rm damp}}{f_{\rm damp}^2 + (f - f_{\rm ring})^2}$$

Total: 2 x 8-10 parameters as functions of (η, χ_{eff})

$$\eta = \frac{m1 m2}{(m1 + m2)^2}, \qquad \chi_{eff} = \frac{m_1 \chi_1 + m_2 \chi_2}{m_1 + m_2} - \frac{76}{113} \frac{1}{2} (\chi_1 + \chi_2) \eta$$

P. Ajith, Phys. Rev. D 84, 084037 (2011)

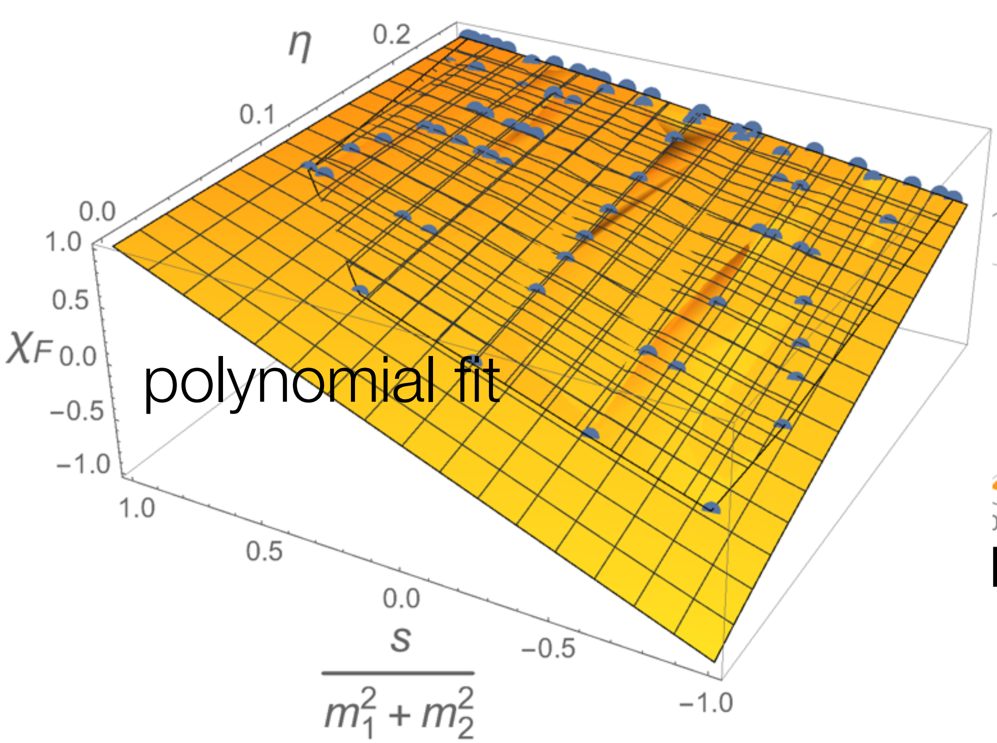

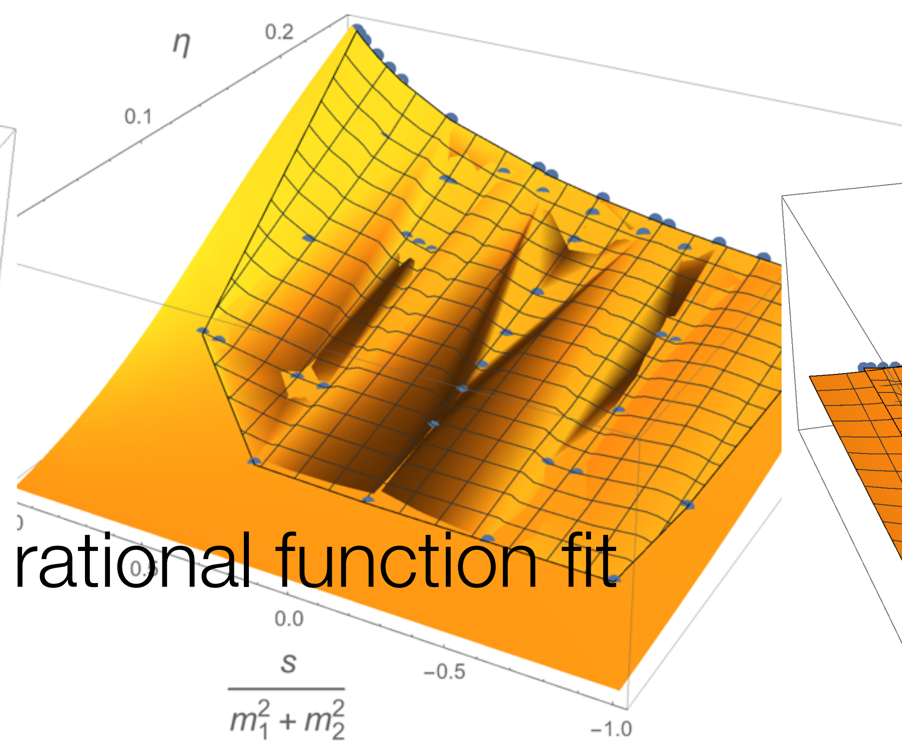

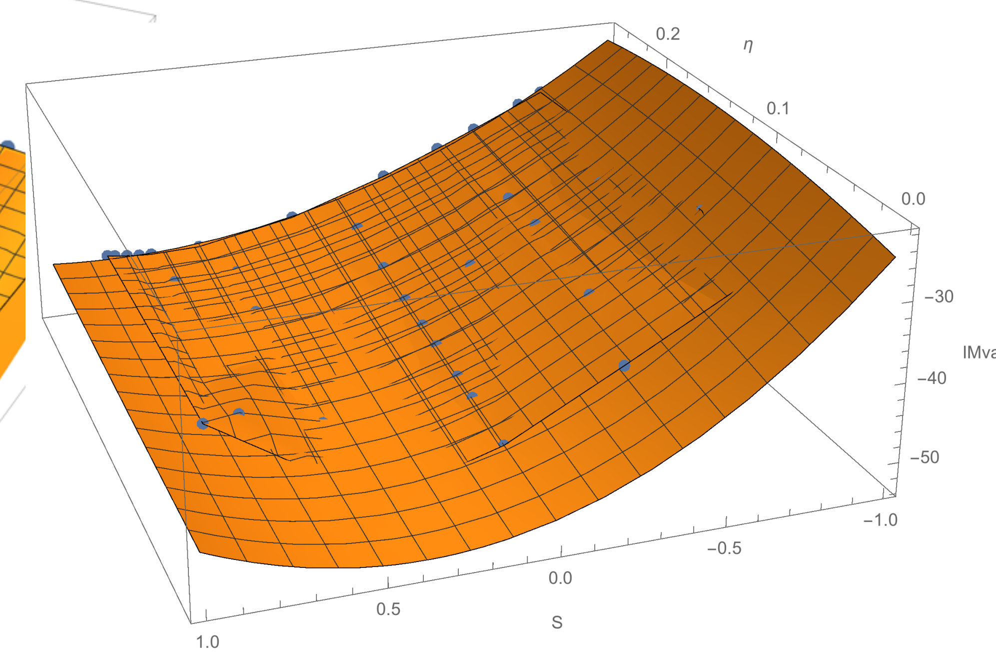

Need more high spin data points

How well does this work: waveforms

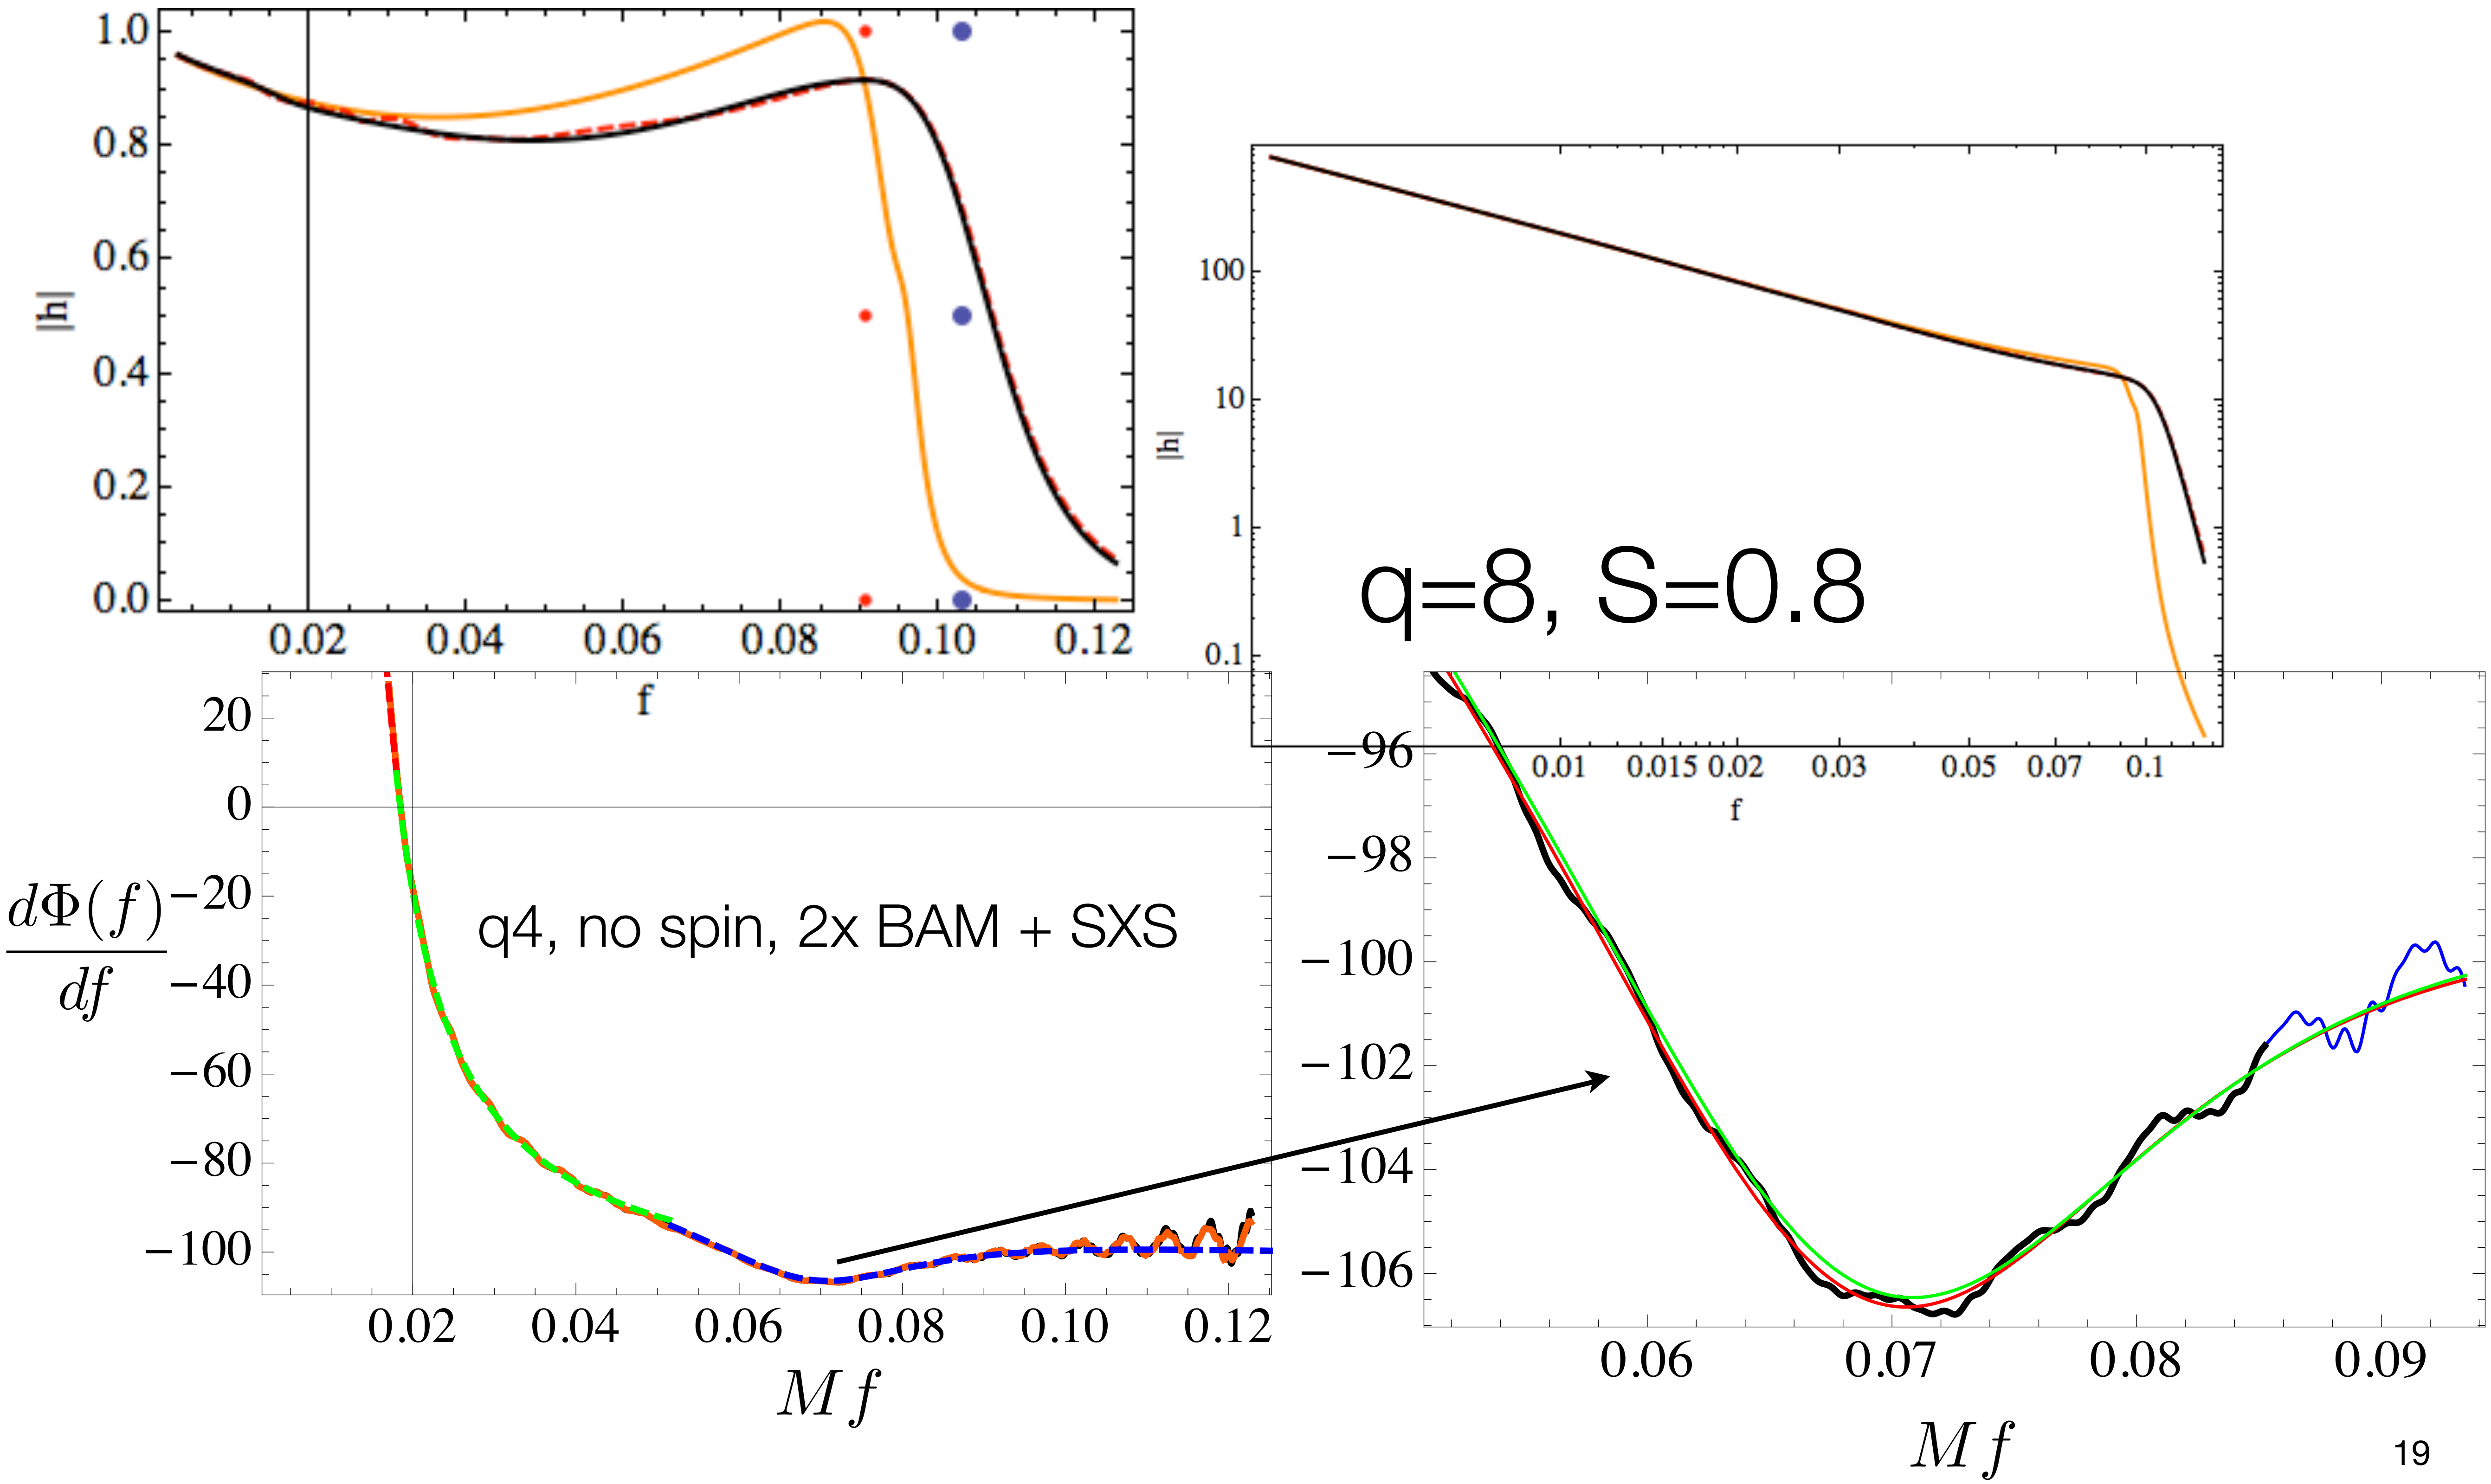

Matches vs. hybrids & between models

Nitz+ Phys. Rev. D 88, 124039 (2013)

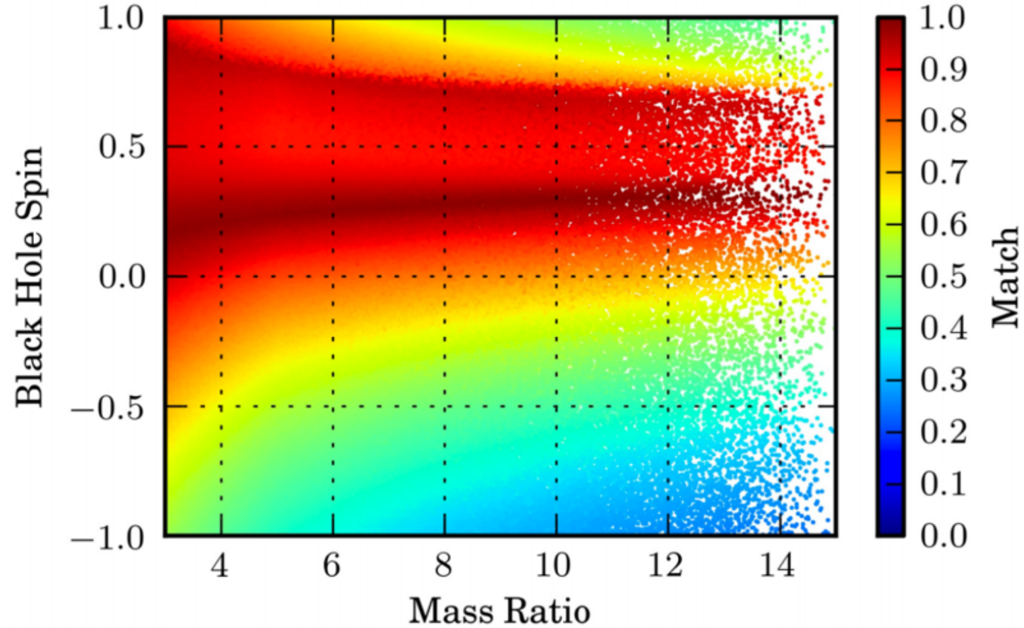

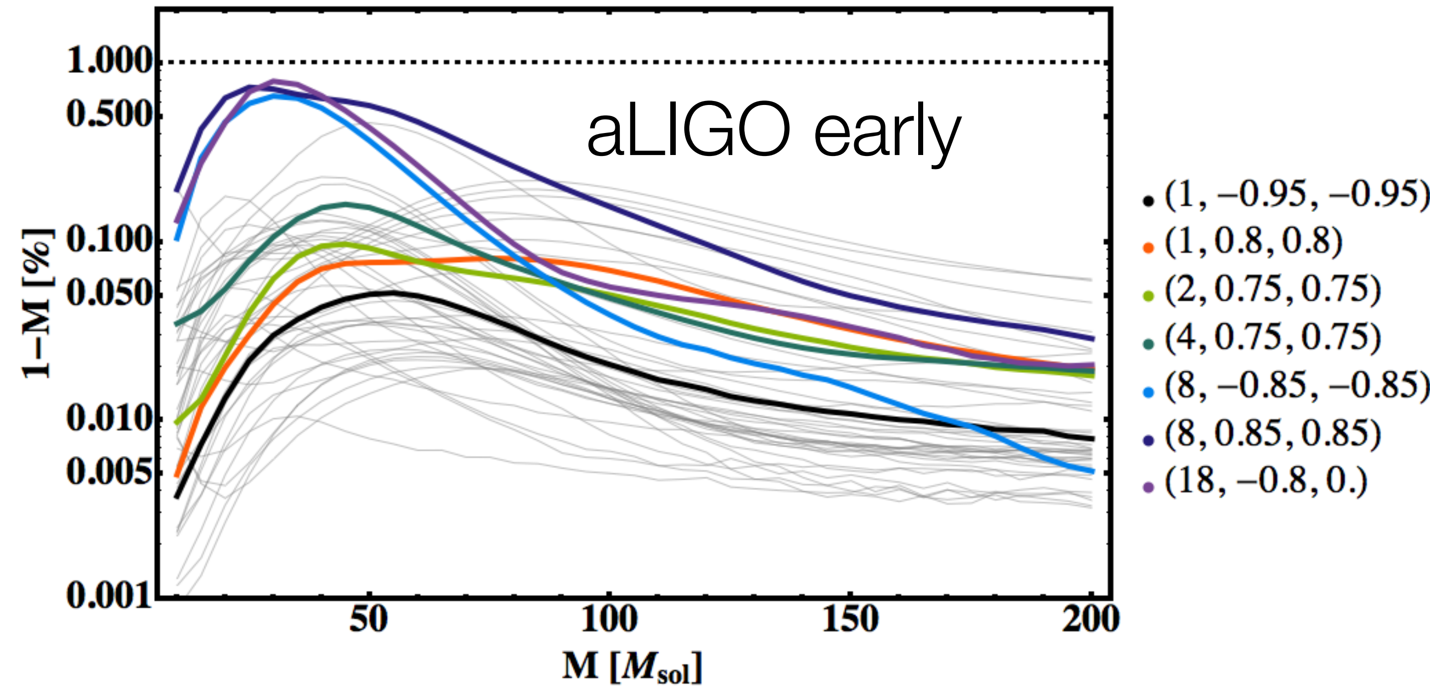

FIG. 16. The match TaylorF2 and TaylorT4 approximants, w
3.5PN spin-orbit and 3.0PN spin-orbit tail corrections includ
a function of the spin of the black hole and the mass ratio
system. The approximants include only the nown spin terms
2.5PN. Matches are calculated using a 30Hz lower frequency
to approximate the sensitivity of the early aLIGO detector. Ir

PhenD vs. SEOBNRv2

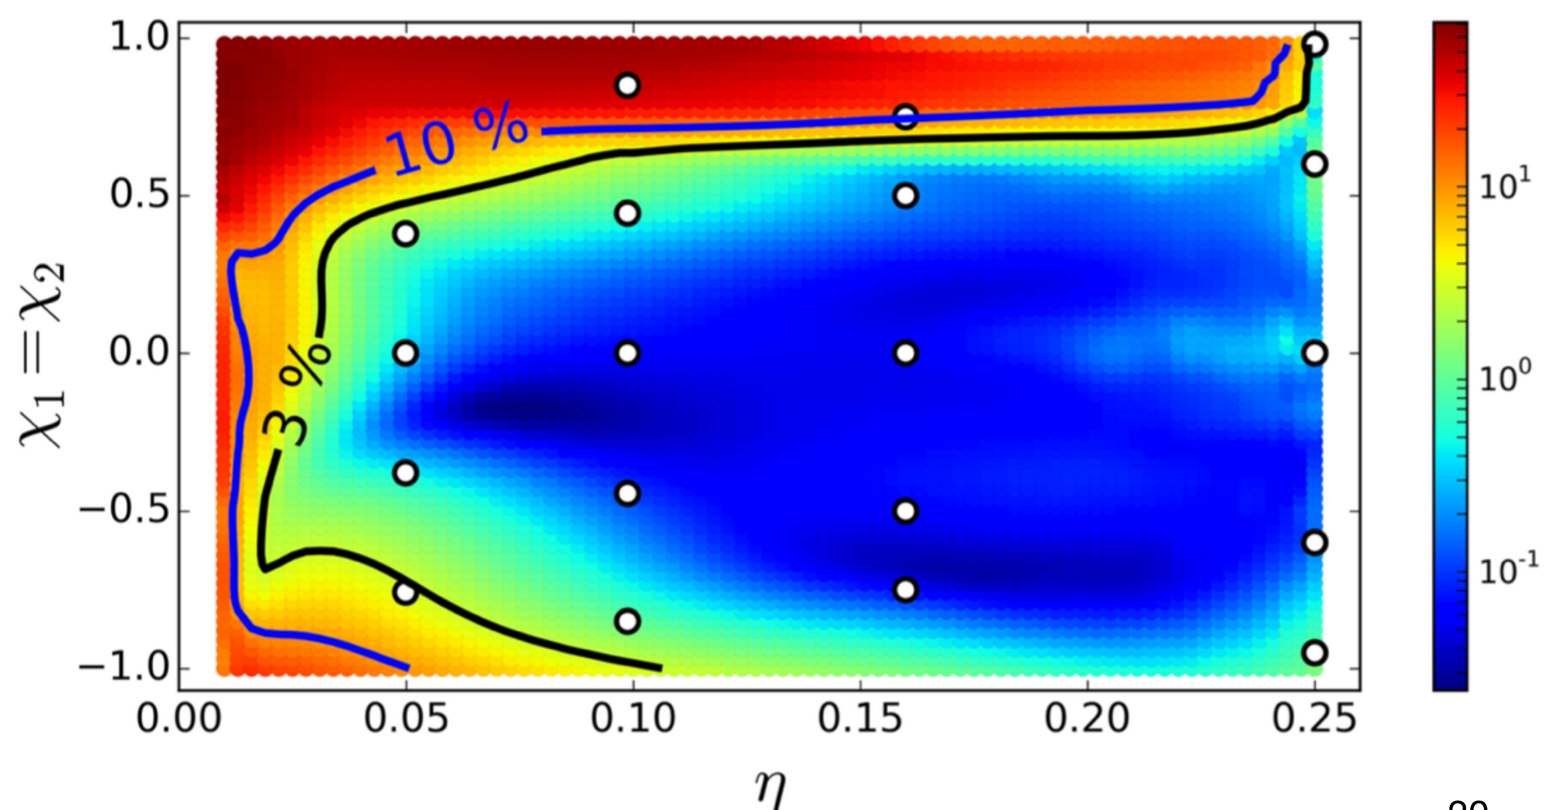

Conclusions

- Systematic study of highest mass ratio spinning BH mergers to date.
- Calibrated most accurate dominant mode non-precessing model to date.
  - Model accuracy drops significantly when BHs have large positive spins, astrophysical likely & “louder”.
  - —> further need for refinement
- Developed a plan to conquer precessing spin space.
  - supported by ERC Consolidating grant to Mark Hannam.
  - Technically ready to run 100s of precessing cases.
- Advanced GW detectors ready for first observing run 09/2015, 6-month run in 2016.
  - Follow up simulations in 2017?

Conclusions

- Systematic study of highest mass ratio spinning BH mergers to date.
- Calibrated most accurate dominant mode non-precessing model to date.
  - Model accuracy drops significantly when BHs have large positive spins, astrophysical likely & “louder”.
  - —> further need for refinement
- Developed a plan to conquer precessing spin space.
  - supported by ERC Consolidating grant to Mark Hannam.
  - Technically ready to run 100s of precessing cases.
- Advanced GW detectors ready for first observing run 09/2015, 6-month run in 2016.
  - Follow up simulations in 2017?