Dipartimento di Fisica "E. Fermi"
Università di Pisa

samsara: Continuous-Time Markov Chain Monte Carlo Sampler for Trans-Dimensional Bayesian Analysis

Gabriele Astorino (PhD, University of Pisa), ***BiCoQ 2026***

Lorenzo Valbusa Dall'Armi (Post-Doc, UniPi), Riccardo Buscicchio (Post-Doc, Unimb),
Joachim Pomper (PhD, UniPi),
Angelo Ricciardone (Prof., UniPi), Walter Del Pozzo (Prof., UniPi)

Trans-dimensional sampling

- Product space [B. P. Carlin and S. Chib, 1995]
  - Search for all combinations of models in the model space
  - Bayes Factor statistics for finding the best combination of models

- RJ-MCMC [P. J. Green, 1995]
  - MCMC extension to explore model space with between-model proposals relying on auxiliary variables $u_l \sim g(\cdot),\ u_k \sim g'(\cdot)$ to match dimensions
  - Acceptance rule similar to the MH

$$\alpha = \min\left\{1, \frac{p(\theta_l|l,d)g'(u_l)}{p(\theta_k|k,d)g(u_k)}|J|\right\}$$

How samsara algorithm works

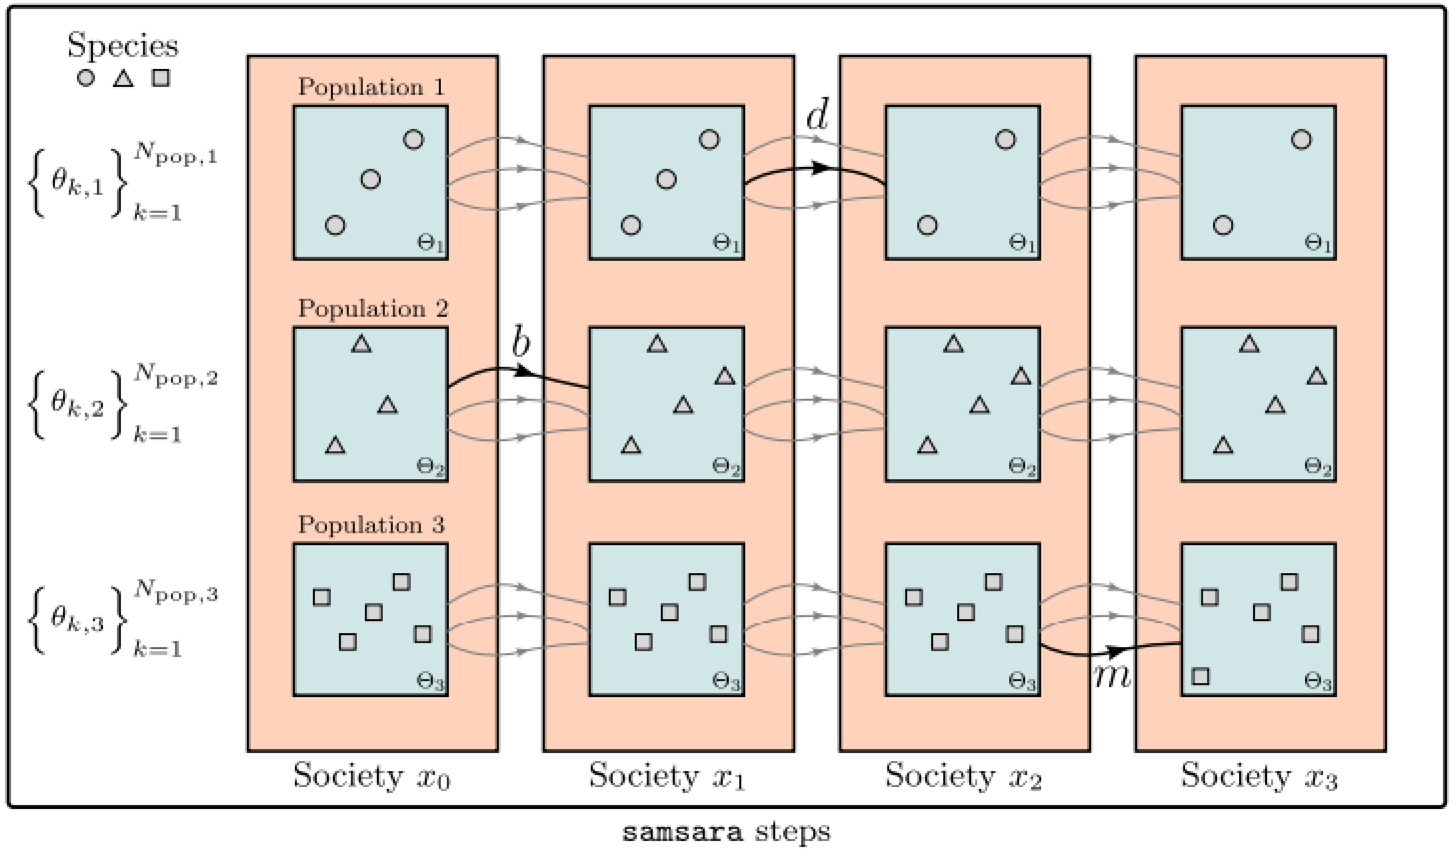

At each step:

- 3 independent Poisson Processes: birth, death, mutation
- Rates contains information for all the possible transitions; computed by imposing detail balance
- Poisson dynamics decide which species will evolve and how (birth, death, mutation of an individual)

How samsara algorithm works

state with one
individual less

$$R_{j,d,\alpha}(y) = \min\left\{1, \frac{\mathcal{Z}(N_{\mathrm{pop},\alpha}, \theta_{j,\alpha})}{N_{\mathrm{pop},\alpha}} \frac{p(x_\alpha|D)}{p(y_\alpha|D)} h(\theta_{j,\alpha}|x_\alpha)\right\},$$

$$R_{b,\alpha}(y) = \min\left\{1, \frac{N_{\mathrm{pop},\alpha}+1}{\mathcal{Z}(N_{\mathrm{pop},\alpha}, \theta_{j,\alpha})} \frac{p(z_\alpha|D)}{p(y_\alpha|D)} \frac{1}{h(\theta_{j,\alpha}|y_\alpha)}\right\}.$$

state with one
individual more

proposal for the birth

Analytic test case

Goal is to asses samsara samples correcly a trans-dimensional distribution from which we know how to sample

$$p(N,\theta_i) = \frac{\lambda^N e^{-\lambda}}{N!}\prod_{i=1}^{N}\sum_{j=1}^{3}\pi_j \mathcal{N}\left(\theta_i \,\middle|\, \mu_j, \Sigma_j^{2\times 2}\right)$$

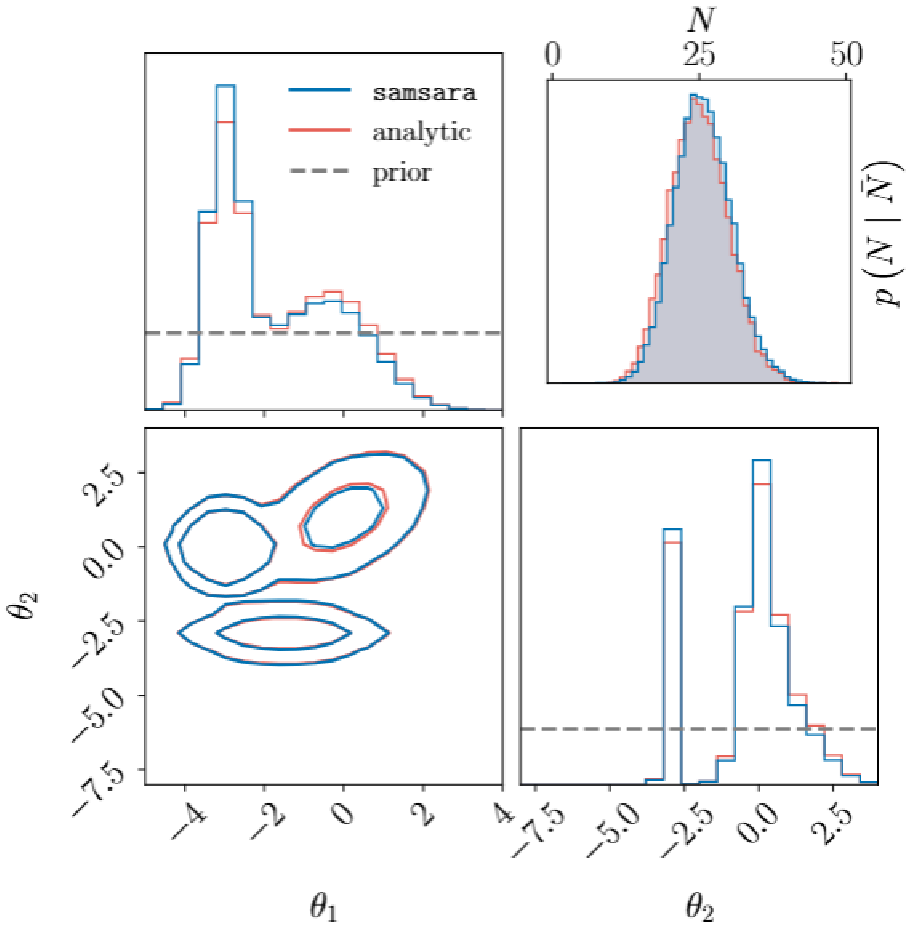

GMM with unknown number of components

Observations are N_s i.i.d. samples x from $f(x) = \sum_{i=1}^{N} \pi_i \mathcal{N}(x|\mu_i, \sigma_i)$

We want to infer: $N,\ \pi,\ \mu,\ \sigma$

$$\begin{aligned} p(\pi, \mu, \sigma|x) &\propto p(\pi, \mu, \sigma) p(x|\pi, \mu, \sigma) \\ &\propto p(\pi, \mu, \sigma) \prod_{j=1}^{Ns} f(x_j) \end{aligned}$$

GMM with unknown number of components

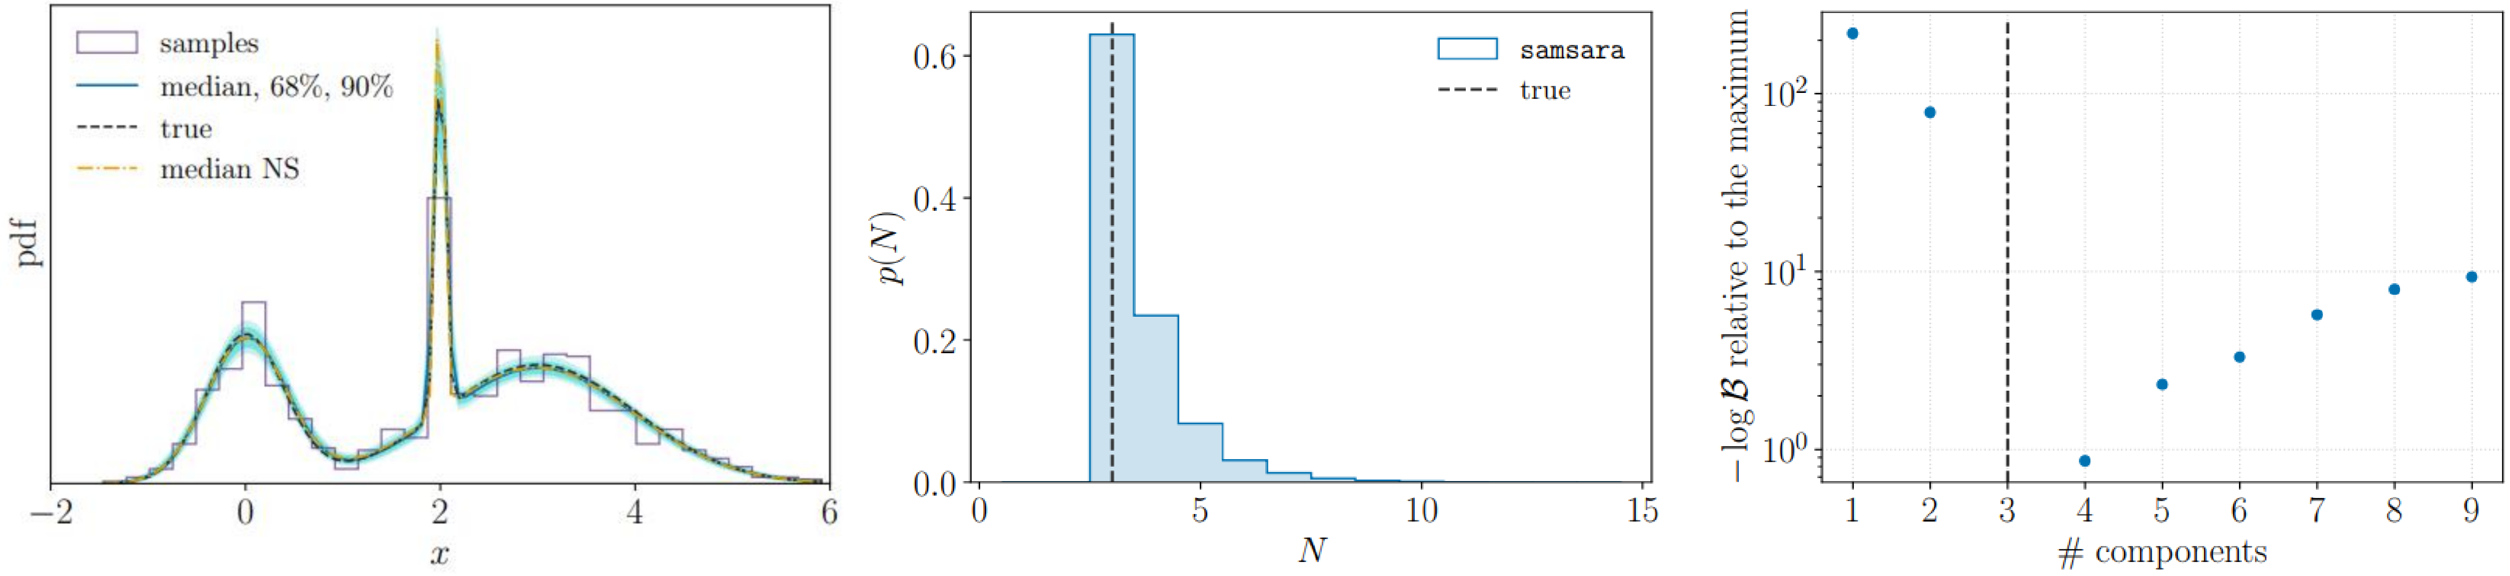

Sine waves + lorentzians

Overlapping signals in a noisy background

$$d(t) = n(t) + \sum_{i=1}^{N_s} A\cos(2\pi f t + \pi f' t^2 + \phi) + \sum_{j=1}^{N_l} \frac{A}{1 + \big((t - t_0)/w\big)^2}$$

- zero-noise test for proof of principle
- test with noise (but without inferring the noise itself)

Sine waves + lorentzians

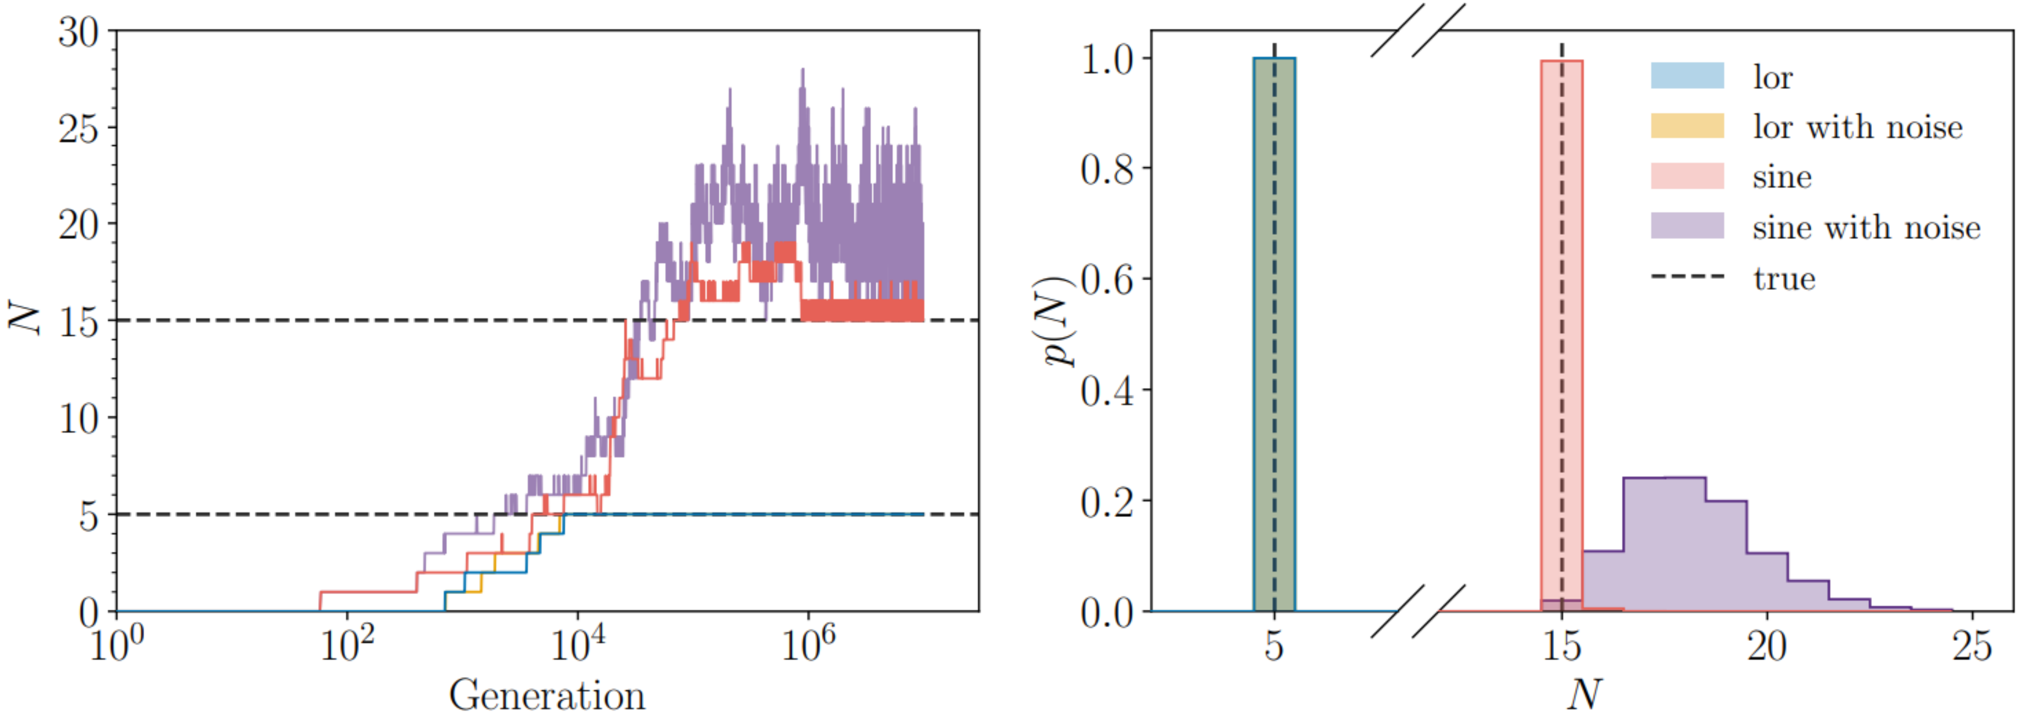

Sine waves + lorentzians

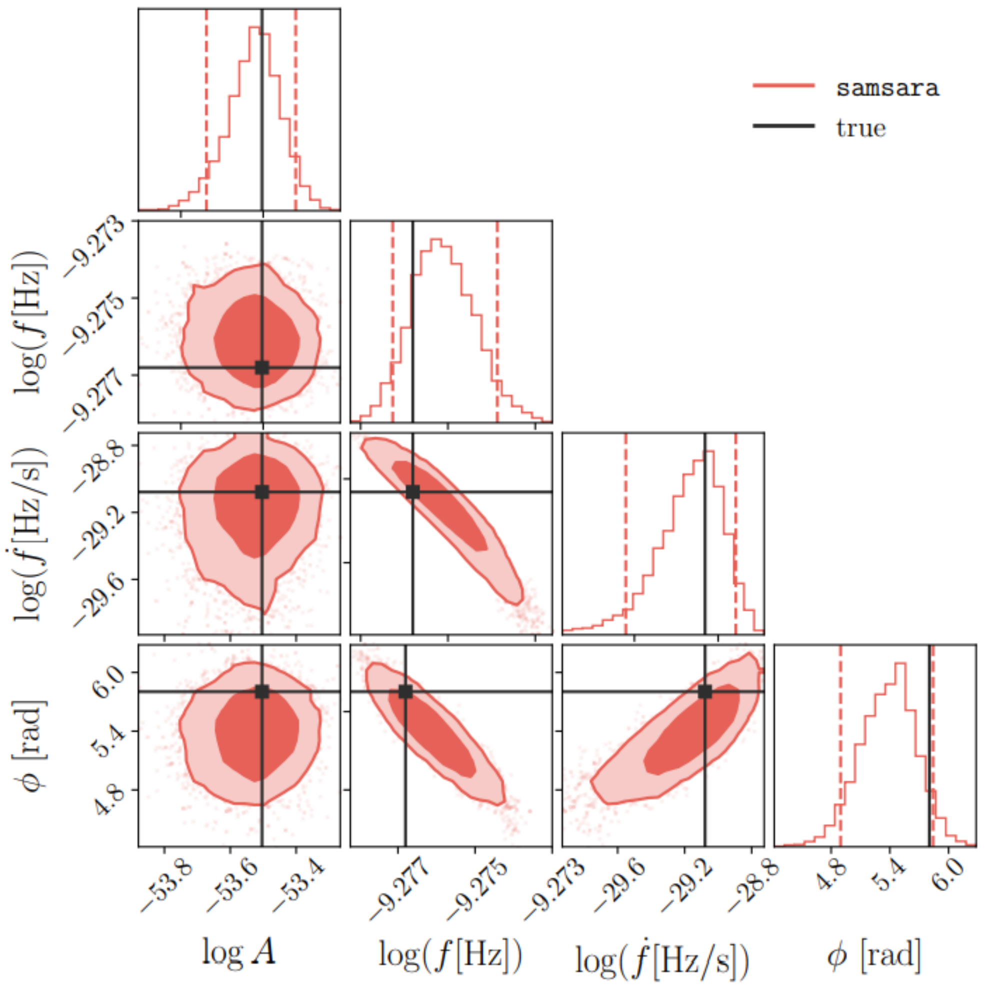

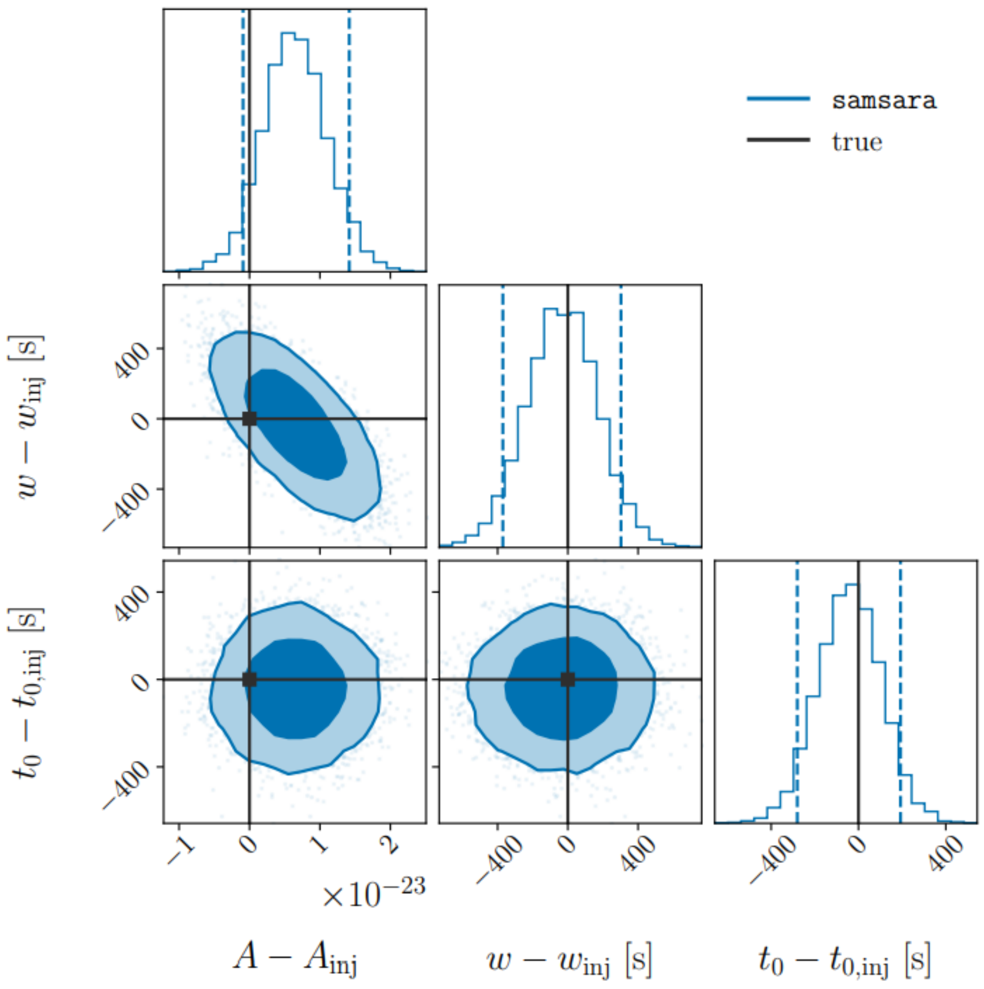

Sine waves + lorentzians

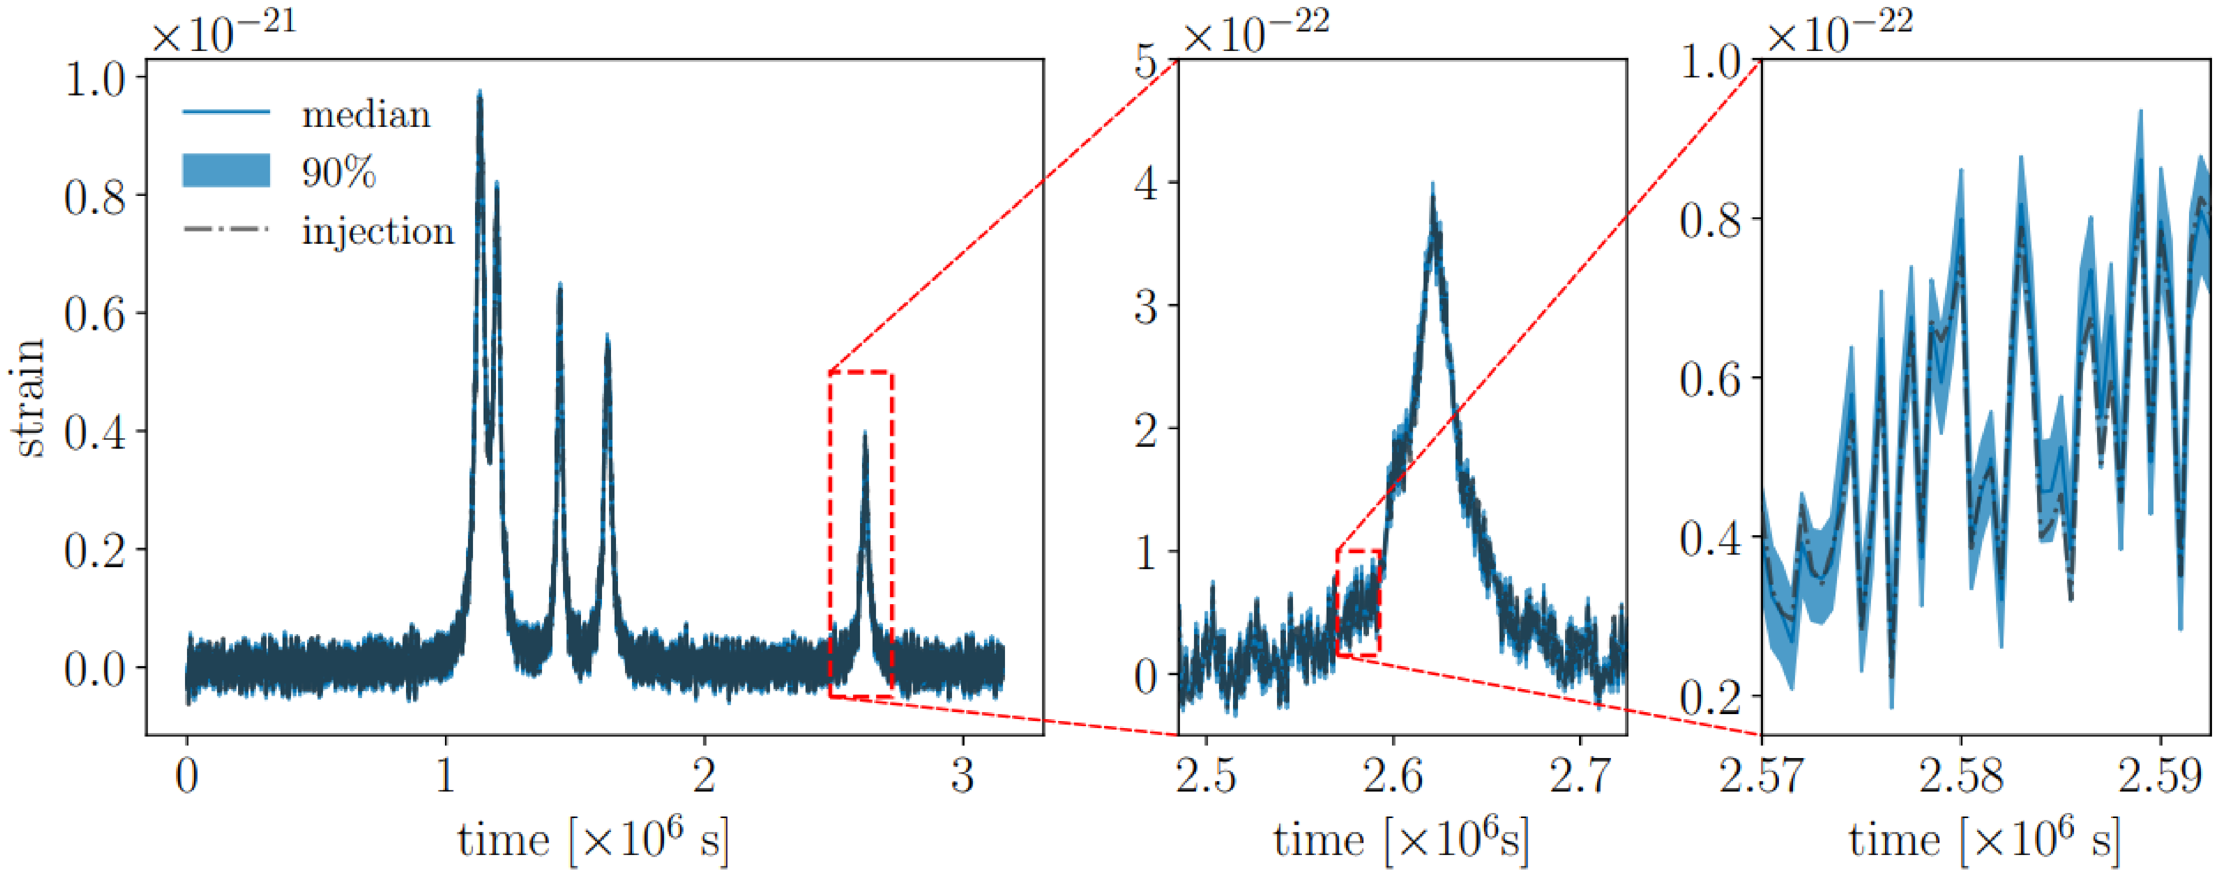

Preliminary GF application

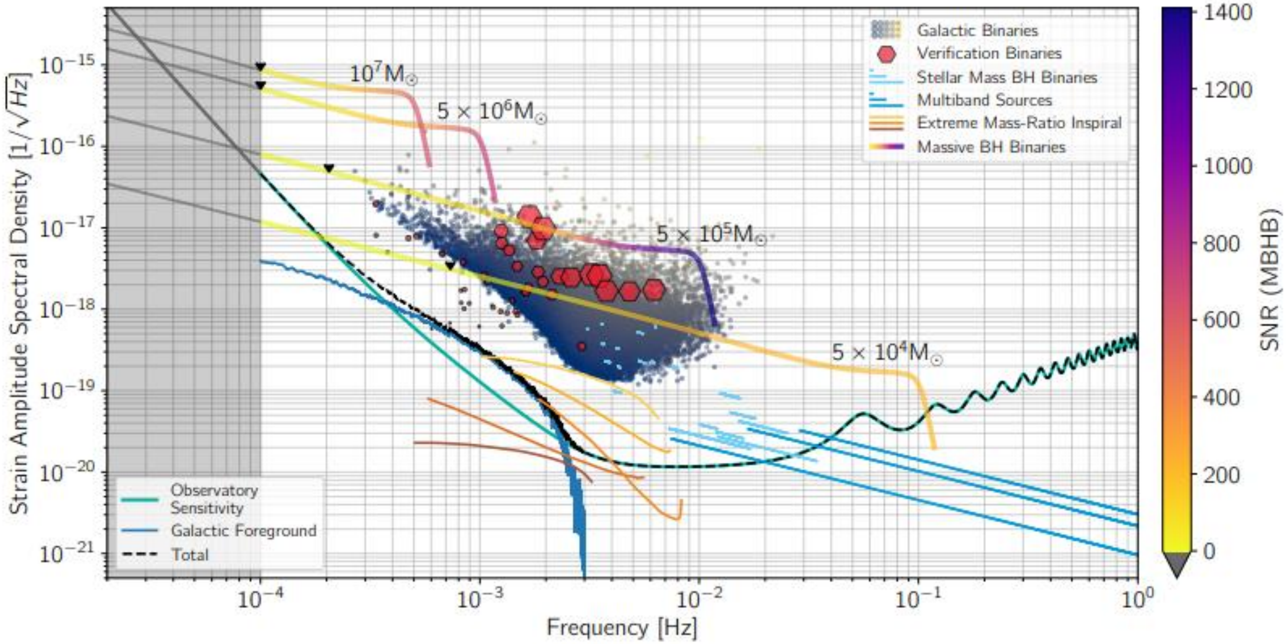

From [Colpi et al., 2024]

Preliminary GF application

- $T_{obs} = 1yr$, zero noise

- 100 DWD spaced by 3 frequency bins

- SNR $\in [5, 25]$

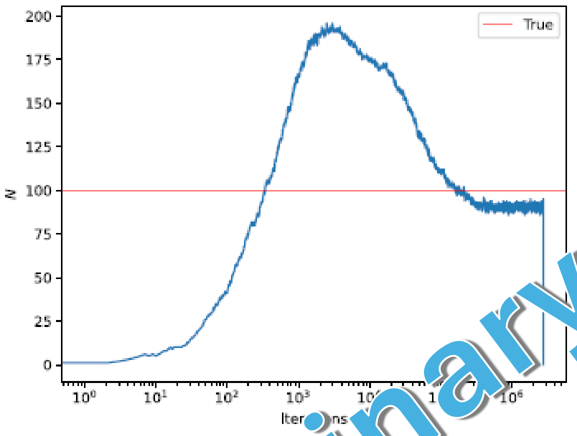

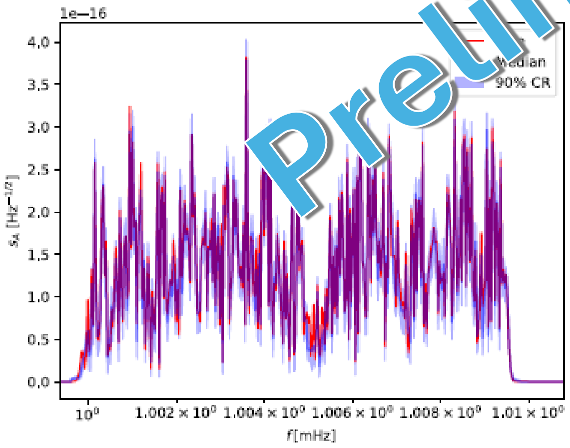

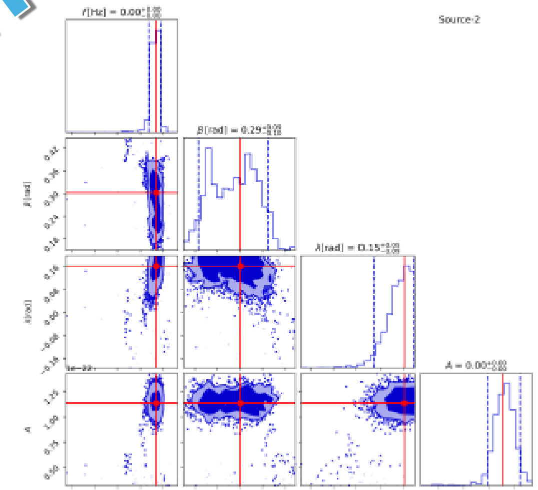

Preliminary

Conclusions

- CTMCMC framework for trans-dimensional sampling based on Poisson Processes dynamics
- Alternative to traditional methods: product space and RJMCMC
- Validated with test cases
- Preliminary GF application with promising results

Backup slides

Detailed balance equations

Detailed balance equations [M. Stephens, 2000]

$$\int dP_k(\mathrm{y}_\alpha|D)\, R_{b,\alpha}(\mathrm{y}_\alpha) I\{y \in F\} = \int dP_{k+1}(z_\alpha|D)\, R_{d,\alpha}(z_\alpha) K_d(z_\alpha; F)$$

$$\int dP_{k+1}(\mathrm{z}|D)\, R_{d,\alpha}(\mathrm{z}_\alpha) I\{z \in G\} = \int dP_k(y_\alpha|D)\, R_{b,\alpha}(y_\alpha) K_b(y_\alpha; G)$$

- $K_{b/d}$ are the probabilities that state y_α/z_α transits to a point in a subspace of +1/-1 dimension.
- Birth are done by sampling the parameter of the new individual from a certain proposal $h(\cdot)$

Transition kernels

$$K_b(y;G) = \int d\theta\ h(\theta|y)\frac{R_b(y,\theta)}{R_b(y)}I(y \cup \theta \in G)$$

$$K_d(z;F) = \sum_{\theta_j \in z\ :\ z\setminus\theta_j \in F} d(\theta_j, z)\frac{R_{j,d,\alpha}(\theta_j, z)}{R_{d,\alpha}(z)}$$

h proposal distribution, in our case prior and gaussian.

Correlation

In principle not a trivial task since samples have different dimension.

Correlation computed via scalar function ρ mapping trans-dimensional sample to real numbers. We use the posterior itself

Given a certain lag δ

$$ACF(\rho,\delta) = \frac{N_s}{N_s - \delta} \frac{\sum_{i=1}^{N_s-\delta} \left(\rho^{(i)} - \bar{\rho}\right)\left(\rho^{(i+\delta)} - \bar{\rho}\right)}{\sum_{i=1}^{N_s} (\rho^{(i)} - \bar{\rho})^2}$$

Convergence

Given C independent chains, is based on computing distance between each sample and a certain subset of samples left fixed.

The idea is that if this distance tends to have a consistent statistic, then the chains have converged.

Actual test: given the ensemble of distances from the fixed points, perform Gelman-Rubin test convergence test