Multimarginal Schrödinger Bridge for Probabilistic Learning of Hardware Resource Usage by Control Software

Abhishek Halder

Department of Aerospace Engineering, Iowa State University
Translational AI Center, Iowa State University
Department of Applied Mathematics, UC Santa Cruz

Joint work with G. Bondar (UCSC), R. Gifford (UPenn), L. Phan (UPenn)

CPS Frontier Workshop, UC Berkeley, November 02, 2023

Motivation

HW resource in CPS are time-varying, stochastic and dynamically correlated

e.g., last-level shared cache (LLC), memory bandwidth, processor availability

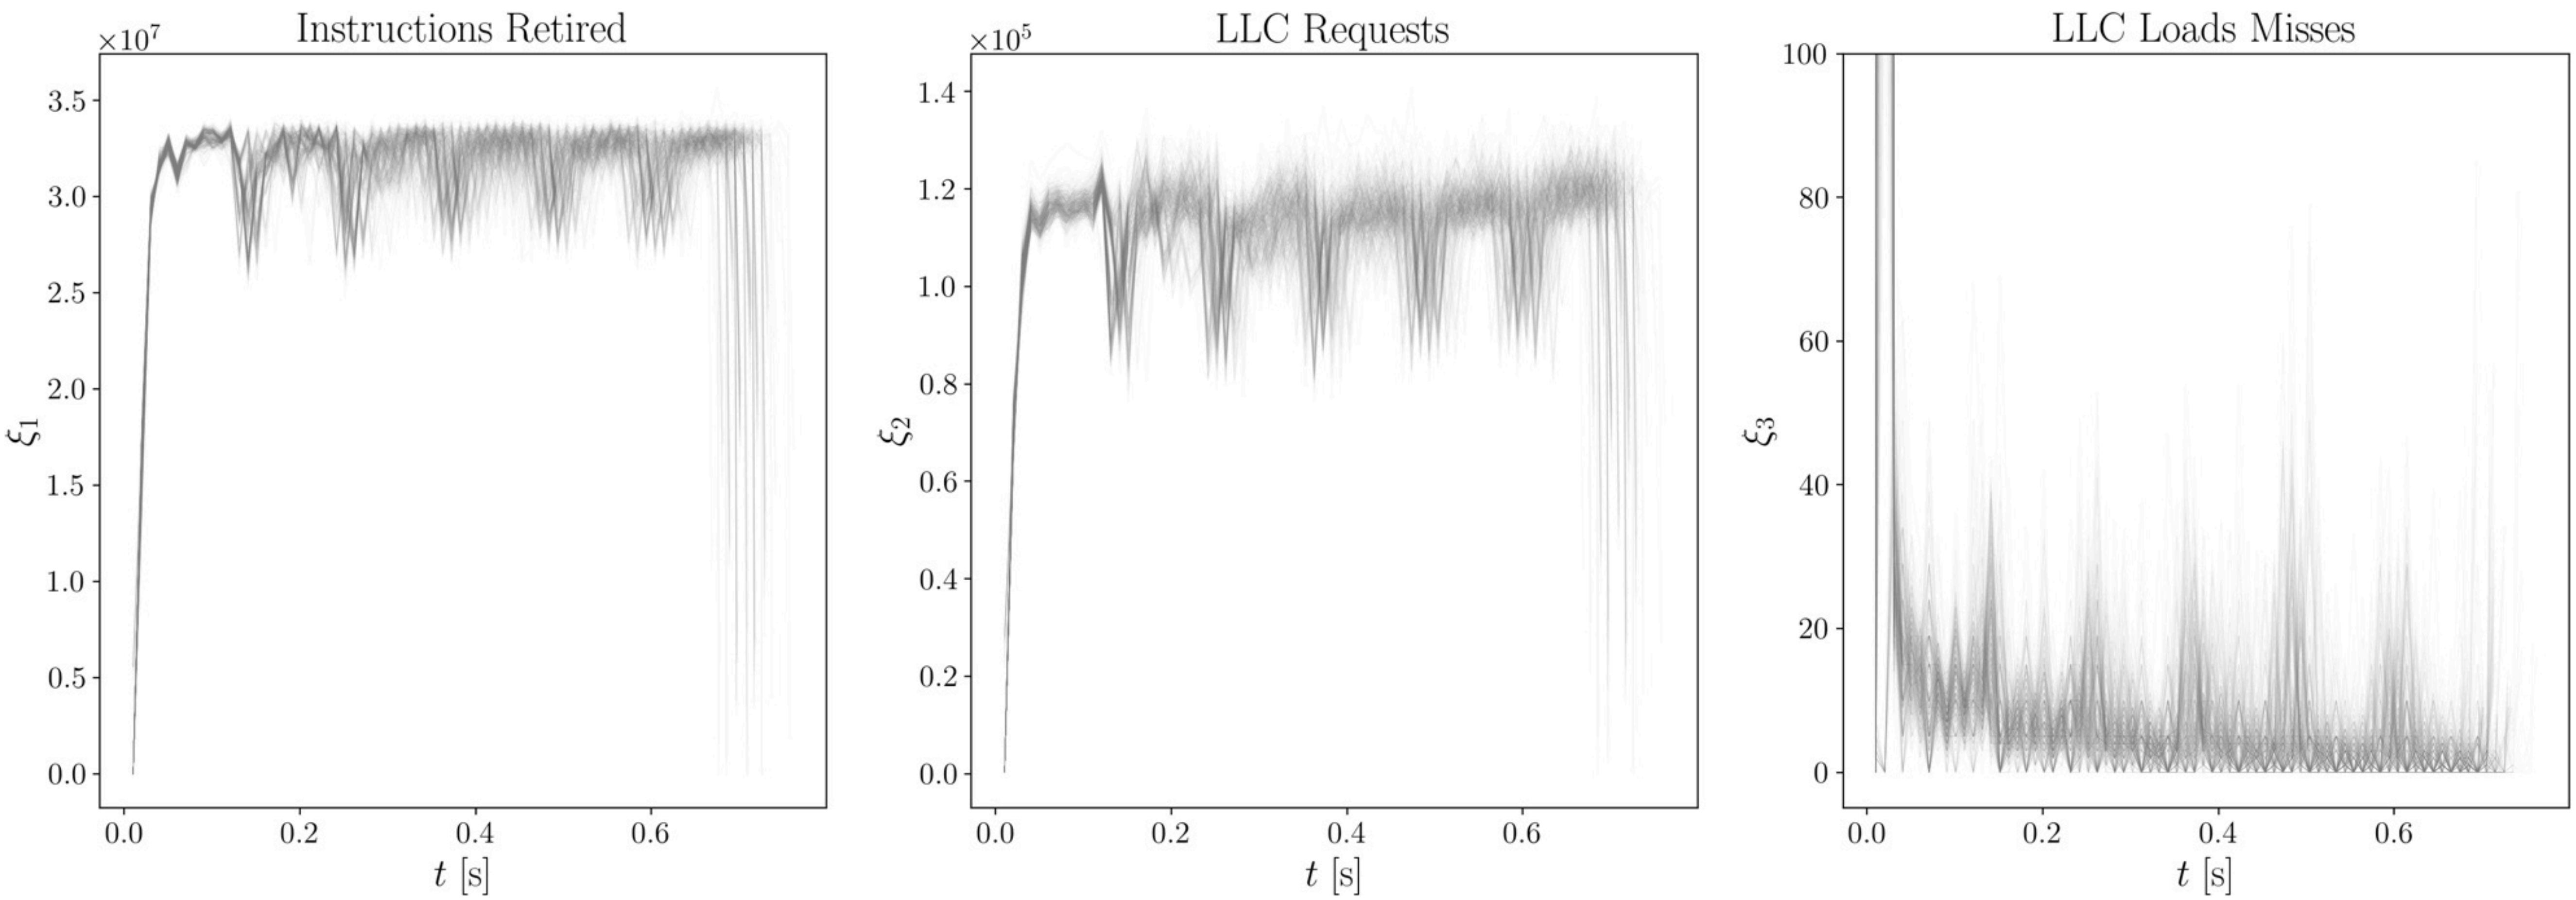

Motivation

HW resource in CPS are time-varying, stochastic and dynamically correlated

e.g., last-level shared cache (LLC), memory bandwidth, processor availability

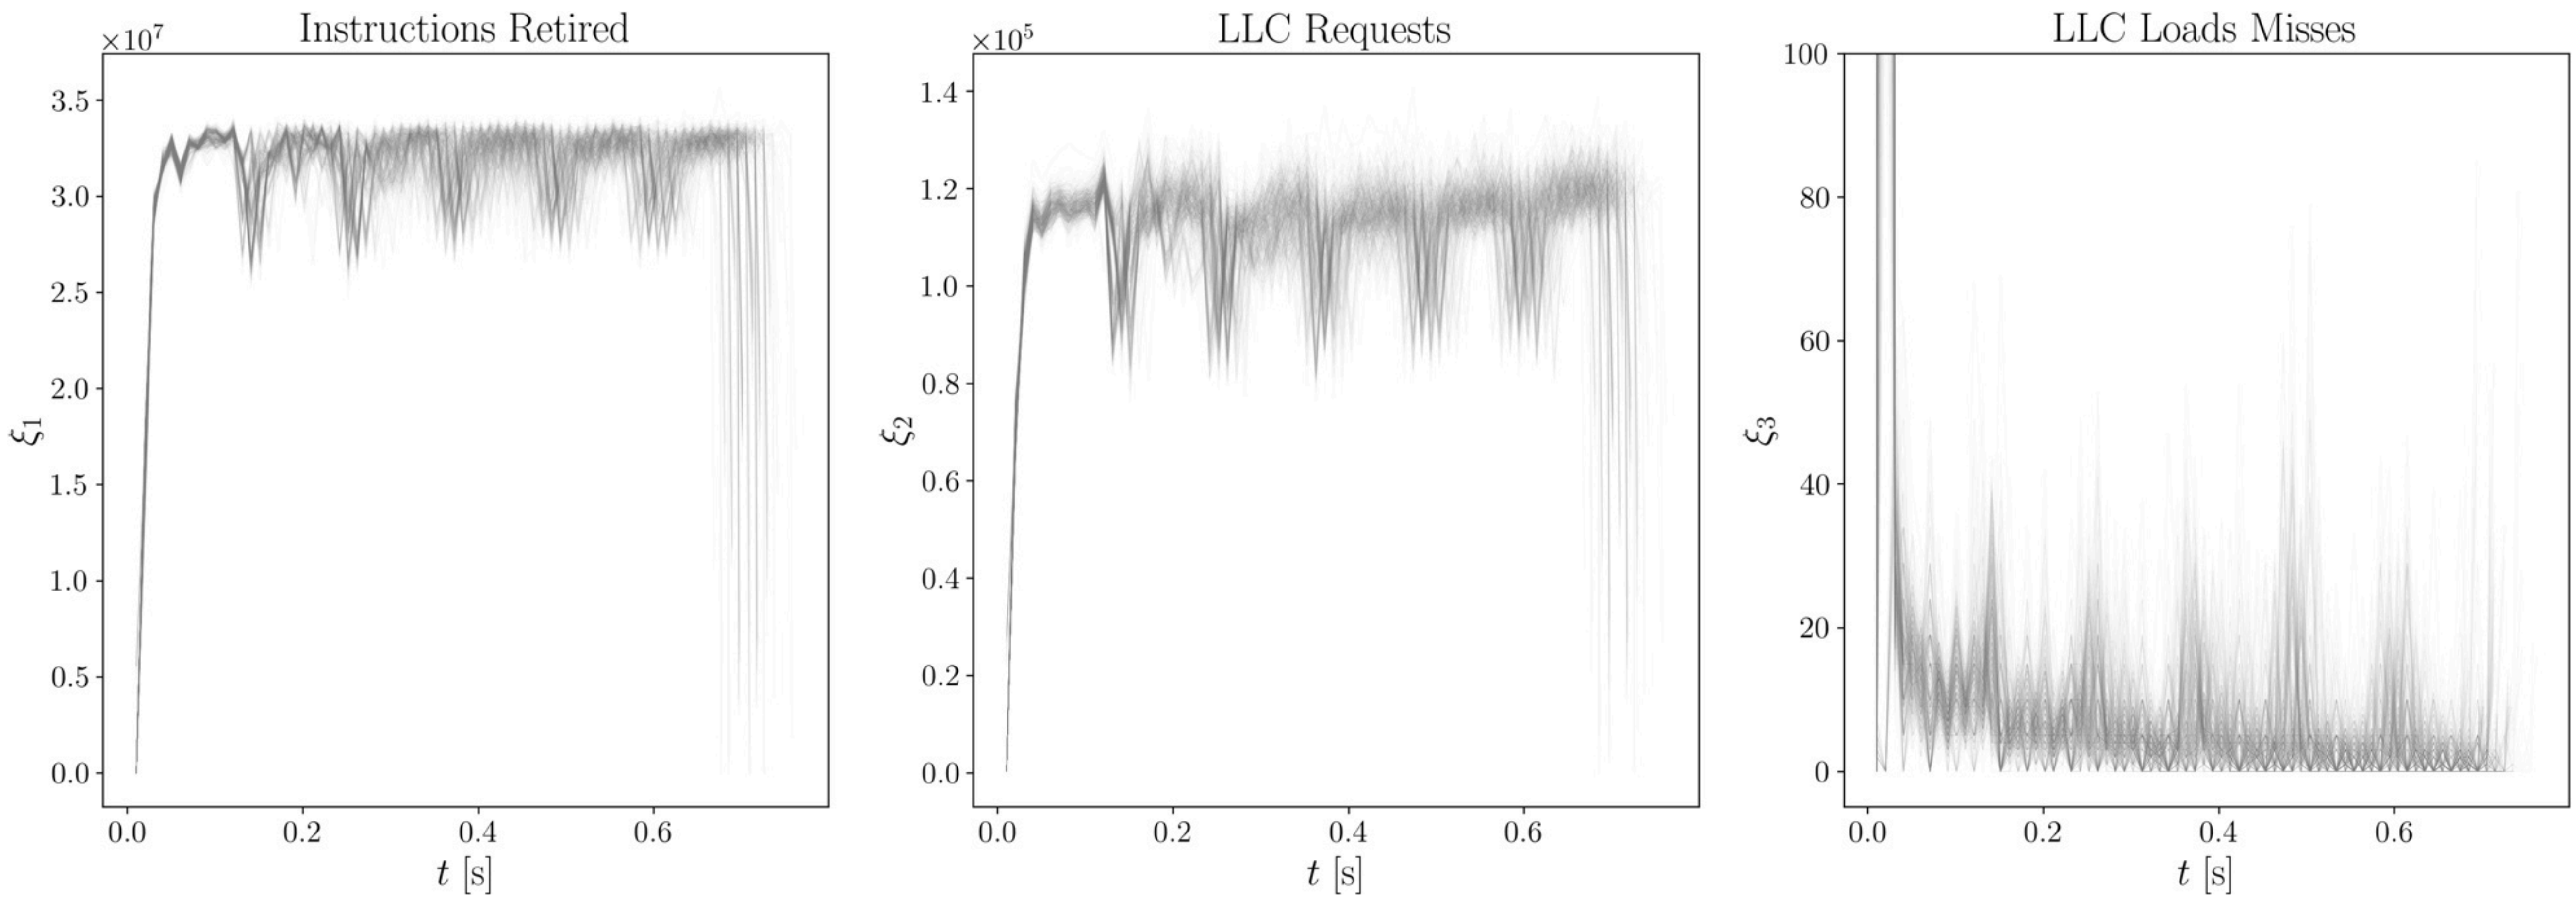

Different resource usage for the same control SW for different runs on same HW

HW-level stochasticity more pronounced for compute-intensive control SW such as MPC (than say PID)

Idea

Learn probabilistic model of HW resource from control SW execution profile

$\rightsquigarrow$ can be used for adaptive scheduling, switching among a bag of controllers

Challenge

Want to predict HW joint stochastic state

But profile data come as scattered …. want to avoid gridding

Difficult to get first principle physics based prior $\rightsquigarrow$ data-driven learning

Need guarantee $\rightsquigarrow$ "most likely HW state consistent with observed snapshots"

Need parsimony $\rightsquigarrow$ nonparametric learning

Need benign computational complexity for learning

Proposed Workflow

Step 1. Implement a control SW case study

Step 2. Profile control SW for different CPS “contexts”

Step 3. Formulate and solve multimarginal Schrödinger bridge problem (MSBP) for the measured profile scattered data snapshots

Step 4. Validate predictions w.r.t. “hold out” data

Step 1: Implement Control SW Case Study

Kinematic bicycle path tracking with NMPC and PID

Implemented in C (needed for Step 2)

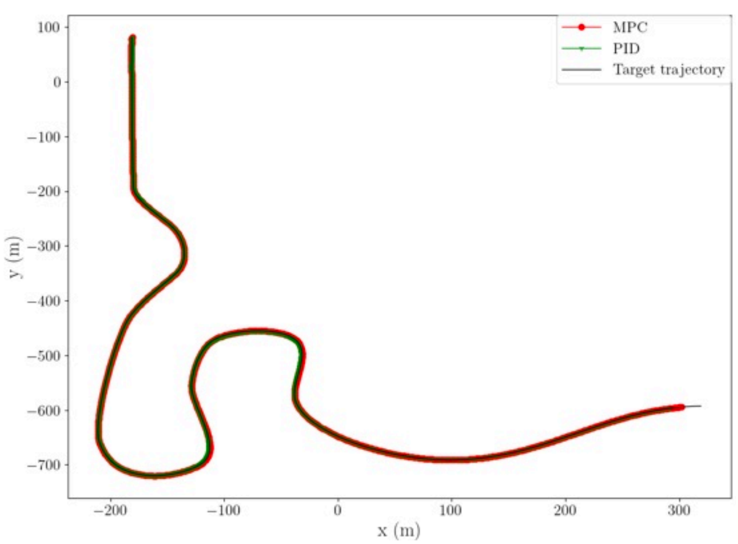

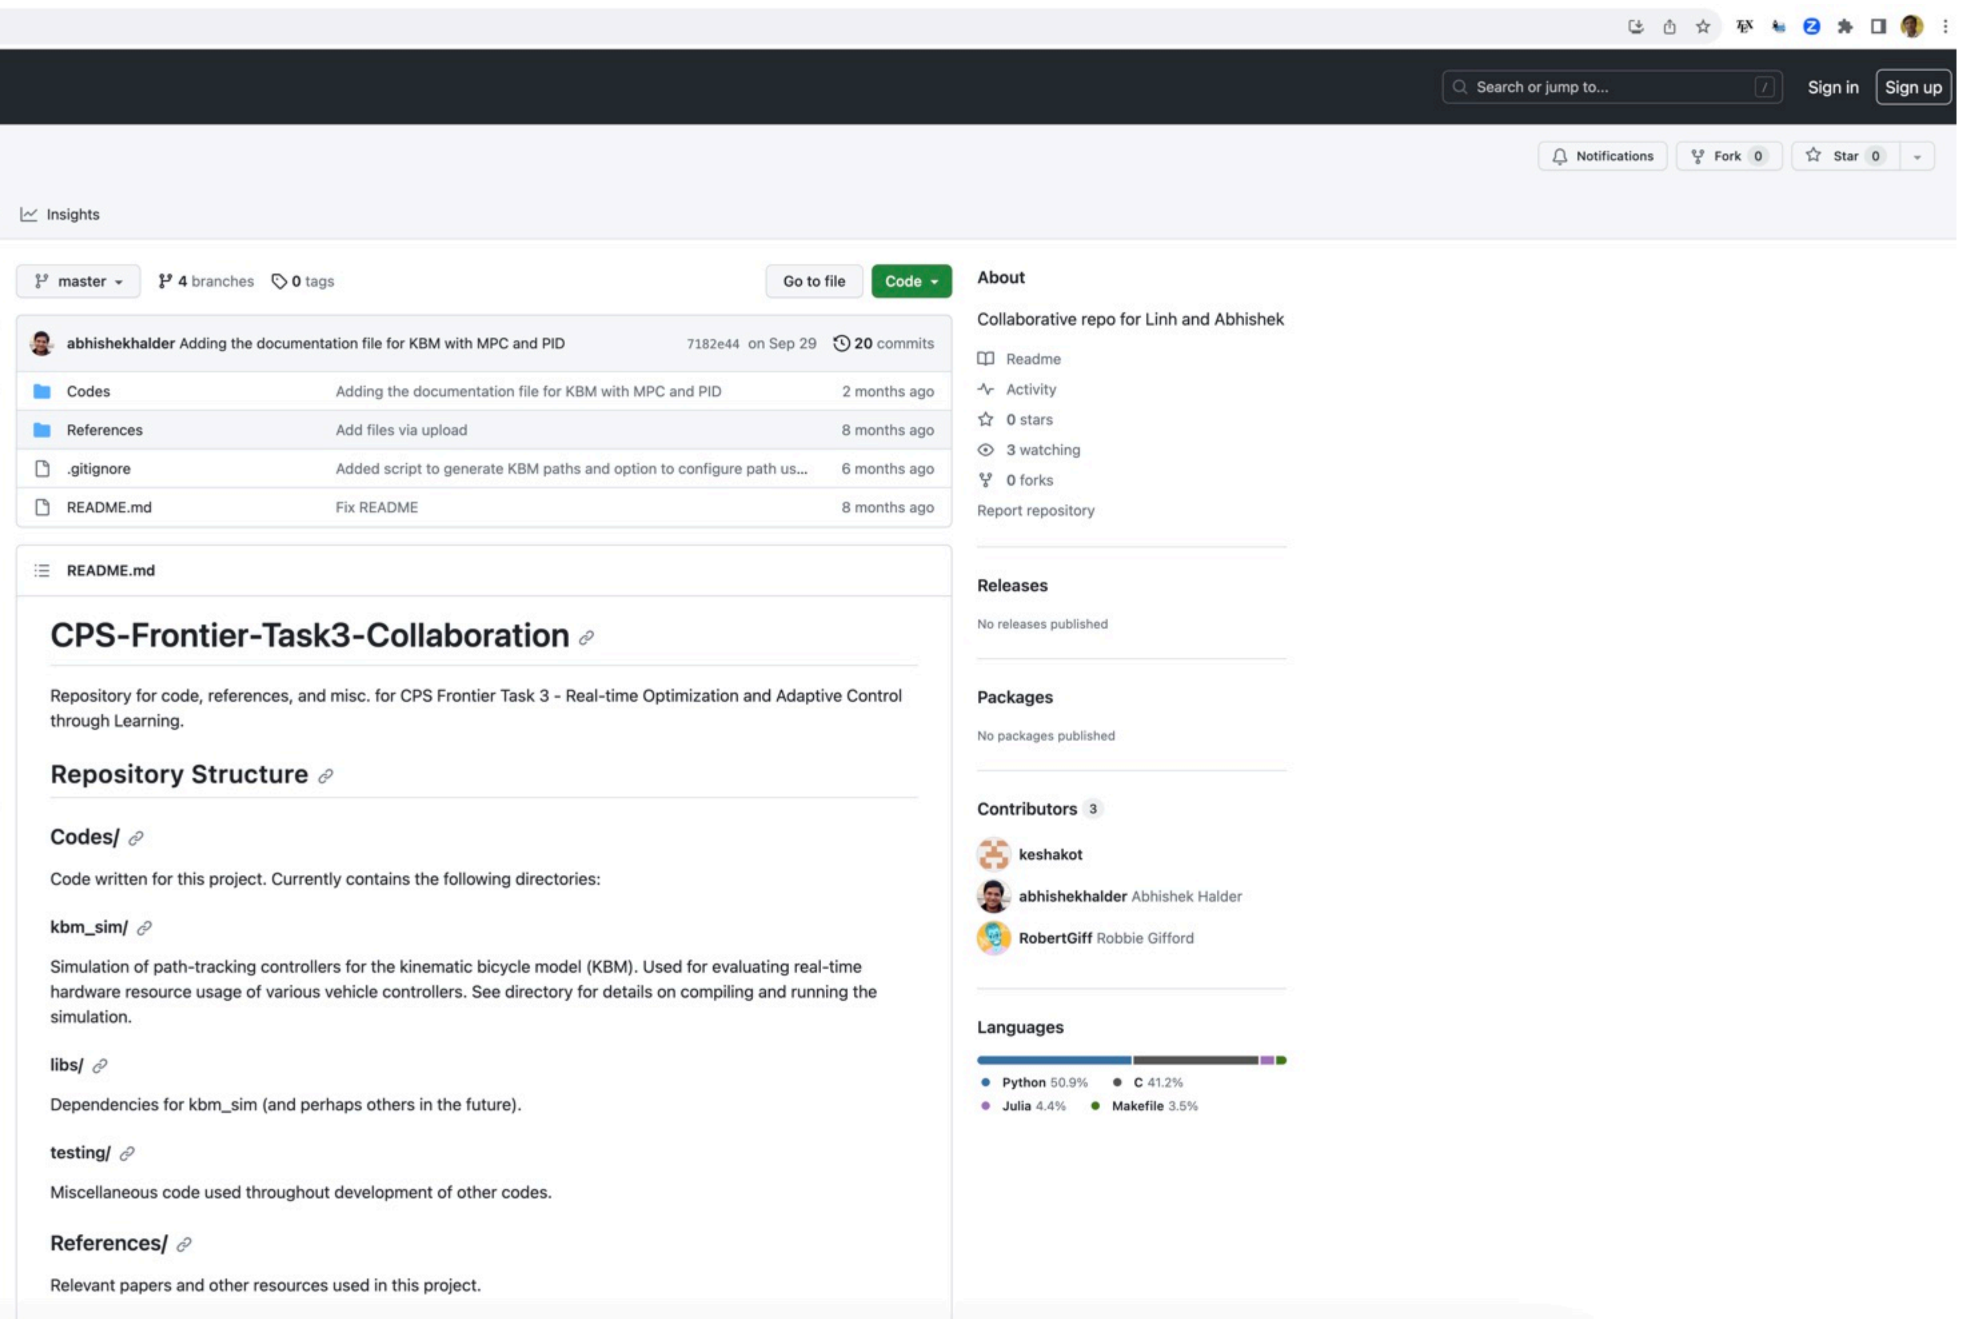

Step 2: Profile Control SW | CPS Contexts

Context vector $\boldsymbol{c} := \begin{pmatrix} \boldsymbol{c}_{\text{cyber}} \\ \boldsymbol{c}_{\text{phys}} \end{pmatrix}.$

a sample of contexts $\{\boldsymbol{c}^i\}_{i=1}^{n_{\text{context}}}$

$$\boldsymbol{c}_{\text{cyber}} = \begin{pmatrix} \text{allocated last-level cache} \\ \text{allocated memory bandwidth} \end{pmatrix},$$

$$\boldsymbol{c}_{\text{phys}} = y_{\text{des}}(x) \in \mathrm{GP}\left([x_{\min}, x_{\max}]\right),$$

In our numerical case study:

$$n_{\text{context}} = 5 \times 12 = 60$$

5 combinations of LLC partitions and memory bandwidth allocated in blocks of 2 MB using Intel CAT and Memguard

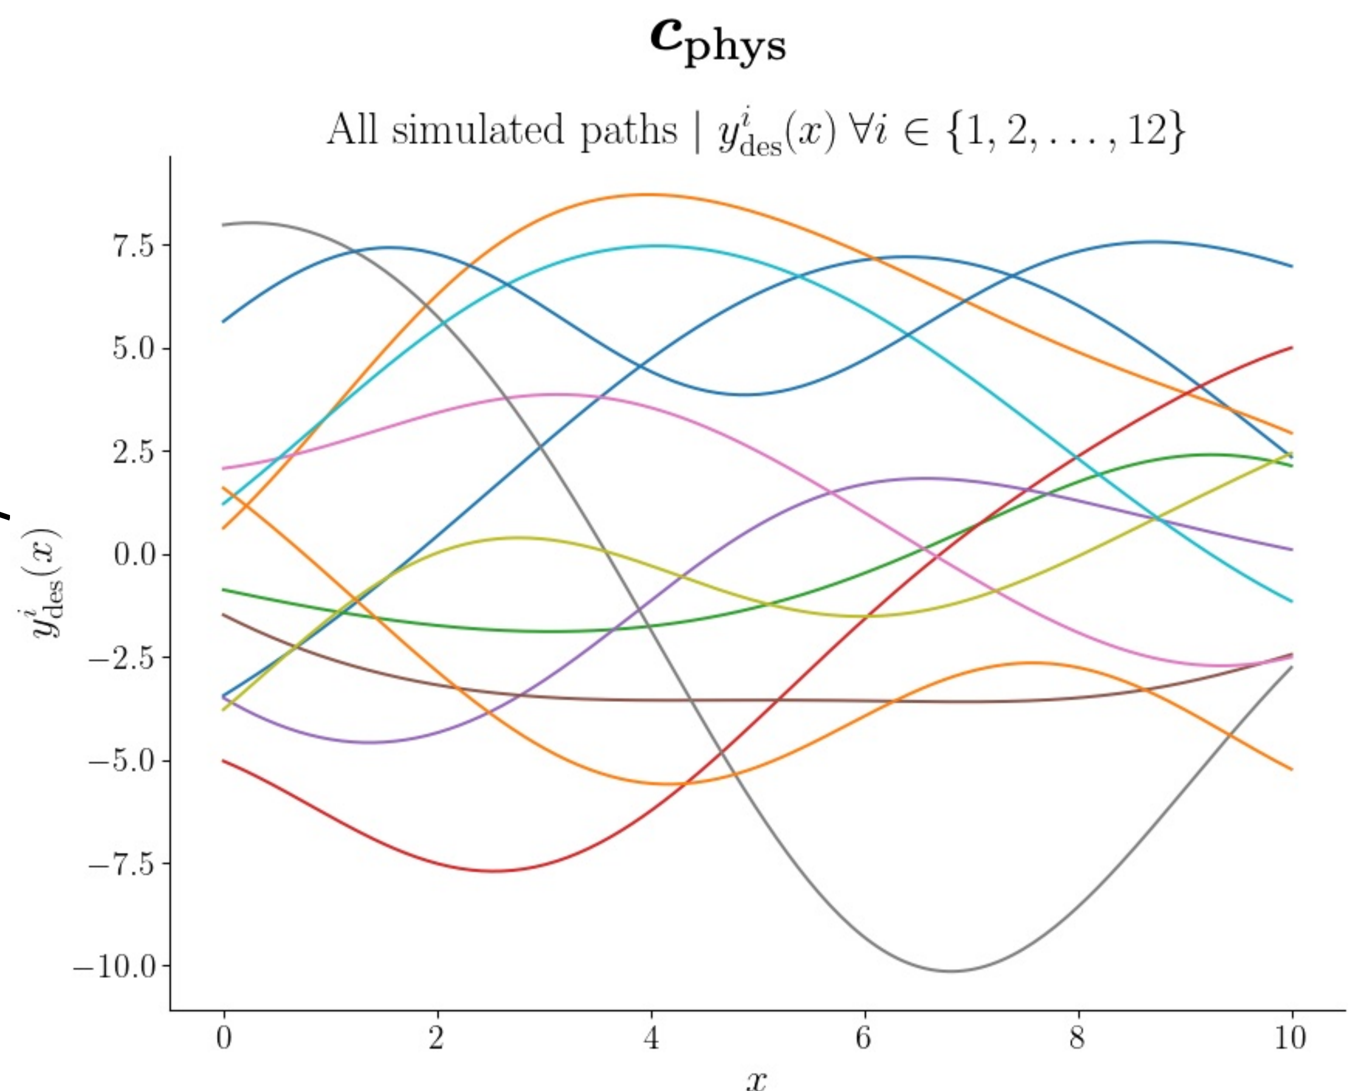

Step 2: Profile Control SW | CPS Contexts

HW resource state $\boldsymbol{\xi} := \begin{pmatrix} \xi_1 \\ \xi_2 \\ \xi_3 \end{pmatrix} = \begin{pmatrix} \text{instructions retired} \\ \text{LLC requests} \\ \text{LLC misses} \end{pmatrix}.$ Profiled every 10 ms

Think of $\boldsymbol{\xi}(\tau)$ as $\mathbb{R}^d$ valued stochastic process in continuous time $\tau \in [0,t]$

Record snapshot data at $\tau_1 \equiv 0 < \tau_2 < \ldots < \tau_{s-1} < \tau_s \equiv t.$

Snapshot index set $[\![s]\!] := \{1, 2, \ldots, s\}$

Snapshot observations $\{\mu_\sigma\}_{\sigma \in [\![s]\!]}$, i.e., $\boldsymbol{\xi}(\tau_\sigma) \sim \mu_\sigma \quad \forall \sigma \in [\![s]\!].$

Empirical measures $\mu_\sigma := \frac{1}{n} \sum_{i=1}^{n} \delta(\boldsymbol{\xi} - \boldsymbol{\xi}^i(\tau_\sigma))$, where $\{\boldsymbol{\xi}^i(\tau_\sigma)\}_{i=1}^{n}$ is scattered data

Want to predict most likely statistics $\boldsymbol{\xi}(\tau) \sim \mu_\tau$ for any $\tau \in [0, t]$.

Step 2: Profile Control SW | CPS Contexts

For each fixed context sample, generate $n = 500$ profiles, i.e., total 30k profiles sampled every 10 ms using Linux perf tool v4.9.3

Simulated $n_c = 5$ "control cycles"

Care needed to account for asynchrony across profiles

Control cycle	Mean	Standard deviation
#1	0.1181	0.0076
#2	0.2336	0.0106
#3	0.3495	0.0127
#4	0.4660	0.0143
#5	0.5775	0.0159

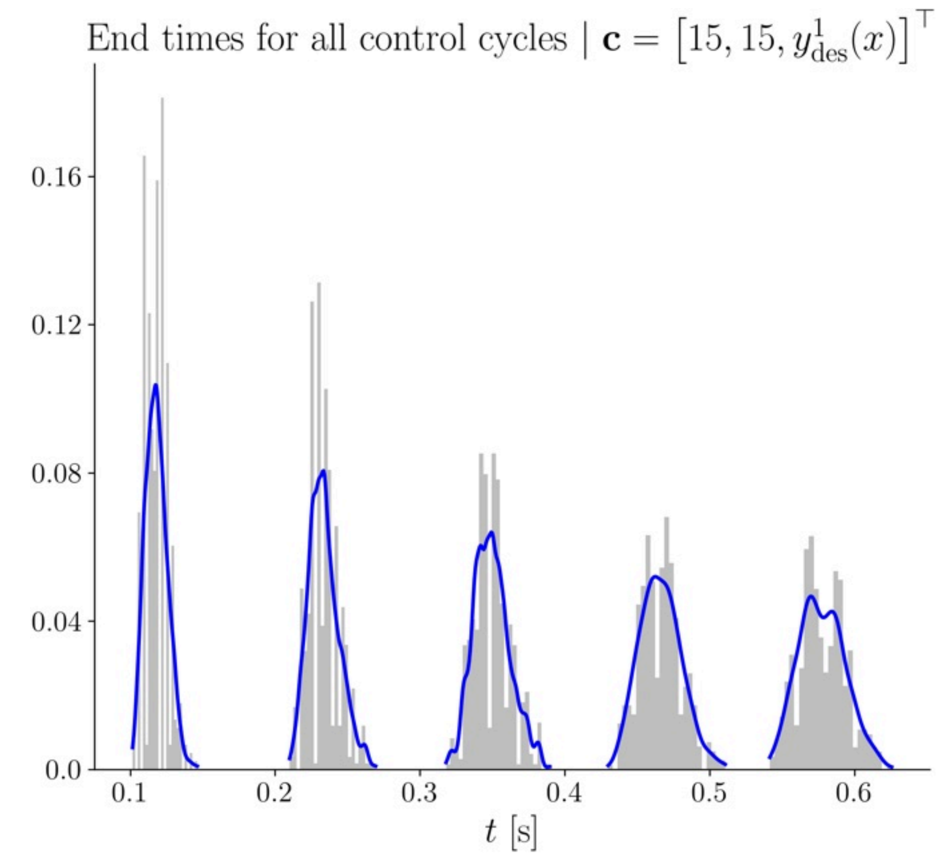

In our experiment $s := 1 + n_c\left(s_{\text{int}} + 1\right) = 1 + 5(4+1) =$ 26 snapshots

where $\tau_{\sigma(s_{\text{int}}+1)+1}$ is the sampled mean end time for the σth control cycle

Step 3: Formulate and Solve MSBP

Classical (bi-marginal) SBP

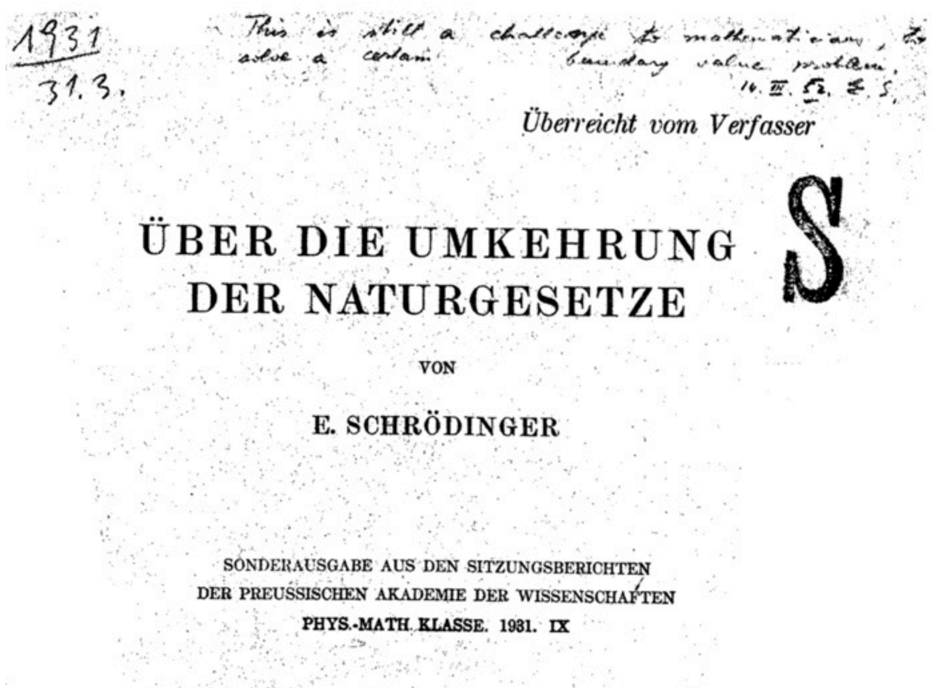
Überreicht vom Verfasser

ÜBER DIE UMKEHRUNG DER NATURGESETZE

VON

E. SCHRÖDINGER

SONDERAUSGABE AUS DEN SITZUNGSBERICHTEN DER PREUSSISCHEN AKADEMIE DER WISSENSCHAFTEN PHYS.-MATH. KLASSE. 1931. IX

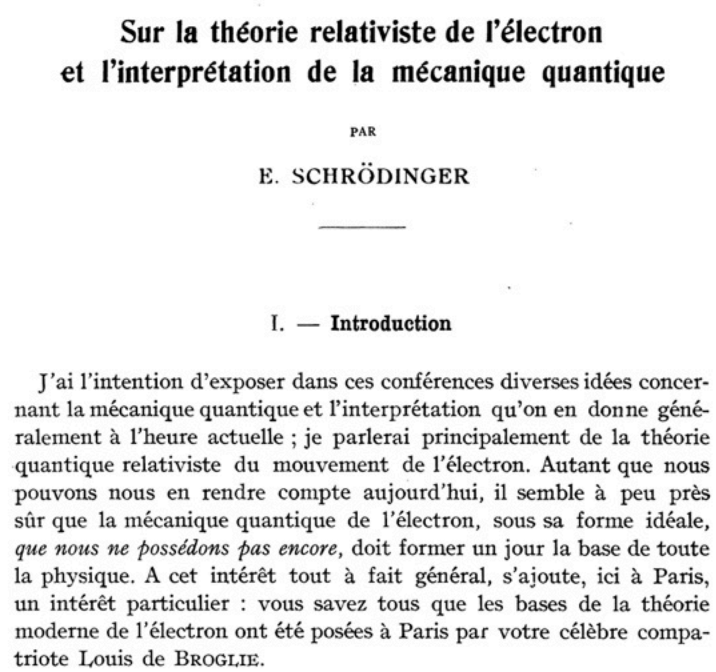
Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique

PAR

E. SCHRÖDINGER

I. — **Introduction**

J'ai l'intention d'exposer dans ces conférences diverses idées concernant la mécanique quantique et l'interprétation qu'on en donne généralement à l'heure actuelle ; je parlerai principalement de la théorie quantique relativiste du mouvement de l'électron. Autant que nous pouvons nous en rendre compte aujourd'hui, il semble à peu près sûr que la mécanique quantique de l'électron, sous sa forme idéale, *que nous ne possédons pas encore*, doit former un jour la base de toute la physique. A cet intérêt tout à fait général, s'ajoute, ici à Paris, un intérêt particulier : vous savez tous que les bases de la théorie moderne de l'électron ont été posées à Paris par votre célèbre compatriote Louis de BROGLIE.

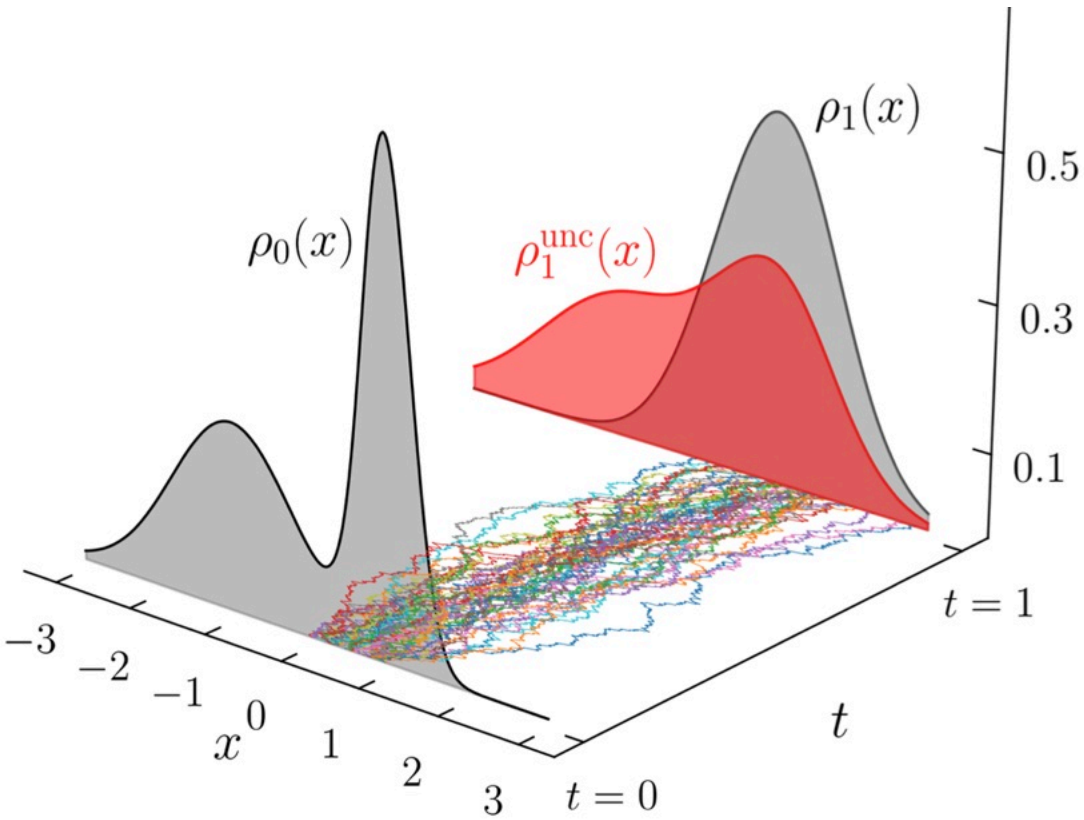

$$\mathcal{P}(\mathrm{AC}([0,1];\mathcal{P}_2(\mathcal{X})))$$

Large deviation principle on path measure

Step 3: Formulate and Solve MSBP

Multi-marginal version: MSBP formulation

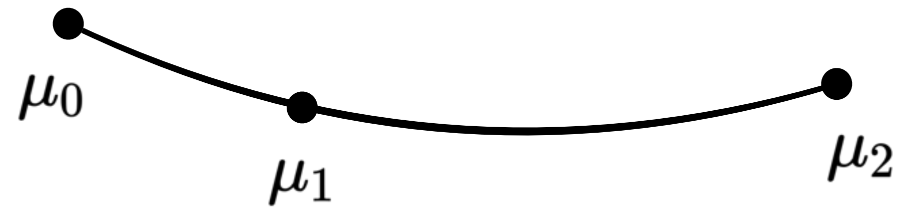

$\mathcal{X}_\sigma := \text{support}\,(\mu_\sigma) \subseteq \mathbb{R}^d \;\forall \sigma \in [\![s]\!]$, $\quad \mathcal{X}_1 \times \mathcal{X}_2 \times \ldots \times \mathcal{X}_s =: \boldsymbol{\mathcal{X}} \subseteq \left(\mathbb{R}^d\right)^{\otimes s}$

$\mathcal{M}(\mathcal{X}_\sigma)$ and $\mathcal{M}(\boldsymbol{\mathcal{X}})$ denote manifold of prob. measures on $\mathcal{X}_\sigma$ and $\boldsymbol{\mathcal{X}}$

Ground cost $\boldsymbol{C} : \boldsymbol{\mathcal{X}} \mapsto \mathbb{R}_{\geq 0}$

Let

$$\mathrm{d}\boldsymbol{\xi}_{-\sigma} := \mathrm{d}\boldsymbol{\xi}(\tau_1) \times \ldots \times \mathrm{d}\boldsymbol{\xi}(\tau_{\sigma-1}) \times \mathrm{d}\boldsymbol{\xi}(\tau_{\sigma+1}) \times \ldots \times \mathrm{d}\boldsymbol{\xi}(\tau_s)$$
$$\boldsymbol{\mathcal{X}}_{-\sigma} := \mathcal{X}_1 \times \ldots \times \mathcal{X}_{\sigma-1} \times \mathcal{X}_{\sigma+1} \times \ldots \times \mathcal{X}_s$$

MSBP:

$$\min_{\boldsymbol{M} \in \mathcal{M}(\boldsymbol{\mathcal{X}})} \int_{\boldsymbol{\mathcal{X}}} \{\boldsymbol{C}(\boldsymbol{\xi}(\tau_1), \ldots, \boldsymbol{\xi}(\tau_s)) + \varepsilon \log \boldsymbol{M}(\boldsymbol{\xi}(\tau_1), \ldots, \boldsymbol{\xi}(\tau_s))\} \boldsymbol{M}(\boldsymbol{\xi}(\tau_1), \ldots, \boldsymbol{\xi}(\tau_s)) \mathrm{d}\boldsymbol{\xi}(\tau_1) \ldots \mathrm{d}\boldsymbol{\xi}(\tau_s)$$
$$\text{subject to} \int_{\boldsymbol{\mathcal{X}}_{-\sigma}} \boldsymbol{M}(\boldsymbol{\xi}(\tau_1), \ldots, \boldsymbol{\xi}(\tau_s)) \mathrm{d}\boldsymbol{\xi}_{-\sigma} = \mu_\sigma \forall \sigma \in [\![s]\!].$$

Step 3: LDP Interpretation of MSBP

Multimarginal Gibbs kernel $\boldsymbol{K}(\boldsymbol{\xi}(\tau_1),\ldots,\boldsymbol{\xi}(\tau_s))\mu_1\otimes\ldots\otimes\mu_s$

$$\boldsymbol{K}(\boldsymbol{\xi}(\tau_1),\ldots,\boldsymbol{\xi}(\tau_s)) := \exp\left(-\frac{\boldsymbol{C}(\boldsymbol{\xi}(\tau_1),\ldots,\boldsymbol{\xi}(\tau_s))}{\varepsilon}\right)$$

Then MSBP is the same as

$$\min_{\pi\in\Pi(\mu_1,\ldots,\mu_s)} \varepsilon D_{\mathrm{KL}}\left(\pi \,\|\, \boldsymbol{K}\left(\boldsymbol{\xi}(\tau_1),\ldots,\boldsymbol{\xi}(\tau_s)\right)\mu_1\otimes\ldots\otimes\mu_s\right)$$

Set of all path measures on $\mathcal{C}([\tau_1,\tau_s],\mathbb{R}^d)$ whose time τ_σ marginal is $\mu_\sigma\ \forall\sigma\in[\![s]\!]$

Step 3: Discrete Formulation of MSBP

Ground cost is order s tensor $\boldsymbol{C} \in (\mathbb{R}^n)^{\otimes s}_{\geq 0}$, with components $[\boldsymbol{C}_{i_1,\ldots,i_s}] = \boldsymbol{C}(\boldsymbol{\xi}_{i_1}, \ldots, \boldsymbol{\xi}_{i_s})$.

Ditto for the discrete mass tensor $\boldsymbol{M} \in (\mathbb{R}^n)^{\otimes s}_{\geq 0}$

Define (marginalized) projection from nonneg tensor to nonneg vector:

$$\left[\mathrm{proj}_\sigma(\boldsymbol{M})_j\right] = \sum_{i_1,\ldots,i_{\sigma-1},i_{\sigma+1},\ldots,i_s} \boldsymbol{M}_{i_1,\ldots,i_{\sigma-1},j,i_{\sigma+1},\ldots,i_s}.$$

Discrete MSBP on scattered data:

$$\begin{aligned} &\min_{\boldsymbol{M} \in (\mathbb{R}^n)^{\otimes s}_{\geq 0}} \quad \langle \boldsymbol{C} + \varepsilon \log \boldsymbol{M}, \boldsymbol{M} \rangle \\ &\text{subject to} \quad \mathrm{proj}_\sigma(\boldsymbol{M}) = \boldsymbol{\mu}_\sigma \quad \forall \sigma \in [\![s]\!]. \end{aligned}$$

Strictly convex program in n^s decision variables

Step 3: Sequential Information Structure

Snapshot observation is a path tree:

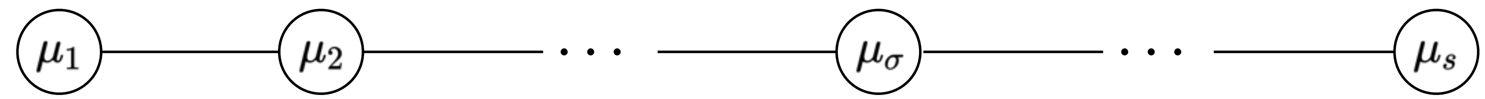

Ground cost admits path structure: $$\boldsymbol{C}(\boldsymbol{\xi}(\tau_1), \ldots, \boldsymbol{\xi}(\tau_s)) = \sum_{\sigma=1}^{s-1} c_\sigma\left(\boldsymbol{\xi}(\tau_\sigma), \boldsymbol{\xi}(\tau_{\sigma+1})\right).$$

KKT: $\boldsymbol{M}_{\text{opt}} = \boldsymbol{K} \odot \boldsymbol{U}$ where $\boldsymbol{K} := \exp(-\boldsymbol{C}/\varepsilon) \in (\mathbb{R}^n)_{>0}^{\otimes s}$, $\boldsymbol{U} := \otimes_{\sigma=1}^{s} \boldsymbol{u}_\sigma \in (\mathbb{R}^n)_{>0}^{\otimes s}$, $\boldsymbol{u}_\sigma := \exp(\boldsymbol{\lambda}_\sigma/\varepsilon)$

where $\boldsymbol{u}_\sigma$ solves multi marginal Sinkhorn contractive fixed point recursions:

$$\boldsymbol{u}_\sigma \leftarrow \boldsymbol{u}_\sigma \odot \boldsymbol{\mu}_\sigma \oslash \operatorname{proj}_\sigma\left(\boldsymbol{K} \odot \boldsymbol{U}\right) \;\; \forall \sigma \in [\![s]\!]$$

But computing $\boldsymbol{K} \odot \boldsymbol{U}$ requires $\mathcal{O}\left(n^s\right)$ operations

Step 3: From Exponential to Linear Complexity

Thm.

$$\mathrm{proj}_\sigma(\boldsymbol{K} \odot \boldsymbol{U}) = \left(\boldsymbol{u}_1^\top K^{1\to 2} \prod_{j=2}^{\sigma-1} \mathrm{diag}(\boldsymbol{u}_j) K^{j\to j+1}\right)^\top \odot \boldsymbol{u}_\sigma \odot \left(\left(\prod_{j=\sigma+1}^{s-1} K^{j-1\to j} \mathrm{diag}(\boldsymbol{u}_j)\right) K^{s-1\to s} \boldsymbol{u}_s\right) \quad \forall \sigma \in [\![s]\!],$$

Recursions become

$$\boldsymbol{u}_\sigma \leftarrow \boldsymbol{\mu}_\sigma \oslash \left(\left(\boldsymbol{u}_1^\top K^{1\to 2} \prod_{j=2}^{\sigma-1} \mathrm{diag}(\boldsymbol{u}_j) K^{j\to j+1}\right)^\top \odot \left(\left(\prod_{j=\sigma+1}^{s-1} K^{j-1\to j} \mathrm{diag}(\boldsymbol{u}_j)\right) K^{s-1\to s} \boldsymbol{u}_s\right)\right) \quad \forall \sigma \in [\![s]\!].$$

Only $s-1$ matrix-vector multiplications: complexity $\mathcal{O}\left((s-1)n^2\right)$

Numerical Case Study: Convergence

$n = 500, s = 26:$ solving for $\sim 1.49 \times 10^{70}$ decision variables in ~ 10 s in MATLAB

Linear convergence of multimarginal Sinkhorn iterates in Hilbert's projective metric

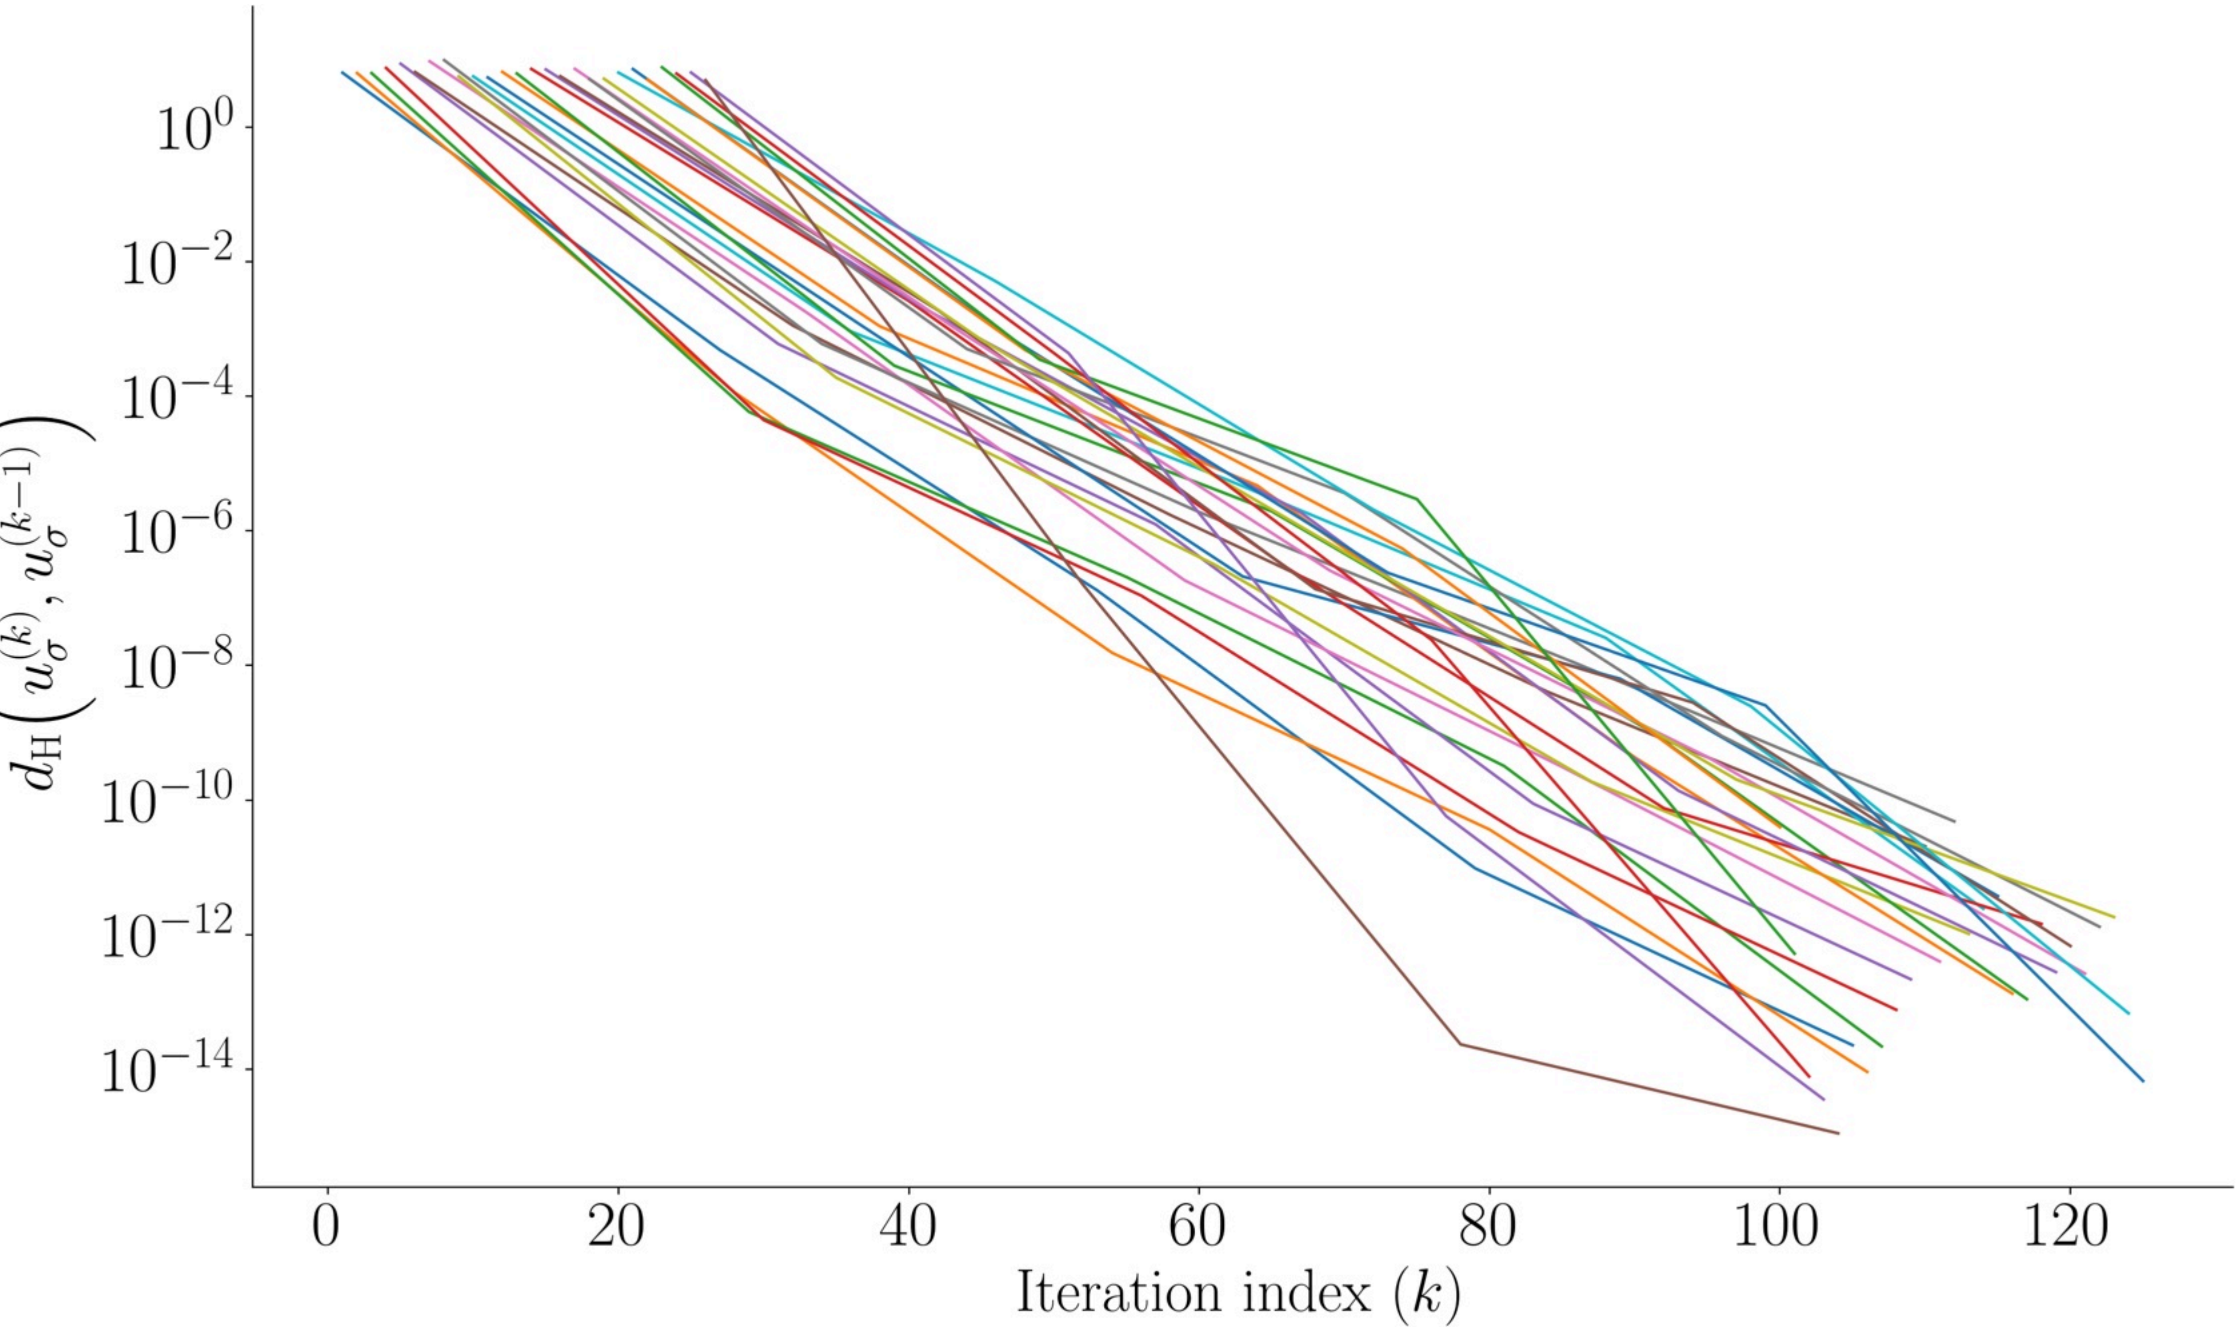

Numerical Case Study: Predicted vs Measured

Blue: predicted, red: measured, black: measured at control cycle boundaries

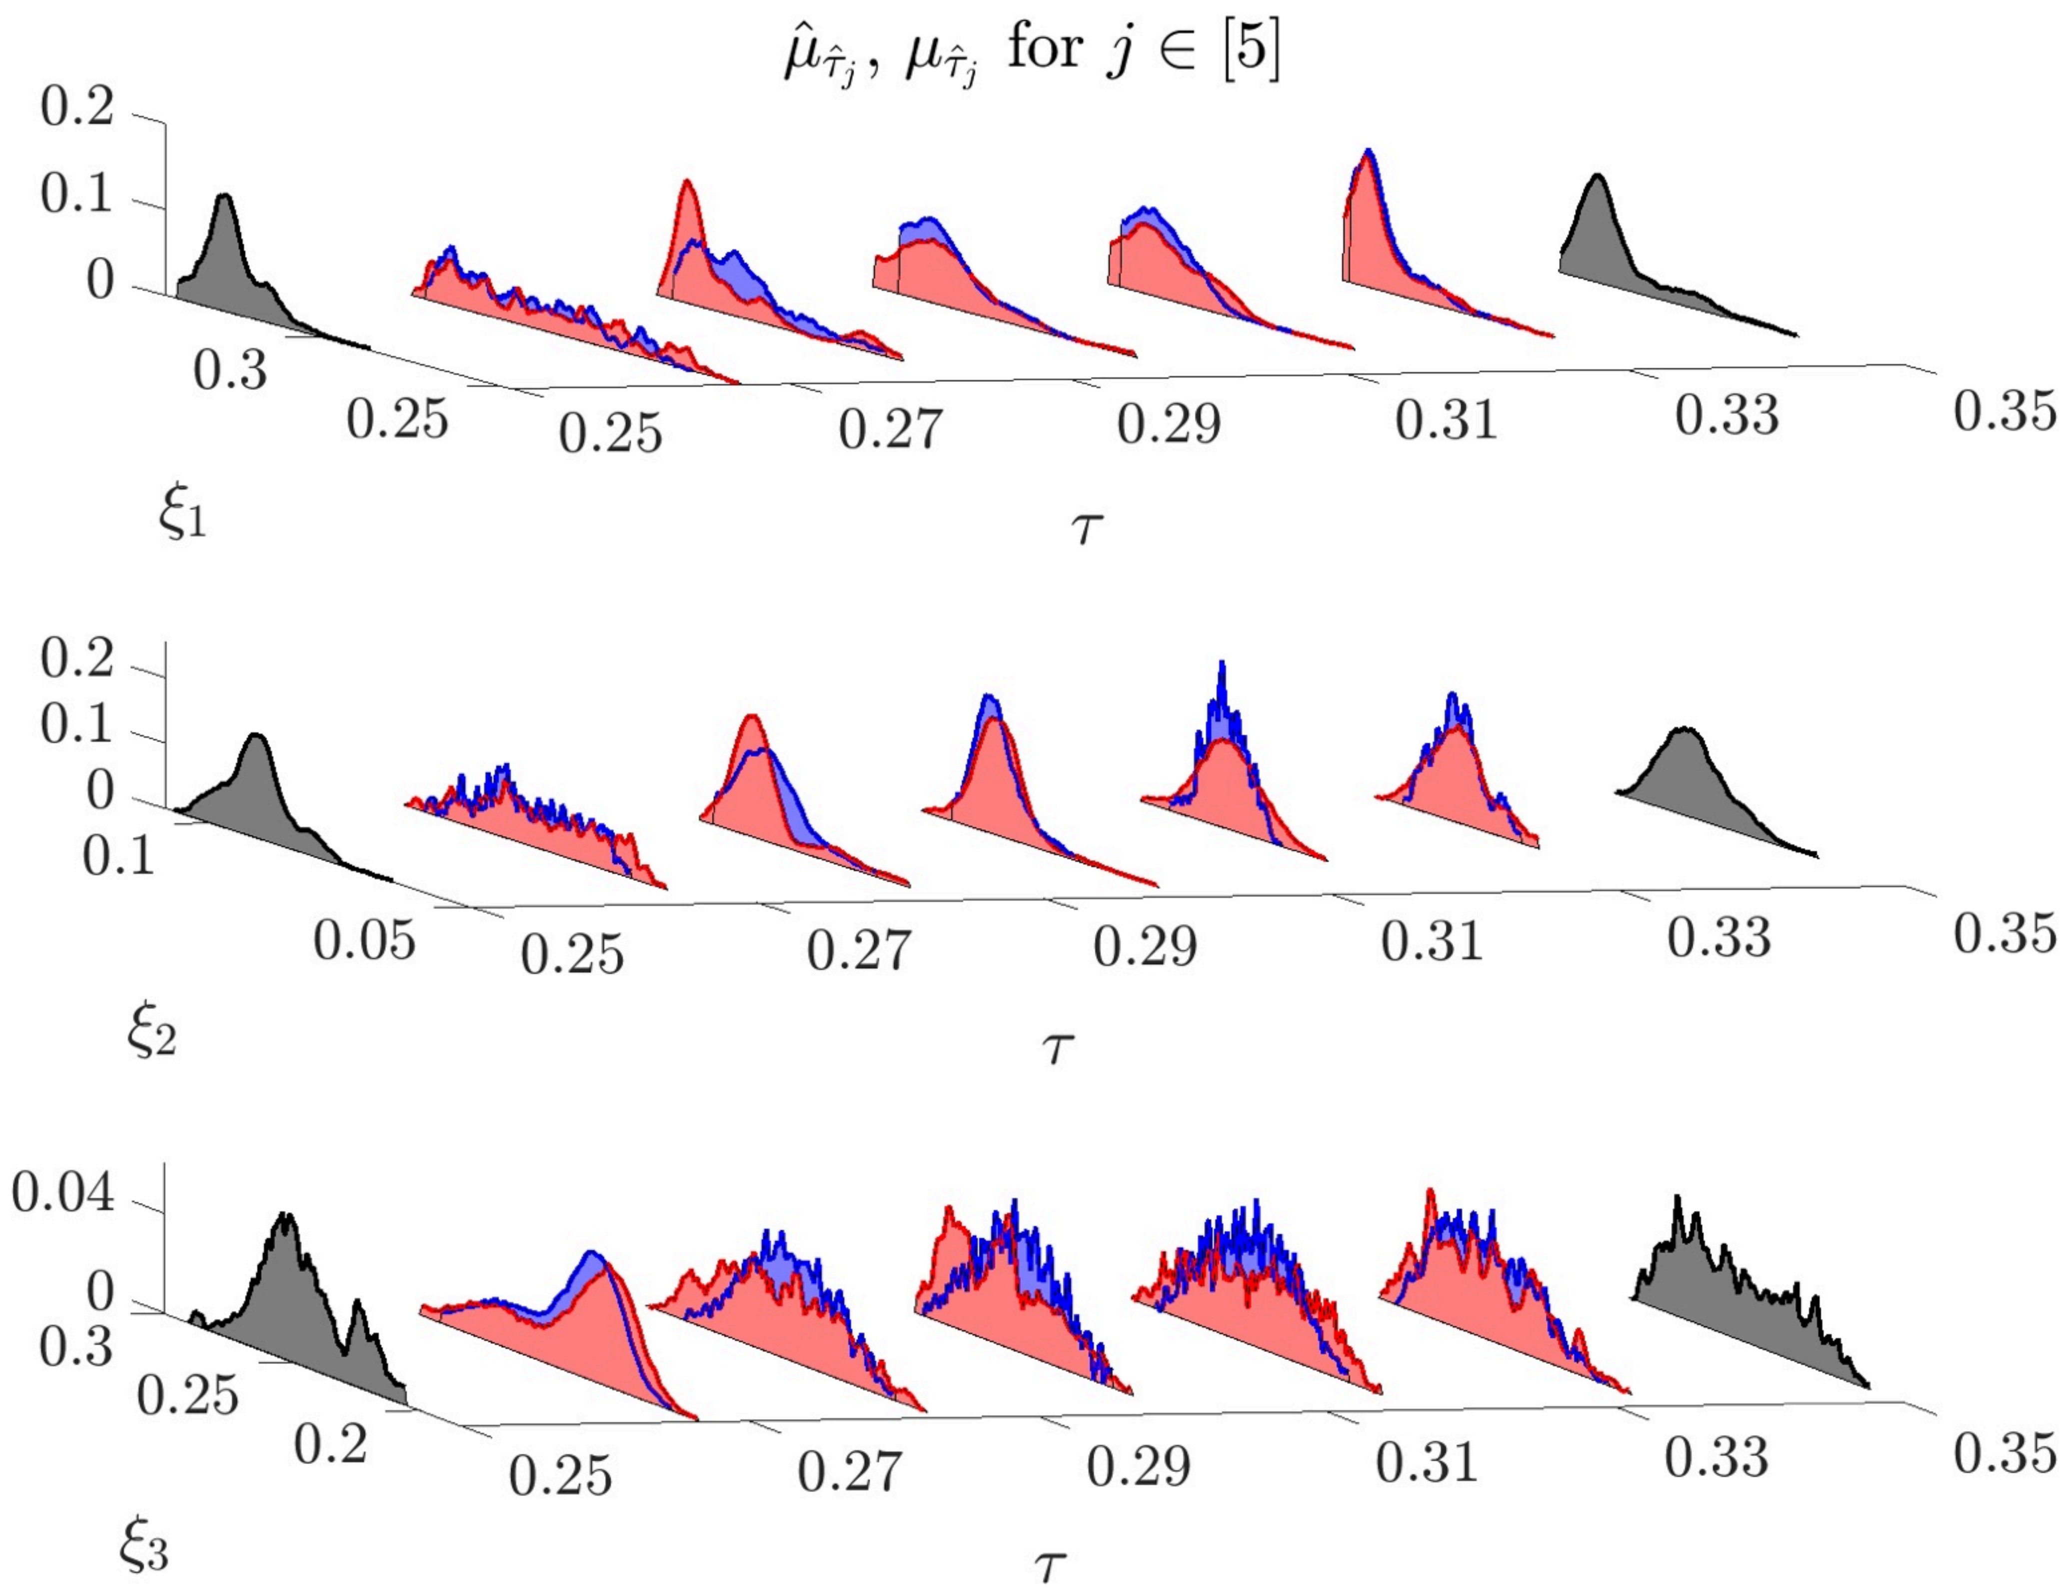

Summary

A data-driven offline learning method to predict most-likely joint HW stochastic state

Computation scales linearly with both dimension and number of snapshots

Ongoing work: multi-core profiles (DAGs that are not paths), adaptive scheduling

Details: arXiv:2310.00604

Thank You

Support: 2112755, 2111688, 1750158