

Tensor Optimization Problems in Optimal Transport

Georgiy Antonovich Bondar

Ph.D. Candidate, Applied Mathematics

University of California, Santa Cruz

February 25, 2025

Transport of a Probability Measure

Random variable $X \sim \mu$ known



Given differentiable $\tau(\cdot)$

Random variable $Y = \tau(X) \sim ?$

Transport of a Probability Measure

Random variable $X \sim \mu$ known



Given differentiable $\tau(\cdot)$

Random variable $Y = \tau(X) \sim ?$

Notation: $\nu = \tau_{\#}\mu$, read as: ν is the pushforward of μ under *transport map* τ

$$\text{Computation: } \mu(dx) = f(x) dx, \quad \nu(dy) = \tau_{\#}\mu = \frac{f(\tau^{-1}(y))}{|\nabla \tau(\tau^{-1}(y))|} dy$$

Example: 1D Transport of a Probability Measure

Measure Density


$$\mu(dx) = f(x) dx, x \in \mathcal{X}$$

$$\downarrow \tau(\cdot) = \tan(\cdot)$$

$$\begin{aligned} \nu(dy) &= g(y) dy, y \in \mathcal{Y} \\ &= \frac{f(\arctan y)}{1 + y^2} dy \end{aligned}$$

In general, $\mu(\tau^{-1}(\mathcal{U})) = \nu(\mathcal{U})$

\forall Borel $\mathcal{U} \subseteq \mathcal{Y}$



What is Optimal Transport (OT)?

Given μ, τ , the new measure ν is unique: *nothing to optimize*

Inverse problem: Given μ, ν , the map τ is *underdetermined*

What is Optimal Transport (OT)?

Given μ, τ , the new measure ν is unique: *nothing to optimize*

Inverse problem: Given μ, ν , the map τ is *underdetermined*

OT map \rightarrow

$$\tau_{\text{opt}} = \arg \inf_{\tau: \mathcal{X} \rightarrow \mathcal{Y}} \int_{\mathcal{X}} c(x, \tau(x)) d\mu(x)$$

subject to $\tau_{\#} \mu = \nu$



Monge formulation, 1781

Ground cost = cost of transporting unit amount of mass from x to $\tau(x)$

Example: $c(x, y) = \|x - y\|_2^2$, squared minimal geodesic length, etc.

Very Brief History




Gaspard Monge: OT map formulation in 1781 with $c(x, y) = \|x - y\|_1$



Erwin Schrödinger: attempts stochastic interpretation of quantum mechanics in 1931-32

Now called the **Schrödinger Bridge (SB):** diffusive version of OT




Leonid Kantorovich: OT plan reformulation in 1941

Wins 1975 Nobel prize in Economics for this work

Math for OT takes shape in late 20th - early 21st century



C. Villani



A. Figalli



Y. Brenier



J-D. Benamou




R.J. McCann



X-N. Ma



N. Trudinger



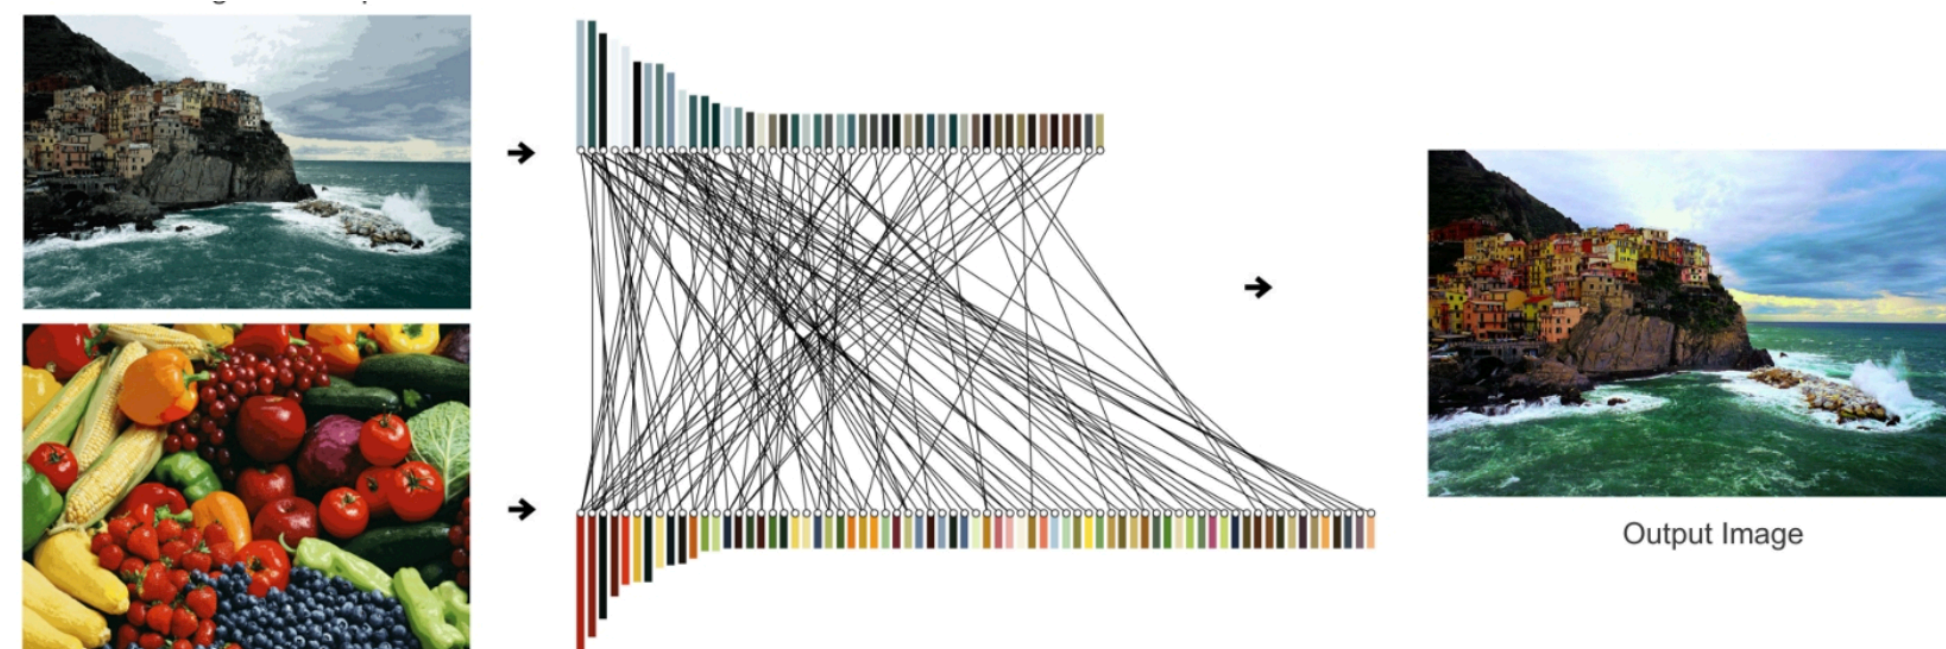
X-J. Wang

AI/ML Applications

OT of color



Credit: https://oriel.github.io/color_transfer.html



OT of style

Picasso


Dürer

Matisse




ν_1

ν_2

ν_3



μ



Credit: Kolkin, Salavon, Shakhnarovich, CVPR 2019

SB in diffusion model generative AI

Stable diffusion, DALL-E



Credit: <https://github.com/Stability-AI/generative-models>

Science Applications of SB


Protein synthesis

UAI 2023


Aligned Diffusion Schrödinger Bridges

Vignesh Ram Somnath^{*1,2} Matteo Pariset^{*1,3} Ya-Ping Hsieh¹
 Maria Rodriguez Martinez² Andreas Krause¹ Charlotte Bunne¹

¹Department of Computer Science, ETH Zürich
²IBM Research Zürich
³Department of Computer Science, EPFL



Material synthesis



Dispersed particles

Ordered structure




Superresolution

NeurIPS 2024

Diffusion Schrödinger Bridge Matching

Yuyang Shi^{*} Valentin De Bortoli^{*} Andrew Campbell Arnaud Doucet
 University of Oxford ENS ULM University of Oxford University of Oxford



Low res

High res

2024 Hugo Schuck Award by American Automatic Control Council

Outline of This Talk

Mathematical Background


Tensor Optimization for OT Regularity

First computational method for OT regularity

Tensor Optimization for Graph-structured Multimarginal SB

Application in learning computational resource usage

Background: OT Formulation



OT map

$$\tau_{\text{opt}} = \arg \inf_{\tau: \mathcal{X} \rightarrow \mathcal{Y}} \int_{\mathcal{X}} c(x, \tau(x)) d\mu(x)$$

subject to $\tau_{\#} \mu = \nu$




Monge formulation, 1781
Nonlinear nonconvex program

↑ ?

OT plan

$$\pi_{\text{opt}} = \arg \inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\pi(x, y)$$



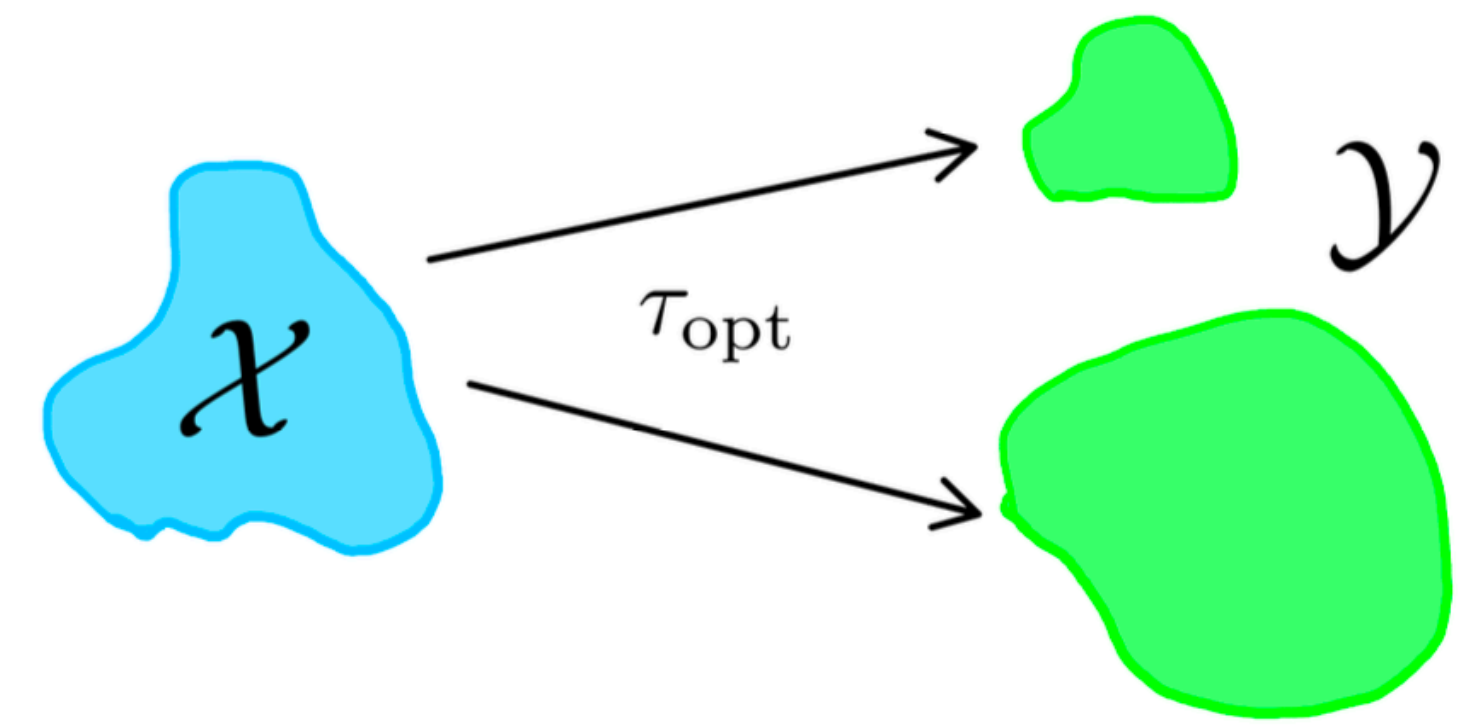
Kantorovich formulation, 1941
Linear program

The inf value is called the squared Wasserstein distance


Background: OT Regularity

Question: is τ_{opt} continuous?

Answer: Yes if μ, ν abs. continuous + extra condition on c and manifold




Background: OT Regularity



Question: is τ_{opt} continuous?

Answer: Yes if μ, ν abs. continuous + **extra condition on c and manifold**

Defn: Ma-Trudinger-Wang (MTW) tensor (2005, 2009)



X-N. Ma N. Trudinger X-J. Wang

$$\mathfrak{S}_{(x,y)}(\xi, \eta) := \sum_{i,j,k,l,p,q,r,s} (c_{ij,p}c^{p,q}c_{q,rs} - c_{ij,rs})c^{r,k}c^{s,l}\xi_i\xi_j\eta_k\eta_l$$

$$\forall x \in \mathcal{X}, y \in \mathcal{Y}, \xi \in T_x\mathcal{X}, \eta \in T_y^*\mathcal{Y}$$

$$c_{ij,kl} = \partial_{x_i}\partial_{x_j}\partial_{y_k}\partial_{y_l}c(x,y), \quad c^{i,j}(x,y) = \left[((\nabla_x \otimes \nabla_y)c)^{-1} \right]_{i,j}$$

Background: OT Regularity

Defn: $\text{MTW}(0)$ and $\text{MTW}(\kappa)$, $\kappa > 0$

If $\mathfrak{S}_{(\cdot, \cdot)}(\xi, \eta) \geq 0 \forall (\xi, \eta)$ s.t. $\eta(\xi) = 0$ then c satisfies $\text{MTW}(0)$

If $\exists k > 0$ s.t. $\mathfrak{S}_{(\cdot, \cdot)}(\xi, \eta) \geq \kappa \|\xi\|^2 \|\eta\|^2$ then c satisfies $\text{MTW}(\kappa)$

Defn: Nonnegative Cost Curvature (NNCC)

If $\mathfrak{S}_{(\cdot, \cdot)}(\xi, \eta) \geq 0 \forall (\xi, \eta)$ then c satisfies NNCC

Difficult to verify analytically. **Our approach:** computational certificate

Background: Schrödinger Bridge Problem (SBP)


Static SBP = Kantorovich OT + entropic regularization

$$\pi_{\text{opt}} = \arg \inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathcal{X} \times \mathcal{Y}} (c(x, y) + \varepsilon \log \pi(x, y)) d\pi(x, y)$$

Strictly convex program

Continuous SBP = optimization over measure-valued path space

Generates the maximum likelihood trajectory on path space



$$\arg \inf_{M \in \mathcal{P}(\mathcal{X} \times \mathcal{Y})} \int_{\mathcal{X} \times \mathcal{Y}} (c(\xi(0), \xi(t)) + \varepsilon \log M(\xi(0), \xi(t))) M(\xi(0), \xi(t)) d\xi(0) d\xi(t)$$

subject to $\int_{\mathcal{X}} M(\xi(0), \xi(t)) d\xi(0) = \nu, \quad \int_{\mathcal{Y}} M(\xi(0), \xi(t)) d\xi(t) = \mu$

Background: Discrete SBP

Bi-marginal a.k.a. classical SBP



$$M_{\text{opt}} = \arg \min_{M \in \mathbb{R}_{\geq 0}^{n \times n}} \langle C + \varepsilon \log M, M \rangle$$

$\mathbb{R}_{\geq 0}^{n \times n}$

$$\text{subject to } \text{proj}_{\sigma} (M) = \mu_{\sigma} \quad \forall \sigma \in \{1, 2\}$$

Δ^{n-1}

Darker distributional path = more likely



Weighted scattered data:

$$\{\xi^i(\tau_{\sigma})\}_{i=1}^n, \mu_{\sigma} = \frac{1}{n} \sum_{i=1}^n \delta(\xi - \xi^i(\tau_{\sigma}))$$

Background: Discrete SBP

Bi-marginal a.k.a. classical SBP

$$M_{\text{opt}} = \arg \min_{M \in \mathbb{R}_{\geq 0}^{n \times n}} \langle C + \varepsilon \log M, M \rangle$$

$\mathbb{R}_{\geq 0}^{n \times n}$

$$\text{subject to } \text{proj}_{\sigma} (M) = \mu_{\sigma} \quad \forall \sigma \in \{1, 2\}$$

Δ^{n-1}

Multimarginal SBP (MSBP)

$$M_{\text{opt}} = \arg \min_{M \in (\mathbb{R}_{\geq 0}^n)^{\otimes s}} \langle C + \varepsilon \log M, M \rangle$$

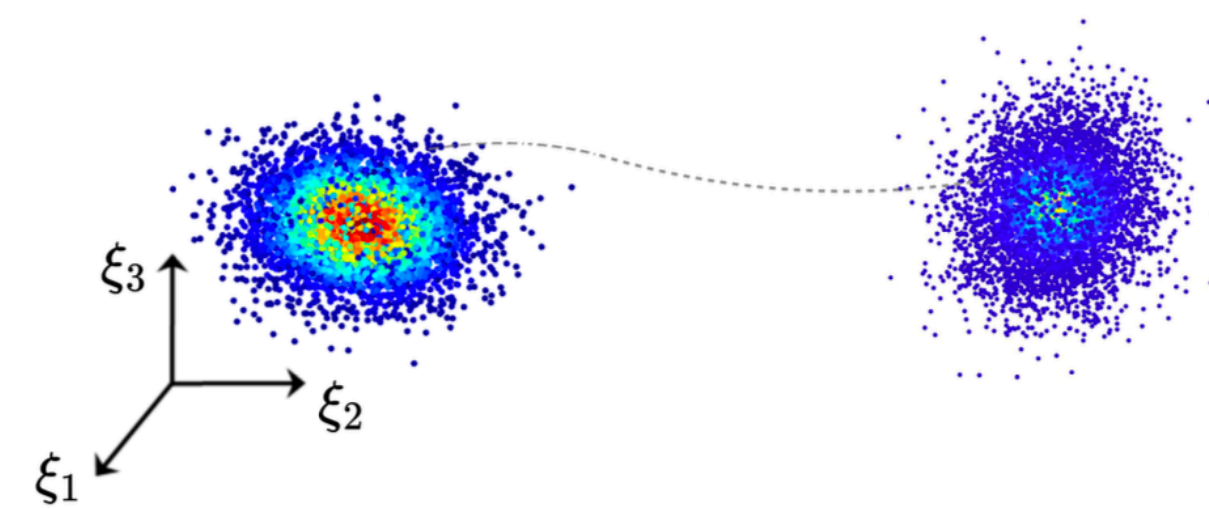
$(\mathbb{R}_{\geq 0}^n)^{\otimes s}$

$$\text{subject to } \text{proj}_{\sigma} (M) = \mu_{\sigma} \quad \forall \sigma \in [s]$$

Δ^{n-1}

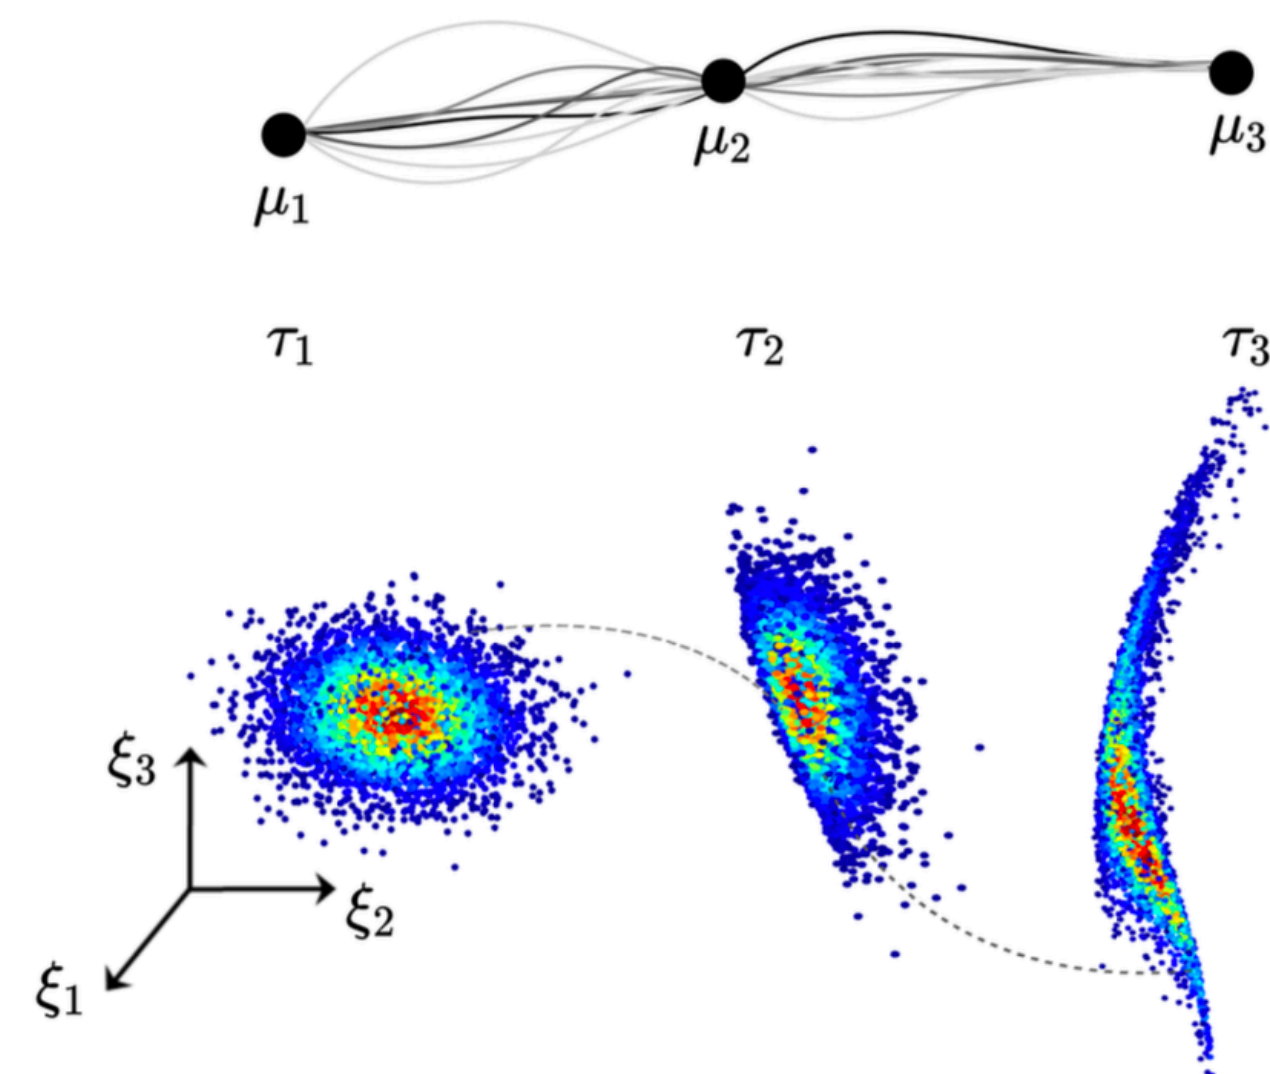
$\mathbb{N}_{\geq 2}$

Darker distributional path = more likely



Weighted scattered data:

$$\{\xi^i(\tau_{\sigma})\}_{i=1}^n, \mu_{\sigma} = \frac{1}{n} \sum_{i=1}^n \delta(\xi - \xi^i(\tau_{\sigma}))$$



Background: Sinkhorn Iteration to Solve MSBP

Step 1: Let $\mathbf{K} := \exp(-\mathbf{C}/\varepsilon) \in (\mathbb{R}^n)_{>0}^{\otimes s}$, initialize $\mathbf{u}_\sigma := \exp(\boldsymbol{\lambda}_\sigma/\varepsilon) \in \mathbb{R}_{>0}^n$

Step 2: Perform Sinkhorn iterations until (linear) convergence

$$\mathbf{u}_\sigma \leftarrow \mathbf{u}_\sigma \otimes \boldsymbol{\mu}_\sigma \oslash \text{proj}_\sigma(\mathbf{K} \odot \mathbf{U}) \quad \forall \sigma \in \llbracket s \rrbracket$$

Step 3: $\mathbf{M}_{\text{opt}} = \mathbf{K} \odot \mathbf{U}$ where $\mathbf{U} := \otimes_{\sigma=1}^s \mathbf{u}_\sigma \in (\mathbb{R}^n)_{>0}^{\otimes s}$

Background: Sinkhorn Iteration to Solve MSBP

Step 1: Let $\mathbf{K} := \exp(-\mathbf{C}/\varepsilon) \in (\mathbb{R}^n)_{>0}^{\otimes s}$, initialize $\mathbf{u}_\sigma := \exp(\boldsymbol{\lambda}_\sigma/\varepsilon) \in \mathbb{R}_{>0}^n$

Step 2: Perform Sinkhorn iterations until (linear) convergence

$$\mathbf{u}_\sigma \leftarrow \mathbf{u}_\sigma \otimes \boldsymbol{\mu}_\sigma \oslash \text{proj}_\sigma(\mathbf{K} \odot \mathbf{U}) \quad \forall \sigma \in [s]$$

Step 3: $\mathbf{M}_{\text{opt}} = \mathbf{K} \odot \mathbf{U}$ where $\mathbf{U} := \otimes_{\sigma=1}^s \mathbf{u}_\sigma \in (\mathbb{R}^n)_{>0}^{\otimes s}$

Trouble: computing $[\text{proj}_\sigma(\mathbf{M})]_j = \sum_{i_1, \dots, i_{\sigma-1}, i_{\sigma+1}, \dots, i_s} \mathbf{M}_{i_1, \dots, i_{\sigma-1}, j, i_{\sigma+1}, \dots, i_s}$

has $\mathcal{O}(n^s)$ complexity more on this later

Tensor Optimization for OT Regularity

Problem Formulation

Assumption A1: MTW tensor is **rational** in $(x, y) \in \mathcal{X} \times \mathcal{Y}$ **semialgebraic**

Sufficient but not necessary: c is **rational** in $(x, y) \in \mathcal{X} \times \mathcal{Y}$ **semialgebraic**

Problem Formulation

Assumption A1: MTW tensor is **rational** in $(x, y) \in \mathcal{X} \times \mathcal{Y}$ **semialgebraic**

Sufficient but not necessary: c is **rational** in $(x, y) \in \mathcal{X} \times \mathcal{Y}$ **semialgebraic**

Forward problem:

Given $c, \mathcal{X}, \mathcal{Y}$ as per **A1**, certify / falsify if the ground cost $c : \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}_{\geq 0}$ satisfies either MTW(0) or MTW(κ) or NNCC condition

Inverse problem:

Given $c, \mathcal{X}, \mathcal{Y}$ as per **A1**, find semialgebraic $\mathcal{U} \times \mathcal{V} \subseteq \mathcal{X} \times \mathcal{Y}$ such that $c : \mathcal{U} \times \mathcal{V} \mapsto \mathbb{R}_{\geq 0}$ satisfies either MTW(0) or MTW(κ) or NNCC condition

Sum-of-Squares (SOS) Polynomials

$$\text{poly}(x) \in \sum_{\text{SOS}} [x] \text{ if } \text{poly}(x) = (\text{poly}_1(x))^2 + (\text{poly}_2(x))^2 + \dots + (\text{poly}_m(x))^2$$


for some $\text{poly}_1(x), \text{poly}_2(x), \dots, \text{poly}_m(x) \in \mathbb{R}[x], x \in \mathbb{R}^n$

SOS decomposition:

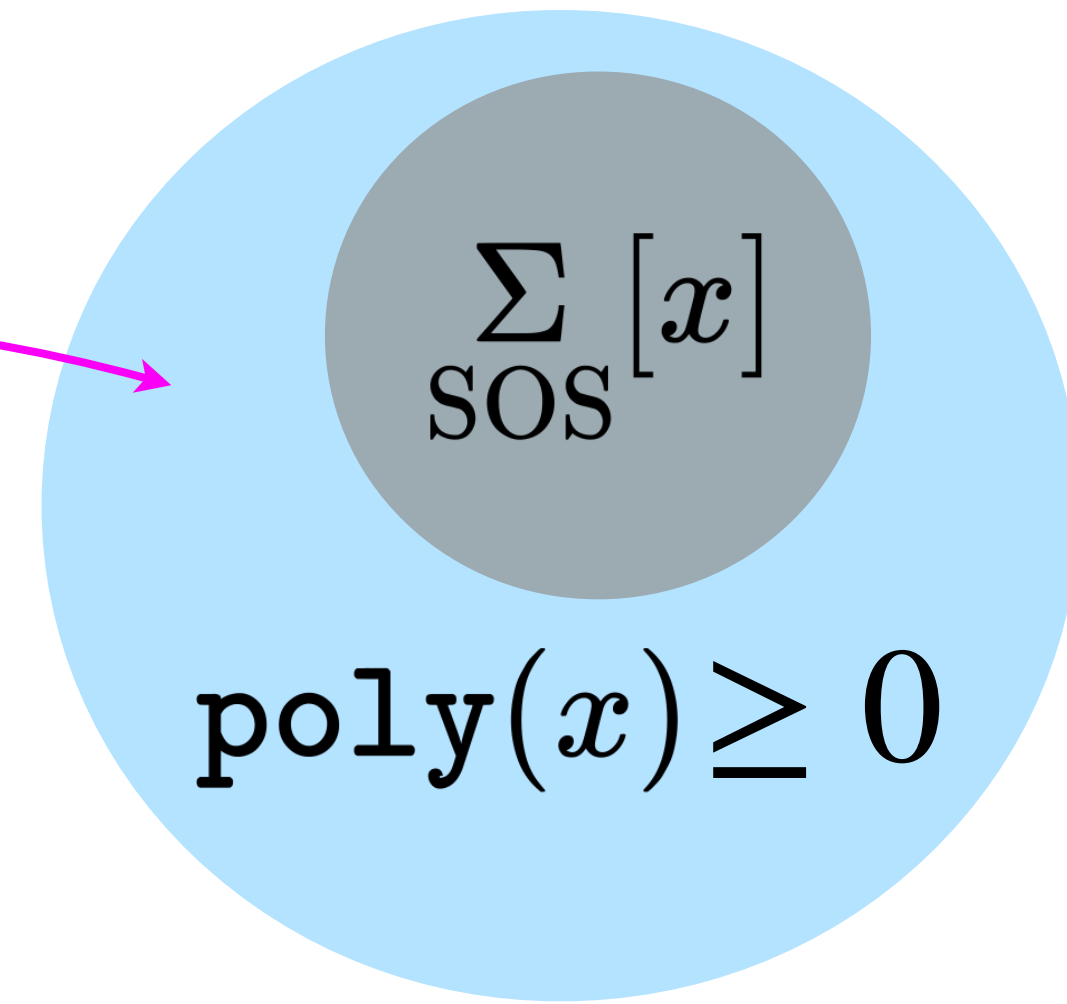
$$s \in \sum_{\text{SOS}} [x] \Rightarrow s = Z_d(x)^\top S Z_d(x)$$

for some $S \succeq 0, d \in \mathbb{N}$

where $Z_d(x)$ with monomial entries $(1, x, \dots, x^d)$ of length $\zeta := \sum_{r=0}^d \binom{n+r-1}{r}$



$$\text{polyMotzkin}(x_1, x_2) := x_1^4 x_2^2 + x_1^2 x_2^4 - 3x_1^2 x_2^2 + 1$$



SOS Programming

Defn: Semialgebraic set

Finite union of sets of the form $\{x \in \mathbb{R}^n \mid g(x) \leq 0, g \in \mathbb{R}_{d_g}[x], d_g \in \mathbb{N}\}$

Defn: Polynomial Optimization

Computationally intractable

For $f, g_{i \in [n_g]} \in \mathbb{R}[x]$

$\min_{x \in \mathbb{R}^n} f(x)$ such that $x \in \mathcal{C} := \{x \in \mathbb{R}^n \mid g_i(x) \leq 0 \forall i \in [n_g]\}$

\Leftrightarrow

$\max_{\gamma \in \mathbb{R}} \gamma$ subject to $f(x) - \gamma \geq 0 \quad \forall x \in \mathcal{C}$ semialgebraic

SOS Programming (contd.)

$$f(x) - \gamma \geq 0 \quad \forall x \in \mathcal{C} \text{ Archimedean}$$

\Downarrow Putinar's Positivstellensatz (1993)

$$\exists s_0, s_1, \dots, s_{n_g} \in \sum_{\text{SOS}} [x] \text{ such that } f(x) - \gamma = s_0(x) - \sum_{i \in [n_g]} s_i(x) g_i(x)$$

Defn: SOS tightening of Polynomial Optimization \rightsquigarrow Semidefinite program

$$\max_{(\gamma, S_0, S_1, \dots, S_{n_g}) \in \mathbb{R} \times \underbrace{\mathbb{S}_+^\zeta \times \dots \times \mathbb{S}_+^\zeta}_{n_g+1 \text{ times}}} \gamma$$

$$\text{subject to } f(x) - \gamma = Z_d(x)^\top S_0 Z_d(x) - \sum_{i \in [n_g]} Z_d(x)^\top S_i Z_d(x) g_i(x)$$

Semidefinite program (SDP) \rightsquigarrow software SOSTOOLS, YALMIP, SOSOPT

Back to Forward Problem

Forward problem:

Given $c, \mathcal{X}, \mathcal{Y}$ as per A1, certify / falsify if the ground cost satisfies either MTW(0) or MTW(κ) or NNCC condition

$$\mathcal{X} \times \mathcal{Y} = \{(x, y) \in \mathbb{R}^n \times \mathbb{R}^n \mid m_i(x, y) \leq 0, m_i(x, y) \in \mathbb{R}_{d_m}[x, y] \forall i \in [\ell]\}$$

NNCC forward problem:

$$\begin{aligned} & \min \quad 0 \\ & \text{subject to} \quad \mathfrak{S}_{(x,y)}(\xi, \eta) \geq 0, \quad \forall (x, y) \in \mathcal{X} \times \mathcal{Y}, \xi \in T_x \mathcal{X}, \eta \in T_y^* \mathcal{Y} \end{aligned}$$

MTW(κ) forward problem:

$$\begin{aligned} & \min \quad 0 \\ & \text{subject to} \quad \mathfrak{S}_{(x,y)}(\xi, \eta) \geq \kappa \|\xi\|^2 \|\eta\|^2, \\ & \quad \quad \quad \forall (x, y) \in \mathcal{X} \times \mathcal{Y}, \xi \in T_x \mathcal{X}, \eta \in T_y^* \mathcal{Y} \text{ s.t. } \eta(\xi) = 0. \end{aligned}$$

Solution to (NNCC) Forward Problem

$$T_x \mathcal{X}, T_y^* \mathcal{Y} \cong \mathbb{R}^n \quad \Longrightarrow \quad \mathfrak{S}_{(x,y)}(\xi, \eta) = (\xi \otimes \eta)^\top \overbrace{F(x,y)}^{\in \mathbb{R}^{n^2 \times n^2}} (\xi \otimes \eta)$$

$$[F(x,y)]_{i+n(j-1), k+n(l-1)} = \sum_{p,q,r,s} (c_{ij,p} c^{p,q} c_{q,rs} - c_{ij,rs}) c^{r,k} c^{s,l}$$

Thm: SOS Tightening of NNCC Forward Problem \rightsquigarrow SDP

For $F = \frac{F_N}{F_D} \in \mathbb{R}_{N,D}[x,y]$, $N, D \in \mathbb{N}$, if $\exists s_0, s_1, \dots, s_\ell \in \sum_{\text{SOS}}^{n^2} [x,y]$ such that

$$(F_N(x,y) + F_N^\top(x,y)) - s_0(x,y)F_D(x,y) + \sum_{i \in [\ell]} s_i(x,y)m_i(x,y) \in \sum_{\text{SOS}}^{n^2} [x,y]$$

then c satisfies NNCC condition on $\mathcal{X} \times \mathcal{Y}$

Computational Complexity: Forward Problem

Parameters: $\omega \in [2.376, 3]$

$$F = \frac{F_N}{F_D} \in \mathbb{R}_{N,D}[x, y], N, D \in \mathbb{N} \quad \Rightarrow \quad [F]_{i,j} \in \mathbb{R}_{(n^4-1)d_D+d_N, n^4d_D}[x, y]$$

$$d_N, d_D = \mathcal{O}(ND)$$

$\ell = \#$ of polynomial constraints defining $\mathcal{X} \times \mathcal{Y}$ semialgebraic

NNCC complexity: $\mathcal{O} \left(\ell^{5/4} n^{9+5d_N/4} + n^{\omega(4+d_N)} + \ell^{\omega/2} n^{\omega(2+d_N/2)} \right)$

MTW(κ) complexity: $\mathcal{O} \left(\ell^{5/4} n^{9d_N/4} + \ell^{\omega/2+1/4} n^{(\omega/2+1/4)d_N} \right)$

Sub-quadratic in ℓ , polynomial in n

Solution to (NNCC) Inverse Problem

Thm: SOS Tightening of NNCC Inverse Problem \rightsquigarrow SDP

For compact $\Lambda := \{(x, y) \in \mathcal{X} \times \mathcal{Y} \mid \lambda(x, y) \leq 0, \lambda(x, y) \in \mathbb{R}_{d_\lambda}[x, y], d_\lambda \in \mathbb{N}\}$
 chosen a priori, let $V_\pm : \Lambda \mapsto \mathbb{R}$ solve

$$\begin{aligned} & \min_{V \in \mathbb{R}_d[x, y]} \int_{\Lambda} V(x, y) dx dy, \\ \text{subject to } & V(x, y) - m_i(x, y) + r_i(x, y)\lambda(x, y) \in \sum_{\text{SOS}} [x, y], \quad \forall i \in \llbracket \ell \rrbracket, \\ & V(x, y) \pm F_D(x, y) + s_0(x, y)\lambda(x, y) \in \sum_{\text{SOS}} [x, y], \\ & V(x, y) \pm f_j(x, y) + s_j(x, y)\lambda(x, y) \in \sum_{\text{SOS}} [x, y], \quad \forall j \in \llbracket |\text{pminor}(F_N)| \rrbracket, \\ & s_0(x, y), s_j(x, y), r_i(x, y) \in \sum_{\text{SOS}} [x, y] \quad \forall i \in \llbracket \ell \rrbracket, j \in \llbracket |\text{pminor}(F_N)| \rrbracket, \end{aligned}$$

then c satisfies NNCC on $\{(x, y) \in \Lambda \mid V_+(x, y) \leq 0\} \cup \{(x, y) \in \Lambda \mid V_-(x, y) \leq 0\}$

Numerical Results: Forward Problem

Example 1: Perturbed Euclidean Cost

$$c(x, y) = \|x - y\|_2^2 - \varepsilon \|x - y\|_2^4, \quad x, y \in \mathbb{R}^n, \varepsilon > 0$$

Lee & Li (2009): for ε small enough, MTW(0) holds on

$$\mathcal{X} \times \mathcal{Y} := \{(x, y) \in \mathbb{R}^n \times \mathbb{R}^n \mid \|x - y\|_2 \leq 0.5\}$$

But how small is small enough?

We used SOS SDP + bisection to find ε_{\max} such that MTW(0) holds

Dimensions, n	1	2
ε_{\max}	$0.67 (\approx \frac{2}{3})$	$1.05 \cdot 10^{-2}$
Residual	$1.19 \cdot 10^{-7}$	$4.18 \cdot 10^{-7}$

No analytical computation known

in Lee & Li (2009) via nontrivial analysis

Numerical Results: Forward Problem

Example 2: Log Partition Cost (Pal & Wong, 2018; Khan & Zhang, 2020)

used in stochastic portfolio theory

$$c(x, y) = \Psi_{\text{IsoMulNor}}(x - y), \quad \Psi_{\text{IsoMulNor}}(x) := \frac{1}{2} \overbrace{\left(-\log x_1 + \sum_{i=2}^n x_i^2 / x_1 \right)}^{\text{not rational!}}$$

Our method still applies because $\mathfrak{S}_{(x,y)}(\xi, \eta) \propto \mathfrak{A}_x(\xi, \eta) = \text{poly}(x, \xi, \eta) / x_1^2 \xi_n^2$

No analytical method exists to verify MTW(0) for $n > 2$

We can verify MTW(0) for $\mathcal{X} = \mathcal{Y} = \{x \in \mathbb{R}^n \mid x_1 > 0\} \forall n \in \mathbb{N}$

Dimensions, n	3	4	5	6
Residual	$1.034 \cdot 10^{-7}$	$4.804 \cdot 10^{-8}$	$4.683 \cdot 10^{-8}$	$3.475 \cdot 10^{-11}$
Total time (s)	0.7220	0.8050	1.2520	1.6690

Numerical Results: Forward Problem

Example 2: Log Partition Cost (Pal & Wong, 2018; Khan & Zhang, 2020)

For $n = 3$, our method discovered SOS decomposition:

$$\text{poly}(x, \xi, \eta) = s(x, \xi, \eta)^\top s(x, \xi, \eta)$$

$$\text{where } s(x, \xi, \eta) = \begin{bmatrix} 0 & -1.4 & 0 & 0.24 & 0 & 0 \\ 2.4 & 0 & -0.17 & 0 & 0 & 0 \\ 0 & 1.4 & 0 & -0.24 & 0 & 0 \\ -2.4 & 0 & 0.17 & 0 & -0.0002 & 0 \\ 0 & -1.4 & 0 & 0.25 & 0 & -1.2 \\ 2.4 & 0 & -0.17 & 0 & -0.0002 & 0 \\ -1.6 & 0 & -1.9 & 0 & 0 & 0 \\ 0 & 0.52 & 0 & 1.3 & 0 & 0 \\ 0 & 1.4 & 0 & -0.25 & 0 & 0 \\ -0.84 & 0 & 2 & 0 & 0 & 0 \\ 0 & -0.52 & 0 & -1.3 & 0 & 0 \end{bmatrix} \begin{bmatrix} \eta_1 \xi_1^2 \xi_2 \\ \eta_1 \xi_1^2 \xi_3 \\ \eta_1 \xi_1 \xi_2 \xi_1 \\ \eta_1 \xi_1 \xi_3 \xi_1 \\ \eta_2 \xi_1 \xi_2 \xi_2 \\ \eta_2 \xi_1 \xi_2 \xi_3 \end{bmatrix}$$

Numerical Results: Inverse Problem

Example 3: Perturbed Euclidean Cost Revisited

$$c(x, y) = \|x - y\|_2^2 - \varepsilon \|x - y\|_2^4, \quad x, y \in \mathbb{R}^n, \varepsilon > 0$$

Lee & Li (2009): for $\varepsilon \gg 0$, MTW(0) fails on


$$\mathcal{X} \times \mathcal{Y} := \{(x, y) \in \mathbb{R}^n \times \mathbb{R}^n \mid \|x - y\|_2 \leq 0.5\}$$

Fix $\varepsilon = 1$, $\Lambda = [-1, 1]^2$, $\mathcal{X} = \{[0, 0]\}$

Parameterize $(\xi, \eta) = ([a, 1]^\top, [-1, a]^\top)$

Solve MTW(0) inverse problem

Exec time 115 sec, solve time 0.97 sec



Inner approx. of region
where MTW tensor ≥ 0

Numerical Results: Inverse Problem

Example 4: Squared Distance Cost for a Surface of Positive Curvature

$$c(x, y) = 3(x_1 - y_1)^2(x_2 + y_2) + 4(x_2^3 + y_2^3) - (4x_2y_2 - (x_1 - y_1)^2)^{\frac{3}{2}}$$


MTW holds around $\{x = y\}$

Fix $\Lambda = [-1, 1] \times [0, 2]$, $\mathcal{X} \times \mathcal{Y} = \{[0, 1]\} \times \{(y_1, y_2) \in [-1, 2] \times [0, 2] \mid 4y_2 - y_1^2 \geq 0\}$

Parameterize $(\xi, \eta) = ([a, 1]^\top, [-1, a]^\top)$

Solve MTW(0) inverse problem

Exec time 119 sec, solve time 19.6 sec




Inner approx. of region
where MTW tensor ≥ 0


Recap: SOS Programming for OT Regularity

NNCC, MTW(κ), MTW(0) conditions \Rightarrow regularity of the OT map τ_{opt}

Assumption A1: MTW tensor is rational in $(x, y) \in \mathcal{X} \times \mathcal{Y}$ semialgebraic



SOS tightening of forward & inverse problems



Solve SDP using SOSTOOLS + YALMIP \rightsquigarrow computational certificates

**Tensor Optimization for Graph-structured
Multi-marginal SB**

&

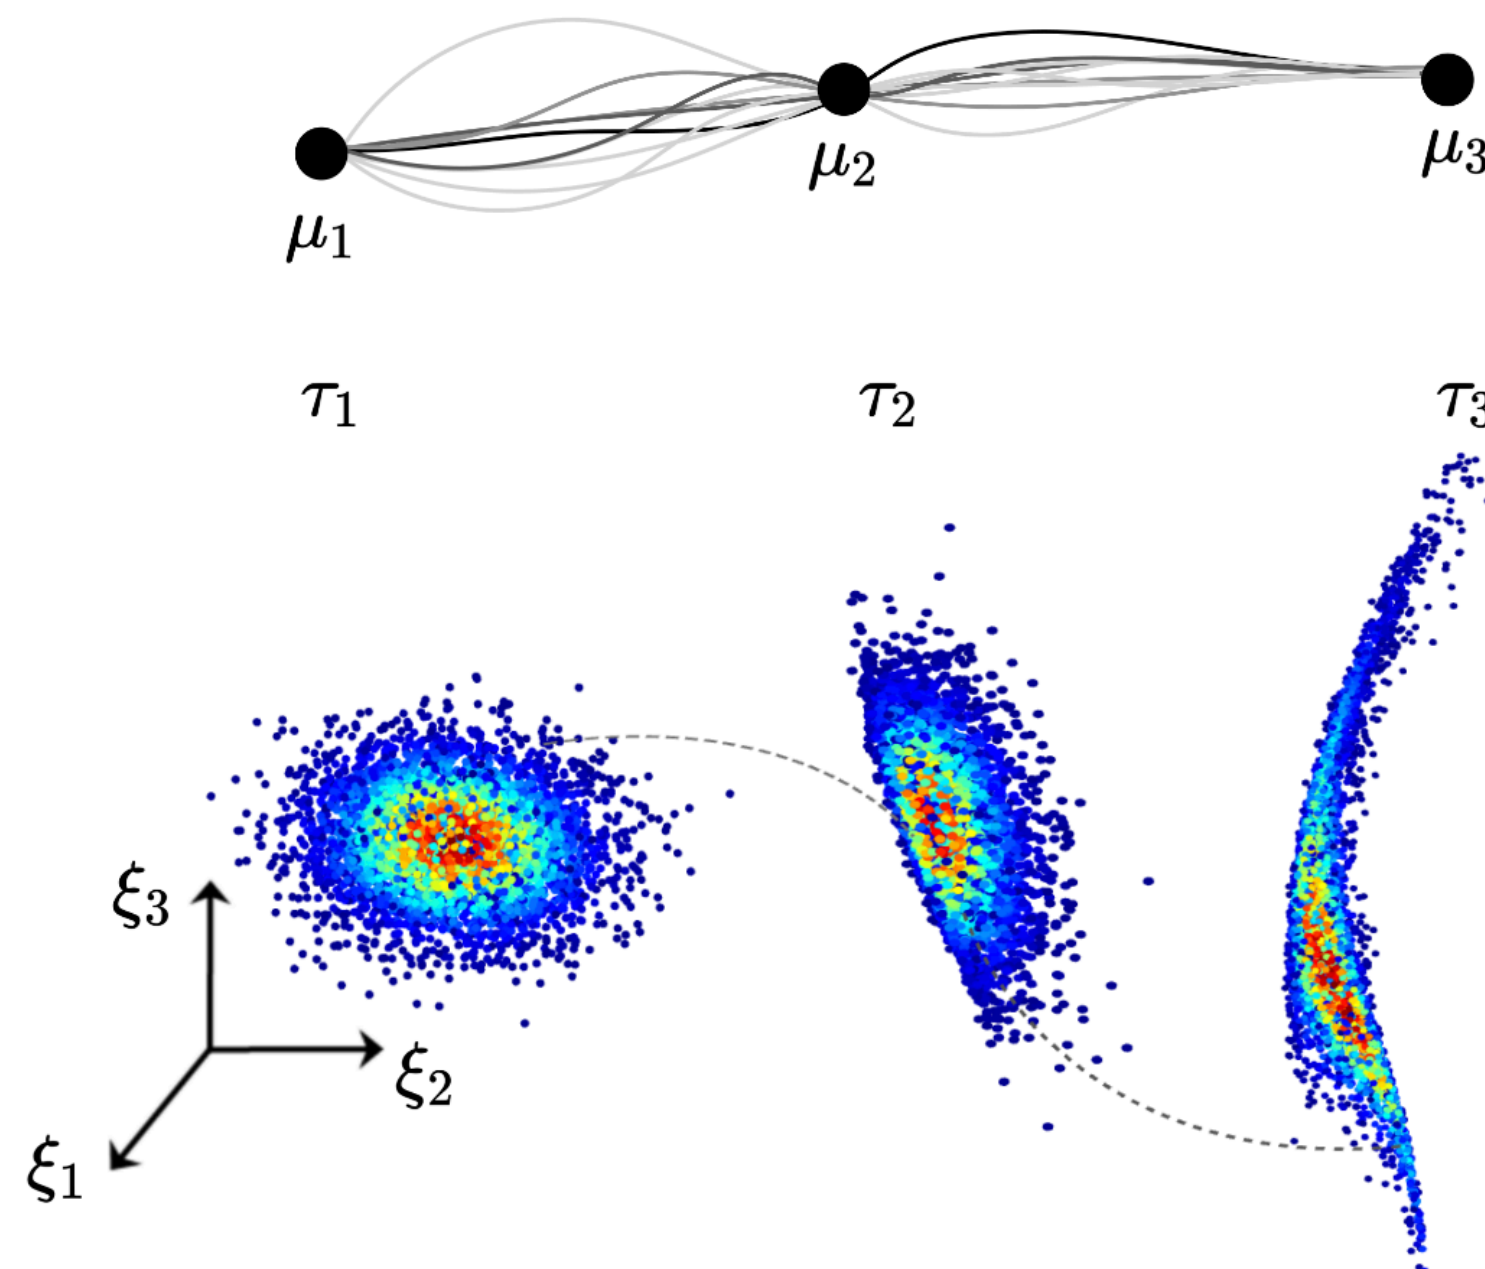
Learning Computational Resource Usage

Motivation: Computational Resource Usage of Multicore Software

Software running on $J \in \mathbb{N}$ CPU cores

Resource usage stochastic process $\xi(\tau) \sim \mu_\tau$

Example: $\xi^j := \begin{pmatrix} \xi_1^j \\ \xi_2^j \\ \xi_3^j \end{pmatrix} = \begin{pmatrix} \text{instructions retired} \\ \text{LLC requests} \\ \text{LLC misses} \end{pmatrix} \quad \forall j \in \llbracket J \rrbracket$



“Profiling” in RTOS community: sample $\underbrace{\{\xi^{i,j}(\tau_\sigma)\}_{i=1}^n}_{\text{scattered data}} \quad \forall \sigma \in \llbracket s \rrbracket$ where

Time/resource intensive!

$$\tau_1 \equiv 0 < \tau_2 < \dots < \tau_{s-1} < \tau_s \equiv t, \quad s \geq 2$$

Problem Formulation




Use (weighted scattered) profile data

$$\{\boldsymbol{\xi}^{i,j}(\tau_\sigma)\}_{i=1}^n, \mu_\sigma^j := \frac{1}{n} \sum_{i=1}^n \delta(\boldsymbol{\xi}^j - \boldsymbol{\xi}^{i,j}(\tau_\sigma)) \quad \forall (j, \sigma) \in [\mathcal{J}] \times [\mathcal{S}]$$

to learn $\hat{\mu}_\tau \quad \forall \tau \in [0, t]$

Typical profiles

$(\tau = 4, J = 1)$:



Challenges

Difficult to have first-principle physics based model for combined $S/W+H/W$ level stochasticity

Learning must be over joint resources (e.g., processor & cache correlated)

Correlation structure among resource states changes with time

Need: nonparametric learning, also desire: learning with optimality

Main Idea

Step 1: Model the spatio-temporal correlation induced by HW+SW architecture by graph structures

Step 2: Solve MSBP over the resulting graph

Step 3: Use the MSBP solution to predict most likely $\hat{\mu}_\tau$

Steps 1,2: Discrete Graph-structured MSBP

Problem template: $\arg \min \langle \mathbf{C} + \varepsilon \log \mathbf{M}, \mathbf{M} \rangle$

$$\mathbf{M} \in (\mathbb{R}^n)_{\geq 0}^{\otimes |\Lambda|}$$

index set capturing graph structure

$$\text{subject to } \text{proj}_{(j,\sigma)}(\mathbf{M}) = \boldsymbol{\mu}_\sigma^j \quad \forall (j,\sigma) \in \Lambda$$

Prop: (Strong duality \rightsquigarrow Sinkhorn recursions, **complexity:** $\mathcal{O}(n^{|\Lambda|})$)

Lagrange multipliers

$$\text{Let } \underbrace{\mathbf{K} := \exp(-\mathbf{C}/\varepsilon)}_{\in (\mathbb{R}^n)_{>0}^{\otimes |\Lambda|}}, \underbrace{\mathbf{u}_\sigma^j := \exp(\boldsymbol{\lambda}_\sigma^j/\varepsilon)}_{\in \mathbb{R}_{>0}^n \quad \forall (j,\sigma) \in \Lambda}, \underbrace{\mathbf{U} := \bigotimes_{(j,\sigma) \in \Lambda} \mathbf{u}_\sigma^j}_{\in (\mathbb{R}^n)_{>0}^{\otimes |\Lambda|}}$$


The multi-marginal Sinkhorn recursions

$$\mathbf{u}_\sigma^j \leftarrow \mathbf{u}_\sigma^j \odot \boldsymbol{\mu}_\sigma^j \oslash \text{proj}_{(j,\sigma)}(\mathbf{K} \odot \mathbf{U}) \quad \forall (j,\sigma) \in \Lambda,$$

converges with linear rate to minimizer $\mathbf{M}^{\text{opt}} = \mathbf{K} \odot \mathbf{U}$

$J = 1$: Single CPU Core: Path-structured MSBP

Correlation induced by **time**



Ground cost tensor decomposes: $\mathbf{C}(\boldsymbol{\xi}(\tau_1), \dots, \boldsymbol{\xi}(\tau_s)) = \sum_{\sigma=1}^{s-1} c_\sigma(\boldsymbol{\xi}(\tau_\sigma), \boldsymbol{\xi}(\tau_{\sigma+1}))$


Discrete version: $[\mathbf{C}]_{i_1, \dots, i_s} = \sum_{\sigma=1}^{s-1} [C^\sigma]_{i_\sigma, i_{\sigma+1}}$

$J > 1$: Multiple CPU Cores

Correlation induced by **time + CPU cores**

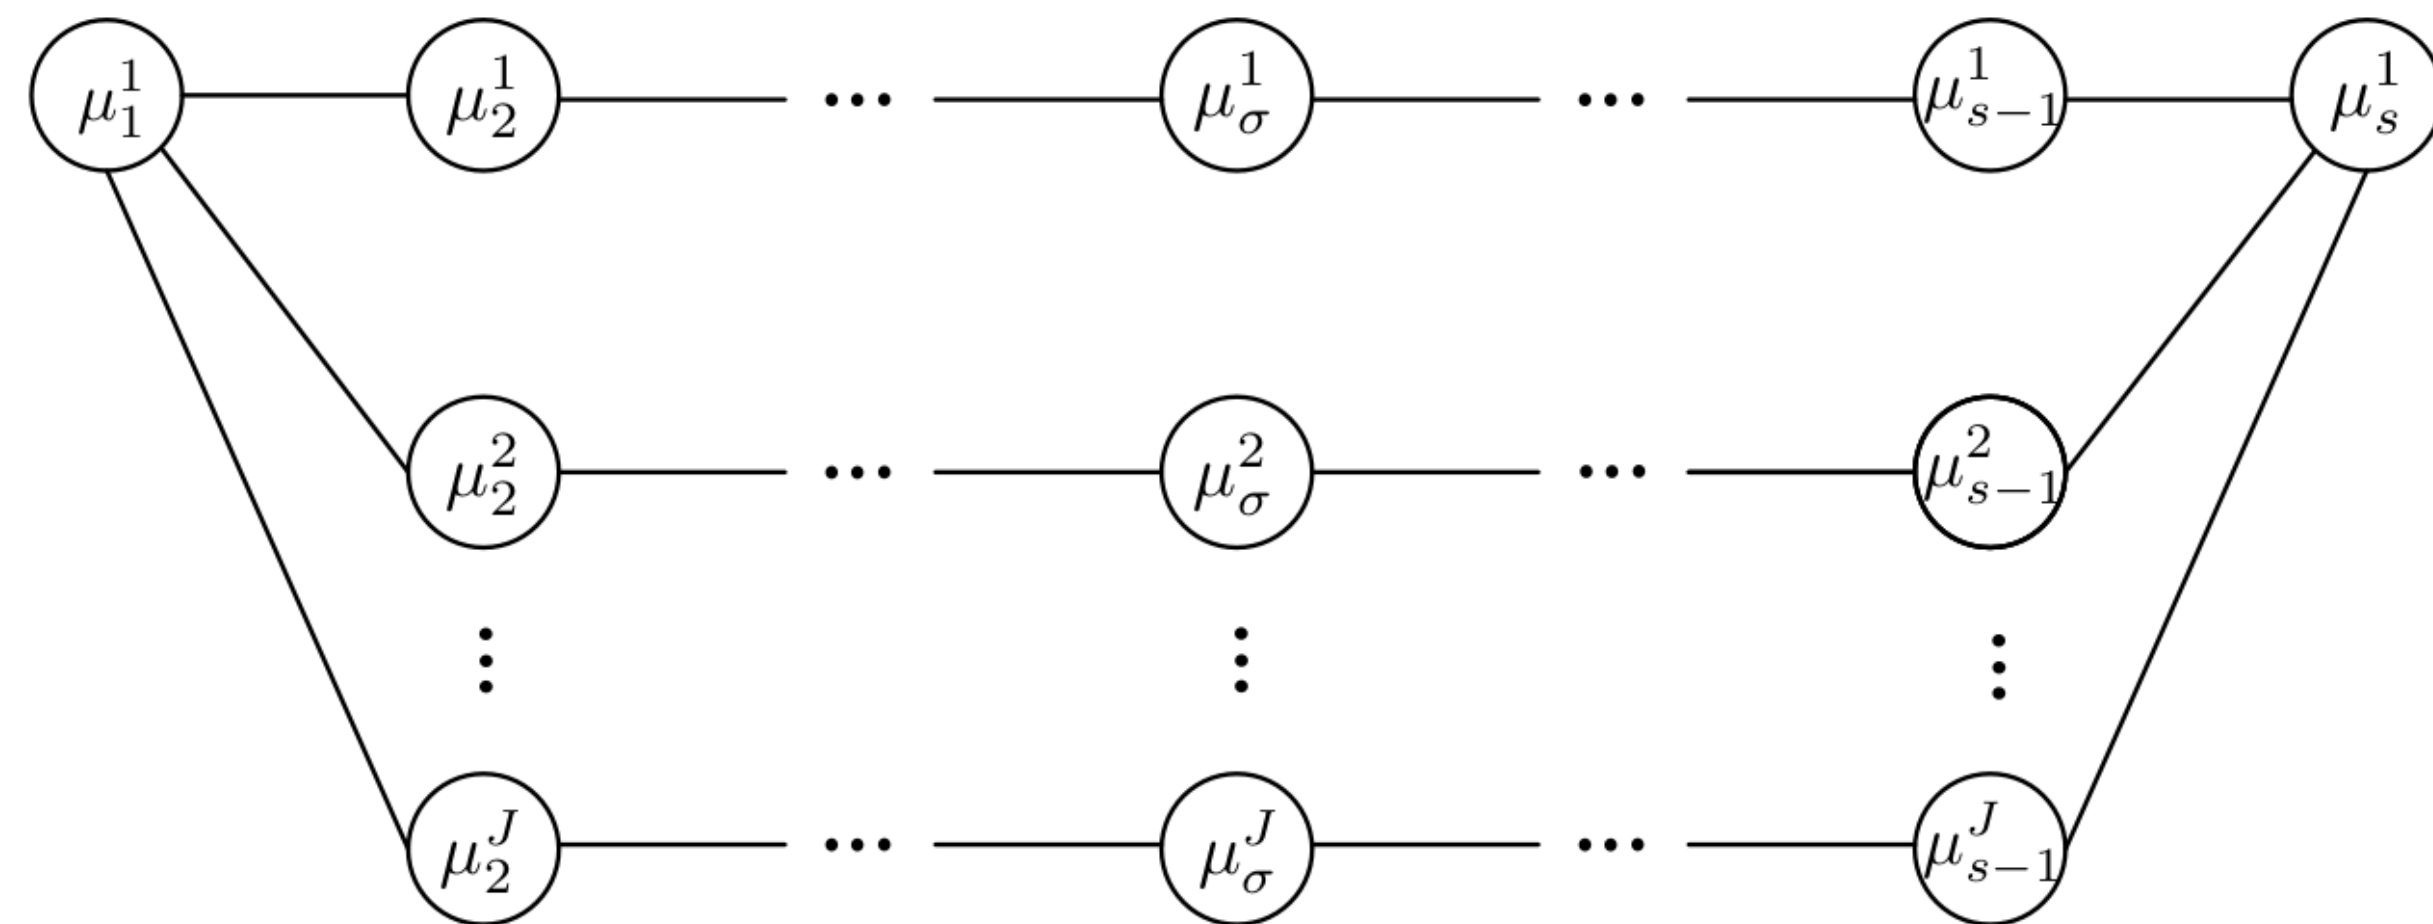
Graph structure:

barycentric (BC)



inter-CPU communication

series-parallel (SP)




parallel execution

$J > 1$: Multiple CPU Cores: Barycentric MSBP

Idea: **phantom CPU** resource statistics $\mu_\sigma^0 = \text{barycenter of } \{\mu_\sigma^j\}_{j \in [J]} \forall \sigma \in \llbracket s \rrbracket$

Index set:

$$\Lambda_{\text{BC}} := (\{0\} \cup \llbracket J \rrbracket) \times \llbracket s \rrbracket$$



Ground cost tensor decomposition:

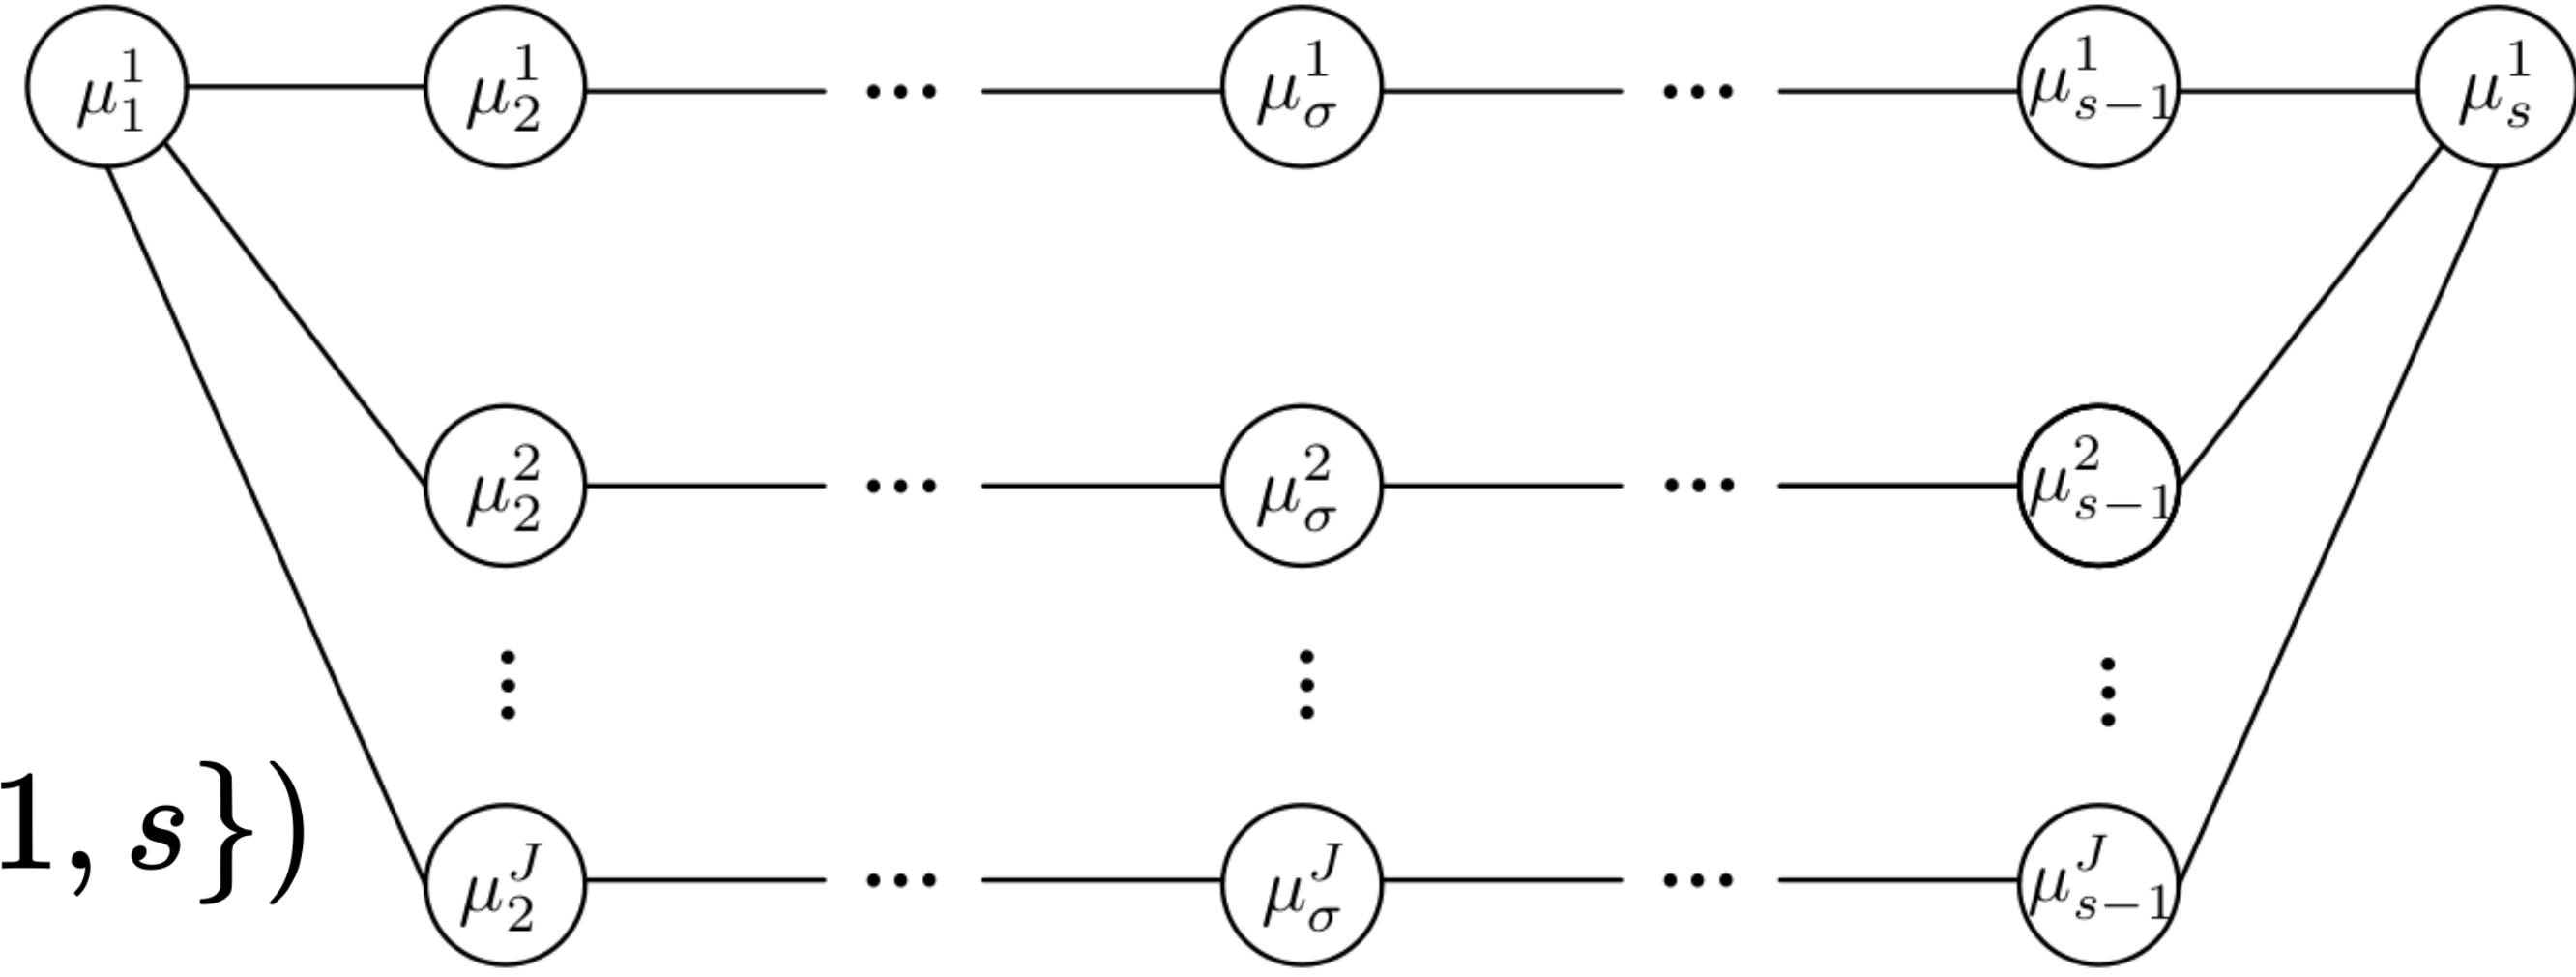
$$\mathbf{C}(\boldsymbol{\xi}(\tau_1), \dots, \boldsymbol{\xi}(\tau_s)) = \sum_{\sigma=1}^{s-1} c_{0,\sigma} (\boldsymbol{\xi}^0(\tau_\sigma), \boldsymbol{\xi}^0(\tau_{\sigma+1})) + \sum_{\sigma=1}^s \sum_{j=1}^J c_{j,\sigma} (\boldsymbol{\xi}^j(\tau_\sigma), \boldsymbol{\xi}^0(\tau_\sigma))$$

$J > 1$: Multiple CPU Cores: Series-parallel MSBP

Idea: fork and merge

Index set:

$$\Lambda_{\text{SP}} := (\llbracket J \rrbracket \times \llbracket s \rrbracket) \setminus ((\llbracket J \rrbracket \setminus \{1\}) \times \{1, s\})$$



Ground cost tensor decomposition:

$$C(\boldsymbol{\xi}(\tau_1), \dots, \boldsymbol{\xi}(\tau_s)) = \sum_{j=1}^J \left\{ c_{j,1}(\boldsymbol{\xi}^j(\tau_1), \boldsymbol{\xi}^j(\tau_2)) + c_{j,s-1}(\boldsymbol{\xi}^j(\tau_{s-1}), \boldsymbol{\xi}^j(\tau_s)) \right\} + \sum_{\sigma=2}^{s-1} \sum_{j=1}^J c_{j,\sigma}(\boldsymbol{\xi}^j(\tau_\sigma), \boldsymbol{\xi}^j(\tau_{\sigma+1}))$$

$J > 1$: Computational Complexity for MSBP

linear in $J, s \rightarrow$

Structure	General	Path	BC	SP
Index set	Λ	$\llbracket s \rrbracket$	Λ_{BC}	Λ_{SP}
# of indices	$ \Lambda $	s	$(J+1)s$	$J(s-2)+2$
$\mathcal{O}(\cdot)$ for $\text{proj}_\sigma(\mathbf{M})$	$n^{ \Lambda }$	$(s-1)n^2$	$(Js)n^2$	$(Js)n^3$

Exact flop count for BC:
$$\underbrace{Js(n_0n + n_0) + (2n_0) + (2s - 2)n_0^2}_{\text{barycenter}}$$

$$\underbrace{Js(n_0n + n_0) + (3n_0 + n + 2n_0n) + (2s - 2)n_0^2}_{\text{other dist.}}$$

Exact flop count for SP:
$$\underbrace{J(1 + 2(s - 2))n^3 + (J + 1)n^2 + n}_{\text{end dist.}}$$

$$\underbrace{(J(1 + 2(s - 2)) + 3)n^3 + (J - 1)n^2 + n}_{\text{other dist.}}$$

Step 3: MSBP Solution to Predicting $\hat{\mu}_\tau$

Given $\tau \in [0, t)$, $j \in \llbracket J \rrbracket$, find $\sigma \in \{\llbracket s \rrbracket \mid \tau_\sigma \leq \tau < \tau_{\sigma+1}\}$

$$M^{j,\sigma} := \text{proj}_{(j,\sigma),(j,\sigma+1)}(\mathbf{M}^{\text{opt}}) : \mu_\sigma^j \rightarrow \mu_{\sigma+1}^j \quad (\in \mathbb{R}_{\geq 0}^{n \times n})$$

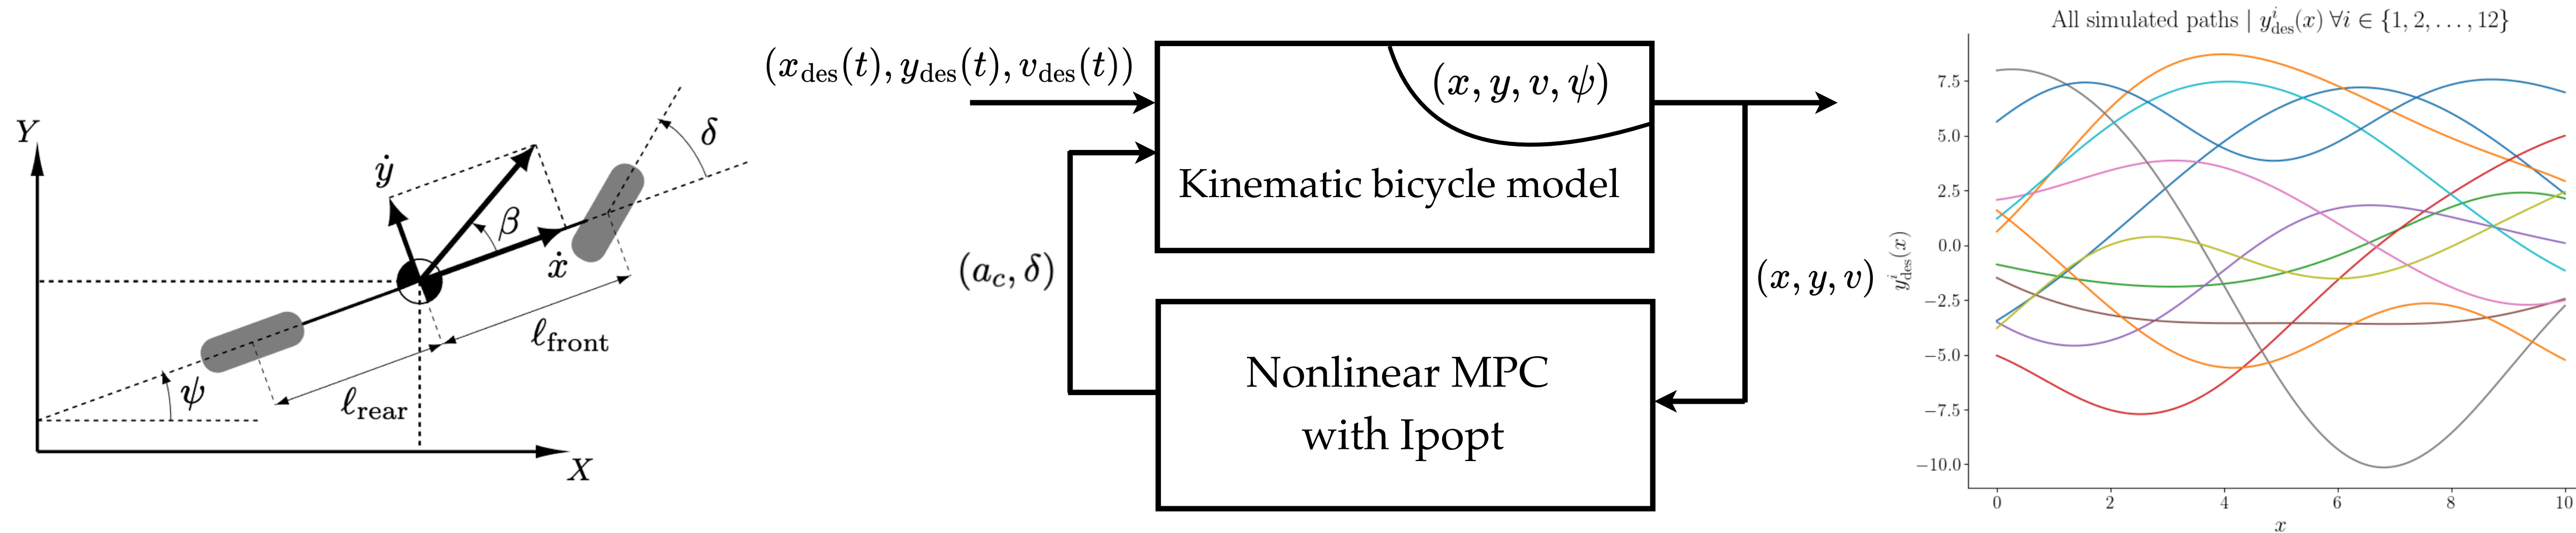
Compute measure interpolating μ_σ^j and $\mu_{\sigma+1}^j$ as:

$$\hat{\mu}_\tau^j := \sum_{r=1}^n \sum_{\ell=1}^n \left[M_{r,\ell}^{j,\sigma} \right] \delta(\boldsymbol{\xi}^j - \hat{\boldsymbol{\xi}}^j(\tau, \boldsymbol{\xi}^{r,j}(\tau_\sigma), \boldsymbol{\xi}^{\ell,j}(\tau_{\sigma+1})))$$

and its support:

$$\hat{\boldsymbol{\xi}}^j(\tau, \boldsymbol{\xi}^{r,j}(\tau_\sigma), \boldsymbol{\xi}^{\ell,j}(\tau_{\sigma+1})) := (1 - \lambda)\boldsymbol{\xi}^{r,j}(\tau_\sigma) + \lambda\boldsymbol{\xi}^{\ell,j}(\tau_{\sigma+1}), \quad \lambda := \frac{\tau - \tau_\sigma}{\tau_{\sigma+1} - \tau_\sigma} \in [0, 1]$$

Case Study: Path Tracking Control Software



$$\mathbf{c}_{\text{cyber}} = \begin{pmatrix} \text{alloc. last-level cache} \\ \text{alloc. memory bandwidth} \end{pmatrix}, \quad \mathbf{c}_{\text{phys}} = \mathbf{y}_{\text{des}}(\mathbf{x}) \in \text{GP}([x_{\min}, x_{\max}])$$

$$\boldsymbol{\xi} := \begin{pmatrix} \xi_1 \\ \xi_2 \\ \xi_3 \end{pmatrix} = \begin{pmatrix} \text{instructions retired} \\ \text{LLC requests} \\ \text{LLC misses} \end{pmatrix}$$

Single core \Rightarrow Path-structured MSBP


Case Study: Path Tracking Control Software

$n = 500$, $\mathbf{c}_{\text{cyber}} = [15 \quad 15]^\top$, $\mathbf{c}_{\text{phys}} = y_{\text{des}}^1(x)$, 30 MB LLC, mem. bandwidth

Each profile with $n_c = 5$ control cycles




Cycle No.	Mean	Std. Dev.
#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

Sampling period = 5 ms



Case Study: Path Tracking Control Software

Profiles:



H/W-level stochasticity, fixed context \mathbf{c}

Case Study: Path Tracking Control Software


MSBP convergence:

of marginals $s := 1 + n_c(s_{\text{int}} + 1)$; Euclidean $C^\sigma \forall \sigma \in \llbracket s - 1 \rrbracket$

$$\text{Cost tensor element: } [\mathbf{C}]_{i_1, \dots, i_s} = \sum_{\sigma=1}^{s-1} [\mathbf{C}]_{i_\sigma, i_{\sigma+1}}^\sigma$$

$$d_{\text{H}}(\mathbf{u}, \mathbf{v}) = \log \left(\frac{\max_{i=1, \dots, n} u_i / v_i}{\min_{i=1, \dots, n} u_i / v_i} \right), \quad \mathbf{u}, \mathbf{v} \in \mathbb{R}_{>0}^n$$

Hilbert projective metric



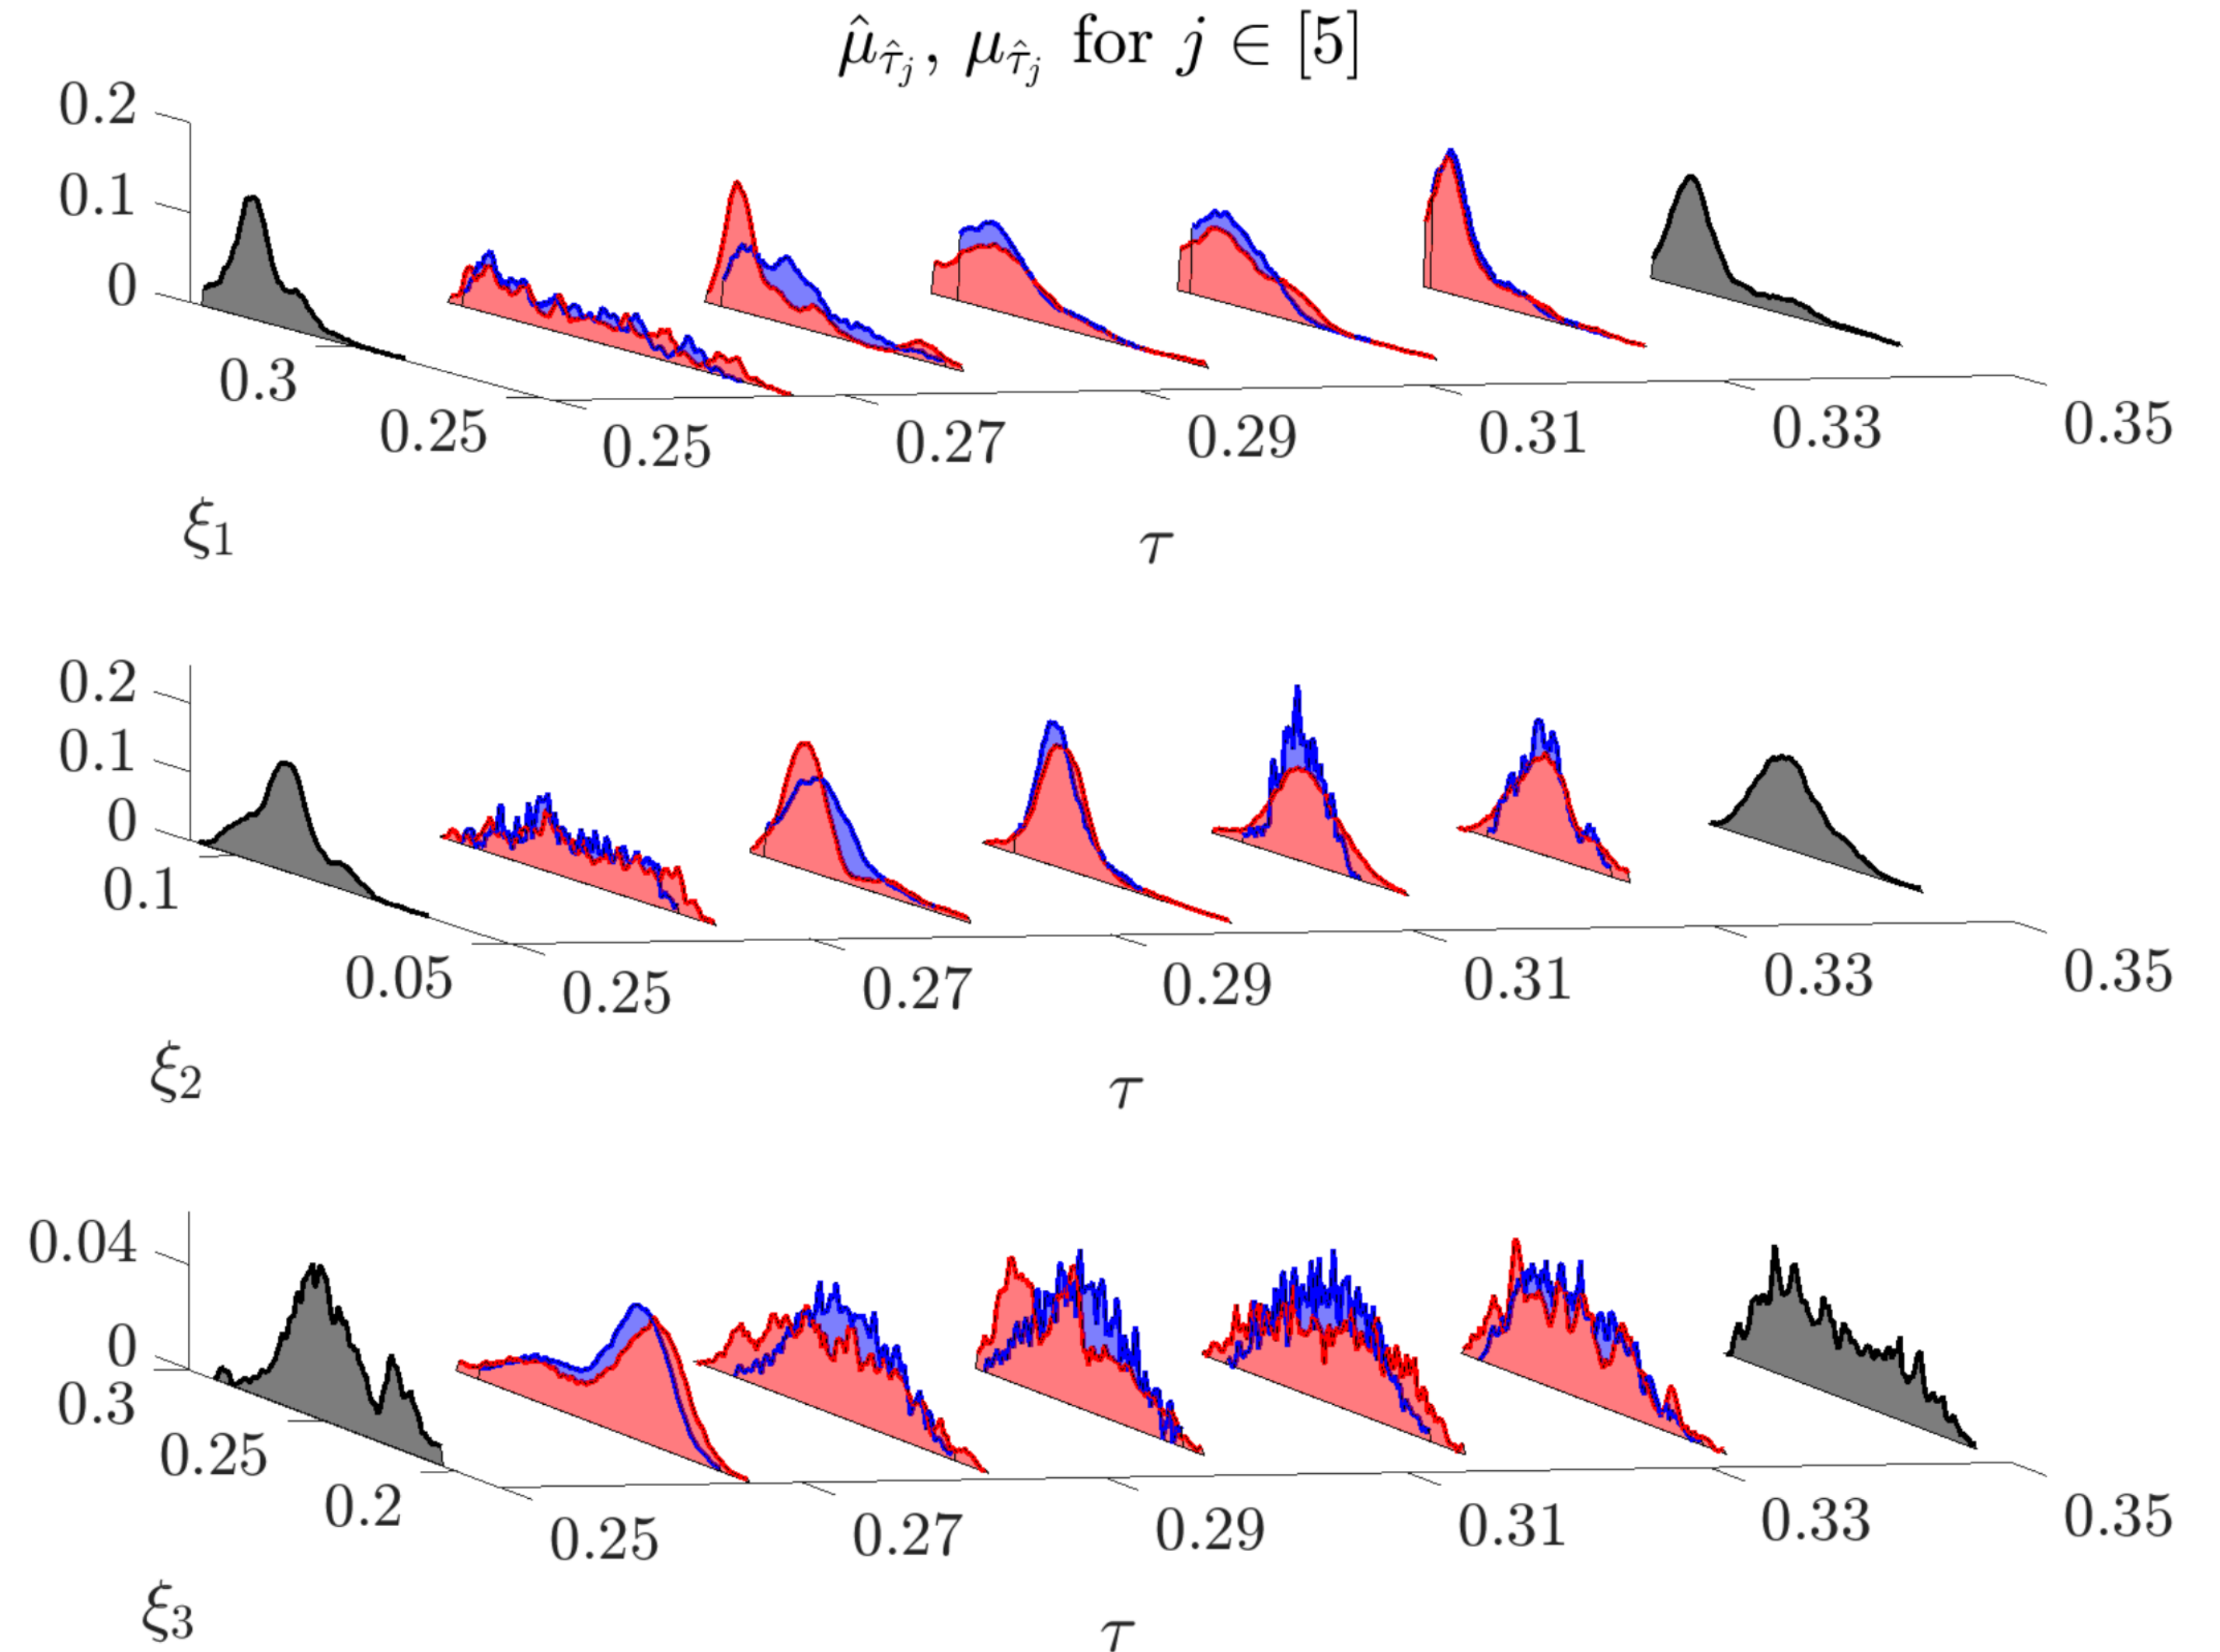
Case Study: Path Tracking Control Software

MSBP prediction vs “hold out” observation, 3rd control cycle, $s_{\text{int}} = 4$:

Predicted $\hat{\mu}$

Measured μ

μ at control-cycle boundaries



Case Study: Path Tracking Control Software

MSBP accuracy:

of intra-cycle marginals Wasserstein distances $W_j := W(\hat{\mu}_{\hat{\tau}_j}, \mu_{\hat{\tau}_j}) \forall j \in \llbracket s_{\text{int}} + 1 \rrbracket$

s_{int}	W_1	W_2	W_3	W_4	W_5
0	2.0489×10^{-4}	-	-	-	-
1	2.2695×10^{-4}	1.1750×10^{-4}	-	-	-
2	5.7717×10^{-4}	0.9163×10^{-4}	0.3794×10^{-4}	-	-
3	2.2413×10^{-4}	1.6432×10^{-4}	1.2345×10^{-4}	0.6010×10^{-4}	-
4	0.6372×10^{-4}	1.2691×10^{-4}	0.9176×10^{-4}	0.6689×10^{-4}	0.2111×10^{-4}

$$\uparrow s_{\text{int}} \implies \downarrow \mathbb{E}[W_j]$$

Case Study: Multicore Benchmark

Canneal: quad-core ($J = 4$) benchmark from PARSEC

$$\mathbf{c}_{\text{cyber}} = \begin{pmatrix} \text{alloc. last-level cache (MB)} \\ \text{alloc. memory bandwidth (MBps)} \end{pmatrix} := \begin{pmatrix} 24 & 10 & 4 & 2 \\ 125 & 25 & 5 & 1 \end{pmatrix}$$


Profiled $n = 400$ times at $\tau_\sigma \in \{0.0, 0.5, 1.5, 2.5, 5.0, 9.5, 10.5\}$, i.e., $s = 7$

Multicore \Rightarrow both BC and SP MSBP

BC: 400^{35} decision variables




SP: 400^{22} decision variables

Convergence in 0.5 s in MATLAB





Case Study: Multicore Benchmark

Profiles:



Case Study: Multicore Benchmark

MSBP convergence:



Case Study: Multicore Benchmark

MSBP accuracy:

BC:

CPU core Wasserstein distances $W_j := W(\hat{\mu}_{\hat{\tau}_j}, \mu_{\hat{\tau}_j}) \forall j \in \llbracket s_{\text{int}} + 1 \rrbracket$

j	W_1^j	W_2^j	W_3^j	W_4^j	W_5^j
1	4.077×10^{-5}	1.009×10^{-7}	2.131×10^{-7}	1.976×10^{-7}	1.509×10^{-7}
2	0	1.135×10^{-7}	2.342×10^{-7}	7.684×10^{-8}	8.805×10^{-8}
3	0	1.149×10^{-7}	1.534×10^{-7}	5.752×10^{-8}	6.538×10^{-8}
4	0	3.647×10^{-8}	2.146×10^{-7}	1.906×10^{-7}	9.713×10^{-8}

SP:

j	W_1^j	W_2^j	W_3^j	W_4^j	W_5^j
1	4.254×10^{-5}	1.020×10^{-7}	2.023×10^{-7}	1.412×10^{-7}	2.589×10^{-7}
2	0	2.386×10^{-7}	2.329×10^{-7}	8.962×10^{-8}	1.908×10^{-7}
3	0	2.392×10^{-7}	1.513×10^{-7}	4.693×10^{-8}	1.100×10^{-7}
4	0	4.868×10^{-8}	2.050×10^{-7}	1.617×10^{-7}	1.204×10^{-7}


Case Study: Context-dependent Resource Usage

Idea: account for software's resource allocation / execution context


$$\beta = (\beta_1, \beta_2, \dots, \beta_b)^\top \in \mathcal{B} \subset \mathbb{R}^b$$

Augment $\eta := [\xi \quad \beta]^\top \in \mathcal{X} \times \mathcal{B} \subset \mathbb{R}^{d+b}$ to form distributions

$$\mu_\sigma := \frac{1}{n_d n_b} \sum_{i=1}^{n_d} \sum_{j=1}^{n_b} \delta(\eta - \eta^{i,j}(\tau_\sigma)), \quad \forall \sigma \in [n_s]$$



Solve path-structured MSBP for μ_τ , $\eta(\tau) \sim \mu_\tau \quad \forall \tau \in [\tau_1, \tau_{n_s}]$



Apply Bayes' theorem to obtain $\xi(\tau) \mid \beta \sim \frac{\mu_\tau}{\int_{\mathcal{X}} \mu_\tau d\xi}$


Profiling: Context-dependent Resource Usage

Benchmarks: dedup, canneal, fft, radiosity


$$N_{\text{ca}} = N_{\text{bw}} = 20 \implies \mathcal{B} = \llbracket N_{\text{ca}} \rrbracket \times \llbracket N_{\text{bw}} \rrbracket$$

Profile over $\mathcal{B}' = \{1, 5, 10, 15, 20\}^2 \subsetneq \mathcal{B}$, $n_b = |\mathcal{B}'| = 25$, $n_d = 10 \quad \forall \beta \in \mathcal{B}'$

$$\tau = 0.05 \cdot (\sigma - 1)$$

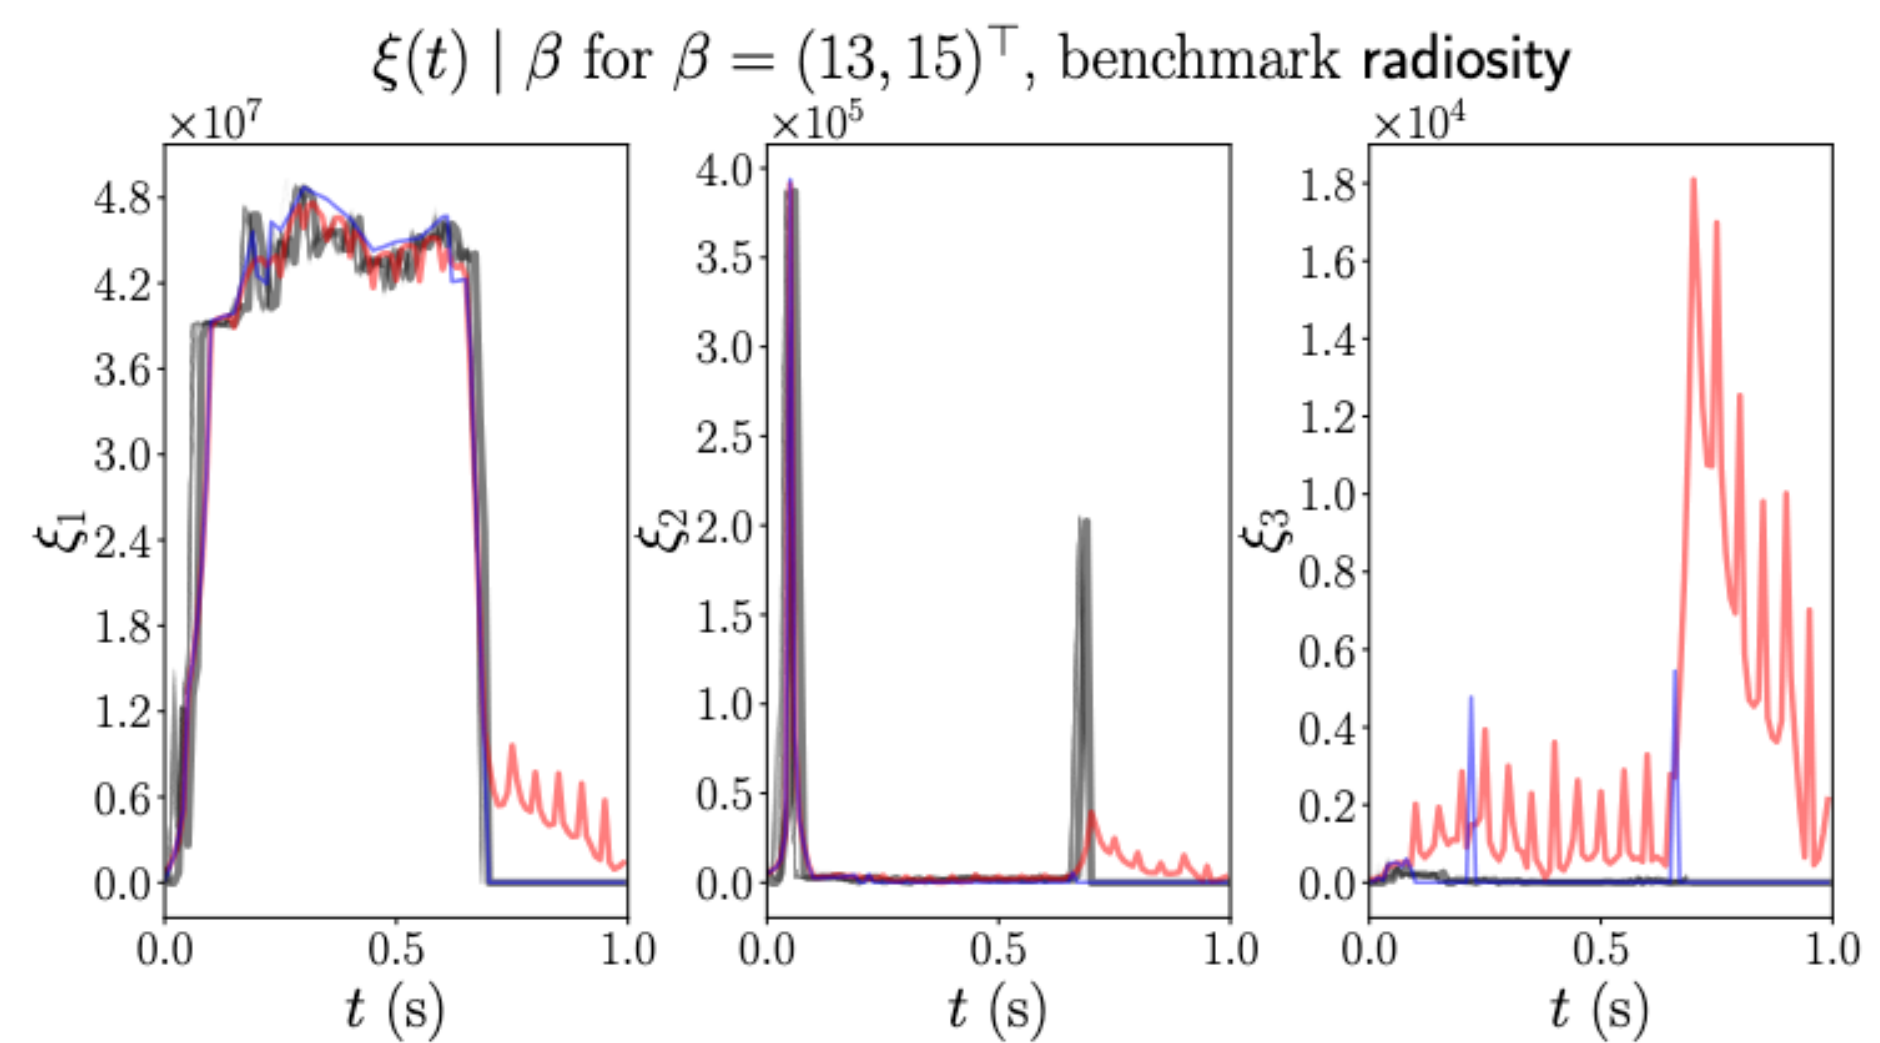
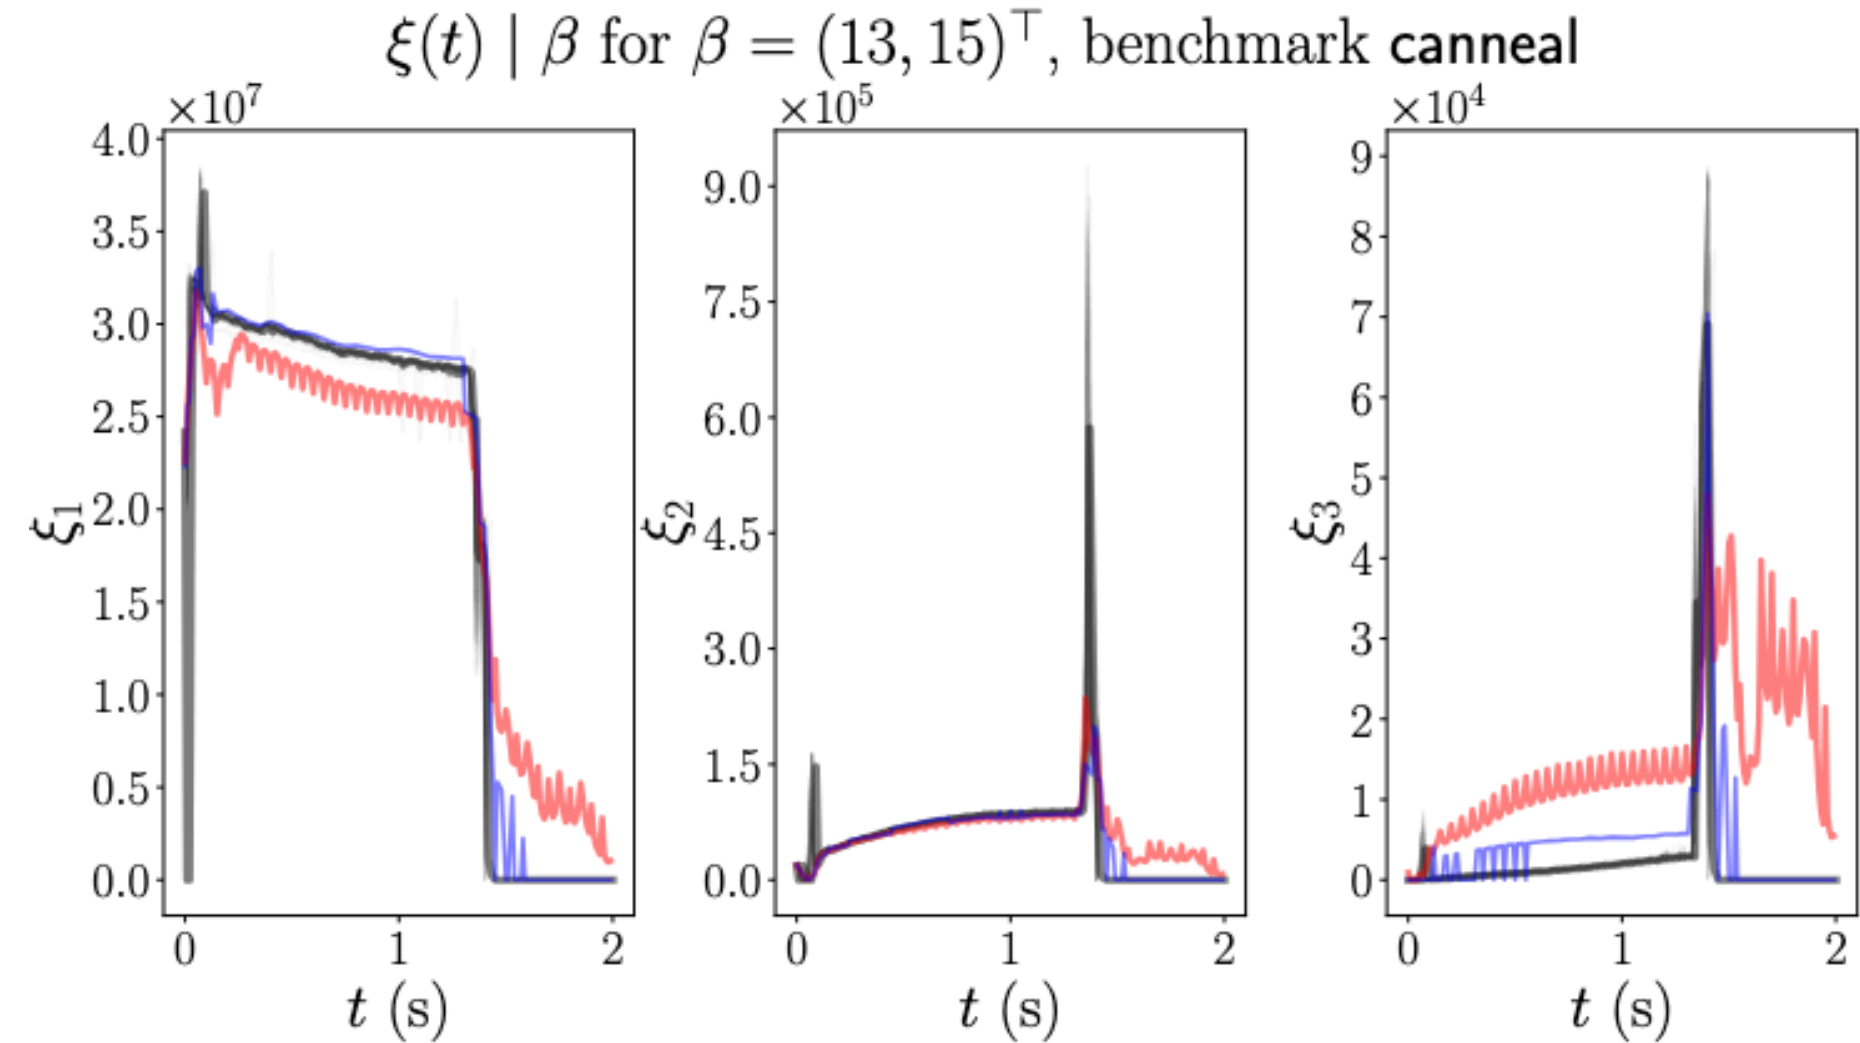
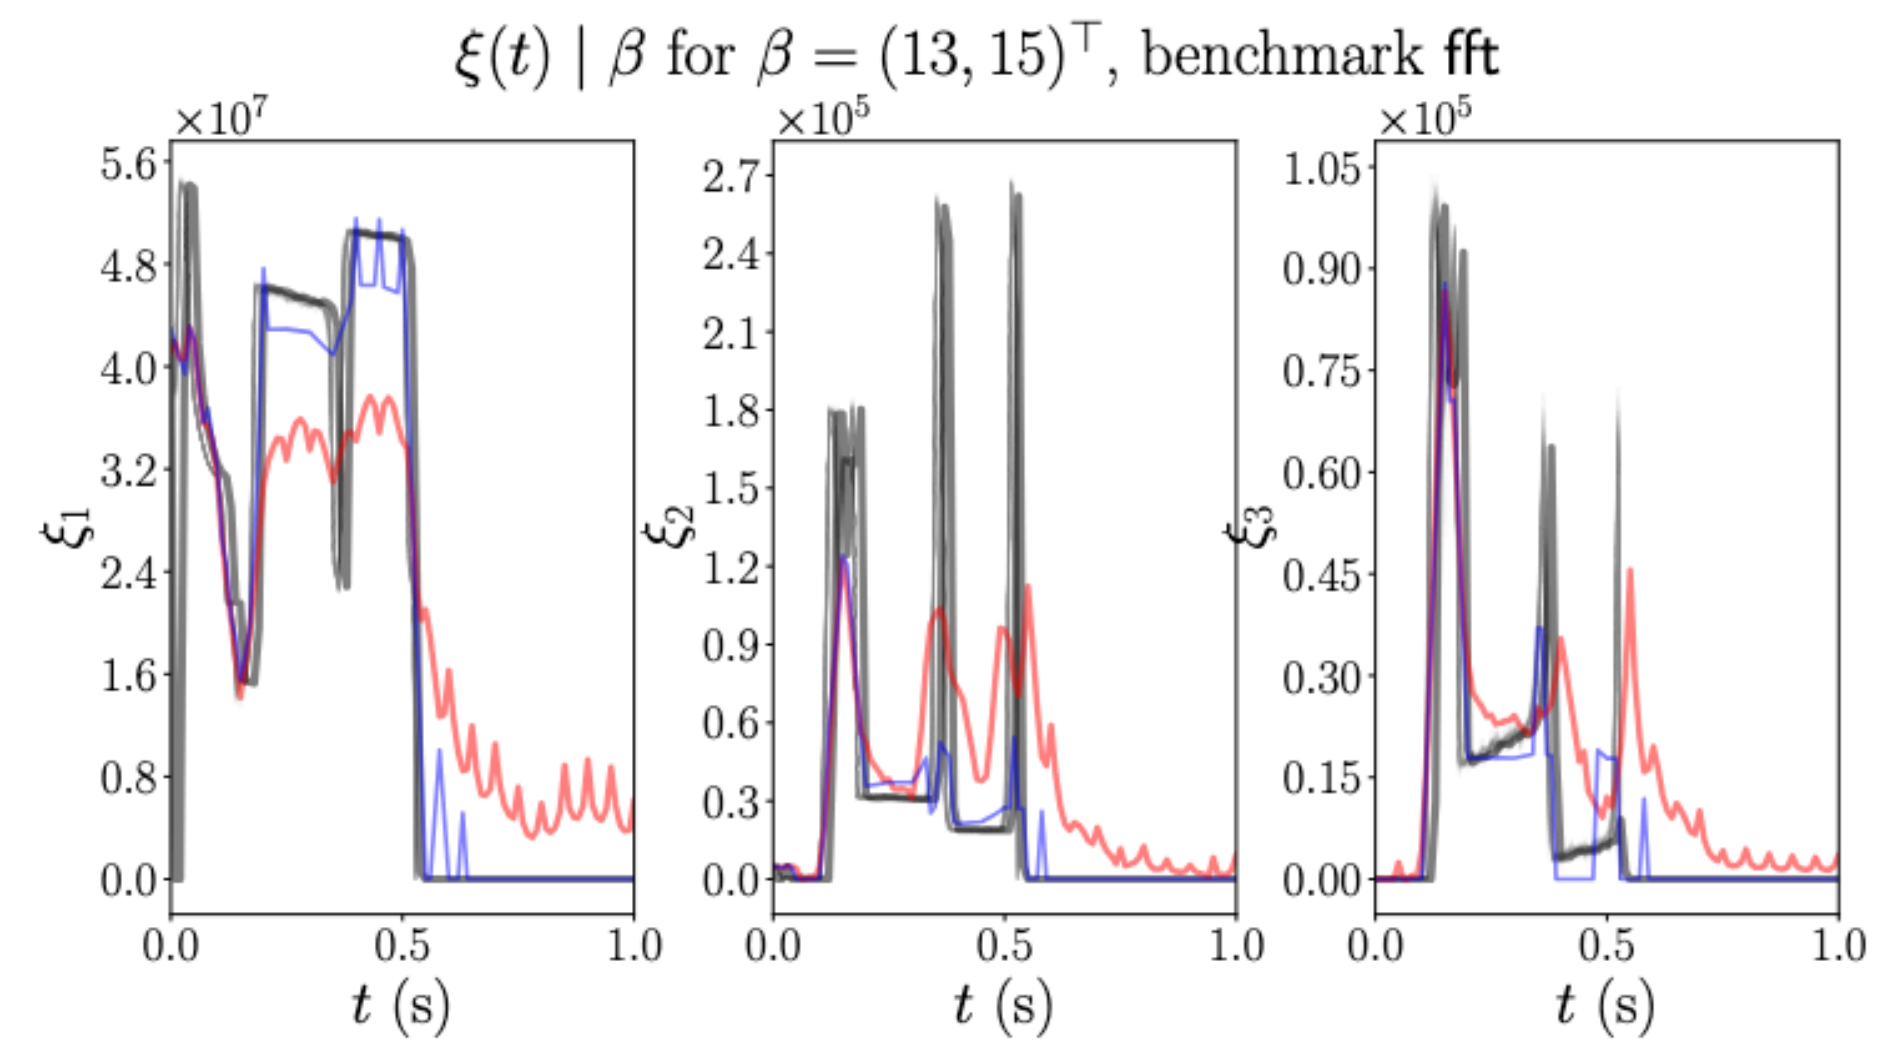
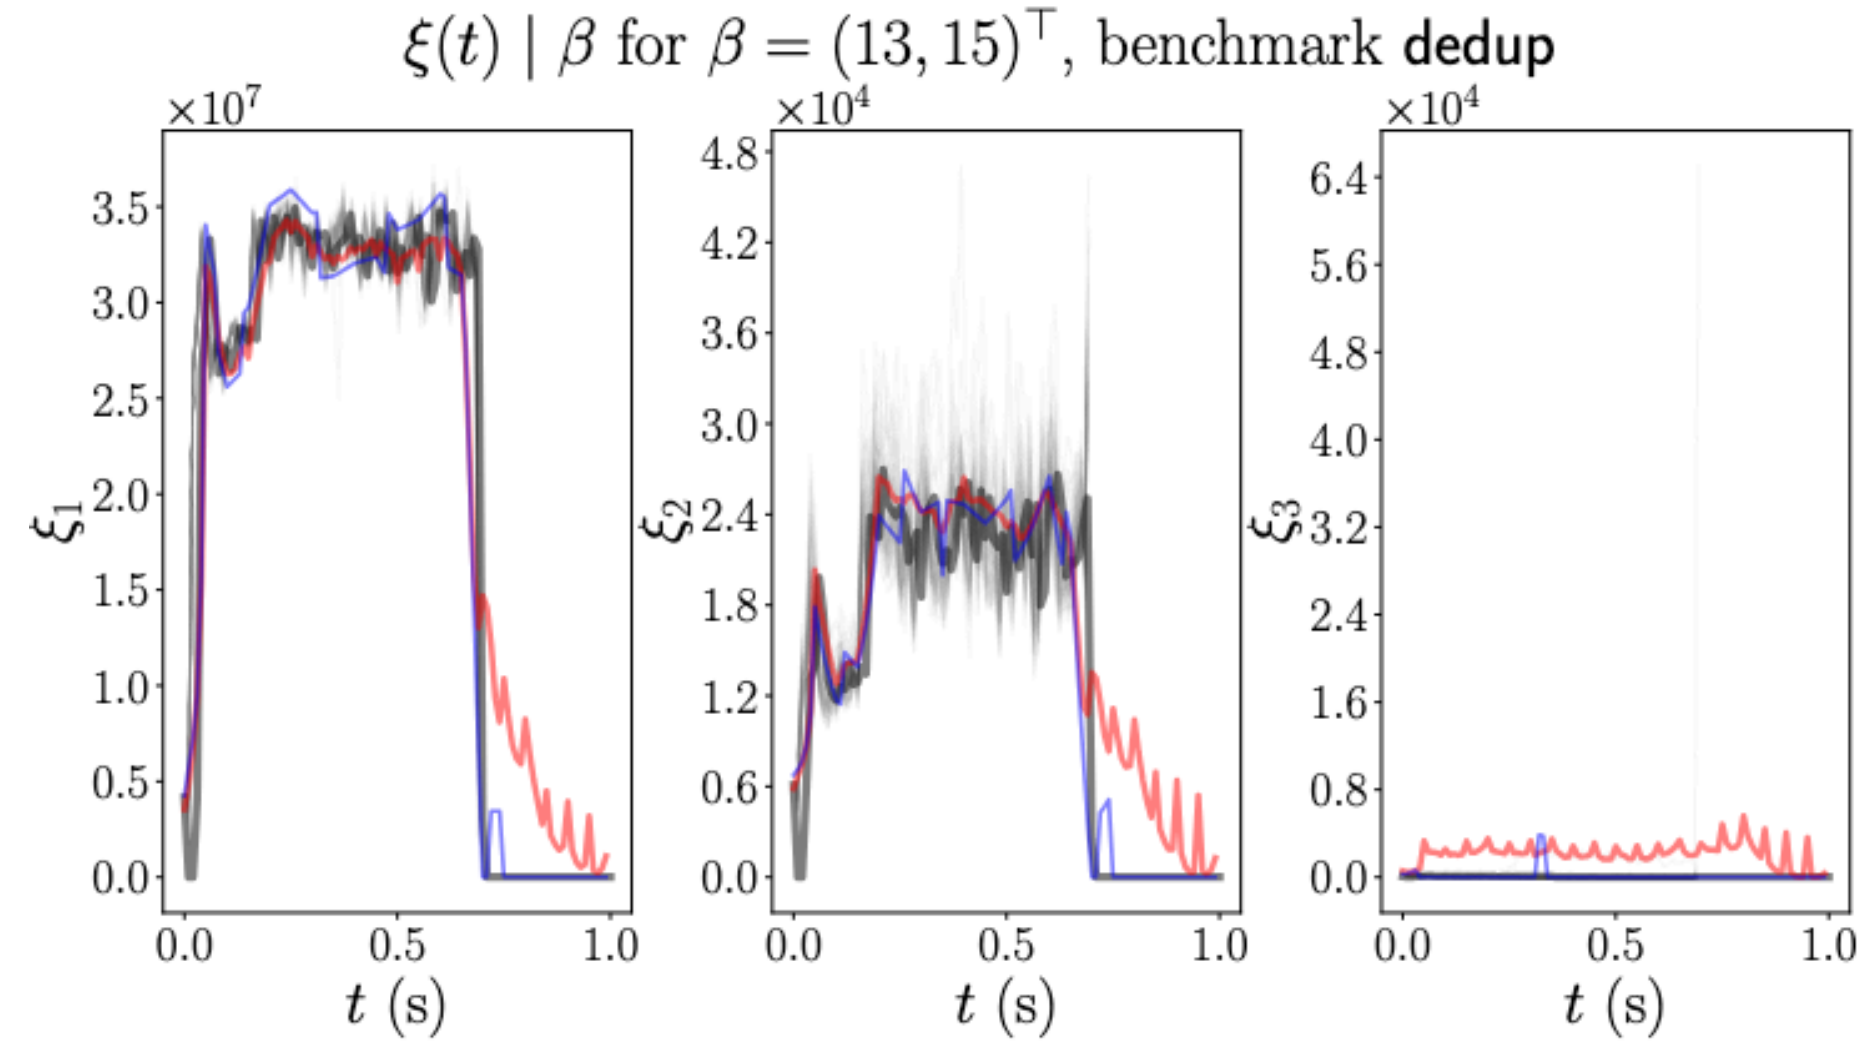


Generate $\xi(\tau) \mid \beta$ for all $\tau \in \{0, 0.01, \dots, \tau_{n_s}\}$, $\beta \in \mathcal{B}$



Generate *mean*, *max-likelihood*, and *avg. empirical* profiles for all $\beta \in \mathcal{B}$

Empirical Profiles for Benchmarks




Maximum-likelihood synthetic profile, mean synthetic profile, mean empirical profile, and all empirical profiles


CORD: A Practical Application

Task scheduling *and* resource allocation

Profiles $\forall \beta \in \mathcal{B}$ required



Generative profiling (MSBP)



Summary



Tensor Optimization for Regularity of OT Maps

Polynomial complexity for forward problem

c rational over $\mathcal{X} \times \mathcal{Y}$ semi algebraic \implies SOS tightening \implies SDP

Forward problem: Computational certificate of NNCC and $\text{MTW}(\kappa)$

Inverse problem: Inner approximation of region of regularity



Tensor Optimization for Graph-structured SB

Graph-structured SBP \implies Solve via Sinkhorn

Complexity $\mathcal{O}(n^{|\Lambda|})$



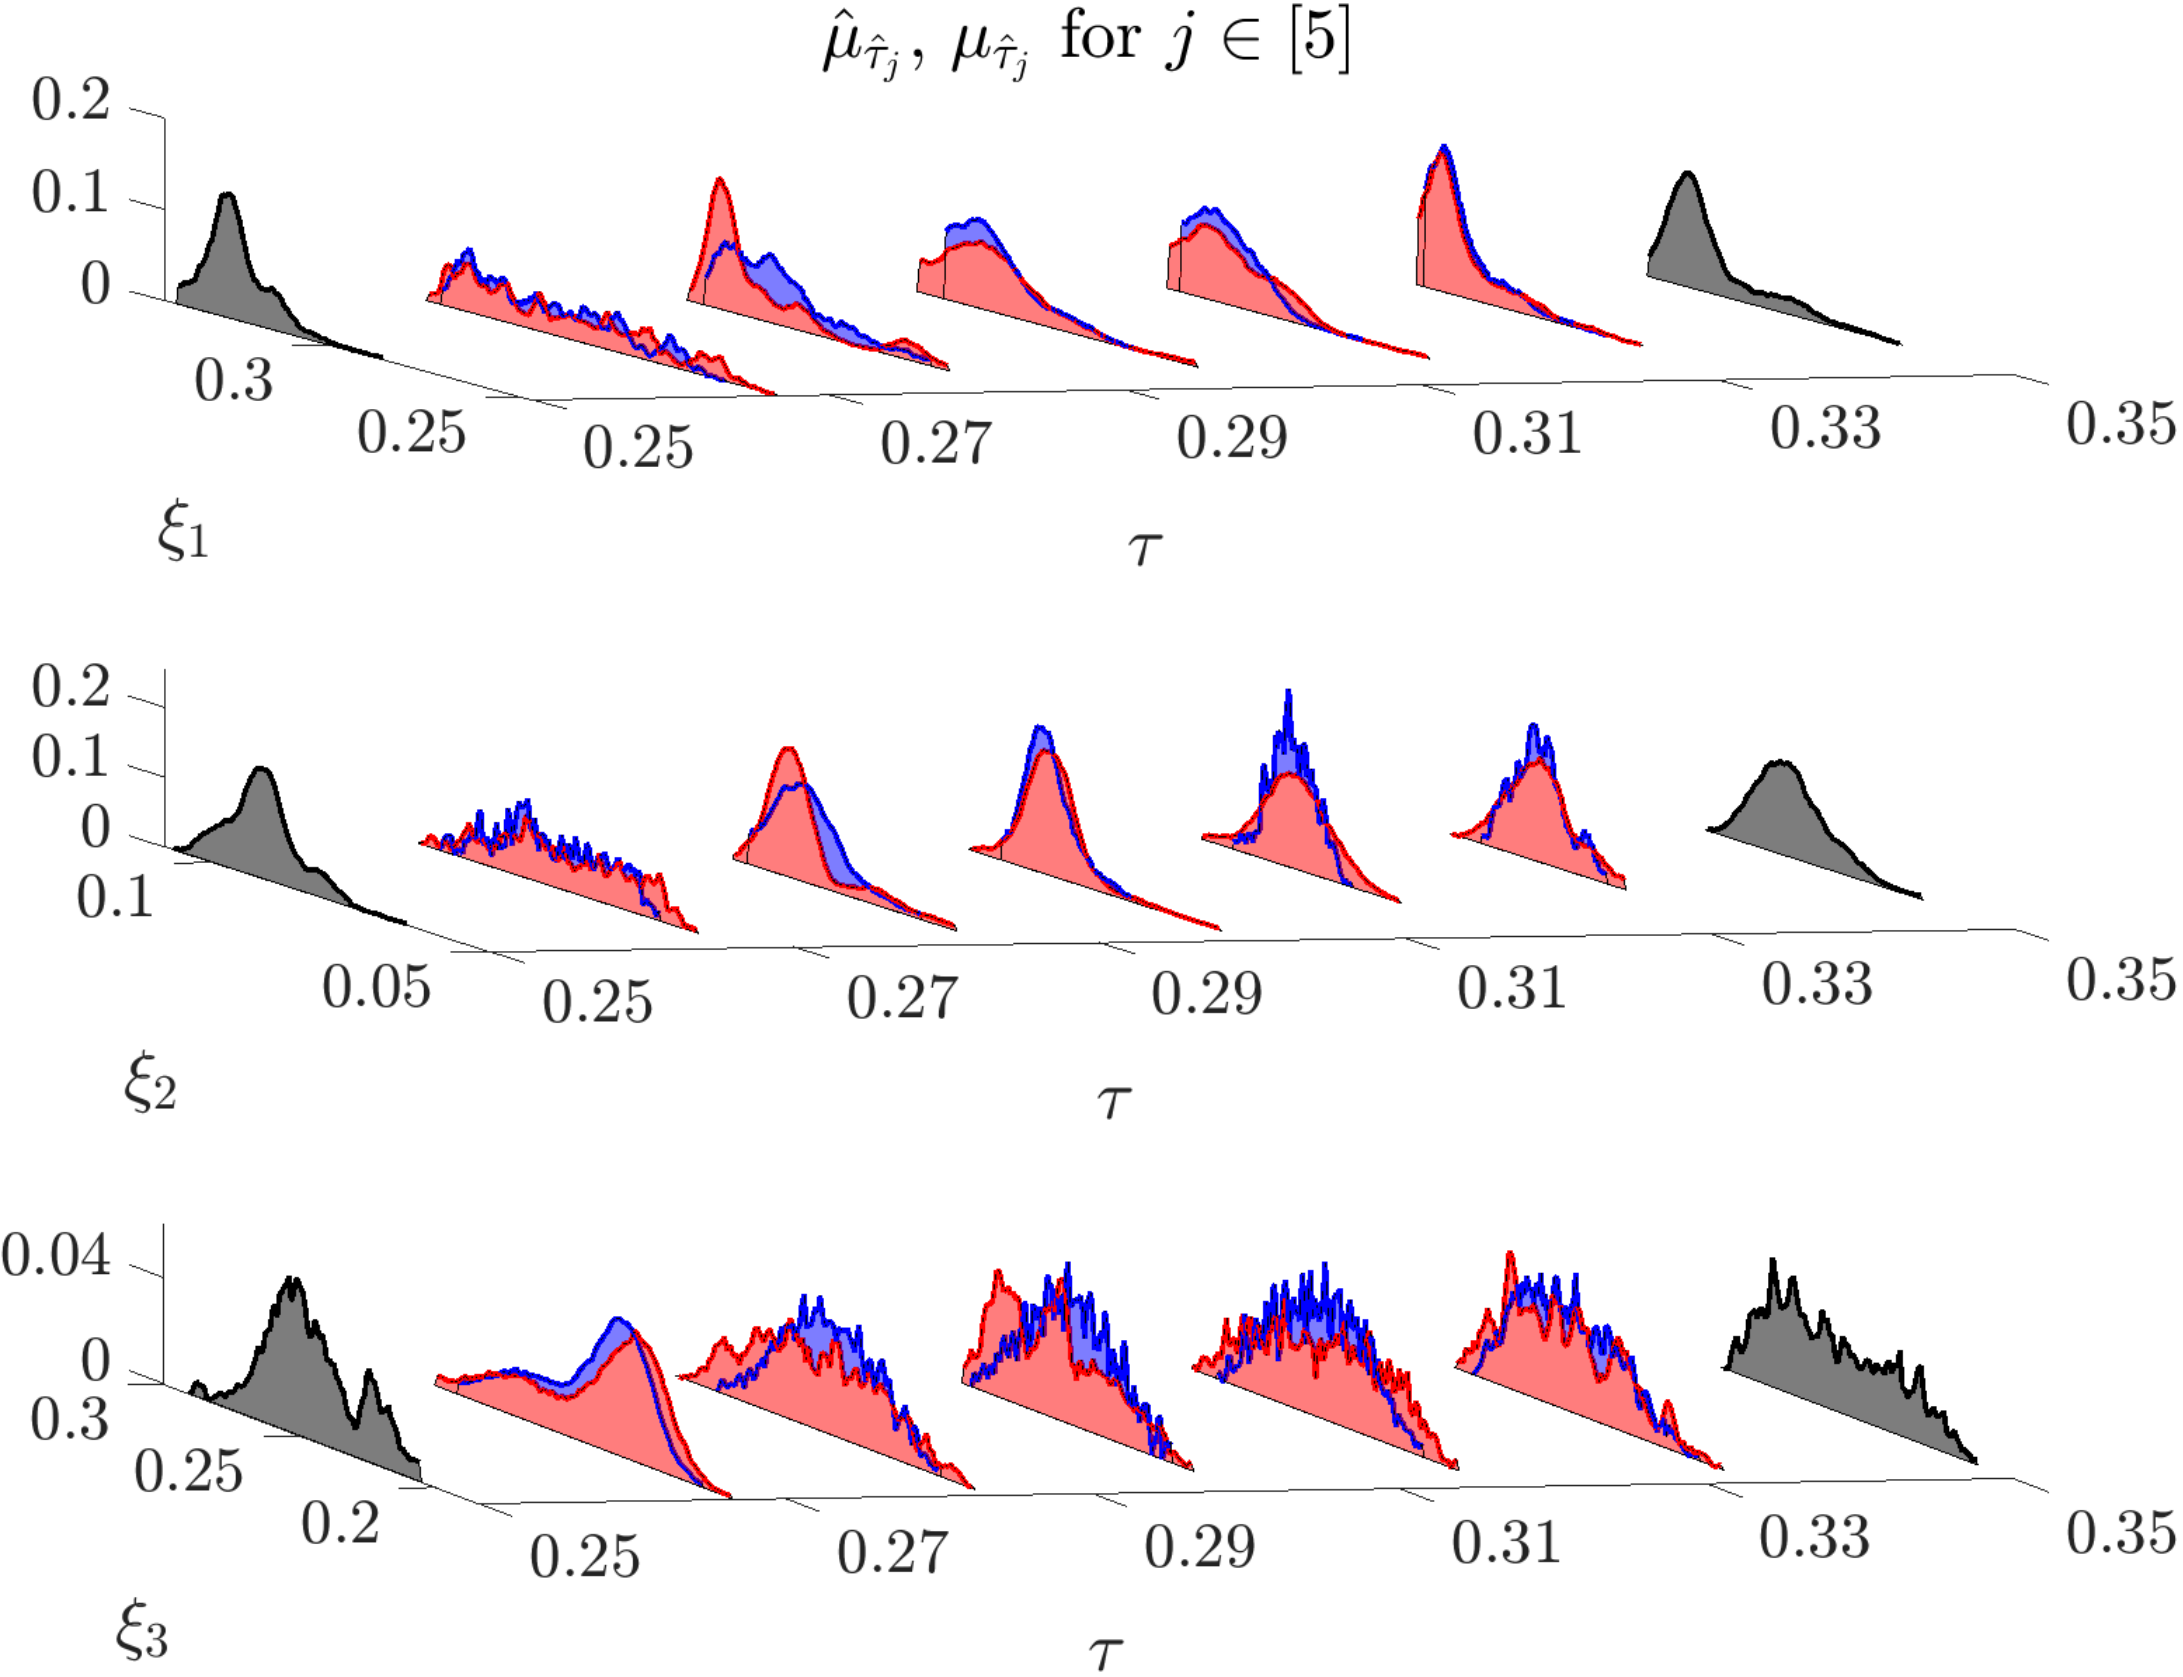
Path, BC, SP graph



Complexity $\mathcal{O}((Js)n^2)$

Linear convergence

Reduce profiling workload



Publications

G.A.B., Gifford, R., Phan, L.T.X., & Halder, A. Path structured multimarginal Schrödinger bridge for probabilistic learning of hardware resource usage by control software. *American Control Conference 2024*

G.A.B., Gifford, R., Phan, L. T. X., & Halder, A. (2024). Stochastic Learning of Computational Resource Usage as Graph Structured Multimarginal Schrödinger Bridge. *Accepted, IEEE Trans. Control Syst. Technology, arXiv:2405.12463.*

Shivakumar, S., **G.A.B.**, Khan, G., & Halder, A. Sum-of-Squares Programming for Ma-Trudinger-Wang Regularity of Optimal Transport Maps. *arXiv:2412.13372.*

Gifford, R., Eisenklam, A., **G.A.B.**, Cai, Y., Sial, T., Phan, L. T. X., & Halder, A. **CORD**: Co-design of Resource Allocation and Deadline Decomposition with Generative Profiling. *arXiv:2501.08484.*

Some Directions for Future Work

Theoretical

- Extend SOS programming approach to higher dimensions
- Learning of optimal tree structures for the MSBP
- MSBP with dynamical constraints

Practical/Applied

- Networked systems
- 3D reconstruction
- Prediction of environmental dynamics (e.g. fire, weather)

Acknowledgements

University of Pennsylvania:




Robert Gifford



Linh T.X. Phan



Abby Eisenklam



Yifan Cai

Iowa State University:



Sachin Shivakumar



Gabriel Khan



Tushar Sial



Abhishek Halder

Funding

UC Santa Cruz Dean's Fellowship (2022-2023), NSF grant 2111688

Thank you!