32nd IEEE Real-Time and Embedded Technology and Applications Symposium

Generative Profiling for Soft Real-Time Systems

and its Applications to Resource Allocation

Georgiy A. Bondar, **Abigail Eisenklam**, Yifan Cai, Robert Gifford,
Tushar Sial, Linh Thi Xuan Phan, Abhishek Halder

Penn UNIVERSITY of PENNSYLVANIA UC SANTA CRUZ IOWA STATE UNIVERSITY

Multi-Core Resource Allocation

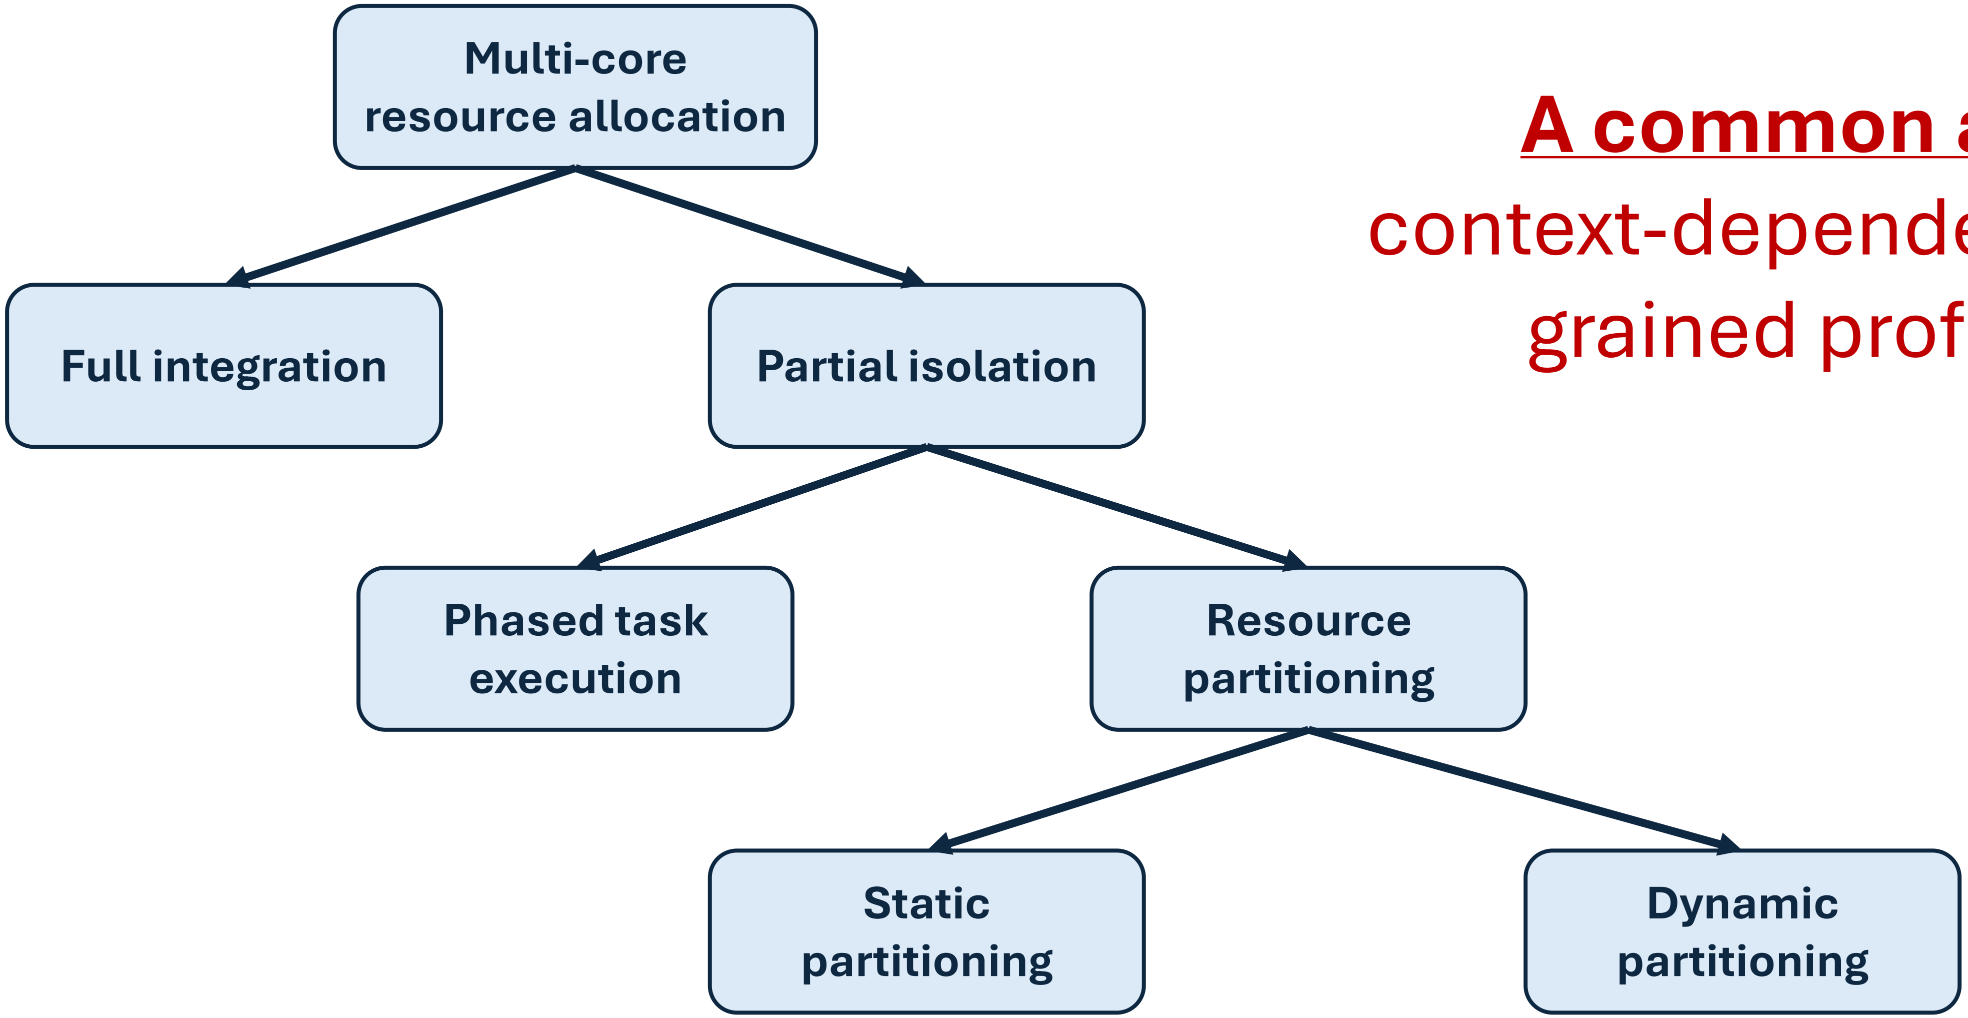

A common aid:

context-dependent fine-grained profiles

Context-Dependent Fine-Grained Profiles

PARSEC benchmark:
`canneal`

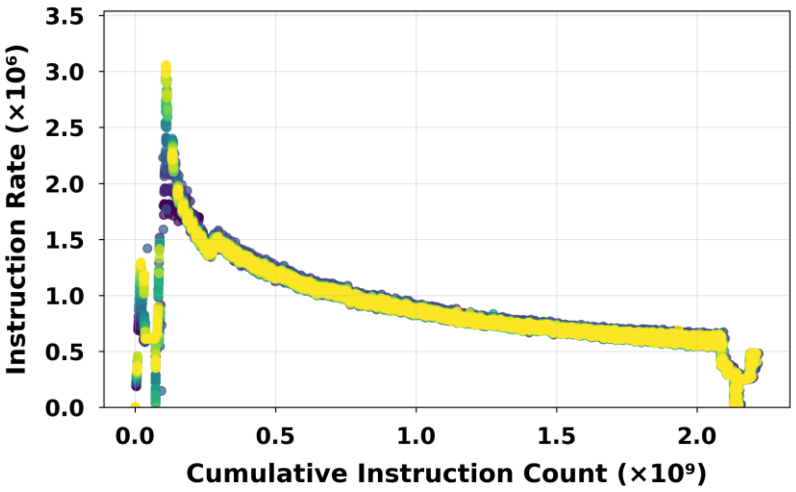

10% of LLC + MemBW

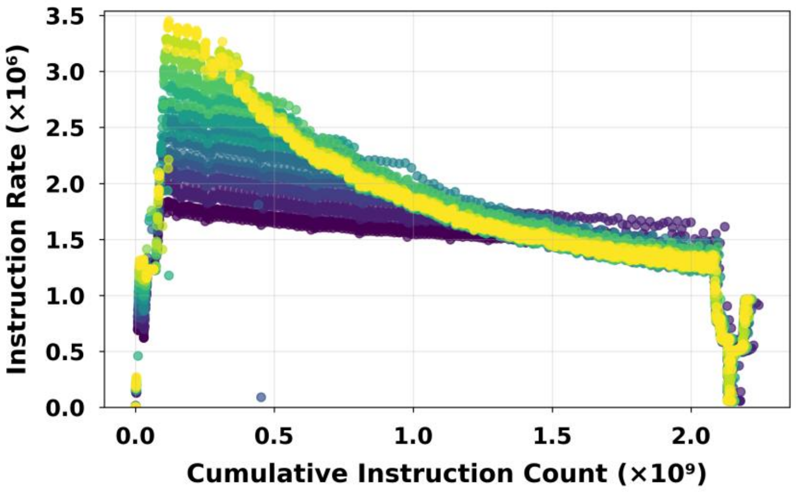

20% of LLC + MemBW

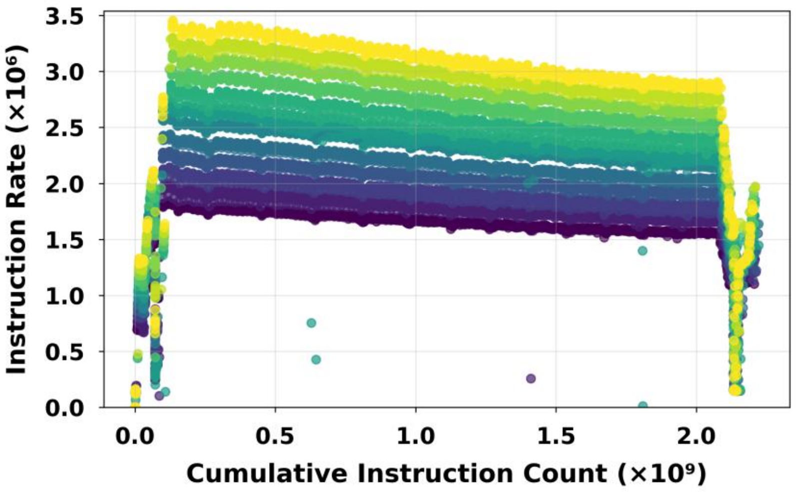

50% of LLC + MemBW

1.2 GHz 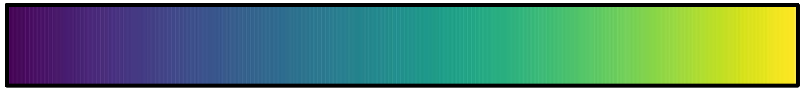 2.3 GHz

Context-Dependent Fine-Grained Profiles

PARSEC benchmark:
`canneal`

Profile for resource context:
(LLC, MemBW, CPUFreq) = (4 MB, 280 MB/s, 2.3 GHz)

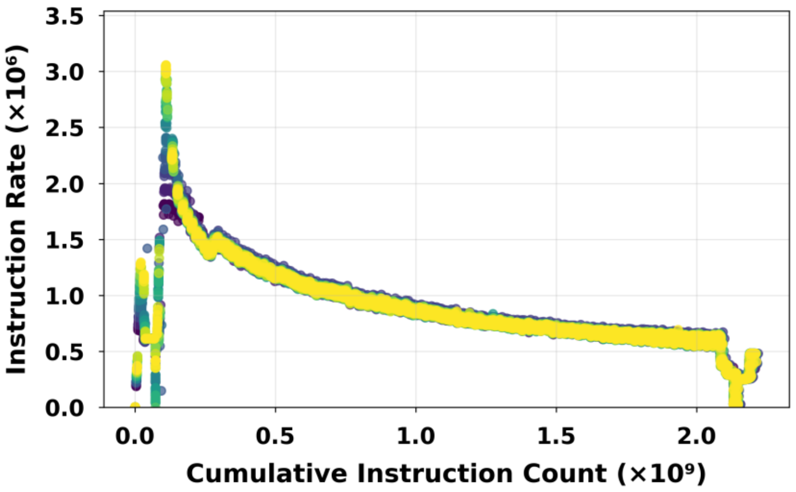

10% of LLC + MemBW

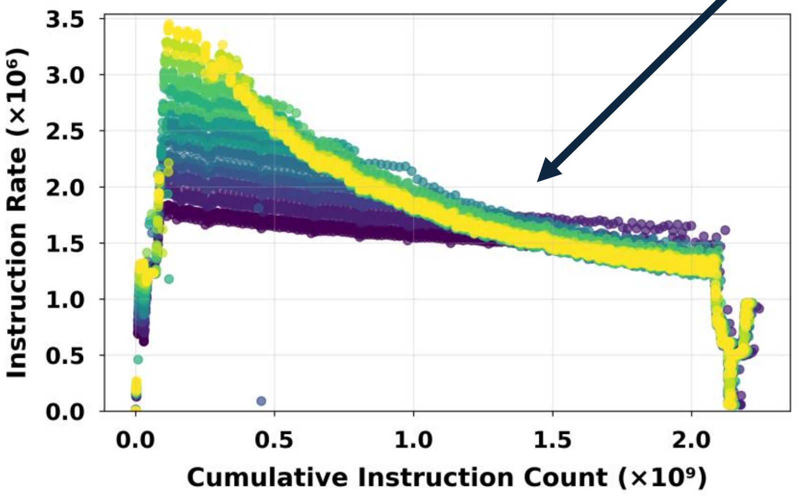

20% of LLC + MemBW

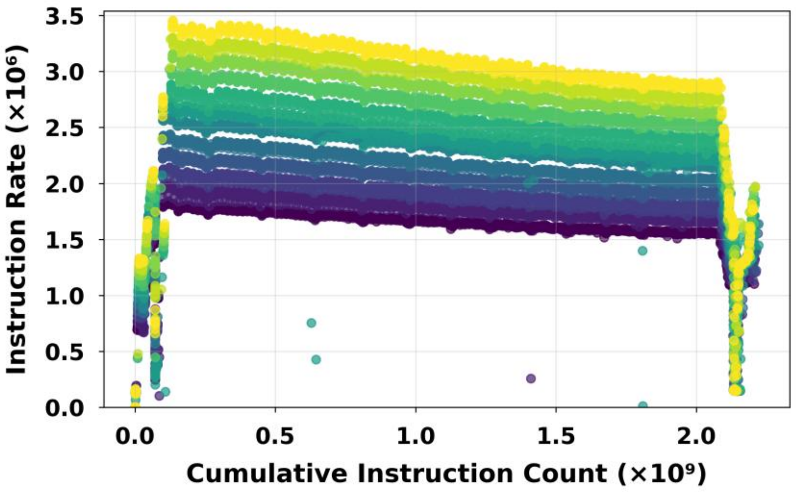

50% of LLC + MemBW

1.2 GHz 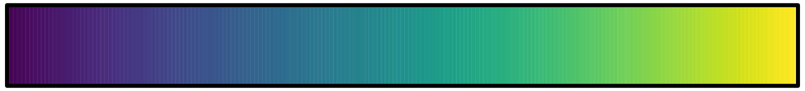 2.3 GHz

Challenges and Limitations of Existing Approaches

Measurement only	Static/hybrid approaches

The space of possible resource contexts is large

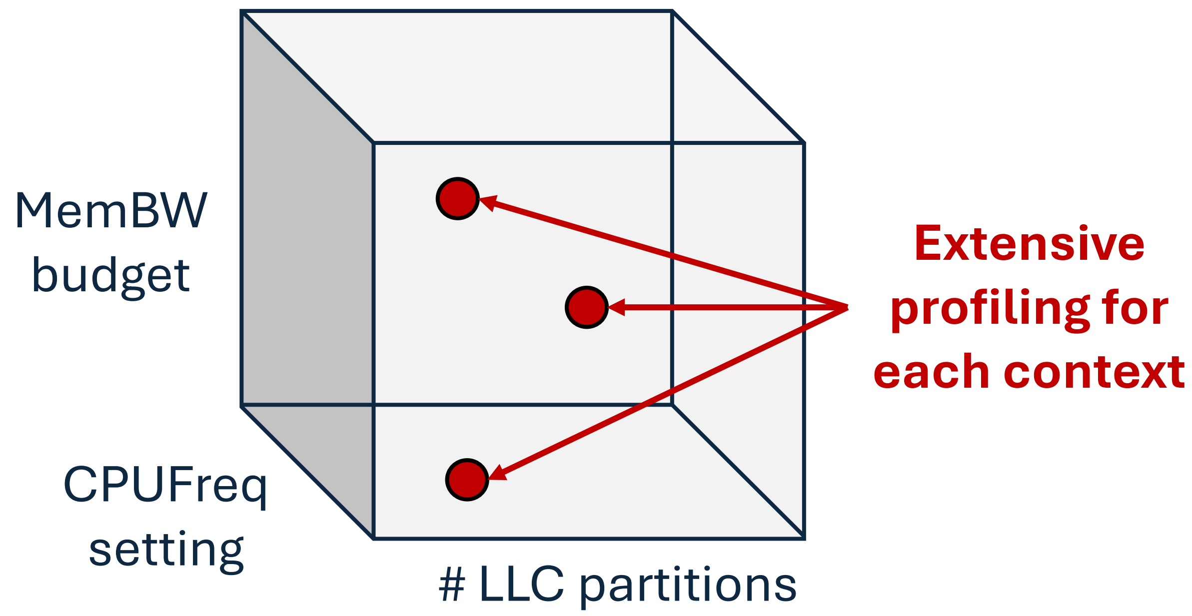

Challenges and Limitations of Existing Approaches

Measurement only

The space of possible resource contexts is large

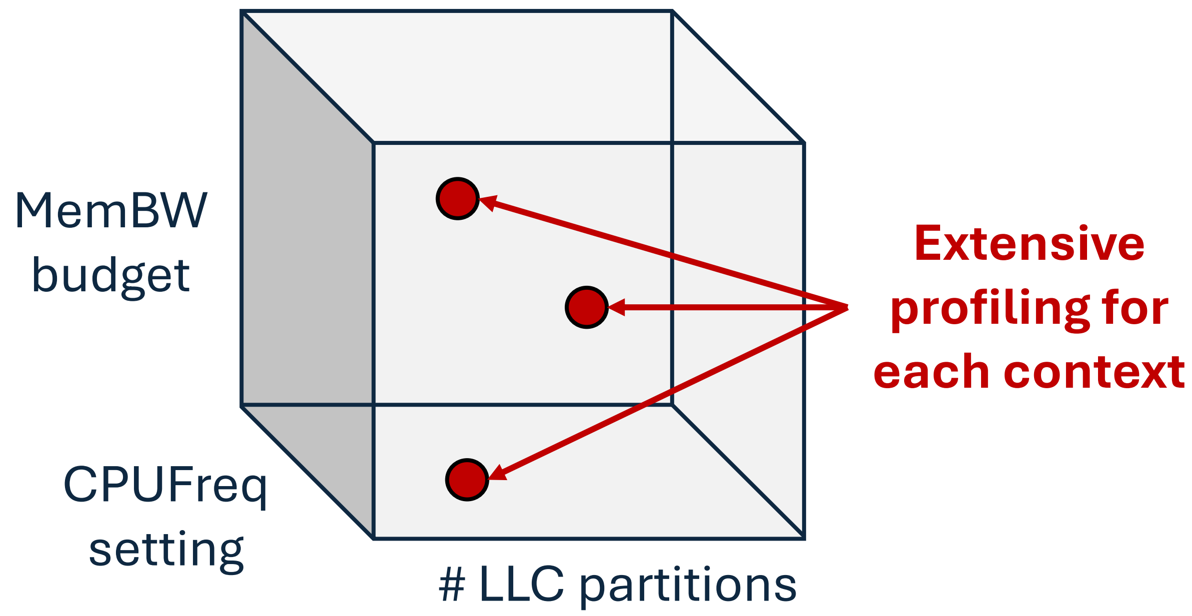

Static/hybrid approaches

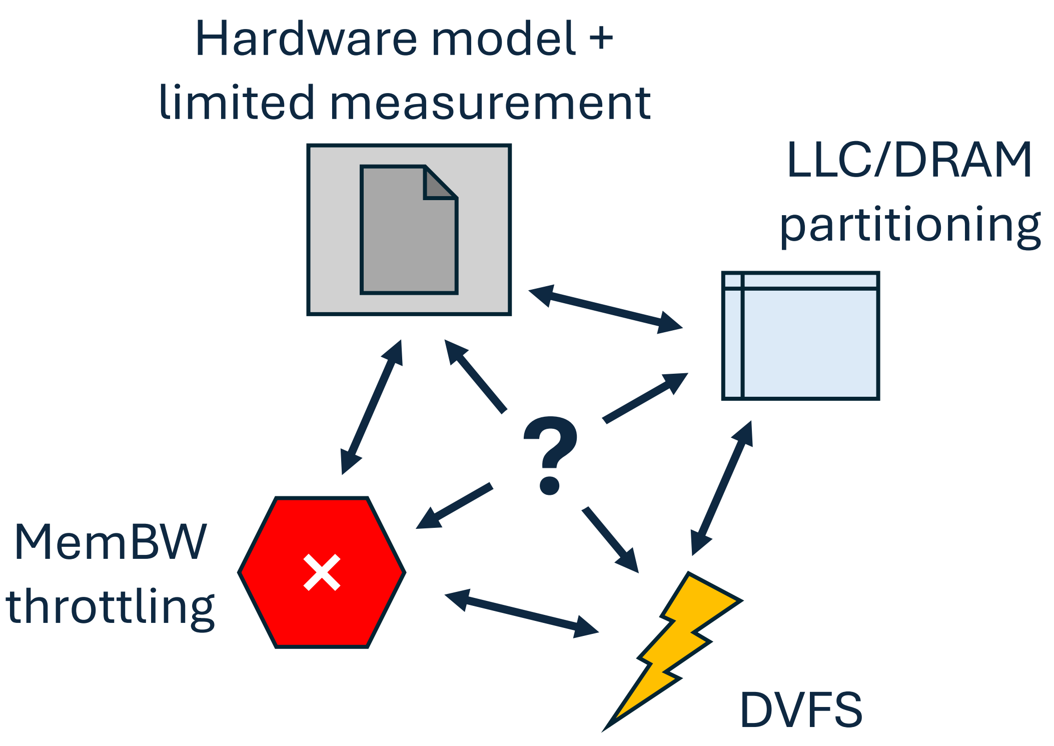

Complex interactions are challenging to model

Our Objective

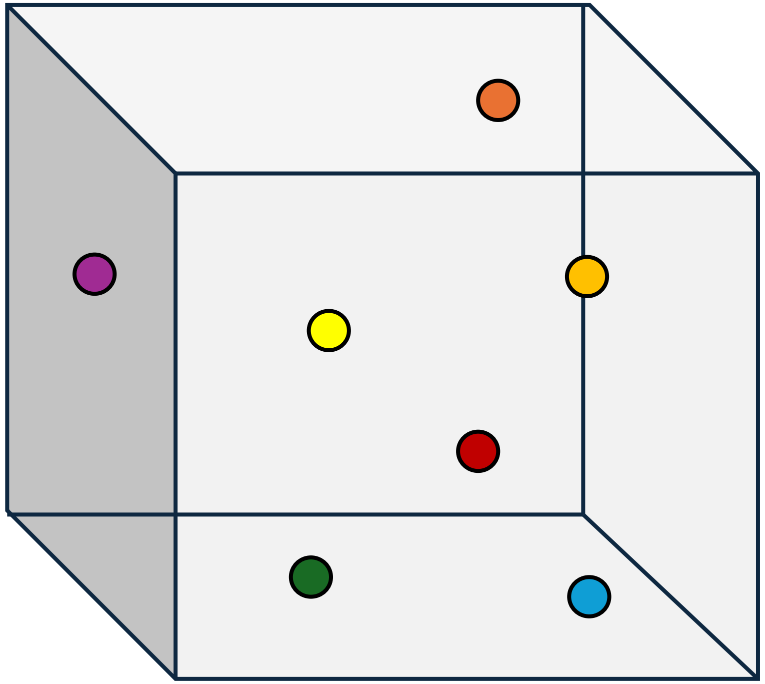

Resource context space

Given:

Sparse set of known resource contexts = { }

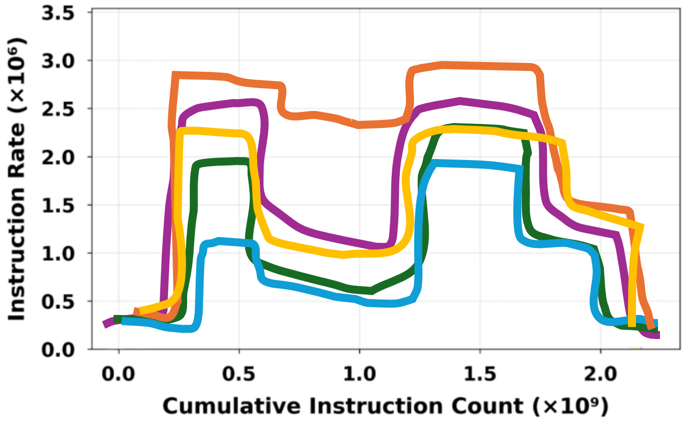

Learn:

Unknown context:

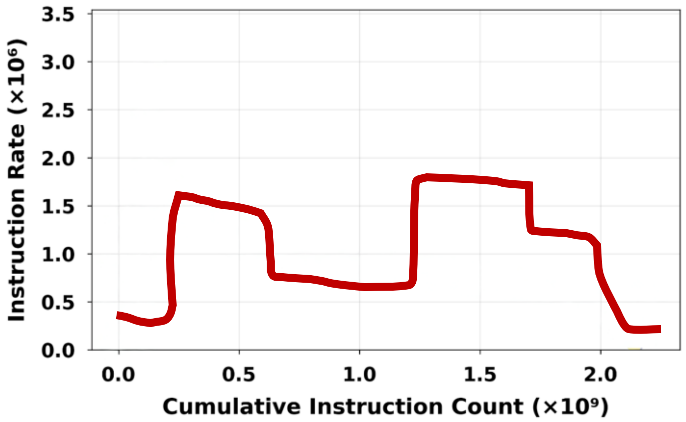

Desired Properties

Given:

Sparse set of known resource contexts = { ● ● ● ● ● ● }

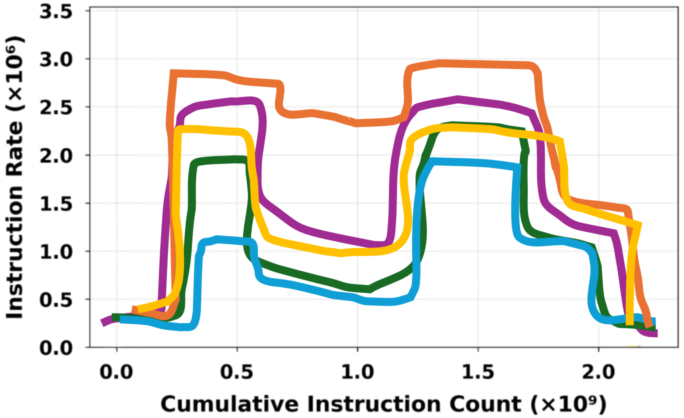

<u>Property 1:</u>
Consistency with known observations

Desired Properties

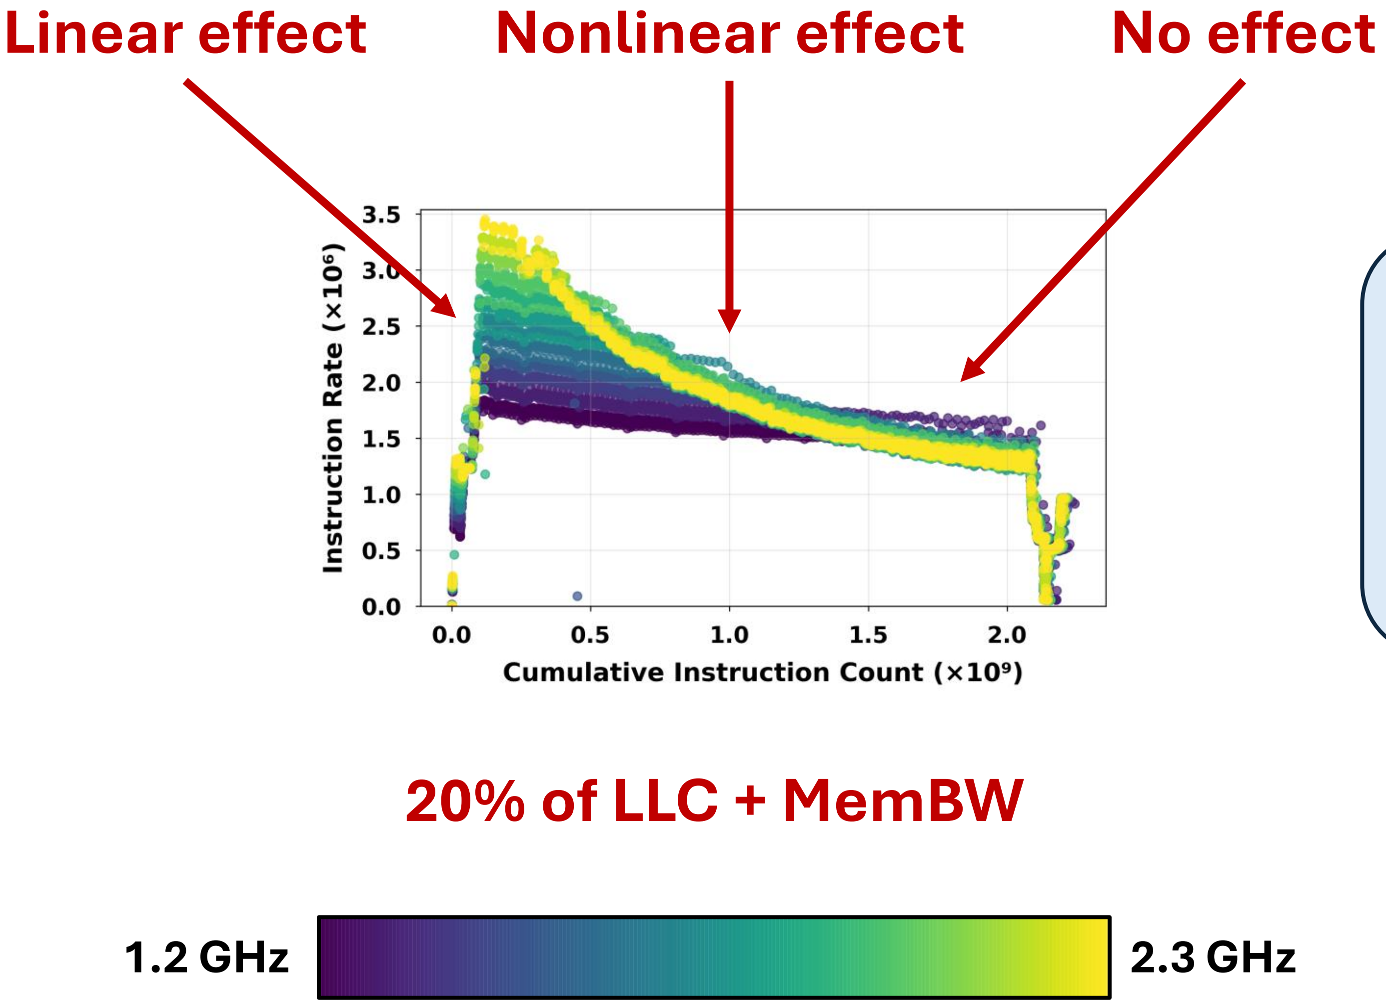

Property 2:
Nonparametric (no distributional assumptions)

Desired Properties

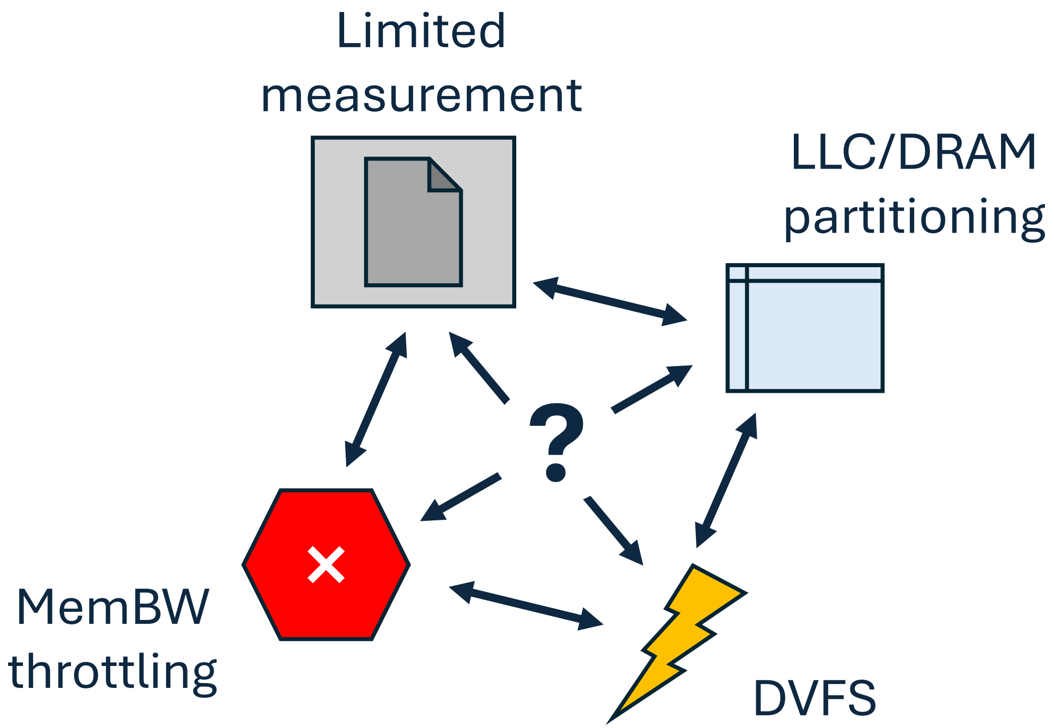

Property 3:
Statistical guarantee of accuracy

Talk Outline

- **Technique Background**
- Solving for MSBs
- Conditional MSBs
- Profile Generation
- Evaluation

The Bimarginal Schrödinger Bridge (SB) Problem

- **Given:** distributions μ_1 and μ_2 corresponding to times t_1 and t_2

The Bimarginal Schrödinger Bridge (SB) Problem

- **Given:** distributions μ_1 and μ_2 corresponding to times t_1 and t_2
- **Desired:** the highest probability <u>distribution-valued</u> path $t \rightarrow \mu_t$ parameterized by $t \in [t_1, t_2]$

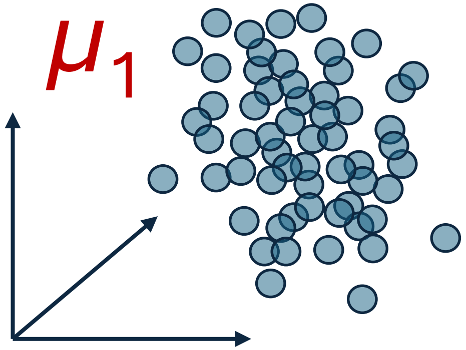

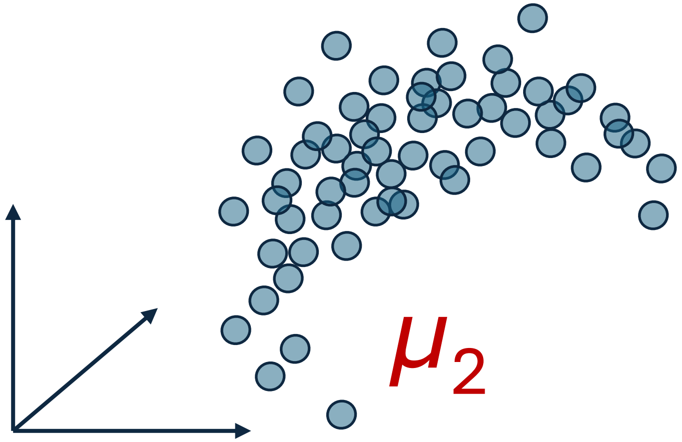

The Bimarginal Schrödinger Bridge (SB) Problem

- **Given:** distributions μ_1 and μ_2 corresponding to times t_1 and t_2
- **Desired:** the highest probability distribution-valued path $t \to \mu_t$ parameterized by $t \in [t_1, t_2]$

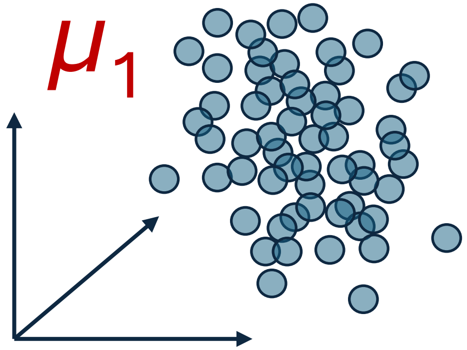

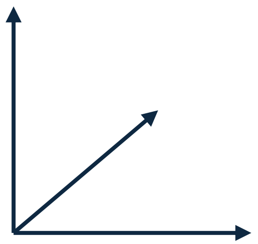

The Bimarginal Schrödinger Bridge (SB) Problem

- **Given:** distributions μ_1 and μ_2 corresponding to times t_1 and t_2
- **Desired:** the highest probability distribution-valued path $t \rightarrow \mu_t$ parameterized by $t \in [t_1, t_2]$

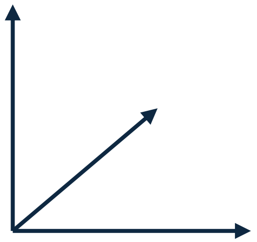

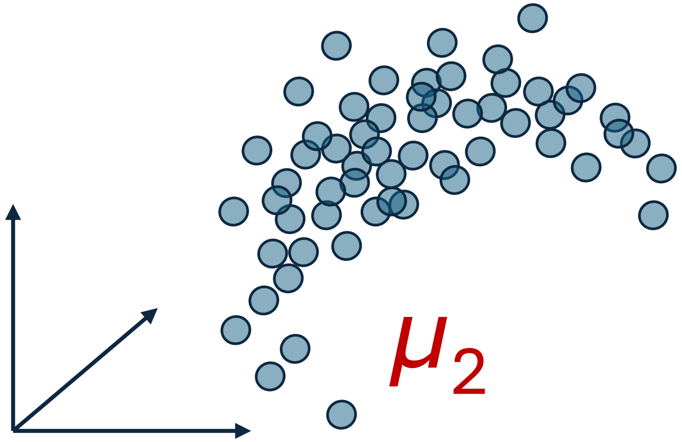

The Bimarginal Schrödinger Bridge (SB) Problem

- **Given:** distributions μ_1 and μ_2 corresponding to times t_1 and t_2
- **Desired:** the highest probability <u>distribution-valued</u> path $t \rightarrow \mu_t$ parameterized by $t \in [t_1, t_2]$

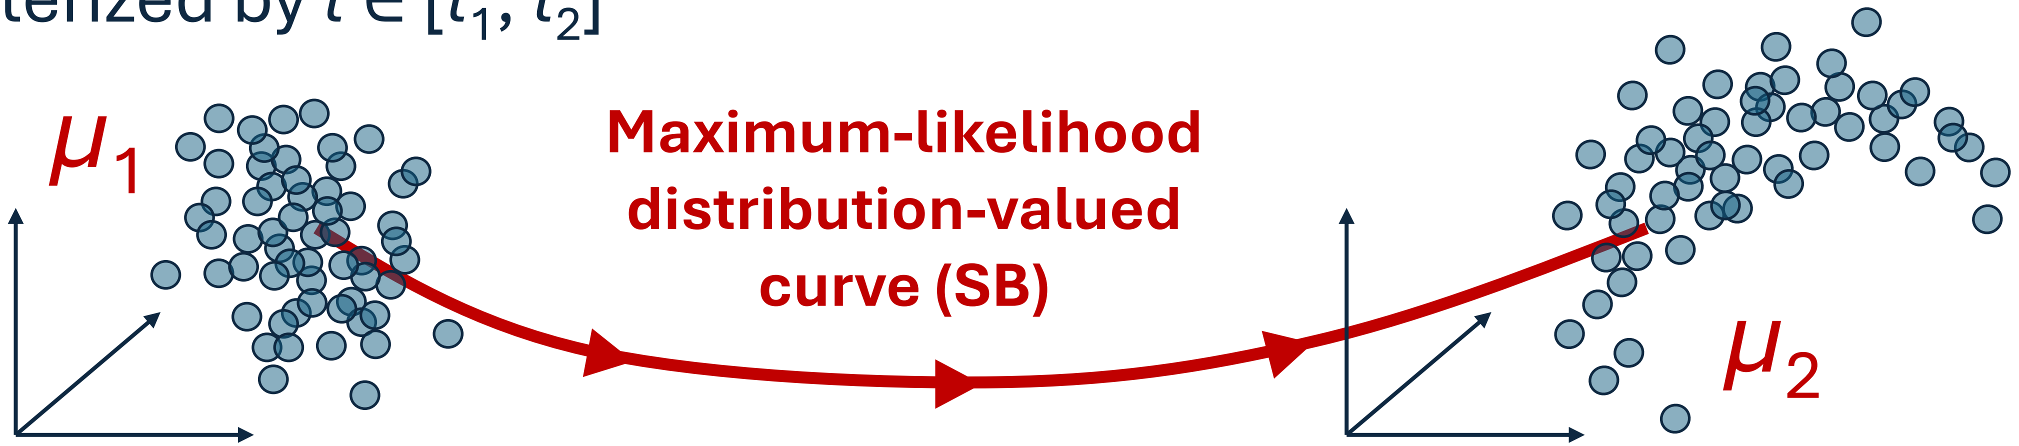

The Bimarginal Schrödinger Bridge (SB) Problem

- **Given:** distributions μ_1 and μ_2 corresponding to times t_1 and t_2
- **Desired:** the highest probability distribution-valued path $t \to \mu_t$ parameterized by $t \in [t_1, t_2]$

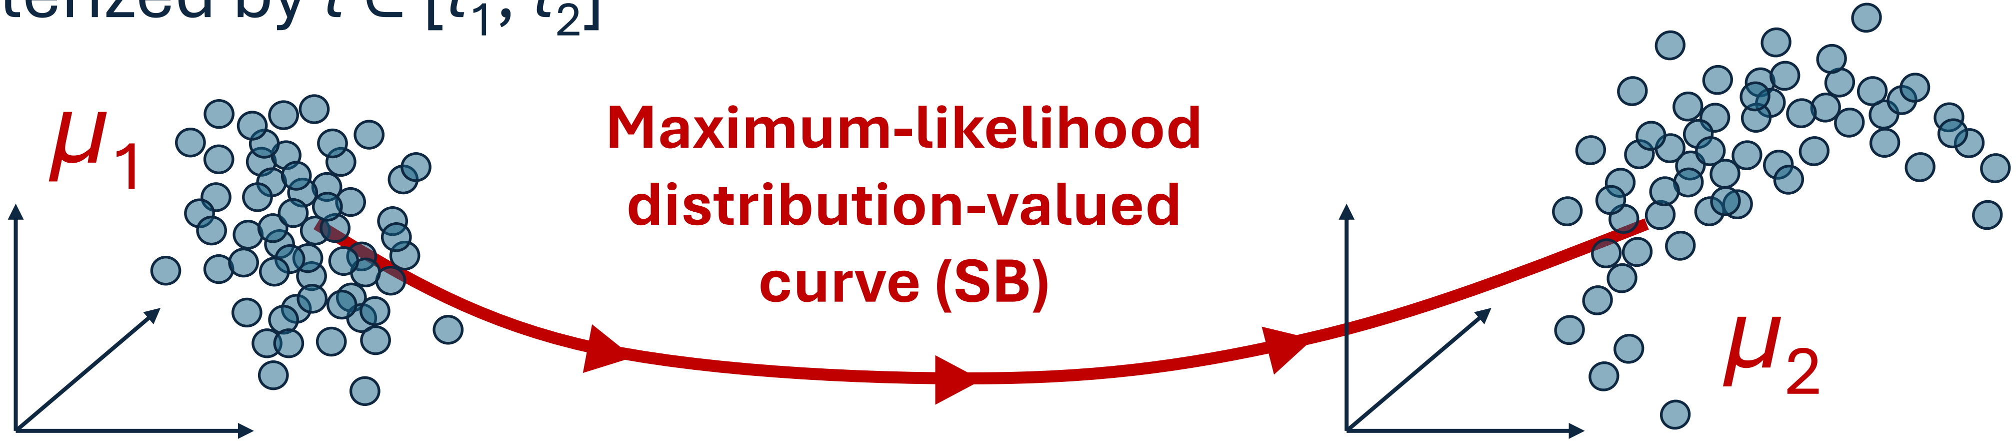

- Originated in 1931-1932 in the works of Schrödinger [1] [2]
- Recent applications in generative AI and stochastic control

[1] E. Schrödinger. Über die Umkehrung der Naturgesetze. Verlag der Akademie der Wissenschaften in Kommission bei Walter De Gruyteru Company, 1931.
[2] E. Schrödinger. Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique. In Annales de l'I. H. P., 1932.

The Bimarginal Schrödinger Bridge (SB) Problem

- **Given:** distributions μ_1 and μ_2 corresponding to times t_1 and t_2
- **Desired:** the highest probability distribution-valued path $t \rightarrow \mu_t$ parameterized by $t \in [t_1, t_2]$

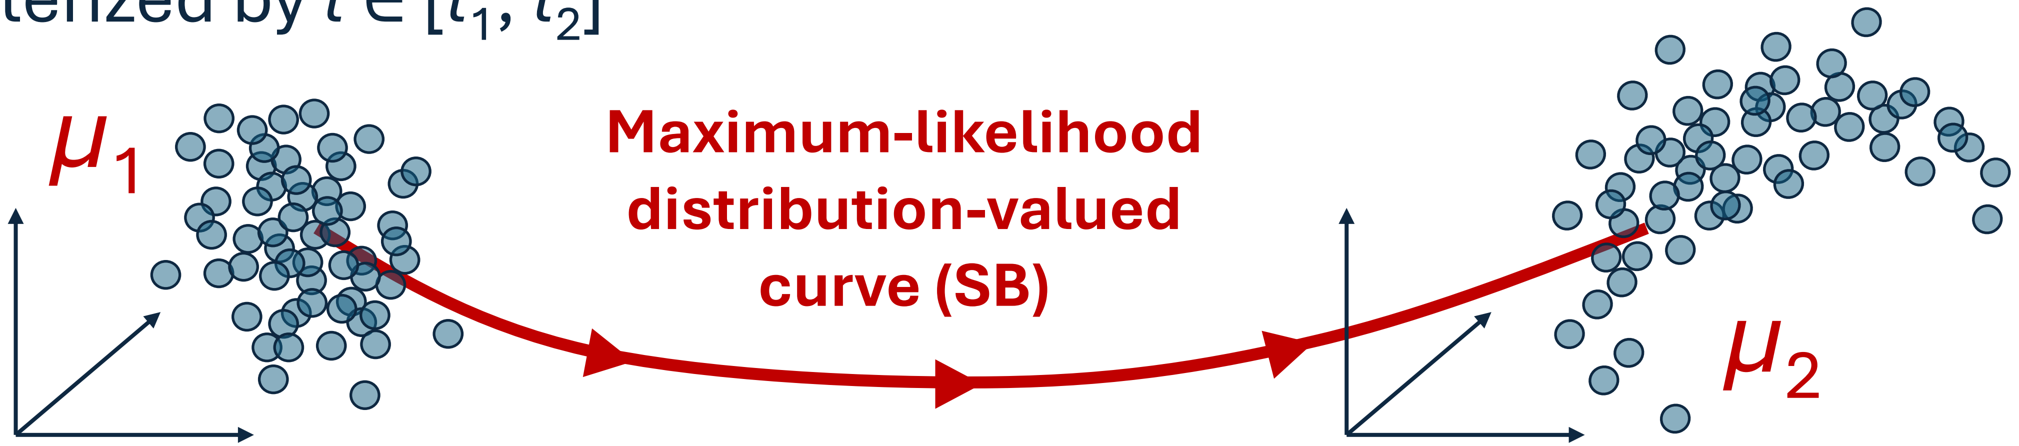

Property 1: Consistency with known observations

Property 2: Nonparametric (no distributional assumptions)

Property 3: Statistical guarantee of accuracy

SB Formulation for Fixed Resource Context β

- Model program behavior as a vectorial stochastic process, conditioned on a resource context

Time-varying execution state $\longrightarrow$ $$\xi(t) \mid \beta$$ $\longleftarrow$ **Resource context**

SB Formulation for Fixed Resource Context β

- Model program behavior as a vectorial stochastic process, conditioned on a resource context

Time-varying execution state → $\xi(t) \mid \beta$ ← **Resource context**

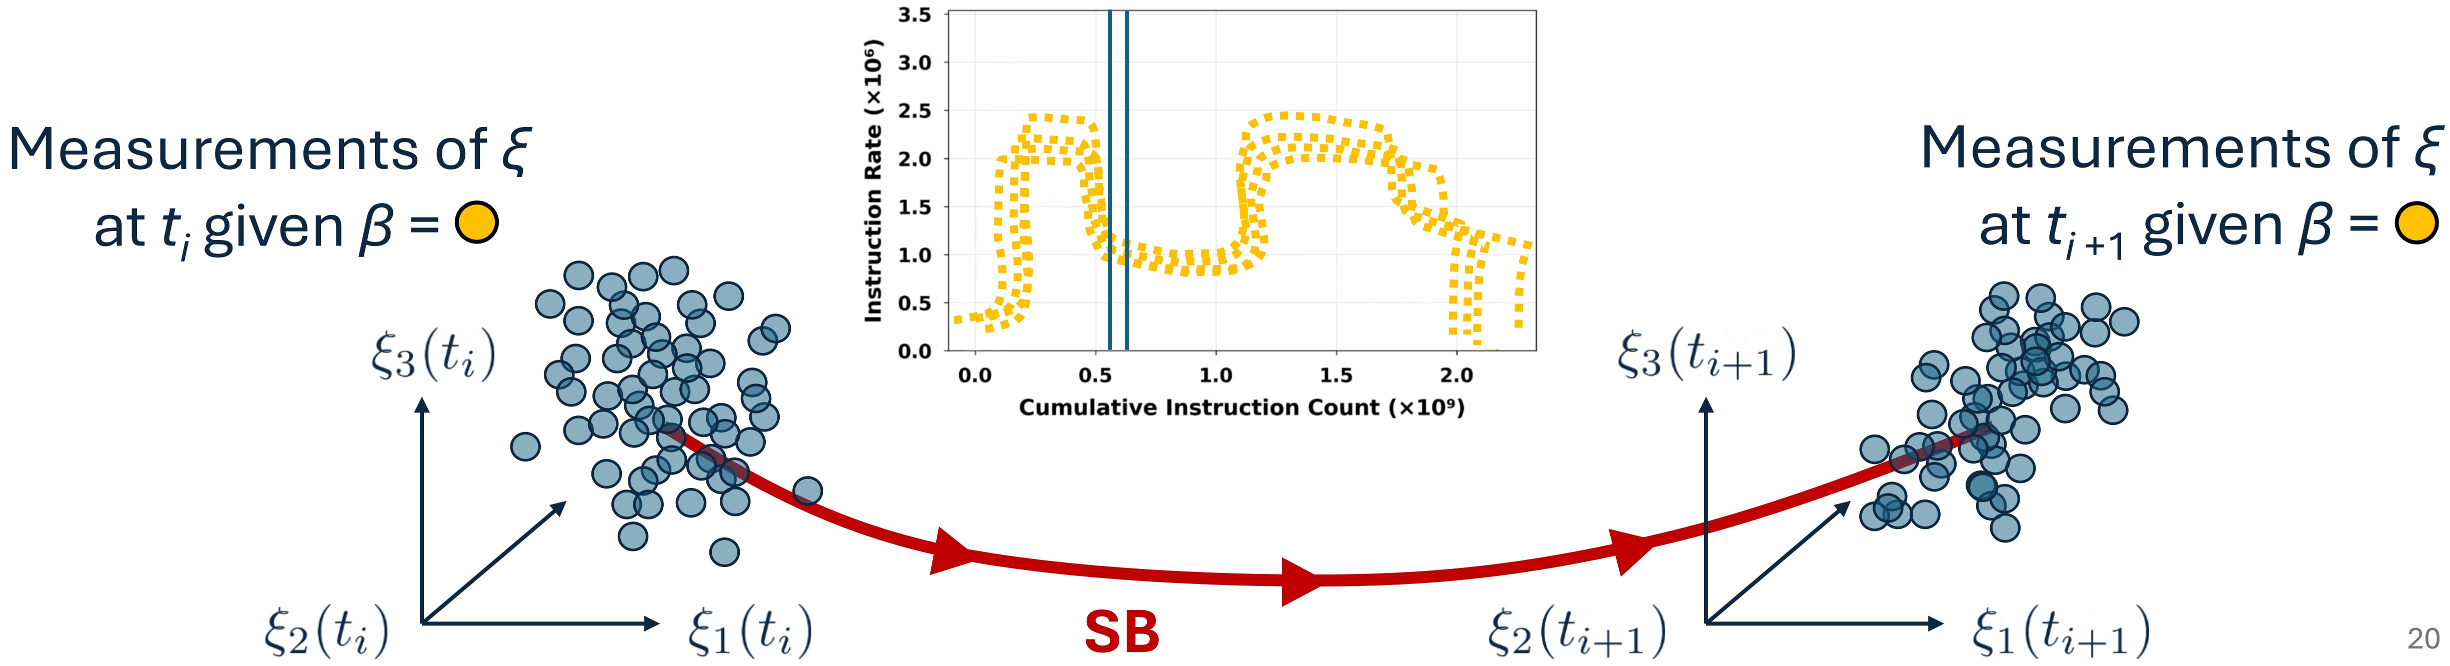

Multimarginal Schrödinger Bridge (MSB)

- SB naturally extends to more than two consecutive distributions $\{\mu_\sigma\}_{\sigma\in[\![s]\!]}$

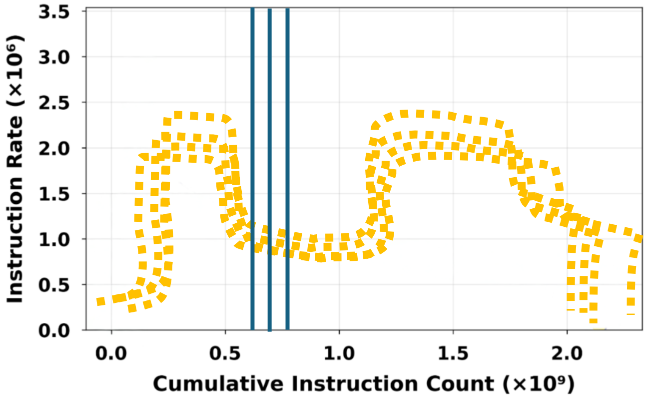

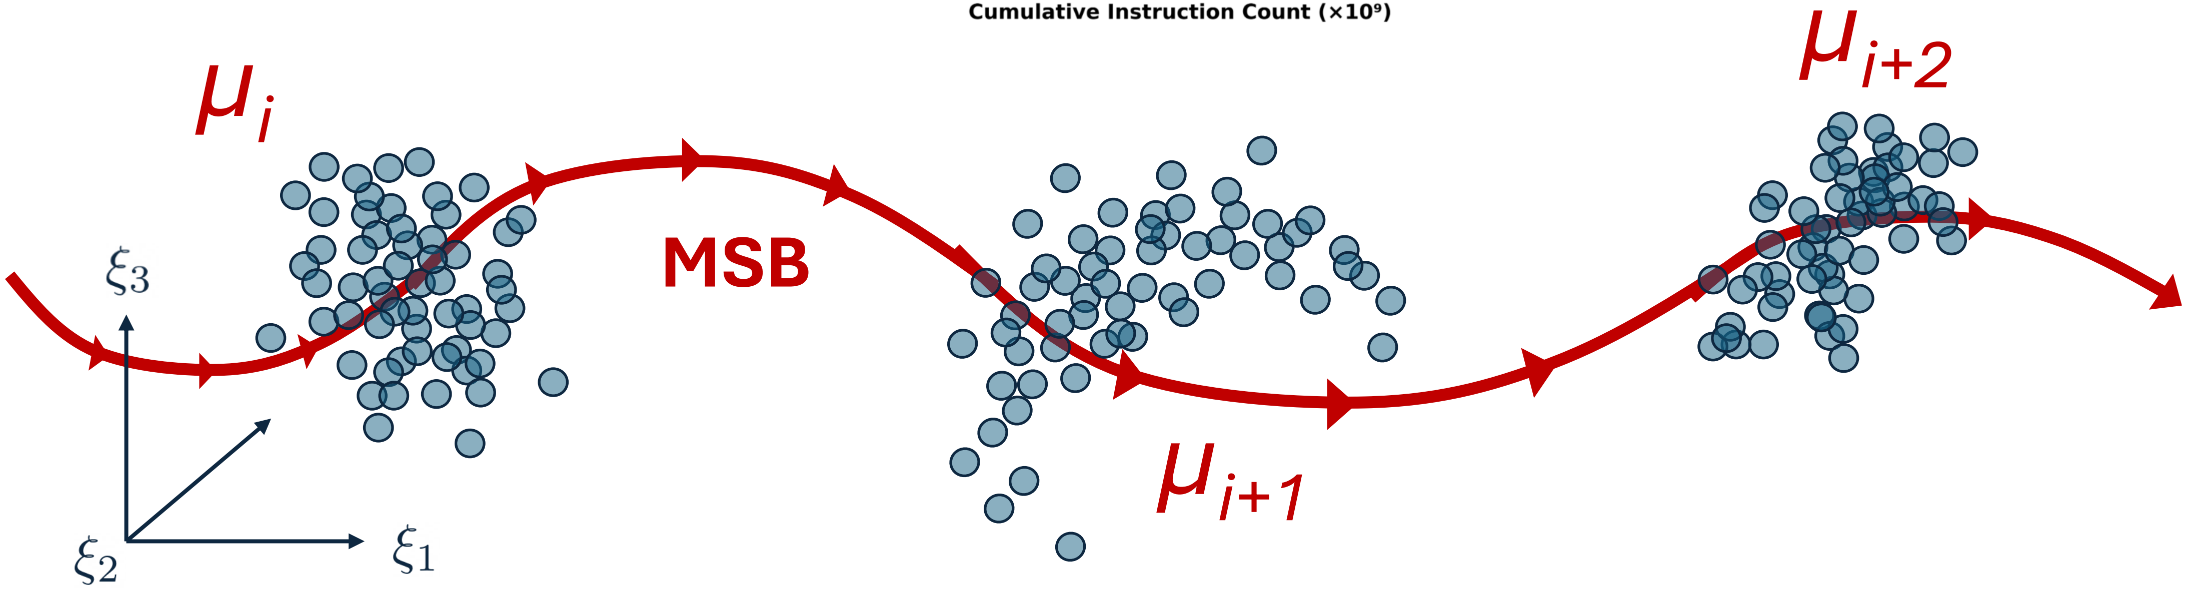

Talk Outline

- SB Background
- **Solving for MSBs**
- Conditional MSBs
- Profile Generation
- Evaluation

Solving for Multimarginal Schrödinger Bridges

- Idea: MSB is a tensor $\boldsymbol{M}$ with $\{\mu_\sigma\}$ as marginals.
- Find cost + entropy minimizing $\boldsymbol{M}_{\text{opt}}$

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

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

$\boldsymbol{C}(\bigcirc, \triangle, \square) = \text{pairwise-Euclidean-cost}(\bigcirc, \triangle) + \text{pairwise-Euclidean-cost}(\triangle, \square)$

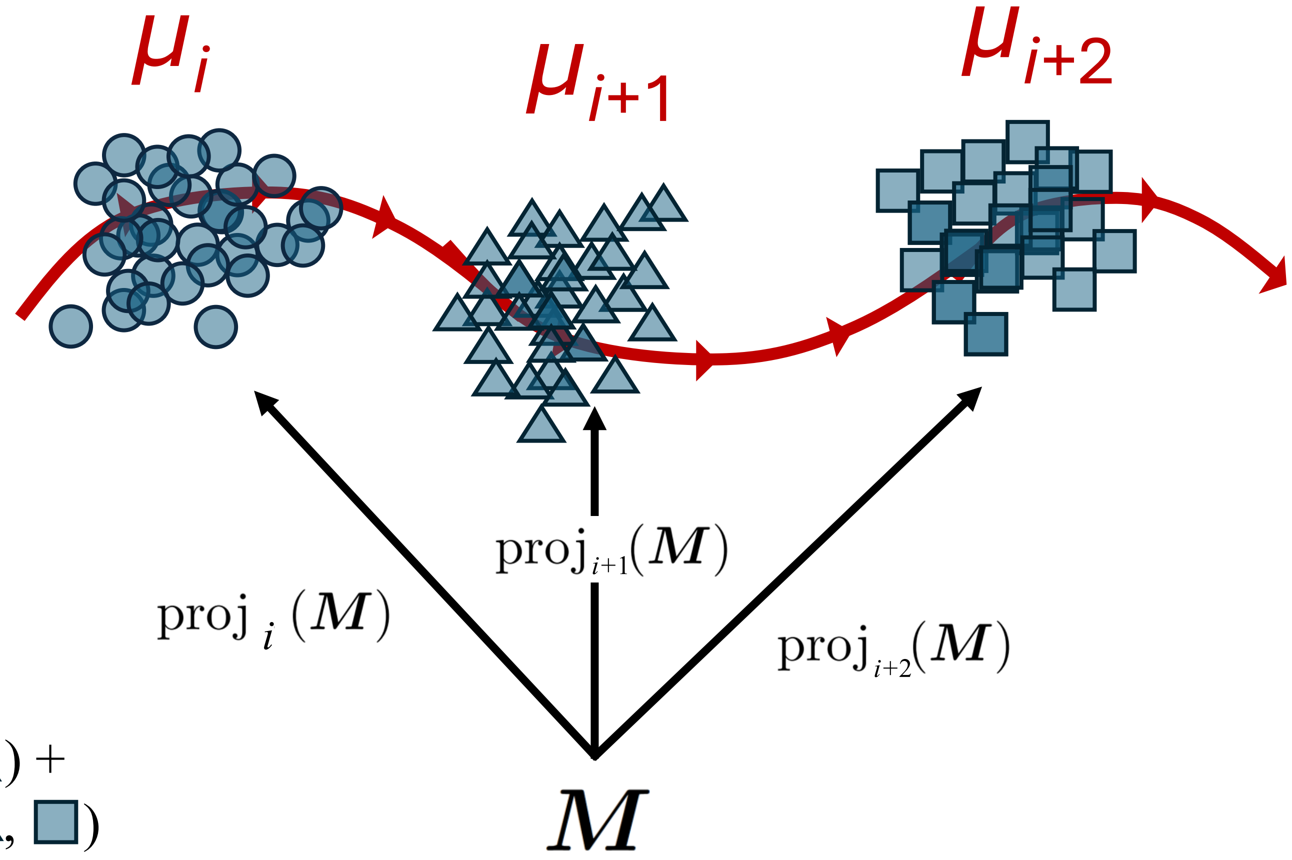

Solving for Multimarginal Schrödinger Bridges

$\underline{\varepsilon > 0}$

- Strictly convex program
- Strong Lagrange duality
- **P3: Statistical guarantee: max. likelihood**

n^s decision variables

$$\min_{\boldsymbol{M} \in (\mathbb{R}^n)^{\otimes s}_{\geq 0}} \quad \langle \boldsymbol{C} + \varepsilon \log \boldsymbol{M}, \boldsymbol{M} \rangle$$

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

P1: Consistency with known observations

P2: Unique, nonparametric solution guaranteed in linear time!

Talk Outline

- SB Background
- Solving for MSBs
- **Conditional MSBs**
- Profile Generation
- Evaluation

Inputs to Conditional MSBP

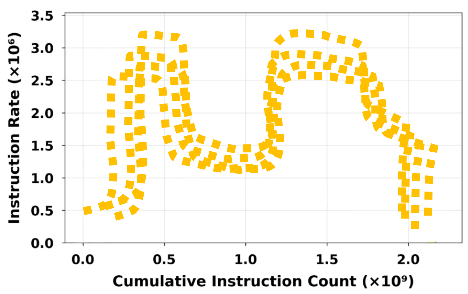

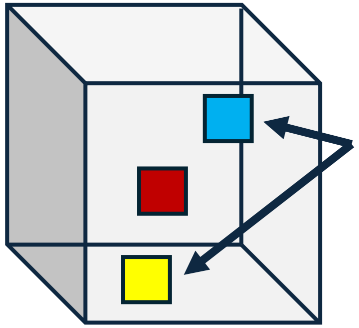

Known resource contexts

Resource context space

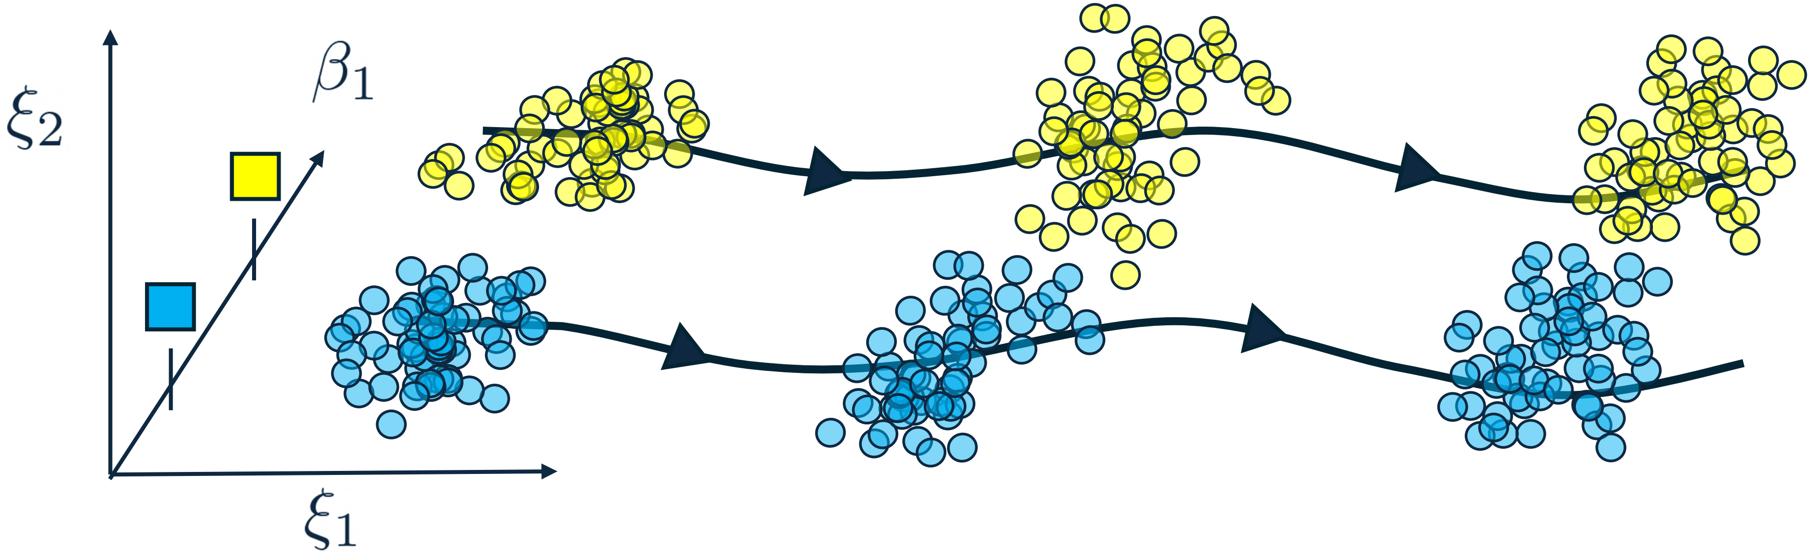

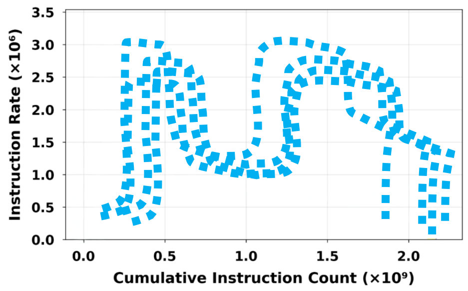

Output of Conditional MSBP

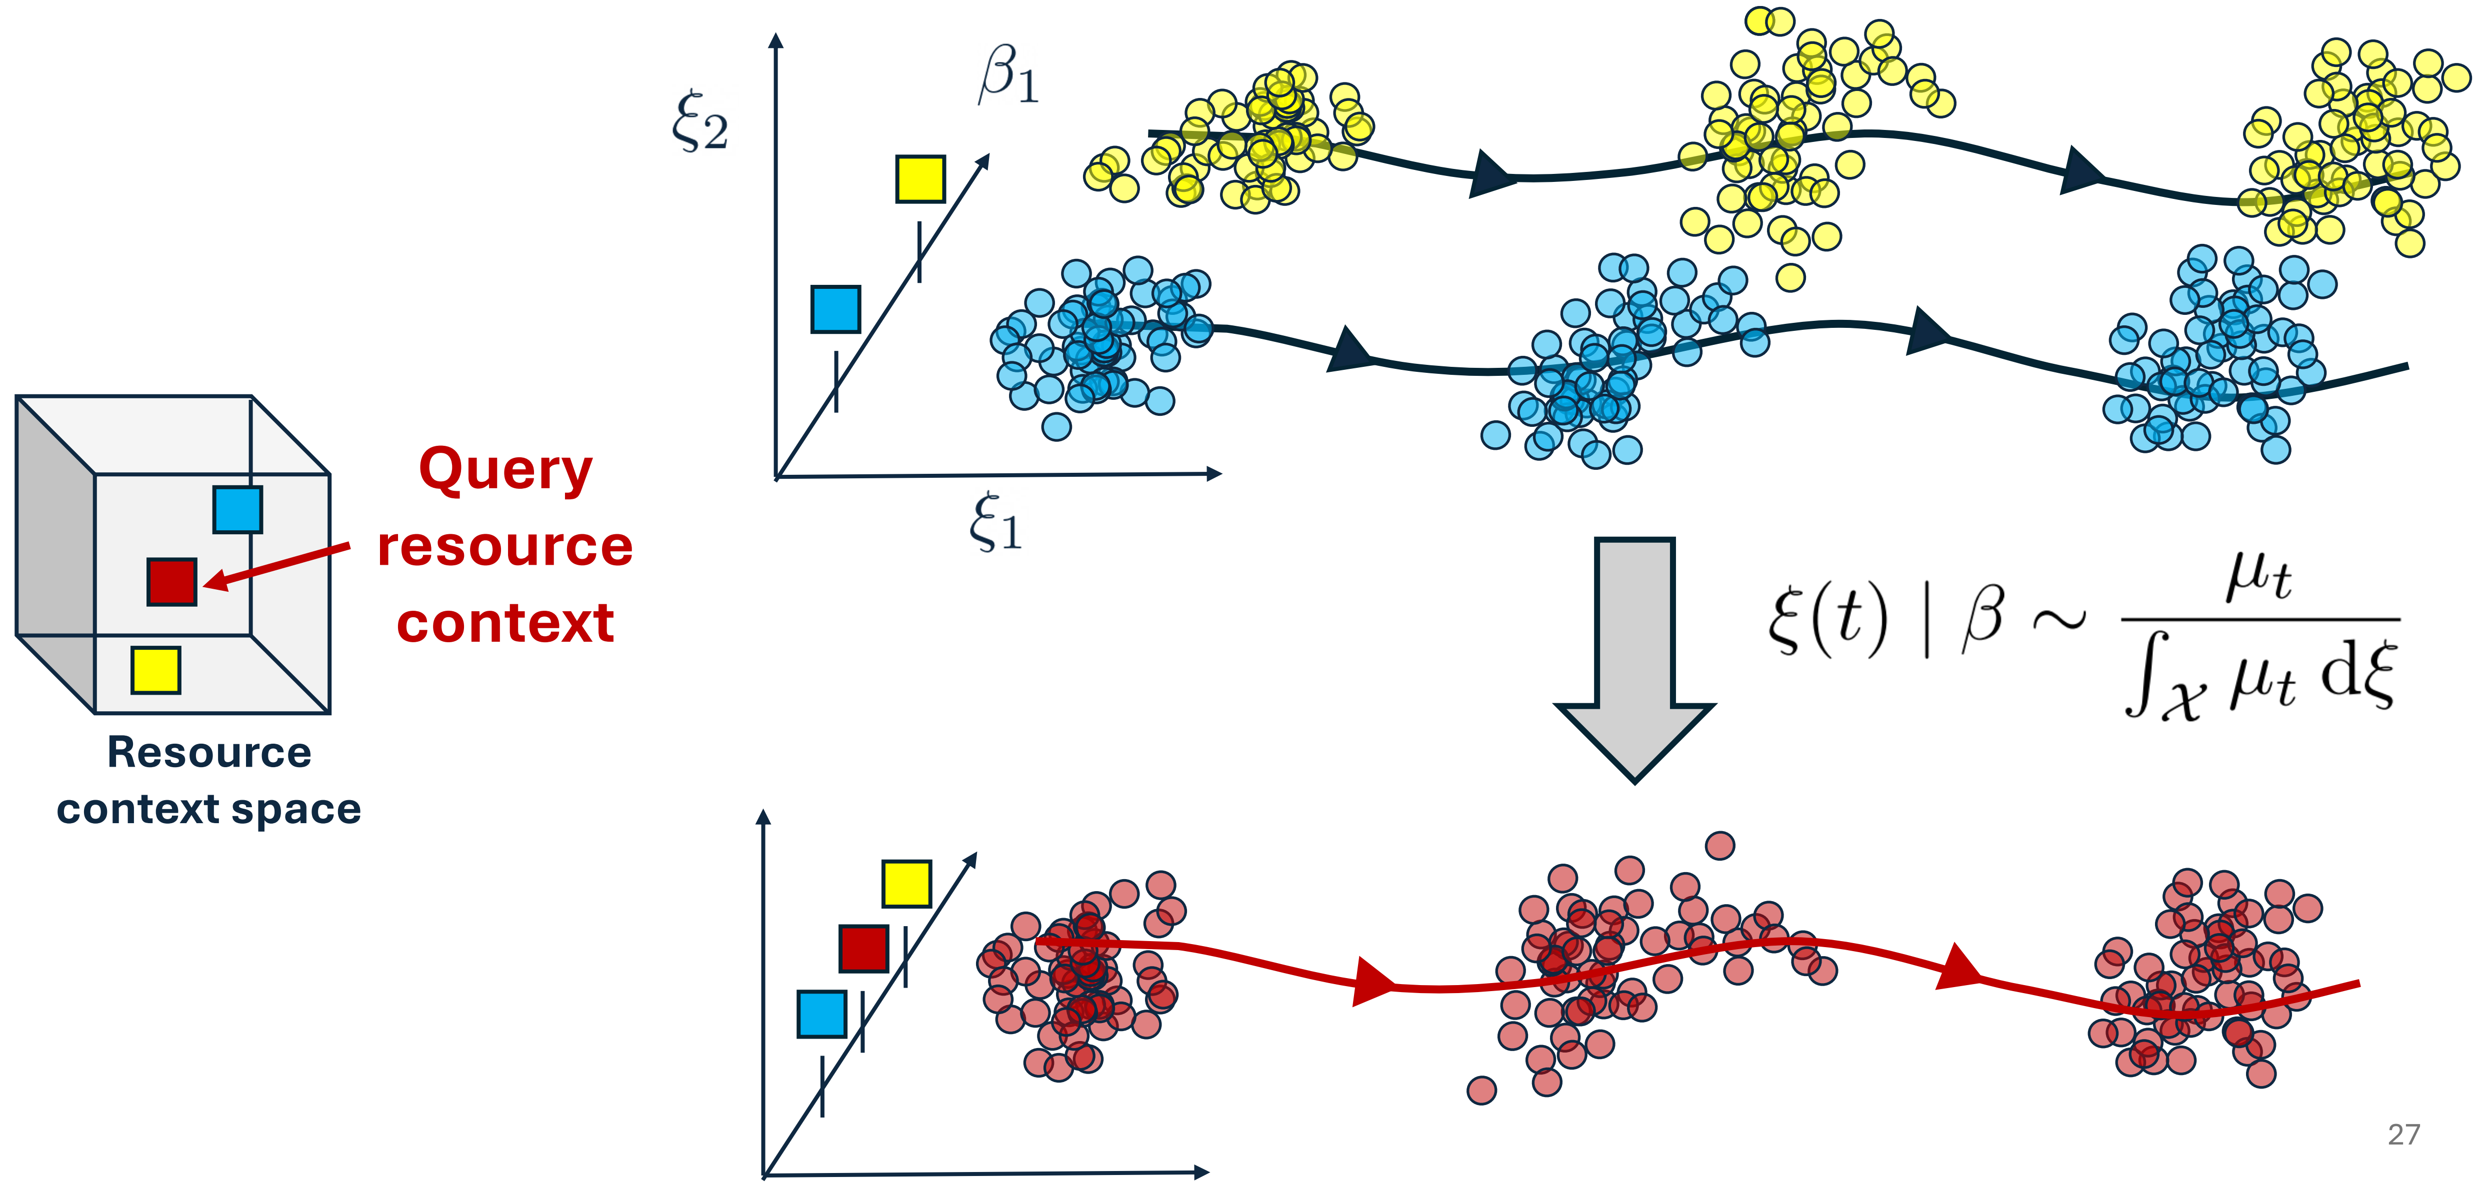

Conditional MSBP Summary

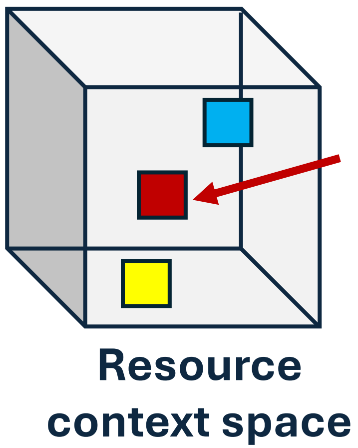

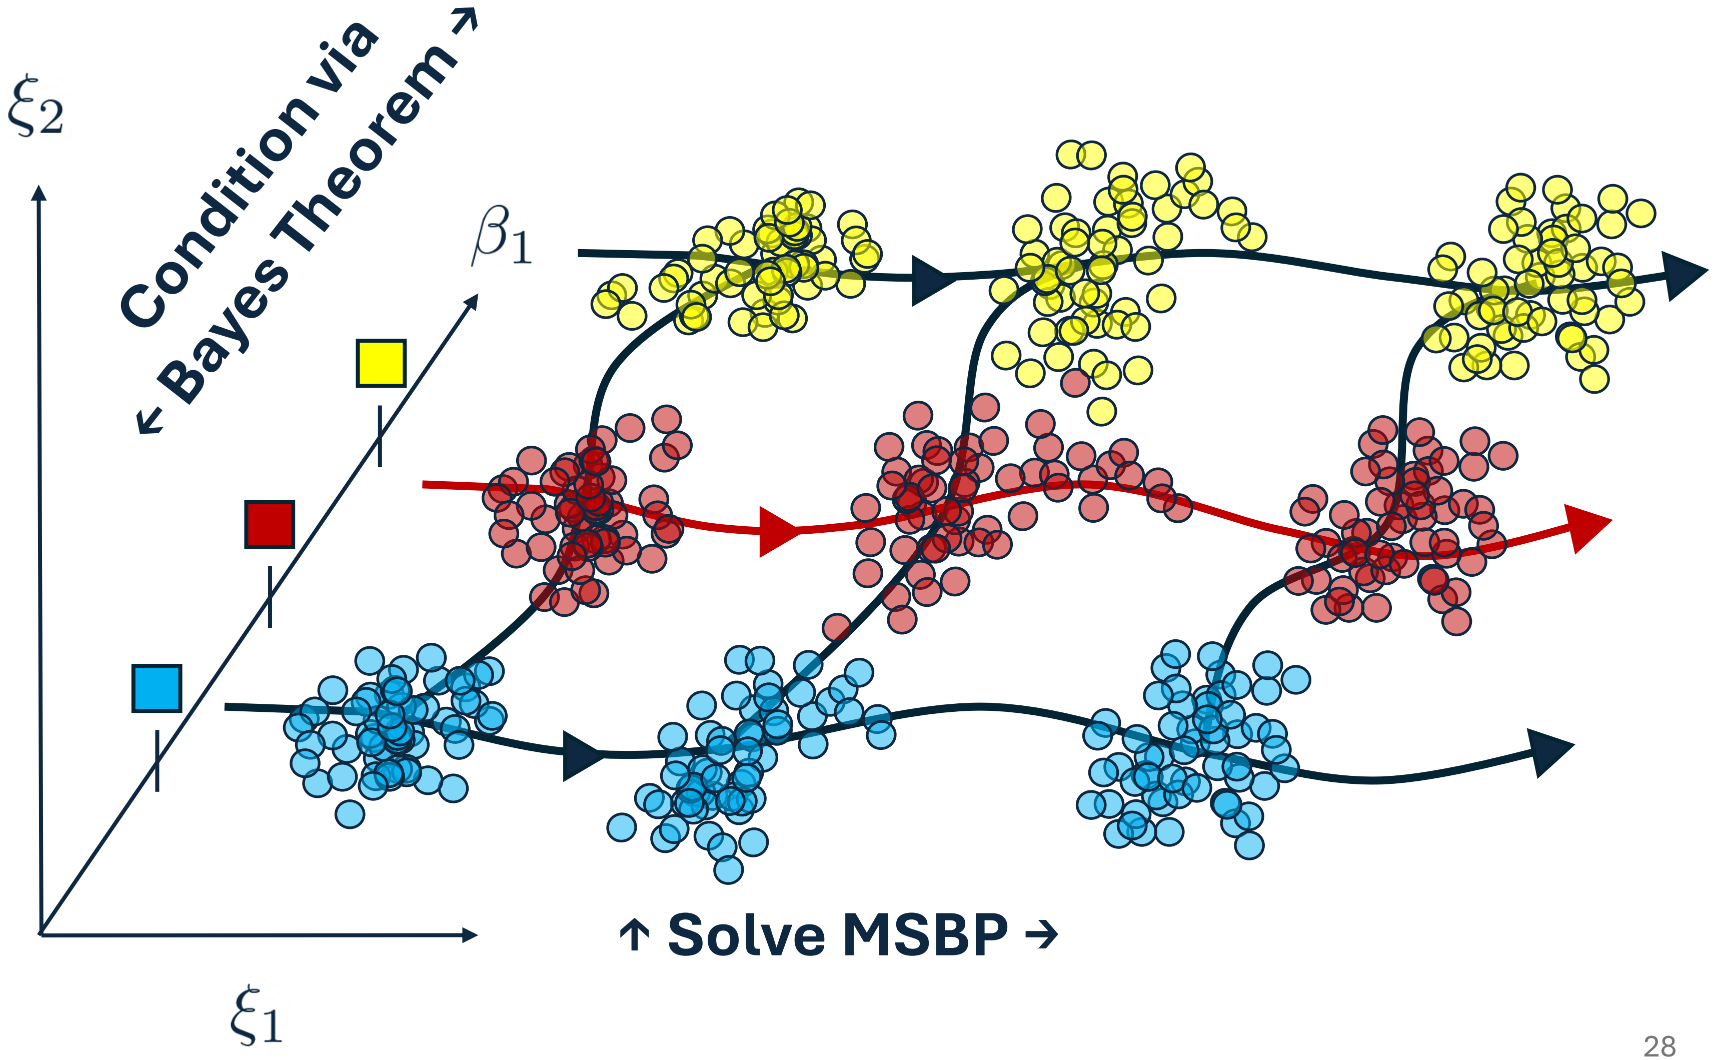

Talk Outline

- SB Background
- Solving for MSBs
- Conditional MSBs
- **Profile Generation**
- Evaluation

Generating Profiles via Sampling

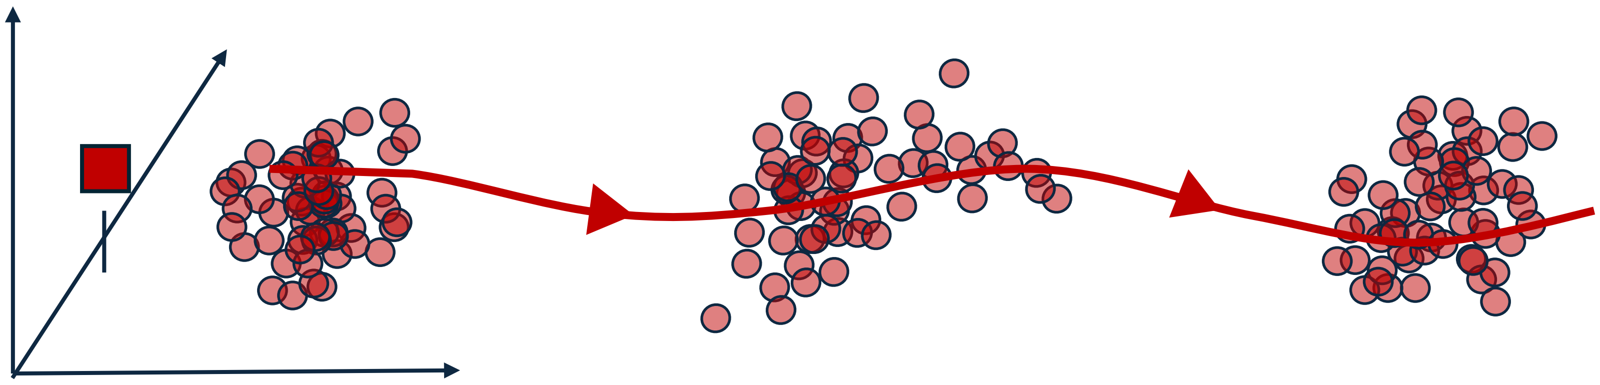

Sample **mean** along **distribution-valued curve**

Sample **maximum-likelihood** along **distribution-valued curve**

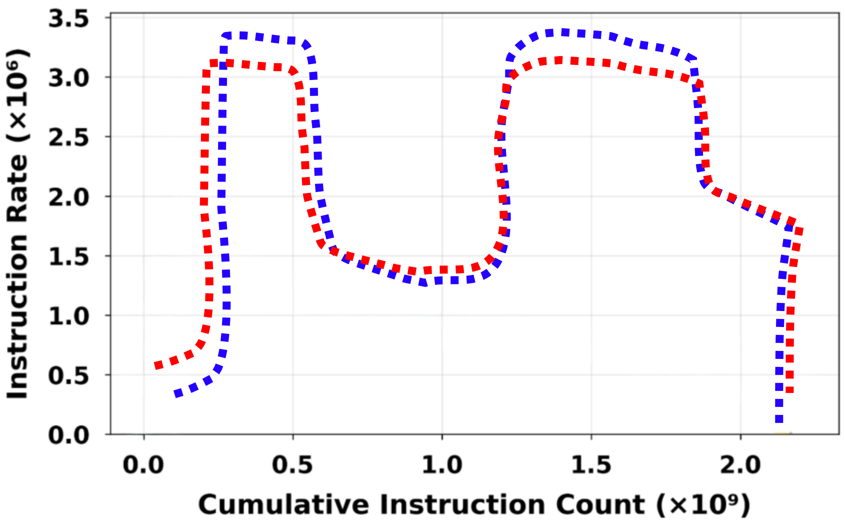

Talk Outline

- SB Background
- Solving for MSBs
- Conditional MSBs
- Profile Generation
- **Evaluation**

Evaluation Setup

- **Benchmarks:** Single-threaded PARSEC with input `simsmall`
- **Hardware:** Intel Xeon E5-2618L v3 processor
- **Resource contexts (4560 unique β):**
  - CAT for LLC partitioning (β_1)
  - MemGuard for memory bandwidth regulation (β_2)
  - DVFS for CPU frequency scaling (β_3)
- **Measured/generated execution state:**
  - Instructions retired (ξ_1)
  - LLC requests (ξ_2)
  - LLC misses (ξ_3)

Accuracy Visualized

- Visualized **mean** and **maximum likelihood** generative profiles overlaid on **empirical profiles**

- Resource context unknown to solver: β = (7, 12, 2.1)

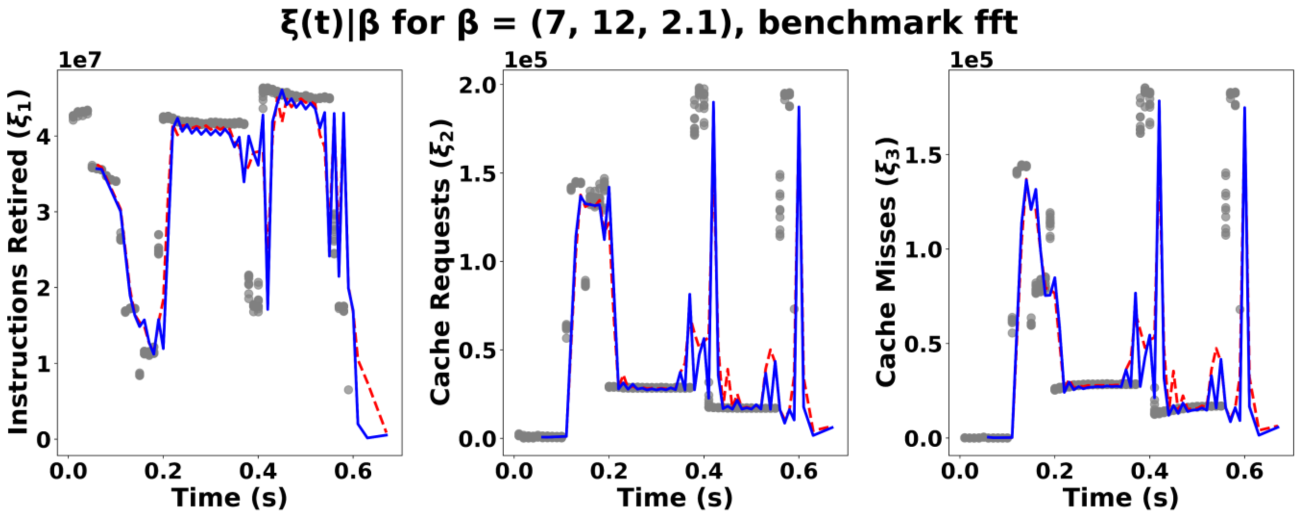

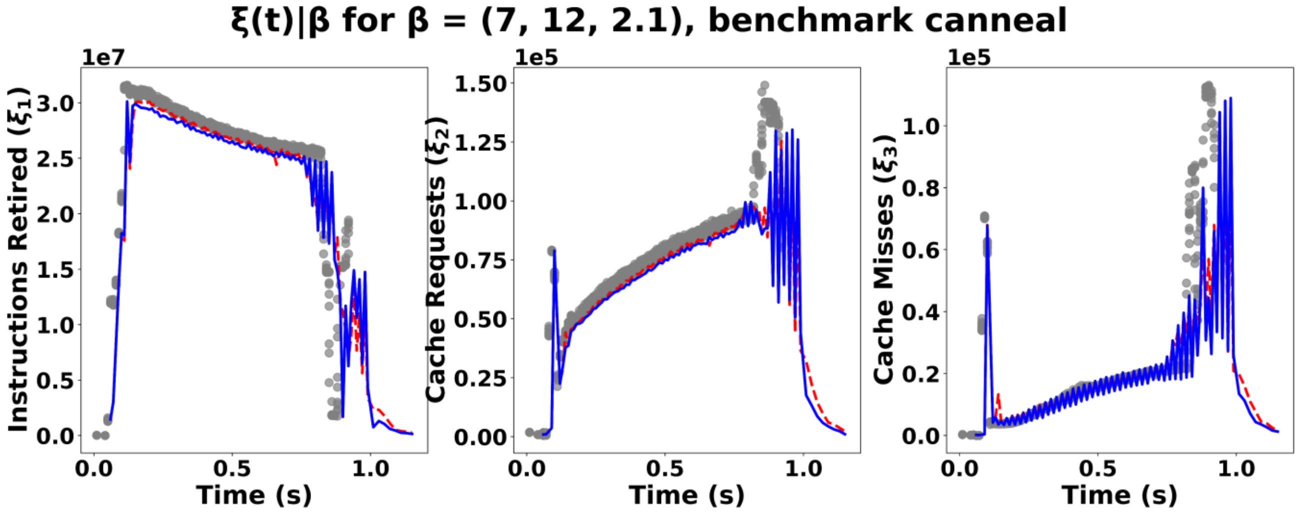

Sparsity Evaluation

- Q: How many resource contexts should be empirically profiled to obtain accurate generative profiles?

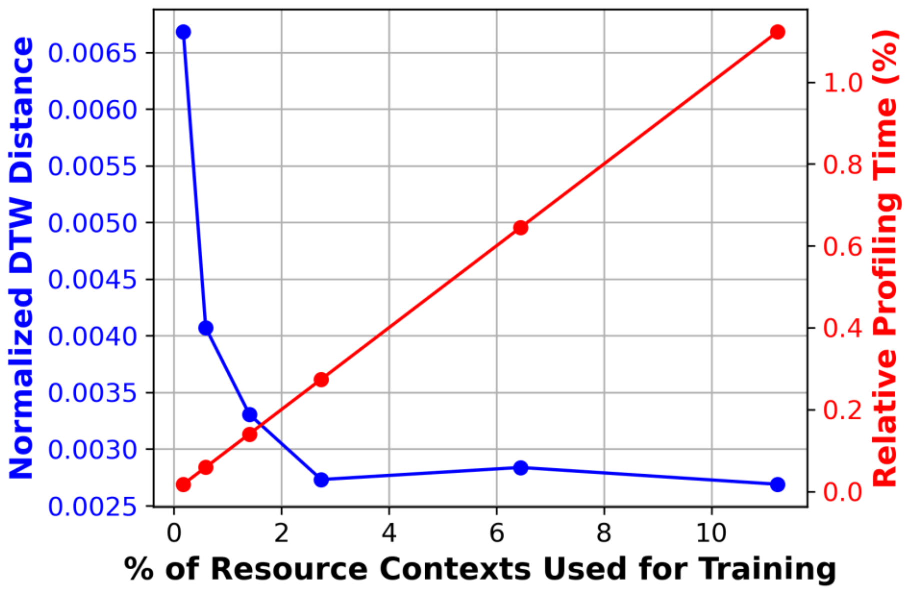

- **Error metric:** normalized Dynamic Time Warping (DTW) distance

Accuracy Evaluation vs. Baseline

- Q: Is a simple solution just as accurate?

	Baseline	Generative	Improvement
blackscholes	0.0429	0.0343	20.0%
bodytrack	0.0437	0.0375	14.2%
canneal	0.0227	0.0027	88.1%
dedup	0.0321	0.0191	40.5%
fft	0.0508	0.0478	5.9%
fluidanimate	0.0357	0.0275	23.0%
radiosity	0.0404	0.0380	5.9%
streamcluster	0.0549	0.0418	23.9%

- **Baseline:** interpolation between known resource contexts

Case Study

- Q: Does this level of improvement in the DTW distance matter?

- **DNA [3]:** phase-aware dynamic cache and memory bandwidth allocator that uses fine-grained context-dependent profiles
- **DVFS-DNA (this case study):** extension of DNA that selects minimal CPU frequency for each phase and resource context

[3] R. Gifford, N. Gandhi, L. T. X. Phan, and A. Haeberlen. DNA: Dynamic resource allocation for soft real-time multicore systems. In RTAS, 2021.

Case Study

- **DNA [3]:** phase-aware dynamic cache and memory bandwidth allocator that uses fine-grained context-dependent profiles
- **DVFS-DNA (this case study):** extension of DNA that selects minimal CPU frequency for each phase and resource context

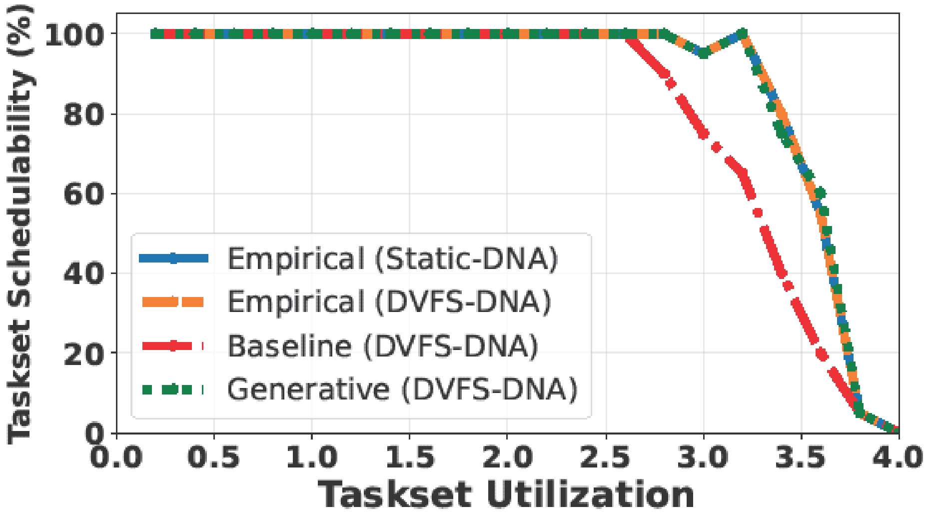

[3] R. Gifford, N. Gandhi, L. T. X. Phan, and A. Haeberlen. DNA: Dynamic resource allocation for soft real-time multicore systems. In RTAS, 2021.

Case Study

- **DNA [3]:** phase-aware dynamic cache and memory bandwidth allocator that uses fine-grained context-dependent profiles
- **DVFS-DNA (this case study):** extension of DNA that selects minimal CPU frequency for each phase and resource context

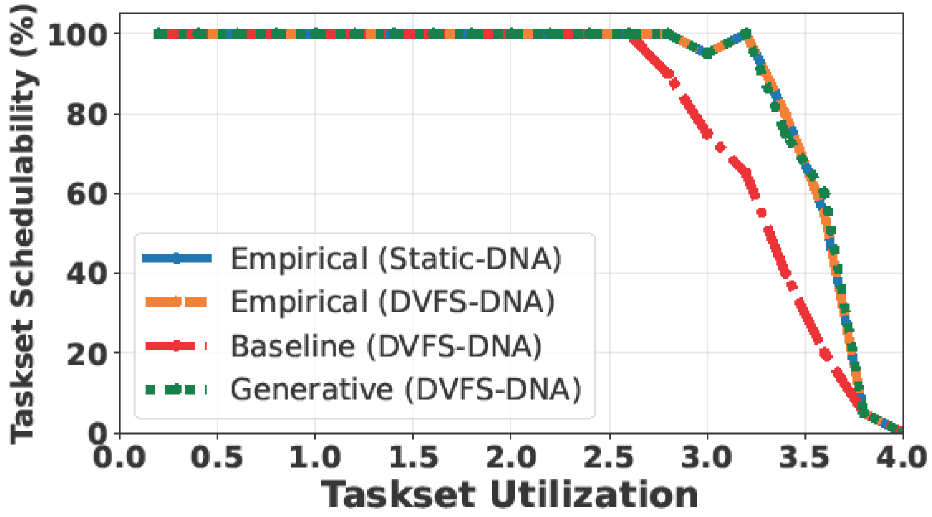

DVFS-DNA energy reduction:
18.1% with empirical profiles
14.8% with generative profiles

[3] R. Gifford, N. Gandhi, L. T. X. Phan, and A. Haeberlen. DNA: Dynamic resource allocation for soft real-time multicore systems. In RTAS, 2021.

Reduction in Profiling Time per Benchmark

- Q: How much profiling time did we actually save?

	Total Time (hrs)	# Empirical Profiles
Empirical	231	456,000
Baseline	0.64	1,250
Generative	1.14	1,250

Conclusion

- We presented context-dependent fine-grained **generative profiling**
- **Properties:**
  - No parametric assumptions
  - Extends to arbitrary resource context and execution state dimensions
  - Consistency with known observations
  - Maximum-likelihood guarantee

Conclusion

- We presented context-dependent fine-grained **generative profiling**
- **Properties:**
  - No parametric assumptions
  - Extends to arbitrary resource context and execution state dimensions
  - Consistency with known observations
  - Maximum-likelihood guarantee
- **Experimental findings:**
  - Orders of magnitude reduction in profiling time
  - Generates profiles for unseen resource contexts with high accuracy
  - Maintains real-time performance in a resource allocation case study

Conclusion

- We presented context-dependent fine-grained **generative profiling**
- **Properties:**
  - No parametric assumptions
  - Extends to arbitrary resource context and execution state dimensions
  - Consistency with known observations
  - Maximum-likelihood guarantee
- **Experimental findings:**
  - Orders of magnitude reduction in profiling time
  - Generates profiles for unseen resource contexts with high accuracy
  - Maintains real-time performance in a resource allocation case study

Thank you!