

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

Georgiy A. Bondar, Robert Gifford, Linh Thi Xuan Phan, Abhishek Halder

Abstract—Solution of the path structured multimarginal Schrödinger bridge problem (MSBP) is the most-likely measure-valued trajectory consistent with a sequence of observed probability measures or distributional snapshots. We leverage recent algorithmic advances in solving such structured MSBPs for learning stochastic hardware resource usage by control software. The solution enables predicting the time-varying distribution of hardware resource availability at a desired time with guaranteed linear convergence. We demonstrate the efficacy of our probabilistic learning approach in a model predictive control software execution case study. The method exhibits rapid convergence to an accurate prediction of hardware resource utilization of the controller. The method can be broadly applied to any software to predict cyber-physical context-dependent performance at arbitrary time.

I. INTRODUCTION

Control software in safety-critical cyber-physical systems (CPS) is often designed and verified based on platform models that do not fully capture the complexity of its deployment settings. For example, it is common to assume that the processor is dedicated to the control software and that overhead is negligible. In practice, hardware resources – such as last-level shared cache (LLC), memory bandwidth, and processor cycles – often vary with time and hardware state, which why we observe varying execution times across different runs of the same control software [1]. This gap can lead to inefficient or unsafe design.

Measurement-based approaches and overhead-aware analysis can reduce the analysis pessimism or ensure safety [2]. The recent work [3] uses fine-grained profiles of the software execution for dynamic scheduling and resource allocation. Supervisory algorithms that dynamically switch among a set of controllers depending on the resource availability also exist [4]. However, the effectiveness of these techniques is contingent on the quality of prediction of future resource availability and on the time horizon of interest.

Hardware resources are not only time-varying and stochastic, but they are also statistically correlated. It is a challenge to predict the *joint stochastic variability* of the hardware resource availability in general, and more so for control

Georgiy A. Bondar is with the Department of Applied Mathematics, University of California, Santa Cruz, CA 95064, USA, gbondar@ucsc.edu.

Robert Gifford and Linh Thi Xuan Phan are with the Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA 19104, USA, {rgif,linhphan}@seas.upenn.edu.

Abhishek Halder is with the Department of Aerospace Engineering, Iowa State University, Ames, IA 50011, USA, ahalder@iastate.edu. This work was supported by NSF grants 2112755, 2111688 and 1750158.

software, where computational cost depends on additional context, e.g., a reference trajectory being tracked.

This work proposes learning a joint stochastic process for hardware resource availability from control software execution profiles conditioned on CPS contexts (to be made precise in Sec. III-A, III-B) based on only a small set of measurements. Our proposed method leverages recent advances in *stochastic control* – specifically in the multimarginal Schrödinger bridge (MSBP) – to allow prediction of time-varying joint statistical distributions of hardware resource availability at any desired time. For safety-critical CPS such predictions, as opposed to those of a lumped variable such as worst-case execution time, can enable the design of improved dynamic scheduling algorithms.

II. NOTATIONS AND PRELIMINARIES

Square braces are used to denote the components. For instance, $[\mathbf{X}_{i_1, \dots, i_r}]$ denotes the (i_1, \dots, i_r) th component of the order r tensor \mathbf{X} , where $(i_1, \dots, i_r) \in \mathbb{N}^r$. We use the r fold tensor product space notation $(\mathbb{R}^d)^{\otimes r} := \underbrace{\mathbb{R}^d \otimes \dots \otimes \mathbb{R}^d}_{r \text{ times}}$.

For two tensors \mathbf{X}, \mathbf{Y} of order r , we define their *Hilbert-Schmidt inner product* as

$$\langle \mathbf{X}, \mathbf{Y} \rangle := \sum_{i_1, \dots, i_r} [\mathbf{X}_{i_1, \dots, i_r}] [\mathbf{Y}_{i_1, \dots, i_r}]. \quad (1)$$

The operators $\exp(\cdot)$ and $\log(\cdot)$ are understood elementwise. We use \odot and \oslash to denote elementwise (Hadamard) multiplication and division, respectively.

For measures μ, ν defined on two Polish spaces, their product measure is denoted by $\mu \otimes \nu$. The *relative entropy* a.k.a. *Kullback-Leibler divergence* $D_{\text{KL}}(\cdot \parallel \cdot)$ between probability measures μ and ν is

$$D_{\text{KL}}(\mu \parallel \nu) := \begin{cases} \int \log \frac{d\mu}{d\nu} d\mu & \text{if } \mu \ll \nu, \\ +\infty & \text{otherwise,} \end{cases} \quad (2)$$

where $\frac{d\mu}{d\nu}$ denotes the Radon-Nikodym derivative, and $\mu \ll \nu$ is a shorthand for “ μ is absolutely continuous w.r.t. ν ”.

The Hilbert (projective) metric (see e.g., [5]) $d_H(\mathbf{u}, \mathbf{v})$ between two vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}_{>0}^n$ is

$$d_H(\mathbf{u}, \mathbf{v}) = \log \left(\frac{\max_{i=1, \dots, n} u_i/v_i}{\min_{i=1, \dots, n} u_i/v_i} \right). \quad (3)$$

We use the term “control cycle” to mean one pass of a feedback control loop. Due to hardware stochasticity, each control cycle completion takes variable amount of time.

III. PROBLEM FORMULATION

A. Context c

We consider a context vector c comprised of separable cyber and physical context vectors

$$c := \begin{pmatrix} c_{\text{cyber}} \\ c_{\text{phys}} \end{pmatrix}. \quad (4)$$

In this work, we consider an instance of (4) where

$$c_{\text{cyber}} = \begin{pmatrix} \text{allocated last-level cache} \\ \text{allocated memory bandwidth} \end{pmatrix}, \quad (5)$$

where both features are allocated in blocks of some size, and

$$c_{\text{phys}} = y_{\text{des}}(x) \in \text{GP}([x_{\min}, x_{\max}]), \quad (6)$$

where GP denotes a Gaussian process over the domain $[x_{\min}, x_{\max}]$. We work with a collection of contexts with cardinality n_{context} , i.e., a sample of contexts $\{c^i\}_{i=1}^{n_{\text{context}}}$.

B. Hardware Resource State ξ

For concreteness, we define a hardware resource state or feature vector used in our numerical case study (Sec. IV):

$$\xi := \begin{pmatrix} \xi_1 \\ \xi_2 \\ \xi_3 \end{pmatrix} = \begin{pmatrix} \text{instructions retired} \\ \text{LLC requests} \\ \text{LLC misses} \end{pmatrix}. \quad (7)$$

The three elements of ξ denote the number of CPU instructions, the number of LLC requests, and the number of LLC misses in the last time unit (10 ms in our profiling), respectively.

We emphasize that our proposed method is not limited by what specific components comprise ξ . To highlight this flexibility, we describe the proposed approach for $\xi \in \mathbb{R}^d$ with suitable interpretations for the specific application.

For a time interval $[0, t]$ of interest, we think of time-varying ξ as a continuous time vector-valued stochastic process over subsets of \mathbb{R}^d . Suppose that $s \in \mathbb{N}, s \geq 2$ snapshots or observations are made for the stochastic state $\xi(\tau), 0 \leq \tau \leq t$, at (possibly non-equispaced) instances

$$\tau_1 \equiv 0 < \tau_2 < \dots < \tau_{s-1} < \tau_s \equiv t.$$

Consider the snapshot index set $\llbracket s \rrbracket := \{1, \dots, s\}$. For a fixed context c , the snapshot observations comprise a sequence of joint probability measures $\{\mu_\sigma\}_{\sigma \in \llbracket s \rrbracket}$ satisfying $\int d\mu_\sigma(\xi(\tau_\sigma)) = 1$. In other words,

$$\xi(\tau_\sigma) \sim \mu_\sigma \quad \forall \sigma \in \llbracket s \rrbracket. \quad (8)$$

In our application, the data $\{\mu_\sigma\}_{\sigma \in \llbracket s \rrbracket}$ comes from control software execution profiles, i.e., by executing the same control software for the same c with all parameters and initial conditions fixed. So the stochasticity in $\xi(\tau_\sigma)$ stems from the dynamic variability in hardware resource availability.

In particular, for finitely many (say n) execution profiles, we consider empirical distributions (a.k.a. atomic measures)

$$\mu_\sigma := \frac{1}{n} \sum_{i=1}^n \delta(\xi - \xi^i(\tau_\sigma)), \quad (9)$$



Fig. 1: The path tree for sequentially observed $\{\mu_\sigma\}_{\sigma \in \llbracket s \rrbracket}$.

where $\delta(\xi - \xi^i(\tau_\sigma))$ denotes the Dirac delta at sample location $\xi^i(\tau_\sigma)$ where $i \in \llbracket n \rrbracket$, $\sigma \in \llbracket s \rrbracket$. At any snapshot index $\sigma \in \llbracket s \rrbracket$, the set $\{\xi^i(\tau_\sigma)\}_{i=1}^n$ is *scattered data*.

Given the data (8)-(9), we would like to predict the *most likely* hardware resource state statistics

$$\xi(\tau) \sim \mu_\tau \quad \text{for any } \tau \in [0, t]. \quad (10)$$

Without the qualifier “most likely”, the problem is over-determined since there are uncountably many *measure-valued continuous curves* over $[0, t]$ that are consistent with the observed data (8)-(9).

C. Multimarginal Schrödinger Bridge

Let $\mathcal{X}_\sigma := \text{support}(\mu_\sigma) \subseteq \mathbb{R}^d \forall \sigma \in \llbracket s \rrbracket$, and consider the Cartesian product $\mathcal{X}_1 \times \mathcal{X}_2 \times \dots \times \mathcal{X}_s =: \mathcal{X} \subseteq (\mathbb{R}^d)^{\otimes s}$. Let $\mathcal{M}(\mathcal{X}_\sigma)$ and $\mathcal{M}(\mathcal{X})$ denote the collection (i.e., manifold) of probability measures on \mathcal{X}_σ and \mathcal{X} , respectively. Define a ground cost $C : \mathcal{X} \mapsto \mathbb{R}_{\geq 0}$.

Following [6, Sec. 3], let

$$d\xi_{-\sigma} := d\xi(\tau_1) \times \dots \times d\xi(\tau_{\sigma-1}) \times d\xi(\tau_{\sigma+1}) \times \dots \times d\xi(\tau_s), \quad (11a)$$

$$\mathcal{X}_{-\sigma} := \mathcal{X}_1 \times \dots \times \mathcal{X}_{\sigma-1} \times \mathcal{X}_{\sigma+1} \times \dots \times \mathcal{X}_s. \quad (11b)$$

For $\varepsilon \geq 0$ (not necessarily small), the multimarginal Schrödinger bridge problem (MSBP) is the following infinite dimensional convex program:

$$\min_{M \in \mathcal{M}(\mathcal{X})} \int_{\mathcal{X}} \{C(\xi(\tau_1), \dots, \xi(\tau_s)) + \varepsilon \log M(\xi(\tau_1), \dots, \xi(\tau_s))\} M(\xi(\tau_1), \dots, \xi(\tau_s)) d\xi(\tau_1) \dots d\xi(\tau_s) \quad (12a)$$

$$\text{subject to } \int_{\mathcal{X}_{-\sigma}} M(\xi(\tau_1), \dots, \xi(\tau_s)) d\xi_{-\sigma} = \mu_\sigma \quad \forall \sigma \in \llbracket s \rrbracket. \quad (12b)$$

In particular, $\mathcal{M}(\mathcal{X})$ is a convex set. The objective (12a) is strictly convex in M , thanks to the ε -regularized negative entropy term $\int_{\mathcal{X}} \varepsilon M \log M$. The constraints (12b) are linear.

In this work, the measures $\{\mu_\sigma\}_{\sigma \in \llbracket s \rrbracket}$ correspond to sequential observation, and we therefore fix the *path structured* (Fig. 1) ground cost

$$C(\xi(\tau_1), \dots, \xi(\tau_s)) = \sum_{\sigma=1}^{s-1} c_\sigma(\xi(\tau_\sigma), \xi(\tau_{\sigma+1})). \quad (13)$$

In particular, we choose the squared Euclidean distance sequential cost between two consecutive snapshot indices, i.e., $c_\sigma(\cdot, \cdot) := \|\cdot - \cdot\|_2^2 \forall \sigma \in \llbracket s \rrbracket$. MSBPs with more general *tree structured* ground costs have appeared in [7].

When the cardinality of the index set $\llbracket s \rrbracket$ equals 2, then (12) reduces to the (bi-marginal) Schrödinger bridge problem (SBP) [8], [9]. In this case, the solution of (12) gives the *most likely evolution* between two marginal snapshots μ_1, μ_2 . This follows from the large deviations [10] interpretation [11, Sec. II] of SBP using Sanov’s theorem [12], [13, Sec. 2.1].

Specifically, let $\mathcal{C}([\tau_1, \tau_2], \mathbb{R}^d)$ denote the collection of continuous functions on the time interval $[\tau_1, \tau_2]$ taking values in \mathbb{R}^d . Let $\Pi(\mu_1, \mu_2)$ be the collection of all path

measures on $\mathcal{C}([\tau_1, \tau_2], \mathbb{R}^d)$ with time τ_1 marginal μ_1 , and time τ_2 marginal μ_2 . Given a symmetric ground cost (e.g., Euclidean distance) $C : \mathcal{X}_1 \times \mathcal{X}_2 \mapsto \mathbb{R}_{\geq 0}$, let

$$K(\cdot, \cdot) := \exp\left(-\frac{C(\cdot, \cdot)}{\varepsilon}\right), \quad (14)$$

and consider the *bimarginal Gibbs kernel*

$$K(\xi(\tau_1), \xi(\tau_2)) \mu_1 \otimes \mu_2. \quad (15)$$

Then, the bimarginal SBP solves

$$\min_{\pi \in \Pi(\mu_1, \mu_2)} \varepsilon D_{\text{KL}}(\pi \| K(\xi(\tau_1), \xi(\tau_2)) \mu_1 \otimes \mu_2), \quad (16)$$

i.e., the most likely evolution of the path measure consistent with the observed measure-valued snapshots μ_1, μ_2 .

Under the stated assumptions on the ground cost c , the existence of minimizer for (16) is guaranteed [14], [15]. The uniqueness of minimizer follows from strict convexity of the map $\pi \mapsto D_{\text{KL}}(\pi \| \nu)$ for fixed ν .

This relative entropy reformulation, and thereby “the most likely evolution consistent with observed measures” interpretation, also holds for the MSBP (12) with $s \geq 2$ snapshots. Specifically, for $\mathbf{C} : \mathcal{X} \mapsto \mathbb{R}_{\geq 0}$ as in (12)-(13), we generalize (14) as

$$\mathbf{K}(\xi(\tau_1), \dots, \xi(\tau_s)) := \exp\left(-\frac{C(\xi(\tau_1), \dots, \xi(\tau_s))}{\varepsilon}\right), \quad (17)$$

and define the *multimarginal Gibbs kernel*

$$\mathbf{K}(\xi(\tau_1), \dots, \xi(\tau_s)) \mu_1 \otimes \dots \otimes \mu_s. \quad (18)$$

Problem (16) then generalizes to

$$\min_{\pi \in \Pi(\mu_1, \dots, \mu_s)} \varepsilon D_{\text{KL}}(\pi \| \mathbf{K}(\xi(\tau_1), \dots, \xi(\tau_s)) \mu_1 \otimes \dots \otimes \mu_s) \quad (19)$$

where $\Pi(\mu_1, \dots, \mu_s)$ denotes the collection of all path measures on $\mathcal{C}([\tau_1, \tau_s], \mathbb{R}^d)$ with time τ_σ marginal $\mu_\sigma \forall \sigma \in \llbracket s \rrbracket$. The equivalence between (12) and (19) can be verified by direct computation. Thus solving (19), or equivalently (12), yields the most likely evolution of the path measure consistent with the observed measure-valued snapshots $\mu_\sigma \forall \sigma \in \llbracket s \rrbracket$.

We propose to solve the MSBP (12) for learning the time-varying statistics of the hardware resource state ξ as in (10). We next detail a discrete formulation to numerically solve the same for scattered data $\{\xi^i(\tau_\sigma)\}_{i=1}^n$ where n is the number of control software execution profiles.

The minimizer of (12), $\mathbf{M}_{\text{opt}}(\xi(\tau_1), \dots, \xi(\tau_s))$ can be used to compute the optimal coupling between snapshot index pairs $(\sigma_1, \sigma_2) \in \{\llbracket s \rrbracket^{\otimes 2} \mid \sigma_1 < \sigma_2\}$ as

$$\int_{\mathcal{X}_{-\sigma_1, -\sigma_2}} \mathbf{M}_{\text{opt}}(\xi(\tau_1), \dots, \xi(\tau_s)) d\xi_{-\sigma_1, -\sigma_2} \quad (20)$$

where

$$d\xi_{-\sigma_1, -\sigma_2} := \prod_{\sigma \in \llbracket s \rrbracket \setminus \{\sigma_1, \sigma_2\}} d\xi(\tau_\sigma), \quad (21a)$$

$$\mathcal{X}_{-\sigma_1, -\sigma_2} := \prod_{\sigma \in \llbracket s \rrbracket \setminus \{\sigma_1, \sigma_2\}} \mathcal{X}_\sigma. \quad (21b)$$

This will be useful for predicting the statistics of $\xi(\tau) \sim \mu_\tau$ at any (out-of-sample) query time $\tau \in [0, t]$.

D. Discrete Formulation of MSBP

For finite scattered data $\{\xi^i(\tau_\sigma)\}_{i=1}^n$ and $\{\mu_\sigma\}_{\sigma \in \llbracket s \rrbracket}$ as in (9), we set up a discrete version of (12) as follows.

With slight abuse of notations, we use the same symbol for the continuum and discrete version of a tensor. The ground cost in discrete formulation is represented by an order s tensor $\mathbf{C} \in (\mathbb{R}^n)_{\geq 0}^{\otimes s}$, with components $[\mathbf{C}_{i_1, \dots, i_s}] = \mathbf{C}(\xi_{i_1}, \dots, \xi_{i_s})$. The component $[\mathbf{C}_{i_1, \dots, i_s}]$ encodes the cost of transporting unit mass for a tuple (i_1, \dots, i_s) .

Likewise, consider the discrete mass tensor $\mathbf{M} \in (\mathbb{R}^n)_{\geq 0}^{\otimes s}$ with components $[\mathbf{M}_{i_1, \dots, i_s}] = \mathbf{M}(\xi_{i_1}, \dots, \xi_{i_s})$. The component $[\mathbf{M}_{i_1, \dots, i_s}]$ denotes the amount of transported mass for a tuple (i_1, \dots, i_s) .

For any $\sigma \in \llbracket s \rrbracket$, the empirical marginals $\mu_\sigma \in \mathbb{R}_{\geq 0}^n$ are supported on the finite sets $\{\xi^i(\tau_\sigma)\}_{i=1}^n$. We denote the projection of $\mathbf{M} \in (\mathbb{R}^n)_{\geq 0}^{\otimes s}$ on the σ th marginal as $\text{proj}_\sigma(\mathbf{M})$. Thus $\text{proj}_\sigma : (\mathbb{R}^n)_{\geq 0}^{\otimes s} \mapsto \mathbb{R}_{\geq 0}^n$, and is given componentwise as

$$[\text{proj}_\sigma(\mathbf{M})_j] = \sum_{i_1, \dots, i_{\sigma-1}, i_{\sigma+1}, \dots, i_s} \mathbf{M}_{i_1, \dots, i_{\sigma-1}, j, i_{\sigma+1}, \dots, i_s}. \quad (22)$$

Likewise, denote the projection of $\mathbf{M} \in (\mathbb{R}^n)_{\geq 0}^{\otimes s}$ on the (σ_1, σ_2) th marginal as $\text{proj}_{\sigma_1, \sigma_2}(\mathbf{M})$, i.e., $\text{proj}_{\sigma_1, \sigma_2} : (\mathbb{R}^n)_{\geq 0}^{\otimes s} \mapsto \mathbb{R}_{\geq 0}^{n \times n}$, and is given componentwise as

$$\begin{aligned} & [\text{proj}_{\sigma_1, \sigma_2}(\mathbf{M})_{j, \ell}] \\ &= \sum_{i_\sigma \mid \sigma \in \llbracket s \rrbracket \setminus \{\sigma_1, \sigma_2\}} \mathbf{M}_{i_1, \dots, i_{\sigma-1}, j, i_{\sigma_1+1}, \dots, i_{\sigma_2-1}, \ell, i_{\sigma_2+1}, \dots, i_s}. \end{aligned} \quad (23)$$

We note that (22) and (23) are the discrete versions of the integrals in (12b) and (20), respectively.

With the above notations in place, the discrete version of (12) becomes

$$\min_{\mathbf{M} \in (\mathbb{R}^n)_{\geq 0}^{\otimes s}} \langle \mathbf{C} + \varepsilon \log \mathbf{M}, \mathbf{M} \rangle \quad (24a)$$

$$\text{subject to } \text{proj}_\sigma(\mathbf{M}) = \mu_\sigma \quad \forall \sigma \in \llbracket s \rrbracket. \quad (24b)$$

The primal formulation (24) has n^s decision variables, and is computationally intractable. Recall that even for the bimarginal ($s = 2$) case, a standard approach [16] is to use Lagrange duality to notice that the optimal mass matrix \mathbf{M}_{opt} is a diagonal scaling of $\mathbf{K} := \exp(-C/\varepsilon) \in \mathbb{R}_{>0}^{n \times n}$, i.e., $\mathbf{M}_{\text{opt}} = \text{diag}(\mathbf{u}_1) \mathbf{K} \text{diag}(\mathbf{u}_2)$ where $\mathbf{u}_1 := \exp(\lambda_1/\varepsilon)$, $\mathbf{u}_2 := \exp(\lambda_2/\varepsilon)$, and $\lambda_1, \lambda_2 \in \mathbb{R}^n$ are the Lagrange multipliers associated with respective bimarginal constraints $\text{proj}_1(\mathbf{M}) = \mu_1$, $\text{proj}_2(\mathbf{M}) = \mu_2$. The unknowns $\mathbf{u}_1, \mathbf{u}_2$ can be obtained by performing the Sinkhorn iterations

$$\mathbf{u}_1 \leftarrow \mu_1 \oslash (K \mathbf{u}_2), \quad (25a)$$

$$\mathbf{u}_2 \leftarrow \mu_2 \oslash (K^\top \mathbf{u}_1), \quad (25b)$$

with guaranteed linear convergence [17] wherein the computational cost is governed by two matrix-vector multiplications.

The duality result holds for the multimarginal ($s \geq 2$) case. Specifically, the optimal mass tensor in (24) admits a structure $\mathbf{M}_{\text{opt}} = \mathbf{K} \odot \mathbf{U}$ where $\mathbf{K} := \exp(-\mathbf{C}/\varepsilon) \in$

$(\mathbb{R}^n)_{>0}^{\otimes s}$, $\mathbf{U} := \otimes_{\sigma=1}^s \mathbf{u}_\sigma \in (\mathbb{R}^n)_{>0}^{\otimes s}$, $\mathbf{u}_\sigma := \exp(\boldsymbol{\lambda}_\sigma/\varepsilon)$, and $\boldsymbol{\lambda}_\sigma \in \mathbb{R}^n$ are the Lagrange multipliers associated with the respective multimarginal constraints (24b). The unknowns \mathbf{u}_σ can, in principle, be obtained from the multimarginal Sinkhorn iterations [18]

$$\mathbf{u}_\sigma \leftarrow \mathbf{u}_\sigma \odot \mu_\sigma \odot \text{proj}_\sigma(\mathbf{K} \odot \mathbf{U}) \quad \forall \sigma \in \llbracket s \rrbracket, \quad (26)$$

which generalize (25). However, computing $\text{proj}_\sigma(\mathbf{K} \odot \mathbf{U})$ requires $\mathcal{O}(n^s)$ operations. Before describing how to avoid this exponential complexity (Sec. III-F), we point out the convergence guarantees for (26).

E. Convergence for Multimarginal Sinkhorn Iterations

The iterations (26) can either be derived as alternating Bregman projections [18] or via block coordinate dual ascent [6]. Following either viewpoints leads to guaranteed linear convergence of (26); see [19], [7, Thm. 3.5]. More recent works have also established [20] guaranteed convergence for the continuous formulation (12) with linear rate of convergence [21].

F. Multimarginal Sinkhorn Iterations for Path Structured \mathbf{C}

We circumvent the exponential complexity in computing $\text{proj}_\sigma(\mathbf{K} \odot \mathbf{U})$ in (26) by leveraging the path structured ground cost (13). This is enabled by a key result from [6], rephrased, and reproved in extended version [22].

Proposition 1. ([6, Prop. 2]) Consider the discrete ground cost tensor \mathbf{C} in (24) induced by a path structured cost (13) so that $[\mathbf{C}_{i_1, \dots, i_s}] = \sum_{\sigma=1}^{s-1} \left[C_{i_\sigma, i_{\sigma+1}}^{\sigma \rightarrow \sigma+1} \right]$ where the matrix $C^{\sigma \rightarrow \sigma+1} \in \mathbb{R}_{\geq 0}^{n \times n}$ encodes the cost of transporting unit mass between each source-destination pair from the source set $\{\boldsymbol{\xi}^i(\tau_\sigma)\}_{i=1}^n$ to the destination set $\{\boldsymbol{\xi}^i(\tau_{\sigma+1})\}_{i=1}^n$.

Let $K^{\sigma \rightarrow \sigma+1} := \exp(-C^{\sigma \rightarrow \sigma+1}/\varepsilon) \in \mathbb{R}_{\geq 0}^{n \times n}$, $\mathbf{K} := \exp(-\mathbf{C}/\varepsilon) \in (\mathbb{R}^n)_{>0}^{\otimes s}$, $\mathbf{U} := \otimes_{\sigma=1}^s \mathbf{u}_\sigma \in (\mathbb{R}^n)_{>0}^{\otimes s}$.

Then (22) and (23) can be expressed as

$$\text{proj}_\sigma(\mathbf{K} \odot \mathbf{U}) = \left(\mathbf{u}_1^\top K^{1 \rightarrow 2} \prod_{j=2}^{\sigma-1} \text{diag}(\mathbf{u}_j) K^{j \rightarrow j+1} \right)^\top \odot \mathbf{u}_\sigma \odot \left(\left(\prod_{j=\sigma+1}^{s-1} K^{j-1 \rightarrow j} \text{diag}(\mathbf{u}_j) \right) K^{s-1 \rightarrow s} \mathbf{u}_s \right) \quad \forall \sigma \in \llbracket s \rrbracket, \quad (27)$$

and

$$\begin{aligned} \text{proj}_{\sigma_1, \sigma_2}(\mathbf{K} \odot \mathbf{U}) &= \text{diag} \left(\mathbf{u}_1^\top K^{1 \rightarrow 2} \prod_{j=2}^{\sigma_1-1} \text{diag}(\mathbf{u}_j) K^{j \rightarrow j+1} \right) \\ &\quad \text{diag}(\mathbf{u}_{\sigma_1}) \prod_{j=\sigma_1+1}^{\sigma_2} (K^{j-1 \rightarrow j} \text{diag}(\mathbf{u}_j)) \\ &\quad \text{diag} \left(\left(\prod_{j=\sigma_2+1}^{s-1} K^{j-1 \rightarrow j} \text{diag}(\mathbf{u}_j) \right) K^{s-1 \rightarrow s} \mathbf{u}_s \right) \\ &\quad \forall (\sigma_1, \sigma_2) \in \{\llbracket s \rrbracket^{\otimes 2} \mid \sigma_1 < \sigma_2\}. \end{aligned} \quad (28)$$

Remark 1. Substituting (27) into (26) further simplifies our multimarginal Sinkhorn recursions to

$$\begin{aligned} \mathbf{u}_\sigma &\leftarrow \mu_\sigma \oslash \left(\left(\mathbf{u}_1^\top K^{1 \rightarrow 2} \prod_{j=2}^{\sigma-1} \text{diag}(\mathbf{u}_j) K^{j \rightarrow j+1} \right)^\top \odot \right. \\ &\quad \left. \left(\left(\prod_{j=\sigma+1}^{s-1} K^{j-1 \rightarrow j} \text{diag}(\mathbf{u}_j) \right) K^{s-1 \rightarrow s} \mathbf{u}_s \right) \right) \quad \forall \sigma \in \llbracket s \rrbracket. \end{aligned} \quad (29)$$

Remark 2. (From exponential to linear complexity in s)

Note that (29) involves $s-1$ matrix-vector multiplications of $\mathcal{O}(n^2)$ complexity. So the computational complexity for (29) becomes $\mathcal{O}((s-1)n^2)$ which is linear in s – a significant reduction from $\mathcal{O}(n^s)$ as mentioned at the end of Sec. III-D. The recent work [23] further reduces this complexity to $\mathcal{O}((s-1)n)$ by approximating the matrix-vector products using nonuniform fast Fourier transform.

Remark 3. (Linear complexity in d) The dimension d of the vector $\boldsymbol{\xi}$ only affects the construction of the time-varying Euclidean distance matrices $C^{\sigma \rightarrow \sigma+1} \forall \sigma \in \llbracket s-1 \rrbracket$ in Prop. 1, which has total complexity $\mathcal{O}(sd)$. Once constructed, the recursions (29) are independent of d .

G. Predicting Most Likely Distribution

For the ground cost (13) resulting from sequential information structure (Fig. 1), we utilize (28) to decompose $\mathbf{M}_{\text{opt}} = \mathbf{K} \odot \mathbf{U}$ of (24) into bimarginal transport plans

$$M^{\sigma_1 \rightarrow \sigma_2} := \text{proj}_{\sigma_1, \sigma_2}(\mathbf{M}_{\text{opt}}) = \text{proj}_{\sigma_1, \sigma_2}(\mathbf{K} \odot \mathbf{U}). \quad (30)$$

Further, when \mathbf{C} is squared Euclidean, as we consider here, the maximum likelihood estimate for μ_τ in (10) for a query point $\tau \in [0, t]$, is (see [6, Sec. 2.2])

$$\hat{\mu}_\tau := \sum_{i=1}^n \sum_{j=1}^n [M_{i,j}^{\sigma \rightarrow \sigma+1}] \delta(\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}(\tau, \boldsymbol{\xi}^i(\tau_\sigma), \boldsymbol{\xi}^j(\tau_{\sigma+1}))) \quad (31)$$

where $\sigma \in \llbracket s \rrbracket$ such that $\tau \in [\tau_\sigma, \tau_{\sigma+1}]$, and

$$\hat{\boldsymbol{\xi}}(\tau, \boldsymbol{\xi}^i(\tau_\sigma), \boldsymbol{\xi}^j(\tau_{\sigma+1})) := (1-\lambda)\boldsymbol{\xi}^i(\tau_\sigma) + \lambda\boldsymbol{\xi}^j(\tau_{\sigma+1}), \quad (32a)$$

$$\lambda := \frac{\tau - \tau_\sigma}{\tau_{\sigma+1} - \tau_\sigma} \in [0, 1]. \quad (32b)$$

H. Overall Algorithm

Our proposed method comprises of following three steps.

Step 1. Given a collection of contexts (Sec. III-A) $\{\mathbf{c}^i\}_{i=1}^{n_{\text{context}}}$, execute the control software over $[0, t]$ to generate hardware resource state sample snapshots (Sec. III-B) $\{\boldsymbol{\xi}^i(\tau_\sigma)\}_{i=1}^n$, and thereby empirical μ_σ as in (9) for all $\sigma \in \llbracket s \rrbracket$, conditional on each of the n_{context} context samples.

Step 2. Using data from **Step 1**, construct Euclidean distance matrices $C^{\sigma \rightarrow \sigma+1}$ from the source set $\{\boldsymbol{\xi}^i(\tau_\sigma)\}_{i=1}^n$ to the destination set $\{\boldsymbol{\xi}^i(\tau_{\sigma+1})\}_{i=1}^n \forall \sigma \in \llbracket s-1 \rrbracket$. Perform recursions (29) until convergence (error within desired tolerance).

Step 3. Given a query context \mathbf{c} and time $\tau \in [0, t]$, return most likely distribution $\hat{\mu}_\tau$ using (31).

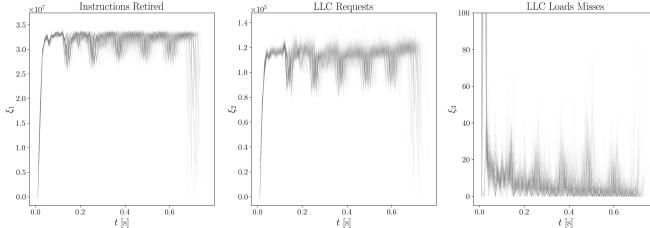


Fig. 2: Components of the measured feature vector ξ in (7) for all of the five control cycles for 500 executions of the NMPC software, where $c = [15, 15, y_{\text{des}}^1(x)]$.

IV. NUMERICAL CASE STUDY

In this Section, we illustrate the application of the proposed method for a vehicle path tracking control software.

Control Software. We implemented¹ in C language a path following nonlinear model predictive controller (NMPC) for a kinematic bicycle model (KBM) [24], [25]. At each control step (≤ 100 ms) the IPOPT nonlinear program solver [26] solved NMPC optimization problem to make a control decision. For formulation details, we refer the readers to [27]; for implementation details see footnote and [22].

While closing the control loop incurs minimal computational overhead, the NMPC is computationally demanding. When multiple vehicle controllers are available it is of practical interest to predict their hardware resource usage for one to several control cycles, conditional on the CPS context c (Sec. III-A) at a given time. For this we ‘profile’ the NMPC, meaning we run the software many times for different values of c as in (4), measuring time evolution of the hardware resource state ξ as in (7). We use these profiles to generate marginals μ_σ as in (9) (**Step 1**, Sec. III-H).

Generating Execution Profiles. For profiling our NMPC control software, we used an Ubuntu 16.04.7 Linux machine with an Intel Xeon E5-2683 v4 CPU. We leveraged Intel’s Cache Allocation Technology (CAT) [28] and Memguard [29] to control allocation of LLC partitions and memory bandwidth to the control software (in blocks of 2MB), respectively. Our application ran on an isolated CPU and used the Linux perf tool [30] to sample ξ every 10 ms.

For each run of our application, we set the cache and memory bandwidth to a static allocation and pass as input a path for the NMPC to follow. We ran the control software for $n_c := 5$ uninterrupted “control cycles”, wherein the NMPC gets the KBM state, makes a control decision, and updates the KBM state. For each of 60 unique contexts c we ran the software for 500 profiles, for a total of 30,000 profiles.

The sample paths $\{y_{\text{des}}^i(x)\}_{i=1}^{12}$ in (6) were all generated using a GP, and the samples $\{c_{\text{cyber}}^i\}_{i=1}^5$ in (5) were $[1, 1]^\top$, $[5, 5]^\top$, $[10, 10]^\top$, $[15, 15]^\top$, and $[20, 20]^\top$, where each entry represents the number of cache/memory bandwidth partitions from 1 to 20; see [22, Fig. 2].

Applying the Proposed Algorithm. Given a query context c , we determine the closest CPS context for which profiling data is available, using the Euclidean distance between cyber context vectors (5), and the Fréchet distance [31] between

¹Git repo: <https://github.com/abhishekhalder/CPS-Frontier-Task3-Collaboration>

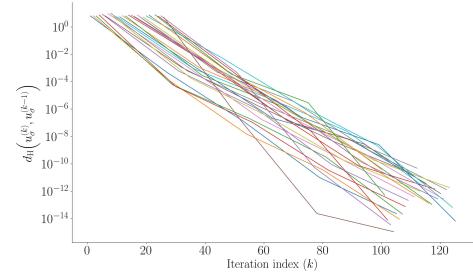


Fig. 3: Linear convergence of Sinkhorn iterations (29) for $s_{\text{int}} = 4$ w.r.t. the Hilbert’s projective metric d_H in (3) between $u_{\sigma \in \llbracket s \rrbracket}$ at iteration indices k and $k - 1$

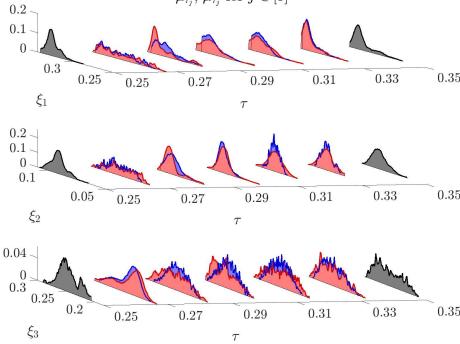


Fig. 4: Predicted $\hat{\mu}_{\hat{\tau}_j}$ (blue) vs. measured $\mu_{\hat{\tau}_j}$ (red) at times $\hat{\tau}_j \in \llbracket 5 \rrbracket$ during the 3rd control cycle with $s_{\text{int}} = 4$. Distributions at the control cycle boundaries are in black.

physical context curves (6). In this case study, we consider a query context with closest $c_{\text{cyber}} = [15, 15]^\top$ and closest $c_{\text{phys}} = y_{\text{des}}^1(x)$. Profiling data for this c is shown in Fig. 2.

We placed marginals at the boundaries of each of the $n_c = 5$ control cycles, using a kernel density estimator (KDE) to find the average end times (see [22, Fig. 4 and Table I] for details). For empirical distributions at times between cycle boundaries, we let s_{int} be the number of marginals equispaced-in-time between each cycle boundary, and we then set $\tau_{\sigma \in \llbracket s \rrbracket}$ to be the control cycle end times, where $s := 1 + n_c(s_{\text{int}} + 1)$ and $\tau_{\sigma(s_{\text{int}}+1)+1}$ is the sampled mean end time for the σ th control cycle.

Our distributions are as per (9), where $\xi^i(\tau_\sigma)$ is the sample of the hardware resource state (7) at time $\approx \tau_\sigma$ (within 5ms) for profile i given context c .

We set $\varepsilon = 0.1$ and solve the discrete MSBP (24) with squared Euclidean cost C using (29). Fig. 3 shows that the Sinkhorn iterations converge linearly (Sec. III-E). We emphasize that the path structure of the information works to minimize the computational cost of the algorithm – we solve the MSBP (24) with $n^s = 500^{26}$ decision variables in approx. 10 s in MATLAB on an Ubuntu 22.04.2 LTS Linux machine with an AMD Ryzen 7 5800X CPU.

Fig. 4 compares predicted versus observed empirical distributions. Specifically, Fig. 4 shows $s_{\text{int}} + 1 = 5$ distributional predictions $\hat{\mu}_{\hat{\tau}_j}$ at times $\hat{\tau}_j$, temporally equispaced throughout the duration of the 3rd control cycle, i.e., between $\tau_{2(s_{\text{int}}+1)+1}$ and $\tau_{3(s_{\text{int}}+1)+1}$, with

$$\hat{\tau}_j = \tau_{2(s_{\text{int}}+1)+1} + \left(\frac{\tau_{3(s_{\text{int}}+1)+1} - \tau_{2(s_{\text{int}}+1)+1}}{s_{\text{int}} + 2} \right) j,$$

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

TABLE I: Number of intracycle marginals s_{int} vs. Wasserstein distances W_j as in (33). All entries are scaled up by 10^4 .

where $j \in \llbracket s_{\text{int}} + 1 \rrbracket$. We used (31) with $\sigma = 2(s_{\text{int}} + 1) + j$, since $\hat{\tau}_j \in [\tau_{2(s_{\text{int}} + 1) + j}, \tau_{2(s_{\text{int}} + 1) + j + 1}]$.

From Fig. 4 it is clear that the measure-valued predictions, while largely accurate, are prone to error in cases where the software resource usage behavior changes in bursts too short to be appear in our observations. Naturally, increasing the number of snapshots should yield an improvement in overall accuracy. We demonstrate this by increasing s_{int} . Table I reports the Wasserstein distances $W(\cdot, \cdot)$ between the corresponding predicted and measured distributions:

$$W_j := W(\hat{\mu}_{\hat{\tau}_j}, \mu_{\hat{\tau}_j}) \quad \forall j \in \llbracket s_{\text{int}} + 1 \rrbracket. \quad (33)$$

We computed each of these W_j as the square root of the optimal value of the corresponding Kantorovich linear program [32, Ch. 3.1] that results from specializing (24) with $s = 2$, $\varepsilon = 0$.

V. CONCLUDING REMARKS

We apply recent algorithmic advances in solving the MSBP to learn stochastic hardware resource usage by control software. The learnt model demonstrates accurate nonparametric measure-valued predictions for the joint hardware resource state at a desired time conditioned on CPS context. The formulation and its solution comes with a maximum likelihood guarantee in the space of probability measures, and the algorithm enjoys a guaranteed linear convergence rate.

REFERENCES

- [1] G. Bernat, A. Colin, and S. Petters, “WCET analysis of probabilistic hard real-time systems,” in *23rd IEEE Real-Time Systems Symposium, 2002 (RTSS)*, 2002, pp. 279–288.
- [2] M. Lv, N. Guan, Y. Zhang, Q. Deng, G. Yu, and J. Zhang, “A survey of WCET analysis of real-time operating systems,” in *2009 International Conference on Embedded Software and Systems*, 2009, pp. 65–72.
- [3] R. Gifford, N. Gandhi, L. T. X. Phan, and A. Haeberlen, “Dna: Dynamic resource allocation for soft real-time multicore systems,” in *2021 IEEE 27th Real-Time and Embedded Technology and Applications Symposium (RTAS)*. IEEE, 2021, pp. 196–209.
- [4] K. Zhang, J. Sprinkle, and R. G. Sanfelice, “Computationally aware switching criteria for hybrid model predictive control of cyber-physical systems,” *IEEE Transactions on Automation Science and Engineering*, vol. 13, no. 2, pp. 479–490, 2016.
- [5] P. J. Bushell, “Hilbert’s metric and positive contraction mappings in a Banach space,” *Archive for Rational Mechanics and Analysis*, vol. 52, pp. 330–338, 1973.
- [6] F. Elvander, I. Haasler, A. Jakobsson, and J. Karlsson, “Multi-marginal optimal transport using partial information with applications in robust localization and sensor fusion,” *Signal Processing*, vol. 171, p. 107474, 2020.
- [7] I. Haasler, A. Ringh, Y. Chen, and J. Karlsson, “Multimarginal optimal transport with a tree-structured cost and the Schrödinger bridge problem,” *SIAM Journal on Control and Optimization*, vol. 59, no. 4, pp. 2428–2453, 2021.
- [8] C. Léonard, “A survey of the Schrödinger problem and some of its connections with optimal transport,” *Discrete and Continuous Dynamical Systems-Series A*, vol. 34, no. 4, pp. 1533–1574, 2014.
- [9] Y. Chen, T. T. Georgiou, and M. Pavon, “Stochastic control liaisons: Richard Sinkhorn meets Gaspard Monge on a Schrödinger bridge,” *Siam Review*, vol. 63, no. 2, pp. 249–313, 2021.
- [10] A. Dembo and O. Zeitouni, *Large deviations techniques and applications*. Springer Science & Business Media, 2009, vol. 38.
- [11] H. Follmer, “Random fields and diffusion processes,” *Ecole d’Ete de Probabilités de Saint-Flour XV-XVII, 1985-87*, 1988.
- [12] I. N. Sanov, *On the probability of large deviations of random variables*. United States Air Force, Office of Scientific Research, 1958.
- [13] M. Pavon, G. Trigila, and E. G. Tabak, “The data-driven Schrödinger bridge,” *Communications on Pure and Applied Mathematics*, vol. 74, no. 7, pp. 1545–1573, 2021.
- [14] I. Csiszár, “I-divergence geometry of probability distributions and minimization problems,” *The annals of probability*, pp. 146–158, 1975.
- [15] J. M. Borwein, A. S. Lewis, and R. D. Nussbaum, “Entropy minimization, DAD problems, and doubly stochastic kernels,” *Journal of Functional Analysis*, vol. 123, no. 2, pp. 264–307, 1994.
- [16] M. Cuturi, “Sinkhorn distances: Lightspeed computation of optimal transport,” *Advances in neural information processing systems*, vol. 26, 2013.
- [17] J. Franklin and J. Lorenz, “On the scaling of multidimensional matrices,” *Linear Algebra and its applications*, vol. 114, pp. 717–735, 1989.
- [18] J.-D. Benamou, G. Carlier, M. Cuturi, L. Nenna, and G. Peyré, “Iterative Bregman projections for regularized transportation problems,” *SIAM Journal on Scientific Computing*, vol. 37, no. 2, pp. A1111–A1138, 2015.
- [19] H. H. Bauschke and A. S. Lewis, “Dykstras algorithm with Bregman projections: A convergence proof,” *Optimization*, vol. 48, no. 4, pp. 409–427, 2000.
- [20] S. D. Marino and A. Gerolin, “An optimal transport approach for the Schrödinger bridge problem and convergence of Sinkhorn algorithm,” *Journal of Scientific Computing*, vol. 85, no. 2, p. 27, 2020.
- [21] G. Carlier, “On the linear convergence of the multimarginal Sinkhorn algorithm,” *SIAM Journal on Optimization*, vol. 32, no. 2, pp. 786–794, 2022.
- [22] G. A. Bondar, R. Gifford, L. T. X. Phan, and A. Halder, “Path structured multimarginal Schrödinger bridge for probabilistic learning of hardware resource usage by control software,” *arXiv preprint arXiv:2310.00604*, 2023.
- [23] F. A. Ba and M. Quellmalz, “Accelerating the sinkhorn algorithm for sparse multi-marginal optimal transport via fast fourier transforms,” *Algorithms*, vol. 15, no. 9, p. 311, 2022.
- [24] J. Kong, M. Pfeiffer, G. Schildbach, and F. Borrelli, “Kinematic and dynamic vehicle models for autonomous driving control design,” in *2015 IEEE intelligent vehicles symposium (IV)*. IEEE, 2015, pp. 1094–1099.
- [25] S. Haddad, A. Halder, and B. Singh, “Density-based stochastic reachability computation for occupancy prediction in automated driving,” *IEEE Transactions on Control Systems Technology*, vol. 30, no. 6, pp. 2406–2419, 2022.
- [26] A. Wächter and L. T. Biegler, “On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming,” *Mathematical Programming*, vol. 106, no. 1, pp. 25–57, 2006.
- [27] [https://github.com/abhishekhalder/CPS-Frontier-Task3-Collaboration/ blob/master/Codes/kbm_sim/Documentation_KinematicBicycle_ Controllers.pdf](https://github.com/abhishekhalder/CPS-Frontier-Task3-Collaboration/blob/master/Codes/kbm_sim/Documentation_KinematicBicycle_Controllers.pdf), accessed: 2023-09-29.
- [28] Intel Corporation, “Improving real-time performance by utilizing cache allocation technology,” Apr. 2015, White Paper.
- [29] H. Yun, G. Yao, R. Pellizzoni, M. Caccamo, and L. Sha, “Memory bandwidth management for efficient performance isolation in multicore platforms,” *IEEE Transactions on Computers*, vol. 65, no. 2, pp. 562–576, Feb 2016.
- [30] “perf(1) — linux manual page,” <https://man7.org/linux/man-pages/man1/perf.1.html>, accessed: 2023-09-29.
- [31] T. Eiter and H. Mannila, “Computing discrete Fréchet distance,” 1994. [Online]. Available: <https://api.semanticscholar.org/CorpusID:16010565>
- [32] G. Peyré and M. Cuturi, “Computational optimal transport: With applications to data science,” *Foundations and Trends® in Machine Learning*, vol. 11, no. 5-6, pp. 355–607, 2019.