

Hidden Markov Models for Non-Well-Mixed Reaction Networks

Nils Napp, David Thorsley, Eric Klavins

Department of Electrical Engineering

University of Washington

Seattle WA 98195

{nnapp, thorsley, klavins}@u.washington.edu

Abstract—The behavior of systems of stochastically interacting particles, be they molecules comprising a chemical reaction network or multi-robot systems in a stochastic environment, can be described using the Chemical Master Equation (CME). In this paper we extend the applicability of the CME to the case when the underlying system of particles is not well-mixed, by constructing an *extended state space*. The proposed approach fits into the general framework of approximating stochastic processes by Hidden Markov Models (HMMs). We consider HMMs where the hidden states are equivalence classes of states of some underlying process. The sets of equivalence classes we consider are refinements of macrostates used in the CME. We construct a series of HMMs that use the CME to describe their hidden states. We demonstrate the approach by building a series of increasingly accurate models for a system of robots that interact in a non-well-mixed manner.

I. INTRODUCTION

Modeling stochastically interacting particles is key to many problems in science and engineering. Chemical reactions [1], gene regulatory networks [2], and stochastic multi-robot systems [3] are examples of systems described in this fashion. Models of these systems typically consist of a set of heterogeneous particles stochastically moving in a reaction volume. The positions, velocities, and internal states of all particles together make up the *microstate* of the system. The stochastic behavior of the microstate captures all system properties, but microstate models are generally difficult to analyze due to the enormous size of their state spaces. However, many system properties of interest do not require detailed knowledge of the microstates. For example, the steady state distribution of particle types or the fluctuation characteristics in the number of particle types do not require detailed knowledge of the microstate positions and velocities. As a result, we attempt to build models with a coarser system description. This is the case in the chemical master equation (CME) [4] where the state describes the copy number of each particle type, but not each particle's position and velocity. In this paper, we call these coarser states *macrostates*. Under the assumption that the particles are well-mixed, that is, they are uniformly distributed on the reaction volume and diffuse quickly, the stochastic process describing the development of the macrostates is a Markov process whose dynamics are described by the CME [5].

The CME is amenable to many mathematical tools because the probability mass function of the system states is governed by linear ordinary differential equations. Further, there exist

efficient exact and approximate algorithms for generating sample paths [6] and systems engineering tools for creating and analyzing reduced order models have been applied in [7]. These modeling and analysis tools have been successfully applied to many of the aforementioned systems of interacting particles.

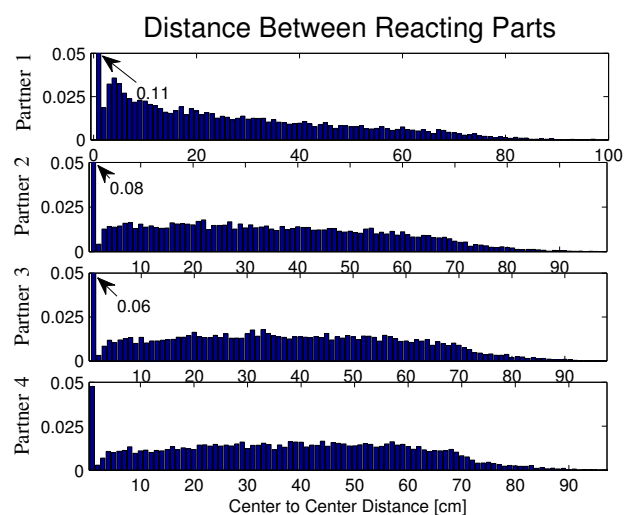


Fig. 1. Histogram of distances between interaction partners. The histograms show the distribution of distances to the next four interaction partners from a distinguished part on the *Programmable Parts Testbed* [8]. The arrows indicate the height of the peak near zero when it is out of range. The data is from simulation of the system described in Section V-D with 24 parts at a density of $\rho = 10 \frac{\text{parts}}{\text{m}^2}$.

However, many interesting real systems are not well-mixed, so the theoretical underpinning of the CME does not apply. One approach for dealing with such systems is to model the whole microstate process. When the stochastic motion of the particles is described by Brownian motion, this is the Smoluchowski model and tools for efficiently simulating sample paths in this model exist [9]. However, this approach introduces a continuum of states for representing each particle's position, which makes tools for the discrete state CME inapplicable. Instead, we propose to augment the *macrostates* of the system using the *extended state space* approach, in which each particle keeps track of the types of previous interaction partners. In this paper we formalize ideas introduced in [10]. The extended extended state space formulation is a compromise between explicitly accounting

for the position, velocity, and identity of each particle, and only using the copy number of particle types as the system state. This formulation approximates the dynamics of the macrostate as a Hidden Markov Model with a CME description for the hidden states.

In this paper we propose a new approach for producing tractable models of non-well-mixed processes. An example system of stochastically interacting robots with this property is described in [8]. The *Programmable Parts Testbed* consists of randomly stirred robots floating on an air table. Specifically, this paper focuses on the case when the well-mixed assumption underlying the CME does not hold because repeat interactions between particles are likely. The robots interact stochastically, but the parts are not well-mixed as Fig. 1 indicates. Each histogram shows the distribution of distances to future interaction partners at the time of an interaction. The peak close to 0 corresponds to the probability of interacting again with the same part. The last histogram represent a uniform distribution of reaction partners over the reaction domain. The peak is roughly $\frac{1}{23}$ high, which is the probability of picking the same partner from the $24-1=23$ options. The differing shape between the first two histograms in Fig. 1 highlights the increased likelihood of reacting with close neighbors. In Section V we construct a model that captures this behavior.

The paper is organized as follows. Section II introduces necessary notation and mathematical background. Section III discusses properties of the Kullback-Leibler divergence rate, which we used to quantify the difference between processes. Section IV introduces a method for approximating arbitrary, stationary, discrete state, discrete time stochastic processes by Markov chains and details the construction of Hidden Markov Models from equivalence classes of the microstate space. In Section V we define a particular set of equivalence classes that are designed to capture repeat reactions and apply this approach to the non-well-mixed Programmable Parts Testbed.

II. MATHEMATICAL FORMALISM

A. Markov Processes

A discrete time *stochastic process* \mathcal{X} is a collection of random variables, $\{X_n\}_{n \in \mathbb{N}}$, parameterized by the natural numbers \mathbb{N} . A stochastic process generates *trajectories*, i.e. functions from the index set \mathbb{N} to the state space X . A particular trajectory $\omega \in \mathbb{N} \rightarrow X$ is an assignment of each random variable X_n to a particular value $x \in X$. We view a stochastic process as a *probability space* on trajectories, i.e. $(\mathbb{N} \rightarrow X, \mathcal{F}, P)$, where \mathcal{F} is generated by cylinder sets [11, Sec. 2].

A discrete state process with the property that

$$P\{X_{n+1} = x_{n+1} | X_n = x_n, X_{n-1} = x_{n-1}, \dots, X_0 = x_0\} = P\{X_{n+1} = x_{n+1} | X_n = x_n\}, \quad (1)$$

where $x_i \in X$ is called a *Markov chain*. Probability measures with property (1) are called *Markov probability measures* and




Fig. 2. Representation of an HMM. The state space of the Markov chain is $X = \{0, 1, 2, 3\}$ and represented circles. The numbers on the edges represent transition probabilities. By adding an output function $f : X \rightarrow Y$ this Markov chain can be turned into an HMM, the boxes represent states that map to the same output.

are memoryless. The *one-step transition probabilities*,

$$A_{ij}(n) = P\{X_{n+1} = j | X_n = i\} \text{ for } i, j \in X,$$

specify the transition dynamics of a Markov chain. In a stationary process matrix A is independent of n . The initial distribution of a stochastic process is described by the vector $\alpha_i \equiv P\{X_0 = i\}$, for $i \in X$. The pair (α, A) , completely describes the probability measure over trajectories of a Markov chain [11, Thm. 8.1].

A *Hidden Markov Model* (HMM) is a stochastic process with state space Y whose probability measure on trajectories can be described the by a Markov chain \mathcal{X} and an *output function* $f : X \rightarrow Y$, where Y is called the *output space*. Since the function f can be many-to-one, the resulting process \mathcal{Y} can be non-Markov although the process \mathcal{X} is. HMMs are a rich class of models and in general can have random output functions, however, the simpler deterministic output description suffices for this paper. Like the transition probabilities, which can be represented as a matrix, f can be written as a $|Y| \times |X|$ matrix B ,

$$B_{ij} = \begin{cases} 1 & \text{if } f(j) = i \\ 0 & \text{otherwise} \end{cases}.$$

Example 1. Fig. 2 is a schematic representation of an HMM, with state space $X = \{0, 1, 2, 3\}$ and output space $Y = \{a, b\}$. The nodes represent states of the Markov chain, the arrows represent transitions, and the two grey boxes represent the two output states. The function f is given by

$$f = \begin{cases} 0 \mapsto a \\ 1 \mapsto a \\ 2 \mapsto b \\ 3 \mapsto b \end{cases}.$$

When the Markov chain is stationary and A is irreducible, the initial distribution α is determined by the relation $A^T \alpha = \alpha$. The parameters for the HMM in Fig. 2 are

$$A = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} & 0 & 0 \\ 0 & 0 & 0 & 1 \\ \frac{1}{4} & 0 & 0 & \frac{3}{4} \\ \frac{1}{3} & 0 & \frac{2}{3} & 0 \end{pmatrix},$$

$$B = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{pmatrix},$$

and

$$\alpha = \left(\frac{6}{19} \frac{3}{19} \frac{4}{19} \frac{6}{19} \right)^T.$$

The Markov chain \mathcal{X} is specified by A and α and together with output function B define an HMM.

B. Equivalence Relations

The remainder of this section is about notation related to equivalence classes, partitions, and quotient spaces. An *equivalence relation* \sim on a given space X , divides X into a set of disjoint subsets called *equivalence classes*. The set of equivalence classes is called the *quotient space* and denoted by $X_{/\sim}$. The map $\pi : X \rightarrow X_{/\sim}$ that maps each element of X to its equivalence class is called the *quotient map*. Every equivalence relation on X produces equivalence classes that form a *partition* of X , a set of disjoint subsets of X whose union is X . Similarly, every partition induces an equivalence relation.

By defining an equivalence relation \sim on the state space of a stochastic process \mathcal{X} , we can form a new stochastic process $\mathcal{X}_{/\sim}$ by composing X_n with π :

$$\mathcal{X}_{/\sim} = \{(X_{/\sim})_n\}_{n \in \mathbb{N}} \quad \text{where} \quad (X_{/\sim})_n = \pi \circ X_n.$$

In particular, if \mathcal{X} is Markov then $\mathcal{X}_{/\sim}$ is an HMM with output space $X_{/\sim}$. If the original process \mathcal{X} is stationary, then so is $\mathcal{X}_{/\sim}$.

Example 2. We can create the same process as in Example 1 by defining an equivalence relation on $X = \{0, 1, 2, 3\}$ such that

$$0 \sim 1 \quad \text{and} \quad 2 \sim 3.$$

This induces the two equivalence classes $\{0, 1\}$ and $\{2, 3\}$. If one identifies $a \in Y$ from Example 1 with $\{0, 1\}$ and $b \in Y$ with $\{2, 3\}$ then the HMM given in Example 1 can be expressed as $\mathcal{X}_{/\sim}$, the output function f is represented by equivalence classes. States with the same output are in the same equivalence class.

The notion of *refinement* gives a *partial order* on the space of all partitions. Given two partitions H and H' of a space X , H' is a refinement of H , written as $H' \sqsubseteq H$, iff

$$\forall h' \in H' \quad \exists h \in H \quad \text{s.t.} \quad h' \subset h.$$

The *top* element \top of this partial order is induced by the equivalence relation in which all elements are equivalent. The *bottom* element \perp is induced by the equivalence relation where every element is equivalent only to itself.

Given a particular partition H of X we construct series of refinements to build increasingly better HMM approximations of \mathcal{X} . For the remainder of this paper assume that we have a process \mathcal{X} with state space X , a partition H of X , and a series of successive refinements H_n such that

$$\perp \sqsubseteq \dots \sqsubseteq H_2 \sqsubseteq H_1 \sqsubseteq H_0 = H \sqsubseteq \top. \quad (2)$$

The partition H is given by the aspect of the system we want to model. It should be chosen so that all elements in a $h \in H$

are equivalent with respect to the quantity of interested. For example, in Section V the partition H is generated by putting all states with the same copy number of each species in the same set $h \in H$. Any coarser partition would not capture this system feature we are trying to model. Associated with each partition H_n is an equivalence relation \sim_n . Given two partitions in this hierarchy H_n and H_m with $n > m$, the *coarser* H_m naturally defines an equivalence relation on the *finer* H_n ,

$$h_n \sim_m h'_n \Leftrightarrow \exists h_m \in H_m \quad \text{s.t.} \quad h_n \subseteq h_m \wedge h'_n \subseteq h_m.$$

Since this equivalence relation on H_n is induced by the same equivalence relation that induces H_m on X , we will use the same symbol to denote both. Further, we identify each element in H_{n/\sim_m} with the element in H_m that creates it. With this identification we can view the quotient map $\pi : H_n \rightarrow H_{n/\sim_m}$ as a mapping to H_m instead of H_{n/\sim_m} . For precision, sub- and super-scripts will be added to the quotient map π to indicate the involved spaces. For example, with $n > m$ $\pi_m^n : H_n \rightarrow H_m$ is given by

$$\pi_m^n(h_n) = h_m \quad \text{where} \quad h_n \subseteq h_m.$$

III. THE KULLBACK-LEIBLER DIVERGENCE RATE

To compare probability measures, P and Q , for a random process \mathcal{X} we use the Kullback-Leibler (KL) divergence rate or relative entropy rate [12], defined as,

$$D_X(P||Q) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{\omega \in X^N} P(\omega) \log \left(\frac{P(\omega)}{Q(\omega)} \right). \quad (3)$$

When this limit exists it is called the divergence rate between the processes defined under the probability measures P and Q . For stationary, ergodic processes this limit always exists, save the absolute continuity requirement, $P \ll Q$ [13, THM 4.2]. This requirement is satisfied by the approximations introduced in the next section.

The KL-divergence rate measures how different the probability measure Q is from a true measure P . Since the problem considered in this paper is the approximation of a given process, we let the true measure on trajectories be the one generated by the original process. The other distribution, Q in (3), is the measure on trajectories generated by the various approximations. Since the remaining sections compare processes on different output spaces, derived from probability measures on trajectories in the hidden states, the output space will be added to D as a subscript to avoid confusion. For example, the expression $D_X(P||Q)$ denotes the divergence rate where $\omega \in \mathbb{N} \rightarrow X$, under two different probability measures P and Q , while $D_{X_{/\sim}}(P||Q)$ denotes the divergence when $\omega \in \mathbb{N} \rightarrow X_{/\sim}$.

Since the KL-divergence rate is an asymptotic property of two stochastic processes it is suitable to measure differences in their steady state behavior. The theorems in the next section all require stationarity, which can be interpreted as a system starting in its steady state. For transient characteristics this approach does not capture the differences in probability measures on trajectories. One possible alternative approach

is to use the KL-divergence on finite length trajectories. The immediate difficulty then is computing with probability measures on the potentially very large discrete space of possible trajectories. We use this approach to estimate the KL-divergence rate for a small example in Section V, but do not explore the theory of creating approximate models that capture transient behavior. Instead, we leave this interesting question as future work.

IV. APPROXIMATION

In this section we construct Markov processes that use equivalence classes of $X_{/\sim}$ as their states. Each such process is an approximation of the process $\mathcal{X}_{/\sim}$ and does not take into account the larger state space X . We think of the equivalence classes as macrostates and the state space X as the microstates. Both the approximation and the original process define probability measures on $\mathbb{N} \rightarrow X_{/\sim}$ and thus the KL-divergence rate can be used to compare $\mathcal{X}_{/\sim}$ and its approximation.

A. Markov Approximations

In this section we construct a Markov chain approximations of an arbitrary stationary process \mathcal{X} . We give the optimal choice for the transition probabilities, in the sense that this choice minimizes the KL-divergence rate between the original process and its Markov approximation.

Definition 1. *Given an arbitrary stationary stochastic process \mathcal{X} , a Markov approximation $\hat{\mathcal{X}}$ is a Markov chain that has the same state space as \mathcal{X} .*

Any choice of (α, A) will define a Markov process with state space X . For $\hat{\mathcal{X}}$ we pick the initial distribution α as

$$\hat{\alpha}_i = \hat{P}\{\hat{X}_0 = i\} \equiv P\{X_0 = i\}. \quad (4)$$

In addition, we choose matching one step transition probabilities. For $i, j \in X$,

$$\hat{A}_{ij} = \hat{P}\{\hat{X}_1 = j \mid \hat{X}_0 = i\} \equiv P\{X_1 = j \mid X_0 = i\}. \quad (5)$$

Assignments (4)-(5) define a probability measure \hat{P} on trajectories of $\hat{\mathcal{X}}$ as described in Section II-A. When the original process \mathcal{X} is Markov, then P and \hat{P} will be the same for all trajectories. However, when the original process is not Markov, then the Markov approximation $\hat{\mathcal{X}}$ will match only the one step transitions of \mathcal{X} , but will define a different probability measure on trajectories.

Example 3. Let \mathcal{Y} be the output process $\mathcal{X}_{/\sim}$. The Markov approximation $\hat{\mathcal{Y}}$ is given by the initial conditions

$$\hat{\alpha} = \begin{pmatrix} \alpha_0 + \alpha_1 \\ \alpha_2 + \alpha_3 \end{pmatrix} = \begin{pmatrix} \frac{9}{19} \\ \frac{10}{19} \end{pmatrix}$$

and one step transition probabilities

$$\hat{A} = \begin{pmatrix} \frac{2}{3} & \frac{1}{3} \\ \frac{3}{10} & \frac{7}{10} \end{pmatrix}.$$

The entries of \hat{A} can be computed from the underlying process \mathcal{X} . For example, the (1,1) entry of A is given by

$$\begin{aligned} P\{X_1 \in \{0, 1\} \mid X_0 \in \{0, 1\}\} &= \\ \sum_{i,j \in \{0,1\}} P\{X_1 = j \mid X_0 = i\} P\{X_0 = i \mid X_0 \in \{0, 1\}\} &= \\ &= \frac{1}{2} \frac{6}{9} + \frac{1}{2} \frac{6}{9} + 0 \frac{3}{9} + 0 \frac{3}{9} = \frac{2}{3} \end{aligned}$$

Theorem 1. [13, Remark 4.7] *Choosing the parameters $\hat{\alpha}$ and \hat{A} for the Markov approximation as in (4)-(5) minimizes the KL-divergence rate, $D_X(P\|\hat{P})$, between a stationary \mathcal{X} and the Markov process $\hat{\mathcal{X}}$.*

The remark in [13] sets up a minimization problem for the KL-divergence rate. The resulting optimization problem can be explicitly solved via Lagrange multipliers, see Appendix A.

Given a partition H_n of the state space X in the hierarchy (2) we use the following notation. The Markov approximation $\hat{\mathcal{X}}_{/\sim_n}$ of $\mathcal{X}_{/\sim_n}$ is a Markov chain with state space H_n . Once again, we specify $\hat{\mathcal{X}}_{/\sim_n}$ by an initial distribution and transition probabilities according to (4)-(5). Denote the resulting probability measure on trajectories $\mathbb{N} \rightarrow H_n$ by \hat{P}_n .

B. HMM Approximations

Recall that for process \mathcal{X} with state space X , we have a partition H that is determined by some quantity we are interested in modeling, and a fixed series of successive refinements H_n (2). Our goal is to build a hierarchy of HMMs that can serve as approximations for a macrostate process $\mathcal{X}_{/\sim_0}$. For each partition, H_n , in the refinement (2), we construct a Markov chain approximation $\hat{\mathcal{X}}_{/\sim_n}$. If $m < n$, we can construct an HMM $(\hat{\mathcal{X}}_{/\sim_n})_{/\sim_m}$ by using π_m^n as the output map.

For example, $(\hat{\mathcal{X}}_{/\sim_3})_{/\sim_1}$ is the HMM that has the Markov approximation of $\mathcal{X}_{/\sim_3}$ as its hidden states and $\pi_1^3 : H_3 \rightarrow H_1$ as its output function. The KL-divergence rate between this process and $\mathcal{X}_{/\sim_1}$ is denoted by

$$D_{H_1}(P\|\hat{P}_3). \quad (6)$$

Theorem 2 and Conjecture 3 relate the divergence rates of the different HMMs that can be formed within the hierarchy (2). First, we relate processes with the same probability measure on its hidden state states, but with different output spaces.

Theorem 2. *Given a stationary stochastic process \mathcal{X} with state space X , a probability measure P and two partitions $H_m \sqsubseteq H_n$ of X in the hierarchy (2) the following relation holds,*

$$D_{H_m}(P\|\hat{P}_k) \geq D_{H_n}(P\|\hat{P}_k)$$

with $n \leq m \leq k$.

Proof: We prove the statement for two arbitrary probability measures P and Q . The proof uses Jensen's inequality to show that this relation holds for each term in the limit for (3). Fixing N , a term from D_{H_m} has the form

$$\sum_{\omega \in H_n^N} P(\omega) \log \left(\frac{P(\omega)}{Q(\omega)} \right) = \sum_{\rho \in H_n^N} \sum_{\pi(\omega)=\rho} P(\omega) \log \left(\frac{P(\omega)}{Q(\omega)} \right) \quad (8)$$

For a fixed $\rho \in H_n^N$, looking only at the inner sum gives

$$- \sum_{\pi(\omega)=\rho} P(\omega) \log \left(\frac{P(\omega)}{Q(\omega)} \right) = \sum_{\pi(\omega)=\rho} P(\omega) \log \left(\frac{Q(\omega)}{P(\omega)} \right). \quad (9)$$

Because log is a concave function, applying Jensen's inequality to (9) gives

$$\frac{1}{\sum_{\omega} P(\omega)} \sum_{\omega} P(\omega) \log \left(\frac{Q(\omega)}{P(\omega)} \right) \leq \log \left(\frac{\sum_{\omega} Q(\omega)}{\sum_{\omega} P(\omega)} \right), \quad (10)$$

where \sum_{ω} stands for $\sum_{\pi(\omega)=\rho}$. Using expression (10) in (9) yields

$$\sum_{\omega} P(\omega) \log \left(\frac{P(\omega)}{Q(\omega)} \right) \geq \sum_{\omega} P(\omega) \log \left(\frac{\sum_{\omega} P(\omega)}{\sum_{\omega} Q(\omega)} \right), \quad (11)$$

where the right hand side is equal to $P(\rho) \log \left(\frac{P(\rho)}{Q(\rho)} \right)$. Finally, establish that

$$\begin{aligned} & \frac{1}{N} \sum_{\omega \in H_m^N} P(\omega) \log \left(\frac{P(\omega)}{Q(\omega)} \right) \\ &= \frac{1}{N} \sum_{\rho \in H_n^N} \sum_{\pi(\omega)=\rho} P(\omega) \log \left(\frac{P(\omega)}{Q(\omega)} \right) \\ &\geq \frac{1}{N} \sum_{\rho \in H_n^N} P(\rho) \log \left(\frac{P(\rho)}{Q(\rho)} \right). \end{aligned}$$

It follows that $D_{H_m}(P||Q) \geq D_{H_n}(P||Q)$ since the relation is true for each term in the limit (3). By inspecting relation (11), we note that equality holds when probabilities have a constant ratio on each equivalence class in H_m , $\frac{\sum_{\omega} P(\omega)}{\sum_{\omega} Q(\omega)} = \frac{P(\omega)}{Q(\omega)}$. \square

The following conjecture relates processes with the same output space, but different hidden states.

Conjecture 3. *Given a stationary stochastic process \mathcal{X} , a hierarchy of partitions (2), and the associated Markov approximations we have*

$$D_{H_n}(P||\hat{P}_m) \leq D_{H_n}(P||\hat{P}_k) \quad (12)$$

with $0 \leq n \leq k \leq m$.

Both the left and right hand side of (12) correspond to the divergence of the original macrostate process $\mathcal{X}_{j \sim n}$ and an HMM approximation. The difference between the left and right hand expressions is that the HMM on the left has as its hidden states a finer Markov approximation than the right hand side.

The reason we believe this conjecture to be true is that the trajectories of hidden states on the left hand side of (12) contain more information about the original process than the hidden states on the right hand side. By Theorem 1, the Markov approximation of the hidden states is optimal for

the trajectories in the hidden states, so it should also lead to a good HMM approximation as well.

Theorem 2 suggests that in order to make a more accurate model of a process one could simply decrease the size of the output space by going up the hierarchy of refinements to a coarser partition. In fact, if the output space is \top , then the divergence rate is zero no matter what the underlying state space is. However, as indicated in (2), H_0 corresponds to some given partition H that is determined by the process we want to model. H specifies the least amount of detail the output space must have. For example, in the next section the coarsest equivalence relation is the one resulting in the macrostates used by the CME.

Fig. 3 illustrates the results of Theorem 2 and Conjecture 3. In order to decrease the KL-divergence rate between the original process and an approximation, one can either increase the size of the hidden state space, or decrease the size of the output space. The approach taken when constructing the *extended state space* for stochastic reaction networks in the next section, is to increase the size of the state space for constructing more accurate models.




Fig. 3. Relation between the various Markov approximations. The arrows indicate the direction of increasing accuracy. The dark shading indicates Markov models, the light shading indicates HMMs.

V. APPLICATION TO REACTION NETWORKS

In this section we introduce a model for non-well-mixed chemical reactions and propose a refinement hierarchy that captures repeat reactions. We then use this hierarchy to define an *extended state space* for chemical reaction networks. Finally, we apply the extended state space approach to the robotic system introduced in Fig. 1.

A. Microstate Model


In the microstate model for chemical reactions particles of different types¹ S move randomly in a *reaction domain*. If two particles encounter one another they may react according to a set of possible reactions, Ψ . The positions, velocities, orientations, and internal states of the particles define the microstates of the system. Denote the set of all possible

¹In the context of chemical reactions the types are chemical species. In our model particles can change their type, so we avoid the name species, since it implies some fundamental description of the particle that does not change spontaneously.

positions, orientations, and velocities for each particle by V . The state space for a microstate process with N particles is

$$X = S^N \times V^N.$$

The particles move in the reaction domain according to the microstate process \mathcal{X} . When two particles are in close proximity, they can react according Ψ and change their type. The details of the interaction depend on the physical system, but reactions generally require proximity. For the detailed reaction mechanism between the robots from Figure I see [3]. The set of reactions Ψ considered for the remainder of the paper is



Reaction network Ψ is cyclic in the sense that reactions occur even when the system is in equilibrium. This is important since the theoretical development in Section IV is applicable only to stationary systems.

B. Macrostate Model

In this section we describe the macrostates used in the CME in terms of an equivalence relation on the microstates. This will form the coarsest partition in a hierarchy of partitions used to define the extended macrostate in the next section.

In the CME the macrostates of the system are the copy number of each particle type. When the system is well-mixed the next reaction only depends on the current macrostate, since the positions, velocities, and orientations of all particles at the time of the last reaction are rapidly mixed to their steady state values and do not influence the next macrostate transition [4].

Let $\Pi_{1,N} : X \rightarrow S^N$ be the projection of X onto the components 1 through N , and S_N the group of permutations of the N indices. The equivalence relation defining the CME macrostates is,

$$x \sim_0 x' \Leftrightarrow \exists \sigma \in S_N \text{ s.t. } \Pi_{1,N}(x) = \sigma(\Pi_{1,N}(x')),$$

for any $x, x' \in X$. The expression on the right hand side highlights two essential features of the macrostate. The projection Π indicates that only the particle types matter and the permutation of indices indicates that the macrostate does not keep track of particle identities.

C. Extended Macrostates

In this section we construct a hierarchy of equivalence relations that is tailored to capturing repeat reactions. The motivation for this particular hierarchy of refinements is that even when repeat interactions between individual particles are likely, they can only affect the macrostate if they are between appropriate types. Instead of keeping track of individual pairs that are likely to re-react, we only keep track of the *type* of previous reaction partners. In a given reaction network only repeat reactions between certain types can change the macrostate. Keeping track of the number

of possible repeat reactions that can change the macrostate should capture this behavior.

In order to facilitate the construction of the *extended state space*, we first augment the microstates such that each particle has a history of the types of previous interaction partners. This does not change the stochastic behavior of the system. To preserve the finite nature of the state space, assume that each particle has a history length of L of the types of particles it has encountered on the reaction domain. Given a set of N particles the state space X for the new microstate process \mathcal{X} describing this system of interacting particles is

$$X = S^N \times (S^N)^L \times V^N.$$

The first $(1+N) \times L$ components are extended types, and the last N components are positions of particles in the reaction domain.


Let \sim_0 be the equivalence relation that $x, x' \in X$ are equivalent iff they have the same copy number of each species, which induces macrostates of the CME. Let the \sim_n be defined as follows. Two states are similar $x \sim_n x'$ iff

$$\forall k \in \{0, n\} \quad \Pi_{1+kN, (k+1)N}(x) = \sigma(\Pi_{1+kN, (k+1)N}(x')).$$


The same σ has to work for all the history states in order for two states to be similar. An immediate consequence of this definition is that

$$x \sim_{n+1} x' \Rightarrow x \sim_n x'. \quad (13)$$

Property (13) implies that this set of equivalence relations induces a hierarchy of partitions as described in section IV-B. Now we can construct a new reaction network Ψ^n that correspond to the admissible state changes in H_n . For example, in H_1




where $*$ $\in \{A, B, C\}$ is a wild card that describes all possible species types. The letters indicate the particle type and the subscripts indicates the history of types. For example, A_B denotes a particle of type A whose last interaction was with a particle of type B . Similarly, higher order extensions of Ψ are written by having more history states, for example




is a reaction from Ψ^3 on H_3 . Each extended reaction network Ψ^n can be thought of as an HMM approximation of the macrostate process, by considering π_0^n as the output function.

The reaction network Ψ^1 has the important feature that



is considered a different reaction than, for example,



This allows the two reactions to have different rates of occurring when repeat reactions are likely. This is not possible in the CME model of Ψ .

D. Numerical Example

The *Programmable Parts Testbed* [8] consists of triangular robots. Every robot has a magnetic latch and an IR transceiver on each side. We refer to these robots as the *programmable parts*. They float on an air-hockey table and are randomly mixed by air jets that are mounted along the perimeter of the table. The following experiments were conducted using a high-fidelity mechanics-based simulation of the system [3]. In all experiments the average kinetic energy per robot, is kept at $5 \times 10^{-4} J$ by adjusting the mixing strength. This is analogous to performing experiments at a constant temperature.

We implemented the reaction network Ψ with the programmable parts, where the type is represented by an internal state of the parts. The state can only change upon temporarily latching with another part of appropriate type. We then ran simulations until the system reached steady state and collected trajectories and estimated rates as in [3]. Additionally, we took into account interactions that did not change the macrostate to estimate the probability of self loops. We used this data to extract the embedded Markov chain of the system.

The trajectories started from a small number of parts (initial macrostate = $(2, 1, 1)'$). As a result there are only a few reachable macrostates. This restriction allowed us to gather enough data to estimate the measure P for fixed length trajectories and compute the KL-divergence. Fig. 1 suggest that after four interactions the system is roughly well-mixed. We considered trajectories of length 8, making the problem computationally feasible yet sufficiently long to mix away spatial aspects.

Specifically, there are only 3 macrostates resulting in $3^8 = 6561$ different trajectories of length 8. However, not all of these trajectories are feasible due the restrictions imposed by Ψ . We used simulations to estimate P on this space and compared it to the probability measures induced by different approximations. The results of this comparison are shown in Fig. 4.

This example shows the extended state space approach producing models with less KL-divergence from the original process than a Markov approximation. Also, the data suggest that the mixing in the system is such that keeping the type of more than two previous reaction partners in the state provides little improvement. This observation is consistent with the data displayed in Fig. 1 and justifies only using short trajectories for estimating the KL-divergence.

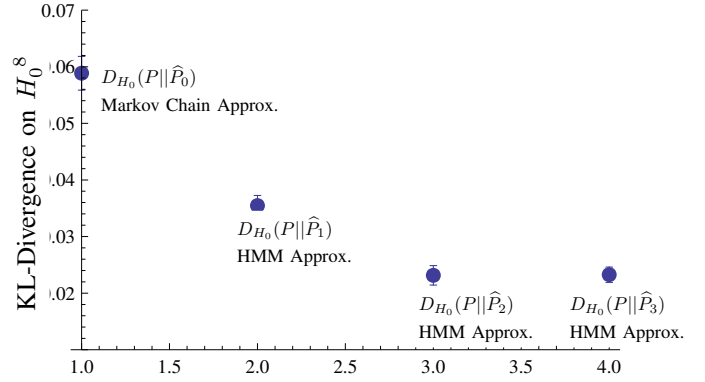


Fig. 4. This figure shows an estimate of $D_{H_0}(P||\hat{P}_0)$, $D_{H_0}(P||\hat{P}_1)$, $D_{H_0}(P||\hat{P}_2)$, and $D_{H_0}(P||\hat{P}_3)$ of the system described in Section V-D. Each estimate has a label indicating the type of approximation the KL-divergence is computed with respect to. The estimate is formed by looking at finite trajectories $\omega \in \{1, \dots, 8\} \rightarrow H$. The error bars result from different sets of Monte Carlo simulations for estimating the probability measures of the approximations.

VI. DISCUSSION

The contributions of this paper are (1) to introduce general approach of state space refinements to construct a series of HMMs, (2) to use HMM tools to analyze reaction networks, and (3) to construct the extended state space, which can capture repeat reactions in reaction networks. The extended state space is able to capture spatial aspects of the particle interactions without explicitly modeling space itself.

The theory developed in Section IV only requires a series of refinements. The hierarchy chosen in Section V is physically motivated, but by no means the only possible choice. One avenue for future research is to compare different refinements schemes, we would also like to address the problem of automatically generating optimal refinement schemes from data. This is different from the question of finding optimal HMM approximations as in [13], since we require the states of the HMM to be equivalence classes of some microstate model. Such analysis could not only yield reduced order models, but also give physical insight into the dynamics of the system.

The utility in expressing the hidden states of an HMM as a reaction network, is that one can apply linear ODE and systems engineering tools to the hidden states of the HMM, while the output function allows the model to capture some of the non-well-mixed aspects of a system.

The optimality of the Markov approximation depends on the measure used to assess the distance between stochastic processes. The KL-divergence rate used here captures asymptotic differences corresponding to the steady state. Using a similar approach with a different distance measure could allow the development of approximate models that are geared toward reproducing the transient behavior of an arbitrary stochastic processes.

We intend to formally prove the claims of Conjecture 3.

We would also like to extend the work to include the continuous time case since this would allow reasoning about rates in reaction networks. We would like to find out how the *basic rates* used in chemical reaction networks are connected to the analysis presented here, which focuses on the Markovian aspect only.

APPENDIX

A. Proof of Theorem 1

Here we show that the transition probabilities and states as defined in Section IV-A define the Markov process minimizing the KL-divergence rate between the original process and its approximation. The approach is to minimize the expression for KL-divergence rate over all possible Markov transition probabilities via Lagrange multipliers. It will be useful to only look at short sections of a trajectories. Let X_n^m denote the set of trajectories that are only specified from time $\{n, n+1, \dots, m-1, m\} \subset T$. Equation (3) can be rewritten as

$$D_X(P\|\hat{P}) = \lim_{n \rightarrow \infty} \frac{1}{n} E_P \log \left(\frac{P(X_0^{n-1})}{\hat{P}(X_0^{n-1})} \right),$$

where E_P is the expectation with respect to the probability measure P .

$$D_X(P\|\hat{P}) = \lim_{n \rightarrow \infty} \frac{1}{n} E_P \log (P(X_0^{n-1})) - \lim_{n \rightarrow \infty} \frac{1}{n} E_P \log (\hat{P}(X_0^{n-1})).$$

The first term on the RHS is finite since the original process is stationary, and does not depend on the choices of transition probabilities so it can be replaced by a constant C .

$$D_X(P\|\hat{P}) = C - \lim_{n \rightarrow \infty} \frac{1}{n} E_P \log (\hat{P}(X_0^{n-1})).$$

Next note that due to the stationary Markov transition structure of \hat{P} the expression simplifies according to arguments in [13] resulting in

$$D_X(P\|\hat{P}) = C - E_P \log (\hat{P}(X_1|X_0)).$$

We want to optimize this expression over all non-negative $\hat{P}(X_1 = j|X_0 = i)$ with the constraints that

$$g_i(\hat{P}(X_1|X_0)) = \sum_{n=1}^N \hat{P}(X_1 = n|X_0 = i) = 1 \quad \forall i \in X.$$

For notational convenience let

$$p_{ji} = P(X_1 = j, X_0 = i)$$

and note that

$$\hat{A}_{ij} \equiv \hat{P}(X_1 = j|X_0 = i)$$

The Lagrangian then is

$$L(\hat{A}, \lambda_1, \dots, \lambda_n) = C - \sum_{i=1, j=1}^N p_{ji} \log(\hat{A}_{ij}) + \lambda_i \left(\sum_{i=1, j=1}^N \hat{A}_{ij} - 1 \right).$$

To find the critical point, differentiates with respect to the \hat{A}_{ij} and obtain

$$\frac{\partial L}{\partial \hat{A}_{ji}} = -\frac{p_{ji}}{\hat{A}_{ij}} + \lambda_i = 0.$$

So that the optimal value is given by

$$\frac{p_{ji}}{\lambda_i} = \hat{A}_{ij}.$$

Choosing $\lambda_i = P(X_0 = i)$ gives that,

$$P(X_1 = j|X_0 = i) = \hat{P}(X_1 = j|X_0 = i),$$

the one step transition probabilities must match to minimize the divergence rate between the original process and its Markov approximation.

REFERENCES

- [1] D. A. McQuarrie, "Stochastic approach to chemical kinetics," *Journal of Applied Probability*, vol. 4, pp. 413–478, Dec 1967.
- [2] H. H. M. Adam Arkin, John Ross, "Stochastic kinetic analysis of developmental pathway bifurcation in phage λ -infected escherichia coli cells," *Genetics*, vol. Vol. 149, pp. 1633–1648, 1998.
- [3] S. Burden, N. Napp, and E. Klavins, "The statistical dynamics of programmed robotic self-assembly," in *Conference Proceedings ICRA 06*, pp. pp. 1469–76, May 2006.
- [4] N. V. Kampen, *Stochastic Processes in Physics and Chemistry*. Elsevier, 3 ed., 2007.
- [5] D. T. Gillespie, "A rigorous derivation of the chemical master equation," *Physica A: Statistical Mechanics and its Applications*, vol. 188, pp. 404–425, Sept. 1992.
- [6] D. T. Gillespie, "Stochastic simulation of chemical kinetics," *Annual Review of Physical Chemistry*, vol. 58, no. 1, pp. 35–55, 2007.
- [7] B. Munsky and M. Khammash, "The finite state projection approach for the analysis of stochastic noise in gene networks," *Automatic Control, IEEE Transactions on*, vol. 53, pp. 201–214, Jan. 2008.
- [8] E. Klavins, "Programmable self-assembly," *Control Systems Magazine, IEEE*, vol. 27, pp. 43–56, Aug. 2007.
- [9] S. S. Andrews and D. Bray, "Stochastic simulation of chemical reactions with spatial resolution and single molecule detail.," *Phys Biol*, vol. 1, pp. 137–151, December 2004.
- [10] N. Napp and E. Klavins, "An extended state-space markov chain model for self-organizing systems in non-well-mixed environments," in *4th Annual Conference on the Foundations of Nanoscience*, (Snowbird, UT), April 2007. Available at: <http://glycine.ce.washington.edu/nnapp/fnano07.ppt>.
- [11] P. Billingsley, *Probability and Measure*. John Wiley & Sons, 3rd ed., 1995.
- [12] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley-Interscience, 2nd ed., July 2006.
- [13] L. Finesso, A. Grassi, and P. Spreij, "Approximation of stationary processes by hidden markov models," 2006.