

Setpoint Regulation for Stochastically Interacting Robots

Nils Napp

nnapp@u.washington.edu
Electrical Engineering
University of Washington
Seattle WA 98195

Samuel Burden

sburden@eecs.berkeley.edu
Electrical Engineering
University of California at Berkeley
Berkeley CA 94720

Eric Klavins

klavins@u.washington.edu
Electrical Engineering
University of Washington
Seattle WA 98195

Abstract—We present an integral feedback controller that regulates the average copy number of a particular assembly in a system of stochastically interacting robots. The mathematical model for the stochastic system is a *tunable reaction network*, which makes this approach applicable to a large class of other systems, including ones that exhibit stochastic self assembly at various length scales. We prove that this controller works for a range of set-points, and how to compute this range. Finally, we demonstrate the approach on a physical testbed.

I. INTRODUCTION

Self-assembly of complex systems and structures promises many new applications, such as easily combining different micro-fabrication technologies [1] or building arbitrary, complex nano-structures [2]. While many natural systems are reliably self-assembled at vastly different length and time scales, engineered self-assembled systems remain comparatively simple. The difficulties of engineering complex self-assembling systems are associated with large configuration spaces, our lack of understanding of the relationship between local and global dynamics, and the stochastic or uncertain nature of their dynamics.

The interplay between uncertainty and sensitivity of global to local behavior can often lead to a profound lack of modularity in design techniques, small unintended local interactions can drastically alter the behavior from what is expected by composition.

In this paper we partially address this problem by designing a feedback controller that can regulate the expected value of the number of some an arbitrary component type. This approach could be used for composition in the sense that other subsystems can rely on the presence of these regulated quantities.

We are guided by the application of *stochastic self-assembly*, in which self-assembling particles interact randomly. Such systems abound in engineered settings, such as in DNA self-assembly [2], micro and meso-scale self-assembly [1, 3, 4], and robotic self-assembly [5, 6].

Self-assembly can be either *passive* or *active*. Designing systems that passively self-assemble is a problem of engineering a favorable free energy landscape in configuration space. Passive self-assembling systems often lack flexibility since a specific design of the energy landscape can be difficult to adapt to new tasks. In addition, there are often physical limitations to

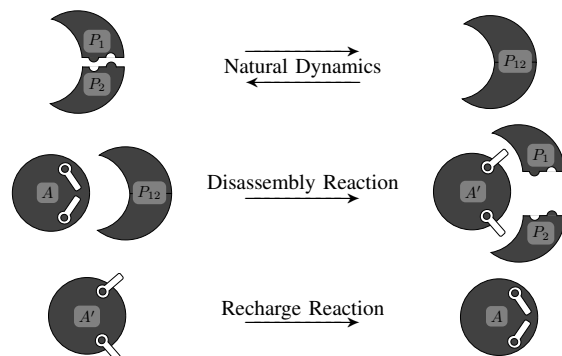


Fig. 1. Schematic representation of the different robot interactions. The passive robots P_1 and P_2 can form heterodimers, which can disassemble spontaneously. The active robot A can expend energy to undo bonds. When the arms of an active robot are retracted, it is charged and can actively disassemble a dimer. If the arms of an active robot are extended (denoted A') then it is not charged, but may become charged via the recharge reaction, the rate of which can be controlled.

how much the energy landscape can be manipulated. However, the yield of the desired output structure is a function of the shape and depth of energy wells, as a result passive self-assembly often leads to low yields.

In *active self-assembly*, energy can be locally injected into the system. In particular, we focus on the situation when we have the ability selectively undo bonds that are formed by the passive dynamics. Active self-assembly can overcome the lack of flexibility in passive self assembling system by making aspects of the system re-programmable and making it clear how these changes affect the system as a whole.

The particular model for active self-assembly we investigate is that of a *tunable reaction network*. We present a system of simple stochastically interacting robots that are well modeled as a tunable reaction network and demonstrate a feedback setpoint regulation scheme. Fig. 1 shows a pictorial representation of the tunable reaction network investigated in this paper. There are three robot types and several instances of each (see Fig. 2(a)). The passive robots P_1 and P_2 are able to bind and form heterodimer complexes P_{12} , which in turn can spontaneously disassemble. The active robots A can dock with heterodimers and disassemble them. The disassembly reaction leaves active robots in an uncharged state, denoted by A' . The last reaction in Fig. 1 recharges uncharged robots at a rate that

is externally controllable. *The control problem for this system is to regulate the number of heterodimers P_{12} in the system by adjusting the recharge rate.* (This problem is stated formally in Sec. IV.) While the tunable reaction network shown in Fig. 1 is comparatively simple, tunable reaction networks in general can describe much more complicated systems.

For example, many biological systems can be viewed as tunable reaction networks. Inside cells, enzymes can be either expressed or not to control the rates of various metabolic reactions. Similar to the problem solved here, one of the many functions of the biochemical processes inside cells is maintaining equilibria of chemical species. Regulating the concentration of chemical species is a particular aspect of *homeostasis*, which can be viewed as a control problem [7].

For the artificial systems depicted in Fig. 1 we propose, analyze, and implement a feedback controller. However, the proposed controller works for tunable reaction networks in general, since the analysis and proof in Sec. IV do not rely on any particular structure of the network. In the context of engineering self-assembling systems, the proposed feedback controller can be used to provide stable operating conditions for other self-assembling processes, much like homeostasis in biological systems. For example, in a hypothetical system with a vat of self-assembling miniature robots, we might care that the relative concentration of robot feet and robot legs is fixed in order to maximize the yield of functioning miniature robots. In general, we envision the self-assembling systems of the future as having metabolisms of their own that regulate the various species of partially assembled objects in the system to maximize the yield of the desired final assembly.

II. EXPERIMENTAL ROBOTIC CHEMISTRY

The robots described here interact stochastically as in [5, 6], however, they are much simpler both mechanically and electronically. Also, while other robotic platforms consist of a homogeneous group of robots, the robotic testbed described here is a heterogeneous mixture of three different robot types, Fig. 2(b)(c). The assembly of the two passive robot types P_1 and P_2 is driven by complementary shape and embedded magnets. The magnetic force creates an energy well that tends to pull two robots together and form a heterodimer. The third, active robot type can expend energy to disassemble a heterodimer into its constituents.

The energy for this disassembly is supplied to the active robots via solar panels. Each active robot can store energy from the solar panel in an on-board capacitor. If the charge in the capacitor reaches a threshold and an active robot A is bound to a heterodimer then the robot activates a motor and disassembles the heterodimer. Disassembling heterodimers depletes the on-board energy storage of active robots, requiring more energy from the solar cells to disassemble additional heterodimers. Adjusting the amount of incident light changes the recharge rate of active robots and thus indirectly affects the rate at which heterodimers are disassembled.

Although this indirect approach may seem unnecessarily complicated, it possesses a key design feature that we believe




Fig. 2. Hardware for test-bed. (a) Picture of the air-table showing the robots, the air-jets, the overhead lamps, and the overhead camera. (b) Picture of the two passive component types showing the tracking targets and the complementary shapes. (c) The active robot showing solar cells, contact sensors, the spinning levers that pull bound passive complexes apart.

justifies this complexity: the structural, energy delivery, and computational functions reside on separate components of the overall system. We think of P_1 and P_2 as the structural components we want to control, the active robots as agents of energy delivery, and the controller implemented on a computer as the computational component. This division of labor is analogous to many biological systems where different cellular functions are largely separated into different types of molecules. We believe that such a separation of functionality in self organization is essential to engineering large scale complex systems. Distributing the functionality in this way can yield much simpler individual components on average. For example, the passive robots contain no electronic components whatsoever, and the active robots only contain a simple circuit made from discrete electrical components, a motor, and a solar panel.

A. Physical Characteristics of Testbed

The body of each robot is machined from polyurethane prototyping foam and painted black to aid the vision system. This material is easy to machine, light, and stiff.

The robots float on the air-table shown in Fig. 2(a), which has a large HVAC blower attached to the bottom of a perforated board (blower not visible in image). The blower is able to maintain a high flow-rate of air through the table surface and allows us to float relatively heavy pieces $\approx 2.5 \frac{\text{g}}{\text{cm}^2}$. The active area of the table is $60\text{cm} \times 60\text{cm}$. Mounted along the perimeter of the table are computer controlled solenoid valves. These valves can deliver short bursts of pressurized air from a compressor (30psi). By randomly activating these air-jets robots on the air-table are driven to perform a random walk. The bursts are randomized and controlled via a MATLAB script, which can also adjust the intensity of four overhead lamps. These lamps determine the amount of incident light to the solar panels, thereby setting the recharge reaction rate.

Images from the overhead camera are used to extract the number and position of *targets*, consisting of small, circular disks with a pattern of concentric light and dark rings, see Fig. 2(b). We detect targets in real time and use the data both in the feedback loop to exert control and open loop to estimate the system reaction rates and diffusion constants.

The number of heterodimers is determined by using the same targets and adjusting the image processing parameters so that only whole targets register in the vision system. Half of a target is attached each passive robot. When two robots bind they create a complete target, which the image processing script detects. Note that the rotational symmetry of the targets simplifies the image processing by reducing the convolution of the target kernel with the image to two dimensions instead of three. This allows us to sample the testbed state at ≈ 1 Hz.

III. MATHEMATICAL MODEL

This section describes *stochastic chemical kinetics* [8] and the associated *chemical master equation* (CME), used to model the discrete configurations of the robotic testbed. This section also describes a *stochastic hybrid system* (SHS) model that extends stochastic chemical kinetics to include continuous state variables, needed to model the closed loop feedback system.

A. Model for Stochastic Chemical Kinetics

The idea is to create a stochastic model that reflects our understanding of how chemical reactions occur at a microscopic level, as opposed to *mass action kinetics*, which is a deterministic model of the evolution of chemical concentrations. When the number of molecules involved in a set of chemical reactions grows, the approximations of mass action kinetics become very good. The large number of molecules averages stochastic effects away [9, Ch. 5.8]. However, when only a few molecules are involved, the stochastic nature of chemical reactions dominates the dynamics and requires explicit modeling.

Let


$$S = \{A, A', P_1, P_2, P_{12}\},$$

denote the set of chemical species, in this case the robot types of the testbed. The symbol A stands for an active robot that is charged, A' is an uncharged active robot. The symbol P_1 and P_2 are the two different types of passive robots and P_{12} is a heterodimer of passive robots, see Fig. 1 and 2. The *copy number* of each species is the number of instances of that particular species and is denoted by a capital N subscripted with the appropriate symbol, i.e. N_A specifies the copy number of species A . The *state* q of the system is described by the vector of copy numbers $\mathbf{q} = (N_A, N_{12}, N_{A'}, N_1, N_2)^T$. The set of all possible states is denoted by Q .

Events that affect the state \mathbf{q} are called *reactions*. This paper considers the set of reactions in Fig. 1. In General, if reactions are indexed by some set L and the state of the system is \mathbf{q} before a reaction l and \mathbf{q}' after the reaction, then we have

$$\mathbf{q}' = \mathbf{q} + \mathbf{a}_l,$$

where \mathbf{a}_l is a vector that is specific to the reaction type. The chemical species that correspond to negative entries in \mathbf{a}_l are called *reactants* and those that correspond to positive entries are called *products*. For example, the reaction



where two different passive robots form a dimer has the associated a vector

$$\mathbf{a} = (0, 0, -1, -1, 1)^T.$$

Both P_1 and P_2 are reactants and P_{12} is a product. The *multiplicity* of a reaction from a given state \mathbf{q} , denoted $M(\mathbf{a}, \mathbf{q})$, specifies the number of different ways the reactants of \mathbf{a} can be chosen from state \mathbf{q} . In addition to the \mathbf{a} vector each reaction has associated with it a *rate constant* k_l , that depends on the underlying stochastic behavior of the interacting species. Determining these rate constants for the system of robots is the topic of Sec. III-D.

Stochastic chemical kinetics defines a discrete state, continuous time *Markov process* with state space Q and the following transitions rates. The transition rate between \mathbf{q} and \mathbf{q}' is given by

$$k_l M(\mathbf{a}_l, \mathbf{q}),$$

when $\mathbf{q}' = \mathbf{q} + \mathbf{a}_l$ and \mathbf{a}_l is applicable in \mathbf{q} (i.e. \mathbf{q}' is non-negative). Given that the process is in state \mathbf{q} at time t , the probability of transitioning to state \mathbf{q}' within the next dt seconds is

$$k_l M(\mathbf{a}_l, \mathbf{q}) dt.$$

This property suffices to define the conditional transition probabilities of the stochastic process and together with an initial distribution over the states defines the Markov process that comprises the stochastic chemical kinetics model. This model is applicable to a set of interacting molecules if the system is *well mixed* [9, 10]. In practice this assumption is difficult to verify. However, in our system of robots we can explicitly check if the assumptions are true, since we can observe the position of all involved particles. A description of the procedures used to verify the well-mixed assumption is given in Sec. III-C.

Conveniently, discrete state Markov Processes can be expressed as linear algebra in the following way. Fix an enumeration of Q and let \mathbf{p}_i denote the probability of being in the i th state $\mathbf{q} \in Q$. The enumeration is arbitrary but assumed fixed for the remainder of this paper. The dynamics of the probability vector \mathbf{p} are governed by the *infinitesimal generator* \mathbf{A} defined as follows: All entries of \mathbf{A} are zero unless:

- If $i \neq j$ and $\mathbf{q}_i + \mathbf{a}_l = \mathbf{q}_j$: $\mathbf{A}_{ij} = k_l M(\mathbf{a}_l, \mathbf{q}_i)$
- If $i = j$: $\mathbf{A}_{ii} = -\sum_m \mathbf{A}_{im}$.

By construction the rows of \mathbf{A} sum to zero and all off-diagonal entries are non-negative. Probability mass functions over Q are expressed as row vectors and real functions of Q , $y : Q \rightarrow \mathbb{R}$



as column vectors. The dynamics an arbitrary probability mass function \mathbf{p} is governed by

$$\dot{\mathbf{p}} = \mathbf{p}\mathbf{A}, \quad (1)$$

the chemical master equation.

B. A Reaction Network for the Testbed

The reaction network description for our robotic testbed consists of four distinct reactions: two describe the spontaneous association and dissociation of passive robots P_1 and P_2 , one describes the disassembly of P_{12} by active robots, and the last reaction describes recharging of active robots. Denote the rate constant for association and dissociation by the natural dynamics by k_1 and k_{-1} , for the disassembly reaction by k_2 , and for the tunable recharge reaction by k_3 . The rate constant for the tunable recharge reaction corresponds to the maximal physically possible rate, in this case highest operating intensity of the overhead lamps. These reactions are summarized in (2)-(4).



Note that the rate constant in (4) depends on u . As a result the infinitesimal generator matrix \mathbf{A} is a function of u .

The discrete state space Q is finite and obeys the conservation equations

$$N_1 + N_{12} = N_2 + N_{12} \doteq C_1, \quad (5)$$

and

$$N_A + N_{A'} \doteq C_2. \quad (6)$$

The first relation (5) holds when the system has the same number of both types of passive robots (C_1 of each, which we ensure in our experiments), while (6) asserts that there are C_2 active robots that can either be in a charged or discharged state. As a consequence of (5) and (6), N_1 , N_2 , and A' can be expressed in terms of N_{12} , A , and the constants C_1 and C_2 . Instead of five different species we can keep track of only two. For the remainder of this paper we will assume that

$$\mathbf{q} = \begin{pmatrix} N_A \\ N_{12} \end{pmatrix} \in \mathbb{N}^2$$

and note that the copy number for the missing species can be reconstructed from this reduced state.

C. Checking the Well-Mixed Condition

There are several equivalent definitions of what it means for a system to be well-mixed. Basically, all definitions are sufficient conditions for guaranteeing that a process is Markov and that each *possible* combination of reactants for a particular reaction \mathbf{a}_i is equally likely to be involved in the next reaction. While being well-mixed in this sense is a strong assumption, it




Fig. 3. Observed distribution of robots on air-table. The left figure corresponds to active robots, A or A' , the middle plot correspond to passive robots P_1/P_2 , and the right figure to heterodimers P_{12} . These plots demonstrate that the occupancy of parts on the air-table is roughly uniform on the table. The area of low occupancy around the perimeter is due to the geometry of the components interacting with the boundary of the air-table.

allows for the characterization of a reaction by a single parameter, the rate constant k_i . For the remainder of this section we use the definition of well-mixedness from [10]. For alternative conditions see [9, Ch. 7.2]. The two conditions that must be checked are that: (a) the reactants are uniformly distributed throughout the environment and (b) that the reactants diffuse through the reaction domain faster than they react.

To estimate the distribution of the different types of robots on the air-table we decomposed it into a 11×11 grid and extracted the occupancy statistics for each grid box from video data. Fig. 3 shows the resulting distributions. The red area in the center of each plot is roughly at the same level and indicates a uniform distribution. The area of low occupancy around the perimeter results from the fact the position of each robot is estimated at its center yet geometric constraints keep the center away from from the air-table border.

The diffusion coefficient for a robot is defined as

$$D = \frac{\mathbb{E} r^2(t)}{4t},$$

where $r(t)$ denotes the displacement of the robot as a function of time. We used the targets described in Sec. II to track the position of different robot types. We averaged over multiple experiments as well as the instances of of each robot type to compute the expected value. The resulting estimates for the diffusion coefficient are given in Tab. I. The subscripts of D indicates what robot type the diffusion coefficient was calculated for. For example, $D_{P_{12}}$ is the diffusion coefficient of heterodimers.

Combined with the rate constants measured in Sec. III-D we conclude that condition (a) and (b) are approximately met. The testbed is well-mixed and the stochastic chemical kinetic model is appropriate.

D. Characterizing Rate Constants

One method to determine rate constants is to measure the average waiting time between reactions from a known state. This quantity, together with the known inverse relationship between the reaction rate and average waiting time, yields an estimate of the rate [5]. Although useful in simulation, one drawback of this method is that one needs to repeatedly re-initialize the system to gather statistical data, which is tedious and time consuming. An exception is k_3 , which was measured




Fig. 4. Curve fitting results used to determine rate constants.

in this way. The reason is that the recharge reaction represents a change in internal state, which is easy to re-initialize.

For the other rate constants we take a different approach. We average multiple longer trajectories all *starting* from the same initial condition. However, the system is allowed to continue evolving for a set amount of time, possibly undergoing many reactions. This has the advantage that each re-initialization gives much more information than a single waiting time. We then fit this empirical average to solutions of the CME (1).

We determined the remaining rate constants k_1, k_{-1} and k_2 in two steps. First, we gathered trajectories starting from $N_{12} = N_A = 0$ with $u = 0$ (lights off). This way the disassembly reaction and recharge reaction do not influence the natural dynamics. We then used MATLAB to numerically fit the CME solution with the two free parameters k_1 and k_{-1} to the empirical average, minimizing the mean squared error, see Figure 4(a).

Using the values previously determined for k_3, k_1 , and k_{-1} we then used the same approach (this time with $u = 1$) for determining the only remaining parameter in the CME solution, k_2 . The resulting curve fit is shown in Fig. 4(b).

Parameter	Estimate	Uncertainty	Units
k_1	0.0046		$\frac{\text{reaction}}{\text{sec number}^2}$
k_{-1}	0.00071		$\frac{\text{reaction}}{\text{sec number}}$
k_2	0.0027		$\frac{\text{reaction}}{\text{sec number}^2}$
k_3	0.08		$\frac{\text{reaction}}{\text{sec number}}$
$D_A/D_{A'}$	0.0018	0.0002	$\frac{\text{m}^2 \text{sec}}{\text{sec}}$
D_{P_1}/D_{P_2}	0.0015	0.0001	$\frac{\text{m}^2 \text{sec}}{\text{sec}}$
$D_{P_{12}}$	0.00083	0.00001	$\frac{\text{m}^2 \text{sec}}{\text{sec}}$

TABLE I

ESTIMATES OF RATE CONSTANTS AND DIFFUSION COEFFICIENTS.

IV. CONTROLLER DESIGN AND ANALYSIS

This section describes an integral feedback controller for the reaction network (2)-(4) and a stochastic hybrid system (SHS) that describes the closed loop system.

The idea of the feedback controller is simple. If we increase the light intensity, the active robots will recharge more quickly and disassemble more heterodimers. An increase in u will decrease N_{12} . If we decrease the amount of light then fewer charged robots will be available to disassemble heterodimers and the average number of N_{12} will increase. With the




Fig. 5. Block diagram of the proposed control system. Only the air-table state and output signal are discrete, all other signals are continuous.

mathematical model for the testbed in place we can state the problem formally.

Design a control system that measures N_{12} and adjusts intensity of the overhead lamps such that $\mathbb{E}N_{12} = y^$ for a given reference value y^* .*

A. Integral Control

The discrete state \mathbf{q} of the closed loop system develops according to (2)-(4). Let \mathbf{y} be the a vector corresponding to an output function $y : Q \rightarrow \mathbb{R}$, in this case $y(\mathbf{q}) = N_{12}$. Define a new continuous part of the state that models the cumulative error from a setpoint y^* as

$$\dot{x} = f(\mathbf{q}, x) = \gamma(y(\mathbf{q}) - y^*) = \gamma(N_{12} - y^*). \quad (7)$$

In order to express saturation of the input, here the physical limitations of the overhead lamps, define $h : \mathbb{R} \rightarrow \mathbb{R}$ by

$$h(x) = \begin{cases} 0, & x \leq 0 \\ x, & 0 < x \leq 1 \\ 1, & 1 < x. \end{cases}$$

With this notation we can define an integral feedback controller by

$$u = h(x). \quad (8)$$

A block diagram of the control system is shown in Fig. 5. The remainder of this section is dedicated to analyzing the closed loop system.

B. Stochastic Hybrid System

Adding a continuous random variable whose dynamics depend on the discrete state \mathbf{q} of a Markov process results in a *stochastic hybrid system* (SHS). This section is a brief description of the notation and some specific mathematical tools available for SHSs, for more information see [11, 12].

The key feature of an SHS is that the dynamics of the system are stochastic and that the state are hybrid, meaning the state space of the system has the form $Q \times X$ where Q is some discrete set and $X \subseteq \mathbb{R}$ is continuous. The set of possible discrete states Q is typically finite or countably infinite. We use $z \in Z = Q \times X$ to denote the pair (q, x) . Denote the stochastic processes describing the system on the various components of the state space by \mathcal{Q} , \mathcal{X} , and \mathcal{Z} .

In each discrete state, the dynamics of \mathcal{X} are governed by a differential equation,

$$\dot{x} = f(\mathbf{q}, x) \quad f : Q \times X \rightarrow TX. \quad (9)$$

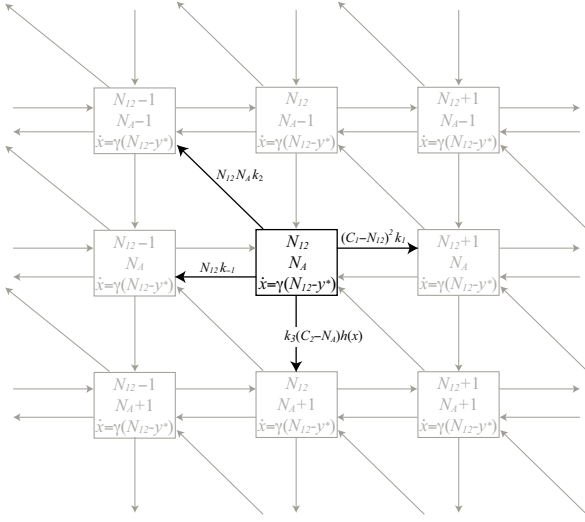


Fig. 6. A schematic representation of the closed loop SHS. The boxes represent discrete states and the arrows represent transitions. Each box shows both the discrete state it represents and the ODE describing the continuous states. An arbitrary state $(N_A, N_{12})^T$ is highlighted in black. The transition intensities for all transitions leaving $(N_A, N_{12})^T$ are shown next to the arrows.

The dynamics of the discrete state Q are governed by a set of transitions, indexed by a finite set L . Each transition $l \in L$ has associated with it an intensity function

$$\lambda_l(\mathbf{q}, x) \quad \lambda_l : Q \times X \rightarrow [0, \infty), \quad (10)$$

and a reset map

$$(\mathbf{q}, x) = \phi_l(\mathbf{q}^-, x^-) \quad \phi_l : Q \times X \rightarrow Q \times X. \quad (11)$$

The intensity function is the instantaneous rate of the transition l occurring, so that $P(l \text{ occurs during } (t, t + dt) | Q = \mathbf{q}, X = x) = \lambda_l(\mathbf{q}, x, t) dt$. The reset map ϕ_l determines where the process jumps after a transition is triggered at (\mathbf{q}^-, x^-) at time t . The minus in the superscript denotes the left hand limit of \mathbf{q} and x at time t . We think of this limit as the state of the process immediately before the jump. Fig. 6 shows part of the system representing the closed loop SHS. The boxes represent discrete states and the arrows represent transitions.

C. Extended Generator

This section describes the *extended generator*, \mathcal{L} , associated with an SHS. This operator is analogous to the *generator matrix* of a discrete state Markov process but in the hybrid case is a partial differential equation describing the dynamics of the expected value of arbitrary *test functions* on the state space. In particular, the extended generator allows us to derive ordinary differential equations (ODEs) that govern the dynamics of the statistical moments of the state variables of an SHS.

Operator \mathcal{L} in (12) is the extended generator for an SHS described by (9)-(11). Let ψ be a real valued function on $Q \times X$ and define

$$\mathcal{L}\psi(z) = \frac{\partial \psi(z)}{\partial x} f(z) + \sum_{l \in L} (\psi(\phi_l(z)) - \psi(z)) \lambda_l(z). \quad (12)$$

The operator \mathcal{L} has the following useful property relating the time derivative of the expected value of a test function ψ to $\mathcal{L}\psi$

$$\frac{d \mathbb{E}\psi}{dt} = \mathbb{E} \mathcal{L}\psi \quad (13)$$

[11]. The extended generator for the closed loop system is given by

$$\begin{aligned} \mathcal{L} \psi(N_{12}, N_A, x) &= \frac{\partial \psi(N_{12}, N_A, x)}{\partial x} \gamma(N_{12} - y^*) \\ &+ (\psi(N_{12} + 1, N_A, x) - \psi(N_{12}, N_A, x)) k_1 (C_1 - N_{12})^2 \\ &+ (\psi(N_{12} - 1, N_A, x) - \psi(N_{12}, N_A, x)) k_{-1} N_{12} \\ &+ (\psi(N_{12} - 1, N_A - 1, x) - \psi(N_{12}, N_A, x)) k_2 N_{12} N_A \\ &+ (\psi(N_{12}, N_A + 1, x) - \psi(N_{12}, N_A, x)) x (C_2 - N_A). \end{aligned} \quad (14)$$

It can be used to find ODEs describing the evolution of the statistical moments of the SHS. Specifically, letting $\psi = x$ we obtain

$$\frac{d \mathbb{E}x}{dt} = \mathbb{E} \gamma(N_{12} - y^*). \quad (15)$$

If the closed loop system is stochastically stable, in the sense that the probability distribution of states approaches a fixed invariant distribution, then by (15) we can conclude that

$$\mathbb{E}N_{12} = y^*.$$

The controller works in expected value when the system is in steady state. Now, the problem of showing correctness of the controller reduces to showing that the system is stochastically stable or *ergodic*, i.e. that the system always approaches a unique steady state distribution.

D. Ergodicity

We use a Lyapunov function argument [13, THM 5.1] to show that the closed loop SHS is ergodic. This allows us to set the LHS in (15) to zero and argue that the controller works in steady state. We show that the system is ergodic for some reference values y^* and give sufficient conditions for ergodicity for a range of y^* .

Denote the generator matrices of minimum and maximum input by $\mathbf{A}_m = \mathbf{A}(0)$, $\mathbf{A}_M = \mathbf{A}(1)$ and the corresponding steady state probability mass functions by \mathbf{p}_m and \mathbf{p}_M respectively.

Theorem: Let $\mathbf{A}(u)$ be the generator of a tunable reaction network and \mathbf{y} the vector corresponding to an output function $y : Q \rightarrow \mathbb{R}$ of the discrete state. The feedback controller proposed in (8) results in a closed loop system with a stationary distribution that has $\mathbb{E}y = y^*$ when y^* is in the *controllable region*, $\mathbf{p}_M \mathbf{y} < y^* < \mathbf{p}_m \mathbf{y}$.

Note: If $\mathbf{p}_M \mathbf{y} > \mathbf{p}_m \mathbf{y}$, then the theorem applies with the sign in (7), and the upper and lower limits of the controllable region flipped.

Proof: Let \mathcal{Z} be the SHS corresponding to the closed loop system. By [13, THM 5.1], \mathcal{Z} is ergodic when there exists a

function $V : Z \rightarrow \mathbb{R}^+$ with the property that $V(z) \rightarrow \infty$ as $|z| \rightarrow \infty$ and

$$\mathcal{L}V(z) \leq -f(z) \quad \forall z \notin C \quad (16)$$

for some compact region C and positive function f .¹

For our system, we define the function \widehat{V} to be

$$\widehat{V}(q, x) = \begin{cases} x + c^+(q) & \text{for } x > 0 \\ -x + c^-(q) & \text{for } x < 0, \end{cases}$$

where c^+ and c^- depend on q . Note that the function \widehat{V} is neither differentiable (required to apply \mathcal{L}) nor strictly positive (required by theorem) since the constants can be negative. To address this problem, let V be some function that agrees with \widehat{V} when x is outside some interval $[v_{min}, v_{max}]$ for all $q \in Q$, and is both non-negative and twice differentiable. This function always exists since the q is assumed to be finite and \widehat{V} increases with $|x|$ so that \widehat{V} is positive for sufficiently large $|x|$.

Let the compact region required by the theorem be $C = Q \times [\min(v_{min}, 0), \max(v_{max}, 1)]$. Since we are only interested in V outside C , we look at the cases when the feedback input is saturated at either $u = 0$ or $u = 1$. This situation simplifies the analysis, since the transition intensities $\lambda(q, x)$ are independent of x in the saturated regions. We now argue that for some range of set points y^* we can find c^+ and c^- to make V a Lyapunov function in the sense of (16).

Choosing $f = \epsilon$ and considering saturation at $u = 1$ first, we can rewrite the conditions of (16) in vector form,

$$\mathbf{y} - y^* \mathbf{1} + \mathbf{A}_M \mathbf{c}^+ \leq -\epsilon \mathbf{1}. \quad (17)$$

Let $\tilde{\epsilon}$ be an arbitrary vector with strictly positive entries, then we can rewrite (17) as

$$\mathbf{y} - y^* \mathbf{1} + \mathbf{A}_M \mathbf{c}^+ = -\tilde{\epsilon}. \quad (18)$$

We want to determine when this equation has a solution for \mathbf{c}^+ . Note that

$$\mathbf{A}_M \mathbf{c}^+ = -\tilde{\epsilon} + y^* \mathbf{1} - \mathbf{y}$$

has a solution only if $(-\tilde{\epsilon} + y^* \mathbf{1} - \mathbf{y})$ is in the column space of \mathbf{A}_M , which we write $(-\tilde{\epsilon} + y^* \mathbf{1} - \mathbf{y}) \in \text{Col} \mathbf{A}_M$. Equivalently

$$(\text{Col} \mathbf{A}_M)^\perp \perp (-\tilde{\epsilon} + y^* \mathbf{1} - \mathbf{y}) \quad (19)$$

$$(\text{Nul} \mathbf{A}_M^T) \perp (-\tilde{\epsilon} + y^* \mathbf{1} - \mathbf{y}) \quad (20)$$

$$(\mathbf{p}_M^*)^T \perp (-\tilde{\epsilon} + y^* \mathbf{1} - \mathbf{y}) \quad (21)$$

$$0 = \mathbf{p}_M^* (-\tilde{\epsilon} + y^* \mathbf{1} - \mathbf{y}) \quad (22)$$

$$0 = -\mathbf{p}_M^* \tilde{\epsilon} + y^* - \mathbf{p}_M^* \mathbf{y}, \quad (23)$$

where Nul denotes the null space and \perp the orthogonal complement. Since $\tilde{\epsilon}$ has arbitrary, strictly positive entries, a solution for \mathbf{c}^+ exists when

$$\mathbf{p}_M^* \mathbf{y} < y^*.$$

¹The theorem has some technical preconditions, which are fulfilled in our case, namely that all compact sets are petite see [13]. This follows from [14, THM 4.1], [12, THM 27.6] and the fact that every Feller process is also a T-process.

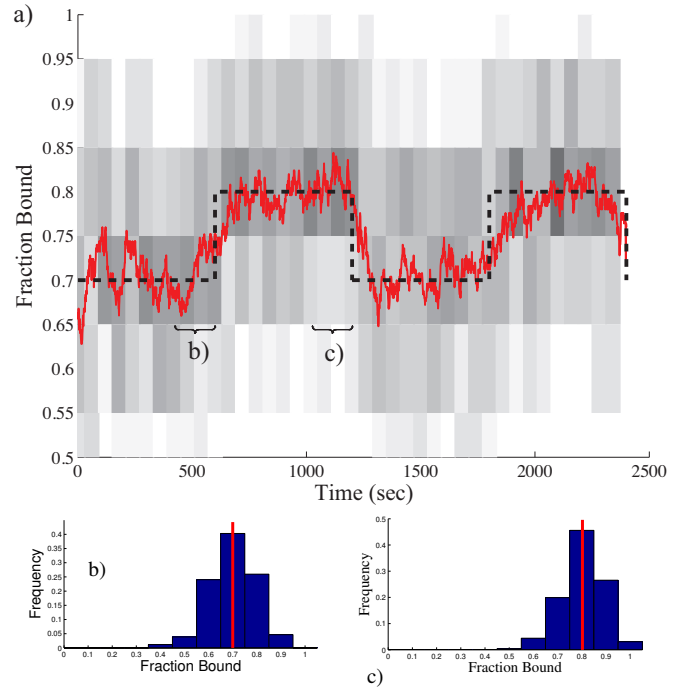


Fig. 7. Tracking data from robotic test-bed. This is the average of 25 different experiments. Each experiment has 10 P_1 and P_2 each, and 4 active robots. The grey shading in the background corresponds to the fraction of trajectories with that number of pairs. Darker regions correspond to higher occupancy. The two histograms on the right show the fraction of possible dimers taken for the last half of each step. The red line is the mean, which demonstrates correct behavior for the controller.

Similarly, for saturation with $u = 0$ we get

$$\mathbf{p}_M^* \mathbf{y} > y^*.$$

Thus the system is ergodic if

$$\mathbf{p}_M^* \mathbf{y} < y^* < \mathbf{p}_m^* \mathbf{y}.$$

Furthermore, by (15) the expected value of N_{12} tracks the reference value y^* when it is in the controllable region. ■

The proof does not rely on any special structure of $\mathbf{A}(u)$ nor the value of γ , as a result the theorem is generally applicable to tunable reaction networks with saturating inputs.

V. EXPERIMENTAL RESULTS

We implemented the proposed controller on the robotic test-bed described in Sec. II. The generator matrix $\mathbf{A}(u)$ is defined by (2)-(4), and the output function is $y(\mathbf{q}) = N_{12}$. To show its tracking capability we tracked two periods of a square wave. The low and high set-points were 0.7 and 0.8 (corresponding to 7 and 8 P_{12}). Both of the setpoints are inside the empirically determined controllable region for this system, 0.60-0.86.

The combined results of 25 trajectories are shown in Fig. 7. We let each trajectory run with a set point of 0.7 for 5 minutes




Fig. 8. Standard deviation of output for different set points and integrator constants γ . The red line at the bottom corresponds to the standard deviation when the system is under open loop control. The green line at the top corresponds to the standard deviation of the system when driven with a bang-bang controller and the input is always saturated.

(a half period) before recording data, which allowed transients resulting from the manual initialization to dissipate. After the warm up period we collected 20 minutes of data for each trajectory.

This experiment demonstrates that the controller tracks the reference signal in mean (Fig. 7(b)(c)). This experiment also demonstrates the fundamental stochasticity in the system. The spread in Fig. 7(b)(c) is not due to measurement uncertainty or noise, but a fundamental property of the stochastic system we are controlling.

We also present some preliminary simulation experiments exploring how the variance of the copy number relates to the integrator gain γ . The proof for tracking in mean did not depend on the value of γ , so the proposed controller will always yield the desired mean steady-state copy number. However, it might differ in the degree of fluctuation around the correct mean.

The relation between set-point, integrator gain, and standard deviation at steady state are shown in Fig. 8. Each data point was collected by adjusting γ and estimating the standard deviation at steady state. This approach limits smallest value of γ we can feasibly simulate, since small values slow down the system dynamics and make simulations prohibitively computationally expensive.

We observe that less aggressive values of γ result in a smaller standard deviation of the output. The upper and lower limits of the standard deviation correspond to open-loop and bang-bang control. Another interesting feature of Fig. 8 is that the standard deviation of the output seems to be less sensitive to γ if the reference y^* is close to the edge of the controllable region.

VI. CONCLUSIONS AND FUTURE WORK

We proposed an integral feedback controller for the average copy number of an arbitrary species in a system modeled by stochastic chemical kinetics. We proved that the controller tracks a reference in mean and demonstrated the approach

on a robotic experimental platform. We also presented some preliminary simulation results regarding the variance of the copy number as a function of the integrator gain. We are currently working on analytical results describing the steady state variance of the control scheme. We are also investigating ways to decentralize the controller by using local estimates of the global output.

Finally, we would like to emphasize the generality of our approach. The functionality of the controller requires no tuning of the integrator gain γ as long as the reference is in the controllable region, which is easy to measure experimentally. The detailed structure of the underlying Markov process does not matter.

ACKNOWLEDGEMENTS

We would like to thank Sheldon Rucker, Michael McCourt, and Yi-Wei Li for their dedication and hard work in designing and building the robots. We would also like to thank Alexandre Mesquita and João Hespanha for their helpful discussions about proving ergodicity.

REFERENCES

- [1] E. Saeedi, S. Kim, H. Ho, and B. A. Parviz, "Self-assembled single-digit micro-display on plastic," in *MEMS/MOEMS Components and Their Applications V. Special Focus Topics: Transducers at the Micro-Nano Interface*, vol. 6885, p. 688509, SPIE, 2008.
- [2] P. W. K. Rothmund, "Folding dna to create nanoscale shapes and patterns," *Nature*, vol. 440, pp. 297–302, Mar. 2006.
- [3] M. Boncheva, D. A. Bruzewicz, and W. G. M., "Millimeter-scale self-assembly and its applications," *Pure and Applied Chemistry*, vol. 75, pp. 621–630, 2003.
- [4] H. Onoe, K. Matsumoto, and I. Shimoyama, "Three-dimensional micro-self-assembly using hydrophobic interaction controlled by self-assembled monolayers," *Journal of Microelectromechanical Systems*, vol. 13, no. 4, pp. 603–611, 2004.
- [5] S. Burden, N. Napp, and E. Klavins, "The statistical dynamics of programmed robotic self-assembly," in *Conference Proceedings ICRA 06*, pp. 1469–76, May 2006.
- [6] P. J. White, K. Kopanski, and H. Lipson, "Stochastic self-reconfigurable cellular robotics," in *IEEE International Conference on Robotics and Automation (ICRA04)*, pp. 2888–2893, 2004.
- [7] H. El-Samad, J. P. Goff, and M. Khammash, "Calcium homeostasis and parturient hypocalcemia: An integral feedback perspective," *Journal of Theoretical Biology*, vol. 214, pp. 17–29, Jan. 2002.
- [8] D. A. McQuarrie, "Stochastic approach to chemical kinetics," *Journal of Applied Probability*, vol. 4, pp. 413–478, Dec 1967.
- [9] N. V. Kampen, *Stochastic Processes in Physics and Chemistry*. Elsevier, 3 ed., 2007.
- [10] D. T. Gillespie, "Exact stochastic simulation of coupled chemical reactions," *Journal of Physical Chemistry*, vol. 81, no. 25, pp. 2340–2361, 1977.
- [11] J. P. Hespanha, "Modeling and analysis of stochastic hybrid systems," *IEE Proc — Control Theory & Applications*, Special Issue on Hybrid Systems, vol. 153, no. 5, pp. 520–535, 2007. Available at <http://www.ece.ucsb.edu/hespanha/published>.
- [12] M. Davis, *Markov Processes and Optimization*. Chapman & Hall, 1993.
- [13] S. P. Meyn and R. L. Tweedie, "Stability of markovian processes III: Foster-lyapunov criteria for continuous-time processes," *Advances in Applied Probability*, vol. 25, no. 3, pp. 518–548, 1993.
- [14] S. P. Meyn and R. L. Tweedie, "Stability of markovian processes II: Continuous-time processes and sampled chains," *Advances in Applied Probability*, vol. 25, no. 3, pp. 487–517, 1993.