Multiplexing Robot Experiments: Theoretical Underpinnings, Conditions for Existence, and Demonstrations

Rachel A. Moan Dylan A. Shell Jason M. O'Kane

Abstract— **Scores of papers show, given some robots, how to improve the useful work they perform. Continuing this line, we consider the efficiency of *robot experiments* by examining the feasibility of conducting several experiments simultaneously, interleaving execution and sharing resources between them. This paper lays theoretical groundwork for that concept and demonstrates its feasibility and utility.**

I. Introduction

Demonstrations of capabilities or claims on robot hardware are more persuasive than the same in pure simulation. But robot hardware is burdensome: it is expensive to build or purchase; considerable time must be expended in maintenance, repair, and also operation. Multiple robot systems are only more so. Thankfully, the community has learned to amortize some of these costs through re-use and sharing: designs, schematics, and configurations for several *platforms* have been released and made available so that others may benefit [1–5]. Increasingly, we see not only hardware platforms, but broader *test-beds* [6–12]. Several existing robot test-beds directly facilitate experiments, culminating in shared, remotely-accessible scientific infrastructure [13, 14]. But we ask: to extract greater use from such resources, can re-use and sharing go further still, be finer-grained and more dynamic?

Decades of computer systems research and technology development have shown how to boost the availability, utilization, and effectiveness of computing devices through multiprogramming, multitasking, and multiprocessing (via multithreading and multicore architectures). Notwithstanding that some later instances grew out of earlier ones, this multitude of *multi*s can all be understood as being flexible in how resources are employed. We contend that a similar approach might be used for robots, with the ultimate aim of improving availability, utilization, and effectiveness of robot hardware. In this paper we are specifically interested in common, general, re-usable robot infrastructure designed to conduct robot hardware experiments [13]. For already existing infrastructure, to maximize their benefit, idle portions may well be put to productive use. The idea explored in this paper, and illustrated in Figure 1, is to leverage hardware resources more effectively by intermingling the execution of more than one robot experiment on the same test-bed.

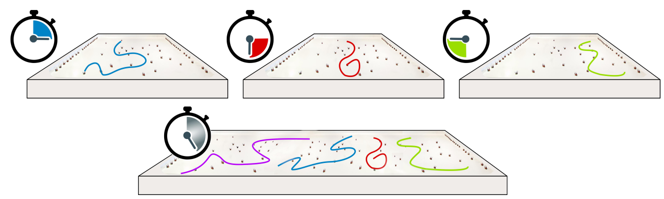

Fig. 1. Robot test-beds currently only time multiplex at a coarse-grained level. For example, to use the Georgia Tech Robotarium for multiple experiments, one schedules them one after another, as illustrated by three arenas. This paper shows that it is feasible to overlap several simultaneous experiments. (One larger, hypothetical, Robotarium allows multiple experiments to proceed in parallel.) The upshot is that tasks previously believed to be inherently sequential, such as tuning parameters/gains, may be effectively speeded up.

The question of what constitutes the essence of a valid robot experiment is a complex, interesting, and even philosophical one. For the present paper the core consideration is how, without losing that essence, to provide an opening for the flexible apportioning of resources. We provide *an* answer: experiments produce a stream of sensor readings contingent on states visited and influenced via selection of actions by robots, such selections being made in light of earlier sensor readings. This point of view, where sensing may be manipulated up the perceptual limits of the agents, has been suggested before in robotics (see [15, 16]), and is now widely employed in biology (see the entire special-issue in [17]). Formalizing this sensing-oriented point of view provides the scaffold over which we build a conception of *multi-experiments*. Beyond the degrees-of-freedom of the robots in the test-bed itself, this context identifies specific additional freedoms which can be used to optimize execution of collections of experiments: freedom to warp time, freedom to distort robot identity.

But how this might bear fruit, concretely? Suppose a roboticist has developed some control software for their favorite mobile robot, say, a controller based on artificial potential fields [18]. This controller has several parameters which manipulate gains and weight various competing factors. Usually these are understood to be empirical parameters, items tuned by running the robot and making adjustments interactively. When some criterion can be provided (such as straightness of travel subject to exemption from collisions) then the time to arrive at a set of suitable parameters depends on completing sufficient runs across the ranges of parameters. If multiple runs can occur in quick succession, or even simultaneously, then overall productivity will be increased.

R. A. Moan is with the Department of Computer Science, Winthrop University, Rock Hill, South Carolina, USA. D. A. Shell is with the Department of Computer Science and Engineering, Texas A&M University, College Station, Texas, USA. J. M. O'Kane is with the Department of Computer Science and Engineering, University of South Carolina, Columbia, South Carolina, USA. moanr2@mailbox.winthrop.edu; dshell@tamu.edu; jokane@cse.sc.edu

This work was supported by the NSF through awards IIS-1453652, IIS-1659514, IIS-1849249, and IIS-1849291.

This is precisely what multi-experiments enable.

The contribution of this paper is to lay a foundation for understanding this notion of multi-experiments, including introducing some fundamental definitions (§ II), establishing conditions on when multi-experiments may be achieved (§ III), examining differences from standard notions of multi-processing (§ IV), and demonstrating a simple but practical multi-experiment (§ V).

II. Definitions

This section introduces the basic definitions from which the subsequent analysis proceeds, starting with this definition of a(n idealized) system in which multiple robots operate.

Definition 1. *A* deterministic multi-robot transition system *[16] is a* 8*-tuple* $(n, X, U, f, d, Y, h, x_0)$*, in which*

1) n *is a positive integer identifying the number of robots,*
2) $X = X^{(1)} \times \cdots \times X^{(n)}$ *denotes a state space, composed of individual state spaces for each robot,*
3) $U = U^{(1)} \times \cdots \times U^{(n)}$ *denotes an action space, composed of individual action spaces for each robot,*
4) $f : X \times U \to X$ *is a state transition function, defined in terms of transition functions* $f^{(1)}, \ldots, f^{(n)}$ *for each robot, so that*
$$f\big((x^{(1)}, \ldots, x^{(n)}), (u^{(1)}, \ldots, u^{(n)})\big) = \big(f^{(1)}(x, u^{(1)}), \ldots, f^{(n)}(x, u^{(n)})\big),$$
5) $d : X \times U \to [0, \infty)$ *is a transition duration function, indicating the amount of time that elapses when executing each action from each state,*
6) $Y = Y^{(1)} \times \cdots \times Y^{(n)}$ *denotes an observation space, composed of individual observation spaces,*
7) $h : X \to Y$ *is an observation function, defined in terms of observation functions* $h^{(1)}, \ldots, h^{(n)}$ *for each robot, so that* $h(x) = \big(h^{(1)}(x), \ldots, h^{(n)}(x)\big)$*,*
8) $x_0 \in X$ *is the system's initial state.*

The execution of such a system proceeds in discrete stages, indexed $k = 1, 2, \ldots$, through a sequence of states $x_0, x_1, \ldots$, influenced by the actions $u_0, u_1, \ldots$ selected by each robot. Actions and states are related by the state transition equation $x_{k+1} = f(x_k, u_k)$. At each stage, the robots' sensors provide observations, determined by the state, so that $y_k = h(x_k)$. We allow the duration of the stages to be non-uniform via the transition duration function d. Writing t_k for the time at the conclusion of action u_k, the times obey the recurrence $t_{k+1} = t_k + d(x_k, u_k)$ with initial condition $t_0 = 0$.

Example 1. A straightforward example, first introduced in a prior paper [16], illustrates the concept. See Fig. 2. Suppose n robots move along a single-lane road. Each robot r_i controls its own speed, within an allowable range $[v_{\min}, v_{\max}]$. Sensors enable each robot to detect the distance to its immediate neighbors both ahead and behind. This scenario is readily modeled as a deterministic multi-robot transition system

$$S_{n,v_{\min},v_{\max}} = (n, \mathbb{R}^n, [v_{\min}, v_{\max}]^n, f, d, \mathbb{R}^+ \times \mathbb{R}^+, h, x_0),$$

in which each state $x \in \mathbb{R}^n$ represents the position along the road of each of the robots, actions represent each robot's chosen velocity at a given time step, and each observation encodes the distance to the nearest robot behind and ahead. We refer the reader to the original paper [16] for more detail.

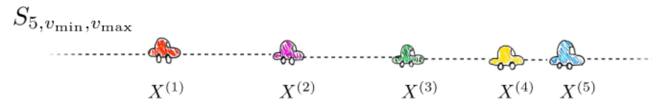

Fig. 2. An illustration of Example 1 for $n = 5$.

Example 2. We note concisely here that public-facing robot multi-robot test-beds such as the Georgia Tech Robotarium [13] can be modeled as deterministic multi-robot transition systems, using the appropriate models of transitions and observations. We used this platform to conduct a number of proof-of-concept executions of the strategies described in this paper —see Figure 3— and utilized a simulator mimicking it to obtain the quantitative results that appear below.

A. (Multi-)illusions

Our primary interest is in the ability of one deterministic multi-robot transition system to present, to some of its robots, the appearance of operating within another system. Particularly relevant is the case in which such illusions are maintained for multiple systems within the same host. The next definition formalizes this idea, generalizing an earlier formulation [16] for single systems.

Definition 2. *For some finite set of deterministic multi-robot transition systems* $\mathcal{S} = S_1, \ldots, S_d$ *where for each* $1 \le j \le d$, $S_j = (n_j,\, X_j,\, U_j,\, f_j,\, Y_j,\, h_j,\, x_{0_j})$ *and* $\widehat{S} = (\widehat{n}, \widehat{X}, \widehat{U}, \widehat{f}, \widehat{Y}, \widehat{h}, \widehat{x}_0)$ *and a positive integer* m_j*, we say that* $\widehat{S}$ *presents a* multi-illusion *for* $\mathcal{S}$ *if there exist*

(i) *robot policies* $\widehat{\pi}^{(1)}, \ldots, \widehat{\pi}^{(\widehat{n})}$ *in* $\widehat{S}$*,*
(ii) *a strictly increasing function* $z_j : \mathbb{Z}^+ \to \mathbb{Z}^+$*, and*
(iii) *an infinite series of functions* $\rho_{j\,k} : \mathbb{Z}_{m_j} \to \mathbb{Z}_{\widehat{n}}$

for any robot policies $\pi_j^{(1)}, \ldots, \pi_j^{(n_j)}$ *in each* S_j*, such that for all* $k \ge 0$ *and all* $1 \le i \le m_j$*, we have*

$$h_j^{(i)}(x_{j\,k}) = \widehat{h}^{(\rho_{j\,k}(i))}\left(\widehat{x}_{z_j(k)}\right). \qquad (\star)$$

We refer to the specific choices of policies and functions satisfying this constraint as a realization *of the illusion.*

Here each of the robots' policies $\widehat{\pi}^{(1)}, \ldots, \widehat{\pi}^{(\widehat{n})}$ in $\widehat{S}$ is a function mapping from information available to that robot

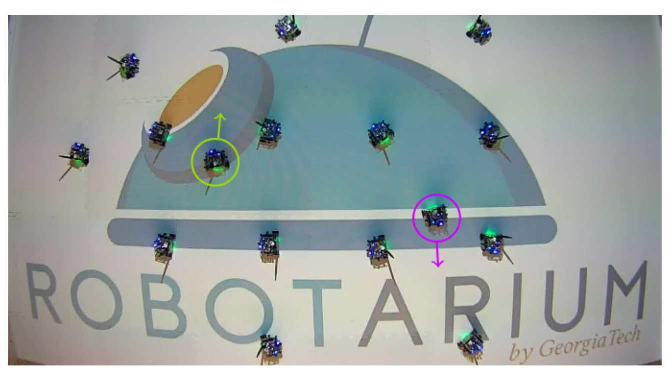

Fig. 3. A collection of 15 robots conspiring to present 2 simultaneous independent illusions, each of a robot moving through an infinite field of obstacles. The illusioned robots are circled.

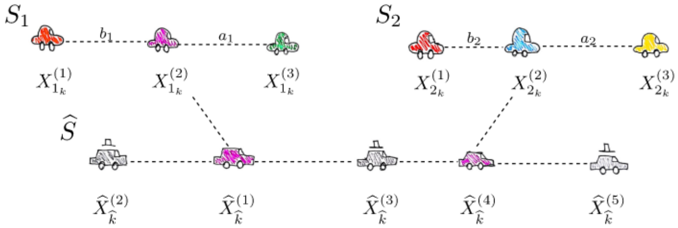

Fig. 4. An illustration of Example 3. A hatted system of 5 vehicles reproduces the expected observations from two separate 3-vehicle systems.

—namely, its own action-observation history in $\widehat{S}$ along with the complete state history in each of the systems in $\mathcal{S}$— to the action that robot should execute.

As already mentioned, what exactly constitutes an experiment is not easy to say. We grant that the notion of presenting an appearance as just formalized retains, operationally, many of the aspects which underlie a robot demonstration. Where the intention is to learn some fact or validate/establish some property, we will be happy to consider experiments as demonstrations up to illusion. And, henceforth, multi-experiments via multi-illusions.

Example 3. Recall the caravan systems from Example 1. For any two such systems $S_1 = S_{3,v_{\min_1},v_{\max_1}}$ and $S_2 = S_{3,v_{min_2},v_{max_2}}$, we can form a multi-illusion of the two systems, with $m_1 = m_2 = 1$, within a host system of the form $\widehat{S} = S_{5,\widehat{v}_{min},\widehat{v}_{max}}$.

One way to construct this illusion is to select policies $\widehat{\pi}$ in which robot 1 moves at a constant speed $(\widehat{v}_{min} + \widehat{v}_{max})/2$. Robots 2 and 3, knowing the desired observation $y^{(1)}_{1_k} = (a^{(1)}_{1_k}, b^{(1)}_{1_k})$ from S_1, position themselves on opposite sides of robot 1, moving as fast as possible at each stage in $\widehat{S}$ toward positions where $\widehat{x}^{(1)}_k - \widehat{x}^{(2)}_k = b^{(1)}_{1_k}$ and $\widehat{x}^{(3)}_k - \widehat{x}^{(1)}_k = a^{(1)}_{1_k}$. Then robots 4 and 5, knowing the desired observation $y^{(1)}_{2_k} = (a^{(1)}_{2_k}, b^{(1)}_{2_k})$ from S_2, position themselves to the right of robot 3 such that $\widehat{x}^{(4)}_k - \widehat{x}^{(3)}_k = b^{(1)}_{2_k}$ and $\widehat{x}^{(5)}_k - \widehat{x}^{(4)}_k = a^{(1)}_{2_k}$. To satisfy the remaining conditions of Definition 2, we define z_1 to return the time when robots 2 and 3 in $\widehat{S}$ have reached their target positions. Similarly, z_2 returns the time when robots 4 and 5 in $\widehat{S}$ have reached their target positions. Figure 4 illustrates the construction, which can, of course, also be generalized to present a multi-illusion of any number d of $S_{3,\cdot,\cdot}$ systems within $\widehat{S} = S_{2d+1,v_{\min},v_{\min}}$.

Notice that Definition 2 allows sufficient freedom that, for a given $\widehat{S}$ capable of presenting an illusion of $\mathcal{S}$ for certain $m_1, \ldots, m_d$, there may be variety of distinct realizations of that illusion, differing in the $\widehat{\pi}^{(i)}$ policies, scheduling functions z_j, and role functions ρ_{j_k}.

Example 4. Significant differences in the realization of a multi-illusion may be realized even in the basic case in which $d = 1$. For example, prior results [16, Example 5] demonstrated that a Robotarium-like system can present a illusion of a mobile robot traversing an arbitrarily large field of disc-shaped obstacles, by holding the recipient robot fixed at the center of the workspace and positioning the other robots at appropriate locations relative to that center, based on the locations of obstacles visible to the recipient robot.

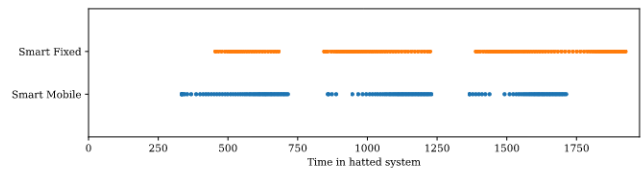

Fig. 5. Progress over time of the two illusions described in Example 4 for identical polices executed in S. Each mark indicates that the system being illusioned has been evolved forward by a step. *Smart Fixed* refers to the original illusion in which the recipient robot stays in the center. *Smart Mobile* refers to the new approach in which the recipient robot moves normally, except for occasional resets to the center. The two data series show the image of the corresponding z functions. That is, each plotted point denotes a stage in $\widehat{S}$ during which S makes a stage of forward progress.

But that same effect might also be achieved by holding the robots playing the role of obstacles fixed in place, and allowing the recipient robot to move normally. Upon nearing the boundary of the available workspace, the illusion can 'reset' to the center, utilizing a suitable z function to halt the execution until the shift back to center is complete. Figure 5 depicts the relative progress of these two illusions as time proceeds.

B. Illusion latency and speedup

A primary question for the applicability of multi-illusions is to understand how efficiently, in terms of time elapsed, the system $\widehat{S}$ can progress through the executions of each of the $S_1, \ldots, S_d$ systems. To that end, we draw inspiration from the computer systems community to define concepts of latency and speedup, tailored for the present setting.

Definition 3. *If $\widehat{S}$ is a multi-illusion for $S_1, S_2, \ldots, S_d$. Then each S_j has* latency *at a stage-concluding time $t_{\widehat{k}}$ defined as*

$$L_j(\widehat{t}_{\widehat{k}}) = \frac{\widehat{t}_{\widehat{k}}}{t_k}, \tag{1}$$

with $k = \max\{\kappa \in \mathbb{Z}^+ \mid z_j(\kappa) \leq \widehat{k}\}$, should such a k exist.

We may compare two illusions via the notion of speedup.

Definition 4. *Given systems $\widehat{S}$ and $\widehat{S}'$ providing illusions for $S_1, \ldots, S_d$, the system $\widehat{S}'$ exhibits* speedup *over $\widehat{S}$ as*

$$\mathcal{S} = \lim_{t\to\infty} \frac{\max\limits_{j=1,\ldots,d} L_j(t)}{\max\limits_{i=1,\ldots,d} L'_i(t)}, \tag{2}$$

where L_j and L'_i denote latencies for $\widehat{S}$ and $\widehat{S}'$, respectively.

Example 5. Recall from Example 4 the different realizations for the illusion of a robot within a large field of disc-shaped obstacles, realized with a host system with a finite workspace. Figure 6 plots the latency of two variations of each of two realizations, varying the number $\widehat{n}$ of robots in $\widehat{S}$. The results are averaged across 15 trials for each, with the motions in S varying randomly across the trials. The figure also shows the speedup in latency of these approaches compared to the original illusion from our prior work [16].

The ensuing sections establish some conditions on the existence of multi-illusions and characterize their performance under those conditions.

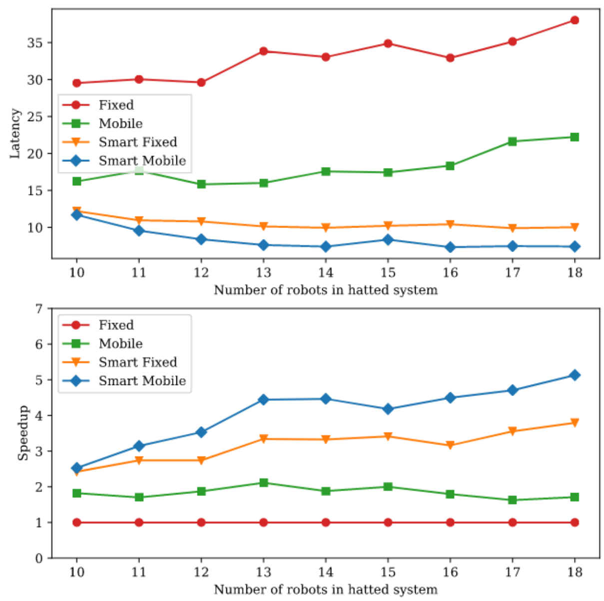

Fig. 6. [top] Latency comparison for four realizations of an illusion in the Robotarium. *Smart Fixed* and *Smart Mobile* are the same as in Figure 5. The non-Smart versions use a more naïve method to plan motions for the complicit robots. [bottom] Speedup for the same scenario, computed relative to the Fixed realization.

III. Sufficient Conditions for Multi-illusions

This section explores of the possibilities for multi-illusions by examining conditions under which pairs of individual illusions can be combined into a single multi-illusion. Those conditions will depend on a certain type of equivalence relation on the states, defined next.

Definition 5. *For a given deterministic multi-robot transition system* $(n, X, U, f, d, Y, h, x_0)$ *and equivalence relation on observations* $\doteqdot \subseteq Y \times Y$, *a relation on its states* $\sim \subseteq X \times X$ *is an* observation-simulation relation *if* $\sim$ *is an equivalence relation, and for any* $x_1, x_2 \in X$ *for which* $x_1 \sim x_2$, *we have*

1) $h(x_1) \doteqdot h(x_2)$, *and*
2) $\forall u_1 \in U, \exists u_2 \in U$ *such that* $f(x_1, u_1) \sim f(x_2, u_2)$.

When the observation equivalence $\doteqdot$ is taken as equality, an observation-simulation equivalence relation has much of the character of a bisimulation relation [19]. Definition 5 generalizes the bisimulation notion, which would have $u_2 = u_1$ in the second condition.

Notice that, given an observation-simulation $\sim_a$ using $\doteqdot_a$ on Y, if $\doteqdot_a$ is a finer relation on observations than $\doteqdot_b$, then $\sim_a$ is an observation-simulation using $\doteqdot_b$ as well.

Definition 6. *For a given deterministic multi-robot transition system* $(n, X, U, f, d, Y, h, x_0)$, *an observation-simulation relation* $\sim \subseteq X \times X$, *the* quotient graph *is a directed graph whose vertex set is the set of equivalence classes of* $\sim$. *Each such class can be represented by an arbitrary representative:* $[x] := \{x' \in X \mid x \sim x'\}$. *Between any two vertices* $[x_1]$ *and* $[x_2]$, *an edge* $[x_1] \to [x_2]$ *exists if and only if there exists some action* $u \in U$ *for which* $[x_2] = [f(x_1, u)]$.

Note that the quotient graph may not necessarily be a finite graph (if the state space X is infinite) nor even locally finite (if the action space U is infinite).

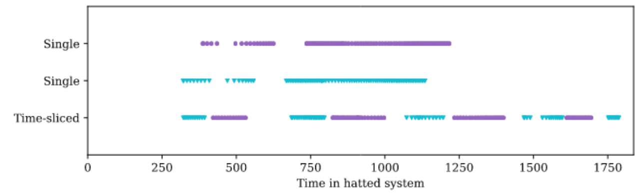

Fig. 7. Execution progress of two single Smart Mobile illusions executed in two separate simulated Robotaria, compared against progress of the time-sliced multi-illusion as constructed for Theorem 1.

Example 6. Referring again to Example 1, consider the system $S_{n,v_{\min},v_{\max}}$. Take $\doteqdot$ to be the identity relation and define $\sim$ such that for any pair of states x_1 and x_2 we have $x_1 \sim x_2$ if and only if $h(x_1) = h(x_2)$, which occurs precisely when the inter-vehicle spacing is identical between the two states, regardless of those states' absolute positions along the roadway. Notice that the construction does indeed produce an observation-simulation relation. Moreover, the quotient graph in this case is strongly connected, since the robots can, in some finite series of stages, adjust positions relative to each other arbitrarily. Note in particular that this remains true even if $v_{\min} > 0$ — because $\sim$ ensures that only the relative positions of the robots are relevant, paths exist between any pair of equivalence classes, without the need for any robot to move backward in absolute position.

A. Time-sliced multi-illusions

We can now use the idea of the quotient graph to give conditions under which a sort of 'time slicing' illusion can be guaranteed to exist.

Theorem 1. *Let* S_1, S_2, *and* $\widehat{S}$ *denote three systems, and suppose there exist illusions of* S_1 *by* $\widehat{S}$ *and of* S_2 *by* $\widehat{S}$. *If there exists an observation-simulation relation for* $\widehat{S}$ *for which the quotient graph is strongly connected, then there exists a multi-illusion of both* S_1 *and* S_2 *by* $\widehat{S}$.

Proof roadmap: Under the presumed conditions, portions of the execution of the individual illusions for S_1 and S_2 may be interleaved. Let $\tau \in \mathbb{Z}^+$ denote an arbitrary positive time slice, to be interpreted as a number of stages elapsed in $\widehat{S}$. We can construct a realization of the multi-illusion for both S_1 and S_2 that cycles through four phases: (1) an execution of the illusion for S_1 for τ stages, (2) an interregnum of actions that transition to a state suitable to begin or continue the S_2 illusion, (3) an execution of the illusion for S_2 for τ stages, and (4) an interregnum of actions that transition to a state suitable to continue the S_1 illusion. The key observation is that the strong connectedness of the quotient graph ensures action sequences always exist that can achieve phases (2) and (4).

Example 7. Figure 7 shows a progress plot for an example in which the construction for Theorem 1 is utilized to time share two of the sorts of illusions introduced in Example 4 within a single simulated Robotarium.

Corollary 1. *Under the same premises as Theorem 1, suppose the construction in its proof is used to provide a multi-illusion for $\mathcal{S} = \{S_1, S_2\}$ by $\widehat{S}$. For $j \in \{1, 2\}$, let $L_j^{\mathcal{S}}(\widehat{t}_{\widehat{k}})$ denote the latency for S_j in this combined illusion. Then*

$$L_j^{\mathcal{S}}(\widehat{t}_{\widehat{k}}) \leq \frac{\widehat{t}_{\widehat{k}}}{\widehat{t}_{\beta(\widehat{k})}} L_j\left(\widehat{t}_{\beta(\widehat{k})}\right)$$

in which $\beta(\widehat{k}) = \tau \left\lfloor \frac{\widehat{k}}{2(\tau + \operatorname{diam}(G))} \right\rfloor$.

Proof roadmap: Observe that after $\widehat{k}$ stages, the number complete cycles of the four phases of the constructed illusion is at least $\lfloor \widehat{k}/(2(\tau + \operatorname{diam}(G)) \rfloor$, since each instance of phases (1) or (3) consumes precisely τ stages, and each instance of phases (2) or (4) consumes at most $\operatorname{diam}(G)$ stages. The illusion for S_j executes for τ stages in each of these cycles, amounting to $\tau \left\lfloor \widehat{k}/(2(\tau + \operatorname{diam}(G)) \right\rfloor$ stages in total for S_j. Let $k_j^{\mathrm{m}} = \max\{\kappa \in \mathbb{Z}^+ \mid z_j^{S_j}(\kappa) \leq \beta(\widehat{k})\}$ and observe that

$$L_j^{\mathcal{S}}(\widehat{t}_{\widehat{k}}) = \frac{\widehat{t}_{\widehat{k}}}{t_{k_j^{\mathcal{S}}}} \leq \frac{\widehat{t}_{\widehat{k}}}{t_{k_j^{\mathrm{m}}}} = \frac{\widehat{t}_{\widehat{k}}}{\widehat{t}_{\beta(\widehat{k})}} L_j\left(\widehat{t}_{\beta(\widehat{k})}\right),$$

applying Definition 3 in the final step to complete the proof.

B. Leveraging role reassignments

Moving now beyond mere time slicing, note that in Definition 2, for the $(\star)$ equation corresponding to system S_j, the i^{th} element may be re-mapped (via the ρ_j function) to be some arbitrary element observed in $\widehat{S}$. Since there is no requirement that ρ_j be injective, duplicate values are entirely superfluous. Thus, though prior discussion of specific cases, as in Example 6, has been for an observation-simulation relation $\sim$ using $\doteqdot$ as equality on elements of Y, equality is stronger than is strictly necessary for illusions. This motivates the following function, π, which repackages data from an n-vector into a set:

$$\begin{array}{ccc} \left(y^{(1)}, \ldots, y^{(n)}\right) & \xmapsto{\pi} & \left\{y^{(1)}, \ldots, y^{(n)}\right\}. \\ \pitchfork & & \cap\!\!\!\!\mid \\ Y^{(1)} \times \cdots \times Y^{(n)} & & \bigcup_i Y^{(i)} \end{array}$$

And, hence, let equivalence relation $\stackrel{s}{=} \subseteq Y \times Y$ be defined as $y' \stackrel{s}{=} y$ if and only if $\pi(y') = \pi(y)$.

Using $\stackrel{s}{=}$ on Y for some deterministic multi-robot transition system $(n, X, U, f, d, Y, h, x_0)$, if we are given some observation-simulation equivalence relation $\underset{s}{\sim}$, we can construct the quotient graph $G_{\underset{s}{\sim}}$. (In this case we may think of the vertices of the graph as labelled by sets.) Theorem 1 holds for $\underset{s}{\sim}$, but observe that the time-slicing used in the construction as part of the proof argument only modifies the z_j functions; essentially it is purely a strict interleaving the illusions for S_1 and S_2. Under $\underset{s}{\sim}$, one might do better: rather than reaching a state $x \in [x]_\sim$ where the ρ_j for the multi-illusion is identical with the original illusion for S_j, instead one reaches $x' \in [x]_{\underset{s}{\sim}}$ and uses ρ_j' constructed to map to *some* robot producing the required observation (different x' will have different ρ_j').

Further, because the identity relation is finer than $\stackrel{s}{=}$, i.e., $\mathrm{id}_Y \subseteq \stackrel{s}{=}$, every observation-simulation relation $\sim$ via the former, is one for $\stackrel{s}{=}$ too. We would expect there to be additional observation-simulation relations $\underset{s}{\sim}$ for $\stackrel{s}{=}$, generally. (This needn't be so always, however, e.g., if every $Y^{(i)}$ and $Y^{(j)}$, with $i \neq j$ are disjoint.) When some $\underset{s}{\sim} \supset \sim$ this has practical advantages because the multi-illusion can do less work transitioning on the graph (for interregnum steps) because the set of acceptable targets has grown. The added flexibility to give more efficient illusions manifests in the bounds on latency because, one might expect generally that $\operatorname{diam}(G_{\underset{s}{\sim}}) \ll \operatorname{diam}(G)$.

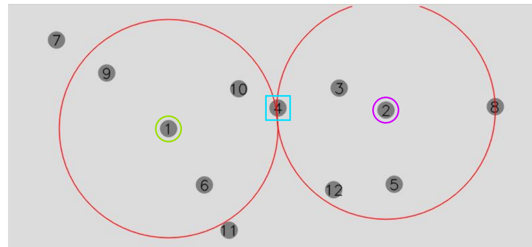

Fig. 8. A multi-illusion of two systems executing simultaneously in a single host. The cyan-marked robot represents a sort as shared resource as it plays different roles at the same in both illusions.

C. Multi-illusions via $\limsup$

To give a different, tighter characterization for sufficiency than the preceding, two additional definitions are needed.

Definition 7. *Given a multi-robot transition system* $S = (n, X, U, f, d, Y, h, x_0)$, *the* observation footprint *of* S *is*

$$F(S) = \underbrace{\bigcup_{\pi^{(1)}, \ldots, \pi^{(n)}}}_{\textit{All policies}} \bigcup_{k \geq 0} \{\, \pi(h(x_k)) \,\}.$$

Intuitively, the observation footprint is a collection all the sets of observations that might be seen by the system at any particular point in time.

Definition 8. *A* walk *on graph* $G = (V, E)$ *starting at* $[x_0] \in V$ *is a function* $w : \mathbb{Z}^+ \to V$, *with* $w(0) = [x_0]$ *and such that, for* $n \geq 1$, *there is an edge* $([x_{n-1}] \to [x_n])$ *in* E.

Now, denoting the powerset by $\mathcal{P}(\cdot)$, we state the theorem.

Theorem 2. *There exists a multi-illusion for* $\mathcal{S} = S_1, \ldots, S_d$, *on* $\widehat{S}$, *if there exists an observation-simulation relation for* $\widehat{S}$, *which induces a quotient graph* G, *and some walk* w *on* G *exists such that*

$$\bigcup_{j=1\ldots d} F(S_j) \subseteq \limsup_{n \to \infty} \mathcal{P}(\pi(w(n))).$$

Proof roadmap: Given such a walk w, one can ensure progress is made in any of the d systems. In S_i, any set of observations you wish to concoct appears in $F(S_i)$. The condition, thus, ensures that, at any point in time, a set containing it will appear in finite time. Constructing the illusion requires choosing z_i so it progresses when such a set appears, and having ρ_i unpack the subset of observations needed and dispatching them to the intended recipients.

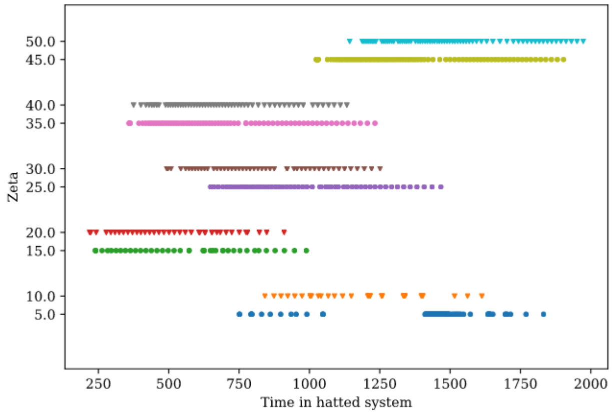

Fig. 9. Progress chart for each of the 10 trials in the parameter tuning experiment, executed in pairs across 5 runs of the host system.

IV. Seemingly paradoxical speed-up

Our investigation of concepts like latency and speedup has been inspired by the longstanding application of those concepts to understand computational processes. Though much has followed by close analogy to existing conceptions, in other ways the present setting holds some surprises. For example, in some purely computational contexts, there are theoretical limits on the speedup that may be obtained.

Example 8. Suppose we want to execute d copies of a system S, i.e., $\mathcal{S} = S_0, S_0, \ldots, S_0 = S_0^d$, within test-bed $\widehat{S}$. It is tempting to define a multi-illusion in which these copies are executed sequentially. But $\widehat{S}$ can illusion all of $\mathcal{S}$ in a single swoop by bookkeeping; each requirement in Definition 2 is satisfied d times by mere copying. The speedup between between the two multi-illusions is $\mathbb{S} = d$. Thus, by selecting a large d, we can obtain arbitrarily large speedup.

Since Example 8 has the appearance of chicanery (and seems perhaps to connote some definitional flaw) it is worth examining what, precisely, accounts for this windfall. In essence, the observations obtained by robots in $\widehat{S}$ may be used productively for any system(s) S_i that make progress after seeing a subset of those observations. Further, because there is indirection between state configurations and observations, a single robot in $\widehat{S}$ may simultaneously service multiple illusions. Indeed, we witness this in practice, as noted below.

V. Case Study: Parameter optimization

To bring our analysis full circle back to the promise of efficient experimentation in multi-robot test-beds, this section describes a simple example realization of that idea. To maintain focus on the novel aspects of the setting, let us consider the straightforward and well-understood task of optimizing a parameter in an artificial potential field controller.

Recall that this sort of controller uses parameters η and ζ to weight the forces pushing the robot away from obstacles and toward the goal, respectively. One may desire to conduct

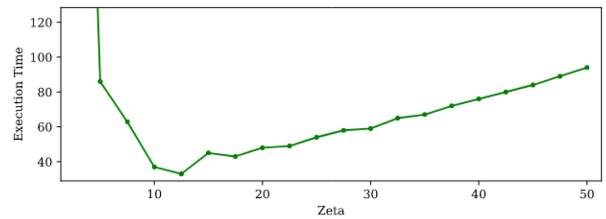

Fig. 10. Results for the parameter optimization experiment, computed via parallel multi-illusions.

an experiment to determine the best values for η and ζ to balance these forces and to achieve the fastest travel time to a goal. In such a scenario, we might hold the sum $\eta + \zeta$ fixed, and conduct a series of trials, varying ζ and measuring the robot's travel time for each trial.

We conducted this experiment via multi-illusion in a simulated Robotarium. The specific form of the multi-illusion divides the workspace into left and right halves, and executes a distinct trial on each side. However, the two sides are sufficiently close together that some robots representing obstacles sometimes play different roles at the same time in both ongoing trials. Figure 8 shows an example: Two illusioned robots (#1 highlighted in green, #2 highlighted in purple) are observing what they believe to be an obstacle, as concocted by robot #4 (in the cyan square) for both of them simultaneously. This resource sharing enables the illusions to proceed more rapidly than might have been expected. Figure 9 shows the overall progress of the process for an experiment with 10 trials, executed in simultaneous pairs. Of particular interest there is the high density of simultaneous execution. We also conducted a larger scale version of this experiment, using more finely-grained selection of ζ values and conducting 10 trials for each value to account for randomness in obstacle placement. Figure 10 shows the results, in which potential field controller parameter $\zeta = 12.5$ gave the best performance.

VI. Conclusion

Identifying the concept multi-experiments as a topic of study, this paper lays the initial foundations on which further development should build. It provides definitions and conditions for the existence of multiplexing, as well as results of a more quantitative nature in the form of performance bounds. Further, the paper has presented demonstrations of the feasibility of the concept in practice. Though the implemented examples are simple, they already showcase the utility of the approach, for instance, in parameter tuning. Perhaps most significant is that different multi-illusions can be seen to have vastly different performance characteristics. Further work would need to explore these aspects more completely, potentially offering tighter bounds on the basis of opportunities for optimization, and expand the theory beyond the basic deterministic model.

References

[1] Serge Kernbach. Swarmrobot.org—Open-hardware Microrobotic Project for Large-scale Artificial Swarms. *arXiv preprint arXiv:1110.5762*, 2011.

[2] James McLurkin, Adam McMullen, Nick Robbins, Golnaz Habibi, Aaron Becker, Alvin Chou, Hao Li, Meagan John, Nnena Okeke, Joshua Rykowski, et al. A robot system design for low-cost multi-robot manipulation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 912–918, 2014.

[3] Michael Rubenstein, Christian Ahler, and Radhika Nagpal. Kilobot: A low cost scalable robot system for collective behaviors. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 3293–3298, 2012.

[4] Gabriel T. Sibley, Mohammad H. Rahimi, and Gaurav S. Sukhatme. Robomote: A tiny mobile robot platform for large-scale ad-hoc sensor networks. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1143–1148, 2002.

[5] Sean Wilson, Ruben Gameros, Michael Sheely, Matthew Lin, Kathryn Dover, Robert Gevorkyan, Matt Haberland, Andrea Bertozzi, and Spring Berman. Pheeno, a versatile swarm robotic research and education platform. *IEEE Robotics and Automation Letters*, 1(2):884–891, 2016.

[6] Liam Paull, Jacopo Tani, Heejin Ahn, Javier Alonso-Mora, Luca Carlone, Michal Cap, Yu Fan Chen, Changhyun Choi, Jeff Dusek, Yajun Fang, et al. Duckietown: an open, inexpensive and flexible platform for autonomy education and research. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1497–1504, 2017.

[7] Adrián Jiménez-González, Jose Ramiro Martinez-De Dios, and Anibal Ollero. Testbeds for Ubiquitous Robotics: A Survey. *Robotics and Autonomous Systems*, 61(12):1487–1501, December 2013.

[8] Dennis Larkin, Matthew Michini, Alexandra Abad, Stephanie Teleski, and M Ani Hsieh. Design of the Multi-Robot Coherent Structure Testbed (mCoSTe) for Distributed Tracking of Geophysical Fluid Dynamic. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, volume 46377. American Society of Mechanical Engineers, 2014.

[9] Ellis King, Yoshi Kuwata, Mehdi Alighanbari, Luca Bertuccelli, and Jonathan How. Coordination and control experiments on a multi-vehicle testbed. In *Proceedings of the IEEE American Control Conference (ACC)*, volume 6, pages 5315–5320, 2004.

[10] Logan E. Beaver, Behdad Chalaki, A. M. Ishtiaque Mahbub, Liuhui Zhao, Ray Zayas, and Andreas A. Malikopoulos. Demonstration of a time-efficient mobility system using a scaled smart city. *Vehicle System Dynamics*, 58(5):787–804, 2020.

[11] Nicholas Hyldmar, Yijun He, and Amanda Prorok. A Fleet of Miniature Cars for Experiments in Cooperative Driving. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 3238–3244, 2019.

[12] Daniel Pickem, Myron Lee, and Magnus Egerstedt. The GRITSBot in its natural habitat: A multi-robot testbed. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 4062–4067, 2015.

[13] Sean Wilson, Paul Glotfelter, Li Wang, Siddharth Mayya, Gennaro Notomista, Mark Mote, and Magnus Egerstedt. The Robotarium: Globally Impactful Opportunities, Challenges, and Lessons Learned in Remote-Access, Distributed Control of Multirobot Systems. *IEEE Control Systems Magazine*, 40:26–44, 2020.

[14] Daniel Pickem, Paul Glotfelter, Li Wang, Mark Mote, Aaron Ames, Eric Feron, and Magnus Egerstedt. The robotarium: A remotely accessible swarm robotics research testbed. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1699–1706, 2017.

[15] Markku Suomalainen, Alexandra Q Nilles, and Steven M LaValle. Virtual reality for robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.

[16] Dylan A. Shell and Jason M. O'Kane. Reality as a simulation of reality: robot illusions, fundamental limits, and a physical demonstration. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 10327–10334, 2020.

[17] Klaudia Witte, Stefanie Gierszewski, and Laura Chouinard-Thuly. (Guest Editorial) Virtual is the new reality. *Current Zoology—Special Issue: Computer Animations and Virtual Reality in Animal Behavior Research*, 63:1–4, February 2017.

[18] Ronald C. Arkin. *Behavior-Based Robotics*. MIT Press, Cambridge, MA, U.S.A., 1998.

[19] Robin Milner. A calculus of communicating systems. *Lecture Notes in Comput. Sci. 92*, 1980.