PNAS RESEARCH ARTICLE | APPLIED PHYSICAL SCIENCES

Adaptive nonequilibrium design of actin-based metamaterials: Fundamental and practical limits of control

Shriram Chennakesavalu[a], Sreekanth K. Manikandan[a], Frank Hu[a], and Grant M. Rotskoff[a,b,1]

Edited by Monica Olvera de la Cruz, Northwestern University, Evanston, IL; received July 5, 2023; accepted November 13, 2023

The adaptive and surprising emergent properties of biological materials self-assembled in far-from-equilibrium environments serve as an inspiration for efforts to design nanomaterials. In particular, controlling the conditions of self-assembly can modulate material properties, but there is no systematic understanding of either how to parameterize external control or how controllable a given material can be. Here, we demonstrate that branched actin networks can be encoded with metamaterial properties by dynamically controlling the applied force under which they grow and that the protocols can be selected using multi-task reinforcement learning. These actin networks have tunable responses over a large dynamic range depending on the chosen external protocol, providing a pathway to encoding "memory" within these structures. Interestingly, we obtain a bound that relates the dissipation rate and the rate of "encoding" that gives insight into the constraints on control—both physical and information theoretical. Taken together, these results emphasize the utility and necessity of nonequilibrium control for designing self-assembled nanostructures.

actin | reinforcement learning | nonequilibrium statistical mechanics | fluctuation theorem

Significance

Building new materials from biological components with exotic and controllable properties requires strategies to assemble these components into nontrivial structures. Cytoskeletal proteins—which transduce force—give structure to a cell, help it move, and enable it to adapt dynamically to growth conditions, resulting in materials with a long-lived memory of how they were built. Based on recent experiments, we develop a minimal model of actin assembly under load and we show how reinforcement learning techniques can be employed to control the growth of actin networks to make structures that have properties distinct from those that arise in the biological context. We also show that this controlled self-organization is constrained by a fundamental thermodynamic limit on the rate of controllability.

Biomolecular materials spontaneously self-assemble in fluctuating, nonequilibrium environments. Despite the complex environment, many cellular assemblies possess a remarkable capacity for adaptation and dynamism (1–5). In contrast, our ability to design materials with nanoscale components remains primitive in comparison with the intricate, hierarchical structures that reliably assemble in living systems (6). The strategies employed by biological self-assembly thus offer a unique lens into the limits of control for synthetic efforts to engineer nanoscale metamaterials.

Cytoskeletal networks are perhaps the most illustrative example of the dramatic effects of assembly dynamics on resulting material properties (7, 8). The key structural proteins that form this scaffolding of the cell must rapidly rearrange to drive fundamental processes such as motility (9) and signaling (3, 10). The remarkable material properties and dynamical assembly of branched actin networks, which self-organize through the interplay between numerous proteins and external forces from the environment, have attracted significant experimental and theoretical scrutiny. The components of these dynamic networks have been purified, and reconstituted systems have been studied extensively (8, 11–14), yielding detailed insight into their nonequilibrium dynamics.

Atomistic simulations and experimental studies have provided insight into the molecular mechanisms of actin polymerization and allostery (15), branching (16), and plasticity (17). In silico studies of mesoscale material properties of cytoskeletal networks have, on the other hand, relied heavily on minimal models (8, 18, 19), that represent the network of actin filaments as elastic networks with filaments that resist bending and stretching. For example, randomly assembled networks have been shown to recapitulate interaction with the membrane (20) and myosin-induced contraction (18, 19). Similarly, interesting dynamical behaviors like treadmilling (21) as well as aster formation and sorting in actin–myosin networks (22, 23) have been captured accurately with simple models. Higher-resolution reaction–diffusion models have also been used to investigate the thermodynamics of actin network growth and phenomena such as cytoquakes (24, 25). Floyd et al. (26), in particular, introduced a systematic scheme for computing dissipation in actomyosin networks and showed that chemical energy is more efficiently stored in dense networks. The aforementioned works show the utility and predictive power of a minimal modeling approach, and we follow suit. We develop a model that closely resembles the experimental geometry of a recent study that images branched actin network growth under load (11, 14).

We show here that the model we have developed captures the essential behavior of actin network growth under varying external loads and, additionally, we show how to control that behavior to tune the response properties of the resulting material. We demonstrate,

Author affiliations: [a]Department of Chemistry, Stanford University, Stanford, CA 94305; and [b]Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA 94305

Author contributions: S.C., S.K.M., and G.M.R. designed research; S.C., S.K.M., F.H., and G.M.R. performed research; S.C., S.K.M., F.H., and G.M.R. analyzed data; and S.C., S.K.M., and G.M.R. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence may be addressed. Email: rotskoff@stanford.edu.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2310238121/-/DCSupplemental.

Published February 15, 2024.

in particular, that an external feedback protocol that dynamically modulates the growth conditions encodes a nonequilibrium memory of the growth conditions in the material. The deviation from a homogeneous density profile of the molecular components leads to distinctive metamaterial properties (5), but driving the system externally to push it toward a target structure incurs a dissipative encoding cost, which, at its core, is a consequence of the second law of thermodynamics; we compute and analyze this cost below. Using reinforcement learning (RL), we optimize external, nonequilibrium protocols that drive the resulting networks into regimes inaccessible to uncontrolled assemblies. We further show that the resulting networks have elastic coefficients that depend strongly on the applied force.

Encoding Material Properties with Nonequilibrium Growth Dynamics

Mechanical response in soft materials is dictated by both intrinsic material properties of the constituents and also their spatial organization (27, 28). When the individual components of a structure can be manipulated in a spatially localized fashion, exquisitely precise control of the resulting material is sometimes possible (27, 29). However, in many materials, we can only reconfigure a fixed set of components into distinct topologies. This motivates an exploration of the possibility of controlling mechanical response without manipulating the components themselves or their intermolecular interaction energy. For example, the reorganization of fixed components in a network has been used to control global response by systematically removing a small set of bonds from disordered networks of springs (5, 30).

Making metamaterials by manipulating a few degrees of freedom in a complex network indeed demonstrates the profound sensitivity of mechanical properties on spatial organization. However, in nanoscale systems self-assembled by stochastic dynamics, surgical reconfiguration is not possible experimentally. Nevertheless, many biomolecular materials have adaptive properties that evince memory of their growth conditions (4, 5, 31, 32), which are ultimately nonequilibrium, kinetic effects. Perhaps the most direct example of this behavior is the force-induced stress stiffening in branched actin networks (12, 13, 24, 33) which can be reversibly saturated to the point of softening (34). When grown under high load forces, actin networks become more rigid than networks assembled in the absence of an applied force, despite consisting of exactly the same protein components.

The metamaterials that assemble under the external control described below retain information about their growth process. Many works investigating memory in nonequilibrium systems assume the existence of a Hamiltonian with multiple metastable states and instead examine the retrieval of specific states in this memory landscape (35, 36). However, the physical processes that use dissipative driving to encode metastability have not been analyzed in detail. Here, we examine the thermodynamic constraints on this class of dissipative processes and generalize this notion to non-Hamiltonian systems. This analysis emphasizes an intuitive trade-off between the extent of a perturbation, which we quantify through a measure of the deviation of a configurational distribution away from its steady state, and dissipation as the system undergoes controlled dynamics.

Consider an encoding process that transforms a system from some initial steady-state ρ_{ss} at $t = 0$ to some new distribution ρ_τ at $t = \tau$, driven by a nonequilibrium feedback-dependent protocol. Using the detailed information fluctuation theorem (37–39), we can relate the log ratio of the path measures of a controlled trajectory and its time-reversal

$$\log \frac{P[\boldsymbol{x}_t, \boldsymbol{m}_t]}{\tilde{P}[\tilde{\boldsymbol{x}}_t, \tilde{\boldsymbol{m}}_t]} = -\beta Q(\boldsymbol{x}_t, \boldsymbol{m}_t) + \Delta S_{\text{sys}} + \mathcal{I}(\boldsymbol{x}_t; \boldsymbol{m}_t), \qquad [1]$$

where $\mathcal{I}(\boldsymbol{x}_t; \boldsymbol{m}_t)$ denotes the trajectory-wise information between the trajectory $\boldsymbol{x}_t$ and the measurement sequence $\boldsymbol{m}_t$, Q denotes the heat along the trajectory, ΔS_{sys} denotes the change in system entropy from ρ_{ss} to ρ_τ, and $\tilde{\cdot}$ denotes a time-reversed quantity. While this result was originally discussed in the context of transformations starting from an equilibrium distribution (38), this assumption is not needed for Eq. **1** (39).

Furthermore, averaging over trajectories, using the second law of thermodynamics, and by employing the data-processing inequality (*SI Appendix*), we get

$$D_{\text{KL}}(\rho_{ss}||\rho_\tau) \leq -\beta \langle Q \rangle + \Delta S_{\text{sys}} + I(\boldsymbol{x}_t; \boldsymbol{m}_t), \qquad [2]$$

where $I[\boldsymbol{x}_t; \boldsymbol{m}_t]$ is the mutual information between $\boldsymbol{x}_t$ and $\boldsymbol{m}_t$ and the brackets $\langle \cdot \rangle$ denotes an average over trajectories. We note here that the information term vanishes when there is no feedback in the protocol. The KL divergence between the initial stationary distribution, ρ_{ss}, and the distribution resulting from the controlled dynamics, ρ_τ, directly measures the deviation from stationarity and hence provides a useful quantification of how much the configurational distribution has been altered by the external protocol.

Interestingly, the rate of memory encoding satisfies a universal "speed-limit." Because the total heat flow out of the system is time extensive, it is more useful to consider the time derivative of Eq. **2**,

$$\partial_t D_{\text{KL}}(\rho_{ss}||\rho_\tau) \leq \partial_t S_{\text{sys}}(\rho_\tau) + \langle \sigma_{\text{env}} \rangle + \dot{I}(\boldsymbol{x}_t; \boldsymbol{m}_t), \qquad [3]$$

where σ_{env} is the rate of entropy production in the medium. In this work, we do not attempt to minimize the rate of dissipation in the medium, meaning that this upper bound will likely be quantitatively uninformative, though the terms in this bound do provide qualitative insight into the relationship between entropy production and structural changes in the networks (*SI Appendix*, Fig. S6). Minimizing dissipation in nonequilibrium transformations is an active area of research (40–43), and we anticipate that exploring the bound Eq. **2** in the context of minimum dissipation protocols will prove a fruitful setting to explore this result.

The interpretation of this speed limit Eq. **3** is straightforward: The rate at which the system can be driven to deviate from its steady-state behavior requires either a positive rate of dissipation of heat to the reservoir or a positive quantity of information must be extracted from the measurement sequence at each point in time. Bounds similar to Eq. **3** have appeared without the feedback term to quantify the cost to maintain a steady state (44, 45) and as a constraint on number fluctuations in self-assembled lattice gases (46). In the special case of an equilibrium target steady-state, this bound can be saturated because there is no steady-state entropy production; note that if the target process is dissipative in the steady state, this upper bound may not be tight. While this upper bound may provide only a weak constraint in some cases, we observe that changes in the quantities that make up the bound are still physically and mechanistically informative, as discussed immediately below.

The entropy production rate, information acquisition rate, and the time derivative of the KL-divergence for actin networks grown under feedback protocols are plotted as a function of protocol duration in *SI Appendix*, Figs. S5–S7. We observe that growing networks with metamaterial-like properties often leads to dramatic changes in the rate of dissipation. The layered networks interestingly compensate changes in the dissipation

rate with an opposite change in the information acquisition rate; while the rate of information acquisition is an order of magnitude lower than the steady-state dissipation rates, both are clearly strongly correlated with structural changes in the network. Finally, as mentioned above, we note that the rate of change of KL divergence is several orders of magnitude below the sum of the dissipation rate and the information acquisition rate demonstrating that our control framework exists in a regime far from saturation of the speed limit. This last point is not a surprising one; our protocol optimization strategy does not explicitly account for any thermodynamic costs, and furthermore, we are targeting highly dissipative nonequilibrium steady states. This is similar to what has been observed with the thermodynamic uncertainty relations (47, 48), where current fluctuations often provide only a weak upper bound on dissipation rates when there is not tight coupling between the observed current and dissipation (49). Strategies for determining protocols that are minimally dissipative (42) could serve as a useful framework for learning protocols that saturate the bound in Eq. **3**.

The general thermodynamic constraint Eq. **3** implies that a finite dissipation rate is required to encode new properties in a material. Indeed, the large range of elasticities observed in branched actin networks that result only from changing the growth conditions raises an important general question: How tunable is response with a given material assuming that a dissipative, feedback protocol is used to control assembly? Here, we seek to answer this question using a minimal model of branched actin networks, described in detail in the following sections. This model provides an attractive platform for our investigation because this system can be reconstituted using purified proteins and the growth conditions can be directly manipulated using the cantilever arm of an atomic force microscope, as depicted in Fig. 1. Our model captures the essential physical features of recent experiments reported in ref. 14 and, because they work with a fixed set of pure protein components, ambiguities related to cellular regulation of protein expression and other nontrivial side effects of biology can be systematically excluded, allowing us to focus on the underlying physical mechanisms of control.

Obtaining a metamaterial by modulating growth conditions requires external control, and the primary experimentally accessible control variable in this system is the applied force as a function of time. Throughout, we consider a setting where the controller interacts with the material, allowing both measurement and feedback, which is a feature of the biological systems and could also be tractable experimentally with appropriate microscopy. Directly tuning elastic coefficients with a feedback protocol is not straightforward because this global property depends on the network structure in a complicated and nonlinear fashion. External growth forces do, however, directly impact the local density at the growth front, which in turn is strongly connected to the response. Throughout, we study the emergent properties of a fixed, final network at the end of its growth trajectory.

Minimal Model of Actin Growth under External Load

A Brownian ratchet has long been speculated as the mechanism through which external force impacts growth in actin networks (21, 50–53). The experiments of Li et al. (14) lend additional support to the Brownian ratchet model by imaging components of the assembly during growth. This set of experiments uses the cantilever arm of an atomic force microscope to precisely apply force to the network in the direction of its growth. A similar experimental setup was previously used to elucidate stress-stiffening and stress-softening in networks growing under large loads and stress-dependent growth dynamics (11, 12, 34). These experiments provide fundamental insight into the nature of branched actin network assembly, and in particular, they highlight the delicate balance between polymerization, capping, and additional branching near nucleation-promoting factors. Additionally, these reconstitution experiments highlight the strong dependence of material properties on growth conditions, even in the absence of biological stimuli.

Because these experiments contain precisely controlled sets of ingredients, we sought to assess if the experimental findings were consistent with a minimalist model of polymerization against a ratchet-like load. The model we develop is partially inspired by and shares many features with those in refs. 54–56. These two-dimensional (2D) models provide useful qualitative insight into the dynamics of self-organization in branched actin networks in the lamellipodium. The network growth involves the interplay between ATP-dependent polymerization, Arp2/3-mediated branching of the network, and quasi-irreversible capping of the growing barbed-ends of individual filaments. These reactions all have experimentally measured rates, and the physical properties of individual actin filaments have been extensively characterized at the single molecule level (*SI Appendix*, Table S1). Remarkably, a Gillespie model that incorporates only these experimentally known parameters predicts emergent, network-level properties that are consistent with experimental results. While the experimental analogue of the actin networks we consider are three-dimensional, we find, as shown below, that a 2D model is sufficient to explain the experimentally observed behavior.

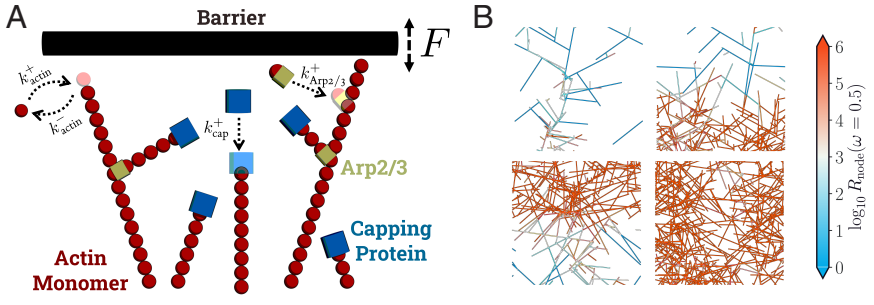

Fig. 1. External control of actin network enables design of actin-based metamaterials. (*A*) A schematic overview of branched actin networks growing against a barrier, where individual filaments can polymerize, depolymerize, branch (via Arp2/3), or be capped. Network growth can be controlled by modulating load force *F*. (*B*) Actin networks grown with layered density profiles (*Top Right* and *Bottom Left*) have heterogeneous responses, while networks grown with constant density profiles (*Top Left* and *Bottom Right*) have more homogeneous responses.

Our model geometry closely mimics the experimental setup described in ref. 14. As depicted in Fig. 1, actin elongation occurs homogeneously throughout the network with a rate k^+_{actin}. Biologically, actin growth is biased toward the plus end (3), but the experimental setup that we compare most directly with includes only profilin-bound actin which exclusively binds the plus end; that is, no minus-end growth occurs in our model. Polymerization is balanced by a reverse reaction in which actin spontaneously unbinds from existing filaments with a rate k^-_{actin}. The imbalance in the forward and reverse growth rates leads to a steady-state positive growth velocity. For actin and all other components of the network, we fix the chemical potential and assume there is sufficient excess of protein that depletion effects can be neglected.

Near the barrier, nucleation-promoting factors encourage branching of actin filaments via binding of the protein complex Arp2/3 (14). We do not explicitly represent nucleation-promoting factors but capture this proximity effect by allowing Arp2/3 to bind to filaments within $10\ell_{\text{actin}}$ of the barrier, where ℓ_{actin} denotes the diameter of an actin monomer. Interactions between the growing end of an actin filament and nucleation-promoting factors can disrupt branching (14), so we mollify the rate of branching based on the barbed end density near the growth front. Explicitly, the branching propensity is given by

$$n_{\text{branching}} = k^+_{\text{Arp2/3}} n_{\text{Arp2/3 binding sites}} \times n_{\text{eff Arp2/3}}[\text{Arp2/3 Monomers}], \quad [4]$$

where n_{free} is the number of free barbed ends,

$$n_{\text{eff Arp2/3}} = \max\left(0, n_{\text{NPF}} - n_{\text{free}}\right), \quad [5]$$

and n_{NPF} is a fixed constant.

The elongation of an individual actin filament can be halted by the so-called capping protein, which binds free barbed ends with high affinity. When a filament binds capping protein, an event that is considered irreversible in our model, it no longer polymerizes or depolymerizes. The only constraint on capping, which occurs homogeneously throughout the network, is steric hindrance with the barrier.

Indeed, sterics dictate the overall growth trajectory of the network via the constraint imposed by the rigid barrier depicted in Fig. 1. The Brownian ratchet picture of actin growth under load asserts that polymerization spontaneously occurs as the barrier fluctuates away from the growth front. In our model, diffusion of the barrier under an applied load is reflected at the point of maximum height of the network. The transition probability density for such a reflected diffusion can be solved analytically (*SI Appendix*, Eq. **S12**), meaning that we can propagate the position of the barrier simply by sampling conditionally on time after a reaction occurs.

Fig. 2 illustrates the model branched actin network growth under a step-wise increasing load force, which is in reliable qualitative agreement with experimental observations (14). First, under an increasing load, the growth velocity of the actin network decreases. This is primarily a consequence of the increased steric hindrance between the barrier; under higher loads, fluctuations in the barrier that enable polymerization are less likely. Second, the density of the materials near the growth front increases. To enable comparisons with experimental observables, we consider a region of interest within 100 nm of the growth-front, a region comparable to the imaging depth afforded by total internal reflection fluorescence microscopy (57). The reduction in growth

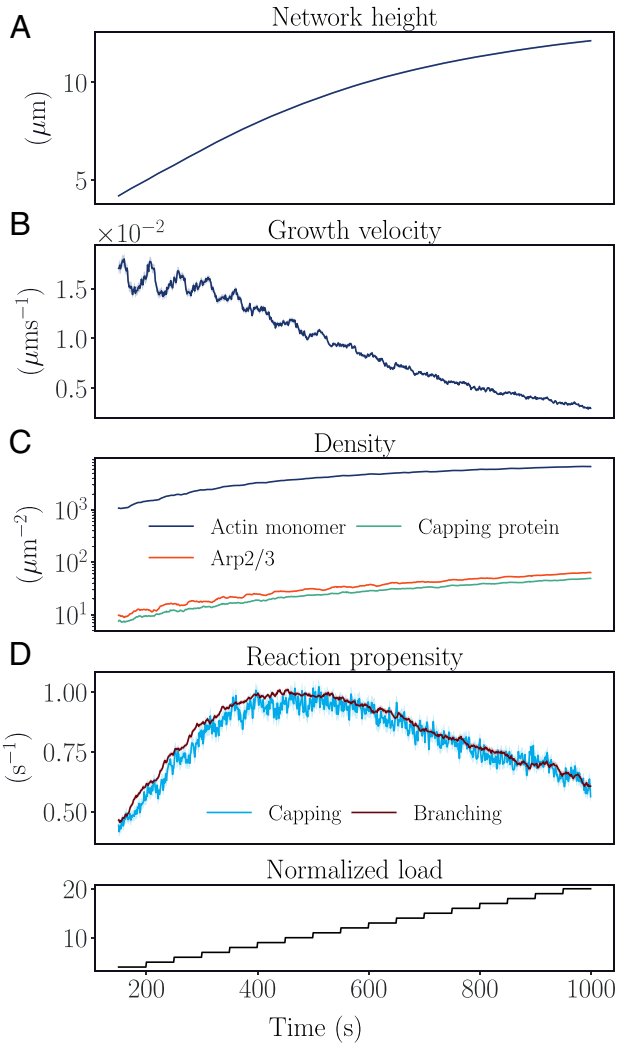

Fig. 2. Average network height, growth velocity, densities of constituent materials, and reaction propensities for 200 actin networks grown under increasing growth loads normalized by f_0 (*SI Appendix*, Table S1). Under increasing loads (*Bottom* panel), (*A*) network height increases with (*B*) a reduced growth velocity, (*C*) density of actin monomers, capping protein, and Arp2/3 within 100-nm increases, and (*D*) reaction propensity of capping and branching decreases but remains balanced.

velocity of the network under an increased load will result in a larger number of uncapped filaments near the growth front of the network, inducing further capping and branching events, resulting in increased densities of actin monomers, capping proteins, and Arp2/3 complexes within this region. Third, we observe that the reaction propensities for branching and capping are balanced across different loads, with both propensities eventually decreasing under an increased load. In this setting, the increased number of free barbed ends interferes with branching events, while the increase in steric hindrance restricts capping of filaments.

Directing Self-Assembly of Actin Structures

The load-induced stiffening by the external load force on a growing actin network suggests the possibility of engineering actin-based metamaterials. The spatial organization of the actin network depends strongly on the load history (12), and furthermore, there is a strong coupling between the response of a network and its spatial organization. Previous work (45) has shown that RL offers a robust, yet straightforward framework

for nonequilibrium control, when the tunable parameters are extrinsic, such as experimental control variables. Following the approach detailed in ref. 45, we devise a multi-task RL strategy (*SI Appendix*) which allows us to spatially tune the network densities and elicit material responses distinct from those available to homogeneous networks.

The response of actin networks grown under a fixed protocol exhibits high variance, limiting protocol-optimization strategies that directly seek to target a particular response. However, determining protocols that control the density of the network allows us to assemble actin networks in a targeted way because the response of a network is tightly coupled to density. In practice, we can only reasonably control the growth front of the actin network as filaments that are not proximal to the barrier are quickly capped, and filament branching requires nucleation promoting factors, which are implicitly present on the barrier. Therefore, we can only regulate branching events on monomers proximal to the barrier. With this in mind, we use RL to design external control protocols that produce specific growth front densities of a network. Here, we define the growth-front density as the density of actin monomers within $10\ell_{\text{actin}}$ of the barrier.

We consider a discrete set of external loads (*SI Appendix*, Table S2) to limit potential experimental challenges of applying load forces with arbitrary resolution. We then train external controllers that can, in a feedback fashion, modulate the external load to achieve a desired density. In Fig. 3, we detail trajectories of protocols to grow layered networks with a low-density, soft,

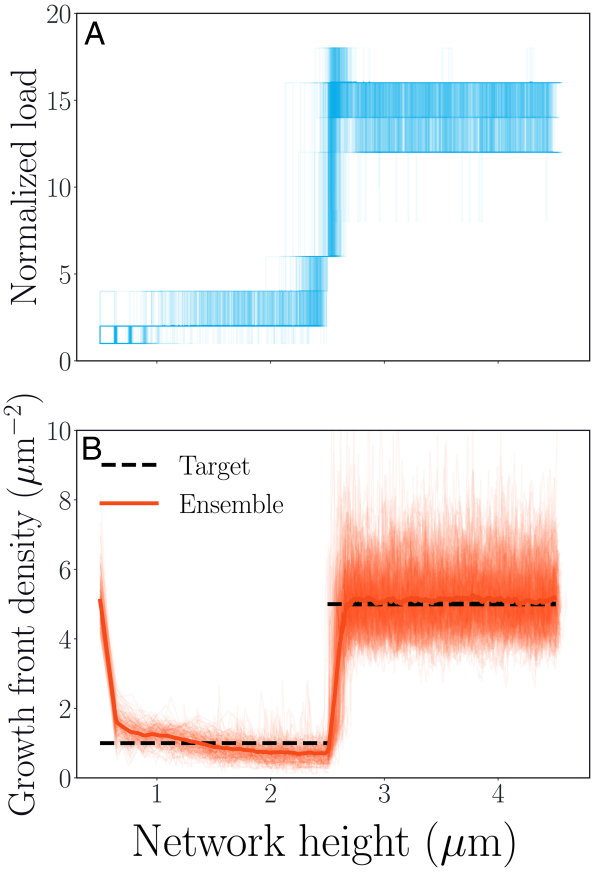

Fig. 3. Protocols and growth front density as a function of time for 200 networks grown to a target layered structure with feedback protocols, where external load is normalized by f_0 (*SI Appendix*, Table S1). The ensemble of state-dependent control protocols (*A*) enable realization of target growth-front density (*B*).

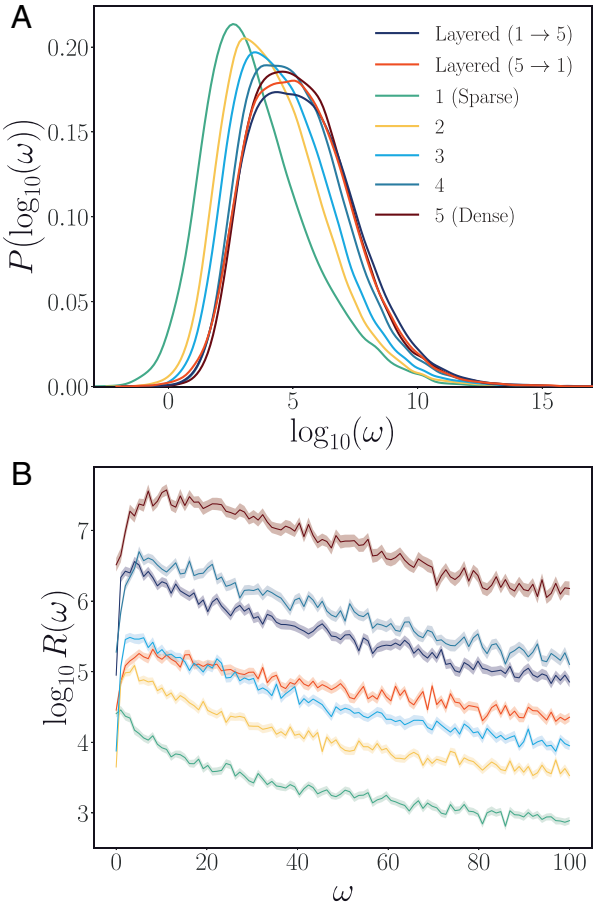

Fig. 4. Actin networks targeted to different densities demonstrate diverse responses. Note that 200 networks were grown to target growth-front monomeric densities (1 to 5), and layered networks were grown with two layers of different growth-front densities (1 → 5 and 5 → 1). (*A*) The distribution of the eigenvalues for the mass-normalized stiffness matrix (*SI Appendix*, Eqs. **S20**–**S28**) is right-shifted under increasing density, with layered densities displaying intermediate distributions. (*B*) The response function—normalized by the number of nodes—$R(\omega)$ *SI Appendix*, Eq. **S32** is higher under increasing density with layered densities displaying intermediate responses.

foundation and a high-density, stiff, growth-front. The protocols to grow target layered structures are nondeterministic; depending on the current state of the system, different loads are required to maintain a particular density and to transition between different densities. Importantly, we are able to target network densities with high fidelity.

Next, we consider the response properties of actin networks to time-periodic forces. We grow networks targeted to different densities, both uniform and layered (Fig. 4). For uniform networks, we observe that the distribution of eigenvalues of the mass-normalized stiffness matrix (*SI Appendix*) is right-shifted under increasing density. Actin networks grown to a higher density will have a larger number of smaller filaments, resulting in higher natural frequencies of the network in comparison to networks with lower density. We further compute the corresponding response functions $R(\omega)$ (*SI Appendix*), which we normalize by the number of nodes. Similarly, for uniform networks, we observe that the denser networks have a higher response as perturbations can more easily be distributed.

Interestingly, we can engineer spatially tuned responses by growing networks with heterogeneous density profiles. We target networks with a low-density foundation and a high-density growth-front (1 → 5) and a high-density foundation with a

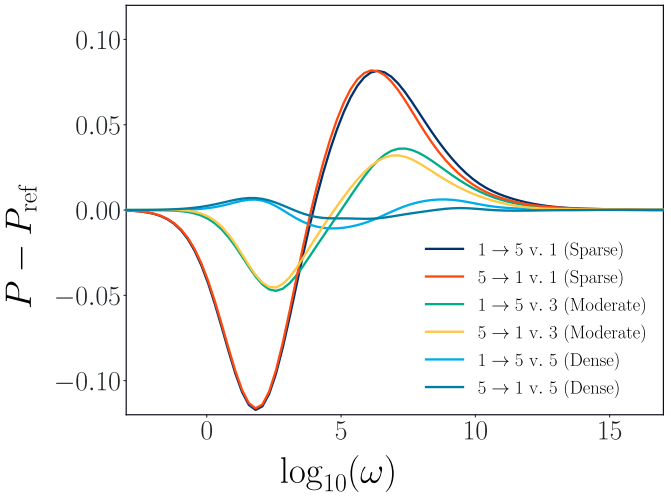

Fig. 5. Differences in the density of states between inhomogeneous networks and homogeneous networks. The layered networks respond to perturbations with a profile distinct from those of any homogeneous network. Compared to the ensemble of sparse networks, the layered networks have a suppressed response at low frequencies and an enhanced response at higher frequencies. Compared to the ensemble of dense networks, the layered networks have an enhanced response at low frequencies and a suppressed response at higher frequencies. These plots indicate the layer networks mix, albeit in a complicated way, the properties of the low-density and high-density actin networks.

low-density growth-front ($5 \rightarrow 1$). We see that the overall response of these layered networks exhibits an intermediate response in comparison to other uniform networks. The response of each node depends strongly on its local environment and the nearby density (Fig. 1*B*). At frequency $\omega = 0.5$, the denser regions have higher responses, while sparser regions of the network have lower responses (Fig. 5). That is, the emergent metamaterial modality of these networks results directly from targeted control of the growth density.

Opportunities for Nonequilibrium Control

Nonequilibrium protocol design provides a compelling route to expanding the capabilities to control and construct nanoscale metamaterials, especially in contexts where modifying the energetic interactions among components is impossible. Here, we take inspiration from reconstitution experiments of dynamically assembling branched actin networks (14) and demonstrate that targeting distinct metamaterial properties is possible without altering the underlying components, just by modulating an external control variable with a feedback protocol. We also show (Eq. **3**) that the entropy production rate and information acquisition rate provide insight into the rate at which we can alter the configurational distribution of a desired material, as discussed in detail in *SI Appendix*. Such fundamental limits shed light on the necessary energetic costs of control, but further investigation of such bounds is required to extract practical design principles for controlling fluctuating materials. The bounds we have derived will provide tighter constraints on control in settings where minimizing dissipation is also an objective (58), a topic we plan to investigate in future work.

The RL control framework we introduce here can readily be extended to an experimental setting, even one without mechanisms for feedback. RL has proved to be flexible and useful toolkit for controlling physical (59) and biomolecular systems (45, 60) evolving under stochastic dynamics. Recent works (61, 62) have also made it increasingly apparent that controlling the kinetics and dynamical pathways of self-assembly expands the scope of control beyond engineering a thermodynamic ground state. Furthermore, the computational and theoretical frameworks we develop motivate a more systematic study into trade-offs between information acquisition, dissipated heat, and the control of a system toward a target state divergent from a nonequilibrium steady state.

Data, Materials, and Software Availability. All code for this project is available on GitHub at the following URL, including scripts and instructions for generating all plots: https://github.com/rotskoff-group/actin-control (63).

ACKNOWLEDGMENTS. We thank Jiawei Yan for helpful discussions. This material is based upon work supported by the United States Department of Energy, Office of Science, Office of Basic Energy Sciences, under Award Number DE-SC0022917. S.K.M. acknowledges the Knut and Alice Wallenberg Foundation for financial support through Grant Number KAW 2021.0328.

1. T. D. Pollard, G. G. Borisy, Cellular motility driven by assembly and disassembly of actin filaments. *Cell* **112**, 453–465 (2003).
2. D. A. Fletcher, P. L. Geissler, Active biological materials. *Annu. Rev. Phys. Chem.* **60**, 469–486 (2009).
3. P. A. Janmey, The cytoskeleton and cell signaling: Component localization and mechanical coupling. *Physiol. Rev.* **78**, 763–781 (1998).
4. P. A. Janmey, C. A. McCulloch, Cell mechanics: Integrating cell responses to mechanical stimuli. *Annu. Rev. Biomed. Eng.* **9**, 1–34 (2007).
5. D. R. Scheff *et al.*, Actin filament alignment causes mechanical hysteresis in cross-linked networks. *Soft Matter* **17**, 5499–5507 (2021).
6. S. Whitelam, R. L. Jack, The statistical mechanics of dynamic pathways to self-assembly. *Annu. Rev. Phys. Chem.* **66**, 143–163 (2015).
7. T. D. Pollard, Regulation of actin filament assembly by Arp2/3 complex and formins. *Annu. Rev. Biophys. Biomol. Struct.* **36**, 451–477 (2007).
8. F. C. MacKintosh, J. Käs, P. A. Janmey, Elasticity of semiflexible biopolymer networks. *Phys. Rev. Lett.* **75**, 4425–4428 (1995).
9. M. L. Gardel, I. C. Schneider, Y. Aratyn-Schaus, C. M. Waterman, Mechanical integration of actin and adhesion dynamics in cell migration. *Annu. Rev. Cell Dev. Biol.* **26**, 315–333 (2010).
10. T. D. Pollard, J. A. Cooper, Actin, a central player in cell shape and movement. *Science* **326**, 1208–1212 (2009).
11. P. Bieling *et al.*, Force feedback controls motor activity and mechanical properties of self-assembling branched actin networks. *Cell* **164**, 115–127 (2016).
12. S. H. Parekh, O. Chaudhuri, J. A. Theriot, D. A. Fletcher, Loading history determines the velocity of actin-network growth. *Nat. Cell Biol.* **7**, 1219–1223 (2005).
13. J. W. Shaevitz, D. A. Fletcher, Load fluctuations drive actin network growth. *Proc. Natl. Acad. Sci. U.S.A.* **104**, 15688–15692 (2007).
14. T. D. Li, P. Bieling, J. Weichsel, R. D. Mullins, D. A. Fletcher, The molecular mechanism of load adaptation by branched actin networks. *eLife* **11**, e73145 (2022).
15. J. W. Chu, G. A. Voth, Allostery of actin filaments: Molecular dynamics simulations and coarse-grained analysis. *Proc. Natl. Acad. Sci. U.S.A.* **102**, 13111–13116 (2005).
16. B. Ding *et al.*, Structure of Arp2/3 complex at a branched actin filament junction resolved by single-particle cryo-electron microscopy. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2202723119 (2022).
17. A. C. Schramm, G. M. Hocky, G. A. Voth, J. L. Martiel, E. M. De La Cruz, Plastic deformation and fragmentation of strained actin filaments. *Biophys. J.* **117**, 453–463 (2019).
18. K. Popov, J. Komianos, G. A. Papoian, MEDYAN: Mechanochemical simulations of contraction and polarity alignment in actomyosin networks. *PLoS Comput. Biol.* **12**, e1004877 (2016).
19. S. L. Freedman, S. Banerjee, G. M. Hocky, A. R. Dinner, A versatile framework for simulating the dynamic mechanical structure of cytoskeletal networks. *Biophys. J.* **113**, 448–460 (2017).
20. A. P. Liu *et al.*, Membrane-induced bundling of actin filaments. *Nat. Phys.* **4**, 789–793 (2008).
21. E. Hohlfeld, P. L. Geissler, Communication: Dominance of extreme statistics in a prototype many-body Brownian ratchet. *J. Chem. Phys.* **141**, 161101 (2014).
22. A. Lamtyugina, Y. Qiu, É. Fodor, A. R. Dinner, S. Vaikuntanathan, Thermodynamic control of activity patterns in cytoskeletal networks. *Phys. Rev. Lett.* **129**, 128002 (2022).
23. Y. Qiu, M. Nguyen, G. M. Hocky, A. R. Dinner, S. Vaikuntanathan, A strong nonequilibrium bound for sorting of cross-linkers on growing biopolymers. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2102881118 (2021).
24. L. Hu, G. A. Papoian, Molecular transport modulates the adaptive response of branched actin networks to an external force. *J. Phys. Chem. B* **117**, 13388–13396 (2013).
25. C. Floyd, H. Levine, C. Jarzynski, G. A. Papoian, Understanding cytoskeletal avalanches using mechanical stability analysis. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2110239118 (2021).
26. C. Floyd, C. Jarzynski, G. A. Papoian, Quantifying dissipation in actomyosin networks. *Biophys. J.* **116**, 254a (2019).
27. H. Ronellenfitsch, N. Stoop, J. Yu, A. Forrow, J. Dunkel, Inverse design of discrete mechanical metamaterials. *Phys. Rev. Mater.* **3**, 095201 (2019).
28. J. W. Rocks *et al.*, Designing allostery-inspired response in mechanical networks. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 2520–2525 (2017).

29. M. J. Falk *et al*., Learning to learn: Non-equilibrium design protocols for adaptable materials. arXiv [Preprint] (2022). https://arxiv.org/abs/2211.02270 (Accessed 3 September 2023).
30. C. P. Goodrich, A. J. Liu, S. R. Nagel, The principle of independent bond-level response: Tuning by pruning to exploit disorder for global behavior. *Phys. Rev. Lett.* **114**, 225501 (2015).
31. D. E. Discher, P. Janmey, Y. l. Wang, Tissue cells feel and respond to the stiffness of their substrate. *Science* **310**, 1139–1143 (2005).
32. M. L. Gardel *et al*., Prestressed F-actin networks cross-linked by hinged filamins replicate mechanical properties of cells. *Proc. Natl. Acad. Sci. U.S.A.* **103**, 1762–1767 (2006).
33. C. Storm, J. J. Pastore, F. C. MacKintosh, T. C. Lubensky, P. A. Janmey, Nonlinear elasticity in biological gels. *Nature* **435**, 191–194 (2005).
34. O. Chaudhuri, S. H. Parekh, D. A. Fletcher, Reversible stress softening of actin networks. *Nature* **445**, 295–298 (2007).
35. W. Zhong, D. J. Schwab, A. Murugan, Associative pattern recognition through macro-molecular self-assembly. *J. Stat. Phys.* **167**, 806–826 (2017).
36. A. K. Behera, Enhancing associative memory recall in non-equilibrium materials through activity. arXiv [Preprint] (2022). https://arxiv.org/abs/2203.03024 (Accessed 3 September 2023).
37. J. M. R. Parrondo, J. M. Horowitz, T. Sagawa, Thermodynamics of information. *Nat. Phys.* **11**, 131–139 (2015).
38. J. M. Horowitz, S. Vaikuntanathan, Nonequilibrium detailed fluctuation theorem for repeated discrete feedback. *Phys. Rev. E* **82**, 061120 (2010).
39. T. Sagawa, M. Ueda, Nonequilibrium thermodynamics of feedback control. *Phys. Rev. E* **85**, 021104 (2012).
40. D. A. Sivak, G. E. Crooks, Thermodynamic metrics and optimal paths. *Phys. Rev. Lett.* **108**, 190602 (2012).
41. G. M. Rotskoff, G. E. Crooks, E. Vanden-Eijnden, Geometric approach to optimal nonequilibrium control: Minimizing dissipation in nanomagnetic spin systems. *Phys. Rev. E* **95**, 012148 (2017).
42. S. Chennakesavalu, G. M. Rotskoff, Unified, geometric framework for nonequilibrium protocol optimization. *Phys. Rev. Lett.* **130**, 107101 (2023).
43. K. Brandner, K. Saito, Thermodynamic geometry of microscopic heat engines. *Phys. Rev. Lett.* **124**, 040602 (2020).
44. J. M. Horowitz, K. Zhou, J. L. England, Minimum energetic cost to maintain a target nonequilibrium state. *Phys. Rev. E* **95**, 042102 (2017).
45. S. Chennakesavalu, G. M. Rotskoff, Probing the theoretical and computational limits of dissipative design. *J. Chem. Phys.* **155**, 194114 (2021).
46. M. Nguyen, S. Vaikuntanathan, Design principles for nonequilibrium self-assembly. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 14231–14236 (2016).
47. A. C. Barato, U. Seifert, Thermodynamic uncertainty relation for biomolecular processes. *Phys. Rev. Lett.* **114**, 158101 (2015).
48. T. R. Gingrich, J. M. Horowitz, N. Perunov, J. L. England, Dissipation bounds all steady-state current fluctuations. *Phys. Rev. Lett.* **116**, 120601 (2016).
49. T. R. Gingrich, G. M. Rotskoff, J. M. Horowitz, Inferring dissipation from current fluctuations. *J. Phys. A: Math. Theor.* **50**, 184004 (2017).
50. C. S. Peskin, G. M. Odell, G. F. Oster, Cellular motions and thermal fluctuations: The Brownian ratchet. *Biophys. J.* **65**, 316–324 (1993).
51. A. Mogilner, G. Oster, Cell motility driven by actin polymerization. *Biophys. J.* **71**, 3030–3045 (1996).
52. A. Mogilner, G. Oster, Force generation by actin polymerization II: The elastic ratchet and tethered filaments. *Biophys. J.* **84**, 1591–1605 (2003).
53. A. Mogilner, G. Oster, Polymer motors: Pushing out the front and pulling up the back. *Curr. Biol.* **13**, R721–R733 (2003).
54. I. Maly, G. Borisy, Self-organization of a propulsive actin network as an evolutionary process. *Proc. Natl. Acad. Sci. U.S.A.* **98**, 11324–11329 (2001).
55. T. E. Schaus, E. W. Taylor, G. G. Borisy, Self-organization of actin filament orientation in the dendritic-nucleation/array-treadmilling model. *Proc. Natl. Acad. Sci. U.S.A.* **104**, 7086–7091 (2007).
56. J. Weichsel, U. S. Schwarz, Two competing orientation patterns explain experimentally observed anomalies in growing actin networks. *Proc. Natl. Acad. Sci. U.S.A.* **107**, 6304–6309 (2010).
57. K. N. Fish, Total internal reflection fluorescence (TIRF) microscopy. *Curr. Protocol.* **2**, e517 (2022).
58. S. Chennakesavalu, G. M. Rotskoff, Unified, geometric framework for nonequilibrium protocol optimization. *Phys. Rev. Lett.* **130**, 107101 (2023).
59. M. Bukov *et al*., Reinforcement learning in different phases of quantum control. *Phys. Rev. X* **8**, 031086 (2018).
60. S. Chennakesavalu, G. M. Rotskoff, "Cooperative multi-agent reinforcement learning for high-dimensional nonequilibrium control" in *Fourth Workshop on Machine Learning and the Physical Sciences*, M. Ranzato *et al*., Eds. (Curran Associates, Inc., 2021), p. 7.
61. C. P. Goodrich, E. M. King, S. S. Schoenholz, E. D. Cubuk, M. P. Brenner, Designing self-assembling kinetics with differentiable statistical physics models. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2024083118 (2021).
62. A. Trubiano, M. F. Hagan, Optimization of non-equilibrium self-assembly protocols using Markov state models. *J. Chem. Phys.* **157**, 244901 (2022).
63. S. Chennakesavalu, S. K. Manikandan, F. Hu, G. M. Rotskoff, Actin Control Codebase. GitHub. https://github.com/rotskoff-group/actin-control. Deposited 20 November 2023.