2018 IEEE Conference on Decision and Control (CDC)
Miami Beach, FL, USA, Dec. 17-19, 2018

Control of chemotaxis through absolute concentration robustness

Sayak Bhattacharya, Debojyoti Biswas, Germán A. Enciso and Pablo A. Iglesias

Abstract— **Chemotaxis, the directed motion of cells in response to chemical gradients, is important for a variety of biological processes ranging from embryogenesis to killing of pathogens. Increasing the speed and efficiency of directed migration is critical in such situations. We provide a control mechanism by which one can minimize the noise-driven firings at the back of the cell, enabling faster motion towards the front. We achieve this through a mechanism called absolute concentration robustness (ACR), which robustly maintains the steady-state concentration of intracellular biochemical species and, at the same time, provides control over the concentration variance. More particularly, by incorporating ACR, we develop a correspondence between the concentration mean and variance — both of which are independent of total concentrations. We show that by incorporating ACR into the back of a moving cell, we can create a mechanism to robustly control the noise variance at the back — thus limiting the deterring firings while the cell moves in the direction of the gradient.**

I. INTRODUCTION

Chemotaxis, or directed migration towards a source, is important in numerous biological and physiological processes — making its regulation and control critical for many phenomena. For example, when neutrophils or macrophages move towards a wound site for healing [1] or towards bacteria for phagocytosis [2], it is essential that the process be efficient and fast. Engineering cells to move faster is thus an important area of ongoing research.

Amoeboid cells move by extending repeated protrusions of the cell cortex. These protrusions — called pseudopods, are the fundamental unit of migration. Pseudopods are formed when waves of actin polymerization [3] push on the cell membrane, creating an extension [4]. Recent advances in cell biology have shown that these waves are generated through an excitable network that resides in the signaling apparatus of the cell [5], [6]. When this network is triggered, an all-or-none type response is generated in signaling molecules such as Ras and $\mathrm{PI}(3,4,5)\mathrm{P}_3$ that in turn signal the downstream actin components. It is this response, followed by a brief refractory period, that generates a wave front. The wave front can propagate along adjacent excitable cortex elements — leading to wave propagation.

A key concept of excitable systems is the notion of a threshold for activation. Supra-threshold inputs result in a large response while sub-threshold inputs are ignored. Even without any gradient or source, signaling and actin waves are seen to be randomly generated on the cell membrane, leading to the conclusion that stochastic perturbations are mainly responsible for the triggering of the excitable network. Cells with high thresholds fire more often while the cells with low thresholds are more quiescent. In the presence of a gradient, a spatial bias in threshold is formed through a local excitation, global inhibition mechanism that creates a lower threshold towards the side of the cell facing the gradient (the front) while the rest of the cell (the back) increases its threshold level [7]. This causes the activity of the cell to be biased towards the direction of the chemoattractant source, resulting in directed migration [8].

In such situations, although the threshold is raised at the back, stochastic firings continue to occur at the back of the cell owing to the natural fluctuations of intrinsic and extrinsic components — albeit at a lower rate than the front. These back firings slow down the motion of the cell. There are two potential ways to control these firings. One could engineer a mechanism to further increase the threshold at the back. Alternatively, one could design a control system that provides some control over the variance in the fluctuations of specific molecules that trigger cell movement.

There are specific biochemical entities known in the cellular signaling cascade, whose levels are critical for the firing of the system. In an earlier paper, we showed how synthetic manipulations in the levels of these proteins/lipids can dramatically alter the migratory mode of the cell [9]. Hence achieving control over the levels and fluctuations of these entities would be ideal for efficient chemotaxis. Here, we propose to do this through a mechanism called absolute concentration robustness (ACR) [10]. ACR comes about through the binding of the target species in a specific reaction scheme, which ensures that the steady-state concentration of the target remains robustly stable to fluctuations in its environment. In this study, we show how ACR can be used to achieve control over the levels of critical lipids in the cellular signaling network which in turn directly affects the firing rate of the underlying excitable system. Together, this creates a control scheme for faster and more efficient directed migration.

II. ABSOLUTE CONCENTRATION ROBUSTNESS

Absolute concentration robustness guarantees that the concentration of a particular species is robustly stable around a fixed point in the steady-state. More particularly, in a deterministic setting, for every possible equilibrium of the dynamical system, the species concentration is constant —

This work was supported by DARPA under contract number HR0011-16-C-0139

SB is with Department Electrical and Computer Engineering, The Johns Hopkins University, Baltimore, MD, USA sbhatt11@jhu.edu

DB is with Department Electrical and Computer Engineering, The Johns Hopkins University, Baltimore, MD, USA dbiswas2@jhu.edu

GAE is with Faculty of Mathematics, University of California, Irvine, CA, USA enciso@uci.edu

PAI is with Faculty of Electrical and Computer Engineering, The Johns Hopkins University, Baltimore, MD, USA pi@jhu.edu

impervious to changes in total molar concentrations. This concept was suggested by Shinar and Feinberg [10], who imposed certain conditions on the structure of a biochemical network to guarantee ACR.

To illustrate this concept, we consider the species P_A as the target species for ACR. To achieve ACR we use another species P_B in the following mass action reaction scheme:

$$P_A + P_B \xrightarrow{\alpha} 2P_B \tag{1a}$$

$$P_B \xrightarrow{\beta} P_A \tag{1b}$$

Solving for the steady-state concentrations of the species for a total molar concentration of $c_M = c_{P_A} + c_{P_B}$, it is not difficult to see that the steady-state is:

$$c_{P_A} = \frac{\beta}{\alpha}$$

$$c_{P_B} = c_M - \frac{\beta}{\alpha}$$

Thus, no matter what the initial or total concentrations are, the concentration of P_A is constant, and depends only on the binding rates. This concentration is independent even of the binding species P_B, which would not be true for regular binding reactions like:

$$P_A \underset{\widetilde{\beta}}{\overset{\widetilde{\alpha}}{\rightleftharpoons}} P_B,$$

where:

$$c_{P_A} = \frac{\widetilde{\beta}}{\widetilde{\alpha} + \widetilde{\beta}} c_M,$$

i.e. the resulting steady state concentration of P_A depends on the total molar concentrations.

Thus the ACR scheme provides a robust mechanism to maintain a constant concentration of a particular reactant. As we motivated briefly in the introduction, when this species concentration acts as an input for the excitable network, besides the mean level of the species, it is the fluctuation around this mean level that triggers the excitable system. Hence, controlling the variance of the species concentration is critical in the context of excitable systems.

The exact relation between the mean and variance of the input signal to the excitable system performance will be described in detail in the next section. For now, it is vital to also consider the structure of the ACR in a stochastic setting, in order to understand the steady-state probability distributions of the target species. The deterministic concentrations give us an idea of the mean of the resulting distribution, but no comments can be made about the variance from the deterministic equations.

Anderson et al. [11] provide a rigorous stochastic analysis of the ACR structure as they present conditions with regards to the long-term and short-term behavior of the system. Here we highlight some key points from that study, which will be pertinent for our analysis of the excitable system.

The stochastic model of the ACR can be approximated as a birth-death process in P_B as the first reaction creates a new molecule of P_B while the second reaction destroys one. Thus the propensities of the two reactions can be written as:

$$a_1(p) = \alpha p(M-p)$$

$$a_2(p) = \beta p,$$

where p denotes the number of molecules of P_B, and M denotes the total number of molecules.

There exists an absorption state in this reaction setup when $c_{P_B} = 0$ and $c_{P_A} = c_M$. Although this extinction state will inevitably be reached, it is not practical for biologically reasonable timescales [11]. For finite time reactions, it can be proved using the chemical master equation, that the ACR settles at a quasi-stationary distribution [12]. For a large M, the steady-state number of molecules of P_A resembles the Poisson distribution:

$$\pi_{P_A}(p) = \frac{e^{-\beta/\alpha}}{p!}\left(\frac{\beta}{\alpha}\right)^p,$$

with parameter $\frac{\beta}{\alpha}$.

The advantage of a Poisson process is that variance is now equal to the mean. Thus, for the ACR scheme, we have that:

$$\mu_{ACR} = \sigma^2_{ACR} = \frac{\beta}{\alpha}$$

The same is not necessarily true for normally distributed systems where the mean and variance can be independent parameters. For large means, Poisson fluctuations are small relative to the mean. Thus, in the case of a Poisson process, if we have a method to control the mean of the system, we can automatically put a restriction on the variance, or fluctuations about the mean. Thus, while the ACR scheme grants us robust control over the mean (μ_{ACR}) of the process, the resulting Poisson distribution automatically provides a control mechanism to modulate the variance (σ^2_{ACR}) of the system by tuning μ_{ACR}.

III. THE EXCITABLE MODEL

We now illustrate how the mean and variance of a particular species concentration can be crucial to controlling the dynamics of the excitable system. For this purpose, we introduce a two-dimensional model of an excitable network [13]. This is given by the following equations that describe the dynamics of a coupled activator-inhibitor system:

$$\frac{\mathrm{d}u}{\mathrm{d}t} = -a_1 u - a_2 u(v-r) + \frac{a_3 u^2}{a_4^2 + u^2} + B = f(u,v) \tag{2a}$$

$$\frac{\mathrm{d}v}{\mathrm{d}t} = \varepsilon(-v + c_{th} u) = g(u,v) \tag{2b}$$

The activator (u) equation (2a) is non-linear owing to a co-operativity term that enables it to stimulate its own production through positive feedback. It simultaneously creates an inhibitor (v) to subdue its response through a delayed negative feedback (Fig. 1A). The delay in the inhibitor is incorporated through the ε term in (2b) where $\varepsilon \ll 1$. This creates two different time-scales for the activator and inhibitor, causing this system to act as a relaxation oscillator.

The dynamics of this system are better visualized in phase space (Fig. 1B). The set point, or equilibrium of the system is the intersection between the nullclines (u_-, v_0). This equilibrium is close to a Hopf bifurcation point — which is at the local minimum of the cubic nullcline ($v_{\min}$, Fig. 1C). If the input is sufficient to displace the state beyond this bifurcation point, then the state undergoes a large excursion in phase space as illustrated (Fig. 1B, past $u^\star$ upto u_+), with the trajectory time-scales typical of any relaxation oscillator.

The input to the system is modeled through the signal r which serves to raise the f-nullcline upwards.

$$r = k_r(r_{in} - m_r),$$

where r_{in} is the input random process and k_r and m_r are constants that are used to scale the input process so as to be compatible with the excitable system.

Near the bifurcation point, a sharp increase in activator concentration occurs resulting in a large excursion in phase space. Clearly, if the input is sufficient enough such that the initial rightward trajectory of the state does not encounter the local minimum, a response will occur. Thus, the threshold for activation can be roughly demarcated as the distance between the set point and the local minimum of the cubic nullcline. We say "roughly", as excitable systems in fact possess a threshold manifold that has to be crossed to elicit a response [14]. However for the purposes of this paper, it suffices to assume that the closer the initial equilibrium is to the bifurcation point, the smaller the threshold is. Other than through the input signal, we can also modulate the threshold of the system using the slope (c_{th}) of the inhibitor nullcline.

Kramers' theory [15], [16] provides a formula relating firing rate (f) with threshold and input variance.

$$f \propto \exp\left(-\frac{2\Delta U}{\sigma}\right),$$

where ΔU is a measure of the threshold barrier, while σ is the variance of the input signal. This agrees with the intuitive reasoning, that the firing rate increases with a larger input variance and decreases with a larger threshold value. Simulations altering the mean and variance of the input process r_{in} show how the firing rate changes, further confirming this intuition (Fig. 1D). For these simulations, we used a normally distributed r_{in} with chosen mean and variance parameters. Owing to the refractory period of the output response, an infinite number of spikes are not possible in finite time. The simulation time for each trial was set such that a maximum of 6 spikes could occur.

In the context of cell migration, many proteins/lipids have been shown to control the threshold of the system, and perturbations in the levels of these molecules have dramatic consequences on the activity of the cell [9]. The signal r_{in} can be thought of as one of those species that are critical for regulating threshold, where we want to establish a control mechanism to achieve robust control over the mean *and* variance of that species concentration.

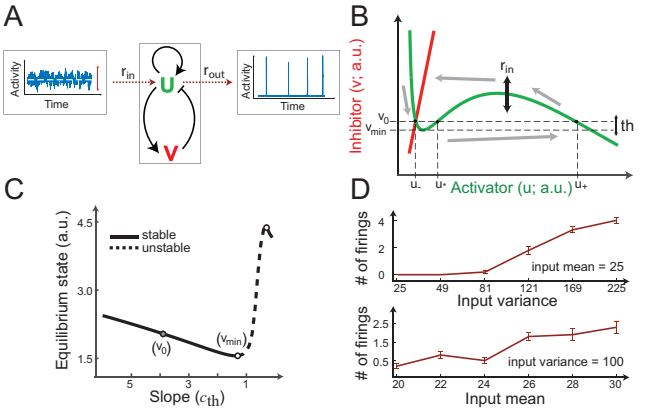

Fig. 1. Model dynamics. A. The input-output system, where the input signal is normally distributed, while the output is obtained in spikes. The system consists of the activator (U, green) and the inhibitor (V, red). B. The excitable system in phase plane, where the green and red curves denote the activator and inhibitor nullclines respectively. The grey arrows show the large excursion in phase plane that produces one spike. The threshold for activation is shown by th. The input signal r_{in} is set up such that it causes the activator nullcline to fluctuate vertically. C. The bifurcation diagram showing the stability and instability regions. D. Plots to illustrate how number of firings can be significantly altered when the input mean (top) or input variance (bottom) is changed. The number of firings are calculated in 200 time units where a maximum of 6 spikes can occur. The parameters from (2) are: $a_1 : 0.167, a_2 : 16.67, a_3 : 167, a_4 : 1.2, B : 1.47, \varepsilon : 0.03, c_{th} : 40, k_r : 0.04, m_r : 25$. The Euler integration time-step (Δt) was 0.01. Errorbars were calculated using 10 simulations.

IV. THE CONTROL MECHANISM

Using the idea that absolute concentration robustness grants us robust control over the mean and variance of the input signal with respect to fluctuations, we suggest that the target species be one of those responsible for the firings of the excitable system. The mean concentration of the protein decides the relative threshold of the system, while the variance controls the firing rate. We analyze two possible input pathways to the excitable system to illustrate the advantage of ACR. In the first case, a normally distributed random variable of a certain mean and variance is directly added as an input to the excitable network. In the second case, we start with a normal process of the same mean and variance, but allow the species to bind as dictated by the ACR scheme. We take the resultant signal and feed it as the input to the excitable network (Fig. 2A).

As mentioned earlier, we set the simulation time so that a maximum of six spikes could occur. We choose the initial mean and variance from Fig. 1D such that we can expect around 2 firings in the same time from the excitable system for a white noise input of constant mean and variance. The processed and un-processed input signals are compared through histograms (Fig. 2B), which reflect the frequency and spread of the concentration levels. The variance of the ACR distribution is significantly reduced for the same mean as compared to the Gaussian process, reflected by a tighter spread of the histogram (Fig. 2B). More particularly, if the chosen initial variance is less than the square root of the mean, then the Gaussian input will fluctuate lesser than the ACR generated process. On the other hand, if this variance

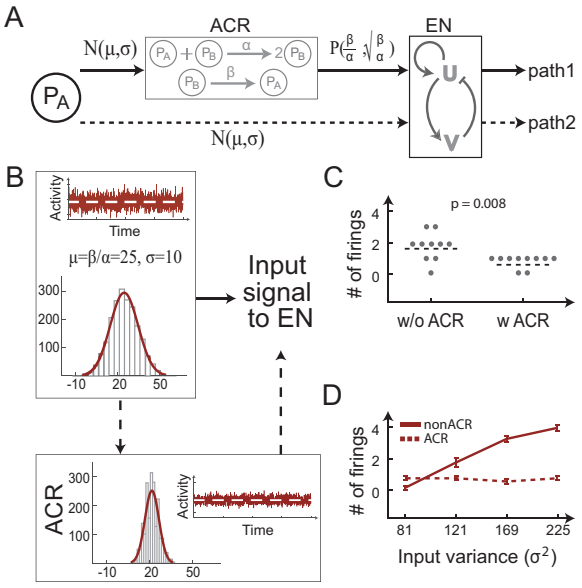

Fig. 2. The control mechanism. A. The two mechanisms for controlling the number of spikes. In both cases, the input concentration P_A is initially normally distributed. In path 1, the concentration level over time is processed through the ACR scheme, converting it to a Poisson process as it is fed into the excitable system (EN). In path 2, the concentration of P_A directly serves as the input for the excitable system. B. Illustrating the two pathways with probability distributions. Both processes are shown through histograms. An example signal is also shown for both cases. For the ACR process, $\alpha = 1, \beta = 25$ were chosen. C. Comparison of the number of firings for both pathways. The p-value was obtained from the student's t-test for 10 simulations. D. Change in the number of firings when the initial variance is altered for both processes, where the dashed plot represents the ACR path, while the solid line represents the normal path. All number of firings are in 200 time units where a maximum of 6 firings can occur. Errorbars were obtained using 10 simulations.

is higher, the fluctuations will be higher which will translate to significantly increased firings of the excitable system (Fig. 2C). Thus, while for the Gaussian input, the variance and, in turn, the number of firings are dependent on the initial condition, the ACR-processed signal will robustly display a constant variance, and thus, in turn, a relatively constant number of firings from the excitable system (Fig. 2D).

In the simulations, the initial concentration of P_A is generated through a normally distributed random number. P_A is then made to bind through the ACR scheme, simulated using the direct method of the stochastic simulation algorithm [17]. As we see from Fig. 2C, per the chosen parameters, while a normally distributed random variable generates two spikes on average, the ACR processed variable robustly generates fewer firings from the excitable system (Fig. 2D), no matter what the input variance is. This is again because, in the ACR scheme, irrespective of the initially chosen variance, the post-processed mean and variance depends only on the binding constants. Note that this means that α and β can be modified to generate the desired number of firings from the excitable system. Why we chose to illustrate the low firing case will become clear when we discuss the case of directed migration.

V. ACR IN DIRECTED MIGRATION

During directed migration towards a gradient or source such as a chemoattractant (cAMP), the cell biases its threshold spatially so as to ensure increased firing in the direction of the source, i.e. the front of the cell [7]. It is imperative that the back of the cell remain relatively quiescent during this process. During chemotaxis, the front of the cell experiences the maximum pull towards the source while the number of firings at the back can affect how fast the cell will go. In an ideal case, zero firings at the back will result in the fastest and most efficient movement of the cell towards the source.

For this purpose, we hypothesized that the fluctuations at the front need not be controlled as a low enough threshold and maximum fluctuations at the front is ideal for the cell. However at the back, we want the threshold to be robustly stable at a high and at the same time minimize its fluctuations, as even for a high threshold large fluctuations can elicit firings from the system. In order to simulate this, (2) was substituted by reaction-diffusion equations across a one-dimensional domain discretized into 314 points, where diffusion was simulated using the central difference approximation. Owing to the diffusion terms, each firing could excite adjacent excitable elements — resembling the signaling and actin wave activity on the cell membrane.

$$\frac{\partial u}{\partial t} = D_u \nabla^2 u + f(u,v) \tag{3a}$$

$$\frac{\partial v}{\partial t} = D_v \nabla^2 v + g(u,v) \tag{3b}$$

The control scheme is illustrated in Fig. 3A, where the P_A concentration at the front (half of the cell) is allowed to fluctuate based on a Gaussian process, while the back is made to bind through ACR. At the back of the cell, we chose a high initial threshold (c_{th}) and applied a normally distributed random process r_{in} of mean μ and variance σ. This was compared to a r_{in} that was processed with ACR. In the same simulation, ACR binding was initiated after a certain time interval using the instantaneous concentration of P_A, which was normally distributed till that point in time (Fig. 3B)). The front of the cell was assumed to have a low c_{th}, pushing the equilibrium close to the Hopf bifurcation point, so as to ensure maximal activity in the direction of the source. The threshold was initialized as a Gaussian function with the lowest threshold at the front and high threshold at the tails (back) (Fig. 3B), as is the case in the presence of a gradient [7].

The ACR parameters were chosen with $\mu = \beta/\alpha$, so that the ACR-processed mean was comparable. For $\sigma^2 < \frac{\beta}{\alpha}$, there would not be any noticeable difference between the back firings of the two cases. However, for $\sigma^2 > \frac{\beta}{\alpha}$, the ACR process showed significantly fewer number of firings at the back than that for the Gaussian input. The simulation results are shown through kymographs (Fig. 3B) where the discretized cell membrane was unwrapped and plotted along the y-axis, with time on the x-axis. As each firing now translated to a wave, more activity at the back of the cell was

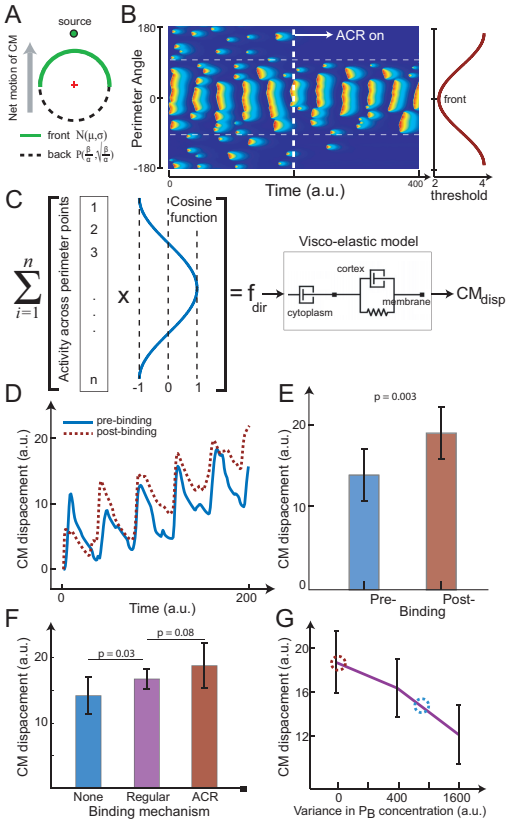

Fig. 3. ACR in directed migration. A. The control scheme, where the front of the cell (marked in green) is the half of the cell facing the source. The front fluctuations are controlled by a Gaussian process. The back (dashed line) is controlled by the Poisson process output from the ACR scheme. The center of mass (CM) is shown by the red cross. B. Kymograph showing the simulated membrane activity. The x-axis shows the perimeter of the cell. The ACR signal is activated after 200 time units (dashed white line). The front of the cell has the lowest threshold, as the threshold gradually increases towards the back. The horizontal dashed white lines demarcate the half of the cell designated as the front. The parameters chosen were the same as in Fig. 1. The diffusion coefficients in (3) were chosen as: $D_u = 0.8, D_v = 1$. The visco-elastic parameters were taken from Yang et al. [18] C. The center-of-mass approximation scheme, where the net force (f_{dir}) calculated using the weighted cosine was fed into the visco-elastic model to obtain the net displacement (CM_{disp}). D. Graph plotting the CM displacement with time, before and after ACR is initiated. E. Bar-graph quantification of the end point of the CM. Errorbars were calculated using 10 simulations. The p-value was obtained from the student's t-test. F. The CM displacement bar graph for all three binding situations, no binding, ACR binding and regular binding for a P_B variance of 100. G. The change in the CM displacement as the variance in P_B is increased. The ACR (brown circle) and no-binding (blue circle) values are marked for reference. The errorbars are calculated using 10 simulations.

reflected on the kymograph through multiple little wavelets which ceased to exist once binding was initiated. After the ACR binding was switched on, almost all the firings were concentrated in the front of the cell (between the horizontal dashed white lines in Fig. 3B).

To understand the implication of ACR on cell speed, we used a center of mass approximation to calculate the net displacement of a cell towards the source. This scheme is illustrated in Fig. 3C. The activity on the membrane was weighted with a cosine function, so that the activity at the front had the highest positive weight, while the activity at the back received a negative weight. The sum of this weighted activity was used as a measure of the net displacement. This number was fed into a visco-elastic cell membrane model that accounted for surface tension and membrane elastic effects [18]. The effect of the elastic nature of the membrane can be seen in the center-of-mass (CM) displacement plot which is oscillating (Fig. 3D).

Note from Fig. 3D that even in the case of the normally distributed input, the net motion of the cell is indeed towards the source. However as hypothesized, after the ACR binding was initiated the cell moved faster towards the source, owing to less hindrance from back firings. Using the center-of-mass approximation, we showed that the net displacement of the cell was significantly more after binding was initiated than that before ACR (Fig. 3E).

We further compared the performance of the ACR scheme with a regular reversible binding reaction. Although the simple binding also provides a correspondence between the mean and variance of the process — the mean still remains susceptible to total molar concentration fluctuations. We illustrate this problem by introducing some randomness in the concentration of P_B and calibrating the center of mass displacement for directed migration. When there is no variance in concentration, the regular binding works just as well as the ACR scheme. However, for increasing P_B fluctuations, while the ACR binding remains unaffected, the regular binding performance steadily declines (Fig. 3F,G).

VI. DISCUSSION

Efficient and fast directed migration is key in numerous processes and tremendous efforts are underway in synthetic biology to engineer faster cells. One method to do this is to control the level of threshold as the cell moves up a gradient and ensure that the back of the cell fires as little as possible. Many signaling proteins/lipids such as $\mathrm{PI}(4,5)\mathrm{P}_2$, RasC and $\mathrm{PI}(3,4,5)\mathrm{P}_3$ have been identified [9], which control the level of threshold in the cellular network. Using the concept of absolute concentration robustness we can robustly control the levels of any of these molecules to stabilize their concentrations at the desired level.

However, in the stochastic setting, it is not enough to control the mean level of the concentrations. As we see even with high means, the fluctuations in the mean can be sufficient to trigger undesirable firings from the excitable system. For that purpose, we want a system that not only renders the mean impervious to fluctuations in molecular concentrations but also gives us control over the variance of the process.

Because the variance of Poisson processes equals the mean, the ACR-processed variable is advantageous over a normally distributed variable as the ACR process approaches a Poisson quasi-stationary distribution. This creates a tight distribution with a controlled variance, and for large means the Poisson variance is small.

Using this idea, we can effectively eliminate the fluctuations at the back of the cell by robustly controlling the threshold at the back and maintaining a low variance of

concentrations using ACR. This would cause the activity of the cell to remain focussed at the front of the cell, resulting in maximum velocity towards the source, enabling faster and more efficient migration. A simple binding scheme, although generates a Poisson process, would still remain susceptible to variations in total molar concentrations — something that can be expected across different cells. Thus, although a regular reversible binding mechanism would be effective for some molecular concentrations, it would not be robust enough to guarantee faster migration for all cells.

A key component of this idea is to find prospective candidates for P_A and P_B. P_A, as noted earlier, will be a threshold modulating lipid like $\mathrm{PI}(4,5)\mathrm{P}_2$. We need for P_B to be a lipid that selectively localizes at the back of the cell during migration. Only then can we initiate binding using the ACR scheme and localize the binding at the back of the cell. There are many proteins/lipids known that spatially organize themselves and thus can be segregated into front and back regions. For example, PTEN [19], CynA [20] or $\mathrm{PI}(3,4)\mathrm{P}_2$ [21] are candidates that fall off the membrane during a protrusion and localize to the back of the cell as inhibitory enzymes. Synthetically engineering the ACR scheme to bind $\mathrm{PI}(4,5)\mathrm{P}_2$ and PTEN, and by choosing the binding constants carefully, one can potentially engineer faster and more efficiently chemotaxing cells.

VII. ACKNOWLEDGMENTS

We are thankful to Dr. Peter N. Devreotes at the Johns Hopkins School of Medicine for many fruitful conversations over the years.

REFERENCES

[1] T. J. Shaw and P. Martin, "Wound repair at a glance," *Journal of cell science*, vol. 122, no. 18, pp. 3209–3213, 2009.

[2] A. Aderem and D. M. Underhill, "Mechanisms of phagocytosis in macrophages," *Annual review of immunology*, vol. 17, no. 1, pp. 593–623, 1999.

[3] T. Bretschneider, K. Anderson, M. Ecke, A. Müller-Taubenberger, B. Schroth-Diez, H. C. Ishikawa-Ankerhold, and G. Gerisch, "The three-dimensional dynamics of actin waves, a model of cytoskeletal self-organization," *Biophys J*, vol. 96, no. 7, pp. 2888–900, Apr 2009.

[4] A. Mogilner and G. Oster, "Cell motility driven by actin polymerization," *Biophysical journal*, vol. 71, no. 6, pp. 3030–3045, 1996.

[5] P. A. Iglesias and P. N. Devreotes, "Biased excitable networks: how cells direct motion in response to gradients," *Curr Opin Cell Biol*, vol. 24, no. 2, pp. 245–53, Apr 2012.

[6] P. N. Devreotes, S. Bhattacharya, M. Edwards, P. A. Iglesias, T. Lampert, and Y. Miao, "Excitable signal transduction networks in directed cell migration," *Annual Review of Cell and Developmental Biology*, vol. 33, pp. 103–125, 2017.

[7] Y. Xiong, C.-H. Huang, P. A. Iglesias, and P. N. Devreotes, "Cells navigate with a local-excitation, global-inhibition-biased excitable network," *Proc Natl Acad Sci U S A*, vol. 107, no. 40, pp. 17 079–86, Oct 2010.

[8] C. Shi, C.-H. Huang, P. N. Devreotes, and P. A. Iglesias, "Interaction of motility, directional sensing, and polarity modules recreates the behaviors of chemotaxing cells," *PLoS Comput Biol*, vol. 9, no. 7, p. e1003122, 2013.

[9] Y. Miao, S. Bhattacharya, M. Edwards, H. Cai, T. Inoue, P. A. Iglesias, and P. N. Devreotes, "Altering the threshold of an excitable signal transduction network changes cell migratory modes," *Nature Cell Biology*, vol. 19, no. 4, pp. 329–340, 2017.

[10] G. Shinar and M. Feinberg, "Structural sources of robustness in biochemical reaction networks," *Science*, vol. 327, no. 5971, pp. 1389–1391, 2010.

[11] D. F. Anderson, G. A. Enciso, and M. D. Johnston, "Stochastic analysis of biochemical reaction networks with absolute concentration robustness," *Journal of The Royal Society Interface*, vol. 11, no. 93, p. 20130943, 2014.

[12] G. A. Enciso, "Transient absolute robustness in stochastic biochemical networks," *Journal of The Royal Society Interface*, vol. 13, no. 121, p. 20160475, 2016.

[13] S. Bhattacharya and P. A. Iglesias, "The threshold of an excitable system serves as a control mechanism for noise filtering during chemotaxis," *PloS one*, vol. 13, no. 7, p. e0201283, 2018.

[14] E. M. Izhikevich, *Dynamical systems in neuroscience: the geometry of excitability and bursting*. Cambridge, Mass.: MIT Press, 2007.

[15] P. Iglesias, "Excitable systems in cell motility," in *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, Florence, Italy, Dec 2013, pp. 757–762.

[16] P. Hänggi, P. Talkner, and M. Borkovec, "Reaction-rate theory: fifty years after kramers," *Reviews of modern physics*, vol. 62, no. 2, p. 251, 1990.

[17] D. T. Gillespie, "Stochastic simulation of chemical kinetics," *Annu Rev Phys Chem*, vol. 58, pp. 35–55, 2007.

[18] L. Yang, J. C. Effler, B. L. Kutscher, S. E. Sullivan, D. N. Robinson, and P. A. Iglesias, "Modeling cellular deformations using the level set formalism," *BMC Syst Biol*, vol. 2, p. 68, 2008.

[19] P. J. Van Haastert and P. N. Devreotes, "Chemotaxis: signalling the way forward," *Nature reviews Molecular cell biology*, vol. 5, no. 8, p. 626, 2004.

[20] K. F. Swaney, J. Borleis, P. A. Iglesias, and P. N. Devreotes, "Novel protein callipygian defines the back of migrating cells," *Proc Natl Acad Sci U S A*, vol. 112, no. 29, pp. E3845–54, Jul 2015.

[21] X. Li, M. Edwards, K. F. Swaney, N. Singh, S. Bhattacharya, J. Borleis, Y. Long, P. A. Iglesias, J. Chen, and P. N. Devreotes, "Mutually inhibitory Ras-PI(3,4)P2 feedback loops mediate cell migration," *Proceedings of the National Academy of Sciences*, 2018, doi: 10.1073/pnas.1809039115.