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

Architecture and Trade-offs in the Heat Shock Response System

Noah Olsman[1], Carmen Amo Alonso[2], and John C. Doyle[3]

***Abstract*— Biological control systems often contain a wide variety of feedforward and feedback mechanisms that regulate a given process. While it is generally assumed that this apparent redundancy has evolved for a reason, it is often unclear how exactly the cell benefits from more complex circuit architectures. Here we study this problem in the context of a minimal model of the Heat Shock Response system in *E. coli* and show, through a combination of theory and simulation, that the complexity of the natural system outperforms hypothetical simpler architectures in a variety of robustness and efficiency tradeoffs. We have developed a significantly simplified model of the system that faithfully captures these rich issues. Because a great deal of biological detail is known about this particular system, we are able to compare simple models with more complete ones and obtain a level of theoretical and quantitative insight not generally feasible in the study of biological circuits. We primarily hope this will inform future analysis of both heat shock and newly studied biological complexity.**

I. INTRODUCTION

Biological control systems, much like engineered ones, are faced with a variety of heterogeneous constraints that shape their design [1], [2]. Because of this, the selective pressures of evolution have likely not only selected for nominal performance, but also for robustness and efficiency [3]. While many different system architectures may be able to perform a given task, it is reasonable to assume that the ones that actually evolve reflect a balance between the myriad trade-offs faced by the cell.

In this paper, we examine such trade-offs in the context of the Heat Shock Response (HSR) system of *E. coli*. When a cell encounters a rapid increase in temperature, there is a corresponding increase in the rate at which its proteins become misfolded. If too many of the cell's proteins are misfolded, the cell will likely die. To prevent this, cells will produce Heat Shock Proteins (HSPs) whose job is to rapidly refold proteins so that the organism can continue functioning normally.

From a systems perspective the HSR clearly must be an extremely fast and robust system, as it is vital to a cell's ability to respond to sudden and unexpected stress. On the other hand, if heat shock is a rare event then it would be wasteful for the cell to constantly produce HSPs at a high rate. Intuitively, we would expect that a strong architecture would be capable not only of refolding proteins, but would do so as quickly, efficiently, and robustly as possible.

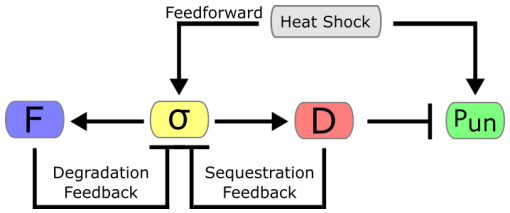

Fig. 1. Abstract block diagram of the full HSR System. This shows a high-level representation of how the HSR system uses a combination of feedforward and feedback mechanisms to regulate the level of unfolded protein in a cell. Arrows signify effects which increase concentration, while flat-headed line signify repression of activity.

We find that the natural HSR system in *E. coli* does an excellent job balancing these various trade-offs. To demonstrate this, we examine several hypothetical alternative architectures for the system, and show that the complexity of the full system results in performance that none of the simpler systems can match. While it is possible for these reduced architecture to perform well on a given metric, e.g. efficiency or speed, they are limited in their ability to do well on many tasks simultaneously. We find that the strong performance of the natural HSR system is due to several elegant mechanisms, for example a layering of planning and control modules, that work together to make the system highly functional.

In section II, we will present a new reduced-order model of the HSR system that faithfully captures the quantitative results of more complex models. In section III we will present hypothetical alternative architectures to the HSR system. In section IV we analyze trade-offs between several response metrics, and compare the relative performance of the architectures discussed in section III.

II. REDUCED-ORDER MODEL

The core of the HSR system involves four classes of proteins: the unfolded proteins which the system is trying to refold, the chaperone whose job it is to refold proteins, the σ factor that regulates the expression of all proteins involved in heat shock response, and the

[1,2,3] Department of Control and Dynamical Systems, California Institute of Technology.
[1] Email: nolsman@caltech.edu
[2] Email: camoalon@caltech.edu
[3] Email: doyle@caltech.edu

proteases that degrades the σ factor. While the real HSR system has a variety of chaperones and proteases that are expressed during heat shock, we follow previous work [4], [5] and substitute a single chaperone (DnaK) denoted $[D]$ and a single protease (FtsH) denoted $[F]$ in our model. Each of these has a corresponding mRNA responsible for its translation ($[M_D]$ and $[M_F]$, respectively), and the system is controlled by a single σ factor (σ^{32}) denoted $[\sigma]$ that regulates RNA transcription and has a constant amount of mRNA ($[M_\sigma]$) at all times. This assumption comes from the fact that $[M_\sigma]$ transcription is regulated by a separate mechanism, independent of what we study in this model. We ignore the effects of temperature on the protein synthesis rates, which is admittedly a simplification but is consistent with prior modeling work by domain experts [4], [5].

Before heat shock occurs, much of the σ factor is inactivated by DnaK. This occurs because σ^{32} contains a region that closely resembles an unfolded protein, so any DnaK not in the process of refolding proteins will bind to σ^{32}, effectively sequestering it. The protease FtsH will target sequestered σ factor for degradation, effectively giving it a high turnover rate (i.e. rapid production and degradation). Both provide feedback control.

Upon heat shock in the cell, two key events occur. The first is an increase in unfolding rate of proteins $k_{un}(T_{high}) = \delta \cdot k_{un}(T_{low})$, and the second is an increase in σ factor production represented by the parameter $\eta(T)$ in translation of σ, which ultimately leads to the increased production of DnaK and FtsH. This change in η is mediated by a temperature-dependent change to the structure of the σ^{32} mRNA, caused by a melting of certain bonds that result in the RNA being more accessible to ribosomes. This acts as a direct temperature sensor in the HSR system. Increases in unfolded protein cause an immediate increase in free σ factor, because the high concentration of unfolded protein displaces any sequestered σ factor bound to DnaK. This increase in free σ factor, along with the increased production rate, cause a fast spike in production of DnaK and FtsH. Because the σ factor is no longer sequestered it will no longer be degraded by FtsH, further boosting its net production. Once DnaK has reduced unfolded protein to pre-shock levels, excess chaperone will re-sequester σ factor and facilitate degradation via protease, reducing production of both DnaK and FtsH. These mechanisms act as two negative feedback loops between σ factor and the proteins it regulates.

In fig. 1 we present a block diagram of the full system containing a single feedforward mechanism and two feedback loops. A more detailed explanation of the biology of the HSR system can be found in [6]. Our deterministic model of the system consists of 6 ODEs:

$$[\dot{\hat{\sigma}}] = k_p\eta(T)[M_\sigma] - k_{pd}[\hat{\sigma}] - k_{pr}[\sigma : D : F] \quad (1)$$

$$[\dot{M}_D] = k_{mD}[\sigma] - k_{md}[M_D] \quad (2)$$

$$[\dot{M}_F] = k_{mF}[\sigma] - k_{md}[M_F] \quad (3)$$

$$[\dot{\hat{D}}] = k_p[M_D] - k_{pd}[\hat{D}] \quad (4)$$

$$[\dot{F}] = k_p[M_F] - k_{pd}[F] \quad (5)$$

$$[\dot{\hat{P}}_{un}] = k_{un}(T)([P_{tot}] - [\hat{P}_{un}]) - k_f[P : D], \quad (6)$$

and 6 algebraic relationships:

$$[\sigma : D] = K_{\sigma D}[\sigma][D] \quad (7)$$

$$[\sigma : D : F] = K_{\sigma F}[\sigma : D][F] \quad (8)$$

$$[P : D] = K_{PD}[P_{un}][D] \quad (9)$$

$$[\hat{\sigma}] = [\sigma] + [\sigma : D : F] + [\sigma : D] \quad (10)$$

$$[\hat{D}] = [D] + [P : D] + [\sigma : D] + [\sigma : D : F] \quad (11)$$

$$[\hat{P}_{un}] = [P_{un}] + [P : D]. \quad (12)$$

Note that $[\dot{\cdot}]$ denotes the total quantity of the protein, $[\cdot]$ represent the free (i.e. unbound) quantity of the protein, and $[\cdot : \cdot]$ represents complexed proteins. We provide tables describing all variables (table I) and parameters (table II) of the system in section V-A for reference. This model is based on a more complex one proposed by El-Samad *et al.* [4], which contains a mix of 31 algebraic and differential equations. One of the core assumptions that makes both our model and previous ones tractable is that fast dynamics (e.g. biochemical interactions) are assumed to be at quasi-steady state and that no delays are explicitly modeled. In section V we discuss how these assumptions make it difficult to study certain properties of the HSR, such as the relationship between feedforward response and delays in the system. While this likely plays an important role in biology, our model likely undersells the role of planning (here feedforward) in the HSR system.

Our model not only has fewer equations, but also has the property that the algebraic relationship can be explicitly approximated, whereas in the original work these constraints had to be implicitly solved. In section V-B we explain in more detail how we derive this model from the original and justify our approximations. The simplicity not only makes the model tractable for theoretical analysis, but also makes it possible to simulate the system using an explicit solver in MATLAB (as opposed to the implicit solver required for the original model).

We believe that this model is relatively simple while still accurately capturing the complexity of the natural HSR system. Our formulation manages to be simple enough to be analytically tractable, yet complicated enough to have interesting performance trade-offs

(which will be analyze in sections III and IV). Due to the model's relative simplicity, it is possible to analytically approximate the steady-state solutions of most quantities in the system. First we note that $[\sigma]_{ss}$ determines $[\hat{D}]_{ss}$ and $[F]_{ss}$:

$$[M_D]_{ss} = \frac{k_{mD}}{k_{md}}[\sigma]_{ss} \implies [\hat{D}]_{ss} = \frac{k_p}{k_{pd}}\frac{k_{mD}}{k_{md}}[\sigma]_{ss},$$
$$[M_F]_{ss} = \frac{k_{mF}}{k_{md}}[\sigma]_{ss} \implies [F]_{ss} = \frac{k_p}{k_{pd}}\frac{k_{mF}}{k_{md}}[\sigma]_{ss}.$$

Next, if we are in the regime where the system is capable of efficient refolding ($[P_{tot}] \gg [P_{un}]_{ss}$) and that feedback results in most DnaK being bound to unfolded protein ($[P:D] \approx [\hat{D}]$), then we get the relationship

$$[\hat{D}]_{ss} \approx \frac{k_{un}(T)}{k_f}[P_{tot}]. \tag{13}$$

This yields the relationship

$$[\sigma]_{ss} \approx \frac{k_{pd}}{k_p}\frac{k_{md}}{k_{mD}}\frac{k_{un}(T)}{k_f}[P_{tot}]. \tag{14}$$

We note that this independence of $[\sigma]_{ss}$ and the protease rate k_{pr} was observed experimentally in [6].

Next we can see from eq. (1), with the assumption that the intrinsic protein degradation rate is much slower than active degradation via FtsH ($k_{pd} \ll k_{pr}$), that

$$[\sigma:D:F]_{ss} \approx \frac{k_p}{k_{pr}}\eta(T)[M_\sigma]. \tag{15}$$

Finally, if we assume that there is enough FtsH such that it will bind to most sequestered σ factor ($[\sigma:D] \ll [\sigma:D:F]$, or equivalently $K_{\sigma F}^{-1} \ll [F]$), we can combine these terms to see that $[\hat{\sigma}]_{ss} \approx [\sigma]_{ss} + [\sigma:D:F]_{ss}$.

The result of this design is a fast system that is able to quickly create new DnaK proteins when heat shock first occurs, and is then able to quickly adapt down once the system has been stabilized (see green curves in figure fig. 2). In the next section we will explore in more depth how and why this architecture performs so well.

III. ARCHITECTURE

Here we can observe some interesting properties of eqs. (14) and (15). First we note that, when the assumptions of the previous section hold, the biochemical parameters $K_{\sigma F}$, $K_{\sigma D}$, and K_{PD} *do not* appear anywhere in the steady state equations. This tells us that the system in some sense abstracts away the particulars of the binding kinetics and makes the system purely dependent on the topology of the network at the biochemical level. This is largely a result of assumptions regarding time scales (i.e. that biochemistry is much faster than gene expression) and concentration scales (i.e. the inequalities that allowed us to derive eqs. (14) and (15)). While it is not obvious from the physics of the system that these

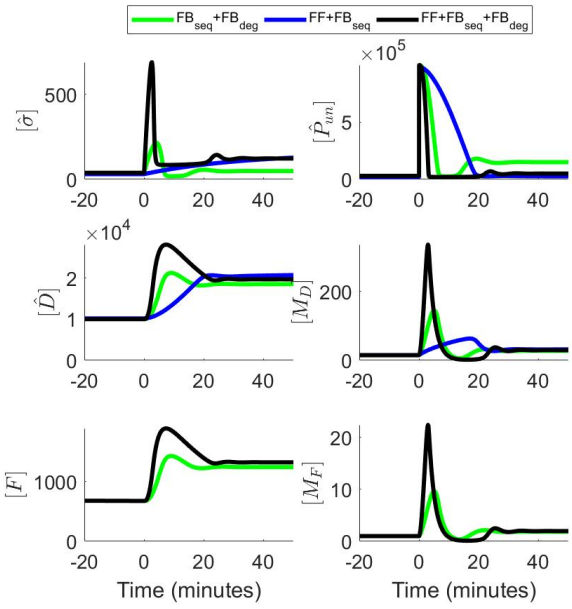

Fig. 2. The dynamics of the architectures described in section III. Here we see that, for a given set of parameters, the full HSR system architecture responds much more quickly than the simpler designs. We choose parameters to match pre-stimulus steady-state values of $[\hat{\sigma}]$ and $[\hat{D}]$ as closely as possible. For all architectures we use parameters as described in table II. We note that the dynamics of $[F]$ match almost exactly to those of $[\hat{D}]$ up to scaling (the same is true of their respective mRNAs), this is because our simple model of protein synthesis leads to these species having the same dependence on $[\sigma]$, up to parameter scaling. We note that the bottom two panels do not have blue trajectories because $[F]$ does not appear in that architecture.

assumptions all must hold simultaneously, it seems to be the case that the HSR system evolved to operate in a regime where they are correct.

Second we note that, since $[\sigma]_{ss}$ and $[\sigma:D:F]_{ss}$ reflect different aspects of the overall concentration of σ factor, we would intuitively expect them to be governed by many of the same parameters. In fact we see that, with the exception of the protein synthesis rate k_p, the parameters that *do* appear in both eqs. (14) and (15) are entirely non-overlapping. When we consider the temperature dependent parameters $\eta(T)$ and $k_{un}(T)$, we see that all of the feedforward architecture (i.e. direct measurement and response to changes in temperature caused by the parameter η) is encapsulated by $[\sigma:D:F]_{ss}$ in eq. (14), and all of the feedback architecture (measurement of and response to changes in $[P_un]$, rather than T directly, mediated by k_{un}), is reflected by $[\sigma]_{ss}$ in eq. (15).

If we now think about the dynamics of the system (see curves in fig. 2), this separation is a nice feature. We note that, to compare different architectures, we choose parameters such that the steady-state values of $[\hat{\sigma}]$ are the same across simulations. When heat shock first occurs,

the fast increase in production via $\eta(T)$ and decrease in degradation via FtsH will cause a spike in σ factor, allowing for a fast initial refolding response. Because these terms do not appear in the steady-value of $[\hat{\sigma}]_{ss}$ in eq. (14), it effectively allows for tuning of dynamics independently of steady-state response in the system. We note that the approximations presented here do not capture the precise value of $[P_{un}]_{ss}$ accurately, because eq. (13) implies $[P_{un}]_{ss} = 0$. The simulations in fig. 2 show that $[P_{un}]_{ss}$ has a small but non-zero steady-state value.

To provide contrast to our analysis so far, we will now examine some alternative designs for the HSR system. We will describe first some qualitative features of the different designs, and then in the next section delve into more quantitative comparisons.

A. Sequestration and Degradation Feedback Architecture

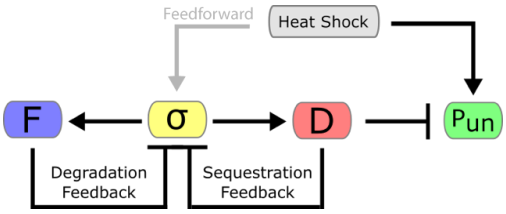

Fig. 3. Block Diagram of the Sequestration and Degradation Feedback Architecture

This architecture assumes that there are two feedback control mechanisms in response to heat shock that respectively sequester σ factor with DnaK, and then degrade it (see fig. 3):

$$\begin{aligned}
[\dot{\hat{\sigma}}] &= k_p\eta[M_\sigma] - k_{pd}[\hat{\sigma}] - k_{pr}[\sigma : D : F] \\
[\dot{M}_D] &= k_{mD}[\sigma] - k_{md}[M_D] \\
[\dot{M}_F] &= k_{mF}[\sigma] - k_{md}[M_F] \\
[\dot{\hat{D}}] &= k_p[M_D] - k_{pd}[\hat{D}] \\
[\dot{F}] &= k_p[M_F] - k_{pd}[F] \\
[\dot{\hat{P}}_{un}] &= k_{un}(T)([P_{tot}] - [\hat{P}_{un}]) - k_f[P : D],
\end{aligned}$$

This system incorporates dynamics analogous to those in eqs. (1), (2), (4), (6), (9), (11) and (12). The only difference is that in this case the translation rate of σ factor, η, does not increase with temperature. This architecture effectively has a lower σ^{32} translation rate than an architecture with the feedforward mechanism.

We see in fig. 2 that this system has a somewhat slower adaptation time to that of the system with full regulation, and the steady state concentration of $[\hat{P}_{un}]$ is higher, implying imprecise adaptation. When feedback, is present there is more than enough total σ factor to produce the requisite amount of DnaK to refold proteins. In the absence of the feedforward response, it is the case that most of the σ factor is needed to produce as much DnaK as possible, so the small amount that is sequestered and degraded has a significant impact on the steady-state level of $[\hat{P}_{un}]$. In this sense, it seems that a significant role of the feedforward loop is to ensure that there is enough total σ factor to balance the effects of degradation feedback. Because $k_{pr} \gg k_{pd}$, even small amounts of sequestered protein can lead to reduced levels of DnaK when the system is saturated, consequently increasing the final level of unfolded protein.

B. Feedforward and Sequestration Feedback Architecture

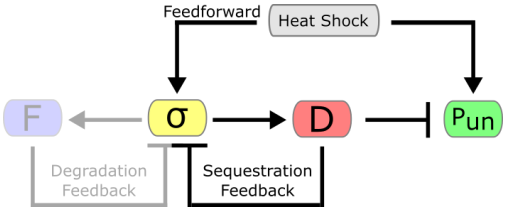

Fig. 4. Block Diagram of the Sequestration Feedback Architecture

This architecture contains only a single feedback loop where DnaK binds to σ factor, sequestering it (see fig. 4). We also add in the effect of the feedforward mechanism, which essentially provides an immediate response to temperature changes via an increase to the translation rate of σ factor. In terms of the model, we simply modify the equation for $[\hat{\sigma}]$ so that η is temperature dependent and $k_{pr} = K_{\sigma F} = 0$,

$$[\dot{\hat{\sigma}}] = k_p\eta(T)[M_\sigma] - k_{pd}[\hat{\sigma}]. \tag{16}$$

This is almost identical to the full HSR system described in section II, except that it lacks the degradation feedback loop mediated by the protease FtsH. This feedback serves to regulate the free amount of $[\sigma]$, but has no direct affect on the total amount $[\hat{\sigma}]$. We can use the same argument from section II to see that $[\sigma]_{ss}$ will be approximately the same as in eq. (14):

$$[\sigma]_{ss} \approx \frac{k_{pd}}{k_p}\frac{k_{md}}{k_{mD}}\frac{k_{un}(T)}{k_f}[P_{tot}],$$

And the complex will have the form

$$\begin{aligned}
[\sigma : D]_{ss} &= [\hat{\sigma}]_{ss} - [\sigma]_{ss} \\
&\approx \frac{k_p}{k_{pd}}\left(\eta(T)[M_\sigma] - \frac{k_{un}(T)}{\gamma}[P_{tot}]\right),
\end{aligned} \tag{17}$$

where

$$\gamma = k_f\frac{k_p^2}{k_{pd}^2}\frac{k_{mD}}{k_{md}}.$$

We can think of this feedback system as serving two purposes. First, we see that it provides the benefit generally seen in feedback systems, namely that the signal level (in this case $[\sigma]$) can be directly coupled to the disturbance that it is trying to compensate for, $k_{un}(T)$. This feedback mechanism can also serve to provide fast response for the system, in that a large amount of σ factor can be stored in the complex $[\sigma : D]$ and quickly released during heat shock (see the blue curves in fig. 2).

The downside to this architecture is that it is, in a sense, leaky. We see from eq. (16) that the time scale for $[\hat{\sigma}]$ to reach steady state is determined by the protein degradation rate k_{pd}, which is typically quite slow (we use $k_{pd} = .03\text{min}^{-1}$, corresponding to a dilution-limited time scale of about 30 minutes). If the system has a low initial amount of σ factor, (corresponding to a small value of $[M_\sigma]$), then the time to refold proteins may be comparable to the scale of k_{pd} (as seen in the blue curves in fig. 2). Since the cell division time of *E. coli* is approximately 20-30 minutes, it would likely be dangerous (if not lethal) for the cell to respond on such a long time scale.

Alternatively, the cell could have a high initial amount of $[\hat{\sigma}]$ (large $[M_\sigma]$), and a correspondingly large amount of $[\sigma : D]$. This would allow for fast response, however $[\hat{\sigma}]$ would still reach steady state on time scale set by k_{pd}. This implies that after the cell has already adapted to heat shock, it would still be synthesizing σ factor unnecessarily and simply sequestering it into the complex $[\sigma : D]$, as can be seen by the fact that eq. (17) is a function of both $\eta(T)$ and $k_{un}(T)$.

This is in contrast from the full model in section II, where sequestered and free σ factor are governed by separate parameters (see eqs. (14) and (15)). Intuitively, this is a result of the protease providing much faster degradation than would intrinsically be seen in a cell (i.e. $k_{pr} \gg k_{pd}$).

C. Natural Architecture

The simpler architectures described in the previous sections serve to motivate the benefits of the natural HSR system (seen in fig. 1). We see that the structure of eqs. (1) to (12) allows for a layering of what we might consider planning (the terms in the feedforward loop $\eta(T)$ and $[M_\sigma]$, contained in the expression for $[\sigma : D : F]_{ss}$ in eq. (15)) and control (the terms in the feedback loop k_{mD}, k_{md}, and k_f, contained in the expression for $[\sigma]_{ss}$ in eq. (14)).

The independence of parameters in these layers allows for evolution to tune the system's dynamics independent of steady-state expression levels, a powerful feature that likely has been taken advantage of over the billions of years the system has been in place. In the next section, we will explore parametric variations in these different architectures to gain a better understand of the trade-offs that constrain their performance.

IV. TRADE-OFF ANALYSIS

In this section we will explore quantitative trade-offs in performance for the HSR system. While the primary goal of the system is to refold proteins, we can see from fig. 2 that simply achieving a good steady state does not fully characterize the performance of the different architectures. We see that, while all of the architectures are capable of achieving a good steady state, they do so on different time scales and with different levels of efficiency. We define the response time to be the amount of time between when heat shock occurs, and when $[\hat{P}_{un}]$ first comes within 5% of $[\hat{P}_{un}]_{ss}$. We define the inefficiency of the response to be:

$$\text{Inefficiency} = 1 - \frac{[P : D]_{ss}}{[\hat{D}]_{ss}},$$

the fraction of $[\hat{D}]$ that is not being utilized for protein folding at steady state. Intuitively we might think of the most efficient response as the one that produces exactly as much $[\hat{D}]$ as there is $[\hat{P}_{un}]_{ss}$. Any excess $[\hat{D}]$ is not refolding proteins, and is thus considered to be in excess. An ideal architecture would not only be able to refold proteins efficiently, but would also minimize response time and inefficiency.

In fig. 5 we explore these trade-offs by randomly varying the parameters $[M_\sigma]$ and $K_{\sigma D}$, and examining the performance of architectures that refold at least 85% of proteins within 100 minutes. The specifics of these thresholds are somewhat arbitrary, but the key idea is to only look at parameter sets where the HSR system is at least somewhat functional. Intuitively, it does not matter how fast or efficient a circuit is if it is not able to perform the primary task of refolding proteins. The reasoning behind our choice of $[M_\sigma]$ and $K_{\sigma D}$ as the varying parameters is that they directly affect σ factor dynamics and are not global parameters of the cell (like k_p and k_{pd}). We omit varying k_{pr} and $K_{\sigma F}$, the respective protease degradation and binding rates, because they are only present in the full architecture and would not affect the simpler architectures.

We see in fig. 5A that there is a clear separation between the architectures with degradation feedback (green and black points) and the one without it (blue points). This is due to the fact that much of the fast response time of the system is tied to the degradation mechanism, so in order for the architecture without it to respond quickly, it must greatly overproduce DnaK and consequently be highly inefficient.

One might also consider efficiency in terms of the rate at which a circuit incurs metabolic overhead. It is likely

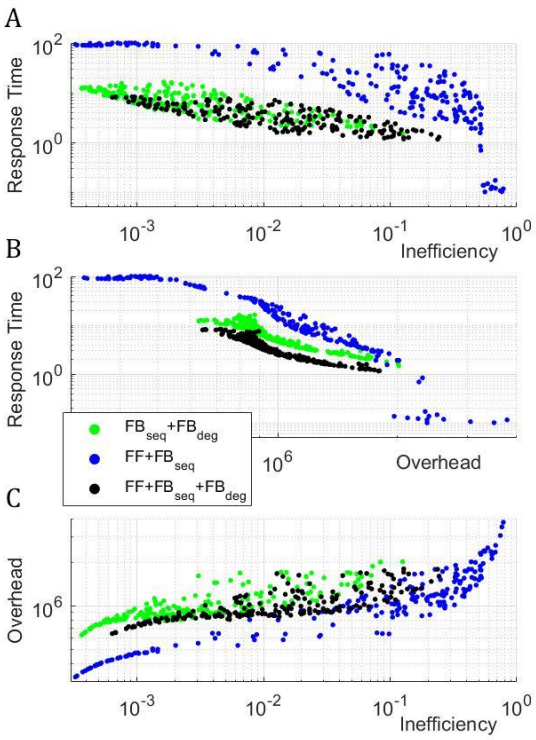

Fig. 5. Trade-offs across architectures. Here we simulate the dynamics of each architecture for 1000 random parameter sets, where we sample values over $M_\sigma \in [1, 100]$ and $K_{\sigma D} \in [10^{-4}, 10^{-2}]$, with all other parameters kept as they are in fig. 2. These two parameters allow us to tune both the total amount of σ factor and the strength of the sequestration feedback loop (for the architectures with feedback). Each panel shows a different pairwise tradeoff between performance metrics, providing insight into the robustness of the qualitative features discussed in section III.

that cells face heat shock rarely, so it makes sense to think about how much ATP (the common unit of energy in the cell) per minute a given architecture uses in the absence of heat shock. While precise numbers for this may be difficult to come by, a reasonable proxy can be calculated easily in terms of how much ATP is spent on protein production, as this is likely the dominant metabolic expense to the cell. Using estimates of ATP use as a function of protein size, we estimate a cost of $\alpha = 2400 \frac{\text{ATP}}{\text{protein}}$ for large proteins (DnaK and FtsH) and $\beta = 1200 \frac{\text{ATP}}{\text{protein}}$ for small proteins (σ factor) [7]. The overhead rate is then computed from the steady-state production rates in the absence of heat shock,

$$\text{Overhead} = \beta k_p \eta(T_{low})[M_\sigma] + \alpha k_p [M_D]_{ss} + \alpha k_p [M_F]_{ss}.$$

We see in fig. 5B that the more complex architectures outperform the simpler ones. First we see again that the systems with degradation outperform the one without it. Because the feedforward mechanism allows for a very fast change to the translation rate of σ factor, it is the case that for a given protein overhead the feedforward system systematically responds more quickly to heat shock. The few cases that respond quickly but with high overhead also likely sacrifice a great deal of resources to overproduce DnaK and are effectively behaving as an open-loop response.

In fig. 5C we see that the full HSR system actually performs worse in terms of the trade-off between overhead and inefficiency. This is because the simpler architecture that perform well in this trade-off (i.e. those that are efficient and cheap), have the worst response time. This tells us that evolution may have selected strongly for architecture that optimizes response time, at the expense of some performance on the other two metrics.

V. DISCUSSION

In this work, we presented a novel theoretical perspective on the architecture and design of the heat shock response system in *E. coli*. We showed that the HSR system seen in nature has many desirable properties, such as layering, speed, and efficiency. Further, we showed that simpler hypothetical architectures for the HSR system are not sufficient to match the response of the full architecture.

More generally, we believe that the HSR system serves as an ideal example of engineering trade-offs in a biological setting. While several architectures are capable of performing the task of refolding proteins, there are many other performance metrics that may be subject to selective pressure. In engineered control systems stability is a primary goal, however properties such as speed, robustness, and cost are equally important for a system to be functional in the real world. Similar constraints are likely pervasive in biology, leading to the apparent complexity of many biological systems beyond what would naively be necessary.

Our analysis shows, through a combination of theory and simulation, that it is possible to systematically investigate the role of complexity in biology. In the future, we plan to expand our model to incorporate stochastic effects and a more nuanced model of heat shock. While our deterministic model was sufficient to ask many interesting questions, it is not able to tell us how noise in the system affects performance. Because σ factor is often at very low copy number, it is likely that stochastic effects play an important role and may present trade-offs that are hidden in the deterministic setting.

Similarly, our model likely hides some of the performance gain from the feedforward system. In our model, σ factor translation and heat shock occur at exactly the same time. In reality it is plausible that heat shock is gradual and that the feedforward mechanism is activated before a large accumulation of unfolded proteins occurs. If this is the case, then the feedforward mechanism likely

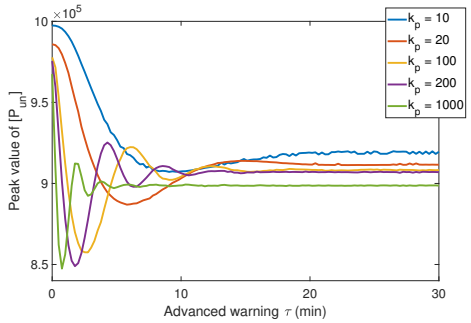

Fig. 6. The effect of advanced warning. Here we model advanced warning as a time delay τ between when σ^32 mRNA first initiates translation and when proteins start to become misfolded. In each simulation, k_p, k_{pd}, and k_{pr} are scaled relative to their nominal respective values of $k_p = 20\text{min}^{-1}$, $k_{pd} = .03\text{min}^{-1}$, and $k_{pr} = 5\text{min}^{-1}$ such that the steady-state concentrations remain the same across simulations. For all systems advanced warning reduces the peak transient value of $[P_{un}]$, however for systems with higher turnover rates the benefit is larger and peaks at a lower value of τ. The oscillations that appear for higher turnover rate systems are likely the related to when the peak value of $[P_{un}]$ occurs relative to the peak level of $[DnaK]$.

serves a more direct planning role for the system than is apparent in the simulations we have presented so far.

To study the effect of this sort of advanced warning, we repeated our previous simulations with addition of a parameter τ that represents how much earlier σ^{32} mRNA initiates translation before unfolded protein begin to accumulate. We observe in fig. 6 that advanced warning does in fact reduce the peak transient levels of $[P_{un}]$, however the benefit of increasing advanced warning levels off roughly around $\tau > 5$ mins. To study how this effect interacts with parameters, we scale k_p, k_{pr}, and k_{pd} such that the steady-state concentrations of all total species remains the same.

While it is difficult to analytically study the role this sort of discrete delay in the dynamics, these simulations result yield insight into how biology might tune parameters to improve transient performance. It appears that when dynamics are slow (low k_p) there are moderate benefits to long advanced warning. For fast dynamics the benefits are larger and require less warning. Since we are essentially tuning the turnover rate of all proteins in the system, this implies that the systems which get the most out of this sort of warning are those that have high gains and thus are less efficient.

We find it encouraging that even a relatively simple model like that of the HSR system yields dynamics are quite complex and difficult to analyze in general. While we derived some approximate results, the nonlinearities in the system make it difficult to write down closed-form solutions to the various steady-states of the system in general. Further, these structural nonlinearities make it challenging to say anything precise about the dynamics of the system. Because the heat shock radically disturbs the system, linearizations that work well locally do a poor job of describing the global dynamics. We hope that future will yield deeper theoretical insights into the types of complex control systems that are pervasive in biology.

APPENDIX

A. Description of Variables and Parameters

Variable	Description	Initial Condition
σ	The σ factor σ^{32}, regulates transcription of HSPs	30
D	The chaperone DnaK, responsible for refolding proteins	10,000
F	The protease FtsH, responsible for degrading σ factor	1,000
M_D	DnaK mRNA, responsible for translation of DnaK	10
M_F	FtsH mRNA, responsible for translation of FtsH	3
P_{un}	Unfolded protein in the cell	30,000

TABLE I

TABLE CONTAINING VARIABLE DESCRIPTIONS

Here we present tables describing the variables and parameters of the HSR system, along with descriptions and initial conditions/parameter values used in typical simulations. We followed parameters as much as was possible from [4], with the only difference being the translation rates k_{mD} and k_{mF} which had to be fit to typical steady-state values seen in their simulations. The reason for this is that these two parameters in our model reflect a large number of processes in their original model (mostly involving the binding of RNA polymerase), which we found could be simplified while still faithfully reproducing the dynamics seen in [4].

B. Model Reduction and Assumptions

In [4] and [5], detailed models of the HSR system are presented which aim to capture all relevant cellular processes at a mechanistic level. These models are quite complex (each containing on the order of 30 equations), and incorporate biochemistry, gene expression, and transcription/translation. Since these processes occur on vastly different time scales, the original models make a quasi-steady state assumption and treat fast processes as if they are algebraic (rather than differential) equations.

This makes simulations tractable, however the algebraic constraints were so complex that they could only reasonably be simulated with an implicit Differential-Algebraic Equation (DAE) solver. The shear number of equations makes it difficult to make an simplifications that would allow for analytic approximations to the algebraic constraints. To simplify the system, we observed that much of the complexity of the original model stems

Parameter	Description	Value
k_p	Translation rate of proteins	20min^{-1}
k_{pd}	Intrinsic degradation rate of proteins	0.03min^{-1}
k_{pr}	Protease degradation rate of σ factor	5min^{-1}
k_{mD}	Transcription rate of DnaK mRNA	0.45min^{-1}
k_{mF}	Transcription rate of FtsH mRNA	0.03min^{-1}
k_{md}	Degradation rate of mRNA	0.5min^{-1}
$K_{\sigma D}$	Dissociation constant for $[\sigma : D]$ binding	$\frac{1}{400}M^{-1}$
$K_{\sigma F}$	Dissociation constant for $[\sigma : D : F]$ binding	$\frac{1}{400}M^{-1}$
K_{PD}	Dissociation constant for $[P : D]$ binding	$\frac{1}{400}M^{-1}$
$[P_{tot}]$	Total amount of protein in the cell	$2 \cdot 10^6$
$[M_\sigma]$	Amount of σ factor mRNA in the cell	10
$\eta(T_{low})$	Pre-shock translation rate of σ factor	0.35min^{-1}
$\eta(T_{high})$	Post-shock translation rate of σ factor	1.75min^{-1}
k_f	Protein refolding rate	15000min^{-1}
$k_{un}(T_{low})$	Pre-shock protein unfolding rate	75min^{-1}
$k_{un}(T_{high})$	Post-shock protein unfolding rate	$\delta \cdot 75\text{min}^{-1}$
δ	Heat shock magnitude	2

TABLE II

TABLE CONTAINING PARAMETER DESCRIPTIONS AND NOMINAL VALUES

from the author's detailed description of transcriptional regulation. Because they explicitly σ factor binding to RNA polymerase and the consequent binding of polymerase to various sites on the genome, about 1/3 of the equations have little to do with the actual heat shock response and mostly govern transcriptional regulation.

We observed that, so long as RNA polymerase and promoter regions are not saturated, all of these relationships could be captured simply by the level of free σ factor in the cell. Additionally we observe that, given the parameter values observed in biology, FtsH is generally far in excess of σ factor. This means that we can safely ignore conservation equations for FtsH, further simplifying the model. these assumptions yield the model described by eqs. (1) to (12). In our model, the only equations that have implicit dependences are eqs. (10) to (12). Combining these equations with eqs. (7) to (9), we can get the relationships:

$$[\sigma] = \frac{[\hat{\sigma}]}{1 + K_{\sigma D}[D](1 + K_{\sigma F}[F])},$$

$$[D] = \frac{[\hat{D}]}{1 + K_{PD}[P_{un}] + K_{\sigma D}[\sigma](1 + K_{\sigma F}[F])},$$

$$[P_{un}] = \frac{[\hat{P}_{un}]}{1 + K_{PD}[D]}.$$

These equations all depend on each other, so analytical expressions for $[\sigma]$, $[D]$, and $[\hat{P}_{un}]$ would require us to solve a system of three nonlinear equations. Fortunately, we can simplify the system by making assumptions based on biological information. First we assume $[\sigma : D] \ll [P_{un} : D]$, or equivalently $K_{\sigma D}[\sigma] \ll K_{PD}[P_{un}]$. This gives us a simpler equation for free DnaK:

$$[D] = \frac{[\hat{D}]}{1 + K_{PD}[P_{un}]}.$$

Now we can solve for $[\sigma]$ if we know $[D]$, and $[D]$ if we know $[P_{un}]$. Using the equations for $[D]$ and $[P_{un}]$, we get a quadratic with the solution:

$$[P_{un}] = \frac{-\alpha + \sqrt{\alpha^2 + 4\beta}}{2},$$

$$\alpha = K_{PD}^{-1} - [\hat{P}_{un}] + [\hat{D}], \beta = \frac{[\hat{P}_{un}]}{K_{PD}}.$$

Since we can solve for $[P_{un}]$ independently of $[D]$ and $[\sigma]$, we can solve the entire system explicitly, give values of $[F]$, $[\hat{D}]$, $[\hat{\sigma}]$, and $[\hat{P}_{un}]$ determined by the ODES in eqs. (1) to (6). With this result, we were able to simulate the model with the MATLAB ode15s function, and have it match almost exactly to the results in [4].

ACKNOWLEDGMENT

We thank Nikolai Matni and Fangzhou Xiao for their helpful feedback and discussion while writing this paper.

The project or effort depicted was or is sponsored by the Defense Advanced Research Projects Agency (Agreement HR0011-17-2-0008) and AFOSR. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

REFERENCES

[1] J. C. Doyle and M. Csete, "Architecture, constraints, and behavior," *Proceedings of the National Academy of Sciences*, vol. 108, no. Supplement 3, pp. 15 624–15 630, 2011.

[2] O. Shoval, H. Sheftel, G. Shinar, Y. Hart, O. Ramote, A. Mayo, E. Dekel, K. Kavanagh, and U. Alon, "Evolutionary trade-offs, pareto optimality, and the geometry of phenotype space," *Science*, vol. 336, no. 6085, pp. 1157–1160, 2012.

[3] F. A. Chandra, G. Buzi, and J. C. Doyle, "Glycolytic oscillations and limits on robust efficiency," *Science*, vol. 333, no. 6039, pp. 187–192, 2011.

[4] H. El-Samad, H. Kurata, J. C. Doyle, C. A. Gross, and M. Khammash, "Surviving heat shock: Control strategies for robustness and performance," vol. 102, no. 8, pp. 2736–2741, 2005.

[5] H. Kurata, H. El-Samad, R. Iwasaki, H. Ohtake, J. C. Doyle, I. Grigorova, C. A. Gross, and M. Khammash, "Module-based analysis of robustness tradeoffs in the heat shock response system," *PLoS Comput Biol*, vol. 2, no. 7, p. e59, 2006.

[6] E. Guisbert, T. Yura, V. A. Rhodius, and C. A. Gross, "Convergence of molecular, modeling, and systems approaches for an understanding of the escherichia coli heat shock response," *Microbiology and Molecular Biology Reviews*, vol. 72, no. 3, pp. 545–554, 2008.

[7] B. Alberts, *Molecular Biology of the Cell: Reference edition*, ser. Molecular Biology of the Cell: Reference Edition. Garland Science, 2008, no. v. 1. [Online]. Available: https://books.google.com/books?id=iepqmRfP3ZoC