Evolving modular genetic regulatory networks with a recursive, top-down approach

Javier Garcia-Bernardo¹ · Margaret J. Eppstein¹

Received: 6 October 2014 / Revised: 1 July 2015 / Accepted: 10 August 2015 / Published online: 21 August 2015
© Springer Science+Business Media Dordrecht 2015

Abstract Being able to design genetic regulatory networks (GRNs) to achieve a desired cellular function is one of the main goals of synthetic biology. However, determining minimal GRNs that produce desired time-series behaviors is non-trivial. In this paper, we propose a ‘top-down’ approach to evolving small GRNs and then use these to recursively boot-strap the identification of larger, more complex, modular GRNs. We start with relatively dense GRNs and then use differential evolution (DE) to evolve interaction coefficients. When the target dynamical behavior is found embedded in a dense GRN, we narrow the focus of the search and begin aggressively pruning out excess interactions at the end of each generation. We first show that the method can quickly rediscover known small GRNs for a toggle switch and an oscillatory circuit. Next we include these GRNs as non-evolvable subnetworks in the subsequent evolution of more complex, modular GRNs. Successful solutions found in canonical DE where we truncated small interactions to zero, with or without an interaction penalty term, invariably contained many excess interactions. In contrast, by incorporating aggressive pruning and the penalty term, the DE was able to find minimal or nearly minimal GRNs in all test problems.

Keywords Genetic regulatory networks · Synthetic biology · Genetic network inference · Differential evolution

Introduction

A genetic regulatory network (GRN) is a collection of genes that interact with each other and with the environment to govern the expression levels of proteins that help to control cell functions. Genes are encoded in the DNA and produce messenger RNA (mRNA) by a process called transcription, while proteins are produced from mRNA by a process called translation. The rate of transcription can be controlled by the binding of proteins, called transcription factors (TFs), to a gene’s promoter region, which is a region of DNA needed to initialize the transcription of the gene. TFs can act as repressors or activators, decreasing or increasing the transcription rate of mRNA, respectively, when bound to the promoter region (Alon 2006). GRNs produce rich, complex behaviors that, when working together, regulate every biological organism. Genetic networks can be mathematically modeled, and the expression patterns recreated computationally (de Jong 2002).

The field of synthetic biology aims to design and construct biological components with the goal of controlling cellular behavior to achieve a specific function (Mukherji and van Oudenaarden 2009). It is a rapidly growing field with many applications including biofuel production (Dunlop 2011), biological waste management (Gilbert et al. 2003), and bio-sensing (Rajendran and Ellington 2008), among many others (Khalil and Collins 2010). Synthetic GRN circuits can be constructed and tuned to achieve a variety of desired functions (Khalil and Collins 2010). However, designing DNA circuits that can achieve a desired dynamical behavior can be a non-trivial task. Building a GRN requires identifying a large number of parameters that represent interactions between genes and proteins, transcription, translation and degradation rates. Evolutionary algorithms are stochastic global search algorithms that enable one to search this large space of

✉ Javier Garcia-Bernardo
javier.garcia.bernardo@gmail.com

Margaret J. Eppstein
Maggie.Eppstein@uvm.edu

¹ Department of Computer Science, University of Vermont,
Burlington, VT 05405, USA

parameters for combinations that produce a desired output. In an evolutionary algorithm, the fitness of a set of candidate solutions (individuals, where each solution contains different parameters of the model) is calculated, and the best performing individuals are combined in subsequent generations in order to find optimal or near-optimal solutions. Since François started evolving GRNs in 2004 (François and Hakim 2004), researchers have been using evolutionary algorithms to evolve network motifs or synthetic networks with a desired behavior (Drennan and Beer 2006; François and Hakim 2004; Noman et al. 2013b; van Dorp et al. 2013a, b); see (Noman et al. 2013a) for a recent review. However, those studies rarely evolved minimal or near-minimal GRNs. Even in methods that attempted to prune away excess parts of the network (Drennan and Beer 2006; van Dorp et al. 2013a), the number of interactions of the evolving networks increased with the number of generations.

Synthetic biology is also used to help scientists understand how natural biological systems are genetically assembled and how they operate in vivo (Mukherji and van Oudenaarden 2009). The ability to understand the behavior of small GRNs makes it easier to design and produce synthetic circuits that will behave in a predictable manner. Some small biological circuits, such as bistable and oscillatory circuits, are well-understood. Bistable GRNs, those that can exhibit two mutually exclusive states depending on initial conditions, are ubiquitous in nature. For instance, stem cells can remain undifferentiated for years and differentiate when needed, moving from one stable state to another (Wang et al. 2009). A minimal GRN for a bistable switch with monomeric factors is well known (Widder et al. 2009). Similarly, oscillatory GRNs, those whose genes are expressed in cycles, are found in every plant and animal and many bacteria, controlling global behaviors including sleep, stress and light responses (Levine et al. 2013). For example, the protein p53 oscillates when mammal cells are exposed to UV or gamma radiation (Purvis and Lahav 2013). The so-called “repressilator” is a well-known minimal GRN that produces oscillations (Elowitz and Leibler 2000). However, the networks for many natural GRNs are not known or fully understood. The development of single-cell microscopy and high-throughput gene expression analysis has allowed the study of simultaneous expression levels of hundreds of genes under different conditions, and GRN inference from gene expression levels is active area of research (Bar-Joseph et al. 2012; Marbach et al. 2010, 2012) where evolutionary computation is proving useful (Cao et al. 2010; Ruskin and Crane 2011). While our primary motivation is for designing GRNs for synthetic biology, the algorithm introduced in this work could be used for the inference of naturally-occurring GRNs from gene expression data.

In this contribution we use differential evolution (DE), an evolutionary algorithm that is explicitly designed for efficient evolution of candidate solutions represented as vectors of real-valued variables (Storn and Price 1997). Using DE, we evolve small GRNs with four different types of desired time-series behaviors. We experiment with three different approaches for minimizing excess interactions between genes during the evolution. We first assess how well these methods can evolve known minimal GRNs that produce bistable or oscillatory behaviors. We then show how one can incorporate previously identified GRNs as non-evolvable subnetworks into the subsequent evolution of larger (but still minimal or nearly-minimal) modular GRNs for two more complex types of dynamical behaviors. We find that a top-down approach, where we start with many more interactions than needed, enables us to rapidly identify a non-minimal network that contains a desired dynamical behavior. Once the target behavior is detected, we switch the strategy of the DE to a more focused search while also pruning away up to five interactions per candidate solution per generation and increasingly penalizing for the number and strength of interactions. Using this approach we are able to evolve small, often minimal, GRNs. In contrast, the methods that did not incorporate aggressive pruning resulted in GRNs that invariably contained many excess interactions, whether or not fitness was increasingly penalized for the number and strength of interactions.

Materials and methods

Model of gene expression

GRNs are commonly modeled as systems of differential equations with degradation and production terms (Ruskin and Crane 2011; Smolen et al. 2000). Promoter dynamics and mRNA concentrations are usually modeled with Hill functions, as shown in Eq. (1), while protein concentration depends on the concentration of mRNA as shown in Eq. (2). Normalizing mRNA and protein concentrations, the pair of differential equations corresponding to expression for each gene i can be expressed as

$$\frac{dm_i}{dt} = -m_i + \frac{\alpha_i}{1 + K_{1,i}p_1^{n_{1,i}} + K_{2,i}p_2^{n_{2,i}} + \dots + K_{G,i}p_G^{n_{G,i}}} + \alpha_0 \quad (1)$$

$$\frac{dp_i}{dt} = -\beta_i(p_i - m_i) \quad (2)$$

where m_i is mRNA level for gene i , p_i is protein level of the protein transcribed from gene i , α_i is the maximum expression of the promoter, β_i is the translation rate, α_0 is

the basal level of gene expression, and G is the number of potentially interacting genes (Elowitz and Leibler 2000). Together, $K_{j,i}$ and $n_{j,i}$, determine the strength of an interaction. Each $n_{j,i}$ determines the sign and strength of the interaction between the protein j and the promoter region of the gene i that turns on expression of m_i . and $K_{j,i}$ is the strength of repression/activation once the protein is bound to the promoter. When $n_{j,i}$ is zero, there is no feedback of protein p_j on gene i and therefore no interaction. Positive $n_{j,i}$ indicate that the protein p_i is a repressor of the gene i , while negative $n_{j,i}$ indicate that the protein p_i is an activator of the gene i . The specific value of $n_{j,i}$ depends on the interaction between transcription factors, the number of protein monomers, the DNA accessibility and the number of promoter binding sites (Sheinman and Kafri 2012). The representation of gene expression with Hill functions accurately models the biological system and has been broadly used (Smolen et al. 2000), but it requires many parameters. Note that Eq. (1) assumes effectively models the action of transcription factors with an OR gate. For example, for a gene regulated only by repressors, mRNA production is low if *any* repressor is bound to promoter. Consequently, activators only exert their action if no repressor is bound. Other models, such a competition between activation and repression are possible, but are not explored here.

Using this model of gene expression for simulating the dynamical behavior of a GRN with G genes one must thus solve $2 \cdot G$ coupled differential equations. In this work, we used Matlab's *ode15s* to solve these potentially stiff systems and obtain a time series of simulated protein expression.

For a GRN with G genes there are potentially G^2 interactions, each with two associated coefficients ($n_{j,i}$ and $K_{j,i}$), and each gene has two other associated coefficients (α_i and β_i). Thus, for the inverse problem of evolving GRNs, there are $M \leq 2 \cdot G^2 + 2 \cdot G$ unknown decision variables that must be identified to fully characterize the GRN (the inequality is because some of the potential interactions may be fixed and not subject to evolution).

Evolutionary algorithm

In preliminary experimentation with two evolutionary algorithms designed for optimizing real-valued decision variables, covariance matrix adaptation evolution strategy [CMA-ES (Hansen 2006)] and differential evolution [DE (Storn and Price 1997)], DE proved to be able to handle boundaries more easily and in general performed better for this application. Thus, all reported results here used DE (we used the open-source Matlab implementation of DE available here: <http://www1.icsi.berkeley.edu/~storn/code.html#matl>) with the modifications described below.

Using an iterative process, DE evolves a population of N candidate solutions, each of which is represented as a real-valued individual vector of length M . By creating new candidate solutions using difference vectors of existing solutions in the population (as described in more detail below), step sizes are large when the initial population is scattered but generally decrease as the population begins to converge in various dimensions. Thus, DE naturally follows the contours of the fitness landscape and automatically shifts from global search (exploration) to local search (exploitation) as the population converges. Compared to competing approaches designed for real-valued optimization (such as CMA-ES), DE is simple to implement, requires relatively small populations, has low computational overhead per generation, requires relatively few generations to converge to global optima in a variety of benchmark problems, and performs well even in the presence of correlated decision variables and noise (Price et al. 2005). Consequently, DE has rapidly gained traction in the evolutionary computation community for real-valued optimization (Das and Suganthan 2011).

There are several variants of DE (Storn and Price 1997). In this work, evolutionary runs started with the DE/rand/1/bin strategy, where, for each target vector $\vec{x}_i$ a mutant solution ($\vec{v}_i$) was created by adding the difference of two random population members ($\vec{x}_{r_2}$ and $\vec{x}_{r_3}$) scaled by a factor F , to another random population member ($\vec{x}_{r_1}$), as shown below. All three random vectors are different from each other and from the target vector.

$$\vec{v}_i = \vec{x}_{r_1} + F(\vec{x}_{r_2} - \vec{x}_{r_3}) \quad (3)$$

The vector $\vec{v}_i$ then underwent uniform crossover with the target vector $\vec{x}_i$, using the default crossover rate of 0.8 (meaning that approximately 80 % of the decision variables come from $\vec{v}_i$), and the resulting vector replaces $\vec{x}_i$ in the population if it has better fitness.

We then implemented the following modifications. We performed up to three restarts when the fitness had not improved for ten generations (note that we did count fitness evaluations of aborted starts in the total for the trial run). As soon as the desired dynamical behavior was found in a solution (when raw fitness fell below a pre-specified threshold; see “Raw fitness functions” section for details), the strategy was changed to DE/rand-to-best/1/bin, where each mutant solution was created by adding to each target vector $\vec{x}_i$ the scaled difference between the *best* population member and the target vector and another scaled difference vector of two other random population members, as shown below.

$$\vec{v}_i = \vec{x}_i + F(\vec{x}_{best} - \vec{x}_i) + F(\vec{x}_{r_1} - \vec{x}_{r_2}) \quad (4)$$

At the same time as this strategy shift occurred, we also changed the DE scaling parameter F from 0.8 (default parameter) to 1.2. The strategy shift enabled the DE to

focus the search around the best GRN to speed up the evolution, but increasing F also allowed larger step sizes to permit exploration of other nearby topologies that may be better.

Each of the decision variables was represented as a bounded real variable, uniformly randomly initialized from the ranges shown in Table 1. Following random initialization, some of the G^2 Hill coefficients $n_{j,i}$, were then randomly set to zero to control the density of the initial interconnection network. In this work, we tried evolving GRNs starting from three different levels of initial interconnection density: *dense*, *medium*, and *sparse* (see “Test problems” section for details.)

As the evolution progressed, parameter boundaries were checked and values outside the allowable range were truncated to the maximum (or minimum) of the ranges specified in Table 1. For all experiments reported here, $\alpha_0 = 0.2$. These ranges for the system parameters are in agreement with previous studies (Elowitz and Leibler 2000; Kim et al. 2008), but could easily be modified to evolve other systems with values outside them (e.g., if one desired an ultrasensitive switch, where the Hill coefficient $n_{j,i}$ is larger than 3).

Since the goal of this study was to evolve minimal GRNs, we tried various ways of reducing excess interactions in the evolving solutions. These included (1) penalizing for the number and strength of interactions, (2) aggressively pruning out interactions, and (3) truncating small interaction coefficients to zero. In this paper, we report on three methods that used various combinations of these techniques, as follows.

In two of the methods reported on here, we included the following penalty term:

$$Pen = \frac{gen}{C} \sum_i \sum_j |n_{j,i}| \quad (5)$$

The penalty is proportional to the sum of the absolute values of Hill coefficients $n_{j,i}$, thus penalizing for the total number and strength of interactions. The penalty term is multiplied by the generation number gen , divided by a positive constant C , so that the penalty would increase slowly as the evolution progressed. This constant was set to allow for small contributions to the fitness at early generations and penalties contributing up to 20 % at later

generations. As a rule of thumb, the expected final number of interactions times the maximum number of generations should equal 10 % of the optimal fitness. Because we were minimizing fitness, the penalty term was added to the raw fitness (see “Raw fitness functions” section for details).

ForcedReduction method

In the main method we are proposing, we used the penalty term along with an aggressive pruning step at the end of each DE generation subsequent to the strategy shift to random-to-best, as follows. For each solution that exhibited the desired dynamics fitness, we iterated through each Hill coefficient $n_{j,i}$, temporarily set the coefficient to zero, and if the raw fitness of the modified solution was no more than 15 % worse than the raw fitness of the solution with that coefficient, the zero was kept in the solution vector. i.e., we accepted as much as 15 % degradation in fitness to increase parsimony. This process was continued until we had replaced five coefficients with zeros, we ran out of Hill coefficients to try or the fitness of the modified solution was 10 % worse than the unmodified one (we accepted a 15 % degradation in an individual step but only a 10 % global degradation). Preliminary studies showed that these two thresholds worked well in combination. Note that higher thresholds increase the aggressiveness of the method, and therefore if thresholds are too high it may contribute to premature convergence to a suboptimum. In contrast, lower thresholds decrease the aggressiveness and therefore decrease the speed of the algorithm and the ability to reduce the number of interactions.

NoPenalty method

For comparison, we implemented a second method in which we ran the DE without using the penalty term and without any aggressive pruning. In this method, we simply truncated very small Hill coefficients $n_{j,i}$ (those with absolute values less than 0.5) to zero. This value was chosen to disrupt weak interactions and create parsimonious, stable networks.

Penalty method

Finally, we implemented a third method that was exactly like the *NoPenalty* method except that, in addition to truncating small values, it also used the penalty term from Eq. (5).

We summarize the combinations of techniques incorporated into these three methods in Table 2. Storn and Price (1997) recommend DE population sizes of between 5 and 10 individuals per decision variable. However, since we had so many decision variables (for most problems tested here we had 84 unknowns), and since our fitness

Table 1 Parameter ranges used

Coefficient	Initial range	Bounded range
$n_{j,i}$	$[-3,3]$	$[-3,3]$
$K_{i,j}$	$[1,2]$	$[0.5,5]$
α_i	$[100,500]$	$[0.5,500]$
β_i	$[0.5,5]$	$[0.5,5]$

Table 2 Parsimony strategies used

Method	Penalty term (Eq. 5)	Truncation of small values	Aggressive pruning
<i>ForcedReduction</i>	Yes	No	Yes
<i>NoPenalty</i>	No	Yes	No
<i>Penalty</i>	Yes	Yes	No

function requires numerical integration and so is relatively slow, we opted to use much smaller population sizes of only $N = 25$ individuals, which was found to be sufficient during preliminary experimentation. The only exception to this was for our most difficult test problem (*DualOscillator*, described in “[Test Problems](#)” section), in which we started with a population of $N = 50$ individuals, but then discarded all but the best 25 individuals as soon as the desired dynamical behavior was found and continued the evolution with these $N = 25$ remaining individuals. We recommend 5–10 individual per expected final number of interactions. Values of problem-specific parameters are summarized in Table 3.

All DE trials were started from one of three levels of initial network interconnection density (*dense*, *medium*, and *sparse*); the initial number of non-zero interactions prescribed for each of these three densities, for each of the four test problems, are shown in Table 4.

Test problems

In this study, we first tested each approach to see how well we could evolve known small GRNs for a simple toggle switch (*Bistable*) and an oscillatory circuit (*Oscillator*). We then included one *Bistable* GRN and/or one *Oscillator* GRN as non-evolvable subnetworks in the subsequent evolution of more complex GRNs, to illustrate how the methods can be recursively applied to bootstrap the evolution of increasingly complex, modular GRNs. In one problem, we evolved GRNs that combined non-evolvable

subnetworks that included a toggle switch and an oscillatory circuit (*ConditionalOscillator*), as in (Thomas and Jin 2012); to make this problem more difficult we precluded direct interactions between these two subnetworks. In the other problem, we evolved GRNs with two mutually exclusive oscillators (*DualOscillation*); to make this problem more difficult we only included one copy of a non-evolvable oscillatory circuit.

Raw fitness functions

Assessing whether or not a GRN exhibits a desired dynamical behavior requires using the GRN to simulate time-series data and then computing some metric (referred to as the raw fitness) for how closely that time-series data meets the desired characteristics. Appropriate raw fitness metrics are dependent on the nature of the particular target dynamical behavior. Thus, each of the four test problems had their own unique raw fitness metric, as described below.

Bistable raw fitness

To detect bistability, raw fitness was calculated as (the negative of) the absolute difference in the mean levels of one pre-specified target protein, between two 50-time step simulations. In one simulation the initial levels of the target protein were set to 1, in the other simulation the initial levels were set to 100. The minimum value of the raw fitness was truncated at -15 , so as not to reward GRNs that

Table 3 Parameters of the four test problems

Test problem	DE popsize N	Penalty divisor C	# GRN genes G	# Decision variables M
<i>Bistable</i>	25	250	6	84
<i>Oscill.</i>	25	250	6	84
<i>Cond. Oscill.</i>	25	250	3 Evolved + 5 fixed	36
<i>Dual Oscill.</i>	50–25	2500	6 Evolved + 3 fixed	84

Table 4 Number of non-zero Hill coefficients $n_{j,i}$ (interactions) in initial solution vectors for the three initialization densities

Test problem	Dense initial network	Medium initial network	Sparse initial network
<i>Bistable</i>	33	18	4
<i>Oscill.</i>	33	18	4
<i>Cond. Oscill.</i>	13	8	3
<i>Dual Oscill.</i>	33	18	4

included excess positive regulation of one of the proteins by extra genes. Such excess positive regulation is not necessary for bistability, but does increase the difference between the two protein levels. It is worth noting that the difference between two stable states can thus be easily increased to a desired level by adding in extra activators. When the raw fitness of an evolving solution fell below the threshold of -9 , the correct behavior for the *Bistable* circuit was considered to be found, triggering the strategy shift described in “[Evolutionary algorithm](#)” section.

The other three test problems (*Oscillator*, *ConditionalOscillator*, and *DualOscillator*) each contained an oscillatory component. The presence of oscillatory behavior was assessed using an autocorrelation-based fitness metric, as follows. Auto-correlation is the correlation between a time series and the same time series with a specified time-lag. Here, we simulated time series data from evolving GRNs, starting from initially random protein levels, for 100 time steps; we then computed a vector of normalized unbiased auto-correlations of target oscillatory protein, for time lags from 0 to 50 time steps, using Matlab’s *xcorr* function. A perfectly oscillatory GRN will have a periodic unbiased auto-correlation vector that oscillates between -1 and $+1$. The autocorrelation-based fitness metric was calculated as the sum of the first local minimum in this autocorrelation vector plus two times the sum of the second through fifth local minima (in the positive domain), where the minima are found using the Matlab *findpeaks* function. The first minimum is weighted less to account for an initial transient period before the protein levels settle into oscillatory behavior, but the first minimum is not completely ignored to account for the case where there is only one minimum. Thus, for a perfectly periodic time series with no transient period, this autocorrelation-based fitness metric will be -9 ; damped oscillatory behaviors will have a smaller (but still negative) magnitude. Details of how the specific raw fitnesses for the oscillatory circuits were calculated from this autocorrelation-based fitness metric follow.

Oscillatory raw fitness

Because the *Oscillator* behavior was found easily with our algorithm, the raw fitness was taken to be the autocorrelation-based fitness measure for one arbitrarily chosen protein. When the raw fitness fell below the threshold of -5.5 , the correct behavior for the *Oscillator* circuit was considered to be found, triggering the strategy shift described in “[Evolutionary algorithm](#)” section. This threshold depends on the problem considered and sets the timing of the algorithm, with higher (less negative) thresholds allowing for longer searching times before the strategy shift.

ConditionalOscillator raw fitness

To assess the raw fitness of the *ConditionalOscillator* GRN solutions, two different simulations were run: in the first simulation, the protein and mRNA levels of one of the proteins in the *Bistable* building block were initially set to 1; in the second simulation they were both set to 100. The raw fitness was computed as the difference in the autocorrelation-based fitness measures of the first and the second simulation. When the raw fitness fell below the threshold of -5.5 , the correct behavior for the *ConditionalOscillator* circuit was considered to be found, triggering the strategy shift described in “[Evolutionary algorithm](#)” section.

DualOscillator raw fitness

To assess the raw fitness of the *DualOscillator* GRN solutions, two different simulations were run; in the first simulation, the hard-coded oscillatory building block was initially oscillating, and in the second simulations its protein levels were initially flat. The raw fitness was computed as the difference between the autocorrelation-based fitness measure of the first and second simulation. When the raw fitness fell below the threshold of -3 , the correct behavior for the *DualOscillator* circuit was considered to be found, triggering the strategy shift described in “[Evolutionary algorithm](#)” section.

Experimental procedures

We conducted 25 repetitions of each the three methods on each of the four test problems, starting from each of three the different densities of interactions in the initial solution vectors, for a total of 900 experimental runs.

Evolutionary runs using *ForcedReduction* were terminated after a maximum of 50 generations; those using the *NoPenalty* or *Penalty* methods were terminated after a maximum of 100 generations to compensate for the fact that *ForcedReduction* had more fitness evaluations per generation when it was actively pruning excess interactions. However, runs were terminated earlier if the number of interactions in the best evolved GRN did not change for five generations when using *ForcedReduction* or for 30 generations when using the other two methods (because the reduction in the number of genes is slower in these), after the desired behavior was found. Our results show that *ForcedReduction* actually required fewer total evaluations to terminate than did either of the other methods, as shown in “[Results](#)”.

Autoregulation (where a gene has a positive or negative feedback on itself) is not necessary for any of the behaviors

we sought, so because we were seeking minimal networks autoregulation was never spontaneously evolved. However, while not strictly necessary to achieve the desired behaviors, autoregulation can help to make certain GRNs more robust (Komiya et al. 2012). Thus, at the end of each evolutionary run, we tested each autoregulation Hill coefficient $n_{i,i}$ with the minimum value of -3 (positive autoregulation) and the maximum value of 3 (negative autoregulation); if either of these values improved the raw fitness then the altered value was kept in the final solution. This only proved beneficial in the case of *Bistable*, given our fitness criteria.

Success criterion

Following the termination of each experimental run, we determined whether or not the run was successful, where only GRNs that were robust in being able to generate the desired dynamical behavior were considered successful. To assess robustness, we performed 100 simulations (50 time steps for the *Bistable* GRNs and 100 time steps for the other three test problems) of the best resulting GRN from a given run. Each simulation started from random concentrations of mRNA and protein (selected from a uniform distribution between 0 and 20 molecules per species). We then assessed whether or not the target behavior was present in each of these simulations. If the number of successful simulations was at or above 90 out of 100, we considered the GRN robust and the run successful.

Results

Of the 900 total runs, 689 (77 %) met our criterion for success, with 422 exhibiting the desired dynamical behavior in all 100 of 100 simulations (Fig. 1). Using the *ForcedReduction* method we were able to successfully evolve small, often minimal, GRNs that produced the four desired behaviors (Fig. 2). For example, in the *Bistable* problem, successful solutions consistently recovered the well-known monomeric bistable switch (Widder et al. 2009), with only two evolved nodes and two evolved interactions. The positive autoregulatory interaction was added in the post-processing step to increase the robustness of the bistability (Fig. 2a). Similarly, in the *Oscillatory* problem, successful solutions invariably found the well-known repressilator (Elowitz and Leibler 2000) shown in Fig. 2b, with three evolved nodes and three evolved interactions. In contrast, in (Drennan and Beer 2006) the repressilator was found as the core of oscillatory genetic networks that were evolved, but those networks contained many excess interactions. We were also able to evolve small GRNs for the two more complex problems. In the

Fig. 1 Histogram of the number of successful simulations (out of 100) of the best evolved GRNs resulting from the 900 experimental runs. Those with at least 90 successful simulations (above the red line) were considered ‘successful’ runs

ConditionalOscillator problem, most successful runs found the nearly-minimal GRN with two evolved nodes and three evolved interactions connecting the hardcoded bistable and oscillatory motifs, shown in Fig. 2c. Results on the *DualOscillator* problem were slightly more variable; two successful evolved solutions that were commonly found are shown in Fig. 1d, one with five evolved nodes and seven evolved interactions, and one with only three evolved nodes and five evolved interactions.

Recall that, in the *Bistable*, *Oscillatory*, and *DualOscillatory* problems, we were starting from six evolvable nodes (so 36 evolvable interactions), with either 33, 18, or 4 of the interaction coefficients initially non-zero, whereas in the *ConditionalOscillator* problem, we were starting from three evolvable nodes (so nine evolvable interactions), with either 13, 8, or 3 of the interaction coefficients initially non-zero (Table 4). The number of genes and interactions in successfully evolved GRNs using the *ForcedReduction* method was nearly insensitive to these initial number of non-zero interaction coefficients, and minimal or near-minimal GRNs were found in all cases (Fig. 3a). The success rate, however, was lower when starting from the sparsest initial interconnection networks, as shown by the shorter white bars in Fig. 3b. In all except the *Bistable* problem, we observed that, when starting from the sparsest initial networks, we often observed an initial increase in the number of interactions before the desired behavior was found, and then the size of the best GRN fluctuated as nearby topologies were explored and found to be better (red solid lines, Fig. 4). Ultimately, the increasing penalty term helps *ForcedReduction* to prune away the excess interactions. In the difficult *DualOscillatory* problem, this also occurred when starting from the initial networks with medium density (e.g., red solid lines, Fig. 4b, c). Thus, runs starting from sparse networks took longer to find the desired behavior than did runs starting from the denser networks; this resulted in fewer successes and, in the case of the two more complex networks, sometimes slightly larger final GRNs (Fig. 3a), although not significantly so ($p > 0.1$).

Fig. 2 Representative evolved small GRNs and their time-series behaviors. *Nodes* represent proteins coded for by specific genes. *Nodes outlined in black* and interactions shown with *solid lines* were evolved. *Nodes with no outline* and interactions shown with *dotted lines* were included as non-evolvable subnetworks. All *nodes* and interactions were evolved in **a** *Bistable*, and **b** *Oscillator*, with the exception that the (non-essential) positive autocorrelation shown in *Bistable* was added in the post-processing phase. In

c *ConditionalOscillator*, two nodes and three interactions were evolved such that protein I_c oscillates when I_a is present but not when I_b is present (here, we externally manipulated I_b at time steps 50 and 100). In **d** *DualOscillator*, two successful solutions are shown, one in which five nodes and seven interactions were evolved, and one in which three nodes and five interactions were evolved. In both solutions, oscillations in protein I_a and C_4 are mutually exclusive (here, we externally manipulated I_c at time steps 50 and 100)

The *NoPenalty* method was also very successful in identifying GRNs with the desired dynamical behavior, with slightly higher success rates than *ForcedReduction* for three of the four test problems, and less sensitivity to the initial sparsity of the networks (compare Fig. 3d to b). However, even though this method truncated small non-zero interaction coefficients to zero, all of the successfully evolved GRNs had many excess interactions, regardless of the initial sparsity of the network (compare Fig. 3c to a, noting the difference in the scale of the y-axis). Note how, on the difficult *DualOscillator* problem, the number of interactions actually increased to about 35 (Fig. 3c; blue dashed lines, Fig. 4) regardless of the initial density of the

network. The robust success rates of this approach indicate that perhaps a hybrid approach that starts with *NoPenalty* and then transitions to *ForcedReduction* may be better than either method alone, yielding high success rates and minimal GRNs, although we have not yet tried this.

When the *Penalty* method was successful, the number of interactions in the evolved GRNs were lower than with the *NoPenalty* method (compare Fig. 3e to c), showing that the penalty term was, indeed, helping to reduce the number of interactions. However, there were still many excess interactions in comparison to the small networks evolved with *ForcedReduction* (compare Fig. 3e to a). Unlike in the other two methods, the number of interactions in successful

Fig. 3 Success rates and final number of interactions in successful trials for each method, initial network density, and test problem. In the left column (a, c, e), the heights of the solid bars show the mean number of interactions in the final GRNs of successful runs and the error bars indicate one standard deviation; note the different scales on

the y-axis. In the right column (b, d, f) we show the percentage of successful trials (out of 25). The first row (a, b) shows results for the *ForcedReduction* method, the second row (c, d) shows results for the *NoPenalty* method, and the third row (e, f) shows results for the *Penalty* method. The legend in a applies to all six plots

GRNs evolved using the *Penalty* method was significantly smaller when starting from the sparsest network in comparison to when starting from the densest network ($p < 0.001$ for all four test problems) because the penalty term slowed down the addition of extra interactions. However, this also made it harder for the *Penalty* method to find successful solutions, resulting in lower success rates than the *NoPenalty* method (compare Fig. 3f to d). Unsuccessful runs occurred when the growing penalty term dominated the raw fitness value and removing interactions was favored over finding and retaining the desired dynamical behavior. To be fair, the dynamically growing penalty term was crafted to work in conjunction with the aggressive pruning of *ForcedReduction*, where there was no need to explicitly prevent the penalty term from dominating raw fitness. One could certainly incorporate such a safeguard for use in the *Penalty* method. However, even in the successful runs with *Penalty*, the method was not capable of removing many of the excess interactions. For example, on the *DualOscillator* problem the number of interactions in successful solutions was high (Fig. 3e) and appears to have plateaued, with no sign of further decrease as the evolution progressed (black dash-dot lines, Fig. 4). This indicates that a less dominant penalty term applied without aggressive pruning may improve the success rate of the *Penalty* method, but would still not be able to achieve the small GRNs evolved by *ForcedReduction*, even when allowed to run for many more generations.

Discussion

Here, we have shown that starting from initially excessively dense interaction networks makes it relatively easy for DE to discover networks that exhibit the desired dynamical behaviors, even when using population sizes that are much smaller than are generally recommended. Excess interactions can then be aggressively pruned away as the evolution continues, even when this entails a temporary degradation in the dynamical behavior. This approach, in conjunction with a dynamic penalty term that increasingly penalizes for the number and strength of interactions as the GRNs evolve, often yields minimal, or nearly-minimal, GRNs that robustly exhibit the desired dynamical behavior when starting from a wide variety of initial conditions. Using this approach, we were able to consistently find well-known GRNs for the oscillatory repressilator (Elowitz and Leibler 2000) and a small bistable switch (Widder et al. 2009). Others who have tried to incorporate pruning of excess interactions (Drennan and Beer 2006; van Dorp et al. 2013a) have not succeeded in evolving minimal networks.

We also demonstrated how one can recursively apply the method to evolve larger, more complex, modular GRNs. This is accomplished by incorporating previously discovered GRN circuits as non-evolvable subnetworks in subsequent runs, so that the evolutionary process may utilize these circuit without having to rediscover them.

Fig. 4 Number of interactions in the best current population member during the evolution for three representative successful runs on the *DualOscillator* problem and three initial network densities. **a** The densest initial networks, **b** initial networks of medium density, and **c** the sparsest initial networks. The legend in the **a** applies to all three plots. Note that the x-axis is number of evaluations, not generations, in order to fairly compare the three methods. In all cases, the final networks exhibited the desired dynamical behavior in 100 out of 100 random simulations

Using this approach, we were able to evolve larger (but still small) GRNs that produced two types of more complex dynamical behaviors by combining oscillatory and/or bistable GRNs. Because we are evolving gene interaction coefficients, the number of decision variables grows

quadratically with the number of evolvable genes. However, by recursively using the method to incrementally build up more complex networks, the DE can focus on only a relatively small number of evolvable genes and interconnections during a given run, thus keeping the method tractable.

We compared this approach to two more canonical versions of DE, using more traditional approaches for removing excess interactions. In both of these methods we truncated small interaction coefficients to zero; one of the methods used the same dynamic penalty term described above and the other used no penalty function. While these methods were often successful in finding GRNs with the desired dynamical behaviors, they invariably yielded non-minimal GRNs that included many excess interconnections, even though there were allowed many more fitness evaluations than the method using aggressive pruning.

In this contribution, we used DE to evolve real-valued Hill coefficients in systems of deterministic differential equations that represented the dynamics of GRNs. However, the concepts of (a) first evolving the correct dynamics in a relatively dense network, (b) subsequently condensing the network by alternating aggressive pruning of excess interactions with continued evolutionary refinement, and (c) using previously identified networks as non-evolvable subnetworks in the subsequent evolution of more complex, modular GRNs are easily generalized to other types of network representations (such as stochastic differential equations or graph-based representations) and other types of evolutionary algorithms.

Our interest is primarily in designing GRN circuits for synthetic biology. We envision using the approach described herein for identifying small GRN circuits that achieve specific desired dynamical behavior, to inform synthetic biologists as they seek to construct actual GRNs from DNA. While in this work we used deterministic simulations and Hill functions based on the independent action of transcription factors through OR gates, other models of gene expression could also be used. For instance, allowing for competition between activators and repressors may facilitate the evolution of other oscillatory motifs, such as the mixed feedback loop (François and Hakim 2005). Similarly, stochastic simulations could be used to find GRNs that are robust to noise or that exhibit noise-driven behaviors. We also believe our approach could prove useful in other domains. For example, using a fitness function that tries to match the behavior of evolving GRNs with real time-series gene expression patterns (e.g., as in (Cao et al. 2010; Ruskin and Crane 2011)) our approach could prove useful in inferring naturally-occurring GRNs. Additionally, our approach could be used for evolving parsimonious topologies for Artificial Neural Networks.

In summary, we present a top-down, evolutionary approach that can be recursively applied to evolve increasingly complex, modular GRNs with few, if any, excess interactions. The approach uses DE to find dense GRNs that exhibit desired dynamical behaviors, interlaced with a procedure that aggressively prunes away excess interactions, favoring parsimonious GRNs even at the expense of small temporary degradations in raw fitness. Using this method we were able to recover known minimal GRNs for a bistable switch and an oscillatory circuit. We subsequently evolved modular GRNs that incorporated these previously identified sub-circuits in various ways to generate more complex target behaviors.

Acknowledgments We thank J. Lindle for contributions to the initial development of the project and M. Dunlop for helpful comments on the manuscript.

References

- Alon U (2006) An introduction to systems biology: design principles of biological circuits. CRC press, Florida
- Bar-Joseph Z, Gitter A, Simon I (2012) Studying and modelling dynamic biological processes using time-series gene expression data. *Nat Rev Genet* 13:552–564. doi:[10.1038/nrg3244](https://doi.org/10.1038/nrg3244)
- Cao H, Romero-Campero FJ, Heeb S, Cámara M, Krasnogor N (2010) Evolving cell models for systems and synthetic biology. *Syst Synth Biol* 4:55–84. doi:[10.1007/s11693-009-9050-7](https://doi.org/10.1007/s11693-009-9050-7)
- Das S, Suganthan PN (2011) Differential evolution: a survey of the state-of-the-art. *IEEE Trans Evol Comput* 15(1):4–31
- de Jong H (2002) Modeling and simulation of genetic regulatory systems: a literature review. *J Comput Biol* 9:67–103. doi:[10.1089/10665270252833208](https://doi.org/10.1089/10665270252833208)
- Drennan B, Beer R (2006) Evolution of repressilators using a biologically-motivated model of gene expression. The MIT Press (Bradford Books), Cambridge
- Dunlop MJ (2011) Engineering microbes for tolerance to next-generation biofuels. *Biotechnol Biofuels* 4:32. doi:[10.1186/1754-6834-4-32](https://doi.org/10.1186/1754-6834-4-32)
- Elowitz MB, Leibler S (2000) A synthetic oscillatory network of transcriptional regulators. *Nature* 403:335–338. doi:[10.1038/35002125](https://doi.org/10.1038/35002125)
- François P, Hakim V (2004) Design of genetic networks with specified functions by evolution in silico. *Proc Natl Acad Sci USA* 101:580–585. doi:[10.1073/pnas.0304532101](https://doi.org/10.1073/pnas.0304532101)
- François P, Hakim V (2005) Core genetic module: the mixed feedback loop. *Phys Rev E* 72(3):031908
- Gilbert ES, Walker AW, Keasling JD (2003) A constructed microbial consortium for biodegradation of the organophosphorus insecticide parathion. *Appl Microbiol Biotechnol* 61:77–81. doi:[10.1007/s00253-002-1203-5](https://doi.org/10.1007/s00253-002-1203-5)
- Hansen N (2006) The CMA evolution strategy: a comparing review, vol 192. Springer, Berlin. doi:[10.1007/3-540-32494-1_4](https://doi.org/10.1007/3-540-32494-1_4)
- Khalil AS, Collins JJ (2010) Synthetic biology: applications come of age. *Nat Rev Genet* 11:367–379. doi:[10.1038/nrg2775](https://doi.org/10.1038/nrg2775)
- Kim J-R, Yoon Y, Cho K-H (2008) Coupled feedback loops form dynamic motifs of cellular networks. *Biophys J* 94:359–365. doi:[10.1529/biophysj.107.105106](https://doi.org/10.1529/biophysj.107.105106)
- Komiya K, Noman N, Iba H (2012) The search for robust topologies of oscillatory gene regulatory networks by evolutionary computation. ACM Press, New York. doi:[10.1145/2330784.2330963](https://doi.org/10.1145/2330784.2330963)
- Levine JH, Lin Y, Elowitz MB (2013) Functional roles of pulsing in genetic circuits. *Science* 342:1193–1200. doi:[10.1126/science.1239999](https://doi.org/10.1126/science.1239999)
- Marbach D, Prill RJ, Schaffter T, Mattiussi C, Floreano D, Stolovitzky G (2010) Revealing strengths and weaknesses of methods for gene network inference. *Proc Natl Acad Sci*. doi:[10.1073/pnas.0913357107](https://doi.org/10.1073/pnas.0913357107)
- Marbach D et al (2012) Wisdom of crowds for robust gene network inference. *Nat Meth* 9:796–804
- Mukherji S, van Oudenaarden A (2009) Synthetic biology: understanding biological design from synthetic circuits. *Nat Rev Genet* 10:859–871. doi:[10.1038/nrg2697](https://doi.org/10.1038/nrg2697)
- Noman N, Palafox L, Iba H (2013a) Evolving genetic networks for synthetic biology. *New Gener Comput* 31:71–88. doi:[10.1007/s00354-013-0201-8](https://doi.org/10.1007/s00354-013-0201-8)
- Noman N, Palafox L, Iba H (2013b) Reconstruction of gene regulatory networks from gene expression data using decoupled recurrent neural network model, vol 6. Springer, Berlin. doi:[10.1007/978-4-431-54394-7_8](https://doi.org/10.1007/978-4-431-54394-7_8)
- Price KV, Storn RM, Lampinen JA (2005) Differential evolution: a practical approach to global optimization. Springer, Berlin
- Purvis JE, Lahav G (2013) Encoding and decoding cellular information through signaling dynamics. *Cell* 152:945–956. doi:[10.1016/j.cell.2013.02.005](https://doi.org/10.1016/j.cell.2013.02.005)
- Rajendran M, Ellington AD (2008) Selection of fluorescent aptamer beacons that light up in the presence of zinc. *Anal Bioanal Chem* 390:1067–1075. doi:[10.1007/s00216-007-1735-8](https://doi.org/10.1007/s00216-007-1735-8)
- Ruskin HJ, Crane M (2011) Stages of gene regulatory network inference: the evolutionary algorithm role. InTech, Rijeka
- Sheinman M, Kafri Y (2012) How does the DNA sequence affect the Hill curve of transcriptional response? *Phys Biol* 9(5):056006
- Smolen P, Baxter DA, Byrne JH (2000) Mathematical modeling of gene networks. *Neuron* 26:567–580. doi:[10.1016/S0896-6273\(00\)81194-0](https://doi.org/10.1016/S0896-6273(00)81194-0)
- Storn R, Price K (1997) Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces. *J Global Optim* 11:341–359. doi:[10.1023/A:1008202821328](https://doi.org/10.1023/A:1008202821328)
- Thomas SA, Jin Y (2012) Combining genetic oscillators and switches using evolutionary algorithms. IEEE, New York. doi:[10.1109/CIBCB.2012.6217207](https://doi.org/10.1109/CIBCB.2012.6217207)
- van Dorp M, Lannoo B, Carlon E (2013a) Evolutionary generation of small oscillating genetic networks, vol 7824. Springer, Berlin. doi:[10.1007/978-3-642-37213-1_13](https://doi.org/10.1007/978-3-642-37213-1_13)
- van Dorp M, Lannoo B, Carlon E (2013b) Generation of oscillating gene regulatory network motifs. *Phys Rev E* 88:012722. doi:[10.1103/PhysRevE.88.012722](https://doi.org/10.1103/PhysRevE.88.012722)
- Wang L, Walker BL, Iannaccone S, Bhatt D, Kennedy PJ, Tse WT (2009) Bistable switches control memory and plasticity in cellular differentiation. *Proc Natl Acad Sci*. doi:[10.1073/pnas.0806137106](https://doi.org/10.1073/pnas.0806137106)
- Widder S, Macía J, Solé R (2009) Monomeric bistability and the role of autoloops in gene regulation. *PLoS ONE* 4:e5399. doi:[10.1371/journal.pone.0005399](https://doi.org/10.1371/journal.pone.0005399)