Cite this: Mol. BioSyst., 2011, 7, 1081–1092

www.rsc.org/molecularbiosystems

PAPER

Evolvability and robustness in a complex signalling circuit⁺

Karthik Raman^{ab} and Andreas Wagner*^{abc}

Received 18th August 2010, Accepted 14th December 2010 DOI: 10.1039/c0mb00165a

Biological systems at various levels of organisation exhibit robustness, as well as phenotypic variability or evolvability, the ability to evolve novel phenotypes. We still know very little about the relationship between robustness and phenotypic variability at levels of organisation beyond individual macromolecules, and especially for signalling circuits. Here, we examine multiple alternate topologies of the Saccharomyces cerevisiae target-of-rapamycin (TOR) signalling circuit, in order to understand the circuit's robustness and phenotypic variability. We consider each of the topological variants a genotype, a set of alternative interactions between TOR circuit components. Two genotypes are neighbours in genotype space if they can be reached from each other by a single small genetic change. Each genotype (topology) has a signalling phenotype, which we define via the concentration trajectories of key signalling molecules. We find that the circuits we study can produce almost 300 different phenotypes. The number of genotypes with a given phenotype varies very widely among these phenotypes. Some phenotypes have few associated genotypes. Others have many genotypes that form genotype networks extending far through genotype space. A minority of phenotypes accounts for the vast majority of genotypes. Importantly, we find that these phenotypes tend to have large genotype networks, greater robustness and a greater ability to produce novel phenotypes. Thus, over a broad range of phenotypic robustness, robustness facilitates phenotypic variability in our study system. Our observations show parallels to studies on macromolecules, suggesting that similar principles might govern robustness and phenotypic variability in biological systems. Our approach points a way towards mapping genotype spaces in complex circuitry, and it exposes some challenges such mapping faces.

Introduction

Biological macromolecules such as proteins and RNA show intriguing properties that increase their ability to withstand perturbations, as well as to evolve novel phenotypes with new functions. Their genotypes-amino acid or nucleotide sequences-exist in vast genotype spaces. Genotypes that form the same phenotype—a secondary or tertiary structure with a specific function-are connected into large neutral networks or genotype networks.¹⁻⁵ Individual genotypes in such a network have many neighbours with the same phenotype. They are therefore to some extent robust to mutations that change single amino acids or nucleotides.^{5,6} At the same time, through evolutionary exploration of a genotype network,

these molecules encounter novel phenotypes in the immediate vicinity of the genotype network, some of which may be useful evolutionary adaptations.^{4–7} The existence of such genotype networks is thus important for the evolvability of molecules.

Although much is known about how different molecular genotypes and their phenotypes are organised in genotype space, much less is known in this regard about biological systems on higher levels of organisation, that is, about biological networks. We increasingly appreciate that biological systems on multiple levels of organisation are robust to perturbations. They remain able to continue performing their functions, even in the face of environmental or genetic perturbations.^{8–14} This feature emerges from how phenotypes form from genotypes, and how phenotypes are distributed in genotype space. However, with few exceptions, 7,15-18 we know little about this organisation. Our ignorance in this area is especially stark for the complex signalling circuits that play key roles in many physiological and developmental processes.¹⁹⁻²³ How do circuit genotypes map onto circuit phenotypes, and how are the resulting signalling phenotypes distributed in genotype space? Does their organisation in genotype space have implications for their robustness, as well

^a Department of Biochemistry, University of Zürich, Winterthurerstrasse 190, 8057 Zürich, Switzerland. E-mail: aw@bioc.unizh.ch

^b Swiss Institute of Bioinformatics, Quartier Sorge, Batiment Genopode, 1015 Lausanne, Switzerland

² The Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501. USA

[†] Electronic supplementary information (ESI) available: Supplementary Figures and Tables. See DOI: 10.1039/c0mb00165a

as for their ability to evolve new phenotypes? We here address these questions for a biochemically realistic model of a eukaryotic signalling circuit, namely the target of rapamycin (TOR) signalling circuit in the budding yeast *Saccharomyces cerevisiae*.^{21,24–26}

Quantitative modelling of biological systems forms the cornerstone of systems biology.^{27,28} Modelling and simulation provide powerful means for developing and testing hypotheses on the function and behaviour of complex biological systems. Quantitative models have been useful in furthering the understanding of many biological systems.^{29–32} Examples include bacterial chemotaxis,³³ early *Drosophila* development,¹⁰ and synthetic oscillator circuits.³⁴ Different models of the same process may differ in their parameters, or in their topology, the qualitative interaction pattern of circuit components. Alternative circuit topologies often represent variants in circuit structures that can evolve through accumulating small genotypic changes on evolutionary time scales.^{35–37}

In this study, we ask how topological changes in the TOR signalling circuit in S. cerevisiae affect the circuit's behaviour. TOR is a highly conserved atypical protein kinase that controls the growth of proliferating yeast, fly and mammalian cells in response to nutrients.^{25,26} It is the target of rapamycin, an immuno-suppressant and anti-cancer drug. Despite several biochemical experiments aimed at characterising the components and the mechanisms of the TOR signalling circuit,^{24,38–43} uncertainty prevails about the qualitative interactions between various circuit components-the circuit topology-and the parameters describing the quantitative dynamics of these components. Recently, Kuepfer and co-workers proposed multiple alternate dynamic models for TOR signalling in S. cerevisiae.²⁴ Specifically, these authors proposed a core topology of TOR signalling, consisting of molecular interactions and reactions that are well-understood, along with 18 extensions to the core (Fig. 1). Each extension can be viewed as a variant of the signalling circuit that affects a set of elementary molecular interactions (Table S1, ESI⁺, ref. 24), and represents a hypothesis about the mechanism of TOR signalling. Individual extensions are based on direct and indirect evidence from biochemical experiments (Table S1, ESI^{\dagger}). For example, variant V_1 is based on a study by Hall and co-workers, who suggest that Tip41p is phosphorylated at multiple sites, based on immuno-precipitation of Tip41p followed by treatment with phosphatase.⁴¹ The complex formation between Tap42p and Pph21/22p or Sit4p (variants V_{17} and V_{18}) has been suggested by Di Como and Arndt,³⁸ based on the co-precipitation of Tap42p with Sit4p, as well as with Pph21/22p. Jiang and Broach⁴⁰ hypothesise that the Tap42p-Pph21/22p complex would protect substrate phosphoproteins from dephosphorylation by PP2A or other phosphatases (variant V_2)⁴⁰; this hypothesis is able to account for many of their observations on the effects of over-expressing or activating Tap42p. Each combination of these and other elementary variants of the TOR signalling mechanism leads to a different topology of the signal transduction circuit.

Any one yeast strain would typically harbour only one circuit topology. We can think of this topology as the signalling circuit's *genotype*, a pattern of molecular interactions that is

ultimately encoded in the strain's genome. Other topologies, formed by different combinations of elementary variants may occur in different yeast strains, or in one of several species closely related to *S. cerevisiae*, although such variants have thus far not been characterised in yeasts. By examining all topologies, we may get insights into different evolutionary trajectories that a signalling circuit may take.

A signalling circuit's *phenotype* includes the concentrations of signalling molecules, and how this concentration changes in response to environmental signals. We here determined the *phenotype* of each TOR circuit topology based on the concentration trajectories of several key proteins/complexes (see Methods). Specifically, we clustered these trajectories to group similar 'signalling behaviours', into different *phenotypes*.

By mapping the topology of a signal transduction circuit to a genotype, and by computing a signalling phenotype from this genotype, we can address a number of questions with implications for circuit evolution: can one alter the topology of a signal transduction circuit without affecting its behaviour? Can signal transduction circuits with significantly different topologies exhibit the same behaviour? Do different topologies have very different or similar behaviours? How diverse is the range of novel behaviours that variation in a single topology can produce? We here address these questions by studying multiple alternate circuit topologies for TOR signalling, and their phenotypes.

Results

The possible combinations of the 18 elementary variants we consider create $2^{18} > 2.6 \times 10^5$ circuit topologies. After eliminating incompatible combinations of such variants (see Methods), 6.9×10^4 topologies remain. To date, only few of them have been examined.²⁴ We here analyse all of them. To facilitate their comparison, and to relate our observations to previous work,^{4,5,7,15,16} we discretize both genotypes and phenotypes. Specifically, we represent each genotype as an 18-bit long Boolean vector (a vector with binary co-ordinates 0 or 1), where each bit indicates the presence or absence of one elementary variant (Table S1, ESI[†]) in a circuit topology.

For the signalling circuits that we study here, it is challenging to map genotypes to different signalling behaviours, because these behaviours are *continuous* in nature. To obtain a classification of phenotypes, and a discrete mapping of genotypes to phenotypes, we therefore clustered the trajectories for the 6.9×10^4 models to obtain 286 well-separated clusters, representative of different phenotypes (see Methods). Each 'trajectory' in this case represents a set of measurements of the (normalised) concentrations of key signalling molecules including Tap42p^(P) (^{P)} indicates phosphorylation), the Tip41p–Tap42p complex and the Tap42p^(P)–Sit4p complex, at a set of time points (see Methods). Our approach provides a clear separation of phenotypes into different clusters (Fig. S1, ESI†).

To facilitate a systematic analysis, we also mapped the circuit topologies onto a graph. Each node in this graph represents a specific topology. Two topologies are neighbours and thus connected by an edge, if they differ in exactly one of the 18 elementary variants (Table S1, ESI[†]). Any such



Fig. 1 Schematic of TOR signal transduction models. Schematic illustration of the connectivity of TOR circuit genotypes in genotype space. The figure illustrates a few circuit topologies (blue circles), and indicates neighbouring topologies through black lines. Each of the displayed genotypes has 18 neighbours in genotype space (stubs emerging from blue circles), only few of which are shown in detail. Each topology is labelled by the elementary topological variants (Table S1, ESI⁺) that it contains. The core circuit (top rectangle), together with three other topologies (middle and bottom rectangles) are illustrated. Differences in topologies are indicated by shaded boxes. Each topology is represented according to a standardised process diagram graphical notation.^{44,45} In this notation, green rectangles represent proteins and protein complexes (with or without phosphorylation), while green ellipses correspond to small molecules, such as rapamycin. On arrows, open squares indicate transitions, while filled circles indicate complex formation. Arrows ending in open circles adjacent to reactions indicate catalysis. The process diagrams were drawn using CellDesigner.⁴⁶

difference can potentially be caused by a small genetic change. This graph is the *genotype space* for the set of TOR circuit topologies we study. A small portion of this space is shown in Fig. 1. The stubs emerging from the blue circles (representing topologies in the figure) indicate that each topology has multiple neighbours, of which only few are shown in detail. Since there are 18 elementary variants, each node in this graph has at most 18 neighbours.

All genotypes that adopt a given phenotype form this phenotype's *genotype set*. A genotype set can consist of one or more connected *genotype networks*, which are connected sub-graphs of genotype space. Such networks have also been called *neutral networks*.^{1,6} We deliberately refrain from using the term neutral network here, because phenotypes in one of the clusters we study may not be neutral variants in the strict sense used by evolutionary biologists.⁴⁷ Genotype networks have very different sizes and differ in many of their properties, which may have implications for the evolutionary dynamics of these systems.⁴⁸ In the following sections, we discuss the genotype sets and genotype networks of the TOR signal transduction circuit.

Most genotypes are contained in few genotype networks

The 6.9×10^4 genotypes of the TOR signalling circuit that we analysed display 286 different phenotypes. The size distribution of the genotype set for these 286 phenotypes is shown in Fig. 2A. The figure shows that a majority of the genotype sets are quite small, but multiple large sets exist as well. The largest genotype set has 21,633 genotypes. Fig. 2B illustrates a rank-ordered size distribution of genotype sets. It shows that most of the genotypes are contained in a minority of genotype sets. Specifically, the 57 largest genotype sets contain over 90% of all genotypes. We report results from most of the following analyses both for all genotype sets, and for those genotype sets that contain 90% of the genotypes, in order to eliminate biases caused by the many smaller networks that collectively contain few genotypes. For brevity, we will refer to the latter genotype sets as large genotype sets.



Fig. 2 (A) Genotype set size distribution. Note the logarithmic scale on the vertical axis. While a majority of the genotype sets are small in size, there are a few genotype sets that are very large. (B) Most genotypes are contained within the largest genotype sets. The horizontal axis (log-scale) represents genotype sets ranked in descending order of size. The vertical axis indicates the fraction of genotypes *contained* in genotype sets within a given size. The dotted line illustrates that 90% of the genotypes are contained within the largest 57 genotype sets. (C) Distribution of the fraction of a genotype set occupied by the largest genotype network. (D) Distribution of maximum genotype distance within the largest genotype network of a set. Note that the maximum genotype distance has been expressed as a fraction of genotype space diameter (18). In both (C) and (D), only the large genotype sets are shown.

Topologies sharing the same phenotype form highly connected networks

Two extreme scenarios of genotype set organisation are possible, and a broad spectrum in between. In one scenario, all genotypes in a genotype set are disconnected. In this case, the entire genotype *set* is a fragmented collection of genotypes that cannot be reached from one another through phenotypepreserving evolutionary change. In the other scenario, the entire genotype set is connected and thus consists of only one genotype network. In this case, one can move from one topology to another on this network, through small evolutionary steps that leave the phenotype unchanged. We next analysed the connectedness of genotype sets.

Fig. 2C shows the distribution of the fraction of a genotype set occupied by its largest genotype network, for phenotypes corresponding to the large genotype sets. Fig. S2A (ESI[†]) shows the distribution for all genotype sets. For large

genotype sets, merely 14% have the majority of their nodes contained in a single genotype network. The largest genotype set contains 21 633 genotypes, 21 307 of which are connected in a single genotype network. On such a large genotype network, one can change the circuit's topology (genotype) through small evolutionary steps dramatically, without altering the behaviour of the signalling circuit.

Dissimilar genotypes can exhibit the same phenotype

The distance between two genotypes can be calculated as the Hamming distance between the binary vectors representing the genotypes. This genotype distance corresponds to the number of elementary variants in which two circuit topologies differ. Within a genotype network, the maximal genotype distance indicates how different two genotypes sharing the same phenotype can be. Fig. 2D shows the distribution of the maximum genotype distance between any two genotypes in a



B

Fig. 3 (A) Phenotypes of larger genotype sets are more robust. (B) Robust genotypes encounter fewer phenotypes in their neighbourhood. (C) Robust phenotypes encounter more novel phenotypes in their neighbourhood. (D) Larger genotype networks encounter a greater number of phenotypes in their neighbourhood. The panel shows the correlation between phenotype evolvability and the size of the largest genotype network for the large genotype sets.

Α

genotype network, for the large genotype sets. We refer to this maximal distance as the diameter of the genotype network.⁴⁹ We express this distance as a fraction of the maximally possible distance, the diameter of genotype space, which is 18 in our model. Fig. S2B (ESI⁺) shows the distribution of maximum genotype distance, for all phenotypes.

The largest genotype set of size 21 633 contains the largest genotype network of size 21 307. This network has a diameter of 18, equal to the genotype space diameter. This network corresponds to a cluster of circuits whose phenotype is closest to the experimentally determined reference signalling phenotype.^{24,38–41} The second and third largest genotype networks have 6913 and 4446 genotypes, and show maximal distances of 17 and 16, respectively. The median distance for the 10 largest genotype networks, containing nearly 60% of all genotypes, is 15.5.

Fig. S3A (ESI[†]) shows, for all phenotypes, the association between the fraction of genotypes contained in the largest genotype network and this network's diameter. When only the large genotype sets are considered (Fig. S3B, ESI[†]), a significant positive correlation is seen between the two quantities: unsurprisingly, phenotypes that have a greater fraction of their genotypes connected in a single network also have a relatively larger diameters. However, this correlation disappears when smaller networks are included (Fig. S3A, ESI⁺). The reason is that many of the phenotypes have fragmented genotype sets with very small networks, which must have small diameters. Taken together, these observations indicate that circuit genotypes with a given phenotype can be very diverse. Phenotypes with large genotype sets can well differ in most of the elementary topological variants we study. For such phenotypes, one can move through the associated genotype network via single changes of elementary topological variants, and change the topology of a signal transduction circuit dramatically, while preserving similar signalling behaviour. However, for many phenotypes with small or fragmented genotype sets, such flexibility is much more limited.

Robust phenotypes have larger genotype sets

The patterns of genotype network connectivity we discussed above (Fig. 2C) also have implications for robust signalling behaviour. In a highly fragmented or small genotype set, many mutations (changes in topology) would lead to a change in the signalling phenotype. The other extreme is a large genotype set consisting of only one large genotype network. In such a network, individual genotypes can have many neighbours circuits differing in only one topology-altering mutation—with the same phenotype. A circuit that is part of such a large genotype network could absorb many mutations without changing its phenotype. We define the robustness of a circuit (genotype) as the fraction of its neighbours with the same phenotype.

Fig. S4 (ESI[†]) shows the distributions of genotype robustness for six genotype networks—the three largest genotype networks, and three genotype networks of smaller sizes. As one might expect, the median robustness is substantially higher for the large genotype sets than for the smaller genotype sets. In the largest genotype network, for example, robustness ranges from 0.056 to 0.94, and the average robustness is 0.49.

Analogously to the robustness of a genotype, we can define the robustness of a phenotype as the average robustness of all genotypes with this phenotype.⁵ Robust phenotypes are less easy to perturb by changing their circuit's topology. In Fig. 3A and Fig. S5A (ESI[†]), we analyse the relationship between genotype set size and phenotype robustness. We observe a high (and highly significant) positive association. In other words, robust phenotypes typically are phenotypes adopted by many genotypes. For example, the largest genotype network (with 21 307 genotypes) corresponds to a phenotype that has the highest average robustness of 0.49. This is also the phenotype displaying the TOR reference signalling behaviour. Every circuit displaying this TOR reference behaviour is on average connected to 9 neighbouring circuit topologies with this behaviour. Fig. S5B (ESI[†]) illustrates a strong positive correlation between genotype network diameter and phenotype robustness. Genotype networks with greater diameters exhibit more robust phenotypes. This can be understood in light of the fact that genotype networks with large diameters harbour circuits that differ widely in their topologies, yet exhibit similar behaviour. Thus, their phenotypes are invariant to many changes in topology, or in other words, show a robust signalling behaviour.

Robust phenotypes have higher evolvability

We will now turn to the ability of signalling circuits to explore new phenotypes in a blind, phenotype-preserving search of their genotype space. We will refer to this ability as evolvability, and focus here on those phenotypes that occur in the immediate neighbourhood of a genotype or a genotype set. Specifically, we define genotype evolvability as the number of different phenotypes found in the 1-neighbourhood of a circuit genotype G. This neighbourhood is the set of genotypes that differ from G in exactly one of the elementary topological variants (Table S1, ESI[†]). Analogously, we define *phenotype evolvability* as the number of different phenotypes found in the immediate neighbourhood of the largest genotype network associated with a phenotype P. Since the majority of genotypes lie on the largest genotype networks, these networks are most appropriate for our analysis. This neighbourhood includes all genotypes that are neighbours of genotypes in the largest genotype network of P, but that are not themselves members of this network. The different phenotypes in these neighbourhoods are precisely the phenotypes that are readily accessible from a genotype or a genotype network, via single topology-changing mutations.

We first analyse the relationship between a genotype's robustness and its evolvability. Fig. 3B shows that genotypes with high robustness can access fewer phenotypes in their neighbourhood. This is not entirely surprising, because robust genotypes have many neighbours with the same phenotype, and thus fewer neighbours with different phenotypes. Perhaps more surprising is that the robustness of a phenotype and its evolvability show a strongly positive association (Fig. 3C). The more robust a phenotype is, the greater the number of novel phenotypes that occur in its phenotypic neighbourhood. This relationship is a joint consequence of two facts. First, robust phenotypes tend to be phenotypes with large genotype networks, as we discussed earlier (Fig. 3A). Second, large genotype networks tend to have a larger number of different phenotypes in their neighbourhood, as shown by the analysis of Fig. 3D and Fig. S5C (ESI[†]). These observations mirror observations made earlier in evolving macromolecules.⁵ A high phenotypic evolvability implies that a large number of unique phenotypes are accessible from the corresponding genotype network; however, a genotype may have to undergo one or more neutral mutations, before it accesses these novel phenotypes, particularly on very large genotype networks.

Evolving networks encounter novel phenotypes in their immediate neighbourhood

In an extended genotype network, genotypes can change substantially without changing their phenotypes. Two different genotypes G_i and G_j on the same genotype network, but at some distance d_{ij} may contain very different phenotypes in their 1-neighbourhoods. How many novel phenotypes are accessible in the immediate neighbourhood of these genotypes? To address this question, we now study the fraction u of phenotypes that occur in the neighbourhood of one but not the other genotype. This fraction $u(G_i,G_j)$ can be calculated as:

$$u(G_i, G_j) = 1 - \frac{|\mathcal{N}_i \cap \mathcal{N}_j|}{|\mathcal{N}_i \cup \mathcal{N}_j|}$$

where \mathcal{N}_i and \mathcal{N}_j represent the sets of unique phenotypes in the 1-neighbourhoods of the genotype G_i and G_j respectively, and $|\mathcal{N}|$ represents the number of elements in the set \mathcal{N} . Note that this fraction may depend on the distance d_{ij} between two genotypes.

This analysis speaks to a circuit's evolvability in a fashion complementary to the last section. If u is large even at small genotype distance d, then a genotype would not have to change by much until it can access a different spectrum of novel phenotypes *via* single mutations. This is exactly the case, as Fig. 4A shows. The figure is based on averages over all 286 phenotypes. Even at the smallest possible distance between genotypes, over 70% of phenotypes occur in the neighbourhood of one but not the other genotype. This percentage does not increase dramatically for larger distances.

Circuits with different signalling behaviours can be close together in genotype space

In addition to the above analysis, we asked how far one must travel in genotype space from one genotype set to find another genotype set with an arbitrary new phenotype. To address this question, we computed the minimal distance between genotypes having different phenotypes (see Methods). If this distance is typically large, then it would be rather difficult to reach a new phenotype from a genotype having a different phenotype through a small series of genetic changes that alter the topology of the signalling circuit. If, in contrast, this distance is typically small, it would be possible to discover new phenotypes through a relatively small number of genetic changes. This distance thus has implications for the evolvability of signalling circuits.

Fig. S6 (ESI[†]) indicates the distribution of minimal genotype distances between different genotype sets. The minimal distances between smaller genotype sets are larger. We do not normalise for the size of the genotype networks here, because we are interested in the *number* of mutations—evolutionary distance—that separate two genotype networks, irrespective of their sizes. For larger genotype sets, most distances are equal to 6% of the genotype space's diameter, corresponding to a single topology-altering mutation (Fig S6B, ESI[†]). Thus, it is possible to access many novel phenotypes on large genotype networks through only a single change in circuit topology. This again reaffirms the earlier observation that robust



Fig. 4 (A) Phenotypic diversity of different neighbourhoods. The horizontal axis shows the distance d_{ij} between two circuits with the same phenotype. The vertical axis shows the mean fraction $u(G_i, G_j)$ of unique phenotypes found in a 1-neighbourhood around these circuits, as defined in the main text. The analysis is based on the largest genotype network for each of the 286 genotype sets. The error bars indicate one standard deviation. (B) Populations evolving on larger genotype networks can access more new phenotypes. The largest genotype network from each of the genotype sets that contained over 90% of the genotypes were binned by size and the mean number of unique phenotypes in the 1-neighbourhood of a population evolving on each network (after 100 generations) is indicated for the different bins. The error bars indicate one standard deviation. Mutation rate per generation per individual $\mu = 0.25$; population size N = 100.

phenotypes are also more evolvable; they have many novel phenotypes in their neighbourhood, and their genotype sets are also located closer to genotype sets of novel phenotypes.

Populations evolving on larger genotype networks can access a wider variety of phenotypes

All evolution occurs in populations of organisms. In order to understand the evolvability of any biological system, a population perspective is thus necessary. To see whether phenotypic robustness also facilitates the evolution of new circuit phenotypes in a population context, we allowed a population of N = 100 initially identical circuits to evolve via repeated cycles of mutation-defined as an elementary change in the bit-string representing circuit topology-and selection confining the circuits to one genotype network. During this process, we recorded the number of unique phenotypes $P_{\rm U}(t)$ in the 1-neighbourhood of the entire population. $P_{\rm U}$ indicates the number of novel phenotypes that are immediately accessible to individuals in an evolving population. Fig. 4B (ESI^{\dagger}) shows how $P_{\rm U}$ after t = 100generations depends on genotype network size, for a mutation rate corresponding to $\mu = 0.25$ (per generation, per individual). Fig. S7 (ESI[†]) shows the same for mutation rates corresponding to $\mu = 0.10, 0.50$. For all mutation rates, populations evolving on larger genotype networks generally have greater access to novel phenotypes. The exceptions are the very largest genotype networks, where access to novel genotypes declines, again similar to earlier observations in macromolecules.⁵⁰ We also observe this trend in Fig. S8 (ESI[†]), where we illustrate the evolution of $P_{\rm U}(t)$ for four different genotype networks. Except for the largest genotype network, we find an increase in the number of unique phenotypes in the 1-neighbourhood of the evolving population. Also, for larger population sizes, we see an increase in the number of unique phenotypes accessible in the 1-neighbourhood of the population (Fig. S8B, ESI[†]). These results highlight the importance of a population-centric view to understand the evolvability of these signalling circuits.

Diverse topologies can describe TOR signal transduction

We next focus on the phenotype that describes the canonical TOR signalling behaviour. The genotype set for this phenotype is the largest set in our genotype space. We will refer to this genotype set as the TOR genotype set. This set comprises 21633 genotypes, with the vast majority of 98.5% (21307 genotypes) connected in a single network, the largest genotype network for all phenotypes. A typical circuit genotype on this network can absorb multiple mutations without losing its phenotype. The maximum genotype distance within this genotype network is 18; this means that topologies that differ maximally in their structure can still preserve the TOR signalling phenotype. For example, the models represented by the genotypes [000001001010011111] and [11111011010100000] differ in all their elementary topological variants, but are still on the same genotype network. In addition, the neighbourhood of this genotype network contains 272 novel phenotypes, >95% of all phenotypes. This observation hints at the high evolvability of this phenotype. We note that some of these properties may be a by-product of our parameter optimisation

procedure; we optimised the parameters to reproduce the canonical TOR signalling behaviour (see Discussion).

We now address the question how unusual the structure of this genotype set is, by comparing it to a specific class of random graphs. These random graphs have the same number of circuits as the genotype set, and two circuits are connected if they differ in exactly one topological variant; however, these circuits are simply drawn at random from genotype space, and thus need not have the same phenotype. Fig. 5 shows the distribution of three different measures describing the structure of these random graphs, and compares them to the TOR genotype set. It can be seen that the structure of the TOR genotype set is dramatically different from that of the random graphs. Specifically, the random graphs are highly fragmented and contain many disconnected networks (Fig. 5A), which are also smaller in size (Fig. 5B). In contrast, the genotype set for TOR is cohesively connected, containing one very large network and few tiny networks (Fig. 5A and 5B). The number of edges in the TOR genotype network is also substantially higher, by almost an order of magnitude, compared to the average number of edges in random graphs (Fig. 5C). Taken together, these observations show that the genotype set characterising the canonical TOR behaviour is highly unusual in its connectivity properties, which also affects its evolvability.

Discussion

We here studied nearly 70000 biochemically realisable genotypic variants of the yeast TOR circuit, in order to understand its robustness, phenotypic variability, and the relationship of phenotypic variants to experimentally characterised signalling behaviour.²⁴ The $\approx 70\,000$ potential TOR genotypes represent alternative interactions between TOR circuit components, and thus alternative TOR circuit topologies. We represent these genotypes using systems of ordinary differential equations, describing mass-action kinetics.²⁴ For each topology, *i.e.* genotype, we compute a phenotype based on the concentration trajectories of key signalling molecules. Mapping a phenotype to each of the signalling circuits enabled us to identify genotype networks, connected sets of genotypes with the same phenotype. Genotype networks have been previously investigated for RNA molecules,^{1,5} regulatory networks,^{7,15,16,18} as well as protein structures.^{3,4} Earlier studies on the evolution of robustness in biological networks, such as circadian oscillators¹⁵ and transcriptional regulatory networks,7,16,18 were based on more abstract models. In contrast, the TOR circuit we study is more biochemically detailed; the core of this model is also experimentally validated.24

Our observations fall into three categories. First, we show that the circuit genotype space can be partitioned into almost 300 sets of genotypes, where genotypes in each set adopt the same signalling phenotype. The size distribution of these sets is highly skewed, with a minority of such sets (phenotypes) encompassing the vast majority of genotypes. Smaller genotype sets are highly fragmented and typically contain multiple small genotype networks. In contrast, large sets contain extended and connected genotype networks that reach far through genotype space. Some of these networks contain



Fig. 5 The TOR genotype set exhibits unusual connectivity properties compared to random graphs of similar size. The distributions of three graph characteristics for 10 000 random graphs, whose nodes correspond to the possible alternate models of TOR, are shown. Each random graph was established by choosing a random sample of genotypes (as many as in the genotype set corresponding to the TOR phenotype, *viz.* 21 633), with the genotypes being connected if they differed by exactly one variant. (A) Number of networks. (B) Size of the largest network. (C) Number of edges. In all three cases, the TOR genotype set shows a significant departure from the characteristics of the random graphs.

thousands of genotypes (circuits). In large and extended genotype networks, circuits with substantially different topologies exhibit similar signalling behaviour. Genotypes on larger genotype networks also have a signalling phenotype that is more robust to changes in circuit topology. Such high robustness of biological circuitry to genotypic change is not unprecedented. For example, it may exist in *Caenorhabditis elegans* vulva development,^{36,51} where despite substantial variation in the underlying pathways in different environments, observed phenotypes are very similar.

The organisation of the genotype space we study shows some similarities to that of macromolecules and gene networks,^{1-5,7,15-18} including a highly non-uniform genotype set size distribution, and the existence of genotype networks. However, it also shows differences, for example a stronger fragmentation of genotype sets into smaller genotype networks of low diameter. In this regard, we note that our genotype space is tiny compared to the astronomical genotype spaces of macromolecules such as proteins and RNAs. The fragmentation we observe may result from this fact.

A second category of observation regards the phenotypic variability of signalling circuits, the ability to explore novel signalling behaviour. We have shown that over a broad range of robustness, robust phenotypes exhibit higher phenotypic variability. The reason is that robust phenotypes typically have large genotype networks, which have more novel phenotypes in their neighbourhood. These observations have interesting parallels to a recent study on RNA,⁵ where phenotypic robustness can lead to higher evolvability, whereas genotypic robustness hinders evolvability. The only exception in our study system regards the most robust phenotypes (Fig. 4B and Fig. S7, ESI⁺), which have access to slightly fewer novel phenotypes. Again, this observation has precedents, for example in models of RNA evolution,⁵⁰ and in population genetic models of evolvability,⁵² which suggest that extreme robustness can hinder access to novel phenotypes.

Thirdly, we analysed the phenotype that represents the canonical TOR signalling behaviour. Among all phenotypes we studied, this phenotype has the largest genotype network. It extends through the entire genotype space, and its neighbourhood contains >95% of all other phenotypes. The large size of this genotype network means that many circuits with different topologies exhibit a signalling behaviour close to the experimentally observed TOR signalling behaviour.^{24,38–41} The large diameter of the network indicates that widely different topologies can exhibit a similar behaviour. Together, these observations show that the TOR signalling phenotype is robust, because the TOR circuit can accommodate multiple changes to its topology without losing this phenotype. Furthermore, the large number of phenotypes in the neighbourhood of this genotype network indicate its ability to access many novel phenotypes.

Our work has several limitations. A serious limitation of our study relates to the complexity of our circuits and the resulting computational requirements for our analysis. We computed phenotypes from a single parameter set, obtained through a lengthy optimisation procedure aimed at finding parameters for which a given topology reproduces the canonical TOR signalling behaviour most faithfully. One might argue that a more sophisticated approach should be pursued. For example, we could have explored the entire parameter space for each topology. However, this is infeasible. First, our parameter space, depending on topology, may have more than 100 dimensions, and sampling it for even just one topology can be difficult. Second, we needed to examine not one but almost 70 000 different topologies. Rigorous sampling of the parameter space for all of them is impossible.

These computational constraints also introduce uncertainty in our phenotypic analysis. Our phenotypes result from only a single parameter set per genotype, while in reality, a large number of different parameter sets might produce widely varying behaviours. Because we cannot explore all possible behaviours a particular topology can produce, we limited ourselves to one behaviour per topology, which arises from a single parameter set. In addition, our choice to search for parameter sets that reproduced the canonical TOR behaviour most faithfully may have introduced artifacts. For example, although our analysis shows that the TOR signalling behaviour is a robust phenotype, because many genotypes can display it, the observation that it is the most robust phenotype may be influenced by our optimisation procedure. To validate whether such artifacts exist, one could in principle optimise for different signalling behaviours, and ask whether these behaviours then become associated with the largest genotype network. However, to do this in a systematic way again exceeds our current computational abilities.

Further, not all of the different topologies examined in this study may be realised in vivo. However, many of the individual variants or hypotheses we considered have varying degrees of support from biochemical experiments, as we mentioned before. Another limitation lies in the definition of a phenotype. To categorise phenotypes is challenging for all systems where phenotypes are continuous and not discrete in nature. This is the case for our system, whose signalling behaviour results from (continuous) concentration changes in signalling molecules. Because categorisation and enumeration of phenotypes is useful in analysing phenotypic variability, we categorised phenotypes according to their similarity in a continuous space. To do so, we used a clustering approach that assigned each phenotype to a cluster of similar phenotypes. Although our approach yielded a clear separation of phenotypic clusters, this may not be the case for all comparable systems. In addition, completely different definitions of phenotypes are conceivable.

In sum, we see two main values of this contribution. First, it highlights a general approach to analyse the genotype space of complex regulatory circuitry with many components and parameters. It provides a framework to systematically analyse a vast space of alternate circuit configurations, and provides insights into the organisation of this space. Also, it permits us to compare and contrast the robustness and evolvability of regulatory circuits to that of macromolecules and other well-studied systems.^{1–5,7,15–18} For example, our observations suggest that robustness of a circuit's signalling phenotype can facilitate the exploration of novel phenotypes. This hypothesis could also be experimentally tested in the laboratory: if we were to evolve two yeast species harbouring signalling circuits with differing robustness, then the species with the more robust circuit should be able to evolve new signalling behaviours and exhibit more diverse signalling behaviours over time. Second, our approach exposes several challenges

that we need to address if we want to understand complex biological circuitry, and the organisation of their genotype space. This organisation will hold the key to understand both robustness and evolvability of regulatory systems.

Methods

Genotype space

We view each topology of the TOR signalling circuit as one genotype. Such a topology comprises the core pathway with one or more of the 18 elementary pathway variants listed in Table S1 (ESI⁺). Since there are 18 different variants of the TOR signalling circuit, and because multiple such variants can occur simultaneously in the circuit, the total number of topologies is equal to $2^{18} \approx 2.6 \times 10^5$. However, some combinations of variants are incompatible with one another. For example, variants V_2 and V_6 represent conflicting hypotheses on the mechanism of TOR signalling: while V_2 proposes that the Tap42p^(P)–Pph21/22p complex forms an anti-phosphatase that protects phosphoproteins from de-phosphorylation, V_6 proposes that the same complex has phosphatase activity. V_8 and V_9 propose opposing roles for the Tap42p^(P)-Sit4 complex. On excluding all such incompatible variants, the total number of possible TOR circuit topologies reduces to 69120. Each topology corresponds to a set of reactions describing TOR signalling and is represented by a set of differential equations describing the dynamics of each of the chemical species involved in the signalling circuit. The number of reactions in any topology varies from 19-72, while the number of differential equations ranges from 24-56. These equations involve a number of parameters, which range from as few as 24 for the core pathway, to as many as 117, for more complex topologies.

The various genotypes form a genotype network, which can be represented as a graph whose nodes are the genotypes, and where two nodes are connected by an edge if they differ in exactly one of the 18 elementary variants. Implicit in this concept is the assumption that single mutations can change a topology and transform it into another, neighbouring topology. For example, a mutation in one or more amino acid residues could cause the loss or gain of a phosphorylation site (e.g. variants V_1 , V_{11} in Table S1, ESI[†]). Indeed, it is known that mutations in residues such as serine, threonine and tyrosine can lead to a loss or gain of phosphorylation sites on proteins.^{53,54} It is also known that point mutations that modify the interface between interacting proteins can affect protein-protein interactions,55-57 and consequently the topology of the signalling network.

The simple representation of genotypes that we use enables the easy calculation of a distance between two genotypes. We here use the Hamming distance between the two binary genotype vectors, which indicates the number of differences between the two genotypes when represented as binary strings. This distance represents the length of the *shortest mutational path* between the two genotypes, in the genotype space. Note that this is different from (and likely to be smaller than) the length of the shortest path between the two genotypes on any one genotype network, because not all the intermediate genotypes on the shortest mutational path may be part of the same genotype network.

Parameter estimation

The equations describing the TOR core pathway contain 24 kinetic parameters,²⁴ but some pathway variants contain over 100 parameters; these parameters capture the rates for the association/dissociation of various protein complexes involved in the signal transduction circuit, as well as the rates of important protein modifications, such as protein phosphorylation and dephosphorylation. For the core pathway and those 18 topologies that can be obtained by incorporating a single variant (Table S1, ESI⁺) into the core, the parameters have already been estimated earlier.²⁴ These estimates have been obtained through a global optimisation method employing an evolutionary algorithm, with the objective of minimising the deviation between predicted behaviour and experimental data. For each of the 6.9×10^4 models, we applied an Evolutionary Strategy (ES) based optimisation procedure, similar to the one described in ref. 24 (for 100 generations), to find an optimal parameter set that minimises the deviation of the model predictions from the available experimental data. In each case, we started with a parameter set based on the published data for the multiple TOR extensions.²⁴ We bounded the parameters in the interval $\{10^{-2}, 10^3\}$, as suggested in ref. 24.

For each topology, this approach aims to finds the set of parameters, **K**, that can best reproduce the reference signalling behaviour. In other words, it aims to find a parameter set that minimises the deviation of the model's predictions from the experimental data. This set of parameters is then used to predict the signalling behaviour of the topology, and to estimate the deviation from the reference behaviour. Owing to the differences in topologies between different circuits, many circuits may not reproduce the reference signalling behaviour, and give rise to multiple signalling phenotypes. We calculate the deviation \mathcal{D}_i from the reference signalling behaviour, for each experimental measurement i (0 < i < 11) as in ref. 24:

$$\mathscr{D}_i = \sum_j \left(\frac{x_j^* - x_j^p(\mathbf{K})}{\delta_j x_j^*} \right)^2$$

where the integer *j* runs through all measurements from the *i*th experiment (see Tables S2 and S3, ESI[†]). The vector \mathbf{x}^* contains the values of those state variables of the pathway that were measured experimentally. The vector $\mathbf{x}^p(\mathbf{K})$ contains the values of the same state variables, but as predicted by integrating the differential equations²⁴ corresponding to a genotype. Note that \mathbf{x}^p is a function of a particular parameter set **K** (see below). The variable δ_j is the estimated accuracy of the measured data point *j* (see Table S3, ESI[†], ref. 24). For the computed optimal parameter set, we computed the predicted trajectories of various species in the model, which we used to determine the phenotypes. Fig. S9 and S10 (ESI[†]) illustrate for two example models, the time-course data for some of the signalling molecules and how the predicted time courses compare with the experimental data.

Clustering phenotypes for TOR signalling

We identify phenotypes for the different models by clustering the 'trajectories' or the time-courses for different species, thereby distinguishing models with different behaviours. In particular, we consider eight species with non-zero concentrations at t = -90, viz. Tap42p, Tor12p, Pph2122p, Cdc55p-Tpd3p, Sit4p, Sap, Tip41p and Fpr1p; the timecourses for these species can be normalised with respect to the initial concentrations. t = 0 represents the steady state, at which rapamycin is introduced into the system. The time course was computed at another 22 time points in the time interval [0, 180]. These also include time-points for which data have not been previously estimated. By performing the parameter estimate as described above, we have already attempted to fit the trajectories to the available data (reference signalling behaviour). By introducing additional time-points and by computing the predicted values of the different species at these time-points, we increase the amount of information available, to better differentiate between different models. For each model, there are $22 \times 8 = 176$ data axes or 'features' to discriminate different model behaviours. We computed these features using a single parameter set for each model, which we had obtained after extensive optimisation, as described above.

Owing to the large size of the data, we performed an approximate on-line unsupervised clustering using BIRCH (balanced iterative reducing and clustering using hierarchies).⁵⁸ The data were grouped into 286 different clusters/phenotypes. Fig. S1 (ESI[†]) shows the distribution of inter-cluster distances and intra-cluster distances for these clusters. The figure illustrates a stark contrast between the inter-cluster distances and the intra-cluster distances, which are much lower, indicating good cluster separation.

Fraction of unique phenotypes in a genotypic neighbourhood

The immediate neighbourhood (1-neighbourhood) of a TOR topology (genotype) *G* consists of all topologies that differ from *G* in exactly one of the 18 elementary variants (Table S1, ESI†). We consider all phenotypes in the 1-neighbourhood of *G*, which are different from the phenotype of *G* itself. For brevity, we refer to these phenotypes as novel phenotypes here. For two genotypes G_i and G_j , we compute the average fraction of novel phenotypes in the 1-neighbourhood of G_i that are different from novel phenotypes in the 1-neighbourhood of G_i , for all pairs of genotypes with a fixed genotypic distance $k (d(G_i,G_j) = k, k = 1,2,...,18)$. Where the number of such pairs was greater than 10^5 , we performed the computations only for 10^5 pairs, chosen at random.

Minimal distance between genotype sets

We performed this analysis for all pairs of genotype sets. The minimal distance between two genotype sets can be defined as:

$$D_{\min}(\mathbf{P_i}, \mathbf{P_j}) = \min_{\forall G_m \in \mathbf{P_i}, \forall G_n \in \mathbf{P_j}} d(G_m, G_n)$$

where P stands for the genotype set whose genotypes have the same phenotype P. The distance between two genotypes is calculated through the Hamming distance of their bit-string representations, as described above.

Population dynamics on genotype networks

For a given genotype network, we chose a random genotype (circuit topology) and seeded a population of size N = 100with this genotype. At time t = 0, the population thus consists of N individuals with the same genotype (and thus the same phenotype). We allowed the population to evolve through repeated cycles ("generations") of replication and "mutation", where each mutation corresponds to a change in a single topological variant of the model (we used various rates of mutation, specifically $\mu = 0.10, 0.25, 0.50$ per circuit and generation). During this process, we confined the population to the genotype network. Specifically, whenever any mutation yielded genotype(s) outside the genotype network, we discarded these genotype(s) and maintained the population size by replacing these individuals with random individuals sampled (with replacement) from the previous generation. At each generation t, we computed the number of unique phenotypes $P_{\rm U}(t)$ in the neighbourhood of the entire population. We allowed the population to evolve for 100 generations. We examined the association between $P_{\rm U}(t = 100)$ and genotype set size.

Acknowledgements

We would like to thank Jörg Stelling and Sotiris Dimopoulos for helpful suggestions and critical reading of the manuscript. We also thank Elias Zamora-Sillero, Alberto Busetto and Mikael Sunnaker for helpful discussions. AW acknowledges support through Swiss National Science Foundation grants 315200-116814, 315200-119697, and 315230-129708. AW and KR are grateful for support through the YeastX project of SystemsX.ch.

References

- P. Schuster, W. Fontana, P. F. Stadler and I. L. Hofacker, *Proc. R. Soc. London, Ser. B*, 1994, 255, 279–284.
- 2 A. Babajide, I. L. Hofacker, M. J. Sippl and P. F. Stadler, *Folding Des.*, 1997, 2, 261–269.
- 3 E. Bornberg-Bauer and H. S. Chan, Proc. Natl. Acad. Sci. U. S. A., 1999, 96, 10689–10694.
- 4 E. Ferrada and A. Wagner, Proc. R. Soc. London, Ser. B, 2008, 275, 1595–1602.
- 5 A. Wagner, Proc. R. Soc. London, Ser. B, 2008, 275, 91-100.
- 6 A. Wagner, Nat. Rev. Genet., 2008, 9, 965–974.
- 7 S. Ciliberti, O. C. Martin and A. Wagner, Proc. Natl. Acad. Sci. U. S. A., 2007, 104, 13591–13596.
- 8 E. van Nimwegen, J. P. Crutchfield and M. Huynen, Proc. Natl. Acad. Sci. U. S. A., 1999, 96, 9716–9720.
- 9 T. M. Yi, Y. Huang, M. I. Simon and J. Doyle, *Proc. Natl. Acad. Sci. U. S. A.*, 2000, 97, 4649–4653.
- 10 G. von Dassow, E. Meir, E. M. Munro and G. M. Odell, *Nature*, 2000, **406**, 188–192.
- 11 N. T. Ingolia, PLoS Biol., 2004, 2, e123.
- 12 J. Stelling, U. Sauer, Z. Szallasi, F. J. Doyle III and J. C. Doyle, *Cell*, 2004, **118**, 675–685.
- 13 H. Kitano, Nat. Rev. Genet., 2004, 5, 826-837.
- 14 A. Wagner, *Robustness and Evolvbability in Living Systems*, Princeton University Press, 2005.
- 15 A. Wagner, Proc. Natl. Acad. Sci. U. S. A., 2005, 102, 11775–11780.
- 16 S. Ciliberti, O. C. Martin and A. Wagner, PLoS Comput. Biol., 2007, 3, e15.
- 17 Y. D. Nochomovitz and H. Li, Proc. Natl. Acad. Sci. U. S. A., 2006, 103, 4180–4185.

- 18 J. Draghi and G. P. Wagner, J. Evol. Biol., 2009, 22, 599-611.
- 19 C. Y. Logan and R. Nusse, Annu. Rev. Cell Dev. Biol., 2004, 20, 781–810.
- 20 S. J. Bray, Nat. Rev. Mol. Cell Biol., 2006, 7, 678-689.
- 21 S. Wullschleger, R. Loewith and M. N. Hall, *Cell (Cambridge, Mass.)*, 2006, **124**, 471–484.
- 22 R. Derynck and Y. E. Zhang, Nature, 2003, 425, 577-584.
- 23 J. Massagu, Nat. Rev. Mol. Cell Biol., 2000, 1, 169–178.
- 24 L. Kuepfer, M. Peter, U. Sauer and J. Stelling, *Nat. Biotechnol.*, 2007, **25**, 1001–1006.
- 25 A. Lorberg and M. N. Hall, *Curr. Top. Microbiol. Immunol.*, 2004, **279**, 1–18.
- 26 E. Jacinto and M. N. Hall, Nat. Rev. Mol. Cell Biol., 2003, 4, 117-126.
- 27 H. Kitano, Science, 2002, 295, 1662-1664.
- 28 R. Albert, Plant Cell, 2007, 19, 3327-3338.
- 29 A. Arkin, J. Ross and H. H. McAdams, *Genetics*, 1998, **149**, 1633–1648.
- 30 N. Barkai and S. Leibler, Nature, 2000, 403, 267-268.
- 31 A. Goldbeter, Proc. R. Soc. London, Ser. B, 1995, 261, 319–324.
- 32 C. Espinosa-Soto, P. Padilla-Longoria and E. R. Alvarez-Buylla, *Plant Cell*, 2004, 16, 2923–2939.
- 33 U. Alon, M. G. Surette, N. Barkai and S. Leibler, *Nature*, 1999, 397, 168–171.
- 34 M. B. Elowitz and S. Leibler, Nature, 2000, 403, 335-338.
- 35 B. Yi and R. J. Sommer, Development, 2007, 134, 3111-3119.
- 36 M.-A. Félix, Curr. Biol., 2007, 17, 103-114.
- 37 M.-A. Félix and A. Wagner, Heredity, 2008, 100, 132-140.
- 38 C. J. Di Como and K. T. Arndt, Genes Dev., 1996, 10, 1904–1916.
- 39 M. M. Luke, F. D. Seta, C. J. D. Como, H. Sugimoto, R. Kobayashi and K. T. Arndt, *Mol. Cell. Biol.*, 1996, 16, 2744–2755.
- 40 Y. Jiang and J. R. Broach, EMBO J., 1999, 18, 2782-2792.
- 41 E. Jacinto, B. Guo, K. T. Arndt, T. Schmelzle and M. N. Hall, *Mol. Cell*, 2001, 8, 1017–1026.
- 42 H. Wang, X. Wang and Y. Jiang, *Mol. Biol. Cell*, 2003, 14, 4342–4351.
- 43 J. R. Rohde, S. Campbell, S. A. Zurita-Martinez, N. S. Cutler, M. Ashe and M. E. Cardenas, *Mol. Cell. Biol.*, 2004, 24, 8332–8341.
- 44 H. Kitano, A. Funahashi, Y. Matsuoka and K. Oda, Nat. Biotechnol., 2005, 23, 961–966.
- 45 N. Le Novere, S. Moodie, A. Sorokin, M. Hucka, F. Schreiber, E. Demir, H. Mi, Y. Matsuoka, K. Wegner, and H. Kitano, Available from Nature Precedings, http://hdl.handle.net/10101/ npre.2008.2320.1.
- 46 A. Funahashi, N. Tanimura, M. Morohashi and H. Kitano, *BIOSILICO*, 2003, **1**, 159–162.
- 47 W.-H. Li, Molecular Evolution, Sinauer Associates, 1997.
- 48 M. Ebner, M. Shackleton and R. Shipman, *Complexity*, 2001, 7, 19–33.
- 49 F. Harary, *Graph Theory*, Addison-Wesley Publishing, Reading, 1969.
- 50 L. W. Ancel and W. Fontana, J. Exp. Zool., 2000, 288, 242-283.
- 51 C. Braendle and M.-A. Félix, Dev. Cell, 2008, 15, 714-724.
- 52 J. A. Draghi, T. L. Parsons, G. P. Wagner and J. B. Plotkin, *Nature*, 2010, **463**, 353–355.
- 53 P. Radivojac, P. H. Baenziger, M. G. Kann, M. E. Mort, M. W. Hahn and S. D. Mooney, *Bioinformatics*, 2008, 24, i241–i247.
- 54 A. Vas, W. Mok and J. Leatherwood, *Mol. Cell. Biol.*, 2001, **21**, 5767–5777.
- 55 S. Jones and J. M. Thornton, Proc. Natl. Acad. Sci. U. S. A., 1996, 93, 13–20.
- 56 J. Berg, M. Lssig and A. Wagner, BMC Evol. Biol., 2004, 4, 51.
- 57 P. Beltrao, J. C. Trinidad, D. Fiedler, A. Roguev, W. A. Lim, K. M. Shokat, A. L. Burlingame and N. J. Krogan, *PLoS Biol.*, 2009, 7, e1000134.
- 58 T. Zhang, R. Ramakrishnan, and M. Livny, Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data, Montreal, Quebec, Canada, June 4–6, 1996, 1996, pp. 103–114.