

Signatures of Arithmetic Simplicity in Metabolic Network Architecture

William J. Riehl¹, Paul L. Krapivsky², Sidney Redner², Daniel Segrè^{1,3,4*}

1 Program in Bioinformatics and Systems Biology, Boston University, Boston, Massachusetts, United States of America, **2** Department of Physics, Boston University, Boston, Massachusetts, United States of America, **3** Department of Biology, Boston University, Boston, Massachusetts, United States of America, **4** Department of Biomedical Engineering, Boston University, Boston, Massachusetts, United States of America

Abstract

Metabolic networks perform some of the most fundamental functions in living cells, including energy transduction and building block biosynthesis. While these are the best characterized networks in living systems, understanding their evolutionary history and complex wiring constitutes one of the most fascinating open questions in biology, intimately related to the enigma of life's origin itself. Is the evolution of metabolism subject to general principles, beyond the unpredictable accumulation of multiple historical accidents? Here we search for such principles by applying to an artificial chemical universe some of the methodologies developed for the study of genome scale models of cellular metabolism. In particular, we use metabolic flux constraint-based models to exhaustively search for artificial chemistry pathways that can optimally perform an array of elementary metabolic functions. Despite the simplicity of the model employed, we find that the ensuing pathways display a surprisingly rich set of properties, including the existence of autocatalytic cycles and hierarchical modules, the appearance of universally preferable metabolites and reactions, and a logarithmic trend of pathway length as a function of input/output molecule size. Some of these properties can be derived analytically, borrowing methods previously used in cryptography. In addition, by mapping biochemical networks onto a simplified carbon atom reaction backbone, we find that properties similar to those predicted for the artificial chemistry hold also for real metabolic networks. These findings suggest that optimality principles and arithmetic simplicity might lie beneath some aspects of biochemical complexity.

Citation: Riehl WJ, Krapivsky PL, Redner S, Segrè D (2010) Signatures of Arithmetic Simplicity in Metabolic Network Architecture. *PLoS Comput Biol* 6(4): e1000725. doi:10.1371/journal.pcbi.1000725

Editor: Jason A. Papin, University of Virginia, United States of America

Received: October 20, 2009; **Accepted:** February 26, 2010; **Published:** April 1, 2010

Copyright: © 2010 Riehl et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: This work was supported by grants from the Office of Science (BER), U.S. Department of Energy (No. DE-FG02-07ER64388), NASA (NASA Astrobiology Institute, NNA08CN84A), and the National Science Foundation (NSF DMR0535503, DMR0906504 and NSF CCF-0829541). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

* E-mail: dsegre@bu.edu

Introduction

The prominent role of metabolism in any biological process and the fact that a large portion of the environmental factors shaping living systems are ultimately metabolic in nature, suggest that strong selective forces have been acting on metabolic networks throughout the history of life. In laboratory evolution experiments [1–3] one can witness mostly short term metabolic adaptations, affecting metabolic enzyme regulation and fine tuning of kinetic parameters. However, especially during major transitions, such as the early stages of life's appearance or the rise of oxygen in the Earth's atmosphere, selective forces must have shaped the metabolic wiring itself [4]. Comparative genomics can provide top-down insight into some long-term evolution of metabolic pathways [5,6]. In addition, studies of prebiotic chemistry scenarios have suggested possible seeds of biochemical organization from a bottom-up perspective [7–10]. Yet, whether the long term evolution of metabolism was dominated by unpredictable frozen accidents, or by inevitable network optimization processes, remains a fundamental open question.

In a 1961 review, Baldwin and Krebs suggested that biochemical network topologies may reflect the adaptation toward optimally efficient metabolic strategies, and that manifold use of certain

molecules may be a crucial element of this adaptation, as “it is indeed a general principle of evolution that multiple use is made of given resources.” [11]. Some computational studies have proposed that the topology of specific metabolic pathways may have evolved towards maximal efficiency [12], minimal number of steps [13], or that network properties may reflect optimal organization [14]. Here we seek to address this problem by exploring a system that can reach a level of complexity comparable to the one observed in the union of all known metabolic pathways, yet is simple enough to allow efficient computation and analytical calculations. In addition, we wish to explore at an ecosystem-level the potential role of “metabolic multi-tasking”, as suggested by Baldwin and Krebs. The increasing evidence of abundant horizontal gene transfer in the history of life suggests that this question may be indeed especially relevant at the ecosystem level, where the interchange of genetic information might have created a free economy of enzymes among simple organisms, allowing for the emergence of species that share common molecular tools [15]. Recent metagenomic studies of microbial consortia [16] also suggest the question of whether metabolic functions, more than individual species distributions, might be directly dependent on environmental conditions. Hence it is possible that hallmarks of metabolic optimality in metabolic network wiring may be observable at the level of global

Author Summary

Metabolism is the network of biochemical reactions that transforms available resources (“inputs”) into energy currency and building blocks (“outputs”). Different organisms have different assortments of metabolic pathways and input/output requirements, reflecting their adaptation to specific environments, and to specific strategies for reproduction and survival. Here we ask whether, beneath the intricate wiring of these networks, it is possible to discern signatures of optimal (i.e., shortest and maximally efficient) pathway architectures. A systematic search for such optimal pathways between all possible pairs of input and output molecules in real organic chemistry is computationally intractable. However, we can implement such a search in a simple artificial chemistry, which roughly resembles a single atom (e.g., carbon) version of real biochemistry. We find that optimal pathways in our idealized chemistry display a logarithmic dependence of pathway length on input/output molecule size. They also display recurring topologies, including autocatalytic cycles reminiscent of ancient and highly conserved cores of real biochemistry. Finally, across all optimal pathways, we identify universally important metabolites and reactions, as well as a characteristic distribution of reaction utilization. Similar features can be observed in real metabolic networks, suggesting that arithmetic simplicity may lie beneath some aspects of biochemical complexity.

(multi-species) metabolic networks [17,18], more than at the individual species level. Do specific molecules, reactions or pathway topologies appear to be universally useful in a biochemical network, i.e., relevant for maximally efficient completion of several possible metabolic tasks, possibly across multiple organisms?

In the present work, we combine the study of an extremely simple artificial chemistry [19–22] with recent systems biology approaches [23,24] to systematically compute pathways that are optimal for an array of elementary metabolic tasks, converting an input molecule into an output one. Behind the apparent complexity of the ensuing pathways, we identify recurring, modularly organized categories of network topologies, and analytically predictable trends in pathway length. In addition, we observe the emergence of “universal metabolic tools” across all optimal pathways. Finally, despite the huge gap in the underlying chemical rules, we find that some properties of real metabolic pathways are consistent with the patterns detected in the model, suggesting that fundamental optimality principles may have played a role in shaping biochemical networks.

Results

Our artificial chemistry consists of a set of \mathcal{N} possible molecules $\{a_1, a_2, a_3, \dots, a_{\mathcal{N}}\}$, that can participate in reversible ligation/cleavage reactions of the form $a_i + a_j \leftrightarrow a_k$, with $i+j=k$. This model could be viewed as the simplest possible string-based artificial chemistry [19]. The reaction network $R_{\mathcal{N}}$ that includes all metabolites up to length \mathcal{N} and all possible reactions (of the order of $\mathcal{N}^2/4$, see Methods) between them (Fig. 1) can be thought of as the underlying chemistry based on which specialized metabolic tasks could emerge. Here we were concerned with pathways, within the $R_{\mathcal{N}}$ network, that can optimally perform a given metabolic task. In particular, we searched for optimal solutions to the problem of producing a specific end-product (e.g., a_j , with output flux v_{out}) from a single available nutrient (e.g., a_i , with input flux v_{in}). We define an optimal pathway as one that satisfies the

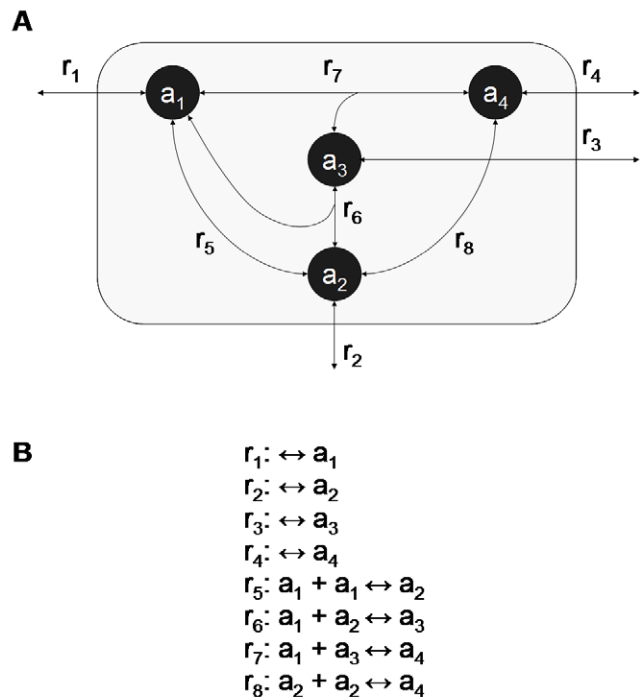


Figure 1. Representation of the R_4 network. (A) This schematic of the R_4 artificial chemistry network is composed of metabolite “strings” of a up to a maximum length of four, and all allowed reactions between them. (B) The reaction list for the R_4 network. There are four reactions that represent the exchange of mass with the environment (r_1 – r_4) – one for each metabolite – and four reactions between the metabolites (r_5 – r_8). doi:10.1371/journal.pcbi.1000725.g001

following conditions: (i) it allows a steady state solution, i.e., a mass-conserving flow from input to output; (ii) it has maximal yield, and no waste [25], such that $v_{out} = v_{in} \cdot j/i$; and (iii) it has the fewest reaction steps possible. A pathway satisfying these conditions is termed a *minimal balanced pathway* (MBP) between a_i and a_j . MBPs (also referred to below as optimal pathways) can be thought of as the pathways that are most efficient for a specific metabolic task, in the sense that they require the smallest possible number of different enzymes for producing the maximal possible yield [12,26,27].

Despite the simplicity of our artificial chemistry, identifying the MBPs between all possible input-output pairs in a given artificial chemistry $R_{\mathcal{N}}$ is a challenge for large \mathcal{N} . We implement three algorithms to approach this problem: a mixed integer linear programming (MILP) akin to flux balance analysis (FBA) [28]; an algorithm that uses enumeration of elementary flux modes [23]; and finally an iterative algorithm that gradually assembles new MBPs from already identified simple ones (see Methods). The three algorithms differ mainly in their scalability, and in their capacity to predict multiple degenerate solutions (see Table S1). A partial overview of the results of our calculations is shown in Fig. 2A and Fig. S1 (see Tables S2 and S3 for a comprehensive list of MBPs).

Behind the apparent complexity of the topologies encountered in each of the different pathways, it is possible to observe the recurrence of three fundamental categories: each MBP functions either as a pure “addition chain” [29], where smaller metabolites are progressively added together to build the target molecule, or as an “addition-subtraction chain”, in which metabolites are both synthesized and degraded within the pathway. Addition and

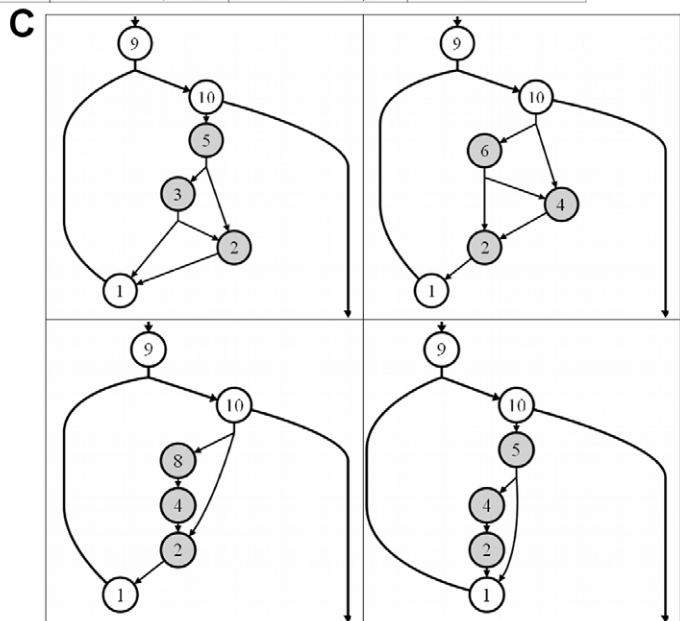

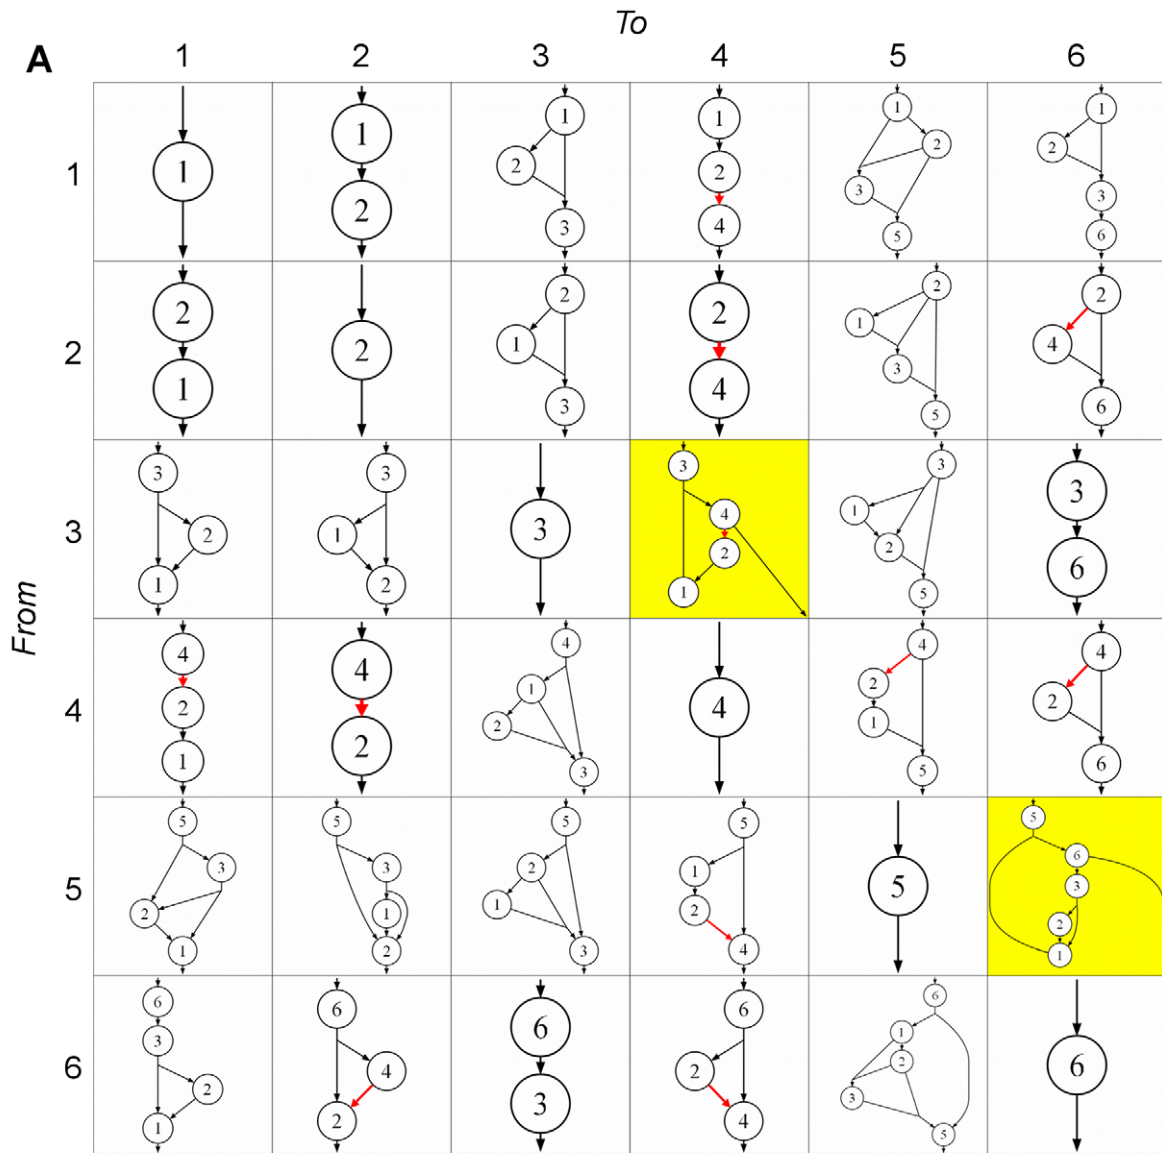


Figure 2. Emergent complexity and modularity in artificial network topology. (A) This set of example MBPs displays the emergent modularity of structure and function, as well as the multiple usage of different reactions. Each row denotes the input metabolite used, and each column the output metabolite. The reaction marked in red interconverts $a_2 + a_2 \leftrightarrow a_4$ and is the most used reaction. MBPs on a yellow background are autocatalytic cycles. For a larger image with more examples, and a more detailed view of the properties observed in these networks, see Figure S1. (B) A single autocatalytic loop is used as a modular backbone for several MBPs: the cycle that constructs a_8 from a_7 is also used to produce a_1 , a_2 , and a_4 . (C) Four examples of the MBP that produces a_{10} from a_9 . Each breaks down a_{10} into a_1 in different but equally optimal ways. Each of these sub-pathways (shaded metabolites) is an MBP in itself, showing the modularity of use of each of these metabolic tools.
doi:10.1371/journal.pcbi.1000725.g002

addition-subtraction chains are concepts borrowed from the field of cryptography, whose relevance to our question will become apparent later. There is also a third, smaller category of cyclical pathways that cannot proceed unless a certain intermediate molecule is already present in the system. These pathways are autocatalytic cycles (Fig. 2B) that very much resemble autocatalytic cycles found in real biochemistry, such as the reverse TCA cycle [7], or the formose reaction [30]. Our results show that autocatalytic cycles can be simultaneously optimal for multiple tasks (Fig. 2B), suggesting that such types of structure may have a fundamental evolutionary advantage in a biological context. In addition to the recurrence of these topological categories among MBPs, we find that some specific structures are used repeatedly, often in a modular fashion (Fig. 2C). Specifically, many simple MBPs are used hierarchically as a toolkit for the construction of progressively more complex MBPs (data not shown), similar to what has been observed in real metabolic networks [15,31,32].

This modular architecture of recurring graph types provides a topological signature of optimally efficient pathways in our idealized chemistry. Since these pathways are chosen based on their minimal length, one may expect that a systematic analysis of all MBP lengths will display additional distinctive properties. Indeed, pathway lengths increase roughly logarithmically with the size of the input (or output) molecule (Fig. 3 and S3), with superimposed sharp jumps. For example, the task $a_9 \Rightarrow a_6$ can be performed in 2 steps, but the neighbor task $a_9 \Rightarrow a_7$ requires a minimum of 6 steps. Moreover, while most MBPs have only one or a few optimal realizations, selected instances display a peak in possible redundant solutions (Fig. S2), usually due to interconversions between molecules of similar size (e.g., $a_x \Rightarrow a_{x+1}$), or to the inherent complexity of a specific molecule (e.g., $a_7 \Rightarrow a_j$). These regular patterns suggest that it may be possible to reproduce the MBP length curves without having to actually compute the MBPs.

A similar search for patterns associated with minimal steps had been previously encountered in the mathematics of addition-subtraction chains, of high importance in cryptography [29]. These are integer sequences, beginning with 1, in which the i -th entry is either the sum or difference of any two previous entries in the sequence. These chains are often used in calculating large exponents of numbers [33]. For example, calculating n^{128} can either be performed in 127 multiplications ($n \times n = n^2$, $n^2 \times n = n^3$, ..., $n^{127} \times n = n^{128}$) or in a chain of 7 exponent multiplications ($n \times n = n^2$, $n^2 \times n^2 = n^4$, $n^4 \times n^4 = n^8$, ..., $n^{64} \times n^{64} = n^{128}$). The latter can be further simplified by tracking the sums of the exponents in each calculation, which form an addition chain (1, 2, 4, 8, 16, 32, 64, 128). Shortest addition-subtraction chains are commonly used to calculate very large numbers in the fewest number of steps, thus speeding up computation time. These are often applied to methods in cryptography where the calculated exponents can have on the order of thousands to tens of thousands of bits [34].

The pathways explored in our model resemble optimal addition-subtraction chains. For example, the problem of obtaining a_{128} from a_1 is formally equivalent to the addition chain example described above. However while typical addition-subtraction chains start with

the number 1, in our MBPs we explore minimal paths starting from any molecule a_i ($i \geq 1$). As described in detail in the Methods, we extended previous work on addition-subtraction chains [29,35] to derive the following analytical estimate of the length of MBPs:

$$L(i,j) \sim \log_2 \frac{i}{\gcd(i,j)} + \log_2 \frac{j}{\gcd(i,j)} \quad (1)$$

where $L(i,j)$ is the number of reactions in the MBP with input a_i and output a_j , and $\gcd(i,j)$ is the greatest common divisor of i and j . As seen in Fig. 3, Eq. (1) reproduces the corresponding pathway lengths obtained by computing individual MBPs. This agreement implies that the number of reaction steps needed to construct an efficient metabolic pathway between two metabolites in our artificial chemistry can be roughly estimated from Eq. (1). The only feature that determines the pathway lengths is the complexity of the input and output molecules.

We can now ask whether similar minimal pathway length signatures are discernible in real metabolic networks. To cope with the gap in complexity between our model and real chemistry, we mapped real metabolic networks onto a single atom backbone [13,14]. For example, the aldolase reaction, which cleaves fructose-1,6-bisphosphate ($C_6H_{14}O_{12}P_2$) into dihydroxyacetone phosphate ($C_3H_7O_6P$) and glyceraldehyde-3-phosphate ($C_3H_7O_6P$), can be mapped onto a carbon atom backbone, becoming simply $C_6 \leftrightarrow C_3 + C_3$ (see Methods). This reaction is now formally analogous to the $a_6 \leftrightarrow a_3 + a_3$ reaction in the idealized chemistry. Upon performing this mapping onto a carbon atom backbone, we ask whether the structure of real metabolic networks allows interconversions that use the minimal, logarithmic number of steps found for the artificial chemistry (Fig. 3 and Eq. (1)). Specifically, we identified all shortest pathways between any two carbon compounds in *Escherichia coli*'s metabolic network. This was performed using two methods. The first was an explicit use of elementary flux modes as done in the artificial chemistry. As with the artificial chemistry method, this has the advantage of finding all of the shortest pathways that connect any two carbon compounds, but is limited in computational scope to a smaller network. Because of this limitation, we used the network of *E. coli*'s central carbon metabolism [36,37], modified to remove cofactors and reactions that do not affect carbon transfer (see Methods). After finding all minimal elementary flux modes that connect every pair of carbon compounds in the network, we reduced those compounds to their carbon content alone, as described above. We determined, for each input compound, the length of the shortest elementary flux mode that reaches its closest molecule with j carbons; then, for each value of i , we averaged these path lengths over all input molecules with i carbons. The results show that the lengths of the *E. coli* elementary flux modes correlate with the lengths of the corresponding artificial chemistry MBPs and with the analytical predictions, though the actual *E. coli* values are overall larger than the artificial chemistry ones (Fig. 4). This last fact, as discussed later, may be due, for example, to energetic constraints, or to the higher complexity of real organic chemistry.

The second method is aimed at identifying all shortest pathways between any two carbon compounds in the whole genome-scale

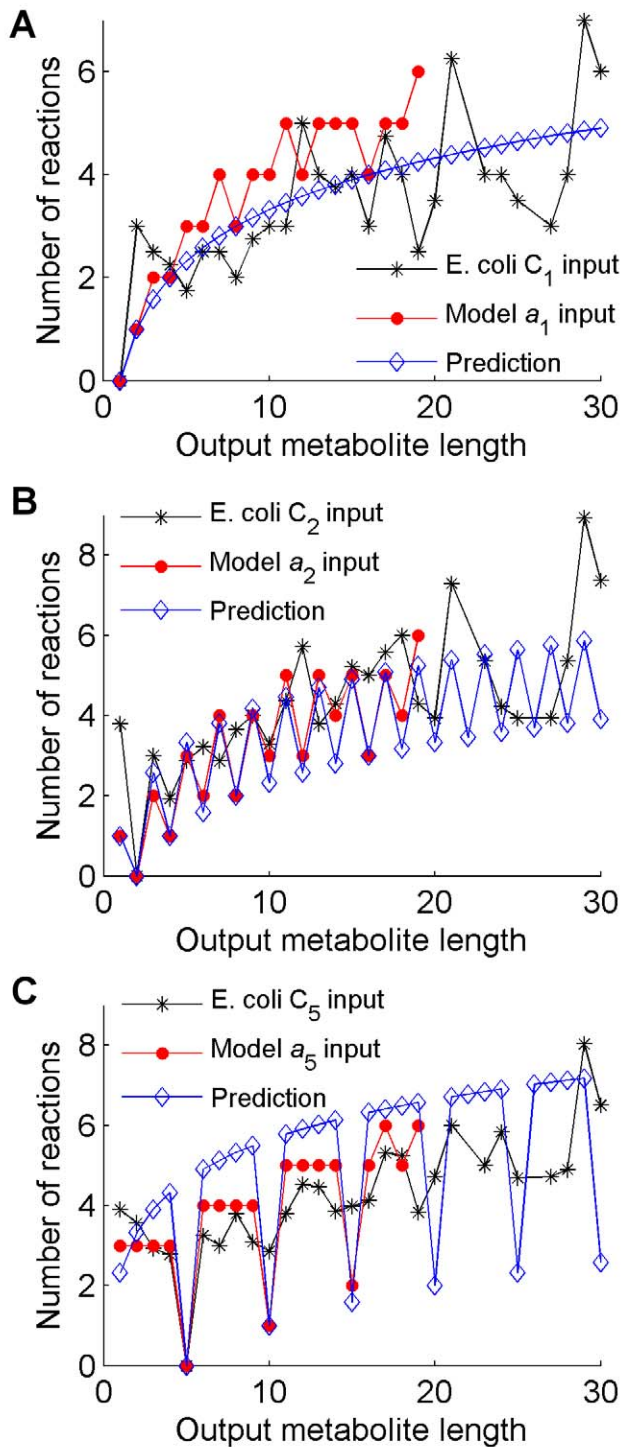


Figure 3. Logarithmic growth is observed among pathways from both the R_{19} and *E. coli* metabolic networks. Each plot has a single starting value i corresponding to either a_i (for the model) or C_i (for *E. coli*). The red lines show the number of reactions in each MBP with different output metabolite size, and the black lines show the average number of reactions used to reach the nearest metabolite of increasing carbon number in *E. coli*. The predicted number of reactions from Equation (1) is also shown for each plot in blue. (A) C_1/a_1 input. (B) C_2/a_2 input. (C) C_5/a_5 input.

doi:10.1371/journal.pcbi.1000725.g003

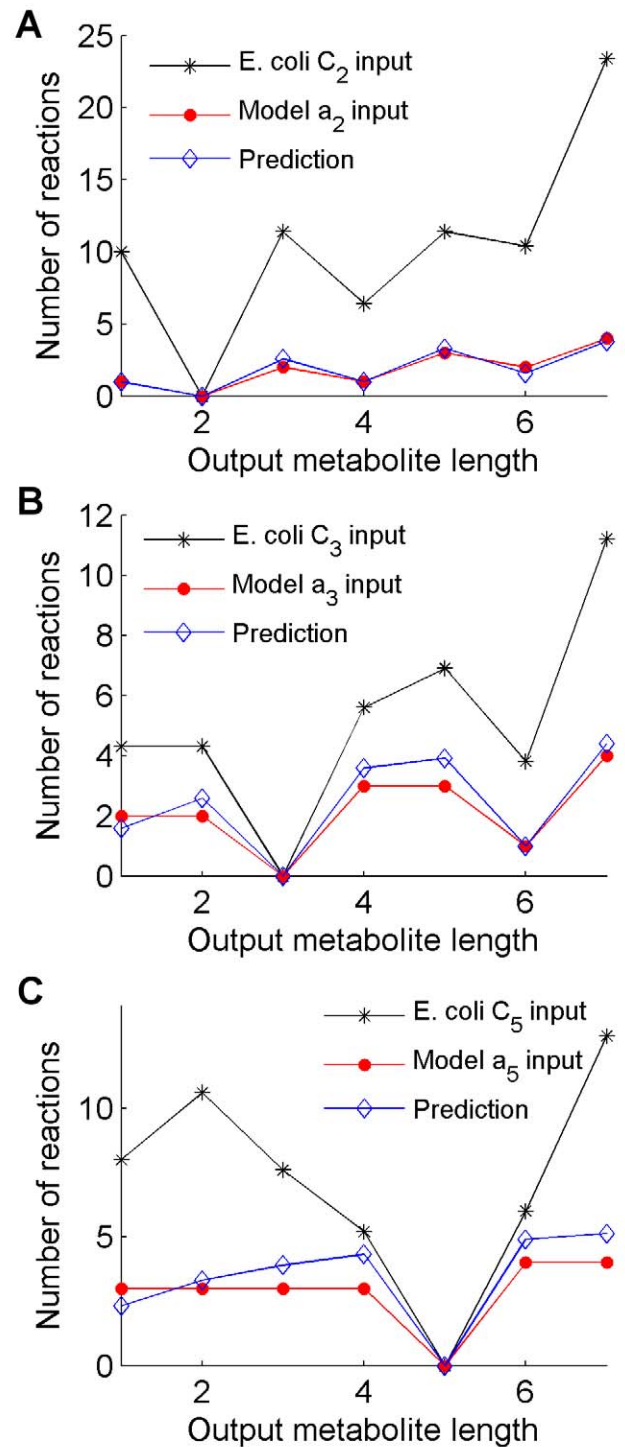


Figure 4. Similar distributions are observed among MBPs from the R_{19} network and minimal elementary modes from the *E. coli* central carbon metabolic network. Each plot has a single starting value i corresponding to either a_i for the model or C_i for *E. coli*. The red lines show the number of reactions in each MBP with different output metabolite size, and the black lines show the average number of reactions used to reach the nearest metabolite of increasing carbon number in *E. coli* central carbon metabolism. The predicted number of reactions from Equation (1) is shown for each plot in blue. (A) C_2/a_2 input. Correlation coefficient = 0.92, p-val = 0.003. (B) C_3/a_3 input. Correlation coefficient = 0.94, p-val = 0.001. (C) C_5/a_5 input. Correlation coefficient = 0.78, p-val = 0.04.

doi:10.1371/journal.pcbi.1000725.g004

metabolic network of *Escherichia coli* [38], for which it is still infeasible to apply the elementary flux mode analysis. For this, we implemented a heuristic approach to analyze the set of shortest pathways between every pair of metabolites. We first determined, for each input compound, the minimal path length to reach its closest molecule with j carbons; then, for each value of i , we averaged these path lengths over all input molecules with i carbons. The results (Fig. 3 and S3) show that these *E. coli* minimal path lengths approximately follow the predicted logarithmic trend. For some curves (e.g. the one with C_5 as an input), the specific peaks and valleys of the predicted function are closely followed by the *E. coli* network. While this does not prove that MBPs are indeed used in real metabolic networks, it suggests that the logarithmic strategy of MBPs is embedded in their architecture. However, because this second method focuses on shortest paths across a metabolite-to-metabolite network rather than on flow-conserving MBPs, the predicted values are likely an underestimate of the number of reactions necessary to construct one metabolite from another in a mass-conserved manner.

So far, we have analyzed the properties of individual MBPs in our idealized chemistry, as well as analogous minimal length pathways in *E. coli* metabolism. However, some fundamental aspects of the architecture of metabolism may be visible only at the ecosystem-level, namely by collectively analyzing the metabolic network obtained as the union of the metabolic maps of known individual species (sometimes called the “meta-metabolome”). For example, previous work using an algorithm of network expansion applied to this meta-metabolome has identified potential signatures of major evolutionary events [39], including the metabolic transition that took place upon the great oxidation event, about two billion years ago [4].

Here, we build a meta-metabolome for our idealized chemistry by considering the collection of MBPs. One could imagine that each task $a_i \Rightarrow a_j$ corresponds to a different organism, which has filled a specific metabolic niche (availability of a_i), and found an optimal solution (the MBP) for its main metabolic task (produce a_j). The question we ask next is whether, in this ecosystem of MBPs, all metabolites and reactions are used in roughly the same number of pathways, or if specific metabolites or reactions seem to be essential for many optimal tasks, hence representing “universal tools”.

For this analysis we used the set of MBPs calculated on the R_{19} network using the MILP method. One first result of this analysis is that every metabolite of an even length is used in many more MBP reactions than their odd length neighbors, compared to the underlying chemistry (Fig. 5B). Thus even-length metabolites are more important in that they can be used for more tasks. A possible explanation for this enhanced importance comes from the logarithmic nature of the MBP path lengths. For example, producing a_8 from a_1 requires only three doubling reactions ($a_1 + a_1 \rightarrow a_2$, $a_2 + a_2 \rightarrow a_4$, and $a_4 + a_4 \rightarrow a_8$). In addition, this same pathway, with one additional reaction, can also be used to optimally produce a_9 and a_{10} (see Tables S2 and S3), overall increasing the number of pathways in which each of those even-length intermediates is used. Indeed, because similar logarithmic pathways can be used as a backbone connecting distant inputs and outputs, we expect metabolites of an even length to appear more often. Similarly, one can address the relevance of each possible reaction across different MBPs. The existence of ubiquitous reactions is visible in Fig. 2A and Fig. S1, and can be more systematically assessed by plotting a usage distribution (Fig. S2). The most abundant reactions – the “universal tools” in this model chemistry – are the ones that ligate two identical molecules (e.g. $a_2 + a_2 \leftrightarrow a_4$, see Tables 1, S4, and S5). The distribution of reaction

utilization follows a long-tailed distribution (Fig. 5A), whose fit to a power law gives an exponent of approximately -1.1 ($R^2 = 0.99$). This value is close to our theoretically predicted value of -1 (See Methods).

As in the case of the artificial chemistry network, we can now search for patterns of metabolites and reactions usage in the collective set of all metabolic reactions known in living systems, obtained from the KEGG database [17]. The presence of such signatures would suggest a long-term selective advantage of molecules and reactions that are useful for multiple tasks across different organisms and environments. By counting how many times each possible carbon backbone reaction is used across this biosphere-level metabolism we obtained a broad distribution, and a fit to a power-law gives the exponent of -0.89 , comparable with the analytically predicted value, and with that in the artificial chemistry model (Fig. 5A and Fig. S4). We found that several reactions that are top ranking in their count across MBPs in the artificial network, are also at the top of the list in the KEGG-derived reactions (Spearman correlation p-value $< 10^{-6}$; see also Tables 1 and S6, S7, S8, S9, S10, and S11). This suggests that the R_N network model, despite its simplified chemical rules, captures some fundamental features of the role of the carbon reaction backbone of real metabolic networks.

In addition to a preference for specific reactions, we can ask whether the spectrum of metabolite usage across the whole KEGG metabolism reflects the possible optimality criteria encountered in the model (Fig. 5C). The metabolite usage in the hydrogen backbone network (see Methods) is similar to that in the artificial chemistry: each even-length hydrogen metabolite is used more often than its odd-length neighbors (Fig. S4A). For the carbon backbone distribution, we see a similar descending periodic behavior, but with a periodicity of approximately 5 (Fig. 5C and Fig. S5B). Hence, molecules containing carbons in a number that is a multiple of 5 are used more abundantly than other molecules across different metabolic reactions. One possible explanation for this C_5 periodicity is the profuse usage of adenine and nicotinamide adenine dinucleotide compounds as energy carriers and redox balance molecules, although the removal of such compounds has little effect on the observed periodicity (Fig. S6). Hence, the prominent usage of compounds with specific numbers of carbons might reflect global network optimization principles for the efficiency of multiple pathways, as observed in the artificial chemistry model. The periodicity of 5 that we observe, together with the evidence displayed in Fig. 3C, may suggest that the evolutionary optimization of metabolism has been partially taking place around building blocks of five carbons, compatible with previous observations of prebiotic abundance of terpenoids [40] and pentoses [41]. It is also interesting to note that an unexplained periodicity of two had been previously observed in the distribution of the number of carbons among known organic compounds [42–44]. While our analysis is based on the distribution of usage of carbon compounds in different reactions, rather than the total count of molecules, future analyses may investigate possible connections between these trends.

Discussion

We have explored the potential existence of general principles underlying the evolution of metabolic network architecture. Specifically, we studied the properties of pathways (the MBPs) that perform elementary metabolic tasks with maximal yield and minimal length in an idealized chemistry. Using the results from the model chemistry, we asked whether similar signatures of optimally efficient organization could be found in real metabolic networks.

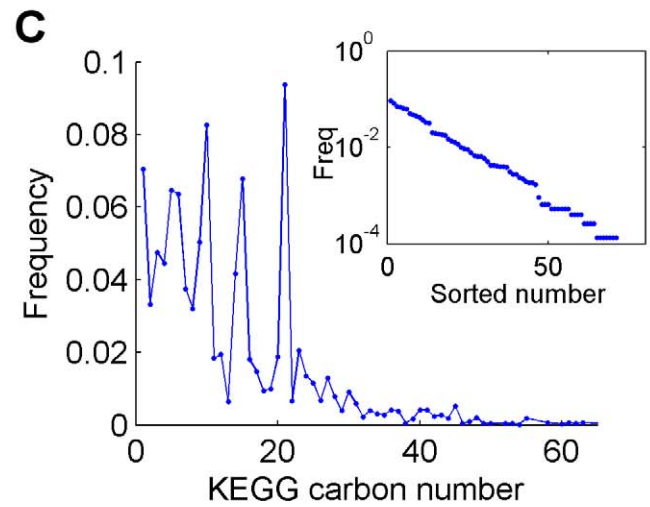
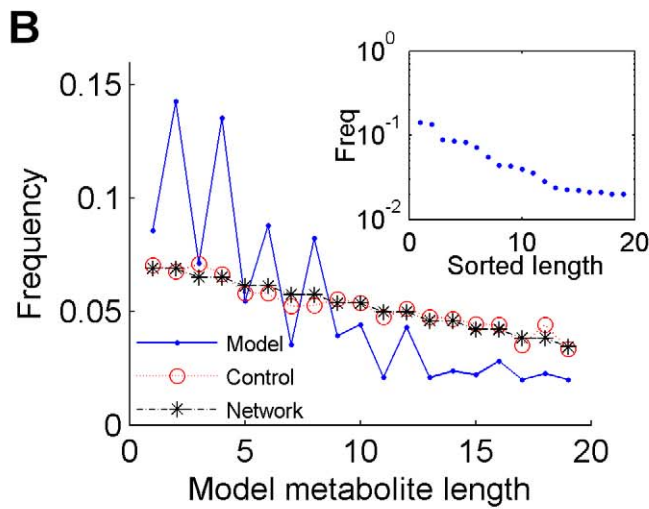



Figure 5. Metabolite and reaction usage frequencies. (A) The frequency of usage of reactions in the artificial chemistry model and in the KEGG-derived carbon reaction set. MBPs for the R_{19} network were calculated using the MILP method while the others (R_{30} through R_{100}) were estimated using the iterative algorithm. We calculated the reaction usage by counting the number of MBPs that use each reaction. These were then ranked in descending order, yielding curves that follow a power law with an average exponent of -1.14 (± 0.03) ($R^2 = 0.99$). The reaction usage in the KEGG-derived carbon dataset was calculated by counting the number of times each equivalent reaction appears, and follows a power law tail distribution, with exponent -0.89 . The curve predicted by the analytical model, with exponent -1 , is shown as a solid line (see Methods). (B) The usage frequency of each metabolite among all MBPs in the R_{19} model. This also shows the frequency of use of each metabolite in the R_{19} network itself, and in a randomly chosen set of reactions (control). In the inset, the metabolite usage was sorted by rank and plotted on a semilog axis. (C) The usage frequency of each metabolite among all reactions in the KEGG carbon reaction set. The inset shows the usage sorted by rank on a semi-log scale. doi:10.1371/journal.pcbi.1000725.g005

In computing possible MBPs, we have focused mostly on identifying modular features, on predicting their lengths, and on the statistics of usage of metabolites and reactions. In the future, it may be interesting to characterize the full spectrum of degenerate MBPs for large artificial chemistries. This would allow us to assess, for example, the density of specific topologies (such as autocatalytic cycles), or the dependence of degeneracy on the numerical properties of input/output pairs. One of our algorithms (the elementary modes one) can find a large number of degenerate solutions, including autocatalytic cycles. This algorithm is currently not scalable, because of the difficulty of computing elementary flux modes, especially in the highly connected artificial chemistry network we have used, though very recent improvements in elementary flux mode calculations [45] might be useful for this enumeration. In addition, while intuitively this approach seems to capture the full degeneracy of MBPs, this still remains to be formally proven. Another intriguing possibility might be to modify our flux balance MILP approach to identifying degenerate solutions by employing integer cuts, as described in [46]. Alternative avenues for optimization using Linear Programming rather than MILP could be in principle devised to reduce the complexity of calculations. For example minimizing the sum of absolute values of fluxes allows for rapid calculation of pathways up to the R_{100} network, though in this case the ensuing pathways are not of minimal length (see Fig. S7). In any case, for the purpose of the current work, we verified that degeneracy does not affect the statistics of usage of different reactions (Fig. S8).

Among the recurrent MBP topologies identified, we encountered numerous autocatalytic cycles. The properties of autocatalytic cycles have been studied previously [8,11], and their self-replication

potential has been theorized to be important in the early evolution of carbon fixation [7,9]. Autocatalytic cycles have also been shown to be kinetically stable, even in the absence of regulatory control [47]. We found that some autocatalytic pathways (e.g. the pathway from a_7 to a_8) are simultaneously optimal for multiple metabolic tasks. In this specific case, the MBP $a_7 \Rightarrow a_8$ is also an MBP for the production of each intermediate in the cycle (Fig. 2B). This special property of autocatalytic cycles in our artificial chemistry may have a parallel in real metabolism. For example, many metabolites in the TCA cycle (which is autocatalytic when run in reverse [7]) are precursors for fundamental anabolic processes [48,49]. Similar properties can be observed in the fundamental autocatalytic cycle known as the formose reaction [30].

Along with the structural details of MBPs, we also used analytical methods to estimate the length of MBPs as a function of the length of input and output molecules. This estimate closely matched the lengths of the artificial chemistry pathways computed with numerical algorithms. These calculations establish a new link between two apparently unrelated disciplines, namely the mathematics of addition-subtraction chains and biochemistry. It will be interesting to explore in the future whether extensions to more realistic artificial chemistries can be formalized in a similar fashion. Conversely, the MBP length estimate obtained for biochemical pathways may suggest new avenues in applied mathematics.

To determine whether predictions of minimal MBP length in our idealized chemistry could have implications in real biochemistry, we searched for pathways of minimal length between compounds with different counts of carbon atoms in the *E. coli* metabolic network. The complexity of real chemistry relative to our idealized system made this comparison difficult to interpret. MBPs in *E. coli* were found to be composed of many more reactions than in the artificial chemistry. This might be due to the additional requirement of energy gradients (e.g. the phosphorylation steps), to the complex interdependence of multiple elements, and to the properties of stability of intermediates.

Previous work had addressed the question of optimality in specific metabolic pathways. For example, Meléndez-Hevia and Torres [50], used optimality criteria to infer that some metabolic pathways, most notably the pentose-phosphate pathway, can be traversed using the fewest number of reactions. Heinrich and Schuster [51], conversely, describe the identification of a series of phosphorylation/dephosphorylation and ATP consumption/production steps that maximize the flux of ATP production in central carbon metabolism pathways. In contrast, our search for MBPs in the artificial chemistry model corresponds to a systematic search for maximal yield, minimal length pathways for all possible input-output relationships (something not yet feasible for real metabolic pathways). This allows us to infer analytical relationships and potential principles that may hold (perhaps in an approximate way) for virtually any evolved metabolism. In real systems, various compounds are often produced as waste byproducts and simply excreted into the environment, leading to suboptimal yields of target metabolites. For example, one could go in one step from a C_6 compound to C_5 , with a yield of $5/6 \sim 83\%$ and an additional

Table 1. Of the top 10 most used reactions in the R_{19} network and carbon-only KEGG network, there are five equivalent reactions that appear in both.

	Model reaction	KEGG carbon reaction
1	$a_2 + a_2 \leftrightarrow a_4$	$C_6 + C_6 \leftrightarrow C_{12}$
2	$a_1 + a_1 \leftrightarrow a_2$	$C_1 + C_5 \leftrightarrow C_6$
3	$a_4 + a_4 \leftrightarrow a_8$	$C_1 + C_3 \leftrightarrow C_4$
4	$a_3 + a_3 \leftrightarrow a_6$	$C_1 + C_4 \leftrightarrow C_5$
5	$a_2 + a_4 \leftrightarrow a_6$	$C_5 + C_5 \leftrightarrow C_{10}$
6	$a_6 + a_6 \leftrightarrow a_{12}$	$C_1 + C_7 \leftrightarrow C_8$
7	$a_1 + a_2 \leftrightarrow a_3$	$C_1 + C_8 \leftrightarrow C_9$
8	$a_8 + a_8 \leftrightarrow a_{16}$	$C_3 + C_3 \leftrightarrow C_6$
9	$a_5 + a_5 \leftrightarrow a_{10}$	$C_4 + C_5 \leftrightarrow C_9$
10	$a_1 + a_4 \leftrightarrow a_5$	$C_1 + C_2 \leftrightarrow C_3$

Recognizing that 74 of the 90 possible reactions from the R_{19} set are found in the 631 carbon-only reactions from KEGG, we can use the Spearman rank-correlation to find that this has a correlation value of 0.54 with p-value $8 \cdot 10^{-7}$. doi:10.1371/journal.pcbi.1000725.t001

C_1 byproduct; yet, obtaining maximal yield in this transformation, would cost at least 3 more reactions. One could argue that the combination of all of these criteria, including the maximization of ATP production and optimization of enzymatic catalysis may have played a key role in the evolution of modern metabolism, leading to compromise solutions. Exploring pathways that produce multiple compounds from multiple inputs with the addition of thermodynamic constraints might constitute an interesting model extension for further investigation.

We found that the statistical properties of the usage of reactions across MBPs recapitulate the statistics of reaction usage in the union of all known metabolic pathways (represented by the KEGG metabolic database). Both across the set of all MBPs for the idealized chemistry, and in the KEGG metabolic map, we observed that a few reactions are used far more often than many of the others in the set. Another way of determining the importance of individual reactions in the context of the global functionalities of a meta-metabolome would be to perform perturbation experiments. We implemented such an experiment in our idealized chemistry, by progressively removing reactions and checking how many metabolites can still be produced. Depending on whether reactions are removed in random order or in the order determined by their usage across MBPs, the outcome is quite different (Fig. S9). This analysis sheds light on the importance of reactions in terms of the capacity to produce a certain output.

Determining to what extent real metabolic networks obey optimality principles like the ones described here will take additional effort. Even if an underlying arithmetic simplicity governs idealized optimal pathways, deviations from ideal behavior should be expected. For example, parallel selection pressures for energy production and biochemical stability would likely sacrifice pathway minimality. However, guiding principles as the ones we are proposing could serve as reference points for future research, including circumstances in which metabolism can be different from what we are used to. Using synthetic biology techniques, for example, it might be possible to redesign metabolic pathways so as to approach predicted ideal efficiencies and minimal enzyme cost [52,53]. From a totally different perspective, in the field of astrobiology, having a prediction of possible signatures of an evolved metabolism might help select, among the molecular spectra of extrasolar planets, those possibly indicative of biogenic processes [54–56].

Methods

Artificial Chemistry Model

We define an artificial chemistry inspired by previous string-based artificial chemistries (see also main text and Fig. 1). One may think of molecules in this artificial chemistry as polymers (up to a given length N) of a monomeric unit a . Since no specific assumption is made in the model about the nature of these molecules, they could equally represent aggregates or branched polymers of different sizes, as well as molecules with different counts of a specific atom. A network $R_N = \{M_N, C_N\}$ is defined by the set of N molecules $M_N = \{a_i \mid \forall i = 1, \dots, N\}$ and the set of all possible uni-bi ligation/lysis reactions between them, $C_N = \{a_i + a_j \leftrightarrow a_k \mid \forall i, j, \text{ and } k, \text{ such that } i \leq j, i+j=k, \text{ and } k=2, \dots, N\}$. The number of reactions in C_N can be shown to grow quadratically with N . Since each reaction describes how to reversibly combine two molecules to make a larger one, there is a fixed number of ways to produce any given metabolite a_i , equal to $\lfloor i/2 \rfloor$. The number of reactions in C_N is then given by $\sum_{i=1}^N \lfloor i/2 \rfloor$, which can be approximated by $N^2/4$.

Flux Balance Analysis

Flux Balance Analysis (FBA) is a steady state constraint-based approach to study the flow of mass through metabolic networks [28,57,58]. Briefly, FBA represents the metabolic network of interest as an $n \times m$ stoichiometric matrix S , whose element S_{ij} indicates the number of molecules of metabolite i ($i = 1, \dots, m$) that participate in reaction j ($j = 1, \dots, n$) (with a positive sign if the metabolite is produced, negative if it is consumed). Each reaction can be associated with a flux, v_j . Under the assumption of a steady state the following set of mass conservation constraints on the fluxes is generated:

$$\sum_{j=1}^n S_{ij} v_j = 0 \quad i = 1, 2, \dots, m \quad (2)$$

Additional constraints (such as availability of nutrients, experimentally observed irreversibility, maximal or minimal rates, etc.) can be imposed on the fluxes as inequalities of the form

$$\alpha_j \leq v_j \leq \beta_j \quad (3)$$

where α_j is the minimal allowed rate of a reaction and β_j is its maximal rate. Taken together, the above constraints define a convex polyhedron (the “feasible space”) in the n -dimensional space of fluxes. Linear programming (LP) can be used to identify, within the feasible space, flux vectors that maximize or minimize a given linear objective function. In microbial systems it has been often hypothesized that a biologically meaningful objective is the maximization of the flux through the reaction that represents cellular growth, or biomass production [2,59]. Hence, LP applied to FBA provides a prediction of all metabolic fluxes in a cell. FBA can be applied at genome scale, and corresponding stoichiometric models are available for a number of organisms. FBA predictions have been experimentally validated most thoroughly in *Saccharomyces cerevisiae* SC288 [60] and *E. coli* K-12 [38].

Minimal Balanced Pathway Discovery Algorithms

Minimal Balanced Pathways (MBPs) are defined as sets of reactions in the R_N network that can optimally perform a given metabolic task. A task is defined as the production of a specific end product (e.g., a_j with output flux v_{out}) from a single available nutrient (e.g., a_i with input flux v_{in}). A pathway between two molecules is a MBP if (i) it satisfies a steady state solution, analogous to Eq. (1); (ii) it produces the final product with maximal yield, i.e., $v_{out} = v_{in} \cdot j/i$; and (iii) it contains the smallest possible number of reaction steps. The MBP between a_i and a_j will be indicated as $a_i \Rightarrow a_j$.

We have developed three different algorithms for computing MBPs, as described below:

Flux Balance Analysis/Mixed Integer LP Algorithm

We use a modified FBA approach to formulate the MBP problem in a constrained optimization framework. Specifically, we impose the same constraints used in an FBA problem, and further require that the maximal yield condition $v_{out} = v_{in} \cdot y/x$ for the MBP $a_x \Rightarrow a_y$ be satisfied. We then search for a solution that minimizes the number of active (nonzero) fluxes. Towards this goal, we use a modification of the LP problem described above to introduce binary variables (b_j) that represent flux activity: $b_j = 0$ if $v_j = 0$, and $b_j = 1$ otherwise. To identify a minimal path, we can then search for the set of fluxes that minimize $\sum_i b_i$. Because of the nature of the

variables involved – the fluxes are continuous, and the number of active fluxes is an integer – this problem must be solved using a mixed integer linear programming (MILP) algorithm. Our MILP problem for the optimal MBP satisfying $a_x \Rightarrow a_y$ is formulated as follows:

$$\begin{aligned} & \text{Minimize} && \sum_{j=1}^N b_j \\ & \text{Subject to:} && \\ & && \sum_{j=1}^N S_{ij} v_j = 0 \quad (4) \\ & && \alpha_j b_j \leq v_j \leq \beta_j b_j \quad \text{for } j = 1, 2, \dots, N \\ & && v_{out} = \frac{y}{x} v_{in} \\ & && b_j \in \{0, 1\} \quad \text{for } j = 1, 2, \dots, N \end{aligned}$$

The optimal solution for this problem will give the flux distribution v that uses the fewest nonzero values to maximize the objective. In our MBP computations, the only flux constraints used were those that limit the uptake of the single nutrient to an arbitrary value of $10 \text{ mmol}\cdot\text{gDW}^{-1}\cdot\text{h}^{-1}$, and the production of the target metabolite to the known maximal yield $v_{out}/v_{in} = j/i$.

Elementary Flux Modes Algorithm

Given a metabolic network defined by a stoichiometric matrix S (as described in the above FBA section), a vector of fluxes v is said to correspond to an elementary flux mode (EFM) if it satisfies the following three conditions [23].

1. It satisfies the steady state condition ($Sv = 0$).
2. It must be feasible within the conditions of the model: if there are known boundaries for the fluxes, then v must fall within them.
3. It must be non-decomposable. There are no two smaller EFMs that can be linearly combined to form the one in question.

Because of these constraints, those EFMs that use the minimal number of reactions satisfy the requirements for being an MBP. We used the METATOOL software package [61] to find all EFMs in the R_{10} network, and then identified all of those EFMs that are also MBPs.

Iterative Additive Algorithm

We designed and implemented an algorithm to produce most MBPs *de novo*, without relying on prior steady state stoichiometric modeling methods. The algorithm works in an iterative manner, producing longer pathways from shorter ones. For example, we can start from two trivial MBPs: $a_1 \Rightarrow a_1$ (which requires no reactions), and $a_1 \Rightarrow a_2$ (requiring one reaction, $a_1 + a_1 \rightarrow a_2$). To compute $a_1 \Rightarrow a_3$, we identify all the ways in which we can decompose 3 into two smaller addends (in this case, only one: $3 = 2+1$). Next we combine together the previously computed MBPs that progress from a_1 to each of these two addends, giving a new putative MBP for the desired new task ($a_1 + a_1 \rightarrow a_2$, and $a_1 + a_2 \rightarrow a_3$). This procedure can be then iterated to give a prediction of MBP $a_i \Rightarrow a_j$ ($i, j \leq N$).

This algorithm is fast and efficient compared to the previous methods, allowing us to apply it to even the R_{100} network. However, it has two main drawbacks. First, it will miss pathways

that “overshoot” the target value then subtract down to it. Second, it may miss MBPs that are not built modularly from smaller ones. From a comparison of the MBPs predicted by the different algorithms, one can see that the approximations introduced in this algorithm cause 18 out of 361 MBPs (5%) in R_{19} to overestimate pathway length by one reaction. Also, this algorithm correctly identifies 204 of the 384 degenerate MBPs that the EFM algorithm finds in R_{10} . The reaction usage using this method is highly correlated with that of the MILP method applied to R_{19} (Pearson correlation = 0.96, p-val = 10^{-51}), and the EFM method applied to R_{10} (Pearson correlation = 0.98, p-val = $2 \cdot 10^{-17}$).

KEGG Reaction Reduction

Data used for the comparison between the R_N metabolic network and real metabolism was gathered from the KEGG LIGAND database (July 26, 2009 release)[17]. This database was parsed to convert its compounds and reactions into a single-atom form, as described in the text. Compounds that carried any uncertainty in their atomic makeup, including non-specific side-chains or variable chain length were removed from the current analysis. We also removed from the analysis reactions with no associated formula, as well as reactions involving non-specific molecules (such as glycans and non-specific nucleotide or peptide chains). Finally, a number of reactions were found to leave the atomic composition of the compounds essentially unchanged on either side of the reaction (e.g., $C_3 \leftrightarrow C_3$). These reactions were ignored as well, without consequences on the results (data not shown).

Metabolite and Reaction Usage

We counted how often each metabolite and reaction was used in the artificial chemistry pathways as well as in the KEGG-derived single-atom networks. In the model pathways, reaction usage was calculated by counting how many times each reaction was used across all pathways. Metabolite usage was similarly calculated by counting the occurrence of reactions in which each metabolite participates. For example, in the pathways that convert a_9 to a_{10} in Fig. 2C, a_9 participates in only one reaction, but a_{10} participates in two.

In the KEGG-derived networks, a similar counting scheme was used. The reaction usage was calculated by counting how many times each reduced reaction appears, and the metabolite usage was calculated by counting how many times each metabolite appears across all reactions.

Shortest Paths in Escherichia coli

The first method used EFMs to find all shortest pathways in the central carbon metabolism of *Escherichia coli* [36,37]. Because we are interested in the pathways that alter the carbon content of different molecular species, we removed those common cofactors that do not alter the carbon content of other metabolites (ATP, ADP, AMP, NADH, NADPH, NAD⁺, NADP⁺, H⁺, and PO₄). Also, we effectively ignored reactions involving transport and exchange. We then used the EFM method described above to find the number of reactions in each of the MBPs for this reduced network.

For each input compound, we listed the lengths of the MBPs for all output compounds containing a number j of carbons. For each j , we select among these paths the shortest one, giving an estimate of the shortest path between any individual compound and the closest j -carbon compound. Finally, for each value of i , we averaged these path lengths over all input molecules with i carbons. The end result is a matrix that provides the average of the shortest paths from any i -carbon compound to its nearest j -carbon compound.

All-Pairs Shortest Paths with Johnson’s Algorithm

The larger, genome-scale metabolic network of *E. coli* [38] has 761 metabolites and 1075 reactions, and is currently infeasible to study with the EFM method described above, limiting us to a graph theoretical approach. Because there are both metabolites and reactions, metabolic networks are inherently bipartite: metabolites connect to each other only through reactions. Pathway lengths were computed by transforming the metabolic reaction network stoichiometric matrix into an adjacency matrix where metabolites and reactions are all represented by the same type of node.

As described above, we are only interested in the connections between carbon compounds, so we removed any non-carbonaceous metabolites (water, phosphate, ammonia, etc.). Also, we removed the following cofactors that are used in many reactions, but do not participate in the transformation of carbons: ATP, ADP, AMP, NAD+, NADH, NADP+, NADPH, coenzyme A, acetyl-CoA, and the acyl carrier protein.

Next, we used Johnson’s all-pairs shortest paths algorithm (available as a Matlab function) to find shortest pathways between any two carbon compounds in *E. coli*’s metabolic network. This set of pathways was collated as in the previous section, to produce a matrix of the average of the shortest paths from any *i*-carbon compound to its nearest *j*-carbon compound.

Analytical Estimate of MBP Lengths in Analogy with Addition-Subtraction Chains

We developed an analytical approximation for the expected numbers of reactions to be found in any MBP $a_i \Rightarrow a_j$. We begin with a simplified version of the artificial chemistry model in which only irreversible addition reactions of the form



are allowed. Under these restrictions, we first ask what is the smallest number of reactions necessary to produce any a_j from a_1 . We shall denote by $l(j)$ the smallest possible number of such reactions (we count the use of each reaction (5) once). This problem is equivalent to the problem of addition chains [29], in which one attempts to compute a positive integer by generating a sequence of integers such that each term in the sequence is the sum of two previous terms. Addition chains have been studied extensively, mainly because of their applications in computer science and cryptography [29]. For addition chains, $l(j)$ grows logarithmically with j :

$$l(j) \propto \log_2(j) \tag{6}$$

Our artificial chemistry represents a generalization, in which a metabolite of any length i can be used to produce an output metabolite of any length j . If we still assume that only addition reactions are possible (i.e. molecules cannot be broken down), a chain from a_i to a_j will exist only when i is a divisor of j . The problem can then be reduced to the case with a_1 input and $a_{j/i}$ output. Therefore, in the irreversible case, we can assume that inputs consist of monomers without loss of generality. Let $L(j)$ be the length of the shortest reaction chain in this case. Because not every reactant exists when dividing by the input length i , we have the obvious inequality

$$l(j) \leq L(j) \tag{7}$$

Sometimes the shortest chain can be found easily. For instance, $\{2^0, 2^1, \dots, 2^k\}$ is obviously the shortest chain from 1 to $j=2^k$ whose length is $k+1$. This suggests the general lower bound on the shortest length $L(j)$ of the addition chain:

$$L(j) \geq \lceil \log_2(j) \rceil + 1 \tag{8}$$

where $\lceil x \rceil$ represents the ceiling of x , or the smallest integer not less than x . Likewise, as seen below, $\lfloor x \rfloor$ represents the floor of x , or the largest integer not greater than x (for example, $\lceil 3.14 \rceil = 4$, and $\lfloor 3.14 \rfloor = 3$). The longest minimal addition chain arises when the output length is $j = 2^m - 1$. From this fact, we have the upper bound [29]

$$L(j) \leq \lfloor \log_2(j) \rfloor + v(j) \tag{9}$$

where $v(j)$ is the number of 1s in the binary representation of j . Since $v(j) \leq \lfloor \log_2(j) \rfloor + 1$, the bound in equation (9) implies a simpler (but weaker) upper bound

$$L(j) \leq 2\lfloor \log_2(j) \rfloor + 1 \tag{10}$$

The above bounds give precise values in some cases and act as bounds in others. For instance, $L(16) = 5$, and $L(17) = L(18) = 6$ for both the lower and upper bounds, while $L(31) = 8$ is between the lower bound and the upper bound (5 and 9, respectively).

There are various conjectures regarding $L(j)$; one of the most famous [35] asserts that computing $L(j)$ is NP-hard. Nonetheless, the computation of $L(j)$ has been pushed up to $n \leq 2^{25}$. Two other conjectures [33] predict the general lower bound

$$L(j) \geq \lfloor \log_2(j) \rfloor + \log_2 v(j) + 1 \tag{11}$$

and the upper bound

$$L(2^k - 1) \leq k + L(k) - 1 \tag{12}$$

While algorithms for generating the shortest addition chains are discussed by Thurber [33], these all hold for the specific case of pure addition where the input is always a_1 .

We are interested in the general case involving both addition and subtraction, and specifically the lengths $l(i, j)$ of the shortest reaction chains (MBPs) with a_i input and a_j output. Addition-subtraction chains have also been studied previously as an expansion of addition chains, although these correspond to MBPs with only a_1 as an input. Sometimes, in these cases, $l(j)$ is readily computable, e.g.

$$l(2^k - 1) = k + 2 \quad \text{for } k \geq 3 \tag{13}$$

while $L(2^k - 1)$ remains unknown for sufficiently large k . Both lengths can also be equal, i.e. $l(j) = L(j)$. For example,

$$l(2^k) = L(2^k) + k = 1 \tag{14}$$

$$l(2^k + 1) = L(2^k + 1) = k + 2 \tag{15}$$

Note also an inequality:

$$l(j) \geq \lceil \log_2(j) \rceil + 1 \tag{16}$$

All of these features explain the growth law in equation (6).

The quantity $l(i, j)$ has a rich behavior, e.g., there is only a trivial lower bound since $l(j, j) = 1$. To ignore this non-interesting effect, let us divide i and j by their greatest common divisor as it never affects the length of the MBP:

$$l(i, j) = l\left(\frac{i}{d}, \frac{j}{d}\right), \quad d = \text{gcd}(i, j) \quad (17)$$

then we can use an obvious inequality

$$l(i, j) \leq l(i, 1) + l(1, j) = l(i) + l(j) \quad (18)$$

Recalling (6) we finally arrive at an approximation for the number of reactions in an MBP that uses a_i to produce a_j :

$$l(i, j) \sim \log_2\left(\frac{i}{d}\right) + \log_2\left(\frac{j}{d}\right) \quad (19)$$

The approximation in (19) can also be used to estimate the rank distribution of reaction usage. Consider all possible MBPs producing a_j from a_i . For each (i, j) pair, take an MBP and mark all reactions. Let the reaction in (5) occur E_{pq} times: that is, there are E_{pq} MBPs that use (5). We now divide E_{pq} by the total number of MBPs and call $e_{pq} = N^{-2}E_{pq}$ the reaction frequency. It is better to order reactions not according to (p, q) but to their ranking j , so that the reaction of rank $j = 1$ is the most frequent, that of rank $j = 2$ is the second in frequency, etc. This gives e_j . How does e_j decrease with rank? To infer the answer we note that

$$\sum_{j=1}^{N^2} e_j = \langle l \rangle \quad (20)$$

From (19) it is clear that the average length $\langle l \rangle$ of the shortest reaction chain scales as $\log N$. This is consistent with (20) if and only if we have $r_j \sim j^{-1}$. Thus we predict the power-law decay in (21).

$$r_j \sim j^{-1} \text{ when } j \gg 1 \quad (21)$$

Supporting Information

Figure S1 A version of Fig. 2, extended to the R19 network. As in Fig. 2, yellow boxes denote autocatalytic cycles; the red reaction is the one that is used most often, $a_2 + a_2 \rightleftharpoons a_4$.
Found at: doi:10.1371/journal.pcbi.1000725.s001 (2.81 MB TIF)

Figure S2 Heatmaps describing MBP length and degeneracy
Found at: doi:10.1371/journal.pcbi.1000725.s002 (0.02 MB PDF)

Figure S3 Logarithmic trend of metabolic pathway length
Found at: doi:10.1371/journal.pcbi.1000725.s003 (0.04 MB PDF)

Figure S4 Metabolite and reaction usage frequencies for KEGG hydrogen, nitrogen, and oxygen sets
Found at: doi:10.1371/journal.pcbi.1000725.s004 (0.09 MB PDF)

Figure S5 Fast Fourier transforms showing periodicity of metabolite usage

Found at: doi:10.1371/journal.pcbi.1000725.s005 (0.02 MB PDF)

Figure S6 KEGG carbon reaction and metabolite usage with cofactors removed PDF

Found at: doi:10.1371/journal.pcbi.1000725.s006 (0.04 MB PDF)

Figure S7 Pathway length comparison between iterative algorithm and algorithm minimizing absolute values of fluxes
Found at: doi:10.1371/journal.pcbi.1000725.s007 (0.06 MB PDF)

Figure S8 Reaction usage is similar between R10 EFM and MILP results

Found at: doi:10.1371/journal.pcbi.1000725.s008 (0.02 MB PDF)

Figure S9 Removing reactions from the R19 network affects the system's ability to yield balanced pathways

Found at: doi:10.1371/journal.pcbi.1000725.s009 (0.04 MB PDF)

Table S1 Comparison of the three different MBP search algorithms

Found at: doi:10.1371/journal.pcbi.1000725.s010 (0.04 MB DOC)

Table S2 List of MBPs found in the R19 network

Found at: doi:10.1371/journal.pcbi.1000725.s011 (0.07 MB XLS)

Table S3 List of all degenerate MBPs found in the R10 network

Found at: doi:10.1371/journal.pcbi.1000725.s012 (0.07 MB XLS)

Table S4 Metabolite and reaction usage from MBPs in the R19 network

Found at: doi:10.1371/journal.pcbi.1000725.s013 (0.03 MB XLS)

Table S5 Metabolite and reaction usage from MBPs in the R10 network

Found at: doi:10.1371/journal.pcbi.1000725.s014 (0.02 MB XLS)

Table S6 Metabolite and reaction usage from the KEGG carbon dataset

Found at: doi:10.1371/journal.pcbi.1000725.s015 (0.08 MB XLS)

Table S7 Metabolite and reaction usage from the KEGG hydrogen dataset

Found at: doi:10.1371/journal.pcbi.1000725.s016 (0.17 MB XLS)

Table S8 Metabolite and reaction usage from the KEGG nitrogen dataset

Found at: doi:10.1371/journal.pcbi.1000725.s017 (0.02 MB XLS)

Table S9 Metabolite and reaction usage from the KEGG oxygen dataset

Found at: doi:10.1371/journal.pcbi.1000725.s018 (0.09 MB XLS)

Table S10 Metabolite and reaction usage from the KEGG phosphorous dataset

Found at: doi:10.1371/journal.pcbi.1000725.s019 (0.02 MB XLS)

Table S11 Metabolite and reaction usage from the KEGG sulfur dataset

Found at: doi:10.1371/journal.pcbi.1000725.s020 (0.01 MB XLS)

Acknowledgments

We are grateful to Oliver Ebenhöf, Tzachi Pilpel, Scott Mohr and members of the Segre lab for helpful feedback and discussions.

Author Contributions

Conceived and designed the experiments: WJR PLK SR DS. Performed the experiments: WJR. Analyzed the data: WJR PLK SR DS. Wrote the paper: WJR DS.

References

- Lenski RE, Winkworth CL, Riley MA (2003) Rates of DNA sequence evolution in experimental populations of *Escherichia coli* during 20,000 generations. *J Mol Evol* 56: 498–508.
- Fong SS, Palsson BO (2004) Metabolic gene-deletion strains of *Escherichia coli* evolve to computationally predicted growth phenotypes. *Nat Genet* 36: 1056–1058.
- Lee M, Chou H, Marx CJ (2009) Asymmetric, bimodal trade-offs during adaptation of *Methylobacterium* to distinct growth substrates. *Evolution* 63: 2816–2830.
- Raymond J, Segrè D (2006) The effect of oxygen on biochemical networks and the evolution of complex life. *Science* 311: 1764–1767.
- Galperin MY, Koonin EV (1999) Functional genomics and enzyme evolution. Homologous and analogous enzymes encoded in microbial genomes. *Genetica* 106: 159–170.
- Hernández-Montes G, Díaz-Mejía JJ, Pérez-Rueda E, Segovia L (2008) The hidden universal distribution of amino acid biosynthetic networks: a genomic perspective on their origins and evolution. *Genome Biol* 9: R95.
- Morowitz HJ, Kostelnik JD, Yang J, Cody GD (2000) The origin of intermediary metabolism. *Proc Natl Acad Sci USA* 97: 7704–7708.
- Gánti T (2003) *The Principles of Life*. New York: Oxford University Press, USA.
- Wächtershäuser G (1990) Evolution of the first metabolic cycles. *Proc Natl Acad Sci USA* 87: 200–204.
- Segrè D, Ben-Eli D, Lancet D (2000) Compositional genomes: prebiotic information transfer in mutually catalytic noncovalent assemblies. *Proc Natl Acad Sci USA* 97: 4112–4117.
- Baldwin JE, Krebs H (1981) The evolution of metabolic cycles. *Nature* 291: 381–382.
- Ebenhöh O, Heinrich R (2001) Evolutionary optimization of metabolic pathways. Theoretical reconstruction of the stoichiometry of ATP and NADH producing systems. *Bull. Math. Biol* 63: 21–55.
- Meléndez-Hevia E, Waddell TG, Montero F (1994) Optimization of Metabolism: The Evolution of Metabolic Pathways Toward Simplicity Through the Game of the Pentose Phosphate Cycle. *J Theor Biol* 166: 201–220.
- Ebenhöh O, Heinrich R (2003) Stoichiometric design of metabolic networks: multifunctionality, clusters, optimization, weak and strong robustness. *Bull Math Biol* 65: 323–357.
- Maslov S, Krishna S, Pang TY, Sneppen K (2009) Toolbox model of evolution of prokaryotic metabolic networks and their regulation. *Proc Natl Acad Sci USA* 106: 9743–9748.
- Venter JC, Remington K, Heidelberg JF, Halpern AL, Rusch D, et al. (2004) Environmental genome shotgun sequencing of the Sargasso Sea. *Science* 304: 66–74.
- Kanehisa M, Goto S, Hattori M, Aoki-Kinoshita KF, Itoh M, et al. (2006) From genomics to chemical genomics: new developments in KEGG. *Nucleic Acids Res* 34: D354–357.
- Caspi R, Foerster H, Fulcher CA, Kaipa P, Krummenacker M, et al. (2008) The MetaCyc Database of metabolic pathways and enzymes and the BioCyc collection of Pathway/Genome Databases. *Nucleic Acids Res* 36: D623–631.
- Kauffman SA (1993) *The Origins of Order: Self-Organization and Selection in Evolution*. 1st ed. New York: Oxford University Press, USA.
- Fontana W, Buss LW (1994) What would be conserved if “the tape were played twice”? *Proc Natl Acad Sci USA* 91: 757–761.
- Hintze A, Adami C (2008) Evolution of complex modular biological networks. *PLoS Comput Biol* 4: e23.
- Dittrich P, Ziegler J, Banzhaf W (2001) Artificial Chemistries—A Review. *Artif Life* 7: 225–275.
- Klamt S, Stelling J (2003) Two approaches for metabolic pathway analysis? *Trends Biotechnol* 21: 64–69.
- Edwards JS, Covert M, Palsson B (2002) Metabolic modelling of microbes: the flux-balance approach. *Environ Microbiol* 4: 133–140.
- Srinivasan V, Morowitz HJ (2009) The canonical network of autotrophic intermediary metabolism: minimal metabolome of a reductive chemoautotroph. *Biol Bull* 216: 126–130.
- Meléndez-Hevia E, Waddell TG, Cascante M (1996) The puzzle of the Krebs citric acid cycle: assembling the pieces of chemically feasible reactions, and opportunism in the design of metabolic pathways during evolution. *J Mol Evol* 43: 293–303.
- Papp B, Teusink B, Notebaart RA (2009) A critical view of metabolic network adaptations. *Hfsp J* 3: 24–35.
- Kauffman KJ, Prakash P, Edwards JS (2003) Advances in flux balance analysis. *Curr Opin Biotechnol* 14: 491–496.
- Knuth DE (1997) *Art of Computer Programming, Volume 2: Seminumerical Algorithms* (3rd Edition). 3rd ed. Reading: Addison-Wesley Professional.
- Breslow R (1959) On the mechanism of the formose reaction. *Tetrahedron Lett* 1: 22–26.
- Ravasz E, Somera AL, Mongru DA, Oltvai ZN, Barabási AL (2002) Hierarchical organization of modularity in metabolic networks. *Science* 297: 1551–1555.
- Shen-Orr SS, Milo R, Mangan S, Alon U (2002) Network motifs in the transcriptional regulation network of *Escherichia coli*. *Nat Genet* 31: 64–68.
- Thurber EG (1999) Efficient Generation of Minimal Length Addition Chains. *SIAM J Comput* 28: 1247.
- ShiJianghong, ZhouTing (2007) A Novel Fast Exponentiation Algorithm for Encryption. In: *Anti-counterfeiting, Security, Identification, 2007 IEEE International Workshop on*. pp 245–248.
- Downey P, Leong B, Sethi R (1981) Computing Sequences with Addition Chains. *SIAM J Comput* 10: 638–646.
- Palsson BO (2006) *Systems Biology: Properties of Reconstructed Networks*. 1st ed. New York: Cambridge University Press.
- Varma A, Palsson BO (1993) Metabolic Capabilities of *Escherichia coli*: I. Synthesis of Biosynthetic Precursors and Cofactors. *J Theor Biol* 165: 477–502.
- Reed JL, Vo TD, Schilling CH, Palsson BO (2003) An expanded genome-scale model of *Escherichia coli* K-12 (JW904 GSM/GPR). *Genome Biol* 4: R54.
- Handorf T, Ebenhöh O, Heinrich R (2005) Expanding metabolic networks: scopes of compounds, robustness, and evolution. *J Mol Evol* 61: 498–512.
- Ourisson G, Nakatani Y (1994) The terpenoid theory of the origin of cellular life: the evolution of terpenoids to cholesterol. *Chem Biol* 1: 11–23.
- Ricardo A, Carrigan MA, Olcott AN, Benner SA (2004) Borate minerals stabilize ribose. *Science* 303: 196.
- Desiraju GR, Jack DDunitz, AshwiniNangia, Jagarlapudi ARPSarma, EngelbertZass (2000) The Even/Odd Disparity in Organic Compounds. *Helv Chim Acta* 83: 1–15.
- Sarma JARP, Nangia A, Desiraju GR, Zass E, Dunitz JD (1996) Even-odd carbon atom disparity. *Nature* 384: 320–320.
- Morowitz HJ (1979) *Energy Flow in Biology*. Ox Bow Press.
- de Figueiredo LF, Podhorski A, Rubio A, Kaleta C, Beasley JE, et al. (2009) Computing the shortest elementary flux modes in genome-scale metabolic networks. *Bioinformatics* 25: 3158–3165.
- Burgard AP, Vaidyaraman S, Maranas CD (2001) Minimal Reaction Sets for *Escherichia coli* Metabolism under Different Growth Requirements and Uptake Environments. *Biotechnol Prog* 17: 791–797.
- Reznik E, Segrè D (2010) On the Stability of Metabolic Cycles. arXiv:1001.1944v1 [q-bio.MN].
- Akashi H, Gojbori T (2002) Metabolic efficiency and amino acid composition in the proteomes of *Escherichia coli* and *Bacillus subtilis*. *Proc Natl Acad Sci USA* 99: 3695–3700.
- Nelson DL, Cox MM (2004) *Lehninger Principles of Biochemistry*, Fourth Edition. Fourth Edition. New York: W. H. Freeman.
- Meléndez-Hevia E, Torres NV (1988) Economy of design in metabolic pathways: Further remarks on the game of the pentose phosphate cycle. *J Theor Biol* 132: 97–111.
- Heinrich R, Schuster S (1998) The modelling of metabolic systems. Structure, control and optimality. *Biosystems* 47: 61–77.
- Murtas G (2009) Artificial assembly of a minimal cell. *Mol Biosyst* 5: 1292–1297.
- Hallenbeck PC, Ghosh D (2009) Advances in fermentative biohydrogen production: the way forward? *Trends Biotechnol* 27: 287–297.
- Cockell CS, Léger A, Fridlund M, Herbst TM, Kaltenecker L, et al. (2009) Darwin-A Mission to Detect and Search for Life on Extrasolar Planets. *Astrobiology*.
- Conrad PG, Nealson KH (2001) A non-earthcentric approach to life detection. *Astrobiology* 1: 15–24.
- Nealson KH, Tsapin A, Storrie-Lombardi M (2002) Searching for life in the Universe: unconventional methods for an unconventional problem. *Int Microbiol* 5: 223–230.
- Edwards JS, Palsson BO (2000) Metabolic flux balance analysis and the in silico analysis of *Escherichia coli* K-12 gene deletions. *BMC Bioinformatics* 1: 1.
- Varma A, Palsson BO (1994) Stoichiometric flux balance models quantitatively predict growth and metabolic by-product secretion in wild-type *Escherichia coli* W3110. *Appl Environ Microbiol* 60: 3724–3731.
- Schuetz R, Kuepfer L, Sauer U (2007) Systematic evaluation of objective functions for predicting intracellular fluxes in *Escherichia coli*. *Mol Syst Biol* 3: 119.
- Mo ML, Palsson BO, Herrgård MJ (2009) Connecting extracellular metabolic measurements to intracellular flux states in yeast. *BMC Syst Biol* 3: 37.
- Kamp AV, Schuster S (2006) Metatool 5.0: fast and flexible elementary modes analysis. *Bioinformatics* 22: 1930–1931.