The biochemical factory that autonomously fabricates itself: A systems-biological view of the living cell

Jan-Hendrik S. Hofmeyr
Dept. of Biochemistry, University of Stellenbosch, South Africa

March 23, 2007

Final, corrected draft of a chapter for the book *Systems Biology: Philosophical Foundations*, (Boogerd, F.C., Bruggeman, F., Hofmeyr, J.-H.S. and Westerhoff, H.V., eds), Chapter 10, pp. 217–242, Elsevier, 2007.

The essence of life must lie somewhere between molecule and autonomously living, unicellular organism. Modern biology generally views organisms as beads along the necklace of lineage; it attempts to explain life from an evolutionary viewpoint, with reproduction (of cells) and replication (of DNA) as defining phenomena. Systems biology, however, studies each bead *per se* as an autonomous entity. I suggest that, for systems biology, the defining difference between a living organism and any non-living object should be that an organism is a system of material components that are organised in such a way that the system can autonomously and continuously fabricate itself, i.e., it can live longer than the lifetimes of all its individual components. Systems biology therefore goes beyond the properties of individual biomolecules, taking seriously their organisation into a living whole.

The concept of autonomous self-fabrication of systems is of course not new; it has a distinguished history. Although Maturana and Varela's concept of autopoietic systems is perhaps most prominent in this history, I find that for the purpose of formalisation it less useful than either Rosen's theory of replicative metabolism-repair systems or Von Neumann's theory of self-reproducing automata based on the concept of

a universal constructor. Rosen in particular has shown, using category theory, how to describe such organisations in terms of relational models, although he never realised his metabolism-repair systems in terms of biochemistry as we know it. I shall show how it is possible to combine these two strands of thought into a relational model that commutes with our current knowledge of cellular biochemical processes. This model, which I call a metabolism-construction-assembly system, also makes explicit the role of information, and identifies unassisted self-assembly as the process that ultimately makes the system self-fabricating (or, using Rosen's words: 'closes the system with respect to efficient causation'). What makes this model even more interesting is that it is consistent with Barbieri's ribotype theory, and, through that, with the body of thought known as biosemiotics.

1 How to be a systems biologist

The aim of this book is to explore the possibility that systems biology may need philosophical foundations of its own. I believe it does, and that systems biology should aim to provide a way of thinking about living organisms that will allow us to understand them as autonomous entities. At present the dominant views of biology are through the glasses of evolutionary biology and of molecular biology. Evolutionary biology seeks to understand life in terms of how natural selection has molded organisms through the millennia. The philosophy of molecular biology is based on the idea that exhaustive knowledge of all the individual molecular components of the cell will afford the best understanding of life. I shall make a case for a philosophy of systems biology that is based on the premise that the living state exists because of a particular organisation of the internal components of cells.

What does systems biology actually entail in practice (Westerhoff and Hofmeyr, 2005)? If one listens to talks at conferences on systems biology (Cornish-Bowden, 2005) or reads editorial introductions to special journal issues devoted to systems biology (Russel and Superti-Furga, 2005), it becomes clear that, although many individual scientists who regard themselves as systems biologists have very clear views on what they regard as the gist of their discipline, there is no clear consensus. It also does not help that the phrase 'systems biology' in a grant proposal has become, or is at least perceived as, a means for ensuring funds. Be it as it may, the different views are all compatible and can be consolidated into something along the lines of 'explaining or understanding the emergence of systemic functional properties of the living cell as a result of the interactions of its components'. How this is to be achieved is usually seen to be by means of two approaches, either on their own or in combination (Westerhoff and Kell, 2007) (Chapter 2 in this volume). Both approaches espouse what I would call the 'system-wide' view: the conviction that one cannot understand the cell if one does not consider it as a whole. One approach comes from the age of 'omics' and proposes that, now that

the new high-throughput techniques have made it possible, we should measure the amount of everything that there is to be measured inside a cell under different conditions (DNA, RNA, proteins, metabolites) and then 'data-mining' will do the rest. The other approach has as its aim the 'silicon cell' (Snoep *et al.*, 2005), a computer simulation of the complete cellular network of reactions and interactions based on the experimentally measured properties of the 'agents' of the cell (enzymes, pumps, receptors, etc.). What is not always so clear is exactly how the results of these two approaches—exhaustive cell-wide data-sets and complete cell models—necessarily lead to deeper understanding. In the rest of this introduction I would like to propose that what needs to be added first is a clear view of what the 'systemic approach' entails.

The geneticist Theodosius Dobzhansky summed up the dominant explanatory modality for biology of the last century in the mantra "Nothing in biology makes sense except in the light of evolution" (Dobzhansky, 1973). Ernst Mayr, however, pointed out that there are two types of explanation in biology, ultimate and proximate, which would respectively follow from evolutionary and functional considerations (Mayr, 1988). I suggest that systems biology seeks to expose general principles that underlie proximate explanations of what governs life. My mantra for systems biology would therefore be *"Nothing in an organism makes sense except in the light of context"*. There are three words here that need elaboration. First, 'sense' emphasises that what is sought is explanation and understanding, not just description. As an example: an indispensable part of the system-wide study of the cell is to make a complete map of all reactions and interactions that comprise the intracellular network. However, in the words of Count Alfred Korzybski (1994) "a map is not the territory"; making the map does not in itself afford understanding; neither does measuring the concentrations of all the nodes on the cellular map. Second, 'organism' emphasises that systems biology studies a particular cell or organism as a material system that is to be explained in terms of itself and its interactions with its environment; in contrast, an evolutionary explanation would be in terms of its history. Third, and most important, 'context' captures the essence of the systems approach. Always taking context into account amounts to using a 'macroscope' (de Rosnay, 1979), a tool for studying the very complex (in contrast to using a microscope for the very small and a telescope for the very large).[1] The macroscope is a 'symbolic instrument' that collects a number of techniques and methods into what De Rosnay calls the 'systemic approach', which, in contrast to the analytical approach, takes into account not only all the elements in the system under study but also all their interactions. A 'system' itself is, in de Rosnay's words 'a set of interacting elements that form an integrated whole'. The living cell is prototypic of such systems.

However, most discussions of the systemic approach remain on a metalevel, and are, for a practising biologist such as myself, ultimately unsatisfactory. I would

[1] A web-edition of this book is freely available from Principia Cybernetica (http://pespmc1.vub.ac.be/LIBRARY.html). It is recommended reading for all systems biologists.

rather prefer to give an example, first to explain what I think it means to be a systems biologist, and second to demonstrate how taking context into account by using the macroscope has helped me understand deeply and in a radically new way some functional properties of the cell that were considered to have been explained long ago. What I also intend showing is that it is not always necessary to take the whole system into account to understand something; doing systems biology does not necessarily entail doing 'system-wide' biology (and *vice versa*).

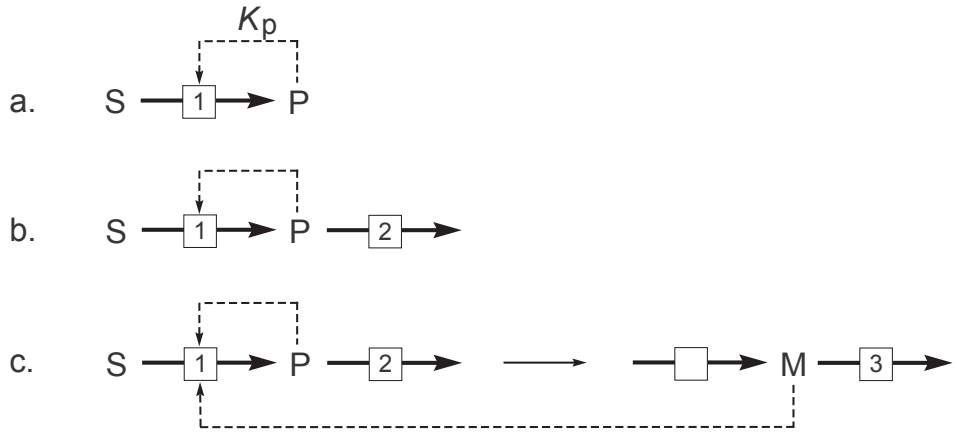

Figure 1: The functional context of an enzyme that converts a substrate S to a product P.

Consider a metabolic enzyme, the systems biologist's favourite molecule (the enzyme marked 1 in Fig. 1a). Consider further a particular kinetic property of this enzyme, say K_p, the Michaelis constant for the product of the enzyme-catalysed reaction, which is an indication of the strength with which the product binds to the active site of the enzyme (the smaller its value, the stronger the binding). This enzyme parameter has a specific value that can be measured experimentally. How can we explain why it has one particular value and not another? There are a number of answers to this question. First, a classical 'enzymological' answer would explain the value as a function of experimentally determined rate constants for the elementary steps in the enzyme mechanism. A second, more modern, 'structural' answer would be based on the three-dimensional conformation of the active site and the complementary structure of the product molecule, which could serve as a point of departure for *ab initio* calculations of the value of the rate constants for binding and dissociation. This type of computational enzymology is fast becoming a reality; it uses the physical and chemical principles of statistical mechanics and quantum mechanics, and they are implemented in computational form using techniques from computational chemistry (Cunningham and Bash, 1997). These two approaches are both reductionistic in that they reduce a property to more elementary properties, here either kinetic or structural; in essence they use the microscope, zooming into the system which here is the isolated enzyme-product complex. The

third answer from molecular biology shows no particularly interested in explaining the value K_p, but would rather explain the enzyme as a gene product. A fourth 'historical' answer would explain the K_p as an evolved enzyme property that can be followed down the phylogenetic tree and across generations; this is essentially a telescopic approach—but instead of looking far out, it looks far back.

Although all of the above approaches are perfectly valid and necessary, I would argue that none of them truly explain why the enzyme has that particular value for its K_p. The only way to find that out is to study the enzyme in the context of the metabolic reaction network in the cell. Let us therefore use the macroscope to zoom out, taking into account the immediate functional context of the enzyme. We find that the product of this enzyme is a substrate for another enzyme (the one marked 2 in Fig. 1b) and therefore couples the two enzymatically catalysed reactions. Note that context here does not mean spatial environment, but rather network environment. Now the question of why K_p has a particular value takes on a whole new meaning. Metabolic control analysis (Kacser and Burns, 1973; Heinrich and Rapoport, 1974) teaches us that the role that an enzyme plays in controlling steady-state flux and concentrations is determined by the elasticity coefficients of the enzymes that catalyse the reaction network (an elasticity coefficient describes the sensitivity of reaction rate with respect to a chemical species that directly affects the enzyme activity, such as a substrate, product or modifier; mathematically it is the fractional change in reaction rate divided by the fractional change in concentration of the chemical species in question). If, for example, our enzyme reaction is far from equilibrium in the forward direction (so that the mass-action ratio is much smaller than the equilibrium constant) and the K_p has a large value relative to the steady-state concentration of P, so ensuring that the ratio of product concentration p to K_p is very small, the elasticity coefficient with respect to P approaches zero: under these conditions the enzyme is kinetically and thermodynamically insensitive to anything that happens downstream from that enzyme (assuming that the coupling product is the only way through which the downstream reactions can communicate with enzyme 1). Under these conditions the enzyme has complete control over the steady-state flux.[2] The smaller K_p, the more sensitive the enzyme becomes to changes in its product concentrations, and the less its controls the flux.

Now imagine that lower down in the pathway there is a metabolite M that saturates the enzyme for which it is a substrate (enzyme 3 in Fig. 1c) so that the enzyme is insensitive to changes in the concentration of M. In addition, M also feeds back allosterically onto enzyme 1 higher up in the pathway. Hofmeyr and Cornish-Bowden (2000) have shown that in this situation enzyme 1 has no control over the overall flux through the full system, but completely determines the degree of homeostasis of M. However, this can only happen if enzyme 1 retains flux control over that part of the system that leads up to M (the supply pathway for M). As before this is partly

[2]This explanation is simplified for the purposes of this discussion; it ignores the elasticity coefficients of the other steps in the system which all play a role in determining the control profile.

determined by the ratio p/K_p, which must be small. When K_p becomes smaller and p/K_p increases concomitantly, the system becomes structurally unstable and exhibits multistationarity (Hofmeyr *et al.*, 2000). This behaviour only obtains when enzyme 1 binds S and P cooperatively, which is the norm for allosteric enzymes, and it is independent of the specific mechanism (Hofmeyr and Cornish-Bowden, 1997). It is therefore conceivable that evolution has selected for large values of K_p in order to avoid this type of pathological behaviour. However, most kinetic studies of allosteric enzymes have ignored the K_p and it may turn out that in some cases K_p is small, which could cause the system to exhibit switching behaviour instead of a smooth response to changes in the concentration of M. Whatever the case may be, it is clear that one can only understand why K_p has the value that it has by analysing it in its functional context, which need not be that of the whole cell.

Having now made and illustrated the claim that nothing in an organism makes sense except in the light of context, it is time to consider the nature of the overall context that the living cell provides for systems biology. This amounts to asking of systems biology how it defines life.

2 The self-fabricating cell: a context for systems biology

> Ironically, biology itself provides a ground upon which epistemology and ontology directly meet. Put simply, organisms are themselves fabricators; they build new things, they make new things, they deploy new things. Hence, an essential part of a theory of organism is precisely a theory of fabrication; a theory of invention and deployment. Thus, a theory of organisms has within itself an ineluctable ontological component; a science of fabrication. Nothing shows more clearly than this the unique character of biology among the sciences, and the unique role that its own theory must play in its own application.
>
> Robert Rosen, *On Theory in Biology* [3]

Biologists, more than ever before, are living in a golden age. The cell, that unit on which all life is based, no longer seems a mystery; in fact, we apparently feel we know and understand it so well and have such advanced technology that we can manipulate life at the molecular level confidently and responsibly. However, keeping in mind E.F. Schumacher's admonition that "the greatest danger invariably arises from the ruthless application, on a vast scale, of partial knowledge" (Schumacher, 1973), should we not be asking ourselves seriously whether we can we really explain life? How successfully can we at present answer the questions 'why *Escherichia coli*?', 'why *Homo sapiens*?', 'why any organism?

[3] http://www.rosen-enterprises.com/RobertRosen/BioTheoryHistoryofBiology.html

Organisms as we know them are material systems, and according to Aristotle (1998, 350 B.C.E) there are four different ways of answering 'why'-questions about material objects, questions that he placed at the heart of science. Put differently, there are four fundamentally different explanatory factors that together explain any object fully.[4] These four *aitia*, as he called them, are now commonly described as material, efficient, formal, and final causes. However, to avoid confounding Aristotelian explanations with 'causation' in the sense of Hume, Cohen's suggestion[5] to replace the noun 'cause' with the verb 'make' is useful:

1. What is an organism *made out of*? (its material cause)
2. What *makes* (in the sense of '*what is it to be*') an organism? (its formal cause)
3. What *makes* (in the sense of '*what produces*') an organism? (its efficient cause)
4. What is an organism *made for*? What is its *purpose* or *function*? (its final cause)

Biochemistry, molecular biology and molecular genetics have been spectacularly successful in providing us with answers to the first two questions: (i) a century's worth of research tells what organisms are composed of and what the structure of their molecular constituents are, and (ii) after Watson and Crick biologists generally ascribe, rightly or wrongly, the essence of an organism to its DNA. These two answers explain life *statically* in terms of *matter* and *form*, and seem, for many, to suffice.

However, Aristotle insisted that all four explanations are needed for full understanding. The other two questions are questions of process and transformation; they explain why *change* occurs and lead to *dynamic* explanations. Currently, biology's answer to the third question of what produces an organism is essentially: 'its parent(s)'. Rudolf Virchow famously summed up this view as *cellula e cellula* (every cell from a pre-existing cell), a phrase actually coined by François-Vincent Raspail, another founder of cell theory[6]. This is of course also the point of departure for the evolutionary view of organisms as beads along the necklace of lineage. The rest of this chapter will argue that there is another, and for systems biology more productive, answer to this question.

The fourth explanation of purpose is generally considered, especially by those of a mechanistic bent, to be outside the realms of science. Contemporary biology has

[4] In contrast with Humean doctrine in which effects and their causes are *events*, Aristotle typically considered the causes of *substances* or *objects*; this approach is particularly applicable to artifacts, whether artificial or natural. Living organisms are the ultimate natural artifacts (Barbieri, 2005).

[5] Lecture on the four causes (http://faculty.washington.edu/smcohen/320/4causes.htm). To quote Cohen: "Aristotle's point may be put this way: if we ask 'what makes something so-and-so?' we can give four very different sorts of answer - each appropriate to a different sense of 'makes.'"

[6] http://en.wikipedia.org/wiki/Rudolf_Virchow

nothing to offer on the question of the final cause of an organism. However, I am of the opinion that pondering precisely this question will lay a path to a philosophy of systems biology. In fact, for organisms it turns out that the answers to the last two questions are one and the same.

Stafford Beer, cyberneticist and systems thinker, said that *"the purpose of a system is what it does"* (an idea now entrenched in the acronym POSIWID)[7]. If one asks "what does an organism do?" the usual reply of biologists since Darwin has been "an organism evolves through natural selection"; related replies are "an organism reproduces", or, post-Dawkins (1989), "an organism replicates its genes". The view of life that leads to these answers is perhaps most clearly enunciated in Dobzhansky's mantra mentioned in the introductory section. However, since the 1960s another answer to the question "what does an organism do?" has been given with increasing frequency: "An organism produces itself", by which is meant that organisms constantly and autonomously rebuild or *fabricate* themselves during their own lifetimes. In the words of Humberto Maturana and Francisco Varela (Maturana and Varela, 1980) organisms are *autopoietic*.[8] It is probably fair to say that, together, the 'evolutionary' and the 'autopoietic' answer, either on their own or together, form the basis for most current definitions of life (Ruiz-Mirazo and Moreno, 2004). Note also the convergence of causes, alluded to in the previous paragraph, in this concept of self-fabrication: an organism is its own efficient cause in that it autonomously fabricates itself; but then, the purpose of an organism is to fabricate itself—it is its own final cause.

Although the term autopoiesis is associated with Maturana and Varela, the concept of self-fabrication has a long and venerable history, and seems to have been first formulated explicitly by Immanuel Kant, who conceived of organisms as dynamic, functional wholes in which all components are made by and for each other, in contrast with a machine in which components exist only for each other but cannot make each other:

> In such a product of nature *every part not only exists by means of the other parts*, but is thought as existing for the sake of the others and the whole, that is as an (organic) instrument. Thus, however, it might be an artificial instrument, and so might be represented only as a purpose that is possible in general; but also *its parts are all organs reciprocally producing each other*. This can never be the case with artificial instruments, but only with nature which supplies all the material for instruments

[7] Beer said this many times, but never more forcefully than in his address to the University of Valladolid, Spain in October 2001, a month after September 11th: "According to the cybernetician, the purpose of a system is what it does. This is a basic dictum. It stands for a bald fact, which makes a better starting point in seeking understanding than the familiar attributions of good intentions, prejudices about expectations, moral judgments, or sheer ignorance of circumstances." (Beer, 2002).

[8] The term 'fabricate' will be used throughout instead of 'build' or 'make' or 'produce', the last being all to often confused with 'reproduce'. 'Autopoiesis' lacks a verb-form, whereas 'fabrication' has one: a system that fabricates itself is self-fabricating.

> (even for those of art). Only a product of such a kind can be called a natural purpose, and this because it is an organised and self-organising being [my italics]. (Kant, 1914, §65)

The philosopher of science George Kampis (1991, p. 345) recently put it this way:

> In a component system [a type of system defined by Kampis which includes living organisms], due to the continual turnover that gives rise to the components and then removes them from the system, no component and no higher structure, organised form of the components, can persist, unless produced and renewed over and over again.

I shall take this view of the cell as the foundation on which systems biology must be built. For this to be possible we must have a formal, abstract language with which to describe the functional organisations that would make autonomous self-fabrication possible. To my knowledge, only two such formalisations have been developed. In the late 1950s, more than a decade before Maturana and Varela invented the term autopoiesis, the theoretical biologist Robert Rosen put forward a formalised treatment in terms of category theory of what he called metabolism-repair or (M,R)-systems, which become self-fabricating when supplemented with a mechanism that he called 'replication' (an unfortunate choice, as it turns out, because it does not agree with modern biology's use of the term) (Rosen, 1958a,b, 1959b, 1972) (for a recent review and exploration of (M,R)-systems see Rosen (1991); Letelier *et al.* (2006)). Later on he would summarise the central property of such systems as being 'closed to efficient causation' (Rosen, 1991). At the same time John von Neumann (1966) was developing his theory of self-replicating automata which centred around the concept of a universal constructor. In the rest of this chapter I shall provide a rather informal version of Rosen's formal language, and then use it to show how the main tenets of these two theories can be merged and mapped onto cell biochemistry. However, we first need to explore the central concept of autonomy, because it goes hand in hand with the idea of self-fabrication.

3 Autonomy of material systems: the need for specific catalysis

A logic of life, at least of earth-bound life as we know it, can be deduced from two basic postulates:

Postulate 1 Living organisms are material objects.

Postulate 2 Living organisms are autonomous.

The first postulate commits us to a view of life that is inextricably linked to chemistry: the science of spontaneous transformation of matter and therefore the

science of creativity and what Stuart Kauffman (2000) calls the 'adjacent possible'. The creative nature of chemistry is captured in the concept of 'component systems' (Kampis, 1991). Whether a chemical transformation will actually occur under specified conditions, and if it does, how fast, is answered from thermodynamic and kinetic considerations. An important generalisation of chemical biology is that covalent chemistry is virtually exclusively enzyme-catalysed, whereas the non-covalent chemistry involved in, for instance, chemical recognition, protein folding, and self-assembly of macromolecular complexes is largely uncatalysed (although we now know that at least folding is often assisted by chaperones). This distinction between molecular (covalent) and supramolecular (non-covalent) chemistry, made by Jean-Marie Lehn (1995), will be seen further on to be crucial in understanding the ability of living cells to fabricate themselves. Supramolecular chemistry refers to the formation of ordered molecular aggregates that are held together by noncovalent binding interactions. Because these forces are weak, the formation of supramolecular assemblies is usually thermodynamically-controlled and therefore a spontaneous process of self-assembly rather than a sequential bond-forming synthesis.

A recent series of papers (Ruiz-Mirazo and Moreno, 2004, 1998; Ruiz-Mirazo *et al.*, 1998, 2004) provide an excellent analysis of the concept of autonomy, not only as a point of departure for a universal definition of life, but also in relation to autopoietic theory (see also Chapter 11 in this volume; Moreno (2007)). They make a strong and convincing argument that the concept of autonomy is multifaceted; living systems exemplify all these facets, whereas the autopoietic perspective only considers an abstract organisation that recursively produces itself; real-world autonomy cannot escape the requirements of chemistry, energetics and kinetics, and the necessity for spatial autonomy by self-bounding.[9]

Living systems are open and can never be fully thermodynamically autonomous; as dissipative structures they depend on an externally-determined Gibbs energy gradient (Nicolis and Prigogine, 1977). However, living systems also create internal non-equilibrium conditions that allow them a degree of thermodynamic autonomy. As an example consider a chemotrophic bacterium that not only grows on glucose by fermenting or oxidising it, but also stores glucose as glycogen. If the external glucose is depleted, i.e., if the external Gibbs energy gradient collapses, this bacterium will still be able to survive due to the internal non-equilibrium condition that it has created. As long as its glycogen store lasts it is thermodynamically autonomous with regard to its carbon source. There is therefore a difference between a dissipative system in which a certain range of external conditions create and maintain the system (so that if outside this range the dissipative system no longer exists), and an autonomous dissipative system that also actively creates and maintains internal non-equilibrium conditions. A Bénard cell would be an example of the first type, a living cell an example of the latter.

[9]However, Ruiz-Mirazo and Moreno emphasise the thermodynamic aspects of autonomy, and virtually ignore the kinetic aspects, which, in my opinion, are just as, if not more, important.

To be kinetically autonomous, the chemical reactions that comprise the system must operate on a faster timescale than the rest of the underlying network of spontaneous mass-action chemical transformations; the greater the separation on the timescale, the smaller the effects of these spontaneous side-reactions and the greater the degree of kinetic autonomy. This can only be achieved by catalysts that are specific with regard to both reactants/products and reaction; kinetic autonomy therefore absolutely requires the existence of catalysts that specifically recognise their substrates and transform them into specific products. If such catalysts are themselves short-lived, the autonomous system must be able to replace them. In short, such a system must itself also be a catalyst factory. However, to fabricate molecular catalysts requires both building blocks and additional machinery, which itself must be made within the system. The building blocks can of course be supplied by the environment, but even if the system has to fabricate them this is not a problem: all it needs is to be able to make the specific catalysts that will accomplish the synthesis. However, the machinery that constructs the catalysts must itself be replaceable by the system, lest it fails; this implies even more additional machinery. It is clearly here that the linear hierarchy of efficient causes followed up to now seems to wander off into an infinite regress that is incompatible with the existence of real autonomous systems. In some way this hierarchy of efficient causation must fold back into itself, must close, must become circular.[10]

The possibility, mentioned above, for internal creation and maintenance of non-equilibrium conditions and their dynamic adaptation in the face of a continually changing internal and external context also depends on catalysts, in this instance on their capacity for being regulated. Hofmeyr and Cornish-Bowden (2000) developed their theory of metabolic supply-demand to describe this aspect of autonomy.

The above logical analysis of the consequences of materiality and autonomy has led inexorably to the need for specific catalysts that are functionally organised in such a way that they form a closed loop of efficient causation. The rest of this chapter explores what this type of functional organisation may look like and how it is realised in living cells as we know them. First, however, we need a way of formally representing a self-fabricating system as an organisation of catalytic components.

4 Fabrication and the logic of life

What is fabrication? Are there basic principles underlying fabrication? Must a fabricator be more 'complicated' than that which it fabricates? Can a fabricator fabricate itself? One would suppose that by now there would have been developed, either by engineers, technologists, or anybody that designs or makes gadgets, a

[10]The kinetic autonomy that is ensured by specific catalysts is essentially what is lacking from Tibor Ganti's chemoton (Gánti, 2003). There is nothing in a chemoton that would prevent its chemical intermediates dissipating into side-reactions. Much closer to the kinetic autonomy of living systems is that of the autocatalytic Belousov-Zhabotinsky reaction system (Field and Burger, 1985) in which the catalytic species are produced within the system itself.

fully-fledged theory of fabrication that answers such basic questions. The theory of self-replicating automata developed by von Neumann (1966, 1951) goes some way towards answering these questions, but other than that I have not been able to find a theory of fabrication.[11]

Seen abstractly, fabrication is a process in which a material object is created either by rearrangement of, or by taking away from, or by adding to an existing object. Usually, one assumes that this process is accomplished by a *fabricator*, which is itself a material object (and, of course, not necessarily alive). However, one has to leave open the possibility that the fabrication process happens spontaneously without assistance from a fabricator. In keeping with my background as (bio)chemist, I take my cue from chemistry (the epitome of a fabrication world). Consider A, B, C and P to be either (i) single molecules in which all the atoms are bonded covalently, or (ii) assemblies of molecules that associate through non-covalent forces (ionic and hydrogen bonds, Van der Waals forces, hydrophobic interactions, etc.). One could consider as an example of the first a single polypeptide and of the second an enzyme consisting of non-covalently associated subunits, each consisting of a single polypeptide. A new molecule can form from existing molecule(s) in a number of ways, shown in Fig. 2. Because it is all too easy to forget that physical laws such as the conservation of mass underlie all fabrication processes I depict them using both symbols and schematic representations. For example, I could have written the process $\mathrm{A} + \mathrm{B} \rightarrow \mathrm{AB}$ as $\mathrm{A} + \mathrm{B} \rightarrow \mathrm{C}$, but that obscures the fact that C must contain exactly the atoms of A and B.

In general, therefore, I consider both the input and output to a fabrication process to be a material object which can be considered a unit with a fixed internal arrangement of components ('atoms'). The fabrication process itself involves either an internal rearrangement within an input object, the combining of objects, the splitting of an object, or the transfer of part of one object to another. Following Rosen (1991), the fabrication process can in the abstract be regarded as a mapping f from a domain (a set A of input objects) to a codomain (a set B of output objects). Such a mapping is usually depicted as

$$f : A \longrightarrow B \qquad \text{or, equivalently,} \qquad A \xrightarrow{f} B \tag{1}$$

Any specific conversion of $a \in A$ to $b \in B$ can be depicted with the 'mapsto' notation

$$a \mapsto f(a) \qquad \text{or, equivalently,} \qquad a \mapsto b \tag{2}$$

This mapping is the fundamental relationship on which Rosen (1991) builds his relational theory of biology.

Whereas the nature of A and B is reasonably clear, that of f is not. Is it just a process or is it itself a physical object? We shall see below that f can be either. Rosen (1991) provided the mapping in Eq. 1 with a natural interpretation in terms

[11] Besides biology, nano-engineering is also a field from which such a theory could arise (Drexler, 1992; Freitas Jr. and Merkle, 2004).

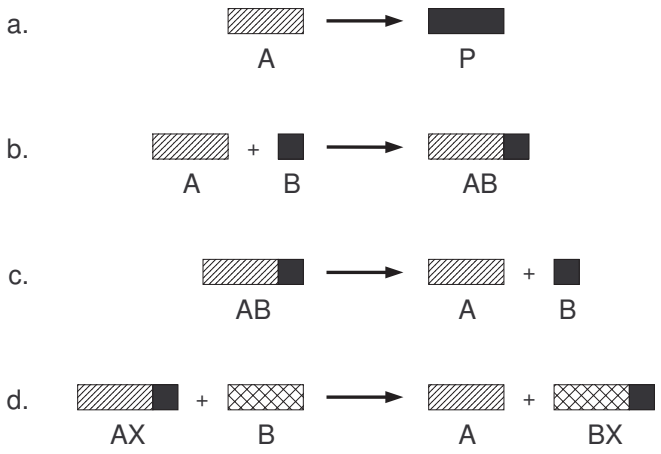

Figure 2: Basic fabrication processes. (a) The atoms of A are *rearranged* into a new configuration P. A and P have the same atomic composition and are either structural or conformational isomers of each other; (b) *Synthesis* of a new molecule AB from A and B, or an *association* of A and B to form a non-covalently bound complex AB; (c) *Degradation* of AB to two fragments A and B, or the *dissociation* of a complex AB into components A and B; (c) The synthesis of two molecules A and BX through the *transfer* of a part of a donor molecule AX to acceptor molecule B.

of Aristotelean causes: the effect B has material cause A and efficient cause f. *In this particular system* A and f have only final cause, namely B; the function of A is to serve as material from which B is made, while the function of f is to fabricate B.

What about formal cause? In the above mapping $f : A \longrightarrow B$ there is nothing that can explicitly be interpreted as formal cause. Here we would have to assume that formal and efficient cause are inseparably part of f (think of sculptor f carving a sculpture according to a vision which exists in her mind only). However, there are clearly situations where formal cause is, at least partly, associated with a separate object (think of an electronic engineer building a circuit board according a design on paper, or a polypeptide being synthesised according to the nucleotide sequence in a particular mRNA).

To account for objects that serve as formal causes of, for instance, macromolecular synthesis, the mapping in Eq. 1 clearly needs an additional entity. Rosen (1989) suggested the more general formulation

$$\begin{aligned} f : A \times I &\longrightarrow B \\ (a, i) &\mapsto b = f(a, i) \end{aligned} \tag{3}$$

where I is a set of templates or blueprints. In this formulation f is the efficient and I the formal cause, although the separation need not be absolute; part of the formal cause can remain associated with f itself.

There are two problems with this formulation. The first is that it leads to a logical paradox when an $i \in I$ is the blueprint for f itself, in the sense that f is an element of its own range (Rosen, 1959a, 1962). No mapping can be defined before its domain and range are stipulated; however, if the range contains the mapping itself as an element, it cannot be stipulated before the mapping is given. Thus, in the words of Rosen (1959a), "neither the mapping f nor its range can be specified until the other is given."

The second problem is that I appears in the mapping with the same status as A, namely as a material cause. However, the role of I is purely informational; logically, any particular $i \in I$ should be associated with f as the pair (f, i). We should rather consider (f, i) as the efficient cause in which the formal part has been made explicit: f is an agent acting on the information contained in i. In these terms the mapping would be:

$$\begin{array}{rcl} (f,i) : A & \longrightarrow & B \\ a & \mapsto & b = (f,i)(a) \end{array} \tag{4}$$

with (f, i) is an element of the Cartesian product $f \times I$.

Readers interested in how these mappings can be formally composed (combined) into fabrication networks are referred to Rosen (1991). In the following section I use a more informal approach to develop an understanding of what it would entail for a collection of material components to become self-fabricating. As mentioned above, the interpretation of objects in the diagrams in terms of Aristotelean causes is also due to Rosen.

5 How to construct a self-fabricating factory

Just as Shakespeare surely found it more profitable to compare his love to a summer's day than to a rock or a whale, we, in order to understand the nature and logical requirements of a self-fabricating system, need to find a useful image to compare it with. From the long line of machine metaphors that have since Descartes been used to describe organisms—through hydraulic automata, clockworks, steam engines, servomechanisms and computers to the vending machine (stick in a gene, pull out a product)—the image of a chemical factory is most useful for the purpose. It embodies the essence of a system that not only consists of fabricators, but, as a whole, is also a fabricator, though, in the case of all factories thus far made by man, not of self.

Let us therefore begin by considering a man-made factory as an generalised example of a fabricator. A bird's-eye view of this factory could be Fig. 3a: the factory L is a black box that transforms raw materials P (its input) into products Q (its output). P and Q can be single objects or collections of objects. Fig. 3a is also a graphical equivalent of the mappings described in Eq. 1. Q can of course be used as input for another 'downstream' factory M that transforms it to R (Fig. 3b). Zooming

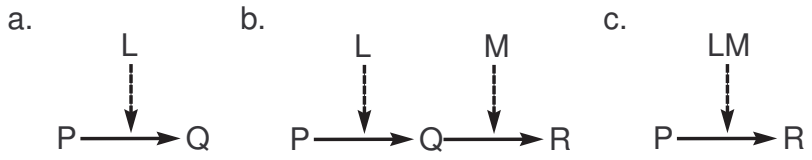

Figure 3: Single (a), linked (b) and lumped (c) fabricators.

our view out even further, L and M may even be viewed as a single factory LM that produces R from P. However, instead of zooming out, let us zoom into the details of our factory. Now, instead of representing the factory as such, L could represent an agent or machine (a simple fabricator) inside the factory; L performs the elementary task of transforming an intermediate widget P somewhere in the production process into the next widget Q. In fact, the production process as a whole can be visualised as a network in which simple fabrication processes are linked as in Fig. 3b. It need not be a linear process: there could be branches where one intermediate widget is used as input for two different processes. Branches can converge; cyclically organised processes can occur. The details of such a transformation network are, however, not important for this discussion. What is important is that, in the picture as painted above, whether it is applied to the factory as a whole or to a machine inside the factory, there is a clear conceptual difference between fabricators on the one hand and their inputs and outputs on the other.

In a computer analogy, P and Q would represent the software which runs on the hardware L and M. However, further on, when our factory becomes more complicated, we shall see that this analogy becomes so ambiguous as to be useless. On the other hand, Aristotelean causal descriptions will prove to be robust. Let us use such a description to explain how the answer to the question 'why L or M?' differs from that to the question 'why P or Q or R?'. For example, the question 'why Q?' is explained by considering Q to be the 'effect' of material cause P and efficient cause L. However, Q can also be considered to have final cause R (the purpose or function of Q is to serve as material cause for R). There is no explicit formal cause for Q in the diagram—it could be considered to be embedded in the properties of L or it could be added to the diagram as information needed by L to fabricate Q. Note that material, efficient and formal cause 'flow forward' to Q, whereas final cause 'flows backwards' to Q. This is always the case. R can be similarly analysed as effect of material cause Q and efficient cause M; unlike Q, R is only effect and plays no functional role within the system. Note that, whereas 'why Q?' and 'why R?' can be answered from within the system (they both have material and efficient causes), 'why P?', 'why L?', and 'why M?' do not have answers within the system. They can only be explained in terms of their final cause: P functions as material cause for Q, while L and M function as efficient causes (for Q and R).

In order to emphasise that from now on we only consider fabricators *inside* our factory, i.e., the components that comprise the factory, I switch to different symbols

(Fig. 4a). In a perfect world where machines do not deteriorate, the factory will run forever as long as enough input material is available. Consider, however, that the 'hardware' of the factory, i.e., its fabricators C, have a limited lifetime; after a while they malfunction and have to be either repaired or replaced in order for the factory to outlive the lifetimes of its machines. Let C in Fig. 4a malfunction. As shown in Fig. 4b, one possibility for overcoming this problem would be to expand the scope of the factory by acquiring a new fabricator D that builds replacement C from material X (or repairs C using spare parts X). It is clear that we have now started

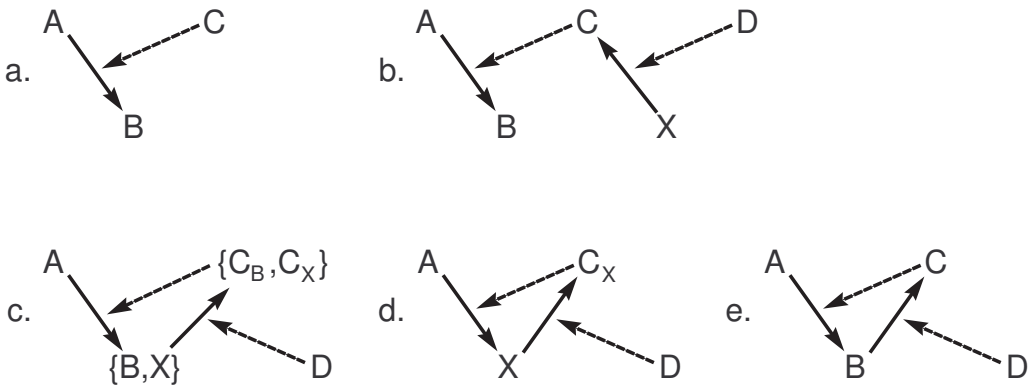

Figure 4: How to build a factory

to create a 'fabrication hierarchy' in that C, which acts as fabricator for the lower transformation level, now also is the target of a fabrication process at a higher level. Furthermore, in this expanded description the factory is less autonomous in that it now depends on not only an external supply of A but also of X. What if the supplier of X becomes unreliable? The most effective measure to counter this would be to incorporate additional machinery C_X that can fabricate X from A into the factory, alongside the original machinery that fabricates B (now distinguished as C_B), thereby increasing the degree of autonomy of the factory at the expense of more machinery (Fig. 4c).

However, there is a problem in Fig. 4c. D is required not only to fabricate C_B from X, but apparently also C_X from something. But from what? We could consider new material Y, but in order to become independent of a possibly unreliable supply of Y that would mean incorporating into the factory even more new machinery C_Y to fabricate Y from A. An infinite regress looms. This regress can be sidestepped by ensuring that both C_B and C_X can be fabricated by D from X (in which case Fig. 4c would be a valid diagram). Furthermore, nothing now prevents us from discarding C_B, so forcing the factory to become a producer of primarily C_X instead of B, X now being an intermediate in the process (Fig. 4d). When, further on, we analyse the living cell as we know it terms of its fabrication hierarchy we shall see that it has adopted both these strategies. Now that we understand the subtleties involved we shall, in the interest of readability, continue to use the diagram in Fig. 4e.

Just as C, fabricator D is also subject to failure and we can imagine safeguarding the factory (so increasing its relative independence even further) by adding another level to the fabrication hierarchy in which D is fabricated by a new fabricator E from either C (Fig. 5a) or B (Fig. 5b). As before, adding this level to the fabrication

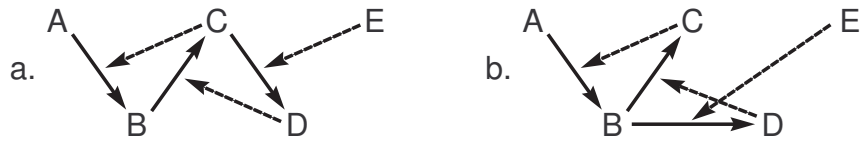

Figure 5: Adding another level to the fabrication hierarchy

hierarchy makes all the lower levels more complicated in that new machines must be fabricated as needed. However, and this is important, the factory still depends on a supply of its input A and the one fabricator, here E, that is not manufactured inside the factory. In fact, as the proportion of endogenously produced fabricators to external fabricator increases the factory becomes more autonomous. We can easily extend the hierarchies in Fig. 5 indefinitely (Fig. 6 shows examples of such extensions). It is left to the reader to imagine hierarchies that are mixtures of these two motifs.

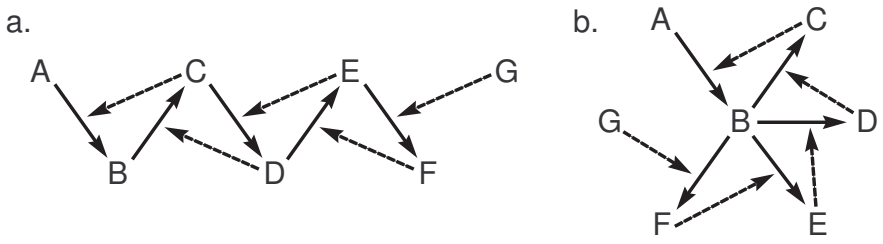

Figure 6: Linear (a) and wheel (b) fabrication hierarchies

Starting with Fig. 4e, and even more so with Figs. 5 and 6, the clean differentiation between hardware and software that obtained in Fig. 3 has become obscured. For example, C, which in Fig. 4 was unambiguously the hardware for software A, has itself become in Fig. 5a the software for E, a higher level of hardware. In addition, hardware C is now seemingly fabricated from its own 'software'. Clearly in Fig. 5a the distinction between hardware and software has been lost and it has become meaningless to discuss the factory in these terms. However, Aristotelean language comes to the rescue, as follows: Consider in turn each object in the complicated situation of Fig. 5a:

- Why A? A functions as material cause for B (A's final cause).
- Why B? B is made from A (its material cause) by C (its efficient cause), and it functions as material cause for C (B's final cause).

- Why C? C is made from B (its material cause) by D (its efficient cause), and it functions as efficient cause for B and as material cause for D (C's two final causes).
- Why D? D is made from C (its material cause) by E (its efficient cause), and it functions as efficient cause for C (D's final cause).
- Why E? E functions as the efficient cause for D (E's final cause).

The beauty of this analysis is that it provides a way of soothing the bugbear of final cause. Consider, for instance, 'why C?' In Fig. 5a there are ostensibly four answers to this question, but upon closer inspection there are only two: 'because B' and 'because D'; in both cases a final cause of C has become identified with either material or efficient cause of C: on the one hand, C functions as fabricator of B, which itself is the material cause for C; on the other hand, C functions as the material cause for its own fabricator, D. The final causes for C have been absorbed into the system.

Thus far, however, there has always still been one fabricator which cannot be replaced or repaired from within the system; one final cause that has not been internalised. If this fabricator fails, it would cause a domino effect down the hierarchy which eventually would bring the whole factory to a standstill. Clearly the problem cannot be solved by adding extra levels of fabrication. Is it possible to internalise this fabrication process so as to make the system completely autonomous (closed) with respect to fabrication, i.e., self-fabricating? (Note that the factory will always remain open to material cause through its dependence on input A from outside.) Consider the arrangement in Fig. 4e. One conceivable way in which the system can rid itself from the necessity for level E and thereby become self-fabricating is if a fabricator lower down in the fabrication hierarchy is able to manufacture D: In Fig. 7a C manufactures D from B, while in Fig. 7b B manufactures D from C. Another

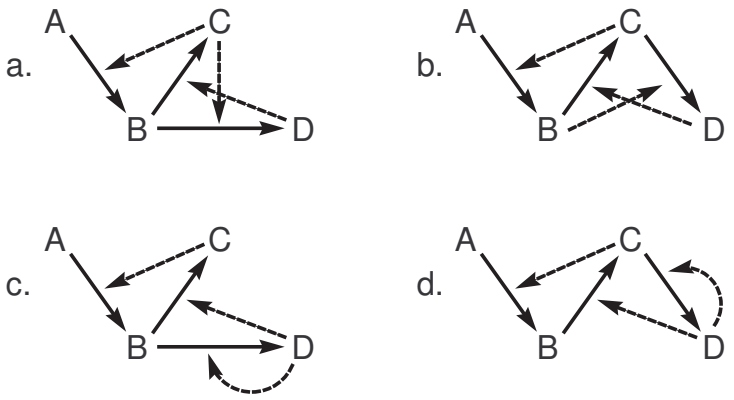

Figure 7: Potential self-fabricating organisations

possibility is that D be able to fabricate itself from either B (Fig. 7c) or C (Fig. 7d).

However, as the following argument shows, there is another incipient regress hidden in our factory: the problem of insufficient numbers of machines. If all the individual steps in all the levels of the fabrication hierarchy are to be performed by a dedicated specific machine, then it is impossible to make the factory self-fabricating using the organisations in Fig. 7. A very simple numerical example demonstrates the problem: Let the set of building blocks B have two members, each fabricated from A in two steps. The set of specific fabricators C therefore needs four members. Let each member of C be made from B in two steps, each step being facilitated by a specific member of the set of D, which therefore needs eight members. If each member of D again needs two specific fabricators then the next level E (Fig. 5) would have contained 16 fabricators. However, in Fig. 7a and 7b these sixteen functions must be performed by members of either B or C. On could conceive of adding extra members to C, but then they would also need to be made from B, which implies that D must be even larger than before, which in turn implies even more C, and so on. The same holds for 7b, but here we also require the members of B to also be fabricators. The organisations in Figs. 7c and 7d lead to the same problem. Clearly this option is a logical impossibility. Another way out would be to require members of B or C to become multi-functional. However, even in this simple case it means that the four members of C in Fig. 7a must share 20 functions between them (four to make B and 16 to make D), or that the two members of B in Fig. 7b share 16 functions between them. Whereas this option is not logically impossible, it confronts us directly with an number of crucial questions: Is there by necessity an increase in the degree of complication of fabricators as one goes up the fabrication hierarchy? Intuitively, one would think so. Therefore, if yes, it is possible for a 'simpler' fabricator to make a more complicated fabricator? In fact, can a fabricator conceivably make itself?

It seems to be generally accepted that von Neumann (1966, 5th Lecture) showed that self-fabrication of a machine (autonomous turnover of self on the basis of a supplied blueprint) and self-reproduction (making a copy of self, including the blueprint) is in principle possible. Von Neumann's so-called kinematic self-reproducing machine consists of a general purpose fabricator $\mathrm{P} + \phi(\mathrm{X})$, which is an automaton consisting of two parts: a constructor P that fabricates a machine X from spare parts according to $\phi(\mathrm{X})$, the blueprint for X. When supplied with its own blueprint $\phi(P)$ the constructor makes itself.[12]

The incorporation of a general purpose fabricator such as $\mathrm{P} + \phi(\mathrm{X})$ into our factory solves the 'insufficient number of machines' dilemma sketched above; in fact, I suspect this to be the only way to circumvent the problem. In Fig. 8 I sketch a self-fabricating factory that shows how incorporating the Von Neumann architec-

[12] To give the entire fabricator $\mathrm{P} + \phi(\mathrm{X})$ the ability to make a copy of itself, von Neumann added a blueprint copier Q and a controller R so that the fabricator becomes $(\mathrm{P} + \mathrm{Q} + \mathrm{R}) + \phi(\mathrm{X})$, which can make not only X but also a copy of $\phi(\mathrm{X})$. When supplied with its own blueprint, $\phi(\mathrm{P} + \mathrm{Q} + \mathrm{R})$, it can make a copy $(\mathrm{P} + \mathrm{Q} + \mathrm{R}) + \phi(\mathrm{P} + \mathrm{Q} + \mathrm{R})$ of itself and of its blueprint, thereby ensuring self-fabrication of the full system. However, for the purpose of analysing self-fabrication we do need to concern ourselves with replication of the blueprint.

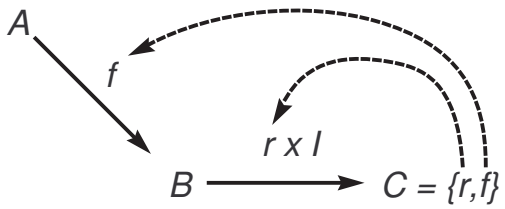

Figure 8: An abstract self-fabricating factory that incorporates the Von Neumann architecture. Raw materials A are converted into building blocks B for the fabricators in the factory. A Von Neumann constructor r uses the information in the set of blueprints I to fabricate the set of machines C. The fabricator (r, i_k) with $i_k \in I$ can make itself directly when supplied with its own blueprint $i_k = i_r$. With $i_k \neq i_r$ it makes all the other machines f required by the factory.

ture makes Fig. 4e self-fabricating. The symbolism used in this figure derives from Eq. 4; the Von Neumann constructor P is symbolised by r, while the blueprint $\phi(X)$ referred to above is now an element i of the set of blueprints I; any element (r, i) of the Cartesian product $r \times I$ is a fabricator of a particular X. Following Rosen (1991) we call f a 'metabolic' mapping; r will be called the 'construction' mapping; this is therefore a metabolism-construction system.

It should be clear that without the ability of constructor r to make itself with the help of its own blueprint, the factory cannot become self-fabricating. Although it seems to be generally accepted that constructors can in principle do this, what if it turns out not to be so? It is now time to turn to life-as-we-know-it and ask how organisms manage to fabricate themselves. Does what we know about metabolism and protein synthesis match the organisation sketched in Fig. 8 or do organisms do it differently?

6 Self-fabrication in living systems

For the purpose of matching known biochemistry to our abstract representation of a self-fabricating factory, consider the diagram in Fig. 9. It has been pointed out many times in the literature that ribosomes are really the only known examples of Von Neumann constructors. They fit the description perfectly: on its own a ribosome can do nothing, but in conjunction with the information embedded in a messenger RNA molecule that has been transcribed from DNA it can (with the help of a plethora of auxiliary enzymes, cofactors and an energy source, GTP) string amino acids together in the specified sequence. However, and this seems to have been universally ignored, the genetic blueprint for a ribosome is made up of a set of individual blueprints for the myriad of protein and ribonucleic acid components that make up a ribosome; there is no contiguous genetic blueprint

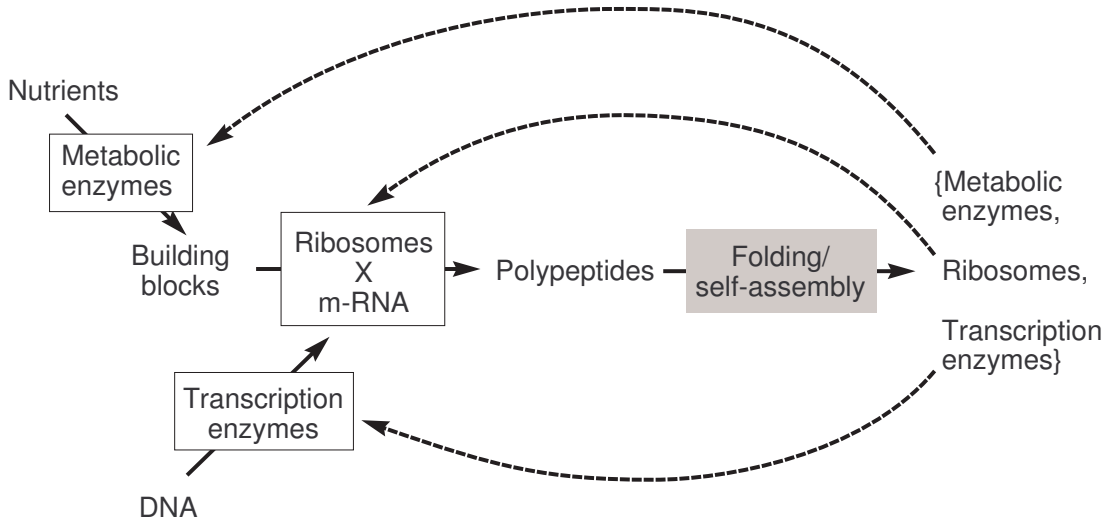

Figure 9: A summary of the biochemistry relevant to self-fabrication. As explained in the text, the assembly of supramolecular complexes (the shaded box) must be a spontaneous, unassisted process if the diagram is to depict a self-fabricating organisation. For the sake of clarity the diagram has been kept simple; for example, ribosomal RNA and its synthesis by transcription enzymes has been omitted. The chaperones that assist in the folding of some polypeptides are also absent on the diagram, but are discussed in the text.

for a complete ribosome. Therefore a ribosome never directly makes a ribosome, only the protein bits from which it is made up (the ribosomal RNAs are of course made by ribosomally-synthesised enzymes). Note that the problem of whether a Von Neumann constructor can fabricate itself directly therefore does not arise in the cell. Nevertheless, we still need to explain how the ribosomal components assemble into a fully functional entity. The fabrication of all ribosomes entails two processes: the *construction* of the parts (here the polypeptide chains and ribosomal RNA), and their subsequent *assembly* into a fully functional entity. In fact, there is another process wedged in between, namely that of the *folding* of newly synthesised polypeptide chains into a functional, three-dimensional conformation.

Above I argued, following Rosen, that for a system to be self-fabricating it must be closed to efficient causation. The existence of the non-covalent, supramolecular processes of folding and assembly therefore forces us to search for their efficient causes inside the system, which immediately confronts us with the 'insufficient number of machines' dilemma described above. Adding extra blueprints does not solve the problem; each addition implies a new polypeptide that has to be accounted for in terms of internal efficient cause for folding and possible association with other proteins. The recently discovered existence of chaperones that assist the folding of some polypeptides also cannot fully solve that part of the

supramolecular problem: chaperones are themselves proteins that need to fold in order to become active. There may be chaperones that assist the folding of other chaperones, but somewhere along the line there must then be chaperones that either fold spontaneously or assist in their own folding (or there must be a group that form a closed autocatalytic system). However, as far as we know chaperones fold spontaneously on their own. Similarly, with regard to assembly we are reasonably certain that supramolecular complexes such as ribosomes, spliceosomes, proteasomes, multimeric and oligomeric enzymes self-assemble spontaneously—the efficient and formal causes of self-assembly are embedded in the properties of the subunits of these complexes and in the properties of the environment. It is possible to dissociate these complexes in vitro and then have them reassemble themselves spontaneously. There appears to be no need for a physical agent to assist in the assembly process. It therefore turns out that, at least for life as we know it, unassisted self-assembly is the process that makes self-fabrication, and therefore life, possible. It is interesting to note that not one of the myriad of definitions of life listed in Barbieri (2003) and Popa (2004), nor the two regularly quoted sets of criteria for life—the Seven Pillars (de Duve, 1991) or PICERAS (Koshland Jr., 2002)—mention self-assembly as a necessary condition for life. In fact, I conjecture that if we discover life elsewhere in the universe, we shall recognise it by two properties: being autonomously self-fabricating by having learnt how to harness supramolecular chemistry and self-assembly, and having the ability to adapt and evolve by means of near-perfect replication and natural selection.

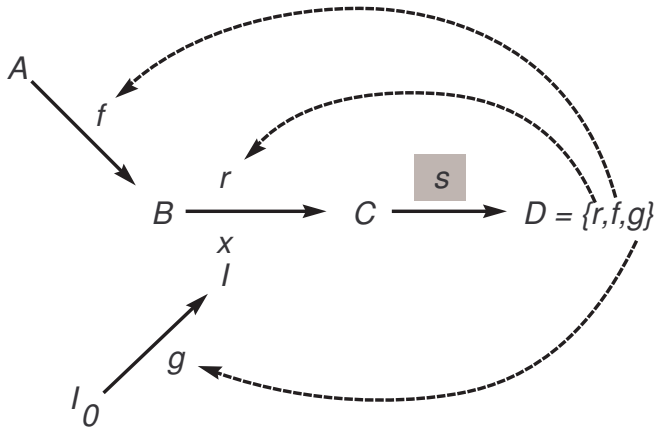

Figure 10: Adding self-assembly and information processing to the metabolism-construction system in Fig. 8. Here the fabricator r cannot make itself directly, but it can make all its own components and of course those of all the other machines; together they form set C. These machine components then self-assemble spontaneously through mapping s to form the set D, which then contains the constructor r and all other machines. This factory also makes a set of information processors g that translate archival information I_0 into blueprints I.

The abstract diagram in Fig. 8 can be extended to one that matches Fig. 9. This abstract representation of the self-fabricating metabolism-construction-assembly (M,C,A) organisation of living cells is given in Fig. 10. I propose this as an alternative to the replicative (M,R)-systems described by Rosen. Both are closed to efficient causation, but the (M,C,A) description has a number of distinct advantages. First, it maps onto the known biochemistry of the cell, whereas neither Rosen nor anybody else has been able to map the replication aspect of (M,R)-systems (which closes these systems to efficient causation) onto biochemical processes. In the language of category theory the replication component of (M,R)-systems is equivalent to an inverse evaluation map, and nobody seems to have been able to interpret this in terms of a physical process or object. The second advantage is that the (M,C,A) organisation reconciles Rosen's and Von Neumann's treatments. As mentioned in Section 2, Rosen's recasting of the Von Neumann architecture led to a logical paradox which has since apparently served to isolate these two views from each other (Rosen, 1959a, 1962). In a separate paper we shall show how to formally avoid this Rosen-Von Neumann paradox by using the formulation in Eq. 4 instead of Eq. 3. A third advantage is that the unassisted self-assembly component of (M,C,A)-systems obviates the need to postulate an agent that directly fabricates itself, a notion that I still find problematical. A fourth advantage is that the (M,C,A) architecture matches the triadic relationship between genotype, phenotype and ribotype suggested by Barbieri (1981, 2003) (see Fig. 11). Barbieri suggested that the machinery for protein synthesis (ribosomes, associated enzymes, tRNA adaptors) forms a logically distinct ontological type besides the phenotype and genotype. As suggested by Fig. 11, this idea fits perfectly with both (M,C,A)-systems and the triadic logic of biosemiotics as discussed by, for example, Hoffmeyer (1996).

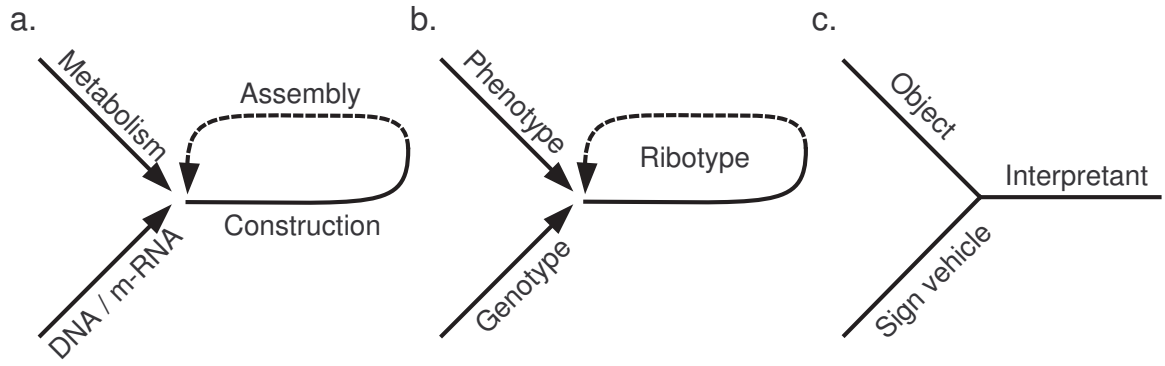

Figure 11: Triadic relationships in a. metabolism-construction-assembly (M,C,A) systems, b. Barbieri's genotype-phenotype-ribotype triad of ontologically distinct types, and c. the logical distinctions between sign vehicle, object and interpretant in biosemiotics.

7 Conclusion

The model of self-fabrication developed above is at this stage rather simple, informal and certainly not yet mathematically rigourous. However, a formal, mathematical exposition of the theory has been worked out and will be presented elsewhere. I have concentrated solely on the functional organisation of processes that make self-fabrication possible, and have purposefully ignored important aspects such as energy requirements, control and regulation, self-bounding, and communication with the environment to name but a few.

In conclusion I therefore argue for an epistemology for systems biology that is essentially relational and views everything that happens inside a living cell in the context of a functional organisation that makes self-fabrication possible. Working out all the implications this has for how we study, how we model and how we attempt to manipulate the cell is one of the tasks that systems biology must tackle if we want to lay claim to a deep understanding of life as we know it.

8 Author details

Jan-Hendrik Hofmeyr is Professor in the Department of Biochemistry at the University of Stellenbosch, South Africa. He obtained his Ph.D. in 1986 at the University of Stellenbosch after collaborating with Henrik Kacser (one of the founders of metabolic control analysis) and the enzymologist Athel Cornish-Bowden. Hofmeyr and his colleagues Jacky Snoep and Johann Rohwer form the Triple-J Group for Molecular Cell Physiology, a research group that studies the control and regulation of cellular processes using theoretical, computer modelling and experimental approaches. He has made numerous fundamental contributions to the development of metabolic control analysis and computational cell biology, and with Athel Cornish-Bowden developed both co-response analysis and supply-demand analysis as a basis for understanding metabolic regulation. A recent interest is to seek a way of expressing formally the functional organisation of the cell in terms of a theory of molecular fabrication, which could form a theoretical foundation for both systems biology and nanotechnology. He has an National Research Foundation A-rating and is a member of the Academy of Science of South Africa and a Fellow of the Royal Society of South Africa. He currently chairs BTK-ISSB (International Study Group for Systems Biology). He won the Harry Oppenheimer Fellowship Award for 2002 and the Beckman Gold Medal of the South African Biochemical Society in 2003.

9 Acknowledgements

This project was funded by the Ernest Oppenheimer Foundation in the form of the Harry Oppenheimer Fellowship Award and by the National Research Foundation of South Africa.

References

Alberghina L, Westerhoff H (Eds.). *Systems Biology: Definitions and Perspectives*, *Topics in Current Genetics*, vol. 13. Springer, Berlin, 2005.

Aristotle. *The Metaphysics*. Penguin Books, London, 1998. (translated by Hugh Lawson-Tancred) A.3.

Aristotle. *Physics*. The Internet Classics Archive (classics.mit.edu//Aristotle/physics.html), 350 B.C.E. (Translated by R. P. Hardie and R. K. Gaye) II.3.

Barbieri M. *The ribotype theory of the origin of life.* J. theor. Biol., 91(4), 545–601, 1981.

Barbieri M. *The Organic Codes: An Introduction to Semantic Biology*. Cambridge University Press, Cambridge, 2003.

Barbieri M. *Life is "artifact-making"*. J. Biosemiotics, 1, 107–134, 2005.

Beer S. *What is cybernetics?* Kybernetes, 31, 209–219, 2002.

Boogerd F, Bruggeman F, Hofmeyr JH, Westerhoff H (Eds.). *Systems Biology: Philosophical Foundations*. Elsevier, Amsterdam, 2007.

Cornish-Bowden A. *Making systems biology work in the 21st century (a report on the Biochemicial Society meeting 'Systems Biology: will it work?')*. Genome Biol., 6, 317, 2005.

Cunningham MA, Bash PA. *Computational enzymology*. Biochimie, 79, 687–689, 1997.

Dawkins R. *The Selfish Gene*. Oxford University Press, Oxford, New ed., 1989.

de Duve C. *Blueprint for a Cell: The Nature and Origin of Life*. Neil Patterson Publishers, Carolina Biological Supply Company, Burlington, 1991.

de Rosnay J. *The Macroscope: a new world scientific system*. Harper and Row, New York, 1979. (Translated by Robert Edwards).

Dobzhansky T. *Nothing in biology makes sense except in the light of evolution*. The American Biology Teacher, 35, 125–129, 1973.

Drexler KE. *Nanosystems: molecular machinery, manufacturing and computation*. John Wiley and Sons, New York, 1992.

Field RJ, Burger M. *Oscillations and Traveling Waves in Chemical Systems*. John Wiley & Sons, New York, 1985.

Freitas Jr RA, Merkle RC. *Kinematic self-replicating machines*. Landes Bioscience, Georgetown, 2004.

Gánti T. *The Principles of Life*. Oxford University Press, Oxford, 2003.

Heinrich R, Rapoport TA. *A linear steady-state treatment of enzymatic chains: General properties, control and effector strength*. Eur. J. Biochem., 42, 89–95, 1974.

Hoffmeyer J. *Signs of meaning in the universe*. Indiana University Press, Bloomington, 1996.

Hofmeyr JHS, Cornish-Bowden A. *The reversible Hill-equation: How to incorporate cooperative enzymes into metabolic models*. Comp. Appl. Biosci., 13, 377–385, 1997.

Hofmeyr JHS, Cornish-Bowden A. *Regulating the cellular economy of supply and demand*. FEBS Lett., 476, 46–51, 2000.

Hofmeyr JHS, Olivier BG, Rohwer JM. *From mushrooms to isolas: surprising behaviour in a simple biosynthetic system subject to feedback inhibition*. In: Hofmeyr JHS, Rohwer JM, Snoep JL (Eds.), *Animating the Cellular Map*, pp. 213–219. Stellenbosch University Press, Stellenbosch, 2000.

Kacser H, Burns JA. *The control of flux*. Symp. Soc. Exp. Biol., 32, 65–104, 1973.

Kampis G. *Self-Modifying Systems in Biology and Cognitive Science: A New Framework for Dynamics, Information and Complexity*. Pergamon Press, Oxford, 1991.

Kant I. *Critique of Judgment (1790)*. Macmillan and Company, London, 2nd ed., 1914. (Translated by J. H. Bernard). Available as e-text at oll.libertyfund.org/ToC/0318.php.

Kauffman SA. *Investigations*. Oxford University Press, New York, 2000.

Korzybski A. *Science and sanity: an introduction to non-Aristotelean systems and general semantics*. Institute of General Semantics, San Francisco, California, 5th ed. ed., 1994.

Koshland Jr DE. *The seven pillars of life*. Science, 295, 2215–2216, 2002.

Lehn JM. *Supramolecular Chemistry : Concepts and Perspectives*. Wiley-VCH, Weinheim, 1995.

Letelier JC, Soto-Andrade J, Abarzúa FG, Cornish-Bowden A, Cárdenas ML. *Organizational invariance and metabolic closure: Analysis in terms of (m,r) systems*. J. theor. Biol., 238, 949–961, 2006.

Maturana HR, Varela FJ. *Autopoiesis and Cognition: The Realisation of the Living*. D. Reidel Publishing Company, Dordrecht, Holland, 1980.

Mayr E. *Toward a New Philosophy of Biology: Observations of an Evolutionist*. Harvard University Press, Cambridge, Massachusetts, 1988.

Moreno A. *A systemic approach to the origin of biological organization*. In: Boogerd *et al.* (2007), pp. 243–268.

Nicolis G, Prigogine I. *Self-organisation in non-equilibrium systems*. Wiley, New York, 1977.

Popa R. *Between necessity and probability: searching for the definition of life*. Springer, Berlin, 2004.

Rosen R. *A relational theory of biological systems*. Bull. Math. Biophys., 20, 245–260, 1958a.

Rosen R. *The representation of biological systems from the standpoint of the theory of categories*. Bull. Math. Biophys., 20, 317–341, 1958b.

Rosen R. *On a logical paradox implicit in the notion of a self-reproducing automaton*. Bull. Math. Biophys., 21, 387–394, 1959a.

Rosen R. *A relational theory of biological systems II*. Bull. Math. Biophys., 21, 109–128, 1959b.

Rosen R. *Self-reproducing automaton*. Bull. Math. Biophys., 24, 243–245, 1962.

Rosen R. *Some relational cell models: The metabolism-repair systems*. In: Rosen R (Ed.), *Foundations of Mathematical Biology*, vol. II (Cellular Systems), chap. 4, pp. 217–253. Academic Press, New York, 1972.

Rosen R. *The roles of necessity in biology*. In: Casti JL, Karlqvist A (Eds.), *Newton to Aristotle*, pp. 11–37. Birkhäuser, New York, 1989.

Rosen R. *Life Itself: A Comprehensive Inquiry into the Nature, Origin, and Fabrication of Life*. Columbia University Press, New York, 1991.

Ruiz-Mirazo K, Moreno A. *Autonomy and emergence: how systems become agents through the generation of functional constraints*. Acta Polytechnica Scandinavica, Ma91 (Complexity, Hierarchy, Organization. Special Issue. Farre, G.L. & Oksala, T. (eds) The Finnish Academy of Technology), 273–282, 1998.

Ruiz-Mirazo K, Moreno A. *Basic autonomy as a fundamental step in the synthesis of life*. Artificial Life, 10, 235–259, 2004.

Ruiz-Mirazo K, Moreno A, Morán F. *Merging the energetic and the relational-constructive logic of life*. In: Adami C, Belew RK, Kitano H, Taylor CE (Eds.), *Artificial Life VI: Proceedings of the Sixth International Conference on Artificial Life*, pp. 448–451. MIT Press, 1998.

Ruiz-Mirazo K, Peretó J, Moreno A. *A universal definition of life: autonomy and open-ended evolution*. Orig. Life Evol. Biosph., 34, 323–346, 2004.

Russel R, Superti-Furga G. *Systems Biology: understanding the biological mosaic*. FEBS Lett., 579, 1771, 2005.

Schumacher EF. *Small is beautiful*. Vintage, London, 1973.

Snoep JL, Olivier BG, Westerhoff HV. *From isolation to integration: a systems biology approach for building the silicon cell*. In: Alberghina and Westerhoff (2005).

Taub AH (Ed.). *John von Neumann: Collected Works. Volume V: Design of Computers, Theory of Automata and Numerical Analysis*. Pergamon Press, Oxford, 1961.

von Neumann J. *The general and logical theory of automata*. In: Taub (1961), chap. 9, pp. 288–328. Delivered at the Hixon Symposium, September 1948; first published 1951 as pages 1–41 of: L. Jeffress, A. (ed), *Cerebral Mechanisms in Behavior*, New York: John Wiley.

von Neumann J. *Theory of Self-Reproducing Automata*. University of Illinois Press, Urbana, Illinois, 1966.

Westerhoff HV, Hofmeyr JHS. *What's systems biology? from genes to function and back*. In: Alberghina and Westerhoff (2005).

Westerhoff HV, Kell DB. *The methodologies of systems biology*. In: Boogerd *et al.* (2007), pp. 23–70.