# On organization[1]

Walter Fontana

Theoretical Chemistry
University of Vienna
Währingerstraße 17
A-1090 Vienna, Austria
and
International Institute for Applied
Systems Analysis (IIASA)
Schloßplatz 1
A-2361 Laxenburg, Austria

## *Biology's missing theory*

Biology has claim to two theories unto itself: Darwin's natural selection and Mendel's transmission rules. Both are correct, their joint operation can be nicely formalized, and together they are insufficient to account for the history of life as we know it.

Consider what is lacking. The formalization of Mendelism and Darwinism, known as the "Modern Synthesis", codified the evolutionary process as a problem in the dynamics of alleles (think genes) governed jointly by fitness (herein Darwinism) and transmission (herein Mendel) relations [7, 12, 21]. The theory, thus, assumes the *prior* existence of the very entities it is meant to explain. Indeed, nowhere in the formalism appears any representation of the organism. The theory cannot, therefore, be reasonably expected to account for changes in the features of organisms (phenotypes), nor for the progression from self-reproducing molecules to self-maintaining metabolisms[2], to modern cells containing organizational elements that once were autonomous simple cells, to multicellular units with cellular differentiation, to the entities with cognition that assemble in an auditorium to talk about this very failing.

[1] This version derives from a combination of two talks. One was given in Milan on December 15th, 1994, at the conference "The future of science has begun. Approaches to Artificial Life and Artificial Intelligence" sponsored by the Carlo Erba Foundation. The other was given on April 19th, 1996, at the University of Chicago on the occasion of the Dean's Symposium 1996 (Division of the Social Siences) "The Dynamic Emergence of Individuals and Cognition".
**This is joint work with Leo W. Buss, Department of Biology and Department of Geology and Geophysics, Yale University, New Haven CT 06520-8104 USA.**

[2] or vice versa

The genius of the crafters of the synthesis was to abstract away the organism – to see that the synthesis of Darwinism and Mendelism need not await a then unavailable "theory of the organism". That theory is still unavailable. Exploring its grounding is our project.

*A representation of chemistry for biology*

Any theory of biological organization must be grounded in a representation of that which organisms are composed of. The theory must be grounded in chemistry. We may picture chemistry as an informally systematized, autonomous body of knowledge at the interface between two very different tales of nature, physics and biology (Figure 1). Chemistry comes, of course, complete with its own theoretical structure. Quantum mechanics is the theory that nicely knits chemistry into the fabric of physics.

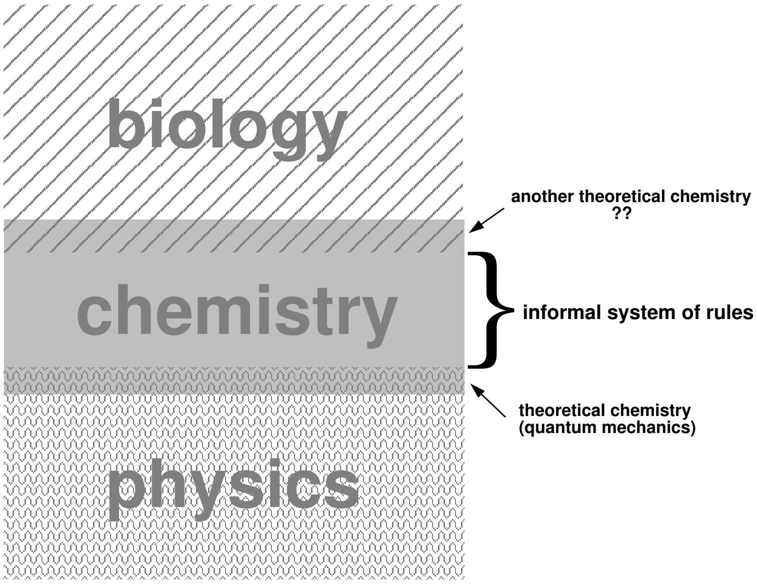


**Figure 1:** Chemistry as a threshold science. Quantum mechanics connects to physics, and sees the trees but not the forest of chemistry. Couldn't there be a different, more coarse-grained and action-focused description suited for understanding how organism arises from chemistry?

Quantum mechanics, however, is not a useful representation in which to ground a theory of biological organization. Quantum mechanics is a description that is oriented "downwards" and "inwards". Its level of detail puts the focus on single molecules and individual reactions, away from their context within a large number of concurrently operating molecules. For biological understanding, we require something different: an outward, coarse-grained, yet formal point of view capable of abstracting network relationships between reactions. The "chemistry for biology" we seek is a high level *specification language* for abstracting molecular actions, for plugging them together, and for generating and analyzing the *network* closures of these actions under a variety of boundary conditions.

Biological organizations are specialized chemical collectives. Biology's missing theory of the organism does not need a theory of chemistry at the level of quantum mechanics. Rather, it needs a formal representation of chemistry, truthful to stylized facts, designed to abstract those features of molecules germane to the origin and alteration of organized molecular collectives.

*Chemistry as object construction and equivalence relations*

What properties of chemistry do we require for the desired *representation for biology*. Consider an idealized chemical situation. In an empty container we mix two substances, say, methanol and acetic acid. We pass it to a friend, asking her to verify that it contains two substances. Upon analysis our friend will tell us that there are four, not two. In a sense we have put two "variables" into the system, and now need four to describe it. What happened is a chemical reaction. The chemist expresses the event with the notation:

$$\mathrm{CH_3OH + CH_3COOH \rightleftharpoons CH_3COOCH_3 + H_2O}$$

Read from left to right the "equation" states that two objects - molecules - interact to produce two new objects. What has changed are the objects themselves and they have done so on the basis of a structure that determines which transformations can occur. Chemistry at the level of the balance equation above is a matter of object construction upon interaction. This leads to equivalences between molecular compositions in the sense that many different reactive combinations yield the same product.

There is no mystery in this, but quantum mechanics does not capture chemistry at this level. In fact, if you go close enough, the molecules don't exist the way the chemist writes them down. No one seriously believes that a test tube contains an alphabet soup with symbol strings such as $\mathrm{CH_3OH}$. What exists depends on how close you get. At some level of resolution the quantitative properties, which are the stuff of quantum mechanics, start popping up again: a molecule is a (probability) distribution of mass and charge in space. Yet, no chemist really uses it in a significant way to plan, say, synthesis of complex compounds like drugs.

Chemists adopt a subtly different point of view. A chemist views the formulas he uses to notate his objects in much the same way that a mathematician views her weird strings of symbols. The strings "$\int x^2 \mathrm{d}x$" or "$4 + 3$" evidently possess a structure, yet it makes no sense at all to express this structure in a spatial coordinate system. To a mathematician's mind "$4+3$" is a syntactical structure that stands for an action; an action that yields "7". Slightly more generally, "$4+x$" stands for an action whose result depends on the nature of "$x$". Similarly, to a chemist's mind $CH_3OH$ is a statement about an action - a chemical action - that depends on the structure of the other object it interacts with. As a physical entity, $CH_3OH$ does, of course, have a spatial structure specifiable by an array of coordinates. *The chemist, however, sees shape as function.* That perspective is not about positions in space, but rather about relationships. To help codify this view, chemists have developed a symbolic notation and an elaborate system of rules to link - more or less sucessfully - the action of molecules to their formulas.

The notion of chemistry in textbooks is an informal one. It is not grounded in a mathematical or logical framework within which one can discover empirically verifiable chemical truths by "calculation". Yet this is what we have claimed that biology needs in a "theory of the organism"; a formal system that mathematizes a combinatorial variety of objects (think molecules) capable of specific actions whose effect is the construction of further objects. At present, we do little more than point in a (very) broad direction by making a leap of abstraction. Consider, we need at the very minimum an axiomatic theory with two ingredients: (i) a grammar to express syntactical structures and (ii) a formal way to connect these structures with actions on syntactical structures, such that (iii) structures bear equivalence relations. *This is exactly what the concept of computation does.*

*Computation as object construction and equivalence relations*

To appreciate the sense in which "computation" can stand as a rough proxy for chemistry, one should resist conflating the mathematical concept of "computation" with the notion of computers chewing numbers. The concept of computation is subtle. Its key aspect is the refinement of the world into "behavior" and "that which behaves". This refinement seems deceptively trivial, but until the 1930s there was no *formalization* of such a distinction. The prevailing view was to consider "behavior" as a mathematical function in the sense of one huge (in fact, infinite) look-up table, which tabulates assignments of outputs to inputs, without considering the process of *how* an output is obtained from an input. Incidentally, the interaction matrix used to specify entities in a traditional dynamical system corresponds to such a look-up table.

In contrast, the core idea of computation is to compress an infinity of possible behaviors (the look-up table) into a finite *rule* describing a *process* that transforms input to output. What would a system look like that enables us to reason about such rules? It clearly must have a syntactical component to express these rules. Rules, thus, are pieces of text (i.e., syntactical objects). By specifying how to construct them, the syntactical component defines a universe of "well-formed" possible objects.

So far this is a prescription for syntactical diversity. Since these rules are supposed to be about "actions", the intuition is that different syntactical objects may actually "mean the same thing". Consider two syntactically different structures, such as $10/2$ and $1 + 4$ (Figure 2). To establish whether they "mean the same thing", something must be *done*. That "something" consists in structural rearrangements of these objects on the basis of schemes defining elementary rewrites of specific configurations. A sequence of such rearrangements is called a "computation". One can think of the objects $10/2$ and $1+4$ as "unstable", because their structure permits actions - transformations - to occur. In a sense "computation" is a process by which such structures are stabilized. In the trivial example considered, both objects, $10/2$ and $1+4$, are transformed under the syntactical laws of arithmetic into the objects 5 and 5. The latter are identical, and, hence, one says that $10/2$ and $1 + 4$ are equal.

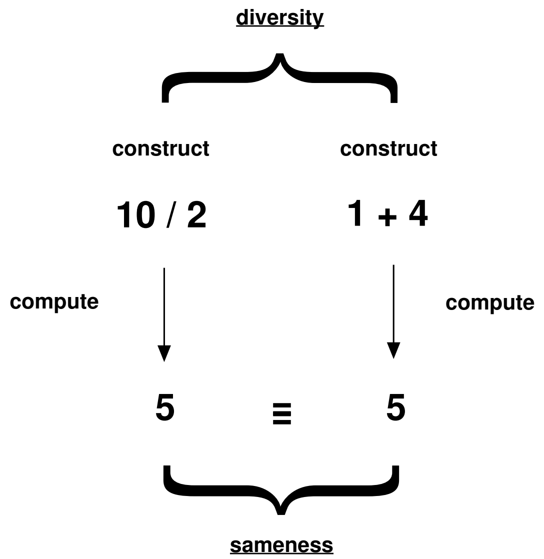


**Figure 2:** "Computation" is a theory about which object combinations are equivalent.

The first component of a system to reason about rules was a syntax for expressing their compositional structure. The second component, then, consists in defining laws that couple this syntax with actions that unfold on the syntax and establish a notion of sameness. Syntax and action are here but two sides of the same coin. Action is expressed by syntactical transformations triggered by specific syntactical configurations, i.e. the syntactical system *is* the theory of action.

Well, but just which action are we talking of? The action that was of interest to logicians in the early 20th century was that of a function. What exactly does the action of a function consist in? It depends on how one looks at it. One approach is to think of a function as a series of recursive "applications" of other simpler constituent functions. The "application" of a function, $A(x)$, to another, $B(y)$, is carried out by literally "substituting" for the "variable" $x$ of $A(x)$ the argument $B(y)$, i.e., $A(x := B(y))$. The formal system that analyzes the quoted words - in particular the notion of "substitution" - is known as the $\lambda$-calculus [3, 4]. It mirrors the informal usage of functions, except that it makes explicit all that we take for granted. An interesting aspect is that within $\lambda$-calculus *everything* is a function (including representations of "numbers"), and so there is no syntactical distinction between a function and the argument(s) it is applied to. Any object can be in either role. We skip further details, as they are not necessary to understand what follows at the "big picture" level.

*A first order approximation: Chemistry as a calculus*

The perspective of computation just sketched is a theory of action defined on the compositional structure of objects. This, we conjecture, is "chemistry" in a nutshell. It corresponds to the way chemists use the notation they have developed. Molecules are stable syntactical objects representing a specific range of possible actions. New "unstable" molecules - so-called "transition states" - are constructed by combining stable molecules when they collide. The molecular transition structures undergo syntactical rearrangements into stable product molecules. This corresponds to "computation" (i.e., the evaluation of a function applied to an argument). Different reactive combinations may yield the same product. This corresponds to the concept of "sameness" introduced above. The emphasis here is on abstracting just these two characteristics, construction and sameness, claiming that whatever deserves the label "chemistry" cannot be without them. Obviously, a number of familiar and important features of chemistry do not appear at all, or at best schematically, in this picture. This chemistry is clearly not what is captured by quantum mechanics at the lower interface in Figure 1. Yet, it is also far from a respectful projection of chemistry on the upper interface in Figure 1. We believe, however, that this approach broadly outlines the spirit of a "theoretical chemistry" that has yet to be invented.

For now we have with $\lambda$-calculus a simple off-the-shelf formalism that is remarkably well-suited to implement and to study the consequences of a first metaphor[3]:

[3] In the following we will use "function" and "operator" synonymously.

| | | |
|---|---|---|
| **chemistry** | $\cdots\cdots$ | **a calculus** ($\lambda$-calculus) |
| physical molecule | $\cdots\cdots$ | expression representing an operator |
| molecule's behavior | $\cdots\cdots$ | operator's action |
| chemical reaction | $\cdots\cdots$ | evaluation of functional application |

Having metaphorized chemical objects into abstract computational entities, we need to recover some of the physical context. Again, we stay minimal. Our interest is in overlaying this calculus/chemistry with a simple kinetics (i.e., a dynamical system). To explore this we turn to computer experiments. Our ultimate goal is a mathematical account, but we first need to know what to look for, and so we proceed in the spirit of an experimental mathematics.

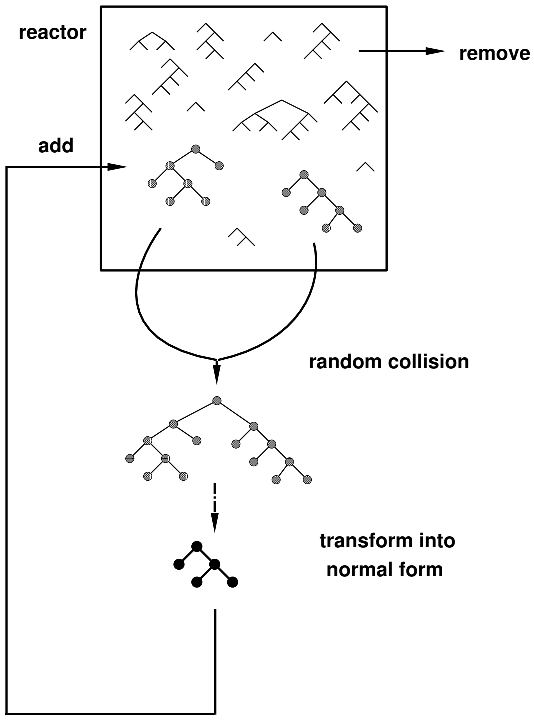


**Figure 3:** $\lambda$-calculus flow reactor. Two expressions $A$ and $B$ are chosen at random and an object, $(A)B$, is constructed by "application". Restructuring $(A)B$ into normal form ("computation") yields the new stable object.

The setting we chose is what chemists call a flow-reactor. It is a well stirred box in which we place a few thousand randomly generated $\lambda$-operators. We consider, thus, a situation in which each operator acts in a context of others. Operators can occur in multiple instances. As for molecular species, we have a notion of concentration. We now think of these operators as floating around in the box where they collide at random. Upon collision one operator is applied to the other. The resulting $\lambda$-expression is obtained according to the rules of the calculus which stand as a proxy for the "laws of chemical transformation". The result is a new operator that is fed back to the system. In this way we implement a simple chemical kinetics, where two object species interact with a probability proportional to their concentration.

Our metaphor does obvious violence to the chemistry we know. We briefly return to this point in conclusion. One problem, however, needs to be mentioned now. In a calculus, the application of an operator to another produces a single result, rather than two or more. Thus, a key feature of many chemical interactions - the feature of producing multiple products in a single event - is not realized. This doesn't jeopardize our exploration of the consequences of joining construction with dynamics. It only means that the dynamical systems component has to be artificially constrained. We do so by not using up the operators immediately at the moment of the reaction, as it would be the case in chemistry. To compensate for keeping the reactants, we introduce a dilution flow by removing one randomly chosen object from the system each time an interaction has happened. This means that every object has effectively a finite lifetime. No operator will persist over a long period of time, unless it is somehow produced by interactions among extant operators. This is schematically summarized in Figure 3.

In the first phases of its history, the system is highly innovative in the sense that almost every interaction produces an operator that is new to the box. The system starts "moving in object space". We leave it moving, and analyze its content after hundreds of thousands of collisions have occurred.

*Level 0*

At first something rather disappointing happens. We find that all operators are of the same kind. It is easy to find out what this kind of operator does: it performs an identity operation on itself. It copies itself, and, therefore, behaves kinetically like a replicator. More complicated instances also occur, such as small hypercycles of functions that mutually copy one another. The hypercycle is a coupling structure invented by Eigen and Schuster [5], to which a key role in the early stages of prebiotic molecular evolution is imputed.

Thus, without further boundary conditions the reactor settles on ensembles of operators, each of which is the fixed point of the interaction with some other operator in the same ensemble. In the case illustrated in Figure 4 no interaction generates products that are not already in the system. The set of objects has become closed with respect to interaction. Under such circumstances we can describe the system without reference to the interal structure of the constituent objects. All that we need are three concentration variables and their couplings (i.e., who interacts with whom to produce whom). When convergence in object space has occurred, the system has become a conventional dynamical system[4].

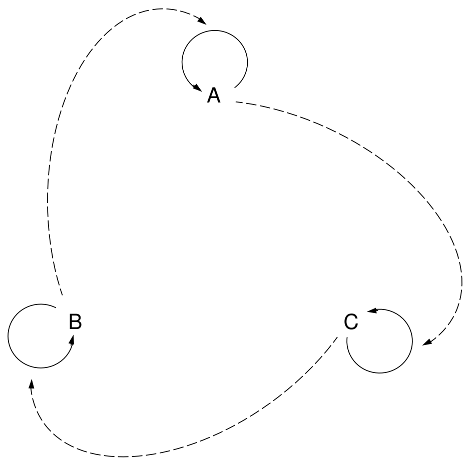


**Figure 4:** A simple three membered hypercycle of $\lambda$-expressions.

*Level 1*

The system permits to test the consequences of constraints on the action of the operators. For example, we can prohibit all identity actions, thus outright barring replicators from the system. Whenever two operators interact to produce an operator that is syntactically identical to one of them, the collision is declared "elastic", that is to say: nothing happens. The analysis reveals a quite different picture from the previous case. Only those main points are sketched that are common to the outcomes of all experiments.

1. The box contains a large diversity of operators, a substantial fraction of which is engaged in a self-maintaining network of mutual production pathways. Like before, we see a fixed-point behavior in object space. This time, however, not at the level of a single operator, but at the level of a collective.

2. At the syntactical level there exist common regularities that characterize the structures of all operators maintained in the system. These regularities define a grammar, i.e., lawful arrangements of identifiable substructures. The grammar is characteristic for the ensemble and defines a subspace of $\lambda$-calculus. Furthermore, when new operators are created from interactions within the system, their structure conforms with the grammar. That is to say, the subspace specified by the grammar is invariant as interactions proceed; closure has been attained.

[4]Notice that this object-less description breaks down, as soon as we perturb the system by injecting a new object...

3. A few laws specify all relationships of transformation among operators whose structure conforms with the grammar. These laws constitute an abstract algebra. The algebraic laws can be read as a $\lambda$-calculus independent definition of the action associated with the building blocks on which the grammar is defined. Grammar and laws remain invariant. They provide a complete, yet $\lambda$-calculus independent, level of description of the system.

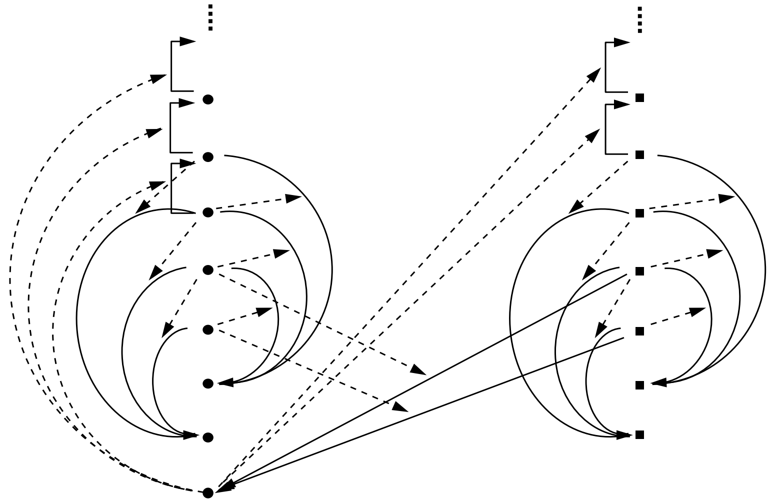

**Figure 5:** A simple self-maintaining organization. The dots (left) and the squares (right) represent two families of $\lambda$-expressions with a particular grammatical structure (whose details are omitted). A solid arrow indicates the transformation of an argument (tail) to a result (tip) by an operator (dotted arrow). For clarity, only a subset of the possible interrelations is shown. Notice the connectivity enables kinetic confinement. Most transformations yield objects at the bottom (leading to an increasing concentration profile from top to bottom). Some operations, however, yield objects up the "ladder", thus establishing self-maintenance. Both syntactical families depend on each other for maintenance as indicated by the "cross-family" connections.

These properties make the system one single object in its own right. It is an object of an alltogether different class than its components. We call such an invariant entity an organization. Figure 5 illustrates the main points of an actual example. The individual components of this entity are constantly being turned-over[5]. That which persists over time is something abstract: the grammar and

[5]When replicators also entertain constructive interactions, we sometimes find that the contents of the flow reactor keep changing indefinitely, thus violating property (1), while properties (2) and (3) remain valid. An example is discussed in [8].

the algebraic laws. They allow to decide whether a given arbitrary operator-particle taken from our formal chemical universe belongs to the organization or not. Grammar and algebra constitute a *boundary* of the new object. This boundary delimits the organization in the abstract object space, not in physical space. An organization is an object made of components that are held together on the basis of specific invariant relationships of transformation. This does not suffice to separate different instances of the same organization. Some form of physical space seems to be required to encapsulate an organization, but physical space is clearly not required to bound it.

The main take-home at this point is a useful definition of what we mean by an "organization": a kinetically self-maintaining algebraic structure. Such algebras are generated spontaneously in this model. The link with algebra allows to clarify some noteworthy stability properties of organizations. They turn out to be extremely robust with respect to both the subtraction and the addition of operators, but for different reasons.

When removing large portions of an organization, it repairs itself. Some organizations are even undestructible: they regenerate themselves from any component-operator. The reason for this robustness with respect to subtraction is the existence of generator sets of the underlying algebra. These are sets of operators whose iterated interactions rebuild step by step the population of operators that carry the organization.

Consider the addition of operators. The integration of an operator that is not a member of a given organization requires the organization to change. The organization must stably sustain the pathways necessary for the maintenance of the perturbing operator, otherwise it will be removed by the dilution flow. These pathways involve direct and indirect products generated by the perturbing agent's actions within the organization. Organizations are very resilient to the addition of new operators. However, when they change, they change in a characteristic way. The original unperturbed organization is conserved, and "on top" of it a further layer of pathways is added. This layer implies the addition of new grammatical elements and algebraic laws without undoing the original. This is exactly what is known as an algebraic extension. It is clear that the relationships of transformation that hold between the operators of an organization severly constrain its variation and possible extensions. At present we don't understand these constraints in any systematic way.

*Level 2*

Self-maintaining organizations can be combined. In some cases they do not drive each other out of existence, but they coexist stably. This coexistence is structural. Cross-interactions between members of different organizations produce new operators that belong to neither organization. These operators and their interactions constitute a "glue" that integrates both self-maintaining organizations into a higher order unit, where they continue to persist as autonomous entities (Figure 6).

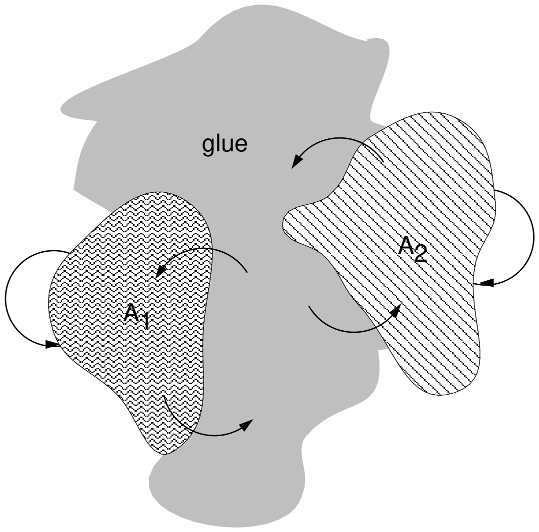


**Figure 6:** A cartoon of a Level-2 situation. Two autonomous self-maintaining organizations are "glued" together within a higher order organization by the products of their "cross-talk". If one of the component organizations disappears, so does the glue. If many organizations were put together, what do you think would happen?

*Concluding remarks*

There are two ways of looking at chemistry, which reflect its threshold position between physics and biology. From the vantage point of physics, molecules appear as nuclear and electronic distributions in space, governed by electromagnetic interactions. This view emphasizes a molecule as a many-body dynamical system. It succeeds (at least in principle) in deriving its quantitative properties, such as spectra, bond lengths and angles. In contrast, from the vantage point of molecular biology, molecules appear as units of syntactical action tied together into organized networks of transformations. This view emphasizes a molecule as a functional entity to which a "semantics" can be attributed on the basis of its relations within a context of other molecules.

The conceptual gap between the two views is worth noting. The physics point of view offers a formal theory, yet its focus on quantitative aspects puts it so close to (into?) individual molecules as to lose the view of molecules as carriers of transformation relations that can be plugged together into systems. In fact, it is the latter view that the bulk of informally codified chemical knowledge used in laboratories is about. The loss of formality constitutes no more a problem for the everyday practice of chemistry, as a growing number of laws and amendments does for the practice of lawyers. Yet, one is left wondering what kind of theory, if any, is going to integrate the findings of molecular biology on a formal platform. An answer to the question "what is life?" may perhaps depend on how we answer the question "what is chemistry?". If we don't frame the basic constituents of life in the "appropriate" way, how are we going to frame their most interesting collective property?

The stance taken in Leo Buss' and my approach is to view molecules as compositional objects capable of constructive interactions. The foundations of computation offer a theory of such objects. In a very first round this stance amounts to treating molecules as if they were symbolic functions expressed in a canonical calculus. The calculus of choice is $\lambda$-calculus. Our intent is *not* to produce an arbitrary toy chemistry, but rather to produce a valid *abstraction of chemistry.* In the end this means a formal translation between chemical syntax and a syntax of computation. New axiomatic systems may have to be invented for this purpose. For now we find it surprising that among the available formalizations of computation there is such a central one that fits the ontology of chemistry at all.

The results of our work even suggest that with respect to biological applications it is not necessary to succeed in a truthful translation of actual chemistry. The crude approximation achieved by simply placing $\lambda$-calculus in a constrained dynamical setting already yields a highly useful notion of functional organization as a special class of fixed points with respect to an endogenous "motion in object space". The abstract organizations generated within our system exhibit properties, such as regeneration, structure-dependent extension and capacity for hierarchical nesting that are akin to those of living organisms. To illuminate the mathematical nature of organized molecular collectives, it suffices, as a first step, to succeed in a proper abstraction of chemical *concepts* (rather than in producing a homomorphism to chemical reality).

There is no doubt that $\lambda$-expressions are far from molecules and our organizations far from organisms. The extent to which chemical realism can be increased without giving up formalism will determine which further properties of concurrent chemical organization become amenable to theory. The major limitations of the present abstraction are the absence of selective interactions (i.e., some notion of "steric constraints"), mass-action or stoichiometry, rate constants and the violation of interaction symmetry (functional application being not commutative), to mention but a few. In the attempt of incorporating interaction specificity (i.e., "shape") at a formal level, we have recently refined our metaphor to $\lambda$-calculus with a type system. Once having arrived at typed $\lambda$-calculus, it is almost inevitable to proceed further to proof-theory. This stance leads one to view a chemical action as a logical proposition and a molecule capable of that action as a proof of that proposition. Chemical reactions, then, appear as rules of inference (such as "cut" or "modus ponens"), with reactants as hypotheses and products as conclusions. While still falling short of chemistry in detail, we begin to see how the connections with recently developed proof-systems (e.g., linear logic [11]) enable us to address each of the limitations mentioned. Formal logic, it seems, is more appropriate for capturing chemistry than human cognition and reasoning. For an interim report the reader may consult [10].

*Further reading*

Related work that provided important motivation has been published by Bagley, Farmer, Kauffman, Packard and Rasmussen [1, 6, 13, 14, 18], in particular Kauffman's pioneering work on reflexively autocatalytic chemical collectives. Maturana and Varela's concept of "autopoiesis" [15, 16, 17, 19, 20] is of particular relevance. Leo Buss' classic [2] was the driving force to go beyond standard formal frameworks in evolutionary biology to address the issue of organization. Our joint work is reported in [8, 9, 10].

*Acknowledgement:* Thanks to Leo Buss for ironing and starching this text.

# References

[1] R. J. Bagley and J. D. Farmer. Spontaneous emergence of a metabolism. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II*, Santa Fe Institute Studies in the Sciences of Complexity, pages 93–141, Redwood City, 1992. Addison-Wesley.

[2] L. W. Buss. *The Evolution of Individuality.* Princeton University Press, Princeton, 1987.

[3] A. Church. A set of postulates for the foundation of logic. *Annals of Math. (2)*, 33:346–366, 1932.

[4] A. Church. *The Calculi of Lambda Conversion.* Princeton University Press, Princeton, 1941.

[5] M. Eigen and P. Schuster. *The Hypercycle.* Springer Verlag, Berlin, 1979.

[6] J. D. Farmer, S. A. Kauffman, and N. H. Packard. Autocatalytic replication of polymers. *Physica D*, 22:50–67, 1982.

[7] R. A. Fisher. *The Genetical Theory of Natural Selection*. Clarendon Press, Oxford, 1930.

[8] W. Fontana and L. W. Buss. 'The arrival of the fittest': Toward a theory of biological organization. *Bull. Math. Biol.*, 56:1–64, 1994.

[9] W. Fontana and L. W. Buss. What would be conserved 'if the tape were played twice'. *Proc. Natl. Acad. Sci. USA*, 91:757–761, 1994.

[10] W. Fontana and L. W. Buss. The barrier of objects: From dynamical systems to bounded organizations. In J. Casti and A. Karlqvist, editors, *Barriers and Boundaries*, pages 56–116. Addison-Wesley, Reading, MA, 1996.

[11] J.-Y. Girard. Linear logic. *Theoretical Computer Science*, 50:1–102, 1987.

[12] J. B. S. Haldane. *The Causes of Evolution*. Longmans Green, New York, 1932.

[13] S. A. Kauffman. Cellular homeostasis, epigenesis and replication in randomly aggregated macromolecular systems. *J. Cybernetics*, 1:71–96, 1971.

[14] S. A. Kauffman. Autocatalytic sets of proteins. *J. Theor. Biol.*, 119:1–24, 1986.

[15] P.-L. Luisi. Defining the transition to life: self-replicating bounded structures and chemical autopoiesis. In W. Stein and F. J. Varela, editors, *Thinking About Biology*, Santa Fe Institute Studies in the Sciences of Complexity, pages 3–23, Redwood City, 1993. Addison-Wesley.

[16] H. Maturana and F. J. Varela. *De Máquinas y Seres Vivos: Una teoría de la organizacíon biológica*. Editorial Universitaria, Santiago de Chile, 1973. Reprinted in: H. Maturana and F. J. Varela, *Autopoiesis and Cognition: The Realization of the Living*, 1980.

[17] H. Maturana and F. J. Varela. *Autopoiesis and Cognition: The Realization of the Living*. D. Reidel, Boston, 1980.

[18] S. Rasmussen, C. Knudsen, R. Feldberg, and M. Hindsholm. The coreworld: Emergence and evolution of cooperative structures in a computational chemistry. *Physica D*, 42:111–134, 1990.

[19] F. J. Varela. *Principles of Biological Autonomy*. North-Holland, New York, 1979.

[20] F. J. Varela, H. R. Maturana, and R. Uribe. Autopoiesis: the organization of living systems, its characterization and a model. *BioSystems*, 5:187–196, 1974.

[21] S. Wright. Evolution in mendelian populations. *Genetics*, 16:97–159, 1931.