A Systematic Approach to Nanotechnology Based on a Small Set of Molecular Building Blocks

Bruce J. MacLennan, *Member, IEEE*

***Abstract*— A small set of molecular building blocks will allow the implementation of "universally programmable intelligent matter," that is, matter whose structure, properties, and behavior can be programmed, quite literally, at the molecular level. The approach is based on combinatory logic, a term-rewriting or graph-substitution system, which can be interpreted as an abstract model of molecular processes. In such systems there are small sets of substitution operations (reactions) that are universal in the sense that they can be used to program any Turing-computable function. Although these operations provide the building blocks for univerally programmable intelligent matter, there are many issues that arise in the context of molecular computation, which we address.**

***Index Terms*— adaptive matter, combinatory logic, DNA computing, functional programming, intelligent matter, molecular computing, nanocomputation, nanotechnology, programmable matter, term-rewriting systems.**

I. Introduction

We use the term *intelligent matter* to refer to any material in which individual molecules or supra-molecular clusters function as agents to accomplish some purpose. Intelligent matter may be solid, liquid, or gaseous, although liquids and membranes are perhaps most typical. *Universally programmable* intelligent matter (UPIM) is made from a small set of molecular building blocks that are universal in the sense that they can be rearranged to accomplish any purpose that can be described by a computer program. In effect, a computer program controls the behavior of the material at the molecular level. In some applications the molecules self-assemble a desired nanostructure by "computing" the structure and then becoming inactive. In other applications the material remains active so that it can respond, at the molecular level, to its environment or to other external conditions. An extreme case is when programmable supra-molecular clusters act as autonomous agents to achieve some end.

Although materials may be nano-engineered for specific purposes, we will get much greater technological leverage by designing a "universal material" which, like a general-purpose computer, can be "programmed" for a wide range of applications. To accomplish this, we must identify a set of molecular primitives that can be combined for widely varying purposes. The existence of such universal molecular operations might seem highly unlikely, but there is suggestive evidence that it may be possible to discover or synthesize them.

II. Approach

Term-rewriting systems [1]-[2] are simple computational systems in which networks are altered according to simple rewrite rules, which describe *substitutions* that have much in common with abstract chemical reactions. (One particular term-rewriting system, the lambda calculus, has been used already to model prebiotic chemical evolution [3]-[5].) Term-rewriting systems have been investigated extensively by mathematicians and computer scientists for several decades [6]-[7].

The author is in the Department of Computer Science, University of Tennessee, Knoxville, TN 37996-3450 USA (e-mail: maclennan@cs.utk.edu).

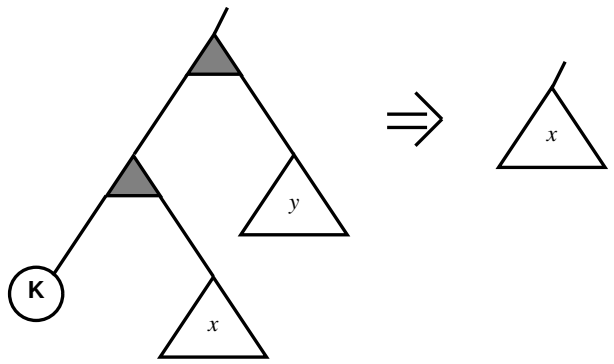

Fig. 1. K-substitution. The network on the left may be replaced by that on the right, wherever and whenever it occurs. The variables x and y represent arbitrary networks.

One attractive feature of term-rewriting systems is that many of them have the *Church-Rosser property* [2], [8], which means, roughly, that substitutions can be done in any order without affecting the computational result [9, ch. 4]. Therefore these term-rewriting systems have been investigated as a possible basis for massively-parallel computer architectures [6]-[7]. This is an important property for a model of molecular computation, in which molecular processes take place stochastically.

A. SK Programming

One class of term-rewriting systems, the *combinatory logic systems* [9]-[12], is very relevant for programmable intelligent matter, for it has been known since the early twentieth century that there are several small sets of substitution operations that can be used to program any Turing-computable function. One such *universal* set is the SK calculus.

The SK calculus is defined by two simple substitution rules. The K-substitution is expressed by this rewrite rule,

$$((\mathsf{K}X)Y) \Longrightarrow X,$$

which describes the transformation shown in Fig. 1, in which X and Y represent any networks. In effect, since the value of $(\mathsf{K}X)$, when applied to any Y, is X, the K operation, when applied to X yields the constant function $(\mathsf{K}X)$. This is the

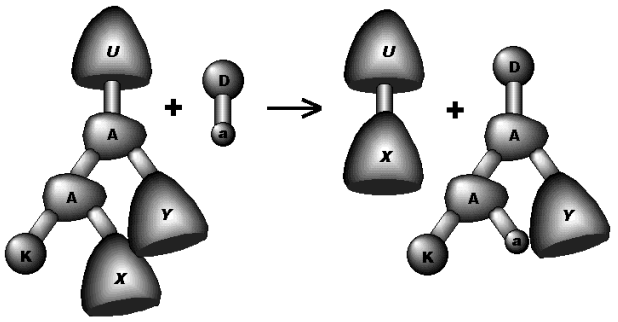

Fig. 2. K-substitution as a molecular process. A spontaneous chemical reaction, which replaces the reactants on the left with the molecular clusters on the right. U, X, and Y represent arbitrary molecular networks; K, A, D, and a are particular funtional groups on which the reaction depends.

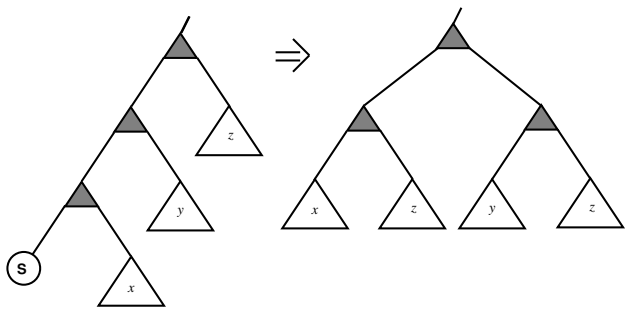

Fig. 3. S-substitution with copying. The vaiables x, y, and z represent arbitrary networks. The S-substitution, $(((\mathsf{S}x)y)z) \Longrightarrow ((xz)(yz))$, may be interpreted as making a duplicate copy of z.

interpretation, but the computational effect is entirely expressed in the substitution in Fig. 1.

It will be apparent that this substitution rule suggests a molecular process, but the equivalent depiction in Fig. 2 makes the similarity more apparent. It can be put in the style of a chemical reaction, including reaction resources and waste products:

$$U\mathsf{A}_2\mathsf{K}XY + \mathsf{D}\mathsf{a} \longrightarrow UX + \mathsf{D}\mathsf{A}_2\mathsf{K}\mathsf{a}Y.$$

Here A, K, D, and a are functional groups, and U, X, and Y represent arbitrary molecular networks. D is a disposal operator and a is a computationally inert place-holding group.

The S operator is only slightly more complicated; it is defined

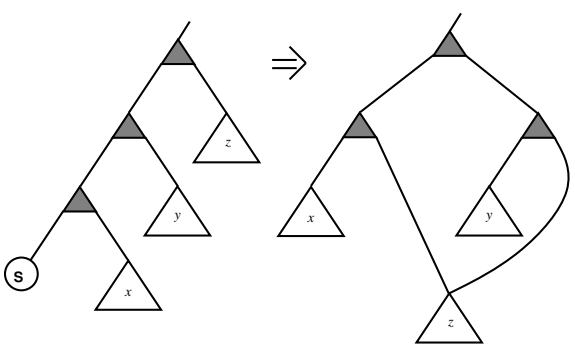

Fig. 4. S-substitution with sharing. The vaiables x, y, and z represent arbitrary networks. The S-substitution, $(((\mathsf{S}x)y)z) \Longrightarrow ((xz)(yz))$, may be interpreted as sharing a single copy of z.

by the rewrite rule,

$$(((\mathsf{S}X)Y)Z) \Longrightarrow ((XZ)(YZ)).$$

There are two ways of interpreting it as a network substitution, depending on whether we make a new copy of Z (Fig. 3) or share a single copy (Fig. 4). However, the Church-Rosser property [2], [8] shows that the two interpretations lead to the same computational result, but the interpretations have practical differences, which are discussed later

It is important to stress the significance of the SK calculus: these two simple operations are capable of computing anything that can be computed on any digital computer. This is certainly remarkable, and so it is surprising that there are quite a few other universal sets of combinators. For example, the set of K with B′ and W is universal, where the latter two operators are defined:

$$\begin{aligned}(((\mathsf{B}'X)Y)Z) &\Longrightarrow (Y(XZ)),\\ ((\mathsf{W}X)Y) &\Longrightarrow ((XY)Y).\end{aligned}$$

A third universal set comprises K and these three combinators:

$$\begin{aligned}(((\mathsf{B}X)Y)Z) &\Longrightarrow (X(YZ)),\\ ((\mathsf{C}^*X)Y) &\Longrightarrow (YX),\\ (\mathsf{W}^*X) &\Longrightarrow (XX).\end{aligned}$$

There are even some general guidelines [9, sec. 5H] for universality (i.e., the combinators must be able to delete, duplicate, and permute). The existence of multiple universal sets is very fortunate, because it implies that when we begin to search for molecular implementations of these operations, we will have a greater probability of finding reactions implementing at least one universal set of substitutions.

The combination of the parallel computation permitted by the Church-Rosser property and the simplicity of the SK calculus has led computer scientists to investigate it as a basis for parallel computer architecture [6]-[7]. There are simple algorithms for translating functional computer programs into SK networks, and considerable effort has been devoted to optimizing them. Therefore, if we can identify molecular processes corresponding to a universal set of combinators (SK, for example), then we can at least see the possibility of writing a computer program and translating it into a molecular process.

To illustrate the idea, we will present, with little explanation, a program for computing a nanotube. A single ring of the nanotube is computed by:

$$\begin{aligned}\mathsf{Ring}(X,Y) &= R,\\ \text{where rec } R &= \mathsf{Aux}(X,Y,R),\\ \mathsf{Aux}(X,\mathsf{nil},R) &= R,\\ \mathsf{Aux}(x:X,y:Y,R) &= (x,y):\mathsf{Aux}(X,Y,R).\end{aligned}$$

The function is given two molecular chains, $X = (x_1,\ldots,x_m)$ and $Y = (y_1,\ldots,y_m)$. It creates a ring of the x_i, each of which is linked to the next x group in the ring, as well as to the corresponding y_i (see Fig. 5). A nanotube is computed by creating a series of linked rings:

$$\begin{aligned}\mathsf{Tube}(\mathsf{nil},X,Y) &= \mathsf{Ring}(X,Y),\\ \mathsf{Tube}(a:N,X,Y) &= \mathsf{Ring}[X,\mathsf{Tube}(N,X,Y)].\end{aligned}$$

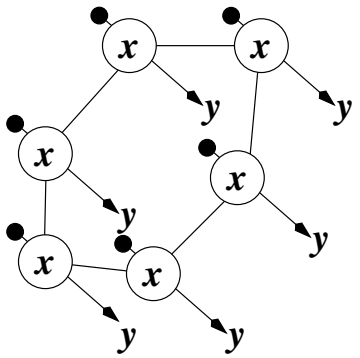

Fig. 5. Example: Single ring of a nanotube computed by a simple functional computer program, which may be translated into an SK network. The variables x and y represent two molecular species provided to the construction operation.

If $N = (a_1, \ldots, a_{n-1})$ is any chain of length $n-1$, then $\mathsf{Tube}(N, X, Y)$ will compute a nanotube of length n, with x groups forming the sides of the tube, and y groups forming its terminus. Our discussion of this example is of necessity short, but it will serve to demonstrate how simple programs can generate useful nanostructures.

B. Replication/Sharing Problem

Intriguing as the SK calculus is as a basis for universally programmable intelligent matter, it also illustrates some of the problems that our research is addressing. As previously mentioned, term-rewriting systems assume that a network can be copied for free, as illustrated in Fig. 3. This is certainly a bad assumption for molecular computation, in which the time to replicate a structure is at least proportional to the logarithm of its size. It is also extremely wasteful, since when programs are compiled to use only the S and K combinators, one often observes that an S operation replicates a structure, which is almost immediately discarded by a K. There are ways to avoid much of this needless replication (at the expense of introducing additional primitive combinators), but considerable replication will remain.

The obvious solution is to use the sharing implementation (Fig. 4), since this does not require any copying. It is the solution adopted in many implementations of SK on conventional computers, in which one may have any number of pointers to a single data structure. However, this option does not seem to be possible in molecular computing, in which each link connects only two groups of atoms.

There are some possible solutions to this problem, but they will not be discussed at this time. We raise the copying/sharing issue to show that the constraints of molecular computing are different from those of electronic computing. Therefore, while term-rewriting systems, combinatory logic, and the SK calculus in particular are suggestive of how universally programmable intelligent matter might be implemented, we must be prepared to develop new models of computation that are compatible with the constraints of molecular processes. Developing such a model is a principal objective of our research.

III. Extensions

A. Sensors and Effectors

To expand the range of application of universally programmable intelligent matter and for other practical purposes, it is advisable to extend the set of primitive operations beyond those minimally necessary for computational universality (e.g., S and K). First, we might want to add *sensor operations* that respond differently in different environmental conditions. For example, they might be sensitive to light or to the presence of some chemical. The results of these tests could be used to control conditional execution of the program. In addition to such external input to the program, it is also useful to have means for external output, which can be accomplished with *effector operations*. These reactions, when they take place, cause some noncomputational effect, such as the release of a chemical, the emission of light, or physical motion. They are one of the ways that intelligent matter can have an effect beyond its own internal computational reconfiguration.

To some extent the sensors and effectors are ad hoc additions to the basic computational framework (e.g., SK). However, they are fundamentally incompatible with it in one sense, for the time when their reactions take place is usually important. They are termed *imperative* operations and do not have the Church-Rosser property. Therefore, programs incorporating them must have means for controlling the time of their execution. Fortunately, these issues have been addressed long ago in the design and implementation of functional programming languages (see references in [13]), and the results of those investigations can be applied to universally programmable intelligent matter.

B. Production of Molecular Networks

Another issue that must be addressed is the production of a molecular combinator network (e.g., an SK tree) from a macroscopic program, and the subsequent replication of a large number of copies. Although the best approach is one of the objectives of our research, a possible method can be presented at this time. Arbitrary combinator trees can be represented uniquely as parenthesized strings, such as "(((SK)S)(SK))." Therefore, such a string could be encoded by a chain of four molecular groups (s, k, p, q), such as "$pppskqsqpskqq$" for the previous example. Thus we proceed in stages. The program is compiled into SK trees (or other combinators); the trees are flattened into parenthesized strings; and the strings are encoded into molecular chain structures (e.g., DNA sequences), which are synthesized and replicated by standard techniques from biotechnology. The replicated program chains are converted back into (now molecular) networks by a simple set of substitution rules, implemented chemically.

C. Developing an Application

To tie the foregoing ideas together we may present the typical process of developing an application of universally programmable intelligent matter:

1) Write a program in an appropriate high-level programming language to create the desired nanostructure or to exhibit the desired interactive behavior at the nanoscale.

Debug and simulate the execution of the program on a conventional computer.
2) Compile the program into a combinator tree (e.g., a network of S, K, and other combinators).
3) Simulate (on a conventional computer) the substitutions on the network, but subject to molecular constraints (e.g., including reactant concentrations, substitution errors, etc.).
4) On a computer, flatten the combinator tree into a string representing a sequence of DNA bases.
5) Use this string to guide the synthesis of a DNA sequence.
6) Replicate the DNA sequence until the required number of copies of the program are produced.
7) Use the translation or tree-building substitutions to construct a molecular combinator tree from each DNA string. (An intermediate RNA stage could be used, if required.)
8) Supply reactants for the computational substitutions (e.g., S, K, A, D groups), and allow the reaction to proceed to equilibrium.
9) If the application is static, wash out or otherwise eliminate any remaining reaction waste products.
10) If the application is static, substitute permanent replacement groups for computational groups by ordinary chemical processes (if required by the application).

IV. Related Work

It is worthwhile to contrast universally programmable intelligent matter with some other more or less related ideas.

Programmable matter [14]-[17] is an approach to computation based on lattice-like arrays of simple computational elements; cellular automata (such as Conway's "game of life") are examples. Although techniques from the "programmable matter paradigm" certainly will be applicable to universally programmable intelligent matter, there are differences in objective: programmable matter seems to be intended primarily for implementation on electronic digital computers, and computational universality does not seem to be a goal.

Complex adaptive matter (CAM) has been under investigation at Los Alamos National Laboratory as an approach to adapting matter or materials to a desired functionality by a quasi-evolutionary process comprising amplification, noise or variation, and filtering. Again, the goals are different, but there are several intersections with our research. First, like ours, their focus is on molecular processes rather than electronic computation. Second, CAM techniques might be used for synthesizing functional groups implementing universal sets of molecular operators. Third, such a universal set might provide building blocks for the quasi-evolutionary CAM process. (Indeed, we have already begun investigating statistical properties of "soups" of SK complexes, which might be used as the raw materials of the CAM process [18]-[19].)

Smart Matter is being developed at the Xerox Palo Alto Research Center. Its long-term goals are quite similar to those of intelligent matter, as defined in this paper. However, the current focus seems to be on small (but not molecular) building blocks that combine sensor, computation, and actuation functions. That is, the goal is MEMS (Micro-Electro-Mechanical Systems) rather than nanotechnology. Certainly, however, intelligent matter, in our sense, will benefit from the more general techniques of distributed control and embedded computation that might come out of the Smart Matter project.

V. Applications

Finally, it will be worthwhile to discuss briefly some of the possible applications of universally programmable intelligent matter, which may be static or dynamic (or interactive). By a *static* application we mean one in which the intelligent matter computes into an equilibrium state, and is inactive thereafter. Therefore static applications are most often directed toward generating some specialized material with a computationally defined nanostructure. On the other hand, *dynamic* or *interactive* applications never terminate, but always remain ready to respond to their environment in some specified way; they are the truly "smart" materials.

A. Static Applications

Programs are ideally suited to creating complex data structures, which can be converted to complex physical structures by means of universally programmable intelligent matter. Networks, chains, tubes, spheres, fibers, and quasi crystalline structures are all straightforward to compute. The network resulting from such a computation will be composed of computational groups (e.g., S, K, A) as well as inert groups, which are manipulated by the computation but do not affect it. Typically, in these applications the computational phase will be followed by a chemical phase in which the computational groups are replaced by substances appropriate to the application (a sort of "petrification"). In addition to the examples already mentioned, such an approach could be used to synthesize membranes with pores or channels of a specified size and arrangement (determined either deterministically by the program or stochastically by molecular processes).

A number of applications are suggested by the requirements of implementing small, autonomous robots. Some of these will be controlled by very dense analog neural networks, but to achieve densities comparable to mammalian cortex (15 million neurons per square cm, with up to several hundreds of thousands of connections each), we will need to be able to grow intricately branching dendritic trees at the nanoscale. Generation of such structures is straightforward with universally programmable intelligent matter (e.g., using L-systems [20]). The sensor and effector organs of microrobots will also require very fine structures, which intelligent matter can be programmed to generate.

Of course, we should not neglect the potential of universally programmable intelligent matter to do conventional computation, such as solving NP-complete problems by massively parallel computation. For example, we might replicate many copies of a program to test a potential solution, then mix them in a reaction vessel with structures representing possible solutions, and wait for equilibrium to determine actual solutions. The advantage of our approach to this kind of search problem

over others, such as DNA computation, is that our nanoscale test molecules are programmable.

B. Dynamic Applications

Dynamic intelligent matter is interactive in the sense that it is continually monitoring its environment and capable of responding according to its program. That is, it is in a state of temporary equilibrium, which can be disrupted by changes in the environment, resulting in further computation and behavior as the material seeks a new equilibrium.

For example, a membrane with channels, such as mentioned above, could be made active by having the channels open or close in response to environmental conditions, including control commands transmitted optically or chemically. The program located in each channel is simple: in response to its sensor state it executes one or the other of two effectors, blocking the channel or not. The sensor and the medium would determine whether the channel is sensitive to global conditions (e.g., overall chemical environment or ambient illumination) or to its local environment (e.g., molecules or light in its immediate vicinity).

Similarly, unanchored or free-floating molecular clusters (e.g., in colloidal suspension) may react to their environment and change their configuration, thus affecting physical properties of the substance, such as viscosity or transparency. Or they might polymerize or depolymerize on command. Unanchored supramolecular networks might also operate as semiautonomous agents to recognize molecules or molecular configurations, and act upon them in some intended way (e.g. binding toxins or pollutants). However, such applications will require the agents to operate in a medium that can supply the reactants needed for computation.

These sorts of active intelligent matter will find many applications in autonomous microrobots. For example, active membranes can serve as sensory transducers, responding to conditions in the environment and generating electrical, chemical, or other signals. They can also be programmed to self-organize into structures capable of preprocessing the input (e.g., artificial retinas or cochleas). Further, it is a simple modification of a membrane with channels to make a membrane with cilia that flex on command. By means of local communication, the cilia may be made to flex in coordinated patterns. Similarly we may fabricate artificial muscles, which contract or relax by the coordinated action of microscopic fibers. Universally programmable intelligent matter may also provide a systematic approach to self-repair of autonomous robots and other systems, since if a robot's "tissues" were created by computational processes, then they can remain potentially active, ready to restore an equilibrium disrupted by damage. Less ambitiously, materials can be programmed to signal damage or other abnormal conditions.

VI. Issues

In this section we will summarize some of the issues that need to be addressed in order to make universally programmable intelligent matter a reality, and which are therefore topics of our research.

We need a model of computation that respects the constraints of molecular processes. For example, as explained above, our model cannot assume that replication is free, or that a functional group can be linked to ("pointed at") by an unlimited number of other groups. Such a model of molecular computation might be a modification of the network-substitution model, or it might be a completely different model. (For example, we are investigating a model based on permutation of link participants.)

Assuming that we stay with something like the combinatory logic network-substitution model, then we need a solution to the replication/sharing problem. One promising solution, which we have been investigating, is to implement a "lazy" replication operation. With such an operation the two "copies" can begin to be used even before the replication is complete. (The Church-Rosser property guarantees the safety of such simultaneous use and replication.) Discarding such a partially replicated structure causes the replication process to be prematurely terminated, thus decreasing wasted resources.

We need one or more universal sets of molecular primitives. Practical experience with combinator programming has shown that it is usually more efficient to use more than the theoretically minimum set of combinators. For example, by including the identity function or I combinator $[(\mathsf{I}X) \Longrightarrow X]$, many of the self-canceling replications and deletions can be eliminated. However, we will have to keep in mind that the efficiency tradeoffs of molecular computing are not the same as those of conventional computing (e.g., substitutions take place asynchronously and in parallel, but require physical resources). Other combinators, which are not necessary from a mathematical perspective, must be included because of the structures they create. For example, the Y combinator is used to construct cyclic structures (i.e., self-sharing structures), whereas its definition in terms of S and K constructs potentially infinite branching structures, which are mathematically equivalent to cycles, but not physically equivalent.

As previously discussed, interactive applications of universally programmable intelligent matter will require sensor and effector operations. The sensors will assume two or more different configurations, or react in two or more different ways, depending on some external condition (presence of light, a chemical species, etc.). For example, a sensor could produce either of two different molecular groups, such as K (representing **true**, which selects the first alternative) and (KI) (representing **false**, which selects the second), or their equivalents. Therefore, we need to determine a general way of incorporating sensors and effectors into molecular programs. However, a more serious problem is controlling the execution time of sensors and effectors in a computational model in which substitutions can take place at any time and in parallel. There are ways of delaying substitutions in combinatory networks (for example, by "abstracting" [9], [21] a dummy variable from them), but it is not clear whether this is the best approach, or whether we should use a different model that provides more direct control over time of execution.

Certainly, one of the strengths of the combinatory logic approach to molecular computation is the Church-Rosser property, which means that substitutions can take place in any order without affecting the result of computation, but there are some

issues that must be resolved. For example, if a program has two alternative branches, selected by a conditional operation, then we may have substitutions taking place in both branches simultaneously. Even though it will not affect the result of computation, it may be wasteful to process a program structure that will not be needed. However, there is a worse problem. If programs are written in the natural recursive way (such as our Ring and Tube examples above), then it is possible that all the reaction resources could go to recursive expansion of conditional arms that will end up being discarded. There are ways to solve this problem within the combinatory logic framework. (For example, we can delay the substitutions within a network by "abstracting" [9], [21] from it an argument, which is provided only when we want the substitutions to proceed.) However, there may be more direct solutions to the problem, such as those used in ordinary programming language implementations, which can be adapted to the molecular context.

Similar problems arise in the implementation of interactive intelligent matter, that is, intelligent matter that does not compute to a stable state, but remains active. A rewriting system, such as the SK calculus, is by its nature "compute once" because the program is consumed in the process of computation. Indeed, this similarity to a chemical reaction is one of its virtues in the context of molecular computation. However, it is a problem in the context of interactive intelligent matter, since it means that a program, once executed, does not exist to execute a second time. There are at least two potential solutions to this problem. The standard solution, used in the network-substitution implementation of programming languages, is to use a combinator such as Y (mentioned above) to replicate a program structure whenever it is needed. Another solution would be a molecular process to "interpret" a fixed program structure, much as in an ordinary digital computer. However, aside from the fact that this gets quite far away from network substitution and its advantages, it has serious problems of its own, which come from having multiple functional groups simultaneously "reading" (and thus linking to) the program structure.

Although there are many similarities between term-rewriting systems and molecular processes, there are also important differences, which must be addressed in a theory of molecular computing. For example, the relative rates of reactions can be controlled by the concentrations of the reactants and other conditions, such as temperature. This can be exploited for nonstandard uses of network substitution. For example, we may have two or more conflicting sets of substitution rules for a set of molecular operators, and we can determine which are applied, or their relative probability of being applied, by controlling the concentrations of the reactants needed for the alternate rule sets. This permits probabilistic control of the nanostructures generated. Or we may have different substitutions performed in different stages of a process. Computation can also be controlled by external fields or other gradients for various purposes, such as creating oriented structures. Such considerations raise problems and potentials that are new for models of computation.

One characteristic of molecular computation which distinguishes it from electronic computation is the high probability of error in molecular processes. Therefore, we will need to develop means (both chemical and computational) for decreasing the probability of such errors, for correcting them when they occur, or for assuring that results are insensitive to them. Further, unpublished preliminary investigations indicate that a certain fraction of random SK trees will result in nonterminating, expansive, chain reactions [18]-[19]. (The probability of termination decreases with increasing random tree size.) This is a potential problem, since it suggests that a sizable fraction of substitution errors could result in runaway chain reactions that could use up all the reaction resources.

Network substitution is based on the mathematical definition of a graph: dimensionless nodes linked by edges; normally the geometrical arrangement of the nodes and length of the edges is irrelevant. However, molecular groups occupy finite volumes, and there are constraints on the locations of bonds and lengths of linking groups, which are some of the constraints that need to be accommodated in a theory of molecular computing (e.g., compare Figs. 1 and 2). Folding of program networks could also interfere with substitution operations, and so we will have to find chemical means of keeping networks extended.

Interactive applications must be provided with an adequate supply of reactants to assure that substitutions can take place when they are supposed to. Certainly, some of the reactants can come from the recycled products of previous reactions, but others will require fresh raw materials. The same considerations will apply in static applications that involve long or complex chain reactions. There are several possible solutions to this problem, but the choice depends on the specifics of the UPIM application. For example, if the networks are attached to solid substrates (e.g., membranes, sponges, or particles), then reactants may be made to flow over them (thus also clearing away waste products). If the networks are in colloidal suspension or free-floating in a fluid, then reactants are easy to add; waste products might be removed by osmosis, filtering, precipitation, etc.

These are just a few of the ways in which universally programmable intelligent matter differs from conventional computation and use of term-rewriting systems. Many problems remain to be solved, but potential of universally programmable intelligent matter makes them worth tackling.

ACKNOWLEDGEMENT

This research is supported by a Nanoscale Exploratory Research grant from the National Science Foundation. Preparation of this article has been facilitated by a grant from the University of Tennessee, Knoxville, Center for Information Technology Research. The author's research in this area was initiated when he was a Fellow of the Institute for Advanced Studies of the Collegium Budapest.

REFERENCES

[1] C. Hoffman and M. J. O'Donnell, "Pattern matching in trees," *Journal of the ACM*, vol. 29, no. 1, pp. 68–95, Jan. 1982.

[2] B. K. Rosen, "Tree manipulation systems and Church-Rosser theorems," *Journal of the ACM*, vol. 20, no. 1, pp. 160–187, January 1973.

[3] W. Fontana and L. W. Buss, " 'The arrival of the fittest': Toward a theory of biological organization," *Bulletin of Mathematical Biology*, vol. 56, no. 1, pp. 1–64, 1994.

[4] W. Fontana and L. W. Buss, "What would be conserved if 'the tape were played twice'?" *Proceedings National Academy of Science USA*, vol. 91, pp. 757–761, 1994.
[5] W. Fontana and L. W. Buss, "The barrier of objects: From dynamical systems to bounded organizations," in *Boundaries and Barriers*, J. Casti and A. Karlqvist, Eds. Reading: Addison-Wesley, 1996, pp. 56–116.
[6] J. H. Fasel and R. M. Keller (Eds.), *Graph Reduction, Proceedings of a Workshop, Santa Fe, New Mexico, USA, September 29 – October 1, 1986*, Berlin: Springer Verlag, 1987.
[7] D. A. Turner, "A new implementation technique for applicative languages," *Software — Practice and Experience*, vol. 9, pp. 31–49, 1979.
[8] A. Church and J. B. Rosser, "Some properties of conversion," *Trans. American Math. Soc.*, vol. 39, pp. 472–482, 1936.
[9] H. B. Curry, R. Feys, and W. Craig, *Combinatory Logic, Volume I*, Amsterdam: North-Holland, 1958.
[10] H. B. Curry, "Grundlagen der kombinatorischen Logik," *American Journal of Mathematics*, vol. 52, pp. 509–536, 789–834, 1930.
[11] J. R. Hindley, B. Lercher, and J. P. Seldin, *Introduction to Combinatory Logic*, Cambridge: Cambridge University Press, 1972.
[12] M. Schönfinkel, "Über die Bausteine der mathematischen Logik," *Math. Annalen*, vol. 92, pp. 305–316, 1924.
[13] B. J. MacLennan, *Functional Programming: Practice and Theory*, Reading: Addison-Wesley, 1990.
[14] B. Mayer, G. Koehler, and S. Rasmussen, "Simulation and dynamics of entropy driven, molecular self-assembly processes," *Physical Review E*, vol. 55, no. 4, pp. 4489–4499, 1997.
[15] B. Mayer and S. Rasmussen, "The lattice molecular automata (LMA): A simulation system for constructive molecular dynamics," *International Journal Modern Physics*, vol. 9, no. 1, 1998.
[16] S. Rasmussen and J. Smith, "Lattice polymer automata" *Ber. Bunsenges. Phys. Chem.*, vol. 98, no. 9, pp. 1185–1193, 1994.
[17] S. Rasmussen, C. Knudsen, and R. Feldberg, "Dynamics of programmable matter" In *Artificial Life II*, volume X of *SFI Studies in the Sciences of Complexity*, C. Langton et al., Eds. Redwood City: Addison-Wesley, 1991, pp. 211–254.
[18] B. J. MacLennan, "Preliminary investigation of random SKI-combinator trees," Dept. of Computer Science, University of Tennessee, Knoxville, Technical Report CS-97-370, 1997. Available: http://www.cs.utk.edu/~library/TechReports/1997/ut-cs-97-370.ps.Z
[19] A. YarKhan, "An investigation of random combinator soups," Dept. of Computer Science, University of Tennessee, Knoxville, unpublished technical report, 2000.
[20] A. Lindenmeyer and P. Prusinkiewicz, "Developmental models of multicellular organisms: A computer graphics perspective," *Artificial Life*, volume VI of *SFI Studies in the Sciences of Complexity*, C. G. Langton, Ed., Redwood City: Addison-Wesley, 1989, pp. 221–250,
[21] S. K. Abdali, "An abstraction algorithm for combinatory logic," *Journal of Symbolic Logic*, vol. 41, no. 1, pp. 222–224, March 1976.

Bruce J. MacLennan (M'82) was born in Teaneck, NJ in 1950. He received the BS degree with honors in mathematics from Florida State University in 1972, and the MS in 1974 and the PhD in 1975, both in computer science from Purdue University.

From 1975 to 1979 he was at Intel Corporation (Senior Software Engineer), where he worked on translator writing systems and on the 8086 and iAPX-432 microprocessors. In 1979 he joined the Computer Science faculty of the Naval Postgraduate School, where he was Assistant Professor, Associate Professor, and Acting Chair; his research focused on functional and object-oriented programming and on neural networks. Since 1987 he has been an Associate Professor in the Department of Computer Science of the University of Tennessee, Knoxville, where his research has been on neural networks, field computation, artificial life, and nanocomputation. He has published two books and more than 30 articles.

Dr. MacLennan was a Fellow of the Institute for Advanced Studies of the Collegium Budapest in 1997.