Solving Q-SAT in bounded space and time by geometrical computation

Denys Duchier, Jérôme Durand-Lose, Maxime Senot⋆

LIFO, University of Orleans, BP 6759, F-45067 ORLEANS Cedex 2, FRANCE.

Abstract. Abstract geometrical computation can solve PSPACE-complete problems efficiently: any quantified boolean formula, instance of Q-SAT – the problem of satisfiability of quantified boolean formula – can be decided in bounded space and time with simple geometrical constructions involving only drawing parallel lines on an Euclidean space-time. Complexity as the maximal length of a sequence of consecutive segments is quadratic. We use the continuity of the real line to cover all the possible boolean valuations by a recursive tree structure relying on a fractal pattern: an exponential number of cases are explored simultaneously by a massive parallelism.

Keywords. Abstract geometrical computation; Signal machine; Q-SAT; Fractal computation; Massive parallelism; Unconventional computation.

1   Introduction

When defining and studying a new model of computation, especially an unconventionnal one, these questions arise naturally: what can we compute (in terms of decidability), how can we compute it, and what does it cost (in terms of complexity)? Answers could be found by taking representative problems of classical complexity classes, e.g. SAT for NP or Q-SAT for PSPACE, and coding them in the new computation model. This was done for NP-problems with active membranes system [Păun, 2001] and with an hyperbolic space of cellular automata [Margenstern and Morita, 2001]. Similarly, some solutions for Q-SAT were proposed with P-systems and membranes [Alhazov and Pérez-Jiménez, 2007], and with closed timelike curves in relativistic computation. We showed in [Duchier et al., 2010] that *signal machines*, a geometrical and abstract model of computation, are capable of solving SAT in bounded space and time. In the present paper, we extend this result to the higher complexity class PSPACE by describing a geometrical construction solving Q-SAT through *fractal parallelization*, still in constant space and time.

We also offer a more pertinent, model-specific, notion of time-complexity, namely *collision depth*, which is quadratic for our proposed construction.

The geometrical context proposed here is the following: dimensionless *particles* move uniformly on the real axis. When a set of particles collide, they are

---

replaced by a new set of particles according to a chosen collection of *collision rules*. We consider the temporal evolution of these systems through their *space-time diagram*, in which traces of the particles are materialized by lines segment that we call *signals*. The space-time diagram of a signal machine constitued a *geometrical computation*.

Models of computation, conventional or not, are frequently based on mathematical idealizations of physical concepts and investigate the consequences, on computational power, of such abstractions (quantum, membrane, closed time-like curves, black holes...). However, oftentimes, the idealization is such that it must be interpreted either as allowing information to have infinite density (e.g. an oracle), or to be transmitted at infinite speed (global clock, no spatial extension...). On this issue, the model of signal machines stands in contradistinction with other abstract models of computation: it respects the principle of causality, density and speed of information are finite, as are the sets of objects manipulated. Nonetheless, it remains a resolutely abstract model with no apriori ambition to be physically realizable; it deals with theoretical issues such as computational power.

It is possible to do Turing-computation with such a system [Durand-Lose, 2005] and even to do analog computation by a systematic use of the continuity of space and time [Durand-Lose, 2009a,b]. Other *geometrical* models of computation exist and allow to compute: colored universes [Jacopini and Sontacchi, 1990], geometric machines [Huckenbeck, 1989], piece-wise constant derivative systems [Bournez, 1997], optical machines [Naughton and Woods, 2001]...

Most of the work to date in this domain, called *abstract geometrical computation* (AGC), has dealt with the simulation of sequential computations even though the model, seen as a continuous extension of cellular automata, is inherently parallel (see Fig. 1). In the present paper, we describe a massively parallel evaluation of all possible valuations for a given propositional formula and we provide a way to collect the results. This is the first time that parallelism is really used in AGC.
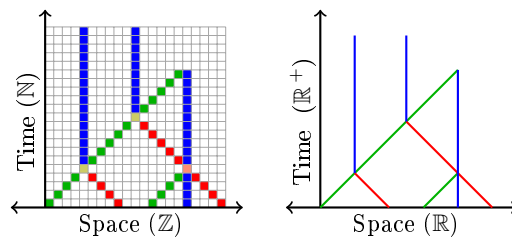

Fig. 1. From cellular automata to signal machines.

To achieve massive parallelism, we follow a fractal pattern to a depth of n (for n propositional variables) in order to partition the space in 2^n regions corresponding to the 2^n possible valuations of the unquantified formula. We call the resulting geometrical construction the *combinatorial comb* of propositional assignments. With a signal machine, such an exponential construction fits in bounded space and time regardless of the number of variables.

Once the combinatorial comb is in place, it is used to implement a binary decision tree for evaluating the formula, where all branches are explored in parallel. Finally, all the results are collected and aggregated respecting the quantifiers of the Q-SAT formula to yield the final answer. Our construction proceeds in stages: we generate and calibrate a beam of signals encoding the formula, making sure that it fits in the combinatorial comb, we propagate it through the binary decision tree, we compute the truth value when reaching each valuation, and finalize the answer at the top of the diagram.

Signal machines are presented in Section 2. Sections 3 to 7 detail step by step our geometrical solution to Q-SAT: splitting the space, coding the formula, broadcasting the formula, evaluating it and finalizing the answer by collecting the results. Complexities are discussed in Section 8 and conclusion and remarks are gathered in Section 9.

2 Definitions

Satisfiability of quantified boolean formulae. Q-SAT is the satisfiability problem for quantified boolean formulae (QBF). A QBF is a closed formula of the form: $\phi = Qx_1Qx_2\ldots Qx_n \quad \psi(x_1, x_2, \ldots, x_n)$ where $Q \in \{\exists, \forall\}$ and ψ is a quantifier-free formula of propositional logic. SAT is the fragment of Q-SAT using only the existential quantifier.

Q-SAT is *PSPACE-complete* [Stockmeyer and Meyer, 1973]: it can be solved by a polynomial-space algorithm and any PSPACE-problem can be reduced in polynomial time to Q-SAT. The classical algorithm is recursive: given a formula $Qx\ \phi(x)$, it recursively determines the satisfiability of $\phi(\texttt{true})$ and $\phi(\texttt{false})$, then aggregates the results with $\vee$ if $Q = \exists$ or with $\wedge$ if $Q = \forall$.

Signal machines. Signal machines are an extension of cellular automata from discrete time and space to continuous time and space. Dimensionless signals/particles move along the real line and rules describe what happens when they collide.

Signals. Each *signal* is an instance of a *meta-signal*. The associated meta-signal defines its *velocity* and what happen when signals meet. Figure 2 presents a very simple space-time diagram. Time is increasing upwards and the meta-signals are indicated as labels on the signals. Meta-signals are listed on the left of Fig. 2.


Meta-Signals	Speed
w	0
$\overrightarrow{\text{div}}$	3
$\overrightarrow{\text{hi}}$	1
$\overrightarrow{\text{lo}}$	3
$\overleftarrow{\text{back}}$	-3

Collision rules
{ w, $\overrightarrow{\text{div}}$ } → { w, $\overrightarrow{\text{hi}}$, $\overrightarrow{\text{lo}}$ }
{ $\overrightarrow{\text{lo}}$, w } → { $\overleftarrow{\text{back}}$, w }
{ $\overrightarrow{\text{hi}}$, $\overleftarrow{\text{back}}$ } → { w }

Fig. 2. Computing the middle

Generally, we use over-line arrows to indicate the direction of propagation of a meta-signal. For example, $\overleftarrow{a}$ and $\overrightarrow{a}$ denote two different meta-signals; but as can be expected, they have similar uses and behaviors. Similarly b_r and b_l are different; both are stationary, but one is meant to be the version for right and the other for left.

Collision rules. When a set of signals collide, they are replaced by a new set of signals according to a matching collision rule $\{\sigma_1, \ldots, \sigma_n\} \to \{\sigma'_1, \ldots, \sigma'_p\}$ where all σ_i and σ'_j are meta-signals. A rule matches a set of colliding signals if its left-hand side is equal to the set of their meta-signals. By default, if there is no exactly matching rule for a collision, the behavior is defined to regenerate exactly the same meta-signals. In such a case, the collision is called *blank*. Collision rules can be deduced from space-time diagram as on Fig. 2. They are also listed on the right of this figure.

Signal machine. A signal machine is defined by a set of meta-signals, a set of collision rules, and an initial configuration, i.e. a set of particles placed on the real line. The evolution of a signal machine can be represented geometrically as a *space-time diagram*: space is always represented horizontally, and time vertically, growing upwards. The example of Fig. 2 computes the middle: the new w is located exactly halfway between the initial two w.

3 Combinatorial comb

In order to determine by brute force whether a unquantified propositional formula with n variables is satisfiable, 2^n cases must be considered. These cases can be recursively enumerated using a binary decision tree.

The intuition is that the decision for variable x_i will be represented by a stationary signal: the space on the left should be interpreted as $x_i = \texttt{false}$, and the space on the right as $x_i = \texttt{true}$. Then we will similarly subdivide the spaces to the left and to the right, with stationary signals for x_{i+1}, and so on recursively for all variables as illustrated in Fig. 3(a).

Starting with two bounding signals w and an initiator $\overrightarrow{\text{start}}$, space is recursively divided as shown in Fig. 3(b). The first step works exactly as in Fig. 2, but then continues on to a depth of n: the counting is realized by using successively $\overrightarrow{m_0}, \overrightarrow{m_1}, \overrightarrow{m_2} \ldots$ The necessary rules and meta-signals are summarized in Tab.1.

Meta-Signal	Speed	Collision rules
$\overrightarrow{\text{start}}, \overrightarrow{\text{start}_{lo}}, \overrightarrow{a}$	3	$\{\ \overrightarrow{\text{start}}, w\ \} \to \{\ w, \overrightarrow{\text{start}_{lo}}, \overrightarrow{m_0}\ \}$
$\overrightarrow{m_0}, \overrightarrow{m_1}, \overrightarrow{m_2} \ldots$	1	$\{\ \overrightarrow{\text{start}_{lo}}, w\ \} \to \{\ \overleftarrow{a}, w\ \}$
$x_1, x_2, x_3 \ldots$	0	$\{\ w, \overleftarrow{a}\ \} \to \{\ w, \overrightarrow{a}\ \}$
$\overleftarrow{m_0}, \overleftarrow{m_1}, \overleftarrow{m_2} \ldots$	-1	$\{\ \overrightarrow{a}, w\ \} \to \{\ \overleftarrow{a}, w\ \}$
$\overleftarrow{a}$	-3	$\{\ \overrightarrow{m_i}, \overleftarrow{a}\ \} \to \{\ \overleftarrow{a}, \overleftarrow{m_{i+1}}, x_i, \overrightarrow{m_{i+1}}, \overrightarrow{a}\ \}$
b_l, b_r	0	$\{\ \overrightarrow{a}, \overleftarrow{m_i}\ \} \to \{\ \overleftarrow{a}, \overleftarrow{m_{i+1}}, x_i, \overrightarrow{m_{i+1}}, \overrightarrow{a}\ \}$
		$\{\ \overrightarrow{m_n}, \overleftarrow{a}\ \} \to \{\ b_r\ \}$
		$\{\ \overrightarrow{a}, \overleftarrow{m_n}\ \} \to \{\ b_l\ \}$

Table 1. Meta-Signals and collision rules to build the comb.

(a) Cases identification   (b) Division process

Fig. 3. Combinatorial comb for 3 variables.

Since each level of the tree is half the height of the previous one, the full tree can be constructed in bounded time regardless of its size. Also, note that the bottom level of the tree is not x_n but b_r and b_l. These are used both to evaluate the formula and to aggregate the results as explained later.

4   Formula encoding

In this section, we will explain how to represent the formula as a set of signals. This is illustrated with a running example:

$$\phi \;=\; \exists x_1 \forall x_2 \forall x_3 \quad x_1 \wedge (\neg x_2 \vee x_3) \quad .$$

We consider the quantifier-free subformula of ϕ: $x_1 \wedge (\neg x_2 \vee x_3)$ which can be viewed as a tree whose nodes are labeled by symbols (connectives and variables). The evaluation of the formula for a given assignment is a bottom-up process that percolates from the leaves toward the root. In order to model that process, we shall represent each node of the tree by a signal. In Fig. 4(a), each node is additionally decorated with a *path* from the root uniquely identifying its position in the tree: thus we are able to conveniently distinguish multiple occurrences of the same symbol. These decorated symbols provide convenient names for the required meta-signals (see Fig. 4(b)). Thus a formula of size t requires the definition of $2t$ meta-signals.

The signals for all subformulae are sent along parallel trajectories and form a *beam*. They are stacked in the diagram in order of nesting, inner-most subformulae first. This order is important for the process of percolation that will take place at the end. The width of the beam must be calibrated to have a proper propagation through the tree: it must be sufficiently narrow to fit in the top level (see [Duchier et al., 2011, App. A] for a detailled explanation and proofs).
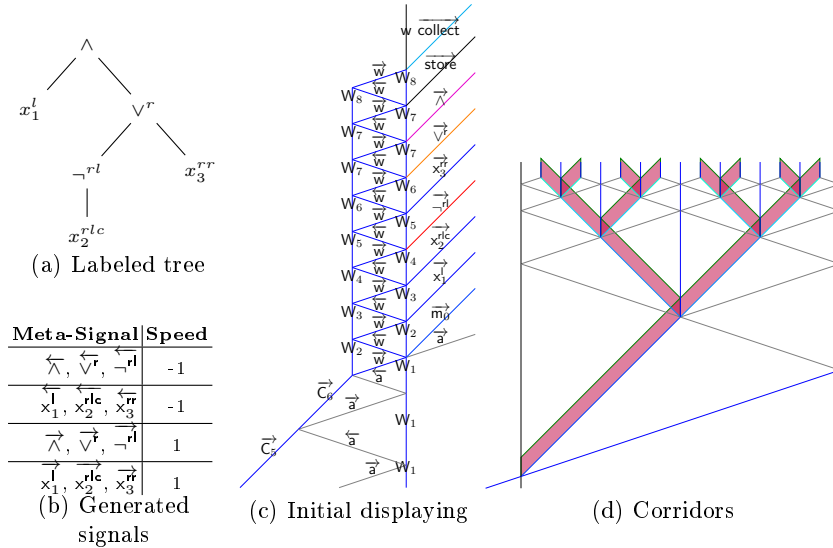
(a) Labeled tree

Meta-Signal	Speed
$\overleftarrow{\wedge}$, $\overleftarrow{\vee^r}$, $\overleftarrow{\neg^{rl}}$	-1
$\overleftarrow{x_1^l}$, $\overleftarrow{x_2^{rlc}}$, $\overleftarrow{x_3^{rr}}$	-1
$\overrightarrow{\wedge}$, $\overrightarrow{\vee^r}$, $\overrightarrow{\neg^{rl}}$	1
$\overrightarrow{x_1^l}$, $\overrightarrow{x_2^{rlc}}$, $\overrightarrow{x_3^{rr}}$	1

(b) Generated signals

(c) Initial displaying

(d) Corridors

Fig. 4. Compiling the formula

5 Propagating the beam

The formula's beam is now propagated down the decision tree. For each decision point, the beam is duplicated: one part goes through, the other is reflected. Thus, by construction, every branch of the beam tree encounters a decision point for every variable at least once. If we make the beam sufficiently narrow, the guarantee become "exactly once," as shown in Fig. 4(d).

When the beam encounters a decision point (a stationary signal for a variable x_i), then a split occurs producing two branches. Except for the sign of their velocity, most signals remain identical in both branches; most, except those corresponding to occurrences of x_i: those become `false` in the left branch and `true` in the right branch. Fig. 5(a) shows the beam intersecting the decision signal for variable x_1. Note how the incident signal $\overrightarrow{x_1^l}$ becomes $\overleftarrow{f^l}$ on the left and $\overrightarrow{t^l}$ on the right; the path decoration is preserved since, as we shall see, it is essential later for the percolation process. This is achieved by the collision rule: $\{\overrightarrow{x_1^l}, x_1\} \rightarrow \{\overleftarrow{f^l}, x_1, \overrightarrow{t^l}\}$. Since a decision point is encountered exactly once for each variable on each branch of the lane, at the bottom of the tree, all signals corresponding to occurrences of variables have been assigned a boolean value.

6 Evaluating the formula

Remember how, at the very bottom of the decision tree, we added an extra division using signals b_l or b_r: their purpose is to initiate the percolation process.

b_l is for starting the percolation process of a left branch, while b_r is for a right branch. Figure 5(c) zooms on one case of our example.

$$\{ \overrightarrow{\vee^r}, \overleftarrow{T^{rl}} \} \rightarrow \{ \overrightarrow{t()^r} \}$$
$$\{ \overrightarrow{t()^r}, \overleftarrow{T^{rr}} \} \rightarrow \{ \overrightarrow{t^r} \}$$
$$\{ \overrightarrow{id^r}, \overleftarrow{T^{rr}} \} \rightarrow \{ \overrightarrow{t^r} \}$$
$$\{ \overrightarrow{\vee^r}, \overleftarrow{F^{rl}} \} \rightarrow \{ \overrightarrow{id^r} \}$$
$$\{ \overrightarrow{t()^r}, \overleftarrow{F^{rr}} \} \rightarrow \{ \overrightarrow{t^r} \}$$
$$\{ \overrightarrow{id^r}, \overleftarrow{F^{rr}} \} \rightarrow \{ \overrightarrow{f^r} \}$$

(a) Split (b) Collision rules to evaluate the disjunction $\vee^r$ (c) Evaluation at the bottom of the comb
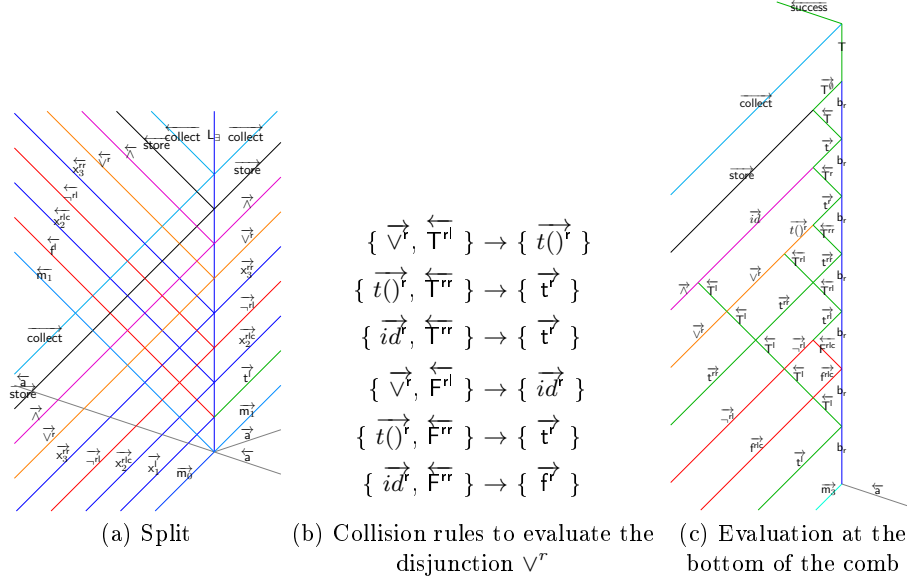
Fig. 5. Split, evaluation process and rules for a connective.

The invariant is that all signals that reach b_r have determined boolean values. When $\overrightarrow{t^l}$ reaches b_r, it gets reflected as $\overleftarrow{T^l}$. The change from lowercase to uppercase indicates that the subformula's signal is now able to interact with the signal of its parent connective. The stacking order ensures that reflected signals of subformulae will interact with the incoming signal of their parent connective before the latter reaches b_r. This enforces the invariant.

A connective is evaluated by colliding with the (uppercased) boolean signals of its arguments. For example, the disjunction collides with its first argument. Depending on its value, it becomes the one-argument function identity or the constant `true`. This is the way the rules of Fig. 5(b) should be understood.

Note how the path decorations are essential to ensure that the right subformulae interact with the right occurrences of connectives. Conjunctions and negations can be handled similarly. Finally, $\overrightarrow{store}$ projects the truth value of the formula's root on b_r where it is temporarilly stored until $\overrightarrow{collect}$ starts the aggregation of the results.

7 Collecting the results

At the end of the propagation phase, the results of evaluating the unquantified formula for all possible assignments have been stored as stationary signals replacing the b_l and b_r signals. We must now evaluate the quantified formula.

Remember that each level of the comb corresponds to a variable: level 1 stands for $x_1, \ldots$ For aggregating all the results, we will "undo" the construction process of the binary tree by mixing two by two the results of evaluations with respect to the initial quantifiers.

At each split of the beam – *i.e.* when it meets the stationary signals coding x_i at the i^{th} level of the tree – the signal $\overrightarrow{\text{collect}}$ (resp. $\overleftarrow{\text{collect}}$) (at the very top of the beam) changes the stationary x_i into a stationary $L^i_{Q_i}$ (resp. $R^i_{Q_i}$), where Q_i denotes the quantifier for x_i in ϕ, and L (resp. R) indicates the direction in which to emit the combined result of trying $x_i = \texttt{false}$ (coming from the left) and $x_i = \texttt{true}$ (coming from the right). The boolean connective to effect the combination depends on the type of the quantifier: $\vee$ for $\exists$ and $\wedge$ for $\forall$. This collection process is illustrated in Fig. 6.

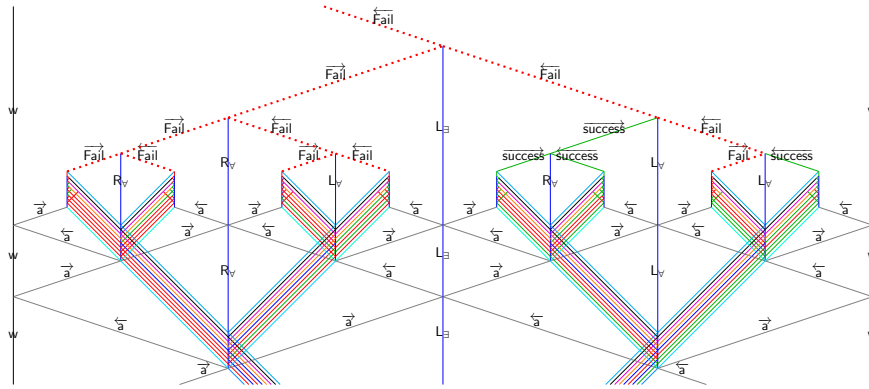The space-time diagram for the entire construction is shown in Fig. 7.

Fig. 6. Collecting the results.

8 Complexities

We now turn to a crucial question: what is the complexity of our construction as a function of the size of the formula? What is a meaningful way to measure this complexity?

The width of the construction measures the space requirement: it is independent of the formula and can be fixed to any value we like. The height measures the time requirement: it is also independent of the formula because of the fractal construction and the continuity of space-time. If more variables are involved, the comb gains extra levels, but its height remains bounded by a fractal, the infinite binary tree.

As a consequence, while width (space) and height (time) are the natural continuous extensions of traditional complexity measures used in the discrete universe of cellular automata, in the context of abstract geometrical computations, they loose all pertinence.

Instead we should regard our construction as a computational device transforming inputs into outputs. The inputs are given by the initial state of the signal machine at the bottom of the diagram. The output is the computed result that comes out at the top. The transformation is performed in parallel by many threads: a thread here is an ascending path through the diagram from an input to the output. The operations that are "performed" by the thread are all the collisions found along the path.

Thus, if we view the diagram as an acyclic graph of collisions (vertices) and signals (arcs), time complexity can then be defined as the maximal size of a chain of collisions and space complexity as the maximal size of an anti-chain *i.e.* a set of signals pairwise un-related.

Let t be the size of the formula and n the number of variables. At the bottom level of the comb, there is an anti-chain of length approximately $t2^n$, making the space complexity exponential. Generation of the comb, initiation, propagation, evaluation and aggregation contribute along any path a number of collisions at most linear in the size of the formula. However, intersections of incident and reflected branches at every level add $O(nt)$ because there are $O(n)$ levels and the beam consists of $O(t)$ parallel signals. Thus the time complexity is $O(nt)$.

Fig. 7. The whole diagram.

It should also be pointed out that the compilation into a rational signal machine is done in polynomial time and only five distinct speeds are involved. Algorithms to generate the machine from the formula and all the collision rules can be found in [Duchier et al., 2011, App. B].

9 Conclusion

In this article, we have shown how to achieve massive parallelism with signal machines, by means of a fractal pattern. We call this *fractal parallelism* and it is a novel contribution in the field of abstract geometrical computation. We have described how this approach is able to solve Q-SAT in bounded space and time.

The complexity is not hidden inside the compilation of the machine (done in quadratic time) nor in the initial configuration (which is very simple) and all the speeds are constant. Since, clearly, time and space are no longer appropriate measures of complexity, we have also proposed to replace them respectively by the maximum size of a chain and an anti-chain in the space-time diagram regarded as a directed acyclic graph. According to these new definitions, our construction has exponential space complexity and quadratic time complexity.

It seems natural to draw a connexion between the construction we propose and Boolean circuits, and we plan to investigate the relation. However, signal machines address an additional concern, which Boolean circuits do not; namely, that computations take place not only over time, but also over physical space, and that, in this respect, they are potentially limited by any bound placed on the speed at which information may travel.

Any formula can be compiled into a signal machine. The next step is to provide a single signal machine that is generic for Q-SAT —i.e. it can solve any instance of Q-SAT— so that the formula would only have to be compiled into an initial configuration. Additionally, we will continue our investigations with complexity classes higher in the hierarchy, such as EXPTIME or EXPSPACE.

Bibliography

A. Alhazov and M. J. Pérez-Jiménez. Uniform solution of Q-SAT using polarization-less active membranes. In J. Durand-Lose and M. Margenstern, editors, *Machines, Computations, and Universality, 5th International Conference (MCU '07)*, number 4664 in LNCS, pages 122–133. Springer, 2007.

O. Bournez. Some bounds on the computational power of piecewise constant derivative systems. In *ICALP '97*, number 1256 in LNCS, pages 143–153, 1997.

D. Duchier, J. Durand-Lose, and M. Senot. Fractal parallelism: Solving SAT in bounded space and time. In O. Cheong, K.-Y. Chwa, and K. Park, editors, *21st International Symposium on Algorithms and Computation (ISAAC '10)*, number 6506 (Part 1) in LNCS, pages 279–290. Springer, 2010.

D. Duchier, J. Durand-Lose, and M. Senot. Solving Q-SAT in bounded space and time by geometrical computation (extended version). Research Report RR-2011-03, LIFO, Université d'Orléans, 2011. URL http://www.univ-orleans.fr/lifo/rapports.php.

J. Durand-Lose. Abstract geometrical computation: Turing computing ability and undecidability. In B. Cooper, B. Löwe, and L. Torenvliet, editors, *New Computational Paradigms, 1st Conf. Computability in Europe (CiE '05)*, number 3526 in LNCS, pages 106–116. Springer, 2005.

J. Durand-Lose. Abstract geometrical computation 3: Black holes for classical and analog computing. *Nat. Comput.*, 8(3):455–572, 2009a.

J. Durand-Lose. Abstract geometrical computation and computable analysis. In J. Costa and N. Dershowitz, editors, *International Conference on Unconventional Computation 2009 (UC '09)*, number 5715 in LNCS, pages 158–167. Springer, 2009b.

U. Huckenbeck. Euclidian geometry in terms of automata theory. *Theoret. Comp. Sci.*, 68(1):71–87, 1989.

G. Jacopini and G. Sontacchi. Reversible parallel computation: an evolving space-model. *Theoret. Comp. Sci.*, 73(1):1–46, 1990.

M. Margenstern and K. Morita. NP problems are tractable in the space of cellular automata in the hyperbolic plane. *Theoret. Comp. Sci.*, 259(1-2):99–128, 2001.

T. J. Naughton and D. Woods. On the computational power of a continuous-space optical model of computation. In M. Margenstern, editor, *Machines, Computations, and Universality (MCU '01)*, number 2055 in LNCS, pages 288–299, 2001.

G. Păun. P systems with active membranes: Attacking NP-Complete problems. *Journal of Automata, Languages and Combinatorics*, 6(1):75–90, 2001.

L. J. Stockmeyer and A. R. Meyer. Word problems requiring exponential time. In *5th ACM Symposium on Theory of Computing (STOC '73)*, volume 16, pages 1–9, 1973.