



## Computing in the fractal cloud: modular generic solvers for SAT and Q-SAT variants.

Denys Duchier, Jérôme Durand-Lose, Maxime Senot

### ► To cite this version:

Denys Duchier, Jérôme Durand-Lose, Maxime Senot. Computing in the fractal cloud: modular generic solvers for SAT and Q-SAT variants.. Theory and Applications of Models of Computation (TAMC 2012), May 2012, Beijing, China. pp.435-447. hal-00673603

HAL Id: hal-00673603

<https://hal.science/hal-00673603>

Submitted on 8 Jun 2012

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Computing in the fractal cloud: modular generic solvers for SAT and Q-SAT variants

Denys DUCHIER, Jérôme DURAND-LOSE, and Maxime SENOT \*

LIFO, Université d'Orléans,  
B.P. 6759, F-45067 ORLÉANS Cedex 2.  
`{denys.duchier, jerome.durand-lose, maxime.senot}@univ-orleans.fr`

**Abstract.** Abstract geometrical computation can solve hard combinatorial problems efficiently: we showed previously how Q-SAT —the satisfiability problem of quantified boolean formulae— can be solved in bounded space and time using instance-specific signal machines and fractal parallelization. In this article, we propose an approach for constructing a particular *generic* machine for the same task. This machine deploys the Map/Reduce paradigm over a discrete fractal structure. Moreover our approach is *modular*: the machine is constructed by combining modules. In this manner, we can easily create generic machines for solving satisfiability variants, such as SAT, #SAT, MAX-SAT.

**Keywords.** Abstract geometrical computation; Signal machine; Fractal; Satisfiability problems; Massive parallelism; Model of computation.

## 1 Introduction

Since their first formulations in the seventies, problems of Boolean satisfiability have been studied extensively in the field of computational complexity. Indeed, the most important complexity classes can be characterized —in terms of reducibility and completeness— by such problems e.g. SAT for NP [4] and Q-SAT for PSPACE [21]. As such, it is a natural challenge to consider how to solve these problems when investigating new computing machinery: quantum, NDA and membrane [20], optical [13], hyperbolic spaces [18], etc. (for an overview of the status of NP-complete problems under several physical assumptions, see [1]).

This is the line of investigation that we have been following with *signal machines*, an abstract and geometrical model of computation. We showed previously how such machines were able to solve SAT [6] and Q-SAT [7] in bounded space and time and in quadratic *collision depth*, a model-specific time complexity measure defined by the maximal number of consecutive collisions, which is better

---

\* This work was partially supported by the ANR project AGAPE, ANR-09-BLAN-0159-03.

suited to the strong parallelism of signal machines. But in both cases, the machines were instance-specific *i.e.* depended on the formula whose satisfiability was to be determined. The primary contribution of the present paper is to exhibit a particular *generic* signal machine for the same task: it takes the instance formula as an input encoded (by a polynomial-time Turing machine) in an initial configuration. This single machine replaces the whole family of machines designed previously (one machine for each formula) and formulae are now represented by inputs (as for classical algorithms) instead of being encoded by signals and rules of machines. We further improve our previous results by describing a *modular* approach that allows us to easily construct generic machines for other variants of SAT, such as #SAT or MAX-SAT. We also introduce a new technical gadget, the *lens device*, which automatically scales by half any beam of signals. These constructions are in cubic collision depth instead of quadratic for the previous instance-specific solutions.

The model of signal machines, called *abstract geometrical computation*, involves two types of fundamental objects: dimensionless *particles* and *collision rules*. We use here one-dimensional machines: the space is the Euclidean real line, on which the particles move with a constant speed. Collision rules describe what happens when several particles collide. By representing continuous time on a vertical axis, we obtain two-dimensional *space-time diagrams*, in which the motion of the particles are represented by lines segment called *signals*. Signal machines can simulate Turing machines and are thus Turing-universal [11]. Under some assumptions and by using the continuity of space and time, signal machines can simulate analog models such as computable analysis [10] and they can even be super-Turing by embedding the black hole model (forecasting a black-hole in signal machines is highly undecidable as shown in [9]).

Other geometrical models of computation exist: colored universes [15], geometric machines [14], optical machines [19], interaction nets [17], piecewise constant derivative systems [2], tilings [16], etc.

All these models, including signal machines, belong to a larger class of models of computation, called *unconventional*, which are more powerful than classical ones (Turing machines, RAM, **while**-programs...). Among all these abstract models, the model of signal machines distinguishes itself by realistic assumptions respecting the major principles of physics —finite density of information, respect of causality and bounded speed of information— which are, in general, not respected all at the same time by other models. Nevertheless, signal machines remain an abstract model, with no a priori ambition to be physically realizable, and is studied for theoretical issues of computer sciences such as computability power and complexity measures.

As signal machines take their origins in the world of cellular automata, not only can they be a powerful tool for understanding the latter's complex behaviors, implementing computations [12] and proving universality [3] but they too can be viewed as a massively parallel computational device. This is the approach proposed here: we put in place a fractal compute grid, then use the Map/Reduce paradigm to distribute the computations, then aggregate the results.

The Map/Reduce pattern, pioneered by Lisp, is now standard in functional programming: a function is applied to many inputs (map), then the results are aggregated (reduce). Google extended this pattern to allow its distributed computation over a grid of possibly a thousand nodes [5]. The idea is to partition the input (petabytes of data) into chunks, and to process these chunks in parallel on the available nodes. When solving combinatorial problems, we are also faced with massive inputs; namely, the exponential number of candidate solutions. Our approach is to distribute the candidates, and thus the computation, over an unbounded fractal grid. In this way, we adapt the map/reduce pattern for use over a grid with fractal geometry.

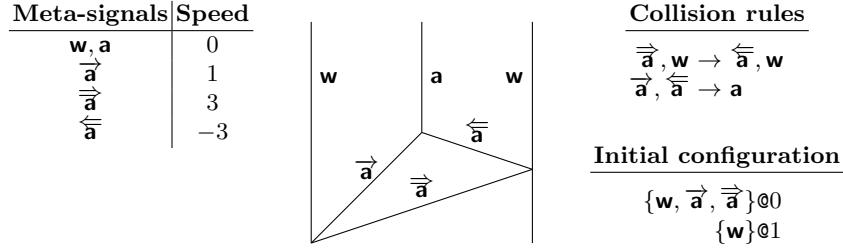
Our contribution in this paper is three fold: first, we show how Q-SAT can be solved in bounded space and time using a *generic machine*, where the input (the formula) is simply compiled into an initial configuration. This improves on our previous result where the machine itself depended on the formula. Second, we propose the first architecture for fractally distributed computing (the *fractal cloud*) and give a way to automatically shrink the data into this structure by means of a *lens device*. Third, we show how generic machines for many variants of SAT can be assembled by composing independent modules, which naturally emerged from the generalization of our previous family of machines into a single machine solving Q-SAT. Each module can be programmed and understood independently. We also discuss notions and choices of complexity measures which strongly depend of the considered model of computation, and we argue that *collision depth*, a time complexity measure introduced in [6], is more relevant to signal machines. The collision depth of the given construction is cubic in the size of the input formula and space complexity is exponential.

The paper is structured as follow. Signal machines are introduced in Section 2. Section 3 presents the fractal tree structure used to achieve massive parallelism and how general computations can be inserted in the tree. Section 4 details this implementation for a Q-SAT solver and Section 5 explains how some variants of satisfiability problems can be solved with the same approach. Complexities are discussed in Section 6 and conclusions and remarks are gathered in Section 7.

## 2 Definitions

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 *speed*. Figure 1 presents a very simple space-time diagram. Time is increasing upwards and the meta-signals are indicated as labels on the signals. Generally, we use over-line arrows to indicate the direction and speed of propagation of a meta-signal. For example,  $\overrightarrow{\mathbf{a}}$  and  $\overleftarrow{\mathbf{a}}$  denote two different meta-signals: the first travels to the right at speed 1, while the other travels to the left at speed  $-3$ .  $\mathbf{w}$  and  $\mathbf{a}$  are both stationary meta-signals.



**Fig. 1.** Geometrical algorithm for computing the middle

*Collision rules.* When a set of signals collide i.e. when they are at the same spatial location at the same time, they are replaced by a new set of signals according to a matching collision rule. A rule has the form:  $\sigma_1, \dots, \sigma_n \rightarrow \sigma'_1, \dots, \sigma'_p$  where all  $\sigma_i$  are meta-signals of distinct speeds as well as  $\sigma'_j$  (two signals cannot collide if they have the same speed and outcoming signals must have different speeds). 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 diagrams as on Fig. 1 where they are also listed on the right.

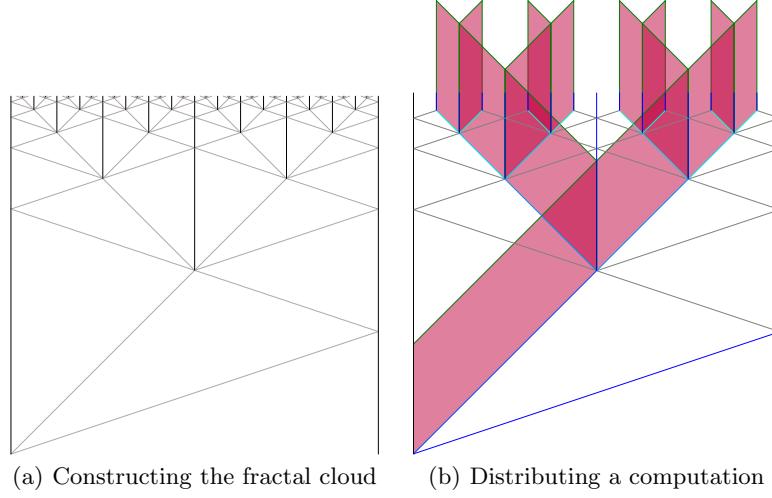
**Definition 1.** A signal machine  $\mathcal{M}$  is a triplet  $\mathcal{M} = (M, S, C)$  where  $M$  is a finite set of meta-signals,  $S : M \rightarrow \mathbb{R}$  is the speed function which assigns a real speed to each meta-signal and  $C$  is the set of collision rules.

An initial configuration  $c_0$  is a finite set  $c_0 = \{(\sigma_i, x_i) \mid \sigma \in M \text{ and } x \in \mathbb{R}\}$ . For a signal  $\sigma_i$  at initial position  $x_i$ , we also note  $\sigma_i @ x_i$ .

A signal machine is executed starting from an initial configuration which corresponds to the input. 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 geometrical algorithm displayed in Fig. 1 computes the middle: the new  $a$  is located exactly halfway between the initial two  $w$ . Constructions given in the present paper are achieved by a *rational signal machine*: all speeds and initial positions are rational values (in particular, speeds take just a few integer values: -3, -1, 0, 1 and 3). It follows that every collision happens at rational coordinates.

### 3 Computing in the fractal cloud

*Constructing the fractal.* The fractal structure that interests us is based on the simple idea of computing the middle illustrated in Fig. 1. We just indefinitely repeat this geometrical construction: once space has been halved, we recursively halve the two halves, and so on.



**Fig. 2.** Computing in the fractal cloud

This is illustrated in Fig. 2(a) and can be generated by the following rules<sup>\*</sup>:

$$\mathbf{w}, \overleftarrow{\mathbf{a}} \rightarrow \mathbf{w}, \overrightarrow{\mathbf{a}} \quad \overrightarrow{\mathbf{a}}, \overleftarrow{\mathbf{a}} \rightarrow \overleftarrow{\mathbf{a}}, \overleftarrow{\mathbf{a}}, \mathbf{a}, \overrightarrow{\mathbf{a}}, \overrightarrow{\mathbf{a}}$$

using  $\{\mathbf{w}, \overrightarrow{\mathbf{a}}, \overleftarrow{\mathbf{a}}\} @ 0$  and  $\{\mathbf{w}\} @ 1$  as the initial configuration. This produces a stack of levels: each level is half the height of the previous one. As a consequence, the full fractal has width 1 and height 1.

*Distributing a computation.* The point of the fractal is to recursively halve space. At each point where space is halved, we position a stationary signal (a vertical line in the space-time diagram). We can use this structure, so that, at each halving point (stationary signal), we split the computation in two: send it to the left with half the data, and also to the right with the other half of the data.

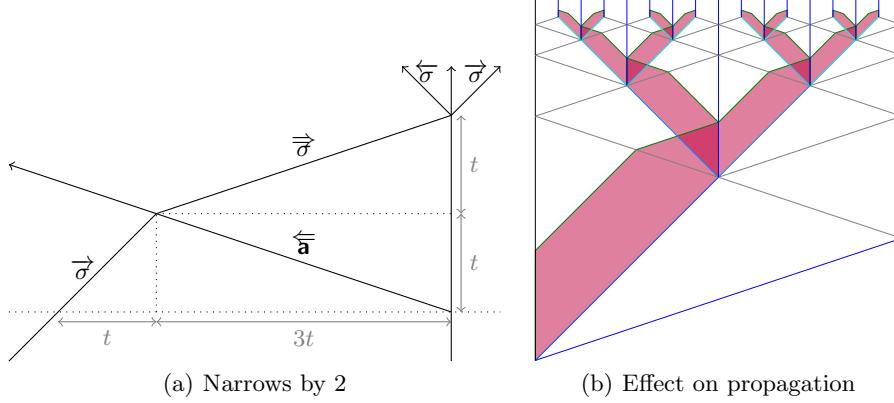
The intuition is that the computation is represented by a *beam* of signals, and that stationary signals split this beam in two, resulting in one beam that goes through, and one beam that is reflected.

Unfortunately, a beam of constant width will not do: eventually it becomes too large for the height of the level. This can be clearly seen in Fig. 2(b).

*The lens device.* The lens device narrows the beam by a factor of 2 at each level, thus automatically adjusting it to fit the fractal (see Fig. 3). It is implemented by the following meta-rule: *unless otherwise specified, any signal  $\vec{s}$  is accelerated by  $\overleftarrow{\mathbf{a}}$  and decelerated and split by any stationary signal  $\mathbf{s}$ .*

---

\* For brevity, we will always omit the rules which can be obtained from the others by symmetry. We refer to the extended version [8] for more details.



**Fig. 3.** The lens device

*Generic computing over the fractal cloud.* With the lens device in effect, generic computations can take place over the fractal by propagating a beam from an initial configuration. We write  $[(\vec{\sigma}_n \dots \vec{\sigma}_1) \text{spawn}]$  for an initial configuration with a sequence  $\vec{\sigma}_n \dots \vec{\sigma}_1$  of signals disposed from left to right on the space line. Geometrically, it can easily be seen that, in order for the beam to fit through the first level, the sequence  $\vec{\sigma}_n \dots \vec{\sigma}_1$  must be placed in the interval  $(-\frac{1}{4}, 0)$ .

*Modules.* A *module* is a set of signals which correspond to a given task. We describe a module by the parametric abstraction defining its instance-specific contribution to the initial configuration in the form  $[\text{module}e] = \vec{\sigma}_n \dots \vec{\sigma}_1$  and by the generic (*i.e.* instance-independent) collision rules describing how this module interacts with other modules.

*Stopping the fractal.* For finite computations, we don't need the entire fractal. The  $[\text{until}(n)]$  module can be inserted in the initial configuration to cut the fractal after  $n$  levels have been generated. We set:  $[\text{until}(n)] = \vec{z} \vec{\zeta}^{n-1}$ .

$$\begin{array}{lll}
 \vec{\zeta}, \overleftarrow{a} \rightarrow \overleftarrow{a}_o, \vec{\zeta} & \vec{\zeta}, a \rightarrow a_o & \vec{\zeta}, a_o \rightarrow \overleftarrow{\zeta}, a_o, \vec{\zeta} \\
 \vec{z}, \overleftarrow{a} \rightarrow \overleftarrow{a}_o, \vec{z} & \vec{z}, \overleftarrow{a}_o \rightarrow \overleftarrow{a}, \vec{z} & \vec{z}, a_o \rightarrow \overleftarrow{z}, a, \vec{z} \\
 \vec{z}, a \rightarrow a, \vec{z} & \vec{z}, \overrightarrow{a} \rightarrow \overrightarrow{a}_o & \overrightarrow{a}, \overleftarrow{a}_o \rightarrow \emptyset
 \end{array}$$

**Table 1.** Stopping the fractal.

The subbeam  $\vec{\zeta}^{n-1}$  are the inhibitors for  $\vec{z}$ . One inhibitor is consumed at each level, after which  $\vec{z}$  takes effect and turns  $\overleftarrow{a}$  into  $\overleftarrow{a}_o$  which finally annihilates the constructing signals via the rule  $\overrightarrow{a}, \overleftarrow{a}_o \rightarrow \emptyset$ , bringing the fractal to a stop. Thus, a computation  $[(\vec{\sigma}_n \dots \vec{\sigma}_1 [\text{until}(n)]) \text{spawn}]$  uses only  $n$  levels. It can be seen geometrically that, for the collision of  $\vec{z}$  with  $\overleftarrow{a}$  to occur before the latter meets with  $\overleftarrow{a}_o$ ,  $\vec{z}$  must initially be placed in  $(-\frac{1}{6}, 0)$ .

## 4 A modular Q-SAT solver

Q-SAT is the satisfiability problem for quantified Boolean formulae (QBF). A QBF is a closed formula of the form  $\phi = Q_1x_1Q_1x_2\dots Q_nx_n \psi(x_1, x_2, \dots, x_n)$  where  $Q_i \in \{\exists, \forall\}$  and  $\psi$  is a quantifier-free formula of propositional logic. A classical recursive algorithm for solving Q-SAT is:

$$\begin{aligned}\text{qsat}(\exists x \phi) &= \text{qsat}(\phi[x \leftarrow \text{false}]) \vee \text{qsat}(\phi[x \leftarrow \text{true}]) \\ \text{qsat}(\forall x \phi) &= \text{qsat}(\phi[x \leftarrow \text{false}]) \wedge \text{qsat}(\phi[x \leftarrow \text{true}]) \\ \text{qsat}(\beta) &= \text{eval}(\beta)\end{aligned}$$

where  $\beta$  is a ground Boolean formula. This is exactly the structure of our construction: each quantified variable splits the computation in 2,  $\text{qsat}(\phi[x \leftarrow \text{false}])$  is sent to the left and  $\text{qsat}(\phi[x \leftarrow \text{true}])$  to the right, and subsequently the recursively computed results that come back are combined (with  $\vee$  for  $\exists$  and  $\wedge$  for  $\forall$ ) to yield the result for the quantified formula. This process can be viewed as an instance of *Map/Reduce*, where the *Map* phase distributes the combinatorial exploration of all possible valuations across space using a binary decision tree, and the *Reduce* phase collects the results and aggregates them using quantifier-appropriate Boolean operations. Our Q-SAT solver is modularly composed as follows (modules `decide`, `map:sat`, and `reduce:qsat` are described below):

```
[([reduce:qsat( $Q_1x_1\dots Q_nx_n$ )] [map:sat( $\psi$ )  
[decide( $n$ )] [until( $n+1$ )]) spawn]
```

### 4.1 Setting up the decision tree

For a QBF with  $n$  variables, we need 1 level per variable, and then at level  $n+1$  we have a ground propositional formula that needs to be evaluated. Thus, the first module we insert is `[until( $n+1$ )`] to create  $n+1$  levels. We then insert `[decide( $n$ )`] because we want to use the first  $n$  levels as decision points for each variable. This is simply achieved by taking `[decide( $n$ )]` =  $\vec{\alpha}^n$  (one signal  $\vec{\alpha}$  per level) with the following rules:  $\vec{\alpha}, \mathbf{a} \rightarrow \mathbf{x}$  (turning stationary signal into assigning ones) and  $\vec{\alpha}, \mathbf{x} \rightarrow \vec{\alpha}, \mathbf{x}, \vec{\alpha}$  (splitting the remaining  $\vec{\alpha}$  for next levels).

### 4.2 Compiling the formula

The intuition is that we want to compile the formula into a form of inverse polish notation to obtain executable code using postfix operators. At level  $n+1$  all variables have been decided, and have become  $\vec{\mathbf{t}}$  or  $\vec{\mathbf{f}}$ . The ground formula, regarded as an expression tree, can be executed bottom up to compute its truth value: the resulting signal for a subexpression is sent to interact with its parent operator. The formula is represented by a beam of signals: each subformula is represented by a (contiguous) subbeam. A subformula that arrives at level  $n+1$  starts evaluating when it hits the stationary  $\mathbf{a}$ . When its truth value has been computed, it is reflected so that it may eventually collide with the incoming signal of its parent connective.

*Compilation.* For binary connectives, one argument arrives first, it is evaluated, and its truth value is reflected toward the incoming connective; but, in order to reach it, it must cross the incoming beam for the other argument and not interact with the connectives contained therein. For this reason, with each subexpression, we associate a beam  $\vec{\gamma}^k$  of inhibitors that prevents its resulting truth value from interacting with the first  $k$  connectives that it crosses. We write  $C[\psi]$  for the compilation of  $\psi$  into a contribution to the initial configuration, and  $\|\psi\|$  for the number of occurrences of connectives in  $\psi$ . The following scheme of compilation produces an initial configuration which has a size—the number of signals—at most quadratic in the size  $s$  of the input formula. Clearly, for each node (*i.e.* symbol) of the formula, only a linear number of inhibitor signals  $\vec{\gamma}$  can be added, so  $C[\psi]$  is composed by at most  $\mathcal{O}(s \cdot s) = \mathcal{O}(s^2)$  signals. The compilation is done by induction on the formula:

$$\begin{aligned} C[\psi] &= C[\psi]^0 \\ C[\psi_1 \wedge \psi_2]^k &= \vec{\wedge} \vec{\gamma}^k C[\psi_1]^0 C[\psi_2]^{\|\psi_1\|} \\ C[\psi_1 \vee \psi_2]^k &= \vec{\vee} \vec{\gamma}^k C[\psi_1]^0 C[\psi_2]^{\|\psi_1\|} \\ C[\neg\psi]^k &= \vec{\neg} \vec{\gamma}^k C[\psi] \\ C[x_i]^k &= [\text{var}(x_i)] \vec{\gamma}^k \end{aligned}$$

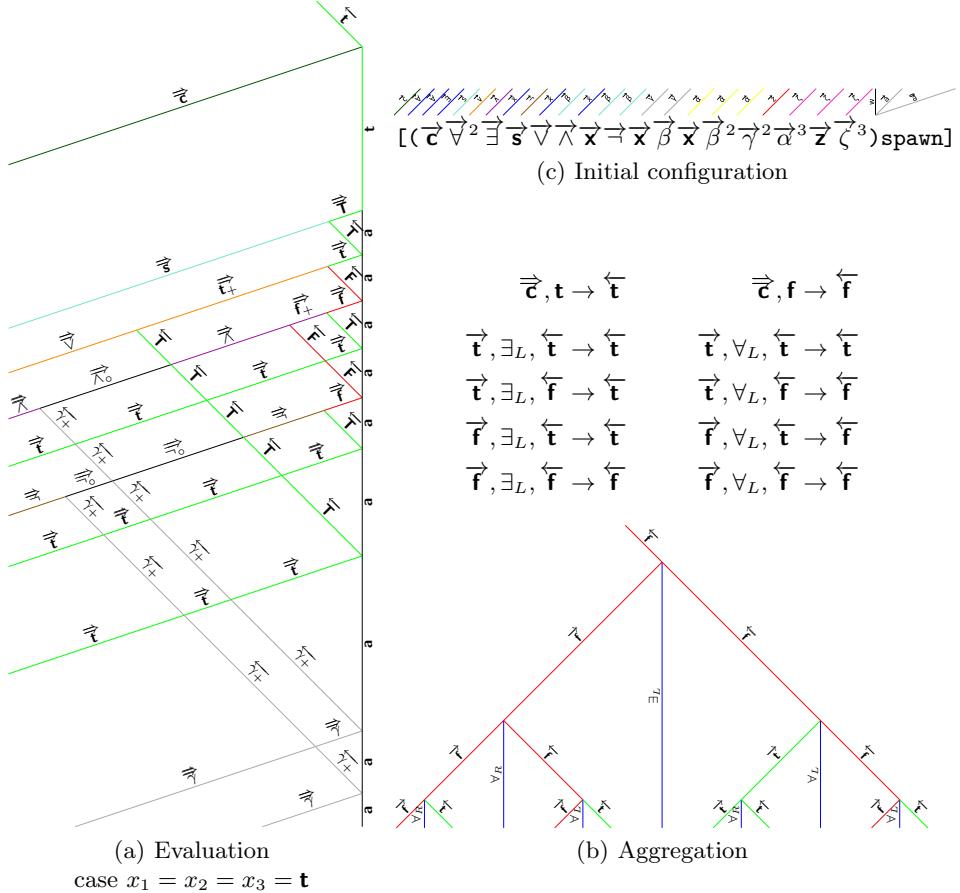
*Variables.* We want variable  $x_i$  to be decided at level  $i$ . This can be achieved using  $i-1$  inhibitors. For variable  $x_i$ , the idea is to protect  $\vec{x}$  from being assigned into  $\vec{f}$  and  $\vec{t}$  until it reaches the  $i^{th}$  level. This is achieved with a stack of  $i-1$  signals  $\vec{\beta}$ : at each level, the first  $\vec{\beta}$  turns the stationary signal  $\mathbf{x}$  into  $\mathbf{x}_o$  (the non-assigning version of  $\mathbf{x}$ ) and disappears. The following  $\vec{\beta}$  and  $\vec{x}$  are simply split,  $\vec{x}$  taking  $\mathbf{x}_o$  back into  $\mathbf{x}$ . After the first  $i-1$  levels, all the  $\vec{\beta}$  have been consumed so that  $\vec{x}$  finally collides directly with  $\mathbf{x}$  and splits into  $\vec{f}$  going left and  $\vec{t}$  going right. The variable  $x_i$  is initially coded by:  $[\text{var}(x_i)] = \vec{x} \vec{\beta}^{i-1}$ .

$\vec{\beta}, \mathbf{x} \rightarrow \mathbf{x}_o$ $\vec{x}, \mathbf{x} \rightarrow \vec{f}, \mathbf{x}, \vec{t}$	$\vec{\beta}, \mathbf{x}_o \rightarrow \vec{\beta}, \mathbf{x}_o, \vec{\beta}$ $\vec{x}, \mathbf{x}_o \rightarrow \vec{x}, \mathbf{x}, \vec{x}$
--	--

**Table 2.** Coding and assigning variables.

*Evaluation.* When hitting  $\mathbf{a}$  at level  $n+1$ ,  $\vec{t}$  is reflected as  $\vec{T}$ , and  $\vec{f}$  as  $\vec{F}$ : these are their *activated* versions which can interact with incoming connectives to compute the truth value of the formula according to the rules given in Tab. 3 for  $\wedge$  (other connectives are similar, *cf* [8]). See Fig. 4(a) for an example.

*Storing the results.* In order to make the result easily exploitable by the *Reduce* phase, we now store it with a signal  $\vec{s}$  as the stationary signal at level



**Fig. 4.** Example  $\exists x_1 \forall x_2 \exists x_3 (x_1 \wedge \neg x_2) \vee x_3$

$n + 1$ ; it replaces **a**, which becomes a signal **t** or **f**. The complete *Map* phase is implemented by:  $[\text{map:sat}(\psi)] = \vec{s} C[\psi]$ .

### 4.3 Aggregating the results

As explained earlier, the results for an existentially (resp. universally) quantified variable must be combined using  $\vee$  (resp.  $\wedge$ ).

*Setting up the quantifiers.* We turn the decision points of the first  $n$  levels into quantifier signals. Moreover, at each level, we must also take note of the direction in which the aggregated result must be sent. Thus  $\exists_L$  represents an existential quantifier that must send its result to the left. Rules are given by table Tab. 4. For this, we set:  $[\text{reduce:qsat:init}(Q_1 x_1 \cdots Q_n x_n)] = \vec{Q}_n \dots \vec{Q}_1$ .

$$\begin{array}{lll}
\overrightarrow{\mathbf{t}}, \mathbf{a} \rightarrow \overleftarrow{\mathbf{T}}, \mathbf{a} & \overrightarrow{\mathbf{f}}, \mathbf{a} \rightarrow \overleftarrow{\mathbf{F}}, \mathbf{a} & \overrightarrow{\gamma}, \mathbf{a} \rightarrow \overleftarrow{\gamma_+}, \mathbf{a} \\
\overrightarrow{\wedge}, \overleftarrow{\mathbf{T}} \rightarrow \overrightarrow{\wedge}_+ & \overrightarrow{\mathbf{f}_+}, \overleftarrow{\mathbf{T}} \rightarrow \overrightarrow{\mathbf{f}} & \overrightarrow{\wedge}_+, \overleftarrow{\mathbf{T}} \rightarrow \overrightarrow{\mathbf{t}} \\
\overrightarrow{\wedge}, \overleftarrow{\mathbf{F}} \rightarrow \overrightarrow{\mathbf{f}_+} & \overrightarrow{\mathbf{f}_+}, \overleftarrow{\mathbf{F}} \rightarrow \overrightarrow{\mathbf{f}} & \overrightarrow{\wedge}_+, \overleftarrow{\mathbf{F}} \rightarrow \overrightarrow{\mathbf{f}} \\
\overrightarrow{\wedge}, \overleftarrow{\gamma_+} \rightarrow \overrightarrow{\wedge}_o & \overrightarrow{\wedge}_o, \overleftarrow{\mathbf{T}} \rightarrow \overleftarrow{\mathbf{T}}, \overrightarrow{\wedge} & \overrightarrow{\wedge}_o, \overleftarrow{\mathbf{F}} \rightarrow \overleftarrow{\mathbf{F}}, \overrightarrow{\wedge}
\end{array}$$

**Table 3.** Evaluation rules for  $\wedge$  connective.

$$\mathbf{x}, \overleftarrow{\exists} \rightarrow \exists_R \quad \overrightarrow{\exists}, \mathbf{x} \rightarrow \exists_L \quad \mathbf{x}, \overleftarrow{\forall} \rightarrow \forall_R \quad \overrightarrow{\forall}, \mathbf{x} \rightarrow \forall_L$$

**Table 4.** Setting up the quantifiers and the direction of the results.

*Aggregating the results.* Actual aggregation is initiated by  $\overrightarrow{\mathbf{c}}$  and then executes according to the rules given in Fig 4(b). We just have  $[\text{reduce:qsat:exec}] = \overrightarrow{\mathbf{c}}$ . The complete *Reduce* phase is implemented by

$[\text{reduce:qsat}(Q_1x_1 \cdots Q_nx_n)] = [\text{reduce:qsat:exec}] [\text{reduce:qsat:init}(Q_1x_1 \cdots Q_nx_n)]$ .

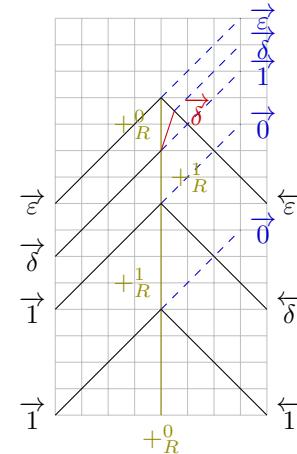
## 5 Machines for SAT variants

Similar machines for variants of SAT can be obtained easily, typically by using different modules for the *Reduce* phase. All the details of modules and rules can be found in [8].

*#SAT.* Counting the number of satisfying assignments for  $\psi$  can be achieved using a module implementing a binary adder with signals as shown in Fig. 5.

*ENUM-SAT.* Returning all the satisfying assignments for a propositional formula  $\psi$  can be achieved easily by adding a module which stores satisfying assignments as stationary beams and which annihilates the non-satisfying ones.

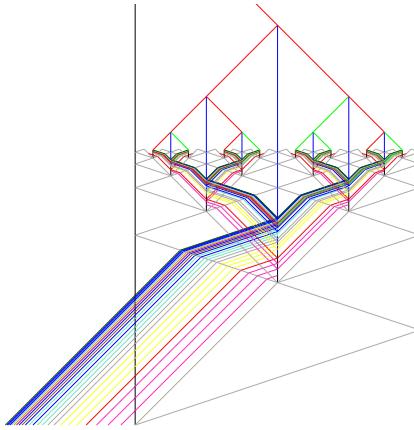
*MAX-SAT.* It consists in finding the maximum number of *clauses* that can be satisfied by an assignment. Here we must count (with the previous adder module) the number of satisfied clauses rather than the number of satisfying assignments, and then stack a module for computing the max of two binary numbers.



**Fig. 5.** Computing  $3 + 1$ .

## 6 Complexities

As mentioned in Sect. 1, we implement algorithms for satisfiability problems on signal machines in order to investigate the computational power of our abstract geometrical model of computation and to compare it to others. As we shall see, for such comparisons to be meaningful, the way complexity is measured is essential and must be adapted to the nature of the computing machine.



**Fig. 6.** The whole diagram.

Since signal machines can be regarded as the extension of cellular automata from discrete to continuous time and space, it might seem natural to measure time (resp. space) complexity of a computation using the height (resp. width) of its space-time diagram. But, in our applications to SAT variants, these are bounded and independent of the formula: the *Map* phase is bounded by the fractal, and, by symmetry, so is the *Reduce* phase. Indeed, in general, by an appropriate scaling of the initial configuration, a finite computation could be made as small as desired. Thus, height and width are no longer pertinent measures of complexity.

Instead, we should regard our construction as a massively parallel computational device transforming inputs into outputs. The input is the initial configuration at the bottom of the diagram, and the output is the truth value signal coming out at the top of the whole construction, as seen in Fig. 6 for formula  $\exists x_1 \forall x_2 \forall x_3 (x_1 \wedge \neg x_2) \vee x_3$  \*\* . 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, and the operations executed by the thread are the collisions occurring on this path.

Formally, we view a space-time diagram as a directed acyclic graph of collisions (vertices) and signals (arcs) oriented according to causality. Time complexity is then defined as the maximal size of a chain of collisions *i.e.* the length of the longest path, and space complexity as the maximal size of an anti-chain *i.e.* the size of the maximal set of signals pairwise un-related. This model-specific measure of time complexity is called *collisions depth*.

For the present construction, if  $s$  is the size of the formula and  $n$  the number of variables, space complexity is exponential: during evaluation,  $2^n$  independent computations are executed in parallel, each one involving less than  $s^2$  signals, so that the total space complexity is in  $\mathcal{O}(s^2 \cdot 2^n)$ .

---

\*\* All the diagrams used as examples in the paper were generated by Durand-Lose's software, implemented in Java, and corresponds to a run of our Q-SAT solver for the running example.

Regarding the time complexity: the initial configuration contains at most  $\mathcal{O}(s^2)$  signals (from the compilation process as explained in Sect. 4, other modules adding only a linear number of signals). The primary contribution to the number of collisions along an ascending path comes, at each of the  $n$  levels, from the reflected beam crossing the incoming beam. Thus a thread involves  $\mathcal{O}(n \cdot s^2)$  collisions, making the collision depth cubic in the size of the formula instead of quadratic for our previous family of machines [7]. So here the measure of the time complexity takes one more polynomial degree (from quadratic to cubic) when we get an algorithm which is independent of the input instead of an instance-dependent one. This gives us an idea of the price for genericity.

## 7 Conclusion

We showed in this paper that abstract geometrical computation can solve Q-SAT in bounded space and time by means of a single generic signal machine. This is achieved through massive parallelism enabled by a fractal construction that we call the *fractal cloud*. We adapted the Map/Reduce paradigm to this fractal cloud, and described a modular programming approach making it easy to assemble generic machines for SAT variants such as #SAT or MAX-SAT.

As we explained in Sect. 6, time and space are no longer appropriate measures of complexity for geometrical computations. This leads us to propose new definitions of complexity, specific to signal machines , and taking in account the parallelism of the model: time and space complexities are now defined respectively by the maximal sizes of a chain and an anti-chain, when the diagram is regarded as a directed acyclic graph. Time complexity thus defined is called *collision depth* and is cubic for the construction given here.

Although the model is purely theoretical and has no ambition to be physically realizable, it is a significant and distinguishing aspect of signal machines that they solve satisfiability problems while adhering to major principles of modern physics —finite density and speed of information, causality— that are typically not considered by other unconventional models of computation. They do not, however, respect the quantization hypothesis, nor the uncertainty principle.

We are now furthering our research along two axes. First, the design and applications of other fractal structures for modular programming with fractal parallelism. Second, the investigation of computational complexity classes, both classical and model-specific for abstract geometrical computation.

## References

1. Aaronson, S.: NP-complete problems and physical reality. SIGACT News. 36(1), 30–52 (2005)
2. Asarin, E., Maler, O.: Achilles and the Tortoise climbing up the arithmetical hierarchy. In: 15th Int. Conf. on Foundations of Software Technology and Theoretical Computer Science (FSTTCS '95). pp. 471–483. No. 1026 in LNCS (1995)

3. Cook, M.: Universality in elementary cellular automata. *Compl. Syst.* 15(1), 1–40 (2004)
4. Cook, S.: The complexity of theorem proving procedures. In: 3rd Symp. on Theory of Computing (STOC '71). pp. 151–158. ACM (1971)
5. Dean, J., Ghemawat, S., Inc., G.: Map/Reduce: simplified data processing on large clusters. In: 6th Symp. on Operating Systems Design & Implementation (OSDI'04). USENIX Association (2004)
6. Duchier, D., Durand-Lose, J., Senot, M.: Fractal parallelism: Solving SAT in bounded space and time. In: Cheong, O., Chwa, K.Y., Park, K. (eds.) 21st Int. Symp. on Algorithms and Computation (ISAAC '10). pp. 279–290. LNCS 6506, Springer (2010)
7. Duchier, D., Durand-Lose, J., Senot, M.: Massively parallel automata in Euclidean space-time. In: IEEE 4th Int. Conf. on Self-Adaptive and Self-Organizing Systems Workshops (SASOW '10). pp. 104–109. IEEE Computer Society (2010)
8. Duchier, D., Durand-Lose, J., Senot, M.: Computing in the fractal cloud: modular generic solvers for SAT and Q-SAT variants (extended version). Arxiv preprint arXiv:1105.3454, available at <http://arxiv.org/abs/1105.3454> (2011)
9. Durand-Lose, J.: Forecasting black holes in abstract geometrical computation is highly unpredictable. In: Cai, J.Y., Cooper, S.B., Li, A. (eds.) 3rd Int. Conf. on Theory and Applications of Models of Computations (TAMC '06). pp. 644–653. No. 3959 in LNCS, Springer (2006)
10. Durand-Lose, J.: Abstract geometrical computation and computable analysis. In: Costa, J., Dershowitz, N. (eds.) 8th Int. Conf. on Unconventional Computation 2009 (UC '09). pp. 158–167. LNCS 5715, Springer (2009)
11. Durand-Lose, J.: Abstract geometrical computation 4: small Turing universal signal machines. *Theoret. Comp. Sci.* 412, 57–67 (2011)
12. Fischer, P.: Generation of primes by a one-dimensional real-time iterative array. *Jour. ACM* 12(3), 388–394 (1965)
13. Goliae, S., Jalili, S.: An optical solution to the 3-SAT problem using wavelength based selectors. *The Journal of Supercomputing* pp. 1–10 (2010)
14. Huckenbeck, U.: Euclidian geometry in terms of automata theory. *Theoret. Comp. Sci.* 68(1), 71–87 (1989)
15. Jacopini, G., Sontacchi, G.: Reversible parallel computation: an evolving space-model. *Theoret. Comp. Sci.* 73(1), 1–46 (1990)
16. Jeandel, E., Vanier, P.:  $\pi_0^1$  sets and tilings. In: Ogihara, M., Tarui, J. (eds.) 8th Int. Conf. on Theory and Applications of Models of Computation (TAMC '11). LNCS, vol. 6648, pp. 230–239. Springer (2011)
17. Mackie, I.: A visual model of computation. In: Kratochvíl, J., Li, A., Fiala, J., Kolman, P. (eds.) 7th Int. Conf. on Theory and Applications of Models of Computation (TAMC '10). LNCS, vol. 6108, pp. 350–360. Springer (2010)
18. Margenstern, M., Morita, K.: NP problems are tractable in the space of cellular automata in the hyperbolic plane. *Theoret. Comp. Sci.* 259(1–2), 99–128 (2001)
19. Naughton, T., Woods, D.: An optical model of computation. *Theoret. Comput. Sci.* 334(1–3), 227–258 (2005)
20. Păun, G.: P-systems with active membranes: Attacking NP-Complete problems. *Jour. of Automata, Languages and Combinatorics* 6(1), 75–90 (2001)
21. Stockmeyer, L., Meyer, A.: Word problems requiring exponential time. In: 5th ACM Symp. on Theory of Computing (STOC '73). vol. 16, pp. 1–9 (1973)