

Abstract geometrical computation 4: small Turing universal signal machines

Jérôme Durand-Lose^a

^a*Laboratoire d'Informatique Fondamentale d'Orléans, Université d'Orléans,
B.P. 6759, F-45067 ORLÉANS Cedex 2.*

Abstract

This article aims at providing signal machines as small as possible able to perform any computation (in the classical understanding). After presenting signal machines, it is shown how to get universal ones from Turing machines, cellular-automata and cyclic tag systems. Finally a halting universal signal machine with 13 meta-signals and 21 collision rules is presented.

Key words: Abstract geometrical computation, cyclic tag systems, Turing universality, signal machines, small universal machines.

1 Introduction

Computation and universality have been defined in the 1930's. In the last five or so decades, it has been unveiled how common they are. The question about the frontier moved from proving the universality of dynamical systems to the complexity of universal machines. This is not only an intellectual challenge, but also important to find niches escaping tons of non decidability results or otherwise assert that these niches are too small to be of any interest.

There have already been a lot of investigations on small Turing machines (Rogozhin, 1982, 1996; Margenstern, 1995; Kudlek, 1996; Baiocchi, 2001), register machines (Korec, 1996), and cellular automata (Ollinger, 2002; Cook, 2004).

Email address: Jerome.Durand-Lose@univ-orleans.fr (Jérôme Durand-Lose).
URL: <http://www.univ-orleans.fr/lifo/Members/Jerome.Durand-Lose>
(Jérôme Durand-Lose).

Moreover, as (Turing) universality (capability to carry out any Turing/classical computation) has been developed in very limited systems, some “adaptations” were made and various notions of universality exist:

- *polynomial time universality* when polynomial computations (as defined in complexity theory) are still done in polynomial time, as opposed to *exponential time universality* provided by e.g. 2-counter automata, and
- *semi-universality* when the computation must be started on an infinite configuration –for example the whole tape of a Turing machine is filled with some ultimately-periodic infinite world– or in many cases like cellular automata which naturally work on infinite configurations.

In the present article, minimal universal machine in the context of *Abstract geometrical computation* (AGC) is investigated. AGC has been introduced as a continuous counterpart of cellular automata. This move is inspired by the way dynamics of CA is often designed or analyzed in an Euclidean space. In AGC, signals are moving with constant speed in an Euclidean space. When they meet, they are replaced/rewritten; interacting in a collision based computing way. A *signal machine* (SM) defines existing kinds of signal, *meta-signals*, and their interactions, *collision rules*. AGC allows, since this is a graphical model, to understand the way the information is moved around and interacts as shown by the various illustrations.

Abstract geometrical computation uses continuous space and time so that the possibility Zeno effect has to be considered. Indeed, it can happen and be used to compute beyond Turing computability and to climb the arithmetical hierarchy by using accumulations of collisions (Durand-Lose, 2009a). But since only Turing computability is addressed here, accumulations are not considered.

The straightforward measure of complexity (or simple/small-ness) of an SM is the number of meta-signals. If there are m meta-signals defined, then there are at most $2^m - m - 1$ possible collision rules (at least two signals are needed in a collision and they must be all different). Since parallel signals cannot interact, the number of rules could be much lower, but nevertheless exceeding by far the number of meta-signals. In many constructions, only a small part is defined, the rest being either undefined or *blank* (i.e. signals just cross each other). So that the number of specially defined rules is also a good complement to the size of an SM.

In Durand-Lose (2005a,b), Turing universality is proven by reduction from 2-counter automata. This result is not interesting here since the universal SM they provide are exponential time universal and not very small. (For an SM computation, the *time complexity* is the longest chain of collisions linked by signals.)

In the present article, universal SM are generated from Turing machines, cellular automata and cyclic tag systems. In each case, special care is taken in order to save signals and collision rules. For Turing machines, only results based on the construction in Durand-Lose (2009a) are presented; this direct simulation is not detailed since it does not provide the best bound. Cellular automata (CA) are massively parallel devices where bounds are also very tight and provide a very small semi-universal SM.

Cyclic tag systems (CTS) work by considering a binary word and a circular list of binary words. At each iteration, the first bit of the word is deleted and if it is 1 the first word of the list is added to the end of the word, then the list is rotated. CTS are polynomial time Turing universal (Neary and Woods, 2006a) and have been used to provide the best known bounds on Turing machines (Woods and Neary, 2006, 2007; Neary and Woods, 2006b) and cellular automata (Cook, 2004). They also provide the best bounds presented here.

Signal machines are presented in Section 2. Each following section deals with a different universal model and presents ways to simulate them with SM: Turing machines in Section 3, cellular automata in Section 4, and cyclic tag systems in Section 5. Section 6 gathers some concluding remarks.

2 Definitions

In *Abstract geometrical computation*, dimensionless objects are moving on the real axis. When a collision occurs they are replaced according to rules. This is defined by the following machines:

Definition 1 A *signal machine* is defined by (M, S, R) where M (*meta-signals*) is a finite set, S (*speeds*) a mapping from M to \mathbb{R} , and R (*collision rules*) a function from the subsets of M of cardinality at least two into subsets of M (all these sets are composed of meta-signals of distinct speed).

Each instance of a meta-signal is a *signal*. The mapping S assigns *speeds* to signals. They correspond to the inverse slopes of the line segments in space-time diagrams. A *collision rule*, $\rho^- \rightarrow \rho^+$, defines what emerges (ρ^+) from the collision of two or more signals (ρ^-). Since R is a function, signal machines are deterministic. The *extended value set*, V , is the union of M and R plus one symbol for void, \emptyset . A *configuration*, c , is a mapping from \mathbb{R} to V such that the set $\{x \in \mathbb{R} \mid c(x) \neq \emptyset\}$ is finite. An *infinite configuration*, is a similar mapping such that the previous set has no accumulation point.

A signal corresponding to a meta-signal μ at a position x , i.e. $c(x) = \mu$, is moving uniformly with constant speed $S(\mu)$. A signal must start (resp. end)

in the initial (resp. final) configuration or in a collision. This corresponds to condition 2 in Def. 2. At a $\rho^- \rightarrow \rho^+$ collision, signals corresponding to the meta-signals in ρ^- (resp. ρ^+) must end (resp. start) and no other signal should be present (condition 3).

Definition 2 The *space-time diagram* issued from an initial configuration c_0 and lasting for T , is a mapping c from $[0, T]$ to configurations (i.e. from $\mathbb{R} \times [0, T]$ to V) such that, $\forall (x, t) \in \mathbb{R} \times [0, T]$:

- (1) each $\{x \in \mathbb{R} \mid c_t(x) \neq \emptyset\}$ is finite,
- (2) if $c_t(x) = \mu$ then $\exists t_i, t_f \in [0, T]$ with $t_i < t < t_f$ or $0 = t_i \leq t < t_f$ or $t_i < t \leq t_f = T$ s.t.:
 - $\forall t' \in (t_i, t_f), c_{t'}(x + S(\mu)(t' - t)) = \mu,$
 - $t_i = 0$ or $(c_{t_i}(x + S(\mu)(t_i - t)) = \rho^- \rightarrow \rho^+ \text{ and } \mu \in \rho^+),$
 - $t_f = T$ or $(c_{t_f}(x + S(\mu)(t_f - t)) = \rho^- \rightarrow \rho^+ \text{ and } \mu \in \rho^-);$
- (3) if $c_t(x) = \rho^- \rightarrow \rho^+$ then $\exists \varepsilon, 0 < \varepsilon, \forall t' \in [t - \varepsilon, t + \varepsilon] \cap [0, T], \forall x' \in [x - \varepsilon, x + \varepsilon],$
 - $(x', t') \neq (x, t) \Rightarrow c_{t'}(x') \in \rho^- \cup \rho^+ \cup \{\emptyset\},$
 - $\forall \mu \in M, c_{t'}(x') = \mu \Leftrightarrow \text{or } \begin{cases} \mu \in \rho^- \text{ and } t' < t \text{ and } x' = x + S(\mu)(t' - t), \\ \mu \in \rho^+ \text{ and } t < t' \text{ and } x' = x + S(\mu)(t' - t). \end{cases}$

On space-time diagrams, time is increasing upward. The traces of signals are line segments whose directions are defined by $(S(\cdot), 1)$ (1 is the temporal coordinate). Collisions correspond to the extremities of these segments. This definition can easily be extended to $T = \infty$ and to infinite initial configuration.

Although speeds may be any real and thus encode information, in the following, only a few integer values are used. Similarly, the distance between signals may be any real but only integer positions are used.

2.1 Time complexity measure

As a computing device, the input is the initial configuration and the output is the final configuration. A SM is *Turing universal* if there exists encodings/representations through which it can go from the code of a Turing machine (or any equivalent model of computation) and an entry to the output (if any). AGC provides dynamical systems with no halting feature. Mainly two approaches exist to settle this:

- an observer checks that the end of the computation is reached according to some property over the configuration, e.g. the presence of a meta-signal, or
- the system reaches a stable state, i.e. no more collision is possible.

The last one is preferred since the halting is then a part of the system; although it is sometimes meaning-less: for example, when simulating a cellular automata, since there is no halting feature in CA, the same discussion arises again.

To consider polynomial time universality, the time complexity of a computation should be defined. Space and time are continuous, by rescaling a finite computation, it can be made as short as wanted, and even worse: any infinite computation starting from a finite configuration can be automatically folded into a finite portion of the space-time diagram (Durand-Lose, 2009a). So that a correct notion of time complexity lies elsewhere.

Collisions are considered as the discrete steps related by signals: a collision is *causally before* another if a signal generated by the first one ends in the second one. This yields a direct acyclic graph.

Definition 3 The *time complexity* of an SM computation is the longest length of a chain in the collision causality DAG.

For *space complexity* one may consider the longest length of an anti-chain or the maximum number of signals present at the same time (which form a cut).

2.2 Generating an infinite periodic signal pattern

The aim is to generate a periodic infinite sequence of signals on the side of a space-time diagram. This is useful to generate, starting from a finite number of signals the infinite data for semi-universality (when it is periodic). For TM, since it happens away from the head, there is no problem. The CA case is not so simple as explained later. In both cases, the increasing of the size allows to fall back into regular universality.

The construction works as follows: a signal bounces way and back between two signals, each time it bounces on the upper signal, a signal is emitted. The boundary signals are used to record the location in the pattern. The construction is simple and straightforward. It is only presented on an example: an infinite sequence of period 3, $(\mu_1\mu_2\mu_3)^\omega$, where each μ_i is already defined. The added signals and rules as well as the resulting space-time diagram are displayed on Fig. 1.

To save a signal, the signal is emitted at the bottom. The number of added signals is one plus the period (the output signals are not counted). The number of rules is twice the period. If the output signals are to be set at unequal distances, the generator uses one plus twice the period meta-signals (one for the two borders and different way and back for the whole period).

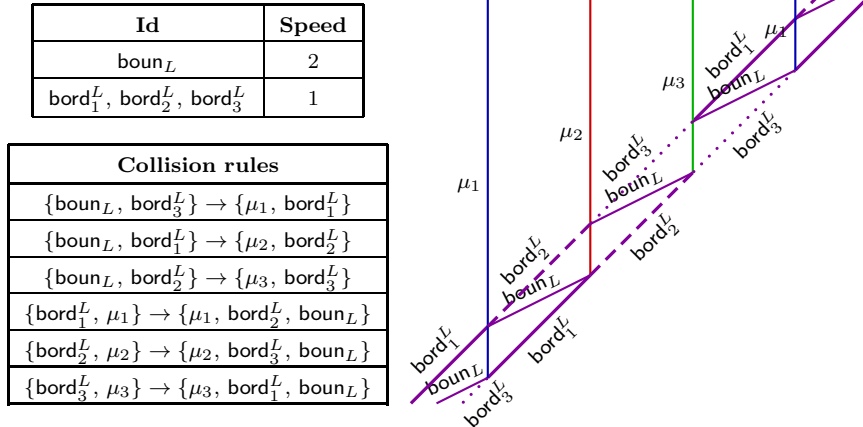


Fig. 1. Generating a periodic pattern.

3 Turing machines

Due to the poor results generated and the lack of space, the presentation of Turing machines and the full construction from Durand-Lose (2009a) are not given here. The number of meta-signals in the signal machine simulating a TM is: 1 for each tape symbol, 2 for each state, and 4 extra signals for enlarging the tape. The number of collision rules is up bounded by: 2 for each entry of the transition table of the TM, and plus 2 for each transition on the blank symbol plus 2 (for enlarging the table). In the cited paper, TM halting corresponds to the head leaving the tape. With a halting state, the corresponding collision rules produce the disappearance of the state signal, at no extra cost.

Considering the curve of Neary and Woods (2006b); Woods and Neary (2006, 2007), this leads to the following values: 18 meta-signals for polynomial time universality. The number of collision rules is bounded by 62. The exact numbers have not been computed since fewer meta-signals is possible.

Semi-universal TM with fewer states exists. Like the one presented in Subsect. 2.2, add hoc constructions to generate the extension of the tape –to achieve full universality– would add too many states and has not been considered. For semi-universality, using Smith (2007), the 2-states 3-symbol TM generates a SM with 7 meta-signals (the ones for enlarging are not needed) and 6 collision rules.

4 Cellular automata

Cellular automata (CA) operate over infinite arrays of cells. Each cell can be in finitely many states. (Infinite array is the only way to ensure unbounded memory.) A CA changes the states of all the cells simultaneously according to

a local function and the states of the two surrounding cells. This is parallel, synchronous, local and uniform process.

Definition 4 A cellular automaton is defined by (Q, f) where: Q is a finite set of states, $f : Q^3 \rightarrow Q$, is the local function. The global function, $\mathcal{G} : Q^{\mathbb{Z}} \rightarrow Q^{\mathbb{Z}}$, is defined by: $\forall i \in \mathbb{Z}, \mathcal{G}(c)_i = f(c_{i-1}, c_i, c_{i+1})$.

Only CA of dimension 1 and radius 1 are considered here. Higher dimensions can be covered similarly by higher dimension signal machines. Radius 1 means that a cell only communicates with its two closest neighbors (one on each side). Broader radii could have been considered, but more signals are needed to convey information at greater distances (e.g. 5 meta-signals per state instead of 3 for radius 2).

Halting is not provided by CA, it can be defined by reaching a stable/periodic configuration, the appearance of a state/pattern in the configuration or on some designated cell. In any case, more meta-signals are needed.

There are two ways to manipulate finite CA-configurations: use some *quiescent state*, $q_{\#}$ (satisfying $f(q_{\#}, q_{\#}, q_{\#}) = q_{\#}$), for undefined cells; or use a periodic spatial extension on both side. This pattern is also (ultimately) time-periodic in the CA evolution. In the simulation, the computation is framed by periodic signals according to oblique discrete lines in the space-time diagram of the CA as illustrated by the example on Fig. 3.

The idea is to locate the cells at integer positions and each time the local function is used, three signals are emitted, one for the cell and one for each of the closest cell on each side. For each state, s , there are three meta-signals: s_L , s and s_R or speeds $-1, 0$ and 1 respectively. A transition is performed when a cell received simultaneously the values from this two neighbors (special care has to be put on locations to ensure exact meetings). The local function is encoded in the collision rules: if $f(s, t, u) = v$ then the following rule is defined $\{s_R, t, u_L\} \rightarrow \{v_L, v, v_R\}$. Rule 110 is presented Fig. 2 as well as a generated rule.

Output	0	1	1	0	1	1	1	0
Input	111	110	101	100	011	010	001	000

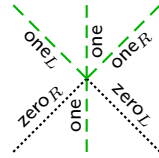


Fig. 2. Local function and the rule implementing $f(0, 1, 0) = 1$.

Evolution and simulation on the configuration 11, framed on the left by ${}^{\omega}(10)$ and on the right by $(011)^{\omega}$ is presented on Fig. 3. As it can be seen, the way the frames are positioned (in boldface) in the evolution is not trivial and neither is the generated periodic pattern on both side of the simulation.




Fig. 3. Evolution and simulation of rule 110 on 11 framed by $\omega(10)$ and $(011)^\omega$.

With the references in Durand-Lose (2009b) and Ollinger (2008), it is hard to go below 30 meta-signals for full universality. Ollinger (2002); Cook (2004) and Richard and Ollinger (2008) provide universal CA with very few states but they use periodic extensions that encode a boolean function or a cyclic tag system which leads to too many meta-signals unless a clever way to generate them is found. Nevertheless, with rule 110 (proved universal in Cook (2004)), a 6 meta-signals 8 rules semi-universal SM is generated. This SM is not halting. To get a halting SM from a CA, one has to know precisely what corresponds to halt and add signals and rules for the SM to halt.

5 Cyclic tag systems simulation

Cyclic tag systems (CTS) are defined by a word and a circular list of appendants. Both the word and the appendants are binary words. The system is updated in the following way: the first bit of the word is removed. If it is 1 then the first appendant is appended at the end of the word (otherwise nothing is done). Then, the list is rotated circularly. The list represents the code and the word the input. Not only are CTS able to compute (Cook, 2004) but also they can do it with polynomial slowdown (Neary and Woods, 2006a).

The simulation is done with two objects: parallel signals encoding the word, and a cycle list: each time it cycles, depending on the signals that started it can deliver a copy of the first appendant or not. Signals encoding the word are placed so that delivered copies automatically enlarge the word on the right.

The initial configuration is presented before the dynamics of the various elements. The signals are defined on Figure 4. The following naming convention is used: meta-signals with no subscript have speed 0 and the ones with subscript LL , R and RR have speed -2 , 1 and 2 respectively.

As illustrated on Figure 5, the initial configuration is composed, left to right, of: **last**, \mathbf{go}_{LL} that starts the dynamics, then **one's** and **zero's** to encode the word, then **first** to indicate the beginning of the cyclic list then alternatively **one's** and **zero's** to encode each appendant and **sep** to separate them and finally **last**. The iteration starts when the \mathbf{go}_{LL} signal erases **last** and bounces as \mathbf{go}_{RR} .

	zero	one	first	sep	last	go _{LL}
go _{RR}	last, zero _R	last, one _R	first			
one _{RR}	zero, zero _R , one _{RR}	one, one _R , one _{RR}		go _{LL} , last, false _R		
zero _{RR}	zero _R , zero _{RR}	one _R , zero _{RR}		go _{LL} , last, false _R		
true _R	—	—			first, false _R	—
false _R	—	—	sep, go _{RR}	—	first, one _R	go _{LL} , true _R
one _R	—	—	first, true _R , one _{RR}	—	—	—
zero _R	—	—	first, false _R , zero _{RR}	—	—	—
go _{LL}	—	—	go _{LL}		go _{RR}	

	zero _R	one _R	false _R	true _R
go _{RR}		last	zero, go _{RR}	one, go _{RR}
one _{RR}		—	—	—
zero _{RR}		—	—	—

“—” means blank

Other blank rules
zero _R , zero, go _{LL}
zero _R , one, go _{LL}
one _R , zero, go _{LL}
one _R , one, go _{LL}

Fig. 6. List of all the collision rules.

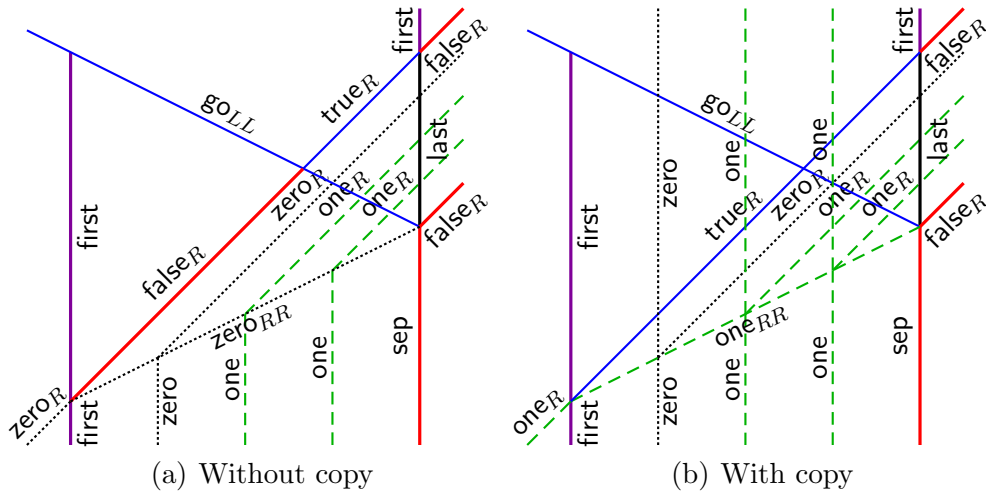


Fig. 7. Starting the rotation and restarting the dynamics.

is straightforward: the parallel $false_R$, $zero_R$ and one_R are crossing sep , $zero$ and one until $last$ is met.

The last part of the rotation could have been the symmetric of the first part; in some ways, it is as presented on Figure 8. To save meta-signals, $false_R$ have been used both to mark the beginning and the end of the appendant. But they have different meanings, so the first one changes $last$ to $first$ and the second to sep and go_{RR} . The problem is that $first$ interacts with $zero_R$ and one_R : $zero_R$ (resp. one_R) is changed to $false_R$ (resp. $true_R$) and one_{RR} (resp. $zero_{RR}$). This generates the lattice in the triangle (the signals do not interact inside it). The bits of the appendant are now encoded with $false_R$'s and $true_R$'s and follow the paths $zero_R$'s and one_R 's would have. It remains to go_{RR} to turn them back into $zero$'s and one 's and to turn the final one_R to $last$. go_{RR} , $zero_{RR}$ and one_{RR} are parallel, all the translating signals are parallel; so that the appendant is

recreated with exactly the same distances.




Fig. 8. Ending the rotation.

Figure 9(a) shows one full iteration of a CTS including a whole rotation. Figure 9(b) shows the halting by reaching an empty word. Figure 9(c) shows the effect of the halting appendant as explained below. Figure 9(d) shows one entire simulation with a halting appendant and cleaning added: blank rules have been modified in order to destroy the garbage signals escaping on the right that nevertheless would never interact with the rest of the configuration nor provoke any collisions since they are parallel.

5.1 Adding halt at the cost of one rule

Three halting conditions exist: on empty word, on cycling or on special halting appendant. The first does not yield any result but nevertheless has to be implemented and already is. The second needs an extra layer of construction to detect cycling and is not considered. In the third case, when the halt appendant is activated, go_{LL} has to meet true_R . The position of this collision is determined by the distance between first and the next sep . It appends exactly at $2/3$ of the distance. If a one signal is set exactly at this position, it gets into the collision. The extra rule

$$\{\text{true}_R, \text{one}, \text{go}_{LL}\} \rightarrow \{\text{true}_R\}$$

is used to destroy both one and go_{LL} . Then the rotation finished and no more collision is possible. The result of the computation is the sequence left of first .

The rotating process ensures that the distances between the signals remain constant, so that a one at $2/3$ remains there (and a one not at $2/3$ cannot get to $2/3$). It should also be ensured that, in the initial position, there is no one at $2/3$ not standing for halt which is straightforward to reach by, e.g., using locations $1/2, 3/4, 7/8, \dots$




Fig. 9. Simulation of a cyclic tag system.

6 Conclusion

Theorem 5 *There is a universal halting signal machine with 13 meta-signals and 21 non-blank rules. There is a non-halting semi-universal signal machine with 6 meta-signals and 8 non-blank rules.*

The first construction uses a collision with 3 signals. A 3-signal collision needs perfect synchrony while 2-signal collisions are more robust to small perturba-

tions. If only 2 signals collisions are allowed, then a halting meta-signals can be added and processed in the circular list as the value 1 (except that it does not generate other signals in the lattice at the end of the rotation). This leads to a signal machine with 15 meta-signals (the same ones plus `halt` and `haltR`) and 24 non blank rules.

The race for small universal devices also runs for restrictions, e.g. to reversible machines. Universality has also been provided for reversible and conservative SM Durand-Lose (2006). We believe that the CTS simulation could be turned reversible with less than two extra meta-signals.

One key feature of AGC is that the continuous space and time can be used to produce accumulations and on top of it to do black hole computations (an accumulation containing a whole, potentially infinite, Turing computation). One may wonder about the minimal number of meta-signals for an accumulation and for the black hole effect. Let us note that 4 signals is enough to make an accumulation and 3 might to be enough, whereas 2 is not. It also seems that the black hole effect can be added to the cyclic tag system simulation with a few extra meta-signals.

References

- Baiocchi, C. (2001). Three small universal Turing machines. In *Machines, Computations, and Universality*, volume 2055, pages 1–10. Springer.
- Cook, M. (2004). Universality in elementary cellular automata. *Complex Systems*, 15:1–40.
- Durand-Lose, J. (2005a). Abstract geometrical computation for black hole computation (extended abstract). In Margenstern, M., editor, *Machines, Computations, and Universality (MCU '04)*, number 3354 in LNCS, pages 176–187. Springer.
- Durand-Lose, J. (2005b). Abstract geometrical computation: Turing computing ability and undecidability. In Cooper, B. S., Löwe, B., and Torenvliet, L., editors, *New Computational Paradigms, 1st Conf. Computability in Europe (CiE '05)*, number 3526 in LNCS, pages 106–116. Springer.
- Durand-Lose, J. (2006). Reversible conservative rational abstract geometrical computation is Turing-universal. In Beckmann, A. and Tucker, J. V., editors, *Logical Approaches to Computational Barriers, 2nd Conf. Computability in Europe (CiE '06)*, number 3988 in LNCS, pages 163–172. Springer.
- Durand-Lose, J. (2009a). Abstract geometrical computation 3: Black holes for classical and analog computing. *Nat. Comput.*
- Durand-Lose, J. (2009b). Cellular automata, Universality of. In Meyers, R. A. and Adamatzky, A., editors, *Encyclopedia of Complexity and System Science*. Springer. Book in press.
- Korec, I. (1996). Small universal register machines. *Theoret. Comp. Sci.*, 168(2):267–301.

- Kudlek, M. (1996). Small deterministic turing machines. *Theoret. Comp. Sci.*, 168(2):241–255.
- Margenstern, M. (1995). Non-erasing Turing machines: A new frontier between a decidable halting problem and universality. In Baeza-Yates, R., Goles, E., and Poblete, P., editors, *LATIN'95*, number 911 in LNCS, pages 386–397. Springer.
- Neary, T. and Woods, D. (2006a). P-completeness of cellular automaton rule 110. In Bugliesi, M., Preneel, B., Sassone, V., and Wegener, I., editors, *International Colloquium on Automata Languages and Programming (ICALP '06)*, number 4051(1) in LNCS, pages 132–143. Springer.
- Neary, T. and Woods, D. (2006b). Small fast universal Turing machines. *Theoret. Comp. Sci.*, 362(1–3):171–195.
- Ollinger, N. (2002). The quest for small universal cellular automata. In *ICALP '02*, number 2380 in LNCS, pages 318–329. Springer.
- Ollinger, N. (2008). Universalities in cellular automata: a (short) survey. In Durand, B., editor, *Journées Automates cellulaires (JAC '08)*, pages 102–118.
- Richard, G. and Ollinger, N. (2008). A particular universal cellular automaton. In Neary, T., Woods, D., Seda, A. K., and Murphy, N., editors, *The Complexity of Simple Programs*. National University of Ireland, Cork.
- Rogozhin, Y. V. (1982). Seven universal Turing machines. In *Systems and theoretical programming*, number 63 in Mat. Issled., pages 76–90. Academia Nauk Moldavskoi SSR. in Russian.
- Rogozhin, Y. V. (1996). Small universal Turing machines. *Theoret. Comp. Sci.*, 168(2):215–240.
- Smith, A. (2007). Universality of Wolfram's 2, 3 Turing machine. *Complex Systems*. to appear.
- Woods, D. and Neary, T. (2006). On the time complexity of 2-tag systems and small universal Turing machines. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS '06)*, Berkeley, CA, USA, pages 439–448. IEEE Computer Society.
- Woods, D. and Neary, T. (2007). Small semi-weakly universal Turing machines. In Durand-Lolse, J. and Margenstern, M., editors, *Machines, Computations and Universality (MCU '07)*, number 4664 in LNCS, pages 246–257. Springer. extended version to appear in *Fundamenta Informaticae*.