

Automata and Program Analysis

Thomas Colcombet¹(✉), Laure Daviaud², and Florian Zuleger³

¹ IRIF, Case 7014 Université Paris Diderot, 75205 Paris Cedex 13, France

thomas.colcombet@irif.fr

² MIMUW, Banacha 2, 02-097 Warszawa, Poland

ldaviaud@mimuw.edu.pl

³ Institut für Informationssysteme 184/4, Technische Universität Wien,
Favoritenstraße 9–11, 1040 Wien, Austria

zuleger@forsyte.at

Abstract. We show how recent results concerning quantitative forms of automata help providing refined understanding of the properties of a system (for instance, a program). In particular, combining the size-change abstraction together with results concerning the asymptotic behavior of tropical automata yields extremely fine complexity analysis of some pieces of code.

This abstract gives an informal, yet precise, explanation of why termination and complexity analysis are related to automata theory.

1 Program Analysis and Termination

Program analysis is concerned with the automatic inference of properties of a chunk of code (or of a full program). Such analysis may serve many purposes, such as guaranteeing that some division by zero cannot occur in the execution of a program, or that types are properly used (when the language is not statically typed), or that there is no memory leakage, etc. Here, to start with, we are concerned with termination analysis, *i.e.*, proving that all executions of a program (or pieces of a program) eventually halt.

Consider the following *code C*:

```
void f() {
  uint x,y;
  x = read_input();
  y = read_input();
  while (x > 0) {
    if (y > 0) { // branch a
      y--;
    }
    else { // branch b
      x--;
      y = read_input();
    }
  }
}
```


It is clear for a human being that this piece of code eventually *terminates* whatever are the input values read during its execution. The question is how this can be automatically inferred? Of course, such a problem is in general undecidable unless restrictions are assumed (using variations around the theorem of Rice). The path we follow here consists in approximating the behavior of this code using size-change abstraction. We shall see that this abstraction transforms the code into a formalism, the size-change abstraction, for which termination is decidable.

2 Size-Change Abstraction

The *size-change abstraction* amounts to abstract a piece of code in the following manner:

- We identify some *size-change variables* that are considered of interest, that range over non-negative integers (in our example x and y). In general, size-change variables can represent any norm on the program state (a function which maps the state to the non-negative integers), such as the length of a list, the height of a tree, the sum of two non-negative variables, etc.
- We construct the *control-flow graph* (possibly simplified) of the code: vertices are positions in the code, and *edges* are steps of computations. We also identify some *entry* and *exit vertices* of the graph as one can expect.
- We abstract tests, *i.e.*, we replace all tests by non-determinism. This means that we consider possible executions independently of whether the tests in **if** statements or in **while** loops are true or not.
- Finally, each edge of the control-flow graph is labeled by *guards* expressing how the values of the size-change variables may evolve while taking the edge. The language for these relations is very restricted: it consists of a conjunction of properties of the form $x \geq y'$ or $x > y'$ where x, y, \dots represent the value of the variables before the edge is taken, while x', y', \dots represent the value of the variables after the edge is taken. We add guards *conservatively* in order to ensure the correctness of the abstraction, *i.e.*, we only add a $x \geq y'$ to some edge guard if we can guarantee that the value of y is not greater than the value of x before edge is taken, and similarly for $x > y'$.

For the code \mathcal{C} , we obtain the following size-change abstraction \mathcal{S} :




We comment on the size-change abstraction \mathcal{S} of code \mathcal{C} . Here, we only consider one position in the code, which is the beginning of the **while** loop. This needs not be the case in general. The two edges a and b correspond respectively to executing the **if** branch and the **else** branch of code \mathcal{C} . In the first case

(edge a), the value of y strictly decreases while the value of x does not change. In the second case (edge b), the value of x strictly decreases while we have no information about the value that y might take after the transition (because of the ‘ $y = \text{read_input}()$ ’ in the code). We see that the two transitions a and b of \mathcal{S} are an *abstraction* of the branches a and b of \mathcal{C} : If the code \mathcal{C} executes branch a resp. b and the variable values (x, y) change to some (x', y') , then transition a resp. b of \mathcal{S} also allows the variable values (x, y) to change to (x', y') . This has the following important consequence for termination analysis: Every execution of \mathcal{C} is also an execution of \mathcal{S} . Thus, if we can show that \mathcal{S} terminates, then we can deduce the termination of \mathcal{C} .

In the following we explain how to reason about the termination of the size-change abstraction \mathcal{S} .

In order to be precise, we have to define the semantics of the model. We call *execution path* of the size-change abstraction a sequence of edges that (1) starts in an entry vertex, and (2) is formed of compatible edges, meaning that for any two consecutive edges in the sequence, the target of the first one should coincide with the source of the second one. Such an execution path is *halting* if it is finite and ends in an exit vertex. Such a definition of an execution path does not yet capture the semantics of variables. For this, we shall consider the traces that realize an execution. Formally, a *trace* of the size-change abstraction is a sequence of *configurations* consisting of a vertex and a valuation of the variables by non-negative integers, that respect the transitions of the size-change abstraction. This is best seen in an example. Consider the execution path *aabaabaa*. One possible trace that realizes this execution is the following one (where the second component represents the value of the variable x and the third one the value of the variable y):



We say that the *size-change abstraction terminates* if every execution path that is realized by a trace is finite. As argued above, this implies that the program also terminates since all the executions of the program are captured by the abstraction. Of course, the converse does not hold, in particular because we threw away a good part of the original semantics, and thus many reasons for the program to terminate are not recovered.

3 Max-Plus Automata

The above definition of size-change abstraction does not make the termination property obviously decidable, yet. What we show now is how this question can be reduced to a problem of universality in automata theory. The key concept behind this last reduction is that the pattern that prevents an execution path to be realizable would be an infinite sequence of variables that are related by the guards via (non-necessarily stricts) inequalities, and infinitely many times by strict inequalities (finite sequences of arbitrarily large number of strict inequalities would also be a witness in some cases). Indeed, such a sequence would mean

the existence in a trace of an infinite decreasing sequence of non-negative integers; a contradiction. For instance, the following infinite execution path of the size-change abstraction:

$$\overbrace{a \dots a}^{n_0 \text{ times}} b \overbrace{a \dots a}^{n_1 \text{ times}} b \overbrace{a \dots a}^{n_2 \text{ times}} b \dots$$

is impossible since the a edge contains the guard $x \leq x'$, and the b relation $x < x'$. Hence, the following infinite sequence of relations is impossible:

$$\overbrace{x_0 \leq x_1 \leq \dots \leq x_{n_0-1} < x_{n_0}}^{n_0 \text{ inequalities}} \overbrace{\leq x_{n_0+1} \leq \dots \leq x_{n_0+n_1} < x_{n_0+n_1+1} \leq \dots}^{n_1 \text{ inequalities}},$$

in which x_i accounts for the value assumed by variable x at time i . No valuation of the x_i 's by *non-negative integers* can fulfill these constraints. For similar reasons, an infinite execution path that would eventually consist only of the edge b would be impossible, this time because of the variable y .

We shall now define an automaton model that can measure these “bad sequences” of dependencies: max-plus automata¹.


We will define a *max-plus automaton* that is able to “count” the maximal length of such sequences of strict inequalities. The recipe is the following:

- The *input alphabet* is the set of edges of the size-change abstraction.
- The *states* of the automaton are the size-change variables, plus two extra states, called \top and \perp .
- The *transitions* of the automaton are labeled by the edges of the size-change abstraction, and there is a transition from state x to state y labeled by δ if the guard of edge δ contains either ‘ $x \geq y'$ ’ or ‘ $x > y'$ ’; Furthermore the state \top is the origin of all possible transitions to every state (including itself), and \perp is the target all possible transitions originating from any states (including itself).
- Some transitions are marked *special* (or *costly*): the ones that arise from the case ‘ $x > y'$ ’.²
- All states are marked both *initial* and *final* (the feature of initial and final states, which is important in the theory of tropical automata, happens to be irrelevant for this application).

In our case, this yields the automaton below, in which the $*$ symbol means ‘all possible labels’, double arrows identify costly transitions, and we omitted to represent initial and final states:

¹ Though we take the principle, we do not use the standard notation of max-plus automata, which are traditionally defined as automata weighted over the max-plus semiring.

² In the standard terminology, non-costly transitions would be given *weight* 0, while costly ones would be attributed weight 1.



In this automaton, a sequence of inequalities over size-change variables can be witnessed by a path (formally, a run). Costly transitions correspond to strict inequalities in the guards.

The semantics of this max-plus automaton is to count the number of costly transitions and to maximize this cost among all runs. Formally, an *input* of the automaton is a sequence of edges of the size-change abstraction, such as the word $aaabaabaa$ (it is not necessarily an execution path so far). A priori, we consider both finite and infinite sequences of edges. A *run* of the automaton over input u is a sequence of transitions which forms a path in the graph of the automaton, and it is accepting if either it is infinite and starts in an initial state or it is finite, starts in an initial state and ends in a final one. The max-plus automaton \mathcal{A} can be used to compute a quantity given some input word u :

$$\llbracket \mathcal{A} \rrbracket(u) = \sup\{\text{cost}(\rho) \mid \rho \text{ accepting run over the input } u\} \in \mathbb{N} \cup \{\infty\},$$

where $\text{cost}(\rho) = \text{number of costly transitions in } \rho$.

For instance, over the input $u = aaabaabaa$, $\llbracket \mathcal{A} \rrbracket(u) = 3$. It corresponds to a run (there are in fact several of them) that assumes state y during the first three letters of the word. This is the maximal one, since in our example the automaton computes the maximum of the number of b -edges with the maximum of the longest block of consecutive a -edges.

The following lemma formalizes the correction of this reduction to automata:

Lemma 1. *The following properties are equivalent:*

1. *The size-change abstraction terminates.*
2. *All infinite execution paths u satisfy $\llbracket \mathcal{A} \rrbracket(u) = \infty$.*

It happens that Item 2 of Lemma 1 is decidable, from which we get:

Corollary 1 [9,10]. *The termination of size-change abstractions is decidable.*

Let us establish the decidability of Item 2. This requires some knowledge about Büchi automata theory. The reader may as well proceed to the next section.

In this proof, we rephrase the second item of Lemma 1 as an inclusion of Büchi automata. The first Büchi automaton \mathcal{E} accepts a sequence of edges if it forms a valid infinite execution path of the size-change abstraction. The second Büchi automaton, \mathcal{B} , is syntactically the max-plus automaton seen as a Büchi automaton, in which special transitions have to be visited infinitely often (the Büchi condition): it accepts an infinite execution path u if there exists a run containing infinitely many costly transitions. In particular, if u is accepted by \mathcal{B} , then $\llbracket \mathcal{A} \rrbracket(u) = \infty$ (property \star). In general the converse is not true, but one can check that if u is ultimately periodic (*i.e.*, of the form $uvvv\dots$), then $\llbracket \mathcal{A} \rrbracket(u) = \infty$ implies $u \in L(\mathcal{B})$ (property $\star\star$).

Let us now show that Item 2 of Lemma 1 is equivalent to:

$$L(\mathcal{E}) \subseteq L(\mathcal{B}).$$

Indeed, $L(\mathcal{E}) \subseteq L(\mathcal{B})$ means that all infinite execution paths u are accepted by \mathcal{B} , and thus $\llbracket \mathcal{A} \rrbracket(u) = \infty$ by \star . For the converse direction, assume that $L(\mathcal{E}) \subseteq L(\mathcal{B})$ does not hold, *i.e.*, there is an input that is accepted by \mathcal{E} but not by \mathcal{B} . It is known from Büchi that in this case there exists such an input u which is ultimately periodic. By $\star\star$, this means that $\llbracket \mathcal{A} \rrbracket(u)$ is finite, contradicting Item 2 of Lemma 1.

4 Complexity Analysis

Counting the number of costly transitions also gives an idea of the worst-case complexity of the program. Indeed, a possible execution of the program corresponds to an execution path of the size-change abstraction that is realized by a trace, and the time complexity of the execution is nothing but the length of the execution path.

To put this idea in action, let us consider the following slightly modified code:

```
void f(uint n) {
    uint x,y;
    x = read_input(n);
    y = read_input(n);
    while (x > 0) {
        if (y > 0) { // branch a
            y--;
        }
        else { // branch b
            x--;
            y = read_input(n);
        }
    }
}
```

This new code takes a non-negative integer as input, and the `read_input(n)` calls now guarantee that the value produced is in the interval $\{0, \dots, n\}$. A

careful look at this code reveals that it terminates in $O(n^2)$ steps. We would like the analysis to reach this level of precision.

In fact, everything is similar to the termination case we have explained up to now. The only change is that the values of the variables are implicitly ranging over the interval $\{0, \dots, n\}$. Under this assumption, the size-change abstraction also terminates within a quadratic bound.

The reduction to a max-plus automaton also remains valid, as shown by this variation around the ideas of Lemma 1:

Lemma 2. *The following properties are equivalent for all n and k :*

- *The size-change abstraction terminates within time bound k , assuming the variable values range in $\{0, \dots, n\}$.*
- *All execution paths u such that $\llbracket \mathcal{A} \rrbracket(u) \leq n$ have length at most k .*

However, we need now a much more delicate result of automata theory than the inclusion of Büchi automata. Here follows what we can do:

Theorem 1 [7]. *One can effectively compute, given as input a max-plus automaton \mathcal{A} , the value*

$$\liminf_{|u| \rightarrow \infty} \frac{\log(\llbracket \mathcal{A} \rrbracket(u) + 1)}{\log |u|}$$

which happens to be a rational in $[0, 1]$ or ∞ .

Now, as a corollary, we get:

Theorem 2 [7]. *The length of the longest execution path realized by a trace in a size-change abstraction is of order $\Theta(n^\alpha)$ if the variables are restricted to take values in $[0, n]$, where $\alpha \geq 1$ is a rational number. Moreover, there is an algorithm that given a terminating size-change abstraction computes such an α .*

It can also be proved that all the rationals $\alpha \geq 1$ can be achieved by a given size-change abstraction.

5 Related Work

The goal of this paper was to illustrate the size-change abstraction (SCA), which is a popular technique for automated termination analysis. The last decade has seen considerable interest in automated techniques for proving the termination of programs. In this short paper we limit ourselves to describing the related work on SCA. SCA has been introduced by Lee et al. [10]. SCA is employed for the termination analysis of functional [10, 11], logical [12] and imperative [1, 6] programs, term rewriting systems [5], and is implemented in the industrial-strength systems ACL2 [11] and Isabelle [8]. Recently, SCA has also been used for resource bound and complexity analysis of imperative programs [14], which motivated the results on complexity analysis presented in this paper. SCA is an

attractive domain for an automated analysis because of several strong theoretical results on termination analysis [10], complexity analysis [7, 13] and the extraction of ranking functions [3, 13]. Further research has investigated the generalization of size-change constraints to richer classes of constraints, including difference constraints [2], gap-order constraints [4] and monotonicity constraints [3].

References

1. Anderson, H., Khoo, S.-C.: Affine-based size-change termination. In: Ogori, A. (ed.) APLAS 2003. LNCS, vol. 2895, pp. 122–140. Springer, Heidelberg (2003). doi:[10.1007/978-3-540-40018-9_9](https://doi.org/10.1007/978-3-540-40018-9_9)
2. Ben-Amram, A.M.: Size-change termination with difference constraints. *ACM Trans. Program. Lang. Syst.* **30**(3), 16 (2008)
3. Ben-Amram, A.M.: Monotonicity constraints for termination in the integer domain. *Logical Methods Comput. Sci.* **7**(3), 1–43 (2011)
4. Bozzelli, L., Pinchinat, S.: Verification of gap-order constraint abstractions of counter systems. In: Kuncak, V., Rybalchenko, A. (eds.) VMCAI 2012. LNCS, vol. 7148, pp. 88–103. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-27940-9_7](https://doi.org/10.1007/978-3-642-27940-9_7)
5. Codish, M., Fuhs, C., Giesl, J., Schneider-Kamp, P.: Lazy abstraction for size-change termination. In: Fermüller, C.G., Voronkov, A. (eds.) LPAR 2010. LNCS, vol. 6397, pp. 217–232. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-16242-8_16](https://doi.org/10.1007/978-3-642-16242-8_16)
6. Codish, M., Gonopolskiy, I., Ben-Amram, A.M., Fuhs, C., Giesl, J.: Sat-based termination analysis using monotonicity constraints over the integers. *TPLP* **11**(4–5), 503–520 (2011)
7. Colcombet, T., Daviaud, L., Zuleger, F.: Size-change abstraction and max-plus automata. In: Csuhaj-Varjú, E., Dietzfelbinger, M., Ésik, Z. (eds.) MFCS 2014. LNCS, vol. 8634, pp. 208–219. Springer, Heidelberg (2014). doi:[10.1007/978-3-662-44522-8_18](https://doi.org/10.1007/978-3-662-44522-8_18)
8. Krauss, A.: Certified size-change termination. In: Pfenning, F. (ed.) CADE 2007. LNCS (LNAI), vol. 4603, pp. 460–475. Springer, Heidelberg (2007). doi:[10.1007/978-3-540-73595-3_34](https://doi.org/10.1007/978-3-540-73595-3_34)
9. Lee, C.S.: Ranking functions for size-change termination. *ACM Trans. Program. Lang. Syst.* **31**(3), 10:1–10:42 (2009)
10. Lee, C.S., Jones, N.D., Ben-Amram, A.M.: The size-change principle for program termination. In: POPL, pp. 81–92 (2001)
11. Manolios, P., Vroon, D.: Termination analysis with calling context graphs. In: Ball, T., Jones, R.B. (eds.) CAV 2006. LNCS, vol. 4144, pp. 401–414. Springer, Heidelberg (2006). doi:[10.1007/11817963_36](https://doi.org/10.1007/11817963_36)
12. Vidal, G.: Quasi-terminating logic programs for ensuring the termination of partial evaluation. In: PEPM, pp. 51–60 (2007)
13. Zuleger, F.: Asymptotically precise ranking functions for deterministic size-change systems. In: Beklemishev, L.D., Musatov, D.V. (eds.) CSR 2015. LNCS, vol. 9139, pp. 426–442. Springer, Cham (2015). doi:[10.1007/978-3-319-20297-6_27](https://doi.org/10.1007/978-3-319-20297-6_27)
14. Zuleger, F., Gulwani, S., Sinn, M., Veith, H.: Bound analysis of imperative programs with the size-change abstraction. In: Yahav, E. (ed.) SAS 2011. LNCS, vol. 6887, pp. 280–297. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-23702-7_22](https://doi.org/10.1007/978-3-642-23702-7_22)