# On the Complexity of Monitoring Orchids Signatures, and Recurrence Equations* 

Jean Goubault-Larrecq ${ }^{1}$ and Jean-Philippe Lachance ${ }^{2}$<br>${ }^{1}$ LSV, ENS Paris-Saclay, CNRS, Université Paris-Saclay<br>ENS Paris-Saclay<br>61, avenue du président Wilson, 94230 Cachan Cedex, France<br>Tel.: +1-47402260<br>Fax: +1-47402464<br>goubault@lsv.fr<br>${ }^{2}$ Coveo Solutions, Inc., Québec City, QC G1W 2K7, Canada jplachance@coveo.com


#### Abstract

Modern monitoring tools such as our intrusion detection tool Orchids work by firing new monitor instances dynamically. Given an Orchids signature (a.k.a. a rule, a specification), what is the complexity of checking that specification, that signature? In other words, let $f(n)$ be the maximum number of monitor instances that can be fired on a sequence of $n$ events: we design an algorithm that decides whether $f(n)$ is asymptotically exponential or polynomial, and in the latter case returns an exponent $d$ such that $f(n)=\Theta\left(n^{d}\right)$. Ultimately, the problem reduces to the following mathematical question, which may have other uses in other domains: given a system of recurrence equations described using the operators + and max, and defining integer sequences $u_{n}$, what is the asymptotic behavior of $u_{n}$ as $n$ tends to infinity? We show that, under simple assumptions, $u_{n}$ is either exponential or polynomial, and that this can be decided, and the exponent computed, using a simple modification of Tarjan's strongly connected components algorithm, in linear time.


## 1 Introduction

Orchids $[13,8]$ is an intrusion detection system. Given a trace $\sigma$ of events, typically obtained in real-time, and a family of so-called signatures (variously otherwise called rules or specifications), Orchids tries to find a subsequence of $\sigma$ that

[^0]satisfies one of the signatures. Each signature is described as an automatonnot a finite-state automaton, though: each state comes with a piece of code, in a simple but expressive imperative language, that is executed whenever control flow enters that state. Orchids then waits for an event matching one of the transitions going out of the state.

See Figure 1 for a slightly edited example of a signature that monitors legal user id (uid) and group id (gid) changes, and reports any system call done with an unexpected uid or gid (at state alert). Events are records with fields such as . syscall or .euid (we have slightly simplified the syntax), and variable names start with a dollar sign. Transitions are introduced by the expect keyword, so, e.g., the start state init has one outgoing transition, and wait has five. Informally, Orchids will wait in state init for an event whose .syscall field is equal to the predefined constant SYS_clone: that is the indication that some process called the Linux system call clone(), creating a new process. Once this is detected, Orchids will then go to state newpid, where it will first record the values of the field .exit in variable \$pid (the process id of the created process), and similarly obtain the values of the user id and group id of the created process. Then it goes to state wait, which is the core monitoring loop: while in state wait, Orchids will wait, concurrently, for five possible conditions, described by so-called expect transitions, on future events. The first three are cases where process number \$pid (as checked by the condition .pid $==\$$ pid, which checks that the current event is one that concerns process \$pid) can legally change its user id or group id: either by calling execve on a executable with the setuid or setgid bit set, or by calling one of the functions setresuid() or setresgid(). The fourth expect transitions catches the case where process \$pid exits, in which case Orchids will cease monitoring it. The last of the transitions of state wait is triggered whenever the process $\$$ pid executes an action with an effective user id .euid that is not the one we expected (in \$uid), or with an unexpected effective group id. Orchids then goes to state alert, where an alert report is produced, and then ceases monitoring that process.

Since Orchids cannot predict which of the five expect transitions will be matched by a subsequent event (and in fact, since, in principle, an event might match several of those transitions), Orchids must monitor all five. To implement that, the Orchids engine simulates so-called threads, and forks a new thread for each pending transition, each new thread waiting for a matching event ${ }^{1}$. That is, on entering state wait, Orchids will create five threads.

This description of the working of Orchids is, of course, oversimplified, but is enough to explain the problem we attack in this paper: evaluating the complexity of detecting a subsequence that matches one of the signatures. A similar question occurs naturally in other modern monitors, such as JavaMOP [11] or the more recent RV-Monitor [12], where signatures are called specifications, and threads are called monitor instances. As the authors argue, and as our own ex-

[^1]```
rule pidtrack {
rule pidtrack {
    expect (.syscall =
    expect (.syscall=
            goto newpid;
    }
    state newpid! {
        $pid = .exit;
        $uid = .euid;
        $gid = .egid;
        goto wait;
    }
    state update_uid_gid! {
        $uid = .euid;
        $gid = .egid;
        goto wait;
    }
    state update_setuid! {
        case (.egid != $gid)
            goto alert;
        else
            goto update_uid_gid;
    }
    state update_setgid! {
        case (.euid != $uid) }
            goto alert;
        else
            goto update_uid_gid;
    }
}
```

```
State wait! {
```

State wait! {
state wait!! {
state wait!! {
state wait! {
state wait! {
state wait!! {
state wait!! {
state wait! {
state wait! {
goto update_uid_gid;
goto update_uid_gid;
goto update_uid_gid;
goto update_uid_gid;
.syscall= =SYS_setresuid)
.syscall= =SYS_setresuid)
goto update_setuid;
goto update_setuid;
expect (.pid = \$pid \&\&
expect (.pid = \$pid \&\&
.syscall =
.syscall =
goto update_setgid;
goto update_setgid;
expect (.pid = \$pid \&\&
expect (.pid = \$pid \&\&
.syscall = SYS_exit)
.syscall = SYS_exit)
goto end;
goto end;
expect (.pid = \$pid \&\&
expect (.pid = \$pid \&\&
expect (.pid = \$pid \&\&
expect (.pid = \$pid \&\&
(.euid != \$uid ||
(.euid != \$uid ||
goto alert;
goto alert;
}
}
state alert! {report(); }
state alert! {report(); }
state end! { }
state end! { }
}

```
}
```





```
|
    *
```

    *
    ```
                    \(\square\)
                    \(\square\)

Figure 1: The pid tracker signature
perience confirms, the main function that has to be estimated is the number of threads that the engine may create after reading \(n\) events.

In the worst case, for a signature \(S\), the Orchids algorithm may create a number of threads \(f_{S}(n)\) that is exponential in \(n\), and that would be untenable. For an intrusion detection system, that would be dangerous, too, as that would open the door to an easy denial-of-service attack on the system itself.

Experience with practical signatures \(S\) shows that \(f_{S}(n)\) is most often a polynomial of low degree. The exponential worst case behavior just means that one could instead craft specific signatures \(S\) such that \(f_{S}(n)\) would be exponential. Most signatures are not of this kind. But can we warn a signature writer of the complexity of his signatures? I.e., how does \(f_{S}(n)\) vary as a function of \(S\) ?

Our main contribution is the design, and proof, of a linear time algorithm that, given a signature \(S\), computes the asymptotic behavior of \(f_{S}(n)\) as \(n\) tends to \(+\infty\). We shall see that \(f_{S}(n)\) is either exponential or polynomial. In the second case, our algorithm computes the unique exponent \(d\) such that \(f_{S}(n)=\Theta\left(n^{d}\right)\).

Our algorithm works in two phases. First, it computes a set of recurrence equations from \(S\), defining sequences \(u_{n}, v_{n}, \ldots\), indexed by \(n \in \mathbb{N}\). Then, and this is the core of our work, it finds the asymptotic behaviors of those sequences. In the example of Figure 1, the recurrence equations are:
\[
\begin{aligned}
& u_{n+1}^{\tau_{\text {Opidtrack }}}=u_{n}^{\tau_{\text {Opidtrack }}}+u_{n}^{\text {init }} \quad \text { one }_{n+1}=\text { one }_{n} \quad u_{n+1}^{\tau_{0}}=u_{n+1}^{\tau_{1}}=\ldots=u_{n+1}^{\tau_{5}}=\text { one }_{n} \\
& u_{n}^{\text {init }}=u_{n}^{\tau_{0}} \quad u_{n}^{\text {wait }}=u_{n}^{\tau_{1}}+\cdots+u_{n}^{\tau_{5}} \\
& u_{n}^{\text {newpid }}=u_{n}^{\text {update_uid_gid }}=u_{n}^{\text {setuid }}=u_{n}^{\text {setgid }}=u_{n}^{\text {alert }}=u_{n}^{\text {end }}=\text { one }_{n} \\
& r_{n}=\max \left(u_{n}^{\tau_{\text {Opidtrack }}}, u_{n}^{\text {init }}, u_{n}^{\text {newpid }}, u_{n}^{\text {update }}{ }_{-}^{\text {uid } \_ \text {gid }}, u_{n}^{\text {setuid }}, u_{n}^{\text {wait }}, u_{n}^{\text {setgid }}, u_{n}^{\text {alert }}, u_{n}^{\text {end }}\right)
\end{aligned}
\]
with the initialization conditions \(u_{0}^{\tau_{\text {Opidtrack }}}=o n e_{0}=u_{0}^{\tau_{0}}=u_{0}^{\tau_{1}}=\ldots=u_{0}^{\tau_{5}}=\) 0 . (We shall explain how those are found in Section 6: to connect with the explanation there, \(\tau_{0}\) is the unique transition out of state init, and \(\tau_{1}, \ldots\), \(\tau_{5}\) are the five transitions out of state wait.) Of the above sequences, \(r_{n}\) is \(f_{\text {pidtrack }}(n)\). Then, our algorithm determines that \(r_{n}=\Theta(n)\), showing that rule pidtrack creates a number of threads that is linear in the number \(n\) of events read in the worst case.

Outline. We review some recent related work in Section 2, describe the Orchids algorithm in Section 3, and define our format of systems of recurrence equations in Section 4. Those systems will be definitions of the sequence \(f_{S}(n)\) by induction on \(n\), together with several other auxiliary sequences, such as \(u_{n}^{\text {newpid }}\) or \(u_{n}^{\tau_{0}}\) above. Any system of recurrence equations \(\Sigma\) in our format gives rise to a graph \(G(\Sigma)\) which will be instrumental in our study, and which we define in Section 5. Translating Orchids signatures to systems of recurrence equations and their associated graph is described in Section 6. The real work begins in Section 7, where we shall examine the possible asymptotic behaviors of our sequences, by looking finely at the structure of the strongly connected components (scc) of \(G(\Sigma)\), carefully distinguishing between trivial and non-trivial
sccs, and so-called cheap and expensive edges. The algorithm quickly follows in Section 8, as an easy adaptation of Tarjan's algorithm, which we describe in full. We also report on our implementation and the result it gives on the ten standard Orchids signatures. We conclude in Section 9.

\section*{2 Related Work}

The question of evaluating the complexity of monitors at this level of detail does not seem to have been addressed already. Efficiency has always been an important subject in the field, and RV-Monitor [12] was recently advocated as a fast implementation of monitors, able to sustain a large number of monitor instances (a.k.a., our threads). This is backed by experimental evidence.

RV-Monitor's algorithm is data-driven. Given a specification with parameters \(x_{1}, x_{2}, \ldots, x_{k}\), RV-Monitor organizes monitor instances inside an indexing tree, and if we agree to call \(N\) the maximal number of different values that parameters can take over an \(n\) event run, there can be at most \(N^{k}\) monitor instances at any given time. If we assume no fixed bound on \(N\), it is however clear that \(N=O(n)\), and that the RV-Monitor analogue of our function \(f_{S}(n)\) above is polynomial in all cases (assuming the specification \(S\) fixed).

The Orchids algorithm is not data-driven, but trace-driven. That is, Orchids does not look merely for values of parameters that make a match, but for a subsequence of the input sequence. This is important for security. See Section 3.1 of [8] for a precise explanation: as we have argued there, Orchids needs to be able to sort matching subsequences (even with different sets of parameter values), so that only the smallest, lexicographically, is eventually reported: this is in most cases the most informative subsequence of events that characterizes a successful attack. Orchids signature matching is therefore necessarily more complex in general, and one can craft signatures that would make Orchids generate exponentially many threads. This is why the algorithm presented here is needed.

Efficiency is also one of the main concerns behind the MonPoly-Reg and MonPoly-Fin tools [4]. The signature language there, MFOTL, is a real-time logic. Each variable varies in a domain of at most \(N\) elements, and it is assumed that there is an upper bound \(m\) on the number of successive events with the same timestamp. Time complexity is always polynomial. By carefully reading Section 5 of [4], one sees that the polynomial degree is linear in the maximal number \(k\) of free variables in the monitored formula and in the number \(c\) of connectives of the formula. When \(m=1\) (i.e., event timestamps are strictly increasing) and there are no temporal future operators, the complexity is comparable with the \(O\left(N^{k}\right)\) bound given for RV-Monitor: if \(t(n)\) is the time to check one RV-Monitor instance, so that RV-Monitor takes time \(O\left(N^{k} t(n)\right)\), the MFOTL-based tools run in time \(O\left(N^{O(k+c)}\right)\).

We will not cite any other paper on the question of monitor complexity. Other (parametric) monitors such as QEA or LogFire are described in [10], with some features in common with Orchids and RV-Monitor respectively. Extra
information can be gleaned by following references from the above papers.
Later, we will argue that our problem reduces to finding asymptotic estimates for sequences \(\left(u_{n}\right)_{n \in \mathbb{N}}\) defined by so-called recurrence equations. Those are (systems of) equations of the form \(u_{n+1}=f\left(u_{n}, v_{n}, w_{n}, \cdots\right), v_{n+1}=g\left(u_{n}\right.\), \(\left.v_{n}, w_{n}, \cdots\right)\), etc., where \(f, g, \ldots\), are some explicitly given functions. There is a huge body of literature, specially in the mathematical literature, on those objects. One of the most relevant source is Flajolet and Sedgwick's book on analytic combinatorics [6]. Unfortunately, we shall need to deal with recurrence equations where \(u_{n+1}\) depend on \(u_{n}, v_{n}, w_{n}\), etc., by using both the + and max operations. The latter seems to be out of scope of what is known in analytic combinatorics.

Estimating the asymptotics of recurrence equations defined using + and max can also be done using the spectral theory of max-plus algebras, see [1]. However, this only handles linear max-plus equations, i.e., equations of the form \(u_{n+1}=\max \left(a_{u u}+u_{n}, a_{u v}+v_{n}, a_{u w}+w_{n}, \cdots\right)\), where \(a_{u u}, a_{u v}, a_{u w}\), \(\ldots\).., are constants. (In max-plus algebra, max takes the rôle of addition and + of multiplication. Hence the latter equation is the max-plus equivalent of \(a_{u u} u_{n}+a_{u v} v_{n}+a_{u w} w_{n}+\cdots\) ) We will allow equations of the form \(u_{n+1}=\) \(\max \left(3+2 u_{n}+v_{n}, 1+u_{n}+w_{n}\right)\), for example, which are not linear, and might be called max-plus polynomial equations. (This example is the max-plus equivalent of \(u_{n+1}=a^{3} u_{n}^{2} v_{n}+a u_{n} w_{n}\), for some unspecified constant \(a\).)

Complexity analysis, and specifically of forms close to our own work, is not limited to monitors. A whole part of the literature is devoted to static analysis of the time and space taken by programs. See [7] for a recent paper on the topic, based on finding upper approximations of solutions to so-called cost equations [2]. Other papers can be found by following references. Probably the piece of work closest to ours is due to Brockschmidt et al. [5], who propose a static analysis framework to infer polynomial upper bounds on the time taken by programs, and (at the same time) on the size of integer values. Both the purpose (complexity evaluation) and the tools of that paper are similar to oursnotably the role of sccs-and this is uncanny: the quantities they evaluate seem to bear no relationship with those we are interested in. Recall that we are interested in evaluating the asymptotic behavior of functions \(f_{S}(n)\) where \(n \in \mathbb{N}\) is some measure of time (number of events received in our case), and which are naturally defined by induction on \(n\). Instead, Brockschmidt et al. compute an over-approximation \(\mathcal{R}(t)(\vec{m})\) of the number of times a given transition \(t\) can be taken in an execution trace that starts with inputs of sizes bounded by the vector of sizes \(\vec{m}\), and to that end they also compute an over-approximation \(\mathcal{S}(t, v)(\vec{m})\) of the size of variable \(v\) after transition \(t\) has just been taken, starting with inputs of sizes bounded by \(\vec{m}\). That is a rather different approach.

As a final note, we have paid special attention to dealing with the issue at the right level of generality. Rather than reporting on a perhaps clever, but highly specialized algorithm devoted to the complexity analysis of a specific tool (Orchids), we identify the core of the problem as finding the asymptotic behavior of sequences defined by recurrence equations. This allows us to study the latter in a proper, well-defined mathematical way, giving not only upper
bounds but also matching lower bounds, and a simple self-contained algorithm that computes them, in linear time. Because it is done at the right level of generality, we believe that the present work may have other uses beyond Orchids.

\section*{3 Orchids}

Let us have a quick look at how Orchids creates and handles threads (a.k.a., monitor instances). We shall ignore most optimizations, both algorithmic and implementation-related, at least those that do not change the worst-case behavior of Orchids.

\subsection*{3.1 The History of Orchids}

Orchids evolved from previous attempts at building intrusion detection systems from a model-checking approach, starting from [16], which presents two approaches. The second approach was a forerunner of the Orchids tool, which was presented in 2005 [13], and whose algorithm and optimizations were described in 2008 [8]. That second approach corrected a few flaws from the first approach. That same first approach (not the second one) is covered by a patent [15], which does not cover Orchids: the main claims of that patent require a means of generating propositional Horn clauses from formulae in a temporal logic for each new event read. The Orchids algorithm is not based on any such mechanism.

We would like to take the 2008 paper [8] as a reference to the Orchids algorithm. However, we need to explain it in a different light in the next subsections, so as to make the complexity analysis clearer.

\subsection*{3.2 A High-Level View}

An Orchids signature consists in finitely many states. We will implicitly refer to Figure 1 to illustrate the notions. One of those states, init, is the initial state. Each state starts with an optional piece of code (for instance, \$pid \(=\). exit; \$uid \(=\).euid; \$gid = .egid; in state newpid), which gets executed on entering the state. That code can contain elementary computations, tests, but no loops: we consider that piece of code irrelevant as far as complexity is concerned.

The second part of the description of a state defines its outgoing transitions, and comes in two flavors. We may either see a block of expect clauses, labeled with conditions that must be satisfied to launch the transition, as in state wait; or a case-delimited multi-way conditional, as in state update_setuid for example (or as in the degenerate case of state update_uid_gid, where there is just one branch, hence no case keyword).

The latter kind of state has an obvious semantics. For example, a thread entering state update_setuid will compare the field .egid with the value of
the variable \(\$ \mathrm{gid}^{2}\), and branch to state alert if they are different (a system call was made with a group id that is not what it was expected to be), or to update_uid_gid otherwise-such transitions were called \(\epsilon\)-transitions in [8].

The former kind of state will wait for a subsequent event matching one of the expect clauses. A same event may match several expect clauses at once, and accordingly Orchids will fork as many threads as needed. We shall ignore the semantics of the tests performed by those clauses, and therefore a state with 5 expect clauses such as wait may fork 5 new threads.

Orchids works slightly differently. Mainly, Orchids threads wait on events, not when they enter a state, rather when they reach an expect clause. That may seem surprising, but look at it this way: expect clauses are the only quiescent places in an Orchids signature, namely the only places where execution does not proceed by itself.

Additional differences occur because of the '!' marker next to the state names, and of the NO_WAIT flag that the Orchids rule compiler uses to decorate transitions, and which is used to avoid forking certain threads while leaving the semantics unchanged. Setting that flag or not is based on the outcome of static analyses, as described in [8]. We will briefly explain those points below.

\subsection*{3.3 The Orchids Algorithm}

Let us give an informal description of how the Orchids algorithm works. We avoid giving a formal description, which would be lengthy and of little interest, and are content with a description that allows us to compute the recurrence equations of Section 6 .

Orchids maintains a queue of threads. Each thread \(\theta\) is a tuple \((\tau, \rho, n)\) where \(\tau\) is an expect transition usually out of some state \(q, \rho\) is an environment, binding variables to values, and \(n\) is a natural number, the thread id. The latter allows us to consider thread groups, namely the collection of threads that have the same thread id. One of the things Orchids will do, at specific points, is committing to a thread: this consists in killing (i.e., removing from the queue) all the threads in the same thread group as the current thread \(\theta\), except \(\theta\) itself.

Committing is done when an Orchids thread enters a state with the '!' marker-so-called commit states. One can see that as a generalization of the final states considered in previous papers on Orchids. In particular, final states are now coded as commit states with no outgoing transition, such as state end or state alert in Figure 1. But it is also interesting to have intermediate commit states, such as states where we can be certain that an attack has been detected, but we wish to proceed with extra actions, typically to analyze what the attacker does once the attack has succeeded (forensics). In general, a commit state \(q\) can be seen as a state that is both final (killing all threads in \(q\) 's thread group) and initial (restarting a thread, by going to state \(q\)-we shall define what "going to" means below).

\footnotetext{
\({ }^{2}\) Variables are thread-local: if an Orchids thread modifies one of its variables, this does not affect any other thread.
}

Any monitoring algorithm needs to delete threads. Apart from the commit mechanism, this is implemented in Orchids by states with no outgoing transition, such as end or state alert. As the following description will make clear, any Orchids thread that enters such a state will just execute the code written there, then be deleted.

Abstracting away from details, we may describe the Orchids algorithm as follows. For each signature \(S\), we assume an additional transition \(\tau_{0 S}\) outside of the signature. Notionally, this transition has the following description:
```

expect (1) goto init;

```
meaning that it will wait for an arbitrary event (1 is always true), and go to the initial state init of \(S\). What is not covered by this description is that going to init will be accompanied with the creation of a new thread group, by changing the thread id to a fresh value. That is unimportant for our purposes here.

The Orchids algorithm starts out by creating one thread waiting on \(\tau_{0 S}\), and with an empty environment. This is used to populate the initial thread queue.

For each new event read \(e\), the Orchids algorithm sweeps through the threads \((\tau, \rho, n)\) in the thread queue. Ordering is important in order to implement the shortest run semantics of [8], and as can be inferred from loc. cit., not an entirely trivial matter. For the purpose of complexity analysis, this is unimportant here again, and we shall ignore this point.

For each thread \((\tau, \rho, n)\) in the thread queue, if the Boolean condition \(b\) labeling the expect clause \(\tau\) is satisfied by \(e\), then Orchids goes to the target state of transition \(\tau\) (relatively to \(\rho, n\) ), possibly inserting new threads into the thread queue. The process of going to a state is described below. If additionally \(\tau\) has the NO_WAIT flag set, then the original thread \((\tau, \rho, n)\) is removed from the queue; otherwise it is kept. If \(b\) is not satisfied, then the thread is kept, unchanged.

This NO_WAIT flag deserves an explanation: a transition of the form expect \((b)\) goto \(q\) will wait for a later event that satisfies condition \(b\), then go to state \(q\). Imagine \(b\) is satisfied at the current event, but also 7 events from now. We must spawn one thread for each of the two possibilities: perhaps the second case will detect a successful run, while the first one will not. This default mechanism is implemented by keeping the original thread \((\tau, \rho, n)\) in the queue, so as to detect such a second case. However, there are situations where one can prove that any such second case will either never lead to a successful run, or will lead to one, but one which will never be shortest. In those situations, which can be detected by an appropriate static analyzer (see [8] again, and replace the notion of final states there by the new notion of commit state), the Orchids compiler flags the transition with NO_WAIT, instructing Orchids to remove ( \(\tau, \rho, n\) ) from the queue at run-time. Then Orchids will simply not look for a second case as above, but that is all right: no shortest run can be found this way, and the shortest run semantics is therefore preserved. In Section 8.2, we shall see an example where this trick decreases the complexity of detection from exponential to polynomial.

We promised we would explain what going to a state \(q^{\prime}\) (relatively to \(\rho, n\) ) would mean. If \(q^{\prime}\) is a commit state, then kill all the threads of the same group,
i．e．，remove all threads with the same \(n\) component from the thread queue， except the current thread．That being done，if \(q^{\prime}\) is a state of the form：

〈code〉
case \(\left(b_{1}\right)\) goto \(q_{1}\) ；
else case \(\left(b_{2}\right)\) goto \(q_{2}\) ；
．．．
else case \(\left(b_{k-1}\right)\) goto \(q_{k-1}\) ；
else goto \(q_{k}\) ；
then run \(\langle\) code \(\rangle\) starting from environment \(\rho\) ，producing a new environment \(\rho^{\prime}\) ， then check the Boolean conditions \(b_{1}, b_{2}, \ldots\) ，in turn in environment \(\rho^{\prime}\)（all Boolean conditions are constrained by the type－checker to be side－effect free）， until one matches，say \(b_{i}\) ，or until the final else clause is reached（in which case we agree that \(i=k)\) ．Then the Orchids algorithm goes to state \(q_{i}\)（relatively to \(\rho^{\prime}, n\) ），recursively．The process must eventually stop：Orchids rejects all signatures that contain cycles of gotos at compile－time；the check is done by a simple graph reachability algorithm．

If \(q^{\prime}\) is a state of the form：
〈code〉
expect（ \(b_{1}\) ）goto \(q_{1}\) ；
expect \(\left(b_{2}\right)\) goto \(q_{2}\) ；
expect \(\left(b_{k}\right)\) goto \(q_{k}\) ；
then run \(\langle\) code \(\rangle\) starting from environment \(\rho\) ，producing a new environment \(\rho^{\prime}\) ． Now do not check any of the conditions \(b_{1}, b_{2}, \ldots, b_{k}\) ，corresponding any of the expect transitions \(\tau_{1}, \tau_{2}, \ldots, \tau_{k}\) out of \(q^{\prime}\) ．Instead，create new threads \(\left(\tau_{i}, \rho^{\prime}, n\right)\) for each \(i, 1 \leq i \leq k\) ，and insert them into the thread queue．

In particular，when \(k=0\)（look at states end and alert in Figure 1），no thread at all will be enqueued，and this is the way threads are eventually deleted．

\section*{4 Systems of Recurrence Equations}

A sequence is an infinite family of natural numbers \(\left(u_{n}\right)_{n \in \mathbb{N}}\) indexed by natural numbers．We say that a property \(P\) holds of \(u_{n}\) for \(n\) large enough if and only if there is an \(n_{0} \in \mathbb{N}\) such that \(P\) holds of \(u_{n}\) for every \(n \geq n_{0}\) ．For a function \(f: \mathbb{N} \rightarrow \mathbb{R}, u_{n}=\Theta(f(n))\) means that there are two real constants \(m, M>0\) such that，for \(n\) large enough，\(m f(n) \leq u_{n} \leq M f(n)\) ．If only the left－hand inequality is assumed，then we write \(u_{n}=\Omega(f(n))\) ，and if only the right－hand inequality is assumed，then we write \(u_{n}=O(f(n))\) ．

We shall say that \(\left(u_{n}\right)_{n \in \mathbb{N}}\) has exponential behavior if and only if \(u_{n}=\Omega\left(a^{n}\right)\) for some constant \(a>1\) ．It has polynomial behavior if and only if \(u_{n}=\Theta\left(n^{k}\right)\) for some constant \(k \in \mathbb{N}\) ．

Let \(Q=\{u, v, \cdots\}\) be a finite non－empty set of symbols．Each symbol \(u \in Q\) is meant to denote a sequence \(\left(u_{n}\right)_{n \in \mathbb{N}}\) of natural numbers．A system of recurrence equations \(\Sigma\) for \(Q\) is，at least informally：
- an initial condition of the form \(u_{0}=a_{u}\), where \(a_{u} \in \mathbb{N} \backslash\{0\}\), one for each \(u \in Q\);
- for each \(u \in Q\), an equation that defines \(u_{n+1}\) in terms of the terms \(v_{n}, v \in Q\), and natural number constants, using the operations max and + . Semantically, since max distributes over + , this means defining \(u_{n+1}\) as \(\max _{i=1}^{m_{u}}\left(\sum_{v \in Q} a_{u i v} v_{n}+b_{u i}\right)\), where \(a_{u i v}\) and \(b_{u i}\) are natural number constants. For reasons explained below, we require \(m_{u} \neq 0\), and for each \(u\) and \(i\), either \(b_{u i} \neq 0\) or \(a_{u i v} \neq 0\) for some \(v \in Q\).

We shall use a slightly different formal definition below (Definition 4.1), and that will be the only authoritative definition for all our mathematical developments, notably in all proofs. We use the above definition for now to illustrate our goals and some inevitable issues, as well as to vindicate the upcoming formal definition itself.

Sticking to the above definition for now, \(\Sigma\) defines a unique family of sequences \(\left(u_{n}\right)_{n \in \mathbb{N}}\), one for each \(u \in Q\), in the obvious way. Our purpose is to show that one can decide, in linear time, which of these sequences have exponential behavior, and which have polynomial behavior; in the latter case, our algorithm will return a natural number \(d\) such that \(u_{n}=\Theta\left(n^{d}\right)\). Note that this will imply that \(\left(u_{n}\right)_{n \in \mathbb{N}}\) has either exponential or polynomial behavior, nothing else e.g., not logarithmic, \(\Theta\left(2^{\sqrt{n}}\right)\) or \(\Theta\left(n^{\log n}\right)\) for example.

Example 4.1 Consider \(Q=\{u\}\), the system \(u_{0}=1, u_{n+1}=2 u_{n}\) is a system of recurrence equations; it defines a unique sequence \(u_{n}=2^{n}\), which has exponential behavior.

Example 4.2 Instead, consider \(Q=\{u, v, w\}\) and the system \(u_{0}=1, v_{0}=1\), \(w_{0}=1, u_{n+1}=v_{n}+1, v_{n+1}=u_{n}+w_{n}, w_{n+1}=w_{n}+2\). Its unique solution is given by \(w_{n}=2 n+1, v_{n}=\frac{1}{2} n^{2}+n+1\) if \(n\) is even, \(v_{n}=\frac{1}{2} n^{2}+n+\frac{1}{2}\) if \(n\) is odd, \(u_{n}=\frac{1}{2} n^{2}+1\) if \(n\) is even, \(u_{n}=\frac{1}{2} n^{2}+\frac{3}{2}\) if \(n\) is odd. In that case, \(w_{n}=\Theta(n), u_{n}=\Theta\left(n^{2}\right), v_{n}=\Theta\left(n^{2}\right)\) all have polynomial behavior.

Notice the slightly oscillating behavior of \(\left(u_{n}\right)_{n \in \mathbb{N}}\) and \(\left(v_{n}\right)_{n \in \mathbb{N}}\). Although those sequences have polynomial behavior, we cannot find an actual, unique polynomial \(p(n)\) such that \(u_{n}=p(n)\) for every \(n \in \mathbb{N}\).

Our recurrence equations have a few constraints attached: \(a_{u}\) is non-zero, \(m_{u}\) is non-zero, and either \(b_{u i} \neq 0\) or \(a_{u i v} \neq 0\) for some \(v \in Q\). This will be the case in all applications. Without this condition, the behaviors of the corresponding sequences might be much wilder, as exemplified below.

Example 4.3 Consider \(Q=\{u, v\}\) with \(u_{0}=0, v_{0}=1, u_{n+1}=2 v_{n}, v_{n+1}=\) \(u_{n}\). This is not a system of recurrence equations in our sense, because the initial value \(a_{u}\) for \(u\) is equal to 0 . Its unique solution is \(u_{n}=0\) if \(n\) is even, \(2^{(n+1) / 2}\) if \(n\) is odd; \(v_{n}=2^{n / 2}\) if \(n\) is even, 0 if \(n\) is odd. Note that \(\left(u_{n}\right)_{n \in \mathbb{N}}\) exhibits neither polynomial nor exponential behavior, as it oscillates between the two. Such a system is forbidden by our definition.

We claimed that, for every vertex \(u \in Q,\left(u_{n}\right)_{n \in \mathbb{N}}\) would either have exponential or polynomial behavior, and that in the latter case, we would be able to find a degree \(d \in \mathbb{N}\) such that \(u_{n}=\Theta\left(n^{d}\right)\). One may wonder whether it would be possible to refine this, and to also find a coefficient \(a\) such that \(u_{n} \sim a n^{d}\) (meaning that \(u /\left(a n^{d}\right)\) would tend to 1 as \(n\) tends to \(\left.+\infty\right)\). This is not possible, as the following example shows.
Example 4.4 Let \(Q=\{u, v, s, t\}\) with \(u_{0}=1, v_{0}=2\), \(s_{0}=1, t_{0}=1\), \(u_{n+1}=\) \(v_{n}, v_{n+1}=u_{n}, s_{n+1}=t_{n}, t_{n+1}=s_{n}+u_{n}\). Its unique solution is: \(u_{n}=1\) if \(n\) is even, 2 if \(n\) is odd; \(v_{n}=2\) if \(n\) is even, 1 if \(n\) is odd; \(s_{n}=n / 2+1\) if \(n\) is even, \(n\) if \(n\) is odd; \(t_{n}=n+1\) if \(n\) is even, \((n+3) / 2\) if \(n\) is odd. Note that both \(s\) and \(t\) exhibit polynomial behavior, as they are \(\Theta(n)\), but we cannot find an a such that \(s_{n} \sim\) an or \(t_{n} \sim\) an: for example, \(s_{n}\) oscillates between \(n / 2+1\) and \(n\).

Let us give a formal definition of systems of recurrence equations. One might do this in the obvious way, using \(Q\) and families of numbers \(a_{u i v}\) and \(b_{u i}\). However, we would also like some equations such as
\[
\begin{equation*}
u_{n+1}=u_{n}+\max \left(v_{n}, \max \left(u_{n}, w_{n}+2\right)+\max \left(2 u_{n}, w_{n}\right)\right) \tag{1}
\end{equation*}
\]
where the max and + operators are freely mixed. Distributing max over + would produce an equivalent system of the right shape, but this transformation takes exponential time and space in the worst case.

Instead, we use the following folklore transform, which works in linear time, at the expense of introducing new symbols to \(Q\) (clearly, only linearly many more). For each non-variable proper subexpression of the term on the right (here, \(u_{n}+\max \left(v_{n}, \max \left(u_{n}, w_{n}+2\right)+\max \left(2 u_{n}, w_{n}\right)\right)\) ), we introduce a fresh symbol. By non-variable we mean any subexpression except the non-constant leaves (here, \(\left.u_{n}, v_{n}, w_{n}\right)\); this includes all non-leaf expressions, such as \(\max \left(u_{n}, w_{n}+2\right)\), and all constant leaves, such as 2 . Let us do so on (1). There are seven nonvariable proper subexpressions there, and we create seven fresh symbols, call them \(a, b, c, d, e, f\) and two. The sequence \(t w o_{n}\) is meant to be the constant sequence equal to 2 , and is defined by \(t w o_{0}=2, t w o_{n+1}=t w o_{n}\). The sequence \(a_{n}\) denotes \(w_{n}+2, b_{n}\) denotes \(\max \left(u_{n}, w_{n}+2\right), c_{n}\) denotes \(2 u_{n}, d_{n}\) denotes \(\max \left(2 u_{n}, w_{n}\right), e_{n}\) denotes \(\max \left(u_{n}, w_{n}+2\right)+\max \left(2 u_{n}, w_{n}\right)\), and \(f_{n}\) denotes \(\max \left(v_{n}, \max \left(u_{n}, w_{n}+2\right)+\max \left(2 u_{n}, w_{n}\right)\right)\). Accordingly, we replace (1) by the following eight equations:
\[
\begin{array}{ccc}
u_{n+1}=u_{n}+f_{n} & f_{n}=\max \left(v_{n}, e_{n}\right) & e_{n}=b_{n}+d_{n} \\
d_{n}=\max \left(c_{n}, w_{n}\right) & c_{n}=2 u_{n} & b_{n}=\max \left(u_{n}, a_{n}\right) \\
a_{n}=w_{n}+t w o_{n} & t w o_{n+1}=t w o_{n} &
\end{array}
\]
plus the initial condition \(t w o_{0}=2\).
Doing so only requires us to be able to state two kinds of recurrence equations: equations of the form \(u_{n+k}=\max \left(v_{n}, w_{n}, \cdots\right)\), and equations of the form \(u_{n+k}=\sum_{v \in Q} a_{u v} v_{n}\), obeying some natural conditions. (Note that constants \(b_{u i}\) have disappeared in the process, being replaced by fresh symbols, such as two in the above example.) This leads us to the following, formal, definition.

Definition 4.1 (System of Recurrence Equations) Let \(Q\) be some set of so-called symbols.
\(A\) recurrence equation on \(Q\) is:
1. either an equation of the form \(u_{n+k}=\max \left(v_{n}, w_{n}, \cdots\right)\), for some nonempty subset of symbols \(v, w, \cdots \in Q\),
2. or an equation of the form \(u_{n+k}=\sum_{v \in Q} a_{u v} v_{n}\), where at least one \(a_{u v}\), \(v \in Q\), is non-zero,
and where in each case \(k=k_{u}\) is equal to 0 or 1 .
\(A\) system of recurrence equations \(\Sigma\) on the set of symbols \(Q\) is a \(Q\)-indexed family of recurrence equations \(E_{u}\), plus initial conditions \(u_{0}=a_{u}\) for each \(u \in Q\) such that \(k_{u} \neq 0\), where \(a_{u} \in \mathbb{N} \backslash\{0\}\) is a constant.

We represent equations \(u_{n+k}=\sum_{v \in Q} a_{u v} v_{n}\) in sparse form, that is, as a list of pairs \(\left(v, a_{u v}\right)\) for each \(v \in Q\) such that \(a_{u v} \neq 0\).

Our formal definition includes strictly more systems than our previous, informal definition. Systems defined per our previous definition always have exactly one solution; in other words, for each \(u \in Q\), they define exactly one sequence \(\left(u_{n}\right)_{n \in \mathbb{N}}\). A contrario, our new definition allows for systems of the form \(u_{n}=u_{n}\), which have infinitely many solutions; or of the form \(u_{n}=u_{n}+\) one \(_{n}\), one \({ }_{n+1}=\) one \(_{n}\), one \(e_{0}=1\), which have no solution. We repair this shortly.

\section*{5 Graphs}

Given a system \(\Sigma\) of recurrence equations on the set of symbols \(Q\), let us define its graph \(G(\Sigma)\) as follows. We write \(s \rightarrow t\) to say there is an edge from \(s\) to \(t\). \(G(\Sigma)\) is a labelled directed graph, and both vertices and edges receive labels. Its vertices are the elements of \(Q\), and are split in two kinds, corresponding to the two kinds of allowed equations:
1. the max vertices \(u\) are those whose associated equation \(E_{u}\) is of the form \(u_{n+k}=\max \left(v_{n}, w_{n}, \cdots\right)\); there is one edge from \(u\) to \(v\), one from \(u\) to \(w\), and so on; \(u\) itself is labeled with \(k\), and the edges receive label 1 ;
2. the plus vertices \(u\) are those whose associated equation \(E_{u}\) is of the form \(u_{n+k}=\sum_{v \in Q} a_{u v} v_{n}\); there is one edge from \(u\) to each \(v \in Q\) such that \(a_{u v} \neq 0\), and it is labeled with \(a_{u v}\); the vertex \(u\) itself is labeled with \(k\);
3. there is no other edge.

Introducing auxiliary symbols as necessary, Example 4.1 is really the system:
\[
u_{n+1}=2 u_{n} \quad u_{0}=1
\]

Its graph is shown on the top left of Figure 2. We distinguish the plus vertices by showing them on a light grey background. We also distinguish the vertex


Figure 2: Three examples of graphs \(G(\Sigma)\)
labels by writing them with a plus sign, viz., +1 , not 1 . The right-hand graph is that of the system of Example 4.2, put into the adequate form:
\[
\begin{aligned}
& u_{n+1}=v_{n}+\text { one }_{n} \quad u_{0}=1 \quad v_{n+1}=u_{n}+w_{n} \quad v_{0}=1 \quad w_{n+1}=w_{n}+t w o_{n} \\
& \text { one }_{n+1}=\text { one }_{n} \quad \text { one }_{0}=1 \quad \text { two }{ }_{n+1}=t w o_{n} \quad t w o_{0}=2 \quad w_{0}=1
\end{aligned}
\]

Similarly for the graph of Example 4.4, shown at the bottom left.
Using the graph \(G(\Sigma)\), we evacuate the problem of those systems \(\Sigma\) that have non-unique solutions, or no solution: we say that \(\Sigma\) is well-formed if and only if there is no cycle in the graph that goes only through vertices labeled +0 .

Proposition 5.1 Every well-formed system \(\Sigma\) has a unique solution, consisting of uniquely-defined sequences \(\left(u_{n}\right)_{n \in \mathbb{N}}\) for each \(u \in Q\), which satisfy all the equations in \(\Sigma\).

Proof. Define a relation \(\succ\) by \(u \succ v\) if and only if there is a non-empty path from \(u\) to \(v\) in \(G(\Sigma)\) whose vertices, including \(u\) but excluding \(v\), are labeled +0 . Write \(v \prec u\) for \(u \succ v\). Since \(\Sigma\) is well-formed, \(\prec\) is irreflexive, hence defines a well-founded strict ordering. Then, \(u_{n}\) is defined by \(\Sigma\), by induction on the pair \((n, u) \in \mathbb{N} \times Q\), ordered by the lexicographic product \((<\times \prec)_{\text {lex }}\) of the usual ordering \(<\) on \(\mathbb{N}\) and of \(\prec\). Since the latter is well-founded, \(u_{n}\) is defined uniquely.
In proofs to come, we shall use the relation \(\succ\) several times, and each time this will be to do well-founded inductions along \((<\times \prec)_{\text {lex }}\). In all the examples shown in Figure 2, \(\succ\) is the trivial relation ( \(u \succ v\) is always false), and such inductions boil down to ordinary inductions on \(n \in \mathbb{N}\).

Note that \(G(\Sigma)\) has no provision for specifying initial conditions such as \(u_{0}=\) 1. They are not needed for Proposition 5.1. They will be useless in subsequent developments as well: the asymptotic behavior of \(u_{n}\) will be independent of \(u_{0}\), provided \(u_{0} \neq 0\).

\section*{6 Generating Recurrence Equations from Orchids Signatures}

Let us describe how we generate the recurrence equations that define the sequences we are interested in from a given Orchids signature. The translation takes linear time in the size of the signature, defined as the sum of the number of its states and its transitions.

We first define the set \(Q\) of symbols, that is, the names \(u\) of sequences \(\left(u_{n}\right)_{n \in \mathbb{N}}: r\) for the final complexity of signature \(S, u^{q}\) for each state \(q, u^{\tau}\) for each expect transition \(\tau\) (including \(\tau_{0 S}\) ), one for a special sequence defining the constant 1. For the latter, generate the recurrence equations:
\[
\begin{equation*}
\text { one }_{n+1}=\text { one }_{n} \tag{2}
\end{equation*}
\]

The initialization condition \(\left(o n e_{0}=1\right)\) is irrelevant to the asymptotic behavior of sequences. In fact, they are discarded when we construct the graph of our system of recurrence equations. We shall therefore omit them.

The meaning of symbol \(u^{\tau}\) is as follows. Imagine the Orchids algorithm is launched on a single thread \((\tau, \rho, i d)\) (with the same \(\tau\) as in \(u^{\tau}\) ), and feed it \(n\) events. Now count how many threads have been created by Orchids once it has read and processed those \(n\) events: \(u_{n}^{\tau}\) is designed so as to be a precise upper bound on that number. The symbol \(u^{q}\) has a similar meaning, starting from state \(q\) instead of transition \(\tau\).

In each state, for each transition \(\tau\), say with target state \(q^{\prime}\), we create the following recurrence equations (3)-(8). If \(q^{\prime}\) is not a commit state, and if \(\tau\) has the NO_WAIT flag set, then we generate:
\[
\begin{equation*}
u_{n+1}^{\tau}=u_{n}^{q^{\prime}} \tag{3}
\end{equation*}
\]

Indeed, given \(n+1\) events, Orchids will simply read the first one, trigger the transition and go to state \(q^{\prime}\), where \(n\) events will remain to be read.

If \(q^{\prime}\) is not a commit state and if \(\tau\) does not have the NO_WAIT flag set, then we generate instead:
\[
\begin{equation*}
u_{n+1}^{\tau}=u_{n}^{\tau}+u_{n}^{q^{\prime}} \tag{4}
\end{equation*}
\]
witnessing the fact that we also keep the current thread, waiting on transition \(\tau\), so that all threads descending from \(\tau\) count (and are accounted through the term \(u_{n}^{\tau}\) ).

If \(q^{\prime}\) is a commit state, then, as we have said earlier, we can consider \(q^{\prime}\) to be a final state, and accordingly we write:
\[
\begin{equation*}
u_{n+1}^{\tau}=\text { one }_{n} \tag{5}
\end{equation*}
\]
if \(\tau\) has the NO_WAIT flag, and:
\[
\begin{equation*}
u_{n+1}^{\tau}=u_{n}^{\tau}+\text { one }_{n} \tag{6}
\end{equation*}
\]
otherwise. We also need to consider \(q^{\prime}\) as an initial state, and this will be taken care of in Equation (9) below.

For each state \(q\) with \(k\) outgoing expect transitions \(\tau_{1}, \tau_{2}, \ldots, \tau_{k}\), we generate:
\[
\begin{equation*}
u_{n}^{q}=u_{n}^{\tau_{1}}+u_{n}^{\tau_{2}}+\cdots+u_{n}^{\tau_{k}} \tag{7}
\end{equation*}
\]

For each state \(q\) with outgoing case transitions, enumerate the target states that are not commit states as \(q_{1}, q_{2}, \ldots, q_{k}\), and generate:
\[
\begin{equation*}
u_{n}^{q}=\max \left(o n e_{n}, u_{n}^{q_{1}}, u_{n}^{q_{2}}, \cdots, u_{n}^{q_{k}}\right) \tag{8}
\end{equation*}
\]

Indeed all commit states only contribute one thread, hence the maximum with one \(_{n}\).

Finally, we write:
\[
\begin{equation*}
r_{n}=\max \left(u_{n}^{\tau_{0 S}}, u_{n}^{q_{1}}, u_{n}^{q_{2}}, \cdots, u_{n}^{q_{m}}\right) \tag{9}
\end{equation*}
\]
where \(q_{1}, q_{2}, \ldots, q_{m}\) are all the commit states of \(S\). This is meant to take into account the fact that each commit state is not only a final state, but also an initial state, and may therefore contribute to \(r_{n}\).

Equation (9) is the only equation presented here whose soundness is not obvious. Indeed, what we have is that the maximum number of threads \(s_{n}\) created by signature \(S\) is the maximum of \(u_{n}^{\tau_{0 S}}\), and of quantities \(u_{n-j}^{q_{i}}\), for \(1 \leq i \leq m\), and \(j\) varying over \(0 \cdots n\) : \(j\) is, intuitively, the event number at which state \(q_{i}\) is entered. We shall see that the asymptotic behavior of those sequences are either exponential or polynomial. For \(i\) fixed, \(u_{n-j}^{q_{i}}\) is an \(\Omega\left(a^{n-j}\right)\) for some \(a>1\), or a \(\Theta\left((n-j)^{d}\right)\) for some degree \(d \in \mathbb{N}\), as we have seen. When \(j\) varies, \(a^{n-j}\) (resp., \((n-j)^{d}\) ) reaches its maximum when \(j\) is minimal. We see that the minimal value of \(j\) is the length of the shortest series of transitions leading from state init to state \(q_{i}\) in \(S\), which is a constant. Since \(a^{n-j}=\Theta\left(a^{n}\right)\) and \((n-j)^{d}=\Theta\left(n^{d}\right), s_{n}\) is, in any case, a \(\Theta\) of \(r_{n}\), as described in (9).

As a final note to this section, observe that all the recurrence equations described here are well-formed in the sense of Proposition 5.1.

The reader is invited to practice, to generate the equations corresponding to the signature of Figure 1, and, finally, to check that those are exactly those given at the end of the introduction. To this end, one needs to know that the Orchids compiler manages to set the NO_WAIT flag on every transition.

We shall give more details, on practical examples, in Section 8.2. The equations we have just described are an upper bound on the actual complexities we are interested in, and a natural question is whether they incur an acceptable loss of precision, compared to the complexities experienced in practice. This is a legitimate concern, and perhaps the most pressing source of worry is as follows. Monitors not only create, but also remove threads (monitor instances) at run-time, whereas our equations only seem to upper-bound the number of created threads, ignoring deletions. This would be understanding our algorithm wrongly. The main mechanism Orchids uses to delete threads is entering committed states. In that case, recall that all threads in the current thread group
except the current thread are killed. Dealing with that case is the purpose of equation (5), (6), and (9). Our complexity evaluation mechanism simulates that bulk removal by considering that all threads of the same group are removed in one go - equations (5), (6) - and that one is born again, as described in equation (9). Hence that particular form of deletion is, in fact, correctly handled. We shall discuss precision issues in more detail Section 8.2, resting on the examples given there.

\section*{7 Sccs, and asymptotics}

We shall see that the key to understanding the asymptotic behavior of sequences defined by a well-formed system of recurrence equations \(\Sigma\) lies in the strongly connected components of the graph \(G(\Sigma)\), introduced in Section 5 .

We fix a well-formed system \(\Sigma\) of recurrence equations for the rest of the section, as well as its set of symbols \(Q\), and the unique sequences \(\left(u_{n}\right)_{n \in \mathbb{N}}\) that it defines. The following trivial lemma is crucial.

Lemma 7.1 Assume two vertices \(u, v\) in \(Q\) such that \(v\) is reachable from \(u\), namely, such that there is a path from \(u\) to \(v\) in \(G(\Sigma)\). There is a constant \(k \in \mathbb{N}\) such that, for every \(n \in \mathbb{N}, u_{n+k} \geq v_{n}\).

More precisely, \(k\) can be taken as the sum of vertex labels on any given path from \(u\) to \(v\), including \(u\) but excluding \(v\).

Proof. The key argument is that for every edge of the form \(s \rightarrow t\), where \(s\) is labeled \(a, s_{n+a} \geq t_{n}\) for every \(n \in \mathbb{N}\). This holds whether \(s\) is a max or a plus vertex.

Before we start doing any real proof, we must mention that the pathological case of a behavior that is neither polynomial nor exponential (see Example 4.3) will not happen. This will rest on the following casual-looking observation, which we shall need in the proof of Proposition 7.5.

Lemma 7.2 For every \(u \in Q\), for every \(n \in \mathbb{N}, u_{n} \geq 1\).
Proof. By well-founded induction on \((n, u)\) along \((<\times \prec)_{\text {lex }}\).
If \(u\) is a max vertex labeled +0 , then \(u_{n}=\max \left(v_{n}, w_{n}, \cdots\right)\) where \(v, w\), \(\ldots\), form a non-empty subset of \(Q\). Note that \(u \succ v, w, \cdots\) By induction hypothesis, \(v_{n} \geq 1, w_{n} \geq 1\), and so on. Since they form a non-empty subset, \(\max \left(v_{n}, w_{n}, \cdots\right) \geq 1\), so \(u_{n} \geq 1\).

If \(u\) is a max vertex labeled +1 , then either \(n=0\) and \(u_{0}=a_{u} \geq 1\), or \(u_{n}=\max \left(v_{n-1}, w_{n-1}, \cdots\right) \geq 1\) by a similar argument as above (but using the fact that \(n>n-1\), not \(u \succ v, w, \cdots\) as a reason to invoke the induction hypothesis).

If \(u\) is a plus vertex labeled +0 , then \(u_{n}=\sum_{v \in Q} a_{u v} v_{n}\), and some \(a_{u v}\) is non-zero. For this \(v, u \succ v\), so by induction hypothesis \(v_{n} \geq 1\). We now have \(u_{n} \geq a_{u v} v_{n} \geq 1\).


Figure 3: Sccs in a graph \(G(\Sigma)\)

If \(u\) is a plus vertex labeled +1 , then either \(n=0\) and \(u_{0}=a_{u} \geq 1\), or \(u_{n}=\sum_{v \in Q} a_{u v} v_{n-1} \geq 1\), by a similar argument as above.

A subset \(A\) of vertices is strongly connected if and only if every vertex from \(A\) is reachable from any other vertex of \(A\). The maximal strongly connected subsets are called the strongly connected components of the graph, and we abbreviate that as scc. They partition the graph, and in particular every vertex \(u\) belongs to a unique scc, which we write \(\operatorname{scc}(u)\).

On Figure 2, the top left graph is its unique scc. The right-hand graph has four sccs, the top cycle \(\{u, v\}\), the middle cycle \(\{w\}\), and the two cycles \(\{o n e\}\) and \(\{t w o\}\). The bottom left graph has two sccs, the top cycle \(\{s, t\}\), and the bottom cycle \(\{u, v\}\).

In general, an scc can be more complex than a mere cycle, but if one wants to picture a non-trivial scc, a cycle is a good first approximation. Sccs can also be trivial, i.e., consist of no cycle at all, just a single vertex with no self-loop. Figure 3 displays a more complex graph, with sccs shown as darker gray rectangles. (We haven't shown any vertex labels on this example, as they would distract us somehow. You can put any vertex labels so long as no cycle goes only through vertices labeled +0 , for example by labeling each vertex +1 .) The topmost scc is an example of an scc that is not just a cycle. There are also trivial sccs: \(\{c\}\) and \(\{f\}\). We let the reader reconstruct a system of recurrence equations associated with this graph. Note that, contrarily to previous examples, this one involves the max operator in order to define \(\left(u_{n}\right)_{n \in \mathbb{N}},\left(v_{n}\right)_{n \in \mathbb{N}},\left(s_{n}\right)_{n \in \mathbb{N}}\), \(\left(t_{n}\right)_{n \in \mathbb{N}},\left(w_{n}\right)_{n \in \mathbb{N}},\left(r_{n}\right)_{n \in \mathbb{N}}\). Although this may seem like a complicated graph,
it will follow from our algorithm that \(u_{n}=\Theta(n), a_{n}=\Theta(n), b_{n}=\Theta(n)\) and \(v_{n}=\Theta(n)\), and all other vertices have constant behavior.

Definition 7.3 (Bad Vertex) Call a vertex \(u \in Q\) bad if and only if it is a plus vertex, and given its associated equation \(u_{n+k}=\sum_{v \in Q} a_{u v} v_{n}\), the sum of the coefficients \(a_{u v}\) where \(v\) ranges over \(\operatorname{scc}(u)\) is at least 2 .

Equivalently, \(u\) is bad if and only if it is a plus vertex, and at least one of the following possibilities occurs:
1. there is an edge \((u, v)\) of label at least 2 to a vertex \(v\) in the same scc as \(u\),
2. or there are at least two edges \(u \rightarrow v\) and \(u \rightarrow w\) to vertices \(v\) and \(w\) that are both in the same scc as \(u\).

In Example 4.1, \(u\) is bad. There is no bad vertex in Example 4.2 or in Example 4.4. There is no bad vertex either in Figure 3: although there are several edges with label at least 2, they all go out of their start scc.

Definition 7.4 (Bad* Vertex) Say that a vertex \(u \in Q\) is bad* if and only if some bad vertex is reachable from it, namely if and only if there is a path from \(u\) to some bad vertex \(v\).

We shall see that the bad \(^{*}\) vertices \(u\) are exactly those such that \(\left(u_{n}\right)_{n \in \mathbb{N}}\) has exponential behavior. In Example 4.1, \(u\) is bad*. There is no bad* vertex in Example 4.2 or in Example 4.4, or in Figure 3.

Proposition 7.5 For every bad* vertex \(u\) in \(Q,\left(u_{n}\right)_{n \in \mathbb{N}}\) has exponential behavior.

Proof. We shall show that this is the case if \(u\) is \(\operatorname{bad}\). If \(u\) is \(\operatorname{bad}^{*}\), then there is a path from \(u\) to some bad vertex \(v\), so that \(\left(v_{n}\right)_{n \in \mathbb{N}}\) will have exponential behavior. By Lemma 7.1, for some \(k \in \mathbb{N}\), for every \(n \in \mathbb{N}, u_{n+k} \geq v_{n}\), so that \(\left(u_{n}\right)_{n \in \mathbb{N}}\) will also have exponential behavior.

So let us assume that \(u\) is bad. In particular, \(u\) is a plus vertex, and we consider two cases.

Case 1. If there is an edge \(u \rightarrow v\) with label \(a_{u v}\) at least 2 for some \(v\) in the same scc as \(u\), then there is also a path \(\pi\) from \(v\) to \(u\). Concatenating \(\pi\) with the edge \(u \rightarrow v\), we obtain a cycle \(u \rightarrow v \xrightarrow{\pi} u\). Let \(a\) be the label of \(u\), and \(k\) be the sum of the vertex labels on \(\pi\), including \(v\) but excluding \(u\). Since \(\Sigma\) is well-formed, \(k+a \geq 1\). By the defining equation for \(u, u_{n+k+a} \geq a_{u v} v_{n+k} \geq\) \(2 v_{n+k}\), and by Lemma \(7.1, v_{n+k} \geq u_{n}\), for every \(n \in \mathbb{N}\). By induction on \(p\), it follows that \(u_{(k+a) p+q} \geq 2^{p} u_{q}\) for all \(p, q \in \mathbb{N}\), and since \(u_{0}, u_{1}, \ldots, u_{k+a-1} \geq 1\) (Lemma 7.2), \(u_{(k+a) p+q} \geq 2^{p}\) for all \(p, q \in \mathbb{N}\). This implies that \(u_{n} \geq m a^{n}\) for every \(n \in \mathbb{N}\), where \(a={ }^{k+a} \sqrt{2}\) and \(m=1 / a^{k+a-1}\). Hence \(\left(u_{n}\right)_{n \in \mathbb{N}}\) has exponential behavior in this case.

Case 2. If instead there are at least two edges \(u \rightarrow v\) and \(u \rightarrow w\) to vertices \(v\) and \(w\) that are both in the same scc as \(u\), then there are two cycles \(u \rightarrow v \xrightarrow{\pi_{1}} u\)
and \(u \rightarrow w \xrightarrow{\pi_{2}} u\). Let \(k_{1}\) be the sum of the vertex labels on \(\pi_{1}\), including \(v\) but excluding \(u\), and similarly for \(k_{2}\) and \(\pi_{2}\). Let also \(a\) be the vertex weight of \(u\). Since \(\Sigma\) is well-formed, \(k_{1}+a\) and \(k_{2}+a\) are both larger than or equal to 1. Using Lemma 7.1, we obtain that for every \(n \in \mathbb{N}, v_{n+k_{1}+a} \geq u_{n+a} \geq v_{n}\) and \(w_{n+k_{2}+a} \geq u_{n+a} \geq w_{n}\). If \(k_{1}\) and \(k_{2}\) were equal, we could reuse the same argument as in Case 1... but we have to work a bit more.

By induction on \(p_{1}\), it follows that \(v_{n+p_{1}\left(k_{1}+a\right)} \geq v_{n}\) and similarly, \(w_{n+p_{2}\left(k_{2}+a\right)} \geq\) \(w_{n}\), for all \(p_{1}\) and \(p_{2}\) in \(\mathbb{N}\). Pick \(p_{1}\) and \(p_{2}\) so that \(\left(p_{1}+1\right)\left(k_{1}+a\right)=\left(p_{2}+1\right)\left(k_{2}+a\right)\) (by taking least common multiples), and let \(k\) be equal to the latter minus \(a\). Then \(v_{n+k}=v_{n+p_{1}\left(k_{1}+a\right)+k_{1}} \geq v_{n+k_{1}}\); by Lemma 7.1 applied to \(\pi_{1}, v_{n+k_{1}} \geq u_{n}\), so \(v_{n+k} \geq u_{n}\). Similarly, \(w_{n+k} \geq u_{n}\). By the defining equation of \(u, u_{n+k+a} \geq\) \(a_{u v} v_{n+k}+a_{u w} w_{n+k} \geq v_{n+k}+w_{n+k} \geq 2 u_{n}\). We obtain the same inequality \(u_{n+k+a} \geq 2 u_{n}\) as in Case 1. It follows that \(\left(u_{n}\right)_{n \in \mathbb{N}}\) has exponential behavior again, by the same reasoning.

If an scc contains a bad* vertex, then all its vertices are bad*. Let us consider the case of sccs \(A\) without any bad* vertex. We shall illustrate the various cases we need to consider on the graph shown in Figure 3. The idea of our algorithm is that we shall iterate on all sccs \(A\), from bottom to top, deducing a characteristic degree \(d_{A}\) such that \(u_{n}=\Theta\left(n^{d_{A}}\right)\) for every vertex \(u\) in \(A\) from the characteristic degrees of sccs below \(A\).

We first deal with the case of trivial sccs, i.e., sccs with only one vertex and no self-loop. By abuse of language, say that \(u\) is a trivial scc iff \(\{u\}\) is.

Proposition 7.6 Assume \(u \in Q\) is a trivial scc, and that for every edge \(u \rightarrow v\), the sequence \(\left(v_{n}\right)_{n \in \mathbb{N}}\) has polynomial behavior, viz., for some \(d_{v} \in \mathbb{N}\), \(v_{n}=\) \(\Theta\left(n^{d_{v}}\right)\). Then \(\left(u_{n}\right)_{n \in \mathbb{N}}\) has polynomial behavior, too, and \(u_{n}=\Theta\left(n^{d_{u}}\right)\), where \(d_{u}=\max \left\{d_{v} \mid v \in Q\right.\) such that \(\left.u \rightarrow v\right\}\).

Proof. This follows from the fact that \(\max \left(\Theta\left(n^{d_{v}}\right), \Theta\left(n^{d_{w}}\right), \cdots\right)=\Theta\left(n^{\max \left(d_{v}, d_{w}, \cdots\right)}\right)\) for max vertices, and that \(\sum_{v \in Q \text { such that } u \rightarrow v} a_{u v} \Theta\left(n^{d_{v}}\right)=\Theta\left(n^{\max \left\{d_{v} \mid u \rightarrow v\right\}}\right)\) for plus vertices, using the fact that \(a_{u v} \geq 1\) for every \(v \in Q\) such that \(u \rightarrow v\).

Now assume \(A\) is a non-trivial scc without any bad vertex. Such an scc must have a special shape, exemplified on the three sccs of the graph on the right of Figure 2, or on the bottom left graph of the same figure, or on the non-trivial sccs of Figure 3: the weights of all edges between vertices of \(A\) must be equal to 1 , and every plus vertex has exactly one successor in \(A\) (all others are outside \(A\) ). That it has at most one successor in \(A\) (and that all its edge labels are equal to 1) is a consequence of the absence of bad vertices. That it has at least one follows from the fact that \(A\) is non-trivial.

Note that plus vertices in \(A\) may have more than one successor; but only one can be in \(A\). For example, \(v\) has two successors in the graph of Example 4.2 (Figure 2, right), but only \(u\) is in the scc \(A=\{u, v\}\) that \(v\) belongs to. Example 4.4 displays a similar situation. Figure 3 does, too: \(b\) has two successors, but only one in its own scc.

Definition 7.7 (Expensive, cheap edges) Say that an edge \(u \rightarrow v\) goes out of \(A\) if and only if \(u \in A\) and \(v \notin A\). If \(u\) is a plus vertex, then we say that it is an expensive edge out of \(A\), otherwise it is a cheap edge out of \(A\).

In Figure 3, there is only one expensive edge out of a non-trivial scc, namely the edge \(b \rightarrow w\), with label 1. There are two other edges going out of non-trivial sccs, namely \(u \rightarrow c\) and \(w \rightarrow f\). They are both cheap. Note that being cheap or expensive is entirely independent of the label it carries. Here the expensive edge has the lowest possible label of the whole graph.

The key argument in Proposition 7.10 below consists in noting that \(\sum_{j=0}^{n} j^{N}=\Theta\left(n^{N+1}\right)\) when \(n\) tends to \(+\infty\)-a wellknown identity. The way we use that is probably best explained on a small test case. Imagine one of the simplest possible nontrivial sccs: a one-vertex loop, as shown on the right, and assume that we know that \(v_{n}=\Theta\left(n^{N}\right)\), and even, to make things simpler, that \(v_{n}=n^{N}\).
 The equation defining \(\left(u_{n}\right)_{n \in \mathbb{N}}\) is \(u_{n+1}=u_{n}+v_{n}=u_{n}+n^{N}\). It follows easily that \(u_{n}=(n-1)^{N}+(n-2)^{N}+\cdots+1^{N}+0^{N}+u_{0}=u_{0}+\sum_{j=0}^{n-1} j^{N}\), and that is \(\Theta\left(n^{N+1}\right)\), as we have just seen. Note that the edge \(u \rightarrow v\) is expensive. The name "expensive" was chosen so as to suggest that those are the edges that are responsible for the increase in the exponent, from \(N\) to \(N+1\). Cheap edges will incur no such increase.

In general, the proof of the bounds in the following proposition uses a similar argument. The case of non-trivial sccs \(A\) without any bad vertex, and with only + vertices (no max vertex) is easier to argue. There is a single elementary cycle from \(u\) back to \(u\), and all its weights are equal to 1 ; all edges out of this path go out of \(A\). Semantically \(u_{n+a}\) can be written as \(u_{n-k}\) (for some constant \(k\) ) plus values \(v_{n}\), for \(v\) outside of \(A\). If \(v_{n}\) has polynomial growth, viz. if \(v_{n}=\Theta\left(n^{d_{v}}\right)\), then \(n^{d_{v}}\) will enter a summation defining \(u_{n+a}\), and a similar argument as above shows that \(u_{n}\) will grow as a polynomial in \(n\) whose degree is the largest exponent \(d_{v}\) plus 1 . The fact that + and max vertices can be used freely inside \(A\) makes the proof of the Proposition 7.10 below slightly more involved.

The key argument in Proposition 7.10 consists in the following well-known bounds. We take the standard convention that \(j^{N}=1\) when \(N=0\), even when \(j=0\). The proof is a standard high school exercise.

Lemma 7.8 Let \(N\) be any non-negative real constant. For every \(p \in \mathbb{N}\),
\[
\frac{1}{N+1} p^{N+1} \leq \sum_{j=0}^{p} j^{N} \leq \frac{1}{N+1}(p+1)^{N+1}
\]

Proof. If \(N>0\), by elementary analysis, \(\sum_{j=0}^{p} j^{N} \leq \int_{0}^{p+1} x^{N} d x=\frac{1}{N+1}(p+\) \(1)^{N+1}\), and \(\sum_{j=0}^{p} j^{N}=\sum_{j=1}^{p} j^{N} \geq \int_{0}^{p} x^{N} d x=\frac{1}{N+1} p^{N+1}\). When \(N=0\), the claim reduces to the inequalities \(p \leq p+1 \leq p+1\), which are obvious.

We will use the left-hand inequality of Lemma 7.8 to prove lower bounds, but we will not use the right-hand inequality for upper bounds. Regarding the latter, we will find it easier to rely on the following inequality-one that can be used to reprove the right-hand inequality of Lemma 7.8, using a telescoping sum, by the way.

Lemma 7.9 Let \(N\) be any non-zero natural number. For every \(x \geq 0, x^{N}+\) \(N x^{N-1} \leq(x+1)^{N}\).

Proof. By the binomial formula, \((x+1)^{N}=\sum_{k=0}^{N}\binom{N}{k} x^{k}=x^{N}+\binom{N}{N-1} x^{N-1}+\) \(\binom{N}{N-2} x^{N-2}+\cdots \geq x^{N}+\binom{N}{N-1} x^{N-1}\) (since \(N \geq 1\), the sum contains at least those first two terms). We conclude since \(\binom{N}{N-1}=N\).

Proposition 7.10 Let \(A\) be a non-trivial scc of \(G(\Sigma)\) without a bad vertex. For each edge \(u \rightarrow v\) going out of \(A\), assume that \(\left(v_{n}\right)_{n \in \mathbb{N}}\) has polynomial behavior. Precisely, assume that \(v_{n}=\Theta\left(n^{d_{v}}\right)\).

Then every vertex of \(A\) has polynomial behavior, with the same degree \(d_{A}\), where \(d_{A}\) is the maximum of:
- all the quantities \(d_{v}\), where \(u \rightarrow v\) ranges over the cheap edges out of \(A\),
- and all the quantities \(d_{v}+1\), where \(u \rightarrow v\) ranges over the expensive edges out of \(A\).

To be completely formal, we agree that the maximum of an empty set of numbers is 0 ; this is needed in case there is no edge going out of \(A\) at all.

Proof. We show that, for every \(u \in A, u_{n}=\Omega\left(n^{d_{A}}\right)\) and that for every \(u \in A, u_{n}=O\left(n^{d_{A}}\right)\), separately.

Lower bounds. Consider any edge going out of \(A\). If this is a cheap edge \(u \rightarrow v\), with \(u\) a max vertex labeled \(a\), then \(u_{n+a} \geq v_{n}\). By assumption, \(v_{n}=\Omega\left(n^{d_{v}}\right)\), hence \(u_{n}=\Omega\left(n^{d_{v}}\right)\). By Lemma 7.1, for every \(t \in A\), there is a constant \(k \in \mathbb{N}\) such that for every \(n \in \mathbb{N}, t_{n+k} \geq u_{n}\), whence \(t_{n}=\Omega\left(n^{d_{v}}\right)\) as well.

If this is an expensive edge \(u \rightarrow w\), then \(u\) is a plus vertex labeled \(a\), and \(u_{n+a}=\sum_{v \in Q} a_{u v} v_{n}\). Among the summands with \(a_{u v} \neq 0\), one is obtained by choosing \(v=w\), because of the existence of the edge \(u \rightarrow w\); and exactly one is a vertex \(v \in Q\) that is in \(A\)-recall the special shape of non-trivial sccs without a bad vertex. Since \(v \in A\) and \(w \notin A, v\) and \(w\) are distinct. It follows that \(u_{n+a} \geq v_{n}+w_{n}\) for every \(n \in \mathbb{N}\). Since \(A\) is an scc, \(u\) is reachable from \(v\). By Lemma 7.1, there is a \(k \in \mathbb{N}\) such that for every \(n \in \mathbb{N}, v_{n+k} \geq u_{n}\). Moreover, since \(\Sigma\) is well-formed, \(k+a \geq 1\). We have obtained: (a) \(u_{n+k+a} \geq u_{n}+w_{n+k}\) for every \(n \in \mathbb{N}\). Since \(w_{n}=\Omega\left(n^{d_{w}}\right)\), there is a constant \(n_{0} \in \mathbb{N}\), and a real constant \(m>0\) such that, for every \(n \geq n_{0}, w_{n} \geq m n^{d_{w}}\). Without loss of generality, we may assume \(n_{0} \geq k+a\). For \(n\) large enough, we may write \(n-a\) uniquely as \(p(k+a)+q\) for some \(p \in \mathbb{N}, q<k+a \leq n_{0}\). Assuming \(n\) sufficiently large, \(p\) will be non-zero. Using \((a), u_{n} \geq w_{p(k+a)+q}+w_{(p-1)(k+a)+q}+\cdots+w_{j(k+a)+q}+\cdots+\) \(w_{q} \geq m \sum_{j=0}^{p}(j(k+a)+q)^{d_{w}}\). Since \(q \geq 0, u_{n} \geq m(k+a)^{d_{w}} \sum_{j=0}^{p} j^{d_{w}}\), and this
is larger than or equal to \(\frac{m(k+a)^{d_{w}}}{d_{w}+1} p^{d_{w}+1}\) by Lemma 7.8. Since \(p \geq \frac{n-a}{k+a}-1\), \(u_{n} \geq \frac{m(k+a)^{d_{w}}}{d_{w}+1}\left(\frac{n-a}{k+a}-1\right)^{d_{w}+1}\), and this is a \(\Theta\left(n^{d_{w}+1}\right)\), since \(m, k, a\), and \(d_{w}\) are constants. Finally, as for cheap edges, we use Lemma 7.1 to conclude that \(t_{n}=\Omega\left(n^{d_{w}+1}\right)\) for every \(t \in A\).

For each \(t \in A\), we have therefore obtained that \(t_{n}=\Omega\left(n^{d_{w}+1}\right)\) for every expensive edge \(u \rightarrow w\) out of \(A\), and that \(t_{n}=\Omega\left(n^{d_{v}}\right)\) for every cheap edge \(u \rightarrow v\) out of \(A\). Since \(d_{A}\) is the largest of those exponents, \(t_{n}=\Omega\left(n^{d_{A}}\right)\).

Upper bounds. There is a constant \(n_{0} \in \mathbb{N}\), and constants \(M_{v}>0\), such that, for every \(n \geq n_{0}\), for every edge \(u \rightarrow v\) out of \(A, v_{n} \leq M_{v}(n+1)^{d_{v}}\). (We take \(n+1\) instead of \(n\) to avoid problems with the case \(n=0\).) Notice that we pick the same \(n_{0}\) for all edges. That does not restrict generality, since we can always take the maximum of all values of \(n_{0}\) collected for each edge.

It will be important to realize what the defining equation is for \(u_{n}\), where \(u\) is a plus vertex in \(A\). In principle, letting \(a\) be the label of \(u\), this is \(u_{n+a}=\) \(\sum_{v \in Q} a_{u v} v_{n}\). Since \(a_{u v} \neq 0\) only when there is an edge \(u \rightarrow v\), this is \(u_{n+a}=\) \(\sum_{v \in Q / u \rightarrow v} a_{u v} v_{n}\). Since \(u\) cannot be bad by assumption, there is a unique \(w \in Q\) such that \(u \rightarrow w\) and for which \(w\) is also in \(A\); moreover, \(a_{u w}=1\). (That this \(w\) exists is because \(A\) is a non-trivial scc.) It follows that the defining equation for such a vertex is of the form \(u_{n+a}=w_{n}+T_{n}\), where \(T_{n}\) is the sum over all (expensive) edges \(u \rightarrow v\) out of \(A\) of \(a_{u v} v_{n}\).

We first deal with the case where \(d_{A}=0\). In that case, there cannot be any expensive edge out of \(A\), and for every cheap edge \(u \rightarrow v\) out of \(A, d_{v}=0\). Let \(B\) be the maximum of all the constants \(M_{v}\) given above, when \(u \rightarrow v\) ranges over all the cheap edges out of \(A\), and of all the values \(u_{n_{0}}, u \in A\). We claim that, for every \(u \in A\), for every \(n \geq n_{0}, u_{n} \leq B(n+1)^{d_{A}}\), namely that \(u_{n} \leq B\). This is by induction on \((n, u)\) with \(n \geq n_{0}\), ordered by \((<\times \prec)_{\text {lex }}\). For \(n=n_{0}\), \(u_{n_{0}} \leq B\) by the definition of \(B\). For \(n>n_{0}\), either \(u\) is a max node (labeled \(a\), say), and \(u_{n}=\max \left(v_{n-a}, w_{n-a}, \cdots\right) \leq B\) by induction hypothesis; or \(u\) is a plus node, but since there is no expensive edge out of \(A\), the term \(T_{n}\) alluded to above is zero, so the defining equation for \(u_{n}\) is \(u_{n}=w_{n-a}\) for some unique vertex \(w\) in \(A\); by induction hypothesis again, this is at most \(B\).

We now deal with the more complicated case \(d_{A} \geq 1\). We deal with that case separately, because we shall need to refer to \(d_{A}-1\), which would not make sense if \(d_{A}=0\).

We shall show that \(u_{n} \leq B(n+1)^{d_{A}}+h(u)(n+1)^{d_{A}-1}\) for every \(n \geq n_{0}\) and \(u \in A\). Here \(B\) and \(h(u)\), for each \(u \in A\), will be non-negative constants that are defined below. What they should be is best found by examining the proof below, and collecting the constraints that they need to obey.

Solving those constraints leads us to define \(h(u)\) by induction along \(\prec\) by:
1. if \(u\) is labeled +1 , then \(h(u)=0\);
2. if \(u\) is a max vertex labeled +0 , then \(h(u)=\max \{h(v) \mid u \rightarrow v, v \in A\}\);
3. if \(u\) is a plus vertex labeled +0 , then recall that the defining equation for \(u\) is \(u_{n}=w_{n}+T_{n}\), where \(w\) is in \(A\) and \(T_{n}=\sum_{v} a_{u v} v_{n}\), where
the sum is taken over all expensive edges \(u \rightarrow v\) out of \(A\) : then we let \(h(u)=h(w)+\sum_{v} a_{u v} M_{v}\).

And we require \(B\) to be so large that:
(a) \(u_{n_{0}} \leq B\left(n_{0}+1\right)^{d_{A}}+h(u)\left(n_{0}+1\right)^{d_{A}-1}\) for every \(u \in A\);
(b) \(B \geq M_{v}\) for every cheap edge \(u \rightarrow v\) out of \(A\);
(c) for every plus vertex \(u\) labeled +1 in \(A, h(w)+\sum_{v} a_{u v} M_{v} \leq d_{A} B\), where \(w\) is the unique vertex in \(A\) such that \(u \rightarrow w\), and the sum is taken over all expensive edges \(u \rightarrow v\) out of \(A\);
(d) for every max vertex \(u\) labeled +1 in \(A\), for every edge \(u \rightarrow v\) with \(v\) in \(A, h(v) \leq d_{A} B\).

Note indeed that all those inequalities are eventually satisfied as \(B\) tends to \(+\infty,(a)\) because \(n_{0}+1>0,(c)\) and \((d)\) because \(d_{A}>0\).

We now prove that \(u_{n} \leq B(n+1)^{d_{A}}+h(u)(n+1)^{d_{A}-1}\) for every \(n \geq n_{0}\) by induction on \((n, u)\) with \(n \geq n_{0}\), ordered by \((<\times \prec)_{\text {lex }}\). The case \(n=n_{0}\) is by (a).

Let us now assume \(n>n_{0}\).
If \(u\) is a plus vertex, labeled \(a\), then \(u_{n}=w_{n-a}+T_{n-a}\) where \(w \in A\) and \(T_{n-a}=\sum_{v} a_{u v} v_{n-a}\), where the sum is taken over all expensive edges \(u \rightarrow v\) out of \(A\). By assumption, \(v_{n-a} \leq M_{v}(n-a+1)^{d_{v}}\) for all the latter vertices \(v\), hence \(v_{n-a} \leq M_{v}(n-a+1)^{d_{A}-1}\). By induction hypothesis, \(w_{n-a} \leq B(n-a+\) \(1)^{d_{A}}+h(w)(n-a+1)^{d_{A}-1}\). If the label \(a\) is +0 , we obtain \(u_{n} \leq B(n+1)^{d_{A}}+\) \(h(w)(n+1)^{d_{A}-1}+\sum_{v} a_{u v} M_{v}(n+1)^{d_{A}-1}=B(n+1)^{d_{A}}+h(u)(n+1)^{d_{A}-1}\) (item 3 of the definition of \(h\) ). If the label \(a\) is +1 , we obtain instead \(u_{n} \leq B n^{d_{A}}+\) \(h(w) n^{d_{A}-1}+\sum_{v} a_{u v} M_{v} n^{d_{A}-1} \leq B n^{d_{A}}+d_{A} B n^{d_{A}-1}\), where the latter inequality is by \((c)\); the last term is less than or equal to \(B(n+1)^{d_{A}}\) by Lemma 7.9, and we therefore obtain \(u_{n} \leq B(n+1)^{d_{A}}+h(u)(n+1)^{d_{A}-1}\) since \(h(u)\) is non-negative. (In fact, \(h(u)=0\), see item 1 of the definition of \(h\). .)

If \(u\) is a max vertex, labeled \(a\), then \(u_{n}=\max \left(v_{n-a}, w_{n-a}, \cdots\right)\). The terms \(v_{n-a}\) obtained from (necessarily cheap) edges \(u \rightarrow v\) out of \(A\) are less than or equal to \(M_{v}(n-a+1)^{d_{v}} \leq M_{v}(n-a+1)^{d_{A}} \leq B(n-a+1)^{d_{A}} \quad\) (by \((b)\) ) \(\leq B(n-a+1)^{d_{A}}+h(u)(n-a+1)^{d_{A}-1}\), since \(h(u)\) is non-negative. Whichever the value of \(a,+0\) or +1 , that is less than or equal to \(B(n+1)^{d_{A}}+h(u)(n+1)^{d_{A}-1}\). The terms \(v_{n-a}\) such that \(v \in A\) are less than or equal to \(B(n-a+1)^{d_{A}}+\) \(h(v)(n-a+1)^{d_{A}-1}\) by induction hypothesis. If \(a=+0\), then this is less than or equal to \(B(n+1)^{d_{A}}+h(u)(n+1)^{d_{A}-1}\), using the fact that \(h(v) \leq h(u)\) in this case (item 2 of the definition of \(h\) ). If \(a=+1\), then this is less than or equal to \(B n^{d_{A}}+h(v) n^{d_{A}-1} \leq B n^{d_{A}}+d_{A} B n^{d_{A}-1}\) by \((d)\). That is less than or equal to \(B(n+1)^{d_{A}}\) by Lemma 7.9 , hence to \(B(n+1)^{d_{A}}+h(u)(n+1)^{d_{A}-1}\), again. It follows that \(u_{n} \leq B(n+1)^{d_{A}}+h(u)(n+1)^{d_{A}-1}\).

Let us use those results to determine the asymptotic behavior of all the sequences defined by the graph of Figure 3. Recall that there is no bad* vertex in that example. We start from the sccs at the bottom, and work our way up:
- The non-trivial sccs \(\{r\}\) and \(\{s, e, t, d\}\) have no outgoing edge at all, hence their associated sequences are \(\Theta\left(n^{0}\right)\) (bounded from below and from above by constants).
- The scc \(\{f\}\) is trivial, and its two successors behave as \(\Theta\left(n^{0}\right)\), hence it itself behaves as \(\Theta\left(n^{0}\right)\), by Proposition 7.6.
- Similarly for the trivial scc \(\{c\}\).
- The scc \(\{w\}\) is not trivial, but it does not have any expensive edge out of it; by Proposition 7.10, it also behaves as \(\Theta\left(n^{0}\right)\).
- The topmost scc is non-trivial, it has one cheap outgoing edge, \(u \rightarrow c\), and one expensive outgoing edge \(b \rightarrow w\). By applying Proposition 7.10, all the sequences associated with vertices in that scc behave as \(\Theta\left(n^{\max (0,0+1)}\right)\), that is, \(\Theta(n)\).

\section*{8 The Algorithm}

To conclude, we need a final, standard ingredient: the condensation of a directed graph \(G\) is the graph whose vertices are the sccs of \(G\), and such that there is an edge from \(A\) to \(B\) if and only if there are vertices \(q \in A\) and \(r \in B\) and an edge \(q \rightarrow r\) in \(G\). The condensation is always acyclic, meaning that working our way up, that is, from the leaves to the roots of the condensation, must terminate (hence terminate in a linear number of steps). In the case of a graph \(G\) of the form \(G(\Sigma)\), we shall say that an edge \(A \rightarrow B\) as above in the condensation is expensive if and only if \(A\) is non-trivial, and we can find a \(q \in A\) and an \(r \in B\) such that the edge \(q \rightarrow r\) is expensive in \(G\); it is cheap otherwise, namely when \(A\) is trivial, or \(A\) is non-trivial and all the edges \(q \rightarrow r\) in \(G\) with \(q \in A\) and \(r \in B\) are cheap.

Theorem 8.1 Given any system \(\Sigma\) of recurrence equations with set \(Q\) of symbols, we can compute a table of numbers \(d_{u} \in \mathbb{N} \cup\{+\infty\}, u \in Q\), in linear time, such that \(d_{u}=+\infty\) iff \(\left(u_{n}\right)_{n \in \mathbb{N}}\) has exponential behavior, and otherwise \(u_{n}=\Theta\left(n^{d_{u}}\right)\).

The algorithm works as follows:
1. Compute \(G(\Sigma)\) and its sccs, building its condensation \(G^{\prime}\).
2. Traverse \(G^{\prime}\) in reverse topological order (i.e., from the bottom up). For each visited scc \(A\), decide whether \(A\) contains a bad vertex. If so, let \(d_{A}:=\) \(+\infty\). Otherwise, for every successor \(B\) of \(A\) in \(G^{\prime}, d_{B}\) has already been computed, and let \(d_{A}:=\max \left(\max _{A \rightarrow B}\right.\) cheap \(d_{B}, \max _{A \rightarrow B \text { expensive }}\left(d_{B}+\right.\) \(1)\) ), where by convention we agree that the maximum of the empty set is zero.
3. Finally, for each \(u \in Q\), let \(d_{u}:=d_{\operatorname{scc}(u)}\).
function \(\operatorname{SCC}(v:\) vertex \()\)
        \(\operatorname{PUSH}(v)\)
        for each successor \(w\) of \(v\) do
            if \(w\).index is undefined then
                \(\operatorname{scc}(w)\)
                if \(w\).low \(<v\).low then
                    \(v\).low \(\leftarrow w\).low
                else if \(w\).onStack then
                    if \(w\).index \(<v\).low then
                    \(v\).low \(\leftarrow w\).index
                end if
                end if
        else
        end if
    end for
    if \(v\).low \(=v\).index then
                \(\triangleright\) pop the scc with root v from the stack
            \(l \leftarrow \mathrm{~S} \quad \triangleright\) sever scc from rest of stack
        \(\mathrm{S} \leftarrow v\).next \(\quad \triangleright l=\) linked list of nodes in scc
        \(v\). next \(\leftarrow\) NULL \(\quad \triangleright\) last node in \(l\) is \(v\)
        \(w \leftarrow l \quad \triangleright\) now loop over scc \(l\)
        while \(w \neq\) NULL do
            \(w\).onStack \(\leftarrow\) false
            \(w\).sccRoot \(\leftarrow v\)
            \(w \leftarrow w\).next
        end while
        COMPUTEComplexities \((l)\)
    end if
end function

Figure 4: Tarjan's algorithm

The correctness of the algorithm is a direct consequence of Proposition 7.5, Proposition 7.6, and Proposition 7.10. That it works in linear time is easy. Notably, the second phase sweeps through all the sccs \(A\) once, and for each, takes time proportional to the number of vertices in \(A\) plus the number of edges that go out of \(A\). The sum over all sccs \(A\) of those values is the size of \(G(\Sigma)\).

In practice, this is implemented by simply modifying Tarjan's scc algorithm [17]. In any description of that algorithm, there is a single line of code where it has just found an scc \(A\), and it must emit it by repeatedly popping a stack. It is enough to compute \(d_{A}\) there, by the formula given in Theorem 8.1, item 2, knowing that at that point, all the values \(d_{B}\) will have been computed earlier.

\subsection*{8.1 A Concrete Implementation}

We make the algorithm more explicit. Figure 4 displays one possible implementation of Tarjan's algorithm. (Comments start with \(\triangleright\).) That algorithm works with the help of a global natural number index, initialized to 0 , and a global stack S, which we implement as a linked list of vertices S, S.next, S.next.next, and so on until we reach the special constant NULL, denoting the end of the list. We assume that each vertex \(v\) of the graph is a record containing the following fields: onStack, a Boolean flag which is true if and only \(v\) is on the
stack; next, used for linking purposes, and in the first place to implement the stack itself; and type, a flag with only two possible values, PLUS or MAX depending on whether \(v\) is a + or max vertex. Initially, \(v\). onStack is false, and \(\mathrm{S}=\) NULL. The vertex \(v\) also contains the following fields, which will be filled in by the algorithm: index (depth-first search index), low (lowest depth-first search index of vertices in the same scc), sccRoot (root vertex of \(v\) 's scc), and degree, all initially undefined. The latter field will eventually contain either a natural number - the degree of the polynomial if the sequence associated with \(v\) has asymptotic polynomial behavior-or the special constant INFTY, denoting exponential behavior. We agree that INFTY is strictly larger than any natural number, and that INFTY \(+1=\) INFTY.

Edges are left implicit, and are handled by loops of the form "for each successor \(w\) " in the code. To obtain the desired degrees, we launch SCC on the root \(r\) of our graph; when it returns, we read off the degree field from \(r\).sccRoot.

The algorithm of Figure 4 is not our contribution. Our contribution is the function computeComplexities, called at the end of the loop in Tarjan's algorithm. At this point, \(l\) is guaranteed to hold a list of all the vertices in the just discovered scc, with root \(v\), that list has been severed off the stack and has \(v\) as last entry, and for each vertex \(w\) in \(l, w . \operatorname{sccRoot}=v\).

The function computeComplexities is shown in Figure 5, right, and follows our informal description: if the scc \(l\) is trivial, then its degree, stored in the degree field of its root \(l\).sccRoot, is the maximum of all degrees of the successor sccs of \(l\); otherwise, its degree is the maximum of the same degrees (computed as variable \(d v\) ), possibly incremented by 1 in case of expensive edges. The helper functions TRIVIALSCC and BADVERTEX determine whether their argument is a trivial scc, resp., a bad vertex.

\subsection*{8.2 Experimental Results}

We have implemented that algorithm inside Orchids, and its ten standard signatures. Execution time was negligible. We had the pleasant surprise of observing that all our signatures had polynomial thread complexity, confirming our intuition that human experts do not write signatures with exponential behavior.

The largest observed complexity is \(\Theta\left(n^{3}\right)\) for lin24_ptrace.rule, a signature that attempts to detect the ptrace attack [14]. The second largest is \(\Theta\left(n^{2}\right)\) for the apachessl.rule, a signature that tries to correlate abnormal variations in message entropy [9] with specific failure events from the Apache server. Out of our ten signatures, seven others have linear behavior, including the pid tracker of Figure 1. We have instrumented our algorithm so that it reports the main causes of complexity. For example, on the pid tracker, our algorithm reports:
```

rule pidtrack may have worst case linear behavior, i.e., O(\#events).
each event may fork a new thread going to 'init'.

```

And indeed, each newly created Unix process (through the clone system call, a Linux abstraction behind the more well-known fork call) may cause the creation of a new Orchids thread for that signature, starting at state init.
```

function TRIVIALScC(l:vertex)
if l.next }\not=\mathrm{ NULL then
| if there is another vertex
in the scc,
then it is not trivial
return false
end if
for each successor w of l do
if}w=l\mathrm{ then
if there is a self-loop,
then the scc
is not trivial
return false
end if
end for
return true
end function
function BADVERTEX(w:vertex)
if w.type }\not=\mathrm{ PLUS then
\triangleright ~ M A X ~ v e r t i c e s ~ a r e ~ n e v e r ~ b a d
return false
end if
v\leftarroww.sccRoot
weight}\leftarrow
for each edge w\xrightarrow{}{a}z do
if z.sccRoot =v then
weight }\leftarrow\mathrm{ weight +a
end if
end for
return weight \geq2
end function

```
function TRIVIALScC(l:vertex)
if \(l\).next \(\neq\) NULL then
\(\triangleright\) if there is another vertex \(\triangleright\) in the scc,
\(\triangleright\) then it is not trivial return false
end if
each successor \(w\) of \(l\) do
\(\triangleright\) if there is a self-loop,
\(\triangleright\) then the scc
\(>\) is not trivial
return false
end if
end for
return true
end function
function \(\operatorname{badV} \operatorname{ErtEx}(w\) :vertex)
if \(w\).type \(\neq\) PLUS then return false
end if
\(v \leftarrow w\).sccRoot
for each edge \(w \xrightarrow{a} z\) do
if \(z\).sccRoot \(=v\) then
weight \(\leftarrow\) weight \(+a\)
end if
end for
end function
function Compute Complexities(l:vertex)
\(v \leftarrow l\).sccRoot
\(v\).degree \(\leftarrow 0\)
if trivialScc \((l)\) then
\(\triangleright\) take max of degrees of successors for each successor \(w\) of \(l\) do
if \(w\).sccRoot.degree \(>v\).degree then
\(v\).degree \(\leftarrow w\).sccRoot.degree
end if
end for
else \(\quad \triangleright\) loop over \(w\) in scc \(l\)
\(w \leftarrow l\)
while \(w \neq\) NULL do
if badVertex \((w)\) then
\(\triangleright\) exponential behavior
\(v\).degree \(\leftarrow\) INFTY
break \(\quad\) exit while loop else
for each successor \(z\) of \(w\) do
if \(z\).sccRoot \(\neq v\) then
\(d v \leftarrow z\).sccRoot.degree if \(w\).type \(=\) PLUS then
\(\triangleright\) expensive edge
\(d v \leftarrow d v+1\)
end if
\(\triangleright\) take \(\max (d v, v\).degree \()\)
if \(d v>v\).degree then
\(v\).degree \(\leftarrow d v\) end if
end if
end for
end if
\(w \leftarrow w\).next
end while
end if
end function

Figure 5: Computing complexities

The last of the ten signatures, taint_auditd.rule, a tainting mechanism for detecting illegal transitive information flows, is not even flagged by our algorithm: it is correctly classified as generating a constant number of threads.

Let us focus on lin24_ptrace.rule in order to understand what is going on in the worst experienced case. We will not describe the actual signature: see [13], where we used a version of that rule as a running example.
\(\Theta\left(n^{3}\right)\) is an overestimation: the actual complexity of that signature is \(\Theta\left(n^{2}\right)\). The source of the overestimation lies in the following piece of Orchids signature:
```

state ptrace_poketext \{
\$counter $=$ \$counter +1 ;
expect (.pid =\$attack_pid \&\&
.syscall =SYS_ptrace \& $\&$
. ptrace_req = POKETEXT \&\&
. ptrace_pid = \$target_pid \&\&
\$counter $<10$ )
goto ptrace_poketext;
expect (.pid =\$attack_pid \&\&
.syscall =SYS_ptrace \& $\&$
. ptrace_req = DETACH \&\&
.ptrace_pid = \$target_pid)
goto ptrace_detach;
\}

```

Let p be the state ptrace_poketext, and d be the state ptrace_detach (not shown). Our complexity analyzer correctly determines that \(u_{n}^{\mathrm{d}}=\Theta(1)\), and produces the equations:
\[
u_{n}^{\mathrm{p}}=u_{n}^{\tau_{1}}+u_{n}^{\tau_{2}} \quad u_{n+1}^{\tau_{1}}=u_{n}^{\mathrm{p}} \quad u_{n+1}^{\tau_{2}}=u_{n}^{\mathrm{d}}
\]
where \(\tau_{1}\) and \(\tau_{2}\) are the two expect transitions. Both have the NO_WAIT flag, as determined by the Orchids signature compiler. Note that, otherwise, the equation defining \(u^{\tau_{1}}\) would have to be \(u_{n+1}^{\tau_{1}}=u_{n}^{\tau_{1}}+u_{n}^{\mathrm{p}}\), in which case \(u_{n}^{\tau_{1}}\) and \(u_{n}^{\mathrm{p}}\) would behave as \(\Theta\left(2^{n}\right)\). This dramatically confirms the claim made in [8] that optimizations found by static analysis are crucial: here they allow us to replace an exponential behavior by a polynomial behavior, preserving the semantics.

The above equations yield an asymptotic estimate for \(u_{n}^{\mathrm{p}}\) of \(\Theta(n)\). There are two expensive edges (not shown) between the root node and vertex \(u^{\mathrm{P}}\), leading to an estimated complexity of \(\Theta\left(n^{3}\right)\) for the whole rule. This result is not optimal: as one can see from the use of the variable \(\$\) counter in state \(p\), and since \$counter starts at 0 , the loop on that state can only be taken at most 10 times. A more precise analysis, taking that into account, would generate the following equation for \(u_{n}^{\mathrm{p}}\) :
\[
u_{n+10}^{\mathrm{p}}=\sum_{i=0}^{9} u_{n+i}^{\mathrm{d}}
\]
which would give a more precise estimate \(u_{n}^{\mathrm{p}}=\Theta(1)\), hence an overall complexity of \(\Theta\left(n^{2}\right)\). That would be doable, given some additional static analyses.

There is a final source of imprecision in our complexity estimation algorithm. Among the static analyses that Orchids implements, the monotonicity analysis allows Orchids to remove threads that are waiting on conditions \(b\) that have failed once, provided we can show that \(b\) is antitonic [8, Section 5]. Consider an expect transition of the form:
```

expect (.syscall= SYS_open \&\&
.time <= $start+$delay) goto q;

```
which one would write to detect a call to the open() system call within \$delay time units, assuming \$start holds the current time. The .time field is trusted to evolve in a non-decreasing way as new events flow in. Then, the Boolean condition .time \(<=\$\) start + \$delay is antitonic in the sense that its value can only decrease (go from true to false), and never increase. Therefore, if it fails once, it will fail forever. An Orchids thread that is waiting on the above expect transition is then reclaimed if .time \(<=\$\) start \(+\$\) delay is found to be false just once. Contrarily to the NO_WAIT flag, Orchids discovers that it can remove the corresponding thread at run-time, not at compile-time. As a consequence, that is ignored by our complexity analyzer, which operates entirely at compiletime: it computes an upper bound on the number of created threads, but ignores deletions (except for the thread killing mechanism arising from commit states, which can be predicted at compile-time).

Should we refine our complexity analyzer to handle such thread deletions? Certainly, if the monotonicity analysis is ever useful, we should do so. We should start by evaluating the gap between the observed number of active threads and the number predicted by our complexity analyzer. However, and perhaps sadly so, the monotonicity analysis may be of limited use. Of the signatures we used, only two would have benefited from that mechanism: the ssh_failed_burst rule and the ssh_failed_long_window rule, which detect bursts of ssh failures (more than \(N\) failed connection attempts during some fixed amount of time, where \(N\) is a constant). The author of these rules, Baptiste Gourdin, preferred to write them in a different style. To explain that on the example of the above expect transition, he would have written it as:
```

expect (.syscall = SYS_open) goto q;
expect (.time > $start+$delay) goto stop;

```
where stop is a commit state. That actually implements the desired timeout by relying on commit states instead: if .time ever exceeds \$start+\$delay, then Orchids will enter state stop, and that will kill the thread waiting on the condition .syscall==SYS_open, among others.

\section*{9 Conclusion}

We have described, and proved, a linear time algorithm that decides the asymptotic complexity of sequences defined by certain forms of systems of recurrence
equations, using both the + and max operators. Our goal was to analyze, automatically, which Orchids signatures have polynomial detection complexity (in terms of numbers of created Orchids threads), and with which exponent.

This turns out to be an extremely reliable and useful tool to Orchids signature writers. Personal experience shows that a high degree in the polynomial, or worse, an estimation of exponential complexity, is indicative of a mistake in the writing of the signature.

Beyond Orchids, it seems obvious that our simple algorithm for estimating the asymptotic complexity of recurrence equations should find applications outside of security or of runtime verification. Mounir Assaf recently proposed a (yet unpublished) static analysis that detects whether leakage of sensitive data in security programs is negligible or not [3]. This is based on estimating the rate of growth of a sequence \(u_{n}\) as the number of steps taken, \(n\), tends to infinity, and we hope that our algorithm, or similar techniques, apply.

\section*{Acknowledgement}

The first author would like to thank Mounir Assaf for drawing his attention to analytic combinatorics, and the anonymous referees for their suggestions.

\section*{References}
[1] Akian M, Bapat R, Gaubert S (2006) Max-plus algebras. In: Hogben L (ed) Handbook of Linear Algebra, Discrete Mathematics and Its Applications, vol 39, Chapman and Hall/CRC, chap 25
[2] Albert E, Arenas P, Genaim S, Puebla G (2011) Closed-form upper bounds in static cost analysis. Journal of Automated Reasoning 46(2):161-203
[3] Assaf M (2015) From qualitative to quantitative program analysis : Permissive enforcement of secure information flow. PhD thesis, Université Rennes I
[4] Basin D, Klaedtke F, Müller S, Zălinescu E (2015) Monitoring metric firstorder temporal properties. Journal of the Association for Computing Machinery \(62(2): 15: 1-15: 45\)
[5] Brockschmidt M, Emmes F, Falke S, Fuhs C, Giesl J (2014) Alternating runtime and size complexity analysis of integer programs. In: Proc. Intl. Conf. Tools and Algorithms for the Construction and Analysis of Systems (TACAS'14), Springer Verlag Lecture Notes in Computer Science 8413, vol 8413
[6] Flajolet P, Sedgwick R (2009) Analytic Combinatorics. Cambridge University Press, iSBN 0521898064. ISBN-13 978-0521898065
[7] Flores-Montoya A, Hähnle R (2014) Resource analysis of complex programs with cost equations. In: Proc. 12th Asian Symposium on Programming Languages and Systems (APLAS'14), Singapore, Singapore, Springer Verlag Lecture Notes in Computer Science 8858
[8] Goubault-Larrecq J, Olivain J (2008) A smell of Orchids. In: Leucker M (ed) Proceedings of the 8th Workshop on Runtime Verification (RV'08), Springer, Budapest, Hungary, Lecture Notes in Computer Science, vol 5289, pp 1-20, DOI 10.1007/978-3-540-89247-2_1
[9] Goubault-Larrecq J, Olivain J (2013) On the efficiency of mathematics in intrusion detection: The NetEntropy case. In: Danger JL, Debbabi M, Marion JY, Garcia-Alfaro J, Zincir-Heywood N (eds) Revised Selected Papers of the 6th International Symposium on Foundations and Practice of Security (FPS'13), Springer, La Rochelle, France, Lecture Notes in Computer Science, vol 8352, pp 3-16, DOI 10.1007/978-3-319-05302-8_1
[10] Havelund K, Reger G (2015) Specification of parametric monitors - quantified event automata versus rule systems. In: Drechsler R, Kuhne U (eds) Formal Modeling and Verification of Cyber-Physical Systems, Springer Verlag, pp 151-189, 1st International Summer School on Methods and Tools for the Design of Digital Systems (SyDe), Bremen, Germany
[11] Jin D, O'Neil Meredith P, Lee C, Roşu G (2012) JavaMOP: Efficient parametric runtime monitoring framework. In: Proceeding of the 34th International Conference on Software Engineering (ICSE'12), IEEE, pp 1427-1430, DOI 10.1109/ICSE.2012.6227231
[12] Luo Q, Zhang Y, Lee C, Jin D, O'Neil Meredith P, Şerbănuţă TF, Roşu G (2014) RV-Monitor: Efficient parametric runtime verification with simultaneous properties. In: Bonakdarpour B, Smolka SA (eds) Proceedings of the 5th International Conference on Runtime Verification (RV'14), Springer Verlag LNCS 8734, Toronto, ON, CA, pp 285-300
[13] Olivain J, Goubault-Larrecq J (2005) The Orchids intrusion detection tool. In: Etessami K, Rajamani S (eds) Proceedings of the 17th International Conference on Computer Aided Verification (CAV'05), Springer, Edinburgh, Scotland, UK, Lecture Notes in Computer Science, vol 3576, pp 286-290, DOI 10.1007/11513988_28
[14] Purczyński W (2003) Linux kernel privileged process hijacking vulnerability. http://www.securityfocus.com/bid/7112, bugTraq Id 7112. Last read: september, 2003
[15] Roger M, Goubault-Larrecq J (1999) Procédé et dispositif de résolution de modèles, utilisation pour la détection des attaques contre les systèmes informatiques. Dépôt français du 13 sep. 1999, correspondant Dyade, demandeurs : 1. INRIA 2. Bull S.A. Numéro de publication: 2798490. Numéro d'enregistrement national: 99 11716. Classification: G 06 F 19/00.

Date de mise à la disposition du public de la demande: 16 mars 2001, bulletin 01/11.
[16] Roger M, Goubault-Larrecq J (2001) Log auditing through model checking. In: Proceedings of the 14th IEEE Computer Security Foundations Workshop (CSFW'01), IEEE Computer Society Press, Cape Breton, Nova Scotia, Canada, pp 220-236
[17] Tarjan RE (1972) Depth-first search and linear graph algorithms. SIAM Journal on Computing 1(2):146-160```


[^0]:    *Partially funded by INRIA-DGA grant 12810312 (2013-2016). The second author also thanks Hydro-Québec and Les Offices jeunesse internationaux du Québec (LOJIQ) for their financial support.

[^1]:    ${ }^{1}$ To dispel a possible misunderstanding, these threads are not system-level threads, rather pairs of a pending transition and an environment holding the values of variables, and which are handled and scheduled by the Orchids engine. Similarly, "forking" an Orchids thread simply means making a copy of the current thread and adding it to the Orchids thread queue.

