New Transience Bounds for Max-Plus Linear Systems

Bernadette Charron-Bost[1] Matthias Függer[2] Thomas Nowak[3]

[1] CNRS, LIX, École polytechnique, 91128 Palaiseau, France
[2] ECS Group, TU Wien, 1040 Wien, Austria
[3] LIX, École polytechnique, 91128 Palaiseau, France

Abstract

Linear max-plus systems describe the behavior of a large variety of complex systems. It is known that these systems show a periodic behavior after an initial transient phase. Assessment of the length of this transient phase provides important information on complexity measures of such systems, and so is crucial in system design. We identify relevant parameters in a graph representation of these systems and propose a modular strategy to derive new upper bounds on the length of the transient phase. By that we are the first to give asymptotically tight and potentially subquadratic transience bounds. We use our bounds to derive new complexity results, in particular in distributed computing.

Keywords: transience bounds; max-plus systems; cyclic scheduling; network synchronizers; link-reversal algorithms

1 Introduction

The behavior of many complex systems can be described by a sequence of N-dimensional vectors $x(n)$ that satisfy a recurrence relation of the form

$$\forall n \geqslant 1 \ \ \forall i \in \{1, \dots, N\} \quad : \quad x_i(n) = \max_{j \in \mathbf{N}_i} \bigl(x_j(n-1) + A_{i,j} \bigr) \tag{1}$$

where the $A_{i,j}$ are real numbers, and the $\mathbf{N}_i$ are subsets of $\{1, \dots, N\}$. For instance, $x_i(n)$ may represent the time of the nth occurrence of a certain event i and the $A_{i,j}$ the required time lag between the $(n-1)$th occurrence of j and the nth occurrence of i. Notable examples are transportation and automated manufacturing systems [18, 11, 14], network synchronizers [24, 16], and cyclic scheduling [19]. Recently, Charron-Bost et al. [9] have shown that the relation (1) also encompasses the behavior of an important class of distributed algorithms, namely *link reversal algorithms* [17], which can be used to solve a variety of problems [29] like routing [17], scheduling [3], distributed queuing [28, 1], or resource allocation [7].

Recurrences of the form (1) are actually linear in the *max-plus algebra* (e.g., [21]). An analog of the Perron-Frobenius theorem in max-plus linear algebra [10] states that the sequence of powers of an irreducible max-plus matrix becomes periodic after a finite index called the *transient* of the matrix. As an immediate corollary, any linear max-plus system with irreducible system matrix is periodic from some index, called the *transient* of the system, which clearly depends on the system's initial vector and is at most equal to the transient of the matrix of the system. For all the above mentioned applications, the study of the transient plays a key role in characterizing the system performances: For example, in the case of link reversal routing, the system transient is equal to the time complexity of the routing algorithm. Besides that, understanding matrix and system transients is of interest on its own for the theory of max-plus algebra.

Hartmann and Arguelles [20] have shown that the transients of matrices and linear systems are computable in polynomial time. However, their algorithms provide no analysis of the transient phase, and do not hint at the parameters that influence matrix and system transients. Conversely, upper bounds involving these parameters help to predict the duration of the transient phase, and to define strategies to reduce transients during system design. From both numerical and methodological viewpoints, it is therefore important to determine accurate transience bounds.

1.1 Our Contributions

In this paper, we present two upper bounds on the transients of linear max-plus systems. Our approach is graph-theoretic in nature: The problem of bounding from above the transient can be reduced to the study of walks in a specific graph. More precisely, for every max-plus matrix A, one considers the weighted directed graph G whose adjacency matrix is A, and its *critical subgraph* which consists of the *critical cycles*, namely those cycles with maximal average weight. The entries of the max-plus matrix power $A^{\otimes n}$ are equal to the maximum weights of walks in G of length n between two fixed nodes, and when redefining the weights of walks in a way that respects initial vector v, the entries of $A^{\otimes n} \otimes v$ are maximum weights of walks of length n starting from a fixed node. The periodicity of matrix powers and linear systems stems from the fact that eventually the weights of critical cycles dominate the maximum weight walks, as their weight-to-length ratio is the highest.

We present a general graph-based strategy whose core idea is a walk reduction $\mathrm{Red}_{d,k}$, which removes cycles from a walk while ensuring that its length remains in the same residue class modulo d, and that node k rests on the walk. The key property of $\mathrm{Red}_{d,k}$ is an upper bound on the length of the reduced walk that is linear both in d and the number of nodes in the graph. The following step in our strategy consists in completing reduced walks with critical cycles of appropriate lengths. For that, we propose two methods, namely the *repetitive* method and the *explorative* method. In the first one, the visit of the critical subgraph is confined to repeatedly follow only one closed walk whereas the second one consists in exploring one whole strongly connected component of the critical subgraph. That leads us to give two upper bounds on the transients of linear systems, namely the *repetitive bound* and the *explorative bound*, which are incomparable in general. We show that in the case of integer matrices, for a given initial vector and a given matrix semi-norm $\|A\|$ – which denotes the difference of the maximum and minimum finite entries of A –, both our transience bounds are $O(\max\{N^2, \|A\| \cdot N^3\})$. We also show that this bound is asymptotically tight.

Another contribution of this paper concerns the relationship between matrix and system transients: We prove that the transient of an $N \times N$ matrix A coincides with the transient of an A-linear system with an initial vector whose norm is at most quadratic in N, provided the latter transient is sufficiently large. In addition to shedding new light on transients, this result provides two upper bounds on matrix transients.

The problem of bounding the transients has already been studied (e.g., [20, 5, 27]), and the best previously known bound has been given by Hartmann and Arguelles [20]. Their bound on system transients is, in general, incomparable with our repetitive and explorative bounds. The significant benefit of our two new bounds is that each of them turns out to be linear in the size of the system in various classes of linear max-plus systems, e.g., those whose critical subgraphs have bounded girth, whereas Hartmann and Arguelles' bound is at least quadratic in the number of nodes. This is mainly due to the introduction of graph parameters that enables a fine-grained analysis of the transient phase. In particular, we introduce the notion of the *exploration penalty* of a graph G as the least integer k with the property that, for every $n \geqslant k$ divisible by the cyclicity of G and every node i of G, there is a closed path starting and ending at i of length n. Interestingly Schwarz [26] and Kim [22] proved two upper bounds on the exploration penalty which are both at most quadratic in the number of nodes of G.

Finally, we demonstrate how our general transience bound enables the performance analysis of a large variety of distributed systems. First, we apply our results to the class of *earliest schedules* in cyclic scheduling: we show that for a large family of sets of tasks, earliest schedules correspond to linear max-plus systems with irreducible matrices. Thus we prove the eventual periodicity of such earliest schedules, and give two upper bounds on their transient phases. Then we derive two transience bounds for a large class of synchronizers, and we quantify how both our synchronizer bounds are better than that given by Even and Rajsbaum [16] in their specific case of integer delays. In the process, we show that our transience bounds are asymptotically tight. Our results also apply to the analysis of the performance of distributed routers and schedulers based on the link-reversal algorithms: We obtain $O(N^3)$ transience bounds, improving the $O(N^4)$ bound established by Malka and Rajsbaum [24], and $O(N)$ bounds for such routers and schedulers when running in trees. For link-reversal routers, eventual periodicity actually corresponds to termination, and an $O(N^2)$ bound on time complexity [6, 9] directly follows from our transience bounds.

The paper is organized as follows. Section 2 introduces basic notions of graph theory and max-plus algebra. In Section 3, we elaborate a graph-based strategy to prove transience bounds. We show an upper bound on lengths of maximum weight walks that do not visit the critical subgraph in Section 4. Section 5 presents a walk reduction that constitutes the core of our strategy and we introduce the notion of *exploration penalty* in Section 6. We derive two transience bounds, namely the explorative and the repetitive bound, in Section 7. We show how to convert upper bounds on the transients of max-plus systems to upper bounds on the transients of max-plus matrices in Section 8. We discuss our results, by comparing them to previous work and by applying them to the analysis of various complex systems, in Section 9.

2 Preliminaries

This section introduces definitions and classical results needed in the rest of the paper. We denote by $\mathbb{N}$ the set of nonnegative integers and by $\mathbb{N}^*$ the set of positive integers.

2.1 Graphs

A *directed graph* G is a pair (V, E) where V is a nonempty finite set and $E \subseteq V \times V$. The elements of V are the *nodes* of G and the elements of E the *edges* of G. In this paper, we refer to directed graphs simply as *graphs*.

A *walk* W in G is a triple $W = (\mathrm{Start}, \mathrm{Edges}, \mathrm{End})$ where Start and End are nodes in G, Edges is a sequence $(e_1, e_2, \ldots, e_n)$ of edges $e_l = (i_l, j_l)$ such that $j_l = i_{l+1}$ if $1 \leqslant l \leqslant n-1$, $i_1 = \mathrm{Start}$ and $j_n = \mathrm{End}$ if the sequence Edges is nonempty, and $\mathrm{Start} = \mathrm{End}$ if the sequence Edges is empty. We define the operators Start, Edges, and End on a walk by setting $\mathrm{Start}(W) = \mathrm{Start}$, $\mathrm{Edges}(W) = \mathrm{Edges}$, and $\mathrm{End}(W) = \mathrm{End}$. We call $\mathrm{Start}(W)$ the *start node* of W and $\mathrm{End}(W)$ the *end node* of W. The *length* $\ell(W)$ of W is the number of edges in W, i.e., the length of the sequence $\mathrm{Edges}(W)$. Walk W is *closed* if $\mathrm{Start}(W) = \mathrm{End}(W)$. Walk W is *empty* if the sequence $\mathrm{Edges}(W)$ is empty. A walk W is empty if and only if $\ell(W) = 0$.

For two walks W and W', we say that W' is a *prefix* of W if $\mathrm{Start}(W) = \mathrm{Start}(W')$ and the sequence $\mathrm{Edges}(W')$ is a prefix of $\mathrm{Edges}(W)$. We say that W' is a *postfix* of W if $\mathrm{End}(W) = \mathrm{End}(W')$ and the sequence $\mathrm{Edges}(W')$ is a postfix of $\mathrm{Edges}(W)$. We call W' a *subwalk* of W if it is the postfix of some prefix of W. A subwalk W' of W is a *proper* subwalk of W if $W' \neq W$. We say a node i is a *node of walk* W if there exists a prefix W' of W with $\mathrm{End}(W') = i$. For two walks W_1 and W_2 with $\mathrm{End}(W_1) = \mathrm{Start}(W_2)$, we define the *concatenation* $W = W_1 \cdot W_2$ by setting $\mathrm{Start}(W) = \mathrm{Start}(W_1)$, $\mathrm{End}(W) = \mathrm{End}(W_2)$, and $\mathrm{Edges}(W)$ to be the juxtaposition of the sequences $\mathrm{Edges}(W_1)$ and $\mathrm{Edges}(W_2)$. If $W = W_1 \cdot C \cdot W_2$ where C is a closed walk, then $W' = W_1 \cdot W_2$ is also a walk with the same start and end nodes as W.

A walk is a *path* if it is non-closed and does not contain a nonempty closed walk as a subwalk. A closed walk is a *cycle* if it does not contain a nonempty closed walk as a proper subwalk. As cycles can be empty, there is a cycle of length 0 at each node of G.

If i and j are two nodes of G, let $\mathcal{W}_G(i,j)$ denote the set of walks W in graph G with $\mathrm{Start}(W) = i$ and $\mathrm{End}(W) = j$, and $\mathcal{W}_G(i{\to})$ the set of walks W in G with $\mathrm{Start}(W) = i$. If n is a nonnegative integer, we write $\mathcal{W}_G^n(i,j)$ (respectively $\mathcal{W}_G^n(i{\to})$) for the set of walks in $\mathcal{W}_G(i,j)$ (respectively $\mathcal{W}_G(i{\to})$) of length n. When no confusion can arise, we will omit the subscript G.

A graph $G' = (V', E')$ is a *subgraph* of G if $V' \subseteq V$ and $E' \subseteq E$. For a nonempty subset E' of E, let the *subgraph of G induced by edge set E'* be the graph (V', E') where $V' = \{i \in V \mid \exists j \in V : (i,j) \in E' \vee (j,i) \in E'\}$. A graph G is *strongly connected* if, for all nodes i and j in G, there exists a walk from i to j. A subgraph H of G is a *strongly connected component* of G if H is maximal with respect to the subgraph relation such that H is strongly connected.

The *girth* $g(G)$ of a graph G is the minimum length of a nonempty cycle in G. For a strongly connected graph G, its *cyclicity* $\gamma(G)$ is the greatest common divisor of cycle lengths in G. If G is not strongly connected, then its cyclicity $\gamma(G)$ is equal to the least common multiple of the cyclicities of its strongly connected components.

2.2 Linear max-plus systems

Let $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty\}$. In max-plus algebra, the set $\overline{\mathbb{R}}$ is equipped with the addition $a \oplus b = \max\{a, b\}$ and the multiplication $a \otimes b = a + b$. With these operations, $-\infty$ is the neutral element for addition and an absorbing element for multiplication, i.e., $(-\infty) \oplus a = a$ and $(-\infty) \otimes a = -\infty$ for all $a \in \overline{\mathbb{R}}$. In this paper, we follow the convention that the empty sum is the neutral element for addition, i.e., $\max \emptyset = -\infty$.

A matrix with entries in $\overline{\mathbb{R}}$ is called a *max-plus matrix*. If A is an $M \times N$ max-plus matrix and B is an $N \times Q$ max-plus matrix, then the *max-plus product* $A \otimes B$ is an $M \times Q$ max-plus matrix defined by

$$(A \otimes B)_{i,j} = \max_{1 \leqslant k \leqslant N} \big(A_{i,k} + B_{k,j} \big) \enspace ,$$

which is an associative operation. If A is an $N \times N$ max-plus matrix and n is a nonnegative integer, we denote by $A^{\otimes n}$ the n times iterated matrix product of A. That is, $(A^{\otimes 0})_{i,i} = 0$ and $(A^{\otimes 0})_{i,j} = -\infty$ if $i \neq j$, and $A^{\otimes n} = A \otimes A^{\otimes (n-1)}$ if $n \geqslant 1$. Given a column vector $v \in \overline{\mathbb{R}}^N$, the corresponding *linear max-plus system* is the sequence of vectors $x(n)$ defined by

$$x(n) = \begin{cases} v & \text{if } n = 0 \\ A \otimes x(n-1) & \text{if } n \geqslant 1 \ . \end{cases} \tag{2}$$

Clearly $x(n) = A^{\otimes n} \otimes v$. We denote with $\langle A, v \rangle$ the A-linear system x with initial vector v.

To an $N \times N$ max-plus matrix A naturally corresponds a graph $G(A)$ with set of nodes $\{1, \dots, N\}$ containing an edge (i,j) if and only if $A_{i,j}$ is finite. The matrix A is said to be *irreducible* if $G(A)$ is strongly connected.

We refer to $A_{i,j}$ as the *A-weight* of edge (i,j) in $G(A)$. If W is a walk in $G(A)$, we abuse notation by writing $A(W)$ for the weight of walk W, i.e., the sum of the weights of its edges. We follow the convention that the value of the empty sum is zero, i.e., $A(W) = 0$ if W is an empty walk. Given a column vector $v \in \overline{\mathbb{R}}^N$, we write $A_v(W) = A(W) + v_j$ where $j = \mathrm{End}(W)$ for W's *A_v-weight*. From these definitions, one can easily establish the following correspondence between the matrix power $A^{\otimes n}$ (respectively the vector $A^{\otimes n} \otimes v$) and the weights of some walks in $G(A)$ by unrolling the definition of max-plus multiplication. It allows us to use both interpretations of maximum weight walks and max-plus algebra.

Proposition 1. *Let i and j be two nodes of $G(A)$, and let n be a nonnegative integer. Then the following equations hold*

$$(A^{\otimes n})_{i,j} = \max\big\{A(W) \mid W \in \mathcal{W}^n_{G(A)}(i,j)\big\}$$
$$(A^{\otimes n} \otimes v)_i = \max\big\{A_v(W) \mid W \in \mathcal{W}^n_{G(A)}(i \to)\big\} \ .$$

For a max-plus vector v, we write $\|v\|$ for the maximum difference of entries in v, except if all entries are $-\infty$, we set it to 0. For a max-plus matrix A, we write $\|A\|$ for the maximum difference of *finite* entries of A.

2.3 The critical subgraph

A nonempty closed walk C in $G(A)$ is said to be *critical* if its average A-weight $A(C)/\ell(C)$ is maximal, i.e., if it is equal to

$$\lambda(A) = \max\big\{A(C)/\ell(C) \mid C \text{ is a nonempty closed walk in } G(A)\big\} \ ,$$

which is easily seen to be finite whenever there is at least one cycle in $G(A)$. A node of $G(A)$ is *critical* if it is a node of a critical closed walk in $G(A)$, and an edge of $G(A)$ is *critical* if it is an edge of a critical closed walk in $G(A)$. The *critical subgraph* of $G(A)$, denoted by $G_{\mathrm{c}}(A)$, is the subgraph of $G(A)$ induced by the set of critical edges of $G(A)$. We recall a useful property of closed walks in $G_{\mathrm{c}}(A)$ (for instance see [21, Lemma 2.6] for a proof).

Proposition 2. *Every nonempty closed walk in $G_{\mathrm{c}}(A)$ is critical in $G(A)$.*

Let us denote $\gamma(A) = \gamma\big(G_{\mathrm{c}}(A)\big)$.

2.4 Eventually periodic sequences

Let I be an arbitrary nonempty index set and $f : \mathbb{N} \to \overline{\mathbb{R}}^I$. Further let π be a positive integer and $\varrho \in \mathbb{R}$. The sequence f is *eventually periodic with period π and ratio ϱ* if there exists a nonnegative integer T such that

$$\forall i \in I \ : \ \forall n \geqslant T \ : \ f_i(n + \pi) = f_i(n) + \pi \cdot \varrho \ . \tag{3}$$

We call such a T a *transient* of f with respect to π and ρ. The ratio is unique if not all component-wise sequences $\big(f_i(n)\big)_n$ are eventually constantly equal to $-\infty$. In all cases, the set of transients of f is independent of the ratio.

Obviously if σ is any multiple of π, then f is also eventually periodic with period σ and ratio ϱ. Hence, there always exists a common period of two eventually periodic sequences.

For every period π, there exists a unique minimal transient T_π. The next lemma shows that these minimal transients do, in fact, not depend on π. We will henceforth call this common value *the* transient of f.

Proposition 3. *Let π and σ be two periods of an eventually periodic sequence f with respective minimal transients T_π and T_σ. Then $T_\pi = T_\sigma$.*

Proof. Denote by $\mathbf{\Pi}_f$ the set of periods of f and by ϱ a ratio of f. Clearly, $\mathbf{\Pi}_f$ is a nonempty subset of $\mathbb{N}^*$ closed under addition. Let $\pi_0 = \min \mathbf{\Pi}_f$ be the minimal period of f; hence $\pi_0\mathbb{N}^* \subseteq \mathbf{\Pi}_f$. Denote by T_0 the minimal transient with respect to period $\pi_0 \in \mathbf{\Pi}_f$. Let $\pi = a\pi_0 + b$ be the Euclidean division of π by π_0. For any integer $n \geqslant \max\{T_\pi, T_0 - b\}$,

$$f(n+\pi) = f(n) + \pi\varrho = f(n+b) + a\pi_0\varrho \ .$$

It follows that either $b = 0$ or b is a period of f. Since $b \leqslant \pi_0 - 1$ and π_0 is the smallest period of f, we have $b = 0$, i.e., π_0 divides π. We have thus shown $\mathbf{\Pi}_f \subseteq \pi_0\mathbb{N}^*$ and thus $\mathbf{\Pi}_f = \pi_0\mathbb{N}^*$. Hence $\pi = a\pi_0$ for some positive integer a.

Since for any $n \geqslant T_0$, $f(n + a\pi_0) = f(n) + a\pi_0\varrho$, we have $T_\pi \leqslant T_0$. We now prove that $T_\pi = T_0$ by induction on a.

1. The base case $a = 1$ is trivial.
2. Let $a \geqslant 2$. Denote by T' the minimal transient with respect to period $(a-1)\pi_0$. By the inductive hypothesis, $T' = T_0$. For any integer $n \geqslant T_\pi$,

   $$f(n + a\pi_0) = f(n) + a\pi_0\varrho \ . \tag{4}$$

   Moreover, if $n + \pi_0 \geqslant T'$ then

   $$f(n + a\pi_0) = f(n+\pi_0) + (a-1)\pi_0\varrho \ . \tag{5}$$

   It follows that for any integer $n \geqslant \max\{T' - \pi_0, T_\pi\}$,

   $$f(n+\pi_0) = f(n) + \pi_0\varrho \ . \tag{6}$$

   Hence $T_0 \leqslant \max\{T' - \pi_0, T_\pi\}$, and by inductive assumption $T_0 \leqslant \max\{T_0 - \pi_0, T_\pi\}$. We derive $T_0 \leqslant T_\pi$, and so $T_0 = T_\pi$, which concludes the proof. □

Cohen et al. proved eventual periodicity of irreducible max-plus matrix powers in the following analog of the Perron-Frobenius theorem in classical linear algebra.

Theorem 1 (Cyclicity Theorem [10])**.** *If A is irreducible, then the sequence of matrix powers $A^{\otimes n}$ is eventually periodic with period $\gamma(A)$ and ratio $\lambda(A)$.*

Consequently, every linear max-plus system with an irreducible matrix A is eventually periodic with period $\gamma(A)$ and ratio $\lambda(A)$.

We call the transient of the sequence of matrix powers $A^{\otimes n}$ the *transient of matrix A*, and the transient of the sequence of vectors $A^{\otimes n} \otimes v$ the *transient of the system $\langle A, v\rangle$*.

For any $\mu \in \mathbb{R}$, let $A + \mu$ denote the matrix obtained by adding μ to each entry of A. Since $(A+\mu)^{\otimes n} = A^{\otimes n} + n\mu$, we easily check that $G_c(A+\mu) = G_c(A)$, $\lambda(A+\mu) = \lambda(A) + \mu$, and the matrix transients of A and $A + \mu$ (resp. the system transients of $\langle A, v\rangle$ and $\langle A+\mu, v\rangle$) are equal.

3 Strategy Outline

This section describes our graph-based strategy to prove upper bounds on the transient of the system $\langle A, v\rangle$, given an irreducible $N \times N$ matrix A and a vector $v \in \mathbb{R}^N$. We also explain how a slight modification of this strategy provides upper bounds on the transient of A.

We start by defining for a set $\mathbf{N}$ of nonnegative integers and a node i, an $\mathbf{N}$-*realizer* for node i to be any walk of maximum A_v-weight in the set of walks in $\mathcal{W}(i \rightarrow)$ with length in $\mathbf{N}$. As shown in the next proposition, of particular interest is the case of sets $\mathbf{N}$ of the form

$$\mathbf{N}_{\geqslant B}^{(n,\pi)} = \{m \in \mathbb{N} \mid m \geqslant B \ \wedge \ m \equiv n \pmod{\pi}\}$$

where B, n, and π are positive integers.

Proposition 4. *Let B and π be positive integers. If for every node i and every integer $n \geqslant B$, there exists an $\mathbf{N}_{\geqslant B}^{(n,\pi)}$-realizer for i of length n, then B is an upper bound on the system transient.*

Proof. Let i be a node. For each integer $n \geqslant B$, let W_n be an $\mathbf{N}_{\geqslant B}^{(n,\pi)}$-realizer for i of length n. Denote by $X(n)$ the set of walks W in $\mathcal{W}(i\to)$ with $\ell(W) \in \mathbf{N}_{\geqslant B}^{(n,\pi)}$, and let $x(n)$ be the maximum of values $A_v(W)$ where $W \in X(n)$. It follows that $x(n) = A_v(W_n)$.

From $n + \pi \equiv n \pmod \pi$ follows $X(n + \pi) = X(n)$ and so $x(n + \pi) = x(n)$. Moreover, we have $\mathcal{W}^n(i\to) \subseteq X(n)$ and $\mathcal{W}^{n+\pi}(i\to) \subseteq X(n + \pi)$, which implies $(A^{\otimes n} \otimes v)_i \leqslant x(n)$ and $(A^{\otimes(n+\pi)} \otimes v)_i \leqslant x(n + \pi)$. Conversely because $W_n \in \mathcal{W}^n(i\to)$, we have $(A^{\otimes n} \otimes v)_i \geqslant A(W_n) = x(n)$. Similarly, $(A^{\otimes(n+\pi)} \otimes v)_i \geqslant A(W_{n+\pi}) = x(n + \pi)$. Since $x(n + \pi) = x(n)$, it follows that $(A^{\otimes n} \otimes v)_i = (A^{\otimes(n+\pi)} \otimes v)_i$. Noting Proposition 3 now concludes the proof. □

Based on Proposition 4, we now define a 5-step strategy for determining upper bounds on system transients. Let n be a nonnegative integer and i be a node. Denote by π the least common multiple of cycle lengths in the critical subgraph G_c. Note that π is a multiple of $\gamma = \gamma(A)$. The strategy includes an additional parameter B to be chosen in step 4.

1. *Normalized matrix.* Because the transients of A and of $\overline{A} = A - \lambda(A)$ are equal, and $\lambda(\overline{A}) = 0$, we can reduce the general case to the case $\lambda(A) = 0$. The condition $\lambda(A) = 0$ guarantees the existence of realizers for every nonempty $\mathbf{N} \subseteq \mathbb{N}$ and yields that adding critical cycles to a walk does not change its A-weight. The rest of the strategy hence considers an irreducible matrix A such that $\lambda(A) = 0$. Let W be an $\mathbf{N}_{\geqslant B}^{(n,\pi)}$-realizer for node i.
2. *Critical bound.* We show that for B large enough, i.e., B greater than or equal to some *critical bound* B_c, the realizer W contains at least one critical node k.
3. *Walk reduction.* Next we show that for every divisor d of π, by removing subcycles, we can *reduce* W to a new walk $\hat{W}$ which starts at node i, contains the critical node k, whose length $\ell(\hat{W})$ is in the same residue class modulo d as $\ell(W)$, and $\ell(\hat{W})$ is upper-bounded by a term linear in the number of nodes in the graph.
4. *Pumping in the critical graph.* Since d divides π, d divides $n - \ell(\hat{W})$, and for two appropriate choices of d and for n sufficiently large ($n \geqslant B_d$), we show how to complete $\hat{W}$ by adding to it a critical closed walk starting from k in order to obtain a new walk of length n starting at node i.

   For $B = \max\{B_c, B_d\}$, this yields an $\mathbf{N}_{\geqslant B}^{(n,\pi)}$-realizer of length n, because removing cycles at most increases the weight and adding a critical closed path does not change the weight.
5. *Conclusion.* Proposition 4 then shows that B is a bound on the transient.

For the transient of the matrix A, we can follow a similar strategy: we consider $\mathcal{W}(i, j)$ instead of $\mathcal{W}(i\to)$, and for a set $\mathbf{N}$ of nonnegative integers we define an $\mathbf{N}$-*realizer* for the pair of nodes i, j to be any walk of maximum A-weight in the set of walks $\mathcal{W}(i, j)$ with length in $\mathbf{N}$. As for walks in $\mathcal{W}(i\to)$, we can show that any walk of maximum A-weight in $\mathcal{W}(i, j)$ with length in $\mathbf{N}_{\geqslant B}^{(n,\pi)}$ contains at least one critical node if B is greater than or equal to some critical bound B'_c. Since the walk reduction described above actually preserves both the starting and ending nodes, then we can derive an upper-bound on the transient of A. In fact, we will not develop this parallel strategy for matrices, but we rather propose a different method, which consists in computing a bound on the transient of matrix A from our bounds on transients of some specific systems $\langle A, v\rangle$.

4 Critical Bound

In this section, we carry out step 2 of our strategy. More precisely, we prove that any $\{n\}$-realizer for node i, i.e., a walk of maximum A_v-weight in the set of walks $\mathcal{W}^n(i\to)$, necessarily contains a critical node if n is large enough.

Let A be an $N \times N$ max-plus matrix, and assume A is irreducible. We write λ for $\lambda(A)$, λ_{nc} for the maximum average A-weight of closed walks without critical nodes, δ for the minimum A-weight, Δ for the maximum A-weight, and Δ_{nc} for the maximum A-weight of edges between non-critical nodes. Further let $\tilde{\Delta}_{\mathrm{nc}} = \max\{\Delta_{\mathrm{nc}}, \lambda\}$. Note that $\tilde{\Delta}_{\mathrm{nc}} = -\infty$ implies $\lambda_{\mathrm{nc}} = -\infty$.

We assume $\|v\|$ to be finite until Section 8, in which we generalize our results to arbitrary v. By comparing the possible A_v-weights of walks that do and do not visit G_c, we can derive an explicit critical bound B_c, which holds for arbitrary λ. In the bound of the following proposition, we impose the interpretation $-\infty/\infty = -\infty$, which can occur if $\tilde{\Delta}_{\mathrm{nc}} = -\infty$. In this case, as is the case as soon as $\lambda_{\mathrm{nc}} = -\infty$, the critical bound is $B_c = N$. Further note that $\tilde{\Delta}_{\mathrm{nc}} - \delta \leqslant \|A\|$.

Proposition 5 (Critical Bound)**.** *Each walk with maximum A_v-weight in $\mathcal{W}^n(i\to)$ contains a critical node if $n \geqslant B_{\mathrm{c}}$ where*

$$B_{\mathrm{c}} = \max\left\{N\ ,\ \frac{\|v\| + (\tilde{\Delta}_{\mathrm{nc}} - \delta)\,(N-1)}{\lambda - \lambda_{\mathrm{nc}}} + 1\right\}\ .$$

Proof. We first reduce to the case $\lambda = 0$. Let $\overline{A}$ be the normalized matrix $\overline{A} = A - \lambda$. The parameters $\overline{\delta}$, $\overline{\tilde{\Delta}}_{\mathrm{nc}}$, and $\overline{\lambda}_{\mathrm{nc}}$ for the matrix $\overline{A}$ are obtained by subtracting λ from the respective parameters of A. Hence $\overline{\lambda} = 0$, and a walk is of maximum A_v-weight in $G(A)$ if and only if it is a walk of maximum $\overline{A}_v$-weight in $G(\overline{A}) = G(A)$.

Given an A with arbitrary λ, we will show that each walk with maximum $\overline{A}_v$-weight in $\mathcal{W}^n(i\to)$ contains a critical node if

$$n \geqslant \max\left\{N\ ,\ \frac{\|v\| + \left(\overline{\tilde{\Delta}}_{\mathrm{nc}} - \overline{\delta}\right)\,(N-1)}{-\overline{\lambda}_{\mathrm{nc}}} + 1\right\} = B_{\mathrm{c}}\ .$$

Note that $\overline{\lambda} = 0$. Together with the facts that (i) walks of maximum $\overline{A}_v$-weight in $G(\overline{A})$ are of maximum A_v-weight in $G(A)$, and (ii) the set of critical nodes is the same in $\overline{A}$ and A, the critical bound on A follows. We can hence assume $\lambda = 0$ right-away in the rest of the proof. We start by distinguishing two cases for λ_{nc}.

If $\lambda_{\mathrm{nc}} = -\infty$, then every nonempty cycle contains a critical node. Because every walk of length greater than or equal to N necessarily contains a cycle as a subwalk and because $B_{\mathrm{c}} \geqslant N$, in particular every walk with maximum A_v-weight in $\mathcal{W}^n(i\to)$ contains a critical node if $n \geqslant B_{\mathrm{c}}$ and $\lambda_{\mathrm{nc}} = -\infty$.

We now consider the case $\lambda_{\mathrm{nc}} \neq -\infty$; in particular $\tilde{\Delta}_{\mathrm{nc}} \neq -\infty$. We proceed by contradiction: Suppose that there exists an integer n with $n \geqslant B_{\mathrm{c}}$ and a node i, and a walk of maximum A_v-weight in $\mathcal{W}^n(i\to)$ with non-critical nodes only; let $\hat{W}$ be any such walk. Let W_0 be the *acyclic part* of $\hat{W}$, defined in the following manner: Starting at $\hat{W}$, we repeatedly remove nonempty subcycles from the walk until we arrive at a path. In general there are several possible choices of which subcycles to remove, but we fix some global choice function to make the construction of W_0 deterministic.

Next choose a pair (k, W_{c}), where k is a critical node and W_{c} a prefix of W_0, such that the distance between k and the end node of W_{c} is minimal. Let W_2 be a path of minimal length from the end node of W_{c} to k. Note that W_2 contains only non-critical nodes by construction. Let W_3 be the walk such that $W_0 = W_{\mathrm{c}} \cdot W_3$. Further let C be a critical cycle starting at k. We next distinguish two cases for n.

1. If $n < \ell(W_{\mathrm{c}}) + \ell(W_2)$ we conclude that $n < \ell(W_{\mathrm{c}}) + \ell(W_2) \leqslant N \leqslant B_{\mathrm{c}}$; a contradiction to $n \geqslant B_{\mathrm{c}}$. The proposition follows.
2. Otherwise, $n \geqslant \ell(W_{\mathrm{c}}) + \ell(W_2)$. Let $m \in \mathbb{N}$ be the quotient in the Euclidean division of $n - \ell(W_{\mathrm{c}}) - \ell(W_2)$ by $\ell(C)$, and choose W_1 to be a prefix of C of length $n - \big(\ell(W_{\mathrm{c}}) + \ell(W_2) + m\cdot\ell(C)\big)$ (cf. Figure 1). Clearly W_1 starts at k. If we set $W = W_{\mathrm{c}} \cdot W_2 \cdot C^m \cdot W_1$, we get $\ell(W) = n$ and

$$A_v(W) \geqslant \min_{1\leqslant j\leqslant N}(v_j) + A(W_{\mathrm{c}}) + A(W_2) + A(W_1) \tag{7}$$

since we assume $\lambda = 0$.

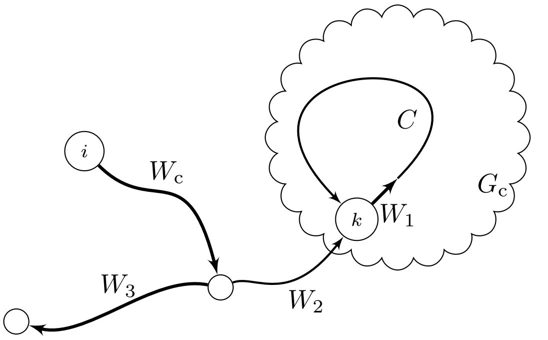

Figure 1: Walk W in proof of Proposition 5

For the A_v-weight of $\hat{W}$, we have

$$A_v(\hat{W}) \leqslant A_v(W_0) + \lambda_{\mathrm{nc}}\cdot\big(\ell(\hat{W}) - \ell(W_0)\big) \leqslant \max_{1\leqslant j\leqslant N}(v_j) + A(W_0) + \lambda_{\mathrm{nc}}\cdot\big(\ell(\hat{W}) - \ell(W_0)\big)\ . \tag{8}$$

By assumption $A_v(W) \leqslant A_v(\hat{W})$, and from (7) and (8) we therefore obtain

$$\min_{1 \leqslant j \leqslant N}(v_j) + A(W_c) + A(W_2) + A(W_1) \leqslant \max_{1 \leqslant j \leqslant N}(v_j) + A(W_0) + \lambda_{\mathrm{nc}} \cdot \big(\ell(\hat{W}) - \ell(W_0)\big)$$

$$-\lambda_{\mathrm{nc}} \cdot \big(\ell(\hat{W}) - \ell(W_0)\big) \leqslant \|v\| - A(W_c) - A(W_2) - A(W_1) + A(W_0) \ ,$$

and further with $\lambda_{\mathrm{nc}} < 0$ and $A(W_0) = A(W_c) + A(W_3)$,

$$\ell(\hat{W}) \leqslant \frac{\|v\| + A(W_3) - A(W_1) - A(W_2)}{-\lambda_{\mathrm{nc}}} + \ell(W_0) \leqslant \frac{\|v\| + \Delta_{\mathrm{nc}}\,\ell(W_3) - \delta\,(\ell(W_1) + \ell(W_2))}{-\lambda_{\mathrm{nc}}} + \ell(W_0).$$

Together with $\lambda_{\mathrm{nc}} \geqslant \delta$,

$$\ell(\hat{W}) \leqslant \frac{\|v\| + \Delta_{\mathrm{nc}}\,\ell(W_3) - \delta\,(\ell(W_1) + \ell(W_2) + \ell(W_0))}{-\lambda_{\mathrm{nc}}} \ .$$

Denote by N_{nc} the number of non-critical nodes. The following inequalities hold: $\ell(W_3) \leqslant N - 1$, and $\ell(W_1) < N - N_{\mathrm{nc}}$. Together with $\Delta_{\mathrm{nc}} \leqslant \tilde{\Delta}_{\mathrm{nc}}$ and $\tilde{\Delta}_{\mathrm{nc}} \geqslant \lambda = 0$, we deduce $\Delta_{\mathrm{nc}}\,\ell(W_3) \leqslant \tilde{\Delta}_{\mathrm{nc}}\,(N-1)$. Moreover from the minimality constraint for the length of W_2, and the fact that W_0 and W_2 contain only non-critical nodes, it follows that $\ell(W_2) + \ell(W_0) \leqslant N_{\mathrm{nc}}$. Thereby

$$\ell(\hat{W}) < \frac{\|v\| + (\tilde{\Delta}_{\mathrm{nc}} - \delta)\,(N-1)}{-\lambda_{\mathrm{nc}}} + 1 \ ,$$

a contradiction to $n \geqslant B_c$. The proposition then follows in the case $n \geqslant \ell(W_c) + \ell(W_2)$.

□

In case A is an integer matrix, i.e., all finite entries of A are integers, the term $\lambda - \lambda_{\mathrm{nc}}$ cannot become arbitrarily small: This is obvious when $\lambda_{\mathrm{nc}} = -\infty$; otherwise, let C_0 be a critical cycle, and let C_1 be a cycle such that $\lambda_{\mathrm{nc}} = A(C_1)/\ell(C_1)$. Then we have

$$\lambda - \lambda_{\mathrm{nc}} = \frac{A(C_0)\ell(C_1) - A(C_1)\ell(C_0)}{\ell(C_0)\ell(C_1)} \ ,$$

and so

$$\frac{1}{\lambda - \lambda_{\mathrm{nc}}} \leqslant (N - N_{\mathrm{nc}}) \cdot N_{\mathrm{nc}} \leqslant \frac{N^2}{4} \ , \tag{9}$$

where N_{nc} denotes the number of non-critical nodes. It follows that given a matrix norm $\|A\|$ and a vector norm $\|v\|$, the critical bound B_c is $O(\max\{N,\, \|v\| \cdot N^2\,,\, \|A\| \cdot N^3\})$.

The proof of Proposition 5 can be modified to show that there *exists* a walk with maximum A_v-weight in $\mathcal{W}^n(i \to)$ that contains a critical node if $n \geqslant B'_c$ where B'_c is slightly smaller than B_c and obtained by removing 1 in its second term. Essentially assumption $A_v(W) \leqslant A_v(\hat{W})$ in the proof is replaced by $A_v(W) < A_v(\hat{W})$, while the remaining part of the proof is analogous.

5 Walk Reduction

This section concerns step 3 of our strategy and constitutes its core. Given a walk W, a positive integer d, and a node k of W, we define a reduced walk, denoted $\mathrm{Red}_{d,k}(W)$, such that (a) it contains node k and has the same start and end nodes as W, (b) its length is in the same residue class modulo d as W's length, and (c) its length is bounded by $(d-1) + 2d\,(N-1)$.

Properties (a) and (b) can be achieved by removing a collection of cycles from W whose combined length is divisible by d and whose removal retains connectivity to k. The key point of the reduction is that we can iterate this cycle removal until the resulting length is at most $(d-1) + 2d\,(N-1)$.

We call a finite, possibly empty, sequence of nonempty subcycles $\mathcal{S} = (C_1, C_2, \ldots, C_n)$ a *cycle pattern of a walk* W if there exist walks $U_0, U_1, \ldots, U_n$ such that

$$W = U_0 \cdot C_1 \cdot U_1 \cdot C_2 \cdots U_{n-1} \cdot C_n \cdot U_n \ . \tag{10}$$

The choice of the U_m's in (10) may be not unique, and we fix some global choice function to make it deterministic. Then we define *the removal of* $\mathcal{S}$ *from* W as

$$\mathrm{Rem}(W, \mathcal{S}) = U_0 \cdot U_1 \cdots U_n \ .$$

The walks W and $\operatorname{Rem}(W,\mathcal{S})$ have the same start and end nodes. Furthermore $\ell\big(\operatorname{Rem}(W,\mathcal{S})\big) = \ell(W) - \ell(\mathcal{S})$ where $\ell(\mathcal{S}) = \sum_{C\in\mathcal{S}} \ell(C)$. In particular, $\operatorname{Rem}(W,\mathcal{S}) = W$ if and only if $\ell(\mathcal{S}) = 0$, i.e., $\mathcal{S}$ is the empty cycle pattern.

Given any node k of a walk W, let $\mathbf{S}_k(W)$ denote the set of cycle pattern $\mathcal{S}$ of W whose removal does not impair connectivity to k, i.e., k is a node of $\operatorname{Rem}(W,\mathcal{S})$. Further for any positive integer d, define $\mathbf{S}_{d,k}(W)$ as the subset of cycle pattern $\mathcal{S} \in \mathbf{S}_k(W)$ that, in addition, leave the length's residue class modulo d intact, i.e., $\ell(\mathcal{S}) \equiv 0 \pmod d$. The set $\mathbf{S}_{d,k}(W)$ is not empty, because k is a node of W and we can hence choose $\mathcal{S}$ to be the empty cycle pattern.

Choose $\mathcal{S} \in \mathbf{S}_{d,k}(W)$ such that $\ell(\mathcal{S})$ is maximal. There may be several possible choices for $\mathcal{S}$, and we again fix some global choice function to make the choice deterministic; then set

$$\operatorname{Step}_{d,k}(W) = \operatorname{Rem}(W,\mathcal{S}) \enspace .$$

The limit

$$\operatorname{Red}_{d,k}(W) = \lim_{t\to\infty} \operatorname{Step}^t_{d,k}(W)$$

exists because the sequence of walks $(\operatorname{Step}^t_{d,k}(W))_{t\geqslant 0}$ is stationary after at most $\ell(W)$ steps, and we call it the *(d,k)-reduction of W*. More specifically, $\operatorname{Red}_{d,k}(W) = W$ if and only if $\mathbf{S}_{d,k}(W)$ is reduced to the sole empty cycle pattern. The walks W and $\operatorname{Red}_{d,k}(W)$ have the same start and end nodes. Also, k is a node of $\operatorname{Red}_{d,k}(W)$ and $\ell\big(\operatorname{Red}_{d,k}(W)\big) \equiv \ell(W) \pmod d$.

Note that, even though $\mathcal{S}$ is chosen maximal in every step, a single step might not suffice to fully reduce the walk. Figure 2 depicts an example with $d = 1$ and k the walk's start node: Cycle C is the only walk that can be removed in the first step, but its removal reveals a second cycle, which is then subsequently removed.

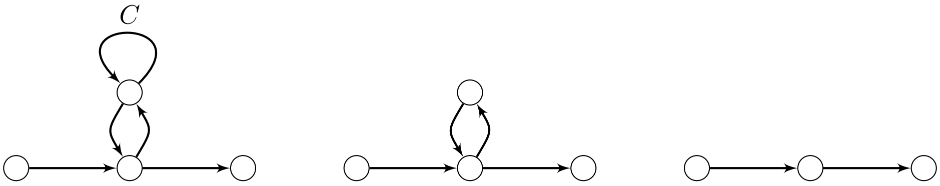

Figure 2: A walk that necessitates several steps in the reduction

Bounding the length of $\operatorname{Red}_{d,k}(W)$ relies on a simple arithmetic lemma which is an elementary application of the pigeonhole principle:

Lemma 1. *Let d be a positive integer and let $x_1,\dots,x_d \in \mathbb{Z}$. Then there exists a nonempty set $I \subseteq \{1,\dots,d\}$ such that $\sum\limits_{i\in I} x_i \equiv 0 \pmod d$.*

Theorem 2. *For each positive integer d and each node k, the length of the (d,k)-reduction of any walk W containing node k is at most equal to $(d-1) + 2d\cdot(N-1)$:*

$$\ell\big(\operatorname{Red}_{d,k}(W)\big) \leqslant (d-1) + 2d\cdot(N-1) \enspace .$$

Proof. We denote $\hat{W} = \operatorname{Red}_{d,k}(W)$. By definition of the (d,k)-reduction, $\operatorname{Red}_{d,k}(\hat{W}) = \hat{W}$. Let $\mathcal{S}$ be any cycle pattern of $\hat{W}$ in $\mathbf{S}_k(\hat{W})$, and let n be the number of cycles of $\mathcal{S}$. We first show that $n \leqslant d-1$. Indeed, suppose for contradiction that $n \geqslant d$. Then Lemma 1 implies that there exists a nonempty subsequence of $\mathcal{S}$ that is in $\mathbf{S}_{d,k}(\hat{W})$, which contradicts $\operatorname{Red}_{d,k}(\hat{W}) = \hat{W}$.

Now let us choose $\mathcal{S}$ in $\mathbf{S}_k(\hat{W})$ with maximal $\ell(\mathcal{S})$. If $\mathcal{S} = (C_1, C_2, \dots, C_n)$, then there exist walks $U_0, U_1, \dots, U_n$ such that

$$\hat{W} = U_0 \cdot C_1 \cdot U_1 \cdot C_2 \cdots U_{n-1} \cdot C_n \cdot U_n \enspace .$$

By definition of $\mathbf{S}_k(\hat{W})$, k is a node of $\operatorname{Rem}(\hat{W},\mathcal{S})$. Hence there exists some index r such that k is a node of U_r. Each U_m with $m \neq r$ is a (possibly empty) path, because otherwise we could add a nonempty subcycle of U_m to $\mathcal{S}$, a contradiction to the maximality of $\ell(\mathcal{S})$. Similarly, if $U_r = W_1 \cdot W_2$ such that k is the end node of W_1, then both W_1 and W_2 are (possibly empty) paths. Hence, apart from the at most $(d-1)$ cycles in $\mathcal{S}$, the reduced walk $\hat{W}$ consists of at most $(d+1)$ subpaths. Noting that each cycle has length at most N and each path has length at most $(N-1)$ concludes the proof. □

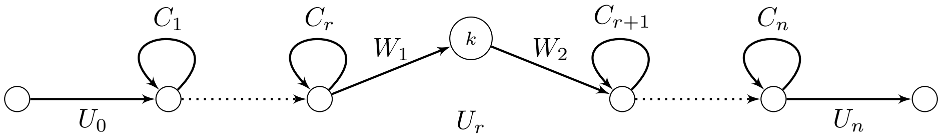

Figure 3: Structure of the reduced walk $\hat{W} = \mathrm{Red}_{d,k}(W)$

6 Exploration Penalty

One of the two pumping techniques that we develop in step 4 of our strategy (*Pumping in the critical graph*) for the construction of arbitrarily long closed walks in the critical graph G_c consists in *exploring* one strongly connected component H of G_c: The closed walks keep inside H, but may visit any node in H.

For that, we first introduce for a strongly connected graph G the *exploration penalty* of G, $ep(G)$, as the smallest integer e such that for any node i and any integer $n \geqslant e$ that is a multiple of G's cyclicity, there is a closed walk of length n starting at i. In the approach developed for the explorative transient bound, the exploration penalty thus constitutes a threshold to pump walk lengths in multiples of the cyclicity.

The exploration penalty can be seen as the transient of diagonal entries in the sequence of Boolean matrix powers of the graph's adjacency matrix. Thereby the Cyclicity Theorem [10] recalled in Section 2.4 demonstrates that the exploration penalty of a strongly connected graph is finite.

In graph theory, the exploration penalty has been studied under the name *index of convergence* of a digraph, or also *exponent* if the digraph is primitive, i.e., if its cyclicity is equal to 1 (see, e.g., [23] for a survey on the index of convergence). Two bounds that do not include the girth g as a parameter were given by Wielandt [30] for primitive graphs and by Schwarz [26] for the general case. Wielandt's bound on the exploration penalty of a primitive strongly connected graph with N nodes is equal to $W(N) = N^2 - 2N + 2$. Then Schwarz generalized this result to arbitrary cyclicities γ and arrived at a bound of $\gamma \cdot W\big(\lfloor N/\gamma \rfloor\big) + (N \bmod \gamma)$.

Denardo [13] and Dulmage and Mendelsohn [15] independently established an upper bound for primitive graphs that include the graph's girth g as a parameter, and is equal to $g \cdot (N-2) + N$. Using the latter upper bound for primitive graphs, Kim [22] then proved a general upper bound which generalized the results of Denardo, and Dulmage and Mendelsohn to arbitrary cyclicities and improved Schwarz' bound.

Theorem 3 ([22])**.** *Let G be a strongly connected graph with N nodes, of girth g and cyclicity γ. The exploration penalty of G satisfies the inequality*

$$ep(G) \leqslant g \cdot \big(\lfloor N/\gamma \rfloor - 2\big) + N \ .$$

Note that Kim's bound as well as Schwarz' bound are smaller than N^2 and thus both are $O(N^2)$.

7 Repetitive and Explorative Transience Bounds

We now follow the five step strategy laid out in Section 3 to prove two new bounds on system transients. They only differ in step 4 of the strategy (*Pumping in the critical graph*), namely, in the way one completes the reduced walk $\mathrm{Red}_{d,k}(W)$ with critical closed walks to reach the desired length n. Naturally this has implications on the appropriate choices for the walk reduction parameters d and k used in step 3 (*Walk reduction*).

7.1 Development of the strategy

Let A be an irreducible $N \times N$ max-plus matrix and let v be a vector in $\mathbb{R}^N$. Recall that π is chosen to be the least common multiple of cycle lengths in the critical subgraph G_c. Let i be any node, and let B and n be two positive integers such that $n \geqslant B \geqslant B_\mathrm{c}$.

1. *Normalized matrix.* We may assume $\lambda(A) = 0$, which guarantees that there exists a walk W that is an $\mathbf{N}^{(n,\pi)}_{\geqslant B}$-realizer for node i.
2. *Critical bound.* By definition of $\mathbf{N}^{(n,\pi)}_{\geqslant B}$ and since $B \geqslant B_\mathrm{c}$, we have

$$\ell(W) \geqslant B_\mathrm{c} \ .$$

Moreover walk W is a $\{\ell(W)\}$-realizer for node i. Proposition 5 shows that W contains at least one critical node k. Let H denote the strongly connected component of G_{c} containing k.

3. *Walk reduction.* Let d to be any divisor of π. By construction, $\hat{W} = \mathrm{Red}_{d,k}(W)$ is obtained by removing a collection of cycles from W, and starts at the same node i as W. Since $\lambda(A) = 0$, this implies
$$A_v(\hat{W}) \geqslant A_v(W) \ . \tag{11}$$
Moreover, walk $\hat{W}$ contains the critical node k, and its length $\ell(\hat{W})$ is in the same residue class modulo d as $\ell(W)$. Besides d divides $n - \ell(\hat{W})$ since d divides π. By Theorem 2, we have
$$\ell(\hat{W}) \leqslant (d-1) + 2d \cdot (N-1) \ . \tag{12}$$

4. *Pumping in the critical graph.*
   (a) *For the repetitive bound,* we choose
$$d = g(H) \ .$$
From (12), we derive that $n \geqslant \ell(\hat{W})$ when $B \geqslant (g(H)-1) + 2g(H) \cdot (N-1)$. Then we use a single critical cycle C to complete $\hat{W}$; see Figure 4(a). Let C be a cycle with length equal to the girth $g(H)$. We can assume that k is a node of C: In case k is not a node of C, we modify W by inserting π copies of a critical closed walk in H that connects W to C. Indeed, the addition of this critical closed walk changes neither the residue class modulo π nor the A_v-weight since $\lambda(A) = 0$. We complete the reduced walk $\hat{W}$ to length n by adding copies of C.
   (b) *For the explorative bound,* we choose
$$d = \gamma(H) \ ,$$
and use the definition of the exploration penalty $ep(H)$. From (12), we derive that $n \geqslant \ell(\hat{W}) + ep(H)$ when $B \geqslant (\gamma(H)-1) + 2\gamma(H) \cdot (N-1) + ep(H)$. By definition of $ep(H)$ and since $n - \ell(\hat{W}) \geqslant ep(H)$, we can complete $\hat{W}$ to length n by a critical closed walk in H; see Figure 4(b).

   Since d divides π, d divides $n - \ell(\hat{W})$, and for two appropriate choices of d and for n sufficiently large ($n \geqslant B_d$), we show how to complete $\hat{W}$ by adding to it a critical closed walk starting from k in order to obtain a new walk of length n starting at node i.

   For $B = \max\{B_{\mathrm{c}}, B_d\}$, this yields an $\mathbf{N}_{\geqslant B}^{(n,\pi)}$-realizer of length n, because removing cycles at most increases the weight and adding a critical closed path does not change the weight.

5. *Conclusion.* In each of the two completions, the resulting walk is of length n, starts at node i, and ends at the same node as $\hat{W}$. With (11), we deduce that its A_v-weight is at least $A_v(W)$. Thereby, it is an $\mathbf{N}_{\geqslant B}^{(n,\pi)}$-realizer for node i of length n. By Proposition 4, the repetitive and explorative completions finally give the following upper bounds on system transients.

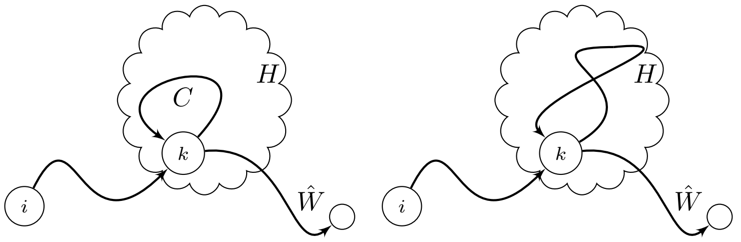

(a) Repetitive realizer (b) Explorative realizer

Figure 4: Repetitive and explorative realizers

Theorem 4 (Repetitive Bound)**.** *Denoting by $\hat{g}$ the maximum girth of strongly connected components of G_{c}, the transient of the linear max-plus system $\langle A, v\rangle$ is at most*
$$\max\left\{ \frac{\|v\| + \left(\tilde{\Delta}_{\mathrm{nc}} - \delta\right) \cdot (N-1)}{\lambda - \lambda_{\mathrm{nc}}} + 1 \ , \ (\hat{g}-1) + 2\,\hat{g} \cdot (N-1) \right\} \ .$$

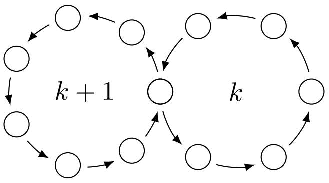

Figure 5: Graphs E_k

Theorem 5 (Explorative Bound). *Denoting by $\hat{\gamma}$ and $\hat{ep}$ the maximum cyclicity and maximum exploration penalty of strongly connected components of G_c, respectively, the transient of the linear max-plus system $\langle A, v\rangle$ is at most*

$$\max\left\{ \frac{\|v\| + \big(\tilde{\Delta}_\mathrm{nc} - \delta\big)\cdot(N-1)}{\lambda - \lambda_\mathrm{nc}} + 1 \ , \ (\hat{\gamma}-1) + 2\,\hat{\gamma}\cdot(N-1) + \hat{ep} \right\} \ .$$

7.2 Discussion

Because $\hat{g}$ is greater than or equal to $\hat{\gamma}$, the two bounds represent a tradeoff between choosing a larger multiplicative term versus the addition of the term $\hat{ep}$. It depends on the critical subgraph G_c which of the two bounds is better, and our two bounds are thus incomparable in general: As an example for which the explorative bound is lower than the repetitive bound, consider the family of graphs E_k depicted in Figure 5: E_k consists of two joint cycles of length k and $k+1$, respectively. All edges have zero weight. Independent of the initial vector v, the critical bound is N, since $\lambda_\mathrm{nc} = -\infty$. With $N = 2k$, $\hat{g} = k$, and $\hat{\gamma} = 1$, the repetitive bound is $4k^2 - k - 1$, and the explorative bound is at most $2k^2 + 4k - 2$. For $k \geqslant 3$ the explorative bound is strictly lower than the repetitive bound. Conversely, the repetitive bound is lower than the explorative bound, if we add a self-loop at the node that is shared by the two cycles in the above example.

Interestingly, the two terms in our transience bounds that are due to the repetitive and explorative completions are both at most quadratic: this is obvious for the repetitive term, and is an immediate corollary of the upper bounds on the exploration penalty in Section 6 (e.g., Theorem 3). In the case of integer matrices, both the repetitive and the explorative bounds are $O(\max\{N^2\ , \ \|A\|\cdot N^3\})$ since the critical bound itself is $O(\max\{N\ , \ \|v\|\cdot N^2\ , \ \|A\|\cdot N^3\})$ in this case.

Hartmann and Arguelles [20] established the best previously known bound on system transients. Their approach includes passing to the max-balancing [25] of G and considering an increasing sequence of threshold graphs which all include the critical subgraph. Their technique to increase the length of maximum weight walks is comparable to our repetitive pumping technique. They proved that the transient of system $\langle A, v\rangle$ is upper-bounded by $\max\big\{(\|v\|+\|A\|\cdot N)/(\lambda-\lambda^0)\ , \ 2N^2\big\}$ where λ^0 is defined in terms of the max-balancing of G. The first term in their bound is in general incomparable with our critical bound, whereas the second term, namely $2N^2$, is always strictly larger than the second term in each of our two bounds and makes their bound at least quadratic in N. Trivially, the minimum of our two bounds, and of Hartmann and Arguelles' bound, yields the best currently known bound.

All our bounds, as well as the bound by Hartmann and Arguelles, are computable in time polynomial in the number N of nodes. Indeed, the critical subgraph can be computed in polynomial time using the max-balancing algorithm by Schneider and Schneider [25]; its girth, cyclicity, and exploration penalty can then be found in polynomial time via Boolean matrix multiplication.

8 Matrix vs. System Transients

As explained in Section 3, we can follow the same strategy as for system transients to bound matrix transients: The proof consists in a copy of the proof developed in the previous sections where the sets $\mathcal{W}^n(i, \rightarrow)$ are replaced by the sets $\mathcal{W}(i,j)$. For an $N\times N$ max-plus matrix A, we thus obtain an upper bound that is $O\big(\max\{N^2\ , \ \|A\|\cdot N^2/(\lambda-\lambda_\mathrm{nc}\}\big)$, and so is $O(\max\{N^2\ , \ \|A\|\cdot N^4\})$ for integer matrices. Unfortunately this proof brings no specific insight on matrix transients and in particular, gives no hint on the relationships between the transient of max-plus matrix A and the transients of the max-plus systems $\langle A, v\rangle$.

In this section, we follow another approach which consists in relating the transient of matrix A to the transient of a specific system $\langle A, v\rangle$ where $\|v\|$ is $O\big(\|A\|\cdot N^2\big)$. While the latter result is interesting on its own, it can be combined with our upper bounds on the system transient established in Theorems 4 and 5 to give two upper bounds on the matrix transient which are also $O\big(\|A\|\cdot N^2/(\lambda-\lambda_{\mathrm{nc}})\big)$.

Let n_A and $n_{A,v}$ denote the transient of matrix A and the transient of system $\langle A, v\rangle$, respectively. Obviously, n_A is an upper bound on the $n_{A,v}$'s. Conversely, the equalities $A^{\otimes n}_{i,j} = \big(A^{\otimes n} \otimes e^j\big)_i$, where the e^j's are the unit vectors defined by $e^j_i = 0$ if $i = j$ and $e^j_i = -\infty$ otherwise, show that $\max\big\{n_{A,e^j} | j \in \{1, \cdots, N\}\big\} \geqslant n_A$. Hence,

$$\sup\big\{n_{A,v} | v \in \overline{\mathbb{R}}^N\big\} = n_A \quad .$$

We now seek a similar expression of n_A, but with finite initial vectors v, i.e., with $v \in \mathbb{R}^N$.

Reusing the notation $\hat{\gamma}$ and $\hat{ep}$ from Theorem 5, we define:

$$\begin{aligned}
&\tilde{B} = 2(N-1) + \hat{ep} + (ep(G) + \hat{\gamma} - 1), \\
&\mu = \sup\Big\{A^{\otimes n}_{i,h} - A^{\otimes n}_{i,j} \mid h, i, j \text{ nodes of } G\ ,\ n \geqslant \tilde{B}\ ,\ A^{\otimes n}_{i,j} \neq -\infty\Big\}
\end{aligned}$$

Clearly μ is finite, i.e., $\mu \in \mathbb{R}$. Then we consider the *μ-truncated* unit vector ε^j obtained by replacing the infinite entries of the e^j's by $-\mu$. Let us observe that for each μ-truncated unit vector ε^j, we have $\|\varepsilon^j\| = \mu$.

In Proposition 6 below, we show that if $B \geqslant \tilde{B}$ and B is a bound on the system transients for all μ-truncated unit vectors, then B is also a bound on the matrix transient. A technical difficulty in the proof lies in the fact that, contrary to the sets $\mathcal{W}^n(i\to)$ which occur in the expression of the i-th component of linear systems, the sets $\mathcal{W}^n(i,j)$ that we consider for matrix powers may be empty. The next two lemmas deal with this technicality.

Lemma 2. *For any pair of nodes i, j of G and any integer $n \geqslant ep(G) + \gamma(G) + N - 2$, there exists a walk W from i to j such that $n - \ell(W) \in \{0, \dots, \gamma(G) - 1\}$.*

Proof. Let i, j be any two nodes, and let W_0 be a path from i to j. For any integer n, consider the residue r of $n - \ell(W_0)$ modulo $\gamma(G)$. By definition of $ep(G)$, if $n - \ell(W_0) - r \geqslant ep(G)$, then there exists a closed walk C starting at node j with length equal to $n - \ell(W_0) - r$. Then, $W_0 \cdot C$ is a walk from i to j with length $n - r$, where $r \in \{0, \dots, \gamma(G) - 1\}$. The lemma follows since $n - \ell(W_0) - r \geqslant ep(G)$ as soon as $n \geqslant ep(G) + (N-1) + \gamma(G) - 1$. □

Lemma 3. *Let n be any integer such that $n \geqslant ep(G) + \gamma(G) + N - 2$. Then $A^{\otimes(n+\gamma(G))}_{i,j} = -\infty$ if and only if $A^{\otimes n}_{i,j} = -\infty$.*

Proof. It is equivalent to claim that $\mathcal{W}^{n+\gamma(G)}(i,j) = \emptyset$ if and only if $\mathcal{W}^n(i,j) = \emptyset$ for any integer $n \geqslant ep(G) + \gamma(G) + N - 2$.

Suppose $\mathcal{W}^{n+\gamma(G)}(i,j) \neq \emptyset$, and let $W_0 \in \mathcal{W}^{n+\gamma(G)}(i,j)$. By Lemma 2, there exists a walk $W \in \mathcal{W}(i,j)$ such that $n = \ell(W) + r$ with $r \in \{0, 1, \dots, \gamma(G) - 1\}$. By definition, $\gamma(G)$ divides $\ell(W_0) - \ell(W) = (n + \gamma(G)) - (n - r) = \gamma(G) + r$; hence $\gamma(G)$ divides r. Therefore, $r = 0$, i.e., $\ell(W) = n$ and thus $\mathcal{W}^n(i,j) \neq \emptyset$.

The converse implication is proved similarly using a walk $W_1 \in \mathcal{W}^n(i,j)$. □

Proposition 6. *If $n \geqslant \tilde{B}$ and $A^{\otimes(n+\gamma)} \otimes \varepsilon^j = A^{\otimes n} \otimes \varepsilon^j$ for each μ-truncated unit vectors ε^j, then $A^{\otimes(n+\gamma)} = A^{\otimes n}$.*

Proof. Let i and j be nodes in G, and let n be an integer such that $n \geqslant \tilde{B}$. Since $\tilde{B} \geqslant ep(G) + \gamma(G) + N - 2$ and $\gamma = \gamma(G_{\mathrm{c}})$ is a multiple of $\gamma(G)$, we derive from Lemma 3 that $A^{\otimes n+\gamma}_{i,j} = -\infty$ if and only if $A^{\otimes n}_{i,j} = -\infty$. There are two cases to consider:

1. $A^{\otimes n}_{i,j} = -\infty$ and $A^{\otimes(n+\gamma)}_{i,j} = -\infty$. In this case, $A^{\otimes(n+\gamma)}_{i,j} = A^{\otimes n}_{i,j}$ trivially holds.
2. $A^{\otimes n}_{i,j} \neq -\infty$ and $A^{\otimes(n+\gamma)}_{i,j} \neq -\infty$. Recall that

$$\Big(A^{\otimes n} \otimes \varepsilon^j\Big)_i = \max_h \big\{A^{\otimes n}_{i,h} + \varepsilon^j_h\big\} = \max\big\{A^{\otimes n}_{i,j}, \max_{h\neq j}\big\{A^{\otimes n}_{i,h} - \mu\big\}\big\} \quad .$$

By definition of μ, it follows that

$$\Big(A^{\otimes n} \otimes \varepsilon^j\Big)_i = A^{\otimes n}_{i,j} \quad .$$

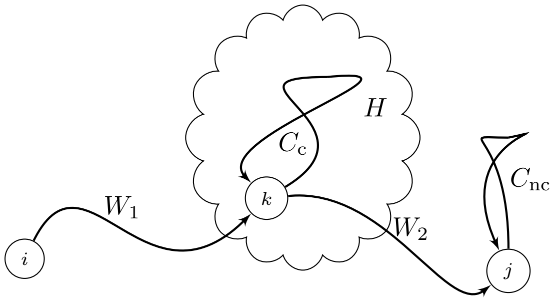

Figure 6: Walk W in proof of Proposition 7

As $n+\gamma \geqslant n$, we similarly have

$$\left(A^{\otimes(n+\gamma)} \otimes \varepsilon^j\right)_i = A^{\otimes(n+\gamma)}_{i,j} \ .$$

Hence if $A^{\otimes(n+\gamma)} \otimes \varepsilon^j = A^{\otimes n} \otimes \varepsilon^j$, then $A^{\otimes(n+\gamma)}_{i,j} = A^{\otimes n}_{i,j}$.

□

The key point for establishing our bound on matrix transients is the following upper bound on μ, which is quadratic in N. The proof uses the pumping technique developed for the explorative bound twice.

Proposition 7. $\mu \leqslant \|A\| \cdot \tilde{B}$

Proof. First, we observe that each term in the inequality to show is invariant under substituting A by $\overline{A}$. Hence we assume that $\lambda = 0$. It follows that

$$A^{\otimes n}_{i,h} \leqslant \Delta \cdot (N-1) \leqslant \Delta \cdot \tilde{B} \ . \tag{13}$$

We now give a lower bound on $A^{\otimes n}_{i,j}$ in the case that it is finite, i.e., if $\mathcal{W}^n(i,j) \neq \emptyset$. Let k be a critical node in the strongly connected component H of G_{c} with minimal distance from i and let W_1 be a shortest path from i to k. Further, let W_2 be a shortest path from k to j. Let r denote the residue of $n - \ell(W_1 \cdot W_2) - ep(G)$ modulo $\gamma(H)$, and let $t = n - \ell(W_1 \cdot W_2) - ep(G) - r$. Since $t \equiv 0 \pmod{\gamma(H)}$, and

$$t \geqslant \tilde{B} - 2(N-1) - ep(G) - \big(\gamma(H) - 1\big) \geqslant \hat{ep} \geqslant ep(H) \ ,$$

there exists a closed walk C_{c} of length t in component H starting at node k. Let $s = ep(G) + r$; then, $s \geqslant ep(G)$. Moreover, $s = n - \ell(W_1 \cdot C_{\mathrm{c}} \cdot W_2)$, and $W_1 \cdot C_{\mathrm{c}} \cdot W_2 \in \mathcal{W}(i,j)$. It follows that $\gamma(G)$ divides s, because $\mathcal{W}^n(i,j) \neq \emptyset$. Hence there exists a closed walk C_{nc} of length s starting at node j.

Now define $W = W_1 \cdot C_{\mathrm{c}} \cdot W_2 \cdot C_{\mathrm{nc}}$. Clearly, $\ell(W) = n$ and

$$n(W) \geqslant \delta \cdot (n-t) \geqslant \delta \cdot \big(2(N-1) + ep(G) + \gamma(H) - 1\big) \ ,$$

and so

$$A^{\otimes n}_{i,j} \geqslant \delta \cdot \big(2(N-1) + ep(G) + \hat{\gamma} - 1\big) \geqslant \delta \cdot \tilde{B} \ . \tag{14}$$

From (13) and (14) follows $\mu \leqslant (\Delta - \delta) \cdot \tilde{B} = \|A\| \cdot \tilde{B}$. □

Combined with our upper bounds on the system transient established in Theorems 4 and 5, Propositions 6 and 7 give a repetitive upper bound and an explorative upper bound on the matrix transient.

Theorem 6. *The transient of an irreducible matrix is at most equal to the minimum of the repetitive bound*

$$\max\left\{\tilde{B}\ ,\ \frac{\|A\| \cdot \tilde{B} + \big(\tilde{\Delta}_{\mathrm{nc}} - \delta\big) \cdot (N-1)}{\lambda - \lambda_{\mathrm{nc}}} + 1\ ,\ (\hat{g} - 1) + 2\,\hat{g} \cdot (N-1)\right\} \ ,$$

and the explorative bound

$$\max\left\{\tilde{B}\ ,\ \frac{\|A\| \cdot \tilde{B} + \big(\tilde{\Delta}_{\mathrm{nc}} - \delta\big) \cdot (N-1)}{\lambda - \lambda_{\mathrm{nc}}} + 1\ ,\ (\hat{\gamma} - 1) + 2\,\hat{\gamma} \cdot (N-1) + \hat{ep}\right\} \ ,$$

where $\tilde{B} = 2(N-1) + \hat{ep} + (ep(G) + \hat{\gamma} - 1)$.

Note that by the bounds of Section 6, the term $\tilde{B}$ in the above bounds is at most quadratic in N. Moreover it can be removed from the maximum when λ_{nc} is finite, since in this case the critical bound dominates the term $\tilde{B}$ as $\lambda - \lambda_{\mathrm{nc}} \leqslant \Delta - \delta = \|A\|$.

Further, from the bounds of Section 6, we immediately obtain that the transient of an irreducible matrix is $O\big(\|A\| \cdot N^2/(\lambda - \lambda_{\mathrm{nc}})\big)$ if λ_{nc} is finite, and is $O(N^2)$, otherwise. In particular, for integer matrices the matrix transient is $O(\max\{N^2\,,\, \|A\| \cdot N^4\})$ for integer matrices.

9 Applications

In this section we demonstrate how our transience bounds enable the performance analysis of various distributed systems, thereby obtaining simple proofs both of known and new results.

In Section 9.1, we discuss properties of optimal cyclic schedules of a set of tasks subject to a set of restrictions. This problem arises, e.g., in manufacturing, time-sharing of processors in embedded systems, and design of compilers for scheduling loop operations for parallel and pipelined architectures. We apply our transience bounds to a naturally arising special case of restrictions, namely those whose direct dependencies in a given iteration reach back at most to the previous one, and not further back (i.e., restrictions with binary heights). For this case, we are able to state explicit upper bounds, and thereby asymptotic upper bounds, on the number of task executions from where on the schedule becomes periodic.

In Section 9.2, we discuss the transient behavior of the α network synchronizer [2]. The α-synchronizer constructs virtually synchronous rounds in a strongly connected network of processes that communicate by message passing with constant transmission delays. Its time behavior can be described by a max-plus linear system. It has hence a periodic behavior and by applying our results, we obtain upper bounds on the time after which the system is periodic. We show that our bounds are strictly better than those by Even and Rajsbaum [16]. In the case of integer matrices considered by Even and Rajsbaum, our bounds are $O(\max\{N^2\,,\, \|A\| \cdot N^3\})$ which we show to be asymptotically tight.

In Section 9.3, we further exemplify the applicability of our results to distributed algorithms by deriving upper bounds on the termination time of the Full Reversal algorithm when used for routing [17], and the time after which it is periodic when used for scheduling [3].

9.1 Cyclic scheduling

Cohen et al. [12] have observed that, in cyclic scheduling, the class of *earliest schedules* can be described as max-plus linear systems. In this section, we show how to use this fact and our general bounds to derive explicit upper bounds on transients of earliest schedules.

If a finite set $\mathcal{T}$ of tasks (each of which calculates a certain function) is to be scheduled repeatedly on different processes, precedence restrictions are implied by the data flow. These restrictions are of the form that task i may start its number n execution only after task j has finished its number $n-h$ execution. A *schedule* t maps a pair $(i,n) \in \mathcal{T} \times \mathbb{N}$ to a nonnegative integer $t(i,n)$, the time the number n execution of task i is started. Formally, if P_i denotes the processing time of task i, then a *restriction* R between two tasks i and j is an inequality of the form

$$\forall n \geqslant h_R : \; t(i,n) \geqslant t(j, n - h_R) + P_j \tag{15}$$

where h_R is called the *height* of restriction R and P_j is its *weight*.

A *uniform graph* [19] describes a set of tasks and restrictions. Formally, it is a quadruple $G^u = (\mathcal{T}, E, p, h)$ such that $(\mathcal{T}, E)$ is a directed (multi-)graph, and $p : E \to \mathbb{N}^*$ and $h : E \to \mathbb{N}$ are two functions, the *weight* and *height* function, respectively. For a walk W in G^u, let $p(W)$ be the sum of the weights of its edges and $h(W)$ the sum of the heights of its edges. An edge from i to j corresponds to a restriction R between i and j of the form (15). All incoming edges of a node j in $\mathcal{T}$ have the same weight, namely P_j.

Call G^u *well-formed* if it is strongly connected and does not contain a nonempty closed walk of height 0. Call a schedule t an *earliest schedule* if it satisfies all restrictions specified by G^u and it is minimal with respect to the point-wise partial order on schedules. Denote the maximum height in G^u by $\hat{h}$. Cohen et al. [12] showed that the earliest schedule t for well-formed G^u is unique and fulfills

$$t(i,n) = (A^{\otimes n} \otimes v)_i \tag{16}$$

for all $i \in \mathcal{T}$ and $n \geqslant 0$, where v is a suitably chosen $(\hat{h} \cdot |\mathcal{T}|)$-dimensional max-plus vector and A a suitably chosen $(\hat{h} \cdot |\mathcal{T}|) \times (\hat{h} \cdot |\mathcal{T}|)$ max-plus matrix.

In the following we will assume that heights in G^u are binary, i.e., either 0 or 1. We can do this without loss of generality [12]: Any uniform graph G^u with $\hat{h} > 1$ can be transformed into a uniform graph with binary heights and at most $\hat{h} \cdot |\mathcal{T}|$ nodes such that the earliest schedules on both are equal for corresponding tasks. However, we cannot directly apply our transience bounds on the graph $G(A)$ obtained from G^u, since $G(A)$ is not necessarily strongly connected. But we present a transformation of G^u that yields a strongly connected graph $G(A)$ in case of binary heights, and has the same earliest schedule as the original graph G^u: For every restriction between tasks i and j in G^u one can add the *redundant restriction* $t(i,n) \geqslant t(j,n-1) + P_j$ without changing the earliest schedule, since $t(j,n) \geqslant t(j,n-1)$ for all tasks j and $n \geqslant 1$. With this transformation we obtain:

Proposition 8. *If G^u is well-formed, has binary heights, and contains all redundant restrictions, then A is irreducible.*

Proof. It suffices to show that whenever there is an edge from i to j in G^u, then it also exists in $G(A)$. Because G^u contains all redundant restrictions, if there exists an edge from i to j, then there also exists an edge of height 1 from i to j. Hence there exists a walk of length 1 from i to j in G^u whose last (and only) edge has height 1. Hence, by definition of A, the entry $A_{i,j}$ is finite. This concludes the proof. □

Bounds in terms of the parameters of the original uniform graph G^u can be derived as well by relating graph parameters of G^u to parameters of $G = G(A)$. For that purpose, we denote by $\delta(G^u)$ and $\Delta(G^u)$ the minimum and maximum weight of an edge in G^u, respectively. From the definition of max-plus matrix A and initial vector v, it immediately follows that in case of binary heights, $N = |\mathcal{T}|$, $\|v\| \leqslant (|\mathcal{T}| - 1) \cdot \Delta(G^u)$, $\Delta(G) \leqslant |\mathcal{T}| \cdot \Delta(G^u)$, $\delta(G) \geqslant \delta(G^u)$,

$$\lambda(G) = \max\{p(C)/h(C) \mid C \text{ is a closed walk in } G^u\} \ ,$$

$\lambda_{\mathrm{nc}}(G)$ is at most the second largest $p(C)/h(C)$ of closed walks C in G^u, and $\hat{g}(G)$ is at most the number of links with height 1 in closed walks C in G^u with maximum $p(C)/h(C)$. As a consequence of the above bounds and the bound stated in (9) for integer matrices, the transient is $O((\Delta(G) - \delta(G)) \cdot |\mathcal{T}|^3) = O(|\mathcal{T}|^4)$, assuming constantly bounded $\delta(G^u)$ and $\Delta(G^u)$. From the reduction of arbitrary heights to binary heights, we thus obtain $O(|\mathcal{T}|^4 \hat{h}^4)$ in the general case. To the best of our knowledge, these are the first asymptotic bounds on the transient of an earliest schedule with tasks $\mathcal{T}$. It is an open problem whether these bounds in $|\mathcal{T}|$ are asymptotically tight.

9.2 Synchronizers

Even and Rajsbaum [16] presented a transience bound for a network synchronizer in a system with constant integer communication delays. They considered a variant of the α-synchronizer [2] in a centrally clocked distributed system of N processes that communicate by message passing over a strongly connected network graph G. Each link has constant transmission delay, specified in terms of central clock ticks. Processes execute the α-synchronizer after an initial boot-up phase: After receiving round n messages from all neighbors, a process proceeds to round $n+1$ and broadcasts its round $n+1$ message. Denote by $t(n)$ the vector such that $t_i(n)$ is the clock tick at which process i broadcasts its round n message. Even and Rajsbaum showed that the synchronizer becomes periodic by time $B_{\mathrm{ER}} = l_0 + 2N^2 + N$, where l_0 is an upper bound on the length of maximum weight walks with only non-critical nodes. It is easily checked that l_0 is always greater than or equal to our critical bound B_{c}.

One can show that $t(n)$ is in fact a max-plus linear system. More precisely, $t(n) = A^{\otimes n} \otimes t(0)$, where A is the adjacency matrix of the network graph G. Our bounds hence directly apply, and we obtain a repetitive bound on the transient of $(t(n))_{n \geqslant 0}$ that is strictly less than $\max\{l_0, 2N^2 - N\}$, and thus strictly less than Even and Rajsbaum's bound B_{ER}.

As an example, let us consider the "ℓ-sized cherry" graph family $H_{\ell,c}$, with $\ell \geqslant 2$ and $c \geqslant 1$, introduced by Even and Rajsbaum [16]. Each weighted graph $H_{\ell,c}$ contains $N = 4\ell$ nodes and is constructed as follows: Let $\hat{C}$ and C be two cycles of length ℓ and $\ell+1$ respectively, with edge weights $3c$, except for one link per cycle with weight $3c+1$. There exists for both $\hat{C}$ and C a path of length ℓ to a distinct node s, and an antiparallel path back. Hereby the edges in the path from s to C and from s to $\hat{C}$ have weight c, the edges in the path from $\hat{C}$ to s have weight $3c$, and from C to s, $4c$. In the following assume simultaneous boot-up of nodes, i.e., $\|v\| = \|t(0)\| = 0$.

Observing that the nodes of $\hat{C}$ are the critical nodes, $\Delta = 4c$, $\delta = c$, $N = 4\ell$, $\lambda = 3c + 1/\ell$, and $l_0 = 112c\ell^3 - 16\ell^3 - 12c\ell^2 + 4\ell - 1$, Even and Rajsbaum's bound is

$$(112c - 16)\ell^3 + (32 - 12c)\ell^2 + 8\ell - 1 \ ,$$

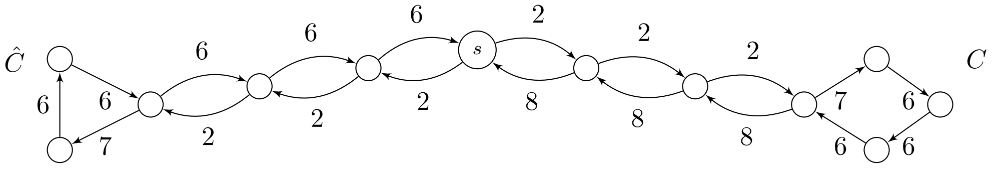

Figure 7: Graph $H_{3,2}$

resulting in an upper bound of 5711 on the transient in case of $H_{3,2}$. Since $\|A\| = 3c$, $\tilde{\Delta}_{\mathrm{nc}} = \Delta_{\mathrm{nc}} = \Delta$ and $\lambda_{\mathrm{nc}} = 3c + 1/(\ell + 1)$, we obtain for the critical bound $B_{\mathrm{c}} = 3c\ell(\ell + 1)(N - 1) + 1 = 12c\ell^3 + 9c\ell^2 - 3c\ell + 1$. Moreover for the critical subgraph G_{c}, the maximum girth of strongly connected components of G_{c} is $\hat{g} = \ell$. Thereby we may bound the transient of $(t(n))_{n\geqslant 0}$ with Theorem 4 by

$$\max\{B_{\mathrm{c}}, 2\ell N - \ell - 1\} = \max\{B_{\mathrm{c}}, 8\ell^2 - \ell - 1\} = 12c\ell^3 + 9c\ell^2 - 3c\ell + 1 \ ,$$

resulting in an upper bound of 793 on the transient in case of $H_{3,2}$.

Since Even and Rajsbaum express transmission delays with respect to a discrete global clock, all weights are integers. For the case of equal delays on all links both transience bounds are $O(N^2)$. In case of different link delays we obtain that the transience is $O(\|A\|\cdot N^3)$, where $\|A\|$ is the difference of the maximum and minimum link delay. The example graph family shows that this is asymptotically tight since Even and Rajsbaum proved that the transient for graph $H_{c,\ell}$ is $\Omega(c \cdot \ell^3) = \Omega(\|A\| \cdot N^3)$ with $\|A\| = 3c$. An adapted example graph family shows asymptotic tightness of our bounds in the general case.

9.3 Full Reversal routing and scheduling

Link reversal [8] is a versatile algorithm design paradigm which was, among others, successfully applied to routing [17] and scheduling [3]. The algorithms are of particular interest for routing in the context of mobile and ad-hoc networks, where communication graphs change frequently, e.g., not allowing for spanning tree constructions to set up routing tables. Charron-Bost et al. [9] showed that the analysis of link reversal algorithms can be reduced to the analysis of *Full Reversal*, a particularly simple algorithm on directed graphs.

The Full Reversal algorithm consists in changing the orientation of links in a directed graph according to a single rule which involves local information only: Each sink, i.e., each node without self-loop and outgoing edge, reverses all its (incoming) edges. For each integer $n \geqslant 0$, we define the *work vector* $W(n)$ by setting $W_i(n)$ to the number of reversals by node i until iteration n, i.e., the number of times node i is a sink in the execution prefix $G(0), \ldots, G(n - 1)$.

Charron-Bost et al. [9] have shown that the sequence of work vectors can be described as a *min-plus* linear dynamical system: Since initially no node has performed any step, $W(0)$ is the null vector and $W(n + 1) = A \otimes' W(n)$ where $\otimes'$ denotes the matrix multiplication in min-plus algebra[1] and

$$A_{i,j} = \begin{cases} 1 & \text{if } i \neq j \text{ and } (i,j) \text{ is an edge of } G(0) \\ 0 & \text{if } (j,i) \text{ is an edge of } G(0) \\ +\infty & \text{otherwise.} \end{cases}$$

The key point is that the min-plus matrix A only depends on the initial graph $G(0)$.

Observe that the above min-plus recurrence is equivalent to the max-plus recurrence

$$-W(n + 1) = (-A) \otimes (-W(n)) \tag{17}$$

where $-A$ is an integer matrix with $\Delta_{\mathrm{nc}} \in \{0, -\infty\}$ and $\delta = -1$. Further $\|W(0)\|= 0$.

9.3.1 Full Reversal routing

In the routing case, the initial graph $G(0)$ contains a nonempty set of *destination nodes*, which are characterized by having a self-loop.The initial graph is required to be weakly connected and without any cycle except the self-loops [9, 17]. It has been proved that for such an initial graph, the execution terminates – that is, eventually all the graphs $G(n)$ are equal – and after termination, the graph is destination-oriented, i.e., every node has a walk to some destination node. We now show how the

[1] Min-plus algebra is a variant of max-plus algebra, using min instead of max.

previously known results that the termination time of Full Reversal routing is quadratic in general [6] and linear in trees [9] directly follow from Theorem 4 or Theorem 5.

Let G be the graph that corresponds to the max-plus matrix in the recurrence relation (17). The set of critical nodes is equal to the set of destination nodes and each strongly connected component of G_{c} consists of a single node. Hence $\lambda = 0$ and $\hat{g} = 1$.

We check that $\lambda_{\mathrm{nc}} \leqslant -1/N_{\mathrm{nc}} \leqslant -1/(N-1)$, i.e., $(N-1)^2+1$ is an upper bound on the critical bound. Since $\hat{g} = 1$, we obtain from Theorem 4 that the termination time is at most $(N-1)^2+1$ if $N \geqslant 3$.

If the graph defined as the undirected support of initial graph $G(0)$ without the self-loop at the destination nodes is a *tree*, we can use our bounds to give a new proof that the termination time of Full Reversal routing is linear in N [9, Corollary 5]. In that particular case either $\lambda_{\mathrm{nc}} = -1/2$ or $\lambda_{\mathrm{nc}} = -\infty$. In both cases the critical bound is at most $2(N-1)+1$. Both Theorem 4 and Theorem 5 yield the linear bound $2(N-1)+1$, whereas Hartmann and Arguelles arrive at $2N^2$.

9.3.2 Full Reversal scheduling

When using the Full Reversal algorithm for scheduling, the undirected support of the weakly connected initial graph $G(0)$ is interpreted as a *conflict graph*: Nodes model processes and an edge between two processes signifies the existence of a shared resource whose access is mutually exclusive. The direction of an edge signifies which process is allowed to use the resource next. A process waits until it is allowed to use all its resources—that is, it waits until it is a sink—and then performs a step, that is, reverses all edges to release its resources. To guarantee liveness, the initial graph $G(0)$ is required to be acyclic; in particular, there is no self-loop in $G(0)$.

Let us consider the graph G that corresponds to the max-plus matrix in the recurrence relation (17). Contrary to the routing case, strongly connected components of the critical subgraph have at least two nodes, because there are no self-loops. By using (9), we get $N^2(N-1)/4+1$ as an upper bound on our critical bound, which shows that the transient for Full Reversal scheduling is at most cubic in the number N of processes. Malka and Rajsbaum [24, Theorem 6.4] proved by reduction to Timed Marked Graphs that the transient is at most in the order of $O(N^4)$. Thus, our bounds allow to improve this asymptotic result by an order of N.

In the case of Full Reversal scheduling on *trees* we even obtain a bound linear in N: In this case it holds that $\lambda = -1/2$, and $\lambda_{\mathrm{nc}} = -\infty$. Thus the critical bound is N. Further, $G_{\mathrm{c}} = G$ and $\hat{g} = 2$. Both Theorem 4 and Theorem 5 thus imply that $4N-3$ is an upper bound on the transient of Full Reversal scheduling on trees, which is linear in N. This was previously unknown. By contrast Hartmann and Arguelles again obtain the quadratic bound of $2N^2$.

Acknowledgments

The authors would like to thank François Baccelli, Anne Bouillard, Philippe Chrétienne, and Sergeĭ Sergeev for helpful discussions, and the anonymous reviewers for their comments, which helped to improve the paper.

References

[1] Attiya, H., V. Gramoli, A. Milani. 2010. A provably starvation-free distributed directory protocol. S. Dolev, J. Cobb, M. Fischer, M. Yung, eds. *Proc. 12th Int. Symp. Stabilization, Safety, and Security of Distributed Systems* (*SSS* 2010). *Springer Lecture Notes in Computer Science*, Vol. 6366. Springer, Heidelberg, Germany, 405–419.

[2] Awerbuch, B. 1985. Complexity of network synchronization. *J. ACM* **32**(4) 804–823.

[3] Barbosa, V. C., E. Gafni. 1989. Concurrency in heavily loaded neighborhood-constrained systems. *ACM Trans. Program. Lang. Syst.* **11**(4) 562–584.

[4] Brauer, A. 1942. On a problem of partitions. *Amer. J. Math.* **64**(1) 299–312.

[5] Bouillard, A., B. Gaujal. 2000. Coupling time of a (max,plus) matrix. INRIA Research Report 4096.

[6] Busch, C., S. Tirthapura. 2005. Analysis of link reversal routing algorithms. *SIAM J. Comput.* **35**(2) 305–326.

[7] Chandy, K. M., J. Misra. 1984. The drinking philosophers problem. *ACM Trans. Program. Lang. Syst.* **6**(4) 632–646.

[8] Charron-Bost, B., A. Gaillard, J. L. Welch, J. Widder. 2013. Link Reversal Routing with Binary Link Labels: Work Complexity. *SIAM J. Comput.* **42**(2) 634–661.

[9] Charron-Bost, B., M. Függer, J. L. Welch, J. Widder. 2015. Time Complexity of Link Reversal Routing. *ACM Transactions on Algorithms* **11**(3) 18:1–18:39.

[10] Cohen, G., D. Dubois, J.-P. Quadrat, M. Viot. 1983. Analyse du comportement périodique de systèmes de production par la théorie des dioïdes. INRIA Research Report 191.

[11] Cohen, G., D. Dubois, J.-P. Quadrat, M. Viot. 1985. A linear system-theoretic view of discrete event processes and its use for performance evaluation in manufacturing. *IEEE Trans. Automat. Control* **30**(3) 210–220.

[12] Cohen, G., P. Moller, J.-P. Quadrat, M. Viot. 1989. Algebraic tools for the performance evaluation of discrete event systems. *P. IEEE* **77**(1) 39–85.

[13] Denardo, E. V. 1977. Periods of connected networks and powers of nonnegative matrices. *Math. Oper. Res.* **2**(1) 20–24.

[14] Dubois, D., K. E. Stecke. 1990. Dynamic analysis of repetitive decision-free discrete event processes: the algebra of timed marked graphs and algorithmic issues. *Ann. Oper. Res.* **26** 151–193.

[15] Dulmage, A. L, N. S. Mendelsohn. 1964. Gaps in the exponent set of primitive matrices. *Illinois J. Math.* **8**(4) 642–656.

[16] Even, S., S. Rajsbaum. 1997. The use of a synchronizer yields maximum computation rate in distributed systems. *Theory Comput. Syst.* **30**(5) 447–474.

[17] Gafni, E. M., D. P. Bertsekas. 1987. Asymptotic optimality of shortest path routing algorithms. *IEEE Trans. Inform. Theory* **33**(1) 83–90.

[18] Goverde, R. M. P., P. H. L. Bovy, G. J. Olsder. 1999. The max-plus algebra approach to transportation problems. H. Meersmann, E. van de Vorde, W. Winkelmans, eds. *Proc. 8th World Conf. Transport Research* (*WCTR* 1999). Pergamon, Amsterdam, The Netherlands, 377–390.

[19] Hanen, C., A. Munier. 1995. Cyclic scheduling on parallel processes: an overview. P. Chrétienne, E. G. Coffman, Jr., J. K. Lenstra, Z. Liu, eds. *Scheduling Theory and its Applications.* John Wiley & Sons, Chichester, NY.

[20] Hartmann, M., C. Arguelles. 1999. Transience bounds for long walks. *Math. Oper. Res.* **24**(2) 414–439.

[21] Heidergott, B., G. J. Olsder, J. van der Woude. 2006. *Max Plus at Work.* Princeton University Press, Princeton, NJ.

[22] Kim, K. H. 1979. An extension of the Dulmage-Mendelsohn theorem. *Linear Algebra Appl.* **27** 187–197.

[23] Li, Q., J. Shao. 1993. The index set problem for Boolean (or nonnegative) matrices. *Discrete Math.* **123** 75–92.

[24] Malka, Y., S. Rajsbaum. 1992. Analysis of distributed algorithms based on recurrence relations. S. Toueg, P. Spirakis, L. Kirousis, eds. *Proc. 5th Int. Workshop on Distributed Algorithms* (*WDAG* 1991). *Springer Lecture Notes in Computer Science*, Vol. 579. Springer, Heidelberg, Germany, 242–253.

[25] Schneider, H., M. H. Schneider. 1991. Max-balancing weighted directed graphs and matrix scaling. *Math. Oper. Res.* **16**(1) 208–222.

[26] Schwarz, Š. 1970. On a sharp estimation in the theory of binary relations on a finite set. *Czechoslovak Math. J.* **20**(4) 703–714.

[27] Soto y Koelemeijer, G. 2003. *On the Behaviour of Classes of Min-Max-Plus Systems.* PhD Thesis, TU Delft.

[28] Tirthapura, S., M. Herlihy. 2006. Self-stabilizing distributed queuing. *IEEE Trans. Parall. Dist. Syst.* **17**(7) 646–655.

[29] Welch, J. L., J. E. Walter. 2012. *Link Reversal Algorithms.* Morgan & Claypool, San Rafael, CA.

[30] Wielandt, H. 1950. Unzerlegbare, nicht-negative Matrizen. *Math. Z.* **52** 642–648.