



Transience Bounds for Distributed Algorithms

Bernadette Charron-Bost¹, Matthias Függer², and Thomas Nowak¹

¹ CNRS, LIX, École polytechnique, France

{charron, nowak}@lix.polytechnique.fr

² ECS Group, TU Wien, Austria

fuegger@ecs.tuwien.ac.at

Abstract. A large variety of distributed systems, like some classical synchronizers, routers, or schedulers, have been shown to have a periodic behavior after an initial transient phase (Malka and Rajsbaum, WDAG 1991). In fact, each of these systems satisfies recurrence relations that turn out to be linear as soon as we consider max-plus or min-plus algebra. In this paper, we give a new proof that such systems are eventually periodic and a new upper bound on the length of the initial transient phase. Interestingly, this is the first asymptotically tight bound that is linear in the system size for various classes of systems. Another significant benefit of our approach lies in the straightforwardness of arguments: The proof is based on an easy convolution lemma borrowed from Nachtigall (Math. Method. Oper. Res. 46) instead of purely graph-theoretic arguments and involved path reductions found in all previous proofs.

1 Introduction

A large variety of distributed systems, like the network synchronizers in [13], or the distributed link reversal algorithms [22] which can be used for routing [14], scheduling [4], distributed queuing [21, 2], or resource allocation [7] have been shown to have a periodic behavior after an initial transient phase. Each of these systems satisfies a recurrence relation [16] that turns out to be linear as soon as one considers max-plus or min-plus algebra. Indeed, the fundamental theorem in these algebras—an analog of the Perron-Frobenius theorem—shows that any linear system with irreducible system matrix is periodic from some index, called the *transient* of the system. For all the above mentioned systems, the study of the transient plays a key role in characterizing performances: For example, in the case of link reversal routing, the system transient is equal to the time complexity of the routing algorithm.

Hartmann and Arguelles [17] have shown that the transients of 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 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.

The problem of bounding the transients has already been studied [1, 5, 8, 13, 17, 20]. A polynomial bound has been established by Even and Rajsbaum [13] for a large class of synchronizers in the specific case of integer delays, and Hartmann and Arguelles [17] have established an upper bound on the transients of general linear systems. More recently, Charron-Bost et al. [8] have given two new transience bounds, called the *repetitive* and the *explorative* bound, which both improve on the Even and Rajsbaum's synchronizer bound, and are incomparable with Hartmann and Arguelles' bound. In each of the above mentioned works [13, 17, 8], the approach is graph-theoretic in nature: The problem of bounding from above the transient is reduced to the study of long walks in a specific weighted graph, and the key technical point is then to design sophisticated walk transformers that do not modify the weights of some walks. On the contrary, our approach here is more algebraic in the sense that instead of walk transformers, we use a convolution lemma borrowed from [18]. Based on simpler and more direct arguments than the classical one, we thus obtain an easy proof that linear systems are eventually periodic, and a new upper bound on the length of the initial transient phase which is in $O(N^4)$ for linear systems of dimension N with integer matrices.

Interestingly, the proof simplification does not damage the resulting bound: Our new transience bound is of the same order as the repetitive and explorative bounds, and so is incomparable with Hartmann and Arguelles' bound, and better than Even and Rajsbaum's bound. Moreover, like the repetitive and explorative bounds by Charron-Bost et al., it is linear in the size of the system in various classes of linear systems whereas both Even and Rajsbaum's bound and Hartmann and Arguelles' bounds are both intrinsically at least quadratic.

Finally, we demonstrate how our general transience bound enables the performance analysis of a large variety of distributed systems. First, we immediately derive a general transience bound for a large class of synchronizers, and we quantify how our synchronizer bound is better than that of Even and Rajsbaum [13] in the specific case of integer delays. From this synchronizer example, we show that our transience bound is asymptotically tight. Our result also applies to the analysis of the performance of distributed routers and schedulers based on the link-reversal algorithms: We obtain an $O(N^3)$ transience bound, improving the $O(N^4)$ bound established by Malka and Rajsbaum [16], and an $O(N)$ bound 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] directly follows from our transience bound.

2 Preliminaries

This section introduces max-plus linear systems, their graph interpretation, and discusses general properties of eventually periodic sequences.

As observed in [16], the behavior of distributed systems like network synchronizers and distributed link reversal algorithms can be described by a sequence $(x(n))_{n \geq 0}$ of N -dimensional vectors that satisfy a recurrence relation of the fol-

lowing form:

$$\forall i \in [N], \forall n \geq 0, \quad x_i(n+1) = \max_{j \in \mathbf{N}_i} (x_j(n) + a_{ij}) \quad (1)$$

where the a_{ij} 's are reals, and \mathbf{N}_i is a subset of $[N] = \{1, \dots, N\}$. Trivially, a system of this form corresponds to the recurrence relation

$$\forall i \in [N], \forall n \geq 0, \quad x_i(n+1) = \max_{j \in [N]} (x_j(n) + A_{ij}) \quad (2)$$

with $A_{ij} = a_{ij}$ if j in \mathbf{N}_i and $-\infty$ otherwise. Recurrence (2) turns out to be linear system in the max-plus algebra, i.e., when replacing '+' by 'max' and ' \times ' by '+', over the set $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty\}$. Let us recall that the *max-plus product* of two matrices A and B is defined by $(AB)_{ij} = \max_k (A_{ik} + B_{kj})$. With this definition, (2) is equivalent to

$$\forall n \geq 0, \quad x(n+1) = Ax(n) . \quad (3)$$

Given an initial vector $v \in \mathbb{R}^N$, recurrence (3) admits a unique solution given by

$$\forall n \geq 0, \quad x(n) = A^n v . \quad (4)$$

Thus the analysis of distributed systems whose behaviors are controlled by a recurrence relation of the form (1) leads to study the systems (4).

To every matrix A naturally corresponds a weighted digraph which we denote by $G(A)$: Its node set is $[N]$ and there exists an edge from node i to node j with weight A_{ij} if and only if A_{ij} is finite. Matrix A is *irreducible* if $G(A)$ is strongly connected. A *walk* consists of a sequence of successive edges, as well as a start and an end node. We denote by $\ell(W)$ the length of walk W . A *path* is a walk without node repetition. A *closed* walk is a walk whose start node is equal to its end node. A *cycle* is a closed walk in which the only node repetition is the start and end node. The *circumference* $\Gamma(G)$ of graph G is the maximum cycle length in G .³ We write $\mathcal{W}(i, j)$ for the set of walks from node i to node j , and $\mathcal{W}(i \rightarrow)$ for the set of walks starting from i . Further, $\mathcal{W}^n(i, j)$ and $\mathcal{W}^n(i \rightarrow)$ denote the set of walks in $\mathcal{W}(i, j)$ and $\mathcal{W}(i \rightarrow)$ of length n , respectively.

We write $p(W)$ for the edge weight of walk W , i.e., the sum of the weights of its edges. Given a vector v , the edge-node weight $p_v(W)$ is equal to $p(W) + v_j$, where j is W 's end node. These definitions yield a correspondence between the matrix power A^n as well as the vector $A^n v$, and the maximum weight of walks of length n in $G(A)$, namely:

$$\begin{aligned} (A^n)_{ij} &= \max\{p(W) : W \in \mathcal{W}^n(i, j)\} \\ (A^n v)_i &= \max\{p_v(W) : W \in \mathcal{W}^n(i \rightarrow)\} \end{aligned} \quad (5)$$

In the following, we write A_{ij}^n instead of $(A^n)_{ij}$, as no confusion can arise. A closed walk is *critical* if its average edge weight is maximum, i.e., if it is equal

³ The computation of the circumference of a graph is NP-hard in the number of nodes. However, the circumference is always upper bounded by the number of nodes.

to the *rate* of A defined by

$$\lambda(A) = \max \{p(C)/\ell(C) : C \text{ is a closed walk in } G(A)\} \quad , \quad (6)$$

and simply denoted λ when no confusion can arise. The set of critical closed walks induces the *critical subgraph* G_c whose nodes are called *critical nodes*. It is well-known (e.g., see [15]) that every closed walk in G_c is necessarily critical.

Let p be a strictly positive integer and ϱ a real number. A sequence $f : \mathbb{N} \rightarrow \overline{\mathbb{R}}$ is *eventually periodic with period p and ratio ϱ* if there exists an integer T such that:

$$\forall n \geq T : f(n+p) = f(n) + p \cdot \varrho \quad (7)$$

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

For every period p , there exists a unique minimal transient T_p that satisfies Equation (7). The next lemma shows that these minimal transients do, in fact, not depend on p . We will henceforth call this common value *the transient* of f . Due to limitation of space, its proof is postponed to the appendix.

Lemma 1 *Let $f : \mathbb{N} \rightarrow \overline{\mathbb{R}}$ be eventually periodic. Let p and q be two periods of f with respective minimal transients T_p and T_q . Then $T_p = T_q$.*

The next two elementary lemmas both play an important role in our approach: They state transience bounds for a sequence obtained by the element-wise composition of two eventually periodic sequences f and g with common ratio in terms of the transients of f and g .

Lemma 2 *Let $f, g : \mathbb{N} \rightarrow \overline{\mathbb{R}}$ be eventually periodic with common ratio ϱ and respective transients T_f and T_g . Then the sequence $\max\{f, g\}$ is eventually periodic with ratio ϱ and transient at most $\max\{T_f, T_g\}$.*

In analogy to classical convolution, one defines the *max-plus convolution* $f \otimes g$ of two sequences f and g as

$$(f \otimes g)(n) = \max_{n_1+n_2=n} (f(n_1) + g(n_2)) \quad . \quad (8)$$

Nachtigall [18] then proved a bound on the transient of the max-plus convolution of two eventually periodic sequences f and g with the same ratio.

Lemma 3 ([18, Lemma 6.1]) *Let $f, g : \mathbb{N} \rightarrow \overline{\mathbb{R}}$ be eventually periodic with common ratio ϱ and respective transients T_f and T_g . Then the convolution $f \otimes g$ is eventually periodic with ratio ϱ and transient at most equal to $T_f + T_g + p - 1$ if p denotes a common period to f and g .*

The notion of eventual periodicity naturally extends to vectors: a sequence of vectors $(x(n))_{n \geq 0}$ is *eventually periodic with period p and ratio ϱ* if each sequence $(x_i(n))_{n \geq 0}$ is eventually periodic with period p and ratio ϱ . Its *transient* is the maximum transient of the $(x_i(n))_{n \geq 0}$'s.

3 Transience Bound

The fundamental theorem in max-plus algebra [11] shows that any linear system as defined in (4) is eventually periodic with ratio equal to the rate λ of A if A is irreducible. In this section, we establish an effective upper bound on the transient of these systems. In this way, we also give an alternative proof of the fundamental theorem that is simpler and more direct than the classical one.

3.1 Proof Strategy

By definition of the max-plus matrix product, the i -th component of $x(n)$ is equal to

$$x_i(n) = \max_{j \in [N]} (A_{ij}^n + v_j) .$$

In view of Lemma 2, our strategy will thus consist in showing that each sequence $(A_{ij}^n + v_j)_{n \geq 0}$ is eventually periodic. In the case either i or j is a critical node, the question is solved by Nachtigall [18, Lemma 3.2] and Even and Rajsbaum [13, Theorem 6]: they independently showed that in this case, the sequence $(A_{ij}^n)_{n \geq 0}$ is eventually periodic, and gave effective upper bounds on the transients. For simplicity, every sequence $(A_{ij}^n)_{n \geq 0}$ will be denoted by A_{ij} in the following.

Lemma 4 *Let A be an irreducible $N \times N$ matrix, and λ the rate of A . If k is a node of a critical cycle C of length ℓ , then both sequences A_{ik} and A_{ki} are eventually periodic with period ℓ , ratio λ , and transient at most $\ell \cdot (N - 1)$.*

In the appendix, we give a proof of this lemma that essentially follows the one by Nachtigall [18].

In the general case where neither i nor j is critical, we observe that for any pair of nonnegative integers (n_1, n_2) such that $n = n_1 + n_2$, we have

$$A_{ij}^n = \max_{k \in [N]} (A_{ik}^{n_1} + A_{kj}^{n_2}) .$$

It trivially follows that

$$A_{ij}^n = \max_{n_1+n_2=n} \left(\max_{k \in [N]} (A_{ik}^{n_1} + A_{kj}^{n_2}) \right) = \max_{k \in [N]} \left(\max_{n_1+n_2=n} (A_{ik}^{n_1} + A_{kj}^{n_2}) \right) ,$$

and so

$$A_{ij}^n = \max_{k \in [N]} ((A_{ik} \otimes A_{kj})(n)) .$$

So we can write

$$x_i(n) = \max_{j \in [N]} \max_{k \in [N]} ((A_{ik} \otimes A_{kj})(n) + v_j) . \quad (9)$$

If k is a critical node which lies on a critical cycle of length ℓ_k , then Lemmas 3 and 4 imply that each sequence $A_{ik} \otimes A_{kj}$ is eventually periodic, with period ℓ_k , ratio λ , and a transient that is bounded from above by $2\ell_k \cdot (N - 1) + \ell_k - 1$.

Thus, if we could show that the range of the second maximum in Equation (9) can be restricted to the critical nodes, i.e.,

$$x_i(n) = \max_{j \in [N]} \max_{k \in G_c} ((A_{ik} \otimes A_{kj})(n) + v_j) , \quad (10)$$

we would readily obtain a bound on the transience of the sequence x_i from Lemma 2. We will, indeed, show in Section 3.2 that Equation (10) holds for n at least equal to some integer n_c . An effective upper bound on n_c finally allows us to conclude with a bound on the transience of x_i .

The following lemma provides a condition ensuring that Equation (10) holds.

Lemma 5 *If there exists a walk of maximum p_v -weight in $\mathcal{W}^n(i \rightarrow)$ that contains a critical node, then*

$$x_i(n) = \max_{j \in [N]} \max_{k \in G_c} ((A_{ik} \otimes A_{kj})(n) + v_j) .$$

Proof. From (9), we trivially derive that

$$x_i(n) \geq \max_{j \in [N]} \max_{k \in G_c} ((A_{ik} \otimes A_{kj})(n) + v_j) .$$

Conversely, let W be a walk of maximum p_v -weight in $\mathcal{W}^n(i \rightarrow)$ that contains a critical node k_0 , and let j_0 be the end node of W . By (5), we have $x_i(n) = p(W) + v_{j_0}$. Let us decompose W into $W = W_1 \cdot W_2$ such that W_1 ends at k_0 . Setting $\ell_1 = \ell(W_1)$ and $\ell_2 = \ell(W_2)$, we get

$$x_i(n) = p(W_1) + p(W_2) + v_{j_0} \leq A_{ik_0}^{\ell_1} + A_{k_0j_0}^{\ell_2} + v_{j_0} .$$

Therefore,

$$x_i(n) \leq \max_{n_1+n_2=n} (A_{ik_0}^{n_1} + A_{k_0j_0}^{n_2} + v_{j_0}) = (A_{ik_0} \otimes A_{k_0j_0})(n) + v_{j_0} .$$

It follows that

$$x_i(n) \leq \max_{j \in [N]} \max_{k \in G_c} ((A_{ik} \otimes A_{kj})(n) + v_j) ,$$

which concludes the proof.

3.2 Critical Bound

We next show that there always exists an integer n_c such that the condition in Lemma 5 holds for all $n \geq n_c$, and we give an effective upper bound B_c on n_c which we call the *critical bound*. For that, we compare a maximum p_v -weight walk \hat{W} in $\mathcal{W}^n(i \rightarrow)$ that does not visit a critical node with a walk W constructed from \hat{W} such that it “pumps” its weight in a critical cycle as often as possible, but still has the same length n as \hat{W} . Since \hat{W} may not use critical cycles, it will, when compared to W , on average lose weight in each of its cycles. We arrive at

a bound B_c , such that, if \hat{W} has length $n \geq B_c$ then \hat{W} has weight less than W and thus cannot be a maximum p_v -weight walk in $\mathcal{W}^n(i \rightarrow)$.

Let $\|v\| = \max_{i \in [N]}(v_i) - \min_{i \in [N]}(v_i)$. Denote by δ the minimal finite entry of A and by Δ_{nc} the maximum A_{ij} where both i and j are non-critical nodes. Further let

$$\lambda_{nc} = \max \{p(C)/\ell(C) : C \text{ is a closed walk with no critical node in } G(A)\} .$$

Lemma 6 (Critical Bound) *For all $i \in [N]$ and $n \geq 0$, each walk with maximum p_v -weight in $\mathcal{W}^n(i \rightarrow)$ contains a critical node if*

$$n \geq B_c = \max \left\{ N, \frac{\|v\| + (\Delta_{nc} - \delta)(N - 1)}{\lambda - \lambda_{nc}} \right\} .$$

Proof. We first reduce the case of a graph G with arbitrary λ to the case where $\lambda = 0$: Let G_0 be the graph constructed from G by subtracting λ from all of G 's edge weights. Clearly then G_0 has $\lambda = 0$ and the (unweighted) critical subgraphs of G and G_0 are the same. Further W is a walk with maximum p_v -weight in G if and only if it is a walk with maximum p_v -weight in G_0 . The graph parameters δ , Δ_{nc} , and λ_{nc} of G_0 are obtained by subtracting G 's λ from the respective graph parameters of G . The lemma's bound B_c thus is the same for graphs G and G_0 , and we may safely assume that $\lambda = 0$.

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

We now consider the case $\lambda_{nc} \neq -\infty$. We proceed by contradiction: Suppose that there exists an integer n such that $n \geq B_c$, a node i and a walk of maximum p_v -weight in $\mathcal{W}^n(i \rightarrow)$ with non-critical nodes only; let \hat{W} be such a 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 critical node k , and then a prefix W_c 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 . Let W_3 be the walk such that $W_0 = W_c \cdot W_3$. Further let C be a critical cycle starting at k .

We distinguish two cases for n , namely (a) $n \geq \ell(W_c) + \ell(W_2)$, and (b) $n < \ell(W_c) + \ell(W_2)$.

Case A: Let $m \in \mathbb{N}$ be the quotient in the Euclidean division of $n - \ell(W_c) - \ell(W_2)$ by $\ell(C)$, and choose W_1 to be a prefix of C of length $n - (\ell(W_c) + \ell(W_2) + m \cdot \ell(C))$ (see Figure 1). Clearly W_1 starts at k . If we set $W = W_c \cdot W_2 \cdot C^m \cdot W_1$, we get $\ell(W) = n$ and

$$p_v(W) \geq \min_{1 \leq j \leq N} (v_j) + p(W_c) + p(W_2) + p(W_1) \quad (11)$$

since we assume $\lambda = 0$.

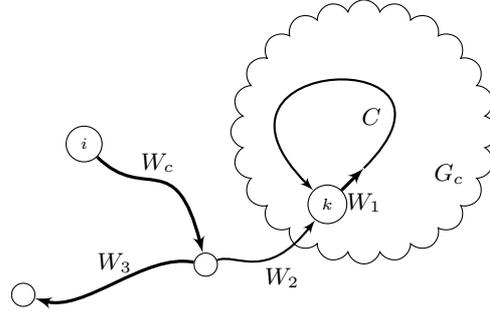


Fig. 1. Walk W in proof of Lemma 6

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

$$\begin{aligned}
 p_v(\hat{W}) &\leq p_v(W_0) + \lambda_{nc} \cdot (\ell(\hat{W}) - \ell(W_0)) \\
 &\leq \max_{1 \leq j \leq N} (v_j) + p(W_0) + \lambda_{nc} \cdot (\ell(\hat{W}) - \ell(W_0))
 \end{aligned} \tag{12}$$

By assumption $p_v(\hat{W}) \geq p_v(W)$, and from (11), (12), and $\lambda_{nc} < 0$ we therefore obtain

$$\begin{aligned}
 \ell(\hat{W}) &\leq \frac{\|v\| + p(W_3) - p(W_1) - p(W_2)}{-\lambda_{nc}} + \ell(W_0) \\
 &\leq \frac{\|v\| + \Delta_{nc} \ell(W_3) - \delta(\ell(W_1) + \ell(W_2))}{-\lambda_{nc}} + \ell(W_0)
 \end{aligned} \tag{13}$$

Denote by N_{nc} the number of non-critical nodes. The following three inequalities trivially hold: $\ell(W_3) \leq N_{nc} - 1$, $\lambda_{nc} \geq \delta$, and $\ell(W_1) < N - N_{nc}$. Since there is at least one critical node, we have $\ell(W_3) < N - 1$. Moreover from the minimality constraint for the length of W_2 follows that $\ell(W_2) + \ell(W_0) \leq N_{nc}$. Thereby

$$\ell(\hat{W}) < \frac{\|v\| + (\Delta_{nc} - \delta)(N - 1)}{-\lambda_{nc}}, \tag{14}$$

a contradiction to $n \geq B_c$. The lemma follows for case A.

Case B: In this case $\ell(W_c) \leq n < \ell(W_c) + \ell(W_2)$, and we set $W = W_c \cdot W_2'$, where W_2' is a prefix of W_2 , such that $\ell(W) = n$. Hence,

$$p_v(W) \geq \min_{1 \leq j \leq N} (v_j) + p(W_c) + p(W_2') . \tag{15}$$

We again obtain (12). By assumption $p_v(\hat{W}) \geq p_v(W)$, and by similar arguments as in case A we derive

$$\ell(\hat{W}) \leq \frac{\|v\| + p(W_3) - p(W_2')}{-\lambda_{nc}} + \ell(W_0)$$

and since W'_2 is a prefix of W_2 with $\ell(W'_2) < \ell(W_2)$,

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

which is less or equal to the bound obtained in (13) of case A. By similar arguments as in case A, the lemma follows in case B.

3.3 Transience Bound

By combination of the above lemmas, we finally obtain the following transience bound.

Theorem 1 *Let A be an irreducible $N \times N$ matrix and let v be a vector in \mathbb{R}^N . Then the sequence of vectors $x(n) = A^n v$ is eventually periodic with ratio λ , and its transient is at most*

$$\max \left\{ \frac{\|v\| + (\Delta_{nc} - \delta)(N - 1)}{\lambda - \lambda_{nc}} , \Gamma_c \cdot (2N - 1) - 1 \right\} ,$$

where Γ_c is the circumference of the critical graph G_c .

Proof. From Lemmas 5 and 6, we know that each i -th component of $x(n)$ equals

$$x_i(n) = \max_{j \in [N]} \max_{k \in G_c} ((A_{ik} \otimes A_{kj})(n) + v_j)$$

when $n \geq B_c$. For each critical node k , let ℓ_k denote the length of a critical cycle containing k . By Lemmas 1 and 4, we obtain that all sequences A_{ik} and A_{kj} are eventually periodic, with period ℓ_k , ratio λ , and a transient less or equal to $\Gamma_c \cdot (N - 1)$ because $\ell_k \leq \Gamma_c$. Lemma 3 shows that the sequence $((A_{ik} \otimes A_{kj})(n) + v_j)_{n \geq 0}$ is eventually periodic, with ratio λ , and a transient less or equal to $2\Gamma_c \cdot (N - 1) + \Gamma_c - 1$. By Lemma 2, the same property holds for the sequence $(\max_{j \in [N]} \max_{k \in G_c} ((A_{ik} \otimes A_{kj})(n) + v_j))_{n \geq 0}$. This proves that each sequence x_i is eventually periodic, with ratio λ , and a transient at most equal to

$$\max \{ B_c , \Gamma_c \cdot (2N - 1) - 1 \} . \quad (16)$$

This concludes the proof.

In the case each finite A 's entry is an integer, the term $\lambda - \lambda_{nc}$ cannot become arbitrarily small: If N_{nc} denotes the number of non-critical nodes, then

$$\frac{1}{\lambda - \lambda_{nc}} \leq (N - N_{nc}) \cdot N_{nc} \leq \frac{N^2}{4} . \quad (17)$$

From that, we immediately derive that our critical bound, and so our transient bound for linear systems with integer matrices is in $O(N^3)$.

3.4 Comparison with Previous Bounds

Hartmann and Arguelles [17] stated the following bound on the transient of the linear system:

$$B_{\text{HA}} = \max \left\{ \frac{\|v\| + N \cdot (\Delta - \delta)}{\lambda - \lambda^0}, 2N^2 \right\} \quad (18)$$

Here, λ^0 is a parameter of the max-balanced reweighted graph [19] of G and Δ is the maximum edge weight in G .

The first term in (18) corresponds to our critical bound, but is incomparable with it in general. The second term in (18), namely $2N^2$, is always greater than the second term in the maximum in our bound, $\Gamma_c \cdot (2N - 1) - 1$, because of the trivial estimate $\Gamma_c \leq N$. As demonstrated in the next section, a major difference to our bound is that (18) is inherently quadratic in N and hence prohibits a subquadratic analysis of the transient.

Charron-Bost et al. [8] gave two transience bounds for linear systems—the repetitive and the explorative bound—which both allow for a subquadratic analysis of the transient. The repetitive bound is equal to $\max\{B_c, \hat{g} \cdot (2N - 1) - 1\}$ where \hat{g} is the maximum girth of strongly connected components of the critical graph G_c . It is always smaller than our bound in Theorem 1 because $\hat{g} \leq \Gamma_c$. The explorative bound is equal to $\max\{B_c, \hat{\gamma} \cdot (2N - 1) - 1 + \hat{e}p\}$ where $\hat{\gamma}$ is the maximum cyclicity and $\hat{e}p$ the maximum *exploration penalty* of strongly connected components of the critical graph G_c . The exploration penalty of a graph is the transient of the sequence of matrix powers of its unweighted adjacency matrix. The explorative bound is, in general, incomparable with our bound in Theorem 1.

4 Applications

As stated in the Introduction, various distributed algorithms correspond to linear max-plus systems. In this section, we explain how our transience bound immediately applies to such distributed algorithms, and allows us to analyze their performances.

4.1 Synchronizers

Even and Rajsbaum [13] presented a transience bound for a type of network synchronizers in a system with constant integer communication delays. They considered a variant of the α -synchronizer [3] 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 its neighbors, a process proceeds to round $n + 1$ and broadcasts its round $n + 1$ message. Denote

The Full Reversal algorithm comprises a single rule: Each sink reverses all its (incoming) edges. Given a weakly connected initial graph G_0 without antiparallel edges, we consider a *greedy* execution of Full Reversal as a sequence $(G_t)_{t \geq 0}$ of graphs, where G_{t+1} is obtained from G_t by reversing the edges of *all* sinks in G_t . As no two sinks in G_t can be adjacent, G_{t+1} is well-defined. For each $t \geq 0$ we define the *work vector* $W(t)$ by setting $W_i(t)$ to the number of reversals of node i until iteration t , i.e., the number of times node i is a sink in the execution prefix G_0, \dots, G_{t-1} .

Charron-Bost et al. [9, Corollary 2] have shown that the sequence of work vectors can be described as a max-plus linear system. More precisely, we have $W(0) = 0$ and $-W(t+1) = (-A) \cdot (-W(t))$, where $A_{i,j} = 1$ and $A_{j,i} = 0$ if (i, j) is an edge of G_0 ; otherwise $A_{i,j} = +\infty$. Observe that $-A$ is a matrix with integer weights, $\Delta_{nc} \in \{0, -1\}$ and $\delta = -1$. Our results, in particular, imply the eventual periodicity of the Full Reversal algorithm.

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 without these self-loops is required to be weakly connected and acyclic [9, 14]. It was shown that for such initial graphs, the execution terminates (eventually all G_t 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 previously known results that the termination time of Full Reversal routing is quadratic in general [6] and linear in trees [9] directly follows from Theorem 1.

The set of critical nodes is equal to the set of destination nodes and each strongly critical component of G_c consists of a single node. Hence $\lambda = 0$ and $\lambda_{nc} \leq -1/N_{nc} \leq -1/(N-1)$, i.e., $(N-1)^2$ is an upper bound on the critical bound. Since $\Gamma_c = 1$, we obtain from Theorem 1 that the termination time is at most $(N-1)^2$, which improves the asymptotic quadratic bound given by Busch and Tirthapura [6].

If the undirected support of initial graph G_0 without the self-loops at the destination nodes is a *tree*, we can use our transience bound to give a new proof that the termination time of Full Reversal routing is linear in N [9, Corollary 5]. Indeed, in that particular case either $\lambda_{nc} = -1/2$ or $\lambda_{nc} = -\infty$, which both give a critical bound at most equal to $2(N-1)$. From Theorem 1 we obtain the linear transience bound of $2(N-1)$, whereas Hartmann and Arguelles' bound is $2N^2$.

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.

Contrary to the routing case, critical components have at least two nodes, because there are no self-loops. By using (17), our critical bound is upper-bounded by $N^2(N-1)/4+1$, which shows that the transient for Full Reversal scheduling is at most cubic in the number N of processes. Malka and Rajsbaum [16, 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 bound allows to improve this asymptotic result by an order of N .

In the case of Full Reversal scheduling on *trees* we even obtain a linear bound in N : In this case it holds that $\lambda = -1/2$, $\lambda_{nc} = -\infty$, and so our critical bound is 1. Further, $G_c = G$ and $\Gamma_c = 2$. Theorem 1 thus implies a linear upper bound on the transient of Full Reversal scheduling on trees of $4N-3$, which establishes a new result for these distributed schedulers. By contrast, Hartmann and Arguelles again can only obtain the quadratic bound of $2N^2$.

References

1. Akian, M., Gaubert, S., Walsh, C.: Discrete Max-Plus Spectral Theory. In: Litvinov, G.L., Maslov, V.P. (eds.) *Idempotent Mathematics and Mathematical Physics*, pp. 53–78. AMS, Providence (2005)
2. Attiya, H., Gramoli, V., Milani, A.: A Provably Starvation-Free Distributed Directory Protocol. In: Dolev, S., Cobb, J., Fischer, M., Yung M. (eds.) *SSS 2010*. LNCS, vol. 6366, pp. 405–419. Springer, Heidelberg (2010)
3. Awerbuch, B.: Complexity of Network Synchronization. *J. ACM* 32, 804–823 (1985)
4. Barbosa, V.C., Gafni, E.: Concurrency in Heavily Loaded Neighborhood-Constrained Systems. *ACM T. Progr. Lang. Sys.* 11, 562–584 (1989)
5. Bouillard, A., Gaujal, B.: Coupling Time of a (max,plus) Matrix. In: *Workshop on Max-Plus Algebra at the 1st IFAC Symposium on System Structure and Control*. Elsevier, Amsterdam (2001)
6. Busch, C., Tirthapura, S.: Analysis of Link Reversal Routing Algorithms. *SIAM J. Comput.* 35, 305–326 (2005)
7. Chandy, K.M., Misra, J.: The Drinking Philosophers Problem. *ACM T. Progr. Lang. Sys.* 6, 632–646 (1984)
8. Charron-Bost, B., Függer, M., Nowak, T.: New Transience Bounds for Long Walks (2012) [arXiv:1209.3342](https://arxiv.org/abs/1209.3342) [cs.DM]
9. Charron-Bost, B., Függer, M., Welch, J.L., Widder, J.: Full Reversal Routing as a Linear Dynamical System. In: Kosowski, A., Yamashita, M. (eds.) *SIROCCO 2011*. LNCS, vol. 6796, pp. 101–112. Springer, Heidelberg (2011)
10. Charron-Bost, B., Függer, M., Welch, J.L., Widder, J.: Partial is Full. In: Kosowski, A., Yamashita, M. (eds.) *SIROCCO 2011*. LNCS, vol. 6796, pp. 113–124. Springer, Heidelberg (2011)
11. Cohen, G., Dubois, D., Quadrat, J.-P., Viot, M.: Analyse du comportement périodique de systèmes de production par la théorie des diodes. *INRIA Research Report 191* (1983)
12. Dubois, D., Stecke, K.E.: Dynamic Analysis of Repetitive Decision-Free Discrete Event Processes: The Algebra of Timed Marked Graphs and Algorithmic Issues. *Ann. Oper. Res.* 26, 151–193 (1990)

13. Even, S., Rajsbaum, S.: The Use of a Synchronizer Yields Maximum Computation Rate in Distributed Systems. *Theor. Comput. Syst.* 30, 447–474 (1997)
14. Gafni, E.M., Bertsekas, D.P.: Asymptotic Optimality of Shortest Path Routing Algorithms. *IEEE T. Inform. Theory.* 33, 83–90 (1987)
15. Heidergott, B., Olsder, G.J., van der Woude, J.: *Max Plus at Work*. Princeton University Press, Princeton (2006)
16. Malka, Y., Rajsbaum, S.: Analysis of Distributed Algorithms Based on Recurrence Relations. In: Toueg, S., Spirakis, P., Kirousis, L. (eds.) *WDAG 1991*. LNCS, vol. 579, pp. 242–253. Springer, Heidelberg (1992)
17. Hartmann, M., Arguelles, C.: Transience Bounds for Long Walks. *Math. Oper. Res.* 24, 414–439 (1999)
18. Nachtigall, K.: Powers of Matrices over an Extremal Algebra with Applications to Periodic Graphs. *Math. Method. Oper. Res.* 46, 87–102 (1997)
19. Schneider, H., Schneider, M.H.: Max-Balancing Weighted Directed Graphs and Matrix Scaling. *Math. Oper. Res.* 16, 208–222 (1991)
20. Soto y Koelemeijer, G.: *On the Behaviour of Classes of Min-Max-Plus Systems*. PhD Thesis, TU Delft (2003)
21. Tirthapura, S., Herlihy, M.: Self-Stabilizing Distributed Queuing. *IEEE T. Parall. Distr.* 17, 646–655 (2006)
22. Welch, J.L., Walter, J.E.: *Link Reversal Algorithms*. Morgan & Claypool, San Rafael (2012)