

On Reducing the Complexity of Matrix Clocks

Lúcia M. A. Drummond

Universidade Federal Fluminense
 Instituto de Computação
 Rua Passo da Pátria, 156
 24210-240 Niterói - RJ, Brazil
 lucia@dcc.ic.uff.br

Valmir C. Barbosa

Universidade Federal do Rio de Janeiro
 Programa de Engenharia de Sistemas e Computação, COPPE
 Caixa Postal 68511
 21941-972 Rio de Janeiro - RJ, Brazil
 valmir@cos.ufrj.br

Abstract

Matrix clocks are a generalization of the notion of vector clocks that allows the local representation of causal precedence to reach into an asynchronous distributed computation's past with depth x , where $x \geq 1$ is an integer. Maintaining matrix clocks correctly in a system of n nodes requires that every message be accompanied by $O(n^x)$ numbers, which reflects an exponential dependency of the complexity of matrix clocks upon the desired depth x . We introduce a novel type of matrix clock, one that requires only nx numbers to be attached to each message while maintaining what for many applications may be the most significant portion of the information that the original matrix clock carries. In order to illustrate the new clock's applicability, we demonstrate its use in the monitoring of certain resource-sharing computations.

Keywords: Vector clocks, matrix clocks.

Address for correspondence and proofs:

Valmir C. Barbosa, valmir@cos.ufrj.br
 Programa de Sistemas, COPPE/UFRJ
 Caixa Postal 68511
 21941-972 Rio de Janeiro - RJ
 Brazil

1. Introduction and background

We consider an undirected graph G on n nodes. Each node in G stands for a computational process and undirected edges in G represent the possibilities for bidirectional point-to-point communication between pairs of processes. A fully asynchronous distributed computation carried out by the distributed system represented by G can be viewed as a set of events occurring at the various nodes. An event is the sending of a message by a node to another node to which it is directly connected in G (a neighbor in G), or the reception of a message from such a neighbor, or yet the occurrence at a node of any internal state transition of relevance (“state” and “relevance” here are highly dependent upon the particular computation at hand, and are left unspecified).

The standard framework for analyzing such a system is the partial order, often denoted by \prec , that formalizes the usual “happened-before” notion of distributed computing [2, 15]. This partial order is the transitive closure of the more elementary relation to which the ordered pair (v, v') of events belongs if either v and v' are consecutive events at a same node in G or v and v' are, respectively, the sending and receiving of a message between neighbors in G .

At node i , we let local time be assessed by the number t_i of events that already occurred. We interchangeably adopt the notations $t_i = \text{time}_i(v)$ and $v = \text{event}_i(t_i)$ to indicate that v is the t_i th event occurring at node i , provided $t_i \geq 1$. Another important partial order on the set of events is the relation that gives the predecessor at node $j \neq i$ of an event v' occurring at node i . We say that an event v is such a predecessor, denoted by $v = \text{pred}_j(v')$, if v is the event occurring at node j such that $v \prec v'$ for which $\text{time}_j(v)$ is greatest. If no such v exists, then $\text{pred}_j(v')$ is undefined and $\text{time}_j(\text{pred}_j(v'))$ is assumed to be zero.

The relation pred_j allows the definition of vector clocks [11–14, 17], as follows. The vector clock of node i at time t_i (that is, following the occurrence of t_i events at node i), denoted by $V^i(t_i)$, is a vector whose j th component, for $1 \leq j \leq n$, is either equal to 0, if $t_i = 0$, or given by

$$V_j^i(t_i) = \begin{cases} t_i, & \text{if } j = i; \\ \text{time}_j(\text{pred}_j(\text{event}_i(t_i))), & \text{if } j \neq i, \end{cases} \quad (1)$$

if $t_i \geq 1$. In other words, $V_j^i(t_i)$ is either the current time at node i , if $j = i$, or is the time at node j that results from the occurrence at that node of the predecessor of the t_i th event of node i , otherwise. If no such event exists (i.e., $t_i = 0$), then $V_j^i(t_i) = 0$.

Vector clocks evolve following two simple rules:

- Upon sending a message to one of its neighbors, node i attaches $V^i(t_i)$ to the message, where t_i is assumed to already account for the message that is being sent.

- Upon receiving a message from node k with attached vector clock V^k , node i sets $V_i^i(t_i)$ to t_i (which is assumed to already reflect the reception of the message) and $V_j^i(t_i)$ to $\max\{V_j^i(t_i - 1), V_j^k\}$, for $j \neq i$.

It is a simple matter to verify that these rules do indeed maintain vector clocks consistently with their definition [11–14, 17]. Under these rules or variations thereof, vector clocks have proven useful in a variety of distributed algorithms to detect some types of global predicates [10, 14].

For large n , attaching a vector clock to every message that is sent is likely to become burdensome, so the question arises whether less costly implementations are possible. Under the very general assumptions we have made concerning the nature of G as a distributed system, the answer is negative: a result similar to Dilworth’s theorem on partially ordered sets [9] establishes that the size- n attachments are necessary [7]. However, it is possible to use more economical attachments if the edges of G provide FIFO communication [22], or if certain aspects of the structure of G can be taken into account [19], or yet if the full capabilities of vector clocks are not needed [1].

One natural generalization of the notion of vector clocks is the notion of matrix clocks [21, 24]. For an integer $x \geq 1$, the x -dimensional matrix clock of node i at time t_i , denoted by $M^i(t_i)$, has $O(n^x)$ components. For $1 \leq j_1, \dots, j_x \leq n$, component $M_{j_1, \dots, j_x}^i(t_i)$ of $M^i(t_i)$ is only defined for $i = j_1 = \dots = j_x$ and for $i \neq j_1 \neq \dots \neq j_x$. As in the definition of $V^i(t_i)$, $M_{j_1, \dots, j_x}^i(t_i) = 0$ if $t_i = 0$. For $t_i \geq 1$, on the other hand, we have

$$M_{j_1, \dots, j_x}^i(t_i) = \begin{cases} t_i, & \text{if } i = j_1 = \dots = j_x; \\ \text{time}_{j_x}(\text{pred}_{j_x} \dots \text{pred}_{j_1}(\text{event}_i(t_i))), & \text{if } i \neq j_1 \neq \dots \neq j_x, \end{cases} \quad (2)$$

which, for $i \neq j_1 \neq \dots \neq j_x$, first takes the predecessor at node j_1 of the t_i th event occurring at node i , then the predecessor at node j_2 of that predecessor, and so on through node j_x , whose local time after the occurrence of the last predecessor in the chain is assigned to $M_{j_1, \dots, j_x}^i(t_i)$. Should any of these predecessors be undefined, the ones that follow it in the remaining nodes are undefined as well, and the local time that results at the end is zero. It is straightforward to see that, for $x = 1$, this definition is equivalent to the definition of a vector clock in (1). Similarly, the maintenance of matrix clocks follows rules entirely analogous to those used to maintain vector clocks [14].

While the j th component of the vector clock following the occurrence of event v' at node $i \neq j$ gives the time resulting at node j from the occurrence of $\text{pred}_j(v')$, the analogous interpretation that exists for matrix clocks requires the introduction of additional notation. Specifically, the definition of a set of events encompassing the possible x -fold compositions of the relation pred_j with itself, denoted by $\text{Pred}_j^{(x)}$, is necessary. If an event v occurs at node j , then we say that $v \in \text{Pred}_j^{(x)}(v')$ for an event v' occurring at node i if one of the following holds:

- $x = 1$, $j \neq i$, and $v = \text{pred}_j(v')$.
- $x > 1$ and there exists $k \neq i$ such that an event \bar{v} occurs at node k for which $\bar{v} = \text{pred}_k(v')$ and $v \in \text{Pred}_j^{(x-1)}(\bar{v})$.

Note that this definition requires $j \neq k$, in addition to $k \neq i$, for $v \in \text{Pred}_j^{(x)}(v')$ to hold when $x = 2$.

For $x > 1$ and $i \neq j_1 \neq \dots \neq j_x = j$, this definition allows the following interpretation of entry j_1, \dots, j_x of the matrix clock that follows the occurrence of event v' at node i , that is, of $M_{j_1, \dots, j_{x-1}, j}^i(\text{time}_i(v'))$. It gives the time resulting at node j from the occurrence of an event that is in the set $\text{Pred}_j^{(x)}(v')$, so long as it is nonempty. Of course, the number of possibilities for such an event is $O(n^{x-1})$ in the worst case, owing to the several possible combinations of j_1, \dots, j_{x-1} .

Interestingly, applications that require the full capabilities of matrix clocks have yet to be identified. In fact, it seems a simple matter to argue that a slightly more sophisticated use of vector clocks or simply the use of two-dimensional matrix clocks (the $x = 2$ case) suffices to tackle some of the problems that have been offered as possible applications of higher-dimensional matrix clocks [14, 18]. What we do in this paper is to demonstrate how the distributed monitoring of certain resource-sharing computations can benefit from the use of matrix clocks and that, at least for such computations, it is possible to employ much less complex matrix clocks (that is, matrix clocks with many fewer components), which nonetheless retain the ability to reach into the computation's past with arbitrary depth.

The key to this reduction in complexity is the use of one single event in place of each set $\text{Pred}_j^{(y)}(v')$, for $1 \leq y \leq x$. In other words, for each node j the matrix clock we introduce retains only one of the $O(n^{y-1})$ components of each of the x y -dimensional original matrix clocks—one component from the vector clock, one from the two-dimensional matrix clock, and so on through one component of the x -dimensional matrix clock. As we argue in Section 3, following a brief discussion of the resource-sharing context in Section 2, this simplification leads to matrix clocks of size nx , therefore considerably less complex than the $O(n^x)$ -sized original matrix clocks.

The single other attempt at reducing the complexity of a matrix clock that we know of was given for the $x = 2$ case specifically and culminated with the introduction of two techniques [20]. The first one requires attachments of size $O(n)$ (a considerable reduction from the original $O(n^2)$), but is only applicable if the full asynchronism we have been assuming does not hold; it is therefore of little interest in our context. The other technique is somewhat closer to our own in spirit, since it aims at approximating the two-dimensional matrix clock by retaining only k of the $O(n)$ components that correspond to each node. However, the criterion to select the components to be

retained is to choose the k greatest components, which seems unrelated to our own criterion, based as it is on the $Pred_j$ sets.

We give concluding remarks in Section 4, after a discussion of how the technique of Section 3 can successfully solve the problem posed in Section 2.

2. Monitoring resource-sharing computations

The resource-sharing computation we consider is one of the classical solutions to the paradigmatic Dining Philosophers Problem (DPP) [8] in generalized form [6]. In this case, G 's edge set is constructed from a given set of resources and from subsets of that set, one for each node, indicating which resources can be ever needed by that node. This construction places an edge between nodes i and j if the sets of resources ever to be needed by i and j have a nonempty intersection. Notice that this construction is consonant with the interpretation of edges as bidirectional communication channels, because it deploys edges between every pair of nodes that may ever compete for a same resource and must therefore be able to communicate with each other to resolve conflicts.

In DPP, the computation carried out by a node makes it cycle endlessly through three states, which are identified with the conditions of being idle, being in the process of acquiring exclusive access to the resources it needs, and using those resources for a finite period of time. While in the idle state, the node starts acquiring exclusive access to resources when the need arises to compute on shared resources. It is a requirement of DPP that the node must acquire exclusive access to all the resources it shares with all its neighbors, so it suffices for the node to acquire a token object it shares with each of its neighbors (the “fork,” as it is called), each object representing all the resources it shares with that particular neighbor. When in possession of all forks, the node may start using the shared resources [2, 6].

The process of collecting forks from neighbors follows a protocol based on the sending of request messages by the node that needs the forks and the sending of the forks themselves by the nodes that have them. More than one protocol exists, each implementing a different rule to ensure the absence of deadlocks and lockouts during the computation. The solution we consider in this section is based on relative priorities assigned to nodes. Another prominent solution is also based on the assignment of relative priorities, but to the resources instead of to the nodes [16, 23].

The priority scheme of interest to us is based on the graph-theoretic concept of an acyclic orientation of G , which is an assignment of directions to the edges of G in such a way that directed cycles are not formed. Such an acyclic orientation is then a partial order of the nodes of G , and is as such suitable to indicate, given a pair of neighbors, which of the two has priority over the other. Most of the details of this priority scheme are not relevant to our present discussion,

but before stating its essentials we do mention that the lockout-freedom requirement leads the acyclic orientation to be changed dynamically (so that relative priorities are never fixed), which in turn leads to a rich dynamics in the set of all the acyclic orientations of G and to important concurrency-related issues [3–5].

Once a priority scheme is available over the set of nodes, what matters to us is how it is used in the fork-collecting protocol. When a request for fork arrives at node j from node i , j sends i the fork they share if j is either idle or is also collecting forks but does not have priority over i . If j is also collecting forks and has priority over i , or if j is using shared resources, then the sending of the fork to i is postponed to until j has finished using the shared resources. Note that two types of wait may happen here. If j is using shared resources when the request arrives, then the wait is independent of n . If j is also collecting forks, then the wait for j to start using the shared resources and ultimately to send i the fork is in the worst case $n - 1$ [2, 4, 6]. The reason for this is simple: j is waiting for a fork from another node, which may in turn be waiting for a fork from yet another node, and so on. Because the priority scheme is built on acyclic orientations of G , such a chain of waits does necessarily end and is $n - 1$ nodes long in the worst case.

Whether such long waits occur or not is of course dependent upon the details of each execution of the resource-sharing computation. But if they do occur, one possibility for reducing the average wait is to increase the availability of certain critical resources so that G becomes less dense [2, 4, 5]. Perhaps another possibility would be to fine-tune some of the characteristics of each node's participation in the overall computation, such as the duration of its idle period, which could be subject to a mandatory lower bound, for example. In any event, the ability to locally detect long waits (a global property, since it relates to the notion of time in fully asynchronous distributed systems [2]) and identify the nodes at which the wait chains end is crucial.

To see how this relates to the formalism of Section 1, suppose our resource-sharing computation consists of the exchange of fork-bearing messages only (that is, request messages and all other messages involved in the computation, such as those used to update acyclic orientations, are ignored). For distinct nodes i and j , and for an event v' occurring at node i at the reception of a fork, the set $Pred_j^{(x)}(v')$ for $x \geq 1$ is either empty or only contains events that correspond to the sending of forks by node j . Now suppose we systematically investigate the sets $Pred_j^{(x)}(v')$ for every appropriate j by letting x increase from 1 through $n - 1$. Suppose also that, for each j , we record the first x that is found such that $Pred_j^{(x)}(v')$ contains the sending of a fork as response to the reception of a request message without any wait for fork collection on the part of j . The greatest such value to be recorded, say x^* , has special significance: it means that the eventual reception of a fork by node i through v' is the result of an x^* -long chain of waits.

If this were the only information of interest, then Lamport's clocks [15] could be used trivially to discover it. However, taking corrective measures may require a wider array of long wait chains to be found (not simply the longest), as well as the nodes at which those chains end. The matrix clock that we introduce in Section 3 is capable of conveying this information to node i succinctly, so long as there exists enough redundancy in the sets $Pred_j^{(x)}(v')$ that attaching only nX integers to forks suffices, where X is a threshold in the interval $[1, n - 1]$ beyond which wait chains are known not to occur, given the structure of G and the initial arrangement of priorities.¹ It so happens that such redundancy clearly exists: the only events in $Pred_j^{(x)}(v')$ that matter are those corresponding to the sending of forks without any wait for fork collection. Detecting any one of them suffices, so we may as well settle for the latest, that is, one single event from the whole set. We return to this in Section 4.

3. A simpler matrix clock

For $x \geq 1$, the new matrix clock we introduce is an $x \times n$ matrix. For node i at time t_i (i.e., following the occurrence of t_i events at node i), it is denoted by $C^i(t_i)$. For $1 \leq y \leq x$ and $1 \leq j \leq n$, component $C_{y,j}^i(t_i)$ of $C^i(t_i)$ is defined as follows. If $t_i = 0$, then $C_{y,j}^i(t_i) = 0$, as for vector clocks and the matrix clocks of Section 1. If $t_i \geq 1$, then we have

$$C_{y,j}^i(t_i) = \begin{cases} t_i, & \text{if } y = 1 \text{ and } j = i; \\ \max_{i \neq j_1 \neq \dots \neq j_y = j} \left\{ time_{j_y} \left(pred_{j_y} \dots pred_{j_1} (event_i(t_i)) \right) \right\}, & \text{if } y > 1 \text{ or } j \neq i. \end{cases} \quad (3)$$

Note, first of all, that for $y = 1$ this definition is equivalent to the definition of a vector clock in (1). Thus, the first row of $C^i(t_i)$ is the vector clock $V^i(t_i)$. For $y > 1$, the definition in (3) implies, according to the interpretation of matrix clocks that follows our definition in (2), that

$$C_{y,j}^i(t_i) = \begin{cases} \max \{ time_j(v) \mid v \in Pred_j^{(y)}(v') \}, & \text{if } Pred_j^{(y)}(v') \neq \emptyset; \\ 0, & \text{if } Pred_j^{(y)}(v') = \emptyset, \end{cases} \quad (4)$$

where v' is the t_i th event occurring at node i . What this means is that, of all the $O(n^{y-1})$ events that may exist in $Pred_j^{(y)}(v')$, only one (the one to have occurred latest at node j) makes it to the matrix clock $C^i(t_i)$. Note also that (3) implies the equality in (4) for $y = 1$ as well, so long as $j \neq i$. In this case, $Pred_j^{(y)}(v')$, if nonempty, is the singleton $\{pred_j(v')\}$.

¹ Discovering the value of X is not necessarily a simple task, but some empirical knowledge has already been accumulated for modestly-sized systems [3]; in any case, in the likely event that X cannot be determined with certainty, there is always the possibility of adaptation as the resource-sharing computation is observed on the system at hand.




Figure 1. A computation fragment on six nodes.

Before we derive the update rules for our new matrix clock, let us pause and examine an example. Consider Figure 1, where a computation on six nodes is illustrated by means of the usual message diagram that forms the basis of the relation \prec . Nodes are numbered 1 through 6, and in the figure local time elapses from left to right independently for each node. Filled circles represent events and the arrows connecting two events indicate messages.

Three events are singled out in Figure 1, namely v' , v_1 , and v_2 . They are related to one another in such a way that $Pred_6^{(4)}(v') = \{v_1, v_2\}$, that is, v_1 and v_2 are the “depth-4” predecessors of v' at node 6. Recalling that $time_1(v') = 3$, there are two components of the four-dimensional matrix clock $M^1(3)$ that reflect this relationship among the three events, namely

$$M_{3,4,5,6}^1(3) = time_6(v_1) = 2$$

and

$$M_{2,3,4,6}^1(3) = time_6(v_2) = 3.$$

These follow directly from (2) with $x = 4$. By (3) with $y = 4$, the same diagram of Figure 1 yields

$$\begin{aligned} C_{4,6}^1(3) &= \max\left\{time_6(pred_6 pred_5 pred_4 pred_3(v')), time_6(pred_6 pred_4 pred_3 pred_2(v'))\right\} \\ &= \max\{time_6(v_1), time_6(v_2)\} \\ &= 3, \end{aligned}$$

which is clearly also in consonance with (4).

Let us now look at the update rules. For $y = 1$, these rules are the same as those given for vector clocks in Section 1. For $y > 1$, we know from the definition of $Pred_j^{(y)}$ that $v \in Pred_j^{(y)}(v')$ if and only if there exists $k \neq i$ such that an event \bar{v} occurs at node k for which $\bar{v} = pred_k(v')$ and $v \in Pred_j^{(y-1)}(\bar{v})$. Thus, if we let $t_k = time_k(\bar{v})$ and consider the matrix clocks resulting from the occurrence of \bar{v} and of v' ($C^k(t_k)$ and $C^i(t_i)$, respectively), then it follows from (4) that

$$C_{y,j}^i(t_i) = \begin{cases} \max_{k \in K^i(t_i, y, j)} C_{y-1,j}^k(t_k), & \text{if } K^i(t_i, y, j) \neq \emptyset; \\ 0, & \text{if } K^i(t_i, y, j) = \emptyset, \end{cases} \quad (5)$$

where $K^i(t_i, y, j)$ is the set comprising every appropriate k . Notice, in (5), that $k = j$ can never occur as the maximum is taken if $y = 2$: aside from i , node j is the only node that cannot possibly be a member of $K^i(t_i, 2, j)$, by (3).

According to (5), and recalling once again the special case of $y = 1$, we are then left with the following two rules for the evolution of our new matrix clocks:

- Upon sending a message to one of its neighbors, node i attaches $C^i(t_i)$ to the message, where t_i is assumed to already account for the message that is being sent.
- Upon receiving a message from node k with attached matrix clock C^k , and assuming that t_i already reflects the reception of the message, node i sets $C_{y,j}^i(t_i)$ to

$$\begin{cases} t_i, & \text{if } y = 1 \text{ and } j = i; \\ \max\{C_{y,j}^i(t_i - 1), C_{y,j}^k\}, & \text{if } y = 1 \text{ and } j \neq i; \\ C_{y,j}^i(t_i - 1), & \text{if } y = 2 \text{ and } j = k; \\ \max\{C_{y,j}^i(t_i - 1), C_{y-1,j}^k\}, & \text{if } 1 < y \leq x, \text{ provided } y > 2 \text{ or } j \neq k. \end{cases}$$

According to these rules, every message carries an attachment of nx integers.

4. Discussion and concluding remarks

We are now in a position to return to the problem posed in Section 2, namely the problem of monitoring executions of the solution to DPP that employs a partial order on G 's set of nodes to establish priorities. As we discussed in that section, the overall goal is to allow nodes to detect locally, upon receiving a fork, whether the delivery of that fork is the result of a chain of fork deliveries that started too far back in the past. In the affirmative case, the wait since the fork was requested will have been too long, in terms of the usual notions of time complexity in asynchronous distributed computations.

More specifically, if v' is the event corresponding to the reception of a fork at node i , then the goal is for i to be able to detect the existence of an event in $Pred_j^{(y)}(v')$ that corresponds to the sending of a fork by node j either immediately upon the reception of the request for that fork or,

if node j was using shared resources when the request arrived, immediately upon finishing. Here $1 \leq y \leq X$ and j is any node, provided $y = 1$ and $j = i$ do not occur in conjunction. The value of X is such that $1 \leq X \leq n - 1$, and is chosen as a bound to reflect the maximum possible chain of waits. As a consequence, the sets $Pred_j^{(y)}(v')$ must include fork-related events only.

The new matrix clocks introduced in Section 3 can be used for this detection with only minimal adaptation. The key ingredients are:

- The only messages sent by node i to be tagged with the matrix clock C^i are forks. If the sending of a fork by node i does not depend on further fork collection by i , then every component of C^i other than $C_{1,i}^i$ is reset to zero before it is attached to the fork. Matrix clocks are $X \times n$, and all further handling of them follows the rules given in Section 3.
- Upon receiving a fork with attached matrix clock, and having updated C^i accordingly, node i looks for components of C^i that contain nonzero values. If $C_{y,j}^i$ is one such component for $y > 1$ or $j \neq i$, then a wait chain of length y that ends at node j has been discovered and can be checked against a certain threshold $X' < X$ representing the maximum allowable chain length. If C^i contains zero components at all positions but $(1, i)$, then it is certain that the value of X has to be revised, as clearly a wait chain exists whose length surpasses X . A greater value for X is then needed.

This strategy reflects the general method of Section 3 when applied to events that relate to the flow of forks only. Whenever the request for a fork is received and the fork can be sent without the need for any forks to be received by the node in question, say node j , zeroes get inserted into the matrix clock at all positions but $(1, j)$ and are sent along with the fork. The sending of this fork may unleash the sending of forks by other nodes in a chain of events, and along the way the original value of $C_{1,j}^j$ may never be reset to zero, reflecting the increasing length of the wait chain rooted at node j . The reception of a fork whose matrix clock has such a nonzero component beyond row X' is then an indication that such reception is part of chains whose lengths are considered too long.

It is now worth returning to Figure 1 in order to view its diagram as representing the fork-bearing messages in a fragment of a DPP computation of the type we have been considering. For such, consider Figure 2, where two six-node graphs are depicted. Figure 2(a) shows the graph G corresponding to a certain resource-sharing pattern for the six nodes; it shows an acyclic orientation of the edges of G as well, indicating the initial arrangement of priorities. In a situation of heavy demand for resources by all six nodes, only node 6 can acquire all the (three) forks it needs and proceed. All others must wait along the length-5 wait chain shown in Figure 2(b): node 1 for node 2, 2 for 3, and so on through node 6.




Figure 2. G oriented acyclically (a) and a possible wait chain (b).

Eventually, it must happen that node 6 sends out its three forks. If we assume that the corresponding request messages arrive when node 6 already holds all three forks, then upon sending them out all components of C^6 are sent as zeroes, except for $C^6_{1,6}$, sent as 1, 2, or 3, depending on the destination. For $X = 5$, at the occurrence of v' node 1 updates its matrix clock in such a way that $C^1_{5,6} = 2$, which characterizes the length-5 wait chain ending at node 6.

In summary, we have in this paper introduced a novel notion of matrix clocks. Similarly to the original matrix clocks [21, 24], whose definition appears in (2), our matrix clock has the potential of reflecting causal dependencies in the flow of messages that stretch as far as depth x into the past. Unlike those original matrix clocks, however, ours increases with x as nx , while in the original case the growth is according to an $O(n^x)$ function, that is, exponentially.

We have illustrated the applicability of the new matrix clocks with an example from the area of resource-sharing problems. What we have demonstrated is a means of collecting information locally during the resource-sharing computation so that exceedingly long global waits can be detected, possibly indicating the need for overall system re-structuring, as described in Section 2.

Acknowledgments

The authors acknowledge partial support from CNPq, CAPES, the PRONEX initiative of Brazil's MCT under contract 41.96.0857.00, and a FAPERJ BBP grant.

References

1. M. Ahuja, T. Carlson, and A. Gahlot, "Passive-space and time view: vector clocks for achieving higher performance, program correction, and distributed computing," *IEEE Trans. on Software Engineering* **19** (1993), 845–855.

2. V. C. Barbosa, *An Introduction to Distributed Algorithms*, The MIT Press, Cambridge, MA, 1996.
3. V. C. Barbosa, *An Atlas of Edge-Reversal Dynamics*, Chapman & Hall/CRC, London, UK, 2000.
4. V. C. Barbosa, “The combinatorics of resource sharing,” in R. Corrêa *et alii* (Eds.), *Models for Parallel and Distributed Computation: Theory, Algorithmic Techniques and Applications*, 27–52, Kluwer Academic Publishers, Dordrecht, The Netherlands, 2002.
5. V. C. Barbosa and E. Gafni, “Concurrency in heavily loaded neighborhood-constrained systems,” *ACM Trans. on Programming Languages and Systems* **11** (1989), 562–584.
6. K. M. Chandy and J. Misra, “The drinking philosophers problem,” *ACM Trans. on Programming Languages and Systems* **6** (1984), 632–646.
7. B. Charron-Bost, “Concerning the size of logical clocks in distributed systems,” *Information Processing Letters* **39** (1991), 11–16.
8. E. W. Dijkstra, “Hierarchical ordering of sequential processes,” *Acta Informatica* **1** (1971), 115–138.
9. R. P. Dilworth, “A decomposition theorem for partially ordered sets,” *Annals of Mathematics* **51** (1950), 161–165.
10. L. M. A. Drummond and V. C. Barbosa, “Distributed breakpoint detection in message-passing programs,” *J. of Parallel and Distributed Computing* **39** (1996), 153–167.
11. C. J. Fidge, “Timestamps in message-passing systems that preserve partial ordering,” *Proc. of the 11th Australian Computer Science Conference*, 56–66, 1988.
12. C. J. Fidge, “Logical time in distributed computing systems,” *IEEE Computer* **24** (1991), 28–33.
13. C. J. Fidge, “Fundamentals of distributed system observation,” *IEEE Software* **13** (1996), 77–83.
14. V. K. Garg, *Principles of Distributed Systems*, Kluwer Academic Publishers, Boston, MA, 1996.
15. L. Lamport, “Time, clocks, and the ordering of events in a distributed system,” *Comm. of the ACM* **21** (1978), 558–565.

16. N. A. Lynch, "Upper bounds for static resource allocation in a distributed system," *J. of Computer and System Sciences* **23** (1981), 254–278.
17. F. Mattern, "Virtual time and global states in distributed systems," in M. Cosnard *et alii* (Eds.), *Parallel and Distributed Algorithms: Proc. of the Int. Workshop on Parallel and Distributed Algorithms*, 215–226, North-Holland, Amsterdam, The Netherlands, 1989.
18. M. Raynal, "Illustrating the use of vector clocks in property detection: an example and a counter-example," in P. Amestoy *et alii* (Eds.), *Euro-Par'99—Parallel Processing*, 806–814, Lecture Notes in Computer Science 1685, Springer-Verlag, Berlin, Germany, 1999.
19. L. E. T. Rodrigues and P. Veríssimo, "Causal separators for large-scale multicast communication," *Proc. of the 15th Int. Conf. on Distributed Computing Systems*, 83–91, 1995.
20. F. Ruget, "Cheaper matrix clocks," in G. Tel and P. Vitányi (Eds.), *Distributed Algorithms: Proc. of the 8th Int. Workshop on Distributed Algorithms*, 355–369, Lecture Notes in Computer Science 857, Springer-Verlag, Berlin, Germany, 1994.
21. S. K. Sarin and L. Lynch, "Discarding obsolete information in a replicated database system," *IEEE Trans. on Software Engineering* **SE-13** (1987), 39–46.
22. M. Singhal and A. Kshemkalyani. "An efficient implementation of vector clocks," *Information Processing Letters* **43** (1992), 47–52.
23. J. L. Welch and N. A. Lynch, "A modular drinking philosophers algorithm," *Distributed Computing* **6** (1993), 233–244.
24. G. T. J. Wu and A. J. Bernstein, "Efficient solutions to the replicated log and dictionary problems," *Proc. of the 3rd Annual ACM Symposium on Principles of Distributed Computing*, 233–242, 1984.