

Keeping Track of the Latest Gossip in a Distributed System*

Madhavan Mukund

SPIC Mathematical Institute

92 G.N. Chetty Road, Madras 600 017, India

E-mail: madhavan@smi.ernet.in

Milind Sohoni

Department of Computer Science and Engineering

Indian Institute of Technology, Bombay 400 076, India

E-mail: sohoni@cse.iitb.ernet.in

Abstract

We tackle a natural problem from distributed computing, involving time-stamps. Let $\mathcal{P} = \{p_1, p_2, \dots, p_N\}$ be a set of computing agents or processes which synchronize with each other from time to time and exchange information about themselves and others. The *gossip problem* is the following: Whenever a set $P \subseteq \mathcal{P}$ meets, the processes in P must decide *amongst themselves* which of them has the latest information, direct or indirect, about each agent p in the system.

We propose an algorithm to solve this problem which is finite-state and local. Formally, this means that our algorithm can be implemented as an asynchronous automaton.

Keywords: Distributed algorithms, synchronous communication, bounded time-stamps, asynchronous automata

*This paper appears in *Distributed Computing*, 10, 3, (1997) 117-127. A preliminary version of this paper appeared in [22]. The first author was partially supported by IFCPAR Project 502-1. This work was done while the second author was with the School of Mathematics, SPIC Science Foundation, now known as the SPIC Mathematical Institute.

Introduction

The aim of this paper is to tackle a natural problem from distributed computing, involving time-stamps. Let $\mathcal{P} = \{p_1, p_2, \dots, p_N\}$ be a set of computing agents or processes which synchronize with each other from time to time and exchange information about themselves and others. The *gossip problem* is the following: Whenever a set $P \subseteq \mathcal{P}$ meets, the processes in P must decide *amongst themselves* which of them has the latest information, direct or indirect, about each agent p in the system.

This is easily accomplished if the agents decide to “time-stamp” every synchronization and pass these time-stamps along with each exchange of information. This does not require that all their clocks be synchronized. For example, each process can use an independent counter. When a set $P \subseteq \mathcal{P}$ meets, the processes in P jointly agree on a new value for their counters which exceeds the maximum of the counter values currently held by them. Thus, for any process p , the time-stamps assigned to synchronization events involving p form a strictly increasing sequence (albeit with gaps between successive time-stamps). So, the problem of deciding who has the latest information about p reduces to that of checking for the largest time-stamp.

This scheme has the following drawback: As the computation progresses these counter values increase without bound and most of the agents’ time would be taken up in passing on large numbers, as opposed to actual gossip.

We propose an algorithm using counters which take on values from a bounded, finite set. We assign an independent counter to each subset of processes which can potentially synchronize. These counters are updated when the corresponding sets of processes meet. The update is performed jointly by the processes which meet.

Since our set of counter values is bounded, time-stamps have to be reused and, in general, different synchronizations involving a particular set of processes will acquire the same time-stamp during a computation. Despite this, our algorithm guarantees that whenever a set $P \subseteq \mathcal{P}$ meets, the processes in P can decide correctly which of them has the best information about any other agent p in the system. Thus, in essence, the processes in \mathcal{P} may be finite state machines and yet manage to keep track of the latest information about other agents. Further, the algorithm itself does not induce any additional communications.

We formalize the gossip problem and our solution to it in terms of asynchronous automata. These machines were first introduced by Zielonka and are a natural generalization of finite-state automata for modelling concurrent systems [30]. An asynchronous automaton consists of a set of finite-state agents which synchronize to process their input. Each letter a in the input alphabet Σ is assigned a subset $\theta(a)$ of processes which jointly update their state when reading a . The processes outside $\theta(a)$ remain unchanged during this move—in fact, they are oblivious to the occurrence of a .

(Calling these automata *asynchronous* is misleading. The automata communicate *synchronously*. Zielonka used the term “asynchronous” to emphasize that different components of the network can proceed independently while processing the input. We continue to use this rather inappropriate terminology for historical reasons.)

A preview of the algorithm

Our algorithm proceeds by having each process maintain different levels of information about the rest of the system. At the *primary* level, a process p records the latest information that it has heard from every other process. Thus, the primary information of process p regarding process q corresponds to the most recent event performed by q which p is aware of, either directly or indirectly.

The *secondary* information of a process consists of its knowledge about the primary information of other processes. Thus, for processes p , q and r , p 's secondary information would refer to events of the form “the latest that p knows about what q knows about r ”. Similarly, the *tertiary information* of a process consists of its knowledge about the secondary information of every other process.

Events are identified by labels assigned to them when they occur—the label of each event is assigned jointly by all the processes that take part in the event. (An event involving only one process corresponds to an internal event.) These labels are best thought of as time-stamps. The goal of the algorithm is to correctly compare and update the primary information of all processes which participate in each synchronization, with the constraint that only a bounded number of labels are used to time-stamp events, regardless of the length of the overall computation.

Let p and q be processes that synchronize. We prove that p 's primary information about r is more recent (and therefore better) than q 's primary information about r if and only if the event corresponding to r recorded in q 's primary information is also present in p 's secondary information. In other words, the primary information of p and q can be compared by just checking for *equality* of labels within their primary and secondary information—we do not have to maintain any order between the labels. This feature is crucial for designing an algorithm which uses only a bounded number of time-stamps: reusing a label will, in general, destroy any a priori order on the set of labels.

It then follows that we can reuse time-stamps provided we maintain the following invariant across the system: At any stage of the computation, if the same label is present in the primary or secondary information of two different processes, then the labels actually point to the same event. In other words, two *different* instances of the same time-stamp should never simultaneously be present in the primary and secondary information of the system. This invariant is difficult to maintain because the processes which take part in an event and assign a time-stamp to it will not, in general, have access to the primary and secondary information of all processes across the system. So, they have no way of knowing precisely which labels are “in use” across the system when the event occurs.

It turns out that tertiary information can be used to resolve the problem of deciding when a time-stamp can be reused. We prove that a time-stamp previously assigned by a process p is currently “in use”—i.e., it currently belongs to the primary or secondary information of some other process q —only if it also belongs to the tertiary information of p . From this, it follows that labels which are not in the tertiary information of a process can be reused. Since the number of events in the tertiary information of each process is bounded, processes in the system can always work with a bounded set of labels large enough to permit each event to be assigned a fresh time-stamp which does not clash with the set of time-stamps currently “in use” across the system.

As we remarked earlier, our algorithm can be described as an asynchronous automaton, where each process is locally a finite-state machine. The automaton that we construct to

solve the gossip problem can be effectively presented in space polynomial in the number of processes in the system. This means that the automaton can be embedded in other distributed algorithms without sacrificing efficiency.

In this paper, we solve the gossip problem for systems with multi-party synchronization. Surprisingly, the problem appears *no easier* when restricted to systems with pairwise synchronization. Though systems with pairwise synchronization are perhaps more representative of “real world” systems, we have chosen to describe our solution in the general setting because it has important applications in theoretical studies of distributed systems based on the asynchronous automaton model [13, 14, 23].

The technique we describe for analyzing synchronizing systems is also applicable to more general classes of distributed systems. For instance, we can adapt our algorithm for maintaining primary, secondary and tertiary information to solve the gossip problem for “well behaved” classes of message-passing systems [21].

The paper is organized as follows. In the next section, we introduce asynchronous automata and formalize the gossip problem in terms of these automata. To do this, we define a natural partial order on events in the system. In Section 2 we introduce ideals and frontiers, both of which play a crucial role in the rest of the paper. Sections 3 and 4 describe how to maintain, compare and update in a local manner the latest information about other processes. The next section puts all these ideas together and formally describes the “gossip automaton” which solves the problem we set out to tackle.

Section 6 briefly examines extensions of the basic gossip automaton and possibilities for optimizing the construction. We also look at applications of the gossip automaton in logic and the theory of asynchronous automata [13, 14, 23, 27].

In the concluding Discussion, we place our results in perspective. We discuss similarities and differences with other work on “gossiping” and bounded time-stamps [1, 2, 3, 5, 8, 9, 12].

1 Preliminaries

Let \mathcal{P} be a finite set of processes which synchronize periodically and let the set of possible synchronizations permitted in the system be denoted \mathcal{C} , where $\mathcal{C} \subseteq (2^{\mathcal{P}} - \{\emptyset\})$. So, each element $c \in \mathcal{C}$ is a non-empty subset of \mathcal{P} . When c occurs, the processes in c share all information about their local states and update their states in synchrony.

We model a computation of the system as a sequence of communications—that is, a word $u \in \mathcal{C}^*$. Let u be of length m . It is convenient to think of u as a function $u : [1..m] \rightarrow \mathcal{C}$, where for natural numbers i and j , $[i..j]$ abbreviates the set $\{i, i+1, \dots, j\}$ if $i \leq j$ and $[i..j] = \emptyset$ otherwise. By this convention, the empty word ε is denoted by the unique function $\emptyset \rightarrow \mathcal{C}$.

Events With $u : [1..m] \rightarrow \mathcal{C}$, we associate a set of *events* \mathcal{E}_u . Each event e is of the form $(i, u(i))$, where $i \in [1..m]$. In addition, it is convenient to include an *initial event* denoted 0. Thus, $\mathcal{E}_u = \{0\} \cup \{(i, u(i)) \mid i \in [1..m]\}$.

The initial event marks an implicit synchronization of all the processes before the start of the actual computation. So, if u is the empty word ε , $\mathcal{E}_u = \{0\}$.

Usually, we will write \mathcal{E} for \mathcal{E}_u . For $p \in \mathcal{P}$ and $e \in \mathcal{E}$, we write $p \in e$ to denote that $p \in u(i)$ when $e = (i, u(i))$. Thus, $p \in e$ implies that process p participated in the

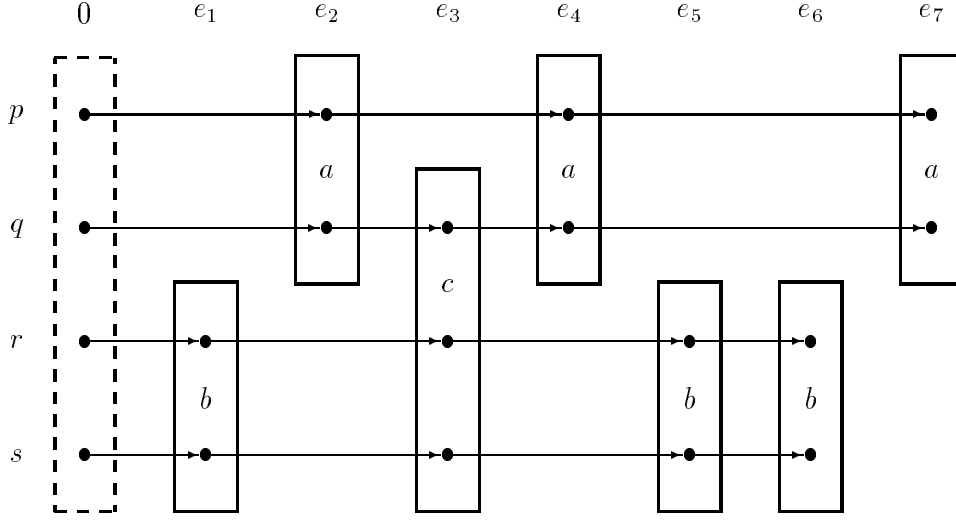


Figure 1: A typical computation of a system of synchronizing processes

synchronization e . For the initial event 0, we define $p \in 0$ to hold for all $p \in \mathcal{P}$. If $p \in e$, then we say that e is a *p-event*.

Ordering relations on \mathcal{E} The word u imposes a total order on events in \mathcal{E} : define $e < f$ if $e \neq f$ and either $e = 0$ or $e = (i, u(i))$, $f = (j, u(j))$, and $i < j$. We write $e \leq f$ if $e = f$ or $e < f$.

However, the temporal order $<$ does not accurately reflect the cause and effect relationship between events in \mathcal{E} . Clearly, synchronizations between disjoint sets of processes can be performed independently. In particular, if two such synchronizations occur consecutively in u , they could also be transposed without affecting the outcome of the computation. To record information about causality and independence, we define a partial order \sqsubseteq^* on \mathcal{E} .

To begin with, we observe that each process p orders the events in which it participates: define \triangleleft_p to be the relation

$$e \triangleleft_p f \triangleq e < f, p \in e \cap f \text{ and for all } e < g < f, p \notin g.$$

The set of all p -events in \mathcal{E} is totally ordered by \triangleleft_p^* , the reflexive, transitive closure of \triangleleft_p .

Define $e \sqsubset f$ if for some p , $e \triangleleft_p f$ and $e \sqsubseteq f$ if $e = f$ or $e \sqsubset f$. Let \sqsubseteq^* denote the transitive closure of \sqsubseteq . If $e \sqsubseteq^* f$ then we say that e is *below* f .

It is not difficult to see that the *causality relation* \sqsubseteq^* accurately models the cause and effect relationship between events in \mathcal{E} . In particular, all rearrangements of the letters in u which arise out of permuting adjacent independent synchronizations will give rise to isomorphic structures $(\mathcal{E}, \sqsubseteq^*)$.

EXAMPLE: Let $\mathcal{P} = \{p, q, r, s\}$ and $\mathcal{C} = \{a, b, c\}$ where $a = \{p, q\}$, $b = \{r, s\}$ and $c = \{q, r, s\}$. Figure 1 shows the events \mathcal{E} corresponding to the word *bacabba*. The dashed box corresponds to the “mythical” event 0, which we insert at the beginning for convenience.

In the figure, the arrows between the events denote the relations \triangleleft_p , \triangleleft_q , \triangleleft_r and \triangleleft_s . From these, we can compute \sqsubseteq and \sqsubseteq^* . Thus, for example, we have $e_1 \sqsubseteq^* e_4$ since $e_1 \triangleleft_r e_3 \triangleleft_q e_4$.

Note that 0 is below every event. Also, for each $p \in \mathcal{P}$, the set of all p -events in \mathcal{E} is totally ordered by \sqsubseteq^* since \triangleleft_p^* is contained in \sqsubseteq^* .

The set of events below e is denoted $e\downarrow$. These represent the only synchronizations in \mathcal{E} which are “known” to the processes in e when e occurs.

Latest information Let \mathcal{E} be the set of events of the communication sequence $u : [1..m] \rightarrow \mathcal{C}$. The \sqsubseteq^* -maximum p -event in \mathcal{E} is denoted $\max_p(\mathcal{E})$. $\max_p(\mathcal{E})$ is the last event in \mathcal{E} in which p has taken part. Since $p \in 0 \in \mathcal{E}$ and all p -events are totally ordered by \sqsubseteq^* , $\max_p(\mathcal{E})$ is well-defined.

Let $p, q \in \mathcal{P}$. The latest information p has about q in \mathcal{E} corresponds to the \sqsubseteq^* -maximum q -event in the subset of events $\max_p(\mathcal{E})\downarrow$. We denote this event by $\text{latest}_{p \rightarrow q}(\mathcal{E})$. Since all q -events are totally ordered by \sqsubseteq^* and $q \in 0 \sqsubseteq^* \max_p(\mathcal{E})$, $\text{latest}_{p \rightarrow q}(\mathcal{E})$ is well-defined.

EXAMPLE: Continuing with our example, in Figure 1, $\max_p(\mathcal{E}) = e_7$ whereas $\max_s(\mathcal{E}) = e_6$. $\text{latest}_{p \rightarrow q}(\mathcal{E}) = e_7$, but $\text{latest}_{p \rightarrow s}(\mathcal{E}) = e_3$. On the other hand, $\text{latest}_{s \rightarrow p}(\mathcal{E}) = e_2$.

For any processes $p, p', q \in \mathcal{P}$, the events $\text{latest}_{p \rightarrow q}(\mathcal{E})$ and $\text{latest}_{p' \rightarrow q}(\mathcal{E})$ are both q -events and are thus always comparable with respect to \sqsubseteq^* . Our goal is to design a scheme whereby each process p maintains a bounded amount of information locally, so that whenever a set $c \subseteq \mathcal{P}$ synchronizes, the processes in c can decide amongst themselves which of them has heard most recently from every process in the system. More formally, for every $q \in \mathcal{P}$, all the processes in c should be able to jointly compute which of the events $\{\text{latest}_{p \rightarrow q}(\mathcal{E})\}_{p \in c}$ is maximum with respect to \sqsubseteq^* .

We make precise the notions of *bounded* and *local* information using asynchronous automata.

Asynchronous automata

Distributed alphabet Let \mathcal{P} be a finite set of processes as before. A *distributed alphabet* is a pair (Σ, θ) where Σ is a finite set of *actions* and $\theta : \Sigma \rightarrow (2^{\mathcal{P}} - \{\emptyset\})$ assigns a non-empty set of processes to each $a \in \Sigma$.

State spaces With each process p , we associate a finite set of states denoted S_p . Each state in S_p is called a *local state*. For $P \subseteq \mathcal{P}$, we use S_P to denote the product $\prod_{p \in P} S_p$. An element \vec{s} of S_P is called a *P-state*. A \mathcal{P} -state is also called a *global state*. Given $\vec{s} \in S_P$, and $P' \subseteq P$, we use $\vec{s}_{P'}$ to denote the projection of \vec{s} onto $S_{P'}$.

Asynchronous automaton An *asynchronous automaton* \mathcal{A} over (Σ, θ) is of the form $(\{S_p\}_{p \in \mathcal{P}}, \{\rightarrow_a\}_{a \in \Sigma}, \mathcal{S}_0, \mathcal{S}_F)$, where $\rightarrow_a \subseteq S_{\theta(a)} \times S_{\theta(a)}$ is the *local transition relation* for a , and $\mathcal{S}_0, \mathcal{S}_F \subseteq S_{\mathcal{P}}$ are sets of *initial* and *final* global states. Intuitively, each local transition relation \rightarrow_a specifies how the processes $\theta(a)$ that meet on a may decide on a joint move. Processes outside $\theta(a)$ do not change their state when a occurs. Thus we define the *global transition relation* $\Rightarrow \subseteq S_{\mathcal{P}} \times \Sigma \times S_{\mathcal{P}}$ by $\vec{s} \xRightarrow{a} \vec{s}'$ if $\vec{s}_{\theta(a)} \rightarrow_a \vec{s}'_{\theta(a)}$ and $\vec{s}_{\mathcal{P}-\theta(a)} = \vec{s}'_{\mathcal{P}-\theta(a)}$.

\mathcal{A} is called *deterministic* if the global transition relation \Rightarrow of \mathcal{A} is a function from $S_{\mathcal{P}} \times \Sigma$ to $S_{\mathcal{P}}$ and the set of initial states \mathcal{S}_0 is a singleton. Notice that \Rightarrow is a function iff each local transition relation $\rightarrow_a, a \in \Sigma$, is a function from $S_{\theta(a)}$ to $S_{\theta(a)}$. We shall only deal with deterministic asynchronous automata in this paper.

Runs Given a word $u : [1..m] \rightarrow \Sigma$, a *run* of \mathcal{A} on u is a function $\rho : [0..m] \rightarrow S_{\mathcal{P}}$ such that $\rho(0) \in \mathcal{S}_0$ and for $i \in [1..m]$, $\rho(i-1) \xrightarrow{u(i)} \rho(i)$. If \mathcal{A} is deterministic, each word u gives rise to a unique run which we denote ρ_u .

The word u is *accepted* by \mathcal{A} if there is a run ρ of \mathcal{A} on u such that $\rho(m) \in \mathcal{S}_F$. $L(\mathcal{A})$, the language recognized by \mathcal{A} , is the set of words accepted by \mathcal{A} . In this paper, we will not look at asynchronous automata as language recognizers. Instead, we shall treat them as devices for locally computing families of functions.

Locally computable functions Let Val be a set (of *values*). A Σ -indexed family of functions is a set $\mathcal{F}_{\Sigma} = \{f_a : \Sigma^* \rightarrow Val\}_{a \in \Sigma}$. So, \mathcal{F}_{Σ} contains a function f_a for each letter $a \in \Sigma$.

\mathcal{F}_{Σ} is *locally computable* if we can find a deterministic asynchronous automaton $\mathcal{A} = (\{S_p\}_{p \in \mathcal{P}}, \{\rightarrow_a\}_{a \in \Sigma}, \mathcal{S}_0, \mathcal{S}_F)$ and a family of *local functions* $\mathcal{G}_{\Sigma} = \{g_a : S_{\theta(a)} \rightarrow Val\}_{a \in \Sigma}$, such that for each word $u : [1..m] \rightarrow \Sigma$, $f_a(u) = g_a(\vec{s}_{\theta(a)})$, where $\vec{s} = \rho_u(m)$ and ρ_u is the unique run of \mathcal{A} over u .

In other words, the processes in $\theta(a)$ can locally compute the value $f_a(u)$ for any $u \in \Sigma^*$ by applying the function g_a to the (unique) $\theta(a)$ -state reached by the automaton after reading u .

Our problem involving the latest information of processes in \mathcal{P} can now be formalized in terms of asynchronous automata.

Given $\mathcal{C} \subseteq (2^{\mathcal{P}} - \{\emptyset\})$, let $\Sigma = \{\hat{c}\}_{c \in \mathcal{C}}$. The distribution function θ is defined in the obvious way—for each $\hat{c} \in \Sigma$, $\theta(\hat{c}) = c$. For convenience, henceforth we shall drop the distinction between a subset $c \in \mathcal{C}$ and the corresponding letter $\hat{c} \in \Sigma$ and refer to both as just c . Thus, we will use S_c to denote the set of $\theta(\hat{c})$ -states and \rightarrow_c to denote the local transition function for \hat{c} .

Let $u : [1..m] \rightarrow \Sigma$ be a communication sequence and $c \subseteq \mathcal{P}$. For each $q \in \mathcal{P}$, we denote by $best_c(u, q)$ the set of processes in c which have the most recent information about q at the end of u —i.e.,

$$best_c(u, q) = \{p \in c \mid \forall p' \in c. latest_{p' \rightarrow q}(\mathcal{E}_u) \sqsubseteq^* latest_{p \rightarrow q}(\mathcal{E}_u)\}.$$

Let $Val = (2^{\mathcal{P}} - \{\emptyset\})^{\mathcal{P}}$. So, each member of Val is a function from \mathcal{P} to non-empty subsets of \mathcal{P} . Our goal is to show that the family of functions $\{latest-gossip_c : \Sigma^* \rightarrow Val\}_{c \in \Sigma}$ is locally computable, where:

$$\forall u \in \Sigma^*. \forall c \in \Sigma. latest-gossip_c(u) \text{ is the function } \{p \mapsto best_c(u, p)\}_{p \in \mathcal{P}}.$$

2 Ideals and Frontiers

For the moment, let us fix a communication sequence $u : [1..m] \rightarrow \Sigma$ and the corresponding set of events \mathcal{E} .

The main source of difficulty in solving the gossip problem is the fact that the processes in \mathcal{P} need to compute global information about the communication sequence u while each process only has access to a local, “partial” view of u . Although partial views of u correspond to subsets of \mathcal{E} , not every subset of \mathcal{E} arises from such a partial view. Those subsets of \mathcal{E} which do correspond to partial views of u are called ideals.

Ideals A set of events $I \subseteq \mathcal{E}$ is called an *order ideal* if I is closed with respect to \sqsubseteq^* —i.e., $e \in I$ and $f \sqsubseteq^* e$ implies $f \in I$ as well. We shall always refer to order ideals as just *ideals*.¹

The requirement that an ideal be closed with respect to \sqsubseteq^* guarantees that the observation it represents is “consistent”—whenever an event e has been observed, so have all the events in the computation which necessarily precede e .

The minimum possible partial view of a word u is the ideal $\{0\}$. This is because of our interpretation of 0 as an event which takes place *before* the actual computation begins. Since 0 lies below every event in \mathcal{E} , $0 \in I$ for every non-empty ideal I . We shall assume that every ideal we consider is non-empty.

Clearly the entire set \mathcal{E} is an ideal, as is $e \downarrow$ for any $e \in \mathcal{E}$. It is easy to see that if I and J are ideals, so are $I \cup J$ and $I \cap J$.

EXAMPLE: Let us look once again at Figure 1. $\{0, e_2\}$ is an ideal, but $\{0, e_2, e_3\}$ is not, since $e_1 \sqsubseteq^* e_3$ but $e_1 \notin \{0, e_2, e_3\}$. $\{0, e_1, e_2, e_3, e_5\}$ is the ideal $e_5 \downarrow$, whereas $\{0, e_1, e_2, e_3, e_4, e_5\}$ is an ideal which is not of the form $e \downarrow$ for any $e \in \mathcal{E}$.

We need to generalize the notion of $\max_p(\mathcal{E})$, the maximum p -event in \mathcal{E} , to all ideals $I \subseteq \mathcal{E}$.

P -views For an ideal I , the \sqsubseteq^* -maximum p -event in I is denoted $\max_p(I)$. The p -view of I is the set $I|_p = \max_p(I) \downarrow$. So, $I|_p$ is the set of all events in I which p can “see”. For $P \subseteq \mathcal{P}$, the P -view of I , denoted $I|_P$, is $\bigcup_{p \in P} I|_p$, which is also an ideal. In particular, we have $I|_{\mathcal{P}} = I$.

EXAMPLE: In Figure 1, let I denote the ideal $\{0, e_1, e_2, e_3, e_4, e_5, e_6\}$. $\max_q(I) = e_4$ and hence $I|_q = \{0, e_1, e_2, e_3, e_4\}$. On the other hand, though $\max_r(I) = e_6$, $I|_r \neq I$; $I|_r = I - \{e_4\}$. The joint view $I|_{\{q,r\}} = I = I|_p$.

For an ideal I , the views $I|_p$ and $I|_q$ seen by two processes $p, q \in \mathcal{P}$ are, in general, incomparable. The events in I where these two views begin to diverge—the *frontier* of $I|_p \cap I|_q$ —play a crucial role in our analysis.

Frontiers Let I be an ideal and $p, q, r \in \mathcal{P}$. We say that an event e is an r -sentry for p with respect to q if $e \in I|_p \cap I|_q$ and $e \triangleleft_r f$ for some $f \in I|_q - I|_p$. Thus e is an event known to both p and q whose r -successor is known only to q . Notice that there need not always be an r -sentry for p with respect to q .

The pq -frontier at I , $\text{frontier}_{pq}(I)$ is defined as follows:

$$\text{frontier}_{pq}(I) = \{e \in I \mid \exists r \in \mathcal{P}. e \text{ is an } r\text{-sentry for } p \text{ with respect to } q\}$$

¹In the theory of partial orders, *order ideals* and *ideals* are distinct concepts. Ideals are normally assumed to be subsets which are \sqsubseteq^* -closed and directed. We shall, however, deal only with order ideals in this paper and so our terminology should cause no confusion.

Observe that this definition is asymmetric—in general, $\text{frontier}_{pq}(I) \neq \text{frontier}_{qp}(I)$.

EXAMPLE: As before, in Figure 1, let I denote the ideal $\{0, e_1, e_2, e_3, e_4, e_5, e_6\}$. $I|_q \cap I|_r = \{0, e_1, e_2, e_3\}$. $\text{frontier}_{rq}(I) = \{e_2, e_3\}$ — e_2 is a p -sentry for r with respect to q whereas e_3 is a q -sentry. On the other hand, $\text{frontier}_{qr}(I) = \{e_3\}$. e_3 is both an r -sentry as well as an s -sentry for q with respect to r .

As the example demonstrates, an event $e \in \text{frontier}_{pq}(I)$ could simultaneously be an r -sentry for p for several different processes r . However, it is not difficult to show that for any process r , there is at most one r -sentry for p with respect to q .

3 Primary and secondary information

For a word u and processes $p, q \in \mathcal{P}$, we have already defined $\text{latest}_{p \rightarrow q}(\mathcal{E})$, the latest information that p has about q after u . We now extend this definition to arbitrary ideals.

Primary information Let I be an ideal and $p, q \in \mathcal{P}$. Then $\text{latest}_{p \rightarrow q}(I)$ denotes the \sqsubseteq^* -maximum q -event in $I|_p$. So, $\text{latest}_{p \rightarrow q}(I)$ is the latest q -event in I that p knows about.

The *primary information* of p after I , $\text{primary}_p(I)$, is the set $\{\text{latest}_{p \rightarrow q}(I)\}_{q \in \mathcal{P}}$. More precisely, $\text{primary}_p(I)$ is an *indexed* set of events—each event $e = \text{latest}_{p \rightarrow q}(I)$ in $\text{primary}_p(I)$ is represented as a triple (p, q, e) . As usual, for $P \subseteq \mathcal{P}$, $\text{primary}_P(I) = \bigcup_{p \in P} \text{primary}_p(I)$.

As we have already remarked, for all $q \in \mathcal{P}$, the set of q -events in $I|_p$ is always nonempty, since $q \in 0 \in I|_p$. Further, since all q -events are totally ordered by \prec_q^* and hence by \sqsubseteq^* , the maximum q -event in $I|_p$ is well-defined. Notice that $\text{latest}_{p \rightarrow p}(I) = \max_p(I)$.

To compare primary events, processes need to maintain additional information. It turns out that it is sufficient for each process to keep track of all the other processes' primary information.

Secondary information The *secondary information* of p after I , $\text{secondary}_p(I)$, is the (indexed) set $\bigcup_{q \in \mathcal{P}} \text{primary}_q(\text{latest}_{p \rightarrow q}(I) \downarrow)$. In other words, this is the latest information that p has in I about the primary information of q , for each $q \in \mathcal{P}$. Once again, for $P \subseteq \mathcal{P}$, $\text{secondary}_P(I) = \bigcup_{p \in P} \text{secondary}_p(I)$.

Each event in $\text{secondary}_p(I)$ is of the form $\text{latest}_{q \rightarrow r}(\text{latest}_{p \rightarrow q}(I) \downarrow)$ for some $q, r \in \mathcal{P}$. This is the latest r -event which q knows about upto the event $\text{latest}_{p \rightarrow q}(I)$. We abbreviate $\text{latest}_{q \rightarrow r}(\text{latest}_{p \rightarrow q}(I) \downarrow)$ by $\text{latest}_{p \rightarrow q \rightarrow r}(I)$.

Just as we represented events in $\text{primary}_p(I)$ as triples of the form (p, q, e) , where $p, q \in \mathcal{P}$ and $e \in I$, we represent each secondary event $e = \text{latest}_{p \rightarrow q \rightarrow r}(I)$ in $\text{secondary}_p(I)$ as a quadruple (p, q, r, e) .

However, we will often ignore the fact that $\text{primary}_p(I)$ and $\text{secondary}_p(I)$ are *indexed* sets of events and treat them, for convenience, as just sets of events. Thus, for an event $e \in I$, we shall write $e \in \text{primary}_p(I)$ to mean that there exists a process $q \in \mathcal{P}$ such that $(p, q, e) \in \text{primary}_p(I)$ —i.e., $e = \text{latest}_{p \rightarrow q}(I)$. Similarly, $e \in \text{secondary}_p(I)$ will indicate that for some $q, r \in \mathcal{P}$, $(p, q, r, e) \in \text{secondary}_p(I)$. We extend this to other set-theoretic operations as well. So, for instance, if we say $e \in \text{primary}_p(I) \cap \text{secondary}_q(I)$, we mean that we can find $p', q', q'' \in \mathcal{P}$ such that $(p, p', e) \in \text{primary}_p(I)$ and $(q, q', q'', e) \in \text{secondary}_q(I)$.

Notice that each primary event $latest_{p \rightarrow q}(I)$ is also a secondary event $latest_{p \rightarrow p \rightarrow q}(I)$ (or, equivalently, $latest_{p \rightarrow q \rightarrow q}(I)$). So, following our convention that $primary_p(I)$ and $secondary_p(I)$ be treated as sets of events, we write $primary_p(I) \subseteq secondary_p(I)$.

Comparing primary information

Our goal is to compare and update the primary information of processes whenever they meet. For this, we need the following observation regarding the significance of events lying on frontiers.

Lemma 1 *Let I be an ideal, $p, q \in \mathcal{P}$ and $e \in frontier_{pq}(I)$ an r -sentry for p with respect to q . Then $e = latest_{p \rightarrow r}(I)$. Also, for some $r' \in \mathcal{P}$, $e = latest_{q \rightarrow r' \rightarrow r}(I)$. So, $e \in primary_p(I) \cap secondary_q(I)$.*

Proof Since e is an r -sentry, for some $f \in I|_q - I|_p$, $e \triangleleft_r f$. Suppose that $latest_{p \rightarrow r}(I) = e' \neq e$. Since all r -events are totally ordered by \triangleleft_r^* , we must have $e \triangleleft_r^+ e'$ (where \triangleleft_r^+ is the transitive closure of the irreflexive relation \triangleleft_r). However, $e \triangleleft_r f$ as well, so we have $e \triangleleft_r f \triangleleft_r^* e'$. Since $e' \in I|_p$, this means that $f \in I|_p$ as well, which is a contradiction.

Next, we must show that $e = latest_{q \rightarrow r' \rightarrow r}(I)$ for some $r' \in \mathcal{P}$. We know that there is a path $e \sqsubset f_1 \sqsubset \dots \sqsubset max_p(I)$, since $e \in I|_p$. This path starts inside $I|_p \cap I|_q$.

If this path never leaves $I|_p \cap I|_q$ then $max_p(I) \in I|_q$. Since $max_p(I)$ is the \sqsubseteq^* -maximum p -event in I , it must be the \sqsubseteq^* -maximum p -event in $I|_q$. So, $e = latest_{q \rightarrow p \rightarrow r}(I)$ and we are done.

If this path does leave $I|_p \cap I|_q$, we can find an event e' along the path such that $e \sqsubseteq^* e' \triangleleft_{r'} f' \sqsubseteq^* max_p(I)$, where $e' \in I|_p \cap I|_q$, $f' \in I|_p - I|_q$ and $r' \in e' \cap f'$. In other words, e' is an r' -sentry for q with respect to p . We know by our earlier argument that $e' = latest_{q \rightarrow r'}(I)$. It must be the case that $e = latest_{r' \rightarrow r}(e' \downarrow)$. For, if $latest_{r' \rightarrow r}(e' \downarrow) = e'' \neq e$, then $e \triangleleft_r^+ e'' \sqsubseteq^* e' \sqsubseteq^* max_p(I)$. Since $e'' \in I|_p$ and $e \triangleleft_r^+ e''$, $e \neq latest_{p \rightarrow r}(I)$, which is a contradiction. So, $e = latest_{r' \rightarrow r}(e' \downarrow) = latest_{q \rightarrow r' \rightarrow r}(I)$ and we are done. \square

Our observation about frontier events immediately gives us a way to compare primary information using both primary and secondary information.

Lemma 2 *Let I be an ideal and $p, q, r \in \mathcal{P}$. Let $e = latest_{p \rightarrow r}(I)$ and $f = latest_{q \rightarrow r}(I)$. Then $e \sqsubseteq^* f$ iff $e \in secondary_q(I)$.*

Proof

(\Leftarrow) Suppose $e \in secondary_q(I)$. Then $r \in e \in I|_q$ and so $e \sqsubseteq^* f \in I|_q$ by the definition of $latest_{q \rightarrow r}(I)$.

(\Rightarrow) If $e = f$, $e \in primary_q(I) \subseteq secondary_q(I)$, and there is nothing to prove. If $e \neq f$, then there exists an event e' such that $e \triangleleft_r e' \triangleleft_r^* f$ and so $e \in I|_p \cap I|_q$. We know that $e' \in I|_q - I|_p$, so e is an r -sentry in $frontier_{pq}(I)$. But then, by our previous lemma, $e \in primary_p(I) \cap secondary_q(I)$ and we are done. \square

Suppose p and q synchronize at an action a after u . At this point they “share” their primary and secondary information. For $r \notin \{p, q\}$, if the event $latest_{p \rightarrow r}(\mathcal{E}_u)$ is also

present in q 's set of secondary events $secondary_q(\mathcal{E}_u)$, both p and q know that q 's latest r -event $latest_{q \rightarrow r}(\mathcal{E}_u)$ is at least as recent as $latest_{p \rightarrow r}(\mathcal{E}_u)$. So, after the synchronization, $latest_{q \rightarrow r}(\mathcal{E}_{ua})$ is the same as $latest_{q \rightarrow r}(\mathcal{E}_u)$, whereas p inherits this information from q —i.e., $latest_{p \rightarrow r}(\mathcal{E}_{ua}) = latest_{q \rightarrow r}(\mathcal{E}_u)$. In this way, for each $r \in \mathcal{P}$, p and q locally update their primary information about r in \mathcal{E}_{ua} . Clearly $latest_{p \rightarrow q}(\mathcal{E}_{ua}) = latest_{q \rightarrow p}(\mathcal{E}_{ua}) = e_a$, where e_a is the new event—i.e., $\mathcal{E}_{ua} - \mathcal{E}_u = \{e_a\}$.

This procedure generalizes to any arbitrary set $P \subseteq \mathcal{P}$ which synchronizes after u . The processes in P share their primary and secondary information and compare this information pairwise. Using Lemma 2, for each $q \in \mathcal{P} - P$ they decide who has the “latest information” about q . Each process then comes away with the best primary information from P .

Once we have compared primary information, updating secondary information is straightforward. Clearly, if $latest_{q \rightarrow r}(I)$ is better than $latest_{p \rightarrow r}(I)$, then every secondary event $latest_{q \rightarrow r \rightarrow r'}(I)$ must also be better than the corresponding event $latest_{p \rightarrow r \rightarrow r'}(I)$. So, secondary information can be locally updated too. In other words, to consistently update primary and secondary information, it suffices to correctly compare primary information, which is achieved by Lemma 2.

After a synchronization involving $P \subseteq \mathcal{P}$, notice that all processes in P will come away with the *same* set of primary and secondary events.

From the preceding argument, it is clear that the new event belongs to the primary (and hence secondary) information of the processes which synchronize at that event. Further, the update procedure reveals that if an event disappears from the secondary information of all the processes, it will never reappear as secondary information at some later stage. This is captured formally in the following proposition.

Proposition 3 *Let $u, w \in \Sigma^*$ such that $w = ua$ for some $a \in \Sigma$. Let e_a denote the new event in w —i.e., $\mathcal{E}_w - \mathcal{E}_u = \{e_a\}$. Then:*

- $e_a \in primary_{\mathcal{P}}(\mathcal{E}_w)$.
- $primary_{\mathcal{P}}(\mathcal{E}_w) \subseteq \{e_a\} \cup primary_{\mathcal{P}}(\mathcal{E}_u)$.
- $secondary_{\mathcal{P}}(\mathcal{E}_w) \subseteq \{e_a\} \cup secondary_{\mathcal{P}}(\mathcal{E}_u)$.

4 Locally updating primary/secondary information

To make Lemma 2 effective, we must make the assertions “locally checkable”—e.g., if $e = latest_{p \rightarrow r}(I)$, processes p and q must be able to decide if $e \in secondary_q(I)$.

Recall that e is represented in $primary_p(I)$ as a triple of the form (p, r, e) . So, to check if $e \in secondary_q(I)$, q has to look for a quadruple of the form $(q, r', r'', e) \in secondary_q(I)$, where $r', r'' \in \mathcal{P}$. This can be checked locally provided events in \mathcal{E}_u are labelled unambiguously while u is being read.

Clearly, labelling each event e as a pair $(i, u(i))$ is impossible since, in general, there is no agent which can consistently supply all processes with the “correct” value of i . Instead, we may naïvely assume that events in \mathcal{E}_u are locally assigned distinct labels—in effect, at each action a , the processes in a together assign a (sequential) time-stamp to the new

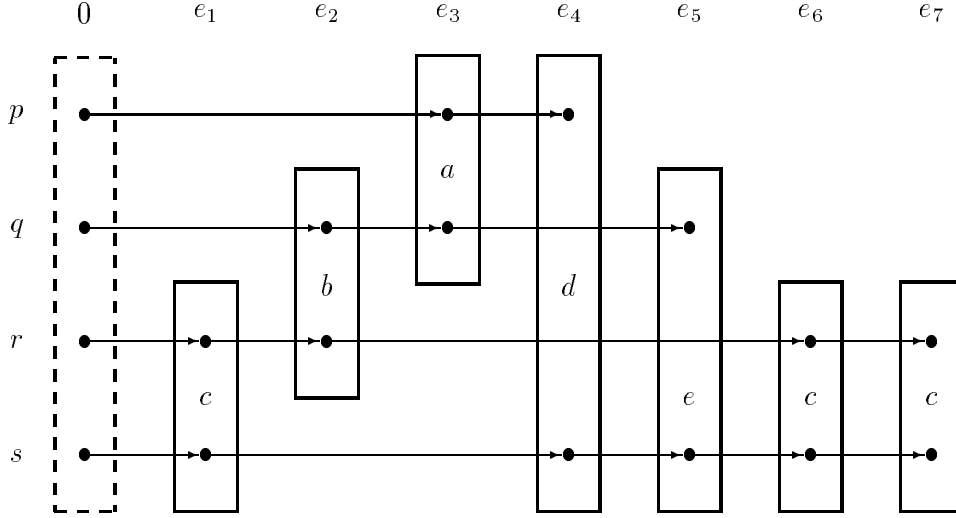


Figure 2: At e_7 , $e_1 \in \text{secondary}_p(e_7\downarrow)$, but r and s cannot “see” e_1 in $\text{secondary}_{\{r,s\}}(e_7\downarrow)$.

occurrence of a .² In this manner, the processes in \mathcal{P} can easily assign consistent local time-stamps for each action which will let them compute the relations \triangleleft_p^* between events.

The problem with this approach is that we will need an unbounded set of time-stamps, since u could get arbitrarily large. Instead we would like a scheme which uses only a finite set of labels to distinguish events. This means that several different occurrences of the same action will eventually get the same label. Since the update of primary and secondary information relies on comparing labels, we must ensure that this reuse of labels does not lead to any confusion.

However, from Lemma 2, we know that to compare primary information, we only need to look at the events which are currently in the primary and secondary sets of each process. So, it is sufficient if the labels assigned to these sets are consistent across the system—i.e., if the same label appears in the current primary or secondary information of different processes, the corresponding event is actually the same.

Notice that we do not need to maintain a global temporal order on labels across the system. Lemma 2 assures us that to compare events of interest to us, it suffices to check for *equality* of labels assigned to the events.

Suppose we have such a labelling on u and we want to extend this to a consistent labelling on $w = ua$ —i.e., we need to assign a label to the new a -event. By Proposition 3, it suffices to use a label which is distinct from the labels of all the a -events currently in the secondary information of \mathcal{E}_u . Since the cardinality of $\text{secondary}_{\mathcal{P}}(\mathcal{E}_u)$ is bounded, such a new label must exist. The catch is to detect which labels are currently in use and which are not.

Unfortunately, the processes in a cannot directly see all the a -events which belong to the secondary information of the entire system. An a -event e may be part of the secondary information of processes *outside* a —i.e., $e \in \text{secondary}_{\mathcal{P}-a}(\mathcal{E}_u) - \text{secondary}_a(\mathcal{E}_u)$.

EXAMPLE: Let $\mathcal{P} = \{p, q, r, s\}$ and $\Sigma = \{a, b, c, d, e\}$ where $a = \{p, q\}$, $b = \{q, r\}$,

²Recall that for each action $\hat{a} \in \Sigma$, $\theta(\hat{a}) = a$, and we use a to denote both the action and the subset of processes which synchronize at that action (see page 7).

$c = \{r, s\}$, $d = \{p, s\}$ and $e = \{q, s\}$. Figure 2 shows the events \mathcal{E} corresponding to the word $cbadecc$.

At the end of this word, $e_1 = \text{latest}_{p \rightarrow q \rightarrow s}(\mathcal{E})$. However, $e_1 \notin \text{secondary}_s(\mathcal{E})$;

$$\begin{aligned} \text{secondary}_s(\mathcal{E}) = \{ & (s, p, p, e_4), (s, p, q, e_3), (s, p, r, e_2), (s, p, s, e_4), \\ & (s, q, p, e_4), (s, q, q, e_5), (s, q, r, e_2), (s, q, s, e_5), \\ & (s, r, p, e_4), (s, r, q, e_5), (s, r, r, e_7), (s, r, s, e_7), \\ & (s, s, p, e_4), (s, s, q, e_5), (s, s, r, e_7), (s, s, s, e_7) \}. \end{aligned}$$

Since $\max_r(\mathcal{E}) = \max_s(\mathcal{E})$, $\text{secondary}_r(\mathcal{E}) = \text{secondary}_s(\mathcal{E})$ when viewed as sets of events. So, $e_1 \notin \text{secondary}_r(\mathcal{E})$ either. Thus, e_1 is a c -event which belongs to $\text{secondary}_p(\mathcal{E}) - \text{secondary}_c(\mathcal{E})$.

To enable the processes in a to know about all a -events in $\text{secondary}_p(\mathcal{E}_u)$, we need to maintain tertiary information.

Tertiary information The tertiary information of p after I , $\text{tertiary}_p(I)$, is the (indexed) set $\bigcup_{q \in \mathcal{P}} \text{secondary}_q(\text{latest}_{p \rightarrow q}(I) \downarrow)$. In other words, this is the latest information that p has in I about the secondary information of q , for all $q \in \mathcal{P}$. As before, for $P \subseteq \mathcal{P}$, $\text{tertiary}_P(I) = \bigcup_{p \in P} \text{tertiary}_p(I)$.

Each event in $\text{tertiary}_p(I)$ is of the form $\text{latest}_{q \rightarrow r \rightarrow s}(\text{latest}_{p \rightarrow q}(I) \downarrow)$ for some $q, r, s \in \mathcal{P}$. We abbreviate $\text{latest}_{q \rightarrow r \rightarrow s}(\text{latest}_{p \rightarrow q}(I) \downarrow)$ by $\text{latest}_{p \rightarrow q \rightarrow r \rightarrow s}(I)$. We represent each event $e = \text{latest}_{p \rightarrow q \rightarrow r \rightarrow s}(I)$ as a quintuple (p, q, r, s, e) in $\text{tertiary}_p(I)$. However, for convenience we will work with $\text{tertiary}_p(I)$ as though it were simply a set of events, rather than an indexed set, just as we have been doing with primary and secondary information.

Just as $\text{primary}_p(I) \subseteq \text{secondary}_p(I)$, clearly $\text{secondary}_p(I) \subseteq \text{tertiary}_p(I)$ since each secondary event $\text{latest}_{p \rightarrow q \rightarrow r}(I)$ is also a tertiary event $\text{latest}_{p \rightarrow p \rightarrow q \rightarrow r}(I)$ (or, equivalently, $\text{latest}_{p \rightarrow q \rightarrow q \rightarrow r}(I)$ and so on).

Lemma 4 Let I be an ideal and $p \in \mathcal{P}$. If $e \in \text{secondary}_p(I)$ then for every $q \in e$, $e \in \text{tertiary}_q(I)$.

Proof Let $e \in \text{secondary}_p(I)$ and $q \in e$. Concretely, let $e = \text{latest}_{p \rightarrow p' \rightarrow p''}(I)$ for some $p', p'' \in \mathcal{P}$. We know that $e \in I|_p \cap I|_q$ and there is a path $e \sqsubset f_1 \sqsubset \dots \sqsubset \max_p(I)$ leading from e to $\max_p(I)$ which passes through $e' = \text{latest}_{p \rightarrow p'}(I)$.

Suppose this path never leaves $I|_p \cap I|_q$. Then $\max_p(I) \in I|_q$ and so $\max_p(I) = \text{latest}_{q \rightarrow p}(I)$. This means that $e \in \text{secondary}_p(\text{latest}_{q \rightarrow p}(I) \downarrow) \subseteq \text{tertiary}_q(I)$ and we are done.

Otherwise, the path from e to $\max_p(I)$ does leave $I|_p \cap I|_q$ at some stage.

If $e' \notin I|_p \cap I|_q$ then for some $f, f' \in \mathcal{E}$ and some $r \in \mathcal{P}$ we have $f \in I|_p \cap I|_q$, $f' \in I|_p - I|_q$ and $e \sqsubseteq^* f \triangleleft_r f' \sqsubseteq^* e'$. This means that $f \in \text{frontier}_{qp}(I)$ is an r -sentry and by our earlier argument we know that $f = \text{latest}_{q \rightarrow r}(I)$. So $e = \text{latest}_{q \rightarrow r \rightarrow p''}(I) = \text{latest}_{q \rightarrow q \rightarrow r \rightarrow p''}(I) \in \text{tertiary}_q(I)$.

On the other hand, if $e' \in I|_p \cap I|_q$ we can find an r -sentry $f \in \text{frontier}_{qp}(I)$ on the path from e' to $\max_p(I)$, for some $r \in \mathcal{P}$. We once again get $f = \text{latest}_{q \rightarrow r}(I)$ and so $e = \text{latest}_{q \rightarrow r \rightarrow p' \rightarrow p''}(I) \in \text{tertiary}_q(I)$. □

We shall use this lemma in the following form.

Corollary 5 *Let I be an ideal, $p \in \mathcal{P}$ and e a p -event in I . If $e \notin \text{tertiary}_p(I)$ then $e \notin \text{primary}_{\mathcal{P}}(I) \cup \text{secondary}_{\mathcal{P}}(I)$.*

So, a process p can keep track of which of its labels are “in use” in the system by maintaining tertiary information. Each p -event e initially belongs to $\text{primary}_e(I)$, and hence to $\text{secondary}_e(I)$ and $\text{tertiary}_e(I)$ as well. (Recall that for an event e , we also use e to denote the subset of \mathcal{P} which meets at e). As the computation progresses, e gradually “recedes” into the background and disappears from the primary and secondary sets of the system. Eventually, when e disappears from $\text{tertiary}_p(I)$, p can be sure that e no longer belongs to $\text{primary}_{\mathcal{P}}(I) \cup \text{secondary}_{\mathcal{P}}(I)$.

Since $\text{tertiary}_p(I)$ is a bounded set, p knows that only finitely many of its labels are in use at any given time. So, by using a sufficiently large finite set of labels, each new event can always be assigned an unambiguous label by the processes which take part in the event.

5 The “gossip” automaton

Using our analysis of primary, secondary and tertiary information of processes, we can now design a deterministic asynchronous automaton to keep track of the “latest gossip”—i.e., to consistently update primary information whenever a set of processes synchronizes. The processes in the automaton maintain this information in terms of time-stamps: when an event e occurs, it is assigned a time-stamp $\tau(e)$ by the processes which take part in e .

For $p \in \mathcal{P}$, a local state of p is defined by three associative arrays \mathbf{prim}_p , \mathbf{sec}_p and \mathbf{ter}_p , indexed by \mathcal{P} , $\mathcal{P} \times \mathcal{P}$ and $\mathcal{P} \times \mathcal{P} \times \mathcal{P}$ respectively. Each entry in these arrays is a pair of the form $\langle P, \ell \rangle$, where P is a subset of \mathcal{P} and ℓ is a label drawn from a finite set \mathcal{L} —it is sufficient for \mathcal{L} to have N^3+1 entries, where N is the number of processes in \mathcal{P} .

After reading a word u , each entry $\langle P, \ell \rangle$ in the arrays \mathbf{prim}_p , \mathbf{sec}_p and \mathbf{ter}_p corresponds to the time-stamp $\tau(e)$ of some event $e \in \mathcal{E}_u$. We shall show that for all $q, r, s \in \mathcal{P}$, the values stored in $\mathbf{prim}_p(q)$, $\mathbf{sec}_p(q, r)$ and $\mathbf{ter}_p(q, r, s)$ are, in fact, the time-stamps of the events $\text{latest}_{p \rightarrow q}(\mathcal{E}_u)$, $\text{latest}_{p \rightarrow q \rightarrow r}(\mathcal{E}_u)$ and $\text{latest}_{p \rightarrow q \rightarrow r \rightarrow s}(\mathcal{E}_u)$, respectively.

The initial state is the global state in which, for each process p , all entries in \mathbf{prim}_p , \mathbf{sec}_p and \mathbf{ter}_p are set to $\langle \mathcal{P}, \ell_0 \rangle$, where ℓ_0 is an arbitrary but fixed label from \mathcal{L} . In other words, the processes jointly assign the time-stamp $\tau(0) = \langle \mathcal{P}, \ell_0 \rangle$ to the initial event 0.

The local transition functions \rightarrow_a modify the local states for processes in a as follows.

- (i) When an a -event e_a occurs, the processes in a first choose a time-stamp $\tau(e_a) = \langle a, \ell \rangle$ such that for some $p \in a$, $\langle a, \ell \rangle$ does not appear in \mathbf{ter}_p

(Recall that \mathcal{L} has N^3+1 labels. Since there are at most N^3 labels in \mathbf{ter}_p for each p , we will always be able to find an unused label in \mathcal{L} to assign to e_a . In general, we have to choose $\langle a, \ell \rangle$ in a canonical way. One possibility is to linearly order the sets \mathcal{P} and \mathcal{L} . Let p_a be the smallest process in a with respect to the ordering on \mathcal{P} . To fix the time-stamp $\langle a, \ell \rangle$ assigned to e_a , choose ℓ such that it is the smallest label with respect to the ordering on \mathcal{L} which does not appear in \mathbf{ter}_{p_a} .)

- (ii) Let $a = \{p_1, p_2, \dots, p_k\}$. Corresponding to each process $r \notin a$, fix a process $\pi_r \in a$ as follows.

$\pi_r := p_1;$
 for each i in $\{2, 3, \dots, k\}$ do
 if the value stored in $\text{prim}_{\pi_r}(r)$ appears in sec_{p_i} then $\pi_r := p_i;$
 od

(iii) For each $p \in a$, update prim_p , sec_p and ter_p in phases.

- First, update $\{\text{prim}_p\}_{p \in a}$ as follows.
 For each $q \in a$, $\text{prim}_p(q) := \langle a, \ell \rangle$.
 For each $r \in \mathcal{P} - a$, $\text{prim}_p(r) := \text{prim}_{\pi_r}(r)$.
- Next, update $\{\text{sec}_p\}_{p \in a}$ as follows.
 For each $q, r \in \mathcal{P}$ such that $q \in a$, $\text{sec}_p(q, r) := \text{prim}_q(r)$.
 For each $q, r \in \mathcal{P}$ such that $q \notin a$, $\text{sec}_p(q, r) := \text{sec}_{\pi_q}(q, r)$.
- Finally, update $\{\text{ter}_p\}_{p \in a}$ as follows.
 For each $q, r, s \in \mathcal{P}$ such that $q \in a$, $\text{ter}_p(q, r, s) := \text{sec}_q(r, s)$.
 For each $q, r, s \in \mathcal{P}$ such that $q \notin a$, $\text{ter}_p(q, r, s) := \text{ter}_{\pi_q}(q, r, s)$.

We now verify that the arrays maintained by each process p in the gossip automaton always record the time-stamps of the primary, secondary and tertiary information of p .

Proposition 6 *Let $u \in \Sigma^*$. For all processes $p \in \mathcal{P}$, the values stored in prim_p , sec_p and ter_p after reading u satisfy the following properties.*

- For all $q \in \mathcal{P}$, $\text{prim}_p(q) = \tau(\text{latest}_{p \rightarrow q}(\mathcal{E}_u))$.
- For all $q, r \in \mathcal{P}$, $\text{sec}_p(q, r) = \tau(\text{latest}_{p \rightarrow q \rightarrow r}(\mathcal{E}_u))$.
- For all $q, r, s \in \mathcal{P}$, $\text{ter}_p(q, r, s) = \tau(\text{latest}_{p \rightarrow q \rightarrow r \rightarrow s}(\mathcal{E}_u))$.

Moreover, for all processes $p, q \in \mathcal{P}$ and events $e \in \text{primary}_p(\mathcal{E}_u) \cup \text{secondary}_p(\mathcal{E}_u)$ and $f \in \text{primary}_q(\mathcal{E}_u) \cup \text{secondary}_q(\mathcal{E}_u)$, if $\tau(e) = \tau(f)$ then $e = f$.

Proof The proof is by induction on n , the length of u .

($n = 0$)

The base case is when u is the empty string. We know that for each $p \in \mathcal{P}$, all entries in prim_p , sec_p and ter_p are initially set to $\langle \mathcal{P}, \ell_0 \rangle$, where $\langle \mathcal{P}, \ell_0 \rangle$ is the time-stamp assigned to the initial event 0. Clearly, these values satisfy all the conditions specified in the statement of the proposition.

($n > 0$)

Suppose $u = wa$. The time-stamp $\langle a, \ell \rangle$ assigned to the new a -event e_a in step (i) of the transition does not appear in ter_p for some $p \in a$. By the induction hypothesis, the values stored in ter_p after reading w are the time-stamps of the corresponding events in $\text{ter}_p(\mathcal{E}_w)$. Lemma 4 then guarantees that no event in $\text{primary}_p(\mathcal{E}_w) \cup \text{secondary}_p(\mathcal{E}_w)$ is time-stamped $\langle a, \ell \rangle$ —if there were such an event, it would be in the tertiary information $\text{ter}_p(\mathcal{E}_w)$ of every process $q \in a$ as a result of which, by the induction hypothesis, the label $\langle a, \ell \rangle$ would be in ter_q after reading w for every $q \in a$.

By the induction hypothesis, we also know that no two distinct events in $\text{primary}_p(\mathcal{E}_w) \cup \text{secondary}_p(\mathcal{E}_w)$ are assigned the same time-stamp. Since $\text{primary}_p(\mathcal{E}_u) \subseteq \{e_a\} \cup \text{primary}_p(\mathcal{E}_w)$ and $\text{secondary}_p(\mathcal{E}_u) \subseteq \{e_a\} \cup \text{secondary}_p(\mathcal{E}_w)$ (Proposition 3), it follows

that the time-stamps assigned to $primary_{\mathcal{P}}(\mathcal{E}_u) \cup secondary_{\mathcal{P}}(\mathcal{E}_u)$ are also distinct. In other words, for all processes $p, q \in \mathcal{P}$ and events $e \in primary_p(\mathcal{E}_u) \cup secondary_p(\mathcal{E}_u)$ and $f \in primary_q(\mathcal{E}_u) \cup secondary_q(\mathcal{E}_u)$, if $\tau(e) = \tau(f)$ then $e = f$.

We now have to verify that the updated values in \mathbf{prim}_p , \mathbf{sec}_p and \mathbf{ter}_p for each $p \in a$ are the time-stamps of the corresponding events in the primary, secondary and tertiary information of p after u .

By the induction hypothesis, for each $p \in a$, the values in \mathbf{prim}_p and \mathbf{sec}_p after reading w are the time-stamps assigned to the corresponding events in $primary_p(\mathcal{E}_w)$ and $secondary_p(\mathcal{E}_w)$ respectively. By the induction hypothesis, we also know that for all $p, q, r, s, s' \in \mathcal{P}$, if the value stored in $\mathbf{prim}_p(r)$ after reading w is the same as the value stored in $\mathbf{sec}_q(s, s')$ after reading w then, in fact, the event $latest_{p \rightarrow r}(\mathcal{E}_w)$ is the same as the event $latest_{q \rightarrow s \rightarrow s'}(\mathcal{E}_w)$.

It then follows that for $r \notin a$, the process $\pi_r \in a$ identified in step (ii) of the transition is one of the processes in a which has the best primary information about r after w —by Lemma 2, for $p, q \in a$ and $r \notin a$, $latest_{p \rightarrow r}(\mathcal{E}_w) \sqsubseteq^* latest_{q \rightarrow r}(\mathcal{E}_w)$ iff the event $latest_{p \rightarrow r}(\mathcal{E}_w)$ appears in $secondary_q(\mathcal{E}_w)$, which is equivalent to checking that the time-stamp $\tau(latest_{p \rightarrow r}(\mathcal{E}_w))$ stored in $\mathbf{prim}_p(r)$ appears in \mathbf{sec}_q .

From this, it is easy to see that for all $p \in a$, the updates made to \mathbf{prim}_p , \mathbf{sec}_p and \mathbf{ter}_p in step (iii) ensure that for all $q, r, s \in \mathcal{P}$, $\mathbf{prim}_p(q) = \tau(latest_{p \rightarrow q}(\mathcal{E}_u))$, $\mathbf{sec}_p(q, r) = \tau(latest_{p \rightarrow q \rightarrow r}(\mathcal{E}_u))$ and $\mathbf{ter}_p(q, r, s) = \tau(latest_{p \rightarrow q \rightarrow r \rightarrow s}(\mathcal{E}_u))$. \square

The gossip automaton does not have any final states, since we do not need to accept any language. Instead, we define for each $a \in \Sigma$ a function $g_a : V_a \rightarrow (2^{\mathcal{P}} - \{\emptyset\})^{\mathcal{P}}$ which checks the arrays \mathbf{prim}_p and \mathbf{sec}_p of each process p in a and computes for each $q \in \mathcal{P}$, the set of processes in a which have the most recent information about q .

(A small technical point: g_a must be defined for all states in V_a . However, not all combinations of local states may be “meaningful”. We can easily assemble local states to form an a -state for which the inductive assertions do not hold, as a result of which our procedure for comparing primary information breaks down. However, since such a -states are unreachable, we can ignore this problem and simply assign a default value to g_a in these cases—for instance, the default value could be the function $\{p \mapsto a\}_{p \in \mathcal{P}}$.)

This immediately yields the result we set out to establish.

Theorem 7 *Let (Σ, θ) be the distributed alphabet corresponding to $\mathcal{C} \subseteq (2^{\mathcal{P}} - \{\emptyset\})$. The family of functions $\{latest_gossip_c : \Sigma^* \rightarrow (2^{\mathcal{P}} - \{\emptyset\})^{\mathcal{P}}\}_{c \in \Sigma}$ is locally computable.*

The size of the gossip automaton

Lemma 8 *In the gossip automaton, the local state of each process $p \in \mathcal{P}$ can be described using $O(N^3 \log N)$ bits, where $N = |\mathcal{P}|$.*

Proof A local state for p consists of the arrays \mathbf{prim}_p , \mathbf{sec}_p and \mathbf{ter}_p . We estimate how many bits are required to store this.

Recall that for any ideal I , each event in $primary_p(I)$ is also present in $secondary_p(I)$. Similarly, each event in $secondary_p(I)$ is also present in $tertiary_p(I)$. So it suffices to store just the labels of tertiary events in the array \mathbf{ter}_p . By fixing an ordering of $\mathcal{P} \times \mathcal{P} \times \mathcal{P}$, these events can be stored as a list with N^3 entries.

Each new event e was assigned a label of the form $\langle P, \ell \rangle$, where P was the set of processes that participated in e and $\ell \in \mathcal{L}$.

We have already seen that it suffices to have $O(N^3)$ labels in \mathcal{L} . As a result, each label $\ell \in \mathcal{L}$ can be written down using $O(\log N)$ bits.

To write down $P \subseteq \mathcal{P}$, we need, in general, N bits. This component of the label is required to guarantee that all secondary events in the system have distinct labels, since the set \mathcal{L} is common across all processes. However, we do not really need to use all of P in the label for e to ensure this property. If we fix a linear order on \mathcal{P} , it suffices to label e by $\langle p_e, \ell \rangle$ where, among the processes participating in e , p_e is the smallest with respect to the ordering on \mathcal{P} . It is easy to verify that Proposition 6 continues to hold with respect to these modified labels.

Thus, we can modify our automaton so that the processes label each event by a pair $\langle p, \ell \rangle$, where $p \in \mathcal{P}$ and $\ell \in \mathcal{L}$. This pair can be written down using $O(\log N)$ bits. Overall there are N^3 such pairs in the array of tertiary events, so the whole state can be described using $O(N^3 \log N)$ bits. \square

The preceding lemma implies that the number of local states of a process could be exponential in N , the number of processes in the system. In general, this would mean that we require an exponential amount of space to specify the entire automaton (though each state can be described using a polynomial number of bits) since the transition table of the automaton would have an exponential number of entries.

However, in this case, we do not need to exhaustively enumerate the entire state space and transition table to completely describe the automaton. Since the states are structured entities and the transition function is presented as an algorithm which manipulates the data in these structured states, the gossip automaton can in fact be effectively presented in space polynomial in N . This ability to build the gossip automaton efficiently “on the fly” is crucial for embedding it in other distributed algorithms.

6 Extensions and Applications

Beyond tertiary information

The gossip automaton consistently labels all primary and secondary events. In other words, at any point, distinct primary and secondary events have distinct time-stamps. By Lemma 2, this is sufficient to correctly compare and update primary information.

However, we may want to keep track of events which are older than secondary events. We can generalize the definition of secondary and tertiary information and inductively define the k -ary information of a process p with respect to an ideal I , for any natural number k .

The case $k = 1$ corresponds to primary information. For $k > 1$, the k -ary information of p after I , $k\text{-ary}_p(I)$, is the (indexed) set $\bigcup_{q \in \mathcal{P}} (k-1)\text{-ary}_p(\text{latest}_{p \rightarrow q}(I) \downarrow)$. In other words, this is the latest information that p has in I about the $(k-1)$ -ary information of q , for all $q \in \mathcal{P}$.

It is not difficult to see that the argument in Lemma 4 can be extended to yield the following result, proved formally in [15].

Lemma 9 *Let I be an ideal and $p \in \mathcal{P}$. If $e \in k\text{-ary}_p(I)$ then for every $q \in e$, $e \in (k+1)\text{-ary}_q(I)$.*

Together with Lemma 2, this implies that the gossip automaton can be modified to maintain consistent time-stamps for all events upto depth k , for any natural number k .

Optimizing the automaton

If we are only concerned with comparing and updating primary information, the gossip automaton can be made more concise. To achieve this, each process maintains its primary information as a directed graph which reflects the underlying \sqsubseteq^* -order between primary events, rather than storing the information as a simple array of labels as we have described here. It turns out that processes can compare and update these *primary graphs* without recourse to secondary information. Secondary information is then required only to ensure that new events get unused labels. Tertiary information is dispensed with altogether. This leads to the following result (details can be found in [15]).

Lemma 10 *The gossip automaton can be modified so that the local state of each process $p \in \mathcal{P}$ is describable using $O(N^2 \log N)$ bits, where $N = |\mathcal{P}|$.*

It is important to note that this optimized automaton consistently labels *only* primary events. Secondary events could get inconsistent labels. This becomes relevant when we come to applications of the gossip automaton.

Further optimizations are possible if the structure of the alphabet (Σ, θ) is such that synchronizations are “well-behaved”. One example is when the processes are located on the vertices of an d -dimensional hypercube (i.e., $|\mathcal{P}| = 2^d$), with synchronizations taking place along faces of the hypercube. Once again, details can be found in [15].

Applications

Building the gossip automaton is a basic step in tackling many problems in the theory of asynchronous automata.

Asynchronous automata play an important role in the theory of distributed systems because of their close connection to *trace theory*. Trace theory, initiated by Mazurkiewicz [19], is a language-theoretic approach to the study of concurrent systems. Traditionally, formal language theory models computations of sequential systems as strings over an abstract alphabet. Instead, trace theory describes computations of concurrent systems in terms of equivalence classes of strings, called *traces*. All strings in a trace are equivalent upto permutation of adjacent letters which belong to an underlying *independence relation* on the alphabet. Thus, the different elements of a trace correspond to different sequential observations of the same concurrent computation.

A deep theorem of Zielonka [30] shows that asynchronous automata are a natural distributed machine model for recognizing trace languages. Zielonka’s original proof of the connection between these automata and recognizable trace languages is widely accepted to be difficult to assimilate. In [23], we show that the gossip automaton can be used to provide a more structured proof of Zielonka’s theorem. Overall, the construction we present is quite similar to the one described in Zielonka’s original paper. However, we feel that by separating out clearly the role played by the gossip automaton, our proof is much easier to digest. (Other new proofs of Zielonka’s theorem exist [2], but these are based on asynchronous *cellular* automata [31], a slightly different—and, in our opinion, less intuitive—machine model.)

In [13], a determinization construction is presented for asynchronous automata using a generalization of the classical subset construction for finite automata. This construction allows us to keep track of the global states which are currently valid at any stage of a computation by a non-deterministic asynchronous automaton.

The key problem to be tackled in the determinization construction is to estimate the amount of information that can be safely “forgotten” by the subset automaton without sacrificing global consistency. It turns out that it is sufficient for each process to maintain “histories” upto the level of secondary events. This is accomplished by using the gossip automaton presented here to assign consistent time-stamps to all secondary events. Notice that the optimized gossip automaton of [15] cannot be used in the determinization construction since it only guarantees consistent time-stamps upto primary events.

In the past few years, there has been a lot of interest in extending asynchronous automata to process infinite inputs. Analogous to Büchi automata on infinite strings, Gastin and Petit [7] have defined Büchi asynchronous automata which operate on infinite traces. Although it has been known for a while that these automata are closed under complementation for algebraic reasons, the only direct complementation construction provided so far has been for Büchi asynchronous *cellular* automata [24]. Safra’s determinization construction for Büchi automata [26] has recently been extended to Büchi asynchronous automata [14]. This provides a natural procedure for complementing these automata. Once again, the gossip automaton plays a crucial role in the construction.

Automata on infinite strings have always been closely linked to logic [28]. This connection holds for Büchi asynchronous automata as well. Recently, Thiagarajan [27] has developed an extension of propositional linear-time temporal logic which is interpreted over infinite traces rather than infinite linear sequences. This logic appears to be quite expressive, while remaining decidable (unlike several other partial-order logics studied in the literature [18]). The decision procedure for this logic is automata-theoretic, in the style of Vardi and Wolper [29], except that it makes use of Büchi asynchronous automata rather than conventional Büchi automata. The gossip automaton plays a crucial role in this construction as well. In fact, *this* was the original motivation for constructing the gossip automaton.

Asynchronous communication

In this paper, we have only dealt with synchronous communication. Synchronous communication achieves coördination among independent agents by periodically permitting subsets of agents to pool information together to make a decision.

Another standard way for agents to exchange information is through message-passing. This mode of exchanging information is usually referred to as asynchronous communication, since there may be an arbitrary delay between the time when a message is sent and the time when it is received.

Synchronous communication is easier to handle, both from a theoretical standpoint as well as from the point of view of programming. It is no coincidence that languages like CCS [20] and CSP [11] which have been developed for specifying communicating systems assume synchronous communication as the basic means of exchanging information.

However, when agents are widely separated in space, asynchronous communication is generally the only practical way of achieving coördination. So, there is a considerable body of literature devoted to distributed algorithms and protocols for asynchronous systems

[16].

We have extended our approach to deal with message-passing systems where the communication medium is reliable, subject to certain restrictions on the number of unacknowledged messages that can be present in the system at any given time [21]. This shows that our analysis based on primary, secondary and tertiary information is applicable to a much wider range of distributed systems than the purely synchronous systems discussed in this paper.

7 Discussion

We now discuss the connections between our results and other work in related areas.

Studies of “gossiping” in networks have traditionally focussed on efficiently disseminating a fixed piece of information (or gossip) from one node to all other nodes in a network [9]. The main aim is to find an optimal sequence of communications to distribute data for a given network topology.

Israeli and Li [12] introduced the notion of “bounded time-stamps” and argued that these were fundamental in solving many problems in distributed systems—notably that of creating what are called “atomic registers” [17]. Their results have been extended and considerably simplified by a number of others [3, 5, 6, 8]. However, this line of work is based on a shared memory model, which is quite different in spirit from the asynchronous automaton model. Though it may be possible to formally translate their results to our framework, we feel that the intuition underlying their time-stamping algorithms is quite different from ours. In general, solutions to time-stamping problems seem to depend heavily on the underlying model of a distributed system that is used and it is difficult to compare such algorithms across models.

More closely connected to our work are the bounded time-stamping algorithms implemented using asynchronous cellular automata in the study of recognizable trace languages [1, 2]. The synchronization mechanism of asynchronous cellular automata is quite different from that of asynchronous automata. Asynchronous cellular automata correspond more closely to shared memory systems than to systems that communicate by direct inter-process communication. Thus, though the overall goal of the constructions in [1, 2] is closely related to the gossip problem which we have studied here, the two time-stamping algorithms appear to be incomparable since they are based on fundamentally different assumptions about how processes interact.

Our definition of locally computable functions is closely linked to the notion of *asynchronous mappings* which is fundamental to the constructions used in [2]. It would be interesting to formally characterize the class of locally computable functions and develop a methodology for automatically synthesizing an asynchronous automaton to compute a given locally computable function, as is done for asynchronous mappings in [2].

Though our algorithm can be implemented as an asynchronous automaton, it correctly computes the latest gossip function locally for *any* input word. In other words, the set of communication sequences generated by the underlying system need not be regular. Our algorithm will also work on sequences generated, for instance, by N communicating Turing machines. Thus, any distributed algorithm which needs to keep track of the flow of information between processes can incorporate the gossip automaton as a “background” routine. Moreover, since the gossip automaton can be constructed efficiently “on the

fly”, embedding it in another algorithm does not involve a very significant computational overhead.

Since our algorithm does not add any extra messages to the underlying computation, it is automatically resilient to failures—even if some processes stop functioning, our algorithm will continue to update primary information correctly for those parts of the system which are still active in the underlying computation. This in-built fault-tolerance is also present in the shared memory time-stamping protocols of [1, 2, 3, 5, 8, 12] and others.

In an asynchronous automaton, each move is *simultaneously* performed by all participating processes. This strong atomicity assumption is not crucial for our algorithm. Notice that the label of an event is *uniquely* fixed by the primary, secondary and tertiary information of the participating processes. Thus, we could implement each multi-way synchronization as a series of synchronous communications where processes exchange their local information pairwise and then proceed independently. The deterministic nature of our algorithm guarantees that each process which participates in an event will locally choose the *same* time-stamp for the new event, since all processes taking part in the event base their choice of time-stamp on identical shared information. Similarly, once each process has collected the primary, secondary and tertiary information of every other participating process, it can update its local information unambiguously. Moreover, it can also compute the updated information of every other process which took part in the event.

The construction of the gossip automaton establishes a non-obvious property for *all* systems with synchronous communication. Suppose an agent p_1 has a private variable X which no other agent can modify, and agents $\{p_2, p_3, \dots, p_N\}$ keep track of the latest value of X that they have heard of from p_1 (either directly or indirectly). Then, along *any* run of the system, bounded time-stamps suffice for determining which of $\{p_2, p_3, \dots, p_N\}$ have the most recent value of X . This is important, for example, for crash recovery. If the system crashes and p_1 fails to come alive after the crash, the other agents can get together and synthesize an optimal “last-known” state of p_1 by comparing their information about p_1 .

There are obvious parallels between the notion of primary, secondary and tertiary information we use in this paper and the concept of levels of knowledge about events in a distributed system [10, 25]. It would be interesting to formally work out how our approach fits in with knowledge-theoretic techniques for analyzing distributed systems. As far as we are aware, none of the work in knowledge theory has addressed the synchronizing model that we consider here, so establishing a precise connection between the two approaches is not straightforward.

Acknowledgments P.S. Thiagarajan suggested the gossip problem. We have benefited greatly from discussions with him. The reformulation of the gossip automaton in terms of ideals and frontiers emerged during work with Nils Klarlund on determinizing asynchronous automata. His suggestions have been a great help in cleaning up the presentation. We also thank Robert de Simone for pointing out some subtle technical errors. Finally, we thank the referees for several useful comments.

References

- [1] R. CORI, Y. METIVIER: Approximations of a trace, asynchronous automata and the ordering of events in a distributed system, *Proc. ICALP '88, LNCS 317* (1988) 147–161.
- [2] R. CORI, Y. METIVIER, W. ZIELONKA: Asynchronous mappings and asynchronous cellular automata, *Inform. and Comput.*, **106** (1993) 159–202.
- [3] R. CORI, E. SOPENA: Some combinatorial aspects of time-stamp systems, *Europ. J. Combinatorics*, **14** (1993) 95–102.
- [4] V. DIEKERT: *Combinatorics on traces*, *LNCS 454* (1990).
- [5] D. DOLEV, N. SHAVIT: Bounded concurrent time-stamps are constructible, *Proc. 21st ACM STOC* (1989) 454–466.
- [6] C. DWORK, O. WAARTS: Simple and efficient bounded concurrent timestamping or bounded concurrent time-stamps are comprehensible, *Proc. 24th ACM STOC* (1992) 655–666.
- [7] P. GASTIN, A. PETIT: Asynchronous cellular automata for infinite traces, *Proc. ICALP '92, LNCS 623* (1992) 583–594.
- [8] R. GAWLICK, N. LYNCH, N. SHAVIT: Concurrent time stamping made simple, *Proc. ISTCS '92, LNCS 601* (1992) 171–183.
- [9] S.M. HEDETNIEMI, S.T. HEDETNIEMI, A.L. LIESTMAN: A survey of gossiping and broadcasting in communication networks, *Networks*, **18** (1988) 319–349.
- [10] J. HALPERN, Y. MOSES: Knowledge and common knowledge in a distributed environment, *J. Assoc. Comput. Mach.*, **37** (1990) 549–578.
- [11] C.A.R. HOARE: *Communicating Sequential Processes*, Prentice-Hall, London (1984).
- [12] A. ISRAELI, M. LI: Bounded time-stamps, *Proc. 28th IEEE FOCS* (1987) 371–382.
- [13] N. KLARLUND, M. MUKUND, M. SOHONI: Determinizing asynchronous automata, *Proc. ICALP 1994, LNCS 820* (1994) 130–141.
- [14] N. KLARLUND, M. MUKUND, M. SOHONI: Determinizing Büchi asynchronous automata, *Proc. FST&TCS 1995, LNCS 1026* (1995) 456–470.
- [15] R. KRISHNAN, S. VENKATESH: Optimizing the gossip automaton, *Report TCS-94-3*, School of Mathematics, SPIC Science Foundation, Madras, India (1994).
- [16] L. LAMPORT, N. LYNCH: Distributed Computing: Models and Methods, in: J. van Leeuwen (ed.), *Handbook of Theoretical Computer Science: Volume B*, North-Holland, Amsterdam (1990) 1157–1200.
- [17] M. LI, P. VITANYI: How to share concurrent asynchronous wait free variables, *Proc. ICALP '89, LNCS 372* (1989) 488–507.
- [18] K. LODAYA, R. PARIKH, R. RAMANUJAM, P.S. THIAGARAJAN: A logical study of distributed transition systems, *Inform. and Comput.*, **119** (1995) 91–118.

- [19] A. MAZURKIEWICZ: Basic notions of trace theory, in: J.W. de Bakker, W.-P. de Roever, G. Rozenberg (eds.), *Linear time, branching time and partial order in logics and models for concurrency*, *LNCS* **354** (1989) 285–363.
- [20] R. MILNER: *Communication and Concurrency*, Prentice-Hall, London (1989).
- [21] M. MUKUND, K. NARAYAN KUMAR, M. SOHONI: Keeping track of the latest gossip in message-passing systems, *Proc. Structures in Concurrency Theory (STRICT)*, Berlin 1995, Workshops in Computing Series, Springer-Verlag (1995) 249–263.
- [22] M. MUKUND, M. SOHONI: Keeping track of the latest gossip: Bounded time-stamps suffice, *Proc. FST&TCS '93*, *LNCS* **761** (1993) 388–399.
- [23] M. MUKUND, M. SOHONI: Gossiping, asynchronous automata and Zielonka's theorem, *Report TCS-94-2*, School of Mathematics, SPIC Science Foundation, Madras, India (1994).
- [24] A. MUSCHOLL: On the complementation of Büchi asynchronous cellular automata, *Proc. ICALP 1994*, *LNCS* **820** (1994) 142–153.
- [25] R. PARIKH, R. RAMANUJAM: Distributed processes and the logic of knowledge, *Proc. Logics of Programs '85*, *LNCS* **193** (1985) 256–268.
- [26] S. SAFRA: On the complexity of ω -automata, *Proc. 29th IEEE FOCS* (1988) 319–327.
- [27] P.S. THIAGARAJAN: TrPTL: A trace based extension of linear time temporal logic, *Proc. 9th IEEE LICS* (1994) 438–447.
- [28] W. THOMAS: Automata on infinite objects, in J. van Leeuwen (ed.), *Handbook of Theoretical Computer Science*, Volume B, North-Holland, Amsterdam (1990) 133–191.
- [29] M. VARDI, P. WOLPER: An automata theoretic approach to automatic program verification, *Proc. 1st IEEE LICS* (1986) 332–345.
- [30] W. ZIELONKA: Notes on finite asynchronous automata, *R.A.I.R.O.—Inform. Théor. Appl.*, **21** (1987) 99–135.
- [31] W. ZIELONKA: Safe executions of recognizable trace languages, *Proc. Logical Foundations of Computer Science*, *LNCS* **363** (1989) 278–289.