

Correctness of Mutual Exclusion Algorithms

Neelam Soundarajan and Ten-Hwang Lai
Computer and Information Science
Ohio State University, Columbus, OH 43210
e-mail: {neelam, lai}@cis.ohio-state.edu

Abstract

One way of characterizing non-token-based mutual exclusion (m.e.) algorithms is in terms of the underlying *information structure*. The information structure for a given m.e. algorithm specifies which particular processes interact with which other processes before entering their critical sections, and which processes they interact with when leaving their critical sections. By focusing on the information structure of an m.e. algorithm, we can compare different algorithms by comparing the respective information structures. Further, we can characterize the correctness of an m.e. algorithm in terms of conditions that its information structure must satisfy. In this paper, we propose a necessary and sufficient condition for ensuring the correctness of the m.e. algorithm and show that this condition is superior to the currently accepted condition.

1 Introduction

Mutual exclusion (m.e.) is a central problem in distributed systems. Over the years, numerous algorithms have been proposed for solving this problem. These algorithms may be classified [4, 8] as token-based and non-token-based. Our focus in this paper is on non-token-based algorithms. In these algorithms [1, 3, 5, 7], each process keeps some information about the status of the system such as which processes currently are waiting to enter their critical sections. When a process wishes to enter its critical section, it sends ‘*request*’ messages to a certain set of processes; it then waits for ‘*grant*’ messages from a certain set of processes granting it permission, at which point it enters its critical section; and when it leaves the critical section, it sends ‘*release*’ messages to a certain set of processes. Information about the set of processes to which each process sends requests to or sends release messages to can be considered as the *information structure* [6] associated with this algorithm. Different non-token-based m.e. algorithms can be considered as special cases of a generalized algorithm, the differences between the different algorithms being captured entirely in the differences between the corresponding information structures.

Our goal in this paper is to consider necessary and sufficient conditions that the information structure of an m.e. algorithm must satisfy in order ensure the correctness of the algorithm. The main contributions of the paper are: First it shows that the standard condition [6, 4, 8] that is accepted in the literature to be the necessary and sufficient condition to guarantee correctness of the m.e. algorithm is, in fact, not necessary. This is proved by considering some simple and correct m.e. algorithms and showing that they do not satisfy the standard condition. Second, it proposes

a new condition and proves that this condition is both necessary and sufficient.

The paper is organized as follows: In Section 2, we specify the system model, describe the information structure associated with an m.e. algorithm, and consider the generalized m.e. algorithm. In Section 3, we present the currently accepted condition and show that it is not necessary for the correctness of the algorithm. In Section 4, we propose a new condition and prove that it is both sufficient and necessary. In the final section, we summarize the work and consider pointers for future work.

2 System model, Information structure, and Generalized M.E. algorithm

2.1 System model

Our system model is a fairly typical one. We assume that the system consists of N processes, each with a unique *process id*. Processes communicate with other processes via asynchronous message passing. Communication is assumed to be reliable; message transit time is arbitrary but finite. Processes exchange three types of messages: a *request* message, which we will denote RQ, indicates that the sending process wishes to enter its critical section (c.s.) and is requesting permission to do so from the receiving process; a *grant* message, denoted GR, means the sending process is granting permission to the receiver to enter the c.s; and a *release* message, RL, indicates that the sending process is ‘releasing’ the c.s., i.e., is leaving its c.s.

2.2 Information structure

The *information structure* consists of two sets for each process P_i ; the *request set* R_i consisting of the id’s of those processes to which P_i sends an RQ message when it wishes to enter its c.s; and the *inform set* I_i consisting of the id’s of the processes to which P_i sends an RL message when it leaves its c.s. It is also useful to define the *status set* S_i as the set $\{j | i \in I_j\}$. This information structure is the one proposed by Sanders [6].

2.3 Generalized m.e. algorithm

The generalized algorithm works as follows: Each process P_i has a boolean variable Free_i which is *true* if according to the information that P_i has, no process is in its critical section and *false* otherwise. When a process P_i wishes to enter its c.s., it sends a request message to each process P_j whose identity j is in its request set R_i ; it then waits for GR messages from each of these processes before entering its critical section. When P_i leaves its critical section, it sends a RL message to every process whose id is in its inform set, I_i .

When a process P_i receives a request, it either immediately grants permission to the requesting process by sending a GR message, or adds the request to its collection of pending requests, depending on whether Free_i is *true* or *false* at the time the request was received. If P_i sends a GR message to another process P_j , then it sets Free_i to *false* if $j \in S_i$, i.e., $i \in I_j$.

When P_i receives a RL message, it does the following: Sets Free_i to *true*. Next, if its collection of

pending requests is non-empty, it picks the highest priority request from the collection and sends a GR message to the requesting process P_j , removes the request from its collection, and if $j \in S_i$, sets Free_i to false; it repeats this until its collection is empty or Free_i is false. A standard time-stamp based notion of priority [2], with ties being resolved by an ordering among process id's, is used. The time-stamp and process id are part of each request.

There is a special case corresponding to $i \in R_i$. In this case, when P_i decides to enter its critical section, before sending its requests to other processes in R_i , either it grants permission to itself if Free_i is true and sets Free_i to false if $i \in I_i$, or it adds the request to its collection of pending requests if Free_i is false. Similarly, if $i \in I_i$, when P_i leaves its critical section, before sending RL messages to the other processes in I_i , P_i does the following: sets Free_i to true and proceeds to grant the highest priority request(s) from its collection of pending requests, as described in the last paragraph.

The generalized algorithm above does not necessarily implement mutual exclusion correctly for arbitrary values of R_i and I_i . For example, if we defined each R_i to be just $\{i\}$, i.e., each process had to just get permission from itself before entering its c.s., m.e. would clearly be violated. At the other extreme, if we defined each R_i and I_i to contain the id's of *all* the processes in the system, m.e. would indeed be guaranteed but far more messages than necessary would be exchanged. It is therefore important to have a condition that is both necessary and sufficient to ensure that m.e. is implemented correctly.

3 Standard Condition for Correctness

In this section we first describe the condition that is accepted in the literature as the necessary and sufficient condition for correctness of the generalized m.e. algorithm. We will show that while this condition is sufficient, it is not necessary; indeed, that many simple and correct algorithms do not satisfy the condition.

Sanders [6] presents the following result: If $(i \in I_i)$ for all i , then the following is a necessary and sufficient condition for the correctness of the m.e. algorithm:

$$[(\forall i. I_i \subseteq R_i) \wedge (\forall i, j. ((I_i \cap I_j \neq \Phi) \vee (j \in R_i \wedge i \in R_j)))] \quad (\text{A})$$

The sufficiency of (A) may be seen as follows: Consider two processes P_i, P_j . Suppose $(I_i \cap I_j \neq \Phi)$ is satisfied, say, $(k \in I_i \cap I_j)$. Then, since $(I_i \subseteq R_i)$ and $(I_j \subseteq R_j)$, P_i and P_j will each seek permission from P_k before entering its c.s.; P_k will, since $(k \in I_i \cap I_j)$, set its Free_k to *false* once it grants permission to either of them; hence P_k will *not* grant permission to the other until it receives a release message from the first.

The argument in the case of the other alternative, $(j \in R_i \wedge i \in R_j)$, is a bit more involved. Suppose P_i and P_j decide to enter their c.s.'s at about the same time. Each will grant itself permission to do so, set its Free to false (since $i \in I_i, j \in I_j$) then receive the request from the other; hence neither will grant permission to the other¹. What if, say, P_j decides to enter its c.s. somewhat before P_i does so, and sends its request to P_i ? In that case, P_i will grant permission to P_j and will *not* set Free_i to *false* (since $j \notin I_i$). Next, P_i decides to enter its c.s., and P_i will grant its own request; but P_j will not grant P_i 's request, hence m.e. will be ensured.

Consider now the necessity of (A). Suppose for a particular i, j , (A)'s second conjunct is not

¹There is a danger of deadlock in this case. That can be handled in a standard manner with a slightly more involved approach. Here we will only concern ourselves with the m.e. issue.

satisfied. Then we can create a scenario in which m.e. is violated, as follows. Suppose $i \notin R_j$. Let P_i first decide to enter its c.s.; it gets permission from itself to enter the c.s., then sends requests to all processes in $(R_i - I_i)$; all of them will grant it permission; moreover, since these processes are not in I_i , their *Free* variables will remain *true*. *Free_i* will be *false* but that will not matter since $i \notin R_j$, so P_j will not seek its permission in the next step; note also that *Free_j* will be *true* although P_i obtained permission from P_j (since $j \notin I_i$). Next, P_j decides to enter its c.s.; it gets permission from itself, then sends requests to all processes in $(R_j - I_j)$; all of them will grant it permission since their *Free* variables are *true*; moreover, their *Free* variables will remain *true* at this point. Finally P_i can request permission from all processes in $(I_i - \{i\})$ and P_j from all processes in $(I_j - \{j\})$. All of these will grant permissions and P_i and P_j will both enter their c.s., thereby violating m.e.

The arguments above for the sufficiency of (A) and the necessity of the second conjunct of (A) are based on the ones in Sanders [6]. Sanders further claims that the first conjunct of (A) is also necessary without providing any detailed justification. However, as we will see next, there are correct m.e. algorithms that do not satisfy this clause so it is not a necessary condition. Further, as we see below, the condition $(i \in I_i)$ that the above result imposes, is not satisfied by many correct m.e. algorithms.

The intuition behind the requirement $(i \in I_i)$ is that, in conjunction with the clause $(I_i \subseteq R_i)$, it represents the idea that it is reasonable for a process wishing to enter its c.s. to get permission from itself (in addition to whichever other processes are in R_i) and similarly to inform itself (in addition to whichever other processes are in I_i) when it leaves its c.s. to help ensure m.e. But this intuition is not well founded. For example, consider a centralized m.e. algorithm in which a single process, say P_0 , makes *all* decisions about when each process may enter its c.s. In other words, $R_i = I_i = \{0\}$ for all i . The requirement $(i \in I_i)$ is not satisfied by this algorithm (except for $i = 0$). But it clearly implements m.e. correctly since no process can enter its c.s. without first sending a request to, and getting permission from, P_0 ; and when P_0 grants permission to a process P_i , it will set *Free₀* to *false* and will not grant permission to any other process to enter its critical section until its *Free₀* becomes *true* which in turn will happen only when P_i finally leaves its c.s. and sends a RL message to P_0 . Thus the requirement $(i \in I_i)$ is not appropriate in general.

But even if $(i \in I_i)$ is satisfied, the first conjunct of (A) need not be. Consider the following:

$$I_i = \{i, 0, 1\}; R_i = \{0\}$$

This structure meets the requirement $i \in I_i$, but does not satisfy the first conjunct of (A). But the corresponding algorithm is correct: Since $0 \in R_i$, before P_i enters its c.s., it will request permission from P_0 . If *Free₀* is *true*, P_0 will grant permission, and since $0 \in I_i$, will set *Free₀* to *false*. Now if P_j wishes to enter its critical section, it too will send a request to P_0 (since $0 \in R_j$); P_0 will not grant permission since *Free₀* is *false*. This situation will change only when P_i leaves its critical section and sends a release message to P_0 (which it will since $0 \in I_i$); at that point, P_0 will set *Free₀* to *true*, and send a grant message to another process, possibly P_j , from which it has received a RQ message. This ensures m.e. Note that P_i will also send RL messages to itself and to P_1 when it leaves its c.s., but those messages will not have any effect because those processes do not get any requests so their collection of pending requests is always empty. Thus, the clause $(I_i \subseteq R_i)$ is not necessary for the correctness of the m.e. algorithm.

4 Necessary and Sufficient Condition

We cannot simply omit $(I_i \subseteq R_i)$, because then we would face the following problem. Suppose $k \in (R_i \cap I_i \cap R_j \cap I_j)$. We would expect P_k to ensure mutual exclusion between the critical sections of P_i and P_j . But there is a problem if there is some i' such that $k \in (I_{i'} - R_{i'})$. In this case, mutual exclusion can be violated as follows: First, $P_{i'}$ sends requests to all processes in $R_{i'}$ (which does not include k), gets grants from them, and enters its c.s. Next, P_i sends a request to P_k which grants the request and sets Free_k to *false*. Next, $P_{i'}$ leaves its c.s., and sends release messages to all processes in $I_{i'}$, including P_k ; P_k , mistaking this for a release message from P_i , sets Free_k to *true*. Next, P_j decides to enter its c.s., and sends a request to P_k ; since Free_k is *true* when P_k receives this request, it grants it (and sets Free_k to *false*). At this stage, both P_i and P_j have received grants from P_k and in the absence of some other process ensuring mutual exclusion between them, each will enter its c.s.

Thus the question is, how do we make sure that for each pair of processes, we have some process that ensures m.e. between them, without imposing a complete ban on such “bad” release messages from “third-party” processes such as $P_{i'}$, i.e., without requiring $(I_i \subseteq R_i)$ for all i ? Let us first introduce some further notation. Define $B_i = I_i - R_i$ and $BB = \cup_i B_i$. We will call processes whose indices are in B_i “bad” processes, these being the processes that may receive spurious release messages from P_i ; BB is the set of *all* bad processes. Let NN be the set of all processes. Let $GG = NN - BB$; thus GG is the set of all “good” processes. We can now state our main result:

Theorem: The necessary and sufficient condition for the correctness of the generalized m.e. algorithm is:

$$\forall i, j. \{ \{ (R_i \cap R_j \cap I_i \cap I_j \cap GG) \neq \Phi \} \vee \{ (i \in (R_i \cap R_j \cap I_i \cap GG)) \wedge (j \in (R_i \cap R_j \cap I_j \cap GG)) \} \} \quad (B) \quad \square$$

In other words, the necessary and sufficient condition for the correctness of the m.e. algorithm is that, for each pair i, j , one of the following must be true:

- (b1). Either there exists k that is in each of R_i, R_j, I_i, I_j , and GG ;
- (b2). Or $[(i \text{ is in each of } R_i, R_j, I_i, GG) \text{ and } (j \text{ is in each of } R_i, R_j, I_j, GG)]$

A special case is when each B_i is empty, i.e., $(I_i \subseteq R_i)$; in other words, there are no bad processes that receive spurious release messages. Then GG will be the set of all processes and $\cap GG$ may be omitted from each clause of (B) .

The intuition behind this condition is as follows: In order to ensure mutual exclusion between P_i and P_j , one of the following must be satisfied:

- b1. There must be a process P_k that sends grants to both P_i and P_j , and keeps track of these grants, i.e., k in $(R_i \cap R_j \cap I_i \cap I_j)$; in addition this k must be a good process, so that it is not misled by spurious release messages from third-party processes. Or,
- b2. P_i and P_j must each be responsible for sending grants to both P_i and P_j ; and each must keep track of grants to itself. In addition, both P_i and P_j must be good so they are not misled by spurious releases.

Lemma: If $k \in (R_i \cap I_i \cap GG)$, and if P_k sends a grant message GR to P_i , then Free_k will not become *true* again until P_i enters its c.s., completes it, leaves the c.s., sends a release message to P_k , and P_k processes it.

Proof: From generalized m.e. algorithm described in Section 2.3, it is clear that Free_k will be set to *true* only when P_k receives a release message RL from some process, say, $P_{i'}$. We will show that

$P_{i'}$ must in fact be the same as P_i . Now the only time that $P_{i'}$ sends a RL is when it leaves its c.s. and at that point, it sends RL messages to all processes in $l_{i'}$. Hence, if P_k receives such a message, k must be in $l_{i'}$ and since $k \in GG$, this implies $k \in R_{i'}$. Before it can leave its c.s., $P_{i'}$ must, of course, have *entered* it, and it can do so only after receiving grant messages from all processes in $R_{i'}$ including P_k . P_k could not have sent its GR message *before* it sent the GR message to P_i because if P_k had done so, it would have set $Free_k$ to *false*, and would not then have sent a GR message to P_i (until $Free_k$ was *true* which in turn would happen only upon receiving, by inductive assumption, the release message from $P_{i'}$). P_k could not have sent the GR to $P_{i'}$ *after* sending its GR to P_i because at that point $Free_k$ would be *false*. The only possibility then is that $P_{i'}$ must in fact be the same as P_i , not some other process. \square

4.1 Proof of sufficiency

Next we will show that if either (b1) or (b2) is satisfied, the algorithm will implement m.e. correctly, thereby proving the sufficiency of the condition (B).

Consider a particular pair i, j of process indices. Suppose for this pair, that (b1) is satisfied. Suppose k is in R_i, R_j, l_i, l_j , and GG . Suppose P_i has got into its c.s. That means P_k must have sent P_i a GR message. When it did that, P_k must have set $Free_k$ to *false* since k is in l_i . Now in order for P_j to enter its c.s., it must receive grant messages from all processes in R_j , including P_k . But P_k will not send a GR to P_j (or to any other process) until it receives a release message from some process. Since $k \in GG$, by our lemma, this release message can only come from P_i , i.e., only when P_i leaves its c.s. Hence P_j cannot get into its critical section while P_i is still in its critical section.

Suppose now (b2) is satisfied for this i, j pair. Suppose P_i has already got its c.s. before P_j decides to enter its c.s. Since $i \in l_i$, $Free_i$ will be set to *false* and will not become *true* until P_i gets a release message; and since $i \in GG$, this will not happen until P_i gets out of its c.s. and “sends itself a release message”, i.e., sets $Free_i$ to *true*. Therefore, given that $i \in R_j$ so P_j needs a GR message from P_i before entering its c.s., P_j will not enter its c.s. until P_i leaves its. We can argue similarly if P_j is already in the c.s. before P_i decides to enter its c.s.

What if P_i and P_j decide to enter their c.s.’s at about the same time? One request will have a higher priority, say P_i ’s request. P_i will process this request before that of P_j since, according to the algorithm in Section 2.3, if $i \in R_i$, as soon as P_i decides to enter its c.s., it will either grant its own request or add it to the pending requests depending on whether $Free_i$ is *true* or *false*. And in the latter case, it will grant its own request before granting P_j ’s since the former has higher priority. So P_i will first grant its own request, set $Free_i$ to *false* and the rest of the argument from the last para will now apply to show that P_j cannot enter its c.s. until P_i leaves its. So mutual exclusion is ensured. \square

Note: Consider again the argument that P_i will process its own request first if that request has higher priority. What if P_i received P_j ’s request *just before* deciding to make its own request? This is not possible because if it has received P_j ’s request, P_i will choose a higher time stamp for its own request, so its request will have a lower priority than that of P_j . So if P_i ’s request has a higher priority, it must be because P_i has not yet received P_j ’s request at the time P_i decides to make its request. Of course if P_j ’s request is higher priority, P_i may receive that request before or after it processes its own; but P_j would have processed its own request before receiving P_i ’s and we can go through the same argument exchanging the roles of i and j to show that mutual exclusion is ensured.

4.2 Proof of necessity:

Suppose for a given i, j pair, that neither (b1) nor (b2) is satisfied. We will then show that it is possible to construct a scenario under which m.e. is violated, thereby proving that (B) is a necessary condition.

First we note the following: $R_i = (R_i - I_i) \cup (R_i \cap I_i)$ (by set theory).

Further, $(R_i \cap I_i) = (R_i \cap I_i \cap BB) \cup (R_i \cap I_i \cap GG)$ (by set theory, given $BB \cup GG = \text{set of all processes}$).

And, $(R_i \cap I_i \cap GG) = [((R_i \cap I_i \cap GG) - \{i\}) \cup ((R_i \cap I_i \cap GG) \cap \{i\})]$ (set theory).

Thus R_i is the union of the four disjoint sets $(R_i - I_i)$, $(R_i \cap I_i \cap BB)$, $((R_i \cap I_i \cap GG) - \{i\})$, and $((R_i \cap I_i \cap GG) \cap \{i\})$. R_j can similarly be written as the union of four disjoint sets. The reason for expressing R_i and R_j in this fashion is that it will help in the construction of the m.e.-violating scenario.

We will describe the construction of the scenario in a series of steps. In each step P_i and P_j will send requests to, and get grants from, the processes in the different sets that make up R_i and R_j respectively. We assume that at the start of the scenario no process is in its critical section, that the Free variables of all the processes are *true*, and that no process has any pending requests.

1. P_i decides to enter its c.s., sends RQ messages to, and gets GRs from all the processes in $(R_i - I_i)$. P_j also decides to enter its c.s., sends RQs to, and gets GRs from all the processes in $(R_j - I_j)$. Note that there is no danger of any of these processes not sending GR messages to P_i or P_j no matter in what order the requests from P_i and P_j reach them because first, all the Free variables are initially *true* and second, since these processes are not in I_i/I_j , they leave their Free variables *true* when they send their GR messages. (We will see in step 4 that either P_i 's or P_j 's requests may have to be postponed until the other's requests have been sent to certain processes and corresponding grants received. The other steps will be unaffected by this consideration.)
2. Next we might consider having P_i send RQs to all the processes in $(R_i \cap I_i \cap BB)$. But this will lead to trouble if there is some k that is in $(R_i \cap I_i \cap BB)$ that is also in $(R_j \cap I_j)$; because in this case, P_k would have set Free_k to *false* when it sends a GR message to P_i , and so will not send a GR when P_j makes its request and we would not be able to achieve our goal of violating m.e.

To take care of this, we will make use of the fact that $k \in B$ and proceed as follows: Suppose k' is such that $k \in (I_{k'} - R_{k'})$; such a k' must exist because $k \in B$. Before P_i (or P_j) sends its RQs, we will start by having $P_{k'}$ send RQs to all processes in $R_{k'}$; these processes will send GRs to $P_{k'}$ (because all Free variables are initially *true*); next, P_i will send a RQ message to P_k ; P_k will respond with a GR (since $k \notin R_{k'}$, so Free_k is *true* before P_k receives P_i 's request); and P_k will set Free_k to *false*; next, we will have $P_{k'}$ get out of its c.s. and send RL messages to all processes in $I_{k'}$, including P_k ; when P_k receives this message, it will set Free_k to *true*. At this point, all Free variables are *true* and P_i has received a GR from P_k . We repeat this process for each $k \in (R_i \cap I_i \cap BB)$; and at the end of this, P_i would have received GR messages from all the processes in $(R_i \cap I_i \cap BB)$, and all their Free variables will still be *true*. Next we go through the same procedure with P_j for all processes in $(R_j \cap I_j \cap BB)$.

There is one special situation: What if the k' considered in the last paragraph turns out to be j ? In other words, suppose k is in $(R_i \cap I_i \cap BB \cap (I_j - R_j))$ (and there is no other k' such that $k \in (R_i \cap I_i \cap BB \cap (I_{k'} - R_{k'}))$). In this case, there is no need to require P_j to do

any of the things we required of $P_{k'}$. The point is that the reason for getting $P_{k'}$ to enter its c.s. and exit it was to make sure that Free_k is *true* after it has sent a GR message to P_i since we may want P_j to be able to get a GR message from P_k . But here k is in $(I_j - R_j)$ which means that P_j does not need permission from P_k to enter its c.s. So there is no need to ensure that Free_k is *true* after P_k has sent the GR message to P_i . Indeed, in general, if $k \in (R_i - R_j)$, we do not need to worry about the value of Free_k after it has sent the GR message to P_i ; similarly if $k \in (R_j - R_i)$, we do not need to worry about the value of Free_k after it has sent the GR message to P_j . Note also that if this situation occurs for P_i , i.e., $k \in (R_i \cap I_i \cap \text{BB} \cap (I_j - R_j))$, then it cannot simulatenously occur for P_j , i.e., we cannot for the same k , have $k \in (R_j \cap I_j \cap \text{BB} \cap (I_i - R_i))$ (since the intersection of these two sets is empty).

3. P_i still needs to get grants from $(R_i \cap I_i \cap \text{GG})$ and P_j needs grants from $(R_j \cap I_j \cap \text{GG})$.
 P_i now sends requests to all processes in $(R_i \cap I_i \cap \text{GG}) - \{i\}$ and receives GR messages from them. And P_j sends requests to all processes in $(R_j \cap I_j \cap \text{GG}) - \{j\}$ and receives GR messages from them. Note that none of the processes in $(R_i \cap I_i \cap \text{GG}) - \{i\}$ can be in $(R_j \cap I_j \cap \text{GG}) - \{j\}$ (because, otherwise, (b1) will be satisfied). Similarly, none of the processes in $(R_j \cap I_j \cap \text{GG}) - \{j\}$ can be in $(R_i \cap I_i \cap \text{GG}) - \{i\}$. So these two sets of requests/grants can proceed (even simultaneously) without difficulty. Note also that a process k that is in $(R_i \cap I_i \cap \text{GG}) - \{i\}$ may well be in $(R_j - I_j)$; but P_j has already received grant messages from such processes, so there is no problem in P_i now receiving a GR from P_k and Free_k being set to *false*. Similarly if $k \in ((R_j \cap I_j \cap \text{GG}) - \{j\}) \cap (R_i - I_i)$, P_i has already obtained its grant from P_k (in step 1), so P_j can get its grant from P_k without problem.
4. Next, suppose i is in $(R_i \cap I_i \cap \text{GG})$; i may also be in $(R_j - I_j)$ (but not $(R_j \cap I_j)$ because then (b1) would be satisfied). In this case, if P_i were to start sending its requests before P_j does, we will have a problem because P_i would grant its own request immediately, and set Free_i to *false*; then it would not grant P_j 's request when that arrives. We cannot get around this by assuming that P_i 's request to itself is postponed while its requests to processes in $(R_i - I_i)$ etc. are sent out and grants received because, as noted before, if $i \in R_i$ then P_i considers its own request to enter the c.s. as soon as it decides to enter the c.s. So in this case, we must have P_j send its RQ message to P_i (as part of step 1 for P_j since $i \in (R_j - I_j)$), P_i sends the GR to P_j ; Free_i remains *true* since $i \in (R_j - I_j)$. Now P_i can decide to enter its c.s.; it will then grant itself permission to enter the c.s., set Free_i to *false*, then start sending out the other requests as described in earlier steps.

Similar considerations apply if $j \in ((R_j \cap I_j \cap \text{GG}) \cap (R_i - I_i))$; i.e., in this case we have to be sure to have P_i send its request to and receive the grant from P_j before P_j decides to enter its c.s. What if both of these situations apply? That cannot happen because then (b2) would be satisfied.

What if $i \in (R_i \cap R_j \cap I_j \cap \text{GG})$ and $j \in (R_i \cap R_j \cap I_i \cap \text{GG})$ are both satisfied? Note first that in this case, i must not be in I_i and j must not be in I_j (else (b1) or both (b1) and (b2) will be satisfied). In this case, P_i will get its grant from itself in the first step (since $i \in (R_i - I_i)$) and similarly P_j will get its grant from itself in the first step. Following that, Free_i and Free_j will still be *true*, and P_i and P_j can get grants from each other in step 3 without any difficulties.

In summary, we expressed each of R_i and R_j as unions of four disjoint sets each, described a scenario in which in a sequence of four steps, one corresponding to each of the four sets making up R_i (and similarly R_j), P_i gets grants from all the processes in R_i and P_j gets grants from all the processes

in R_j . At this point, both processes can enter their respective critical sections since they have all the grant messages they need, and mutual exclusion will be violated. This shows that condition (B) is a necessary condition for ensuring mutual exclusion.

5 Discussion

The importance of mutual exclusion was, of course, recognized from the earliest days of concurrency. Over the years, many non-token-based m.e. algorithms have been proposed. The correctness of these algorithms is often quite difficult to establish because of the complex interactions between the processes, with each process getting permission from a specific set of processes before entering its critical section and informing another set of processes when leaving it, etc. By treating the different algorithms as special cases of a general algorithm, with the differences between the individual algorithms being captured in terms of the differences between their corresponding information structure, we reduce the task of establishing correctness of an m.e. algorithm to checking whether its information structure satisfies a particular condition.

But in order for this approach to be useful, the condition we check for must be both necessary and sufficient. Sufficiency is needed because otherwise there would be no guarantee of correctness of the corresponding algorithm even if its information structure satisfied the condition. Necessity is needed because without it, the information structure of correct m.e. algorithms may not satisfy the condition. In this paper, we showed that the condition that is accepted in the literature as necessary and sufficient is in fact not necessary; proposed a new condition for the information structure; and showed that the proposed condition is necessary and sufficient.

We conclude with a pointer to future work. Throughout the paper, and in the literature on information structures, the information structure is assumed to be *static*. One important extension to consider would be the possibility of a *dynamic* information structure [7], i.e., one that changes as the system executes, perhaps as the relative frequencies with which the different processes need to access their critical sections change. Analyzing these algorithms by direct operational reasoning, for example via the construction of scenarios, is likely to be even more difficult than in the case of algorithms with static structures; hence having a necessary and sufficient condition to guarantee correctness would be essential. One important point to note in such algorithms is that even if in the steady state they satisfy the requirement $I_i \subseteq R_i$ so that in the steady state the information structure would satisfy the condition of theorem 1, *during* the time that the information structure is changing, we cannot expect this requirement to be met. Hence we believe that conditions similar to that of theorem 2 might have to be satisfied during this time. We hope to address this question in our future work.

References

- [1] O. Carvalho and G. Roucairol. On mutual exclusion in computer networks. *Comm. ACM*, 26:3–5, 1983.
- [2] L. Lamport. Time, clocks, and ordering of events in a distributed system. *Comm. ACM*, 21(7):558–564, 1978.

- [3] M. Maekawa. A \sqrt{n} algorithm for mutual exclusion in decentralised systems. *ACM Trans. on Compu. Sys.*, 3:145–159, 1985.
- [4] M. Raynal. A simple taxonomy for distributed mutual exclusion algorithms. *Operating Sys. Rev.*, 25:47–49, 1991.
- [5] G. Ricart and A.K. Agrawala. An optimal algorithm for mutual exclusion in computer networks. *Comm. ACM*, 24:9–17, 1981.
- [6] B. Sanders. Information structure of distributed mutual exclusion algorithms. *ACM Trans. on Computer Systems*, 5(3):284–299, 1987.
- [7] M. Singhal. A dynamic information structure mutual exclusion algorithm. *IEEE Trans. on Par. and Dist. Sys.*, 3(1):121–125, 1992.
- [8] M. Singhal. A taxonomy of distributed mutual exclusion. *J. Par. and Dist. Computing*, 18:94–101, 1993.