On the Complexity of Buffer Allocation in Message Passing Systems

Alex Brodsky [a], Jan Bækgaard Pedersen [b], Alan Wagner [c]

[a] *Department of Computer Science,*
University of Toronto,
10 King's College Road,
Toronto, Ontario, M5S 3G4,
Canada

[b] *School of Computer Science,*
University of Nevada, Las Vegas,
4505 Maryland Parkway,
Las Vegas, Nevada, 89154-4019,
United States of America

[c] *Department of Computer Science,*
University of British Columbia,
201-2366 Main Mall,
Vancouver, British Columbia, V6T 1Z4,
Canada

Abstract

Message passing programs commonly use message buffers to avoid unnecessary synchronizations and to improve performance by overlapping communication with computation. Unfortunately, using buffers can introduce portability problems and can lead to deadlock problems on systems without a sufficient number of message buffers.

We explore a variety of problems related to buffer allocation for safe and efficient execution of message passing programs. We show that determining the minimum number of message buffers or verifying that each process has a sufficient number of message buffers are intractable problems. However, we give a polynomial time algorithm to determine the minimum number of message buffers needed to ensure that no send operation is unnecessarily delayed due to lack of message buffers. We extend these results to several different buffering schemes, which in some cases make the problems tractable.

Key words: Message passing systems, Buffer allocation, Complexity, Parallel and distributed programming

Email addresses: abrodsky@cs.utoronto.ca (Alex Brodsky), matt@cs.unlv.edu

1 Introduction

In the last decade MPI [8] and PVM [14] have become the de facto standards for message passing programs. PVM and libraries for MPI provide a degree of portability for message passing programs. Unfortunately, whether the standard mode blocking communication primitives (`MPI_Send` and `MPI_Recv`) behave asynchronously or not, is implementation dependent. Thus, an application that executes successfully on one system, may deadlock on another due to lack of message buffers. As a result, the MPI standard defines a program to be *safe* and *portable* [24, Section 9.2] if no buffering is required for the program to successfully execute, which is equivalent to requiring that the program only uses synchronous message passing; a send (receive) call cannot complete before the corresponding receive (send) call begins. Such programs can be ported to machines with differing amounts of message buffers.

If a program takes advantage of the asynchronous nature of a specific implementation of these primitives, the program is no longer portable, and thus no longer safe. For example, the cyclic shift example in [24, Section 9.2.1] where each process first performs a send to its right neighbour followed by a receive from its left neighbour is not a safe program. While MPI provides a variety of alternative modes, the most natural one remains asynchronous send and synchronous receive since often the sender can proceed whereas the receiver often must wait for a message to continue computing. Furthermore, using nonstandard mode communication primitives or synchronous communication may require the application programmers to manage the resources themselves and/or to restructure the application. Additionally, use of these modes require an understanding of the underlying buffer requirements of the application.

As an alternative, we investigate the possibility of determining the minimum number of message buffers necessary to ensure a deadlock-free execution of an application that uses the standard mode blocking primitives. We define a more general notion of a safe program (with respect to message buffer availability) that includes both the MPI notion of a safe program (when no buffers are available) and the k-buffer correctness (when each process has k buffers), introduced by Bruck *et al.* [2]. Our notion of a safe program is with respect to minimizing the total number of buffers for the entire system. Namely, for given a program, do k buffers, distributed among the processes (not necessarily evenly), suffice to prevent deadlock.

We investigate the complexity of determining a minimum value of k for programs using asynchronous buffered communication with a static communication pattern and a bounded message size. We consider the following three problems:

- The **Buffer Allocation Problem (BAP)**, which is the problem of determining

(Jan Bækgaard Pedersen), `wagner@cs.ubc.ca` (Alan Wagner).

the minimum number of buffers required to ensure deadlock free execution.
- The **Buffer Sufficiency Problem (BSP)**, which is to determine whether a given buffer assignment is sufficient to avoid deadlock.
- The **Delay-free Buffer Allocation Problem (DBAP)**, which is to determine the minimum number of buffers needed to ensure that every send operation can be performed asynchronously, that is, ensure that a buffer is always available when a send occurs.

The complexity of these questions also depends on the buffer scheme used by the underlying system. We consider four buffering schemes. In the first three schemes the buffers are (1) pre-allocated on the receive side only, (2) the send side only, or (3) mixed and pre-allocated on both sides. Finally, we also consider a scheme that pre-allocates buffers on a per channel basis, where each communication channel can buffer a fixed number of messages.

We show that the Buffer Allocation Problem is intractable under all four buffer allocation schemes. The Buffer Sufficiency Problem is intractable for the receive side buffer and for the mixed buffer allocation schemes, tractable for the channel scheme and conjectured tractable for sender side buffers. Finally, the Delay-free Buffer Allocation Problem is tractable for all buffer placement schemes, except the mixed buffering scheme.

2 The Model

A multiprocess system S is a fixed number of simultaneously executing independent asynchronous processes, where each process has access to local memory only, executes at its own speed, and has no access to a global clock. These processes compute by interspersing local computation and point-to-point communication, where communication between two processes occurs exclusively using message passing.

Processes communicate by invoking *send* and *receive* primitives. The *sender* invokes the send primitive with a reference to a fixed-size message stored in local memory and a destination identifier. A send blocks until the space occupied by the message can be reused safely, that is, until the message is injected into the message passing system. The *receiver* invokes the receive primitive with a source identifier and a reference to fixed-size location in local memory, where the message is to be delivered. A receive blocks until the message has been delivered. The send and receive primitives correspond to the *standard* `MPI_Send` and `MPI_Recv` operations defined in [20, Section 3.4]. As in PVM and MPI, messages are ordered and cannot overtake each other [20, Section 3.5]. We do not consider send/receive operations with wildcards, such as the `ANYTAG` or `ANYSOURCE` specifiers in MPI.

A **message buffer** (or simply a **buffer**) is a fixed-size memory segment used to hold

a single message within the message passing system. The process of temporarily utilizing a buffer for a message that has not yet been delivered is referred to as **buffering**. When a message has been copied into a buffer or delivered to the receiving process, the corresponding send will return. Furthermore, when a message has been delivered from a buffer, the buffer can be reused.

Buffers are maintained by the underlying communication system, and are preallocated at the start of each process in accordance with one of the four following schemes:

Receiver Side Buffering
A fixed number of buffers are allocated by each receiver's underlying message passing system. These buffers may only be used to store messages that have not yet been delivered to the receiver.

Sender Side Buffering
A fixed number of buffers are allocated by each sender's underlying message passing system. These buffers may only be used to store messages that have not yet been delivered to the receiver's message passing system, that is, messages that have not yet been placed on the network.

Sender and Receiver Side Buffering
A fixed number of buffers are allocated by each process' underlying message passing system. These buffers may used interchangeably as either receiver side buffers or as sender side buffers.

Channel Buffering
A **channel** is a process pair with one sender and one receiver. A fixed number of buffers are allocated per channel. Since these channels are unidirectional, the buffers can be physically allocated by either the sender's or the receiver's message passing system.

Messages are buffered in their entirety; this includes both the envelope and message contents. We only consider fixed-sized messages. The earlier work of Burns and Daoud on Guaranteed Envelope Resources [3] uses an envelope queue which corresponds to our receive-side buffering scheme. As noted in [3], many MPI implementations aggressively send short messages, which may cause buffer problems when using a large number of nodes or when space is limited, such as on network interface cards.

Finally, we only consider programs that are *repeatable* [6,7] when executed in an unrestricted environment, that is, programs with *static communication patterns*. While this narrows the class of programs we consider, the class of applications with static communication patterns is still considerable.

3 Definitions

In this paper, a send (receive) is **synchronous** if it cannot complete before the corresponding receive (send) begins. This definition is comparable to the one given by Magee and Kramer [19].

Let S be a multiprocess system with n processes, $P_1, \ldots, P_n$, with $E_i \geq 0$, communication events occurring in process P_i; a **communication event** is either a send or a receive. We say that the system **deadlocks** if at least one process blocks during a communication event and cannot make progress; that is, the program blocks on a send or a receive. Such deadlocks can occur if there is an insufficient number of buffers, resulting in some asynchronous sends becoming synchronous. We say that system S is **unsafe** if a deadlock can occur due to an insufficient number of available buffers; if S is not unsafe, then S is said to be **safe**. Note that this definition of safety differs from that of MPI [20, Section 3.5].

A **buffer assignment** for a multiprocess system S with n processes is an n-tuple $B = (b_1, b_2, \ldots, b_n)$ of nonnegative integers representing the number of buffers that are allocated by each process. Similarly, a **channel buffer assignment** is a q-tuple $B = (b_1, b_2, \ldots, b_q)$, $q = \binom{n}{2}$, representing the number of buffers that each channel in the system allocates.

We say that buffer is **utilized** when it contains a message and a buffer is **free** when it does not contain a message and can be utilized by the underlying message passing system. Initially all buffers are free. A buffer is utilized when a message that has been injected into the system cannot be delivered, and becomes free once the message it contains is delivered to the receiving process

We say that a send is **delay-free** if a free buffer is always available whenever a send is invoked.

3.1 Problem Definitions

Buffers take up memory that may be needed by the application, Therefore, in an ideal scenario, as few buffers as possible should be allocated. However, allocating too few buffers results in an unsafe system.

Two natural decision problems arise from this optimization problem:

BUFFER ALLOCATION PROBLEM (BAP)
INSTANCE: A multiprocess system S and a nonnegative integer k.
QUESTION: Does a buffer assignment B exist such that S is safe and $\sum b_i \leq k$?

In order to solve this problem we need to solve a simpler one:

BUFFER SUFFICIENCY PROBLEM (BSP)
INSTANCE: A multiprocess system S and a buffer assignment B.
QUESTION: Is a buffer assignment B sufficient to make S safe?

Additionally, we can require that no process in system S should never be delayed by a send call:

DELAY-FREE BUFFER ALLOCATION PROBLEM (DBAP)
INSTANCE: A multiprocess system S and a nonnegative integer k.
QUESTION: Does there exists a buffer assignment B, such that all sends in S are delay-free, and $\sum b_i \leq k$.

4 Related Work

4.1 Communication

The model defined in the previous section has been studied extensively in the past. For example, in [4] each process is referred to as an *A-computation* [4], and is equivalent to a process with three different events, such as the one defined by Lamport [18]: *send events*, *receive events* and *internal events*.

MPI [20, Section 3.4] defines a number of different send primitives. These primitives are divided into two categories: The *blocking* and the *nonblocking*. Each category can be divided into four modes: *synchronous* (`MPI_Ssend` / `MPI_Issend`), *buffered* (`MPI_Bsend` / `MPI_Ibsend`), *ready* (`MPI_Rsend` / `MPI_Irsend`), and *standard* (`MPI_Send` / `MPI_Isend`). In the blocking category, the synchronous mode is buffer-independent, because it either returns an error or blocks until the message is delivered. The buffered mode allows the user to take over responsibility for allocating buffers, requiring knowledge about the underlying buffer requirements of the program. The ready mode requires that the corresponding receive be already posted, thus its semantics fall outside the scope of the formalism.

The standard mode, which corresponds to the mode of the send primitive of our model, is either synchronous or has system-managed buffers. In our model, if no buffers are available, the send call is synchronous, but if buffers are available, the send call can return before the message is delivered to the receiver. MPI's blocking receive (`MPI_Recv`) blocks until the message is delivered to the receiver, exactly like the receive primitive in our model. Although we discuss the primitives in our model with respect to MPI, many of the same issues exist for PVM [14] as well.

In the terminology of Cypher and Leu [6,7], the send primitive in our model is for-

mally defined as a *POST-SEND* immediately followed by a *WAIT-FOR-BUFFER-RELEASE* and receive primitive as a *POST-RECEIVE* immediately followed by a *WAIT-FOR-RECEIVE-TO-BE-MATCHED*.

4.2 Safety

According to the MPI standard [20, Section 3.5], "a program is *safe* if no message buffering is required for the program to complete. One can replace all sends in such program with synchronous sends, and the program will still run correctly." Furthermore, according to [24, Section 9.2], safe is synonymous with *portable*. This notion of a portable program corresponds to our notion of a safe program with respect to a buffer assignment of $B = (0, \ldots, 0)$, that is, 0 buffers per process. Such programs require no buffers in order to execute safely, and are thus *buffer independent*.

Determining whether a system is buffer independent was investigated in [6,7]. In our model, the interesting systems are buffer-dependent, and require (an unknown number of) buffers to avoid deadlock.

In a similar vein, Bruck *et al.* [2] introduced the definition of k-buffer correctness to describe programs that complete without deadlock in a message passing environment with k buffers per process. This corresponds to determining safety with respect to a buffer assignment of the form $B = (k, k, \ldots, k)$, that is, k buffers per process. A corollary of one of our results (see Corollary 7.7) is that determining whether a program is k-buffer correct (with respect to point-to-point communication) is intractable. Similarly, Burns and Daoud [3] introduced *guaranteed envelope resources* into LAM [12], a public domain version of MPI. Guaranteed envelope resources was used in LAM to reserve a guaranteed number of message header slots on the receiver side.

4.3 Motivation and Related Approaches

More recently in modern clusters, greater overlap of computation and communication is possible by off-loading communication onto the network interface cards (NIC). Unfortunately, most NICs have orders of magnitude less memory than the average host, which makes message buffers a limited resource. Thus, programs that use asynchronous message passing, and that execute correctly otherwise, might deadlock when executing on a system where parts of the message passing system have been off-loaded to the NIC. These issues have been investigated in several papers [8,9,11,17].

To determine the minimum number of buffers, the execution of a system can be

modeled using a (coloured) Petri net [16]. In order to determine whether the system can reach a state of deadlock, the Petri net occurrence graph [15] is constructed, and a search for dead markings is performed. Methods exist to reduce the size of these graphs, but in many cases the size remains exponential in the size of the original Petri net.

Variations of these problems have been investigated by the operations research community [1,22,23]. In these models, events or products are buffered between various stations in the production process, however, the arrival of these events is governed by probability distributions, which are specified a priori. In our model, since processes are asynchronous, the time for a message to arrive is nondeterministic; that is, a message may take an arbitrarily long time to arrive and a process may take an arbitrarily long time to perform a send or a receive.

5 The Graph Based Framework

We model systems by using communication graphs, and executions of systems by colouring games on these graphs. Communication graphs can be derived from execution traces of a program. The following subsection formally defines the graph based framework used throughout this paper.

A **communication graph** of S is a directed acyclic graph $G = G(S) = (V,A)$ where the set of vertices $V = \{v_{i,c} \mid 1 \leq i \leq n, 0 \leq c \leq E_i\}$ corresponds to events in the system. An event is either the start of a process or a communication event: vertex $v_{i,0}$ represents the **start** of process i, vertex $v_{i,c}$, $1 \leq c \leq E_i$, represents either a **send** or a **receive** event.

The arc set A consists of two disjoint arc sets: the computation arc set P and the communication arc set C. An arc, $(v_{i,c}, v_{i,c+1}) \in P$, $0 \leq c < E_i$, represents a computation within process i and an arc $(v_{i,c}, v_{j,d}) \in C$ represents a communication between different processes, i and j, where $v_{i,c}$ is a send vertex, and $v_{j,d}$ is a receive vertex (e.g. Figure 1). Note, the process arcs are drawn without orientation for clarity; they are always oriented downwards. Communication graphs are comparable to the time-space diagrams—without internal events—noted in [18].

The ith **process component** G_i of G is the subgraph $G_i = (V_i, A_i)$ where $V_i = \{v_{i,c} \in V \mid 0 \leq c \leq E_i\}$ and $A_i = \{(v_{i,c}, v_{i,c+1}) \in A \mid 0 \leq c < E_i\}$. The process component corresponds to a process in S. We construct communication graphs by connecting process components with arcs. Hence, it is more intuitive to treat a process component as a chain of send and receive vertices beginning with a start vertex. A channel is represented by a **channel pair** (G_i, G_j) of process components. An example of a communication graph is depicted in Figure 1.

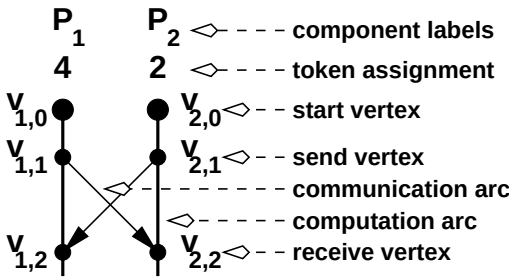

Fig. 1. An example of a communication graph representing two processes. Component P_1 is assigned 4 tokens and component P_2 is assigned 2 tokens.

To model the execution of a system S, we define a colouring game that simulates the execution of the system with respect to $G(S)$.

5.1 Colouring the Communication Graph

Given a communication graph $G(S)$, an execution of a corresponding system S is represented by a colouring game where the goal is to colour all vertices green; a green vertex corresponds to the completion of an event. We use three colours to denote the state of each event in the system: a red vertex indicates that the corresponding event has not yet started, a yellow vertex indicates that the corresponding event has started but not completed, and a green vertex indicates that the corresponding event has completed. Hence, a red vertex must first be coloured yellow before it can be coloured green; this corresponds to a traffic light changing from red, to yellow, to green [1] .

We use a token to represent a message buffer. A buffer assignment of a process (or channel) is represented by a pool of tokens associated with the corresponding process component (respectively, the channel component). An instance of buffer utilization is represented by removing a token from a token pool and placing it on the corresponding communication arc.

The colouring game represents an execution via the following rules. Initially, the start vertices of G are coloured green and all remaining vertices are coloured red; this is called the **initial colouring**.

[1] Naturally, we refer to a European traffic light.

send$\rightarrow$*yel*	A red send vertex may be coloured yellow if the preceding vertex is green—the send is ready.
recv$\rightarrow$*yel*	A red receive vertex may be coloured yellow if the corresponding send vertex is yellow, and the preceding vertex (in the same process component) is green—both the send and the receive are ready.
recv$\overset{\bullet}{\rightarrow}$*yel*	A red receive vertex may be coloured yellow if the corresponding send vertex is yellow, and a token from the corresponding token pool is placed on the incident communication arc—the send is ready and a buffer is utilized.
send$\rightarrow$*grn*	A yellow send vertex may be coloured green if the corresponding receive vertex is coloured yellow—the communication has completed from the sender's perspective.
recv$\rightarrow$*grn*	A yellow receive vertex may be coloured green if both of its preceding vertices are green. If the incident communication arc has a token, the token is returned to its token pool—a receive completes after the send completes.

Buffer utilization is represented by placing a token from the token pool on the selected arc, and colouring the corresponding receive vertex yellow. If no tokens are available, the rule cannot be invoked.

A **token assignment**, also denoted by B, is a list of nonnegative integers, denoting the number of tokens assigned to each token pool; the token assignment is the abstract representation of a buffer assignment. The number of tokens required depends on the number of times that rule *recv*$\overset{\bullet}{\rightarrow}$*yel* can be invoked. If a token pool is empty, this means all buffers are in use. An example of how a token assignment is represented is depicted in Figure 1.

A **colouring** of G, denoted by χ, is a colour assignment to all vertices, which can be obtained by repeatedly applying the colouring rules, starting from the initial colouring. A **colouring sequence** $\Sigma = (\chi_1, \chi_2, \ldots)$ is a sequence of colourings such that each colouring is derived from the preceding one by a single application of a colouring rule. An execution of a multiprocess system S with buffer assignment B is represented by a colouring sequence on $G(S)$. Each transition, from one colouring to the next, within a colouring sequence, corresponds to a change of state of an event in the corresponding execution. Assuming that all events in the system are ordered, there is a correspondence between the colouring sequences on $G(S)$ and the executions of system S. Using the correspondence between colouring sequences on $G(S)$ and executions of system S, we reason about system S by reasoning about colouring sequences on $G(S)$.

We say that a colouring sequence **is complete** if and only if the last colouring in

the sequence comprises only green vertices. We say that a colouring sequence **can complete** if and only if the a colouring sequence can be a suffixed with zero or more colourings such that the last colouring in the sequence comprises only green vertices. A colouring sequence **deadlocks** if and only if the last colouring in the sequence has one or more nongreen vertices and the sequence cannot be extended via the application of the colouring rules. A system S is safe if and only if every colouring sequence on the graph $G(S)$ can complete.

We say that a colouring sequence **delays** if the sequence has a prefix that ends with a colouring containing a yellow send vertex and which cannot be extended by applying rule $recv \overset{\bullet}{\rightarrow} yel$ to the corresponding receive vertex. A colouring sequence is **delay-free** if every prefix of the sequence does not delay; a communication graph G, is delay-free if all colouring sequences on it are also delay-free. If $G(S)$ is delay-free, then no send in S will ever delay during an execution.

5.2 Nondeterminism: An Example of the Colouring Game

The system we are modeling not only lacks a global clock, but the rate of progress of each process can differ and change during the course of the computation. Given this asynchrony, the total order of communication events within the system can differ from execution to execution. Our framework models this by the fact that at many points in a colouring sequence there exists a choice of which vertex to colour next; this choice is a form of nondeterminism in the framework. In fact, there are many communication graphs on which both deadlocking and nondeadlocking colouring sequences exist. A brief example of this phenomenon as well as the colouring game is described below.

Consider a system comprising three processes, whose communication pattern is represented by the left-most communication graph in Figure 2. We shall consider two colouring sequences on this graph: the first, depicted in Figure 2, deadlocks; and the second, depicted in Figure 3, completes. Both sequences begin at the initial colouring, in which the start vertices are green and the remaining vertices are red.

In the first sequence, the send vertex in process component P_1 is coloured yellow ($send \rightarrow yel$), then the corresponding received vertex is coloured yellow ($recv \overset{\bullet}{\rightarrow} yel$). This uses up the single token assigned to process component P_2. The send vertex is then coloured green ($send \rightarrow grn$); however, until the corresponding receive vertex is coloured green, the token cannot be reused. Finally, the two send vertices in the remaining two processes can be coloured yellow ($send \rightarrow yel$). At this point the sequence is deadlocked because, without an available token, neither of the two remaining receive vertices can be coloured yellow until one of the preceding send vertices is coloured green; neither of which can be coloured yellow until one of the succeeding receive vertices is coloured yellow.

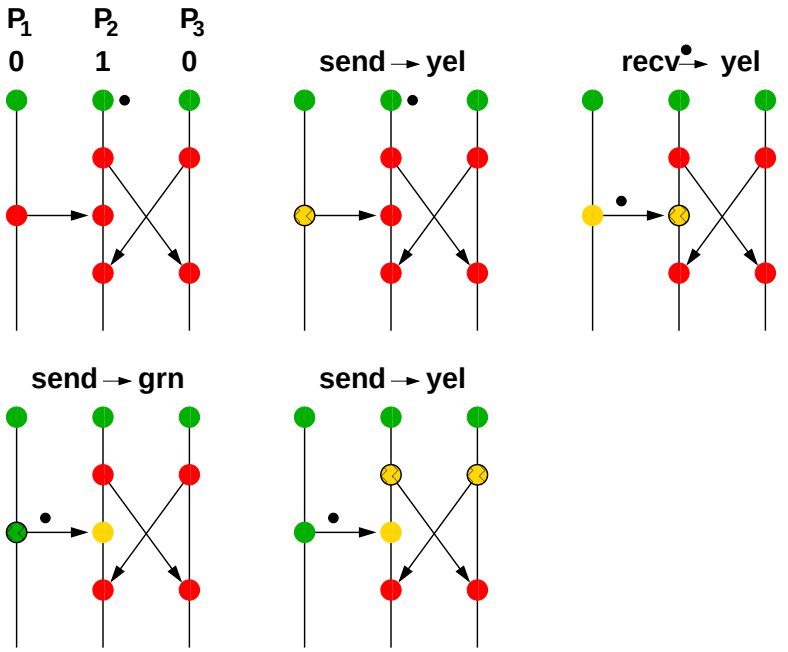

Fig. 2. A colouring sequence that cannot complete.

The second colouring sequence first colours the send vertex in process P_3 yellow ($send \rightarrow yel$), and then colours the corresponding receive vertex yellow by utilizing the token ($recv \overset{\bullet}{\rightarrow} yel$). At this point, this colouring sequence can complete regardless of the order in which remaining vertices are coloured; figure 3 depicts one such colouring sequence. The key point is that the token is utilized to break the deadlock described in the preceding paragraph, and utilizing the token to colour another receive vertex (as in the first sequence) ensures deadlock. Thus, sequence completion, can depend on which vertex is coloured next.

6 Useful Constructions and Lemmas

A **t-ring** is a subgraph of a communication graph $G(S)$, consisting of $t > 1$ process components, such that in each of the t process components there is a send vertex s_{i_j,c_j} and a receive vertex r_{i_j,d_j}, $c_j < d_j$, $1 \leq j \leq t$ such that the arcs $(s_{i_1,c_1}, r_{i_t,d_t})$ and $(s_{i_{j+1},c_{j+1}}, r_{i_j,d_j})$, $1 \leq j < t$ are in A. This definition is equivalent to the definition of a *crown* in [4].

The colouring game defined in the preceding section induces a dependence on the vertices of the communication graph. Namely, the rules induce a partial order according to which the vertices can be coloured. For example, a send vertex may only be coloured yellow if its predecessor is green. Formally, we say that vertex v **depends** on vertex u if vertex v cannot be coloured green until vertex u is coloured yellow. By transitivity, if vertex v depends on vertex u, and vertex w depends on vertex v, then vertex w also depends on vertex u. A communication graph has a **cir-**

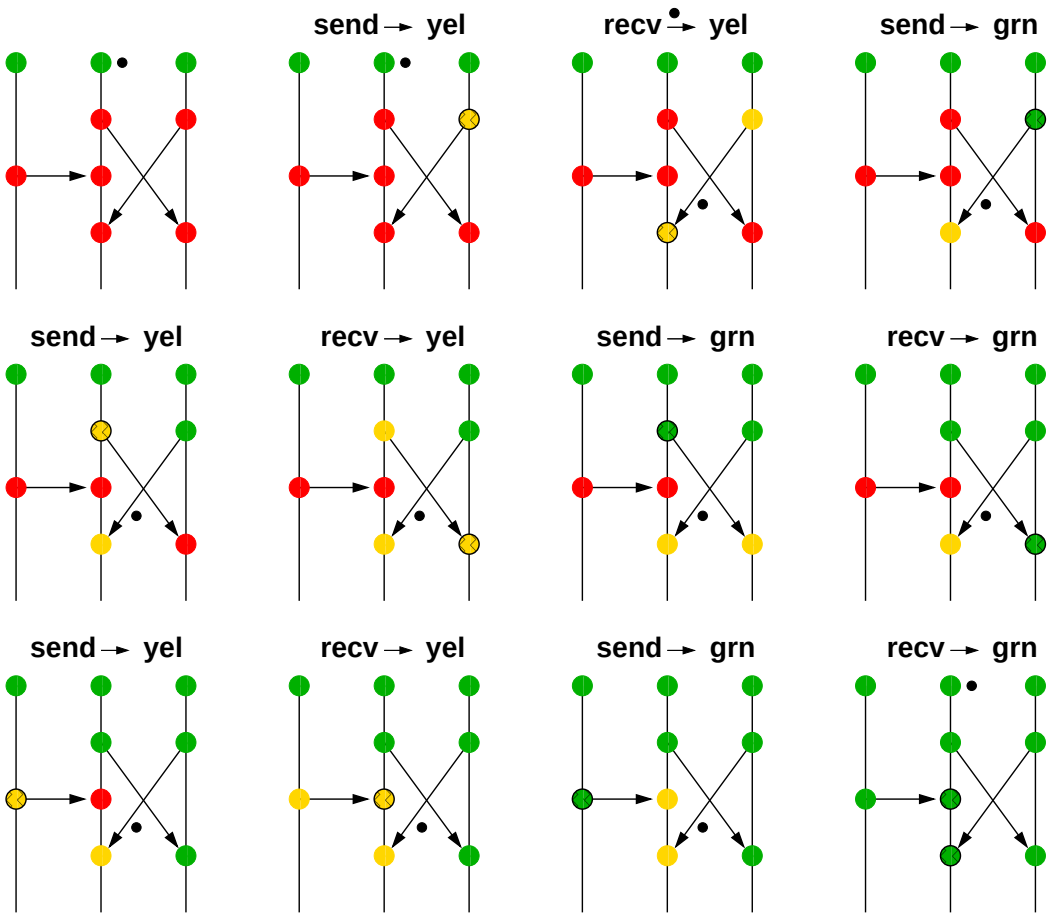

Fig. 3. A colouring sequence that does complete.

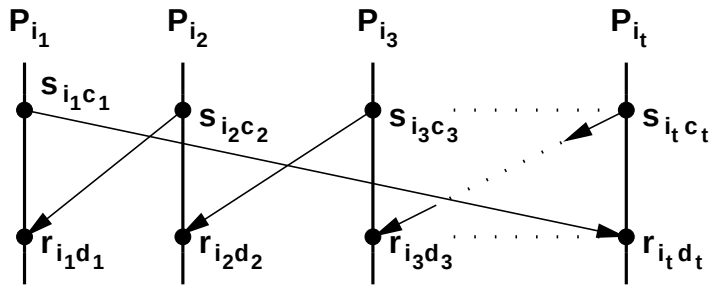

Fig. 4. Dependency cycle in $G(S)$.

cular dependence if it contains four or more vertices, $v_1, v_2, \ldots, v_m \in G(S)$, such that vertex v_i depends on vertex v_{i+1}, $1 \leq i < m$ and vertex v_m depends on vertex v_1.

The t-ring construction, depicted in in Figure 4, contains a circular dependence, which as we shall show, cannot be coloured green without invoking rule $recv \overset{\bullet}{\rightarrow} yel$ at least once. Thus, a t-ring represents a common communication pattern that can lead to deadlock when there is insufficient number of buffers available. This pattern is exhibited by many programs and is explicitly mentioned in [24, Section 9.2.1].

The following lemmas are used throughout our proofs. Lemma 6.1 characterizes the conditions under which a colouring sequence will deadlock. Lemma 6.2 character-

izes conditions under which a single colouring sequence may represent all possible colouring sequences. Finally, Lemma 6.3 characterizes a class of communication graphs on which no colouring sequence will deadlock.

Lemma 6.1 (The t-Ring Lemma) *Let G be a communication graph comprising a single t-ring. Any complete colouring sequence must have at least one application of the rule recv$\stackrel{\bullet}{\rightarrow}$yel.*

Proof: Assume by contradiction that there exists a complete colouring sequence Σ that does not make use of rule *recv*$\stackrel{\bullet}{\rightarrow}$*yel*. Consider the first colouring in Σ where one of the send vertices is green; call the vertex s_i. Let r_j be the corresponding receive vertex. According to rule *send*$\rightarrow$*grn*, the vertex r_j must be yellow. Since rule *recv*$\stackrel{\bullet}{\rightarrow}$*yel* has not been applied, rule *recv*$\rightarrow$*yel* must have been invoked earlier in the sequence. By the definition of a t-ring, the send vertex s_j must be the predecessor of r_j. Since the rule *recv*$\rightarrow$*yel* was applied to r_j, s_j must be green. Hence, there is an earlier colouring in Σ with a green send vertex. This is a contradiction.

In the other direction, if rule *recv*$\stackrel{\bullet}{\rightarrow}$*yel* is invoked on receive vertex r_j, then rule *send*$\rightarrow$*grn* can be invoked on the corresponding send vertex s_j, breaking the circular dependency. ■

Define the **dependency graph** of communication graph $G = (V, A)$ to be $H = (V, E)$ where all process arcs in A are reversed in E and all communication arcs in A are bidirectional in E. Define the **depth** $d(v)$ of a vertex $v \in V$ to be the maximum length path [2] in H from v to a start vertex.

Lemma 6.2 *Let G be a communication graph with a token assignment of size* 0*. For any vertex v in G, if there exists a colouring sequence that colours vertex v green, there does not exist a colouring sequence that deadlocks before colouring v green.*

Proof: Proof by contradiction. Assume that there exist two colouring sequences, such that one colouring sequence colours a vertex green and the other deadlocks and does not colour the vertex green. Let $v \in V$ be such a vertex of minimum depth; that is, all vertices of lesser depth will be coloured green eventually by any colouring sequence on G. In order for a vertex to be coloured green, its component predecessor must be green. Since the depth of the predecessor is less than the depth of v, it can always be coloured green. Furthermore, since a send and its corresponding receive vertex are adjacent to each other, their depths differ by at most 1.

Since v must be a communication vertex, by rules *send*$\rightarrow$*grn* and *recv*$\rightarrow$*grn*, the adjacent communication vertex t must be coloured yellow before v can be coloured green. If vertex t is of a lesser depth than v, then t must be green colourable in all colouring sequences; hence, v must also be green colourable. If t is at the same

[2] Each vertex may appear at most once in a path.

depth as v, then its component predecessor is at a lesser depth and must be green colourable, hence t is yellow colourable, and v is green colourable. If t is at a greater depth than v, the component predecessor of t, say u, is at the same or a lesser depth than t. If the latter, then u is green colourable and t is yellow colourable, otherwise, we apply the same argument to u first. Since there is no path from u to v in H—because $d(u) \leq d(v)$—we need only recurse a finite number of times. ■

Lemma 6.3 *If G is a communication graph whose dependency graph is acyclic, then no colouring sequence on G will deadlock.*

Proof: Proof by contradiction. Assume that a colouring sequence deadlocks on G. Let v be a vertex of minimum depth that cannot be coloured green. If v is a send (receive) vertex, let u be the corresponding receive (send) vertex. Let vertex t be the component predecessor of vertex u and let vertex w be the component predecessor of vertex v. Since the dependency graph is acyclic, it follows from the definition of depth that the depths of both t and w are less than the depths of both u and v. Hence, both t and w may be coloured green based on our minimality assumption. However, then both u and v may be coloured green; this is a contradiction! ■

7 Buffer Allocation in Systems with Receive Side Buffers

In systems with receive side buffers, messages are buffered only by the receiver. Buffers are utilized by the receiving process when a message arrives, but cannot be delivered, and are freed when the message is delivered to the application. Analogously, when colouring a receive vertex of the corresponding communication graph, only a token belonging to the same process component may be used. We call this the **receive side utilization scheme**.

7.1 The Buffer Allocation Problem

In order to prevent deadlock in distributed applications, the underlying system needs to allocate a sufficient number of buffers. Ideally, it should be the minimum number required. Unfortunately, determining the required number of buffers, such that the system is safe, is intractable.

The corresponding graph-based decision problem, denoted BAP_r, is this: given a communication graph G and a positive integer k, determine if there is a token assignment of size k such that no colouring sequence deadlocks on G. We show that BAP_r is **NP**-hard by a reduction of the well known 3SAT problem [5] to BAP_r. Recall the definition of 3SAT: determine if there exists a satisfying assignment to $\bigwedge_{i=1}^{c}(a_i \vee b_i \vee c_i)$, where a_i, b_i, and c_i are Boolean literals in $\{x_1, \bar{x}_1, x_2, \bar{x}_2, \ldots, x_n, \bar{x}_n\}$.

Theorem 7.1 *The Buffer Allocation Problem (*BAP$_r$*) is* **NP***-hard.*

Proof: Proof by reduction of 3SAT to BAP$_r$. For any 3SAT instance F we construct a corresponding communication graph G such that there exists a token assignment of size n that allows for any colouring sequence to complete on G if and only if the corresponding variable assignment satisfies F.

Let F be an instance of 3SAT with n variables and c clauses; the variables are denoted $x_1, x_2, \ldots, x_n$, and the jth clause is denoted $(a_j \vee b_j \vee c_j)$, where $a_j, b_j, c_j \in \{x_1, \bar{x}_1, \ldots, x_n, \bar{x}_n\}$. The corresponding communication graph G comprises $2n+1$ process components: $2n$ of the components, called **literal** components, are labeled P_{x_i} and $P_{\bar{x}_i}$, $i = 1 \ldots n$, and correspond to the literals of F. The last component—called the **barrier** component—is labeled P_{barrier}. The variables and clauses of F are represented by the widgets that are depicted in Figure 5.

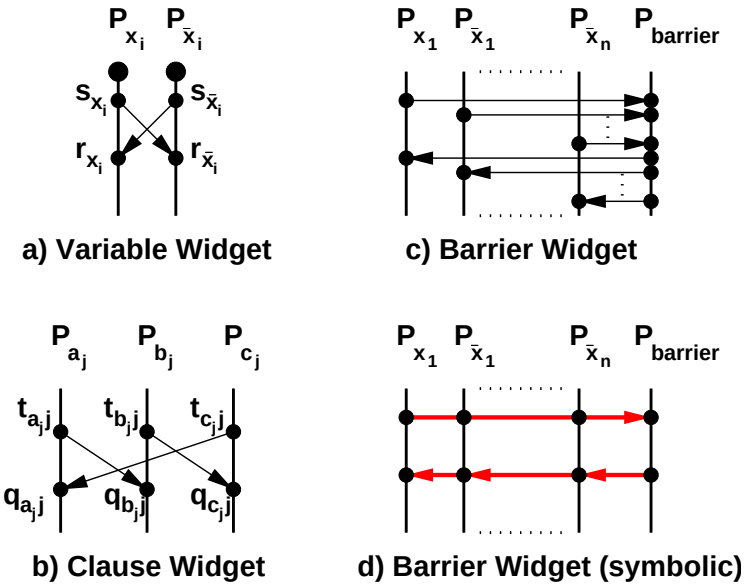

Fig. 5. The three widgets used in construction of G.

Each process component is divided into $c+1$ epochs, where each epoch is a consecutive sequence of zero or more vertices within the component. All epochs are synchronized, that is, the vertices of one epoch must be coloured green before any of the vertices in the next epoch may be coloured. To ensure this we use a barrier component. The jth epoch of the barrier component, $j = 0, \ldots, c$, comprises $2n$ receive vertices, labeled $q_{l,j}$, and $2n$ send vertices, labeled $t_{l,j}$, $l \in \{x_1, \bar{x}_1, \ldots, x_n, \bar{x}_n\}$. At the end of each epoch there is an arc from each of the literal components P_l, $l \in \{x_1, \bar{x}_1, \ldots, x_n, \bar{x}_n\}$, to the barrier component. Each arc emanates from vertex $s_{l,j}$, called a barrier send vertex, and is incident on vertex $q_{l,j}$, where $l \in \{x_1, \bar{x}_1, \ldots, x_n, \bar{x}_n\}$ and $j = 0 \ldots c$. These arcs are followed by arcs emanating from the barrier component to the literal components; the arcs emanate from vertices $t_{l,j}$ and are incident on vertices $r_{l,j}$, called barrier receive vertices. The barrier widget has no cyclic dependencies. Hence, by Lemma 6.3, no colouring sequence will deadlock on a barrier widget.

Epoch 0 fixes a token assignment corresponding to a variable assignment in 3SAT. Each pair of process components, P_{x_i} and $P_{\bar{x}_i}$, $i = 1 \ldots n$, forms a variable widget, which corresponds to a variable. The two process components of a pair share a 2-ring; see Figure 6. By Lemma 6.1, at least one token must be assigned to either process component P_{x_i} or $P_{\bar{x}_i}$ to prevent all colouring sequences from deadlocking on G. Since only n tokens are available, each component pair can be assigned exactly one token. Finally, assigning the token to either process component, P_{x_i} or $P_{\bar{x}_i}$, corresponds to fixing variable x_i to either true or false. Epoch 0 terminates with a barrier send vertex $s_{l_i,0}$, followed by a barrier receive vertex $r_{l_i,0}$, $l_i \in \{x_i, \bar{x}_i\}$.

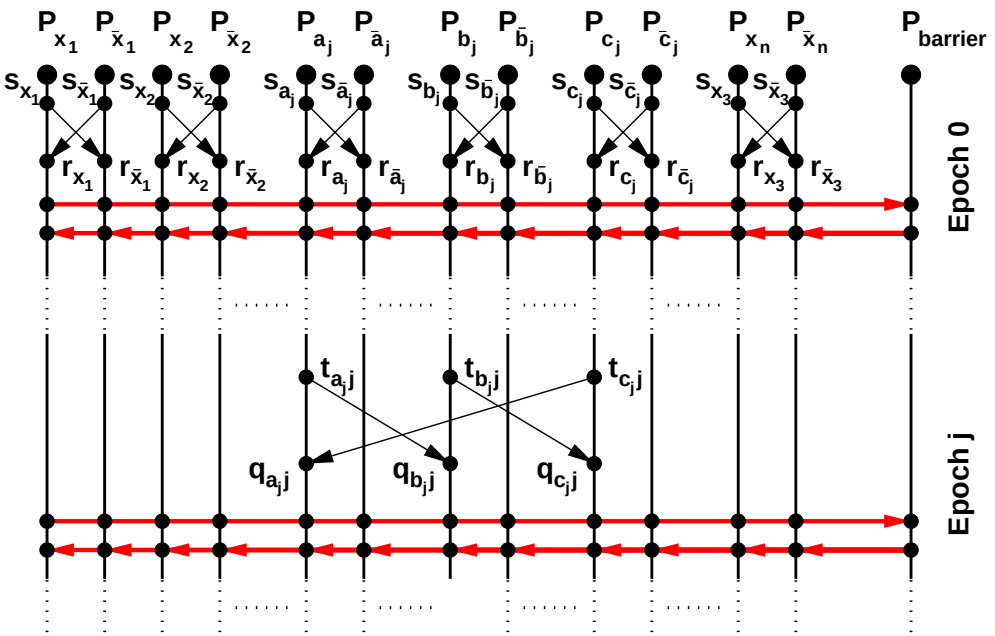

Fig. 6. Construction of G.

Epoch j of each process component corresponds to the jth clause of F. The epoch of a process component P_l, $l \neq a_j, b_j, c_j$—not labeled by a literal of the jth clause—contains only two vertices: the barrier send vertex $s_{l,j}$ and the barrier receive vertex $r_{l,j}$. The three process components, P_{a_j}, P_{b_j}, P_{c_j}, whose labels correspond to the literals in the jth clause share a 3-ring in the jth epoch; see Figure 6. By Lemma 6.1, to avoid deadlock, at least one of the three process components must have a token. If none of the three process components are assigned a token, all literals in the jth clause are false. The epoch is terminated by the barrier send and the barrier receive vertices.

A satisfying assignment on F satisfies at least one literal in every clause. A corresponding token assignment assigns a token to the corresponding process component in each 3-ring—corresponding to the true literal. Hence, by Lemma 6.1 none of the colouring sequences will deadlock on any of the clause widgets and any colouring sequence on G can complete.

For a falsifying assignment of F, there exists at least one clause comprising false literals. The corresponding token assignment fails to assign any tokens to the process

components that share the corresponding 3-ring. Thus, by Lemma 6.1 all colouring sequences will deadlock in that clause widget.

Hence, for a token assignment of size n, any colouring sequence on G can complete if and only if the corresponding assignment satisfies F. Since finding a token assignment of size n such that no colouring sequence on G deadlocks is as hard as finding a satisfying assignment for F, BAP_r is **NP**-hard. ■

7.2 The Buffer Sufficiency Problem

A potentially simpler problem is to verify whether a given buffer assignment is sufficient to prevent deadlock. Formally, given a graph G and a token assignment on G, determine if none of the colouring sequences on G deadlock. This problem turns out to be intractable as well.

We show that BSP_r is **coNP**-complete by a reduction from the TAUTOLOGY problem [13, Page 261] to BSP_r. Given an instance of a formula in disjunctive normal form (DNF), $\bigvee_{i=1}^{t} \bigwedge_{j=1}^{l_i} a_{i,j}$ where $a_{i,j} \in \{x_1, \bar{x}_1, \ldots, x_n, \bar{x}_n\}$, the formula is a tautology if it is satisfied by all assignments. An assignment that falsifies F is a concise proof that the formula is not a tautology. We shall restrict our attention to 3DNF formulas, where each term has three literals: $\bigvee_{i=1}^{t}(a_i \wedge b_i \wedge c_i)$.

Theorem 7.2 *The Buffer Sufficiency Problem (*BSP_r*) is* **coNP***-complete.*

Proof: Let G be a communication graph along with a token assignment. If there exists a deadlocking colouring sequence on G, then the sequence itself is a certificate. The sequence is at most twice the number of vertices in G. Hence, BSP_r is in **coNP**.

Let F be a 3DNF formula with t terms where each term has three literals. For any 3DNF formula F, we construct a communication graph G and fix a token assignment such that there is a colouring sequence on G that deadlocks if and only if the corresponding assignment falsifies F. The construction consists of widgets that correspond to fixing an assignment, a term in the disjunction, the disjunction, and the interconnects between widgets.

Each variable in F is represented by a variable widget comprising three process components that are labeled P_{x_i}, $P_{\bar{x}_i}$, and Q_i. The latter, called the **arbitrator component**, comprises three receive vertices, labeled q_i, r_{x_i}, and $r_{\bar{x}_i}$. The former two process components, called **variable components**, each comprise two send vertices. The first, labeled s_{x_i} ($s_{\bar{x}_i}$), is adjacent to the corresponding receive vertex r_{x_i} ($r_{\bar{x}_i}$) in the arbitrator component. The second, labeled t_{x_i} ($t_{\bar{x}_i}$), is adjacent to receive vertices in widgets called **dispersers**, which are described later. The vertex q_i in the arbitrator component is similarly adjacent to a vertex in a disperser widget. The

corresponding token assignment for each variable widget assigns one token to Q_i and no tokens to the other two components; see Figure 7a.

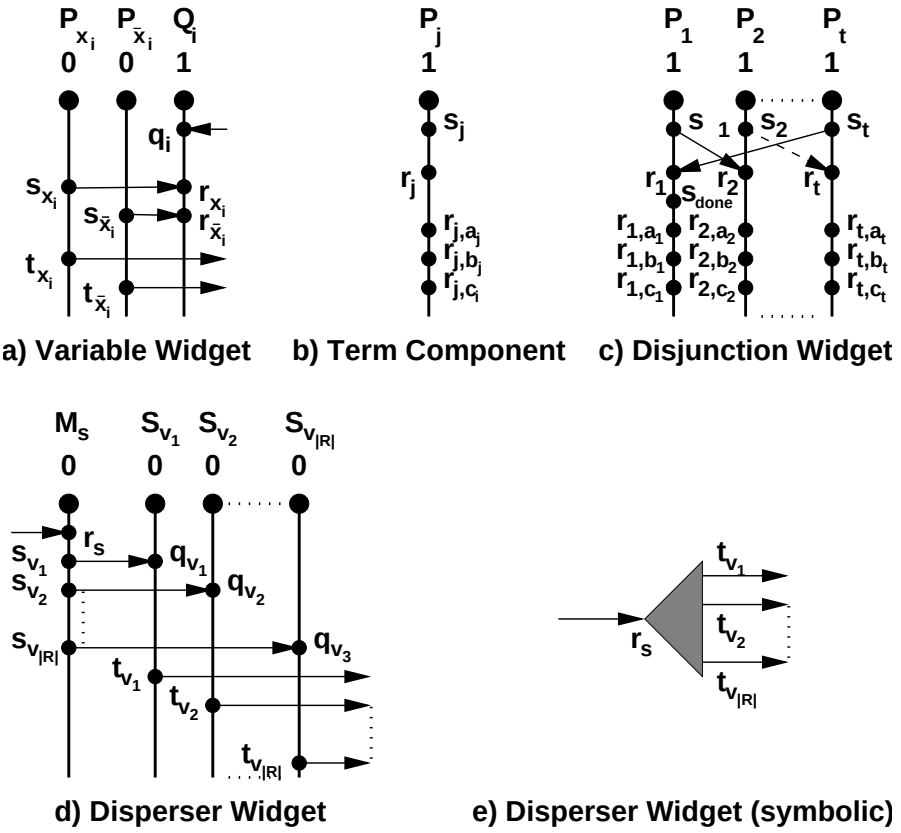

Fig. 7. The widgets used in the construction.

The widget has the following property:

Lemma 7.3 *Let G be a communication graph that contains a variable widget. Any colouring sequence on G may colour exactly one of the two vertices t_{x_i} or $t_{\bar{x}_i}$ yellow before vertex q_i is coloured green.*

Proof: By rule *send*$\rightarrow$*yel*, in order for t_{x_i} ($t_{\bar{x}_i}$) to be coloured yellow, vertex s_{x_i} ($s_{\bar{x}_i}$) must be coloured green. Hence, by rule *send*$\rightarrow$*grn*, vertex r_{x_i} ($r_{\bar{x}_i}$) must first be coloured yellow. Since vertex q_i is not green, vertex r_{x_i} ($r_{\bar{x}_i}$) can only be coloured yellow via rule *recv*$\stackrel{\bullet}{\rightarrow}$*yel*. However, there is only one token assigned to process component Q_i, hence rule *recv*$\stackrel{\bullet}{\rightarrow}$*yel* may only be invoked once. ■

The jth term in the disjunction is represented by a term widget comprising a process component, which is called the **term component** and labeled P_j. The first part of each term component consists of a send vertex s_j and a receive vertex r_j; these vertices are part of a t-ring. In the first term component, P_1, there is an additional send vertex labeled s_{done}, which will be described later. The second part of each term component consists of three receive vertices labeled r_{j,a_j}, r_{j,b_j}, and r_{j,c_j}, where $a_j, b_j, c_j \in \{x_1, \bar{x}_1, \ldots, x_n, \bar{x}_n\}$ correspond to the literals in the jth term; see Figure 7b. These receive vertices are adjacent to send vertices in widgets called

dispersers, which are described later. The term components are used to construct a disjunction widget.

The **disjunction widget** comprises t term components, where the first two vertices, s_j and r_j, are part of a t-ring spanning all t components. Specifically, each send vertex s_j, $j < t$, is adjacent to receive vertex r_{j+1} and vertex s_t is adjacent to receive vertex r_1; see Figure 7c. Each term component is assigned one token. The disjunction widget has the following property.

Lemma 7.4 *Let G be a communication graph that contains a disjunction widget. Any colouring sequence on G can colour r_j, $j \in [1,t]$, green if and only if at least one of r_k, $k \in [1,t]$, is coloured yellow before any r_{k,a_k}, r_{k,b_k}, or r_{k,c_k} are coloured yellow.*

Proof: By Lemma 6.1, vertex r_j can be coloured green, if and only if rule $recv \overset{\bullet}{\rightarrow} yel$ is invoked, colouring one of the receive vertices r_k, $k \in [1,t]$, yellow. The rule may only be invoked when a token is available. Since each term component only has one token assigned and since vertex r_k precedes vertices r_{k,a_k}, r_{k,b_k}, and r_{k,c_k}, a token is available if and only if none of the vertices r_{k,a_k}, r_{k,b_k}, and r_{k,c_k}, are coloured yellow via rule $recv \overset{\bullet}{\rightarrow} yel$, before vertex r_k is coloured yellow. ■

Once vertex r_k, $k \in [1,t]$, is coloured yellow, all r_j, $j = 1 \ldots t$ may be coloured green, and vertex s_{done} may be coloured yellow. We now describe how the widgets are connected together using disperse widgets. Let s be a send vertex and R be a set of receive vertices. An (s,R)-disperser comprises $|R|+1$ process components: one **master** component, labeled M_s, and $|R|$ **slave** components labeled S_r, $r \in R$. The master component comprises one receive vertex labeled r_s, followed by $|R|$ send vertices labeled s_r, $r \in R$. Each send vertex is adjacent to the receive vertex on the corresponding slave component S_r. Each slave component has two vertices: a receive vertex q_r, followed by a send vertex t_r, $r \in R$; see Figure 7d and Figure 7e. The latter vertex is adjacent to the receive vertex r in some other widget. None of the components are assigned any tokens. Since the dependency graph of a disperser is acyclic, the following property follows from Lemma 6.3.

Property 7.5 *Let G be a communication graph containing an (s,R)-disperser. If a colouring sequence colours vertex r_s yellow, then the colouring sequence can be extended to colour all vertices t_r, $r \in R$ yellow.*

Let R_{x_i}, $i = 1 \ldots n$, be the set of receive vertices labeled $r_{j,x_i} \in P_j$, $j \in [1,t]$, and let $R_{\bar{x}_i}$ be similarly defined; recall that a_j, b_j, c_j are simply literal place holders in the vertex labels r_{j,a_j}, r_{j,b_j}, r_{j,c_j}. Hence, a (t_{x_i}, R_{x_i})-disperser connects send vertex $t_{x_i} \in P_{x_i}$ to vertices in R_{x_i}—belonging to the term components. Furthermore, let Q be the set of receive vertices q_i (in the variable widgets), $i = 1 \ldots n$; a (s_{done}, Q)-disperser connects vertex s_{done} to all variable widgets via receive vertices q_i. The construction of G comprises n variable widgets and one disjunction widget, composed of t term

widgets; these are connected together by a (s_{done}, Q)-disperser, and $2n$ (t_a, R_a)-dispersers, where $a \in \{x_1, \bar{x}_1, \ldots, x_n, \bar{x}_n\}$. We claim that there exists a colouring sequence that deadlocks on G if and only if there is a falsifying assignment for formula F, that is, F is not a tautology.

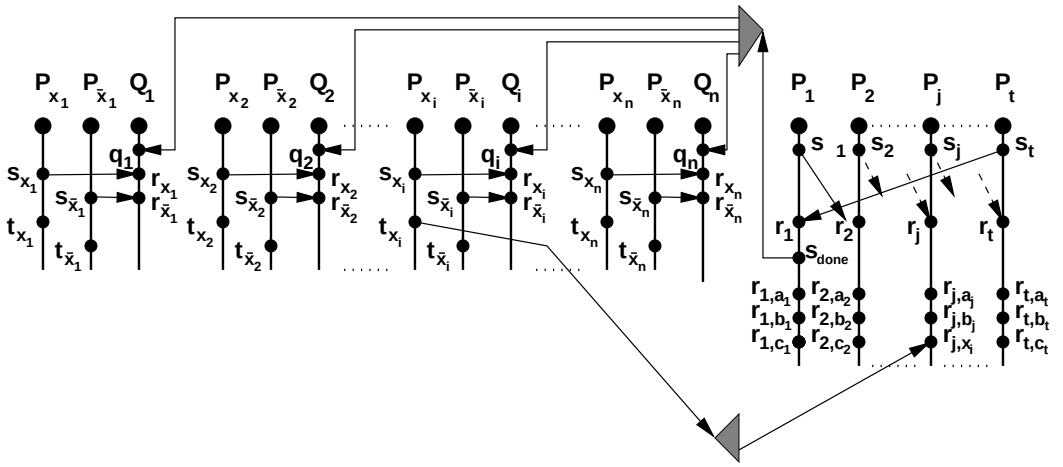

Fig. 8. The construction.

Suppose that F has a falsifying assignment x, that is every term in the disjunction is false because each term has a literal x_i or $\bar{x}_i$, which is false. To construct a colouring sequence on G that deadlocks, we construct a set of vertices U. The first half of the colouring sequence is a maximal colouring sequence involving only the vertices of U. The second half of the sequence may involve all vertices in G. The resulting colouring sequence will always deadlock.

Let $X = \{a \in \{x_1, \bar{x}_1, \ldots, x_n, \bar{x}_n\} \mid a|_x = 0\}$, which is the set of literals that are false, and let $Z = \{s_a \in P_a \mid a \notin X\} \cup \{s_j \mid j = 1, \ldots, t\}$, which contains the set of send vertices from the variable components that are labeled by a true literal and the numbered send vertices in the disjunction widget; the set Z contains the vertices that are not initially (re-)coloured in our construction. Let $U = V \backslash Z$ be the rest of the vertex set.

Consider a colouring sequence involving only vertices in U. By property 7.3 any maximal colouring sequence will colour the vertices t_a yellow (in the arbitrator component), where $a \in X$. Hence, by property 7.5 the vertices t_r (in the dispersers) will be coloured yellow, where $r \in \bigcup_{a \in X} R_a$—the send vertices t_r in the dispersers are adjacent to the receive vertices in R_a. Since x is a falsifying assignment, every term contains a literal, which is falsified by x. Without loss of generality, let a_j denote a literal that is false in the jth term; therefore, process component P_j contains a receive vertex r_{j,a_j}, which is adjacent to the yellow send vertex $t_{r_{j,a_j}}$ (in the disperser). Since none of the vertices of the t-ring (in the disjunction widget) are U—they are still coloured red—the token belonging to component P_j is used to apply rule $recv \xrightarrow{\bullet} yel$ to colour vertex r_{j,a_j} yellow. Since every term has a false literal, the colouring sequence colours a receive vertex r_{j,a_j}, $j = 1 \ldots t$ in every term component P_j.

After the sequence cannot be extended, allow all vertices to be coloured. Since vertices r_{j,a_j} (in the term components), $j = 1 \ldots t$, have been coloured yellow before vertex r_j (in term component P_j), according to property 7.4, the sequence will deadlock.

If a colouring sequence on G deadlocks, according to property 7.4, deadlock occurs only if there is a yellow vertex labeled r_{k,a_k}, r_{k,b_k}, or r_{k,c_k} in each of the term components. Their predecessors—vertices t_l, $l \in \{x_1, \bar{x}_1, \ldots, x_n, \bar{x}_m\}$, in the dispersers—must be green. Since the colouring sequence is maximal, by property 7.3 exactly one of t_{x_i} or $t_{\bar{x}_i}$ is red, thus this corresponds to a valid assignment: setting $x_i = 0$ if t_{x_i} is green, or $x_i = 1$ if $t_{\bar{x}_i}$ is green yields an assignment that falsifies F.

Thus, a colouring sequence on G deadlocks if and only if the corresponding assignment falsifies F. Hence, BSP_r is **coNP**-complete. ■

Therefore, just determining whether a buffer assignment is sufficient is intractable, even one as simple as in the preceding example. Intuitively, the buffers of a process are assigned based on the behaviour of other processes; thus, buffer utilization is not locally decidable. Further, the order in which buffers are assigned is nondeterministic, exploding the search space of possible buffer utilizations. This phenomenon, which our proofs rely on, is what we call *buffer stealing*. For example, in a system corresponding to the variable widget (see Figure 8), the first process to send its message gets the buffer, and the other process remains blocked until the arbitrator performs the receives. This stealing corresponds to fixing a value of a variable. Similarly, the system corresponding to the disjunction widget allocates buffers for each of the term processes. However, if the buffer is stolen in all terms, corresponding to a falsifying assignment, then the system will deadlock within the t-ring.

For completeness, we note the following corollary:

Corollary 7.6 *The Buffer Allocation Problem (*BAP_r*) is in* $\Sigma_2\mathbf{P}$.

Proof: By Theorem 7.2, verifying that a token assignment is sufficient to prevent deadlock (BSP_r) is **coNP**-complete. Since we can nondeterministically guess a sufficient token assignment, the result follows. ■

Another corollary is that determining if a program is k-buffer correct [2], with respect to point-to-point communication, is intractable—even if $k = 1$.

Corollary 7.7 *Determining if a program is* 1*-buffer correct is* **coNP***-complete.*

Proof: Observe that determining if a program is 1-buffer correct is analogous to solving BSP_r with the added requirement that the token assignment be homogeneous, that is, all components must be assigned the same number of tokens. In the reduction described in Theorem 7.2 all process components that are not allocated any tokens are either part of the disperser widget, or do not contain any receive

vertices. Assigning each of these components a token does not alter their function. Thus, the same reduction from TAUTOLOGY to BSP_r works with the homogeneous token assignment requirement. ■

7.3 The Delay-free Buffer Allocation Problem

In addition to the system being safe, we can require that no sending process is ever delayed due to insufficient buffers in the receiving process. The Delay-free Buffer Allocation Problem (DBAP_r) is to determine the minimum number of message buffers needed to achieve a delay-free execution.

Formally, the corresponding decision problem is this: given a communication graph G and an integer k, determine if there exists a token assignment of size k such that all colouring sequences on G are delay-free. Recall that a colouring sequence is delay-free if, whenever a send vertex is coloured yellow, rule $recv \overset{\bullet}{\rightarrow} yel$ may be applied to the corresponding receive vertex.

Let P_i and P_j, $i \neq j$, be two process components. Given two vertices, $v_{i,c}$ and $v_{i,t}$, in P_i, $t > c$, vertex $v_{i,t}$ is **communication dependent** on vertex $v_{i,c}$ if $v_{i,c}$ is the start vertex or if there exists a vertex $v_{j,d} \in P_j$, such that there is a path from $v_{i,c}$ to $v_{j,d}$ and the arc $(v_{j,d}, v_{i,t})$ is in A; see Figure 9. Vertex $v_{i,t}$ is **terminally communication dependent** on vertex $v_{i,c}$ if $v_{i,t}$ is communication dependent on $v_{i,c}$ and is not communication dependent on the vertices $v_{i,l}$, $c < l < t$. The algorithm depicted in Figure 10 computes an optimal token assignment such that all colouring sequences on G are delay-free.

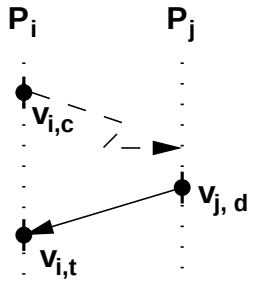

Fig. 9. $v_{i,t}$ is communication dependent on $v_{i,c}$.

Remark 7.8 *In a system corresponding to communication graph G, the time between a message arriving at process i and its receipt corresponds to the interval $I_{i,t}$. Each interval must have a buffer to ensure delay-free sends. Hence, the minimum number of buffers, b_i, is the maximum overlap density over all intervals within process P_i.*

Computing the terminal communication dependencies for G can be done via dynamic programming in $O(|V|n)$ time, where V is the vertex set of G and n is the number of process components. Associate with each vertex $v_{j,d}$ an integer vector

(1) For each receive vertex $v_{i,t}$ determine its terminal communication dependency, vertex $v_{i,c}$, where $t > c$.
(2) Set $I_{i,t} =]c,t]$ to be the interval between vertex $v_{i,c}$ and vertex $v_{i,t}$.
(3) For each process component G_i, compute

$$b_i = \max_j \sum_t [j \in I_{i,t}],$$

which is the maximum overlap density over all intervals $I_{i,t}$, where $[condition] = 1$ if *condition* is true and 0 otherwise.
(4) $B = \{b_1, b_2, ..., b_n\}$ is the optimal delay-free token assignment.

Fig. 10. Algorithm for computing an optimal delay-free buffer assignment.

$a_{j,d}$ of size n, where $a_{j,d}[i] = c$ means that there exists a path from vertex $v_{i,c}$ to vertex $v_{j,d}$; initially, $a_{j,d}[k] = -1$ for $k \neq d$ and $a_{j,d}[j] = d$, otherwise. The vector $a_{j,d}$, is computed by taking the elementwise maximums of the vectors associated with the parent vertices of vertex $v_{j,d}$. This is simply a depth-first traversal of G that begins at the last vertex of each component and traverses the arcs in the opposite direction. Since the number of arcs is bounded by $3|V|/2$ and the pairwise comparison takes n steps, the traversal takes $O(|V|n)$ time.

Next, computing the $O(|V|)$ intervals, $I_{i,t}$, requires one table lookup per interval. To compute the maximum overlap density we sort the intervals and perform a sweep, keeping track of the current and maximum overlap density; this takes $O(|V|\log|V|)$ time. Thus, the total complexity is $O(|V|n + |V|\log|V|)$ time. In the worst case, where $n \approx |V|$, this algorithm is quadratic. However, in practice n is usually fixed, in which case the $|V|\log|V|$ term dominates.

7.3.1 Proof of Correctness of the Delay-free Buffer Allocation Algorithm

Lemma 7.9 *Let G be a communication graph. For all vertices $v_{i,c}, v_{j,d} \in G$; if $v_{j,d}$ is a send vertex and there exists a path from the vertex $v_{i,c}$ to vertex $v_{j,d}$, then vertex $v_{j,d}$ cannot be coloured yellow until vertex $v_{i,c}$ is coloured green.*

Proof: By rule *send→yel*, the predecessor of $v_{j,d}$ must first be coloured green before $v_{j,d}$ can be coloured yellow. Since rules *send→grn*, and *recv→grn* imply that the predecessors of a green vertex must be green, the result follows. ■

Corollary 7.10 *Let G, $v_{i,c}$, and $v_{j,d}$ be as in Lemma 7.9 and let $v_{i,t}$ be the receive vertex corresponding to the send vertex $v_{j,d}$. Rule recv$\stackrel{\bullet}{\rightarrow}$yel will never be applied to vertex $v_{i,t}$ before vertex $v_{i,c}$ is coloured green.*

The preceding corollary implies that a token, which is needed to colour the receive vertex $v_{i,t}$ yellow, need not be available until the vertex on which $v_{i,t}$ is terminally communication dependent is coloured green. Hence, it is sufficient to ensure to-

ken availability just before colouring the respective send vertex green; this is also necessary.

Theorem 7.11 *Given G, let $v_{i,c}$ be a send vertex and $v_{i,t}$ be a receive vertex that is terminally communication dependent on vertex $v_{i,c}$. A token for the application of rule recv$\overset{\bullet}{\rightarrow}$yel on arc $(v_{j,d}, v_{i,t})$ must be available as soon as vertex $v_{i,c}$ is coloured green.*

Proof: Let $v_{j,d}$ be the send vertex corresponding to the receive vertex $v_{i,t}$ and let $Q = \{v_{i,q} \mid c < q < t\}$ be the set of vertices that are predecessors of $v_{i,t}$, but on which $v_{i,t}$ is not communication dependent.

Since $v_{i,t}$ is not communication dependent on the vertices in Q, we can construct a colouring sequence on G that fixes the vertices in Q to be red, and colours vertex $v_{j,d}$ yellow, making the application of rule *recv*$\overset{\bullet}{\rightarrow}$*yel* possible in the next step. Since no progress is made in the ith process component after colouring vertex $v_{i,c}$ green, the state of the associated token pool does not change until the application of rule *recv*$\overset{\bullet}{\rightarrow}$*yel* to vertex $v_{i,t}$. Hence, when vertex $v_{i,c}$ is coloured green, the token pool must have a token destined for arc $(v_{j,d}, v_{i,t})$. ■

Thus, if a receive vertex r is terminally communication dependent on a send vertex s, then it is necessary and sufficient that a token, which is used to apply rule *recv*$\overset{\bullet}{\rightarrow}$*yel* to receive vertex r, must be available as soon as the send vertex s is coloured green; the start vertex may be thought of as a special send vertex. Since the interval corresponding to r begins when s is coloured green, and ends when r is coloured green, a token must be available for the *recv*$\overset{\bullet}{\rightarrow}$*yel* rule, which can occur during this interval. Computing the maximum overlap density of intervals yields the required number of tokens.

7.3.2 Example Use of the DBAP$_r$ *Algorithm*

To demonstrate the DBAP$_r$ algorithm we have implemented it, and analyzed the pipe-and-roll parallel matrix multiplication algorithm [10]. The program has one control process and a number of worker processes arranged in a 2 dimensional mesh. We ran the DBAP$_r$ algorithm on meshes of size 2×2, 3×3 and 4×4. The communication graph for the smallest example, comprising four workers ordered in a 2×2 mesh, is depicted in Figure 11. The corresponding optimal buffer assignment is listed in the second column of Table 1.

In this example, process 0 is the control process and processes 1 through 4 are the workers. The control process needs 4 buffers and the workers each need 3 to execute without delaying. The results obtained when executing the DBAP$_r$ algorithm on a 3×3 worker system is 9 buffers for the control process and between 4 and 5 buffers for the worker processes. For the 4×4 system the numbers are 16 for the control

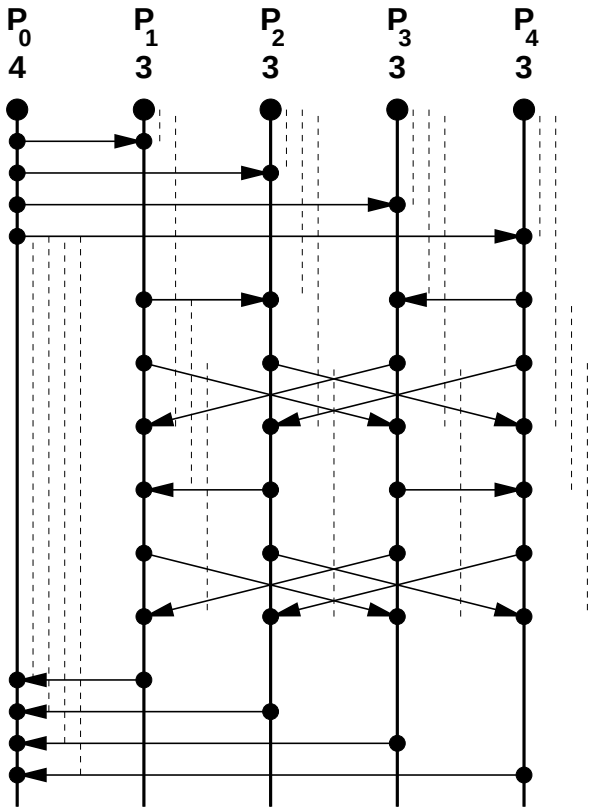

Fig. 11. The communication system for a 2×2 worker process mesh. Note, the dashed lines represent the intervals $I_{i,t}$.

process and between 5 and 7 buffers for the workers.

Proc.		Overlap density of intervals between $v_{i,t-1}$ and $v_{i,t}$							
i	b_i	$j=1$	$j=2$	$j=3$	$j=4$	$j=5$	$j=6$	$j=7$	$j=8$
0	4	0	0	0	0	4	3	2	1
1	3	2	1	2	3	2	1	1	0
2	3	3	2	1	2	1	1	1	0
3	3	3	2	1	2	1	1	1	0
4	3	2	1	2	3	2	1	1	0

Table 1. The result of running the DBAP$_r$ algorithm on the 2×2 worker example.

7.3.3 Approximating BAP$_r$ *with* DBAP$_r$

The DBAP$_r$ algorithm is useful for determining a token assignment that prevents deadlock, that is, approximating BAP$_r$. Since a delay-free colouring sequence does not deadlock, a token assignment determined by the DBAP$_r$ algorithm ensures that the graph is deadlock free. However, the token assignment may be far from optimal. A simple example of this phenomenon is a two process component graph comprised of n arcs emanating from the first component and incident on the second. Such a

graph requires zero tokens to avoid deadlock, but requires n tokens to be delay-free. Thus, the aforementioned token assignment may entail many more tokens than required.

8 Buffer Allocation in Systems with Send Side Buffers

In this section we consider the second of the four buffer placement strategies: send side buffers. Buffers are now allocated on the sending process side if the receive is not ready to accept the message. Correspondingly, the token pool used when applying rule $recv \overset{\bullet}{\rightarrow} yel$ to the receive vertex of arc (s,r) belongs to the process component containing the send vertex s. We call this the **send side utilization scheme**.

The Buffer Allocation Problem (BAP_s) remains intractable. The problem is conjectured to be **NP**-complete (see the following paragraph). The **NP**-hardness follows from the observation that each t-ring in the construction in Theorem 7.1 has to have a token assigned to a process component pair in order to prevent deadlock. It does not matter if the token is allocated from the token pool of the sending or the receiving process component. Hence, the reduction used in Theorem 7.1 can be applied with no modification.

We conjecture that the corresponding Buffer Sufficiency Problem (BSP_s) is in **P**. This is because the relative order in which tokens from a particular token pool are utilized is invariant with respect to the colouring sequences. Hence, we believe that determining the sufficiency is similar to the Delay-free Buffer Allocation Problem and hence is in **P**. If this is the case, BAP_s is **NP**-complete.

The Delay-free Buffer Allocation Problem (DBAP_s) remains in **P**. The problem can be solved by first reversing all arcs in the communication graph and running the algorithm described in Figure 10.

To see this, consider the interval, which is associated with a send vertex $v_{i,c}$, during which a token is utilized when colouring the corresponding receive vertex, $v_{j,d}$, yellow. Interval $I_{i,c}$ is the smallest interval $[c,d[$ such that there is a path from vertex $v_{j,t}$ to vertex $v_{i,d}$, or if vertex $v_{i,d}$ does not exist, the interval $[c,E_i[$. Formally, we add vertex v_{i,E_i}, via arc (v_{i,E_i-1}, v_{i,E_i}) to the end of each component i.

A token may be utilized when receive vertex $v_{j,t}$ is coloured yellow, and freed at some point prior to either when vertex $v_{i,d}$ is coloured yellow, or when all vertices are coloured green; interval $I_{i,c}$ represents this period, that is, the period during vertex $v_{j,t}$ is yellow. The maximum overlap density of the intervals associated with component i is the number of tokens needed for a delay-free colouring.

Finally, observe that if we reverse all arcs—treating all receive vertices as send

vertices and vice-versa, and the added vertices as start vertices—then, by definition, vertex $v_{i,c}$ becomes terminally communication dependent on vertex $v_{i,d}$ and we can use the algorithm in Figure 10 to compute the maximum overlap density of each component and hence the required token assignment.

9 Buffer Allocation in Systems with Send and Receive Side Buffers

So far we have considered systems exclusively with either send side or receive side buffers. In this section we investigate systems with buffers on both the send and the receive sides; many communication systems use per-host buffer pools for both receiving and sending messages. The choice of where to buffer the message—on the sender or on the receiver—increases the difficulty of determining the system's properties.

We assume a lazy mechanism for utilizing buffers: first use a buffer from the sender's pool. If none is available, use a buffer from the receiver's pool. If neither is available, attempt to free a send side buffer by transferring its contents to a buffer belonging to the corresponding receiver. Intuitively, the system attempts to maximize buffer use, without attempting to predict the future.

The corresponding colouring game allows tokens to be allocated from the pools belonging to both the sending component and the receiving component. Correspondingly, a lazy token utilization scheme is used: let (s_i, r_j) be a communication arc from process component P_i to process component P_j. The following rules apply during the application of rule $recv \overset{\bullet}{\rightarrow} yel$ to vertex r_j:

(1) If a token belonging to component P_i is available, use it.
(2) Otherwise, if a token belonging to component P_j is available, use it.
(3) Otherwise, if a token belonging to component P_i is currently placed on arc (t_i, r_k), $t_i \in P_i$, $r_k \in P_k$, and a token belonging to component P_k is available. Then the token on arc (t_i, r_k) may be replaced with the one belonging to P_k, freeing a token to be used in the current application of rule $recv \overset{\bullet}{\rightarrow} yel$.

We call this the **mixed utilization scheme**.

Unsurprisingly, the Buffer Allocation Problem (BAP_{sr}) remains intractable within the mixed utilization scheme. This is because the receive side utilization scheme, which provides no choice of token pools, can be simulated within the mixed utilization scheme. Concretely consider the receive side utilization scheme analyzed in Section 7: to simulate the receive side utilization scheme on communication graph G, within the mixed utilization scheme, each communication arc in G is replaced by the widget illustrated in Figure 12. Since vertex q cannot be coloured green until vertex r is coloured yellow, and component P' has no tokens, applying

rule $recv \overset{\bullet}{\rightarrow} yel$ to r requires that P_j has an available token, regardless of whether P_i has an available token.

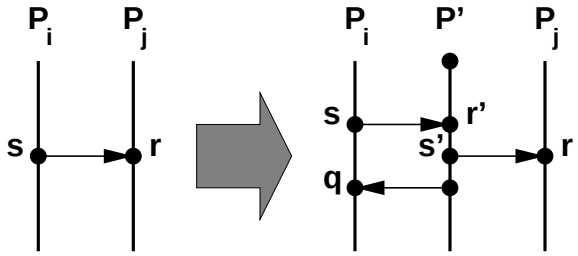

Fig. 12. Nullifying send side token pools.

Similarly, the Buffer Sufficiency Problem (BSP_{sr}), within the mixed utilization scheme, is also **coNP**-complete. The hardness follows from Theorem 7.2 and the preceding argument. Since a colouring sequence also serves as a deadlock certificate in this case, the **coNP**-completeness result follows.

The interesting property of the mixed utilization scheme is that the Delay-free Buffer Allocation Problem ($DBAP_{sr}$) is intractable; the choice of token pools increases the search space of solutions exponentially! The reduction is from 3SAT.

Theorem 9.1 *The Delay-free Buffer Allocation Problem (*$DBAP_{sr}$*) is* **NP***-hard.*

Proof: Let F be an instance of 3SAT, comprising n variables, labeled x_i, $i = 1 \ldots n$, and c clauses. We construct a communication graph G such that there exists a token assignment of $n+2$ tokens that prevents any colouring sequence from delaying on G if and only if the corresponding assignment satisfies F.

The graph G comprises $2n+3$ process components: the first $2n$ are labeled P_{x_i} and $P_{\bar{x}_i}$, $i = 1 \ldots n$, and the remaining three process components are labeled P, Q_0 and Q_1, respectively. The graph is divided into $c+1$ epochs: epoch 0 corresponds to the variable assignment, and epochs 1 through c correspond to clause evaluation.

In epoch 0 each process component P_{x_i} contains a single send vertex s_i that is adjacent to the receive vertex r_i located in epoch 0 of process component $P_{\bar{x}_i}$. Process component Q_0 (and Q_1) contains four vertices: two receive vertices $q_{0,1}$ and $q_{0,2}$ (respectively $q_{1,1}$ and $q_{1,2}$), followed by two send vertices $t_{0,1}$ and $t_{0,2}$ (respectively $t_{1,1}$ and $t_{1,2}$). Finally, process component P contains eight vertices: two send vertices, $s_{0,1}$ and $s_{0,2}$, that are adjacent to vertices $q_{0,1}$ and $q_{0,2}$; two receive vertices, $r_{0,1}$ and $r_{0,2}$, that are adjacent to $t_{0,1}$ and $t_{0,2}$; two more send vertices, $s_{1,1}$ and $s_{1,2}$, that are adjacent to $q_{1,1}$ and $q_{1,2}$; and two more receive vertices, $r_{1,1}$ and $r_{1,2}$, that are adjacent to $t_{1,1}$ and $t_{1,2}$. See Figure 13. Epoch 0 has two important properties.

Property 9.2 *Any token assignment must assign at least one token to either component P_{x_i} or $P_{\bar{x}_i}$ to prevent the colouring sequence from delaying after colouring vertex s_i yellow.*

Lemma 9.3 *A token assignment on G having only $n+2$ tokens must assign two tokens to process component P to prevent a colouring sequence from delaying after yellow colouring one of the send vertices $s_{0,1}$, $s_{0,2}$, $s_{1,1}$ or $s_{1,2}$.*

Proof: Since n tokens must be allocated to the process components P_{x_i} or $P_{\bar{x}_i}$, $i = 1,\ldots,n$, this leaves only two tokens to be allocated. Since the colouring rule sequence *send*→*yel*, *recv*$\dot{\rightarrow}$*yel*, *send*→*grn*, *send*→*yel*, *recv*$\dot{\rightarrow}$*yel* can colour send vertices $s_{0,1}$ and $s_{0,2}$, or send vertices $s_{1,1}$ and $s_{1,2}$, component pairs (P,Q_0) and (P,Q_1) must each have two tokens between them. This can only happen by assigning the tokens to P. ■

A corollary of these properties is that once a legal token assignment is made, no colouring sequences will delay in epoch 0. The choice of allocating the token on P_{x_i} versus $P_{\bar{x}_i}$ corresponds to fixing the variable assignment.

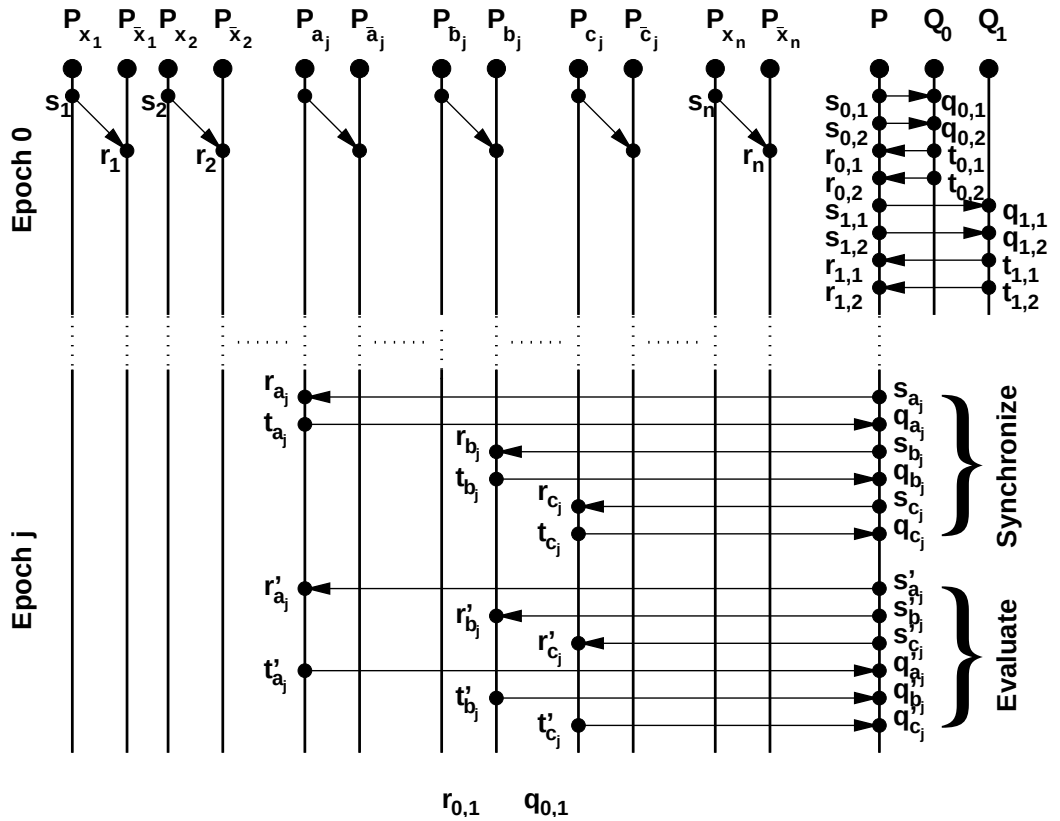

Fig. 13. Reduction from 3SAT to DBAP$_{sr}$.

For $j = 1\ldots c$, epoch j corresponds to the jth clause. Each epoch comprises two parts of six arcs each: the synchronization part and the evaluation part. Four process components are involved in an epoch: the three components, P_{a_j}, P_{b_j}, and P_{c_j}, whose labels are the literals in the jth clause, where $a_j, b_j, c_j \in \{x_1, \bar{x}_1, \ldots, x_n, \bar{x}_n\}$, and component P, which is involved in every epoch. Epoch j of component P_{a_j} comprises four vertices: receive vertex $r_{a_j,j}$, send vertex $t_{a_j,j}$, receive vertex $r'_{a_j,j}$, and send vertex $t'_{a_j,j}$. Process components P_{b_j} and P_{c_j} are analogously formed.

In epoch j component P has 12 vertices, the first six are these: send vertex $s_{a_j,j}$, receive vertex $q_{a_j,j}$, send vertex $s_{b_j,j}$, receive vertex $q_{b_j,j}$, send vertex $s_{c_j,j}$, and receive vertex $q_{c_j,j}$. These are followed by three send vertices: $s'_{a_j,j}$, $s'_{b_j,j}$, and $s'_{c_j,j}$, and three receive vertices: $q'_{a_j,j}$, $q'_{b_j,j}$, and $q'_{c_j,j}$.

Each vertex $s_{l,j}$ is adjacent to vertex $r_{l,j}$, each vertex $t_{l,j}$ is adjacent to vertex $q_{l,j}$, each vertex $s'_{l,j}$ is adjacent to vertex $r'_{l,j}$, and each vertex $t'_{l,j}$ is adjacent to vertex $q'_{l,j}$; see Figure 13. For conciseness we drop the last index, j, if it is obvious from the context. Epoch j has three important properties:

Lemma 9.4 *If vertex q'_{c_j} (in epoch j) is coloured green and vertex $s_{a_{j+1}}$ (in epoch $j+1$) is still red, then no tokens that belong to component P are assigned to arcs. The same applies to vertex pairs (q_{a_j}, s_{b_j}), (q_{b_j}, s_{c_j}), and (q_{c_j}, s'_{a_j}), also in epoch j.*

Proof: All ancestors of q'_{c_j} must be coloured green and all descendants of $s_{a_{j+1}}$ must be coloured red. This includes all vertices in G, except some vertices s_i and r_i in epoch 0, which are not adjacent to vertices in component P. Hence, the tokens belonging to P are not assigned to any arc. The same argument applies to the other vertex pairs. ■

Lemma 9.5 *A colouring sequence on G can delay only when yellow colouring receive vertices r'_{a_j}, r'_{b_j}, r'_{c_j}, q'_{a_j}, q'_{b_j}, or q'_{c_j}.*

Proof: As a corollary of properties 9.2 and 9.3, no colouring sequence can delay in epoch 0. Thus, we need only check that no colouring sequence can delay in the first part of epoch j, $j = 1 \ldots c$.

By property 9.4, if s_{a_j} is red and its predecessor is green, then no tokens of P are in use. Hence, to colour s_{a_j} green, a token is available to colour r_{a_j} yellow. Since vertex r_{a_j} is a predecessor of t_{a_j}, vertex r_{a_j} must be coloured green before t_{a_j} may be coloured yellow. Thus the token is freed before t_{a_j} is coloured green, and may be used to colour vertex q_{a_j} yellow after t_{a_j} is coloured yellow. A similar argument applies to the vertices r_{b_j}, q_{b_j}, r_{c_j}, and q_{c_j}. ■

Property 9.6 *A colouring sequence can delay in epoch j if and only if none of the three process components, P_{a_j}, P_{b_j}, and P_{c_j}, have a token assigned.*

Proof: For the 'if' direction consider a colouring sequence that colours vertex q_{c_j} green, but has not yet coloured vertex s'_{a_j} yellow. By definition, no delays occur, if rule $recv \xrightarrow{\bullet} yel$ may always be applied to colour a receive vertex yellow. To colour the send vertices s'_{a_j}, s'_{b_j}, and s'_{c_j} yellow and then green, the receive vertices r'_{a_j}, r'_{b_j}, and r'_{c_j}, must be coloured yellow via rule $recv \xrightarrow{\bullet} yel$. Since the receive vertices r'_{a_j}, r'_{b_j}, and r'_{c_j} are not ancestors of the send vertices s'_{a_j}, s'_{b_j}, and s'_{c_j}, none of the receive vertices need be coloured green before the send vertices are coloured

yellow. However, component P has only two tokens, and none of components P_{a_j}, P_{b_j}, P_{c_j} have any. Hence, rule $recv \xrightarrow{\bullet} yel$ can only be invoked twice, instead of the requisite three times. Thus, a colouring sequence can delay in epoch j.

For the 'only if' direction we claim that if a literal component P_{a_j}, P_{b_j}, or P_{c_j} has a token, rule $recv \xrightarrow{\bullet} yel$ can be invoked on any of the six receive vertices r'_{a_j}, r'_{b_j}, r'_{c_j}, q'_{a_j}, q'_{b_j}, and q'_{c_j}. Since r'_{a_j} is a predecessor of t'_{a_j}, r'_{a_j} must be coloured green before t'_{a_j}, and hence before q'_{a_j} is coloured yellow. Thus, the same token that was allocated upon the application of rule $recv \xrightarrow{\bullet} yel$ to vertex r'_{a_j}, may also be allocated upon the application of rule $recv \xrightarrow{\bullet} yel$ to vertex q'_{a_j}; the same argument is applicable to vertices q'_{b_j} and q'_{c_j}. Applying rule $recv \xrightarrow{\bullet} yel$ to vertices r'_{a_j} and r'_{b_j}, uses the two tokens from component P. To colour vertex r'_{c_j} yellow there are three possible scenarios:

(1) the colouring sequence has already freed one of the tokens, allowing it to be reused,
(2) component P_{c_j} has a token, in which case it is used, or
(3) component P_{a_j} (or P_{b_j}) has a token, in which case it replaces the token used to yellow colour vertex r'_{a_j} (or r'_{b_j}) and the freed token is used to colour vertex r'_{c_j}.

Since at least one component P_{a_j}, P_{b_j}, or P_{c_j} have a token, the claim is proven. ■

By property 9.6 a colour sequence will delay in epoch j if and only if none of the process components P_{a_j}, P_{b_j}, or P_{c_j} has a token, which corresponds to the jth clause having no literals that are true. Thus, a token assignment of size $2n+2$ prevents any colouring sequence on G from delaying if and only if the corresponding assignment satisfies F. ■

10 Buffer Allocation in Channel Based Systems

In channel based systems processes communicate via pairwise connections that are created at start-up. Each connection, called a channel, is specified by its end-points and is used by one process to send messages to the other. Each channel functions independently of other channels in the system, and resources such as buffers are allocated on a per channel basis, rather than per process. Finally, channels behave like queues, that is, messages are removed from the channel in the same order that they are inserted.

Channels may either be unidirectional, comprising source and destination end-points, or bidirectional, comprising two symmetric end-points. In the former case, only the source process may insert messages into the channel and only the desti-

nation process may remove messages from the channels. A bidirectional channel is equivalent to two unidirectional channels, allowing both processes to insert and remove messages from the channel. Here we only consider unidirectional channels.

Except for buffer allocation, channel based communication does not differ from the previously described send/receive mechanism. In fact, an unbuffered channel communication is just a synchronous send/receive communication. Thus, we can derive similar results for channel based systems.

In the corresponding colouring game tokens, are allocated to channels (component pairs) instead of to components. This change does not change the properties used in our proofs. In fact, Lemma 6.1 may be used unchanged. We call this the **per channel utilization scheme**.

10.1 The Buffer Allocation Problem

The corresponding Buffer Allocation Problem (BAP_{pc}) is this: given a communication graph G and an integer k, determine whether there exist a token assignment of size k, such that no colouring sequence deadlocks on G. Even though token utilization, during the colouring of a communication graph, is only dictated by the communication arcs within a process component pair, determining the number of tokens needed remains **NP**-hard. The proof is similar in spirit to Theorem 7.1.

Theorem 10.1 *The Buffer Allocation Problem (*BAP_{pc}*) is* **NP***-hard.*

Proof: We prove this by reducing 3SAT to BAP_{pc}. For any 3SAT instance F we construct a corresponding communication graph G—polynomial in size of F—such that for a token assignment of size n, any colouring sequence can complete on G if and only if the corresponding variable assignment satisfies F.

Let F be an instance of 3SAT on n variables and comprising c clauses. The construction is nearly identical to that in Theorem 7.1, except for the widgets representing the clauses of F. The graph G has $2n$ process components that are labeled by the literals of F, P_{x_i} and $P_{\bar{x}_i}$, $i = 1 \ldots n$. Each component comprises $c+1$ epochs, where each epoch contains zero or two vertices.

As in Theorem 7.1, epoch 0 fixes a variable assignment. In epoch 0 each component has two vertices: a send vertex, labeled s_{x_i} (or $s_{\bar{x}_i}$), and a receive vertex r_{x_i}, (respectively $r_{\bar{x}_i}$), $i = 1 \ldots n$. Vertex s_{x_i} is adjacent to vertex $r_{\bar{x}_i}$, and vertex $s_{\bar{x}_i}$ is adjacent to vertex r_{x_i}; this is a 2-ring, identical to epoch 0 in Theorem 7.1. Epoch 0 has the the following property:

Property 10.2 *Any colouring sequence on G will deadlock in epoch* 0 *unless each process component pair has a token assigned to the token pool of either* $(P_{x_i}, P_{\bar{x}_i})$*, or*

$(P_{\bar{x}_i}, P_{x_i})$, $i = 1 \ldots n$. *Thus, the token assignment must be of at least size n. (Follows from Lemma 6.1.)*

Property 10.2 yields the following correspondence between assignments on F and token assignments of size n.

Property 10.3 *The corresponding token assignment of a variable assignment on F assigns a token to the channel* $(P_{x_i}, P_{\bar{x}_i})$ *if* x_i *is true, or to* $(P_{\bar{x}_i}, P_{x_i})$ *if* x_i *is false.*

The jth epoch represents the jth clause, (a_j, b_j, c_j), of formula F, where $a_j, b_j, c_j \in \{x_1, \bar{x}_1, \ldots, x_n, \bar{x}_n\}$. The process components P_{a_j}, $P_{\bar{a}_j}$, P_{b_j}, $P_{\bar{b}_j}$, P_{c_j}, and $P_{\bar{c}_j}$ form a 6-ring, while the remaining components have no vertices in the jth epoch. Process component P_{a_j} has two vertices in the jth epoch: a send vertex, $s_{a_j,j}$, and a receive vertex $r_{a_j,j}$; similarly, the other five components have a send and receive vertex that are correspondingly named. The arcs linking the 6 components are $(s_{a_j,j}, r_{\bar{a}_j,j})$, $(s_{\bar{a}_j,j}, r_{b_j,j})$, $(s_{b_j,j}, r_{\bar{b}_j,j})$, $(s_{\bar{b}_j,j}, r_{c_j,j})$, $(s_{c_j,j}, r_{\bar{c}_j,j})$, and $(s_{\bar{c}_j,j}, r_{a_j,j})$, and form a 6-ring, as illustrated in Figure 14. The key property of the jth epoch is this:

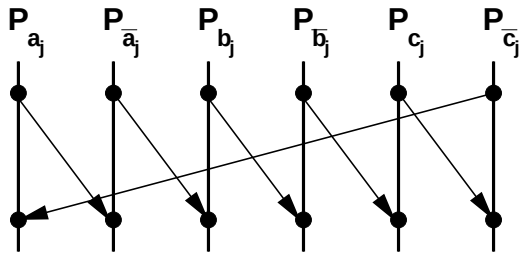

Fig. 14. The clause representation in epoch j.

Property 10.4 *No colouring sequence on G will deadlock in the jth epoch if and only if at least one of the channels has a token:* $(P_{a_j}, P_{\bar{a}_j})$, $(P_{\bar{a}_j}, P_{b_j})$, $(P_{b_j}, P_{\bar{b}_j})$, $(P_{\bar{b}_j}, P_{c_j})$, $(P_{c_j}, P_{\bar{c}_j})$, $(P_{\bar{c}_j}, P_{a_j})$. *(Follows from Lemma 6.1.)*

A refined version of property 10.4 is more useful:

Lemma 10.5 *For any token assignment of size n such that no colouring sequence deadlocks on G in epoch* 0*, no colouring sequence on G will deadlock in the jth epoch if and only if at least one of the channels* $(P_{a_j}, P_{\bar{a}_j})$, $(P_{b_j}, P_{\bar{b}_j})$, *and* $(P_{c_j}, P_{\bar{c}_j})$, *has a token.*

Proof: By property 10.2, all token assignments that do not cause deadlock in epoch 0 only assign tokens to channels of the form $(P_{x_i}, P_{\bar{x}_i})$ or $(P_{\bar{x}_i}, P_{x_i})$. Hence, only channels $(P_{a_j}, P_{\bar{a}_j})$, $(P_{b_j}, P_{\bar{b}_j})$, and $(P_{c_j}, P_{\bar{c}_j})$ can have a token. By property 10.4, no colouring sequence on G will deadlock in epoch j if one of these channels has a token. ■

We claim that given a token assignment of size n, any colouring sequence can complete on G if and only if the corresponding variable assignment satisfies F.

If an assignment x satisfies F, then every clause has at least one literal that evaluates

to true. By Property 10.3, in each of the j epochs at least one of the channels listed in Property 10.5 will be allocated a token. Hence, by Property 10.5 no colouring sequence will deadlock on G.

If an assignment x does not satisfy F then there is at least one clause in which all literals are false. Let (a_j, b_j, c_j) be the unsatisfied clause. By property 10.3, the corresponding token assignment will not assign a token to $(P_{a_j}, P_{\bar{a}_j})$, $(P_{b_j}, P_{\bar{b}_j})$, or $(P_{c_j}, P_{\bar{c}_j})$, hence, by Property 10.5, all colouring sequences will deadlock.

Thus, DBAP$_{pc}$ is **NP**-hard. ■

Since tokens are assigned on a per channel basis, token usage depends only on the two process components that comprise the channel. Consequently, the sufficiency of a token assignment can be verified in linear time. Thus, the easier problem BSP$_{pc}$ is in **P**, implying that BAP$_{pc}$ is **NP**-complete. We describe the verification algorithm and prove its correctness.

To verify the sufficiency of a token assignment, perform a colouring on G: at each step of the colouring a vertex of G is coloured according to the rules in section 3. Using a queue to keep track of colourable vertices, means that determining which vertex to colour next takes $O(1)$ time. Since each vertex changes colour at most twice—the maximum length of any colouring sequence is $2|V|$ colourings—colouring a graph takes $O(|V|)$ time. The token assignment is sufficient if and only if the colouring sequence can complete. The algorithm's correctness follows immediately from the following theorem: any colouring sequence on G can complete if and only if some colouring sequence on G can complete. Thus, a token assignment is sufficient if and only if some colouring sequence on G can complete.

Theorem 10.6 *Let G be a communication graph and B a token assignment on G. Any colouring sequence on G can complete if and only if a colouring sequence on G can complete.*

Proof: For any communication graph G, we construct a new graph G' where a channel's pool of m tokens is simulated by m additional process components, the size of the corresponding token assignment is zero, and every colour sequence on G corresponds to a colouring sequence on G', such that a colouring sequence on G can complete if and only if the corresponding colouring sequence on G' can complete. Since the token assignment on G' is zero, by Lemma 6.2 a colouring sequence on G' can complete if and only if every colouring sequence on G' can complete. Hence, every colouring sequence on G can complete if and only if a colouring sequence on G can complete.

To simulate an m token channel (a channel that has been assigned m tokens) m process components are chained together. For each channel (P, Q) with m tokens, m process components $P_1, P_2, \ldots, P_m$ are inserted between P and Q. The channel (P, Q) is replaced with these channels: $(P, P_1), (P_1, P_2), \ldots, (P_{m-1}, P_m), (P_m, Q)$. Each arc

from P to Q is replaced by a chain of arcs from $P \to P_1 \to P_2 \to \ldots \to P_{m-1} \to P_m \to Q$. The replacement is illustrated in Figure 15.

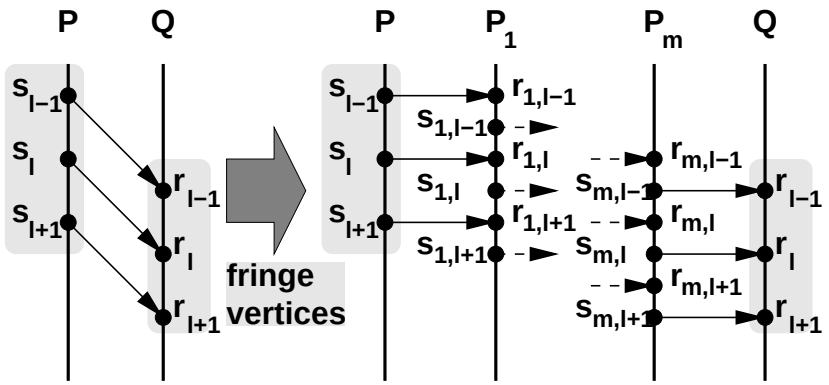

Fig. 15. Simulating m tokens by m components.

We claim that a colouring sequence, Σ, on G will deadlock if and only if the corresponding colouring sequence, Σ', on G' deadlocks. First, we construct the correspondence and argue its correctness. Second, we argue that sequence Σ deadlocks on G if and only if the corresponding sequence Σ' deadlocks on G'. Finally, we apply Lemma 6.2 to prove our result.

Since the transformation is iterative—each m token channel is independent of the other channels—it is sufficient to derive the correspondence between the colouring sequence on G and the graph G' where a single m token channel has been replaced. Let (P,Q) denote the channel in G that is replaced in G'.

Let $(s_l, r_l) \in G$, $l = 1,2,\ldots$, denote the arcs from process component P to Q. The corresponding paths in G' are

$$(s_l, \underbrace{r_{1,l}, s_{1,l}}_{P_1}, \underbrace{r_{2,l}, s_{2,l}}_{P_2}, \ldots, \underbrace{r_{m,l}, s_{m,l}}_{P_m}, r_l),$$

where each arc $(r_{k,l}, s_{k,l})$ is within process component P_k and each arc $(s_{k,l}, r_{k+1,l})$ is between process components P_k and P_{k+1}; the vertices s_l and r_l, $l = 1,2,\ldots$ are called the **fringe** vertices.

A colouring sequence Σ can be represented as a sequence of differences (or moves), δ_i, between every two consecutive colourings χ_i and χ_{i+1}. The sequence $\Delta_\Sigma = \delta_1\delta_2\ldots$ is a sequence of colouring game moves $\delta_i = \langle v, \text{colour}\rangle$ such that applying δ_i to colouring χ_i yields χ_{i+1}, the next colouring in Σ; Δ_Σ can be derived from Σ and, Σ can be derived from Δ_Σ and G. The sequence Δ_Σ comprises two types of moves: those that colour fringe vertices, called **fringe moves**, and those that do not, called **normal moves**.

Given a colouring sequence Σ on G, we transform it into the corresponding colouring sequence Σ' on G'. The transformation replaces some fringe moves in sequence

Δ_Σ with sequences of moves, resulting in the corresponding move sequence $\Delta_{\Sigma'}$. This sequence comprises normal moves and **added moves**; added moves are a mixture of fringe moves and moves on the vertices within the added components P_i. There are four types of fringe moves in Δ_Σ: colour a send vertex s_l yellow ($\langle s_l, \text{yel}\rangle$), colour a send vertex s_l green ($\langle s_l, \text{grn}\rangle$), colour a receive vertex r_l yellow ($\langle r_l, \text{yel}\rangle$), and colour a receive vertex r_l green ($\langle r_l, \text{grn}\rangle$). The transformation is performed in the order that the moves occur in sequence Δ_Σ.

- If δ_i is a normal move, it is not modified.
- If $\delta_i = \langle s_l, \text{yel}\rangle$, then no action is taken.
- If $\delta_i = \langle s_l, \text{grn}\rangle$, we replace it with the sequence

  $$\langle r_{1,l}, \text{yel}\rangle, \langle s_l, \text{grn}\rangle, \langle r_{1,l}, \text{grn}\rangle, \langle s_{1,l}, \text{yel}\rangle,$$

  suffixed by the sequences

  $$\langle r_{j,l}, \text{yel}\rangle, \langle s_{j-1,l}, \text{grn}\rangle, \langle r_{j,l}, \text{grn}\rangle, \langle s_{j,l}, \text{yel}\rangle, \quad j = 2, \ldots, k-1$$

  where k is the smallest integer such that the move $\langle s_{k,l-1}, \text{grn}\rangle$ has not yet been inserted into the move sequence Δ_Σ, that is, vertex $s_{k,l-1}$ has not yet been coloured green. If $k < 3$, then the suffix is empty.
- If $\delta_i = \langle r_l, \text{yel}\rangle$, we remove it from the sequence; it is restored when we replace the move $\langle r_l, \text{grn}\rangle$.
- If $\delta_i = \langle r_l, \text{grn}\rangle$ we replace this move with the sequence

  $$\langle r_l, \text{yel}\rangle, \langle s_{m,l}, \text{grn}\rangle, \langle r_l, \text{grn}\rangle,$$

  suffixed with the sequences

  $$\langle r_{g_j,h_j}, \text{yel}\rangle, \langle s_{g_j-1,h_j}, \text{grn}\rangle, \langle r_{g_j,h_j}, \text{grn}\rangle, \langle s_{g_j,h_j}, \text{yel}\rangle, \quad j = 0, \ldots, k-1,$$

  where $g_j = m - j$, $h_j = l + 1 + j$, and k is the smallest integer such that the move $\langle s_{m-k-1,l+1+k}, \text{yel}\rangle$ has not yet been inserted into the sequence, that is, vertex $s_{m-k-1,l+1+k}$ has not yet been coloured yellow. Since the head of this sequence colours $s_{m,l}$ green, $r_{m,l+1}$ could be coloured yellow, if $s_{m-1,l+1}$ is yellow, then $s_{m-1,l+1}$ could be coloured green followed by $r_{m,l+1}$ and finally, $s_{m,l+1}$ could be coloured yellow; this colouring cascades down the added process components. If $k < 1$, then the suffix is empty.

For all added moves $\langle r_{g,h}, \cdot\rangle$ and $\langle s_{g,h}, \cdot\rangle$, $h \geq l$ and g is not a function of l. Hence, the case of $l = 1$ is no different from the case of $l > 1$. Furthermore, it is important to note that each of the replacement sequences is maximal, that is, no additional valid colouring moves on the chain process components P_i, $i = 1, \ldots, m$, may be

suffixed to them. The new sequence looks like this:

$$\Delta_{\Sigma'} = \overbrace{\delta_1 \ldots \delta_{h_1}}^{\text{normal moves}} \underbrace{\delta'_1 \ldots \delta'_{g_1}}_{\text{added moves}} \overbrace{\delta_{h_1+1} \ldots \delta_{h_2}}^{\text{normal moves}} \underbrace{\delta'_{g_1+1} \ldots \delta'_{g_2}}_{\text{added moves}} \ldots$$

Since G is a contraction of G', all normal vertices are coloured by $\Delta_{\Sigma'}$ in the same order as in Δ_Σ. Recall that normal vertices are not adjacent to the process component chain, and hence, are not affected by the transformation. While normal vertices within process components P and Q may depend on the order that the fringe vertices are coloured, the dependence is via process arcs, not communication arcs. Consequently, the normal vertices only depend on the order that the fringe vertices are coloured green. Fortunately, this order is preserved. By inspection, the replacement sequences of moves are valid. Thus, the transformed sequence $\Delta_{\Sigma'}$ is valid. Additionally, all green colouring moves on fringe vertices are preserved by the transformation; a vertex is coloured green by Δ_Σ if and only if the corresponding vertex is coloured green by $\Delta_{\Sigma'}$. The following property is key:

Lemma 10.7 *Δ_Σ deadlocks on G if and only if $\Delta_{\Sigma'}$ deadlocks on G'.*

Proof: By contradiction, suppose that Δ_Σ deadlocks on G while $\Delta_{\Sigma'}$ can be extended, that is, another vertex colouring move may be suffixed to $\Delta_{\Sigma'}$. Let v be the vertex that can be coloured by the extension. Vertex v may either be a normal vertex, a fringe vertex, or a vertex belonging to a process chain. The latter is impossible because every replacement sequence of moves is maximal.

If v is a normal vertex, then its predecessors are either a normal vertex or a fringe vertex that has been coloured green. Since the transformation preserves the colourings of normal vertices and the order in which vertices are coloured green, if $\Delta_{\Sigma'}$ can be extended by colouring v, then so can Δ_Σ, which is a contradiction.

If v is a fringe vertex, there are four cases: either v is a send vertex s_l being coloured yellow or green, or v is a receive vertex r_l being coloured yellow or green. The transformation does not affect moves that colour send vertices yellow and such a colouring only depends on its process component predecessor being green. Hence, if the colouring can be suffixed to $\Delta_{\Sigma'}$, it can also be suffixed to Δ_Σ; resulting in a contradiction. If the extension colours the send vertex green, this means that the original sequence Δ_Σ can be extended by either adding the colourings $\langle s_l, \text{grn}\rangle$ or $\langle r_l, \text{yel}\rangle\langle s_l, \text{grn}\rangle$, depending on whether r_l has been coloured yellow or not in the original sequence Δ_Σ; thus, it is a contradiction.

Similarly, if v is a fringe receive vertex being coloured green, this is not possible, because the transformation colours fringe receive vertices yellow, then green, by a single replacement sequence. Finally, if v is a fringe receive vertex r_l that can be coloured yellow, the original sequence Δ_Σ can be extended by the move $\langle r_l, \text{grn}\rangle$,

because in the original sequence the corresponding send vertex s_l has already been coloured green. Thus, we have another contradiction.

In the other direction, if the original sequence can be extended, then transforming the extension of the sequence Δ_Σ yields an extension to the presumably deadlocked sequence $\Delta_{\Sigma'}$. Thus, Δ_Σ deadlocks on G if and only if $\Delta_{\Sigma'}$ deadlocks on G'. ■

A corollary of Property 10.7 is that the colouring sequence Σ deadlocks if and only if the colouring sequence Σ' deadlocks.

By Lemma 6.2 a colouring sequence on G' can complete if and only if all colouring sequences on G' can complete. Hence, a colouring sequence on G can complete if and only if all colouring sequences on G can complete. ■

Corollary 10.8 *A colouring sequence on G can complete if and only if the token assignment is sufficient.*

10.2 The Delay-free Buffer Allocation Problem

For the Delay-free Buffer Allocation Problem, the algorithm derived in section 7.3 suffices with a small modification. Since the token pools are per channel, rather than per process component, the computation must be performed on a per pool basis. Hence, there is an additional factor of n in the runtime. Since each process may be using up to n channels, the runtime of the algorithm becomes $O(|V|n^2 + |V|n\log(|V|n))$; the cost increases because the number of allocations to be made becomes quadratic in n.

11 Conclusion

As message passing and message passing standards become increasingly popular, the problem of determining the number of message buffers needed for a program to execute safely, of even whether a program can execute safely using a particular implementation of a message passing system becomes more relevant. The relevance of this problem increases as more and more functionality of message passing systems is off-loaded to the network interface card, where limited buffer space is a serious issue. Even if message passing is kept in main memory, buffer space can still be limited due to the sometimes very large data sets used in many parallel and distributed programs. Unfortunately, determining the message buffer requirements of a program, even one that has a static communication pattern, is intractable.

We have shown that in the receiver side buffering scheme, determining the number of buffers needed to assure safe execution of a program is **NP**-hard, and that even

verifying whether a number of assigned buffers is sufficient is **coNP**-complete. On the positive side, if we require that no send delays, we provide a polynomial time algorithm for computing the minimum number of buffers. By allocating this number of buffers, safe execution is guaranteed. In addition, we have implemented the DBAP$_r$ algorithm, and it is now part of the Millipede debugging system [21].

For systems that use the sender side buffering scheme, the Buffer Allocation Problem remains **NP**-hard. In addition, we conjecture that the Buffer Sufficiency Problem can be solved in polynomial time because the order of the sends in each process is fixed. The Delay-free Buffer Allocation problem for systems with only send side buffers can be solved in polynomial time.

For systems that use the sender and receiver side buffering scheme, the Buffer Allocation Problem as well as the Buffer Sufficiency Problem remain intractable. More interestingly, the Delay-free Buffer Allocation problem is also intractable.

For systems that use a unidirectional channel buffering scheme, both the Buffer Sufficiency Problem as well as the Delay-free Buffer Allocation Problem have polynomial time algorithms. However, the Buffer Allocation Problem still remains an **NP**-complete problem. The results (conjectures) are summarized below.

	Buffer Placement			
Problem	Receive	Send	Send & Receive	Channel
BAP	**NP**-hard	**NP**-hard	**NP**-hard	**NP**-complete
BSP	**coNP**-complete	(**P**)	**coNP**-complete	**P**
DBAP	**P**	**P**	**NP**-hard	**P**

11.1 Strategies for Reducing Buffer Requirements

There are several strategies that a programmer can use to reduce the likelihood of deadlock when only a few buffers are available.

The obvious solution is to use synchronous communication, which does not require any buffers at all. However, this is not always desirable.

For efficiency reasons asynchronously buffered communication is often preferred. To decrease the risk of deadlock the programmer can introduce epochs that are separated by barrier synchronizations. This might reduce the number of buffers needed for each epoch, as no buffer requirement spans an epoch boundary. If each epoch only needs a small number of buffers, the risk of deadlock due to buffer insufficiency is reduced.

References

[1] V. Anantharm. The optimal buffer allocation problem. *IEEE Transactions on Information Theory*, 35(4):721–725, 1989.

[2] J. Bruck, D. Dolev, C. Ho, M. Rosu, and R. Strong. Efficient Message Passing Interface (MPI) for Parallel Computing on Clusters of Workstations. In *7th Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 64 – 73, Santa Barbara, California, July 1995.

[3] G. Burns and R. Daoud. Robust mpi message delivery with guaranteed resources. June 1995.

[4] B. Charron-Bost, F. Mattern, and G. Tel. Synchronous, asynchronous, and causally ordered communication. *Journal of Distributed Computing*, 9(4):173–191, 1996.

[5] S. A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the 3rd Annual ACM Symposium on the Theory of Computing*, pages 151–158, 1971.

[6] R. Cypher and E. Leu. Repeatable and portable message-passing programs. In *Proc. of The Symposium on the Principles of Distributed Computing (PODC)*, pages 22–31, 1994.

[7] R. Cypher and E. Leu. The semantics of blocking and nonblocking send and receive primitives. In *Proceedings of 8th IEEE International parallel processing symposium (IPPS)*, pages 729–735, 1994.

[8] J. Dongarra. MPI: A message passing interface standard. *The International Journal of Supercomputers and High Performance Computing*, 8:165–184, 1994.

[9] J. Dongarra, R. Hempel, A. Hey, and D. Walker. A proposal for a user-level, message-passing interface in a distributed memory environment. Technical Report TM-12231, ORNL, June 1993.

[10] G. Fox, M. Johnson, G. Lyzenga, S. Otto J. Salmon, and D. Walker. *Solving problems on concurrent processors. General techniques and regular problems*, volume 1. Prentice-Hall, Inc., 1988.

[11] D. Frye, R. Bryant, H. Ho, R. Lawrence, and M. Snir. An external user interface for scalable parallel systems. Technical report, IBM highly parallel supercomputing systems laboratory, November 1992.

[12] G. Burns, R. Daoud and J. Vaigl. LAM: An Open Cluster Environment for MPI. In *Supercomputing Symposium '94*, Toronto, Canada, June 1994.

[13] M. R. Garey and D. S. Johnson. *Computers and Intractibility: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, New York, 1979.

[14] A. Geist, A. Beguelin, J. Dongarra, W. Jiang, R. Manchek, and V. Sunderam. *PVM: Parallel Virtual Machine: A Users' Guide and Tutorial for Networked Parallel Computing*. Scientific and engineering computation. MIT Press, 1994.

[15] P. Huber, A. M. Jensen, L. O. Jepsen, and K. Jensen. Reachability trees for high-level Petri nets. *Theoretical Computer Science*, 45:261–292, 1985.

[16] K. Jensen. *Coloured Petri nets. Basic Concepts, Analysis Methods and Practical use*, volume 1. Springer Verlag, 1992.

[17] C. Keppitiyagama and A. Wagner. Asynchronous MPI messaging on myrinet. In *Proceedings 15th International Parallel and Distributed Processing Symposium (IPDPS'01)*. IEEE, 2001.

[18] L. Lamport. Time, clocks and the orderings of events in a distributed system. *Communications of the ACM*, 21:558–565, 1978.

[19] Jeff Magee and Jeffrey Kramer. *Concurrency: State Models & Java Programs*. Wiley, 1999.

[20] Message Passing Interface Forum. MPI: A Message-Passing Interface Standard. http://www.mpi-forum.org/docs/mpi-11-html/mpi-report.html, June 1995.

[21] J. B. Pedersen. *Multi-level Debugging of Parallel Message Passing Programs*. PhD thesis, University of British Columbia, Canada, 2003. In preparation.

[22] M. Reiman. The optimal buffer allocation problem in light traffic. In *IEEE Conference on Decision and Control*, 1987.

[23] T. Sheskin. Allocation of interstage storage along an automatic production line. *AIEE Transactions*, 8(1), 1975.

[24] M. Snir, S. Otto, S. Huss-Lederman, D. Walker, and J. Dongarra. *MPI: The Complete Reference*. Scientific and Engineering Computation Series. MIT Press, 1996.