



**A Characterization of
Deadlock Free Resource Contentions**

**Marina Chen
Ronald L Graham
Martin Rem**

**Computer Science Department
California Institute of Technology**

4684:TR:82

A Characterization of Deadlock Free Resource Contentions

by

Marina Chen
Ronald L. Graham
Martin Rem

Technical Report #4684

Computer Science Department
California Institute of Technology
Pasadena, California 91125

1 January 1982

Sponsored by
Office of Naval Research
Contract No. N00014-76-C-0367

and

Defense Advanced Research Contracts Agency
ARPA Order Number 3771

Monitored by
Office of Naval Research
Contract No. N00014-79-C-0597

Copyright, California Institute of Technology, 1982

A CHARACTERIZATION OF DEADLOCK FREE RESOURCE CONTENTIONS*

12 December 1981

Marina C. Chen[†]
Ronald L. Graham[‡]
Martin Rem[§]

1. INTRODUCTION

Suppose we have a system consisting of some finite number of concurrent processes. The processes are said to be *synchronized* whenever the progress of some process may have to be delayed because of conditions caused by the other processes in the collection. Whenever processes are *synchronized* there exists the possibility that they may maneuver themselves into a deadlock [see, Habermann 69]. A *deadlock* is a state in which there is at least one non-terminated processes indefinitely delayed, i.e., such that no state can be reached from which the process can make further progress.

An important class of deadlock problems arises from conflicts over shared resources, each of which can be used by only one process at a time. For example, one process while using resource *A* may require resource *B*. However resource *B* may already be in use by a second process which, in turn requires resource *A* in order to proceed. Both processes are therefore stalemated and will remain so indefinitely: they have reached a deadlock.

The problem of deadlock arising from shared resources was first noticed in operating systems having quasi-parallel processes [see, Dijkstra 68]. These processes, although logically parallel, are sequentialized when implemented by a traditional sequential machine. Thus a natural solution for

*This work is sponsored by Defense Advanced Research Projects Agency ARPA Order #3771 and monitored by Office of Naval Research Contract #N00014-79-C-0597.

[†]Computer Science Department, California Institute of Technology, Pasadena, California.

[‡]Bell Laboratories, Murray Hill, New Jersey.

[§]Department of Mathematics and Computing Science, Eindhoven University of Technology, Eindhoven, The Netherlands and Computer Science Department, California Institute of Technology.

preventing deadlock is by a global scheduler which controls the sequence of resource assignments. A well-known scheme for deadlock prevention is the "banker's algorithm" [Dijkstra 68]. In this scheme the maximum amount of each resource each process ever requires is given. There is one scheduler for all the resources. It keeps track of which resources each process is using, i.e., the "state" of the system. The banker's algorithm entails a method of computing the "safety" of states. Each time a process requires a resource the scheduler will allow the use of that resource only if the ensuing state is still safe. The banker's algorithm is an example of a global scheduler.

In this paper we look at resource sharing without a global scheduler. In fact, the only scheduling we allow is the delaying of a process requiring a resource already in use. The acquisition of resources is concurrent and asynchronous, and no global information about it is available. A process is always allowed to free a resource; the resource returned is then available for use again. Such "mutual exclusion scheduling" can, for example, be realized with semaphores [Dijkstra 68]. Each resource has a binary semaphore and an arbitration mechanism. We do not assume any particular arbitration scheme. This type of "minimal scheduling" is important, for example, in VLSI systems [Mead & Conway 80], in which the requirement to have a global scheduler would seriously degrade its computing potential.

In this paper, we prove necessary and sufficient conditions for collections of synchronized processes to be free from deadlock. In Section 3 we prove a general theorem which characterizes those collections which have deadlock potential by applying combinatorial arguments. In Section 4 we examine the computational complexity for testing the deadlock condition and derive a polynomial time algorithm for the problem. Finally in Section 5 we discuss the special case that all resources are of different types.

2. RESOURCE CONTENTIONS

Definition 2-1:

A *resource contention system* is a quadruple $(N, M, \{\rho_p\}_{p \in M}, \alpha)$ in which:

- (i) N is a set of *resource types*, $\{r_0, r_1, \dots, r_{n-1}\}$.
 - (ii) M is a set of *processes*, $\{p_0, p_1, \dots, p_{m-1}\}$.
 - (iii) $\{\rho_p\}_{p \in M}$ is a family of *request functions* $\rho_p : N \mapsto \mathcal{N}$.
 - (iv) α is an *availability function* $\alpha : N \mapsto \mathcal{N}$.
- (\mathcal{N} is the set $\{0, 1, 2, \dots\}$ of natural numbers.)

There are a fixed number of resources available, each one having a type. Resources of the same type are indistinguishable in the sense that processes cannot request a particular resource; they can only request a resource of a particular type. The availability function α gives for each resource type the number of resources available. Each request function ρ_p specifies how many resources of each type process p requests to have in use simultaneously. In particular, process p requires $\rho_p(r)$ resources of type r . A process *proceeds* if and only if it has acquired all the resources it requested. The process uses these resources until it is completed, at which time the resources are returned to the system and can be re-used by other processes.

Definition 2-2:

For $T \subseteq N$, let $P(T)$ denotes the set $\{p \in M \mid \rho_p(r) > 0 \text{ for some } r \in T\}$. Thus, $P(T)$ is just the set of processes requesting a resource of a type in T .

Definition 2-3:

For $Q \subseteq M$, let $R(Q)$ denotes the set $\{r \in N \mid \rho_p(r) > 0 \text{ for some } p \in Q\}$. Similarly, $R(Q)$ is just the set of resource types requested by processes in Q . Notice that $Q \subseteq P(R(Q))$ and $T \subseteq R(P(T))$.

Definition 2-4:

Call $p \in M$ and $q \in M$ *competing processes* if $\rho_p(r) > 0$ and $\rho_q(r) > 0$ for some $r \in N$.

Definition 2-5:

Two elements in the transitive closure of the relation "competing" are called *related*. This is an equivalence relation on the set $P(N)$. Its equivalence classes are called *components*. Denote the components by C_0, C_1, \dots, C_{c-1} .

The components C_i partition the set $P(N)$. It is easy to verify that the sets $R(C_i)$, $0 \leq i < p$, partition $R(M)$ and that $C_i = P(R(C_i))$.

Definition 2-6:

$\delta : N \mapsto N$, the *discrepancy function*, is defined by

$$\delta(r) = \max(0, \sum_{p \in P(T)} \rho_p(r) - \alpha(r)).$$

The discrepancy function gives for each resource type the number of resources the system lacks in order to accommodate all requests simultaneously.

If for any resource type r ,

$$\sum_{p \in P(T)} \rho_p(r) - \alpha(r) \leq 0.$$

then type r resources are abundant and thus r is irrelevant in the resource contention system. Therefore, without loss of generality, we assume from now on that

$$\delta(r) = \sum_{p \in P(T)} \rho_p(r) - \alpha(r) > 0.$$

In fact, it is easy to check that all the theorems that follow will still hold when these abundant resources are taken into account. Note that this assumption implies $\delta(r) \geq 1$.

Definition 2-7:

A resource contention system $\langle N, M, \{\rho_p\}_{p \in M}, \alpha \rangle$ has *danger of deadlock* when there exists

a set $T \subseteq N$, $T \neq \phi$, and a family $\{\mu_p\}_{p \in P(T)}$ of *acquisition functions* $\mu_p : T \mapsto N$ with

$$\sum_{p \in P(T)} \mu_p(r) = \sum_{p \in P(T)} \rho_p(r) - \delta(r) = \alpha(r) \text{ for all } r \in T, \quad (1)$$

$$\mu_p(r) \leq \rho_p(r) \text{ for all } p \in P(T), r \in T \text{ and} \quad (2)$$

$$\text{for all } p \in P(T), \text{ there exists } r \in T \text{ such that } \mu_p(r) < \rho_p(r). \quad (3)$$

The value of $\mu_p(r)$ may be interpreted as the number of resources of type $r \in T$ that are in use (or acquired) by process $p \in P(T)$. Equation (1) says that the situation we are interested in is when all available resources in T are fully used by all processes in $P(T)$. Inequality (2) says that the number of resources for each type each process uses can not exceed what it requires. Condition (3) says that each process in $P(T)$ has acquired less than it requested. For all processes $p \in M \setminus P(T)$, the set of resource types required by the processes in $M \setminus P(T)$ is just $R(M \setminus P(T))$. Since $R(M \setminus P(T)) \cap T = \phi$, then even though some of the processes in the system may still be able to proceed, their returned resources are not required by the processes in $P(T)$, and we therefore have a deadlock situation.

Definition 2-8:

A resource contention system $\langle N, M, \{\rho_p\}_{p \in M}, \alpha \rangle$ is called *deadlock free* when there is no danger of deadlock, i.e., for all $T \subseteq N$, $T \neq \phi$, and all $\{\mu_p\}_{p \in P(T)}$ satisfying (1) and (2), we have:

$$\text{For some } p \in P(T), \mu_p(r) = \rho_p(r) \text{ for all } r \in T. \quad (4)$$

Condition (4) says that for all T and any distribution of the resources over the processes in which processes together use as many resources as allowed, there will always be a process in $P(T)$ that has all resources it requests.

3. GENERAL THEOREM

We will state a theorem which characterizes deadlock situations and prove it by applying a theorem of P. Hall, the so-called "Marriage theorem". First we need some definitions. (see also [Mirsky 71] and [Ryser 63])

Definition 3-1:

A family of (not necessarily distinct) subsets S_1, S_2, \dots, S_m of a set S is said to have a *system of distinct representatives* (SDR) if there exists a set of distinct a_i such that $a_i \in S_i, 1 \leq i \leq m$. In this case we say a_i represents S_i and $\{a_1, \dots, a_m\}$ is an SDR for the family $\{S_1, \dots, S_m\}$.

Theorem 3-1 (P. Hall):

A family of subsets S_1, \dots, S_m has an SDR if and only if for any subset V' of the set $V = \{1, 2, \dots, m\}$, $|\bigcup_{i \in V'} S_i|$ contains at least $|V'|$ elements, i.e., for all $V' \subseteq V$,

$$\left| \bigcup_{i \in V'} S_i \right| \geq |V'|. \quad (5)$$

Definition 3-2:

We denote the $\delta(r)$ "units" of discrepancy of resource type r by $e_{r0}, e_{r1}, \dots, e_{r(\delta(r)-1)}$. For each $r \in T$, define

$$E = \{e_{rk} | 0 \leq k < \delta(r), r \in T\},$$

and define for each $p \in P(T)$

$$E_p = \{e_{rk} | \rho_p(r) > 0, 0 \leq k < \delta(r), r \in T\}.$$

$\{E_p\}_{p \in P(T)}$ is a family of subsets of E . For a process p , any element of E_p can potentially prevent it from being able to proceed. Note that each "unit" of discrepancy e_{rk} , is an element of the set E . For each process, the associated E_p is a subset of E .

Theorem 3-2:

A resource contention system $\langle N, M, \{\rho_p\}_{p \in M}, \alpha \rangle$ has danger of deadlock if and only if

$$\text{for some nonempty } T \subseteq N, \quad \sum_{r \in T} \delta(r) \geq |P(T)|. \quad (6)$$

Proof:

We first show the necessity. Assume there is danger of deadlock. Then for some $\{\mu_p\}_{p \in P(T)}$, for all $p \in P(T)$, there exists $r \in T$ such that $\mu_p(r) < \rho_p(r)$. Therefore,

$$\begin{aligned} \sum_{r \in T} \delta(r) &= \sum_{r \in T} \left(\sum_{p \in P(T)} (\rho_p(r) - \mu_p(r)) \right) \\ &= \sum_{p \in P(T)} \left(\sum_{r \in T} (\rho_p(r) - \mu_p(r)) \right) \geq \sum_{p \in P(T)} 1 = |P(T)| \end{aligned}$$

as required.

Now assume that (6) holds. We show that there is danger of deadlock. Since (6) holds, we can choose a T which is minimal in the sense that

$$\text{for all } T' \subset T, \quad \sum_{r \in T'} \delta(r) < |P(T')|. \quad (7)$$

For this T ,

$$\sum_{r \in T} \delta(r) = \left| \bigcup_{p \in P(T)} E_p \right| \geq |P(T)|. \quad (8)$$

Note that (8) is identical to the Hall condition (5) in the case of $V' = V = P(T)$. We prove that $\{E_p\}_{p \in P(T)}$ satisfies the Hall condition (5) for the existence of an SDR. For a given V' , let $G = T \setminus R(V')$. Then $P(G) \cap V' = \phi$. Also $P(G) \subseteq P(T)$ and $V' \subseteq P(T)$. Therefore, we have

$$|V'| \leq |P(T)| - |P(G)|. \quad (9)$$

Thus, for all $V' \subseteq P(T)$,

$$\begin{aligned}
\left| \bigcup_{p \in V'} E_p \right| &= \sum_{r \in (R(V') \cap T)} \delta(r) = \sum_{r \in T} \delta(r) - \sum_{r \in G} \delta(r) \\
&\geq |P(T)| - \sum_{r \in G} \delta(r) && \text{(by(8))} \\
&> |P(T)| - |P(G)| && \text{(by(7))} \\
&\geq |V'|. && \text{(by(9))}
\end{aligned}$$

By the Marriage theorem, $\{E_p\}_{p \in P(T)}$ has an SDR. What remains to be shown is that the existence of such an SDR results in a deadlock, i.e., a family $\{\mu_p\}_{p \in P(T)}$ satisfying conditions (1), (2) and (3) can be constructed.

Let $F_r = \{E_p \mid \text{for some } k, e_{rk} \text{ is the representative of } E_p\}$. We have $|F_r| \leq \delta(r)$. Now assign the value of the functions μ_p in $\{\mu_p\}_{p \in P(T)}$ in the following way: For all $p \in P(T)$, if $E_p \in F_r$, let $\mu'_p(r) = \rho_p(r) - 1$. r is the resource type in which process p is deficient. Let $\mu'_p(r) = \rho_p(r)$ for all other r . Already we have (2) and (3) holding for $\{\mu'_p\}_{p \in P(T)}$. For each $r \in T$, there are $\delta(r) - |F_r|$ "units" of discrepancy left to be distributed. These can be distributed arbitrarily to all $p \in P(\{r\})$ so that (1) is satisfied. Let $\{\mu_p\}_{p \in P(T)}$ be the set of acquisition functions after this further distribution. Then

$$\begin{aligned}
\sum_{p \in P(T)} \mu_p(r) &= \sum_{p \in P(T)} \mu'_p(r) - (\delta(r) - |F_r|) \\
&= \sum_{p \in F_r} (\rho_p(r) - 1) + \sum_{p \notin F_r} \rho_p(r) - (\delta(r) - |F_r|) \\
&= \sum_{p \in P(T)} \rho_p(r) - \delta(r).
\end{aligned}$$

Note that since $\delta(r) - |P(T)| \geq 0$, $\mu_p(r) \leq \mu'_p(r)$, for all $p \in P(T)$ and $r \in T$. Hence (2) and (3) are unchanged. This proves the theorem. \square

4. COMPLEXITY OF THE CRITERION

In Section 3, the condition for the existence of danger of deadlock in a resource contention system was given. We can also state the theorem in the form of deadlock free contention.

Theorem 4-1:

A resource contention system $\langle N, M, \{\rho_p\}_{p \in M}, \alpha \rangle$ is deadlock free if and only if

$$\text{for all nonempty } T \subseteq N, \quad |P(T)| > \sum_{r \in T} \delta(r). \quad (10)$$

To test whether a system is deadlock free, it might appear that $2^{|N|}$ of tests of (10) are necessary. Fortunately, the bipartite nature of the relation between the processes and resources enables us to formulate Theorem 4-1 in terms of an SDR of the "dual" system described by Definition 3-2. The fact that there exist polynomial algorithms for testing the Hall condition (5) [Lawler 76] also makes it possible to test the deadlock free condition in polynomial time.

Lemma 4-2:

Inequality (10) is equivalent to

$$\text{for all } x \in M, \quad \text{for all nonempty } T \subseteq N, \quad |P(T) \setminus \{x\}| \geq \sum_{r \in T} \delta(r). \quad (11)$$

Proof:

(10) \Rightarrow (11) is clear. Assume (10) does not hold. We show that (11) also does not hold. By hypothesis, there exists $T \subseteq N$ such that $|P(T)| \leq \sum_{r \in T} \delta(r)$. Since $P(T) \neq \phi$, there exists $x \in P(T)$, such that

$$|P(T) \setminus \{x\}| = |P(T)| - 1 < \sum_{r \in T} \delta(r).$$

This completes the proof. \square

Definition 4-1:

Let $V = \{(r, k) | 0 \leq k < \delta(r), \quad r \in N\}$. For $(r, k) \in V$, define

$$D_x(r, k) = \{p | p \in M \setminus \{x\}, \quad \rho_p(r) > 0\}.$$

$\{D_x(r, k)\}_{(r, k) \in V}$ is a family of subsets of $M \setminus \{x\}$. Each $D_x(r, k)$ has associated with it a "unit" of discrepancy $e_{r, k}$. Any element p of $D_x(r, k)$ can potentially absorb this discrepancy and therefore be unable to proceed. Note that in contrast with Definition 3-2, each process plays the role of being an element of the set $M \setminus \{x\}$, whereas for each "unit" of discrepancy, the associated $D_x(r, k)$ is a subset of $M \setminus \{x\}$.

Definition 4-2:

For $T \subseteq N$, define

$$V^*(T) = \{(r, k) | 0 \leq k < \delta(r), \text{ for all } r \in T\}$$

Note that $V = V^*(N)$. and for all $V' \subseteq V$, there exists a T such that

$$V' \subseteq V^*(T) \text{ and for all } T' \subset T, V' \not\subseteq V^*(T')$$

Lemma 4-3:

Inequality (11) holds if and only if for all $x \in M$, and all $V' \subseteq V$,

$$\left| \bigcup_{(r, k) \in V'} D_x(r, k) \right| \geq |V'|. \quad (12)$$

Proof:

$$\begin{aligned} \left| \bigcup_{(r, k) \in V'} D_x(r, k) \right| &= \left| \bigcup_{(r, k) \in V^*(T)} D_x(r, k) \right| \\ &= |P(T) \setminus \{x\}| \geq \sum_{r \in T} \delta(r) = |V^*(T)| \geq |V'|. \quad \square \end{aligned}$$

Corollary 4-4:

A resource contention $\langle N, M, \{\rho_p\}_{p \in M}, \alpha \rangle$ is deadlock free if and only if for every $x \in M$ the set $M \setminus \{x\}$ and the family of subsets $\{D_x(r, k)\}_{(j, k) \in V}$ in Definition 4-1 has an SDR.

Corollary 4-4 says that a resource contention system is deadlock free if and only if after deleting any one process from the system, the discrepancy family $\{D_x(r, k)\}_{(j,k) \in V}$ has an SDR. This means that the discrepancies will all be absorbed by the rest of the processes, but there is at least one process (the deleted one) that can proceed. This deadlock free condition can therefore be tested by deleting one process at a time from M , then using the algorithm for testing the existence of an SDR. The best algorithm known for testing for SDR's is of order $O(N^{2.5})$ [Hopcroft & Karp 73], where $N = m + n$ (i.e., the sum of the number of processes and resource types). It follows that the deadlock free condition can be tested in order $O(N^{3.5})$ steps.

If the structure of $\langle N, M, \{\rho_p\}_{p \in M}, \alpha \rangle$ is such that there is more than one component, the following theorem will occasionally enable us to reduce further the number of steps in testing the condition.

Theorem 4-5:

A resource contention system $\langle N, M, \{\rho_p\}_{p \in M}, \alpha \rangle$ with components C_1, C_2, \dots, C_{c-1} is deadlock free if and only if

$$\langle R(C_k), C_k, \{\rho_p\}_{p \in C_k}, \alpha \downarrow R(C_k) \rangle \text{ are deadlock free for all } k, 0 \leq k < c. \quad (13)$$

($\alpha \downarrow R(C_k)$ denotes the restriction of α to $R(C_k)$.)

Proof:

By Theorem 4-1, (13) is equivalent to:

$$|P(T^j)| > \sum_{r \in T^j} \delta(r), \text{ for all nonempty } T^j \subseteq R(C_k), \quad 0 \leq k < c. \quad (14)$$

We show the equivalence of (14) and (10). Since every subset of $R(C_k)$ is a subset of N , (10) obviously implies (14). Assume (14) holds. Let $T \subseteq N$. We show

$$|P(T)| > \sum_{r \in T} \delta(r).$$

Consider the sets $T \cap R(C_k), 0 \leq k < c$, which form a partition of the set T . Similarly, the sets $P(T \cap R(C_k))$ are disjoint since $P(T \cap R(C_k)) \subseteq C_k$ and partition $P(T)$. For sets $P(T \cap R(C_k))$ that are empty, both sides of (15) are 0. We have, because of (14),

$$|P(T \cap R(C_k))| \geq \sum_{r \in (T \cap R(C_k))} \delta(r). \quad (15)$$

Since $P(T) \neq \phi$, there is at least one k for which $P(T \cap R(C_k))$ is nonempty. For these sets, (15) holds with strict inequality (by (14)). Adding (15) for all $k, 0 \leq k < c$, yields (10). \square

Theorem 4-5 shows that in order to prove a resource contention system to be free of deadlock we do not have to check (10) for all subsets of resource types, but we may instead restrict our checking to the subsets of all its components.

5. SIMPLE RESOURCE CONTENTION SYSTEMS

From now on we shall consider the following special case,

$$\begin{cases} \alpha(r) = 1 \text{ for all } r \in N \\ \rho_p : N \mapsto \{0, 1\} \end{cases}$$

Then

$$\delta(r) = \sum_{p \in M} \rho_p(r) - 1 = \sum_{p \in P(\{r\})} \rho_p(r) - 1$$

and so,

$$\delta(r) = |P(\{r\})| - 1. \quad (16)$$

Again, without loss of generality, we assume $\delta(r) > 0$.

Lemma 5-1:

Let $P(N)$ consist of a single component. Then for each $T \subseteq N, \phi \neq T \neq N$,

there exists an $x \in N \setminus T$ such that $|P(T \cup \{x\})| \leq |P(T)| + \delta(x)$.

Proof:

If $P(T) = P(N)$ any $x \in N \setminus T$ will do. We consider the case $P(N) \setminus P(T) \neq \emptyset$. $P(T)$ and $P(N) \setminus P(T)$ together constitute one component. Therefore, there must be a $q \in P(T)$ which is competing with some $p \in P(N) \setminus P(T)$, i.e., for some $x \in N$, $\rho_p(x) > 0$ and $\rho_q(x) > 0$. Since $\rho_p(x) > 0$ and $p \notin P(T)$, we have $x \in N \setminus T$. Since $q \in P(\{x\}) \cap P(T)$ we have

$$|P(\{x\}) \setminus P(T)| \leq |P(\{x\})| - 1. \quad (17)$$

However, by (16) and (17),

$$|P(T \cup \{x\})| = |P(T)| + |P(\{x\}) \setminus P(T)| \leq |P(T)| + \delta(x)$$

which proves the lemma. \square

Lemma 5-2:

A simple resource contention system $\langle N, M, \{\rho_p\}_{p \in M}, \alpha \rangle$ in which $P(N)$ consists of a single component is deadlock free if and only if

$$|P(N)| > \sum_{r \in N} \delta(r) \quad (18)$$

Proof:

Obviously, (10) implies (18). We show that (18) implies (10). (Note that (18) is a special case of (10) with $T = N$). We need to show for all other $T \subset N$, the same inequality holds. Assume for some maximal $T, T \neq N$, (10) does not hold, i.e.,

$$|P(T)| \leq \sum_{r \in T} \delta(r). \quad (19)$$

Then by Lemma 5-1, there exists an $x \in N \setminus T$ such that

$$|P(T \cup \{x\})| \leq |P(T)| + \delta(x). \quad (20)$$

Hence by (19) and (20),

$$\begin{aligned} \sum_{r \in T \cup \{x\}} \delta(r) &= \sum_{r \in T} \delta(r) + \delta(x) \\ &\geq |P(T)| + (|P(T \cup \{x\})| - |P(T)|) = |P(T \cup \{x\})|. \end{aligned}$$

However, this contradicts the maximality assumption in the choice of T and the lemma is proved. \square

Lemma 5-3:

Let $P(N)$ consist of a single component. Then

$$|P(N)| \leq \sum_{r \in N} \delta(r) + 1$$

Proof:

According to (16), for any $y \in N$, $|P(\{y\})| = \delta(y) + 1$. For $T = \{y\}$ we have proved

$$|P(T)| \leq \sum_{r \in T} \delta(r) + 1. \quad (21)$$

For $T \neq N$ there exists, by Lemma 5-1, an $x \in N \setminus T$ such that

$$|P(T \cup \{x\})| \leq |P(T)| + \delta(x)$$

With (21) this yields

$$|P(T \cup \{x\})| \leq \sum_{r \in T \cup \{x\}} \delta(r) + 1$$

By induction, (21) holds for all $T \subseteq N$, and in particular for $T = N$. \square

Corollary 5-4: [of Lemmas 5-2 and 5-3]

A simple resource contention system $\langle N, M, \{\rho_p\}_{p \in M}, \alpha \rangle$ in which $P(N)$ has one component is deadlock free if and only if

$$|P(N)| = \sum_{r \in N} \delta(r) + 1. \quad (22)$$

Theorem 5-5:

A simple resource contention system $\langle N, M, \{\rho_p\}_{p \in M}, \alpha \rangle$ in which $P(N)$ consists of c components is deadlock free if and only if

$$|P(N)| = \sum_{r \in N} \delta(r) + c. \quad (23)$$

Proof:

According to Theorem 4-5 the resource contention system is deadlock free if and only if for each component $C_k, \langle R(C_k), C_k, \{\rho_p\}_{p \in C_k}, \alpha \downarrow C_k \rangle$ is deadlock free, i.e., if and only if

$$\text{for all } 0 \leq i < c, |C_i| = \sum_{r \in R(C_i)} \delta(r) + 1$$

Since $|C_i| = \sum_{r \in R(C_i)} \delta(r) + 1$ the resource contention system is deadlock free if and only if

$$\sum_{i=0}^{c-1} |C_i| = \sum_{i=0}^{c-1} \sum_{r \in R(C_i)} \delta(r) + c$$

i.e., if and only if

$$|P(N)| = \sum_{r \in R(M)} \delta(r) + c. \quad (24)$$

Since we always assume $\delta(r) > 0$, $N = R(M)$, and therefore, (24) is equivalent to (23). \square

Example: "The Dining Chinese Philosophers"

This example of a simple resource contention system is a slight variation of Dijkstra's well-known problem of the five philosophers [Dijkstra 71]. There are n ($n \geq 2$) Chinese philosophers sitting at a round-table, each one with a plate in front of him. There are also n chopsticks, each chopstick placed between two adjacent plates. A philosopher either thinks or eats. When eating he uses the two chopsticks that are at the sides of his plate. In this case, the chopsticks are the shared resources. This is an example of a simple resource contention system. Under what circumstances is it deadlock free?

We first look at the case that all n philosophers are wishing to eat. N is the set of all chopsticks.

We have

$$|P(N)| = n, \quad c = 1 \text{ and}$$

$$\text{for all } r \in N, \quad \delta(r) = 1.$$

Hence $\sum_{r \in N} \delta(r) = n$. Since $|P(N)| \neq \sum_{r \in N} \delta(r) + 1$, by Theorem 5-5, the resource contention system is not deadlock free.

Now assume there are only k , $0 < k < n$, philosophers wishing to eat. Suppose there are l chopsticks which are between two adjacent plates of philosophers wishing to eat ($l < k$). Then

$$|P(N)| = k, \quad \sum_{r \in N} \delta(r) = l, \quad c = k - l.$$

Thus, since $|P(N)| = \sum_{r \in N} \delta(r) + c$, the resource contention is deadlock free.

6. CONCLUSIONS

In this paper we have been looking at resource contention systems. Rather than posing the question how the processes involved have to be scheduled in order to avoid deadlock have we postulated the scheduling and characterized all ensembles of processes that will avoid deadlock. For the scheduling we have chosen the simplest scheme we could think of: scheduling by mutual exclusion on the individual resources.

Comparing the general theorem with that for the special case of single resources we notice that the determination of absence of deadlock is much simpler in the latter case: testing only one equality suffices. In the general case, we have produced a polynomial algorithm of order $O(n^{3.5})$ to determine the absence of deadlock. The second version of the general theorem can further reduce the number of steps to be tested if the processes can be non-trivially partitioned into disjoint components.

REFERENCES

- [Dijkstra 68] Dijkstra, E. W.
Cooperating sequential processes.
In F.Genuys, editor, **Programming Languages**, pages 43–112.
Academic Press, New York, 1968.
- [Dijkstra 71] Dijkstra, E. W.
Hierarchical ordering of sequential processes.
Acta Informatica 1: 115–138, 1971.
- [Habermann 69] Habermann, A. N.
Prevention of system deadlocks.
Communications of ACM 12(7):373–377, 385, 1969.
- [Hopcroft & Karp 73] Hopcroft, J. E. and Karp, R. M.
An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs.
SIAM J. Comput. 2(4):225–231, 1973.
- [Lawler 76] Lawler, E. L.
Combinatorial Optimization: Networks and Matroids.
Holt, Rinehart and Winston, New York, 1976.
- [Mead & Conway 80] Mead, C. and Conway, L.
Introduction to VLSI Systems.
Addison Wesley, Reading, 1980.
- [Mirsky 71] Mirsky, L.
Transversal Theory
Academic Press, New York, 1971.
- [Ryser 63] Ryser, H. J.
Combinatorial Mathematics.
Mathematical Association of America, Buffalo, 1963.