On Quorum Systems for $(m, 1, k)$-Resource Allocation*

Yuh-Jzer Joung
joung@ccms.ntu.edu.tw
Department of Information Management
National Taiwan University
Taipei, Taiwan

Abstract

We present a problem called $(m, 1, k)$-*resource allocation* to model group mutual exclusion with bounded capacity. Specifically, the problem concerns the scheduling of a resource among m groups of processes. The resource can be used by at most k processes of the same group at a time, but no two processes of different groups can use the resource simultaneously. The problem reduces to group mutual exclusion when k is equal to the group size. We then generalize quorum systems for mutual exclusion to the problem. We show that the study of quorum systems for $(m, 1, k)$-resource allocation is closely related to some classical problems in combinatorics and in finite projective geometries. By applying the results there, we are able to obtain some optimal/near-optimal quorum systems. We also exploit the properties of these systems pertaining to distributed computing.

*Research supported in part by the National Science Council, Taipei, Taiwan, Grants NSC 91-2213-E-002-072.

1 Introduction

l-exclusion [6] and group mutual exclusion [13, 16, 10] generalize mutual exclusion in two orthogonal directions. In mutual exclusion, every two processes conflict with each other when both are using a shared resource at the same time. So the resource must be accessed in an exclusive style. In l-exclusion, conflicts occur only when more than l processes are using the resource simultaneously. So up to l processes can concurrently use the resource. In group mutual exclusion, processes are divided into groups, and a conflict occurs only when two processes of different groups are attempting to use the shared resource. Any number of processes are allowed to use the resource concurrently so long as they belong to the same group.

Some applications, however, have properties captured by both l-exclusion and group mutual exclusion. For example, consider a multi-head CD jukebox in which up to l discs can be loaded for access simultaneously. When a disc is loaded, users interested in the disc can concurrently access the disc. By defining the set of users interested in the same disc as a group, we see that up to l groups of users can concurrently use the CD jukebox, and for each group, any number of users can concurrently access the disc they are interested in. Furthermore, to guarantee some quality of service, some system may impose a limit on the number of processes that can concurrently access a disc. A similar problem also occurs in wireless communication, where the communication channel consists of l sub-channels, each of which can allow up to k users to communicate.

To model the above resource allocation problems, we propose a family of problems called (m, l, k)-*resource allocation*. They concern the scheduling of l copies of a resource among m groups of processes. Each copy of the resource can be used by at most k processes of the same group at a time, but no two processes of different groups can use the same copy simultaneously. We replace k with ∞ when there is no limit on the number of processes that may concurrently access a copy of the resource. By assigning different values to the parameters, (m, l, k)-resource allocation can be used to model many existing resource allocation problems, as well as new problems illustrated above. For example, both $(1, 1, 1)$-resource allocation and $(n, 1, 1)$-resource allocation can be used to model n-process mutual exclusion, while both $(1, 1, l)$-resource allocation and $(n, l, 1)$-resource allocation are able to model l-exclusion among n processes. Moreover, $(m, 1, \infty)$-resource allocation reduces to group mutual exclusion.

To solve (m, l, k)-resource allocation, the concept of quorum systems comes to our mind, as they have been successfully applied to mutual exclusion (e.g., [20, 1, 27]) and l-exclusion (e.g., [15, 18, 23]) to reduce system load and to cope with site failures. Informally, a quorum system is a set of sets of processes with some intersection property. Each set in a quorum system is called a *quorum*. The quora provide some guard against conflicts in using a shared resource. To acquire the resource, a process must obtain permission from every member of a quorum. The quorum is usually chosen arbitrarily from the system, and a quorum member can give permission to only one process at a time. So for mutual exclusion, every two quora in a quorum system must intersect; while for l-exclusion, any collection of $l + 1$ quora must contain a pair of intersecting quora. Note that a quorum usually involves only a subset of the processes in the underlying network of computation. So system load is reduced because when a process requests the resource, not all processes are involved in the scheduling. Failure resilience is increased as well because the resource can be accessed so long as not all quora are hit (a quorum is *hit* if one of its members has failed).

For a distinguishing purpose, we call quorum systems for mutual exclusion and l-exclusion 1-*coteries* and l-*coteries*, respectively. We can see that 1-coteries are not general enough to solve $(m, 1, k)$-resource allocation. A direct use of them would result in a degenerate solution in which only one copy of the resource can be used at a time and only one process can use the copy. For l-coteries, it is not clear how they can be applied to $(m, 1, k)$- or (m, l, k)-resource allocation.

The goal of this research is therefore to lay some groundwork for quorum systems for (m, l, k)-resource allocation. This paper presents the results for the $l = 1$ case. This case corresponds to group mutual exclusion with and without bounded capacity, i.e., the $(m, 1, k)$-resource allocation problem and the $(m, 1, \infty)$-resource allocation problem. We begin by establishing some basic and

general results for quorum systems for $(m, 1, k)$-resource allocation, based on which quorum systems for the more general case, (m, l, k)-resource allocation, can be understood and constructed. We show that the study of quorum systems for $(m, 1, k)$-resource allocation is related to some classical problems in combinatorics and in finite projective geometries. By applying the results there, we are able to obtain some quorum systems that can provide optimal/near-optimal concurrency.

The paper is organized as follows. Section 2 defines quorum systems for $(m, 1, k)$-resource allocation, and presents composition methods for the quorum systems. Section 3 presents some optimal/near-optimal quorum systems for $(m, 1, k)$-resource allocation. Section 4 discusses failure resilience. Conclusions and future work are offered in Section 5.

2 Fundamentals

2.1 Basic Definitions

***Definition* 2.1** *Let P be a finite set of elements. An $(m, 1)$-**coterie** over P, is a tuple $\mathfrak{C} = (C_1, \ldots, C_m)$, where each $C_i \subseteq 2^P$, such that the following conditions are satisfied:*

intersection: $\forall\, 1 \leq i \neq j \leq m, \forall\, Q_1 \in C_i, \forall\, Q_2 \in C_j : Q_1 \cap Q_2 \neq \emptyset$.

minimality: $\forall\, 1 \leq i \leq m, \forall\, Q_1, Q_2 \in C_i : Q_1 \not\subset Q_2$.[1]

*We call each C_i a **cartel**, and each $Q \in C_i$ a **quorum**.*

Note that by the intersection condition, no cartel can contain an empty set if $m > 1$; and by the minimality condition, no cartel can contain $\emptyset$ unless $\mathfrak{C} = (\{\emptyset\})$—the ***empty*** $(1, 1)$-coterie.

***Definition* 2.2** *The **degree** of a cartel C, denoted as $\mathsf{deg}(C)$, is the maximum number of pairwise disjoint quora in C. The **degree** of an $(m, 1)$-coterie $\mathfrak{C}$, denoted as $\mathsf{deg}(\mathfrak{C})$, is defined to be the minimum degree of its cartels. $\mathfrak{C}$ is of **uniform degree** k if all its cartels have the same degree k.*

We shall use $(m, 1, k)$-***coteries*** to refer to $(m, 1)$-coteries of uniform degree k. For example, $\big(\{\{1, 2\}, \{3, 4\}\}, \{\{1, 3\}, \{2, 4\}\}, \{\{2, 3\}, \{1, 4\}\}\big)$ is a $(3, 1, 2)$-coterie over $\{1, 2, 3, 4\}$.

For comparison, we present the definition of 1-coteries used for standard mutual exclusion. Formally, a 1-coterie over P is a set $C \subseteq 2^P$ of quora satisfying the following requirements:

intersection: $\forall\, Q_1, Q_2 \in C : Q_1 \cap Q_2 \neq \emptyset$.

minimality: $\forall\, Q_1, Q_2 \in C, Q_1 \not\subset Q_2$.

So each cartel in an $(m, 1, 1)$-coterie is a 1-coterie. Conversely, a 1-coterie C can be straightforwardly converted to an $(m, 1, 1)$-coterie as follows:

$$\mathfrak{T}_m(C) = (C, \ldots, C)$$

To see how $(m, 1, k)$-coteries connect to the $(m, 1, k)$-resource allocation problem, let P be the set of processes in the problem, and let $\mathfrak{C} = (C_1, \ldots, C_m)$ be an $(m, 1, k)$-coterie over P. Each group of processes in P is assigned a cartel so that when a process of group j wishes to use the shared resource, it must acquire an arbitrary quorum $Q \in C_j$ by locking every member of the quorum. Suppose a quorum member can be locked by one process at a time. Then the intersection property of $\mathfrak{C}$ ensures that no processes of different groups can access the resource simultaneously. The degree of $\mathfrak{C}$ ensures that k processes (and no more) can concurrently access the resource. The minimality property is used rather to enhance efficiency. As is easy to see, if $Q_1 \subset Q_2$, then a process that can lock Q_2 can also lock Q_1.

The above quorum-acquiring concept is essentially from Maekawa's well-known algorithm [20] for standard mutual exclusion. The number of messages needed for a process to access a shared

[1] For notational simplicity, throughout the paper, unless stated otherwise, by "$\forall\, a_1, \ldots, a_k \in S$" we assume that the k elements $a_1, \ldots, a_k$ are distinct. Similarly for "$\exists\, a_1, \ldots, a_k \in S$".

resource is $O(|Q|)$, where Q is the quorum the process chooses. The minimum synchronization delay (i.e., the minimum time, in message transmission delay, for a process to access a shared resource) is 2. So by using Maekawa's algorithm, an $(m, 1, k)$-coterie corresponds directly to a distributed solution to $(m, 1, k)$-resource allocation. Also notice that the coterie is usually formed from the processes in the resource allocation problem.

By definition, an $(m, 1, k)$-coterie over an n-set guarantees that every cartel contains at least one unhit quorum even if $k - 1$ processes have failed. So high degree coteries provide better protection against faults. However, every quorum in an $(m, 1, k)$-coterie must have size at least k (unless $m = 1$). So the higher the degree, the larger the quorum size. Because every cartel in an $(m, 1, k)$-coterie contains k pairwise disjoint quora, and because every quorum has size at least k, the system must contain at least k^2 processes. So k is bounded by $\sqrt{n}$:

Theorem 2.1 *For every $(m, 1, k)$-coterie $\mathfrak{C}$ over an n-set, $k \leq \sqrt{n}$ if $m > 1$.*

The above theoretical limit poses a problem to quorum-based solutions to $(m, 1, k)$-resource allocation with large value of k, in particular, the $(m, 1, \infty)$-resource allocation (group mutual exclusion). We shall return to this problem after we have constructed some $(m, 1, k)$-coteries with the maximum possible k in Section 3.3.

Some $(m, 1, k)$-coteries have a structure that is not only mathematically beautiful, but has also an important meaning in distributed computing. Below we define this structure. Section 3 will present construction of the coteries.

***Definition* 2.3** *Let $\mathfrak{C} = (C_1, \ldots, C_m)$ be an $(m, 1, k)$-coterie over P. $\mathfrak{C}$ is **balanced** if all cartels have the same size. $\mathfrak{C}$ is **regular** if all elements in P are involved in the same number of quora; i.e., $\forall p, q \in P : |n_p| = |n_q|$, where n_p is the multiset $\{Q \mid \exists 1 \leq i \leq m : Q \in C_i \text{ and } p \in Q\}$, and similarly for n_q. $\mathfrak{C}$ is **uniform** if all quora have the same size.*

For example, the $(3, 1, 2)$-coterie $\big(\{\{1, 2\}, \{3, 4\}\}, \{\{1, 3\}, \{2, 4\}\}, \{\{2, 3\}, \{1, 4\}\}\big)$ is balanced, uniform, and regular.

When used for $(m, 1, k)$-resource allocation, a balanced $(m, 1, k)$-coterie ensures that each group has an equal chance in competing for the resource. The regular property ensures that each process shares the same responsibility. The uniform property ensures that the number of messages needed per access to the resource is independent of the quorum a process chooses. Thus the three properties are desirable for a truly distributed algorithm for the problem [20].

2.2 Compositions of Coteries

In this section we provide some useful lemmas for constructing "large" $(m, 1, k)$-coteries from "small" ones. We begin with a composition that takes an $(m, 1, k_1)$-coterie over an n_1-element set, and an $(m, 1, k_2)$-coterie over an n_2-element set, and then constructs an $(m, 1, k_1 \cdot k_2)$-coterie over an $n_1 \cdot n_2$-element set. The composition is borrowed from MacNeish's technique for composing Latin Squares [19]. The same technique has also been used in the context of standard quorum systems [24, 21].

The composition needs the following notation. Let $\mathfrak{C} = (C_1, \ldots, C_m)$ be an $(m, 1, k)$-coterie over an n-element set $\{1, \ldots, n\}$. Then, $\mathfrak{C}(P)$ denotes the $(m, 1, k)$-coterie $\mathfrak{C}$ obtained by replacing $\{1, \ldots, n\}$ with another n-element set P (using some arbitrary bijection between $\{1, \ldots, n\}$ and P). Note that $\mathfrak{C}(P)$ is essentially the same as $\mathfrak{C}$ subject to renaming of the elements. Similarly, quorum $Q(P)$ corresponds to the quorum obtained from Q by replacing the elements in Q with that in P.

Lemma 2.2 *Let $P_1, \ldots P_r$ be r pairwise disjoint sets, each of size s. Let $\mathfrak{C} = (C_1, \ldots, C_m)$ be an $(m, 1)$-coterie over an r-element set $\{1, \ldots, r\}$, and $\mathfrak{D} = (D_1, \ldots, D_m)$ be an $(m, 1)$-coterie over an s-element set $\{1, \ldots, s\}$. Let $\mathfrak{C} \otimes \mathfrak{D} = (E_1, \ldots, E_m)$, where each cartel E_i, $1 \leq i \leq m$, is defined*

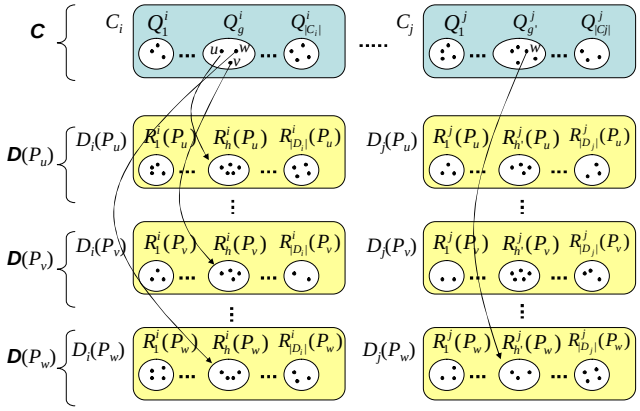

Figure 1: Illustration of the composition in Lemma 2.2.

as follows: Let $C_i = \{Q_1^i, \ldots, Q_{|C_i|}^i\}$, *and let* $D_i = \{R_1^i, \ldots, R_{|D_i|}^i\}$. *Then* $E_i = \{S_{g,h}^i \mid 1 \le g \le |C_i|, 1 \le h \le |D_i|\}$, *where*

$$S_{g,h}^i = \bigcup_{j \in Q_g^i} R_h^i(P_j)$$

Then, $\mathfrak{C} \otimes \mathfrak{D}$ *is an* $(m,1)$*-coterie over* $\bigcup_{1 \le j \le r} P_j$, *and* $\mathsf{deg}(E_i) = \mathsf{deg}(C_i) \cdot \mathsf{deg}(D_i)$.

Proof. See Appendix. □

Figure 1 illustrates the composition. To see how a quorum $S_{g,h}^i \in E_i$ is constructed, consider the quorum $Q_g^i \in C_i$ which consists of three nodes u, v, w. Then $S_{g,h}^i$ is the union of $R_h^i(P_u), R_h^i(P_v)$, and $R_h^i(P_w)$. To see the intersection property of $\mathfrak{C} \otimes \mathfrak{D}$, let $S_{g',h'}^j$ be a quorum in E_j, $j \ne i$. By the intersection property of $\mathfrak{C}$, there is a common node between C_i and C_j, in this figure, w. Then $S_{g',h'}^j$ contains all nodes in $R_{h'}^j(P_w)$. By the intersection property of $\mathfrak{D}(P_w)$, there is a common node between $R_h^i(P_w)$ and $R_{h'}^j(P_w)$, say, x. So $x \in S_{g,h}^i \cap S_{g',h'}^j$.

The above composition allows one to build up a large degree $(m,1,k)$-coterie from ones with smaller degree. The next two compositions we will present build up a coterie with larger number of cartels. The compositions need the following lemma.

Lemma 2.3 *Let* $P_1, \ldots P_r$ *be* r *pairwise disjoint sets, and let* $P = \bigcup_{1 \le j \le r} P_j$. *Let* $\mathfrak{C}^j = (C_1^j, C_2^j, \ldots, C_m^j)$ *be an* $(m,1)$*-coterie over* P_j, $1 \le j \le r$. *Let* u *and* v *be two vectors in* $\{0,1,\ldots,m\}^r$. *Define* $C_u \subset 2^P$ *as follows:*

$$C_u = \left\{ \bigcup_{1 \le j \le r,\ u[j] \ne 0} Q_{u[j]} \;\middle|\; Q_{u[j]} \in C_{u[j]}^j \right\} \tag{1}$$

That is, each set in C_u *is of the form:* $Q_{u[1]} \cup Q_{u[2]} \cup \ldots \cup Q_{u[r]}$, *where each* $Q_{u[j]}$ *is an arbitrary quorum taken from* $C_{u[j]}^j$ *if* $u[j]$ *(the* j*th element of vector* u*) is not zero, or is* $\emptyset$ *otherwise. Similarly for* C_v. *If there is an* i *such that* $u[i] \ne 0, v[i] \ne 0$, *and* $u[i] \ne v[i]$ *(i.e.,* u *and* v *have a distinct nonzero* i*th element), then* C_u *and* C_v *satisfy the following two conditions:*

intersection: $\forall S \in C_u, \forall T \in C_v : S \cap T \ne \emptyset$.

minimality: $\forall S, T \in C_u : S \not\subset T$. *Similarly for* C_v.

Moreover, $\mathsf{deg}(C_u) = \min\{\mathsf{deg}(C_{u[j]}^j) \mid u[j] \ne 0, 1 \le j \le r\}$. *Similarly for* C_v.

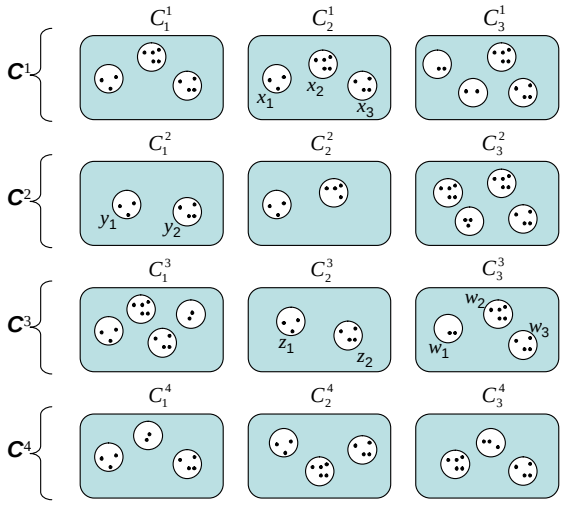

Figure 2: Illustration of the composition in Lemma 2.3.

Proof. See Appendix. □

To illustrate the lemma, consider Figure 2. Assume that $u = [2, 1, 2, 0]$. Then C_u consists of the following 12 sets: $x_i \cup y_j \cup z_k$, $i = 1, 2, 3$, $j = 1, 2$, and $k = 1, 2$. To see the intersection property, let $v = [2, 0, 3, 0]$. Then, the third elements of u and v are both nonzero and distinct. We see that every set in C_u must contain all the nodes in z_1, or all the nodes in z_2; and every set in C_v must contain all the nodes in w_1, w_2, or w_3. z_i's and w_j's are quora from different cartels of the same coterie $\mathfrak{C}^4$. By the intersection property of $\mathfrak{C}^4$, every z_i intersects every w_j. So every set in C_u intersects every set in C_v.

Let us say that two vectors u and v of size r 'hit' if they have a distinct nonzero ith element for some $i \leq r$. Lemma 2.3 then implies that given a collection of q pairwise hit vectors, we can construct a $(q, 1)$-coteries over P. In particular, if the vectors are from $\{1, \ldots, m\}^r$, then any two different vectors must hit. So we immediately have the following lemma:

Lemma 2.4 *Let $\{P_1, \ldots P_r\}$ be a partition of P, i.e., $P = \bigcup_{1 \leq j \leq r} P_j$. Let $\mathfrak{C}^j = (C_1^j, C_2^j, \ldots, C_m^j)$ be an $(m, 1, k)$-coterie over P_j, $1 \leq j \leq r$. Let u_i, $1 \leq i \leq q$, be q different vectors in $\{1, \ldots, m\}^r$, and C_{u_i} be defined as in Equation (1). Then $\mathfrak{C} = (C_{u_1}, C_{u_2}, \ldots, C_{u_q})$ is a $(q, 1, k)$-coterie over P.*

Another way to construct a collection of q pairwise hit vectors is that if two vectors u and v from $\{0, 1, \ldots, m\}^r$ both have $\lfloor r/2 \rfloor + 1$ nonzero elements, then they must have a nonzero element in the same field. By requiring their elements in the field to be distinct, we can obtain such a collection. The following lemma follows from this observation:

Lemma 2.5 *Let $\{P_1, \ldots P_r\}$ be a partition of P. Let $\mathfrak{C}^j = (C_1^j, C_2^j, \ldots, C_m^j)$ be an $(m, 1, k)$-coterie over P_j, $1 \leq j \leq r$. Let u_i, $1 \leq i \leq q$ be q different vectors in $\{0, 1, \ldots, m\}^r$ such that $\left|\{g \mid u_i(g) \neq 0\}\right| = \lfloor r/2 \rfloor + 1$, and for any two different vectors u_i and u_j, there is an h such that $u_i(h) \neq 0$, $u_j(h) \neq 0$, and $u_i(h) \neq u_j(h)$. Let C_{u_i} be defined as in Equation (1). Then $\mathfrak{C} = (C_{u_1}, C_{u_2}, \ldots, C_{u_q})$ is a $(q, 1, k)$-coterie over P.*

The difference between the compositions in the above two lemmas is that in Lemma 2.4, every quorum S in C_{u_i} is formed by r sub-quora, one from $C^1_{u[1]}$, one from $C^2_{u[2]}$, ..., and one from $C^r_{u[r]}$. Thus, if each sub-quorum has size s, then S has size $r \cdot s$. In Lemma 2.5, the size of a quorum in

C_{u_i} is reduced approximately by half by letting S be composed by only $\lfloor r/2 \rfloor + 1$ sub-quora. Note, however, that the maximum number of cartels one can obtain in the first composition is m^r, while in the second composition the number is reduced to $O(m^{r/2})$.

3 Constructions of $(m, 1, k)$-coteries

Recall Theorem 2.1 that the degree of $(m, 1)$-coteries over an n-set is bounded by $\sqrt{n}$. In this section we show that the bound is tight by constructing an $(m, 1, \sqrt{n})$-coterie over an n-set. We also show that the coterie is balanced, uniform, and regular. Based on this coterie and the composition methods presented in the previous section, we show how other coteries with near-optimal degree can be constructed. Note that the existence of an $(m, 1, k)$-coterie implies the existence of $(m, 1, k')$-coteries for all $k' < k$ (and the existence of $(m', 1, k)$-coteries for all $m' < m$ as well). This is because for every cartel C in an $(m, 1, k)$-coterie, removing any subset of quora from C does not affect the intersection and minimality properties. So we start the construction of $(m, 1, k)$-coteries with some optimal degree.

3.1 $(m, 1)$-coteries of Maximal Degree

The study of $(m, 1, \sqrt{n})$-coteries over an n-set is related to the study of finite projective geometries (see, e.g., [25]). The following definition and theorem can be found in the literature:

***Definition* 3.1** *An **affine plane** is an ordered pair $(\mathcal{P}, \mathcal{L})$, where $\mathcal{P}$ is a nonempty set of elements called **points**, and $\mathcal{L}$ is a nonempty collection of subsets of $\mathcal{P}$ called **lines** satisfying the following properties:*

- *Every two points lie on exactly one line. (A point i **lies** on a line L iff $i \in L$).*
- *Given a line L and a point i not on L, there is exactly one line L' such that i is on L', and L and L' are **parallel** (i.e., $L \cap L' = \emptyset$).*
- *Each line has at least two points, and there are at least two lines.*

If each line of an affine plane contains exactly n points, the plane is said to have ***order*** n.

Theorem 3.1 *An affine plane $(\mathcal{P}, \mathcal{L})$ of order n has the following properties:*

- *$\mathcal{P}$ has n^2 points.*
- *$\mathcal{L}$ has $n^2 + n$ lines.*
- *Each point is on $n + 1$ lines.*
- *$\mathcal{L}$ can be partitioned into $n+1$ classes such that each class contains n parallel lines, and every two lines of different classes intersect.*

For example, let the nine points $1, \ldots, 9$ be arranged as follows:

$$\begin{array}{ccc} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{array}$$

Then we can construct an affine plane of order 3 consisting of 4 classes $C_1, \ldots, C_4$, as follows:

$$\begin{aligned} C_1 &= \big\{\{1,4,7\},\ \{2,5,8\},\ \{3,6,9\}\big\} \\ C_2 &= \big\{\{1,2,3\},\ \{4,5,6\},\ \{7,8,9\}\big\} \\ C_3 &= \big\{\{3,5,7\},\ \{1,6,8\},\ \{2,4,9\}\big\} \\ C_4 &= \big\{\{1,5,9\},\ \{3,4,8\},\ \{2,6,7\}\big\} \end{aligned}$$

To picture this, C_1 corresponds to the three vertical lines, C_2 corresponds to the three horizontal lines, and C_3 and C_4 correspond to the three "rounded" lines with slope 1 and -1, respectively.

It is known that an affine plane of order n exists if n is a power of a prime. So Theorem 3.1 immediately implies the following.

Theorem 3.2 *Let $n = p^{2k}$, where p is a prime, and k is a positive integer. Then there is an $(m, 1, \sqrt{n})$-coterie $\mathfrak{C}$ over an n-set for every $m \leq \sqrt{n} + 1$. In particular, $\mathfrak{C}$ is balanced, uniform, and regular, and each quorum has size $\sqrt{n}$.*

By Theorem 2.1, the above construction obtains $(m, 1)$-coteries with optimal degree. Moreover, by Theorem 2.2, given $n = p_1^{2c_1} \ldots p_l^{2c_l}$, and $m \leq \min\{p_1^{c_1}, \ldots, p_l^{c_l}\} + 1$, where $p_1, \ldots, p_l$ are primes and $c_1, \ldots, c_l$ are positive integers, a degree-optimal $(m, 1)$-coterie over an n-element set can also be constructed as well. The construction of finite affine planes (and thus the construction of $(m, 1)$-coteries with optimal degree) can be found in finite projective geometries.

Note that in the above construction m is limited to $\sqrt{n} + 1$. By a combinatorial analysis, we can prove that the bound is also tight.

Theorem 3.3 *Let $\mathfrak{C} = (C_1, \ldots, C_m)$ be an $(m, 1, k)$-coterie over an n-set. If $k = \sqrt{n}$, then $m \leq k + 1$.*

For $m > \sqrt{n} + 1$, we propose deterministic and randomized constructions. The deterministic constructions make use of the composition methods presented in Section 2.2. For example, by Lemma 2.4 and Theorem 3.2 we immediately have the following:

Theorem 3.4 *Given $|P| = r \cdot k^2$, where k is a power of a prime and r an integer, there is an $(m, 1, k)$-coterie $\mathfrak{C}$ over P for every $m \leq (k + 1)^r$.*

So, taking r to be a constant, we can construct an $(m, 1)$-coterie that is near optimal in degree by only a constant factor of $\sqrt{r}$ for any m up to $\approx r^{-\frac{r}{2}} n^{\frac{r}{2}}$. In particular, in the $(m, 1, \infty)$-resource allocation problem, when groups are disjoint, m cannot be greater than the total number of processes n, and $m \leq n/2$ if each group consists of more than one process. In this case, we can let $r = 2$ and obtain an $(m, 1, \sqrt{n/2})$-coterie. Each quorum in the coterie has size only $\sqrt{2n}$.

When groups are not disjoint, m can be larger than n in the $(m, 1, \infty)$-resource allocation problem. In this case, the above construction gives an $(m, 1)$-coterie of degree $O(\sqrt{\frac{n \log n}{\log m}})$. Notice that the above construction requires s to be a power of a prime. However, by the distribution of primes in number theory (see, e.g., [11]), for any given $\epsilon > 0$, there is a prime p such that $x < p \leq (1 + \epsilon)x$ for all $x > x_0$, where x_0 is some constant depending on ϵ. So the construction can be generalized to arbitrary n when n is sufficiently large.

Similarly, Lemma 2.5 and Theorem 3.2 can be used to obtain some near-optimal constructions.

The above constructions require some knowledge about primes and projective geometries. In contrast, the following randomized construction is much simpler and requires less restriction on n and m.

Theorem 3.5 *Given any n, m, and k such that n/k is an integer, the following construction guarantees to generate, with probability one, an $(m, 1, k)$-coterie over $P = \{1, \ldots, n\}$ if $k \leq \sqrt{\frac{n}{2 \log(nm)}}$:*

1. **repeat**
2. *randomly choose m sets $C_1, \ldots, C_m$, each of which is a random partition of P of k equal-size parts;*
3. **until** $\mathfrak{C} = \{C_1, \ldots, C_m\}$ *is an $(m, 1)$-coterie over P;*
4. *output $\mathfrak{C}$;*

Proof. See Appendix. □

Theorem 3.4 constructs an $(m, 1)$-coterie of a fixed degree determined by n. In contrast, the degree of the $(m, 1)$-coterie constructed in Theorem 3.5 is part of the input parameters. Thus, although an $(m, 1, k')$-coterie can be obtained from an $(m, 1, k)$-coterie for any given $k' < k$ (by removing additional quora), the construction in Theorem 3.5 provides a more direct way in constructing an $(m, 1)$-coterie of a given degree.

3.2 Related Work in Combinatorics

As can be seen from the previous section, there is a close connection between $(m, 1, k)$-coteries and combinatorial designs. In fact, the structure of $(m, 1, k)$-coteries can be simplified by additionally requiring quora within a cartel to be pairwise disjoint. Each cartel then is simply a partition of P—the set of all processes. The intersection property then requires that any two blocks of different partitions intersect. In combinatorics, two partitions with such an intersection property are said to be *qualitatively independent*.[2] So by restricting an $(m, 1, k)$-coterie to be a collection of pairwise qualitatively independent partitions, many interesting results in combinatorial designs can be applied to $(m, 1, k)$-coteries, and vice versa.

For example, Rényi raised the following problem [31]: given n and k, what is the maximal size $N(n, k)$ of a collection of pairwise qualitatively independent k-partitions of an n-element set?[3] The exact value of $N(n, 2)$ has been determined [31]

$$N(n, 2) = \binom{n-1}{\lfloor \frac{1}{2}n \rfloor - 1}$$

For $k > 2$, however, the problem becomes very complex, and only bounds have been established. Specifically, Poljak and Rödl [29] established the following lower bound

$$N(n, k) \geq (k+1)^{\frac{n}{k^2}}$$

for k that is a power of a prime. As commented in the previous section on the distribution of primes, the lower bound for general k is not much different from the above. The upper bound established by Poljak and Tuza [30] uses Bollobás' inequality [3], and has the following value

$$N(n, k) \leq \frac{1}{2}\binom{\lfloor \frac{2n}{k} \rfloor}{\lfloor \frac{n}{k} \rfloor} = O(4^{\frac{n}{k}}/\sqrt{n})$$

As can be seen, there is a gap between the two bounds. Although later results have improved these two bounds [28, 17, 9], the difference is not much.

The bounds on $N(n, k)$ certainly highlight the relationship between the degree of an $(m, 1)$-coterie over P, the given m, and the size of P. Specifically, from the above bounds we see that given arbitrary n and m, an $(m, 1)$-coterie over an n-element set of degree $O(\sqrt{\frac{n \log n}{\log m}})$ can be constructed, but no construction is possible for degree higher than $O(\frac{n}{\log \sqrt{n} + \log m})$. The construction presented in Theorem 3.4 achieves the lower bound. In fact, our construction and the construction presented by Poljak and Rödl in establishing the lower bound of $N(n, k)$ are both based on MacNeish's method for composition of Latin Squares [19]. Any improvement on the construction will also improve the lower bound on $N(n, k)$, and vice versa.

[2] This definition has an intuitive meaning in probability theory: two partitions can be generated by two independent random variables if and only if they are qualitatively independent [31].

[3] This is a typical problem in the research of *extremal sets*. The research is stimulated by Sperner's theorem [32] on the maximum possible size of a family of pairwise unrelated (with respect to inclusion) subsets of a finite subset. See [7] for a survey.

	FPP + Maekawa_M	AP + Maekawa	AP + Maekawa_M
$k \leq \sqrt{n}$	$O(\sqrt{n} \cdot k)$	$O(\sqrt{n})$	–
$k > \sqrt{n}$	$O(\sqrt{n} \cdot k)$	–	$O(k)$

Table 1: Comparison of 1-coteries vs. $(m, 1)$-coteries for $(m, 1, k)$-resource allocation. 'AP' denotes the affine plane coterie.

Note that although requiring quora within a cartel to be pairwise disjoint helps connect the problem to combinatorial designs, the restriction also limits the availability of quora. For example, the $(2, 1, 2)$-coterie $(\{\{1, 2\}, \{3, 4\}\}, \{\{1, 3\}, \{2, 4\}\})$ can be 'expanded' to $(\{\{1, 2\}, \{3, 4\}, \{1, 4\}, \{2, 3\}\}, \{\{1, 3\}, \{2, 4\}\})$. The latter provides more quora, and so provides a better protection against faults. Failure resilience is an important metric in quorum systems. So in the definition of $(m, 1, k)$-coteries we generally do not require quora within a cartel to be pairwise disjoint. Section 4 will address the failure resilience issue.

3.3 $(m, 1, k)$-coteries versus 1-coteries

As noted in Section 2, by using Maekawa's algorithm, an $(m, 1, k)$-coterie corresponds directly to a distributed solution to the $(m, 1, k)$-resource allocation problem, and the coterie is formed from the processes in the problem. Assume that there are n processes in the problem. Theorem 2.1, however, says that we cannot construct an $(m, 1, k)$-coterie out of the n processes for any $k > \sqrt{n}$ (unless $m = 1$). This means that a different approach needs to be used to solve the problem when $k > \sqrt{n}$. We discuss two of them.

The first approach introduces auxiliary processes to act as quorum nodes. Since k cannot be greater than n, in the worst case (when $k = n$) we can add $n^2 - n$ auxiliary processes to construct an $(m, 1, n)$-coterie over these n^2 actual and auxiliary processes, and then use Maekawa's algorithm to solve the problem. The message complexity is $O(n)$, and the space overhead (for the auxiliary processes) is $O(n^2)$. Note that we can let the n actual processes act also as auxiliary processes, so that no extra physical process is added to the system.

The second approach is to adopt a different quorum-acquiring algorithm. Two algorithms, Maekawa_M and Maekawa_S, have been proposed in [14]. Both algorithms are a modification of Maekawa's algorithm, but they remove the restriction that a quorum node can be locked by only one process at a time. So, multiple processes (of the same group) may all acquire a quorum simultaneously, regardless of whether their quora intersect or not. Maekawa_M imposes no order on the locking sequence within a quorum, while a global ordering is used in Maekawa_S. As a result, the two algorithms trade off between message complexity and synchronization delay. Maekawa_M has message complexity $O(c \cdot k/d)$ and minimum synchronization delay 2, while Maekawa_S has message complexity and minimum synchronization delay both of $O(c)$, where c is the quorum size and d is the degree of the $(m, 1)$-coteries. When the affine plane $(m, 1, \sqrt{n})$-coterie constructed in Theorem 3.2 is used, Maekawa_M has message complexity $O(n)$, while Maekawa_S has message complexity and minimum synchronization delay both of $O(\sqrt{n})$.

At this point one may have observed that 1-coteries can also be used in Maekawa_M and Maekawa_S to solve $(m, 1, k)$-resource allocation. A natural question then is whether $(m, 1, k)$-coteries are beneficial over 1-coteries. The answer depends on applications. If applications need fast response time, then Maekawa_M should be chosen. If 1-coteries are used in the algorithm, then as commented above, the message complexity is $O(c \cdot k)$, because 1-coteries have degree one when converted to $(m, 1)$-coteries. For example, the FPP 1-coterie in [20] (which also supports a truly distributed solution) has $O(\sqrt{n} \cdot k)$ complexity. On the other hand, when $(m, 1, k)$-coteries are used, then Maekawa's original algorithm can be used for k up to $\sqrt{n}$. In this case, the affine plane coterie yields $O(\sqrt{n})$ message complexity. For $k > \sqrt{n}$, Maekawa_M together with the affine plane coterie yields $O(n)$ message complexity. Table 3.3 summarizes the comparison.

If applications can tolerate long synchronization delay, then Maekawa_S can be used. In this case, there is not much difference in choosing between $(m, 1, k)$-coteries and 1-coteries. However, 1-coteries have been extensively studied in the literature, and they have been optimized in many possible ways; while $(m, 1, k)$-coteries are new concept and many of their properties remain to be exploited. We hope that our studies in the paper can initiate further research on a more general type of quorum systems.

4 Failure Resistance vs. Full Distributedness

Since there are many $(m, 1, k)$-coteries over a given set, some criteria are needed to evaluate them. In this section we compare $(m, 1, k)$-coteries based on their failure resilience. The notion we shall be using is *dominance*, which determines if an $(m, 1, k)$-coterie can be extended to tolerate more failures. We then present some methods to determine dominance, and to convert dominated $(m, 1, k)$-coteries to nondominated ones. Finally, we show that $(m, 1, k)$-coteries constructed from affine planes, although are very suitable for a fully distributed implementation of the resource allocation problem, are unfortunately dominated. We then propose a remedy to cope with the conflict between full distributedness and failure resilience.

4.1 Dominance

The notion of *dominance* was first proposed by Garcia-Molina and Barbara [8] to compare the failure resilience of 1-coteries. Intuitively, a 1-coterie C *dominates* another 1-coterie D if whenever a quorum in D can survive some failures, then some quorum in C can certainly survive as well. Thus in this sense C is said to be superior to D because C provides more protection against failures. Formally, a 1-coterie D is nondominated if there is no other 1-coterie C such that $\forall Q \in D, \exists R \in C : R \subseteq Q$. Similarly, domination of $(m, 1, k)$-coteries can be defined as follows:

***Definition* 4.1** *An $(m, 1, k)$-coterie $\mathfrak{C} = (C_1, \ldots, C_m)$ over P **dominates** an $(m, 1, k')$-coterie $\mathfrak{D}$ over P if*

1. $\mathfrak{C} \neq \mathfrak{D}$,
2. $\forall 1 \leq i \leq m, \forall Q \in D_i, \exists R \in C_i : R \subseteq Q$.

*$\mathfrak{D}$ is **(strongly) nondominated** if there is no $(m, 1, k)$-coterie that dominates $\mathfrak{D}$. $\mathfrak{D}$ is **weakly nondominated** if it is not dominated by any $(m, 1, k')$-coterie of the same degree.*

To illustrate dominance, the $(2, 1, 2)$-coterie $\mathfrak{D} = (\{\{1, 2\}, \{3, 4\}\}, \{\{1, 3\}, \{2, 4\}\})$ is dominated by the $(2, 1, 2)$-coterie $\mathfrak{C} = (\{\{1, 2\}, \{3, 4\}, \{1, 4\}, \{2, 3\}\}, \{\{1, 3\}, \{2, 4\}\})$. By a simple enumeration, it can be proved that $\mathfrak{C}$ is nondominated.

If $\mathfrak{C} = (C_1, \ldots, C_m)$ dominates $\mathfrak{D} = (D_1, \ldots, D_m)$, then $\mathsf{deg}(C_i) \geq \mathsf{deg}(D_i)$ for all $1 \leq i \leq m$. For example, the $(2, 1, 1)$-coterie $\mathfrak{D} = (\{\{1, 2, 3\}, \{2, 3, 4\}\}, \{\{1, 2, 3\}, \{2, 3, 4\}\})$ is dominated by the $(2, 1, 2)$-coterie $\mathfrak{C} = (\{\{1, 2\}, \{3, 4\}, \{1, 4\}, \{2, 3\}\}, \{\{1, 3\}, \{2, 4\}\})$. So nondominated $(m, 1)$-coteries not only are more failure-resilient, but have degree no less than the ones they dominate. On the other hand, weak nondominance can be used to evaluate $(m, 1)$-coteries of a fixed degree.

It is often useful to discuss dominance with respect to a cartel, as defined below.

***Definition* 4.2** *Let $\mathfrak{D} = (D_1, \ldots, D_m)$ be an $(m, 1)$-coterie over P. $\mathfrak{D}$ is **dominated** w.r.t. D_i if there exists some $\mathfrak{C} = (C_1, \ldots, C_m)$ such that $\mathfrak{C}$ dominates $\mathfrak{D}$ and $C_i \neq D_i$; otherwise $\mathfrak{D}$ is **nondominated** w.r.t. D_i.*

It follows that if $\mathfrak{D}$ is nondominated w.r.t. every D_i, then $\mathfrak{D}$ must be nondominated as well. Similar to the theorem proposed by Garcia-Molina and Barbara for checking dominance of 1-coteries [8], the following can be used to check dominance of $(m, 1)$-coteries.

Lemma 4.1 *An $(m, 1)$-coterie $\mathfrak{D} = (D_1, \ldots, D_m)$ over P is dominated w.r.t. D_i if, and only if, there exists a set $H \subseteq P$ such that*

1. $\forall Q \in D_i, Q \not\subseteq H$.
2. $\forall Q \in D_j, j \neq i : Q \cap H \neq \emptyset$.

Proof. See Appendix. □

The following corollary follows immediately from Lemma 4.1.

Corollary 4.2 *An $(m,1)$-coterie $\mathfrak{D} = (D_1, \ldots, D_m)$ over P is dominated if, and only if, there exists a set $H \subseteq P$ and some $1 \leq i \leq m$ such that*

1. $\forall Q \in D_i, Q \not\subseteq H$.
2. $\forall Q \in D_j, j \neq i : Q \cap H \neq \emptyset$.

By Corollary 4.2, it is easy to see that when $m = 1$, there is only one nondominated $(1,1)$-coterie over P: $(\{\{i\} \mid i \in P\})$. When $m > 1$, any $(m,1)$-coterie of the form $(\{\{i\}\}, \ldots, \{\{i\}\})$ for some $i \in P$ is nondominated.

Corollary 4.3 *Let C be a nondominated 1-coterie over P. Then, the $(m,1,1)$-coterie $\mathfrak{T}_m(C) = (C, \ldots, C)$ transformed from C is nondominated if $m > 1$, and is dominated if $m = 1$ (unless $|P| = 1$).*

Proof. The case $m = 1$ follows directly from the fact that $(\{\{i\} \mid i \in P\})$ dominates all other $(1,1)$-coteries over P. For $m > 1$, by Corollary 4.2, if $\mathfrak{T}_m(C)$ is dominated, then there exists some $H \subseteq P$ such that $\forall Q \in C, Q \not\subseteq H \wedge Q \cap H \neq \emptyset$. Then by Theorem 2.1 in [8], C must also be dominated, contradiction. □

In contrast to the above corollary, we note here that if $\mathfrak{C} = (C_1, \ldots, C_m)$ is a nondominated $(m,1)$-coterie, then an $(m',1)$-coterie $\mathfrak{C}' = (C_{i_1}, \ldots, C_{i_{m'}})$ constructed from $\mathfrak{C}$ by taking m' cartels $C_{i_1}, \ldots, C_{i_{m'}}$ in $\mathfrak{C}$ as the cartels of $\mathfrak{C}'$ may not preserve $\mathfrak{C}$'s nondominance property. For example, the $(3,1,2)$-coterie $(\{\{1,2\},\{3,4\}\},\{\{1,3\},\{2,4\}\},\{\{1,4\},\{2,3\}\})$ is nondominated, but the $(2,1,2)$-coterie $(\{\{1,2\},\{3,4\}\},\{\{1,3\},\{2,4\}\})$ is dominated by the $(2,1,2)$-coterie $(\{\{1,2\},\{3,4\},\{1,4\},\{2,3\}\},\{\{1,3\},\{2,4\}\})$.

***Definition* 4.3** *Let $\mathfrak{C} = (C_1, \ldots, C_m)$ be an $(m,1)$-coterie over P, $m > 1$. A $\overline{C_i}$-**transversal** of $\mathfrak{C}$ is a set $T \subseteq P$ such that for every C_j, $i \neq j$, $T \cap Q \neq \emptyset$ for all $Q \in C_j$. T is **minimal** if no proper subset of it is a $\overline{C_i}$-transversal of $\mathfrak{C}$.*

So if $\mathfrak{C} = (C_1, \ldots, C_m)$ is an $(m,1)$-coterie, then every $Q \in C_i$ is a $\overline{C_i}$-***transversal*** of $\mathfrak{C}$. Minimal transversals play an important role in nondominated $(m,1)$-coteries. First, $\mathfrak{C}$ is nondominated w.r.t. C_i if, and only if, C_i is the set of minimal $\overline{C_i}$-transversals of $\mathfrak{C}$.

Lemma 4.4 *Let $\mathfrak{C} = (C_1, \ldots, C_m)$ be an $(m,1)$-coterie over P such that $m > 1$. Then $\mathfrak{C}$ is nondominated w.r.t. C_i if, and only if, C_i consists of all minimal $\overline{C_i}$-transversals of $\mathfrak{C}$.*

Proof. See Appendix. □

The lemma implies that given $\mathfrak{C} = (C_1, \ldots, C_m)$, we can obtain another $\mathfrak{C}' = (C_1, \ldots, C_{i-1}, C_i', C_{i+1}, \ldots, C_m)$ that is nondominated w.r.t. to C_i' by replacing C_i with the set of minimal $\overline{C_i}$-transversals of $\mathfrak{C}$. As we shall see shortly, we can apply the same procedure to other cartels of $\mathfrak{C}$, one after another, to obtain a nondominated $(m,1)$-coterie. For this property, we need the following lemma.

Lemma 4.5 *Let $\mathbb{C} = (C_1, \ldots, C_m)$ be an $(m,1)$-coterie over P, $m > 1$. Suppose $R \in C_i$ is a minimal $\overline{C_i}$-transversal of $\mathbb{C}$ for some $1 \leq i \leq m$. Let T be a minimal $\overline{C_j}$-transversal of $\mathbb{C}$ for some $j \neq i$. Let $\mathbb{C}' = (C_1, \ldots, C_{j-1}, C'_j, C_{j+1}, \ldots, C_m)$, where $C'_j = \{Q \mid Q \in C_j, T \not\subseteq Q\} \cup \{T\}$. Then $\mathbb{C}'$ is an $(m,1)$-coterie over P, and R is also a minimal $\overline{C_i}$-transversal of $\mathbb{C}'$.*

Proof. It is easy to verify that $\mathbb{C}'$ satisfies the intersection and minimality properties of Definition 2.1. So $\mathbb{C}'$is an $(m,1)$-coterie. (Note that the minimality property requires the fact that T is minimal.) To see that R is also a minimal $\overline{C_i}$-transversal of $\mathbb{C}'$, observe that R intersects T (because T intersects every quorum in C_i and $R \in C_i$). So R is a $\overline{C_i}$-transversal of $\mathbb{C}'$. So if R is not minimal, then some proper subset S of R is a minimal $\overline{C_i}$-transversal of $\mathbb{C}'$. So S intersects every quorum in C_k, $k \neq i$, and every quorum in C'_j. Let $Q \in C_j - C'_j$. Since $T \subset Q$ and S intersects T, S intersects Q. So S intersects every quorum in C_j. So S is also a $\overline{C_i}$-transversal of $\mathbb{C}$, contradicting the fact that R is minimal. □

Corollary 4.6 *Let $\mathbb{C} = (C_1, \ldots, C_m)$ be an $(m,1)$-coterie over P such that $m > 1$ and $\mathbb{C}$ is nondominated w.r.t. C_i. Let T be a minimal $\overline{C_j}$-transversal of $\mathbb{C}$ for some $j \neq i$. Let $\mathbb{C}' = (C_1, \ldots, C_{j-1}, C'_j, C_{j+1}, \ldots, C_m)$, where $C'_j = \{Q \mid Q \in C_j, T \not\subseteq Q\} \cup \{T\}$. Then the $(m,1)$-coterie $\mathbb{C}'$ is also nondominated w.r.t. C_i.*

Proof. Since $\mathbb{C}$ is nondominated w.r.t. C_i, by Lemmas 4.4 and 4.5, every $R \in C_i$ is a minimal $\overline{C_i}$-transversal of $\mathbb{C}'$. So if $\mathbb{C}'$ is dominated w.r.t. C_i, then by Lemma 4.4, some $S \subset P$ is a minimal $\overline{C_i}$-transversal of $\mathbb{C}'$ and $S \notin C_i$. By definition, S intersects every quorum in C_k, $k \neq i$, and every quorum in C'_j. Let $Q \in C_j - C'_j$. Since $T \subset Q$ and S intersects T, S intersects Q. So S intersects every quorum in C_j. So S is also a $\overline{C_i}$-transversal of $\mathbb{C}$. But since $\mathbb{C}$ is nondominated w.r.t. C_i, by Lemma 4.4 every minimal $\overline{C_i}$-transversal of $\mathbb{C}$ is in C_i. Since $S \notin C_i$, some $U \in C_i$, $U \subsetneq S$, is a minimal $\overline{C_i}$-transversal of $\mathbb{C}$. Then by Lemma 4.5, U is also a minimal $\overline{C_i}$-transversal of $\mathbb{C}'$. This then contradicts the fact that S is a minimal $\overline{C_i}$-transversal of $\mathbb{C}'$. □

Based on the above lemmas and corollary, we can use the two simple algorithms presented in Figure 3 to convert dominated $(m,1)$-coteries to nondominated. The first algorithm adds the set of minimal $\overline{D_i}$-transversals of $\mathfrak{D}$ to D_i all at once. The second algorithm adds minimal transversals to the cartels in a round robin fashion, and each time only one transversal is added. The correctness of the first algorithm follows directly from Corollary 4.6, while the correctness of the second algorithm follows from Lemmas 4.4 and 4.5.

To illustrate the difference, let $\mathfrak{D} = (\{\{1,2\},\{3,4\}\}, \{\{1,3\},\{2,4\}\})$. Then the first algorithm yields $(\{\{1,2\},\{3,4\},\ \{1,4\},\{2,3\},\},\{\{1,3\},\{2,4\}\})$, while the second algorithm yields $(\{\{1,2\},\{3,4\},$ $\{1,4\}\},\ \{\{1,3\},\ \{2,4\},\ \{1,4\}\})$. It can be seen that the second algorithm results in a more balanced $(m,1)$-coterie.

Both of the above algorithms require a step to compute minimal transversals. Minimal transversals can be computed using the simple algorithm described in [2]. The algorithm may run in time exponential in the input and output size, however [5]. As also pointed out in [5], computing minimal transversals is important in several other areas in computer science, including database theory, switching theory, logic, and artificial intelligence. Unfortunately, it remains open whether the problem can be solved in time polynomial in the input and output size. Note that quora used in Maekawa's algorithm and in Maekawa_M and Maekawa_S are not constructed on the fly. So the time complexity of the algorithms does not depend on the procedure used to construct the quora, but on their finished construction.

4.2 Distributed Cores

Recall from Theorems 3.1, 3.2, and 3.3 that an affine plane $\mathrm{AP} = (\mathcal{P}, \mathcal{L})$ of order n can be used to construct an $(m,1,n)$-coterie over $\mathcal{P}$, for any $m \leq n+1$, by taking m classes of parallel lines

Input. An $(m,1)$-coterie $\mathfrak{D} = (D_1, \ldots, D_m)$, $m > 1$.
Output. A nondominated $(m,1)$-coterie.

Algorithm 1.

1. for $i \longleftarrow 1$ to m do {
2. let Γ be the set of minimal $\overline{D_i}$-transversals of $\mathfrak{D}$.
3. $D_i \longleftarrow \Gamma$ }
4. output $\mathfrak{D}$

Algorithm 2.

1. repeat
2. for $i \longleftarrow 1$ to m do
3. if there is a minimal $\overline{D_i}$-transversal T of $\mathfrak{D}$ such that $T \notin D_i$, then
4. $D_i \longleftarrow D_i - \{Q \mid Q \in D_i, T \subseteq Q\} \cup \{T\}$
5. until no more new transversal T can be found in step 3 for all i
6. output $\mathfrak{D}$

Figure 3: Two algorithms for converting dominated $(m,1)$-coteries to nondominated.

from $\mathcal{L}$ as the m cartels of the coterie. All $(m,1,n)$-coteries constructed in this way are identical subject to isomorphisms. As we do not need to distinguish isomorphic $(m,1,n)$-coteries, we shall use $\mathrm{AP}(n,m)$ to denote an $(m,1,n)$-coterie of this form.

As also noted in Theorem 3.2, $\mathrm{AP}(n,m)$ is balanced, uniform, and regular. The three properties are important to realize a truly distributed implementation of the $(m,1,k)$-resource allocation problem. Unfortunately, $\mathrm{AP}(n,m)$ is, in general, dominated.

Theorem 4.7 $\mathrm{AP}(n,m)$ *is dominated for all* $n > 2$ *and* $m > 1$.

Proof. See Appendix.

Because $\mathrm{AP}(n,m)$ has optimal degree for all $m > 1$, any $(m,1)$-coterie that dominates $\mathrm{AP}(n,m)$ must be of the same degree. So $\mathrm{AP}(n,m)$ cannot be weakly nondominated, either. Theorem 4.7 excludes the case $n = 2$. For this case, $\mathrm{AP}(2,3)$ is nondominated. To see this, let $\mathcal{P} = \{1,2,3,4\}$. Then $\mathrm{AP}(2,3) = (\{\{1,2\},\{3,4\}\},\{\{1,3\},\{2,4\}\},\{\{1,4\},\{2,3\}\})$ (subject to isomorphisms). It is easy to verify that $\mathrm{AP}(2,3)$ is nondominated.

In [14] we presented another $(m,1)$-coterie that is also balanced, uniform, and regular. The coterie minimizes processes' loads by letting each process be included in at most two quora. It has degree $\sqrt{\frac{2n}{m(m-1)}}$, and is also dominated. Except for some special cases, it remains open whether or not there exists a nondominated $(m,1,k)$-coterie that is also balanced, uniform, and regular for any given k.

To cope with the apparently conflicting nature in between full distributedness and failure resilience, we observe that in real systems failures do not occur often. So distributedness and failure resilience need not be considered at the same time. Rather, we can focus on distributedness when failures do not occur, and turn into fault tolerance when failures do occur. In this case, we can design a quorum system such that some quora are used to facilitate a truly distributed implementation of the resource scheduling, while other quora are used to "back up" the system when failures occur. The following definition is used to realize this concept.

Definition* 4.4** *Let* $\mathfrak{C} = (C_1, \ldots, C_m)$ *be an* $(m,1)$*-coterie over* P. *A* ***distributed core *of* $\mathfrak{C}$ *is an* $(m,1)$*-coterie* $\mathfrak{D} = (D_1, \ldots, D_m)$ *such that* $D_i \subset C_i$, $1 \le i \le m$, *and* $\mathfrak{D}$ *is balanced, uniform, and regular.*

Thus, rather than designing a balanced, uniform, regular, and nondominated $(m, 1, k)$-coterie (if it exists!), we can instead construct a nondominated quorum system with a distributed core. We know already how to construct a balanced, uniform, and regular $(m, 1, k)$-coterie. The following lemma shows how such a coterie can be extended to meet our goal.

Lemma 4.8 *Let $\mathfrak{D} = (D_1, \ldots, D_m)$ be a balanced, uniform, and regular $(m, 1, k)$-coterie over P. If for every D_i, $1 \leq i \leq m$, every quorum $Q \in D_i$ is a minimal $\overline{D_i}$-transversal of $\mathfrak{D}$, then $\mathfrak{D}$ can be extended to an $(m, 1, k)$-coterie $\mathfrak{C}$ of which $\mathfrak{D}$ is a distributed core.*

Proof. We can use the algorithms presented in Figure 3 to convert $\mathfrak{D}$ to nondominated. By Lemma 4.5, every quorum in $\mathfrak{D}$ remains in the converted coterie. □

By Lemma 4.8 and the property of $\mathrm{AP}(n, m)$, we see that a nondominated $(m, 1, k)$-coterie with $\mathrm{AP}(n, m)$ as a distributed core can be easily constructed.

To summarize, the concept of distributed cores separates failure resilience from fully distributedness. So when designing quorum systems, we can first design a balanced, uniform, and regular $(m, 1, k)$-coterie $\mathfrak{D}$ such that every quorum in a cartel is minimal. The by applying Lemma 4.8, the constructed $(m, 1, k)$-coterie $\mathfrak{D}$ can be easily enlarged to a nondominated $(m, 1, k)$-coterie $\mathfrak{C}$ that 'contains' $\mathfrak{D}$ as a distributed core. The same concept can also be applied to the design of standard quorum systems for mutual exclusion and l-exclusion.

5 Conclusions and Future Work

We have presented the (m, l, k)-resource allocation problem that concerns the scheduling of l copies of a resource among m groups of processes. Each copy of the resource can be used by at most k processes of the same group at a time, but no two processes of different groups can use the copy simultaneously. The problem can be used to model many existing resource allocation problems, as well as new problems that cope with group resource allocation.

We then studied quorum systems for the $l = 1$ case. This case corresponds to group mutual exclusion with and without bounded capacity, i.e., the $(m, 1, k)$-resource allocation problem and the $(m, 1, \infty)$-resource allocation problem. To lay the groundwork, we have presented some fundamental results for $(m, 1)$-coteries, including basic definitions and composition methods. Based on these results we have then presented some $(m, 1)$-coteries that are optimal/near-optimal in degree, and $(m, 1)$-coteries that can support a fully distributed implementation of the problem An important characteristic we have observed is that all "fully-distributed" $(m, 1)$-coteries we can construct so far are, in general, not optimal in failure resilience. It remains an interesting open problem to see if such an $(m, 1)$-coterie exists. Nevertheless, we have proposed the concept of "distributed cores" to cope with the conflict between full distributedness and failure resilience.

The main benefit of our studies is that they reduced the design of distributed solutions for group resource allocation to combinatorial problems. Many interesting results in combinatorics can then be applied. For example, by using Maekawa's algorithm, the design of an $(m, 1, k)$-coterie provides an immediate solution to $(m, 1, k)$-resource allocation. To our knowledge, no solution for the problem has been proposed before. As we have also proved, no $(m, 1, k)$-coterie can be constructed out of an n-set if $k > \sqrt{n}$. For the corresponding $(m, 1, k)$-resource allocation problem involving n processes, the $(m, 1, \sqrt{n})$-coterie $\mathrm{AP}(m, \sqrt{n})$ constructed from an affine plane of order $\sqrt{n}$ in combination with a modification of Maekawa's algorithm, the Maekawa_M algorithm presented in [14], can be used. In this case, the algorithm has message complexity $O(n)$ and synchronization delay 2 message transmission time. For comparison, the message-passing solutions in [12, 33, 4] for group mutual exclusion all have message complexity $O(n)$, and synchronization delay 2 or $O(n)$ (if they use a ring architecture). Moreover, unlike quorum-based algorithms, none of them can tolerate a single process failure.

In light of the rich literature for quorum systems for mutual exclusion, l-exclusion, and replicated databases, there are many possible directions for future work, including constructions of other possible $(m, 1)$-coteries, availability analysis (cf., [26, 22]), and case studies/performance measurement.

References

[1] Daniel Barbara and Hector Garcia-Molina. Mutual exclusion in partitioned distributed systems. *Distributed Computing*, 1:119–132, 1986.

[2] Claude Berge. *Hypergraphs*, volume 45 of *North-Holland Mathematical Library*. Elsevier-North Holland, Amsterdam, 1989.

[3] Bollobás. On generalized graphs. *Acta Mathematica Academiae Scientiarum Hungaricae*, 16:447–452, 1965.

[4] Sebastien Cantarell, Ajoy K. Datta, Franck Petit, and Vincent Villain. Token based group mutual exclusion for asynchronous rings. In *Proceedings of the 21st International Conference on Distributed Computing Systems (ICDCS)*, pages 691–694. IEEE Computer Society Press, 2001.

[5] Thomas Eiter and Georg Gottlob. Identifying the minimal transversals of a hypergraph and related problems. *SIAM Journal on Computing*, 24(6):1278–1304, December 1995.

[6] Michael J. Fischer, Nancy A. Lynch, James E. Burns, and Allan Borodin. Resource allocation with immunity to limited process failure (preliminary report). In *20th Annual Symposium on Foundations of Computer Science*, pages 234–254, San Juan, Puerto Rico, 29–31 October 1979. IEEE.

[7] Peter Frankl. Extremal set systems. In Ronald L. Graham and Martin Grötschel and L. Lovász (eds.), *Handbook of Combinatorics*, volume 2. Elsevier and The MIT Press, 1995.

[8] Hector Garcia-Molina and Daniel Barbara. How to assign votes in a distributed system. *Journal of the ACM*, 32(4):841–860, October 1985.

[9] Luisa Gargano, János Körner, and Ugo Vaccaro. Qualitative independence and Sperner problems for directed graphs. *Journal of Combinatorial Theory, Series A*, 61:173–192, 1992.

[10] Vassos Hadzilacos. A note on group mutual exclusion. In *Proceedings of the 20th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, Newport, Rhode Island, August 2001. ACM Press.

[11] Godfrey H. Hardy and Edward M. Wright. *An introduction to the theory of numbers*. Oxford: Clarendon Press, 1965.

[12] Yuh-Jzer Joung. The congenial talking philosophers problem in computer networks (extended abstract). In *Proceedings of the 13th International Symposium on DIStributed Computing (DISC99)*, Lecture Notes in Computer Science 1693, pages 195–209. Springer, 1999.

[13] Yuh-Jzer Joung. Asynchronous group mutual exclusion. *Distributed Computing*, 13(4):189–206, 2000.

[14] Yuh-Jzer Joung. Quorum-based algorithms for group mutual exclusion. *IEEE Transactions on Parallel and Distributed Systems*, 14(5):1–14, May 2003. To appear.

[15] Hirotsugu Kakugawa, Satoshi Fujita, Masafumi Yamashita, and Tadashi Ae. Availability of k-coterie. *IEEE Transactions on Computers*, 42(5):553–558, May 1993.

[16] Patrick Keane and Mark Moir. A simple local-spin group mutual exclusion algorithm. In *Proceedings of the 18th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 23–32. ACM Press, 1999.

[17] János Körner and Gábor Simonyi. A Sperner-type theorem and qualitative independence. *Journal of Combinatorial Theory, Series A*, 59:90–103, 1992.

[18] Yu-Chen Kuo and Shing-Tsaan Huang. A simple scheme to construct k-coteries with $O(\sqrt{N})$ uniform quorum sizes. *Information Processing Letters*, 59(1):31–36, 8 July 1996.

[19] Harris F. MacNeish. Euler squares. *The Annals of Mathematics*, 2nd series, 23(3):221–227, March 1922.

[20] Mamoru Maekawa. A $\sqrt{N}$ algorithm for mutual exclusion in decentralized systems. *ACM Transactions on Computer Systems*, 3(2):145–159, May 1985.

[21] Dahlia Malkhi, Michael K. Reiter, and Avishai Wool. The load and availability of Byzantine quorum systems. *SIAM Journal on Computing*, 29(6):1889–1906, December 2000.

[22] Moni Naor and Avishai Wool. The load, capacity, and availability of quorum systems. *SIAM Journal on Computing*, 27(2):423–447, March 1998.

[23] Mitchell L. Neilsen. Properties of nondominated K-coteries. *Journal of Systems and Software*, 37(1):91–96, April 1997.

[24] Mitchell L. Neilsen and Masaaki Mizuno. Coterie join algorithm. *IEEE Transactions on Parallel and Distributed Systems*, 3(5):582–590, September 1992.

[25] Daniel Pedoe. *An Introduction to Projective geometry.* International series of monographs in pure and applied mathematics; v. 33. Pergamon Press, 1963.

[26] David Peleg and Avishai Wool. The availability of quorum systems. *Information and Computation*, 123(2):210–223, December 1995.

[27] David Peleg and Avishai Wool. Crumbling walls: A class of practical and efficient quorum systems. *Distributed Computing*, 10(2):87–97, 1997.

[28] Svatopluk Poljak, Aleš Pultr, and Voljtěch Rödl. On qualitatively independent partitions and related problems. *Discrete Applied Mathematics*, 6:193–205, 1983.

[29] Svatopluk Poljak and Voljtěch Rödl. Orthogonal partitions and covering of graphs. *Czechoslovak Mathematical Journal*, 30:475–485, 1980.

[30] Svatopluk Poljak and Zsolt Tuza. On the maximum number of qualitatively independent partitions. *Journal of Combinatorial Theory, Series A*, 51:111–116, 1989.

[31] Alfred Rényi. *Probability Theory.* Elsevier/North-Holland, Amsterdam, London, New York, 1970.

[32] Emanuel Sperner. Ein Satz über die Untermengen einer endlichen Menge. *Mathematische Zeitschrift*, 27:544–548, 1928.

[33] Kuen-Pin Wu and Yuh-Jzer Joung. Asynchronous group mutual exclusion in ring networks. *IEE Proceedings–Computers and Digital Techniques*, 147(1):1–8, 2000.

Appendix: Proofs

Proof of Lemma 2.2:

We first show that $\mathfrak{C} \otimes \mathfrak{D}$ satisfies the intersection condition of Definition 2.1. Let $S^i_{a,b}$ be a quorum in E_i, and $S^j_{c,d}$ be a quorum in E_j, where $i \neq j$. By the intersection property of $\mathfrak{C}$, there is a w such that $w \in Q^i_a$ and $w \in Q^j_b$. So $S^i_{a,b}$ contains a subset $R^i_b(P_w)$, and $S^j_{c,d}$ contains a subset $R^j_d(P_w)$. By the intersection property of $\mathfrak{D}(P_w)$, $R^i_b(P_w)$ and $R^j_d(P_w)$ contain a common element. So $S^i_{a,b} \cap S^j_{c,d} \neq \emptyset$.

For the minimality condition, let $S^i_{a,b}$ and $S^i_{c,d}$ be two quora in E_i. If $a \neq c$, then by the minimality property of $\mathfrak{C}$, there is some $w \in Q^i_a - Q^i_c$, and some $w' \in Q^i_c - Q^i_b$. So $S^i_{a,b}$ contains $R^i_b(P_w)$ and $S^i_{c,d} \cap R^i_b(P_w) = \emptyset$. Similarly, $S^i_{c,d}$ contains $R^i_d(P_{w'})$ and $S^i_{a,b} \cap R^i_d(P_{w'}) = \emptyset$. So $S^i_{a,b} \not\subset S^i_{c,d}$ and $S^i_{c,d} \not\subset S^i_{a,b}$. If $a = c$, then clearly $S^i_{a,b} \subset S^i_{c,d}$ implies $R^i_b \subset R^i_c$. By the minimality property of $\mathfrak{D}$, we have $S^i_{a,b} \not\subset S^i_{c,d}$; and similarly $S^i_{c,d} \not\subset S^i_{a,b}$.

For the degree of $\mathfrak{C} \otimes \mathfrak{D}$, observe that if $Q^i_a \cap Q^i_c = \emptyset$, then $S^i_{a,b} \cap S^i_{c,d} = \emptyset$ for all $1 \leq c, d \leq |D_i|$. Moreover, if $R^i_b \cap R^i_d = \emptyset$, then $S^i_{a,b} \cap S^i_{c,d} = \emptyset$ for all $1 \leq a, c \leq |C_i|$. So $\mathsf{deg}(E_i) \geq \mathsf{deg}(C_i) \cdot \mathsf{deg}(D_i)$. On the other hand, if $R^i_b \cap R^i_d \neq \emptyset$, then $S^i_{a,b} \cap S^i_{c,d} \neq \emptyset$ if $Q^i_a \cap Q^i_c \neq \emptyset$. So $\mathsf{deg}(E_i) = \mathsf{deg}(C_i) \cdot \mathsf{deg}(D_i)$. □

Proof of Lemma 2.3:

Proof. For the intersection condition, let i be such that $u[i] \neq 0, v[i] \neq 0$, and $u[i] \neq v[i]$. Let $S \in C_u$ and $T \in C_v$. Since $u[i] \neq 0$, by the definition of C_u, S contains a subset $Q_{u[i]} \in C^i_{u[i]}$. Similarly T contains a subset $Q_{v[i]} \in C^i_{v[i]}$. Since $u[i] \neq v[i]$, by the intersection property of $\mathfrak{C}^i$, $Q_{u[i]} \cap Q_{v[i]} \neq \emptyset$. So $S \cap T \neq \emptyset$.

For the minimality condition, suppose otherwise $S, T \in C_u$ and $S \subset T$. Since $|C_u| \geq 2$, there is a j such that $u[j] \neq 0$ and $C^j_{u[j]}$ contains Q_1 and Q_2 such that $Q_1 \subset S$ and $Q_2 \subset T$. Since $\{P_1, \ldots P_r\}$ is a partition of P, no $\mathfrak{C}^i$ can involve an element in Q_1 and Q_2 if $i \neq j$. So $S \subset T$ implies that $Q_1 \subset Q_2$, contradicting the fact that quora in $C^j_{u[j]}$ also satisfy the minimality condition. The case for C_v is similar.

Finally, let $k = \min\{\mathsf{deg}(C^j_{u[j]}) \,|\, u[j] \neq 0, 1 \leq j \leq r\}$, and let h be such that $u[h] \neq 0$ and $\mathsf{deg}(C^h_{u[h]}) = k$. By the definition of C_u, every $S \in C_u$ contains some subset $Q \in C^h_{u[h]}$. So $\mathsf{deg}(C_u) \leq k$. On the other hand, it is easy to see that C_u contains k pairwise disjoint sets, as each $C^j_{u[j]}$, $u[j] \neq 0$, contains at least k pairwise disjoint sets. So $\mathsf{deg}(C_u) = k$. Similarly, $\mathsf{deg}(C_v) = \min\{\mathsf{deg}(C^j_{v[j]}) \,|\, v[j] \neq 0, 1 \leq j \leq r\}$. □

Proof of Theorem 3.5:

Proof. It suffices to show that the probability that step 2 of the construction does not generate an $(m,1)$-coterie over P is strictly less than one. For each $1 \leq i \leq m$, let $C_i = \{Q^j_i,\ 1 \leq j \leq k\}$ be the random partition of P chosen in step 2. Then the probability that step 2 does not generate an $(m,1)$-coterie is equal to

$$
\begin{aligned}
& Pr[\,\exists\, i,j,a,b, i \neq j, \text{ such that } Q^a_i \cap Q^b_j = \emptyset\,] \\
\leq & \sum_{\substack{i,j,a,b \\ 1 \leq i \neq j \leq m \\ 1 \leq a,b \leq k}} Pr[\,Q^a_i \cap Q^b_j = \emptyset\,] \\
= & \; k^2 \binom{m}{2} \frac{\binom{n-\frac{n}{k}}{\frac{n}{k}}}{\binom{n}{\frac{n}{k}}} \;=\; k^2 \binom{m}{2} \left(\frac{n-\frac{n}{k}}{n}\right)\left(\frac{n-\frac{n}{k}-1}{n-1}\right)\cdots\left(\frac{n-\frac{n}{k}-(\frac{n}{k}-1)}{n-(\frac{n}{k}-1)}\right) \\
\leq & \; k^2 \binom{m}{2} \left(\frac{n-\frac{n}{k}}{n}\right)^{\frac{n}{k}} \;=\; k^2 \binom{m}{2} \left((1-\frac{1}{k})^k\right)^{\frac{n}{k^2}} \\
\leq & \; k^2 \binom{m}{2} e^{-\frac{n}{k^2}} \;<\; e^{2\log(km) - \frac{n}{k^2}} \;\leq\; e^{2\log(km) - 2\log(nm)} < 1
\end{aligned}
$$

□

Proof of Lemma 4.1:

For the if-direction, there are two cases to consider. If there exists some $R \in D_i$ such that $H \subsetneq R$, then let D'_i be $D_i - \{S \in D_i \,|\, H \subseteq S\} \cup \{H\}$. It is easy to see that $\mathfrak{D}' = (D_1, \ldots, D_{i-1}, D'_i, D_{i+1}, \ldots, D_m)$ is an $(m,1)$-coterie and it dominates $\mathfrak{D}$. If there is no $R \in D_i$ such that $H \subsetneq R$, then let D'_i be $D_i \cup \{H\}$. Again, it is easy to see that $\mathfrak{D}' = (D_1, \ldots, D_{i-1}, D'_i, D_{i+1}, \ldots, D_m)$ is an $(m,1)$-coterie and it dominates $\mathfrak{D}$.

For the only-if direction, assume that $\mathfrak{C} = (C_1, \ldots, C_m)$ dominates $\mathfrak{D}$ and $C_i \neq D_i$. There are two cases to consider. If $D_i \subsetneq C_i$, then let H be one of the elements in $C_i - D_i$. Then set H must satisfy conditions 1 and 2, or otherwise $\mathfrak{C}$ would not be an $(m,1)$-coterie. If $D_i - C_i \neq \emptyset$, then let $Q \in D_i - C_i$. Since $\mathfrak{C}$ dominates $\mathfrak{D}$, there exists some $H \in C_i$ such that $H \subsetneq Q$. We claim that H

satisfies both conditions 1 and 2. For condition 1, if the condition does not hold, then $Q' \subseteq H$ for some $Q' \in D_i$. But then we have $Q' \subseteq H \subsetneq Q$, contradicting the fact that $\mathfrak{D}$ satisfies the minimality property. For condition 2, suppose otherwise that $Q' \cap H = \emptyset$ for some $Q' \in D_j$, $j \neq i$. Then since $\mathfrak{C}$ dominates $\mathfrak{D}$, there exists some $R \in C_j$ such that $R \subseteq Q'$. So $R \cap H = \emptyset$. Since $R \in C_j$ and $H \in C_i$, we have a contradiction that $\mathfrak{C}$ does not satisfy the interaction property of $(m,1)$-coteries. □

Proof of Lemma 4.4:

For the only-if direction, observe that every $Q \in C_i$ is a C_i-transversal of $\mathfrak{C}$. If Q is not minimal, then there exists some $T \subsetneq Q$ such that T is a C_i-transversal of $\mathfrak{C}$ and T is not a superset of any quorum in C_i. Then by Lemma 4.1 $\mathfrak{C}$ is dominated w.r.t. C_i; contradiction. Moreover, if some minimal C_i-transversal T is not in C_i, then again T cannot be a superset of any quorum in C_i. So by Lemma 4.1 $\mathfrak{C}$ is dominated w.r.t. C_i; contradiction again. So C_i consists of all minimal C_i-transversals of $\mathfrak{C}$.

For the if-direction, if $\mathfrak{C}$ is dominated w.r.t. C_i, then by Lemma 4.1 there exists a minimal C_i-transversal H of $\mathfrak{C}$ not belonging to C_i. This contradicts the fact that C_i consists of all minimal C_i-transversals of $\mathfrak{C}$. □

Proof of Theorem 4.7:

Let $\mathrm{AP}(n,m) = (D_1, \ldots, D_m)$, and let P be the underlying n^2-element set. Without loss of generality, we show that $\mathrm{AP}(n,m)$ is dominated w.r.t. D_1. By Lemma 4.1, it suffices to find a set $H \subseteq P$ such that

1. $\forall Q \in D_1, Q \not\subseteq H$.
2. $\forall Q \in D_j, j \neq 1 : Q \cap H \neq \emptyset$.

Let $D_i = \{Q_i^1, \ldots, Q_i^n\}$, $1 \leq i \leq m$. Note that D_i is a partition of P, and each block has size n. Let a_1 be an element in Q_1^1. By the incidence properties of affine planes, $Q_1^1 - \{a_1\}$ intersects all but one quorum in D_j for every $1 < j \leq m$. Without loss of generality, let Q_j^1 be the quorum in D_j that does not intersect $Q_1^1 - \{a_1\}$. Again, by the incidence properties, $Q_j^1 \cap Q_1^h \neq \emptyset$ for all $h \leq n$. Let a_k be an element in $Q_k^1 \cap Q_1^k$, $1 < k \leq \min(m,n)$. Let $H = Q_1^1 - \{a_1\} \cup \{a_2, \ldots, a_{\min(m,n)}\}$. Then, set H intersects $Q_2^1, \ldots, Q_{\min(m,n)}^1$ (and all other Q_i^j for $1 < i \leq m$ and $1 < j \leq n$). Since $n > 2$, $Q_1^i \not\subset H$ for all $i \leq n$. So if $m \leq n$, then H satisfies the above two conditions and we are done.

If $m = n+1$, then let a_{n+1} be an element in $Q_{n+1}^1 \cap Q_1^2$. Let $H = Q_1^1 - \{a_1\} \cup \{a_2, \ldots, a_{n+1}\}$. (Note that in this case H contains $n-1$ elements from Q_1^1, at most two elements from Q_1^2, and one element from each of $Q_1^3, \ldots, Q_1^{n+1}$.) Then the set H again satisfies the above two conditions. □