SIAM J. DISCRETE MATH.
Vol. 10, No. 2, pp. 223–245, May 1997

006

LOAD BALANCING IN QUORUM SYSTEMS*

RON HOLZMAN†, YOSI MARCUS‡, AND DAVID PELEG‡

Abstract. This paper introduces and studies the question of balancing the load on processors participating in a given quorum system. Our proposed measure for the degree of balancing is the ratio between the load on the least frequently referenced element and on the most frequently used one.

We give some simple sufficient and necessary conditions for perfect balancing. We then look at the balancing properties of the common class of voting systems and prove that every voting system with odd total weight is perfectly balanced. (This holds, in fact, for the more general class of ordered systems.)

We also give some characterizations for the balancing ratio in the worst case. It is shown that for any quorum system with a universe of size n, the balancing ratio is no smaller than $1/(n-1)$, and this bound is the best possible. When restricting attention to nondominated coteries (NDCs), the bound becomes $2/\big(n-\log_2 n+o(\log n)\big)$, and there exists an NDC with ratio $2/\big(n-\log_2 n-o(\log n)\big)$.

Next, we study the interrelations between the two basic parameters of load balancing and quorum size. It turns out that the two size parameters suitable for our investigation are the size of the largest quorum and the *optimally weighted average quorum size* (OWAQS) of the system. For the class of ordered NDCs (for which perfect balancing is guaranteed), it is shown that over a universe of size n, some quorums of size $\lceil (n+1)/2 \rceil$ or more must exist (and this bound is the best possible). A similar lower bound holds for the OWAQS measure if we restrict attention to voting systems. For nonordered systems, perfect balancing can sometimes be achieved with much smaller quorums. A lower bound of $\Omega(\sqrt{n})$ is established for the maximal quorum size and the OWAQS of any perfectly balanced quorum system over n elements, and this bound is the best possible.

Finally, we turn to quorum systems that cannot be perfectly balanced, but have some balancing ratio $0 < \rho < 1$. For such systems we study the trade-offs between the required balancing ratio ρ and the quorum size it admits in the best case. It is easy to get an analogue of the result for perfect balancing, yielding a lower bound of $\sqrt{n\rho}$. We actually get a better estimate by a refinement of the argument.

Key words. quorum, coterie, intersecting hypergraphs, load balancing

AMS subject classifications. 68R05, 68Q22, 68M10, 05C65, 05099

PII. S0895480193260303

1. Introduction.

1.1. Motivation. *Quorum systems* serve as a basic tool providing a uniform and reliable way for achieving coordination between processes in a distributed system. Quorum systems are defined as follows. Suppose that the system is composed of n elements $u_1, \ldots, u_n$, taken from a universe U, representing sites, nodes, processors, or other abstract entities. A *set system* is a collection $\mathcal{S}$ of sets over the universe U. A set system $\mathcal{S}$ is said to satisfy the *quorum intersection* requirement if for every two sets S_i and S_j in $\mathcal{S}$, the intersection $S_i \cap S_j$ is not empty. A *quorum system* is a collection of sets that enjoys the quorum intersection property. The sets of $\mathcal{S}$ are referred to as the *quorums* of the system.

* Received by the editors December 15, 1993; accepted for publication (in revised form) April 23, 1996. The research of the third author was supported in part by a Walter and Elise Haas Career Development Award and by an Israel Science Foundation grant. Part of the work was done while he was visiting the IBM T. J. Watson Research Center.

http://www.siam.org/journals/sidma/10-2/26030.html

† Department of Mathematics, Technion–Israel Institute of Technology, Haifa 32000, Israel (holzman@techunix.technion.ac.il).

‡ Department of Applied Mathematics and Computer Science, The Weizmann Institute of Science, Rehovot 76100, Israel (marcus@wisdom.weizmann.ac.il, peleg@wisdom.weizmann.ac.il).

Applications for quorum systems in distributed systems include control and management problems such as *mutual exclusion* (cf. [R86]), *name servers* (cf. [MV88]), and *replicated data management* (cf. [H84]). In all of these cases, the use of quorum systems is centered on the following basic idea. The application requires that certain information items be stored in the network in a reliable and consistent way. Storing the information at a single central site is problematic in case that site crashes. Storing the information at one particular *set of sites* may overcome this problem, but will prevent working in the system if a communication failure causes a partition in the network, since if users at different parts of the network continue working separately, the information can no longer be guaranteed to be consistent.

The conceptual solution based on quorum systems is to make use of a large *collection* of possible sets of sites in the system. Each such set forms a *quorum* in the sense that any query or update operation concerning the information at hand can be performed by accessing the elements of this single set alone, and the choice of the particular quorum to be used can be made arbitrarily (i.e., all quorums are equally adequate).

In particular, in order to perform an update to the information, the user selects one quorum S_i in the quorum system $\mathcal{S}$, and records the update in every one of the elements that compose S_i. Likewise, a potential consumer of this information may choose any quorum $S_j \in \mathcal{S}$, and query the elements of S_j for the needed information. Note that the consumer must query *each* of the elements of S_j in order to be certain of obtaining the latest version. The reason for this is that a sequence of k updates, performed by a number of different users, may make use of different quorums $S_{i_1}, \ldots, S_{i_k}$, and therefore the elements of a quorum S_j used in a subsequent query may contain different information. Specifically, if the element $x \in S_j$ does not belong to S_{i_k} then the information stored in it will not be the most recent one. Moreover, it is impossible to tell, just by inspecting the data stored at x, whether this is the last version. Luckily, since the intersection of every two quorums in a quorum system is not empty, the consumer is guaranteed to encounter at least one element that is able to supply the most up-to-date version (namely, the element at the intersection of S_j and S_{i_k}).

This type of solution is capable of withstanding crashes and network partitions (up to a point), due to the greater degree of freedom the user has in choosing the quorum. In particular, in the case of crashes, the consumer can choose a quorum that does not include the crashed elements, and in the case of a partition, it may still be possible for one part of the network to contain a complete quorum. (Of course, it is quite impossible for two disconnected parts of the system to both contain complete quorums!)

Considerable attention is given in the literature to a special type of quorum system called a *coterie* (see [GB85] and [IK90]). A *coterie* is a quorum system in which the quorums are not allowed to fully contain each other. A subclass of special interest is that of *nondominated coteries* (or *NDCs*), which are better than other coteries in terms of fault tolerance and communication cost. This subclass is defined as follows. Given two coteries $\mathcal{S}_1$ and $\mathcal{S}_2$ over the same universe U, we say that $\mathcal{S}_2$ *dominates* $\mathcal{S}_1$ if $\mathcal{S}_2 \neq \mathcal{S}_1$ and for every quorum $S \in \mathcal{S}_1$ there is a quorum $T \in \mathcal{S}_2$ such that $T \subseteq S$. An NDC is a coterie which is not dominated by any other coterie (see [GB85]).

1.2. Load balancing. There are many types of quorum systems, and many parameters of quorum systems affecting the applications using them. Such parameters include quorum sizes (affecting communication costs) and the number of quorums (affecting immunity to partitioning).

Of special interest are parameters for evaluating the distribution of workload over the system, and measuring the degree of balancing possible for a given quorum system. If all the users of the system prefer to use one particular quorum while possible (e.g., in a failure-free execution), then the elements participating in this quorum will be overloaded compared to others. So it makes sense to try to use a more uniform distribution for selecting the quorum to be accessed. Formally, given a quorum system $\mathcal{S} = \{S_1, \ldots, S_m\}$, a *quorum load vector* (QLV) is a vector $\mathbf{v} = (v_1, v_2, \ldots, v_m)$ expressing the distribution of relative loads placed on the quorums of $\mathcal{S}$. (That is, in a long sequence of quorum accesses, a v_i fraction of the accesses is directed at quorum S_i.)

This distribution induces an *access rate* for each element u_j, which is the sum of the access frequencies of the quorums it belongs to, $a_j = \sum_{u_j \in S_i} v_i$. Thus the *element load vector* (ELV) $\mathbf{a} = (a_1, a_2, \ldots, a_n)$ induced by the QLV $\mathbf{v}$ expresses the relative loads placed on the elements of U when using the QLV $\mathbf{v}$.

Our proposed measure for the degree of balancing is the ratio between the rate of accesses to the least frequently used element in the quorum system and the rate of accesses to the most frequently used one. Formally, given a QLV $\mathbf{v}$ for $\mathcal{S}$ and the induced ELV $\mathbf{a}$, we let $\rho_{\mathcal{S},\mathbf{v}} = \min\{a_j\}/\max\{a_j\}$, and define the *balancing ratio* $\rho_{\mathcal{S}}$ of $\mathcal{S}$ as the maximum $\rho_{\mathcal{S},\mathbf{v}}$ over all QLVs $\mathbf{v}$. A system is said to be *perfectly balanced* if all the elements are accessed at the same rate, namely, $\rho_{\mathcal{S}} = 1$.

This paper focuses on issues related to balancing. In current technologies, a common and promising way to increase computing power is by connecting many fast processors together into compound systems. Quorum systems can be used for coordination in such systems. For small systems, the effect of the particular quorum system used on the communication cost is not significant. However, when the systems become larger, the importance of choosing a good quorum system may significantly increase. In particular, some quorum systems may be well adapted to the demand of load balancing, while for others, such a demand may impose heavy communication costs. Worse yet, certain types of quorum systems may be incapable of providing perfect or even partial balancing, regardless of the cost.

In this paper we introduce and address this issue, defining the fundamental notions and concepts relevant to load balancing, and developing some basic results on the balancing properties of a variety of quorum system classes.

Let us remark that to the best of our knowledge, currently existing systems do not address the issue of load balancing at all. Consequently, the quorum selection mechanisms used in existing systems typically do not provide such balancing, as they base the selection on some arbitrary choice, or worse, on a fixed search pattern, perpetuating the imbalance.

However, even though current quorum systems do not provide any means for balancing the load on the processors, it should be clear that there is no *inherent* reason that prevents them from doing so. In fact, given a desirable QLV $\mathbf{v}$ for the quorum system $\mathcal{S}$ at hand (i.e., a QLV $\mathbf{v}$ for which $\rho_{\mathcal{S},\mathbf{v}} = \rho_{\mathcal{S}}$), it is rather straightforward to develop a simple randomized protocol for quorum selection, based on interpreting $\mathbf{v}$ as a probability distribution over the quorums, and drawing a quorum at random according to $\mathbf{v}$. Such a protocol will in fact enforce an actual load distribution close to the optimal one, with high probability. For many natural quorum system classes (including most of the specific classes discussed in what follows), this protocol will also enjoy an efficient (and fast) distributed implementation.

1.3. Related work. Synchronization and coordination are central issues in the area of distributed systems. Many types of synchronization protocols rely on variants of quorum systems. In [H84] quorum intersection is defined between *read quorums* and *write quorums*, and also between other abstract types of quorums. In [MV88] aspects of distributed control are examined and lower bounds are presented for certain types of quorum systems. The issues of fault tolerance and availability of quorum systems are studied in [PW93]. For more on the applicability of quorum-based techniques in distributed systems, and on the examples mentioned above, the reader is referred to [H84, GB85] and the references therein. We are unaware of previous discussion of load balancing issues in the context of quorum systems in the literature.

Set systems in general (including intersecting hypergraphs in particular) were studied extensively in recent years (cf. [B86]). The terms *coterie* and *nondominated coterie* (NDC) are defined in [GB85], and many properties of coteries and NDCs are presented. Some interesting properties of NDCs are derived in [L73]. In [IK90] a relationship is established between coteries and boolean functions. Properties of coteries and NDCs are derived from properties of the appropriate functions.

1.4. Contributions. This paper focuses on a number of questions related to the issue of balancing the load on processors participating in a given quorum system.

We begin by giving some simple sufficient and necessary conditions for perfect balancing. (One trivial necessary condition is that the system is nonredundant; namely, that every element participates in some quorum.)

We then look at the balancing properties of the common class of voting systems. (A voting system is based on assigning a number of "votes" to each element of the universe; the votes induce a quorum system by taking as a quorum any collection of elements that holds a "minimal" majority of all the votes.) We define the class of ordered NDCs, which is an extension of voting systems, and prove that every ordered NDC is perfectly balanced. It follows, in particular, that every voting system with odd total number of votes is perfectly balanced.

Next we turn to characterizations for the balancing ratio in the worst case. We show that for any quorum system with a universe of size n, the balancing ratio is no smaller than $1/(n-1)$, and this bound is the best possible. When restricting attention to NDCs, the bound becomes $2/\big(n-\log_2 n+o(\log n)\big)$, and there exists an NDC with ratio $2/\big(n-\log_2 n-o(\log n)\big)$.

Next, we study the interrelationships between the two basic parameters of load balancing and quorum size. It turns out that the two size parameters suitable for our investigation are the size of the largest quorum and the *optimally weighted average quorum size* (OWAQS) of the system (corresponding to an optimal load vector).

For the class of ordered NDCs (for which perfect balancing is guaranteed), it is shown that over a universe of size n, some quorums of size $\lceil (n+1)/2\rceil$ or more must exist (and this bound is the best possible). A similar lower bound holds for the OWAQS measure if we restrict attention to voting systems.

For nonordered systems, perfect balancing can sometimes be achieved with much smaller quorums. A lower bound of $\Omega(\sqrt{n})$ is established for the maximal quorum size and the OWAQS of any perfectly balanced quorum system over n elements, and this bound is the best possible.

Finally, we turn to quorum systems that cannot be perfectly balanced, but have some balancing ratio $0<\rho<1$. For such systems we study the trade-offs between the required balancing ratio ρ and the quorum size it admits in the best case. It is easy to get an analogue of the result for perfect balancing, yielding a lower bound of

$\sqrt{n\rho}$. We actually get a better estimate, by a refinement of the argument.

2. Basic notions.

DEFINITION. *A* quorum system *is a pair* $(U, \mathcal{S})$*, where* U *is a nonempty finite set and* $\mathcal{S}$ *is a set of nonempty subsets of* U *such that the intersection of every two sets in* $\mathcal{S}$ *is nonempty. We refer to the set* U *as the* universe *and to the sets in* $\mathcal{S}$ *as the* quorums *of the system.*

It is sometimes convenient to represent a quorum system by a matrix of 0's and 1's.

DEFINITION. *The* quorum matrix *of a quorum system* $(U, \mathcal{S})$ *is the* $m \times n$ *matrix* $\hat{\mathcal{S}} = (\hat{s}_{ij})$ *obtained as follows: the elements of* U *are enumerated as* $u_1, u_2, \ldots, u_n$*, the quorums in* $\mathcal{S}$ *are enumerated as* $S_1, S_2, \ldots, S_m$*, and*

$$\hat{s}_{ij} = \begin{cases} 1 & \textit{if } u_j \in S_i, \\ 0 & \textit{otherwise.} \end{cases}$$

We shall usually be interested in quorum systems in which no quorum contains another, since in the case of containment the larger quorum is redundant for our purposes.

DEFINITION. *A* coterie *is a quorum system in which no quorum contains another quorum.*

In order to describe and analyze a coterie, it is often convenient to refer to the set of subsets of the universe which contain some quorum. This is facilitated by the following definition.

DEFINITION. *A* monotone quorum system *(MQS) is a quorum system* $(U, \mathcal{M})$ *such that* $S \in \mathcal{M}$ *and* $S \subseteq T \subseteq U$ *imply* $T \in \mathcal{M}$*. Given a coterie* $(U, \mathcal{S})$*, a* superquorum *is any subset of* U *that contains a quorum of* $\mathcal{S}$*. The* MQS generated by $(U, \mathcal{S})$ *is the collection of superquorums of* $(U, \mathcal{S})$*, namely, the system* $(U, \bar{\mathcal{S}})$*, where* $T \in \bar{\mathcal{S}}$ *if and only if* $T \supseteq S$ *for some* $S \in \mathcal{S}$*. Conversely, if we are given a MQS* $(U, \bar{\mathcal{S}})$ *then the coterie* $(U, \mathcal{S})$ *is determined uniquely (*$S \in \mathcal{S}$ *if and only if* $S \in \bar{\mathcal{S}}$ *and no proper subset of* S *is in* $\bar{\mathcal{S}}$*) and is called the* coterie derived from $(U, \bar{\mathcal{S}})$.

Example 2.1. Minimal Majority Coterie. Let $|U| = n$ and let $\bar{\mathcal{S}} = \{S \subseteq U\colon\ |S| > \frac{n}{2}\}$; that is, the superquorums are the sets containing a majority of elements. The coterie derived from $(U, \bar{\mathcal{S}})$ is that in which the quorums are all subsets of U of size $\lceil \frac{n+1}{2} \rceil$. □

Notation. When $U = \{u_1, u_2, \ldots, u_n\}$ and $x_1, x_2, \ldots, x_n$ are real numbers, we denote the *x-weight* of a subset $S \subseteq U$ by

$$x(S) = \sum_{u_j \in S} x_j.$$

Example 2.2. Voting Coterie. Let $U = \{u_1, u_2, \ldots, u_n\}$ and assume that to each $u_j \in U$ we assign a nonnegative integer w_j, called the *weight* of u_j. Then we define the MQS $\bar{\mathcal{S}} = \{S \subseteq U\colon\ w(S) > \frac{w(U)}{2}\}$. The coterie derived from $(U, \bar{\mathcal{S}})$ is that in which the quorums are those subsets of U which carry a majority of the total weight and are inclusion-minimal with respect to this property. A coterie $(U, \mathcal{S})$ obtained in this manner is called a *voting system*. Observe that the minimal majority coterie of Example 2.1 is a special case of a voting system, in which all weights are equal. □

Example 2.3. Star Coterie. Let $U = \{u_1, u_2, \ldots, u_n\}$ and let $\mathcal{S}$ consist of the $n-1$ quorums $\{u_1, u_2\}, \{u_1, u_3\}, \ldots, \{u_1, u_n\}$. Then $(U, \mathcal{S})$ is a coterie. We call such a coterie a *star*. Observe that a star is also a voting system (take $w_1 = n-1$, $w_2 = \cdots = w_n = 1$). ◻

Voting systems play a distinguished role in the study of quorum systems because of the natural and simple way in which they are specified. The defining weights also supply a ranking of the elements of U in terms of their importance for forming quorums. This notion is captured by the following definition.

Notation. Let $U = \{u_1, u_2, \ldots, u_n\}$, and let $S \subseteq U$ with $u_i \notin S$, $u_j \in S$. We denote by S^i_j the *replacement set* $\big(S \setminus \{u_j\}\big) \cup \{u_i\}$.

DEFINITION. *Let $(U, \mathcal{S})$ be a coterie. We say that $(U, \mathcal{S})$ is* ordered *if it is possible to enumerate the elements of U as $u_1, u_2, \ldots, u_n$ so that the following holds: if $1 \le i < j \le n$ and S is a superquorum with $u_i \notin S$, $u_j \in S$, then S^i_j is also a superquorum.*

Intuitively, the above property means that if $i < j$ then u_i is at least as useful as u_j for forming quorums. The reason that the definition refers to superquorums rather than quorums is that it may happen that S is a quorum but S^i_j is a nonminimal superquorum. It is straightforward to check, and we will do so now.

FACT 2.4. *Every voting system is ordered.*

Proof. This is proved by enumerating the elements so that $w_1 \ge w_2 \ge \cdots \ge w_n$. ◻

The converse is known to be false; that is, there exist ordered coteries that cannot be obtained as a voting system [Os85]. There are also coteries that are not ordered, as witnessed by the following class of examples.

Example 2.5. FPP. Let U and $\mathcal{S}$ be the sets of points and lines, respectively, of a *finite projective plane* (see [H86]). We recall that in a finite projective plane of order q (abbreviated FPP(q)) there are n points and n lines, where $n = q^2 + q + 1$. Each line contains $q+1$ points and there are $q+1$ lines going through each point. Any two lines have exactly one point in common, and through any two points there is exactly one line. A FPP(q) is known to exist for every q which is a prime power. Clearly, if $(U, \mathcal{S})$ is a FPP(q), $q \ge 2$, then $(U, \mathcal{S})$ is not ordered, since no point can replace another in a line. ◻

A special class of coteries arises from a concept of domination among coteries (see [GB85]).

DEFINITION. *Let $(U, \mathcal{S}_1)$ and $(U, \mathcal{S}_2)$ be coteries. We say that $(U, \mathcal{S}_2)$* dominates *$(U, \mathcal{S}_1)$ if $\mathcal{S}_2 \ne \mathcal{S}_1$ and for every quorum $S \in \mathcal{S}_1$ there is a quorum $T \in \mathcal{S}_2$ such that $T \subseteq S$. A* nondominated coterie *(NDC) is a coterie which is not dominated by any other coterie.*

The following fact (cf. Cor. 2.1 in [IK90]) can be used as a convenient alternative definition of an NDC.

PROPOSITION 2.6. *Let $(U, \mathcal{S})$ be a coterie. Then $(U, \mathcal{S})$ is an NDC if and only if for every partition of U into two parts S_1 and S_2, one of the S_i $(i = 1, 2)$ is a superquorum.*

We now record a simple but useful property of NDCs.

PROPOSITION 2.7. *Let $(U, \mathcal{S})$ be an NDC, and let $u \in U$ be in $\cup\mathcal{S}$ (that is, u belongs to at least one quorum). Then:*

(a) *There exist two quorums S and T such that $S \cap T = \{u\}$.*

(b) *If, moreover, $(U, \mathcal{S})$ is ordered with corresponding enumeration $u_1, u_2, \ldots, u_n$ of U and $u = u_j$, then there are two quorums S and T such that $S \cap T = \{u_j\}$ and $S \cup T \supseteq \{u_1, \ldots, u_j\}$.*

Proof. Let S be a quorum containing u. Applying the property given in Proposition 2.6 to the partition $S \setminus \{u\}$, $(U \setminus S) \cup \{u\}$, we conclude that $(U \setminus S) \cup \{u\}$ is a superquorum. Let T be a quorum contained in it. Then $S \cap T \subseteq \{u\}$, and since the intersection of two quorums is nonempty we have $S \cap T = \{u\}$, establishing part (a).

To prove part (b), assume that $i < j$ and $u_i \notin S \cup T$. By the property of an ordered coterie it follows that the replacement set T_j^i is a superquorum. This, however, is a contradiction since T_j^i is disjoint from S. ∎

Let us examine the above examples of coteries to see whether they are NDCs.

FACT 2.8.

(a) *The minimal majority coterie of Example* 2.1 *is an NDC if and only if* n *is odd.*

(b) *A sufficient condition for a voting system (Example* 2.2*) to be an NDC is that the total weight be odd.*

(c) *A star coterie (Example* 2.3*) is dominated.*

(d) *A finite projective plane FPP*(q) *(Example* 2.5*) is an NDC for* $q = 2$ *but is dominated for all* $q > 2$*.*

Proof. Parts (a) and (b) [GB85] are seen easily from Proposition 2.6. Part (c) follows since neither $\{u_1\}$ nor $\{u_2, \ldots, u_n\}$ is a superquorum. For Part (d) see [P70, C93]. ∎

We remark that despite Fact 2.8(d), FPP(q) satisfies the property of Proposition 2.7(a) for all q.

The central concept of this research deals with load balancing.

DEFINITION. *Let* $(U, \mathcal{S})$ *be a quorum system with quorum matrix* $\hat{\mathcal{S}} = (\hat{s}_{ij})$, $i = 1, \ldots, m$, $j = 1, \ldots, n$. *A* quorum load vector *(QLV) is a vector* $\mathbf{v} = (v_1, v_2, \ldots, v_m)$ *whose components are real nonnegative numbers (not all zero) expressing the relative loads that are to be placed on the quorums of* $\mathcal{S}$*. The* element load vector *(ELV) induced by the QLV* $\mathbf{v}$ *is the vector* $\mathbf{a} = \mathbf{a}(\mathcal{S}, \mathbf{v}) = (a_1, a_2, \ldots, a_n)$ *computed by* $\mathbf{a} = \mathbf{v}\hat{\mathcal{S}}$ *and expressing the relative loads placed on the elements of* U *when using the QLV* $\mathbf{v}$.

DEFINITION. *Let* $(U, \mathcal{S})$ *be a quorum system. Given a QLV* $\mathbf{v} = (v_1, v_2, \ldots, v_m)$ *which induces the ELV* $\mathbf{a} = (a_1, a_2, \ldots, a_n)$*, we define the* balancing ratio *for* $\mathcal{S}$ *and* $\mathbf{v}$ *as*

$$\rho_{\mathcal{S},\mathbf{v}} = \frac{\min_{j=1,\ldots,n}\{a_j\}}{\max_{j=1,\ldots,n}\{a_j\}}.$$

The balancing ratio *of* $(U, \mathcal{S})$ *is defined as*

$$\rho_{\mathcal{S}} = \max\{\rho_{\mathcal{S},\mathbf{v}}\colon\ \mathbf{v} \text{ is a QLV}\}.$$

A straightforward continuity and compactness argument shows that $\rho_{\mathcal{S}}$ is well defined. We have associated with each quorum system $(U, \mathcal{S})$ a parameter $0 \le \rho_{\mathcal{S}} \le 1$, which tells us how evenly we can spread the load among the elements of U if we are allowed to assign the relative loads to the quorums optimally. The higher the $\rho_{\mathcal{S}}$, the better behaved the quorum system is from the point of view of load balancing.

We note the following basic fact regarding the balancing ratio.

FACT 2.9. *If* $U \neq \cup\mathcal{S}$ *then* $\rho_{\mathcal{S}} = 0$.

Proof. If $U \neq \cup\mathcal{S}$, then there is some element $u_i \in U$ that does not participate in any quorum of $\mathcal{S}$. Hence, no matter which QLV $\mathbf{v}$ we choose, a_i will be zero, and thus the balancing ratio $\rho_{\mathcal{S},\mathbf{v}}$ will be zero too. ∎

Consequently, in studying the balancing ratio it is natural to make the assumption that each element appears in some quorum.

DEFINITION. *A quorum system* $(U, \mathcal{S})$ *is* nonredundant *if each element of* U *appears in some quorum*; *i.e.,* $U = \cup\mathcal{S}$.

Once this assumption holds, we have $\rho_{\mathcal{S}} > 0$. The most pleasing situation is when all element loads can be made equal; that is, $\rho_{\mathcal{S}} = 1$.

DEFINITION. *A quorum system* $(U, \mathcal{S})$ *is* perfectly balanced *if* $\rho_{\mathcal{S}} = 1$.

3. Perfect balancing. We begin with a simple sufficient condition for perfect balancing.

DEFINITION. *Let* $(U, \mathcal{S})$ *be a quorum system and let* $u \in U$. *The* degree *of* u *in* $\mathcal{S}$ *is* $d_{\mathcal{S}}(u) = \big|\{S \in \mathcal{S}\colon\ u \in S\}\big|$. *We say that* $(U, \mathcal{S})$ *is* regular *if all elements of* U *have the same degree in* $\mathcal{S}$.

PROPOSITION 3.1. *Every regular quorum system is perfectly balanced.*

Proof. The proposition is proved by assigning equal loads to all quorums. ∎

As an application of Proposition 3.1, we note that the minimal majority quorum systems of Example 2.1 and the FPP coterie of Example 2.5 are regular, and hence perfectly balanced. The star coterie (Example 2.3), on the other hand, is not perfectly balanced (when $n \geq 3$), since it can be seen that the load on the center of the star is the sum of the loads on the other elements.

In trying to determine when a given quorum system is perfectly balanced, the following characterization is useful.

PROPOSITION 3.2. *Let* $(U, \mathcal{S})$ *be a quorum system, with* $U = \{u_1, u_2, \ldots, u_n\}$. *Then* $(U, \mathcal{S})$ *is perfectly balanced if and only if there exists no* $\mathbf{x} = (x_1, x_2, \ldots, x_n) \in \mathbb{R}^n$ *satisfying*

$$x(S) \geq 0 \ \textit{for all} \ \ S \in \mathcal{S}, \tag{1}$$

$$x(U) < 0. \tag{2}$$

(Recall the x-weight notation.)

Proof. The quorum system $(U, \mathcal{S})$ is perfectly balanced if there exists a real nonnegative vector $\mathbf{v}$ solving the equation system $\mathbf{v}\hat{\mathcal{S}} = \mathbf{1}$, where $\mathbf{1}$ denotes the n-dimensional vector of 1's. By the Minkowski–Farkas Lemma ([F01]; cf. [C83]), this is equivalent to the condition that the system of inequalities $\mathbf{x}\hat{\mathcal{S}}^\top \geq \mathbf{0}$, $\mathbf{x}\cdot\mathbf{1}^\top < 0$ has no solution. ∎

Our main result in this section is concerned with ordered NDCs. It will be derived from the following lemma.

LEMMA 3.3. *Let* $(U, \mathcal{S})$ *be a nonredundant NDC. Suppose that* $(U, \mathcal{S})$ *is ordered with corresponding enumeration* $u_1, u_2, \ldots, u_n$ *of* U. *Let* $\mathbf{x} = (x_1, x_2, \ldots, x_n) \in \mathbb{R}^n$ *and* $\alpha \in \mathbb{R}$ *satisfy*

$$x(S) \geq \alpha \ \textit{for all} \ \ S \in \mathcal{S}, \tag{3}$$

$$x(U) \leq 2\alpha. \tag{4}$$

Then $x_j \geq 0$ *for* $j = 1, 2, \ldots, n$.

Proof. Suppose, for contradiction, that $x_j < 0$ for some j, and let J be the largest such j. By Proposition 2.7(b) there exist two quorums S and T such that $S \cap T = \{u_J\}$ and $S \cup T \supseteq \{u_1, \ldots, u_J\}$. By the choice of J, we have $x_i \geq 0$ for every $u_i \in U \setminus (S \cup T)$, and hence $x(U) - x(S \cup T) = x\big(U \setminus (S \cup T)\big) \geq 0$. Therefore, using (3) and $x_J < 0$, we get

$$x(U) \geq x(S \cup T) = x(S) + x(T) - x_J > 2\alpha,$$

which contradicts (4). ∎

THEOREM 3.4. *Every ordered nonredundant NDC is perfectly balanced.*

Proof. For the sake of contradiction, let $(U,\mathcal{S})$ have the properties stated, but fail to be perfectly balanced. By Proposition 3.2 there exists $\mathbf{x}\in\mathbb{R}^n$ satisfying (1) and (2). We may apply Lemma 3.3 with $\alpha=0$ and conclude that $x_j\geq 0$ for $j=1,2,\ldots,n$. But this is inconsistent with (2). ∎

By Facts 2.4 and 2.8(b) we have the following corollary.

COROLLARY 3.5. *Every nonredundant voting system (Example* 2.2*) with odd total weight is perfectly balanced.* ∎

We remark that none of the assumptions made in Theorem 3.4 is superfluous. Indeed, the nonredundancy assumption is necessary for perfect balancing by Fact 2.9. If we drop the assumption of nondomination, the star coterie is an example that satisfies the other assumptions but not the conclusion. A class of examples indicating that the assumption of being ordered cannot be dispensed with will be presented in the following section (Example 4.3).

4. The balancing ratio in the worst case.

4.1. Characterization for the balancing ratio. The following proposition gives a dual formulation for the balancing ratio in the case when it is less than 1; it complements Proposition 3.2, which dealt with the case when the balancing ratio is 1.

We shall use the following notation: if $U=\{u_1,u_2,\ldots,u_n\}$ and $\mathbf{x}=(x_1,x_2,\ldots,x_n)\in\mathbb{R}^n$ then

$$P=\{u_j\in U:\ x_j>0\},$$
$$N=\{u_j\in U:\ x_j<0\}.$$

The expressions $x(P)$ and $x(N)$ will be used following our x-weight notation.

PROPOSITION 4.1. *Let $(U,\mathcal{S})$ be a quorum system, with $U=\{u_1,u_2,\ldots,u_n\}$ and $\rho_{\mathcal{S}}<1$. Then*

$$\rho_{\mathcal{S}}=\min\{x(P)\},$$

where the minimum is taken over all $\mathbf{x}=(x_1,x_2,\ldots,x_n)\in\mathbb{R}^n$ satisfying

$$x(S)\geq 0\quad \textit{for all}\ \ S\in\mathcal{S},\tag{5}$$

$$x(N)=-1.\tag{6}$$

Proof. The balancing ratio $\rho_{\mathcal{S}}$ can be defined as the optimal value of ρ in the linear programming problem

$$\rho_{\mathcal{S}}=\max_{\mathbf{v},\rho}\{\rho\},$$

subject to

$$\begin{aligned}\mathbf{v}\hat{\mathcal{S}}&\leq\mathbf{1},\\ \boldsymbol{\rho}-\mathbf{v}\hat{\mathcal{S}}&\leq\mathbf{0},\\ \mathbf{v}&\geq\mathbf{0},\\ \rho&\geq 0,\end{aligned}$$

where $\hat{\mathcal{S}}$ is the quorum matrix, $\boldsymbol{\alpha}$ denotes (for $\alpha\in\mathbb{R}$) the vector of appropriate dimension with all components equal to α, vector inequalities are understood componentwise, and the maximum is taken over all QLVs $\mathbf{v}$ and $\rho\in\mathbb{R}$. Note that this formulation is equivalent to the definition of $\rho_{\mathcal{S}}$, since the QLV $\mathbf{v}$ can always be normalized so that the largest component of the induced ELV $\mathbf{a}$ becomes 1. By linear programming duality, we can express $\rho_{\mathcal{S}}$ in the form

$$\rho_{\mathcal{S}} = \min_{\mathbf{y},\mathbf{z}}\{y(U)\}$$

subject to

$$z(U) \geq 1, \tag{7}$$

$$y(S) - z(S) \geq 0, \text{ for all } S \in \mathcal{S}, \tag{8}$$

$$\mathbf{y} \geq \mathbf{0}, \tag{9}$$

$$\mathbf{z} \geq \mathbf{0}, \tag{10}$$

where the minimum is taken over all vectors $\mathbf{y}, \mathbf{z} \in \mathbb{R}^n$.

We begin by showing that there exists a vector $\mathbf{x}$ satisfying (5) and (6) and also $x(P) \leq \rho_{\mathcal{S}}$. The inequality $\rho_{\mathcal{S}} \geq \min\{x(P)\}$ then follows. Suppose now that $\mathbf{y} = (y_1, y_2, \ldots, y_n)$ and $\mathbf{z} = (z_1, z_2, \ldots, z_n)$ satisfy (7)–(10), and yield the optimal value in the dual linear programming problem. We may assume that $z(U) = 1$, since we can achieve this by decreasing the values of the components of $\mathbf{z}$ without affecting the value of the solution or the validity of the constraints. Let $\mathbf{x} = \mathbf{y} - \mathbf{z}$, and let $N \subseteq U$ be defined with respect to $\mathbf{x}$ as in the statement of the proposition.

We observe first that $x(U) < 0$. Indeed,

$$x(U) = y(U) - z(U) = \rho_{\mathcal{S}} - 1$$

and $\rho_{\mathcal{S}} < 1$ by assumption. It follows in particular that $N \neq \emptyset$, and therefore $x(N) < 0$. On the other hand, by (9) and (10),

$$x(N) = \sum_{u_j \in N} y_j - z_j \geq -\sum_{u_j \in N} z_j \geq -\sum_{u_j \in U} z_j = -1.$$

Therefore, we can find a real number $\alpha \geq 1$ so that $\alpha x(N) = -1$; hence the vector $\alpha\mathbf{x}$ satisfies (6). It follows from (8) that $\mathbf{x}$, and hence also $\alpha\mathbf{x}$, satisfies (5). Thus $\alpha\mathbf{x}$ satisfies both (5) and (6). We have

$$\alpha x(P) = \alpha x(U) - \alpha x(N) = \alpha x(U) + 1 \leq x(U) + z(U) = y(U) = \rho_{\mathcal{S}}$$

(where the inequality relies on $x(U) < 0$ and $\alpha \geq 1$, and on (7)).

It remains to show, in the other direction, that any $\mathbf{x} = (x_1, x_2, \ldots, x_n) \in \mathbb{R}^n$ which satisfies (5) and (6) has $x(P) \geq \rho_{\mathcal{S}}$. The inequality $\rho_{\mathcal{S}} \leq \min\{x(P)\}$ follows immediately. Let $\mathbf{x}$ be such a vector. We define the vectors $\mathbf{y} = (y_1, y_2, \ldots, y_n)$ and $\mathbf{z} = (z_1, z_2, \ldots, z_n)$ by

$$y_j = \begin{cases} x_j & \text{if } x_j > 0, \\ 0 & \text{otherwise}, \end{cases}$$

$$z_j = \begin{cases} -x_j & \text{if } x_j < 0, \\ 0 & \text{otherwise}. \end{cases}$$

It can be checked that $\mathbf{y}$ and $\mathbf{z}$ satisfy (7)–(10). It follows that $y(U) \geq \rho_{\mathcal{S}}$. Since $y(U) = x(P)$, we are done. $\square$

4.2. A lower bound for the balancing ratio. We now address the following question: within the class of all nonredundant quorum systems with a universe of size n, how low can the balancing ratio be in the worst case?

THEOREM 4.2. *Let $(U, \mathcal{S})$ be a nonredundant quorum system with $U = \{u_1, u_2, \ldots, u_n\}$, $n \geq 2$. Then $\rho_{\mathcal{S}} \geq 1/(n-1)$. This bound is the best possible.*

Proof. We may assume that $\rho_{\mathcal{S}} < 1$. By Proposition 4.1, we have to prove that any $\mathbf{x} = (x_1, x_2, \ldots, x_n) \in \mathbb{R}^n$ which satisfies (5) and (6) has $x(P) \geq 1/(n-1)$. Since $x(N) = -1$ (by (6)) and $|N| \leq n-1$ (due to (5)), there exists some $u_j \in N$ with $x_j \leq -1/(n-1)$. Using the nonredundancy assumption, let S be a quorum with $u_j \in S$. Then

$$x(P) \geq x(S \cap P) \geq -x(S \cap N) \geq -x_j \geq 1/(n-1)$$

(where the second inequality is due to (5) again).

A candidate for attaining the worst case is the star coterie of Example 2.3, whose balancing ratio is easily seen to be $1/(n-1)$. $\square$

4.3. A lower bound for the balancing ratio on NDCs. The worst case for the balancing ratio occurs for the star, which is a dominated coterie. What happens if we restrict attention to NDCs? The following construction, taken from [EL74], exhibits a low balancing ratio.

Example 4.3. Nucleus Coterie. Let $r \geq 2$ be an integer and let U be the disjoint union of the sets K and L, where $|K| = 2r-2$ and $|L| = \binom{2r-2}{r-1}/2$. Let the elements of L be put in a one-to-one correspondence with the halvings of K. That is, to every unordered pair A, B of disjoint subsets of K of size $r-1$ each there corresponds an element $u_{A,B}$ of L. Let $\mathcal{S}$ consist of all sets of the form $A \cup \{u_{A,B}\}$ and $B \cup \{u_{A,B}\}$, where A, B is a halving of K, as well as all subsets of K of size r. It is easy to verify that $(U, \mathcal{S})$ is an NDC (using Proposition 2.6) and it is nonredundant. The number of elements is $n = 2r - 2 + \binom{2r-2}{r-1}/2$. The balancing ratio is 1 when $r = 2$ and is $\rho_{\mathcal{S}} = 4/\binom{2r-2}{r-1}$ when $r \geq 3$. The latter can be verified by noting that (a) the QLV $\mathbf{v}$ assigning zero load to the quorums contained in K and load 1 to every other quorum satisfies $\rho_{\mathcal{S},\mathbf{v}} = 4/\binom{2r-2}{r-1}$, and (b) the vector $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ defined by

$$x_j = \begin{cases} \frac{2}{(r-1)\binom{2r-2}{r-1}} & \text{if } u_j \in K, \\ -\frac{2}{\binom{2r-2}{r-1}} & \text{if } u_j \in L, \end{cases}$$

satisfies (5) and (6) and $x(P) = 4/\binom{2r-2}{r-1}$. $\square$

We observe that for $r = 3$ the above construction gives a nonredundant NDC $(U, \mathcal{S})$ with $|U| = 7$ which has balancing ratio $\rho_{\mathcal{S}} = 2/3$. It is therefore an example showing that Theorem 3.4 does not remain true if the assumption of being ordered is removed. No such example with universe of size smaller than 7 exists. Indeed, for $n \leq 5$ it is known that every NDC is a voting system and hence ordered [GB85]. For $n = 6$, an exhaustive search shows that all nonredundant NDCs are perfectly balanced.

For large r, the above construction gives almost the worst case as will be proved next.

THEOREM 4.4. *For every nonredundant NDC $(U, \mathcal{S})$ with $U = \{u_1, u_2, \ldots, u_n\}$, $\rho_{\mathcal{S}} \geq 2/\big(n - \log_2 n + o(\log n)\big)$. Furthermore, there exists such an NDC $(U, \mathcal{S})$ with $\rho_{\mathcal{S}} = 2/\big(n - \log_2 n - o(\log n)\big)$.*

Proof. Let $(U, \mathcal{S})$ satisfying the assumptions be given, and let us write

$$\rho_{\mathcal{S}} = \frac{2}{n-\alpha}$$

for a suitable real number α. We have to prove that $\alpha \geq \log_2 n - o(\log n)$.

Let $\mathbf{x} = (x_1, x_2, \ldots, x_n) \in \mathbb{R}^n$ be a vector which satisfies (5) and (6) and has

$$x(P) = \frac{2}{n-\alpha}. \tag{11}$$

Given any $u_j \in N$ we can find, using Proposition 2.7(a), two quorums S_j and T_j such that $S_j \cap T_j = \{u_j\}$. We have then (relying on (5) for the second inequality)

$$\begin{aligned} x(P) &\geq x(S_j \cap P) + x(T_j \cap P) \\ &\geq -x(S_j \cap N) - x(T_j \cap N) \geq -2x_j. \end{aligned} \tag{12}$$

It follows now from (11) and (12) that

$$x_j \geq -\frac{1}{n-\alpha} \quad \text{for all} \quad u_j \in N. \tag{13}$$

Before continuing the proof, let us note that at this stage we could easily deduce that $\alpha \geq 2$. Indeed, in (12) it must be the case that $S_j \cap P$ and $T_j \cap P$ are nonempty (since both S_j and T_j contain $u_j \in N$, yet by (5) both $x(S_j), x(T_j) \geq 0$). As these sets are disjoint, we know that $|P| \geq 2$ and hence $|N| \leq n-2$. It follows by (6) that there exists $u_j \in N$ with $x_j \leq -1/(n-2)$. In view of (13), this implies that $\alpha \geq 2$. Thus we have a simple proof of the estimate $\rho_{\mathcal{S}} \geq 2/(n-2)$. In order to get the slightly better estimate stated in the theorem, some more work is needed.

Assume without loss of generality (w.l.o.g.) that $x_1 \geq x_2 \geq \cdots \geq x_n$. Split U into three disjoint parts by setting the boundary values

$$M_1 = -\frac{1}{\sqrt{\log_2 n}\,(n-\alpha)} \qquad \text{and} \qquad M_2 = -\frac{2}{3(n-\alpha)},$$

and defining

$$\begin{aligned} A &= \{u_1, u_2, \ldots, u_\ell\} &&= \Big\{u_j \in U\colon\ x_j \geq M_1\Big\}, \\ B &= \{u_{\ell+1}, u_{\ell+2}, \ldots, u_p\} &&= \Big\{u_j \in U\colon\ M_2 \leq x_j < M_1\Big\}, \\ C &= \{u_{p+1}, u_{p+2}, \ldots, u_n\} &&= \Big\{u_j \in U\colon\ x_j < M_2\Big\}. \end{aligned}$$

Note that $P \subseteq A$. Hence using these definitions plus (13) and (6), we can deduce

$$M_1\ell + M_2(p-\ell) - \frac{1}{n-\alpha}(n-p) \leq M_1 \cdot |P| + x(A \cap N) + x(B) + x(C) \leq x(N) = -1.$$

This can be rewritten as

$$\alpha \geq \frac{1}{3}p + \left(\frac{2}{3} - \frac{1}{\sqrt{\log_2 n}}\right)\ell. \tag{14}$$

For each $u_j \in C$, let us choose as above two quorums S_j and T_j such that $S_j \cap T_j = \{u_j\}$. Let us write $S'_j = S_j \setminus \{u_j\}$, $T'_j = T_j \setminus \{u_j\}$. We now establish some properties of these sets.

First, we claim that

$$S'_j, T'_j \subseteq A \cup B \quad \text{for} \quad j = p+1, \ldots, n. \tag{15}$$

To see this, suppose for instance that $u_k \in S'_j \cap C$ for some $k \neq j$. Then we may sharpen (12) to get

$$x(P) \geq -2x_j - x_k > \frac{2}{n-\alpha},$$

which contradicts (11).

Second, we estimate the B portion of each set S'_j by

$$|S'_j \cap B| < \frac{2}{3}\sqrt{\log_2 n} \quad \text{for} \quad j = p+1, \ldots, n. \tag{16}$$

This is seen again by sharpening (12) in the form

$$x(P) \geq -2x_j - x(S'_j \cap B) > \frac{4}{3(n-\alpha)} + \frac{|S'_j \cap B|}{\sqrt{\log_2 n}\ (n-\alpha)}$$

and comparing with (11).

Third, we argue that

$$S'_j \neq S'_k \quad \text{for} \quad j \neq k,\ p+1 \leq j,\ k \leq n. \tag{17}$$

Indeed, if $S'_j = S'_k$ then $S'_j \cap T'_k = \emptyset$ which implies, by (15), that $S_j \cap T_k = \emptyset$, in contradiction to the quorum intersection property.

It follows from (15)–(17), by considering the mapping $j \mapsto S'_j$, that

$$n - p \leq 2^{\ell} \sum_{i < \frac{2}{3}\sqrt{\log_2 n}} \binom{p-\ell}{i}. \tag{18}$$

Going back to (14) we see that if $p \geq 3\log_2 n$ we are done. So we assume that $p < 3\log_2 n$ and then obtain from (18) that

$$n - 3\log_2 n < 2^{\ell}(3\log_2 n)^{\frac{2}{3}\sqrt{\log_2 n}}.$$

Taking logarithms we get $\ell > \log_2 n - o(\log n)$. Using (14) and $p \geq \ell$ we have

$$\alpha \geq \left(1 - \frac{1}{\sqrt{\log_2 n}}\right)\ell > \left(1 - \frac{1}{\sqrt{\log_2 n}}\right)\bigl(\log_2 n - o(\log n)\bigr) = \log_2 n - o(\log n)$$

as required.

An example of a coterie nearly matching the bound is the nucleus coterie of Example 4.3, which for large r has $\rho_{\mathcal{S}} = 2/\bigl(n - \log_2 n - o(\log n)\bigr)$. $\square$

We add two comments concerning Theorem 4.4 and its proof:

1. By some finer tuning of the proof it is possible to replace the $o(\log n)$ term by $\frac{4}{3}\sqrt{\log_2 n}$. Details are omitted.
2. The theorem remains true, with the same proof, if instead of an NDC we consider any quorum system having the property given in Proposition 2.7(a).

5. Load balancing and quorum size.

5.1. Measures for quorum size. In this section we study the extent of compatibility of two desirable goals: having a high balancing ratio and having small quorum sizes. The general theme will be that a high balancing ratio cannot be obtained with small quorum sizes.

DEFINITION. *A quorum system is* r-uniform *if every quorum has* r *elements.*

If a quorum system is r-uniform then clearly we should use r as the parameter describing the quorum size. But for more general quorum systems, the question arises as to which parameter should be used for evaluating quorum sizes. Two conceivable parameters that do not serve our purposes well are the minimum quorum size and the average quorum size. This is illustrated by the following example.

Example 5.1. Wheel Coterie. Let $U = \{u_1, u_2, \ldots, u_n\}$ and let $\mathcal{S}$ consist of the $n-1$ quorums $\{u_1, u_2\}, \{u_1, u_3\}, \ldots, \{u_1, u_n\}$ and the additional quorum $\{u_2, u_3, \ldots, u_n\}$. This differs from the star coterie (Example 2.3) only in the addition of the last quorum. It is easy to check that $(U, \mathcal{S})$ is perfectly balanced. Yet the minimum quorum size is 2 and the average quorum size is $3(n-1)/n$, both low numbers. We remark also that $(U, \mathcal{S})$ is a voting system and an NDC.

It turns out that two other parameters are more suitable for our investigation.

DEFINITION. *Let* $(U, \mathcal{S})$ *be a quorum system. The* rank *of* $(U, \mathcal{S})$ *is defined as*

$$r_{\mathcal{S}} = \max\bigl\{|S|\colon\ S \in \mathcal{S}\bigr\}.$$

DEFINITION. *Let* $(U, \mathcal{S})$ *be a quorum system with quorum matrix* $\hat{\mathcal{S}} = (\hat{s}_{ij})$, $i = 1, \ldots, m$, $j = 1, \ldots, n$. *Let* $\mathbf{v} = (v_1, v_2, \ldots, v_m)$ *be a QLV. The* weighted average quorum size *(WAQS) of* $(U, \mathcal{S})$ *corresponding to* $\mathbf{v}$ *is*

$$g_{\mathcal{S},\mathbf{v}} = \frac{1}{\sum_{i=1}^m v_i} \sum_{i=1}^m v_i |S_i| = \frac{\sum_{j=1}^n a_j}{\sum_{i=1}^m v_i},$$

where $\mathbf{a} = (a_1, a_2, \ldots, a_n)$ *is the ELV induced by* $\mathbf{v}$, *that is,* $\mathbf{a} = \mathbf{v}\hat{\mathcal{S}}$. *In the case when* $\mathbf{v}$ *is an optimizing QLV (that is,* $\rho_{\mathcal{S},\mathbf{v}} = \rho_{\mathcal{S}}$*), we refer to* $g_{\mathcal{S},\mathbf{v}}$ *as an* optimally weighted average quorum size *(OWAQS).*

As an illustration, let us apply these notions to the wheel coterie of Example 5.1. The rank there is $n-1$. The unique (up to proportionality) optimizing QLV is $\mathbf{v} = (1, 1, \ldots, 1, n-2)$, which gives the OWAQS $g_{\mathcal{S},\mathbf{v}} = n(n-1)/(2n-3)$, which is slightly more than $n/2$.

In our context of load balancing, it seems that the notion of an OWAQS is the suitable way to measure quorum size. The rank is also interesting as a worst case measure. If the quorum system is r-uniform then all approaches give r as the answer. In general, the WAQS and even the OWAQS are not unique, as they depend on $\mathbf{v}$. Clearly, for every QLV $\mathbf{v}$ we have $g_{\mathcal{S},\mathbf{v}} \le r_{\mathcal{S}}$.

5.2. Quorum size bounds for ordered NDCs. In the first part of our analysis we shall focus on ordered nonredundant NDCs. This is a natural class of quorum systems for which we know that perfect balancing is guaranteed (Theorem 3.4). So it is interesting to ask what quorum sizes this class admits, or more precisely, how low we can make the rank and the OWAQS within this class.

THEOREM 5.2. *Let* $(U, \mathcal{S})$ *be an ordered nonredundant NDC with universe of size* n. *Then* $r_{\mathcal{S}} \ge \lceil (n+1)/2 \rceil$. *This bound is the best possible.*

Proof. Let $u_1, u_2, \dots, u_n$ be an enumeration of U with respect to which $(U, \mathcal{S})$ is ordered. Applying Proposition 2.7(b) with $u = u_n$, we obtain two quorums S and T such that $S \cap T = \{u_n\}$ and $S \cup T = U$. Then $|S| + |T| = n + 1$, so at least one of them has size $\geq \lceil (n+1)/2 \rceil$.

For odd n, the optimality of the bound is shown by the minimal majority coterie of Example 2.1. For even n this is shown by a slight modification of that example. ∎

We note that no assumption of the theorem is redundant. The nucleus coterie of Example 4.3 is an r-uniform nonredundant NDC with $r \sim \frac{1}{2}\log_2 n$. The star (Example 2.3) is a 2-uniform ordered nonredundant coterie. If the nonredundancy assumption is removed then n may be made arbitrarily large without affecting anything else.

A similar lower bound on the OWAQS holds if we restrict attention to voting systems, a subclass of ordered coteries.

THEOREM 5.3. *Let $(U, \mathcal{S})$ be a perfectly balanced voting system with universe of size n. Then for every optimizing QLV $\mathbf{v}$, the OWAQS is greater than $n/2$.*

Proof. Let $\hat{\mathcal{S}} = (\hat{s}_{ij})$, $i = 1, \dots, m$, $j = 1, \dots, n$, be the quorum matrix, and let $\mathbf{v} = (v_1, v_2, \dots, v_m)$ be a QLV such that $\rho_{\mathcal{S},\mathbf{v}} = 1$. Then the ELV induced by $\mathbf{v}$ is $\mathbf{a} = \mathbf{v}\hat{\mathcal{S}} = (a_1, a_2, \dots, a_n)$ with all a_j equal, say, to the common value a. Let $\mathbf{w}^\top$ be a column vector whose components $w_1, w_2, \dots, w_n$ are weights which determine the voting system $(U, \mathcal{S})$. Then it follows from the definition of a voting system that every component of $\hat{\mathcal{S}}\mathbf{w}^\top$ is greater than $w(U)/2$. Therefore

$$\mathbf{v}\hat{\mathcal{S}}\mathbf{w}^\top > \frac{w(U)}{2} \sum_{i=1}^{m} v_i. \tag{19}$$

On the other hand, since every component of $\mathbf{v}\hat{\mathcal{S}}$ equals a, we have

$$\mathbf{v}\hat{\mathcal{S}}\mathbf{w}^\top = a w(U). \tag{20}$$

Combining (19) and (20) we get $a > \frac{1}{2}\sum_{i=1}^{m} v_i$. Therefore

$$g_{\mathcal{S},\mathbf{v}} = \frac{\sum_{j=1}^{n} a_j}{\sum_{i=1}^{m} v_i} = \frac{na}{\sum_{i=1}^{m} v_i} > \frac{n}{2}. \qquad \square$$

Comparing the last two theorems, it is natural to ask whether the (stronger) conclusion of Theorem 5.3 holds under the conditions of Theorem 5.2. The question involves the class of ordered NDCs that are not voting systems (and therefore Theorem 5.3 does not apply to them). It is not easy to construct examples for this class, but this has been done: two such examples with universe of size 13 are given in [Os85]. In the following theorem we show not only that there is a member of this class for which the conclusion of Theorem 5.3 fails, but that it fails for every member of this class.

THEOREM 5.4. *Let $(U, \mathcal{S})$ be an ordered nonredundant NDC with universe of size n. Suppose further that $(U, \mathcal{S})$ is not a voting system. Then there exists an optimizing QLV $\mathbf{v}$ whose OWAQS is equal to $n/2$.*

Proof. Let $(U, \mathcal{S})$ satisfy the assumptions of the theorem and assume that $(U, \mathcal{S})$ is ordered with corresponding enumeration $u_1, u_2, \dots, u_n$ of U.

As the first step in the proof, we claim that there is no pair (x, α), where $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{R}^n$ and α is a real number, such that

$$x(S) > \alpha > x(U \setminus S) \quad \text{for all} \quad S \in \mathcal{S}. \tag{21}$$

To prove the claim, suppose that such $\mathbf{x}$ and α exist. Then we may change the value of α, if necessary, to be $x(U)/2$, and (21) will still hold. Indeed, $x(S) > x(U \setminus S)$ implies that $x(S) > x(U)/2 > x(U \setminus S)$. So we shall assume that $\alpha = x(U)/2$. Applying Lemma 3.3 we deduce that $x_j \geq 0$ for $j = 1, 2, \ldots, n$. Now, let T be any subset of U. If T is a superquorum, say $T \supseteq S \in \mathcal{S}$, then it follows from (21) and the nonnegativity of the components of $\mathbf{x}$ that $x(T) \geq x(S) > \alpha$. If T is not a superquorum, then it follows from Proposition 2.6 that $U \setminus T$ is a superquorum, and therefore $x(T) = x(U) - x(U \setminus T) < x(U) - \alpha = \alpha$. We have shown that $\bar{\mathcal{S}} = \{T \subseteq U\colon\ x(T) > \alpha\}$. This indicates that $(U, \mathcal{S})$ is a voting system (strictly speaking, our definition of a voting system requires the weights to be integers, but this can be arranged by taking good enough rational approximations of the x_j's and clearing denominators). As this contradicts our assumption, we have proved the claim.

Let $\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_m \in \{0,1\}^n$ be the characteristic vectors of the quorums $S_1, S_2, \ldots, S_m$ ($\mathcal{S} = \{S_1, \ldots, S_m\}$). Let

$$\begin{aligned} Y &= \{\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_m\}, \\ Z &= \{\mathbf{1} - \mathbf{y}_1, \mathbf{1} - \mathbf{y}_2, \ldots, \mathbf{1} - \mathbf{y}_m\}, \end{aligned}$$

where $\mathbf{1}$ is the all-1 n-dimensional vector. The claim asserts that there is no hyperplane that separates the points of Y from those of Z. It follows that

$$A = \operatorname{conv}(Y) \cap \operatorname{conv}(Z) \neq \emptyset,$$

where $\operatorname{conv}(X)$ denotes the convex closure of X. The set A is convex and symmetric about $\frac{\mathbf{1}}{\mathbf{2}}$ (that is, $\mathbf{a} \in A$ implies $\mathbf{1} - \mathbf{a} \in A$). Hence $\frac{\mathbf{1}}{\mathbf{2}} \in A$, and in particular $\frac{\mathbf{1}}{\mathbf{2}} \in \operatorname{conv}(Y)$. The latter means that there exists a QLV $\mathbf{v} = (v_1, v_2, \ldots, v_m)$ with $\sum_{i=1}^m v_i = 1$ which induces the ELV $\mathbf{a} = \frac{\mathbf{1}}{\mathbf{2}}$. For this $\mathbf{v}$ we get

$$g_{\mathcal{S},\mathbf{v}} = \frac{\sum_{j=1}^n a_j}{\sum_{i=1}^m v_i} = \frac{n \cdot \frac{1}{2}}{1} = \frac{n}{2}. \qquad \square$$

Theorems 5.3 and 5.4 yield the following characterization of voting systems within ordered NDCs.

COROLLARY 5.5. *Let $(U, \mathcal{S})$ be an ordered nonredundant NDC with universe of size n. Then the following are equivalent:*

(a) *$(U, \mathcal{S})$ is a voting system.*
(b) *Every OWAQS is greater than $n/2$.*
(c) *No OWAQS is equal to $n/2$.*

Before leaving the ordered world, we want to mention without details two examples that we have constructed:

1. An ordered coterie which is perfectly balanced but whose rank is less than $n/2$. (This shows that the nondomination assumption in Theorems 5.2 and 5.4 cannot be removed, even if we add the assumption of perfect balancing. It also shows that relaxing "voting system" to "ordered" in Theorem 5.3 admits examples where the theorem's conclusion fails in a more essential sense than indicated by Theorem 5.4.)
2. A quorum system satisfying all the assumptions of Theorem 5.4 for which there is an OWAQS which is less than $n/2$. (This shows that the existential quantifier in the theorem's conclusion cannot be made universal.)

5.3. Quorum size bounds for (nonordered) perfectly balanced quorum systems. The foregoing theorems indicate that certain methods for constructing quorum systems or certain properties of quorum systems which guarantee perfect balancing are costly in terms of quorum size. But perfect balancing can be achieved with considerably smaller quorums. Indeed, a FPP(q) (Example 2.5) is $(q+1)$-uniform and has a universe of size $n = q^2 + q + 1$, so its rank is roughly $\sqrt{n}$. It is perfectly balanced by Proposition 3.1.

Our next goal is to prove the optimality (in terms of quorum size) of the finite projective planes among all perfectly balanced quorum systems. For this purpose, we first review some known concepts and results on fractional matchings in hypergraphs. We express them using the terminology of the current paper.

DEFINITION. *Let $(U, \mathcal{S})$ be a quorum system with quorum matrix $\hat{\mathcal{S}} = (\hat{s}_{ij})$, $i = 1, \ldots, m$, $j = 1, \ldots, n$. A* fractional matching *in $(U, \mathcal{S})$ is a QLV $\mathbf{v} = (v_1, v_2, \ldots, v_m)$ such that the induced ELV $\mathbf{a} = \mathbf{v}\hat{\mathcal{S}} = (a_1, a_2, \ldots, a_n)$ satisfies $a_j \le 1$, $j = 1, \ldots, n$. The* size *of a fractional matching $\mathbf{v} = (v_1, v_2, \ldots, v_m)$ is defined as*

$$|\mathbf{v}| = \sum_{i=1}^{m} v_i.$$

The fractional matching number *of $(U, \mathcal{S})$ is defined as*

$$\nu^*_{\mathcal{S}} = \max\big\{|\mathbf{v}|\colon \mathbf{v} \text{ is a fractional matching in } (U, \mathcal{S})\big\}.$$

It is easy to deduce the following from the quorum intersection property.

PROPOSITION 5.6. *Let $(U, \mathcal{S})$ be a quorum system. Then for every quorum $S \in \mathcal{S}$ we have $\nu^*_{\mathcal{S}} \le |S|$. As a consequence, $\nu^*_{\mathcal{S}} \le g_{\mathcal{S},\mathbf{v}}$ for every WAQS $g_{\mathcal{S},\mathbf{v}}$.*

The following finer estimate for $\nu^*_{\mathcal{S}}$ is due to Füredi.

PROPOSITION 5.7 (see [F81]). *Let $(U, \mathcal{S})$ be a quorum system of rank $r_{\mathcal{S}} = r$. Then $\nu^*_{\mathcal{S}} \le r - 1 + 1/r$.*

A FPP(r–1), if it exists, is an r-uniform quorum system with universe of size $r^2 - r + 1$ and fractional matching number $r - 1 + 1/r$. Thus Füredi's bound is attained for those values of r such that a FPP(r–1) exists. The following corollary of Proposition 5.7 had been proved earlier by Lovász.

PROPOSITION 5.8 (see [L75]). *Let $(U, \mathcal{S})$ be an r-uniform, regular quorum system. Then $|U| \le r^2 - r + 1$.*

Note that this bound too is attained for those values of r such that a FPP(r–1) exists.

We now return to our investigation of quorum size in perfectly balanced quorum systems.

THEOREM 5.9. *Let $(U, \mathcal{S})$ be a perfectly balanced quorum system with $|U| = n$. Then every OWAQS is at least $\sqrt{n}$.*

Proof. Let $\mathbf{v} = (v_1, v_2, \ldots, v_m)$ be a QLV with $\rho_{\mathcal{S},\mathbf{v}} = 1$. Then the ELV induced by $\mathbf{v}$ is $\mathbf{a} = (a_1, a_2, \ldots, a_n)$ with all a_j equal. By a suitable normalization, which does not affect $g_{\mathcal{S},\mathbf{v}}$, we may assume that $a_1 = a_2 = \cdots = a_n = 1$. With this assumption, $\mathbf{v}$ is a fractional matching. We have

$$g_{\mathcal{S},\mathbf{v}} = \frac{\sum_{j=1}^{n} a_j}{\sum_{i=1}^{m} v_i} = \frac{n}{|\mathbf{v}|} \ge \frac{n}{\nu^*_{\mathcal{S}}},$$

and therefore

$$n \le g_{\mathcal{S},\mathbf{v}} \nu^*_{\mathcal{S}}. \tag{22}$$

Using Proposition 5.6 this implies

$$n \leq g_{\mathcal{S},\mathbf{v}}^2,$$

which yields the desired lower bound on $g_{\mathcal{S},\mathbf{v}}$. ∎

The foregoing theorem establishes the asymptotic optimality (in terms of OWAQS) of the finite projective planes among all perfectly balanced quorum systems. We can get exact optimality in terms of the rank, as follows.

THEOREM 5.10. *Let $(U,\mathcal{S})$ be a perfectly balanced quorum system of rank $r_{\mathcal{S}} = r$. Then $|U| \leq r^2 - r + 1$.*

Proof. We obtain (22) as in the proof of the previous theorem. Then, from $g_{\mathcal{S},\mathbf{v}} \leq r$ and Proposition 5.7, we get $|U| = n \leq r(r-1+1/r) = r^2 - r + 1$. ∎

The last theorem is seen to be a generalization of the result of Lovász (Proposition 5.8): the uniformity assumption is dispensed with, as the rank suffices, and the regularity assumption is relaxed to perfect balancing.

5.4. Size-balancing trade-offs for unbalanced quorum systems. We have seen that if we insist on perfect balancing then the best we can do is to use quorums of size $\sim \sqrt{n}$. What if we relax perfect balancing and are willing to accept a balancing ratio not worse than some number ρ, $0 < \rho < 1$? Is there a trade-off between the required level ρ and the quorum size it admits in the best case?

It is easy to get an analogue of Theorem 5.9 (or 5.10) by observing that when $\rho_{\mathcal{S},\mathbf{v}} \geq \rho$ one obtains an adaptation of (22) in the form $n\rho \leq g_{\mathcal{S},\mathbf{v}}\nu_{\mathcal{S}}^*$. From this it follows that $g_{\mathcal{S},\mathbf{v}} \geq \sqrt{n\rho}$. We shall get a better estimate by a refinement of the argument, based on the following lemma.

LEMMA 5.11. *Let $\rho, a_1, a_2, \ldots, a_n$ be real numbers such that $0 < \rho \leq 1$ and $\rho \leq a_j \leq 1$ for $j = 1, \ldots, n$. Then*

$$\frac{\sum_{j=1}^n a_j^2}{\left(\sum_{j=1}^n a_j\right)^2} \leq \frac{(1+\rho)^2}{4n\rho}.$$

Proof. Given ρ and n, consider the problem of maximizing

$$f(a_1, a_2, \ldots, a_n) = \frac{\sum_{j=1}^n a_j^2}{\left(\sum_{j=1}^n a_j\right)^2}$$

subject to $\rho \leq a_j \leq 1$, $j = 1, \ldots, n$. For any $1 \leq i \leq n$ we have

$$\frac{\partial f}{\partial a_i} = \frac{2\sum_{\substack{j=1 \\ j \neq i}}^n (a_i - a_j)a_j}{\left(\sum_{j=1}^n a_j\right)^3}.$$

Since the numerator in the above expression is an increasing function of a_i, it follows that the maximum under consideration is attained when $a_i = \rho$ or $a_i = 1$. Indeed, if $\rho < a_i < 1$ and $\frac{\partial f}{\partial a_i} = 0$, then $\frac{\partial f}{\partial a_i}$ is negative for smaller values of a_i and positive for larger values of a_i, so we are looking at a minimum of f as a function of a_i.

Thus, it suffices to consider points $(a_1, a_2, \ldots, a_n)$ where k of the a_j's equal ρ and the other $n-k$ equal 1. Letting $x = k/n$ we have for such points

$$f(a_1, a_2, \ldots, a_n) = \frac{\rho^2 k + n - k}{(\rho k + n - k)^2} = \frac{1-(1-\rho^2)x}{n\big(1-(1-\rho)x\big)^2}.$$

One can show by elementary analysis that this expression is maximized in the interval $0 \le x \le 1$ when $x = 1/(1+\rho)$, and attains there the value $(1+\rho)^2/4n\rho$. ∎

THEOREM 5.12. *Let $(U, \mathcal{S})$ be a quorum system with $|U| = n$. Let $0 < \rho \le 1$ and let $\mathbf{v}$ be a QLV such that $\rho_{\mathcal{S},\mathbf{v}} \ge \rho$. Then*

$$g_{\mathcal{S},\mathbf{v}} \ge \frac{2\sqrt{n\rho}}{1+\rho}.$$

Proof. Let $\hat{\mathcal{S}} = (\hat{s}_{ij})$, $i = 1, \ldots, m$, $j = 1, \ldots, n$, be the quorum matrix. By a normalization which does not affect $\rho_{\mathcal{S},\mathbf{v}}$ or $g_{\mathcal{S},\mathbf{v}}$, we may assume that the ELV $\mathbf{a} = \mathbf{v}\hat{\mathcal{S}} = (a_1, a_2, \ldots, a_n)$ induced by $\mathbf{v}$ satisfies $\rho \le a_j \le 1$, $j = 1, \ldots, n$. We observe that by the quorum intersection property we have $\hat{\mathcal{S}}\hat{\mathcal{S}}^\top \ge \hat{1}$, where $\hat{1}$ denotes the $m \times m$ all-1 matrix, and the inequality holds entry-by-entry. Therefore,

$$\sum_{j=1}^{n} a_j^2 = \mathbf{a} \cdot \mathbf{a}^\top = \mathbf{v}\hat{\mathcal{S}}\hat{\mathcal{S}}^\top\mathbf{v}^\top \ge \mathbf{v}\hat{1}\mathbf{v}^\top = |\mathbf{v}|^2.$$

Using this and Lemma 5.11 we have

$$g_{\mathcal{S},\mathbf{v}}^2 = \frac{\left(\sum_{j=1}^{n} a_j\right)^2}{|\mathbf{v}|^2} \ge \frac{\left(\sum_{j=1}^{n} a_j\right)^2}{\sum_{j=1}^{n} a_j^2} \ge \frac{4n\rho}{(1+\rho)^2}.$$

Upon taking square roots we obtain the required result. ∎

We now describe a construction showing that the bound given in Theorem 5.12 is rather tight.

Example 5.13. Ext-FPP. Let $0 < \rho < \frac{1}{2}$ and let r be a positive integer such that a FPP(r–1) exists. Let P and $\mathcal{L}$ be the sets of points and lines, respectively, of a FPP(r–1). Let K be a set of size $\left[(1-2\rho)/\rho\right](r^2 - r + 1)$, disjoint from P, and let $\mathcal{M}$ be the set of all subsets of K of size $(1-2\rho)r$. (We ignore adjustments that need to be made when these numbers are not integers. The effect of such adjustments is negligible when r is large.) Let $U = P \cup K$ and let $\mathcal{S}$ consist of all sets of the form $L \cup M$, where $L \in \mathcal{L}$ and $M \in \mathcal{M}$. Then $\mathcal{S}$ satisfies the quorum intersection requirement, because any two lines in $\mathcal{L}$ intersect. Since P has $r^2 - r + 1$ points and each line in $\mathcal{L}$ contains r points, we see that $|U| = n = [(1-\rho)/\rho]\,(r^2 - r + 1)$ and each quorum in $\mathcal{S}$ has size $2(1-\rho)r$.

Let $\mathbf{v}$ be a QLV assigning equal load to all the quorums in $\mathcal{S}$. Then it can be verified that $\rho_{\mathcal{S},\mathbf{v}} = \rho$. Indeed, it follows from considerations of symmetry that the induced ELV is constant over K and over P, and the ratio between the two constants can be computed as

$$\frac{|M| \cdot |P|}{|L| \cdot |K|} = \frac{(1-2\rho)r \cdot (r^2 - r + 1)}{r \cdot \frac{1-2\rho}{\rho}(r^2 - r + 1)} = \rho$$

(here $M \in \mathcal{M}$ and $L \in \mathcal{L}$). To evaluate the performance of this construction, we have to compare

$$g_{\mathcal{S},\mathbf{v}} = 2(1-\rho)r$$

with the bound of Theorem 5.12:

$$\frac{2\sqrt{n\rho}}{1+\rho} = \frac{2\sqrt{1-\rho}\sqrt{r^2 - r + 1}}{1+\rho}.$$

It is readily seen that the ratio between the two quantities approaches 1 as $\rho \to 0$ and $r \to \infty$. The ratio is in general less than $\big(1+\rho/2)(1+1/(2r)\big)$. $\square$

The theorem and the construction delineate with a good degree of precision a trade-off between the required level of balancing ρ (when $0 < \rho < \frac{1}{2}$) and the quorum size it admits in the best case. We remark that we do not know how to handle profitably the case when $\frac{1}{2} \le \rho < 1$: if the required level of balancing is in this interval, the construction with smallest quorum size that we know is the same as for perfect balancing (namely, the finite projective plane).

5.5. Size-balancing trade-offs for NDCs. In view of the distinguished role played by NDCs among quorum systems, it is interesting to investigate the relation between the level of balancing and the quorum size within this special class. We start by describing a construction, borrowed from [EL74], of an NDC with quorums of size $O(\sqrt{n})$ which is, as we shall show, perfectly balanced.

The method of construction is inductive. In the inductive step, we are given an $(r-1)$-uniform quorum system $(U', \mathcal{S}')$. We take a set R of size r, disjoint from U', and form the new universe $U = U' \cup R$. We define the collection $\mathcal{S}$ by: $S \in \mathcal{S}$ if and only if $S = S' \cup \{u\}$ for some $S' \in \mathcal{S}'$ and $u \in R$, or $S = R$. We thus obtain a new system $(U, \mathcal{S})$.

PROPOSITION 5.14. *Let $(U, \mathcal{S})$ be obtained from the quorum system $(U', \mathcal{S}')$ as above. Then:*

(a) $(U, \mathcal{S})$ *is an r-uniform quorum system.*
(b) *If $(U', \mathcal{S}')$ is an NDC then so is $(U, \mathcal{S})$.*
(c) *If $(U', \mathcal{S}')$ is perfectly balanced then so is $(U, \mathcal{S})$.*

Proof. Part (a) is straightforward. Part (b) can be verified using Proposition 2.6. Indeed, let S_1, S_2 be a partition of U. Since $(U', \mathcal{S}')$ is an NDC and S'_1, S'_2 is a partition of U' (where $S'_i = S_i \cap U'$), we may assume that S'_1, say, contains some $S' \in \mathcal{S}'$. Then, if $S_1 \cap R \neq \emptyset$ we conclude that S_1 contains a quorum of the form $S' \cup \{u\}$; if, on the other hand, $S_1 \cap R = \emptyset$ then S_2 contains the quorum R.

To prove part (c), let $\mathbf{v}'$ be a QLV for $(U', \mathcal{S}')$ which induces a load of 1 on each element of U'. Let $\mathbf{v}$ be the QLV for $(U, \mathcal{S})$ defined by: the load of $S' \cup \{u\}$, where $S' \in \mathcal{S}'$ and $u \in R$, is $1/r$ of the load of S' in $\mathbf{v}'$, and the load of R is $1 - |\mathbf{v}'|/r$ (this quantity is positive by virtue of Proposition 5.6). Then $\mathbf{v}$ induces a load of 1 on each element of U. $\square$

Example 5.15. Triangular. Let $(U^r, \mathcal{S}^r)$ be an r-uniform quorum system obtained by successive applications of the inductive step described above, starting from a system of one element. We call $(U^r, \mathcal{S}^r)$ a *triangular system.* It follows from Proposition 5.14 that $(U^r, \mathcal{S}^r)$ is an NDC and is perfectly balanced. The size of its universe is $|U^r| = n = (r+1)r/2$. $\square$

We observe that the quorum size achieved in the above construction is about $\sqrt{2n}$ and is thus within a multiplicative constant factor of the lower bound of $\sqrt{n}$ given in Theorem 5.9 (for all perfectly balanced quorum systems, not just NDCs). It seems plausible that the lower bound can be improved for the class of NDCs, but we are unable to do this. On the other hand, we can achieve a (very) slight improvement on the construction. Let

$$n(r) = \max\big\{|U|\colon (U, \mathcal{S}) \text{ is an } r\text{-uniform, perfectly balanced NDC}\big\}.$$

Then the above construction gives $n(r) \ge (r+1)r/2$. For $r = 3$ this becomes $n(3) \ge 6$, but the Fano plane (FPP(2)) shows that $n(3) \ge 7$ (in fact, we can deduce from Theorem 5.10 that $n(3) = 7$). When the inductive method described above is applied starting from the Fano plane, we obtain that $n(r) \geq (r+1)r/2+1$ for $r \geq 3$. In order to introduce a further improvement we need the following definition and easy facts.

DEFINITION. *Let $(U, \mathcal{S})$ be a quorum system, with $U = \{u_1, u_2, \ldots, u_n\}$. Let $(U_j, \mathcal{S}_j)$, $j = 1, \ldots, n$, be quorum systems, with the U_j's pairwise disjoint. The* composite quorum system *(CQS) formed by substituting $(U_j, \mathcal{S}_j)$, $j = 1, \ldots, n$, for the elements of $(U, \mathcal{S})$, denoted $CQS(\mathcal{S}, \{\mathcal{S}_j\})$, has as its universe $\bigcup_{j=1}^{n} U_j$ and as its quorums all sets obtained as follows: take any $S = \{u_{j_1}, u_{j_2}, \ldots, u_{j_k}\} \in \mathcal{S}$ and for each j_i, $i = 1, \ldots, k$, take any $S_{j_i} \in \mathcal{S}_{j_i}$, and form the (composite) quorum $\bigcup_{i=1}^{k} S_{j_i}$.*

PROPOSITION 5.16.

(a) *If $(U, \mathcal{S})$ is r-uniform and each $(U_j, \mathcal{S}_j)$ is s-uniform, then the CQS is rs-uniform.*

(b) *If $(U, \mathcal{S})$ is uniform and regular, and all of the $(U_j, \mathcal{S}_j)$ are regular with the same common degree and the same number of quorums, then the CQS is regular.*

(c) *If $(U, \mathcal{S})$ and each of the $(U_j, \mathcal{S}_j)$ are NDCs, then the CQS is an NDC.*

Now, consider the CQS formed by substituting seven copies of the Fano plane for the seven points of a Fano plane. By Proposition 5.16, this is a 9-uniform, regular (hence perfectly balanced) NDC. This shows that $n(9) \geq 49$, whereas $(r+1)r/2 = 45$ for $r = 9$. When the inductive method is applied successively starting from this CQS, we obtain that $n(r) \geq (r+1)r/2+4$ for $r \geq 9$.

CONJECTURE 5.17. *$n(r) = (r+1)r/2 + O(1)$.*

We observe that if $(U, \mathcal{S})$ is an r-uniform, perfectly balanced quorum system with $|U| = n$, then $\nu^*_{\mathcal{S}} = n/r$. Thus, the above conjecture can be reformulated as saying that if $(U, \mathcal{S})$ is an r-uniform, perfectly balanced NDC then $\nu^*_{\mathcal{S}} \leq (r+1)/2 + O(1/r)$. We believe, in fact, that this holds even without the assumptions of uniformity and perfect balancing. That is, we believe that the assumption of nondomination alone should permit the following improvement on Füredi's bound (Proposition 5.7).

CONJECTURE 5.18. *Let $(U, \mathcal{S})$ be an NDC of rank $r_{\mathcal{S}} = r$. Then $\nu^*_{\mathcal{S}} \leq (r+1)/2 + O(1/r)$.*

If we do not insist on perfect balancing, but continue to consider only NDCs, how low can we make the quorum size? It follows from a more general result in [T85] that any nonredundant NDC having rank r has universe of size smaller than $\binom{2r}{r}$. Recalling the nucleus coterie of Example 4.3, where the size of the universe is larger than $\binom{2r-2}{r-1}/2$, we see that Tuza's bound is within a multiplicative constant factor of being best possible. Stating the result differently, we can say that the smallest possible rank among all nonredundant NDCs with universe of size n is $\frac{1}{2}\log_2 n + \frac{1}{4}\log_2\log_2 n + O(1)$.

Suppose we require some level of balancing; that is, we consider NDCs with balancing ratio not worse than some number ρ, $0 < \rho < 1$. How low can we make the quorum size then? We are unable to improve on the lower bound given in Theorem 5.12 (which is not restricted to NDCs). A construction that attempts to approach that lower bound using NDCs is given next. It is not as good as the one (using dominated coteries) given by the Ext-FPP coterie of Example 5.13.

Example 5.19. CQS (Triangular, Nucleus). Let $(U^r, \mathcal{S}^r)$ be an r-uniform triangular NDC with $|U^r| = (r+1)r/2$ which is perfectly balanced, as in Example 5.15. Let $(U_s, \mathcal{S}_s)$ be an s-uniform nucleus NDC with $|U_s| = 2s - 2 + \binom{2s-2}{s-1}/2$ and $\rho_{\mathcal{S}_s} = 4/\binom{2s-2}{s-1}$, as in Example 4.3. Let $(U^r_s, \mathcal{S}^r_s)$ be the CQS formed by substituting $(r+1)r/2$ copies of $(U_s, \mathcal{S}_s)$ for the elements of $(U^r, \mathcal{S}^r)$. It follows from Proposition 5.16 that $(U^r_s, \mathcal{S}^r_s)$ is an rs-uniform NDC. It is easy to see that $\rho_{\mathcal{S}^r_s} = \rho_{\mathcal{S}_s}$. A rough computation shows that in terms of the size n of the universe and the balancing ratio

TABLE 1

Quorum systems	Hypergraph theory	Game theory
element	vertex	player
universe	vertex-set	player-set
quorum	hyperedge	minimal winning coalition
superquorum	nonindependent set of vertices	winning coalition
quorum system	intersecting hypergraph	proper simple game
monotone quorum system	intersecting filter	monotone proper simple game
coterie	intersecting antichain	(minimal winning coalitions of a) monotone proper simple game
nondominated coterie	3-chromatic intersecting antichain	(minimal winning coalitions of a) constant-sum game
rank	rank	size of a largest minimal winning coalition
voting system	threshold	weighted majority game
ordered	shifted	having a complete desirability relation
nonredundant	no isolated vertices	without dummies
quorum load vector	fractional matching (up to normalization)	weight assignment to the minimal winning coalitions
perfectly balanced	quasi-regularizable	having a balanced collection of minimal winning coalitions

ρ, the quorums of the CQS have size $\sim \sqrt{n\rho} \cdot \frac{1}{2} \log_2 \frac{1}{\rho}$. The ratio of this to the lower bound of Theorem 5.12 is $O\left(\log \frac{1}{\rho}\right)$. ∎

Appendix: A polyglot dictionary. The motivation for the research reported in this paper came from computer science, but the concepts involved have also been studied in other areas of science, under various interpretations and using various systems of terminology. To help overcome the language barriers, we thought it useful to provide here translations of the concepts into the languages of two other areas: hypergraph theory and game theory. Our little dictionary (Table 1) is only schematic, and for more information we refer the reader to books such as [B89] and [Ow82]. We should also mention that, due to scope limitations, our dictionary leaves out several other areas in which these concepts have come up. These include Boolean functions theory, reliability theory, neural networks, percolation theory, etc.

As an illustration of the possible appeal of our work to researchers in other areas, we rephrase Theorem 3.4 in game-theoretic terms and Conjecture 5.18 in hypergraph terms.

THEOREM 3.4*. *Let G be a constant-sum game without dummies, having a complete desirability relation. Then G has a balanced collection of minimal winning coalitions.*

CONJECTURE 5.18*. *Let $\mathcal{H}$ be a* 3*-chromatic intersecting hypergraph of rank r. Then $\nu^*_{\mathcal{H}} \leq (r+1)/2 + O(1/r)$.*

Acknowledgment. We thank the anonymous referees for their helpful comments.

REFERENCES

[B89] C. Berge, *Hypergraphs*, North-Holland, Amsterdam, 1989.

[B86] B. Bollobás, *Combinatorics*, Cambridge University Press, Cambridge, UK, 1986.

[C83] V. Chvátal, *Linear Programming*, W.H. Freeman, New York, 1983.

[C93] H. Cohen, *Quorum Systems: Domination, Fault-Tolerance, Balancing*, MS Thesis, The Weizmann Institute of Science, Rehovot, Israel, 1993.

[EL74] P. Erdős and L. Lovász, *Problems and results on* 3*-chromatic hypergraphs and some related questions*, in Infinite and Finite Sets, Colloq. Math. Soc. János Bolyai 10, North-Holland, Amsterdam, 1974, pp. 609–627.

[F01] J. Farkas, *Uber die Theorie der Einfachen Ungleichungen*, J. Reine Angew. Math., 124 (1901), pp. 1–27.

[F81] Z. Füredi, *Maximum degree and fractional matchings in uniform hypergraphs*, Combinatorica, 1 (1981), pp. 155–162.

[GB85] H. Garcia-Molina and D. Barbara, *How to assign votes in a distributed system*, J. Assoc. Comput. Mach., 32 (1985), pp. 841–860.

[H86] M. Hall, *Combinatorial Theory*, John Wiley, New York, 1986.

[H84] M. P. Herlihy, *Replication methods for abstract data types*, Ph.D. thesis, MIT, Cambridge, MA, 1984.

[IK90] T. Ibaraki and T. Kameda, *Theory of coteries*, Tech. Rep. CSS/LCCR TR90-09, Simon Fraser University, Burnaby, BC, Canada, 1990.

[L73] L. Lovász, *Coverings and colorings of hypergraphs*, in Proc. 4th Southeastern Conf. on Combinatorics, Graph Theory and Computing, Florida Atlantic University, Utilitas Mathematica, Winnipeg, Canada, 1973, pp. 47–56.

[L75] L. Lovász, *On the minimax theorems of combinatorics*, Mat. Lapok, 26 (1975), pp. 209–264 (in Hungarian).

[MV88] S. J. Mullender and P.M.B. Vitányi, *Distributed match-making*, Algorithmica, 3 (1988), pp. 367–391.

[Os85] A. Ostmann, *Decisions by players of comparable strength*, Z. Nationalökonom., 45 (1985), pp. 267–284.

[Ow82] G. Owen, *Game Theory*, Academic Press, Boston, MA, 1982.

[PW93] D. Peleg and A. Wool, *The availability of quorum systems*, Tech. Rep. CS93-17, The Weizmann Institute of Science, Rehovot, Israel, 1993.

[P70] J. Pelikán, *Properties of balanced incomplete block designs*, in Combinatorial Theory and Its Applications, Colloq. Math. Soc. János Bolyai 4, North-Holland, Amsterdam, 1970, pp. 869–889.

[R86] M. Raynal, *Algorithms for mutual exclusion*, MIT Press, Cambridge, MA, 1986.

[T85] Zs. Tuza, *Critical hypergraphs and intersecting set-pair systems*, J. Combin. Theory Ser. B, 39 (1985), pp. 134–145.