一個以一般化階級式為基礎的分散式K-互斥方法
A Generalized Hierarchical Quorum Strategy for K-Mutual Exclusion in Distributed Systems

張玉盈，陳柏旭
Ye-In Chang and Bor-Hsu Chen
國立中山大學應用數學系
Dept. of Applied Mathematics
National Sun Yat-Sen University
Kaohsiung, Taiwan, R.O.C
{E-mail: changyi@math.nsysu.edu.tw}
{Tel: 886-7-5252000 (ext. 3819)}
{Fax: 886-7-5253809}

摘 要

在這篇論文中，我們提出一般化階級式的分散式k-互斥方法。其邏輯架構乃以一般化階級架構為基礎。其所得的法定節點集皆為$(\frac{NS}{k})^{0.63}$，其中NS為實際節點的個數。而且，這方法總是可以容忍$NS-k(\frac{NS}{k})^{0.63}$個節點錯誤。從我們的效率分析中顯示，與$K$-majority，Cohorts以及$DIV$方法比起來，在某些時候，我們的方法有較高的可靠度。而且，在沒有節點錯誤下以及當$NS>15$情形下，我們所得的法定節點集總是這四種方法中最小的。而在最差的情況下，當節點錯誤數目大於17時，我們所得的法定節點集總是這四種方法中最小的。

(關鍵詞: K-互斥，可靠性，分散式系統，容錯，法定一致。)

Abstract

In this paper, we propose a generalized hierarchical quorum strategy for distributed k-mutual exclusion. This strategy is based on a logical generalized hierarchy structure. The quorum size constructed from the strategy is always $(\frac{NS}{k})^{0.63}$, *where NS is the number of physical nodes in the system. Moreover, this strategy can be always fault-tolerant up to* $NS-k(\frac{NS}{k})^{0.63}$ *node failures. From our performance analysis, we show that the generalized hierarchical quorum strategy can provide a higher availability than k-majority, cohorts, and DIV strategies, sometimes. Moreover, the quorum size of this strategy is always the smallest one among these four strategies, when no node failure occurs and* $NS > 15$. *While in the worst case, the quorum size of this strategy is always smallest one among these four strategies, when* $NS > 17$.

(*Key Words:* K-mutual exclusion, availability, distributed systems, fault tolerance, quorum consensus.)

[1] This research was supported by National Science Council of the Republic of China, NCS-81-0408-E-110-508

1 Introduction

A distributed system consists of a collection of geographically dispersed autonomous nodes connected by a communication network. The nodes have no shared memory, no global clock, and communicated with one another by passing messages. Message propagation delay is finite but unpredictable. In the problem of k-mutual exclusion, concurrent access to shared resource or the critical section (CS) must be synchronized such that at any time at most k processes can access the CS, where $k \geq 1$.

To make distributed k-mutual exclusion protocols fault-tolerant to node and communication failures, many protocols based on the replica control strategies, for example *coterie*, have been proposed. In [5], they extended the *majority quorum* strategy to *k-majority quorum* strategy; any permission from $\lceil\frac{n+1}{k+1}\rceil$ nodes would form a quorum for k-mutual exclusion, when n is the number of nodes in the system. In [4], they proposed a *cohorts quorum* for k-mutual exclusion based on a *cohorts structure*, $Coh(k,l)$, which has l pairwise disjoint cohorts with first cohort having k members and the others having more than $(2k-2)$ members. In [2], they partition n nodes into k classes with each class using any traditional approach to enforce 1-mutual exclusion. When the traditional approach is the majority quorum strategy, the constructed quorums will be called *DIV of majority* quorums.

To reduce the overhead of achieving k-mutual exclusion while supporting fault tolerance, in this paper, we propose a strategy called *generalized hierarchical quorum* for k-mutual exclusion, which imposes a logical generalized hierarchy structure on the network. The quorum size constructed from the strategy is always $(\frac{NS}{k})^{0.63}$, where NS is the number of physical nodes in the system. Moreover, this strategy can be always fault-tolerant up to $NS-k(\frac{NS}{k})^{0.63}$ node failures. From our performance analysis, we show that the generalized hierarchical quorum strategy can provide a higher availability than k-majority, cohorts,

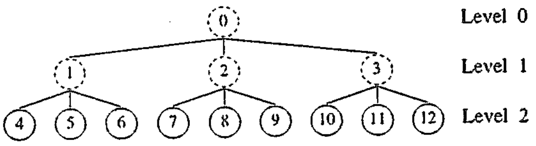

Figure 1: A generalized hierarchy of level 3

and DIV strategies, sometimes. Moreover, the quorum size of this strategy is always the smallest one among these four strategies, when no node failure occurs and $NS > 15$. While in the worst case, the quorum size of this strategy is always smallest one among these four strategies, when $NS > 17$.

2 Generalized Hierarchical Quorums

2.1 Definitions

In this section, we first define the generalized hierarchy and give the definition of the hierarchical quorum for 1-mutual exclusion [6]. Next, based on the hierarchical quorum for 1-mutual exclusion, we present the generalized hierarchical quorum for k-mutual exclusion.

Definition 1. A Hierarchical Quorum [6]. *The hierarchical quorum strategy is based on logically organizing a set of copies of an object in a system into multilevel tree with the root as level 0. The physical copies of an object are stored only in the leaves of this tree, while the higher level nodes of the tree correspond to logical groups. A hierarchical quorum (recursively) for 1-mutual exclusion consists of the hierarchical quorum of the* $\lceil \frac{s+1}{2} \rceil$ *subhierarchies, where* s *is the number of the subhierarchies.*

Note that, here we let each site in the distributed system be mapped to a physical node in the hierarchy, and the number of sites be denoted as NS.

Example 1: For a hierarchy of level 3 as shown in Figure 1, the set R of hierarchical quorum for 1-mutual exclusion is as follows: R = { {4, 5, 7, 8}, {4, 5, 7, 9}, {4, 5, 8, 9}, {4, 5, 10, 11}, {4, 5, 10, 12}, {4, 5, 11, 12}, {4, 6, 7, 8}, {4, 6, 7, 9}, {4, 6, 8, 9}, {4, 6, 10, 11}, {4, 6, 10, 12}, {4, 6, 11, 12}, {5, 6, 7, 8}, {5, 6, 7, 9}, {5, 6, 8, 9}, {5, 6, 10, 11}, {5, 6, 10, 12}, {5, 6, 11, 12}, {7, 8, 10, 11}, {7, 8, 10, 12}, {7, 8, 11, 12}, {7, 9, 10, 11}, {7, 9, 10, 12}, {7, 9, 11, 12}, {8, 9, 10, 11}, {8, 9, 10, 12}, {8, 9, 11, 12} }. Totally, R contains 27 quorums.

Definition 2. A Degree-Three Hierarchy. *A degree-three hierarchy is a finite set of one or more nodes such that*

1. *there is a specially designated node called the root in level 0.*
2. *the remaining nodes are partitioned into three subsets, where each of these subsets is a degree-three hierarchy.*
3. *the nodes without any child are physical nodes, the others are virtual nodes.*

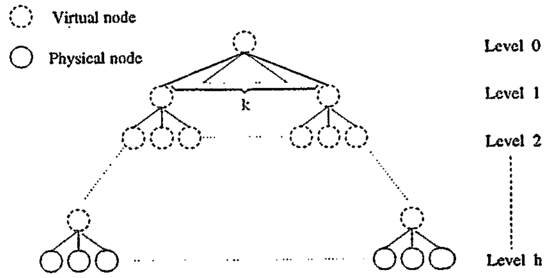

Figure 2: A generalized hierarchy of level (h+1)

Definition 3. A Generalized Hierarchy. *A generalized hierarchy is a finite set of one or more nodes such that*

1. *there is a specially designated node called the root in level 0.*
2. *the remaining nodes are partitioned into* S_1, S_2, *...., and* S_k, *where each of these sets is a degree-three hierarchy.*

Therefore, there are $k(3^{i-1})$ *nodes in level i. Consequently, for a hierarchy of level (h+1), there are totally* $(k(\frac{3^h-1}{2})+1)$ *nodes which include* $(k(\frac{3^{h-1}-1}{2})+1)$ *virtual nodes and* $k(3^{h-1})$ *physical nodes. Moreover, each node in the hierarchy of level (h+1) is numbered from top to down and left to right as 0, 1, 2, ...,* $(k(\frac{3^h-1}{2}))$ *as shown in Figure 2.*

Definition 4. A Generalized Hierarchical Quorum. *A hierarchy of level (h+1) is a collection of interconnected nodes arranged by levels with n nodes where* $h \geq 0$, $n = k(\frac{3^h-1}{2}) + 1$ *and* $NS = k(3^{h-1})$. *A set Q is said to be a generalized hierarchical quorum for k-mutual exclusion if the following conditions is satisfied: Q contains the hierarchical quorum for 1-mutual exclusion for any one of the subhierarchies* S_1, S_2, *....,* S_k.

Example 2: For a hierarchical of level 3 as shown in Figure 1, the set R of a generalized hierarchical quorums for 3-mutual exclusion is as follows: R = { {4, 5}, {4, 6}, {5, 6}, {7, 8}, {7, 9}, {8, 9}, {10, 11}, {10, 12}, {11, 12} }. Totally, R contains 9 quorums.

Example 3: For a hierarchical of level 4 as shown in Figure 3, the set R of a generalized hierarchical quorums for 2-mutual exclusion is as follows: R = { {9, 10, 12, 13}, {9, 10, 12, 14}, {9, 10, 13, 14}, {9, 10, 15, 16}, {9, 10, 15, 17}, {9, 10, 16, 17}, {9, 11, 12, 13}, {9, 11, 12, 14}, {9, 11, 13, 14}, {9, 11, 15, 16}, {9, 11, 15, 17}, {9, 11, 16, 17}, {10, 11, 12, 13}, {10, 11, 12, 14}, {10, 11, 13, 14}, {10, 11, 15, 16}, {10, 11, 15, 17}, {10, 11, 16, 17}, {12, 13, 15, 16}, {12, 13, 15, 17}, {12, 13, 16, 17}, {12, 14, 15, 16}, {12, 14, 15, 17}, {12, 14, 16, 17}, {13, 14, 15, 16}, {13, 14, 15, 17}, {13, 14, 16, 17}, {18, 19, 21, 22}, {18, 19, 21, 23}, {18, 19, 22, 23}, {18, 19, 24, 25}, {18, 19, 24, 26}, {18, 19, 25, 26}, {18, 20, 21, 22}, {18, 20, 21, 23}, {18, 20, 22, 23}, {18, 20, 24, 25}, {18,

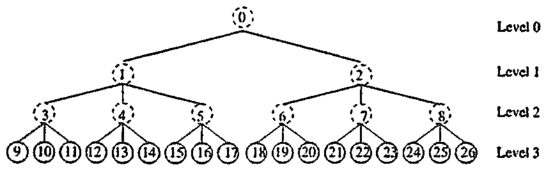

Figure 3: A hierarchy of level 4 for 2-mutual exclusion with $NS = 18$

20, 24, 26}, {18, 20, 25, 26}; {19, 20, 21, 22}, {19, 20, 21, 23}, {19, 20, 22, 23}, {19, 20, 24, 25}, {19, 20, 24, 26}, {19, 20, 25, 26}, {21, 22, 24, 25}, {21, 22, 24, 26}, {21, 22, 25, 26}, {21, 23, 24, 25}, {21, 23, 24, 26}, {21, 23, 25, 26}, {22, 23, 24, 25}, {22, 23, 24, 26}, {22, 23, 25, 26} }. Totally, R contains 54 quorums.

2.2 Correctness

In this section, we prove that the set of the generalized hierarchical quorum for k-mutual exclusion is a k-coteries. Here, we will refer to such a k-coteries as generalized hierarchical quorums.

Definition 5. *A* **k-coteries** *C is a family of non-empty subsets of an underlying set U, which is a set containing all system nodes $1, 2, ..., n$. Each member Q in C is called a quorum, and the following properties should hold for the quorums [5].*

1. **The non-intersection Property.** *For any $h(< k)$ pairwise disjoint quorums $Q_1, ..., Q_h$ in C, there exists one quorum Q_{h+1} in C such that $Q_1, ..., Q_{h+1}$ are pairwise disjoint.*

2. **The intersection Property.** *There are no m, $m > k$, pairwise disjoint quorums in C (i.e., there are at most k pairwise disjoint quorums in C).*

3. **The minimality Property.** *There are no two quorums Q_i and Qj in C such that Q_i is a super set of Q_j where $i \neq j$.*

Definition 6. *Let* **C** *and* **D** *be two k-coteries.* **D dominates C** *if and only if (*$\mathbf{C} \neq \mathbf{D}$*) and (*$\forall R \in \mathbf{C}, \exists S \in \mathbf{D}, S \subseteq R$*). A coterie is said to be* **nondominated** *(ND) if no coterie can dominate it [5].*

Lemma 1. *Let U_1, U_2, ..., and U_k be nonempty sets of nodes such that $U_i \cap U_j = \emptyset$, $i \neq j$. Suppose that C_i is a 1-coterie under U_i, $1 \leq i \leq k$. Let $C = \{ Q_j^i \mid Q_j^i \in C_i, 1 \leq i \leq k, 1 \leq j \leq |C_i| \}$ and $U = \bigcup_{i=1}^{k} U_i$. Then C is a k-coteries under U.*

Proof. First, we show that the non-intersection property is satisfied. Let Q_1, Q_2, ..., and Q_h be quorums in C and they are pairwise disjoint, where $1 \leq h < k$. Note that such quorums exist as long as for any two quorums, Q_i, Q_j, they satisfy the conditions, $Q_i \in C_{p_i}$, $Q_j \in C_{p_j}$, and $p_i \neq p_j$, where $1 \leq i, j < h$, $1 \leq p_i, p_j \leq k$. The reason is that $C_{p_i} \cap C_{p_j} = \emptyset$, which is resulted from $C_{p_i} \in U_{p_i}$, $C_{p_j} \in U_{p_j}$, and $U_{p_i} \cap U_{p_j} = \emptyset$; let W^* be the union of such C_{p_i} (C_{p_j}). Moreover, let C^* be the union of all the C_s, where $1 \leq s \leq k$, and let $W = C^* - W^*$. Note that $W \neq \emptyset$, since there are h different C_{p_i} in W^*, k different C_s in C^*, and $h < k$. Then, we let quorum $Q_{h+1} \in C_{p_{h+1}}$, where $C_{p_{h+1}} \in W$. Since $C_{p_{h+1}} \cap C_{p_i} = \emptyset$, $\forall C_{p_i} \in W^*$, which is resulted from $C_{p_{h+1}} \in U_{p_{h+1}}$, $C_{p_i} \in U_{p_i}$, and $U_{p_{h+1}} \cap U_{p_i} = \emptyset$, we have $Q_{h+1} \cap Q_i = \emptyset$. Therefore, the non-intersection property is satisfied.

Next, we show that the intersection property is satisfied. Assume there exist m pairwise disjoint quorums Q_1, Q_2, ..., and $Q_m \in C$, where $m > k$. If for any two quorums, Q_i, Q_j, they satisfy the conditions, $Q_i \in C_{p_i}$, $Q_j \in C_{p_j}$ and $p_i \neq p_j$, we have $Q_i \cap Q_j = \emptyset$, where $1 \leq p_i, p_j \leq k$. The reason is that $C_{p_i} \cap C_{p_j} = \emptyset$, which is resulted from $C_{p_i} \in U_{p_i}$, $C_{p_j} \in U_{p_j}$, and $U_{p_i} \cap U_{p_j} = \emptyset$. Since there are at most k different C_i under U, there exists quorums, Q_i, Q_j, they satisfy the condition, $Q_i \in C_{p_i}$, $Q_j \in C_{p_j}$, and $p_i = p_j$, for $m > k$. For such quorums Q_i, Q_j, we have $Q_i \cap Q_j \neq \emptyset$, since they belong to the same C_{p_i} (C_{p_j}) that must satisfy the intersection property, which contradicts the assumption that Q_i and Q_j are disjoint. Consequently, there are at most k pairwise quorums in C, and the intersection property is satisfied.

Finally, we show that the minimality property is satisfied. Let Q_i, $Q_j \in C$. There are two cases to be considered: (1) If Q_i, $Q_j \in C_s$, we have $Q_i \not\subset Q_j$ and $Q_j \not\subset Q_i$ due to that C_s satisfies the minimality property, where $1 \leq s \leq k$. (2) If $Q_i \in C_{p_i}$ and $Q_j \in C_{p_j}$, then $Q_i \cap Q_j = \emptyset$, where $1 \leq p_i, p_j \leq k$, and $p_i \neq p_j$. The reason is that $C_{p_i} \cap C_{p_j} = \emptyset$, which is resulted from $C_{p_i} \in U_{p_i}$, $C_{p_j} \in U_{p_j}$, and $U_{p_i} \cap U_{p_j} = \emptyset$, where $1 \leq p_i, p_j \leq k$ and $p_i \neq p_j$. Therefore, we have $Q_i \not\subset Q_j$ and $Q_j \not\subset Q_i$. Consequently, the minimality property is satisfied. □

Theorem 1. *The set of the generalized hierarchical quorums for k-mutual exclusion is a k-coteries.*

Proof. Let C_i be the set of the hierarchical quorums of 1-mutual exclusion for the subhierarchy S_i, $1 \leq i \leq k$. Based on the definition of the hierarchical quorum strategy, C_i is a 1-coterie [6]. Based on **Lemma 1**, the set of the generalized hierarchical quorums for k-mutual exclusion is a k-coteries. □

Lemma 2. *Let C be a 1-coterie under a nonempty set U. The C is dominated iff there exists a set $H \subseteq U$ such that (1) $\forall Q \in C \Rightarrow Q \not\subseteq H$ and (2) $\forall Q \in C \Rightarrow Q \cap H \neq \emptyset$ [7].*

Lemma 3. *Let C be a 1-coterie under a nonempty set U. If C is a nondominated 1-coterie, then for every set $H \subseteq U$, it satisfies the conditions (1) $\exists Q \in C$ and $Q \subseteq H$ or (2) $\exists Q \in C$ and $Q \cap H = \emptyset$.*

Proof. Based on **Lemma 2**, we have that C is dominated if there exists a set $H \subseteq U$ such that (1) $\forall Q \in C \Rightarrow Q \not\subseteq H$ and (2) $\forall Q \in C \Rightarrow Q \cap H \neq \emptyset$. Therefore, if C is a nondominated 1-coterie, then $\nexists H \subseteq U$ such that (1) $\forall Q \in C \Rightarrow Q \not\subseteq H$ and (2) $\forall Q \in C \Rightarrow Q \cap H \neq \emptyset$. That is, if C is a nondominated 1-coterie, then for every $H \subseteq U$, both of the conditions, (1) $\forall\, Q \in C \Rightarrow Q \not\subseteq H$ and (2) $\forall\, Q \in C \Rightarrow Q \cap H \neq \emptyset$, can not be true. In other words, the condition, (1) $\forall\, Q \in C \Rightarrow Q \not\subseteq H$, can not be true, which means that the condition, $\exists\, Q \in C$ and $Q \subseteq H$, is true. Moreover, the condition, (2) $\forall\, Q \in C \Rightarrow Q \cap H \neq \emptyset$, can not be true, which means that the condition, $\exists\, Q \in C$ and $Q \cap H = \emptyset$, is true. Consequently, If C is a nondominated 1-coterie, then for every set $H \subseteq U$, it satisfies the conditions, (1) $\exists\, Q \in C$ and $Q \subseteq H$ or (2) $\exists\, Q \in C$ and $Q \cap H = \emptyset$. □

Lemma 4. *Let C be a k-coteries under a nonempty set U. The C is dominated iff there exists a set $H \subseteq U$ such that (1) $\forall\, Q \in C \Rightarrow Q \not\subseteq H$ and (2) for any collection of k pairwise disjoint quorums $\{ Q_1, Q_2, ..., Q_k \} \subseteq C$, $H \cap Q_i \neq \emptyset$ for some $1 \leq i \leq k$ [7].*

Lemma 5. *Let U_1, U_2, ..., and U_k be nonempty sets of nodes such that $U_i \cap U_j = \emptyset$, $i \neq j$. Suppose that C_i is a nondominated 1-coterie under U_i, $1 \leq i \leq k$. Let $C = \{ Q^i_j \mid Q^i_j \in C_i, 1 \leq i \leq k, 1 \leq j \leq |C_i| \}$ and $U = \bigcup_{i=1}^{k} U_i$. Then C is a nondominated k-coteries under U.*

Proof. Assume that C is dominated. Based on **Lemma 4**, there exist a set $H \subseteq U$ such that (1) $\forall\, Q \in C \Rightarrow Q \not\subseteq H$ and (2) for any collection of k pairwise disjoint quorums $\{ Q_1, Q_2, ..., Q_k \} \subseteq C$, $H \cap Q_i \neq \emptyset$, $1 \leq i \leq k$. Note that such k disjoint quorums exist as long as for any two quorums, Q_i, Q_j, they satisfy the conditions, $Q_i \in C_{p_i}$, $Q_j \in C_{p_j}$, and $p_i \neq p_j$, where $1 \leq i,j \leq k$, and $1 \leq p_i, p_j \leq k$. The reason is that $C_{p_i} \cap C_{p_j} = \emptyset$, which is resulted from $C_{p_i} \in U_{p_i}, C_{p_j} \in U_{p_j}$, and $U_{p_i} \cap U_{p_j} = \emptyset$.

Let $H_i = H \cap U_{p_i}$, where $Q_i \in C_{p_i}$, $C_{p_i} \in U_{p_i}$ and $1 \leq i \leq k$. Since C_i is a nondominated 1-coterie, there exists $Q^*_i \in C_{p_i}$ such that $H_i \cap Q^*_i = \emptyset$ for $1 \leq i \leq k$, which is resulted from the second condition of **Lemma 3**. (Note that, here, the first condition of **Lemma 3** will not be true; that is, the condition, $\exists Q^*_i \in C_{p_i}$ and $Q^*_i \subseteq H_i$, will not be true. The reason is that if $Q^*_i \in C_{p_i}$ and $Q^*_i \subseteq H_i$, then $Q^*_i \in C$ and $Q^*_i \subseteq H_i$ is true, which contradicts the assumption, $\forall\, Q \in C \Rightarrow Q \not\subseteq H$.) Let $P = \{ Q^*_1, Q^*_2, ..., Q^*_k \}$, which is a collection of k pairwise disjoint quorums. Such k disjoint quorums exist as long as for any two quorums Q^*_i, Q^*_j, they satisfy the conditions, $Q^*_i \in C_{p_i}$, $Q^*_j \in C_{p_j}$, and $p_i \neq p_j$, where $1 \leq i,j \leq k$ and $1 \leq p_i, p_j \leq k$. The reason is that $C_{p_i} \cap C_{p_j} = \emptyset$, which is resulted from $C_{p_i} \in U_{p_i}$, $C_{p_j} \in U_{p_j}$ and $U_{p_i} \cap U_{p_j} = \emptyset$. Furthermore, $H \cap Q^*_i = \emptyset$ for all $Q^*_i \in P$ due to $H = \bigcup_{i=1}^{k} H_i$, $H_i \cap Q^*_i = \emptyset$, and $H_i \cap H_j \neq \emptyset, i \neq j, 1 \leq i,j \leq k$, which contradicts the assumption that for any collect of k pairwise disjoint quorums $\{ Q_1, Q_2, ..., Q_k \} \subseteq C$, $H \cap Q_i \neq \emptyset$, where $1 \leq i \leq k$. (Note that, the condition, $H = \bigcup_{i=1}^{k} H_i$ and $H_i \cap H_j = \emptyset$, $i \neq j$, is true due to that $H = H \cap U = H \cap (U_{p_1} \cup U_{p_2} \cup ... \cup U_{p_k}) = (H \cap U_{p_1}) \cup (H \cap U_{p_2}) \cup ... \cup (H \cap U_{p_k}) = H_1 \cup H_2 \cup ... \cup H_k = \bigcup_{i=1}^{k} H_i$ and $H_i \cap H_j = (H \cap U_{p_i}) \cap (H \cap U_{p_j}) = H \cap U_{p_i} \cap U_{p_j} = \emptyset$, which is resulted from $U_{p_i} \cap U_{p_j} = \emptyset$, $i \neq j$, $1 \leq i,j \leq k$.) Consequently, C is a nondominated k-coteries. □

Theorem 2. *The set of the generalized hierarchical quorums for k-mutual exclusion is a nondominated k-coteries.*

Proof. Let C_i be the set of the hierarchical quorum of 1-mutual exclusion for the subhierarchy S_i, $1 \leq i \leq k$. Based on the definition of the hierarchical quorum strategy, C_i is a 1-coterie and ND [6]. Based on **Lemma 5**, the set of the generalized hierarchical quorums for k-mutual exclusion is a nondominated k-coteries. □

2.3 Availability of the Generalized Hierarchical Quorum

In this section, we analyze the availability of the hierarchical quorum for 1-mutual exclusion [6] and the generalized hierarchical quorum for k-mutual exclusion. Here, we assume that all nodes have the same up-probability p, which is the probability that a single node is up operational. The availability of a coterie is defined as the probability that a quorum can be successfully formed.

For the hierarchical quorum strategy, the availability of a hierarchy is the probability that at least $\lceil \frac{s+1}{2} \rceil$ subhierarchical quorum can be formed from the hierarchy, where s is the number of the subhierarchies. Therefore, the availability of the hierarchical quorum under a degree-three hierarchy structure is the probability that at least 2 subhierarchical quorums can be formed the degree-three hierarchy structure.

Let $AVH(h)$ be the function evaluating the probability of a degree-three hierarchy of level $(h+1)$. If a degree-three hierarchy consists of only one node, it degenerates to a central controller and the availability of the availability of itself, i.e., $AVH(0) = p$. Thus we can get the condition

$$AVH(h) = \sum_{j=2}^{3} AVH(h-1)^j (1-AVH(h-1))^{3-j}$$

Next, for the availability of the generalized hierarchical quorum strategy, let (k,l)-availability, $1 \leq l \leq k$, be the probability that l pairwise disjoint quorums of a k-coteries can be formed successfully; it is used as a measure for the fault-tolerant ability of a solution using k-coteries.

Let $AV(h,l)$ be the function evaluating the probability that l pairwise disjoint quorums under generalized hierarchical quorum of the level $(h+1)$ can be formed simultaneously. If $NS = k(3^{h-1})$, the function $AV(h,l)$ has the following two boundary conditions:

1. $AVH(0) = p$ and $AVH(h) = \sum_{j=2}^{3} AVH(h-1)^j (1-AVH(h-1))^{3-j}$
2. $AV(h,l) = \sum_{m=l}^{k} \binom{k}{m} AVH(h-1)^m (1-AVH(h-1))^{k-m}$.

3 A Comparison

In this section, we make a comparison of the generalized hierarchical quorum, k-majority, cohorts, and DIV strategies in terms of availability and quorum size.

For the availability of the k-majority strategy, function $AV(k,h)$ has the following condition:

$AV(k,h) = \sum_{i=h\times\lceil\frac{n+1}{k+1}\rceil}^{n} C(n,i) \times p^i \times (1-p)^{n-i}$.

For the availability of the cohorts strategy, let $AV(h,l)$ be the function evaluating the probability that h pairwise disjoint quorums under $Coh(k,l)$ can be formed simultaneously. Function $AV(h,l)$ has the following three conditions:

1. $AV(0,l) = 1$.
2. $AV(h,1) = PR(S_1,h,S_1)$. (Note that a quorum takes only one member from the first cohort to make it the primary cohort because $S_1 - k + 1 = k - k + 1 = 1$. We also use S_i to denote $|C_i|$ for $1 \le i \le l$, where C_i is the ith item of $Coh(k,l) = (C_1,...,C_l)$ and we use $PR(s,a,b)$ to denote $\sum_{i=a}^{b} [C(s,i) \times p^i \times (1-p)^{s-i}]$).
3. $AV(h,l) = AV(h-1,l-1) \times PR(S_l, S_l - k + h, S_l) + AV(h,l-1) \times PR(S_l, h, S_l - k + h - 1)$.

For the availability of DIV of majority quorum, function $AV(k,h)$ has the following two conditions:

1. $AVM = \sum_{i=\lfloor\frac{n}{2k}\rfloor+1}^{\lceil\frac{n}{k}\rceil} C(\lceil\frac{n}{k}\rceil, i) \times p^i \times (1-p)^{\lceil\frac{n}{k}\rceil - i}$.
2. $AV(k,h) = \sum_{i=h}^{k} C(k,i) \times AVM^i \times (1 - AVM)^{k-i}$.

We make a comparison of the availability of the generalized hierarchical quorum strategy, k-majority, cohorts, and DIV strategies with $NS = 108$, and $l = 1$, 2, 3, and 4, respectively, where the generalized hierarchical quorum strategy is denoted as GHQC. For this comparison, the generalized hierarchical strategy is of level 5, and we let $Coh(4,14) = (C_1, C_2, ..., C_{14})$, where $|C_1| = 4$, and $|C_i| = 8$, $2 \le i \le 14$. The results are summarized in Table 1. From this table, we show that the availability of the generalized hierarchical quorum strategy is always better than that of DIV strategy, when $l = 1$; the availability of the generalized hierarchical quorum strategy is always better than that of cohorts and DIV strategies, when l=2; the availability of the generalized hierarchical quorum strategy is always the highest one among these four strategies, when $l = 3$ and $p < 0.5$; the availability of the generalized hierarchical quorum strategy is always better than that of cohorts and 4-majority strategies, when l=3 and $p \ge 0.5$; the availability of the generalized hierarchical quorum strategy is always better than that of cohorts and 4-majority strategies, when l=4.

Table 2 shows a comparison of these four k-mutual exclusion strategies in terms of quorum size. The first two criteria are the quorum sizes in the best and worst cases. The number of messages required to construct a quorum is proportional to the size of the quorums. The quorum size of generalized hierarchical quorum strategy is always $(\frac{NS}{k})^{0.63}$. Note that, in the generalized hierarchical quorum strategy, NS nodes are divided into k subhierarchies. In each subhierarchy, the quorum size is always $(\frac{NS}{k})^{0.63}$ [6]. Therefore, the quorum size of generalized hierarchical quorum strategy is always $(\frac{NS}{k})^{0.63}$. The quorum size of the cohorts strategy varies from 2 (when $k = 1$) or k (when $k > 1$) to n as the number of node failures in increased [4]. The quorum size of the k-majority strategy is always $\lceil\frac{n+1}{k+1}\rceil$ [5], and the quorum size of the DIV strategy is always $\lceil\frac{n+k}{2k}\rceil$ [2]. Note that, since k-majority, cohorts, and DIV strategies have no virtual nodes, the number of nodes is equal the number of physical nodes. That is $n = NS$.

l	P	The Availability
1	$0 < p < 1$	4-majority > cohorts > GHQC > DIV
2	$p < 0.59$	4-majority > GHQC > cohorts > DIV
	$p \ge 0.59$	4-majority > GHQC > DIV > cohorts
3	$p < 0.5$	GHQC > DIV > cohorts > 4-majority
	$p \ge 0.5$	DIV > GHQC > cohorts > 4-majority
4	$0 < p < 1$	DIV > GHQC > cohorts > 4-majority

Table 1: A comparison of the availability of the generalized hierarchical quorum, k-majority, cohorts, and DIV strategies with $NS = 108$

The third criteria in Table 2 is whether the strategy is a fully distributed one. These four strategies are fully distributed ones. The last two criteria are the number of failed nodes which do not halt the system and such that at most k nodes can simultaneously access their critical section in the best and worst cases. While in the best case, all these strategies can be fault-tolerant up to all node failures except those nodes which have already been constructed in k quorums. In the best case, the cohorts strategy can be fault-tolerant up to $(n - ks + \frac{k(k-1)}{2})$ node failures when $Coh(k,l) = (C_1, C_2, ..., C_l)$, $|C_1| = k$, and $|C_i| = s$, $i > 1$ [4]. Note that, in the best case, $|Q_1| = s - (k-1)$, $|Q_2| = s - (k-2)$, ..., $|Q_{k-1}| = s - 1$, and $|Q_k| = s$ in the cohorts strategy. While in the worst case, the generalized hierarchical quorum strategy can be fault-tolerant up to $k((\frac{n}{k})^{0.63}) - 1)$ node failures and the cohorts strategy can be fault-tolerant up to $(s - k + 1)$ node failures. While in the worst case, the k-majority strategy can not be fault-tolerant to any node failure and the DIV strategy can be fault-tolerant up to $(n - k\lceil\frac{n+k}{2k}\rceil)$ node failures.

Figure 4 shows a comparison of the quorum size of these four strategies for 4-mutual exclusion when no node failure occurs. From this figure, we observe that the quorum size of these four strategies in a decreasing order is 4-majority > cohorts > DIV > generalized hierarchical quorum, when $NS > 15$. That is, the quorum size of the generalized hierarchical is always the smallest one among these four strategies, when $NS > 15$. Moreover, Figure 5 shows a comparison of the quorum size of these four strategies for 4-mutual exclusion when node failures occur in the worst case with $NS = 124$ ($l = 4$). From this figure, we observe that the quorum size of the generalized hierarchical quorum is always the smallest one among these four strategies, when $NS > 17$.

	GHQC	Cohorts*
(1)	$(\frac{NS}{k})^{0.63}$	2 or k**
(2)	$(\frac{NS}{k})^{0.63}$	NS
(3)	yes	yes
(4)	$NS - k(\frac{NS}{k})^{0.63}$	$NS - ks + \frac{k(k-1)}{2}$
(5)	$NS - k(\frac{NS}{k})^{0.63}$	$s - k + 1$

	K-majority	DIV
(1)	$\lceil \frac{NS+1}{k+1} \rceil$	$\lceil \frac{NS+k}{2k} \rceil$
(2)	$\lceil \frac{NS+1}{k+1} \rceil$	$\lceil \frac{NS+k}{2k} \rceil$
(3)	yes	yes
(4)	$NS - k\lceil \frac{NS+1}{k+1} \rceil$	$NS - k\lceil \frac{NS+k}{2k} \rceil$
(5)	0	$NS - k\lceil \frac{NS+k}{2k} \rceil$

* $Coh(k, l) = (C_1, C_2, \ldots, C_l), |C_1| = k, |C_i| = s, i > 1$.
** 2 when $k = 1$, or k when $k > 1$.

Table 2: A comparison of four k-mutual exclusion strategies in terms of quorum size

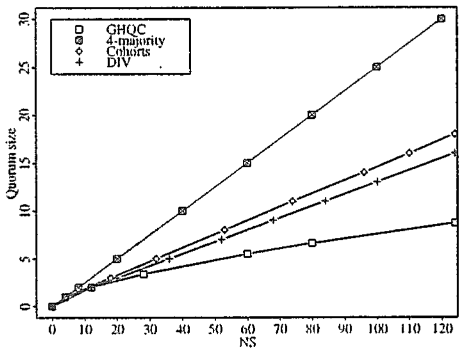

Figure 4: A comparison of the quorum size of the generalized hierarchical quorum, cohorts, k-majority, and DIV strategies for 4-mutual exclusion when no node failure occurs

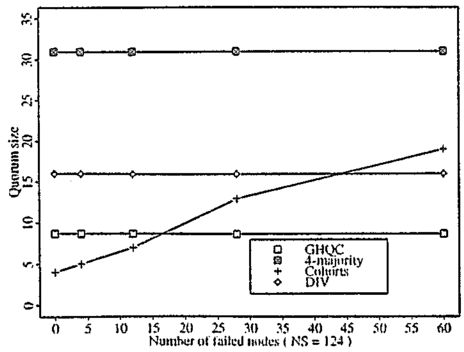

Figure 5: A comparison of the quorum size of the generalized hierarchical quorum, cohorts, k-majority, and DIV strategies for 4-mutual exclusion when node failures occur in the worst case which $NS = 124$ ($l = 4$)

4 Conclusion

In this paper, we have proposed a strategy called generalized hierarchical quorum for k-mutual exclusion, which imposes a logical generalized hierarchy structure on the network. In general, the generalized hierarchical quorum, a quorum contains the hierarchical quorum for 1-mutual exclusion for any one of subtree S_1, S_2, ..., S_k. The quorum size constructed from the strategy is always $(\frac{NS}{k})^{0.63}$, where NS is the number of physical nodes in the system. Moreover, this strategy can be always fault-tolerant up to $NS - k(\frac{NS}{k})^{0.63}$ node failures. From our performance analysis, we show that the generalized hierarchical quorum strategy can provide a higher availability than k-majority, cohorts, and DIV strategies, sometimes. Moreover, the quorum size of this strategy is always the smallest one among these four strategies, when no node failure occurs and $NS > 15$. While in the worst case, the quorum size of this strategy is always smallest one among these four strategies, when $NS > 17$. How to extend the generalized hierarchical quorum strategy to tolerate even more node failures is the future research direction.

References

[1] D. Agrawal and A. EL Abbadi, "The Generalized Tree Quorum Protocol: An Efficient Approach for Managing Replicated Data," *ACM Trans. on Database Systems*, Vol. 17, No. 4, pp. 689-717, Dec. 1992.

[2] D. Agrawal, O. Egecioglu, and A. E. Abbadi, "Analysis of Quorum-Based Protocols for Distributed (k+1)-Exclusion", *IEEE Trans. on Computers*, Vol. 8, No. 5, pp. 533-537, May 1997.

[3] Y. I. Chang, "A Simulation Study on Distributed Mutual Exclusion," *Journal of Parallel and Distributed Computing*, Vol. 33, No. 2, pp. 107-121, March 1996.

[4] J. R. Jiang, S. T. Huang, and Y. C. Kuo, "Cohorts Structures for Fault-Tolerant K Entries to a Critical Section," *IEEE Trans. on Computers*, Vol. 46, No. 2, pp. 222-228, Feb. 1997.

[5] H. Kakugawa, S. Fujita, M. Yamashita, and T. Ae, "Availability of K-Coterie," *IEEE Trans. on Computers*, Vol. 42, No. 5, May 1993.

[6] A. Kumar, "Hierarchical Quorum Consensus: A New Algorithm for Managing Replicated Data," *IEEE Trans. on Computers*, Vol. 40, No. 9, pp. 996-1004, Sep. 1991.

[7] M. L. Neilsen and M. Mizuno, "Nondominated k-coteries for multiple mutual exclusion," *Information Processing Letters*, Vol. 50, pp. 247-255, 1994.