Probabilistic Biquorums

Andreas Klappenecker and Hyunyoung Lee

Texas A&M University, USA

Example

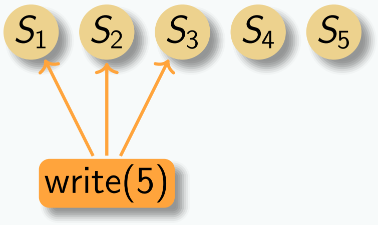

A **quorum system** is a fundamental primitive in distributed computing that allows one to implement distributed storage systems, solve consensus, realize atomic registers, location services, and more.

Example

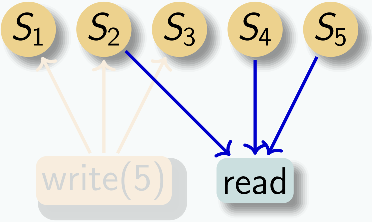

A **quorum system** is a fundamental primitive in distributed computing that allows one to implement distributed storage systems, solve consensus, realize atomic registers, location services, and more.

Strict Biquorum Systems

A **strict biquorum system** $(\mathcal{R}, \mathcal{W})$ over a set S is given by a family $\mathcal{R}$ of **read quorums** and a family $\mathcal{W}$ of **write quorums** such that

$$R \cap W \neq \emptyset$$

for each $R \in \mathcal{R}$ and for each $W \in \mathcal{W}$.

Generalization: quorums can be multisets with elements of S, intersection just with high probability.

Access Strategy

S a set of servers. Quorums are now multisets. Read and write quorums may fail to intersect.

The access strategy of read quorums

$$A_r(R) = \Pr[R \in \mathcal{R}].$$

The access strategy of write quorums

$$A_w(W) = \Pr[W \in \mathcal{W}].$$

Separable Access Strategy

Read: Choose r servers from S with rep. allowed
Write: Choose w servers from S with rep. allowed

A **δ-intersecting probabilistic biquorum system** $(\mathcal{R}, \mathcal{W})$ over a set S is given by

- a family $\mathcal{R}$ of read quorums and
- a family $\mathcal{W}$ of write quorums

such that

$$\sum_{\substack{R \in \mathcal{R}, W \in \mathcal{W}, \\ R \cap W \neq \emptyset}} \Pr[R \in \mathcal{R}] \Pr[W \in \mathcal{W}] \geqslant 1 - \delta$$

The probability to select disjoint read/write quorum pairs is δ or less.

Probabilistic Biquorums

A **probabilistic (r, w)-biquorum** is a probabilistic biquorum such that

- each write quorum contains w **servers** that are independently selected with repetitions allowed from a set S of servers according to a probability distribution p,
- each read quorum contains r **servers** that are independently selected with repetions allowed from the set S of servers according to a probability distribution q.

What is the probability that read and write quorums in a probabilistic (r, w)-biquorum system share a common server?

Theorem

Suppose that a probabilistic (r, w)-biquorum is selected from a set of n servers. This is a δ-intersecting probabilistic biquorum with

$$\delta = \sum \binom{w}{w_1, w_2, \ldots, w_n} \binom{r}{r_1, r_2, \ldots, r_n} \prod_{k=1}^{n} p_k^{w_k} q_k^{r_k},$$

where the sum extends over all vectors $(w_1, \ldots, w_n, r_1, \ldots, r_n)$ of nonnegative integers such that

(a) $w_1 + \cdots + w_n = w$,

(b) $r_1 + \cdots + r_n = r$,

(c) $w_k r_k = 0$ *for all k in the range* $1 \leqslant k \leqslant n$.

Special Case of Uniform Distributions

Proposition

If the quorums are selected uniformly at random, that is, if $p_k = q_k = 1/n$ holds for all k, then a probabilistic (r, w)-biquorum is δ-intersecting with

$$\delta = \sum_{k=2}^{n} \frac{n!}{(n-k)!n^{r+w}} \sum_{x+y=k} \left\{ {w \atop x} \right\} \left\{ {r \atop y} \right\},$$

where $\left\{ {a \atop b} \right\}$ is the Stirling number of the second kind, that is, the number of ways to partition a set of a elements into b nonempty subsets.

Asymptotic Result

Proposition

Let δ be a real number in the range $0 < \delta < 1$ and α a real number in the range $0 \leqslant \alpha \leqslant 1/2$. Suppose that the servers for the read and write quorums are chosen uniformly at random with repetitions allowed.

For large n, any probabilistic (r, w)-biquorum on n servers with quorum sizes

$$r = (n \log(1/\delta))^{\alpha} \quad \textit{and} \quad w = (n \log(1/\delta))^{1-\alpha}$$

is a δ-intersecting probabilistic biquorum.

Application

Example

Suppose that we have $n = 365$ servers.

r	w	δ
16	16	$1/2$
41	41	$1/100$
28	69	$1/100$

Expected Number of Shared Servers

Proposition

Consider a probabilistic (r, w)-biquorum. Let X_k be the indicator random variable of the event that the read and write quorums both contain the server k. Let $X = \sum_{i=1}^{n} X_i$ be the random variable denoting the number of shared servers. Then the expected number of shared servers is

$$\mathrm{E}[X] = \sum_{i=1}^{n} \left(1 - (1 - p_i)^w\right)\left(1 - (1 - q_i)^r\right).$$

Uniform Case

If one of the probability distributions is uniform, then the expected number of servers is at most as large as in the case of uniform distribution, as our next proposition shows.

Proposition

Let us keep the notation of the previous proposition. Suppose that the access probability distribution for read quorums is uniform, but the probability distribution of write quorums is arbitrary. Then the expected number of shared servers satisfies

$$\mathrm{E}[X] \leqslant n(1-(1-1/n)^{w})(1-(1-1/n)^{r}).$$

Conclusions

We derived the exact probability that read and write quorums intersect in a probabilistic (r, w)-biquorum when a separable access strategy is used.

We showed that allowing repetitions in the quorum selection still permits one to use quorums of small size that are the hallmark of probabilistic quorum and biquorum systems that do not allow repetitions.

We determined the expected number of servers that can be found in the intersection of a read quorum and a write quorum.