Randomized Load Balancing by Joining and Splitting Bins

James Aspnes[*†]   Yitong Yin[‡ §]

1 Introduction

Consider the following load balancing scenario: a certain amount of work load is distributed among a set of machines that may change over time as machines join and leave the system. Upon an arrival of a new machine, one of the existing machines gives some of its load to the new machine; and upon a departure of a machine, it gives all its load away to one of the existing machines in the system. Such load balancing schemes can be modeled as a simple game of joining and splitting weighted bins. Each bin corresponds to a machine in the system, and the weight of the bin represents the amount of load assigned to the machine. The arrival of a new machine corresponds to a split of a bin, and the departure of an existing machine is represented by joining two bins.

We consider what happens when the joins and splits are randomized. When the bins are split with probability proportional to their weights, it is not hard to see that this gives the same behavior as uniformly cutting a ring as in [3], which yields an $O(\log n)$ load factor for n bins. Where it is infeasible to bias the random choice with the weights, it is natural to implement uniform splits, where the split bin is sampled uniformly. This is a natural choice, for example, in a peer-to-peer system where a uniform sampling mechanism is available (e.g., [1, 4]). Despite its simple definition, analyzing the performance of this natural load-distribution mechanism is a nontrivial task.

In this paper, we apply a novel technique based on vector norms to analyze the load balancing performance of uniform random joins and splits. We show that if only splits (with no joins)

[*]Department of Computer Science, Yale University.

[†]Supported in part by NSF grants CNS-0435201 and CCF-0916389. Email: aspnes@cs.yale.edu.

[‡]State Key Laboratory for Novel Software Technology, Nanjing University, China.

[§]Supported by the National Science Foundation of China under Grant No. 61003023 and No. 61021062. Email: yinyt@nju.edu.

are applied, the expected **load factor**, the ratio between the maximum weight and the average weight of the bins, is between $\Omega\left(n^{0.5}\right)$ and $O\left(n^{0.742}\right)$. We then study the performance of mixed joins and splits, and show that the expected load factor is $O\left(n^{1/\sqrt{\frac{1}{2}\log_2 n}}\right)$ after alternatively applying sufficiently many joins and splits to an arbitrary initial load assignment of n bins. These results demonstrate that the good load factor obtained by [3] depends strongly on the ability to preferentially split heavily-loaded bins.

2 Load balance in the split-only process

In this section, we analyze the performance of a split-only process.

If we let N_t denote the number of bins at time t, the procedure can be described as follows. Initially, $N_0 = 1$: there is only one bin, whose weight is 1, and in the absence of joins we will always have $N_t = t+1$. Let $X_i^{(t)}$ denote the weight of bin i after t splits, for each $i \in [N_t] = \{0 \dots N_t - 1\}$. The weights of the N_t bins after t splits are inductively defined as follows: Initially, t=0 and $X_0^{(0)} = 1$. For later times t, uniformly choose a bin r from $[N_{t-1}]$. Let $X_i^{(t)} = X_i^{(t-1)}$ for $0 \leq i < r$; let $X_i^{(t)} = \frac{1}{2}X_r^{(t-1)}$ for $r \leq i \leq r+1$; and let $X_i^{(t)} = X_{i-1}^{(t-1)}$ for $r+1 < i \leq N_{t-1}$. In other words, each split chooses a bin uniformly at random and splits it into two equally-weighted bins.

2.1 Lower bound

Due to the definition of the splitting process, every bin in the system has a weight of the form $1/2^i$, where i is some integer. There are one or more **heaviest** bins, whose weights are equal to the maximum weight. The maximum load decreases by half once the last heaviest bin is split, and that split directly creates two new heaviest bins (and may indirectly create more if the maximum load drops to the weight already present in some bins). The maximum load will not decrease further until both of these two bins have been split, and it will not decrease again until the last of the second bin's children have been split, and so on. Intuitively, by concentrating on the rightmost path in this tree, we can obtain a lower bound on the size of the actual largest bin while keeping track of at most two bins at a time. This is formalized in the following theorem.

Theorem 1 *Let $X_0^{(n)}, X_1^{(n)}, \ldots, X_n^{(n)}$ be the weights of the $(n+1)$ bins after n splits. It holds that* $\mathrm{E}\left[\max_{0\le i\le n} X_i^{(n)}\right] \ge \frac{1}{\sqrt{2n}}$.

Proof: Consider the following procedure. Initially, the only bin $X_0^{(0)}$ is marked as special. When a special bin is split, if it is the only special bin, then the two new bins created by the split are special, and if otherwise the two new bins are non-special. After any number of splits, there is either one special bin, or two special bins with the same weight. Let $W^{(n)}$ be the weight of any special bin after n splits. Note that the special bin is not necessarily the heaviest bin, but its weight is certainly a lower bound of the maximum load, i.e. $\max_i X_i^{(n)} \ge W^{(n)}$. We now proceed to obtain a lower bound on $\mathrm{E}\left[W^{(n)}\right]$.

Let $w(n) = \mathrm{E}\left[W^{(n)}\right]$. For $b \in \{1,2\}$, let $w_b(n)$ be the part of $w(n)$ contributed by the cases that there are b special bin(s). Formally, denoting by $B^{(n)}$ the set of special bins after n splits, it holds that $|B^{(n)}| \in \{1,2\}$ and $w_b(n) = \sum_v v \cdot \Pr[W^{(n)} = v \wedge |B^{(n)}| = b]$. By total probability, $w(n) = w_1(n) + w_2(n)$. Let r denote the n^{th} split bin, which is chosen uniformly from $[n]$. If there is only one special bin, then either there was previously only one special bin and it was not split, or there were previously two special bins and one of them was split, therefore

$$
\begin{aligned}
w_1(n) &= \sum_v v \cdot \Pr\left[\begin{array}{l} \left(W^{(n-1)} = v \wedge \left|B^{(n-1)}\right| = 1 \wedge r \notin B^{(n-1)}\right) \\ \vee \left(W^{(n-1)} = v \wedge \left|B^{(n-1)}\right| = 2 \wedge r \in B^{(n-1)}\right) \end{array}\right] \\
&= \sum_v v\left(1 - \frac{1}{n}\right) \Pr\left[W^{(n-1)} = v \wedge \left|B^{(n-1)}\right| = 1\right] + \sum_v v \cdot \frac{2}{n} \Pr\left[W^{(n-1)} = v \wedge \left|B^{(n-1)}\right| = 2\right] \\
&= \left(1 - \frac{1}{n}\right) w_1(n-1) + \frac{2}{n} w_2(n-1).
\end{aligned}
$$

Similarly, for the case that there are two special bins,

$$
\begin{aligned}
w_2(n) &= \sum_v v \cdot \Pr\left[\begin{array}{l} \left(W^{(n-1)} = 2v \wedge \left|B^{(n-1)}\right| = 1 \wedge r \in B^{(n-1)}\right) \\ \vee \left(W^{(n-1)} = v \wedge \left|B^{(n-1)}\right| = 2 \wedge r \notin B^{(n-1)}\right) \end{array}\right] \\
&= \sum_v \frac{v}{2} \cdot \frac{1}{n} \Pr[W^{(n-1)} = v \wedge |B^{(n-1)}| = 1] + \sum_v v\left(1 - \frac{2}{n}\right) \Pr[W^{(n-1)} = v \wedge |B^{(n-1)}| = 2] \\
&= \frac{1}{2n} w_1(n-1) + \left(1 - \frac{2}{n}\right) w_2(n-1).
\end{aligned}
$$

Recall that $w(n) = w_1(n) + w_2(n)$, thus we have the following recursion:

$$w(n) = w_1(n) + w_2(n) = \left(1 - \frac{1}{2n}\right) w_1(n-1) + w_2(n-1) \geq \left(1 - \frac{1}{2n}\right) w(n-1).$$

Solving this recursive inequality with the obvious initial condition that $w(0) = 1$, we have that $w(n) \geq \frac{1}{2}\prod_{k=2}^{n}\left(1 - \frac{1}{2k}\right) \geq \frac{1}{\sqrt{2n}}$. Therefore $\mathrm{E}\left[\max_i X_i^{(n)}\right] \geq w(n) \geq \frac{1}{\sqrt{2n}}$. ∎

2.2 Upper bound

Let X be the vector denoting the weights of the bins after some number of splits. Let $|X|_p$ be the ℓ_p-norm of X. The maximum weight can be therefore represented as $|X|_\infty$. Note that $|X|_\infty \leq |X|_p$ for any $p \geq 1$. This means that we can get an upper bound on the maximum weight by bounding the ℓ_p norm for any such p. The following theorem is developed using this idea.

Theorem 2 *Let $X_0^{(n)}, X_1^{(n)}, \ldots, X_n^{(n)}$ be the weights of the $(n+1)$ bins after n splits. For any $\alpha \geq 1$, it holds that* $\mathrm{E}\left[\max_{0\leq i\leq n} X_i^{(n)}\right] = O\left(n^{-\left(1-2^{(1-\alpha)}\right)/\alpha}\right)$.

In order to bound the norm of the load vector, we first bound another quantity. Let $w(n) = \mathrm{E}\left[\sum_{i=0}^{n}(X_i^{(n)})^\alpha\right]$, where $\alpha \geq 1$ is a parameter. The following lemma gives a recursion for $w(n)$.

Lemma 3 *It holds that* $w(n) = \left(1 - \frac{1}{n}\left(1 - 2^{(1-\alpha)}\right)\right) \cdot w(n-1)$.

Proof: Suppose that r is the bin to split, which is uniformly sampled from $[n]$. From the inductive definition of $X^{(t)}$ from $X^{(t-1)}$, it holds that

$$\sum_{i=0}^{n}\left(X_i^{(n)}\right)^\alpha = \sum_{i\neq r}\left(X_i^{(n-1)}\right)^\alpha + 2\left(\frac{X_r^{(n-1)}}{2}\right)^\alpha = \sum_{i=0}^{n-1}\left(X_i^{(n-1)}\right)^\alpha + \left(2^{(1-\alpha)} - 1\right)\left(X_r^{(n-1)}\right)^\alpha.$$

By total probability,

$$w(n) = \sum_{i=0}^{n-1}\frac{1}{n}\cdot \mathrm{E}\left[\sum_{j=0}^{n}\left(X_j^{(n)}\right)^\alpha \,\middle|\, r = i\right] = \sum_{i=0}^{n-1}\frac{1}{n}\cdot \mathrm{E}\left[\sum_{j=0}^{n-1}\left(X_j^{(n-1)}\right)^\alpha + \left(2^{(1-\alpha)} - 1\right)\left(X_i^{(n-1)}\right)^\alpha\right],$$

which by linearity of expectation, implies that

$$w(n) = \mathrm{E}\left[\sum_{j=0}^{n-1}\left(X_j^{(n-1)}\right)^\alpha\right] + \frac{1}{n}\left(2^{(1-\alpha)}-1\right)\sum_{i=0}^{n-1}\mathrm{E}\left[\left(X_i^{(n-1)}\right)^\alpha\right] = \left(1-\frac{1}{n}\left(1-2^{(1-\alpha)}\right)\right)w(n-1).$$

∎

We now proceed to prove Theorem 2

Proof of Theorem 2: Let $w(n) = \mathrm{E}\left[\sum_{i=0}^{n}(X_i^{(n)})^\alpha\right]$, where $\alpha \geq 1$. According to Lemma 3, it holds that $w(n) = \left(1-\frac{1}{n}\left(1-2^{(1-\alpha)}\right)\right)w(n-1)$. Obviously $w(0)=1$. Solving the recursion,

$$\begin{aligned} w(n) &= \prod_{k=1}^{n}\left(1-\frac{1}{k}\left(1-2^{(1-\alpha)}\right)\right) \\ &= \exp\left(\sum_{k=1}^{n}\ln\left(1-\frac{1}{k}\left(1-2^{(1-\alpha)}\right)\right)\right) \\ &= \exp\left(O(1)-\left(1-2^{(1-\alpha)}\right)\sum_{k=1}^{n}\frac{1}{k}\right) \\ &= \exp\left(O(1)-\left(1-2^{(1-\alpha)}\right)\ln n\right) \\ &= O(n^{-\left(1-2^{(1-\alpha)}\right)}). \end{aligned}$$

For $\alpha \geq 1$, the function $f(x) = x^{1/\alpha}$ is concave. According to Jensen's inequality,

$$\begin{aligned} \mathrm{E}\left[\max_{0\leq i\leq n} X_i^{(n)}\right] &\leq \mathrm{E}\left[\left(\sum_{i=0}^{n}\left(X_i^{(n)}\right)^\alpha\right)^{1/\alpha}\right] \\ &\leq \left(\mathrm{E}\left[\sum_{i=0}^{n}\left(X_i^{(n)}\right)^\alpha\right]\right)^{1/\alpha} \\ &= O\left(n^{-\left(1-2^{(1-\alpha)}\right)/\alpha}\right). \end{aligned}$$

∎

We minimize the above bound numerically using standard methods. By setting $\alpha = 2.4$, the expected maximum load is $O(n^{-0.258})$ and the corresponding expected load factor is $O(n^{0.742})$.

3 Alternation of joins and splits

It is a bit surprising to see that joins can actually optimize the load factor. In this section we study a natural case mixed joins and splits: alternately applying join and split operations to initial population of $n-1$ bins. We will see that this sequence yields an $n^{o(1)}$ expected load factor, which is much better than the split-only case.

We assume that initially there are $(n-1)$ bins $\{X_i^{(0)}\}_{i\in[n-1]}$ with an arbitrary load distribution. At each time t, where the bins are $\{X_i^{(t)}\}_{i\in[n-1]}$, a bin chosen uniformly at random is split to produce weights $\{Y_i^{(t+1)}\}_{i\in[n]}$; and then a pair of bins chosen uniformly are joined to give new weights $\{X_i^{(t+1)}\}_{i\in[n-1]}$. The transitions are formally defined as follows.

- From $\{X_i^{(t)}\}_{i\in[n-1]}$ to $\{Y_i^{(t+1)}\}_{i\in[n]}$: Exactly the same as the split procedure defined in Section 2, except that here we write $Y_i^{(t+1)}$ instead of the original notation $X_i^{(t+1)}$.
- From $\{Y_i^{(t+1)}\}_{i\in[n]}$ to $\{X_i^{(t+1)}\}_{i\in[n-1]}$: Uniformly choose two bins $\{r,s\}$ from $\binom{[n]}{2}$, where $r<s$. Let $X_i^{(t+1)} = Y_r^{(t+1)} + Y_s^{(t+1)}$ if $i=r$; let $X_i^{(t+1)} = Y_{i+1}^{(t+1)}$ if $s \le i \le n-2$; and let $X_i^{(t+1)} = Y_i^{(t+1)}$ if otherwise.

We prove an upper bound on the maximum load after applying sufficiently many alternating joins and splits. The theorem is proved with a norm-based technique which is similar to the one used in the proof of Theorem 2.

Theorem 4 *The expected load factor following a sequence of alternating joins and splits is* $O(\exp(\sqrt{2\ln 2\ln n})) = O\left(n^{1/\sqrt{\frac{1}{2}\log_2 n}}\right)$ *in the limit. That is, for sufficiently large* n*, with any choice of (non-negative)* $X^{(0)}$ *with* $|X^{(0)}|_1 = 1$*, it holds that*

$$\lim_{t\to\infty} \mathrm{E}\left[\max_i X_i^{(t)}\right] = O\left(\frac{\exp\left(\sqrt{2\ln 2\ln n}\right)}{n}\right).$$

Proof: First, we need the following technical lemma.

Lemma 5 *Let* $x \in [0,1]^n$ *be a vector such that* $\sum_{i=1}^n x_i = 1$*. For any integers* k, l *such that*

$k > l > 0$, *it holds that*

$$\sum_{i=1}^{n} x_i^{k-l} \cdot \sum_{j=1}^{n} x_j^{l} \;\le\; \sum_{i=1}^{n} x_i^{k-1}.$$

Proof: For any x_i, x_j, and any integers $k, \ell > 0$, it holds that

$$\left(x_i^{k-l}x_j^l + x_i^l x_j^{k-l}\right) - \left(x_i^{k-1}x_j + x_i x_j^{k-1}\right) = x_i x_j \left(x_j^{k-l-1} - x_i^{k-l-1}\right)\left(x_i^{l-1} - x_j^{l-1}\right) \le 0,$$

therefore $\left(x_i^{k-l}x_j^l + x_i^l x_j^{k-l}\right) \le \left(x_i^{k-1}x_j + x_i x_j^{k-1}\right)$. It holds that,

$$\begin{aligned}
\sum_{i=1}^{n} x_i^{k-l} \cdot \sum_{j=1}^{n} x_j^{l} &= \sum_{i=1}^{n} x_i^k + \sum_{i=1}^{n-1}\sum_{j=i+1}^{n} \left(x_i^{k-l}x_j^l + x_i^l x_j^{k-l}\right) \\
&\le \sum_{i=1}^{n} x_i^k + \sum_{i=1}^{n-1}\sum_{j=i+1}^{n} \left(x_i^{k-1}x_j + x_i x_j^{k-1}\right) \\
&= \sum_{i=1}^{n} x_i^{k-1} \cdot \sum_{j=1}^{n} x_j \\
&= \sum_{i=1}^{n} x_i^{k-1}
\end{aligned}$$

∎

We then proceed to prove the theorem. We write $X_i = X_i^{(t)}$ and $Y_i = Y_i^{(t)}$ if no ambiguity is introduced. Let $k > 1$ be an integer. Let $w_k^{(t)}$ and $u_k^{(t)}$ be the $(\ell_k)^k$-norm of $X^{(t)}$ and $Y^{(t)}$ respectively. Specifically,

$$w_k^{(t)} = \mathrm{E}\left[\sum_{i=0}^{n-2} X_i^k\right] \quad \text{and} \quad u_k^{(t)} = \mathrm{E}\left[\sum_{i=0}^{n-1} Y_i^k\right].$$

Note that $X^{(t)}$ is formed by joining two uniformly random bins in $Y^{(t)}$. Conditioning on that the two random bins are $\{r, s\} \in \binom{[n]}{2}$, it holds that

$$\sum_{i=0}^{n-2} X_i^k = (Y_r + Y_s)^k + \sum_{i \notin \{r,s\}} Y_i^k = \sum_{i=0}^{n-1} Y_i^k + \sum_{l=1}^{k-1} \binom{k}{l} Y_r^{k-l} Y_s^l. \tag{1}$$

By total probability, $w_k^{(t)} = \sum_{\{r,s\}\in\binom{[n]}{2}} \frac{1}{\binom{n}{2}} \cdot \mathrm{E}\left[\sum_{i=0}^{n-2} X_i^k \,\middle|\, \{r,s\}\right]$. Due to (1),

$$\begin{aligned}
w_k^{(t)} &= \sum_{\{r,s\}\in\binom{[n]}{2}} \frac{1}{\binom{n}{2}} \cdot \mathrm{E}\left[\sum_{i=0}^{n-1} Y_i^k + \sum_{l=1}^{k-1} \binom{k}{l} Y_r^{k-l} Y_s^l\right] \\
&= \mathrm{E}\left[\sum_{i=0}^{n-1} Y_i^k\right] + \frac{\sum_{l=1}^{k-1}\binom{k}{l}}{2\binom{n}{2}} \cdot \mathrm{E}\left[\sum_{r=0}^{n-1}\sum_{s\neq r} Y_r^{k-l} Y_s^l\right] \\
&= u_k^{(t)} + \frac{\sum_{l=1}^{k-1}\binom{k}{l}}{2\binom{n}{2}} \cdot \left(\mathrm{E}\left[\sum_{r=0}^{n-1}\sum_{s=0}^{n-1} Y_r^{k-l} Y_s^l\right] - \mathrm{E}\left[\sum_{r=0}^{n-1} Y_r^k\right]\right) \\
&= \left(1 - \frac{2^{k-1}-1}{\binom{n}{2}}\right) u_k^{(t)} + \frac{\sum_{l=1}^{k-1}\binom{k}{l}}{2\binom{n}{2}} \cdot \mathrm{E}\left[\sum_{r=0}^{n-1} Y_r^{k-l} \cdot \sum_{s=0}^{n-1} Y_s^l\right] \\
\text{(due to Lemma 5)} &\leq \left(1 - \frac{2^{k-1}-1}{\binom{n}{2}}\right) u_k^{(t)} + \frac{\sum_{l=1}^{k-1}\binom{k}{l}}{2\binom{n}{2}} \cdot \mathrm{E}\left[\sum_{r=0}^{n-1} Y_r^{k-1}\right] \\
&= \left(1 - \frac{2^{k-1}-1}{\binom{n}{2}}\right) u_k^{(t)} + \frac{2^{k-1}-1}{\binom{n}{2}} u_{k-1}^{(t)}. \qquad (2)
\end{aligned}$$

Due to Lemma 3, it holds for the split operation that

$$u_k^{(t)} = \left(1 - \frac{1}{n}\left(1 - 2^{(1-k)}\right)\right) w_k^{(t-1)}. \qquad (3)$$

Combining (2) and (3), we obtain the recursion

$$\begin{aligned}
w_k^{(t)} &\leq \left(1 - \frac{2^{k-1}-1}{\binom{n}{2}}\right)\left(1 - \frac{1}{n}\left(1 - 2^{(1-k)}\right)\right) w_k^{(t-1)} \\
&\quad + \frac{2^{k-1}-1}{\binom{n}{2}}\left(1 - \frac{1}{n}\left(1 - 2^{(2-k)}\right)\right) w_{k-1}^{(t-1)} \\
&\leq \left(1 - \frac{1}{2n}\right) w_k^{(t-1)} + \frac{2^k}{n(n-1)} w_{k-1}^{(t-1)}.
\end{aligned}$$

Since we seek an upper bound, we can therefore assume that $w_k^{(t)} = \left(1 - \frac{1}{2n}\right) w_k^{(t-1)} + \frac{2^k}{n(n-1)} w_{k-1}^{(t-1)}$ without loss of generality. We show by induction that when $k = o(\log n)$, $w_k^{(t)} \leq \frac{2^{(k^2+3k-4)/2}}{(n-1)^{k-1}} + \frac{(2n+k)^k}{2^n}$ for all sufficiently large t. It is trivial to see this is true for $k=1$ as $w_1^{(t)} = 1$ for any t. Suppose the hypothesis is true for $k-1$. We can rewrite the recursion as $w_k^{(t)} = \left(1 - \frac{1}{2n}\right) w_k^{(t-1)} + \frac{2^{(k^2+3k-6)/2}}{n(n-1)^{k-1}} + \frac{(2n+k-1)^{k-1}}{2^n}$ for sufficiently large t. Let w_k be the fixed point that $w_k = \left(1 - \frac{1}{2n}\right) w_k + \frac{2^{(k^2+3k-6)/2}}{n(n-1)^{k-1}} +$

$\frac{(2n+k-1)^{k-1}}{2^n}$. Then $\frac{|w_k^{(t)}-w_k|}{|w_k^{(t-1)}-w_k|} = \left(1-\frac{1}{2n}\right)$, thus $w_k^{(t)} \leq w_k + 2^{-n}$ in finite steps of t. Therefore for all sufficiently large t, $w_k^{(t)} \leq w_k + 2^{-n} \leq \frac{2^{(k^2+3k-4)/2}}{(n-1)^{k-1}} + \frac{(2n+k)^k}{2^n}$.

According to Jensen's inequality [2], for $k = o(\log n)$,

$$\lim_{t\to\infty} \mathrm{E}\left[\max_i X_i^{(t)}\right] \leq \lim_{t\to\infty}\left(\mathrm{E}\left[\sum_{i=0}^{n-1} X_i^k\right]\right)^{1/k} \leq (w_k)^{1/k} = O\left(2^{\frac{k}{2}} n^{-(1-\frac{1}{k})}\right).$$

By setting $k = \lfloor\sqrt{2\log_2 n}\rfloor$, which is indeed within $o(\log n)$, we have that

$$\lim_{t\to\infty} \mathrm{E}\left[\max_i X_i^{(t)}\right] = O\left(\exp\left(\sqrt{2\ln 2\ln n}\right)\cdot n^{-1}\right) = O\left(n^{-1+1/\sqrt{\frac{1}{2}\log_2 n}}\right).$$

∎

4 Conclusion

We analyze the performance of a very natural randomized load balancing scheme: uniformly joining and splitting weighted bins. We develop a norm-based technique for analyzing this simple procedure. By applying the technique, we prove several bounds for the expected load factor. Specifically, if we keep uniformly splitting the bins without joining them, the expected load factor is between $\Omega\left(n^{0.5}\right)$ and $O\left(n^{0.742}\right)$, however, if we alternatively join and split bins, the expected load factor converges to $O\left(n^{1/\sqrt{\frac{1}{2}\log_2 n}}\right)$. These bounds justify the intuition that the power of being adaptive to the current loads is essential for load balancing tasks, and they also show a somehow surprising phenomenon that joins can actually help load balancing if such power is not available.

References

[1] S. Datta and H. Kargupta. Uniform data sampling from a peer-to-peer network. In *27th IEEE International Conference on Distributed Computing Systems (ICDCS 2007)*, page 50, 2007.

[2] G. Hardy, J. Littlewood, and G. Pólya. *Inequalities*. The University press, 1934.

[3] D. R. Karger, E. Lehman, F. T. Leighton, R. Panigrahy, M. S. Levine, and D. Lewin. Consistent hashing and random trees. In *STOC*, pages 654–663, 1997.

[4] D. Stutzbach, R. Rejaie, N. G. Duffield, S. Sen, and W. Willinger. On unbiased sampling for unstructured peer-to-peer networks. *IEEE/ACM Trans. Netw.*, 17(2):377–390, 2009.