The power of thinning in balanced allocation

Ohad N. Feldheim[*]   Ori Gurel-Gurevich[†]

July 3, 2018

Abstract

Balls are sequentially allocated into n bins as follows: for each ball, an independent, uniformly random bin is generated. An overseer may then choose to either allocate the ball to this bin, or else the ball is allocated to a new independent uniformly random bin. The goal of the overseer is to reduce the load of the most heavily loaded bin after $\Theta(n)$ balls have been allocated. We provide an asymptotically optimal strategy yielding a maximum load of $(1 + o(1))\sqrt{\frac{8 \log n}{\log \log n}}$ balls.

Keywords. Thinning, two-choices, one-retry, $(1+\alpha)$-choice, load balancing, balls and bins, balanced allocation, subsampling.

1   Introduction and Results

Fix $\rho > 0$ and consider a model in which an overseer is monitoring the sequential allocation of $\lfloor \rho n \rfloor$ balls into n bins. Each ball is assigned a *primary allocation*, i.e., an independent, uniformly chosen random bin. Then, the overseer is given the choice to reject this primary allocation, in which case the ball is assigned a *secondary allocation* instead, that is, a new, independent, uniformly chosen random bin. The set of all resulting allocations is called a *two-thinning* of the balls-and-bins process.

A *two-thinning strategy*, is a function determining whether to accept or reject each suggested allocation, depending on all previous allocations. Denote by $\text{MaxLoad}_t^f([n])$ the load of the most heavily loaded bin after the player allocates $\lfloor t \rfloor$ balls into n bins, following the strategy f. A strategy is *asymptotically optimal* if $\text{MaxLoad}_{\rho n}^f([n]) \leq (1 + o(1)) \text{MaxLoad}_{\rho n}^g([n])$, for any strategy g, with high probability.

Here we describe and analyse an optimal two-thinning strategy which we call the ℓ-*threshold* strategy. This is the two-thinning strategy which rejects a ball whenever the number of primary allocations to the suggested bin is at least ℓ. Our main result is the following,

Theorem 1. *Let f be the $\sqrt{\frac{2 \log n}{\log \log n}}$-threshold strategy for the allocation of $\lfloor \rho n \rfloor$ balls into n bins. Then f is asymptotically optimal and, with high probability,*

$$\text{MaxLoad}_{\rho n}^f([n]) = (1 + o(1))\sqrt{\frac{8 \log n}{\log \log n}}.$$

---

[*]Hebrew University of Jerusalem, email:`Ohad.Feldheim@mail.huji.ac.il`.

[†]Hebrew University of Jerusalem, email:`Ori.Gurel-Gurevich@mail.huji.ac.il`, research was supported by the Israel Science Foundation (grant No. 1707/16).

1.1 Discussion

Balls-and-bins, two-choices and two-thinning. It is well known that if each of $\lfloor \rho n \rfloor$ balls is allocated independently to a uniformly chosen random bin in $[n] = \{1, \ldots, n\}$, then the most heavily loaded bin contains $\frac{\log n}{\log \log n} + O(1)$ balls with high probability. In their seminal paper, Azar, Broder, Karlin and Upfal [1] have shown that a significantly lower maximum load of $\log_2 \log n + O(1)$ balls could be achieved, with high probability, in a *two-choices* setting, i.e., if the allocation of each ball is governed by an overseer who is offered a choice between two independent, uniformly chosen random bins. Moreover, the overseer can achieve this simply by following a naïve strategy of always selecting the less loaded of the two bins.

The two-thinning setting, considered in this paper, is intermediate between two-choices and no-choice, as it is equivalent to a two-choices setting in which the overseer is oblivious of the location of one of the two available bins. The name "two-thinning" is due to yet another point of view on this setting. According to this view an infinite sequence of allocations has been drawn independently and uniformly at random, and the overseer is allowed to thin it on-line (i.e., delete some of the allocations depending only on the past), as long as at most one of every two consecutive entries is deleted (for a more thorough discussion of the model see joint work with Ramdas and Dwivedi [4], where the model was introduced).

From Theorem 1 we see that the optimal maximum load under two-thinning is indeed intermediate between the maximum load without thinning and the maximum load in the two-choices setting.

More choice. Already in [1], Azar et al. showed that allowing the overseer choice between $k > 2$ choices, reduces the asymptotic maximal load by merely a factor of $\log(k)$. Nonetheless, we make the following conjecture.

Conjecture 1. *In the two-thinning setting, allowing the overseer to iteratively reject up to k suggested allocations for each ball will result in an improved asymptotically optimal maximum load of*

$$\Theta \left(\left(\frac{\log n}{\log \log n} \right)^{1/(k+1)} \right).$$

More balls. Berenbrink, Czumaj, Steger and Vöcking [2] have considered the power of two choices in the heavily loaded case of the balls and bins model, that is, when $\omega(n)$ balls are allocated into n bins. They showed that in this case under the power of k-choices, the deviation of the maximum load from the average load is asymptotically almost surely $\log_k \log n + O(1)$ (see Talwar and Wieder [9], for a simpler proof). We conjecture that the same phenomenon will occur for two-thinning. Namely,

Conjecture 2. *In the two-thinning setting, where $m = \Omega(n)$ balls are two-thinned, the asymptotically optimal maximum load is $\frac{m}{n} + \Theta \left(\sqrt{\frac{\log n}{\log \log n}} \right)$.*

1+β-thinning. In his thesis [5], Mitzenmacher suggested considering a variant of the power of two-choices in which, for each allocation independently, there is some small probability that a decision opposite to that made by the overseer will be executed. This notion was recently formulated and studied by Peres, Talwar and Wieder [7], viewing it as having two-choices with probability β and no-choice with probability $(1-\beta)$, independently for every ball. Once errors of this nature are introduced to the model, two-choices and one-retry are equivalent up to a parameter change, and in lightly loaded case of $\lfloor \rho n \rfloor$ balls allocated into n bins, both offer no improvement over having no-choice at all (see [4] for more details).

2   Preliminaries

We take advantage of a comparison lemma of Mitzenmacher and Upfal [6, Corollary 5.11], which we reproduce here, relating the balls-and-bins model with independent Poisson random variables. Denote by $\mathbb{N}_0$ the set of natural numbers together with 0. Given two vectors $x, y \in (\mathbb{N}_0)^n$ we write $x \leq y$ if $x_i \leq y_i$ for all $i \in [n]$. A set $S \subset (\mathbb{N}_0)^n$ is called *monotone decreasing (increasing)* if $x \in S$ implies $y \in S$ for all $y \leq x$ ($y \geq x$).

Lemma 2.1 (Mitzenmacher and Upfal). *Let $(X_m)_{m \in [n]}$ be the number of balls in the m-th bin when t balls are independently and uniformly allocated into n bins. Further let $(Y_m)_{m \in [n]}$ be independent Poisson($\frac{t}{n}$) random variables, and let S be a monotone set (either increasing or decreasing). Then*

$$\mathbb{P}\Big((X_1, \ldots, X_n) \in S\Big) \leq 2\mathbb{P}\Big((Y_1, \ldots, Y_n) \in S\Big).$$

We also utilise two corollaries of this lemma.

Lemma 2.2. *Let $(X_m)_{m \in [n]}$ be the number of balls in the m-th bin when (θn)-balls are independently and uniformly allocated into n-bins, for $\theta \in [0, 1]$. Then, for any $a \in [\theta n]$ and $S \subset [n]$ we have*

$$\mathbb{P}\left(\max_{m \in S}(X_m) < a\right) \leq 2\exp\left(-\frac{\theta^a |S|}{ea!}\right)$$

Proof. Let $\{Y_m\}_{m \in [n]}$ be i.i.d. Poisson(θ) random variables. By Lemma 2.1 we have

$$\mathbb{P}\left(\max_{m \in S}(X_m) < a\right) \leq 2\,\mathbb{P}\left(\max_{m \in S}(Y_m) < a\right) = 2\mathbb{P}\left(Y_1 < a\right)^{|S|} \leq 2\left(1 - e^{-\theta}\frac{\theta^a}{a!}\right)^{|S|}$$

$$\leq 2\left(1 - \frac{\theta^a}{ea!}\right)^{|S|} \leq 2\exp\left(-\frac{\theta^a |S|}{ea!}\right).$$

$\square$

Lemma 2.3. *Let $(X_m)_{m \in [n]}$ be the number of balls in the m-th bin when (θn)-balls are independently and uniformly allocated into n-bins, for $\theta \in [0, 1]$. Then, for any $S \subset [n]$ we have*

$$\mathbb{P}\left(|\{m \in S : X_m > 0\}| \leq \frac{\theta |S|}{2e}\right) \leq 2\exp\left(-\frac{\theta^2 |S|}{2e^2}\right).$$

Proof. Let $\{Y_m\}_{m\in[n]}$ be i.i.d. Poisson(θ) random variables. By Lemma 2.1

$$\mathbb{P}\left(|\{m \in S : X_m > 0\}| \leq \frac{\theta|S|}{2e}\right) \leq 2\mathbb{P}\left(|\{m \in S : Y_m > 0| \leq \frac{\theta|S|}{2e}\right).$$

We observe that $\mathbb{P}(Y_1 > 0) \geq \frac{\theta}{e}$. Using Hoeffding bound for the tail of binomial distributions (see, e.g. [3, Proposition 1.12]), we obtain,

$$\mathbb{P}\left(|\{m \in S : Y_m > 0\}| \leq \frac{|S|\theta}{2e}\right) \leq \exp\left(-2|S|\left(\frac{\theta}{e} - \frac{\theta}{2e}\right)^2\right) = \exp\left(-\frac{\theta^2|S|}{2e^2}\right).$$

$\square$

3 Notation

Given a thinning strategy f, generate $\{Z_t\}_{t\in\mathbb{N}_0}$, the sequence of allocations, in the following way. Let $\{Z_t^0\}_{t\in\mathbb{N}_0}$ and $\{Z_t^1\}_{t\in\mathbb{N}_0}$ be two sequences of independent random variables uniformly distributed in $[n]$. Here Z_t^0 represents the primary allocation of the t-th ball, while $\{Z_t^1\}_{t\in\mathbb{N}_0}$ is used as a pool of secondary allocations. Denote by r_t the number of rejections among the first $(t-1)$ primary allocations. For the t-th allocation, look at the history of the process up to time $(t-1)$ and at Z_t^0 and apply f to determine whether to accept or reject the primary allocation. If the primary allocation is accepted then set Z_t to be Z_t^0 while if it is rejected, then set Z_t to be $Z_{r_t}^1$.

We introduce the following notation. For any $t \leq \rho n$, $m \in [n]$ denote

$$F_t(m) = |\{1 \leq i \leq t \ : \ Z_i = m\}|,$$
$$A_t(m) = |\{1 \leq i \leq t \ : \ Z_i^0 = m\}|,$$
$$B_t(m) = |\{1 \leq i \leq t \ : \ Z_i^1 = m\}|.$$

In addition, for $\ell \in \mathbb{N}_0$ and for $S \subset [n]$ denote

$$\phi_t^\ell(S) = |\{m \in S \ : \ F_t(m) \geq \ell\}|$$
$$\alpha_t^\ell(S) = |\{m \in S \ : \ A_t(m) \geq \ell\}|$$
$$\beta_t^\ell(S) = |\{m \in S \ : \ B_t(m) \geq \ell\}|,$$

setting $\phi_t^\ell = \phi_t^\ell([n])$, $\alpha_t^\ell = \alpha_t^\ell([n])$, $\beta_t^\ell = \beta_t^\ell([n])$. Finally, denote

$$\mathrm{MaxLoad}_t^f(S) = \max_{m\in S} F_t(m)$$
$$\mathrm{MaxLoad}_t^f = \mathrm{MaxLoad}_t^f([n]).$$

4   Upper bound on $\mathrm{MaxLoad}^f_{\rho n}([n])$

For $n \geq 3$, denote $L = \left\lceil \sqrt{2\log n / \log\log n} \right\rceil$. Let f be the L-threshold strategy, i.e., the one for which $f_i = 1$ if and only if $A_i(Z^0_i) \geq L$. The main statement of this section is the following.

Proposition 4.1. *For any $n \geq n(\rho)$ sufficiently large and any $\eta > 0$ the strategy f satisfies*

$$\mathbb{P}\big(\mathrm{MaxLoad}_{\rho n} > (2+\eta)L \big) \leq 2n^{-\frac{\eta}{4} + \frac{2\log\log\log n}{\log\log n}} + 2e^{-\sqrt{n}}.$$

Let us begin by reducing the upper bound in theorem 1 to this proposition.

Proof of the upper bound in Theorem 1. We apply Proposition 4.1 with $\eta = \frac{9\log\log\log n}{\log\log n}$. Observe that $\eta = o(1)$ and by the proposition we have

$$\mathbb{P}\big(\mathrm{MaxLoad}_{\rho n} > (2+\eta)L \big) \leq \exp\left(-\frac{\log n}{4\log\log n} \right) + 2e^{-\sqrt{n}} = o(1).$$

$\square$

Proof of Proposition 4.1. Denote $r = r_{\lfloor \rho n \rfloor}$. Our strategy f guarantees that $A_{\rho n}(m) \leq L$ for all $m \in [n]$. Hence, under this strategy

$$\mathbb{P}\big(\mathrm{MaxLoad}_{\rho n} \geq (2+\eta)L \big) \leq \mathbb{P}\left(\max_{m\in[n]} B_r(m) \geq L + \eta L \right) = \mathbb{P}\big(\beta_r^{L+\eta L} > 0\big) . \tag{1}$$

Let $L' = \lceil L + \eta L \rceil$. Notice that if $\beta_r^{L'} > 0$ then for any $0 \leq k \leq \rho n$, either $r > k$ or $\beta_k^{L'} > 0$. Hence, for any $0 \leq k \leq \rho n$ we get

$$\mathbb{P}\big(\mathrm{MaxLoad}_{\rho n} \geq (2+\eta)L \big) \leq \mathbb{P}\big(r > k\big) + \mathbb{P}\big(\beta_k^{L'} > 0\big) . \tag{2}$$

We now bound the two probabilities on the right hand side.

To bound $\mathbb{P}(r > k)$, let $\{Y_m\}_{m\in[n]}$ be i.i.d. Poisson(ρ) random variables and write

$$Y := \sum_{m\in[n]} \max(Y_m - L, 0).$$

By Lemma 2.1 we have

$$\mathbb{P}(r > k) \leq 2\mathbb{P}\big(Y > k\big). \tag{3}$$

For a single Poisson(ρ) random variable we have for $n \geq 100$,

$$\mathbb{E}\big(e^{\max(Y_1 - L, 0)} \big) \leq 1 + e^{-\rho} \sum_{\ell=1}^{\infty} \frac{\rho^\ell e^\ell}{(L+\ell)!} \leq 1 + \frac{1}{L!} < \exp\left(\frac{1}{L!} \right) .$$

Hence, by Markov's inequality, for $k \geq \frac{2n}{L!}$ we have

$$\mathbb{P}\big(Y > k\big) = \mathbb{P}\big(e^Y > e^k\big) \leq \exp\left(\frac{n}{L!} - k \right) < \exp\left(-\frac{n}{L!} \right) < e^{-\sqrt{n}}. \tag{4}$$

5

putting together (3) and (4) we obtain

$$\mathbb{P}(r > k) < 2e^{-\sqrt{n}}. \tag{5}$$

Next we bound $\mathbb{P}\left(\beta_k^{L'} > 0\right)$. Let $\{Y_m\}_{m\in[n]}$ be i.i.d. Poisson(k/n) random variables. By Lemma 2.1 we have,

$$\mathbb{P}\left(\beta_k^{L'} > 0\right) \leq 2\mathbb{P}\left(\max_{m\in[n]}(Y_m) > L'\right).$$

For $k \leq \frac{3n}{L!}$ and $n \geq 100$ we have

$$\mathbb{P}\left(Y_1 > L'\right) = e^{-k/n} \sum_{\ell=L'+1}^{\infty} \frac{(k/n)^\ell}{\ell!} \leq \sum_{\ell=L'+1}^{\infty} \left(\frac{3}{L!}\right)^\ell \leq \left(\frac{3}{L!}\right)^{L'}.$$

Taking a union bound, we obtain

$$\mathbb{P}\left(\beta_k^{L'}(L') > 0\right) \leq 2n\left(\frac{3}{L!}\right)^{L'}. \tag{6}$$

By Stirling's approximation for all $\ell > 1$ we have $\ell! > 3\left(\frac{\ell}{e}\right)^\ell$. Hence,

$$\begin{aligned}
\mathbb{P}\left(\beta_k^{L'} > 0\right) &\leq 2n\left(\frac{3}{L!}\right)^{L'} \leq 2n\left(\frac{L}{e}\right)^{-LL'} \leq 2n\left(\frac{L}{e}\right)^{-(1+\eta)L^2} \\
&\leq 2\exp\left(\log n - (1+\eta)L^2\left(\log L - 1\right)\right) \\
&\leq 2\exp\left(\log n - (1+\eta)\frac{2\log n}{\log\log n}\left(\frac{1}{2}\log\log n - \frac{1}{2}\log\log\log n - 1\right)\right) \\
&\leq 2\exp\left(-\eta\log n + (1+\eta)\frac{2\log n}{\log\log n}\left(\frac{1}{2}\log\log\log n + 1\right)\right) \\
&\leq 2\exp\left(-\eta\log n + (1+\eta)\frac{2\log n\log\log\log n}{\log\log n}\right) \\
&\leq 2n^{-\frac{\eta}{4}+\frac{2\log\log\log n}{\log\log n}}, \tag{7}
\end{aligned}$$

for any $n \geq 100$. Putting (5) and (7) into (2), the proposition follows. $\qquad\square$

5   Lower bound on $\mathrm{MaxLoad}_{\rho n}^g([n])$ for any strategy g

Let $\ell = \ell(n) = \left\lfloor\sqrt{2\log n/\log\log n}\right\rfloor$. In this section we prove the following proposition, from which the lower bound in Theorem 1 is an immediate corollary.

Proposition 5.1. *Let $\varepsilon, \rho > 0$ and n sufficiently large (depending on ρ and ε). For any strategy g we have*

$$\mathbb{P}\left(\mathrm{MaxLoad}_{\rho n}^g < (2 - \varepsilon)\ell\right) \leq \exp\left(-n^{\varepsilon/5}\right).$$

To prove Proposition 5.1 we use the following lemma.

Lemma 5.2. *Let $\varepsilon, \rho > 0$ and n sufficiently large (depending on ρ and ε) and denote $\zeta = \rho/8e\ell$. For any $1 \le k \le 2\ell$, $t > \rho n/2\ell$ and $S \subset [n]$ such that $|S| \ge n\zeta^k$ and any strategy g, we have*

$$\mathbb{P}(E, F) \le \exp(-n^{\varepsilon/4}),$$

where $E = \{\phi_t^1(S) < n\zeta^{k+1}\}$ and $F = \{\mathrm{MaxLoad}_t^g(S) < (2-\varepsilon)\ell - k\}$.

Proof. Write $T = n\zeta^{k+1}$ and denote $E' = \{\alpha_t^1(S) < 2T\}$ and $F' = \left\{\beta_T^{(2-\varepsilon)\ell - k}(S) = 0\right\}$. By applying Lemma 2.3 with $\theta = \frac{\rho}{2\ell}$ and observing that $2T \le 2\zeta|S| \le \frac{\theta|S|}{2e}$, we obtain

$$\mathbb{P}\left(E'\right) \le 2\exp\left(-\frac{\theta^2|S|}{2e^2}\right) \le 2\exp\left(-\frac{8n\rho^{k+2}}{(8e\ell)^{k+2}}\right) \le 2\exp\left(-n^{1+o(1)}\right),$$

Where in the rightmost inequality we used the fact that $k < 2\ell$. By applying Lemma 2.2 with $a = (2-\varepsilon)\ell - k$ and $\theta = \zeta^{k+1}$,

$$\mathbb{P}(F') \le 2\exp\left(-\frac{\zeta^{(k+1)a}|S|}{ea!}\right) \le 2\exp\left(-\frac{\zeta^{(k+1)a}\zeta^k n}{ea^a}\right) \le 2\exp\left(-\frac{\zeta^{(k+1)(a+1)}n}{ea^a}\right).$$

Letting n be large enough, and observing that for such n we have $(a+1)(k+1) \le (1-\varepsilon/2)\ell^2$ we obtain

$$\mathbb{P}(F') \le 2\exp\left(-\frac{(\rho/8e\ell)^{(1-\varepsilon/2)\ell^2}n}{e(2\ell)^{2\ell}}\right) \le 2\exp\left(-\frac{n}{\ell^{(1-\varepsilon/3)\ell^2}}\right) \le 2\exp(-n^{\varepsilon/3}),$$

where the two rightmost inequalities use the fact that $\ell^{\ell^2} \ge n$, while c^{ℓ^2} and ℓ^ℓ are both sub-polynomial in n for any $c > 0$.

We claim that $E'^c \cap F'^c \subseteq E^c \cup F^c$. Indeed, we observe that $\{r_t \le T\} \cap E'^c \subset E^c$, while $\{r_t > T\} \cap F'^c \subset F^c$. Hence $E \cap F \subset E' \cup F'$. From our bounds on $\mathbb{P}(E')$ and $\mathbb{P}(F')$ the proposition follows. $\square$

Proof of Proposition 5.1. Fix $\varepsilon, \rho > 0$ and let g be a thinning strategy. We divide our process into $s = \lceil(2-\varepsilon)\ell\rceil$ stages each consisting of the allocation of $w = \left\lceil\frac{\rho n}{2\ell}\right\rceil$ balls so that the k-th stage process consists of $Z_{(k-1)w+1}, \ldots, Z_{kw}$. These are followed by a final stage in which the remaining balls are allocated.

Denote $S_k = \{m \in [n] : A_{kw}(m) \ge k\}$. For $\zeta = \rho/8e\ell$, we define $E_k = \{|S_k| < n\zeta^k\}$ and $F_k = \{\mathrm{MaxLoad}_{kw}^g < (2-\varepsilon)\ell\}$.

By applying Proposition 5.2 to the k-th stage process with $S = S_k$ we obtain that

$$\mathbb{P}(E_{k+1} \cap F_{k+1} \mid E_k^c) \le \exp(-n^{\varepsilon/4}).$$

The see this, observe that the size of S_{k+1} is at least the number of bins in S_k which were allocated at least one ball in the k-th stage process and that $\mathrm{MaxLoad}_{(k+1)w}^g$ is at least k plus the maximum number of balls that were allocated in the k-th stage process to a single bin in S_k.

7

Observe that $F_{k+1} \subseteq F_k$ we use the law of total probability to obtain

$$\mathbb{P}(E_{k+1} \cap F_{k+1}) = \mathbb{P}(E_{k+1} \cap F_{k+1} \cap E_k) + \mathbb{P}(E_{k+1} \cap F_{k+1} \cap E_k^c) \leq \mathbb{P}(E_k \cap F_k) + \mathbb{P}(E_{k+1} \cap F_{k+1} \mid E_k^c)$$

Since $E_0 \cap F_0 = \emptyset$, we may use induction to deduce that for sufficiently large n we have,

$$\mathbb{P}\left(E_s \cap F_s\right) \leq \sum_{k=1}^{s} \mathbb{P}(E_k \cap F_k \mid E_{k-1}^c) \leq s \exp(-n^{\varepsilon/4}) \leq \exp(-n^{\varepsilon/5}).$$

Since $\{\mathrm{MaxLoad}_{\rho n}^g < (2-\varepsilon)\ell\} \subset E_s \cap F_s$, this concludes the proof.

$\square$

References

[1] Y. Azar, A. Broder, A. Karlin and E. Upfal, *Balanced allocations*. SIAM Journal of Computing 29 no. 1 (1999), pp. 180-200.

[2] P. Berenbrink, A. Czumaj, A. Steger and B. Vöcking, *Balanced allocations: The heavily loaded case*. SIAM Journal on Computing 35 no. 6 (2006), pp. 1350-1385.

[3] R. M. Dudley, *Uniform central limit theorems*, Cambridge university press, 1999.

[4] R. Dwivedi, O. N. Feldheim, O. Gurel-Gurevich, and A. Ramdas, *The power of online thinning in reducing discrepancy*. Probability Theory and Related Fields, to appear. arXiv preprint arXiv:1608.02895.

[5] M. Mitzenmacher, *The Power of Two Choices in Randomized Load Balancing*, PhD thesis, University of California, Berkeley, CA, 1996.

[6] M. Mitzenmacher and E. Upfal. *Probability and computing: Randomized algorithms and probabilistic analysis*. Cambridge University Press, 2005.

[7] Y. Peres, K. Talwar and U. Wieder, *Graphical Balanced Allocations and the (1 + α)-Choices Process*. Random Structures & Algorithms 47, no. 4 (2015), pp. 760-775 157-163.

[8] A. W. Richa, M. Mitzenmacher and R. Sitarman, *The power of two random choices: A survey of techniques and results*, Combinatorial Optimization 9 (2001), pp. 255-304.

[9] K. Talwar and U. Wieder, *Balanced allocations: A simple proof for the heavily loaded case*. International Colloquium on Automata, Languages, and Programming. Springer Berlin Heidelberg, 2014.