Balanced Allocations in Batches: Simplified and Generalized

Dimitrios Los[1], Thomas Sauerwald[1]

[1]University of Cambridge, UK

Balanced allocations: Background

Balanced allocations setting

Allocate m tasks (balls) sequentially into n machines (bins).

Balanced allocations setting

Allocate m tasks (balls) sequentially into n machines (bins).
Goal: minimise the **maximum load** $\max_{i \in [n]} x_i^m$, where x^t is the load vector after ball t.

Balanced allocations setting

Allocate m tasks (balls) sequentially into n machines (bins).
Goal: minimise the **maximum load** $\max_{i \in [n]} x_i^m$, where x^t is the load vector after ball t.

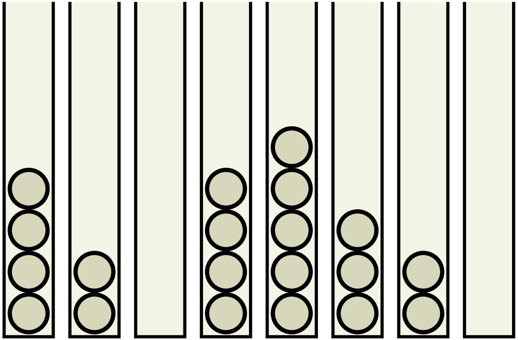

Balanced allocations setting

Allocate m tasks (balls) sequentially into n machines (bins).
Goal: minimise the **maximum load** $\max_{i \in [n]} x_i^m$, where x^t is the load vector after ball t.

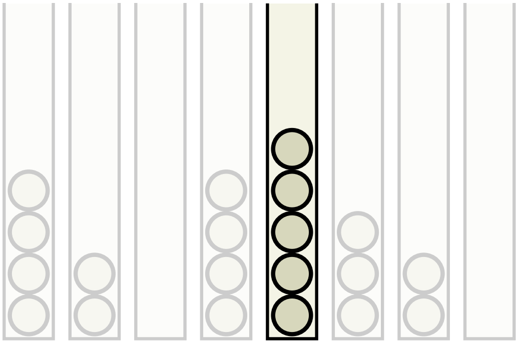

Balanced allocations setting

Allocate m tasks (balls) sequentially into n machines (bins).
Goal: minimise the **maximum load** $\max_{i \in [n]} x_i^m$, where x^t is the load vector after ball t.
$\Leftrightarrow$ minimise the **gap**, where $\mathrm{Gap}(m) = \max_{i \in [n]}(x_i^m - m/n)$.

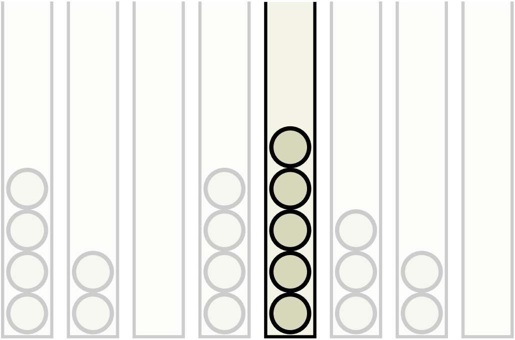

Balanced allocations setting

Allocate m tasks (balls) sequentially into n machines (bins).

Goal: minimise the **maximum load** $\max_{i \in [n]} x_i^m$, where x^t is the load vector after ball t.

$\Leftrightarrow$ minimise the **gap**, where $\text{Gap}(m) = \max_{i \in [n]}(x_i^m - m/n)$.

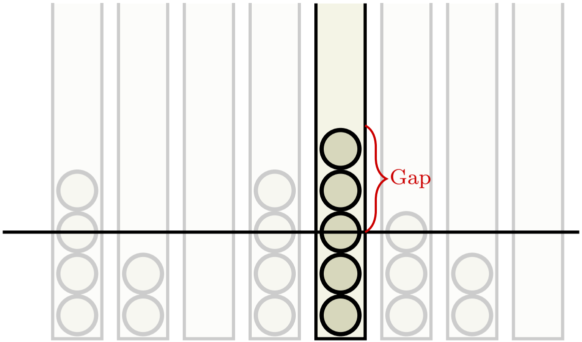

Balanced allocations setting

Allocate m tasks (balls) sequentially into n machines (bins).
Goal: minimise the **maximum load** $\max_{i \in [n]} x_i^m$, where x^t is the load vector after ball t.
$\Leftrightarrow$ minimise the **gap**, where $\text{Gap}(m) = \max_{i \in [n]}(x_i^m - m/n)$.

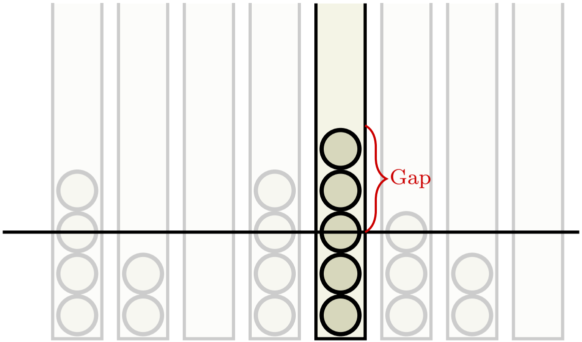

- Applications in hashing, load balancing and routing.

One-Choice and Two-Choice processes

<u>One-Choice Process:</u>
Iteration: For each $t \geq 0$, sample **one** bin uniformly at random (u.a.r.) and place the ball there.

ONE-CHOICE and TWO-CHOICE processes

<u>ONE-CHOICE PROCESS:</u>
Iteration: For each $t \geq 0$, sample **one** bin uniformly at random (u.a.r.) and place the ball there.

- In the lightly-loaded case ($m = n$), w.h.p. $\text{Gap}(n) = \Theta\left(\frac{\log n}{\log \log n}\right)$ [Gon81].

One-Choice and Two-Choice processes

<u>One-Choice Process:</u>
Iteration: For each $t \geq 0$, sample **one** bin uniformly at random (u.a.r.) and place the ball there.

- In the lightly-loaded case ($m = n$), w.h.p. $\text{Gap}(n) = \Theta\left(\frac{\log n}{\log \log n}\right)$ [Gon81].

Meaning with probability at least $1 - n^{-c}$ for constant $c > 0$.

One-Choice and Two-Choice processes

<u>One-Choice Process:</u>
Iteration: For each $t \geq 0$, sample **one** bin uniformly at random (u.a.r.) and place the ball there.

- In the lightly-loaded case ($m = n$), w.h.p. $\text{Gap}(n) = \Theta\left(\frac{\log n}{\log\log n}\right)$ [Gon81].
- In the heavily-loaded case ($m \gg n$), w.h.p. $\text{Gap}(m) = \Theta\left(\sqrt{\frac{m}{n} \cdot \log n}\right)$ (e.g. [RS98]).

One-Choice and Two-Choice processes

<u>One-Choice Process:</u>
Iteration: For each $t \geq 0$, sample **one** bin uniformly at random (u.a.r.) and place the ball there.

- In the lightly-loaded case ($m = n$), w.h.p. $\mathrm{Gap}(n) = \Theta\left(\frac{\log n}{\log\log n}\right)$ [Gon81].
- In the heavily-loaded case ($m \gg n$), w.h.p. $\mathrm{Gap}(m) = \Theta\left(\sqrt{\frac{m}{n} \cdot \log n}\right)$ (e.g. [RS98]).

<u>Two-Choice Process:</u>
Iteration: For each $t \geq 0$, sample **two** bins independently u.a.r. and place the ball in the least loaded of the two.

One-Choice and Two-Choice processes

<u>One-Choice Process:</u>
Iteration: For each $t \geq 0$, sample **one** bin uniformly at random (u.a.r.) and place the ball there.

- In the lightly-loaded case ($m = n$), w.h.p. $\mathrm{Gap}(n) = \Theta\left(\frac{\log n}{\log\log n}\right)$ [Gon81].
- In the heavily-loaded case ($m \gg n$), w.h.p. $\mathrm{Gap}(m) = \Theta\left(\sqrt{\frac{m}{n} \cdot \log n}\right)$ (e.g. [RS98]).

<u>Two-Choice Process:</u>
Iteration: For each $t \geq 0$, sample **two** bins independently u.a.r. and place the ball in the least loaded of the two.

- In the lightly-loaded case ($m = n$), w.h.p. $\mathrm{Gap}(n) = \log_2 \log n + \Theta(1)$ [KLMadH96, ABKU99].

One-Choice and Two-Choice processes

<u>One-Choice Process:</u>
Iteration: For each $t \geq 0$, sample **one** bin uniformly at random (u.a.r.) and place the ball there.

- In the lightly-loaded case ($m = n$), w.h.p. $\mathrm{Gap}(n) = \Theta\left(\frac{\log n}{\log \log n}\right)$ [Gon81].
- In the heavily-loaded case ($m \gg n$), w.h.p. $\mathrm{Gap}(m) = \Theta\left(\sqrt{\frac{m}{n} \cdot \log n}\right)$ (e.g. [RS98]).

<u>Two-Choice Process:</u>
Iteration: For each $t \geq 0$, sample **two** bins independently u.a.r. and place the ball in the least loaded of the two.

- In the lightly-loaded case ($m = n$), w.h.p. $\mathrm{Gap}(n) = \log_2 \log n + \Theta(1)$ [KLMadH96, ABKU99].

One-Choice and Two-Choice processes

<u>One-Choice Process:</u>
Iteration: For each $t \geq 0$, sample **one** bin uniformly at random (u.a.r.) and place the ball there.

- In the lightly-loaded case ($m = n$), w.h.p. $\text{Gap}(n) = \Theta\left(\frac{\log n}{\log \log n}\right)$ [Gon81].
- In the heavily-loaded case ($m \gg n$), w.h.p. $\text{Gap}(m) = \Theta\left(\sqrt{\frac{m}{n} \cdot \log n}\right)$ (e.g. [RS98]).

<u>Two-Choice Process:</u>
Iteration: For each $t \geq 0$, sample **two** bins independently u.a.r. and place the ball in the least loaded of the two.

- In the lightly-loaded case ($m = n$), w.h.p. $\text{Gap}(n) = \log_2 \log n + \Theta(1)$ [KLMadH96, ABKU99].
- In the heavily-loaded case ($m \gg n$), w.h.p. $\text{Gap}(m) = \log_2 \log n + \Theta(1)$ [BCSV06].

One-Choice and Two-Choice processes

<u>One-Choice Process:</u>
Iteration: For each $t \geq 0$, sample **one** bin uniformly at random (u.a.r.) and place the ball there.

- In the lightly-loaded case ($m = n$), w.h.p. $\mathrm{Gap}(n) = \Theta\left(\frac{\log n}{\log\log n}\right)$ [Gon81].
- In the heavily-loaded case ($m \gg n$), w.h.p. $\mathrm{Gap}(m) = \Theta\left(\sqrt{\frac{m}{n} \cdot \log n}\right)$ (e.g. [RS98]).

<u>Two-Choice Process:</u>
Iteration: For each $t \geq 0$, sample **two** bins independently u.a.r. and place the ball in the least loaded of the two.

- In the lightly-loaded case ($m = n$), w.h.p. $\mathrm{Gap}(n) = \log_2 \log n + \Theta(1)$ [KLMadH96, ABKU99].
- In the heavily-loaded case ($m \gg n$), w.h.p. $\mathrm{Gap}(m) = \log_2 \log n + \Theta(1)$ [BCSV06].

Probability allocation vectors

Probability allocation vectors

- **Probability allocation vector** p^t, where p_i^t is the prob. of allocating to i-th most loaded bin.

Probability allocation vectors

- **Probability allocation vector** p^t, where p_i^t is the prob. of allocating to i-th most loaded bin.
- For ONE-CHOICE,

$$p_{\textsc{One-Choice}} = \left(\frac{1}{n}, \frac{1}{n}, \ldots, \frac{1}{n}\right).$$

Probability allocation vectors

- **Probability allocation vector** p^t, where p_i^t is the prob. of allocating to i-th most loaded bin.
- For One-Choice,

$$p_{\textsc{One-Choice}} = \left(\frac{1}{n}, \frac{1}{n}, \ldots, \frac{1}{n}\right).$$

- For Two-Choice,

$$p_{\textsc{Two-Choice}} = \left(\frac{1}{n^2}, \frac{3}{n^2}, \ldots, \frac{2i-1}{n^2}, \ldots, \frac{2n-2}{n^2}\right).$$

Probability allocation vectors

- **Probability allocation vector** p^t, where p_i^t is the prob. of allocating to i-th most loaded bin.
- For One-Choice,

$$p_{\textsc{One-Choice}} = \Big(\frac{1}{n}, \frac{1}{n}, \ldots, \frac{1}{n}\Big).$$

- For Two-Choice,

$$p_{\textsc{Two-Choice}} = \Big(\frac{1}{n^2}, \frac{3}{n^2}, \ldots, \frac{2i-1}{n^2}, \ldots, \frac{2n-2}{n^2}\Big).$$

- [PTW15] studied ϵ-biased processes that:

Probability allocation vectors

- **Probability allocation vector** p^t, where p_i^t is the prob. of allocating to i-th most loaded bin.
- For One-Choice,

$$p_{\textsc{One-Choice}} = \Big(\frac{1}{n}, \frac{1}{n}, \ldots, \frac{1}{n}\Big).$$

- For Two-Choice,

$$p_{\textsc{Two-Choice}} = \Big(\frac{1}{n^2}, \frac{3}{n^2}, \ldots, \frac{2i-1}{n^2}, \ldots, \frac{2n-2}{n^2}\Big).$$

- [PTW15] studied ϵ-biased processes that:
  - Have p is *non-decreasing*,

Probability allocation vectors

- **Probability allocation vector** p^t, where p_i^t is the prob. of allocating to i-th most loaded bin.
- For One-Choice,

$$p_{\textsc{One-Choice}} = \Big(\frac{1}{n}, \frac{1}{n}, \ldots, \frac{1}{n}\Big).$$

- For Two-Choice,

$$p_{\textsc{Two-Choice}} = \Big(\frac{1}{n^2}, \frac{3}{n^2}, \ldots, \frac{2i-1}{n^2}, \ldots, \frac{2n-2}{n^2}\Big).$$

- [PTW15] studied ϵ-biased processes that:
  - Have p is *non-decreasing*,
  - For some constant $\delta \in (0,1)$, satisfy

$$p_{\delta n} \leq \frac{1-\epsilon}{n}.$$

Probability allocation vectors

- **Probability allocation vector** p^t, where p_i^t is the prob. of allocating to i-th most loaded bin.
- For One-Choice,

$$p_{\textsc{One-Choice}} = \left(\frac{1}{n}, \frac{1}{n}, \ldots, \frac{1}{n}\right).$$

- For Two-Choice,

$$p_{\textsc{Two-Choice}} = \left(\frac{1}{n^2}, \frac{3}{n^2}, \ldots, \frac{2i-1}{n^2}, \ldots, \frac{2n-2}{n^2}\right).$$

- [PTW15] studied ϵ-biased processes that:
  - Have p is *non-decreasing*,
  - For some constant $\delta \in (0, 1)$, satisfy

$$p_{\delta n} \leq \frac{1-\epsilon}{n}.$$

- They showed such processes achieve w.h.p. an $\mathcal{O}(\log n)$ gap, for constant $\epsilon > 0$.

$(1+\beta)$ process

$(1+\beta)$ Process:
Parameter: A *mixing factor* $\beta \in (0,1]$.
Iteration: For each $t \geq 0$, with probability β allocate one ball via the Two-Choice process, otherwise allocate one ball via the One-Choice process.

$(1+\beta)$ process

$(1+\beta)$ Process:
Parameter: A *mixing factor* $\beta \in (0,1]$.
Iteration: For each $t \geq 0$, with probability β allocate one ball via the Two-Choice process, otherwise allocate one ball via the One-Choice process.

- Introduced by Mitzenmacher [Mit96] as a *faulty setting* for Two-Choice.

$(1+\beta)$ process

$(1+\beta)$ Process:
Parameter: A *mixing factor* $\beta \in (0,1]$.
Iteration: For each $t \geq 0$, with probability β allocate one ball via the Two-Choice process, otherwise allocate one ball via the One-Choice process.

- Introduced by Mitzenmacher [Mit96] as a *faulty setting* for Two-Choice.
- Its probability vector is given by,

$$p_{(1+\beta)} = \beta \cdot p_{\text{Two-Choice}} + (1-\beta) \cdot p_{\text{One-Choice}}.$$

$(1+\beta)$ process

$(1+\beta)$ Process:
Parameter: A *mixing factor* $\beta \in (0,1]$.
Iteration: For each $t \geq 0$, with probability β allocate one ball via the TWO-CHOICE process, otherwise allocate one ball via the ONE-CHOICE process.

- Introduced by Mitzenmacher [Mit96] as a *faulty setting* for TWO-CHOICE.
- Its probability vector is given by,

$$p_{(1+\beta)} = \beta \cdot p_{\text{TWO-CHOICE}} + (1-\beta) \cdot p_{\text{ONE-CHOICE}}.$$

- In the heavily-loaded case, [PTW15] proved that the gap is w.h.p. $\mathcal{O}\Big(\frac{\log n}{\beta} + \frac{\log(1/\beta)}{\beta}\Big)$ for any $\beta \in (0,1]$.

Settings

Batched Setting

- Two-Choice assumes that the reported bin loads are *up-to-date*.

Batched Setting

- Two-Choice assumes that the reported bin loads are *up-to-date*.
- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] relaxed this assumption, by allocating b balls in parallel.

Batched Setting

- Two-Choice assumes that the reported bin loads are *up-to-date*.
- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] relaxed this assumption, by allocating b balls in parallel.
- They showed that for $b = n$, the gap is w.h.p. $\mathcal{O}(\log n)$.

Batched Setting

- Two-Choice assumes that the reported bin loads are *up-to-date*.
- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] relaxed this assumption, by allocating b balls in parallel.
- They showed that for $b = n$, the gap is w.h.p. $\mathcal{O}(\log n)$.
- The authors [LS22b] showed that for $b = n$, the gap is w.h.p. $\Theta\left(\frac{\log n}{\log \log n}\right)$

Batched Setting

- Two-Choice assumes that the reported bin loads are *up-to-date*.
- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE$^+$12] relaxed this assumption, by allocating b balls in parallel.
- They showed that for $b = n$, the gap is w.h.p. $\mathcal{O}(\log n)$.
- The authors [LS22b] showed that for $b = n$, the gap is w.h.p. $\Theta\left(\frac{\log n}{\log \log n}\right)$ and for $b \in [n, n \log n]$ that it follows the gap of One-Choice for b balls.

Batched Setting

- TWO-CHOICE assumes that the reported bin loads are *up-to-date*.
- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] relaxed this assumption, by allocating b balls in parallel.
- They showed that for $b = n$, the gap is w.h.p. $\mathcal{O}(\log n)$.
- The authors [LS22b] showed that for $b = n$, the gap is w.h.p. $\Theta\left(\frac{\log n}{\log \log n}\right)$ and for $b \in [n, n \log n]$ that it follows the gap of ONE-CHOICE for b balls.

What happens for TWO-CHOICE when $b \geq n \log n$?

Batched Setting

- Two-Choice assumes that the reported bin loads are *up-to-date*.
- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] relaxed this assumption, by allocating b balls in parallel.
- They showed that for $b = n$, the gap is w.h.p. $\mathcal{O}(\log n)$.
- The authors [LS22b] showed that for $b = n$, the gap is w.h.p. $\Theta\left(\frac{\log n}{\log \log n}\right)$ and for $b \in [n, n \log n]$ that it follows the gap of One-Choice for b balls.

What happens for Two-Choice when $b \geq n \log n$?

What happens for other processes (e.g., the ϵ-biased processes)?

Batched Setting

- Two-Choice assumes that the reported bin loads are *up-to-date*.
- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] relaxed this assumption, by allocating b balls in parallel.
- They showed that for $b = n$, the gap is w.h.p. $\mathcal{O}(\log n)$.
- The authors [LS22b] showed that for $b = n$, the gap is w.h.p. $\Theta\left(\frac{\log n}{\log \log n}\right)$ and for $b \in [n, n \log n]$ that it follows the gap of One-Choice for b balls.

What happens for Two-Choice when $b \geq n \log n$?

What happens for other processes (e.g., the ϵ-biased processes)?

Open in Visualiser.

Weighted Setting

Weighted Setting

- Balls have weights sampled from a distribution $\mathcal{W}$

Weighted Setting

- Balls have weights sampled from a distribution $\mathcal{W}$ with $\mathbf{E}[\mathcal{W}] = 1$

Weighted Setting

- Balls have weights sampled from a distribution $\mathcal{W}$ with $\mathbf{E}\left[\,\mathcal{W}\,\right] = 1$ and $\mathbf{E}\left[\,e^{\zeta\mathcal{W}}\,\right] < c$ for constants $\zeta, c > 0$.

Weighted Setting

- Balls have weights sampled from a distribution $\mathcal{W}$ with $\mathbf{E}[\,\mathcal{W}\,] = 1$ and $\mathbf{E}\left[\,e^{\zeta \mathcal{W}}\,\right] < c$ for constants $\zeta, c > 0$.
- [PTW15] showed that ϵ-biased processes achieve w.h.p. $\mathcal{O}(\log n)$ gap.

Open in Visualiser.

Two-Choice in the Graphical Setting

Two-Choice in the Graphical Setting

- Given a graph $G = (V, E)$, where the vertices are bins. For each ball:

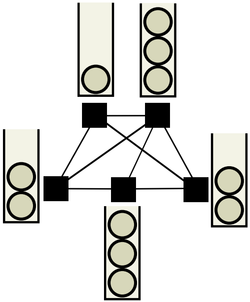

Two-Choice in the Graphical Setting

- Given a graph $G = (V, E)$, where the vertices are bins. For each ball:
  - Sample an edge u.a.r.

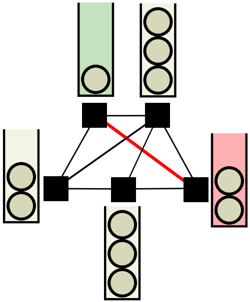

Two-Choice in the Graphical Setting

- Given a graph $G = (V, E)$, where the vertices are bins. For each ball:
  - Sample an edge u.a.r.
  - Allocate the ball to the least loaded of its *two* adjacent bins.

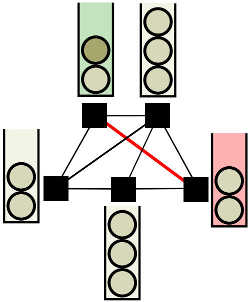

Two-Choice in the Graphical Setting

- Given a graph $G = (V, E)$, where the vertices are bins. For each ball:
  - Sample an edge u.a.r.
  - Allocate the ball to the least loaded of its *two* adjacent bins.

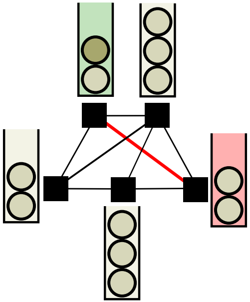

- For any *d-regular* graph with *conductance* Φ, the gap is w.h.p. $\mathcal{O}\left(\frac{\log n}{\Phi}\right)$ [PTW15].

Two-Choice in the Graphical Setting

- Given a graph $G = (V, E)$, where the vertices are bins. For each ball:
  - Sample an edge u.a.r.
  - Allocate the ball to the least loaded of its *two* adjacent bins.

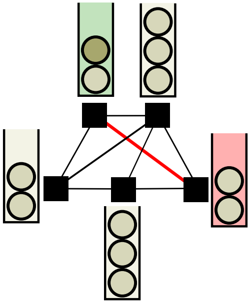

- For any *d-regular* graph with *conductance* Φ, the gap is w.h.p. $\mathcal{O}\left(\frac{\log n}{\Phi}\right)$ [PTW15].

Do similar bounds hold for the weighted graphical setting? (Open Question 1, [PTW15])

Results for Batching

Results (I): Batching

Results (I): Batching

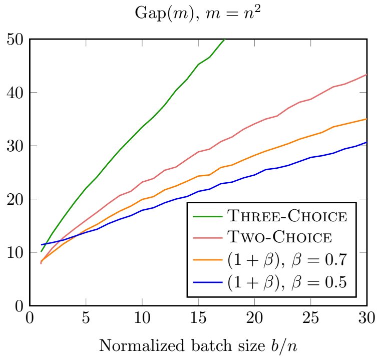

Results (I): Batching

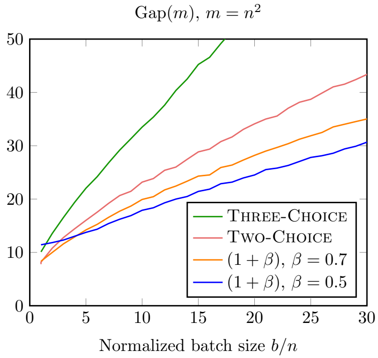

- For $b \leq n \log n$,

Results (I): Batching

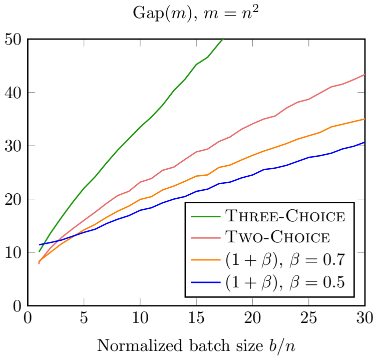

- For $b \leq n \log n$,
  - $(1+\beta)$, Quantile(δ) have w.h.p. $\Omega(\log n)$ gap.

Results (I): Batching

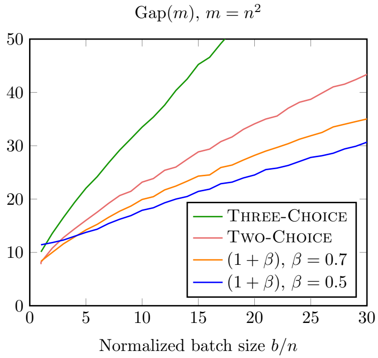

- For $b \leq n \log n$,
  - $(1+\beta)$, Quantile(δ) have w.h.p. $\Omega(\log n)$ gap.
  - Two-Choice "follows" One-Choice with b balls [LS22b].

Results (I): Batching

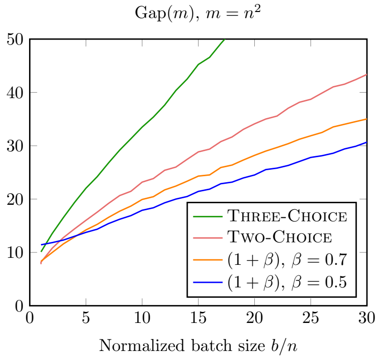

- For $b \leq n \log n$,
  - $(1+\beta)$, Quantile(δ) have w.h.p. $\Omega(\log n)$ gap.
  - Two-Choice "follows" One-Choice with b balls [LS22b].
- For any ϵ-biased process with $p_n \leq \frac{C}{n}$, for constant $\epsilon, C > 0$:

Results (I): Batching

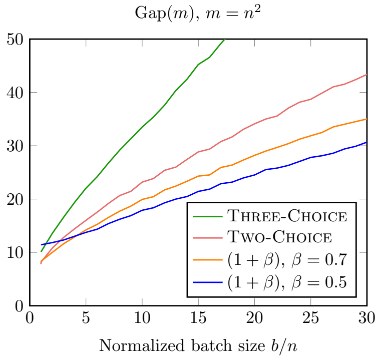

- For $b \leq n \log n$,
  - $(1+\beta)$, Quantile(δ) have w.h.p. $\Omega(\log n)$ gap.
  - Two-Choice "follows" One-Choice with b balls [LS22b].
- For any ϵ-biased process with $p_n \leq \frac{C}{n}$, for constant $\epsilon, C > 0$:
  - For any $b \geq n$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} \cdot \log n)$.

Results (I): Batching

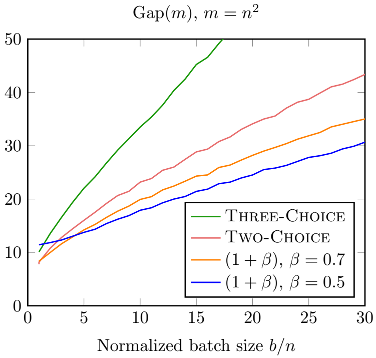

- For $b \leq n \log n$,
  - $(1+\beta)$, Quantile(δ) have w.h.p. $\Omega(\log n)$ gap.
  - Two-Choice "follows" One-Choice with b balls [LS22b].
- For any ϵ-biased process with $p_n \leq \frac{C}{n}$, for constant $\epsilon, C > 0$:
  - For any $b \geq n$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} \cdot \log n)$.
  - For any $b \in [n, n^3]$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} + \log n)$.

Results (I): Batching

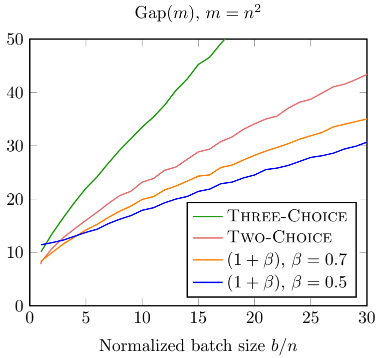

- For $b \leq n \log n$,
  - $(1+\beta)$, Quantile(δ) have w.h.p. $\Omega(\log n)$ gap.
  - Two-Choice "follows" One-Choice with b balls [LS22b].
- For any ϵ-biased process with $p_n \leq \frac{C}{n}$, for constant $\epsilon, C > 0$:
  - For any $b \geq n$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} \cdot \log n)$.
  - For any $b \in [n, n^3]$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} + \log n)$.
  - Same bounds hold for weighted balls.

Results (I): Batching

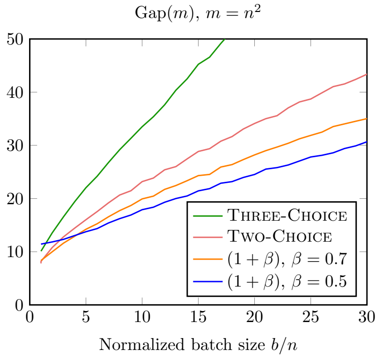

- For $b \leq n \log n$,
  - $(1+\beta)$, Quantile(δ) have w.h.p. $\Omega(\log n)$ gap.
  - Two-Choice "follows" One-Choice with b balls [LS22b].
- For any ϵ-biased process with $p_n \leq \frac{C}{n}$, for constant $\epsilon, C > 0$:
  - For any $b \geq n$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} \cdot \log n)$.
  - For any $b \in [n, n^3]$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} + \log n)$.
  - Same bounds hold for weighted balls.
- For $b \in [n \log n, n^3]$, for any process with $p_n \geq \frac{1+C'}{n}$ for constant $C' > 0$, we prove a lower bound of $\Omega(C' \cdot \frac{b}{n})$.

Results (I): Batching

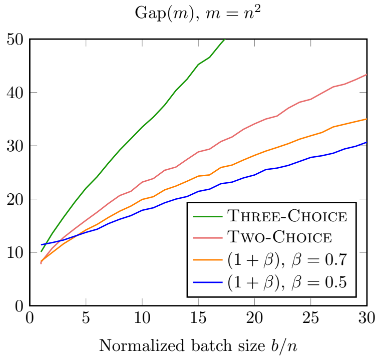

p_n is: $\approx \frac{3}{n}$

- For $b \leq n \log n$,
  - $(1 + \beta)$, Quantile(δ) have w.h.p. $\Omega(\log n)$ gap.
  - Two-Choice "follows" One-Choice with b balls [LS22b].
- For any ϵ-biased process with $p_n \leq \frac{C}{n}$, for constant $\epsilon, C > 0$:
  - For any $b \geq n$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} \cdot \log n)$.
  - For any $b \in [n, n^3]$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} + \log n)$.
  - Same bounds hold for weighted balls.
- For $b \in [n \log n, n^3]$, for any process with $p_n \geq \frac{1+C'}{n}$ for constant $C' > 0$, we prove a lower bound of $\Omega(C' \cdot \frac{b}{n})$.

Results (I): Batching

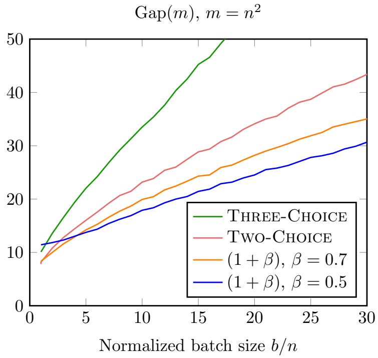

p_n is: $\approx \frac{3}{n}$, $\approx \frac{2}{n}$

- For $b \leq n \log n$,
  - $(1+\beta)$, Quantile(δ) have w.h.p. $\Omega(\log n)$ gap.
  - Two-Choice "follows" One-Choice with b balls [LS22b].
- For any ϵ-biased process with $p_n \leq \frac{C}{n}$, for constant $\epsilon, C > 0$:
  - For any $b \geq n$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} \cdot \log n)$.
  - For any $b \in [n, n^3]$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} + \log n)$.
  - Same bounds hold for weighted balls.
- For $b \in [n \log n, n^3]$, for any process with $p_n \geq \frac{1+C'}{n}$ for constant $C' > 0$, we prove a lower bound of $\Omega(C' \cdot \frac{b}{n})$.

Results (I): Batching

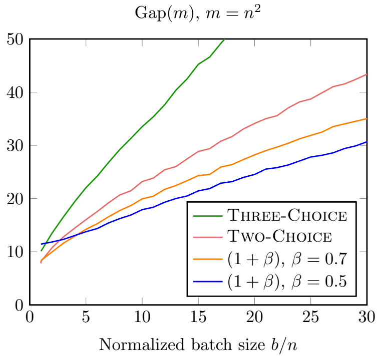

p_n is: $\approx \frac{3}{n}$, $\approx \frac{2}{n}$, $\approx \frac{1+0.7}{n}$

- For $b \leq n \log n$,
  - $(1+\beta)$, Quantile(δ) have w.h.p. $\Omega(\log n)$ gap.
  - Two-Choice "follows" One-Choice with b balls [LS22b].
- For any ϵ-biased process with $p_n \leq \frac{C}{n}$, for constant $\epsilon, C > 0$:
  - For any $b \geq n$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} \cdot \log n)$.
  - For any $b \in [n, n^3]$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} + \log n)$.
  - Same bounds hold for weighted balls.
- For $b \in [n \log n, n^3]$, for any process with $p_n \geq \frac{1+C'}{n}$ for constant $C' > 0$, we prove a lower bound of $\Omega(C' \cdot \frac{b}{n})$.

Results (I): Batching

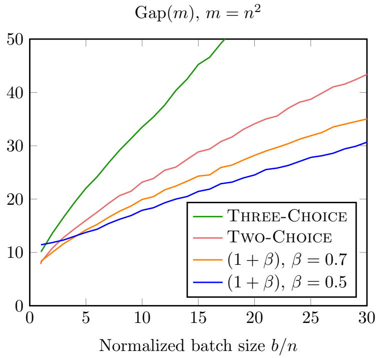

p_n is: $\approx \frac{3}{n}$, $\approx \frac{2}{n}$, $\approx \frac{1+0.7}{n}$ and $\approx \frac{1+0.5}{n}$.

- For $b \leq n \log n$,
  - $(1+\beta)$, Quantile(δ) have w.h.p. $\Omega(\log n)$ gap.
  - Two-Choice "follows" One-Choice with b balls [LS22b].
- For any ϵ-biased process with $p_n \leq \frac{C}{n}$, for constant $\epsilon, C > 0$:
  - For any $b \geq n$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} \cdot \log n)$.
  - For any $b \in [n, n^3]$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} + \log n)$.
  - Same bounds hold for weighted balls.
- For $b \in [n \log n, n^3]$, for any process with $p_n \geq \frac{1+C'}{n}$ for constant $C' > 0$, we prove a lower bound of $\Omega(C' \cdot \frac{b}{n})$.

Results (I): Batching

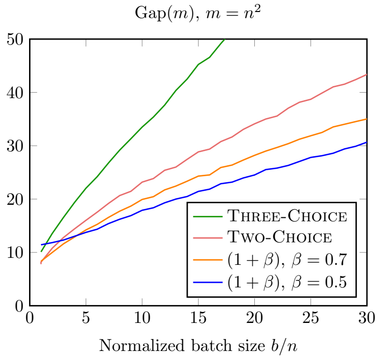

p_n is: $\approx \frac{3}{n}$, $\approx \frac{2}{n}$, $\approx \frac{1+0.7}{n}$ and $\approx \frac{1+0.5}{n}$.

- For $b \leq n \log n$,
  - $(1+\beta)$, Quantile(δ) have w.h.p. $\Omega(\log n)$ gap.
  - Two-Choice "follows" One-Choice with b balls [LS22b].
- For any ϵ-biased process with $p_n \leq \frac{C}{n}$, for constant $\epsilon, C > 0$:
  - For any $b \geq n$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} \cdot \log n)$.
  - For any $b \in [n, n^3]$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} + \log n)$.
  - Same bounds hold for weighted balls.
- For $b \in [n \log n, n^3]$, for any process with $p_n \geq \frac{1+C'}{n}$ for constant $C' > 0$, we prove a lower bound of $\Omega(C' \cdot \frac{b}{n})$.
  - More choices **do not** always help.

Results (I): Batching

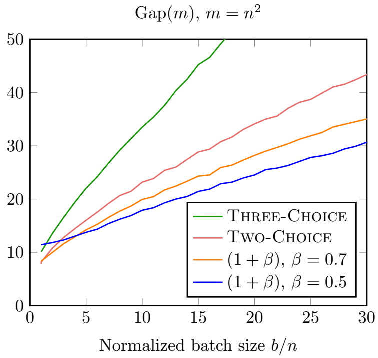

p_n is: $\approx \frac{3}{n}$, $\approx \frac{2}{n}$, $\approx \frac{1+0.7}{n}$ and $\approx \frac{1+0.5}{n}$.

- For $b \leq n \log n$,
  - $(1+\beta)$, Quantile(δ) have w.h.p. $\Omega(\log n)$ gap.
  - Two-Choice "follows" One-Choice with b balls [LS22b].
- For any ϵ-biased process with $p_n \leq \frac{C}{n}$, for constant $\epsilon, C > 0$:
  - For any $b \geq n$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} \cdot \log n)$.
  - For any $b \in [n, n^3]$, w.h.p. $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} + \log n)$.
  - Same bounds hold for weighted balls.
- For $b \in [n \log n, n^3]$, for any process with $p_n \geq \frac{1+C'}{n}$ for constant $C' > 0$, we prove a lower bound of $\Omega(C' \cdot \frac{b}{n})$.
  - More choices **do not** always help.
  - For some values of d, $(1+\beta)$ has a better gap.

A closer look at a single batch

A closer look at a single batch

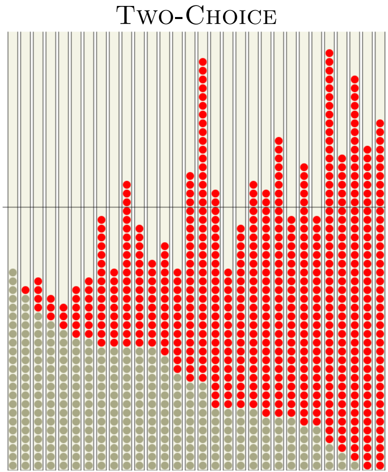

$$p_i = \frac{2i-1}{n^2}$$

A closer look at a single batch

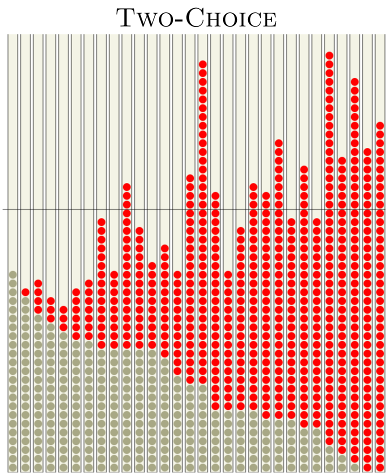

$$p_i = \frac{2i - 1}{n^2}$$

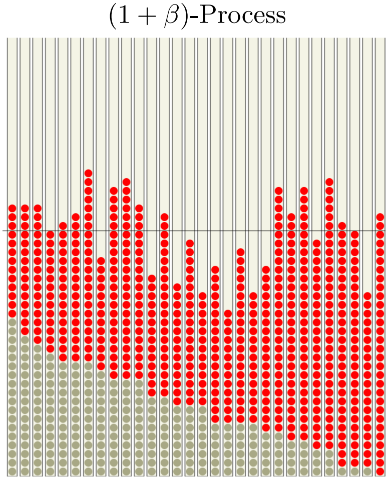

$$p_i = \beta \cdot \frac{2i - 1}{n^2} + (1 - \beta) \cdot \frac{1}{n}$$

A closer look at a single batch

Two-Choice

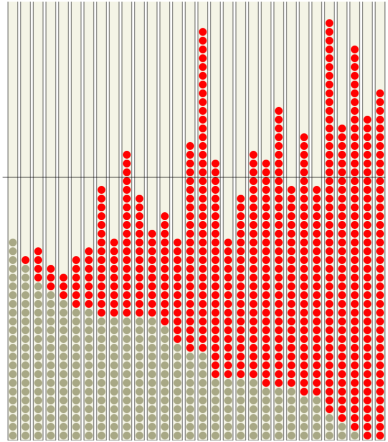

$$p_i = \frac{2i-1}{n^2}$$

$(1+\beta)$-Process

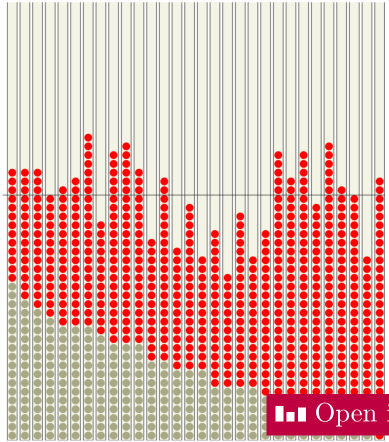

Open in Visualiser.

$$p_i = \beta \cdot \frac{2i-1}{n^2} + (1-\beta) \cdot \frac{1}{n}$$

Further Results

Upper Bound Tools: Hyperbolic Cosine Potential

Upper Bound Tools: Hyperbolic Cosine Potential

- [PTW15] used the **hyperbolic cosine potential**

$$\Gamma^t := \Gamma(\gamma) := \underbrace{\sum_{i=1}^{n} e^{\gamma(x_i^t - t/n)}}_{\text{Overload potential}} + \underbrace{\sum_{i=1}^{n} e^{-\gamma(x_i^t - t/n)}}_{\text{Underload potential}} .$$

Upper Bound Tools: Hyperbolic Cosine Potential

- [PTW15] used the **hyperbolic cosine potential**

$$\Gamma^t := \Gamma(\gamma) := \underbrace{\sum_{i=1}^{n} e^{\gamma(x_i^t - t/n)}}_{\text{Overload potential}} + \underbrace{\sum_{i=1}^{n} e^{-\gamma(x_i^t - t/n)}}_{\text{Underload potential}} .$$

- For the $(1+\beta)$-process, $\gamma = \Theta(\beta)$.

Upper Bound Tools: Hyperbolic Cosine Potential

- [PTW15] used the **hyperbolic cosine potential**

$$\Gamma^t := \Gamma(\gamma) := \underbrace{\sum_{i=1}^{n} e^{\gamma(x_i^t - t/n)}}_{\text{Overload potential}} + \underbrace{\sum_{i=1}^{n} e^{-\gamma(x_i^t - t/n)}}_{\text{Underload potential}} .$$

- For the $(1+\beta)$-process, $\gamma = \Theta(\beta)$.
- [PTW15] show that $\mathbf{E}\left[\Gamma^{t+1} \mid \mathfrak{F}^t\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma}{n}\right) + c_2$.

Upper Bound Tools: Hyperbolic Cosine Potential

- [PTW15] used the **hyperbolic cosine potential**

$$\Gamma^t := \Gamma(\gamma) := \underbrace{\sum_{i=1}^{n} e^{\gamma(x_i^t - t/n)}}_{\text{Overload potential}} + \underbrace{\sum_{i=1}^{n} e^{-\gamma(x_i^t - t/n)}}_{\text{Underload potential}} .$$

- For the $(1+\beta)$-process, $\gamma = \Theta(\beta)$.
- [PTW15] show that $\mathbf{E}\left[\Gamma^{t+1} \mid \mathfrak{F}^t\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma}{n}\right) + c_2$.
- By *induction*, this implies $\mathbf{E}\left[\Gamma^t\right] \leq \frac{c_2}{c_1 \gamma} \cdot n$ for any $t \geq 0$.

Upper Bound Tools: Hyperbolic Cosine Potential

- [PTW15] used the **hyperbolic cosine potential**

$$\Gamma^t := \Gamma(\gamma) := \underbrace{\sum_{i=1}^{n} e^{\gamma(x_i^t - t/n)}}_{\text{Overload potential}} + \underbrace{\sum_{i=1}^{n} e^{-\gamma(x_i^t - t/n)}}_{\text{Underload potential}} .$$

- For the $(1+\beta)$-process, $\gamma = \Theta(\beta)$.
- [PTW15] show that $\mathbf{E}\left[\,\Gamma^{t+1} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1\gamma}{n}\right) + c_2$.
- By *induction*, this implies $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1\gamma} \cdot n$ for any $t \geq 0$.
- By *Markov's inequality*, we get $\mathbf{Pr}\left[\,\Gamma^m \leq \frac{c_2}{c_1\gamma} n^3\,\right] \geq 1 - n^{-2}$,

Upper Bound Tools: Hyperbolic Cosine Potential

- [PTW15] used the **hyperbolic cosine potential**

$$\Gamma^t := \Gamma(\gamma) := \underbrace{\sum_{i=1}^{n} e^{\gamma(x_i^t - t/n)}}_{\text{Overload potential}} + \underbrace{\sum_{i=1}^{n} e^{-\gamma(x_i^t - t/n)}}_{\text{Underload potential}} .$$

- For the $(1+\beta)$-process, $\gamma = \Theta(\beta)$.
- [PTW15] show that $\mathbf{E}\left[\,\Gamma^{t+1} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma}{n}\right) + c_2$.
- By *induction*, this implies $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1 \gamma} \cdot n$ for any $t \geq 0$.
- By *Markov's inequality*, we get $\mathbf{Pr}\left[\,\Gamma^m \leq \frac{c_2}{c_1\gamma} n^3\,\right] \geq 1 - n^{-2}$, which implies

$$\mathbf{Pr}\left[\,\mathrm{Gap}(m) \leq \frac{1}{\gamma}\left(3 \cdot \log n + \log\left(\frac{c_2}{c_1 \gamma}\right)\right)\right] \geq 1 - n^{-2}.$$

Upper Bound Tools: Hyperbolic Cosine Potential

- [PTW15] used the **hyperbolic cosine potential**

$$\Gamma^t := \Gamma(\gamma) := \underbrace{\sum_{i=1}^{n} e^{\gamma(x_i^t - t/n)}}_{\text{Overload potential}} + \underbrace{\sum_{i=1}^{n} e^{-\gamma(x_i^t - t/n)}}_{\text{Underload potential}} .$$

- For the $(1+\beta)$-process, $\gamma = \Theta(\beta)$.
- [PTW15] show that $\mathbf{E}\left[\Gamma^{t+1} \mid \mathfrak{F}^t\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma}{n}\right) + c_2$.
- By *induction*, this implies $\mathbf{E}\left[\Gamma^t\right] \leq \frac{c_2}{c_1 \gamma} \cdot n$ for any $t \geq 0$.
- By *Markov's inequality*, we get $\mathbf{Pr}\left[\Gamma^m \leq \frac{c_2}{c_1 \gamma} n^3\right] \geq 1 - n^{-2}$, which implies

$$\mathbf{Pr}\left[\mathrm{Gap}(m) \leq \frac{1}{\gamma}\left(3 \cdot \log n + \log\left(\frac{c_2}{c_1 \gamma}\right)\right)\right] \geq 1 - n^{-2}.$$

- This gives the $\mathcal{O}(\frac{\log n}{\beta} + \frac{\log(1/\beta)}{\beta})$.

Results (II): Implications of the Upper Bound

- **Extension 1:** Improve the *additive term* in the recurrence inequality

$$\mathbf{E}\left[\,\Gamma^{t+1} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma}{n}\right) + c_2.$$

Results (II): Implications of the Upper Bound

- **Extension 1:** Improve the *additive term* in the recurrence inequality
$$\mathbf{E}\left[\,\Gamma^{t+1} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1\gamma}{n}\right) + c_2 \cdot \gamma.$$

Results (II): Implications of the Upper Bound

- **Extension 1:** Improve the *additive term* in the recurrence inequality
$$\mathbf{E}\left[\, \Gamma^{t+1} \mid \mathfrak{F}^{t} \,\right] \leq \Gamma^{t} \cdot \left(1 - \frac{c_1 \gamma}{n}\right) + c_2 \cdot \gamma.$$

Results (II): Implications of the Upper Bound

- **Extension 1:** Improve the *additive term* in the recurrence inequality

$$\mathbf{E}\left[\,\Gamma^{t+1} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1\gamma}{n}\right) + c_2 \cdot \gamma.$$

  - Implies that $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1} \cdot n$.

Results (II): Implications of the Upper Bound

- **Extension 1:** Improve the *additive term* in the recurrence inequality

$$\mathbf{E}\left[\,\Gamma^{t+1} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma}{n}\right) + c_2 \cdot \gamma.$$

  - Implies that $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1} \cdot n$.
  - Implies w.h.p. an $\mathcal{O}(\frac{\log n}{\beta})$ gap for the $(1+\beta)$-process.

Results (II): Implications of the Upper Bound

- **Extension 1:** Improve the *additive term* in the recurrence inequality

$$\mathbf{E}\left[\,\Gamma^{t+1} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1\gamma}{n}\right) + c_2 \cdot \gamma.$$

  - Implies that $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1} \cdot n$.
  - Implies w.h.p. an $\mathcal{O}(\frac{\log n}{\beta})$ gap for the $(1+\beta)$-process.

$\frac{1}{\gamma} \cdot \left(3\log n + \log(\frac{c_2}{c_1})\right)$

Results (II): Implications of the Upper Bound

- **Extension 1:** Improve the *additive term* in the recurrence inequality

$$\mathbf{E}\left[\,\Gamma^{t+1} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma}{n}\right) + c_2 \cdot \gamma.$$

  - Implies that $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1} \cdot n$.
  - Implies w.h.p. an $\mathcal{O}(\frac{\log n}{\beta})$ gap for the $(1+\beta)$-process.
- **Extension 2:** Extend to $b \geq n$ steps for $\gamma = \Theta(n/b)$,

$$\mathbf{E}\left[\,\Gamma^{t+1} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma}{n}\right) + c_2 \cdot \gamma.$$

Results (II): Implications of the Upper Bound

- **Extension 1:** Improve the *additive term* in the recurrence inequality

$$\mathbf{E}\left[\,\Gamma^{t+1} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1\gamma}{n}\right) + c_2 \cdot \gamma.$$

  - Implies that $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1} \cdot n$.
  - Implies w.h.p. an $\mathcal{O}(\frac{\log n}{\beta})$ gap for the $(1+\beta)$-process.
- **Extension 2:** Extend to $b \geq n$ steps for $\gamma = \Theta(n/b)$,

$$\mathbf{E}\left[\,\Gamma^{t+b} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1\gamma b}{n}\right) + c_2 \cdot \gamma \cdot b.$$

Results (II): Implications of the Upper Bound

- **Extension 1:** Improve the *additive term* in the recurrence inequality
$$\mathbf{E}\left[\,\Gamma^{t+1} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma}{n}\right) + c_2 \cdot \gamma.$$
  - Implies that $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1} \cdot n$.
  - Implies w.h.p. an $\mathcal{O}(\frac{\log n}{\beta})$ gap for the $(1+\beta)$-process.
- **Extension 2:** Extend to $b \geq n$ steps for $\gamma = \Theta(n/b)$,
$$\mathbf{E}\left[\,\Gamma^{t+b} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma b}{n}\right) + c_2 \cdot \gamma \cdot b.$$

Results (II): Implications of the Upper Bound

- **Extension 1:** Improve the *additive term* in the recurrence inequality

$$\mathbf{E}\left[\,\Gamma^{t+1} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma}{n}\right) + c_2 \cdot \gamma.$$

  - Implies that $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1} \cdot n$.
  - Implies w.h.p. an $\mathcal{O}(\frac{\log n}{\beta})$ gap for the $(1+\beta)$-process.

- **Extension 2:** Extend to $b \geq n$ steps for $\gamma = \Theta(n/b)$,

$$\mathbf{E}\left[\,\Gamma^{t+b} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma b}{n}\right) + c_2 \cdot \gamma \cdot b.$$

  - Implies that $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1} \cdot n$.

Results (II): Implications of the Upper Bound

- **Extension 1:** Improve the *additive term* in the recurrence inequality

$$\mathbf{E}\left[\,\Gamma^{t+1} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma}{n}\right) + c_2 \cdot \gamma.$$

  - Implies that $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1} \cdot n$.
  - Implies w.h.p. an $\mathcal{O}(\frac{\log n}{\beta})$ gap for the $(1+\beta)$-process.

- **Extension 2:** Extend to $b \geq n$ steps for $\gamma = \Theta(n/b)$,

$$\mathbf{E}\left[\,\Gamma^{t+b} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma b}{n}\right) + c_2 \cdot \gamma \cdot b.$$

  - Implies that $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1} \cdot n$.
  - So, Two-Choice, $(1+\beta)$-process, Quantile(δ), with batches (and weights): $\mathrm{Gap}(m) = \mathcal{O}(\frac{b}{n} \cdot \log n)$.

Results (II): Implications of the Upper Bound

- **Extension 1:** Improve the *additive term* in the recurrence inequality

$$\mathbf{E}\left[\,\Gamma^{t+1} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma}{n}\right) + c_2 \cdot \gamma.$$

  - Implies that $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1} \cdot n$.
  - Implies w.h.p. an $\mathcal{O}(\frac{\log n}{\beta})$ gap for the $(1+\beta)$-process.
- **Extension 2:** Extend to $b \geq n$ steps for $\gamma = \Theta(n/b)$,

$$\mathbf{E}\left[\,\Gamma^{t+b} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma b}{n}\right) + c_2 \cdot \gamma \cdot b.$$

  - Implies that $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1} \cdot n$.
  - So, Two-Choice, $(1+\beta)$-process, Quantile(δ), with batches (and weights): $\mathrm{Gap}(m) = \mathcal{O}(\frac{b}{n} \cdot \log n)$.
  - For $b \in [n, n^3]$, using that $\Gamma^t = \mathcal{O}(n)$, we can improve the bound to $\mathcal{O}(\frac{b}{n} + \log n)$.

Results (II): Implications of the Upper Bound

- **Extension 1:** Improve the *additive term* in the recurrence inequality

$$\mathbf{E}\left[\, \Gamma^{t+1} \mid \mathfrak{F}^t \,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma}{n}\right) + c_2 \cdot \gamma.$$

  - Implies that $\mathbf{E}\left[\, \Gamma^t \,\right] \leq \frac{c_2}{c_1} \cdot n$.
  - Implies w.h.p. an $\mathcal{O}(\frac{\log n}{\beta})$ gap for the $(1+\beta)$-process.

- **Extension 2:** Extend to $b \geq n$ steps for $\gamma = \Theta(n/b)$,

$$\mathbf{E}\left[\, \Gamma^{t+b} \mid \mathfrak{F}^t \,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma b}{n}\right) + c_2 \cdot \gamma \cdot b.$$

  - Implies that $\mathbf{E}\left[\, \Gamma^t \,\right] \leq \frac{c_2}{c_1} \cdot n$.
  - So, Two-Choice, $(1+\beta)$-process, Quantile(δ), with batches (and weights): $\mathrm{Gap}(m) = \mathcal{O}(\frac{b}{n} \cdot \log n)$.
  - For $b \in [n, n^3]$, using that $\Gamma^t = \mathcal{O}(n)$, we can improve the bound to $\mathcal{O}(\frac{b}{n} + \log n)$.

Number of bins with load $\geq \frac{t}{n} + z$: at most $\mathcal{O}(n \cdot e^{-\gamma z})$.

Results (II): Implications of the Upper Bound

- **Extension 1:** Improve the *additive term* in the recurrence inequality

$$\mathbf{E}\left[\Gamma^{t+1} \mid \mathfrak{F}^t\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1\gamma}{n}\right) + c_2 \cdot \gamma.$$

  - Implies that $\mathbf{E}\left[\Gamma^t\right] \leq \frac{c_2}{c_1} \cdot n$.
  - Implies w.h.p. an $\mathcal{O}(\frac{\log n}{\beta})$ gap for the $(1+\beta)$-process.
- **Extension 2:** Extend to $b \geq n$ steps for $\gamma = \Theta(n/b)$,

$$\mathbf{E}\left[\Gamma^{t+b} \mid \mathfrak{F}^t\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1\gamma b}{n}\right) + c_2 \cdot \gamma \cdot b.$$

  - Implies that $\mathbf{E}\left[\Gamma^t\right] \leq \frac{c_2}{c_1} \cdot n$.
  - So, Two-Choice, $(1+\beta)$-process, Quantile(δ), with batches (and weights): $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} \cdot \log n)$.
  - For $b \in [n, n^3]$, using that $\Gamma^t = \mathcal{O}(n)$, we can improve the bound to $\mathcal{O}(\frac{b}{n} + \log n)$.
- **Extension 3:** Analysis works for a *prefix sum condition* on p.

Results (II): Implications of the Upper Bound

- **Extension 1:** Improve the *additive term* in the recurrence inequality
$$\mathbf{E}\left[\,\Gamma^{t+1} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma}{n}\right) + c_2 \cdot \gamma.$$
  - Implies that $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1} \cdot n$.
  - Implies w.h.p. an $\mathcal{O}(\frac{\log n}{\beta})$ gap for the $(1+\beta)$-process.
- **Extension 2:** Extend to $b \geq n$ steps for $\gamma = \Theta(n/b)$,
$$\mathbf{E}\left[\,\Gamma^{t+b} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma b}{n}\right) + c_2 \cdot \gamma \cdot b.$$
  - Implies that $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1} \cdot n$.
  - So, Two-Choice, $(1+\beta)$-process, Quantile(δ), with batches (and weights): $\mathrm{Gap}(m) = \mathcal{O}(\frac{b}{n} \cdot \log n)$.
  - For $b \in [n, n^3]$, using that $\Gamma^t = \mathcal{O}(n)$, we can improve the bound to $\mathcal{O}(\frac{b}{n} + \log n)$.
- **Extension 3:** Analysis works for a *prefix sum condition* on p.
  - For d-regular expanders with weights and batches $b \in [n, n^3]$: $\mathrm{Gap}(m) = \mathcal{O}(\frac{b}{n} + \log n)$.

Results (II): Implications of the Upper Bound

- **Extension 1:** Improve the *additive term* in the recurrence inequality
$$\mathbf{E}\left[\,\Gamma^{t+1} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma}{n}\right) + c_2 \cdot \gamma.$$
  - Implies that $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1} \cdot n$.
  - Implies w.h.p. an $\mathcal{O}(\frac{\log n}{\beta})$ gap for the $(1+\beta)$-process.
- **Extension 2:** Extend to $b \geq n$ steps for $\gamma = \Theta(n/b)$,
$$\mathbf{E}\left[\,\Gamma^{t+b} \mid \mathfrak{F}^t\,\right] \leq \Gamma^t \cdot \left(1 - \frac{c_1 \gamma b}{n}\right) + c_2 \cdot \gamma \cdot b.$$
  - Implies that $\mathbf{E}\left[\,\Gamma^t\,\right] \leq \frac{c_2}{c_1} \cdot n$.
  - So, Two-Choice, $(1+\beta)$-process, Quantile(δ), with batches (and weights): $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} \cdot \log n)$.
  - For $b \in [n, n^3]$, using that $\Gamma^t = \mathcal{O}(n)$, we can improve the bound to $\mathcal{O}(\frac{b}{n} + \log n)$.
- **Extension 3:** Analysis works for a *prefix sum condition* on p.
  - For d-regular expanders with weights and batches $b \in [n, n^3]$: $\text{Gap}(m) = \mathcal{O}(\frac{b}{n} + \log n)$.
  - For d-regular graphs with conductance Φ and weights: $\text{Gap}(m) = \mathcal{O}(\frac{\log n}{\Phi})$.

Future work

Future work

- Apply the refined analysis to other processes.

Future work

- Apply the refined analysis to other processes.
- Relax the synchronization assumption for batching (as in τ-Delay for Two-Choice [LS22b]).

Future work

- Apply the refined analysis to other processes.
- Relax the synchronization assumption for batching (as in τ-Delay for Two-Choice [LS22b]).
- Determine bounds that are tight up to *lower-order terms*.

Future work

- Apply the refined analysis to other processes.
- Relax the synchronization assumption for batching (as in τ-DELAY for TWO-CHOICE [LS22b]).
- Determine bounds that are tight up to *lower-order terms*.

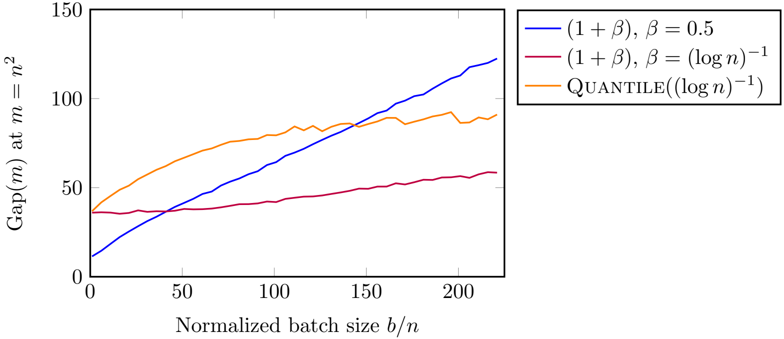

Questions?

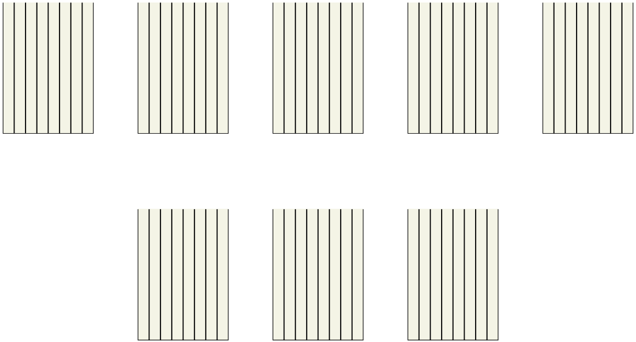

More visualisations: `dimitrioslos.com/spaa22`

Appendix A: Summary of Results

Process	Graphical	Batch Size	Weights	Gap Bound	Reference
Two-Choice	–	$b = n$	–	$\mathcal{O}(\log n)$	[BCE+12, Thm 1]
$\mathcal{C}_1, \mathcal{C}_2$	–	$b \geq n$	random	$\mathcal{O}(\frac{b}{n} \cdot \log n)$	Thm 4.2
$\mathcal{C}_1, \mathcal{C}_2$	–	$b \in [n, n^3]$	random	$\mathcal{O}(\frac{b}{n} + \log n)$	Thm 5.1
$(1+\beta),\ \beta \leq 1 - \Omega(1)$	–	$b \geq 1$	–	$\Omega(\log n)$	Prop 7.3
Two-Choice, $(1+\beta), \beta = \Omega(1)$	–	$b \geq n \log n$	–	$\Omega(\frac{b}{n})$	Prop 7.4
Two-Choice	d-reg., conduct. Φ	–	–	$\mathcal{O}(\frac{\log n}{\Phi})$	[PTW15, Thm 3.2]
Two-Choice	d-reg., conduct. Φ	–	random	$\mathcal{O}(\frac{\log n}{\Phi})$	Thm 6.2
Two-Choice	d-reg., conduct. Φ	$b \geq n$	random	$\mathcal{O}(\frac{b}{n} \cdot \frac{\log n}{\Phi})$	Thm 6.3
Two-Choice	d-reg., conduct. Φ $\Phi = \Theta(1)$	$b \in [n, n^3]$	random	$\mathcal{O}(\frac{b}{n} + \log n)$	Thm 6.3
$(1+\beta),\ \beta \leq 1 - \Omega(1)$	–	–	–	$\Omega(\frac{\log n}{\beta})$	[PTW15, Sec 4]
$(1+\beta)$	–	–	random	$\mathcal{O}(\frac{\log n}{\beta} + \frac{\log(1/\beta)}{\beta})$	[PTW15, Cor 2.12]
$(1+\beta)$	–	–	random	$\mathcal{O}(\frac{\log n}{\beta})$	Thm 6.4

Improved on arxiv version: no dependence on d.

Appendix B: Outline for Tighter Bound

- By the refined analysis, for $\gamma = \Theta(n/b)$, for any $t \geq 0$, $\mathbf{E}\left[\, \Gamma^t \,\right] \leq cn$.
- Using the techniques in [LS22a], w.h.p. $\Gamma^s \leq cn$ for *all* $s \in [m - bn \log^5 n, m]$.
- Hence, the number of bins with normalized load $\Omega(b/n)$ is at most
$$cn \cdot e^{-\gamma \Omega(b/n)} \leq \delta n.$$
- Hence, by looking at the potential for constant $\widetilde{\gamma} > 0$ and with offset $\Omega(b/n)$,
$$\Lambda^t := \sum_{i : x_i^t \geq \frac{t}{n} + \Omega(b/n)} e^{\widetilde{\gamma} \cdot (x_i^t - \frac{t}{n} - \Omega(b/n))},$$
every bin i contributing to the potential has $p_i \leq \frac{1-\epsilon}{n}$, so
$$\mathbf{E}\left[\, \Lambda^{t+1} \mid \mathfrak{F}^t, \Gamma^t \leq cn \,\right] \leq \Lambda^t \cdot \left(1 - \frac{c_1 \widetilde{\gamma}}{n}\right) + c_2 \widetilde{\gamma}.$$
- By induction, this implies that $\mathbf{E}\left[\, \Lambda^m \,\right] = \mathcal{O}(n)$.
- And by Markov's inequality that w.h.p. $\mathrm{Gap}(m) = \mathcal{O}(\frac{b}{n} + \log n)$.

Appendix C: Drift Inequality Statement

Theorem (Corollary 3.2)

Consider any allocation process and probability vector p satisfying condition $\mathcal{C}_1$ for constant $\delta \in (0,1)$ and $\epsilon > 0$. Further assume that it satisfies for some $K > 0$ and some $R > 0$, for any $t \geq 0$,

$$\sum_{i=1}^{n} \mathbf{E}\left[\, \Delta\Phi_i^{t+1} \,\middle|\, \mathfrak{F}^t \,\right] \leq \sum_{i=1}^{n} \Phi_i^t \cdot \left(\left(p_i - \frac{1}{n} \right) \cdot \kappa \cdot \gamma + K \cdot R \cdot \frac{\gamma^2}{n} \right),$$

and

$$\sum_{i=1}^{n} \mathbf{E}\left[\, \Delta\Psi_i^{t+1} \,\middle|\, \mathfrak{F}^t \,\right] \leq \sum_{i=1}^{n} \Psi_i^t \cdot \left(\left(\frac{1}{n} - p_i \right) \cdot \kappa \cdot \gamma + K \cdot \kappa \cdot \frac{\gamma^2}{n} \right).$$

Then, there exists a constant $c := c(\delta) > 0$, such that for $\gamma \in \left(0, \min\left\{1, \frac{\epsilon\delta}{8K}\right\}\right)$

$$\mathbf{E}\left[\, \Delta\Gamma^{t+1} \,\middle|\, \mathfrak{F}^t \,\right] \leq -\Gamma^t \cdot R \cdot \frac{\gamma\epsilon\delta}{8n} + R \cdot c\gamma\epsilon,$$

and

$$\mathbf{E}\left[\, \Gamma^t \,\right] \leq \frac{8c}{\delta} \cdot n.$$

Appendix D: Proof Outline (I)

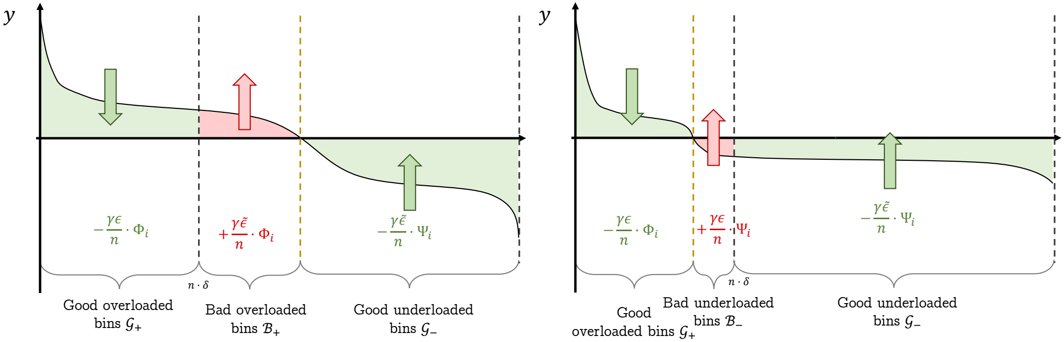

Figure: The two cases of bad bins in a configuration ($\mathcal{B}_+ \neq \emptyset$ or $\mathcal{B}_- \neq \emptyset$) and their *dominating terms* in $\Delta\overline{\Gamma}$ for each of the set of bins.

Appendix D: Proof Outline (II)

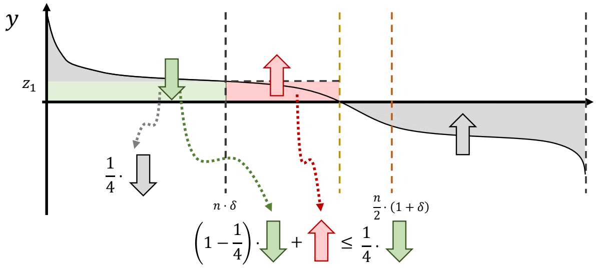

Figure: Case A $[|\mathcal{B}_+| \leq \frac{n}{2} \cdot (1-\delta)]$: The positive (increase) dominant term in the contribution of bins in $\mathcal{B}_+$ is counteracted by a fraction of the negative (decrease) dominant term of the good bins $\mathcal{G}_+$.

Appendix D: Proof Outline (III)

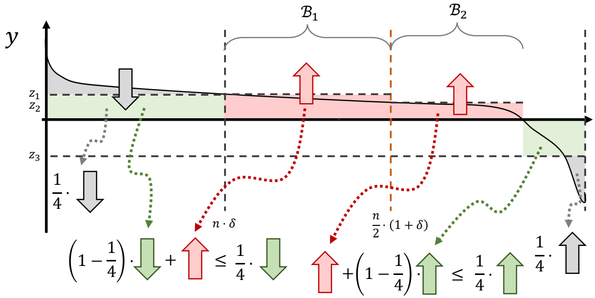

Figure: Case B $[|\mathcal{B}_+| > \frac{n}{2} \cdot (1-\delta)]$: The dominant increase of the bins in $\mathcal{B}_1$ is counteracted by a fraction of the decrease of the bins in $\mathcal{G}_+$ as in Case A. The dominant increase of the bins in $\mathcal{B}_2$ is counteracted by a fraction of the decrease of the bins in $\mathcal{G}_-$, when $z_2 = y_{n(1+\delta)/2}$ is sufficiently large.

Bibliography I

- Y. Azar, A. Z. Broder, A. R. Karlin, and E. Upfal, *Balanced allocations*, SIAM J. Comput. **29** (1999), no. 1, 180–200. MR 1710347
- P. Berenbrink, A. Czumaj, M. Englert, T. Friedetzky, and L. Nagel, *Multiple-choice balanced allocation in (almost) parallel*, 16th International Workshop on Randomization and Computation (RANDOM'12), Springer-Verlag, 2012, pp. 411–422.
- P. Berenbrink, A. Czumaj, A. Steger, and B. Vöcking, *Balanced allocations: the heavily loaded case*, SIAM J. Comput. **35** (2006), no. 6, 1350–1385. MR 2217150
- G. H. Gonnet, *Expected length of the longest probe sequence in hash code searching*, J. Assoc. Comput. Mach. **28** (1981), no. 2, 289–304. MR 612082
- R. M. Karp, M. Luby, and F. Meyer auf der Heide, *Efficient PRAM simulation on a distributed memory machine*, Algorithmica **16** (1996), no. 4-5, 517–542. MR 1407587

Bibliography II

- D. Los and T. Sauerwald, *Balanced Allocations with Incomplete Information: The Power of Two Queries*, 13th Innovations in Theoretical Computer Science Conference (ITCS'22), vol. 215, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022, pp. 103:1–103:23.
- ______, *Balanced allocations with the choice of noise*, 41st Annual ACM-SIGOPT Principles of Distributed Computing (PODC'22), ACM, 2022, p. 164–175.
- M. Mitzenmacher, *The power of two choices in randomized load balancing*, Ph.D. thesis, University of California at Berkeley, 1996.
- Y. Peres, K. Talwar, and U. Wieder, *Graphical balanced allocations and the* $(1+\beta)$*-choice process*, Random Structures & Algorithms **47** (2015), no. 4, 760–775. MR 3418914
- M. Raab and A. Steger, *"Balls into bins"—a simple and tight analysis*, 2nd International Workshop on Randomization and Computation (RANDOM'98), vol. 1518, Springer, 1998, pp. 159–170. MR 1729169