Balanced Allocations in Batches: The Tower of Two Choices

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.

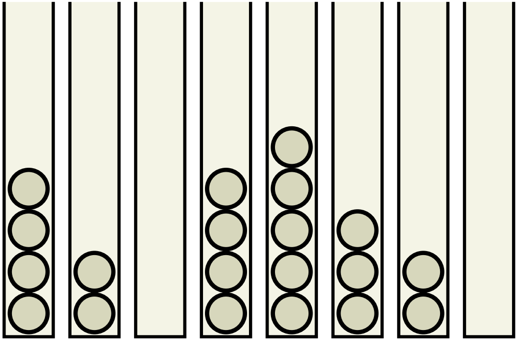

Balanced allocations setting

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.

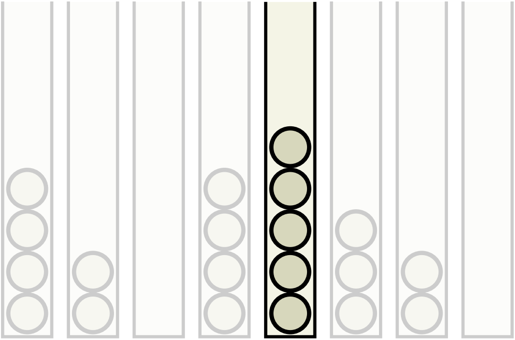

Balanced allocations setting

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 $\text{Gap}(m) = \max_{i \in [n]}(x_i^m - m/n)$.

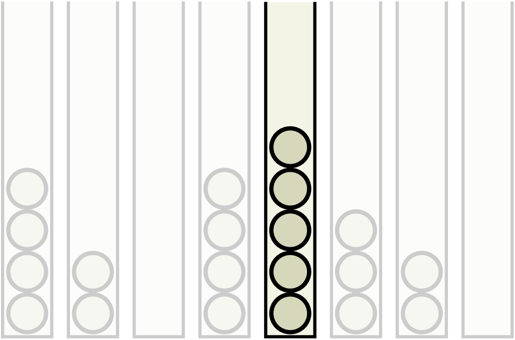

Balanced allocations setting

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 $\mathrm{Gap}(m) = \max_{i \in [n]} (x_i^m - m/n)$.

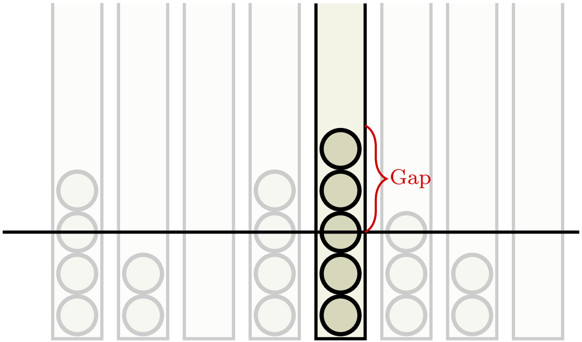

Balanced allocations setting

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 $\text{Gap}(m) = \max_{i \in [n]}(x_i^m - m/n)$.

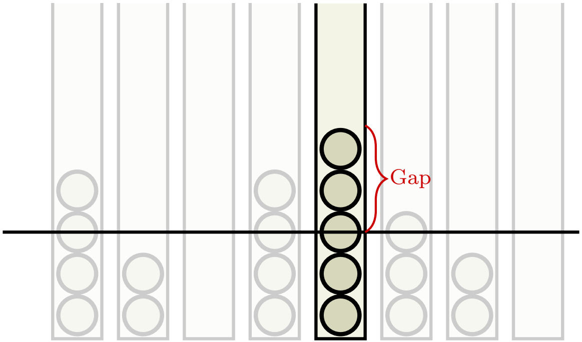

- Applications in hashing, load balancing and routing.

One-Choice and Two-Choice processes

One-Choice Process:
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. $\mathrm{Gap}(n) = \Theta\left(\frac{\log n}{\log \log n}\right)$ [Gon81].

One-Choice and Two-Choice processes

One-Choice Process:
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].

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. $\mathrm{Gap}(n) = \Theta\left(\frac{\log n}{\log \log n}\right)$ [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]).

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

ONE-CHOICE Process:
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]).

TWO-CHOICE Process:
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

One-Choice Process:
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]).

Two-Choice Process:
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].

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].

ACM Theory and Practice Award

ACM Theory and Practice Award

The practical significance of the "power of two choices" was recognised in the 2020 ACM Paris Kanellakis award [ABK+20]:

ACM Theory and Practice Award

The practical significance of the “power of two choices” was recognised in the 2020 ACM Paris Kanellakis award [ABK$^+$20]:

> *“[...] it is not surprising that the power of two choices that requires only a local decision rather than global coordination has led to a wide range of practical applications. These include i-Google’s web index, Akamai’s overlay routing network, and highly reliable distributed data storage systems used by Microsoft and Dropbox, which are all based on variants of the power of two choices paradigm.”*

ACM Theory and Practice Award

The practical significance of the "power of two choices" was recognised in the 2020 ACM Paris Kanellakis award [ABK+20]:

> *"[...] it is not surprising that the power of two choices that requires only a local decision rather than global coordination has led to a wide range of practical applications. These include i-Google's web index, Akamai's overlay routing network, and highly reliable distributed data storage systems used by Microsoft and Dropbox, which are all based on **variants** of the power of two choices paradigm."*

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

$(1+\beta)$-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.
- In the heavily-loaded case, [PTW15] proved that the gap is w.h.p. $\mathcal{O}\left(\frac{\log n}{\beta}\right)$ for any $\beta \in (0,1]$.
- It has been used to analyze

$(1+\beta)$-Process

<u>$(1+\beta)$-Process:</u>
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.
- In the heavily-loaded case, [PTW15] proved that the gap is w.h.p. $\mathcal{O}\left(\frac{\log n}{\beta}\right)$ for any $\beta \in (0,1]$.
- It has been used to analyze *population protocols* [AAG18, AGR21],

$(1+\beta)$-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.
- In the heavily-loaded case, [PTW15] proved that the gap is w.h.p. $\mathcal{O}\big(\frac{\log n}{\beta}\big)$ for any $\beta \in (0,1]$.
- It has been used to analyze *population protocols* [AAG18, AGR21], *distributed data structures* [ABK$^+$18, AKLN17, Nad21]

$(1+\beta)$-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.
- In the heavily-loaded case, [PTW15] proved that the gap is w.h.p. $\mathcal{O}\left(\frac{\log n}{\beta}\right)$ for any $\beta \in (0,1]$.
- It has been used to analyze *population protocols* [AAG18, AGR21], *distributed data structures* [ABK$^+$18, AKLN17, Nad21] and *online carpooling* [GKKS20].

$(1+\beta)$-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.
- In the heavily-loaded case, [PTW15] proved that the gap is w.h.p. $\mathcal{O}\left(\frac{\log n}{\beta}\right)$ for any $\beta \in (0,1]$.
- It has been used to analyze *population protocols* [AAG18, AGR21], *distributed data structures* [ABK$^+$18, AKLN17, Nad21] and *online carpooling* [GKKS20].

Question: Why choose a $\beta < 1$?

Settings

Two-Choice with outdated information

TWO-CHOICE **with outdated information**

- Mitzenmacher [Mit00] and Dahlin [Dah00] empirically observed that TWO-CHOICE does not perform well with outdated information.

TWO-CHOICE with outdated information

- Mitzenmacher [Mit00] and Dahlin [Dah00] empirically observed that TWO-CHOICE does not perform well with outdated information.

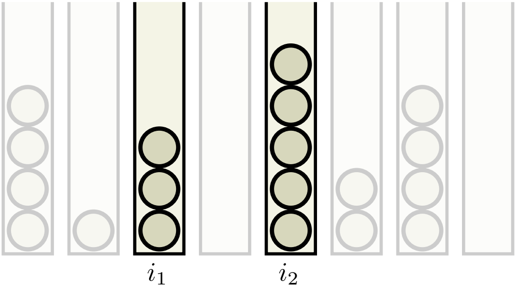

Two-Choice with outdated information

- Mitzenmacher [Mit00] and Dahlin [Dah00] empirically observed that Two-Choice does not perform well with outdated information.

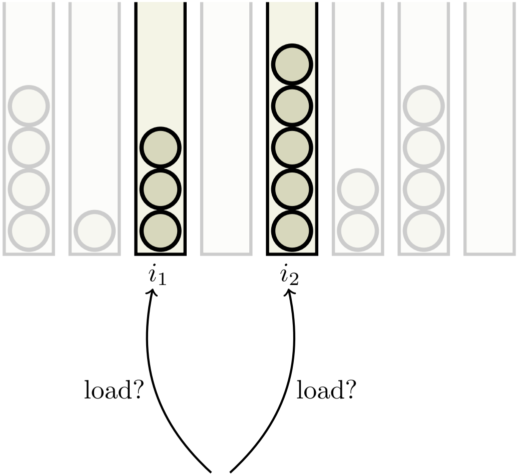

Two-Choice with outdated information

- Mitzenmacher [Mit00] and Dahlin [Dah00] empirically observed that Two-Choice does not perform well with outdated information.

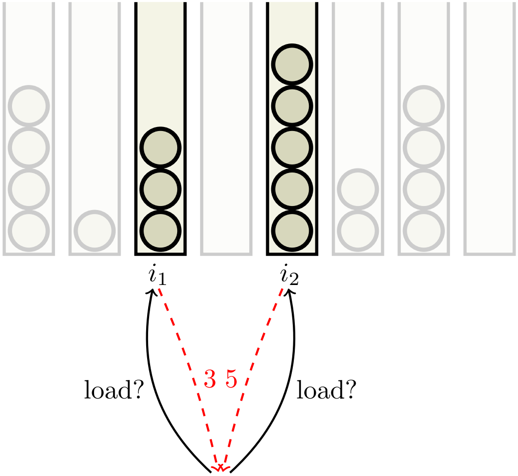

TWO-CHOICE with outdated information

- Mitzenmacher [Mit00] and Dahlin [Dah00] empirically observed that TWO-CHOICE does not perform well with outdated information.

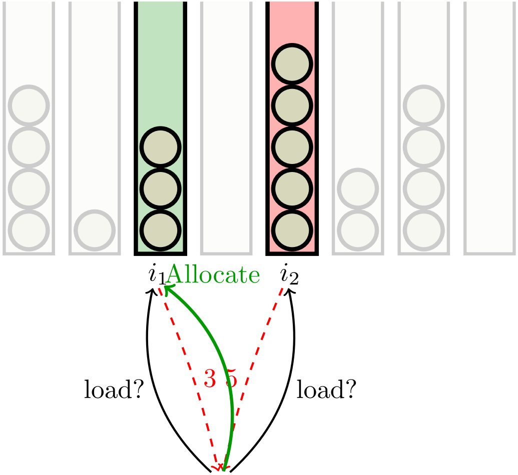

Two-Choice with outdated information

- Mitzenmacher [Mit00] and Dahlin [Dah00] empirically observed that Two-Choice does not perform well with outdated information. ⇝ herd phenomenon

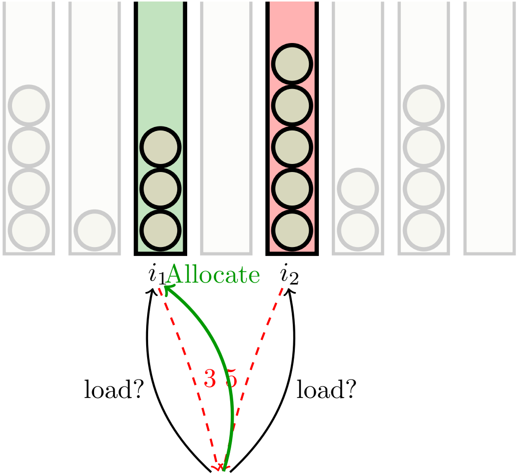

Two-Choice with outdated information

- Mitzenmacher [Mit00] and Dahlin [Dah00] empirically observed that Two-Choice does not perform well with outdated information. $\rightsquigarrow$ herd phenomenon
- Several *low-latency schedulers* use variants of Two-Choice (Eagle [DDDZ16], Hawk [DDKZ15], Peacock [KG18]).

Two-Choice with outdated information

- Mitzenmacher [Mit00] and Dahlin [Dah00] empirically observed that Two-Choice does not perform well with outdated information. $\rightsquigarrow$ herd phenomenon
- Several *low-latency schedulers* use variants of Two-Choice (Eagle [DDDZ16], Hawk [DDKZ15], Peacock [KG18]). Sparrow [OWZS13] remarks

Two-Choice with outdated information

- Mitzenmacher [Mit00] and Dahlin [Dah00] empirically observed that Two-Choice does not perform well with outdated information. $\rightsquigarrow$ herd phenomenon
- Several *low-latency schedulers* use variants of Two-Choice (Eagle [DDDZ16], Hawk [DDKZ15], Peacock [KG18]). Sparrow [OWZS13] remarks

> *The power of two choices suffers from two remaining performance problems: first, server queue length is a poor indicator of wait time, and second, due to messaging delays, multiple schedulers sampling in parallel may experience race conditions.*

Two-Choice with outdated information

- Mitzenmacher [Mit00] and Dahlin [Dah00] empirically observed that Two-Choice does not perform well with outdated information. $\rightsquigarrow$ herd phenomenon
- Several *low-latency schedulers* use variants of Two-Choice (Eagle [DDDZ16], Hawk [DDKZ15], Peacock [KG18]). Sparrow [OWZS13] remarks

> *The power of two choices suffers from two remaining performance problems: first, server queue length is a poor indicator of wait time, and second, due to messaging delays, multiple schedulers sampling in parallel may experience race conditions.*

- In the *queuing setting*, Whitt [Whi86] remarks:

Two-Choice with outdated information

- Mitzenmacher [Mit00] and Dahlin [Dah00] empirically observed that Two-Choice does not perform well with outdated information. $\rightsquigarrow$ herd phenomenon
- Several *low-latency schedulers* use variants of Two-Choice (Eagle [DDDZ16], Hawk [DDKZ15], Peacock [KG18]). Sparrow [OWZS13] remarks

> *The power of two choices suffers from two remaining performance problems: first, server queue length is a poor indicator of wait time, and second, due to messaging delays, multiple schedulers sampling in parallel may experience race conditions.*

- In the *queuing setting*, Whitt [Whi86] remarks:

> *We have shown that several natural selection rules are not optimal in various situations, but we have not identified any optimal rules. Identifying optimal rules in these situations would obviously be interesting, but appears to be difficult.*

b-Batched **Setting**

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).

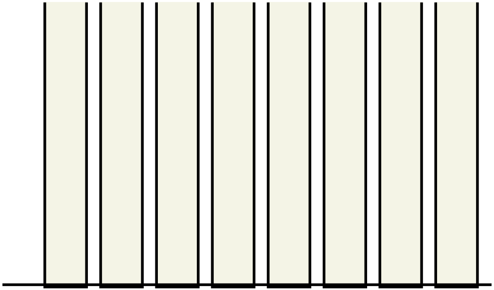

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).

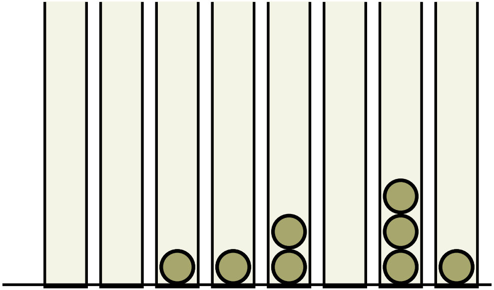

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).

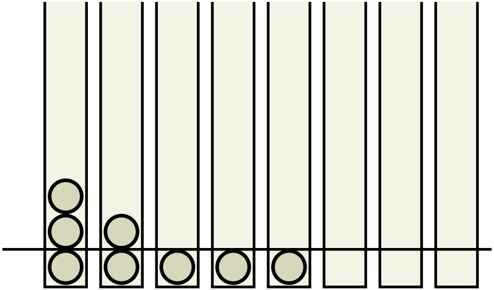

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).

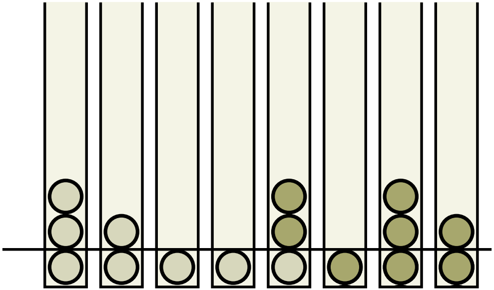

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).

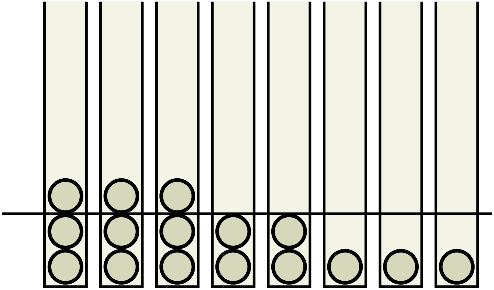

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).

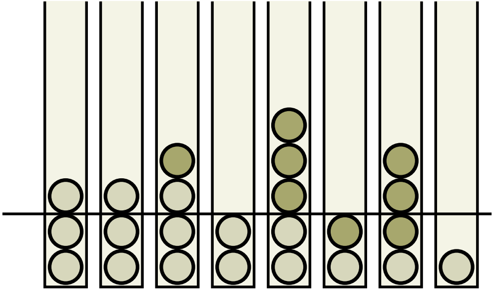

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).

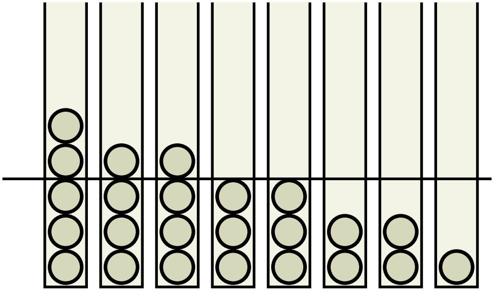

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).

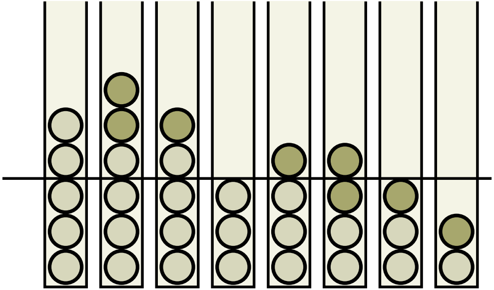

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).

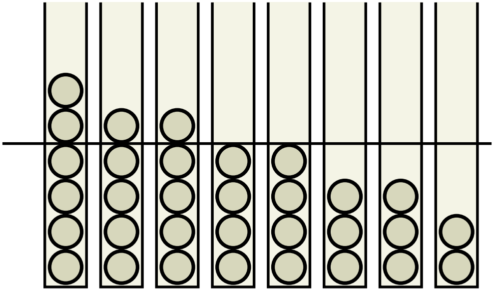

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).

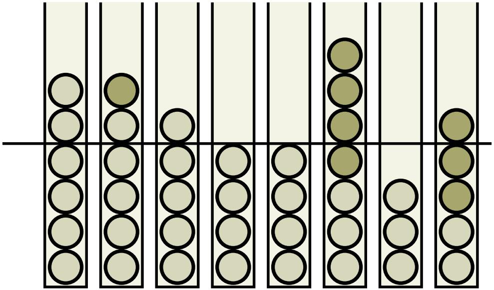

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).

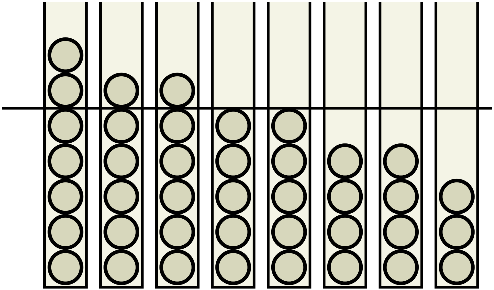

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE⁺12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).
- For $b = n$, they showed that w.h.p. $\text{Gap}(m) = \mathcal{O}(\log n)$.

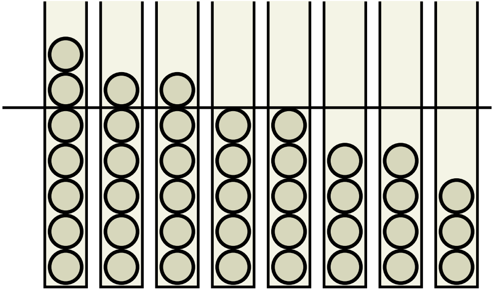

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).
- For $b = n$, they showed that w.h.p. $\text{Gap}(m) = \mathcal{O}(\log n)$.
- In [LS22c] bound was improved to $\Theta(\log n / \log \log n)$.

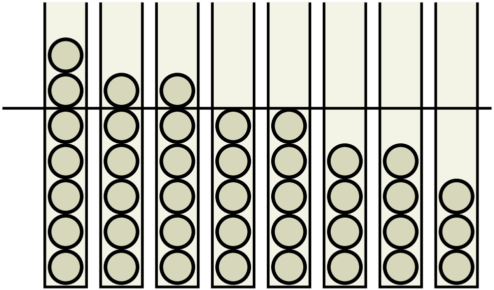

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE$^+$12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).
- For $b = n$, they showed that w.h.p. $\text{Gap}(m) = \mathcal{O}(\log n)$.
- In [LS22c] bound was improved to $\Theta(\log n / \log \log n)$. And more generally, for $b \in \left[\frac{n}{\text{polylog}(n)}, n \log n\right]$ it follows One-Choice with b balls.

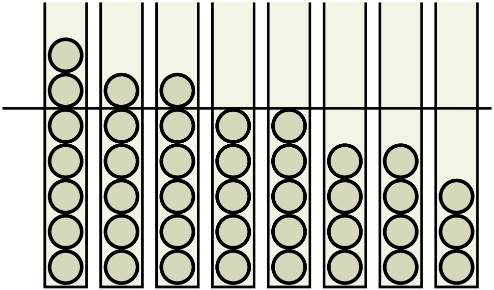

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE$^+$12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).
- For $b = n$, they showed that w.h.p. $\mathrm{Gap}(m) = \mathcal{O}(\log n)$.
- In [LS22c] bound was improved to $\Theta(\log n / \log\log n)$. And more generally, for $b \in \left[\frac{n}{\mathrm{polylog}(n)}, n \log n\right]$ it follows One-Choice with b balls.
- In [LS22a], for $b \in [n \log n, \mathrm{poly}(n)]$, w.h.p. $\mathrm{Gap}(m) = \Theta(b/n)$.

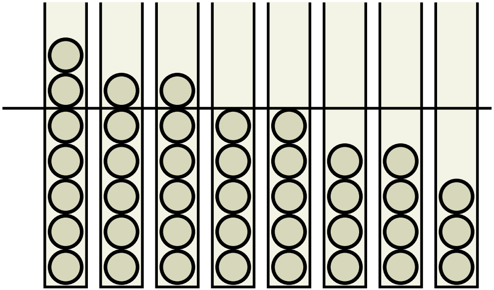

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE+12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).
- For $b = n$, they showed that w.h.p. $\mathrm{Gap}(m) = \mathcal{O}(\log n)$.
- In [LS22c] bound was improved to $\Theta(\log n / \log\log n)$. And more generally, for $b \in \left[\frac{n}{\mathrm{polylog}(n)}, n \log n\right]$ it follows One-Choice with b balls.
- In [LS22a], for $b \in [n \log n, \mathrm{poly}(n)]$, w.h.p. $\mathrm{Gap}(m) = \Theta(b/n)$.

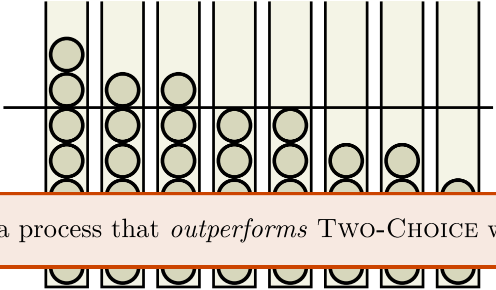

Question: Is there a process that *outperforms* Two-Choice when $b \geq n \log n$?

b-Batched Setting

- Berenbrink, Czumaj, Englert, Friedetzky and Nagel [BCE$^+$12] studied Two-Choice where balls are allocated in *batches* of size b (b-Batched setting).
- For $b = n$, they showed that w.h.p. $\mathrm{Gap}(m) = \mathcal{O}(\log n)$.
- In [LS22c] bound was improved to $\Theta(\log n / \log\log n)$. And more generally, for $b \in \left[\frac{n}{\mathrm{polylog}(n)}, n \log n\right]$ it follows One-Choice with b balls.
- In [LS22a], for $b \in [n \log n, \mathrm{poly}(n)]$, w.h.p. $\mathrm{Gap}(m) = \Theta(b/n)$.

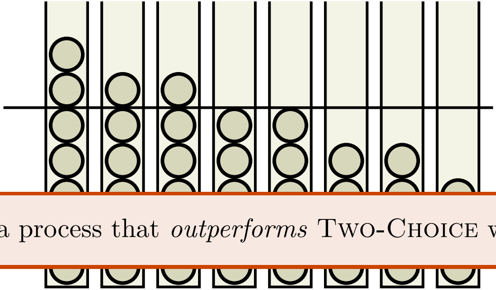

Question: Is there a process that *outperforms* Two-Choice when $b \geq n \log n$?

Open in Visualiser.

Our result

Our result

Our result

In the b-BATCHED setting with any $b \geq n \log n$,

Our result

In the b-Batched setting with any $b \geq n \log n$, the $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$

Our result

In the b-Batched setting with any $b \geq n \log n$, the $(1 + \beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\text{Gap}(m) = \mathcal{O}\big(\sqrt{(b/n) \cdot \log n}\big)$.

Our result

In the b-Batched setting with any $b \geq n \log n$, the $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\mathrm{Gap}(m) = \mathcal{O}\big(\sqrt{(b/n) \cdot \log n}\big)$.

- This is almost a *quadratic improvement* over the $\Theta(b/n)$ gap of Two-Choice.

Our result

In the b-Batched setting with any $b \geq n \log n$, the $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\mathrm{Gap}(m) = \mathcal{O}\big(\sqrt{(b/n) \cdot \log n}\big)$.

- This is almost a *quadratic improvement* over the $\Theta(b/n)$ gap of Two-Choice.
- Mixing One-Choice and Two-Choice results in a process that is *better* than both.

Our result

In the b-Batched setting with any $b \geq n \log n$, the $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\mathrm{Gap}(m) = \mathcal{O}\big(\sqrt{(b/n) \cdot \log n}\big)$.

- This is almost a *quadratic improvement* over the $\Theta(b/n)$ gap of Two-Choice.
- Mixing One-Choice and Two-Choice results in a process that is *better* than both.
- Asymptotically optimal over all processes that take at most a *constant* number of samples for each allocation.

Our result

In the b-Batched setting with any $b \geq n \log n$, the $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\mathrm{Gap}(m) = \mathcal{O}\left(\sqrt{(b/n) \cdot \log n}\right)$.

- This is almost a *quadratic improvement* over the $\Theta(b/n)$ gap of Two-Choice.
- Mixing One-Choice and Two-Choice results in a process that is *better* than both.
- Asymptotically optimal over all processes that take at most a *constant* number of samples for each allocation.
- The upper bound holds in the presence of *weights* and for a more *general* family of processes.

Our result

In the b-Batched setting with any $b \geq n \log n$, the $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\mathrm{Gap}(m) = \mathcal{O}\left(\sqrt{(b/n) \cdot \log n}\right)$.

- This is almost a *quadratic improvement* over the $\Theta(b/n)$ gap of Two-Choice.
- Mixing One-Choice and Two-Choice results in a process that is *better* than both.
- Asymptotically optimal over all processes that take at most a *constant* number of samples for each allocation.
- The upper bound holds in the presence of *weights* and for a more *general* family of processes.
- *Easy* to implement

Our result

In the b-Batched setting with any $b \geq n \log n$, the $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\mathrm{Gap}(m) = \mathcal{O}\big(\sqrt{(b/n) \cdot \log n}\big)$.

- This is almost a *quadratic improvement* over the $\Theta(b/n)$ gap of Two-Choice.
- Mixing One-Choice and Two-Choice results in a process that is *better* than both.
- Asymptotically optimal over all processes that take at most a *constant* number of samples for each allocation.
- The upper bound holds in the presence of *weights* and for a more *general* family of processes.
- *Easy* to implement (≤ 5 lines in `nginx`, `HAProxy`, `Finagle`).

Our result

In the b-Batched setting with any $b \geq n \log n$, the $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\mathrm{Gap}(m) = \mathcal{O}\big(\sqrt{(b/n) \cdot \log n}\big)$.

- This is almost a *quadratic improvement* over the $\Theta(b/n)$ gap of Two-Choice.
- Mixing One-Choice and Two-Choice results in a process that is *better* than both.
- Asymptotically optimal over all processes that take at most a *constant* number of samples for each allocation.
- The upper bound holds in the presence of *weights* and for a more *general* family of processes.
- *Easy* to implement ($\leq$ 5 lines in `nginx`, `HAProxy`, `Finagle`).

    $\rightsquigarrow$ serve $\approx 30\%$ of websites.

Intuition

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,

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

- For Two-Choice,

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

- For $(1+\beta)$-process,

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

Intuition: The tower of Two-Choice

Intuition: The tower of Two-Choice

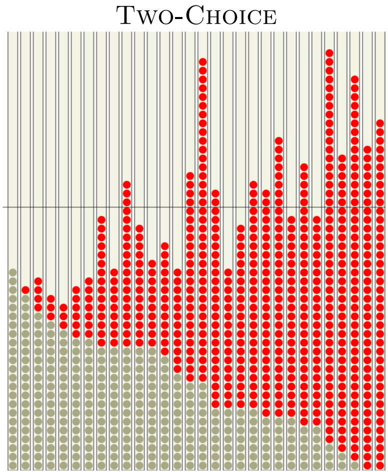

$$p_i = \frac{2i-1}{n^2}$$

Intuition: The tower of Two-Choice

Two-Choice

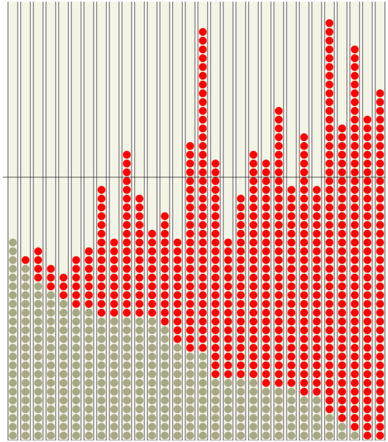

$$p_i = \frac{2i-1}{n^2}$$

$(1+\beta)$-Process

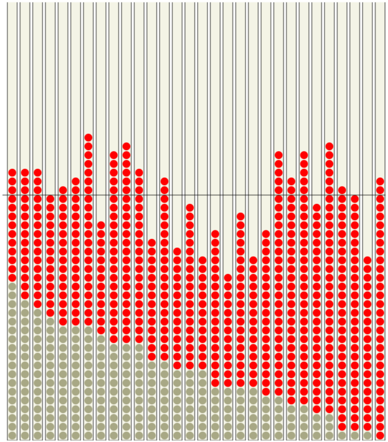

$$p_i = \beta \cdot \frac{2i-1}{n^2} + (1-\beta) \cdot \frac{1}{n}$$

Intuition: The tower of Two-Choice

Two-Choice

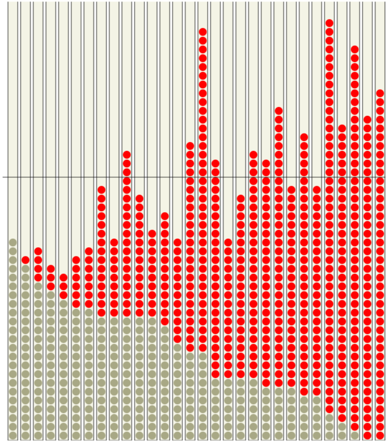

$$p_i = \frac{2i - 1}{n^2}$$

$(1 + \beta)$-Process

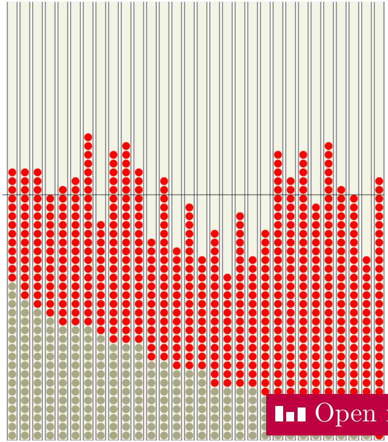

Open in Visualiser.

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

Empirical results for different processes

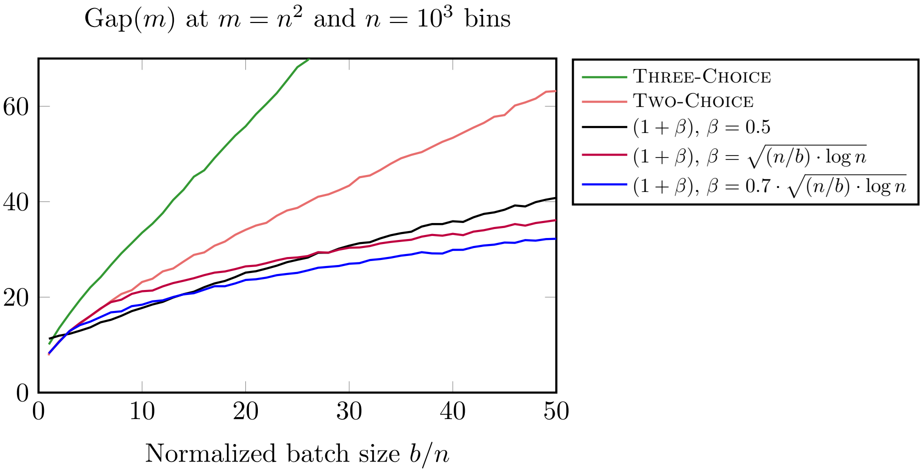

Empirical results for different processes

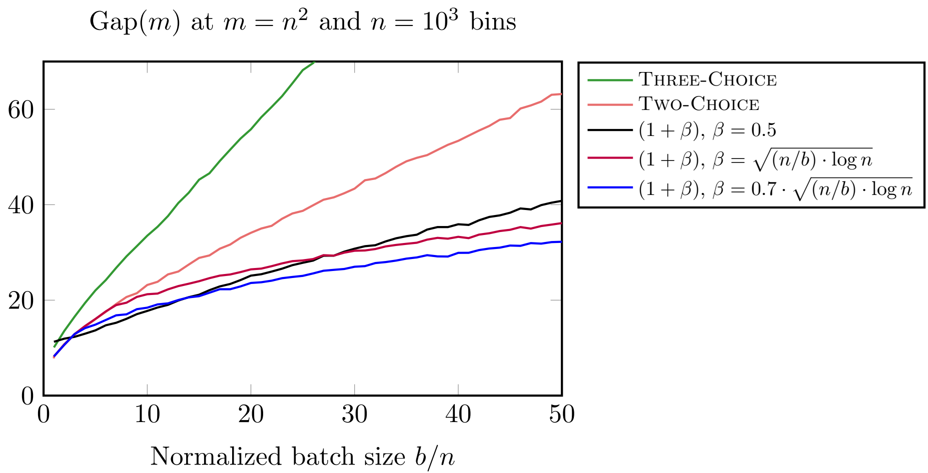

- The gaps are decreasingly ordered by p_n: $\approx \frac{3}{n}$ (for Three-Choice),

Empirical results for different processes

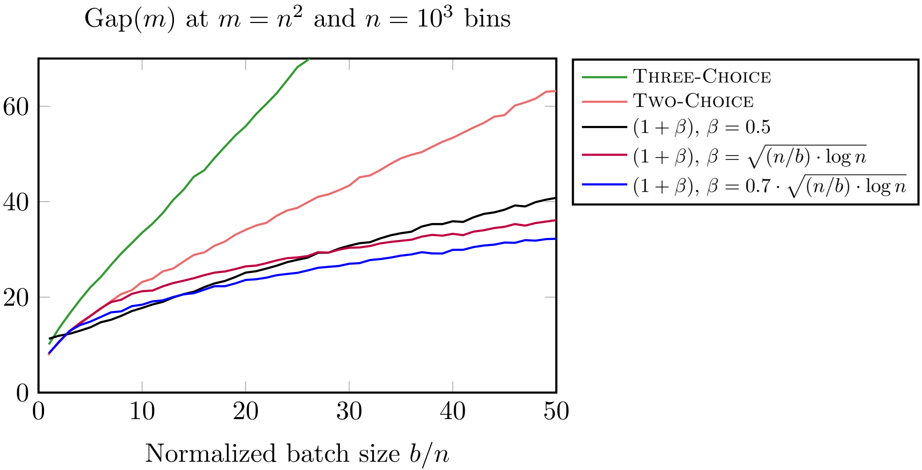

- The gaps are decreasingly ordered by p_n: $\approx \frac{3}{n}$ (for THREE-CHOICE), $\approx \frac{2}{n}$ (for TWO-CHOICE)

Empirical results for different processes

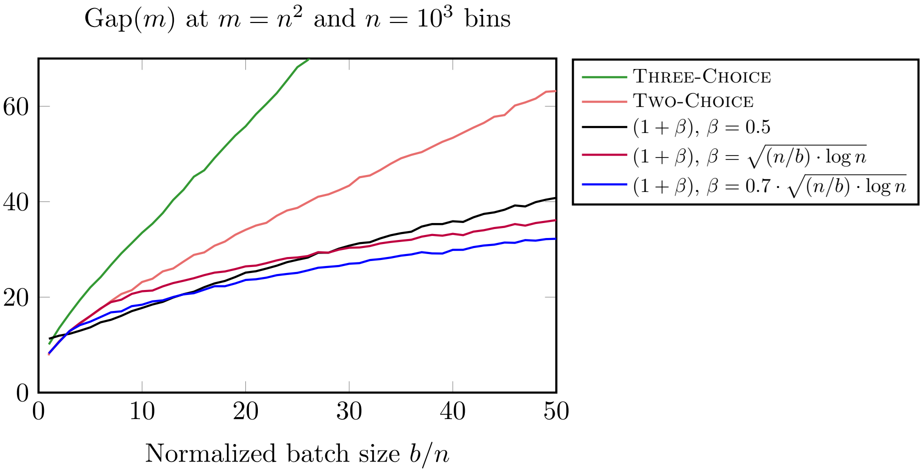

- The gaps are decreasingly ordered by p_n: $\approx \frac{3}{n}$ (for Three-Choice), $\approx \frac{2}{n}$ (for Two-Choice) and $\approx \frac{1+\beta}{n}$ (for the $(1+\beta)$-processes).

Empirical results for different β's

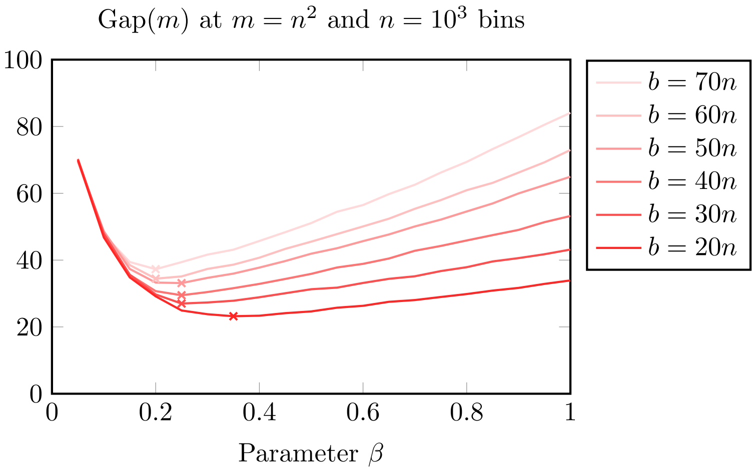

Analysis

Conditions on probability allocation vectors

Conditions on probability allocation vectors

- **Condition $\mathcal{C}_1$:** [PTW15] analyzed processes with (i) p being *non-decreasing*

Conditions on probability allocation vectors

- **Condition $\mathcal{C}_1$:** [PTW15] analyzed processes with (i) p being *non-decreasing* and (ii) which for some constant $\delta \in (0,1)$, satisfy

$$p_{\delta n} \leq \frac{1-\epsilon}{n} \qquad \left(\text{and } p_{\delta n+1} \geq \frac{1+\Omega(\epsilon)}{n} \right).$$

Conditions on probability allocation vectors

- **Condition $\mathcal{C}_1$:** [PTW15] analyzed processes with (i) p being *non-decreasing* and (ii) which for some constant $\delta \in (0,1)$, satisfy
$$p_{\delta n} \leq \frac{1-\epsilon}{n} \qquad \left(\text{and } p_{\delta n+1} \geq \frac{1+\Omega(\epsilon)}{n} \right).$$
- Proved that such processes achieve w.h.p. an $\mathcal{O}\left(\frac{\log n}{\epsilon}\right)$ gap (sequential).

Conditions on probability allocation vectors

- **Condition $\mathcal{C}_1$:** [PTW15] analyzed processes with (i) p being *non-decreasing* and (ii) which for some constant $\delta \in (0,1)$, satisfy
$$p_{\delta n} \leq \frac{1-\epsilon}{n} \qquad \left(\text{ and } p_{\delta n+1} \geq \frac{1+\Omega(\epsilon)}{n}\right).$$
- Proved that such processes achieve w.h.p. an $\mathcal{O}\left(\frac{\log n}{\epsilon}\right)$ gap (sequential).
- **Condition $\mathcal{C}_2$:**

Conditions on probability allocation vectors

- **Condition $\mathcal{C}_1$:** [PTW15] analyzed processes with (i) p being *non-decreasing* and (ii) which for some constant $\delta \in (0,1)$, satisfy

$$p_{\delta n} \leq \frac{1-\epsilon}{n} \qquad \left(\text{and } p_{\delta n+1} \geq \frac{1+\Omega(\epsilon)}{n} \right).$$

- Proved that such processes achieve w.h.p. an $\mathcal{O}\left(\frac{\log n}{\epsilon}\right)$ gap (sequential).
- **Condition $\mathcal{C}_2$:** There exists a $C > 1$, such that

$$\max_{i \in [n]} p_i^t \leq \frac{C}{n}.$$

Conditions on probability allocation vectors

- **Condition $\mathcal{C}_1$:** [PTW15] analyzed processes with (i) p being *non-decreasing* and (ii) which for some constant $\delta \in (0,1)$, satisfy
$$p_{\delta n} \leq \frac{1-\epsilon}{n} \qquad \left(\text{and } p_{\delta n+1} \geq \frac{1+\Omega(\epsilon)}{n} \right).$$
- Proved that such processes achieve w.h.p. an $\mathcal{O}\big(\frac{\log n}{\epsilon}\big)$ gap (sequential).
- **Condition $\mathcal{C}_2$:** There exists a $C > 1$, such that
$$\max_{i \in [n]} p_i^t \leq \frac{C}{n}.$$

Two-Choice satisfies $\mathcal{C}_2$ with $C = 2$

Conditions on probability allocation vectors

- **Condition $\mathcal{C}_1$:** [PTW15] analyzed processes with (i) p being *non-decreasing* and (ii) which for some constant $\delta \in (0,1)$, satisfy
$$p_{\delta n} \leq \frac{1-\epsilon}{n} \qquad \left(\text{and } p_{\delta n+1} \geq \frac{1+\Omega(\epsilon)}{n} \right).$$
- Proved that such processes achieve w.h.p. an $\mathcal{O}\left(\frac{\log n}{\epsilon}\right)$ gap (sequential).
- **Condition $\mathcal{C}_2$:** There exists a $C > 1$, such that
$$\max_{i \in [n]} p_i^t \leq \frac{C}{n}.$$
Two-Choice satisfies $\mathcal{C}_2$ with $C = 2$ and $(1+\beta)$-process for $C = 1 + \beta$.

Conditions on probability allocation vectors

- **Condition $\mathcal{C}_1$:** [PTW15] analyzed processes with (i) p being *non-decreasing* and (ii) which for some constant $\delta \in (0,1)$, satisfy
$$p_{\delta n} \leq \frac{1-\epsilon}{n} \qquad \left(\text{and } p_{\delta n+1} \geq \frac{1+\Omega(\epsilon)}{n} \right).$$
- Proved that such processes achieve w.h.p. an $\mathcal{O}\left(\frac{\log n}{\epsilon}\right)$ gap (sequential).
- **Condition $\mathcal{C}_2$:** There exists a $C > 1$, such that
$$\max_{i \in [n]} p_i^t \leq \frac{C}{n}.$$
Two-Choice satisfies $\mathcal{C}_2$ with $C = 2$ and $(1+\beta)$-process for $C = 1 + \beta$.
- Conditions $\mathcal{C}_1$ (const ϵ) and $\mathcal{C}_2$ sufficient to prove $\text{Gap}(m) = \mathcal{O}(b/n)$ (b-Batched).

Conditions on probability allocation vectors

- **Condition $\mathcal{C}_1$:** [PTW15] analyzed processes with (i) p being *non-decreasing* and (ii) which for some constant $\delta \in (0,1)$, satisfy
$$p_{\delta n} \leq \frac{1-\epsilon}{n} \qquad \left(\text{and } p_{\delta n+1} \geq \frac{1+\Omega(\epsilon)}{n} \right).$$
- Proved that such processes achieve w.h.p. an $\mathcal{O}\left(\frac{\log n}{\epsilon}\right)$ gap (sequential).
- **Condition $\mathcal{C}_2$:** There exists a $C > 1$, such that
$$\max_{i \in [n]} p_i^t \leq \frac{C}{n}.$$
Two-Choice satisfies $\mathcal{C}_2$ with $C = 2$ and $(1+\beta)$-process for $C = 1 + \beta$.
- Conditions $\mathcal{C}_1$ (const ϵ) and $\mathcal{C}_2$ sufficient to prove $\text{Gap}(m) = \mathcal{O}(b/n)$ (b-Batched).
- **Condition $\mathcal{C}_3$:**

Conditions on probability allocation vectors

- **Condition $\mathcal{C}_1$:** [PTW15] analyzed processes with (i) p being *non-decreasing* and (ii) which for some constant $\delta \in (0,1)$, satisfy

$$p_{\delta n} \leq \frac{1-\epsilon}{n} \qquad \left(\text{and } p_{\delta n+1} \geq \frac{1+\Omega(\epsilon)}{n} \right).$$

- Proved that such processes achieve w.h.p. an $\mathcal{O}\left(\frac{\log n}{\epsilon}\right)$ gap (sequential).
- **Condition $\mathcal{C}_2$**: There exists a $C > 1$, such that

$$\max_{i \in [n]} p_i^t \leq \frac{C}{n}.$$

Two-Choice satisfies $\mathcal{C}_2$ with $C = 2$ and $(1+\beta)$-process for $C = 1+\beta$.
- Conditions $\mathcal{C}_1$ (const ϵ) and $\mathcal{C}_2$ sufficient to prove $\mathrm{Gap}(m) = \mathcal{O}(b/n)$ (b-Batched).
- **Condition $\mathcal{C}_3$**: There exists a $C > 1$, such that

$$\max_{i \in [n]} \left| p_i^t - \frac{1}{n} \right| \leq \frac{C-1}{n}.$$

Conditions on probability allocation vectors

- **Condition $\mathcal{C}_1$:** [PTW15] analyzed processes with (i) p being *non-decreasing* and (ii) which for some constant $\delta \in (0,1)$, satisfy

$$p_{\delta n} \leq \frac{1-\epsilon}{n} \qquad \left(\text{and } p_{\delta n+1} \geq \frac{1+\Omega(\epsilon)}{n} \right).$$

- Proved that such processes achieve w.h.p. an $\mathcal{O}\left(\frac{\log n}{\epsilon}\right)$ gap (sequential).
- **Condition $\mathcal{C}_2$:** There exists a $C > 1$, such that

$$\max_{i \in [n]} p_i^t \leq \frac{C}{n}.$$

Two-Choice satisfies $\mathcal{C}_2$ with $C = 2$ and $(1+\beta)$-process for $C = 1 + \beta$.
- Conditions $\mathcal{C}_1$ (const ϵ) and $\mathcal{C}_2$ sufficient to prove $\text{Gap}(m) = \mathcal{O}(b/n)$ (b-Batched).
- **Condition $\mathcal{C}_3$:** There exists a $C > 1$, such that

$$\max_{i \in [n]} \left| p_i^t - \frac{1}{n} \right| \leq \frac{C-1}{n}.$$

Condition $\mathcal{C}_3$ implies condition $\mathcal{C}_2$.

Conditions on probability allocation vectors

- **Condition $\mathcal{C}_1$:** [PTW15] analyzed processes with (i) p being *non-decreasing* and (ii) which for some constant $\delta \in (0,1)$, satisfy

$$p_{\delta n} \leq \frac{1-\epsilon}{n} \qquad \left(\text{and } p_{\delta n+1} \geq \frac{1+\Omega(\epsilon)}{n} \right).$$

- Proved that such processes achieve w.h.p. an $\mathcal{O}\big(\frac{\log n}{\epsilon}\big)$ gap (sequential).
- **Condition $\mathcal{C}_2$:** There exists a $C > 1$, such that

$$\max_{i \in [n]} p_i^t \leq \frac{C}{n}.$$

Two-Choice satisfies $\mathcal{C}_2$ with $C = 2$ and $(1+\beta)$-process for $C = 1 + \beta$.
- Conditions $\mathcal{C}_1$ (const ϵ) and $\mathcal{C}_2$ sufficient to prove $\text{Gap}(m) = \mathcal{O}(b/n)$ (b-Batched).
- **Condition $\mathcal{C}_3$:** There exists a $C > 1$, such that

$$\max_{i \in [n]} \left| p_i^t - \frac{1}{n} \right| \leq \frac{C-1}{n}.$$

Condition $\mathcal{C}_3$ implies condition $\mathcal{C}_2$. $(1+\beta)$-process satisfies $\mathcal{C}_3$ for $C = 1 + \beta$.

Conditions on probability allocation vectors

- **Condition $\mathcal{C}_1$:** [PTW15] analyzed processes with (i) p being *non-decreasing* and (ii) which for some constant $\delta \in (0,1)$, satisfy
$$p_{\delta n} \leq \frac{1-\epsilon}{n} \qquad \left(\text{and } p_{\delta n+1} \geq \frac{1+\Omega(\epsilon)}{n} \right).$$
- Proved that such processes achieve w.h.p. an $\mathcal{O}\left(\frac{\log n}{\epsilon}\right)$ gap (sequential).
- **Condition $\mathcal{C}_2$:** There exists a $C > 1$, such that
$$\max_{i \in [n]} p_i^t \leq \frac{C}{n}.$$
Two-Choice satisfies $\mathcal{C}_2$ with $C = 2$ and $(1+\beta)$-process for $C = 1+\beta$.
- Conditions $\mathcal{C}_1$ (const ϵ) and $\mathcal{C}_2$ sufficient to prove $\text{Gap}(m) = \mathcal{O}(b/n)$ (b-Batched).
- **Condition $\mathcal{C}_3$:** There exists a $C > 1$, such that
$$\max_{i \in [n]} \left| p_i^t - \frac{1}{n} \right| \leq \frac{C-1}{n}.$$
Condition $\mathcal{C}_3$ implies condition $\mathcal{C}_2$. $(1+\beta)$-process satisfies $\mathcal{C}_3$ for $C = 1+\beta$.
- Conditions $\mathcal{C}_1$ ($\epsilon = \Theta(\sqrt{n/b})$) and $\mathcal{C}_3$ sufficient to prove $\text{Gap}(m) = \Theta(\sqrt{(b/n) \cdot \log n})$ (b-Batched).

The hyperbolic cosine potential

The hyperbolic cosine potential

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

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

The hyperbolic cosine potential

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

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

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

The hyperbolic cosine potential

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

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

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

The hyperbolic cosine potential

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

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

- For the $(1+\beta)$-process in the sequential setting, $\gamma = \Theta(\beta)$.
- [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 $t \geq 0$.

The hyperbolic cosine potential

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

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

- For the $(1+\beta)$-process in the sequential setting, $\gamma = \Theta(\beta)$.
- [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 $\mathbf{Pr}\left[\Gamma^m \leq \frac{c_2}{c_1\gamma} n^3\right] \geq 1 - n^{-2}$,

The hyperbolic cosine potential

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

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

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

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

The hyperbolic cosine potential

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

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

- For the $(1+\beta)$-process in the sequential setting, $\gamma = \Theta(\beta)$.
- [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 that w.h.p. $\mathrm{Gap}(m) = \mathcal{O}(\frac{\log n}{\beta})$.

Drift inequality statement

Theorem ([LS22a, 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 $t \geq 0$,

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

and

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

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

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

and

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

Drift inequality statement

Theorem ([LS22a, 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 $K = 2C$ and some $R = 1$, for any $t \geq 0$,

$$\mathbf{E}\left[\,\Phi^{t+1} \mid \mathfrak{F}^t\,\right] \leq \sum_{i=1}^{n} \Phi_i^t \cdot \left(1 + \left(p_i - \frac{1}{n}\right) \cdot \gamma + 2C \cdot \frac{\gamma^2}{n}\right),$$

and

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

Sequential setting with condition $\mathcal{C}_2$ for const $C > 1$

Then, there exists a constant $c := c(\delta) > 0$, such that for $\gamma \in \left(0, \min\left\{1, \frac{\epsilon\delta}{16C}\right\}\right)$

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

and

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

Drift inequality statement

Theorem ([LS22a, 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 $K = 2C$ and some $R = 1$, for any $t \geq 0$,

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

and

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

Sequential setting with condition $\mathcal{C}_2$ for const $C > 1$

Then, there exists a constant $c := c(\delta) > 0$, such that for $\gamma \in \left(0, \min\left\{1, \frac{\epsilon\delta}{16C}\right\}\right)$

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

Implies $\text{Gap}(t) = \mathcal{O}\left(\frac{\log n}{\epsilon}\right)$

and

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

Drift inequality statement

Theorem ([LS22a, 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 $K = 5(C-1)^2 \cdot \frac{b}{n}$ and some $R = b$, for any $t \geq 0$,

$$\mathbf{E}\left[\, \Phi^{t+b} \,\middle|\, \mathfrak{F}^t \,\right] \leq \sum_{i=1}^{n} \Phi_i^t \cdot \left(1 + \left(p_i - \frac{1}{n}\right) \cdot b \cdot \gamma + \frac{5(C-1)^2 b}{n} \cdot b \cdot \frac{\gamma^2}{n}\right),$$

and

$$\mathbf{E}\left[\, \Psi^{t+b} \,\middle|\, \mathfrak{F}^t \,\right] \leq \sum_{i=1}^{n} \Psi_i^t \cdot \left(1 + \left(\frac{1}{n} - p_i\right) \cdot b \cdot \gamma + \frac{5(C \cdots}{\cdots} \cdots \gamma^2 \cdots\right)$$

b-Batched setting with $\mathcal{C}_3$ for $C = 1 + \epsilon = 1 + \Theta(\sqrt{n/b})$

Then, there exists a constant $c := c(\delta) > 0$, such that for $\gamma \in \left(0, \min\left\{1, \frac{\epsilon\delta}{40(C-1)^2} \cdot \frac{n}{b}\right\}\right)$

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

and

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

Drift inequality statement

Theorem ([LS22a, 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 $K = 5(C-1)^2 \cdot \frac{b}{n}$ and some $R = b$, for any $t \geq 0$,

$$\mathbf{E}\left[\, \Phi^{t+b} \,\middle|\, \mathfrak{F}^t \,\right] \leq \sum_{i=1}^{n} \Phi_i^t \cdot \left(1 + \left(p_i - \frac{1}{n}\right) \cdot b \cdot \gamma + \frac{5(C-1)^2 b}{n} \cdot b \cdot \frac{\gamma^2}{n}\right),$$

and

$$\mathbf{E}\left[\, \Psi^{t+b} \,\middle|\, \mathfrak{F}^t \,\right] \leq \sum_{i=1}^{n} \Psi_i^t \cdot \left(1 + \left(\frac{1}{n} - p_i\right) \cdot b \cdot \gamma + \frac{5(C-1)^2 b}{} \quad \gamma^2 \right)$$

b-Batched setting with $\mathcal{C}_3$ for $C = 1 + \epsilon = 1 + \Theta(\sqrt{n/b})$

Then, there exists a constant $c := c(\delta) > 0$, such that for $\gamma \in \left(0, \min\left\{1, \frac{\epsilon\delta}{40(C-1)^2} \cdot \frac{n}{b}\right\}\right)$

$$\mathbf{E}\left[\, \Gamma^{t+b} \,\middle|\, \mathfrak{F}^t \,\right] \leq \Gamma^t \cdot b \left(1 - \frac{\gamma\epsilon\delta}{}\right) \ldots$$

Implies $\text{Gap}(t) = \mathcal{O}\left(\sqrt{n/b} \cdot \log n\right)$

and

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

Drift inequality statement

Theorem ([LS22a, 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 $K = 5(C-1)^2 \cdot \frac{b}{n}$ and some $R = b$, for any $t \geq 0$,

$$\mathbf{E}\left[\, \Phi^{t+b} \mid \mathfrak{F}^t \,\right] \leq \sum_{i=1}^{n} \Phi_i^t \cdot \left(1 + \left(p_i - \frac{1}{n}\right) \cdot b \cdot \gamma + \frac{5(C-1)^2 b}{n} \cdot b \cdot \frac{\gamma^2}{n}\right),$$

and

$$\mathbf{E}\left[\, \Psi^{t+b} \mid \mathfrak{F}^t \,\right] \leq \sum_{i=1}^{n} \Psi_i^t \cdot \left(1 + \left(\frac{1}{n} - p_i\right) \cdot b \cdot \gamma + \ldots\right)$$

b-Batched setting with $\mathcal{C}_3$ for $C = 1 + \epsilon = 1 + \Theta(\sqrt{n/b})$

Then, there exists a constant $c := c(\delta) > 0$, such that for $\gamma \in \left(0, \min\left\{1, \frac{\epsilon\delta}{40(C-1)^2} \cdot \frac{n}{b}\right\}\right)$

$$\mathbf{E}\left[\, \Gamma^{t+b} \mid \mathfrak{F}^t \,\right] \leq \Gamma^t \cdot b \cdot \left(1 - \frac{\gamma\epsilon\delta}{8n}\right) + b \cdot c\gamma\epsilon,$$

For more applications, see "Balanced Allocations: A Refined Drift Theorem with Applications".

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

Summary and Future work

Summary and Future work

Summary of results:

Summary and Future work

Summary of results:

- The $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\mathrm{Gap}(m) = \mathcal{O}(\sqrt{(b/n) \cdot \log n})$ in the b-Batched setting with $b \geq n \log n$.

Summary and Future work

Summary of results:

- The $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\mathrm{Gap}(m) = \mathcal{O}(\sqrt{(b/n) \cdot \log n})$ in the b-Batched setting with $b \geq n \log n$.
- This is almost a *quadratic improvement* over Two-Choice

Summary and Future work

Summary of results:

- The $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\text{Gap}(m) = \mathcal{O}(\sqrt{(b/n) \cdot \log n})$ in the b-Batched setting with $b \geq n \log n$.
- This is almost a *quadratic improvement* over Two-Choice and is asymptotically *optimal*.

Summary and Future work

Summary of results:

- The $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\mathrm{Gap}(m) = \mathcal{O}(\sqrt{(b/n) \cdot \log n})$ in the b-Batched setting with $b \geq n \log n$.
- This is almost a *quadratic improvement* over Two-Choice and is asymptotically *optimal*.
- The upper bound applies to a general family of processes (satisfying $\mathcal{C}_1$ and $\mathcal{C}_3$).

Summary and Future work

Summary of results:

- The $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\text{Gap}(m) = \mathcal{O}(\sqrt{(b/n) \cdot \log n})$ in the b-Batched setting with $b \geq n \log n$.
- This is almost a *quadratic improvement* over Two-Choice and is asymptotically *optimal*.
- The upper bound applies to a general family of processes (satisfying $\mathcal{C}_1$ and $\mathcal{C}_3$).

Several avenues for future work:

Summary and Future work

Summary of results:

- The $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\mathrm{Gap}(m) = \mathcal{O}(\sqrt{(b/n) \cdot \log n})$ in the b-Batched setting with $b \geq n \log n$.
- This is almost a *quadratic improvement* over Two-Choice and is asymptotically *optimal*.
- The upper bound applies to a general family of processes (satisfying $\mathcal{C}_1$ and $\mathcal{C}_3$).

Several avenues for future work:

- Investigate its performance in *practice*.

Summary and Future work

Summary of results:

- The $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\mathrm{Gap}(m) = \mathcal{O}(\sqrt{(b/n) \cdot \log n})$ in the b-Batched setting with $b \geq n \log n$.
- This is almost a *quadratic improvement* over Two-Choice and is asymptotically *optimal*.
- The upper bound applies to a general family of processes (satisfying $\mathcal{C}_1$ and $\mathcal{C}_3$).

Several avenues for future work:

- Investigate its performance in *practice*.
- Is the $(1+\beta)$-process supperior in other settings such as τ-Delay or g-Adv-Comp?

Summary and Future work

Summary of results:

- The $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\mathrm{Gap}(m) = \mathcal{O}(\sqrt{(b/n) \cdot \log n})$ in the b-Batched setting with $b \geq n \log n$.
- This is almost a *quadratic improvement* over Two-Choice and is asymptotically *optimal*.
- The upper bound applies to a general family of processes (satisfying $\mathcal{C}_1$ and $\mathcal{C}_3$).

Several avenues for future work:

- Investigate its performance in *practice*.
- Is the $(1+\beta)$-process supperior in other settings such as τ-Delay or g-Adv-Comp?
- Are there any other attractive processes with similar guarantees?

Summary and Future work

Summary of results:

- The $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\text{Gap}(m) = \mathcal{O}(\sqrt{(b/n) \cdot \log n})$ in the b-Batched setting with $b \geq n \log n$.
- This is almost a *quadratic improvement* over Two-Choice and is asymptotically *optimal*.
- The upper bound applies to a general family of processes (satisfying $\mathcal{C}_1$ and $\mathcal{C}_3$).

Several avenues for future work:

- Investigate its performance in *practice*.
- Is the $(1+\beta)$-process supperior in other settings such as τ-Delay or g-Adv-Comp?
- Are there any other attractive processes with similar guarantees?
- Apply the *mixing operation* to other algorithms and setting.

Summary and Future work

Summary of results:

- The $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\mathrm{Gap}(m) = \mathcal{O}(\sqrt{(b/n) \cdot \log n})$ in the b-Batched setting with $b \geq n \log n$.
- This is almost a *quadratic improvement* over Two-Choice and is asymptotically *optimal*.
- The upper bound applies to a general family of processes (satisfying $\mathcal{C}_1$ and $\mathcal{C}_3$).

Several avenues for future work:

- Investigate its performance in *practice*.
- Is the $(1+\beta)$-process supperior in other settings such as τ-Delay or g-Adv-Comp?
- Are there any other attractive processes with similar guarantees?
- Apply the *mixing operation* to other algorithms and setting.
- Improve the bounds on the gap to be tight up to *lower order* terms.

Summary and Future work

Summary of results:

- The $(1+\beta)$-process with $\beta = \Theta(\sqrt{(n/b) \cdot \log n})$ achieves w.h.p. $\mathrm{Gap}(m) = \mathcal{O}(\sqrt{(b/n) \cdot \log n})$ in the b-BATCHED setting with $b \geq n \log n$.
- This is almost a *quadratic improvement* over TWO-CHOICE and is asymptotically *optimal*.
- The upper bound applies to a general family of processes (satisfying $\mathcal{C}_1$ and $\mathcal{C}_3$).

Several avenues for future work:

- Investigate its performance in *practice*.
- Is the $(1+\beta)$-process supperior in other settings such as τ-DELAY or g-ADV-COMP?
- Are there any other attractive processes with similar guarantees?
- Apply the *mixing operation* to other algorithms and setting.
- Improve the bounds on the gap to be tight up to *lower order* terms.
- Investigate settings with *non-homogeneous* machines.

Questions?

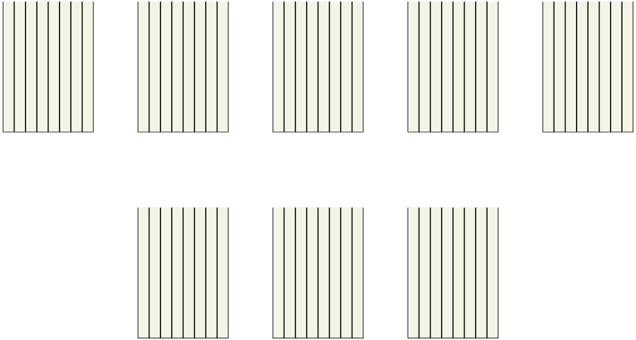

More visualisations: `dimitrioslos.com/spaa23`

Appendix A: Empirical results for Quantile(δ) process

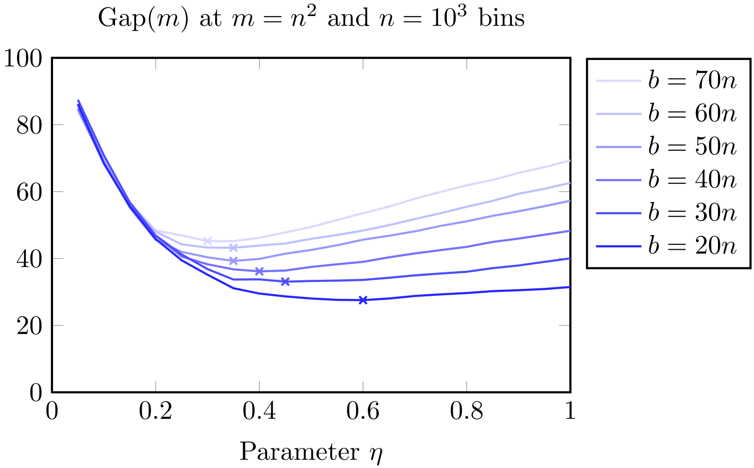

- Results for mixing the Quantile(δ) and the One-Choice process with probability $\eta \in [0, 1]$.

Appendix B: 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 $\zeta, c > 0$.

- [PTW15] showed that processes satisfying $\mathcal{C}_1$ achieve w.h.p. $\mathcal{O}\left(\frac{\log n}{\epsilon}\right)$ gap.

Open in Visualiser.

Appendix C: Empirical results for Weighted setting

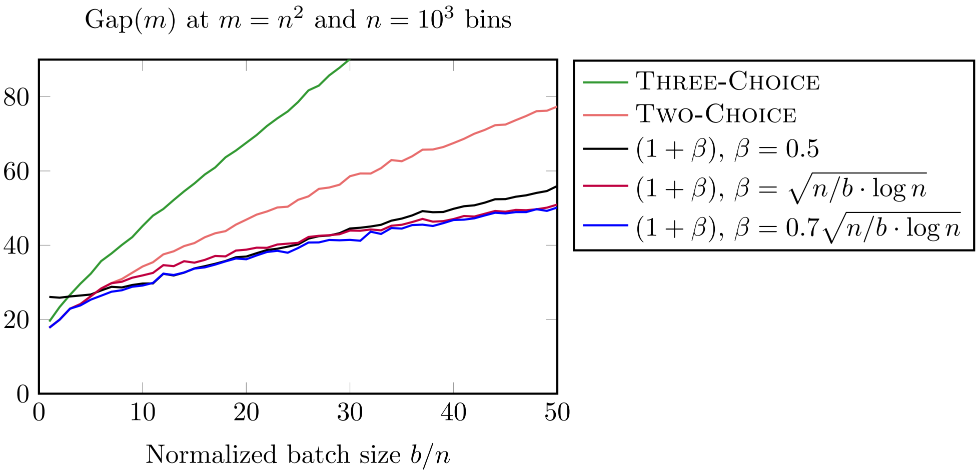

- Weights sampled from an $\mathsf{Exp}(1)$ distribution.

Appendix D: Preconditions for b-Batched setting

Appendix D: Preconditions for b-BATCHED setting

- Consider the $(1+\beta)$-process with $\beta = \Theta(\sqrt{n/b})$,

Appendix D: Preconditions for b-Batched setting

- Consider the $(1+\beta)$-process with $\beta = \Theta(\sqrt{n/b})$, and potentials Φ, Ψ, Γ with $\gamma = \Theta(\sqrt{n/b})$.

Appendix D: Preconditions for b-Batched setting

- Consider the $(1+\beta)$-process with $\beta = \Theta(\sqrt{n/b})$, and potentials Φ, Ψ, Γ with $\gamma = \Theta(\sqrt{n/b})$.
- Consider the expected change of Φ_i^t for bin $i \in [n]$, over one batch:

Appendix D: Preconditions for b-Batched setting

- Consider the $(1+\beta)$-process with $\beta = \Theta(\sqrt{n/b})$, and potentials Φ, Ψ, Γ with $\gamma = \Theta(\sqrt{n/b})$.
- Consider the expected change of Φ_i^t for bin $i \in [n]$, over one batch:

$$\mathbf{E}\left[\,\Phi_i^{t+b} \mid \mathfrak{F}^t\,\right] \quad = \quad \sum_{z \in \{0,1\}^b} \prod_{j=1}^{b} \Phi_i^t \cdot (p_i)^{z_j} (1-p_i)^{1-z_j} (\mathbf{E}[\, e^{\gamma W (1-\frac{1}{n})}\,])^{z_j} (\mathbf{E}[\, e^{-\gamma W/n}\,])^{1-z_j}$$

Appendix D: Preconditions for b-Batched setting

- Consider the $(1+\beta)$-process with $\beta = \Theta(\sqrt{n/b})$, and potentials Φ, Ψ, Γ with $\gamma = \Theta(\sqrt{n/b})$.
- Consider the expected change of Φ_i^t for bin $i \in [n]$, over one batch:

$$\begin{aligned}
\mathbf{E}\left[\,\Phi_i^{t+b} \mid \mathfrak{F}^t\,\right] &= \sum_{z \in \{0,1\}^b} \prod_{j=1}^{b} \Phi_i^t \cdot (p_i)^{z_j} (1-p_i)^{1-z_j} \left(\mathbf{E}[\,e^{\gamma W (1-\frac{1}{n})}\,]\right)^{z_j} \left(\mathbf{E}[\,e^{-\gamma W/n}\,]\right)^{1-z_j} \\
&\leq \sum_{z \in \{0,1\}^b} \prod_{j=1}^{b} \Phi_i^t \cdot \left(p_i \cdot \left(1 + \gamma \cdot \left(1 - \frac{1}{n}\right) + S\gamma^2\right)\right)^{z_j} \cdot \left((1-p_i) \cdot \left(1 - \frac{\gamma}{n} + \frac{S\gamma^2}{n^2}\right)\right)
\end{aligned}$$

Appendix D: Preconditions for b-Batched setting

- Consider the $(1+\beta)$-process with $\beta = \Theta(\sqrt{n/b})$, and potentials Φ, Ψ, Γ with $\gamma = \Theta(\sqrt{n/b})$.
- Consider the expected change of Φ_i^t for bin $i \in [n]$, over one batch:

$$\begin{aligned}
\mathbf{E}\left[\, \Phi_i^{t+b} \mid \mathfrak{F}^t \,\right] &= \sum_{z \in \{0,1\}^b} \prod_{j=1}^{b} \Phi_i^t \cdot (p_i)^{z_j} (1-p_i)^{1-z_j} \left(\mathbf{E}[\, e^{\gamma W (1-\frac{1}{n})} \,]\right)^{z_j} \left(\mathbf{E}[\, e^{-\gamma W/n} \,]\right)^{1-z_j} \\
&\leq \sum_{z \in \{0,1\}^b} \prod_{j=1}^{b} \Phi_i^t \cdot \left(p_i \cdot \left(1 + \gamma \cdot \left(1 - \frac{1}{n}\right) + S\gamma^2\right)\right)^{z_j} \cdot \left((1-p_i) \cdot \left(1 - \frac{\gamma}{n} + \frac{S\gamma^2}{n^2}\right)\right) \\
&= \Phi_i^t \cdot \left(p_i \cdot \left(1 + \gamma \cdot \left(1 - \frac{1}{n}\right) + S\gamma^2\right) + (1-p_i) \cdot \left(1 - \frac{\gamma}{n} + \frac{S\gamma^2}{n^2}\right)\right)^b
\end{aligned}$$

Appendix D: Preconditions for b-Batched setting

- Consider the $(1+\beta)$-process with $\beta = \Theta(\sqrt{n/b})$, and potentials Φ, Ψ, Γ with $\gamma = \Theta(\sqrt{n/b})$.
- Consider the expected change of Φ_i^t for bin $i \in [n]$, over one batch:

$$
\begin{aligned}
\mathbf{E}\left[\,\Phi_i^{t+b} \mid \mathfrak{F}^t\,\right] &= \sum_{z\in\{0,1\}^b} \prod_{j=1}^{b} \Phi_i^t \cdot (p_i)^{z_j} (1-p_i)^{1-z_j} (\mathbf{E}[\,e^{\gamma W(1-\frac{1}{n})}\,])^{z_j} (\mathbf{E}[\,e^{-\gamma W/n}\,])^{1-z_j} \\
&\leq \sum_{z\in\{0,1\}^b} \prod_{j=1}^{b} \Phi_i^t \cdot \left(p_i \cdot \left(1+\gamma\cdot\left(1-\frac{1}{n}\right)+S\gamma^2\right)\right)^{z_j} \cdot \left((1-p_i)\cdot\left(1-\frac{\gamma}{n}+\frac{S\gamma^2}{n^2}\right)\right) \\
&= \Phi_i^t \cdot \left(p_i \cdot \left(1+\gamma\cdot\left(1-\frac{1}{n}\right)+S\gamma^2\right) + (1-p_i)\cdot\left(1-\frac{\gamma}{n}+\frac{S\gamma^2}{n^2}\right)\right)^b \\
&\leq \Phi_i^t \cdot \left(1+\gamma\cdot\left(p_i-\frac{1}{n}\right)+2\cdot p_i\cdot S\gamma^2\right)^b
\end{aligned}
$$

Appendix D: Preconditions for b-Batched setting

- Consider the $(1+\beta)$-process with $\beta = \Theta(\sqrt{n/b})$, and potentials Φ, Ψ, Γ with $\gamma = \Theta(\sqrt{n/b})$.
- Consider the expected change of Φ_i^t for bin $i \in [n]$, over one batch:

$$
\begin{aligned}
\mathbf{E}\left[\, \Phi_i^{t+b} \mid \mathfrak{F}^t \,\right] &= \sum_{z \in \{0,1\}^b} \prod_{j=1}^{b} \Phi_i^t \cdot (p_i)^{z_j} (1-p_i)^{1-z_j} (\mathbf{E}[\, e^{\gamma W (1-\frac{1}{n})} \,])^{z_j} (\mathbf{E}[\, e^{-\gamma W/n} \,])^{1-z_j} \\
&\leq \sum_{z \in \{0,1\}^b} \prod_{j=1}^{b} \Phi_i^t \cdot \left(p_i \cdot \left(1 + \gamma \cdot \left(1 - \frac{1}{n}\right) + S\gamma^2\right)\right)^{z_j} \cdot \left((1-p_i) \cdot \left(1 - \frac{\gamma}{n} + \frac{S\gamma^2}{n^2}\right)\right) \\
&= \Phi_i^t \cdot \left(p_i \cdot \left(1 + \gamma \cdot \left(1 - \frac{1}{n}\right) + S\gamma^2\right) + (1-p_i) \cdot \left(1 - \frac{\gamma}{n} + \frac{S\gamma^2}{n^2}\right)\right)^b \\
&\leq \Phi_i^t \cdot \left(1 + \gamma \cdot \left(p_i - \frac{1}{n}\right) + 2 \cdot p_i \cdot S\gamma^2\right)^b \\
&\leq \Phi_i^t \cdot \left(1 + \left(p_i - \frac{1}{n}\right) \cdot b \cdot \gamma + \frac{5(C-1)^2 b}{n} \cdot b \cdot \frac{\gamma^2}{n}\right).
\end{aligned}
$$

Appendix D: Preconditions for b-Batched setting

- Consider the $(1+\beta)$-process with $\beta = \Theta(\sqrt{n/b})$, and potentials Φ, Ψ, Γ with $\gamma = \Theta(\sqrt{n/b})$.
- Consider the expected change of Φ_i^t for bin $i \in [n]$, over one batch:

$$
\begin{aligned}
\mathbf{E}\left[\, \Phi_i^{t+b} \mid \mathfrak{F}^t \,\right] &= \sum_{z \in \{0,1\}^b} \prod_{j=1}^{b} \Phi_i^t \cdot (p_i)^{z_j} (1-p_i)^{1-z_j} (\mathbf{E}[\, e^{\gamma W (1-\frac{1}{n})} \,])^{z_j} (\mathbf{E}[\, e^{-\gamma W/n} \,])^{1-z_j} \\
&\leq \sum_{z \in \{0,1\}^b} \prod_{j=1}^{b} \Phi_i^t \cdot \left(p_i \cdot \left(1 + \gamma \cdot \left(1 - \frac{1}{n} \right) + S\gamma^2 \right) \right)^{z_j} \cdot \left((1-p_i) \cdot \left(1 - \frac{\gamma}{n} + \frac{S\gamma^2}{n^2} \right) \right) \\
&= \Phi_i^t \cdot \left(p_i \cdot \left(1 + \gamma \cdot \left(1 - \frac{1}{n} \right) + S\gamma^2 \right) + (1-p_i) \cdot \left(1 - \frac{\gamma}{n} + \frac{S\gamma^2}{n^2} \right) \right)^b \\
&\leq \Phi_i^t \cdot \left(1 + \gamma \cdot \left(p_i - \frac{1}{n} \right) + 2 \cdot p_i \cdot S\gamma^2 \right)^b \\
&\leq \Phi_i^t \cdot \left(1 + \left(p_i - \frac{1}{n} \right) \cdot b \cdot \gamma + \frac{5(C-1)^2 b}{n} \cdot b \cdot \frac{\gamma^2}{n} \right).
\end{aligned}
$$

- Similarly, for the Ψ^t potential.

Appendix E: Outline for tighter bound

- By the refined analysis, for $\gamma = \Theta(\sqrt{n/(b \cdot \log n)})$, for any $t \geq 0$, $\mathbf{E}\left[\, \Gamma^t \,\right] \leq cn$.

Appendix E: Outline for tighter bound

- By the refined analysis, for $\gamma = \Theta(\sqrt{n/(b \cdot \log n)})$, for any $t \geq 0$, $\mathbf{E}\left[\,\Gamma^t\,\right] \leq cn$.
- Using the techniques in [LS22b], w.h.p. $\Gamma^s \leq cn$ for *all* $s \in [m - bn\log^5 n, m]$.

Appendix E: Outline for tighter bound

- By the refined analysis, for $\gamma = \Theta(\sqrt{n/(b \cdot \log n)})$, for any $t \geq 0$, $\mathbf{E}[\,\Gamma^t\,] \leq cn$.
- Using the techniques in [LS22b], 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 $\Omega\big(\sqrt{(b/n) \cdot \log n}\big)$ is at most

$$cn \cdot e^{-\gamma\Omega\left(\sqrt{(b/n)\cdot\log n}\right)} \leq \delta n.$$

Appendix E: Outline for tighter bound

- By the refined analysis, for $\gamma = \Theta(\sqrt{n/(b \cdot \log n)})$, for any $t \geq 0$, $\mathbf{E}[\,\Gamma^t\,] \leq cn$.
- Using the techniques in [LS22b], 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 $\Omega\big(\sqrt{(b/n) \cdot \log n}\big)$ is at most

$$cn \cdot e^{-\gamma\Omega\left(\sqrt{(b/n)\cdot\log n}\right)} \leq \delta n.$$

- Hence, by looking at the potential for constant $\widetilde{\gamma} > 0$ and an offset,

Appendix E: Outline for tighter bound

- By the refined analysis, for $\gamma = \Theta(\sqrt{n/(b \cdot \log n)})$, for any $t \geq 0$, $\mathbf{E}\left[\,\Gamma^t\,\right] \leq cn$.
- Using the techniques in [LS22b], 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 $\Omega\big(\sqrt{(b/n) \cdot \log n}\big)$ is at most

$$cn \cdot e^{-\gamma\Omega\left(\sqrt{(b/n)\cdot\log n}\right)} \leq \delta n.$$

- Hence, by looking at the potential for constant $\widetilde{\gamma} > 0$ and an offset,

$$\Lambda^t := \sum_{i : x_i^t \geq \frac{t}{n} + \Omega(\sqrt{(b/n)\cdot\log n})} e^{\widetilde{\gamma}\cdot(x_i^t - \frac{t}{n} - \Omega(\sqrt{(b/n)\cdot\log n}))},$$

Appendix E: Outline for tighter bound

- By the refined analysis, for $\gamma = \Theta(\sqrt{n/(b \cdot \log n)})$, for any $t \geq 0$, $\mathbf{E}[\, \Gamma^t \,] \leq cn$.
- Using the techniques in [LS22b], 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 $\Omega\big(\sqrt{(b/n) \cdot \log n}\big)$ is at most

$$cn \cdot e^{-\gamma\Omega\left(\sqrt{(b/n)\cdot\log n}\right)} \leq \delta n.$$

- Hence, by looking at the potential for constant $\widetilde{\gamma} > 0$ and an offset,

$$\Lambda^t := \sum_{i : x_i^t \geq \frac{t}{n} + \Omega(\sqrt{(b/n)\cdot \log n})} e^{\widetilde{\gamma}\cdot (x_i^t - \frac{t}{n} - \Omega(\sqrt{(b/n)\cdot\log n}))},$$

every bin i contributing to the potential has $p_i \leq \frac{1-\epsilon}{n}$, so

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

Appendix E: Outline for tighter bound

- By the refined analysis, for $\gamma = \Theta(\sqrt{n/(b \cdot \log n)})$, for any $t \geq 0$, $\mathbf{E}[\, \Gamma^t \,] \leq cn$.
- Using the techniques in [LS22b], 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 $\Omega\big(\sqrt{(b/n) \cdot \log n}\big)$ is at most

$$cn \cdot e^{-\gamma \Omega\left(\sqrt{(b/n) \cdot \log n}\right)} \leq \delta n.$$

- Hence, by looking at the potential for constant $\widetilde{\gamma} > 0$ and an offset,

$$\Lambda^t := \sum_{i : x_i^t \geq \frac{t}{n} + \Omega(\sqrt{(b/n) \cdot \log n})} e^{\widetilde{\gamma} \cdot (x_i^t - \frac{t}{n} - \Omega(\sqrt{(b/n) \cdot \log n}))},$$

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 $\mathbf{E}[\, \Lambda^m \,] = \mathcal{O}(n)$.

Appendix E: Outline for tighter bound

- By the refined analysis, for $\gamma = \Theta(\sqrt{n/(b \cdot \log n)})$, for any $t \geq 0$, $\mathbf{E}\left[\, \Gamma^t \,\right] \leq cn$.
- Using the techniques in [LS22b], 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 $\Omega\big(\sqrt{(b/n) \cdot \log n}\big)$ is at most

$$cn \cdot e^{-\gamma\Omega\left(\sqrt{(b/n) \cdot \log n}\right)} \leq \delta n.$$

- Hence, by looking at the potential for constant $\widetilde{\gamma} > 0$ and an offset,

$$\Lambda^t := \sum_{i : x_i^t \geq \frac{t}{n} + \Omega(\sqrt{(b/n) \cdot \log n})} e^{\widetilde{\gamma} \cdot (x_i^t - \frac{t}{n} - \Omega(\sqrt{(b/n) \cdot \log n}))},$$

  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 $\mathbf{E}\left[\, \Lambda^m \,\right] = \mathcal{O}(n)$.
- Finally by Markov's inequality that w.h.p. $\text{Gap}(m) = \mathcal{O}(\sqrt{(b/n) \cdot \log n})$.

Bibliography I

- D. Alistarh, J. Aspnes, and R. Gelashvili, *Space-optimal majority in population protocols*, 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'18), SIAM, 2018, pp. 2221–2239.
- D. Alistarh, t. Brown, J. Kopinsky, J. Z. Li, and G. Nadiradze, *Distributionally linearizable data structures*, 30th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA'18), ACM, 2018, pp. 133–142.
- Y. Azar, A. Z. Broder, A. R. Karlin, M. Mitzenmacher, and E. Upfal, *The ACM Paris Kanellakis Theory and Practice Award*, 2020, `https://www.acm.org/media-center/2021/may/technical-awards-2020`.
- Y. Azar, A. Z. Broder, A. R. Karlin, and E. Upfal, *Balanced allocations*, SIAM J. Comput. **29** (1999), no. 1, 180–200.
- D. Alistarh, R. Gelashvili, and J. Rybicki, *Fast graphical population protocols*, 25th International Conference on Principles of Distributed Systems (OPODIS'21), vol. 217, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, pp. 14:1–14:18.

Bibliography II

- D. Alistarh, J. Kopinsky, J. Li, and g. Nadiradze, *The power of choice in priority scheduling*, 36th Annual ACM-SIGOPT Principles of Distributed Computing (PODC'17), ACM, 2017, pp. 283–292.
- 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.
- M. Dahlin, *Interpreting stale load information*, IEEE Trans. Parallel Distributed Syst.**11** (2000), no. 10, 1033–1047.
- P. Delgado, D. Didona, F. Dinu, and W. Zwaenepoel, *Job-aware scheduling in eagle: Divide and stick to your probes*, 7th ACM Symposium on Cloud Computing (SoCC'16), ACM, 2016, pp. 497–509.

Bibliography III

- P. Delgado, F. Dinu, A. M. Kermarrec, and W. Zwaenepoel, *Hawk: Hybrid datacenter scheduling*, 2015 USENIX Annual Technical Conference (USENIX'15), USENIX, 2015, pp. 499–510.
- A. Gupta, R. Krishnaswamy, A. Kumar, and S. Singla, *Online carpooling using expander decompositions*, 40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'20), vol. 182, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020, pp. 23:1–23:14.
- G. H. Gonnet, *Expected length of the longest probe sequence in hash code searching*, J. Assoc. Comput. Mach. **28** (1981), no. 2, 289–304.
- M. Khelghatdoust and V. Gramoli, *Peacock: Probe-based scheduling of jobs by rotating between elastic queues*, 24th International Conference on Parallel and Distributed Computing (Euro-Par'18), vol. 11014, Springer, 2018, pp. 178–191.
- 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.

Bibliography IV

- D. Los and T. Sauerwald, *Balanced allocations in batches: Simplified and generalized*, 34th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA'22), ACM, 2022, p. 389–399.
- ______, *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.
- ______, *How useful is old information?*, IEEE Trans. Parallel Distributed Syst. **11** (2000), no. 1, 6–20.
- G. Nadiradze, *On achieving scalability through relaxation*, Ph.D. thesis, IST Austria, 2021.

Bibliography V

- K. Ousterhout, P. Wendell, M. Zaharia, and I. Stoica, *Sparrow: distributed, low latency scheduling*, 24th ACM SIGOPS Symposium on Operating Systems Principles (SOSP'13), ACM, 2013, pp. 69–84.
- 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.
- 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.
- W. Whitt, *Deciding which queue to join: Some counterexamples*, Oper. Res. **34** (1986), no. 1, 55–62.