Load Balancing via Randomized Local Search

Petra Berenbrink, Peter Kling, Chris Liaw,
and Abbas Mehrabian

University of British Columbia and Simon Fraser University

University of Waterloo, 20 May 2016

Load balancing

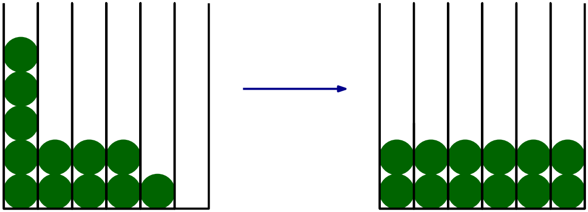

Want to re-allocate the balls into bins in order to achieve perfect balance quickly.

Load balancing

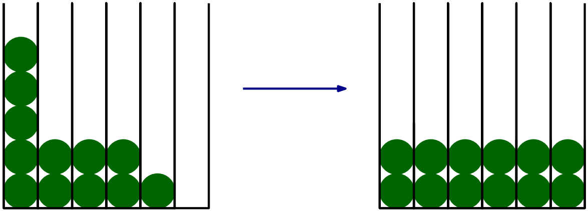

Want to re-allocate the balls into bins in order to achieve perfect balance quickly.

- ✓ Desirable: distributed protocol
- ✓ Bin-controlled vs. ball-controlled (selfish) protocols
- ✓ Synchronous vs. asynchronous protocols
- ✓ Desirable: no global knowledge
- ✓ Desirable: simplicity
- ✓ Randomization often helps!

Informal protocol description

Randomized local search: Each ball acts independently; at random times, it chooses a random bin and moves there if its own load is improved by doing so. [Paul Goldberg'04]

Informal protocol description

Randomized local search: Each ball acts independently; at random times, it chooses a random bin and moves there if its own load is improved by doing so. [Paul Goldberg'04]

Simple, distributed, asynchronous, randomized, ball-controlled

Comparison with other ball-controlled protocols

n = number of bins, m = number of balls

- ✓ Synchronous protocol, balls know m/n: $O(\ln \ln m + \ln n)$
  Even-Dar and Mansour'05
- ✓ Synchronous protocol, no global knowledge: $O(\ln \ln m + n^4)$
  Berenbrink, Friedetzky, Goldberg, Goldberg, Hu, Martin'07
- ✓ Synchronous protocol, no global knowledge: $O(\ln m + n \cdot \ln n)$
  Berenbrink, Friedetzky, Hajirasuliha, Hu'12

Comparison with other ball-controlled protocols

n = number of bins, m = number of balls

- ✓ Synchronous protocol, balls know m/n: $O(\ln \ln m + \ln n)$
  Even-Dar and Mansour'05
- ✓ Synchronous protocol, no global knowledge: $O(\ln \ln m + n^4)$
  Berenbrink, Friedetzky, Goldberg, Goldberg, Hu, Martin'07
- ✓ Synchronous protocol, no global knowledge: $O(\ln m + n \cdot \ln n)$
  Berenbrink, Friedetzky, Hajirasuliha, Hu'12
- ✓ Randomized local search (asynchronous, no global knowledge)
  $O(n^2)$ Goldberg'04
  $O\left(\ln(n)^2 + \ln(n) \cdot n^2/m\right)$
  Ganesh,Lilienthal,Manjunath,Proutiere,Simatos'12

Comparison with other ball-controlled protocols

n = number of bins, m = number of balls

- ✓ Synchronous protocol, balls know m/n: $O(\ln \ln m + \ln n)$

  Even-Dar and Mansour'05
- ✓ Synchronous protocol, no global knowledge: $O(\ln \ln m + n^4)$

  Berenbrink, Friedetzky, Goldberg, Goldberg, Hu, Martin'07
- ✓ Synchronous protocol, no global knowledge: $O(\ln m + n \cdot \ln n)$

  Berenbrink, Friedetzky, Hajirasuliha, Hu'12
- ✓ Randomized local search (asynchronous, no global knowledge)

  $O(n^2)$ Goldberg'04

  $O\left(\ln(n)^2 + \ln(n) \cdot n^2/m\right)$

  Ganesh,Lilienthal,Manjunath,Proutiere,Simatos'12

  We show the balancing time is indeed $O(\ln n + n^2/m)$

Formal protocol description

Randomized local search:

1. Each ball has an exponential clock of rate 1. When the clock rings, the ball is activated.
2. On activation, the ball chooses a random bin and moves there if its own load is improved by doing so.

Formal protocol description

Randomized local search:

1. Each ball has an exponential clock of rate 1. When the clock rings, the ball is activated.
2. On activation, the ball chooses a random bin and moves there if its own load is improved by doing so.

Let $X_1, \ldots, X_m$ be independent exponentials with rate 1, and let Z be their minimum.

- ✓ Z is exponential with rate m, so $\mathbb{E}Z = 1/m$.
- ✓ $\mathbf{Pr}(Z = X_1) = \mathbf{Pr}(Z = X_2) = \cdots = 1/m$

If you start looking at the process at any time, the waiting time for the next ball to be activated is exponential and has mean $1/m$, and the next activated ball is a uniformly random ball.

Formal protocol description

Randomized local search:

1. Each ball has an exponential clock of rate 1. When the clock rings, the ball is activated.
2. On activation, the ball chooses a random bin and moves there if its own load is improved by doing so.

Theorem (Berenbrink, Kling, Liaw and M'16+)

Consider a system of n bins and m balls in an arbitrary initial configuration. Let T be the time to reach a perfectly balanced configuration. We have $\mathbb{E}T \leq O(\ln n + n^2/m)$ and with probability at least $1 - 1/n$, we have $T \leq O(\ln n + \ln n \cdot n^2/m)$.

Tight modulo constants in the big Oh

Tightness of our bound: $\mathbb{E}T \leq O(\ln n + n^2/m)$

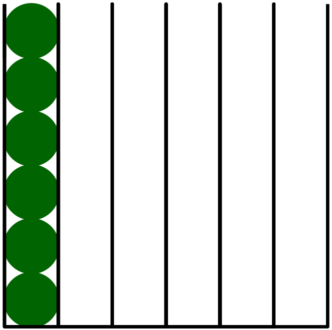

Tightness of our bound: $\mathbb{E}T \leq O(\ln n + n^2/m)$

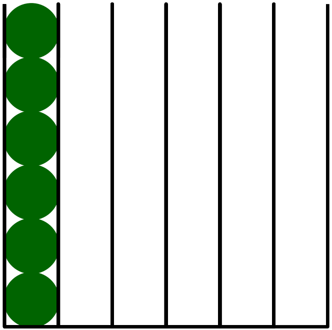

At least $m - m/n$ balls from bin 1 need to be activated:

$$\frac{1}{m} + \frac{1}{m-1} + \cdots + \frac{1}{\frac{m}{n}+1} = H_m - H_{m/n} \approx \ln m - \ln(m/n) = \ln n$$

This shows $\mathbb{E}T \geq \ln n$.

Tightness of our bound: $\mathbb{E}T \leq O(\ln n + n^2/m)$

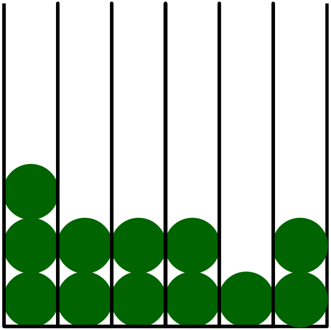

Tightness of our bound: $\mathbb{E}T \leq O(\ln n + n^2/m)$

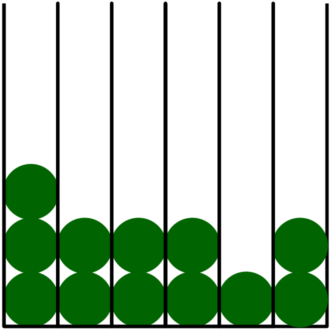

Some ball in bin 1 needs to be activated and choose bin 5.

Tightness of our bound: $\mathbb{E}T \leq O(\ln n + n^2/m)$

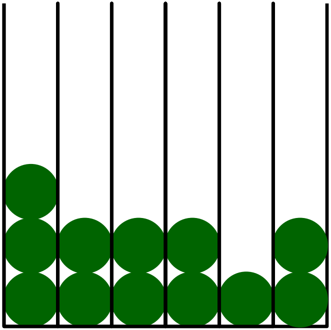

Some ball in bin 1 needs to be activated and choose bin 5.

expected waiting time for some ball in bin 1 to be activated $= \frac{1}{\frac{m}{n}+1}$
expected number of attempts to choose the correct bin $= n$

Tightness of our bound: $\mathbb{E}T \leq O(\ln n + n^2/m)$

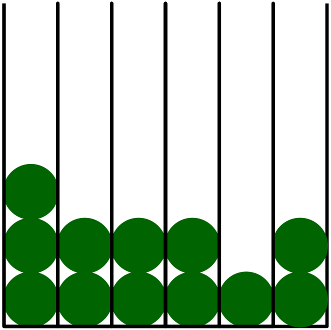

Some ball in bin 1 needs to be activated and choose bin 5.

expected waiting time for some ball in bin 1 to be activated $= \frac{1}{\frac{m}{n}+1}$
expected number of attempts to choose the correct bin $= n$ This

shows $\mathbb{E}T \geq \frac{n}{\frac{m}{n}+1} \geq n^2/2m$

Proof of main result

expected balancing time of any initial configuration

$$\leq O(\ln n + n^2/m)$$

A key lemma

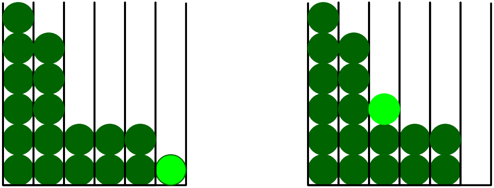

Balancing time of left configuration $\overset{st}{\leq}$ Balancing time of right configuration

A key lemma

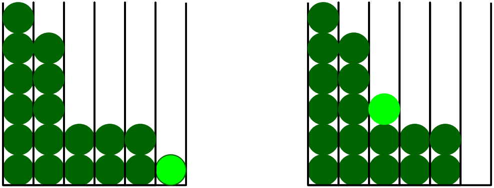

Balancing time of left configuration $\overset{st}{\leq}$ Balancing time of right configuration

$\varnothing := m/n$

Discrepancy of a configuration = maximum difference between load of a bin and the average load $= \max\{\ell_{\max} - \varnothing, \varnothing - \ell_{\min}\}$.

Perfect balance $\equiv$ discrepancy zero

A key lemma

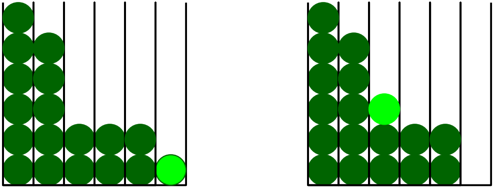

Balancing time of left configuration $\overset{st}{\leq}$ Balancing time of right configuration

$\varnothing := m/n$

Discrepancy of a configuration = maximum difference between load of a bin and the average load $= \max\{\ell_{\text{max}} - \varnothing, \varnothing - \ell_{\text{min}}\}$.

Perfect balance $\equiv$ discrepancy zero

Lemma (The key lemma)

For any $t \geq 0$, consider the configuration $\boldsymbol{\ell}(t)$ resulting from our protocol at time t. Let $\tilde{\boldsymbol{\ell}}(t)$ denote the configuration resulting from our protocol at time t under the presence of an adversary who performs an arbitrary number of destructive moves at arbitrary times. Then $\text{disc}(\boldsymbol{\ell}(t)) \overset{st}{\leq} \text{disc}(\tilde{\boldsymbol{\ell}}(t))$.

A key lemma

$\varnothing = m/n$

Discrepancy of a configuration = maximum difference between load of a bin and the average load $= \max\{\ell_{\max} - \varnothing, \varnothing - \ell_{\min}\}$.

Perfect balance $\equiv$ discrepancy zero

Lemma (The key lemma)

For any $t \geq 0$, consider the configuration $\ell(t)$ resulting from our protocol at time t. Let $\tilde{\ell}(t)$ denote the configuration resulting from our protocol at time t under the presence of an adversary who performs an arbitrary number of destructive moves at arbitrary times. Then $\mathrm{disc}(\ell(t)) \overset{st}{\leq} \mathrm{disc}(\tilde{\ell}(t))$.

Helps in two ways: (1) we may do some destructive moves to make “well-shaped” configurations that are simpler to analyse.

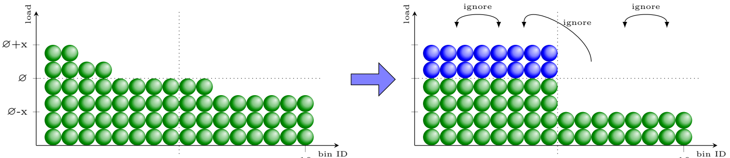

A key lemma

$\varnothing = m/n$

Discrepancy of a configuration = maximum difference between load of a bin and the average load $= \max\{\ell_{\max} - \frac{m}{n}, \frac{m}{n} - \ell_{\min}\}$.

Perfect balance $\equiv$ discrepancy zero

Lemma (The key lemma)

For any $t \geq 0$, consider the configuration $\ell(t)$ resulting from our protocol at time t. Let $\tilde{\ell}(t)$ denote the configuration resulting from our protocol at time t under the presence of an adversary who performs an arbitrary number of destructive moves at arbitrary times. Then $\text{disc}(\ell(t)) \overset{st}{\leq} \text{disc}(\tilde{\ell}(t))$.

Helps in two ways: (2) we may “ignore” certain (at the moment unwanted) moves made by the protocol

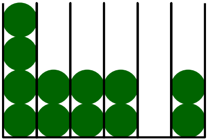

Analysis outline

- ✓ Initial discrepancy is $m - \varnothing$, want to reduce to 0
- ✓ In the zeroth phase, discrepancy is reduced to $\varnothing$

  What happens: $m - \varnothing$ balls leave bin 1

  Running time of zeroth phase $\leq O(\ln n)$
- ✓ In the first phase, discrepancy is reduced to $O(\ln n)$

  What happens: in each subphase, all loads get much closer to $\varnothing$ simultaneously

  Running time of first phase $\leq O(\ln n)$
- ✓ In the second phase, discrepancy is reduced to 0

  What happens: in each step, we get just one ball closer to perfect balance

  Running time of second phase $\leq O(n^2/m)$
- ✓ Total running time $\leq O(\ln n + n^2/m)$

First phase

Lemma

For any $x \geq 4 \ln n$, the expected time to reduce discrepancy from x to $\sqrt{4x \ln n}$ is $\leq \ln\left(\frac{\varnothing + x}{\varnothing - x}\right)$

First phase

Lemma

For any $x \geq 4\ln n$, the expected time to reduce discrepancy from x to $\sqrt{4x\ln n}$ is $\leq \ln\left(\frac{\varnothing+x}{\varnothing-x}\right)$

We have

$$\ln(\varnothing + x) - \ln(\varnothing - x) = \ln\left(1 + \frac{2x}{\varnothing - x}\right) \leq \frac{2x}{\varnothing - x} \leq 4x/\varnothing$$

Straightforward calculations give

$$\begin{aligned} \frac{4x}{\varnothing} + \frac{4\sqrt{4x\ln n}}{\varnothing} + \ldots &\leq \frac{16\ln n}{\varnothing}\left(x + \sqrt{x} + \sqrt{\sqrt{x}} + \ldots\right) \\ &\leq O\left(x \cdot \ln n/\varnothing\right) = O(\ln n) \end{aligned}$$

is the expected time to bring discrepancy down to $O(\ln n)$, hence completing the first phase.

First phase

(Chernoff bound)

Let X be a sum of independent $0,1$-random variables. For any $\varepsilon \in [0,1]$ we have

$$\mathbf{Pr}\left(|X - \mathbb{E}X| > \varepsilon \mathbb{E}X\right) < 2e^{-\varepsilon^2 \mathbb{E}X/3}$$

In particular, if $\mathbb{E}X \geq 4 \ln n$,

$$\mathbf{Pr}\left(|X - \mathbb{E}X| > \sqrt{4 \ln n \cdot \mathbb{E}X}\right) < n^{-2}$$

First phase

Lemma

For any $x \geq 4 \ln n$, the expected time to reduce discrepancy from x to $\sqrt{4x \ln n}$ is $\leq \ln\left(\frac{\varnothing + x}{\varnothing - x}\right)$

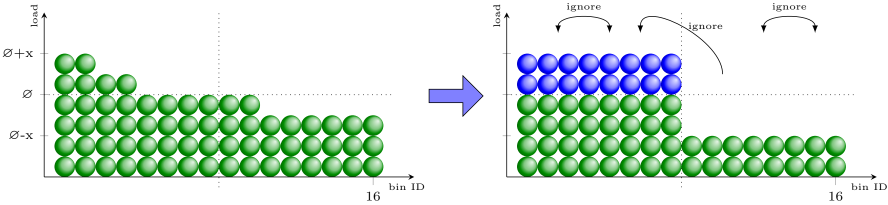

First phase

Lemma

For any $x \geq 4 \ln n$, the expected time to reduce discrepancy from x to $\sqrt{4x \ln n}$ is $\leq \ln\left(\frac{\varnothing + x}{\varnothing - x}\right)$

Let $t := \ln(\varnothing + x) - \ln(\varnothing - x)$ and consider time interval $[0, t]$.

Activation probability of a ball $= 1 - \exp(-t) =: p$.

First phase

Lemma

For any $x \geq 4 \ln n$, the expected time to reduce discrepancy from x to $\sqrt{4x \ln n}$ is $\leq \ln\left(\frac{\varnothing + x}{\varnothing - x}\right)$

Let $t := \ln(\varnothing + x) - \ln(\varnothing - x)$ and consider time interval $[0, t]$.

Activation probability of a ball $= 1 - \exp(-t) =: p$.

Consider a heavy bin.

Each of its balls is activated with probability p, and moves to a light bin with probability $1/2$.

So the number of balls it loses is a sum of independent $\{0, 1\}$-random variables and has mean $= (\varnothing + x) \times p/2 = x$.

By Chernoff, with probability $\geq 1 - n^{-2}$ this bin loses between $[x - \sqrt{4x \ln n}, x + \sqrt{4x \ln n}]$ balls, and so will have between $\varnothing - \sqrt{4x \cdot \ln n}$ and $\varnothing + \sqrt{4x \cdot \ln n}$ balls at time t.

First phase

Lemma

For any $x \geq 4 \ln n$, the expected time to reduce discrepancy from x to $\sqrt{4x \ln n}$ is $\leq \ln\left(\frac{\varnothing + x}{\varnothing - x}\right)$

Let $t := \ln(\varnothing + x) - \ln(\varnothing - x)$ and consider time interval $[0, t]$.

Activation probability of a ball $= 1 - \exp(-t) =: p$.

Consider a heavy bin.

Each of its balls is activated with probability p, and moves to a light bin with probability $1/2$.

So the number of balls it loses is a sum of independent $\{0, 1\}$-random variables and has mean $= (\varnothing + x) \times p/2 = x$.

By Chernoff, with probability $\geq 1 - n^{-2}$ this bin loses between $[x - \sqrt{4x \ln n}, x + \sqrt{4x \ln n}]$ balls, and so will have between $\varnothing - \sqrt{4x \cdot \ln n}$ and $\varnothing + \sqrt{4x \cdot \ln n}$ balls at time t.

A similar reasoning works for a light bin.

Second phase

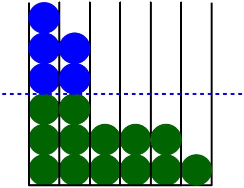

Lemma

Assuming discrepancy is $O(\ln n)$, the average time to reduce the number of overloaded balls to n is $\leq O\left(n(\ln n)^2/m\right)$.

Lemma

Assuming the number of overloaded balls is n, the average time to reduce the discrepancy to 1 is $\leq O\left(n^2/m\right)$.

Lemma

Assuming the discrepancy is 1, the average time to reduce discrepancy to 0 is $\leq O\left(n^2/m\right)$.

Second phase

Lemma

Assuming the discrepancy is 1, the average time to reduce discrepancy to 0 is $\leq O\left(n^2/m\right)$.

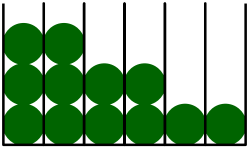

Second phase

Lemma

Assuming the discrepancy is 1, the average time to reduce discrepancy to 0 is $\leq O\left(n^2/m\right)$.

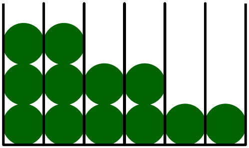

There are A bins of load $> \varnothing$, and so there are also A bins of load $< \varnothing$.

If $A \geq 1$ there are $(\varnothing + 1) \cdot A$ balls that, when activated, find an underloaded bin with probability A/n.

The expected time for such a move to happen is $\frac{1}{A \cdot (\varnothing+1)} \cdot \frac{1}{A/n} \leq \frac{n}{\varnothing \cdot A^2}$

The expected total time to balance out is less than

$$\sum_{A=1}^{\infty} \frac{n}{\varnothing \cdot A^2} = \frac{\pi^2}{6} \times \frac{n}{\varnothing} = O(n^2/m)$$

So we have

expected balancing time of any initial configuration

$$\leq O(\ln n + n^2/m)$$

We now briefly consider the case where bins have different speeds...

Non-uniform-speed setting

Load of a bin = number of balls divided by bin's speed

Non-uniform-speed setting

Load of a bin = number of balls divided by bin's speed

Let $\boldsymbol{s}$ denote the speeds vector, let $\boldsymbol{p} \in [0,1]^n$ be a probability distribution, and let $\boldsymbol{q} \in [0,1]^n$ be arbitrary.

GRLS($m, n, \boldsymbol{s}, \boldsymbol{p}, \boldsymbol{q}$)
(code executed by an activated ball in bin i)
sample random bin i' according to $\boldsymbol{p}$
with probability $q_{i'}$ do the following:
  $\ell_{\text{cur}} \leftarrow$ current load of bin i
  $\ell_{\text{new}} \leftarrow$ load of bin i' in case the ball moved to bin i'
  if $\ell_{\text{new}} \leq \ell_{\text{cur}}$: move to bin i'

Non-uniform-speed setting

Load of a bin = number of balls divided by bin's speed
Let $\boldsymbol{s}$ denote the speeds vector, let $\boldsymbol{p} \in [0,1]^n$ be a probability distribution, and let $\boldsymbol{q} \in [0,1]^n$ be arbitrary.

GRLS$(m, n, \boldsymbol{s}, \boldsymbol{p}, \boldsymbol{q})$
(code executed by an activated ball in bin i)
sample random bin i' according to $\boldsymbol{p}$
with probability $q_{i'}$ do the following:

- $\ell_{\text{cur}} \leftarrow$ current load of bin i
- $\ell_{\text{new}} \leftarrow$ load of bin i' in case the ball moved to bin i'
- if $\ell_{\text{new}} \le \ell_{\text{cur}}$: move to bin i'

$$\mathcal{T}(RLS(m,n,\boldsymbol{s})) \overset{st}{\le} \mathcal{T}\left(GRLS\left(m,n,\boldsymbol{s},\frac{\mathbf{1}}{n},\frac{\boldsymbol{s}}{s_{\max}}\right)\right)$$
$$\overset{d}{\equiv} \frac{ns_{\max}}{S} \cdot \mathcal{T}\left(GRLS\left(m,n,\boldsymbol{s},\frac{\boldsymbol{s}}{S},\mathbf{1}\right)\right)$$
$$\overset{st}{\le} \frac{ns_{\max}}{S}\mathcal{T}(RLS(m,n,\mathbf{1}))$$

Non-uniform-speeds setting

Theorem (Berenbrink, Kling, Liaw and M'16+)

Consider a system of n bins (with speeds) and m identical balls in an arbitrary initial configuration. Assume that the minimum speed is 1, and let $s_{\max}$ and S denote the maximum speed and the sum of speeds, respectively. Let T be the time to reach a perfectly balanced configuration. We have

$$\mathbb{E}T = O\left(\ln(S) \cdot n s_{\max}/S + s_{\max} S \cdot n/m\right)$$

and with probability $\geq 1 - 1/n$ we have
$T = O\left(\ln(S) \cdot n s_{\max}/S + \ln(n) \cdot s_{\max} S \cdot n/m\right)$.

If bins are identical, $s_{\max} = 1$ and $S = n$, so $\mathbb{E}T = O\left(\ln n + n^2/m\right)$.