

MEMORYLESS WORKER-TASK ASSIGNMENT

WITH POLYLOGARITHMIC SWITCHING COST

Aaron Berger · William Kuszmaul · **Adam Polak** · Jonathan Tidor · Nicole Wein



PSEUDORANDOMNESS

**DISTRIBUTED
COMPUTING**

**METRIC
EMBEDDINGS**

**MEMORYLESS
WORKER-TASK
ASSIGNMENT**

COMBINATORICS

**RAMSEY
THEORY**

**PROBABILISTIC
METHOD**

WORKER-TASK ASSIGNMENT




WORKER-TASK ASSIGNMENT




$$v = (1, 2, 0, 2)$$

WORKER-TASK ASSIGNMENT




$$v = (1, 2, 0, 2)$$

WORKER-TASK ASSIGNMENT




$$v = \text{~~(1, 2, 0, 2)~~}$$
$$(0, 2, 0, 3)$$

WORKER-TASK ASSIGNMENT



$$v = \begin{matrix} \text{---} (1, 2, 0, 2) \\ (0, 2, 0, 3) \end{matrix}$$

WORKER-TASK ASSIGNMENT




$$v = \text{---}(1, 2, 0, 2)\text{---}$$

$$(0, 2, 0, 3)$$

switching cost = 2

MEMORYLESS WORKER-TASK ASSIGNMENT



$$v = \text{---}(1, 2, 0, 2)\text{---}$$

$$(0, 2, 0, 3)$$

switching cost = 2


Memoryless workers

[SSDL'17]

- know only current v

[SSDL'17] Su, Su, Dornhaus, Lynch

MEMORYLESS WORKER-TASK ASSIGNMENT



$$v = \underline{(1, 2, 0, 2)}$$

$$(0, 2, 0, 3)$$

switching cost = 2


Memoryless workers

[SSDL'17]

- know only current v
- implicitly communicationless

[SSDL'17] Su, Su, Dornhaus, Lynch

MEMORYLESS WORKER-TASK ASSIGNMENT



$$v = \underline{(1, 2, 0, 2)}$$

$$(0, 2, 0, 3)$$

switching cost = 2

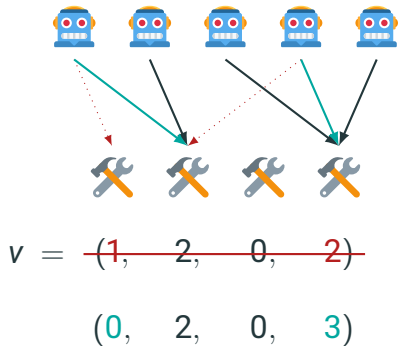
Memoryless workers

[SSDL'17]

- know only current v
- implicitly communicationless
- e.g., internet of things, insects, ...

[SSDL'17] Su, Su, Dornhaus, Lynch

MEMORYLESS WORKER-TASK ASSIGNMENT



switching cost = 2

Memoryless workers

[SSDL'17]

- know only current v
- implicitly communicationless
- e.g., internet of things, insects, ...
- fault-tolerant, can reboot

[SSDL'17] Su, Su, Dornhaus, Lynch

FORMAL PROBLEM STATEMENT

Problem: design a *worker-task assignment function* f that

FORMAL PROBLEM STATEMENT

Problem: design a *worker-task assignment function* f that

given: demand vector $\mathbf{v} \in \mathbb{Z}_{\geq 0}^{\#\text{tasks}}$ with $\sum_i v_i = \#\text{workers}$

outputs: assignment $f(\mathbf{v})$ of workers to tasks that satisfies \mathbf{v}

FORMAL PROBLEM STATEMENT

Problem: design a *worker-task assignment function* f that

given: demand vector $\mathbf{v} \in \mathbb{Z}_{\geq 0}^{\#\text{tasks}}$ with $\sum_i v_i = \#\text{workers}$

outputs: assignment $f(\mathbf{v})$ of workers to tasks that satisfies \mathbf{v}

Goal: minimize the switching cost = $\max_{\substack{v_1, v_2 \\ |v_1 - v_2| = 1}} |f(v_1) - f(v_2)|$

FORMAL PROBLEM STATEMENT

Problem: design a *worker-task assignment function* f that

given: demand vector $\mathbf{v} \in \mathbb{Z}_{\geq 0}^{\#\text{tasks}}$ with $\sum_i v_i = \#\text{workers}$

outputs: assignment $f(\mathbf{v})$ of workers to tasks that satisfies \mathbf{v}

Goal: minimize the switching cost = $\max_{\substack{v_1, v_2 \\ |v_1 - v_2| = 1}} |f(v_1) - f(v_2)|$

...then $|f(v_1) - f(v_2)| \leq \text{switching cost} \cdot |v_1 - v_2|$, for every v_1, v_2

For i from 1 to $\#$ workers

worker $i \mapsto$ the first task with yet unmet demand

For i from 1 to $\#$ workers

worker $i \mapsto$ the first task with yet unmet demand



$$v_1 = (2, 1, 2, 0)$$

For i from 1 to $\#$ workers

worker $i \mapsto$ the first task with yet unmet demand



$$v_1 = (2, 1, 2, 0)$$

For i from 1 to $\#$ workers


worker $i \mapsto$ the first task with yet unmet demand



$$v_1 = (2, 1, 2, 0)$$

For i from 1 to $\#$ workers


worker $i \mapsto$ the first task with yet unmet demand



$$v_1 = (2, 1, 2, 0)$$

For i from 1 to $\#$ workers


worker $i \mapsto$ the first task with yet unmet demand



$$v_1 = (2, 1, 2, 0)$$

For i from 1 to $\#$ workers


worker $i \mapsto$ the first task with yet unmet demand



$$v_1 = (2, 1, 2, 0)$$

For i from 1 to $\#$ workers

worker $i \mapsto$ the first task with yet unmet demand




$$v_1 = (2, 1, 2, 0)$$

$$v_2 = (1, 1, 2, 1)$$

When demands change...

For i from 1 to $\#$ workers

worker $i \mapsto$ the first task with yet unmet demand




$$v_1 = (2, 1, 2, 0)$$

$$v_2 = (1, 1, 2, 1)$$

When demands change...

For i from 1 to $\#$ workers

worker $i \mapsto$ the first task with yet unmet demand




$$v_1 = (2, 1, 2, 0)$$

$$v_2 = (1, 1, 2, 1)$$

When demands change...

For i from 1 to $\#$ workers

worker $i \mapsto$ the first task with yet unmet demand




$$v_1 = (2, 1, 2, 0)$$

$$v_2 = (1, 1, 2, 1)$$

When demands change...

For i from 1 to $\#$ workers

worker $i \mapsto$ the first task with yet unmet demand




$$v_1 = (2, 1, 2, 0)$$

$$v_2 = (1, 1, 2, 1)$$

When demands change...

For i from 1 to $\#$ workers

worker $i \mapsto$ the first task with yet unmet demand




$$v_1 = (2, 1, 2, 0)$$

$$v_2 = (1, 1, 2, 1)$$

When demands change...

For i from 1 to $\#workers$

worker $i \mapsto$ the first task with yet unmet demand



$$v_1 = (2, 1, 2, 0)$$

$$v_2 = (1, 1, 2, 1)$$


When demands change...

at most one worker moves
between tasks j and $j + 1$,

for every $j = 1, 2, \dots, \#tasks - 1$

For i from 1 to $\#workers$

worker $i \mapsto$ the first task with yet unmet demand



$$v_1 = (2, 1, 2, 0)$$

$$v_2 = (1, 1, 2, 1)$$

When demands change...

at most one worker moves
between tasks j and $j + 1$,

for every $j = 1, 2, \dots, \#tasks - 1$

\implies

switching cost $\leq \#tasks - 1$

SWITCHING COST BOUNDS

$$t = \# \text{tasks} \quad w = \# \text{workers}$$

$$2 \leq \text{switching cost} \leq t - 1$$

[Su, Su, Dornhaus, Lynch '17]

SWITCHING COST BOUNDS

$t = \# \text{tasks}$ $w = \# \text{workers}$

$2 \leq \text{switching cost} \leq t - 1$

[Su, Su, Dornhaus, Lynch '17]

switching cost ≤ 2 when $t \leq 4$ and $w \leq 6$

[Su, Su, Dornhaus, Lynch '17]

$3 \leq \text{switching cost}$ when $t \geq 5$ and $w \geq 3$

[Su, Wein '20]

$4 \leq \text{switching cost}$ when $t \geq 2 \uparrow\uparrow (w - 1)$

[Su, Wein '20]

SWITCHING COST BOUNDS

$t = \# \text{tasks}$ $w = \# \text{workers}$

$2 \leq \text{switching cost} \leq t - 1$

[Su, Su, Dornhaus, Lynch '17]

switching cost ≤ 2 when $t \leq 4$ and $w \leq 6$

[Su, Su, Dornhaus, Lynch '17]

$3 \leq \text{switching cost}$ when $t \geq 5$ and $w \geq 3$

[Su, Wein '20]

$4 \leq \text{switching cost}$ when $t \geq 2 \uparrow\uparrow (w - 1)$

[Su, Wein '20]

switching cost $\leq O(\log^2(wt))$

[Berger, Kuszmaul, P., Tidor, Wein '22]

SWITCHING COST BOUNDS

$t = \# \text{tasks}$ $w = \# \text{workers}$

$2 \leq \text{switching cost} \leq t - 1$

[Su, Su, Dornhaus, Lynch '17]

switching cost ≤ 2 when $t \leq 4$ and $w \leq 6$

[Su, Su, Dornhaus, Lynch '17]

$3 \leq \text{switching cost}$ when $t \geq 5$ and $w \geq 3$

[Su, Wein '20]

$4 \leq \text{switching cost}$ when $t \geq 2 \uparrow\uparrow (w - 1)$

[Su, Wein '20]

$w \leq \text{switching cost} \leq O(\log^2(wt))$

[Berger, Kuszmaul, P., Tidor, Wein '22]

\uparrow when $t \geq 2 \uparrow\uparrow \Omega(w)$

Remaining gap: $\Omega(\log^*(wt)) \leq \text{switching cost} \leq O(\log^2(wt))$

OUR ALGORITHM

WITH POLYLOGARITHMIC SWITCHING COST

OUR ALGORITHM

Baby case: $v_i \leq 1$

(reduce to it via dummy w copies of each task)

OUR ALGORITHM

Baby case: $v_i \leq 1$ (reduce to it via dummy w copies of each task)

Each round assigns $1/100$ of yet unassigned workers $\implies O(\log w)$ rounds

OUR ALGORITHM

Baby case: $v_i \leq 1$ (reduce to it via dummy w copies of each task)

Each round assigns $1/100$ of yet unassigned workers $\implies O(\log w)$ rounds

Each round needs two *hash* functions: $h_w : [w] \rightarrow [b]$ and $h_t : [t] \rightarrow [b]$
 \uparrow $b = \#$ unassigned workers

OUR ALGORITHM

Baby case: $v_i \leq 1$ (reduce to it via dummy w copies of each task)

Each round assigns $1/100$ of yet unassigned workers $\implies O(\log w)$ rounds

Each round needs two *hash* functions: $h_w : [w] \rightarrow [b]$ and $h_t : [t] \rightarrow [b]$

\uparrow $b = \#$ unassigned workers

We use **random** functions h_w, h_t

\implies **probabilistic** construction of a **deterministic** assignment function

OUR ALGORITHM

Baby case: $v_i \leq 1$ (reduce to it via dummy w copies of each task)

Each round assigns $1/100$ of yet unassigned workers $\implies O(\log w)$ rounds

Each round needs two *hash* functions: $h_w : [w] \rightarrow [b]$ and $h_t : [t] \rightarrow [b]$

$\lleftarrow b = \#$ unassigned workers

We use **random** functions h_w, h_t

\implies **probabilistic** construction of a **deterministic** assignment function


\lleftarrow **succeeds w.h.p.**, can be derandomized with switching cost $\log^{O(1)}(wt)$

\lleftarrow using *strong dispersers*

SINGLE ROUND OF OUR ALGORITHM




SINGLE ROUND OF OUR ALGORITHM

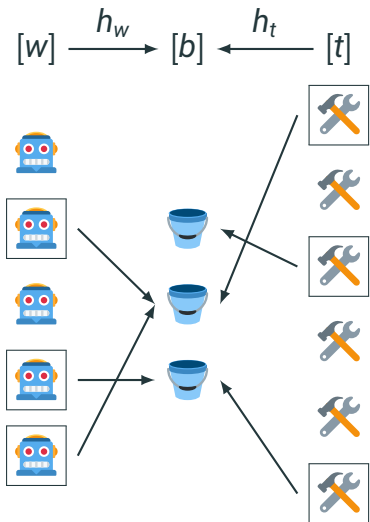


SINGLE ROUND OF OUR ALGORITHM

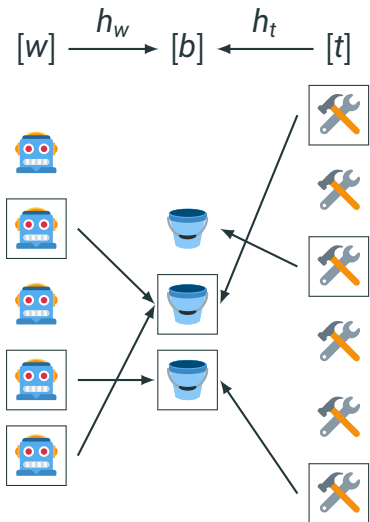
$$[w] \xrightarrow{h_w} [b]$$




SINGLE ROUND OF OUR ALGORITHM



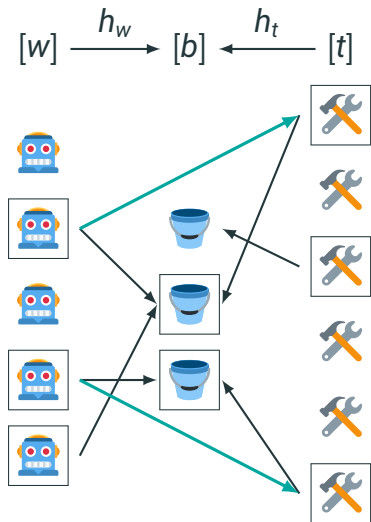
SINGLE ROUND OF OUR ALGORITHM



SINGLE ROUND OF OUR ALGORITHM



SINGLE ROUND OF OUR ALGORITHM




$$\mathbb{P}(\text{const fraction of buckets used}) \geq e^{-b}$$

repeat $O(\log wt)$ times to amplify

\implies can union bound over $\binom{w}{b} \binom{t}{b}$ inputs

SINGLE ROUND OF OUR ALGORITHM




$$\mathbb{P}(\text{const fraction of buckets used}) \geq e^{-b}$$

repeat $O(\log wt)$ times to amplify

\implies can union bound over $\binom{w}{b} \binom{t}{b}$ inputs

switching cost of a round ≤ 4

SINGLE ROUND OF OUR ALGORITHM



$$\mathbb{P}(\text{const fraction of buckets used}) \geq e^{-b}$$

repeat $O(\log wt)$ times to amplify

\implies can union bound over $\binom{w}{b} \binom{t}{b}$ inputs

switching cost of a round ≤ 4

total switching cost $\leq O(\#\text{rounds})$

$\leq O(\log(w) \log(wt))$

APPLICATION TO
METRIC EMBEDDINGS

METRIC EMBEDDINGS

Worker-task assignment function f **embeds**

$$t \gg w$$

w -sparse binary vectors from
high-dimensional L1 space

into

low-dimensional Hamming space

METRIC EMBEDDINGS

Worker-task assignment function f **embeds**

$$t \gg w$$

w -sparse binary vectors from
high-dimensional L1 space

$$v = (0, 1, 0, 1, 1) \in \{0, 1\}^t, \|v\|_0 = w$$

into

low-dimensional Hamming space

METRIC EMBEDDINGS

Worker-task assignment function f **embeds**

$t \gg w$

w -sparse binary vectors from
high-dimensional **L1** space

into

low-dimensional **Hamming** space

$$v = (0, 1, 0, 1, 1) \in \{0, 1\}^t, \|v\|_0 = w$$

$$f(v) = (5, 2, 4) \in [t]^w$$


METRIC EMBEDDINGS

Worker-task assignment function f **embeds**

$$t \gg w$$

w -sparse binary vectors from
high-dimensional **L1** space

into

low-dimensional **Hamming** space

$$v = (0, 1, 0, 1, 1) \in \{0, 1\}^t, \|v\|_0 = w$$

$$f(v) = (5, 2, 4) \in [t]^w$$


$$\|x - y\|_1 \leq \text{Ham}(f(x), f(y)) \leq \underbrace{\text{switching cost}}_{= \text{distortion of the embedding}} \cdot \|x - y\|_1$$

METRIC EMBEDDINGS

Worker-task assignment function f **embeds**

$$t \gg w$$

w -sparse binary vectors from
high-dimensional L1 space

into

low-dimensional **Hamming** space

$$v = (0, 1, 0, 1, 1) \in \{0, 1\}^t, \|v\|_0 = w$$

$$f(v) = (5, 2, 4) \in [t]^w$$


$$\|x - y\|_1 \leq \text{Ham}(f(x), f(y)) \leq \underbrace{\text{switching cost}}_{= \text{distortion of the embedding}} \cdot \|x - y\|_1$$

THANK YOU!