Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

FACULTY
OF MATHEMATICS, INFORMATICS
AND NATURAL SCIENCES

DISCRETE LOAD BALANCING : POWER OF MATCHINGS

Collaborators:

- Petra Berenbrink. University of Hamburg, Germany.
- Tom Friedetzky. Durham University, England.
- Robert Elsässer. University of Salzburg, Austria.
- Hamed Hosseinpour. University of Hamburg, Germany.
- Dominik Kaaser. Hamburg University of Technology, Germany.
- Peter Kling. University of Hamburg.
- Thomas Sauerwald. University of Cambridge, England.

Problem Definition:

- Graph $G = (V, E)$ with n nodes.
- $X_u(t)$ is the discrete load of node u at the beginning of round t.
- In each round a matching is given. Nodes u and v balance their loads if edge (u, v) is in the matching:
  - with probability $1/2$, u receives $\left\lceil \frac{X_u(t)+X_v(t)}{2} \right\rceil$ and v receives $\left\lfloor \frac{X_u(t)+X_v(t)}{2} \right\rfloor$, or vise versa.

Objective: Minimize the maximum load difference between any pair of nodes.

Preliminaries:

- Based on the given matching of round t, the *balancing matrix* $M^{(t)} \in \mathbb{R}^{n\times n}$ is defined as follows:

$$M^{(t)}_{u,v} = \begin{cases} \frac{1}{2} : & u \neq v \text{ are matched in step } t\,; \\ 0 : & u \neq v \text{ are not matched in step } t\,; \\ 1 : & u = v \text{ and } u \text{ is not matched in step } t. \end{cases}$$

Continues Setting:

- The load can be arbitrarily divided. At round t, each of two matched nodes u, v get $(X_u(t) + X_v(t))/2$ load and we have $X(t + 1) = X(t) \cdot M(t + 1)$ ([1]).
- $\tau_{BC}(K)$: the time it takes to reduce the initial discrepancy K to $1/2n$ in the continues setting.
- [1] : $\tau_{BC}(K) = O(\log(Kn)/(1 - \lambda))$, in which $1 - \lambda$ is the spectral gap of the diffusion matrix.

- [1]: The difference of discrete and continues load, for a fixed node, at round t is a weighted sum of *errors* on all edges from all rounds $s \leq t$.
- [2]: For any $\delta > 1/2n$: Pr[discrepancy at round $\tau_{BC}(K) \geq \delta] \leq e^{-c\delta^2}$.
- $\bar{x}$ is the average load. Node u has $\max\{X_u(t) - \bar{x}, 0\}$ *excess* tokens in round t.
- [2]: Discrepancy is at most two times the maximum number of excess tokens on a node.

Result:

$\Pr[$Discrepancy after $O(\tau_{BC}(K))$ rounds is $O(1)] \geq 1 - \exp(-\log(n)/\log\log n)$.

Analysis:

- We establish a tail concentration bound for a weighted sum of loads after $\tau_{BC}(k)$.
- For a vector $(a_w)_{w\in V}$ with $|a|_1 = 1$ and round $t \geq \tau_{BC}(K)$, it holds

$$\Pr\left[\sum_{w\in V} |a_w \cdot X_u(t) - \bar{x}| \geq 1/2n + \delta\right] \leq \exp(-c\delta^2/\|a\|_2^2).$$

- We consider an *auxiliary* process which has the same load distribution as the original process.

Auxiliary Process:

- The auxiliary process in round t has two steps called **moving** and **shuffling**. Assume the edge $\{u, v\} \in M(t)$ and $x_u(t - 1) \geq x_v(t - 1)$.
  - At the beginning of the process the tokens on each node are ordered in an arbitrary way. The **height** of a token is its position in that order, starting at 1.
  - In the moving step the process moves the top $\lfloor (x_u(t - 1) - x_v(t - 1))/2 \rfloor$ tokens from u to v.
  - We call tokens located on u and v with equal height as **buddies**. In the shuffling step, each token on u swaps its location with its buddy with probability $1/2$, independently from all other tokens.
- The height of a token never increases over rounds.
- The location of tokens are **Negatively Correlated**.
  - The probability that a subset of tokens are located at a subset of nodes is at most the product of individual probabilities.

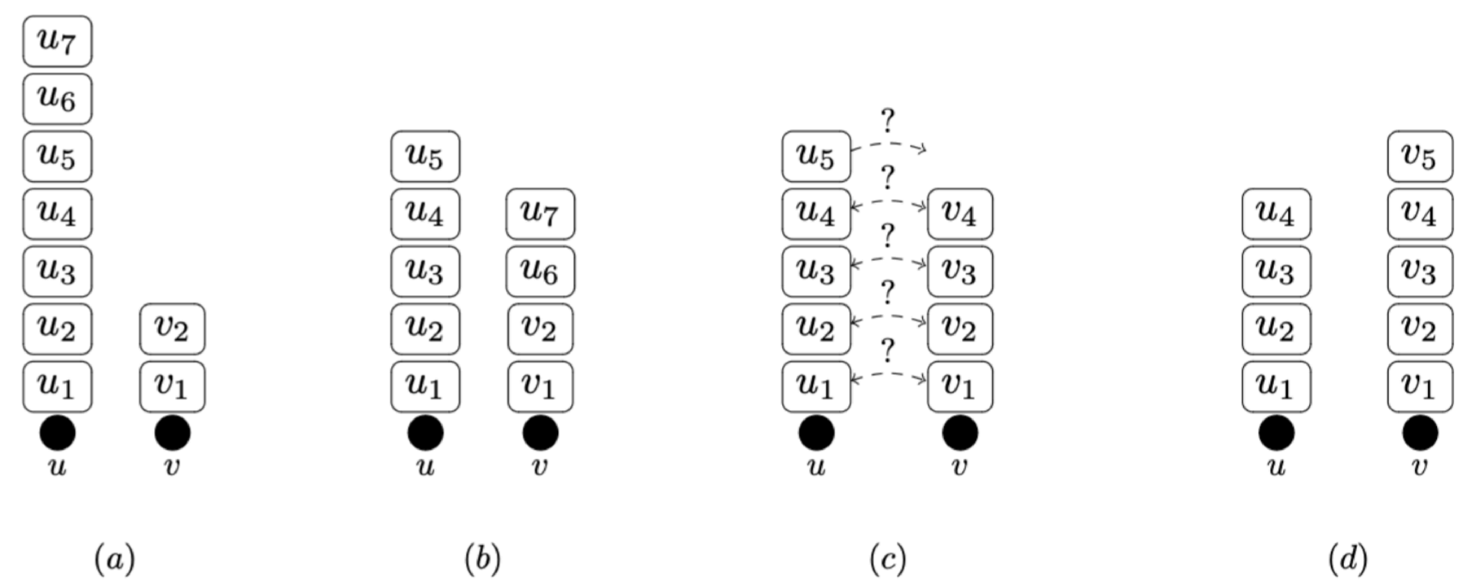

Nodes u and v balance their load. The tokens on u are of height $1, 2, 3, 4, 5, 6, 7$ and the tokens on v are of height $1, 2$. In the moving step, u sends $\lfloor (7 - 2)/2 \rfloor$ top tokens to v ((b)). Then, any token on u may be randomly swapped with its buddy with probability $1/2$ ((c)). The result of this is shown in (d).

Lemma:

Assume there are $(1 - \epsilon) \cdot n$ tokens with height at least c, for some positive constant ϵ. Then after $O(\tau_{BC}(K))$ rounds, no token has height $c + 2$ w.p. $1 - \exp(\epsilon \cdot \log(n)/5 \log\log n)$.

O(1) Discrepancy:

- We split the analysis into 3 parts.

1. There are $12n$ excess tokens in round $\tau_{BC}(K)$ and the maximum height is 24.
   - Each node has $O(\sqrt{\log n})$ excess tokens.
   - The number of excess tokens, for a fixed node, is double exponentially decreasing.
2. Then, after another $O(\tau_{BC}(K))$ round, there are at most $1.8n$ excess tokens.
   - We show as long as there are at least $1.8n$ excess tokens, after one $\tau_{BC}(K)$ rounds, $\Omega(n)$ excess tokens disappear.
3. After another $O(\tau_{BC}(K))$ rounds, the discrepancy is 6 at most.
   - The number of tokens at maximum height is $o(n)$.
   - After another $\tau_{BC}(K)$ rounds, all tokens at maximum height disappear.
   - It works for height at least 4.

References:

[1] Y. Rabani, A. Sinclair, and R. Wanka, "Local divergence of markov chains and the analysis of iterative load-balancing schemes," in *Proceedings of the 39th Annual Symposium on Foundations of Computer Science*, ser. FOCS '98, USA: IEEE Computer Society, 1998, p. 694, ISBN: 0818691727.

[2] T. Sauerwald and H. Sun, "Tight bounds for randomized load balancing on arbitrary network topologies," in *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, 2012, pp. 341–350. DOI: 10.1109/FOCS.2012.86.