Ergodic Effects in Token Circulation

Adrian Kosowski
Inria Paris
adrian.kosowski@inria.fr

Przemysław Uznański
ETH Zürich
przemyslaw.uznanski@inf.ethz.ch

Abstract
We present an *extremely simple* local rule of token propagation with optimal $\widetilde{\Theta}(m/k)$ idle time, for a wide range of parameters, after an initial grace period. Presented at SODA 2018.

Distributed Token Propagation

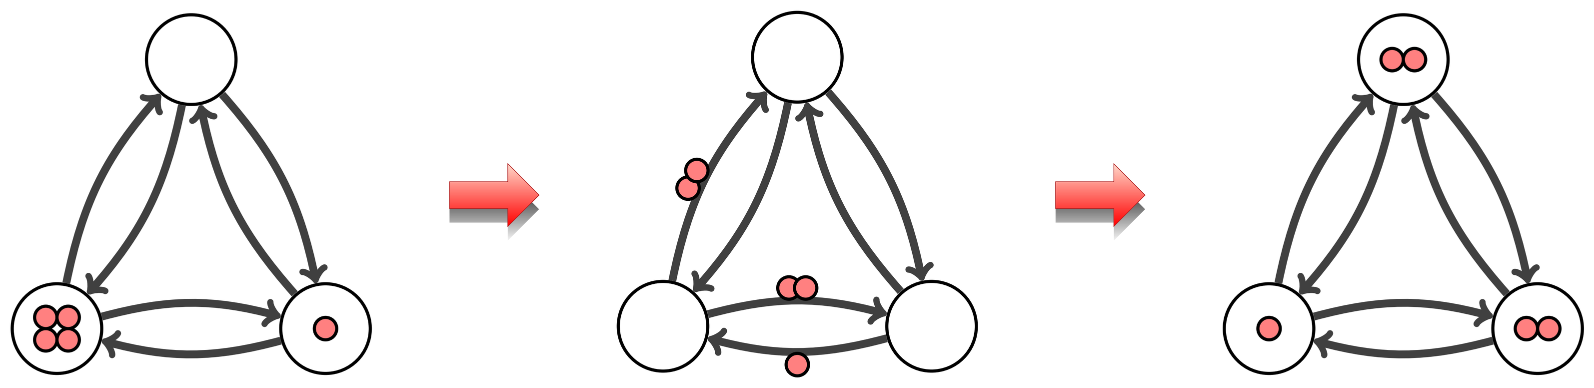

Notation

n	total number of **vertices**	k	total number of **tokens**
m	total number of **edges**	$L_t(v)$	**load of vertex** at step t
$2m$	total number of **arcs**	$L_t(e)$	load **traversing arc** e after step t

Randomized Schemes

Random walk diffusion: $\mathbb{E}[L_t(e)] = \frac{L_t(v)}{\deg(v)}$

Cover time bounds: [AAK+11, ER09, ES11].

Randomized rounding: $|L_t(e) - \frac{L_t(v)}{\deg(v)}| \le 1$

Load balancing bounds: [RSW98, SS12, ABS12, FS09, FGS10].

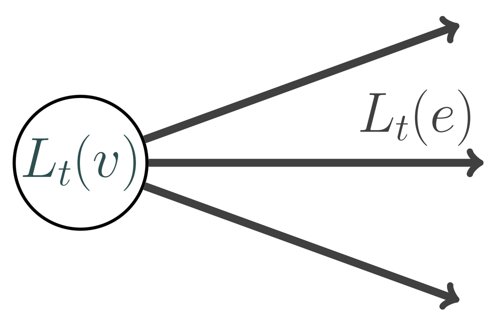

Focus: Deterministic Scheme

RR dynamics (round-robin, rotor-router, eulerian walker, ant walk or Propp machine), introduced in [PDDK96]

Satisfies stronger property:

$$\left|\sum_{t\le T} L_t(e) - \frac{\sum_{t\le T} L_t(v)}{\deg(v)}\right| \le 1$$

Rotor-router:
While there is a token at node v, **do**:
1. Send the token to $pointer_v$,
2. Set $pointer_v = next(pointer_v)$.

Deterministic approach almost as good as randomized approaches for *cover time* [DKPU14, KKPS13, KP14] and *load balancing* [DF09, AB13, KP14, BKK+15, SYKY16].

Idle time

Deterministic token propagation processes are useful beyond load balancing. We look at long-term averaging properties. Specifically, we are interested in bounding the idle time after some *initialization time* T_{init}.

Definition. Idle time *of a token propagation scheme is the smallest value T such that in any consecutive T steps every edge/arc is traversed.*

References

[AAK+11] N. Alon, C. Avin, M. Koucký, G. Kozma, Z. Lotker, and M. R. Tuttle. Many random walks are faster than one. *Combinatorics, Probability & Computing*, 20(4):481–502, 2011.

[AB13] H. Akbari and P. Berenbrink. Parallel rotor walks on finite graphs and applications in discrete load balancing. In *SPAA*, pages 186–195, 2013.

[ABS12] H. Akbari, P. Berenbrink, and T. Sauerwald. A simple approach for adapting continuous load balancing processes to discrete settings. In *PODC*, pages 271–280, 2012.

[BKK+15] P. Berenbrink, R. Klasing, A. Kosowski, F. Mallmann-Trenn, and P. Uznański. Improved analysis of deterministic load-balancing schemes. In *PODC*, pages 301–310, 2015.

[CDG+15] J. Chalopin, S. Das, P. Gawrychowski, A. Kosowski, A. Labourel, and P. Uznański. Limit behavior of the multi-agent rotor-router system. In *DISC*, pages 123–139, 2015.

[DF09] B. Doerr and T. Friedrich. Deterministic random walks on the two-dimensional grid. *Combinatorics, Probability and Computing*, 18(1-2):123–144, 2009.

[DKPU14] D. Dereniowski, A. Kosowski, D. Pająk, and P. Uznański. Bounds on the cover time of parallel rotor walks. In *STACS*, pages 263–275, 2014.

[ER09] K. Efremenko and O. Reingold. How well do random walks parallelize? In *APPROX-RANDOM*, pages 476–489, 2009.

[ES11] R. Elsässer and T. Sauerwald. Tight bounds for the cover time of multiple random walks. *Theor. Comput. Sci.*, 412(24):2623–2641, 2011.

[FGS10] T. Friedrich, M. Gairing, and T. Sauerwald. Quasirandom load balancing. In *SODA*, pages 1620–1629, 2010.

[FS09] T. Friedrich and T. Sauerwald. Near-perfect load balancing by randomized rounding. In *STOC*, pages 121–130, 2009.

[KKPS13] R. Klasing, A. Kosowski, D. Pająk, and T. Sauerwald. The multi-agent rotor-router on the ring: a deterministic alternative to parallel random walks. In *PODC*, pages 365–374, 2013.

[KP14] A. Kosowski and D. Pająk. A case study of cover time for the rotor-router. In *ICALP*, pages 544–555, 2014.

[PDDK96] V.B. Priezzhev, D. Dhar, A. Dhar, and S. Krishnamurthy. Eulerian walkers as a model of self-organized criticality. *Phys. Rev. Lett.*, 77(25):5079–5082, Dec 1996.

[RSW98] Y. Rabani, A. Sinclair, and R. Wanka. Local divergence of markov chains and the analysis of iterative load-balancing schemes. In *FOCS*, pages 694–703, 1998.

[SS12] T. Sauerwald and H. Sun. Tight bounds for randomized load balancing on arbitrary network topologies. In *FOCS*, pages 341–350, 2012.

[SYKY16] T. Shiraga, Y. Yamauchi, S. Kijima, and M. Yamashita. Total variation discrepancy of deterministic random walks for ergodic markov chains. In *ANALCO*, pages 138–148, 2016.

[TV06] T. Tao and V. H. Vu. *Additive Combinatorics*. Cambridge University Press, 2006. Cambridge Books Online.

Issue

Random-walk type processes tend to behave poorly w.r.t. idle time.

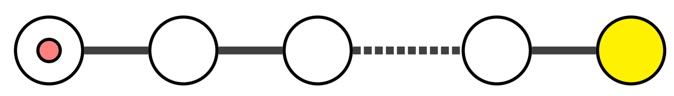

For example, **idle time** of a path (with a single token) is $\Omega(n^2)$ for the random walk, but $\Theta(n)$ for RR dynamics.

Main result

Theorem. *For time $t \ge T_{init} = poly(n, \log k)$ following bounds hold on idle time of RR:*
- $\widetilde{\mathcal{O}}(\gcd(k, 2m)\frac{m}{k})$
- $\mathcal{O}(\frac{m}{k}) = \mathcal{O}(1)$ *for* $k \ge (\frac{1}{2} + \varepsilon)m$
- $\mathcal{O}(Diam \cdot \frac{m}{k})$
- $\widetilde{\mathcal{O}}(\sqrt{n} \cdot \frac{m}{k})$
- $\widetilde{\mathcal{O}}(\sqrt{k} \cdot \frac{m}{k})$
- $\mathcal{O}(\frac{m}{k})$ *for trees*

Techniques

Eulerian circulation

Theorem ([CDG+15]). *Recurrent state of RR $\Leftrightarrow$ there is a bijection $\varphi : \vec{E} \to \vec{E}$ such that*
- $\varphi(e)$ *always starts where e ends*
- $L_{t+1}(\varphi(e)) = L_t(e)$

Recurrent state is reached in time $poly(n, \log k)$.

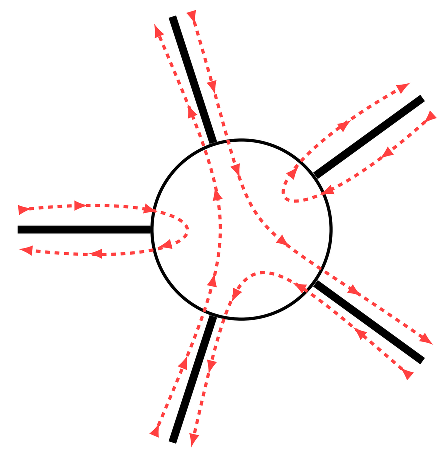

The gcd connection

$$\lim_{T\to\infty} \frac{\sum_{t\le T} L_t(e)}{T} = \frac{k}{2m}$$

therefore

$$\frac{\text{\# tokens on a cycle}}{\text{\# arcs on a cycle}} = \frac{k}{2m}$$

$\#\text{cycles} \mid \gcd(k, 2m)$

$\gcd(k, 2m) = 1 \implies$ single cycle

Skewed distances

$\delta_0(e_1, e_2) = 1$
$\delta_1(e, \varphi(e)) = 0$

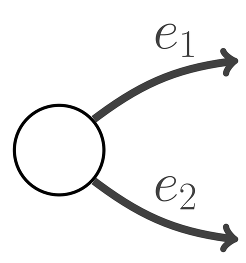

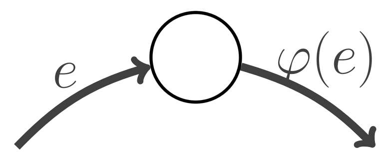

$$\delta_{\Delta t}(e, e') \stackrel{\text{def}}{=} \max_{t_1, t_2} \left| \sum_{t_1 \le t \le t_2} \Big(L_{t+\Delta t}(e_1) - L_t(e') \Big) \right|$$

Self-intersections

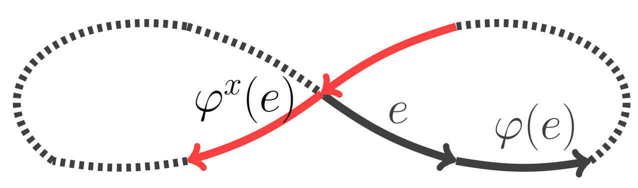

Definition. *x is a **self-intersection** of the cycle: for some e, e and $\varphi^x(e)$ share starting vertex.*

Discrepancy parameter δ:

$$\delta_x \stackrel{\text{def}}{=} \delta_x(e, e) = \delta_0(e, \varphi^x(e))$$

- $\delta_0 = 0$
- $\delta_x \le 1$ for x self-intersections
- $\delta_{x+y} \le \delta_x + \delta_y$

$$\left| \sum_{t\le T} L_t(e) - \frac{k}{2m} T \right| \ge \max_{0\le x<2m} \delta_x$$

Idle time$(e) = \mathcal{O}\left(\frac{m}{k}\right) \cdot \max_{0\le x<2m} \delta_x$

Additive combinatorics

Sumset:
For $A, B \subseteq \mathbb{Z}_{2m}$, $A + B \stackrel{\text{def}}{=} \{a + b : a \in A, b \in B\}$.

Multiplication:
$\kappa \cdot A \stackrel{\text{def}}{=} \underbrace{A + A + \ldots + A}_{\kappa \text{ times}}$.

$\mathcal{X}$ is the set of all self-intersections. What is the smallest κ, such that $\kappa \cdot \mathcal{X} = \mathbb{Z}_{2m}$?

Spectral property

Lemma.

$$\forall_{f \ge 1} \quad \exists_{x\in\mathcal{X}} \quad (f \cdot x) \in \left[\frac{2}{3}m, \frac{4}{3}m\right]$$

Proof:

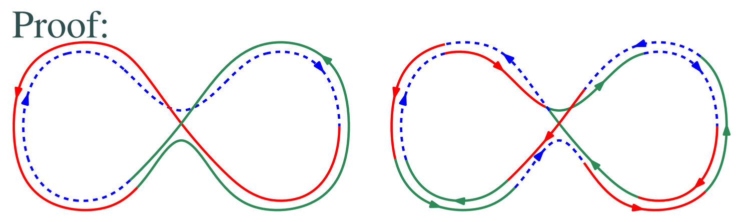

Bohr sets

Restating: $\forall_{f\ge 1} \quad \exists_{x\in\mathcal{X}} \quad (f \cdot x) \in \left[\frac{2}{3}m, \frac{4}{3}m\right] \quad \Leftrightarrow \quad \text{Bohr}(\mathcal{X}, 1/6) = \{0\}$

Theorem. *For any $A \subseteq \mathbb{Z}_{2m}$, if $Bohr(A, 1/6) = \{0\}$, then $\kappa \cdot X = \mathbb{Z}_{2m}$ for some $\kappa = \mathcal{O}(\log^2 m)$.*

Proof: [main technical contribution, generalizes [TV06] Proposition 4.40]

$\Downarrow$

$$\max_{0\le x<2m} \delta_x = \mathcal{O}(\log^2 m)$$