Construction of Lyapunov functions via relative entropy with application to caching

Nicolas Gast[1]

ACM MAMA 2016, Antibes, France

[1] Inria

Outline

1. Why?
2. How to make the fixed point method work (sufficient condition)
3. What: application to caching policy
4. Conclusion

State space explosion and mean-field method

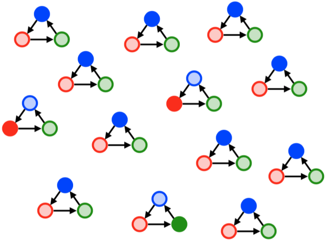

$3^{13} \approx 10^6$ states.

We need to keep track

$$\mathbb{P}(X_1(t) = i_1, \dots, X_n(t) = i_n)$$

State space explosion and mean-field method

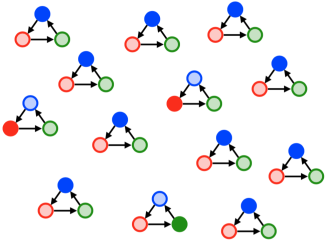

We need to keep track

$$\mathbb{P}(X_1(t) = i_1, \ldots, X_n(t) = i_n)$$

$3^{13} \approx 10^6$ states.

The decoupling assumption is

$$\mathbb{P}(X_1(t) = i_1, \ldots, X_n(t) = i_n) \approx \mathbb{P}(X_1(t) = i_1) \ldots \mathbb{P}(X_n(t) = i_n)$$

Problem: is this valid?

Decoupling assumption: (always) valid in transient regime

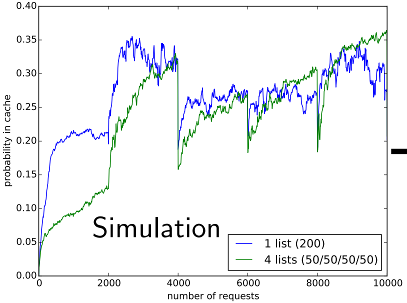

→

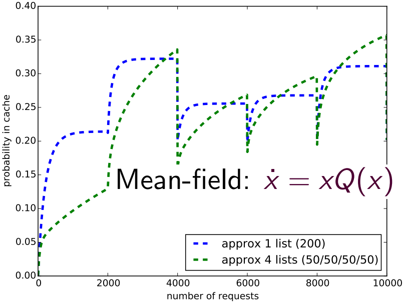

Decoupling assumption: (always) valid in transient regime

Theorem (Kurtz (70'), Benaim, Le Boudec (08),...)

For many systems and any fixed t, if $x \mapsto xQ(x)$ is Lipschitz-continuous then, as the number of objects N goes to infinity:

$$\lim_{N\to\infty} \mathbb{P}(X_k(t) = i) = x_{k,i}(t),$$

where x satisfies $\dot{x} = xQ(x)$.

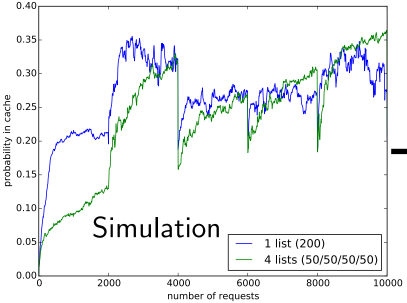

→

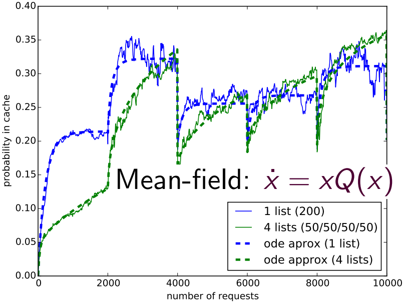

The fixed point method

We know that $x_i(t) \approx \mathbb{P}(X(t) = i)$ satisfies $\dot{x} = xQ(x)$.

Does $\mathbb{P}(X = i)$ satisfies $xQ(x) = 0$?

Method was used in many papers:

- Bianchi 00[2]
- Ramaiyan et al. 08[3]
- Kwak et al. 05[4]
- Kumar et al 08[5]

---

[2] Performance analysis of the IEEE 802.11 distributed coordination function. – G. Bianchi. – IEEE J. Select. Areas Commun. 2000.

[3] Fixed point analys is of single cell IEEE 802.11e WLANs: Uniqueness, multistability. – V. Ramaiyan, A. Kumar, and E. Altman. – ACM/IEEE Trans. Networking. Oct. 2008.

[4] Performance analysis of exponenetial backoff. – B.-J. Kwak, N.-O. Song, and L. Miller. – ACM/IEEE Trans. Networking. 2005.

[5] New insights from a fixed-point analysis of single cell IEEE 802.11 WLANs. – A. Kumar, E. Altman, D. Miorandi, and M. Goyal. – ACM/IEEE Trans. Networking 2007

It does not always work[6][7]

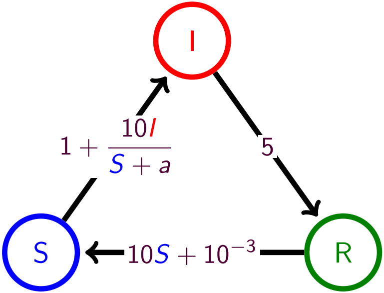

- Markov chain is irreducible.
- Unique fixed point $xQ(x) = 0$.

[6] Benaim Le Boudec 08

[7] Cho, Le Boudec, Jiang, On the Asymptotic Validity of the Decoupling Assumption for Analyzing 802.11 MAC Protoco. 2010

It does not always work[6][7]

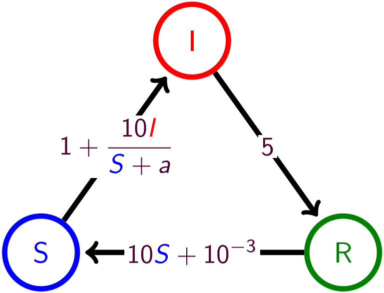

- Markov chain is irreducible.
- Unique fixed point $xQ(x) = 0$.

	Fixed point $xQ(x) = 0$		Stat. measure $N = 1000$	
	x_S	x_I	π_S	π_I
$a = .3$	0.209	0.234	0.209	0.234

[6] Benaim Le Boudec 08

[7] Cho, Le Boudec, Jiang, On the Asymptotic Validity of the Decoupling Assumption for Analyzing 802.11 MAC Protoco. 2010

It does not always work[6][7]

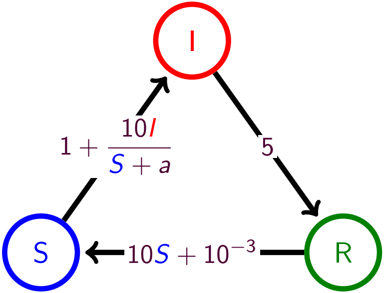

- Markov chain is irreducible.
- Unique fixed point $xQ(x) = 0$.

	Fixed point $xQ(x) = 0$		Stat. measure $N = 1000$	
	x_S	x_I	π_S	π_I
$a = .3$	0.209	0.234	0.209	0.234
$a = .1$	0.078	0.126	0.11	0.13

[6]Benaim Le Boudec 08

[7]Cho, Le Boudec, Jiang, On the Asymptotic Validity of the Decoupling Assumption for Analyzing 802.11 MAC Protoco. 2010

It does not always work

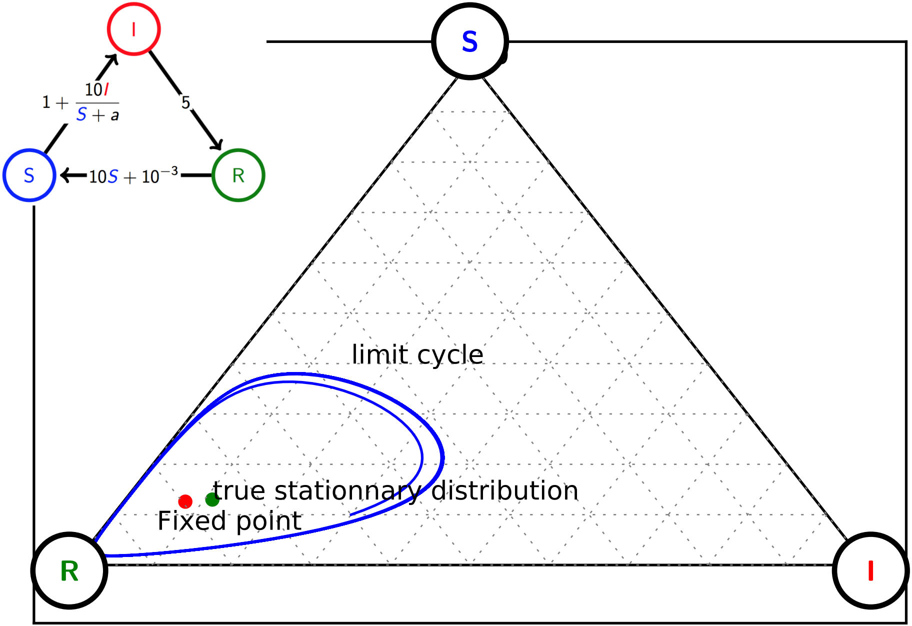

Outline

1. Why?
2. How to make the fixed point method work (sufficient condition)
3. What: application to caching policy
4. Conclusion

Outline

Link between the decoupling assumption and $\dot{x} = xQ(x)$

$$\mathbb{P}(X_1(t) = i_1, \dots, X_n(t) = i_n) \approx \underbrace{\mathbb{P}(X_1(t) = i_1)}_{=x_{1,i_1}(t)} \dots \underbrace{\mathbb{P}(X_n(t) = i_n)}_{=x_{n,i_n}(t)}$$

When we zoom on one object

$$\mathbb{P}(X_1(t+dt) = j | X_1(t) = i) \approx \mathbb{E}\left[\mathbb{P}(X_1(t) = j | X_1 = i \wedge X_2 \dots X_n)\right]$$
$$\approx Q^{(1)}_{i,j}(\mathbf{x}) := \sum_{i_2 \dots i_n} K_{(i,i_2\dots i_n)\to(j,j_2\dots j_n)} x_{2,i_2} \dots x_{n,i_n}$$

We then get: $\frac{d}{dt} x_{1,j}(t) \approx \sum_i x_{1,i} Q^{(1)}_{i,j}(\mathbf{x})$

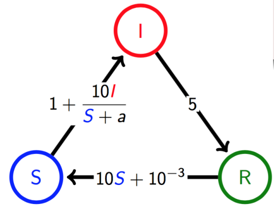

Exchangeability of limits

Markov chain

Transient regime $\dot{p} = pK$

$\Big\downarrow$ $t \to \infty$

Stationary $\pi K = 0$

Exchangeability of limits

	Markov chain		Mean-field
Transient regime	$\dot{p} = pK$	$N \to \infty$ →	$\dot{x} = xQ(x)$
	↓ $t \to \infty$		
Stationary	$\pi K = 0$	? →	$xQ(x) = 0$ fixed points

Exchangeability of limits

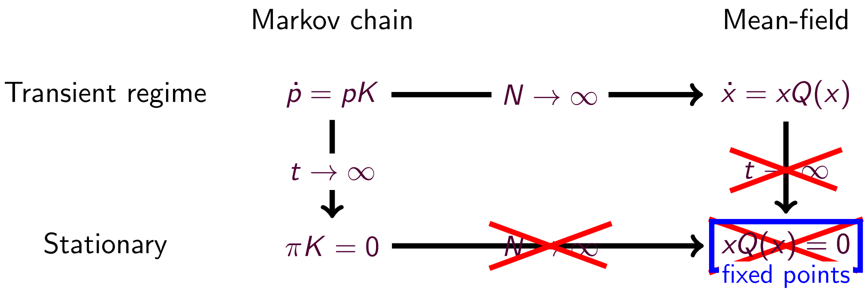

Exchangeability of limits

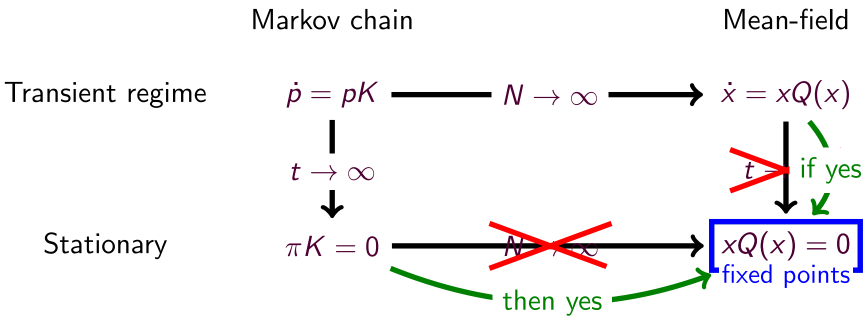

Theorem ((i) Benaim Le Boudec 08,(ii) Le Boudec 12)

The stationary distribution π^N concentrates on the fixed points if :

(i) *All trajectories of the ODE converges to the fixed points.*

(ii) *(or) The markov chain is reversible.*

Lyapunov functions

A solution of $\frac{d}{dt}x(t) = xQ(x(t))$ converges to the fixed points of $xQ(x) = 0$, if there exists a Lyapunov function f, that is:

- Lower bounded: $\inf_x f(x) > +\infty$
- Decreasing along trajectories:

$$\frac{d}{dt}f(x(t)) < 0,$$

  whenever $x(t)Q(x(t)) \neq 0$.

Lyapunov functions

A solution of $\frac{d}{dt}x(t) = xQ(x(t))$ converges to the fixed points of $xQ(x) = 0$, if there exists a Lyapunov function f, that is:

- Lower bounded: $\inf_x f(x) > +\infty$
- Decreasing along trajectories:

$$\frac{d}{dt}f(x(t)) < 0,$$

whenever $x(t)Q(x(t)) \neq 0$.

How to find a Lyapnuov function

- Energy? Distance? Entropy? Luck?

The relative entropy is a Lyapunov function for Markov chains

Let Q be the generator of an irreducible Markov chain and π be its stationary distribution. Let $P(t)$ be the solution of $\frac{d}{dt}P(t) = P(t)Q$.

Theorem (e.g. Budhiraja et al 15, Dupuis-Fischer 11)

The relative entropy

$$R(P\|\pi) = \sum_i P_i \log \frac{P_i}{\pi_i}$$

is a Lyapunov function:

$$\frac{d}{dt}R(P(t)\|\pi) < 0,$$

with equality if and only if $P(t) = \pi$.

Relative entropy for mean-field models

Assume that $Q(x)$ be a generator of an irreducible Markov chain and let $\pi(x)$ be its stationary distribution. Let $P(t)$ be the solution of $\frac{d}{dt}P(t) = P(t)Q(P(t))$. Then

$$\frac{d}{dt}R(P(t)\|\pi(t)) = \underbrace{\frac{d}{dt}P(t)\frac{\partial}{\partial P}R(P(t),\pi(t))}_{\le 0} + \underbrace{\frac{d}{dt}\pi(t)\frac{\partial}{\partial \pi}R(P(t),\pi(t))}_{=-\sum_i x_i(t)\frac{d}{dt}\log \pi_i(t)}$$

$$\le -\sum_i x_i(t)\frac{d}{dt}\log \pi_i(t)$$

Relative entropy for mean-field models

Assume that $Q(x)$ be a generator of an irreducible Markov chain and let $\pi(x)$ be its stationary distribution. Let $P(t)$ be the solution of $\frac{d}{dt}P(t) = P(t)Q(P(t))$. Then

$$\frac{d}{dt}R(P(t)\|\pi(t)) = \underbrace{\frac{d}{dt}P(t)\frac{\partial}{\partial P}R(P(t),\pi(t))}_{\leq 0} + \underbrace{\frac{d}{dt}\pi(t)\frac{\partial}{\partial \pi}R(P(t),\pi(t))}_{=-\sum_i x_i(t)\frac{d}{dt}\log \pi_i(t)}$$

$$\leq -\sum_i x_i(t)\frac{d}{dt}\log \pi_i(t)$$

Theorem

If there exists a lower bounded integral $F(x)$ of $-\sum_i x_i(t)\frac{d}{dt}\log \pi_i(t)$, then $x \mapsto R(x\|\pi(x)) + F(x)$ is a Lyapunov function for the mean-field model.

Outline

I consider a cache (virtually) divided into lists

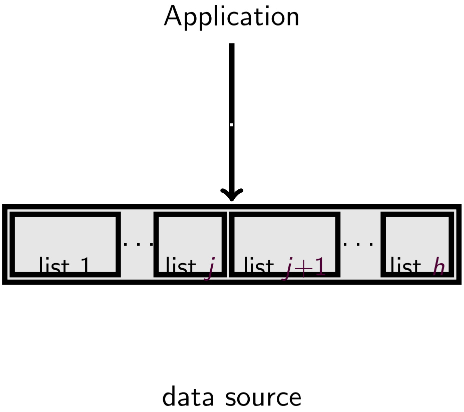

IRM Probability request p_i

RAND Upon hit/miss: Exchanged with random from next list.

I consider a cache (virtually) divided into lists

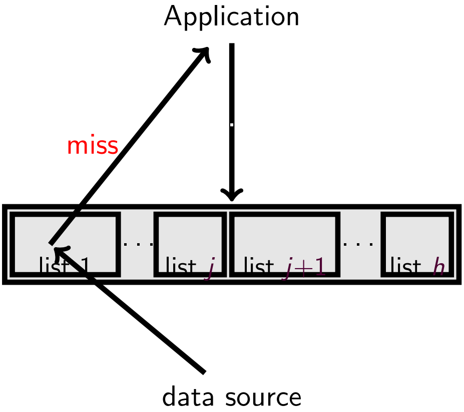

IRM Probability request p_i

RAND Upon hit/miss: Exchanged with random from next list.

I consider a cache (virtually) divided into lists

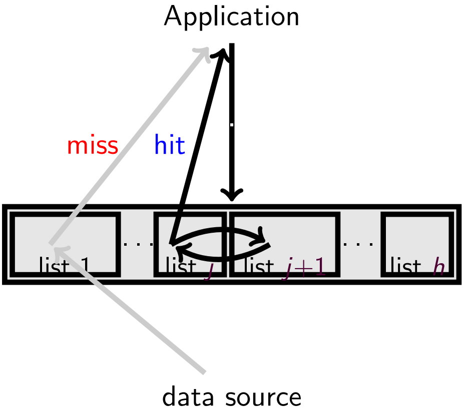

IRM Probability request p_i

RAND Upon hit/miss: Exchanged with random from next list.

We construct the ODE by assuming independence

Let $H_i(t)$ be the popularity in list i.

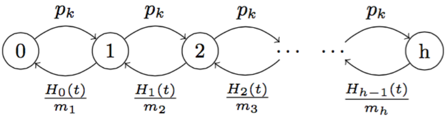

We construct the ODE by assuming independence

Let $H_i(t)$ be the popularity in list i.

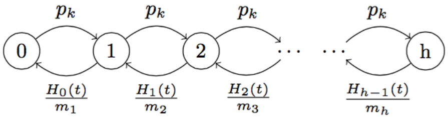

If $x_{k,i}(t)$ is the probability that item k is in list i at time t:

Popularity in cache i-1

$$\dot{x}_{k,i}(t) = p_k x_{k,i-1}(t) - \sum_j p_j x_{j,i-1}(t) \frac{x_{k,i}(t)}{m_i} + \mathbf{1}_{\{i<h\}} \left(\sum_j p_j x_{j,i}(t) \frac{x_{k,i+1}(t)}{m_{i+1}} - p_k x_{k,i}(t) \right)$$

Popularity in cache i

ODE of the type $\dot{x} = xQ(x)$.

Transient regime: this approximation is accurate

THEOREM 6. *For any $T > 0$, there exists a constant $C > 0$ that depends on T such that, for any probability distribution over n items and list sizes $m_1 \dots m_h$, we have:*

$$\mathbf{E}\left[\sup_{t\in\{0\dots\tau\}, i\in\{0\dots h\}} |H_i(t) - \delta_i(t)|\right] \le C\sqrt{\max_{k=1}^{n} p_k + \max_{i=0}^{h} \frac{1}{m_i}},$$

where $\tau := \lceil T/(\max_{k=1}^{n} p_k + \max_{i=0}^{h} \frac{1}{m_i})\rceil$.

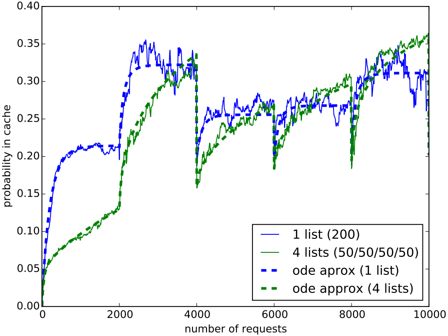

Stationary distribution: uniqueness of the fixed point

THEOREM 7. *The mean-field model (8) has a unique fixed point. For this fixed point, the probability that item k is part of list i, for $k = 1, \dots, n$ and $i = 0, \dots, h$, is given by*

$$x_{k,i} = \frac{p_k^i z_i}{1 + \sum_{j=1}^h p_k^j z_j},$$

where $\mathbf{z} = (z_1, \dots, z_h)$ is the unique solution of the equation

$$\sum_{k=1}^n \frac{p_k^i z_i}{1 + \sum_{j=1}^h p_k^j z_j} = m_i. \tag{14}$$

Stationary distribution: uniqueness of the fixed point

THEOREM 7. *The mean-field model (8) has a unique fixed point. For this fixed point, the probability that item k is part of list i, for $k = 1, \dots, n$ and $i = 0, \dots, h$, is given by*

$$x_{k,i} = \frac{p_k^i z_i}{1 + \sum_{j=1}^h p_k^j z_j},$$

where $\mathbf{z} = (z_1, \dots, z_h)$ is the unique solution of the equation

$$\sum_{k=1}^n \frac{p_k^i z_i}{1 + \sum_{j=1}^h p_k^j z_j} = m_i. \qquad (14)$$

- By simulation: very accurate

m_1	m_2	m_3	m_4	exact	mean field
2	2	96	–	0.3166	0.3169
10	30	60	–	0.3296	0.3299
20	2	78	–	0.3273	0.3276
90	8	2	–	0.4094	0.4100
1	4	10	85	0.3039	0.3041
5	15	25	55	0.3136	0.3139
25	25	25	25	0.3345	0.3348
60	2	2	36	0.3514	0.3517

Relative entropy for the caching model

The stationary measure $\pi_{k,i}$ satisfy: $\pi_{k,i}(\mathbf{x}) = \dfrac{\prod_{j=1}^{i-1} \lambda_{k,j}/\mu_j(\mathbf{x})}{\sum_{j'=1}^{h} \prod_{j=1}^{j'-1} \lambda_{k,j}/\mu_j(\mathbf{x})}$.

$$\sum_{i=1}^{h}\sum_{k=1}^{n} x_{k,i}(t)\frac{d}{dt}\log(\pi_{k,i}(\mathbf{x}(t)))$$

$$= \underbrace{\sum_{i=1}^{h}\sum_{k=1}^{n} x_{k,i}(t)}_{=m_i} \frac{d}{dt}\log\left(\prod_{j=1}^{i-1}\mu_j(\mathbf{x})\right)$$

$$\underbrace{\sum_{k=1}^{n}\sum_{i=1}^{h} x_{k,i}(t)}_{=1} \frac{d}{dt}\log\left(\sum_{j'=1}^{h}\prod_{j=1}^{j'-1}\lambda_{k,j}/\mu_j(\mathbf{x})\right)$$

- Very similar to (Fricker,G. 14), (Fricker et al. 12) (Tibi 11).

Outline

Conclusion

- Decoupling assumption: OK in transient.
- The fixed point method is not always valid. We need either:
  - Reversibility
  - Lyapunov function
- To find Lyapunov functions: we need problem-specific.
  - Physics: energy.
  - Markov chains: relative entropy (since it it decrease along trajectories)

Yet... the method is not robust (*e.g.*: non-IRM, LRU instead of RAND)

Thank you!

http://mescal.imag.fr/membres/nicolas.gast

nicolas.gast@inria.fr

G. Van Houdt 15 — *Transient and Steady-state Regime of a Family of List-based Cache Replacement Algorithms.*, Gast, Van Houdt., ACM Sigmetrics 2015

G. 16 — *Construction of Lyapunov functions via relative entropy with application to caching*, Gast, N., ACM MAMA 2016

Benaïm, Le Boudec 08 — *A class of mean field interaction models for computer and communication systems*, M.Benaïm and J.Y. Le Boudec., Performance evaluation, 2008.

Le Boudec 10 — *The stationary behaviour of fluid limits of reversible processes is concentrated on stationary points.*, J.-Y. L. Boudec. , Arxiv:1009.5021, 2010

Budhiraja et al. 15 — *Limits of relative entropies associated with weakly interacting particle systems.*, A. S. Budhiraja, P. Dupuis, M. Fischer, and K. Ramanan. , Electronic journal of probability, 20, 2015.

Fricker-Gast 14 — *Incentives and redistribution in homogeneous bike-sharing systems with stations of finite capacity.*, C. Fricker and N. Gast. , Euro journal on transportation and logistics:1-31, 2014.

Fricket et al. 13 — *Mean field analysis for inhomogeneous bike sharing systems*, Fricker, Gast, Mohamed, Discrete Mathematics and Theoretical Computer Science DMTCS

Outline

Increasing the number of lists is not always better[8]

The scheme seems to sort the number of items from least popular to most popular:

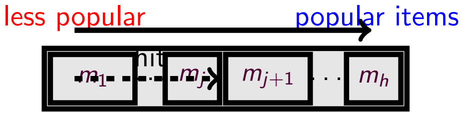

$?\geq?$

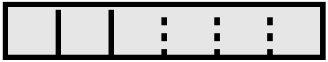

Six lists: $\mathbf{m} = (1, 1, 1, 1, 1, 1)$ Three lists: $\mathbf{m} = (1, 1, 4)$.

[8] contrary to the conjecture of O. I. Aven, E. G. Coffman, Jr., and Y. A. Kogan. Stochastic Analysis of Computer Storage. Kluwer Academic Publishers, Norwell, MA, USA, 1987.

Increasing the number of lists is not always better

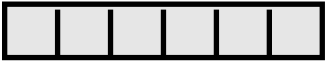

Six lists: $\mathbf{m} = (1, 1, 1, 1, 1, 1)$

? $\geq$ **?**

Three lists: $\mathbf{m} = (1, 1, 4)$.

policy	**m**	$M(\mathbf{m})$	lower bound
Optimal	RAND(1,1,4)	0.005284	0.004925
	RAND(1,1,3,1)	0.005299	0.004884
	RAND(1,1,2,2)	0.005317	0.004884
	RAND(1,1,2,1,1)	0.005321	0.004879
	RAND(1,1,1,3)	0.005338	0.004884
	RAND(1,1,1,2,1)	0.005343	0.004879
	RAND(1,1,1,1,2)	0.005347	0.004879
CLIMB	RAND(1,1,1,1,1,1)	0.005348	0.004878
	RAND(1,2,3)	0.005428	0.004925
	RAND(1,2,2,1)	0.005439	0.004884
LRU	LRU(6)	0.005880	–
RANDOM	RAND(6)	0.015350	0.015350

Table 1: CLIMB is not optimal for IRM model: $p = (49, 49, 49, 49, 7, 1, 1)/205$ **and** $m = 6$**.**

Having 3 lists of sizes $(1, 1, 4)$ is better than 6 lists of size 1. The same holds for the mean-field approximation.