

Performance analysis methods for list-based caches with non-uniform access (Supplementary)

Giuliano Casale, Nicolas Gast

Abstract—This document includes supplemental material for the paper “Performance analysis methods for list-based caches with non-uniform access”.

SUPPLEMENTAL MATERIAL

D. Proof of product-form solution for FIFO-C($\mathbf{m}$)

The global balance equations are

$$\left(\sum_{j=0}^{h-1} \sum_{\substack{j'=1..h \\ j' \neq j}} \sum_{k=1}^{m_j} \left(\sum_{v=1}^u \lambda_{vs(k,j)}(j) c_{vs(k,j)}(j, j') \right) \right) \pi(\mathbf{s}) \\ = \sum_{j=0}^{h-1} \sum_{\substack{j'=1..h \\ j' \neq j}} \sum_{k=1}^{m_j} \left(\sum_{v=1}^u \lambda_{vs(1,j')}(j) c_{vs(1,j')}(j, j') \right) \pi(\mathbf{s}_{kjj'}) \quad (1)$$

where $\mathbf{s}$ is reachable from the initial state and where $\mathbf{s} \equiv s(i, j)$ and $\mathbf{s}$ is obtained from $\mathbf{s}_{kjj'} \in \mathcal{S}$ by promoting item $s(1, j')$, assumed in $\mathbf{s}_{kjj'}$ at position k of list j , to the front of list $j' \neq j$. We now note that the innermost brackets in the global balance equations if positive may be seen as ratios of access factors, e.g., if all the terms are positive

$$\sum_{j=0}^{h-1} \sum_{\substack{j'=1..h \\ j' \neq j}} \sum_{k=1}^{m_j} \frac{\gamma_{s(k,j)j'}}{\gamma_{s(k,j)j}} \pi(\mathbf{s}) = \sum_{j=0}^{h-1} \sum_{\substack{j'=1..h \\ j' \neq j}} \sum_{k=1}^{m_j} \frac{\gamma_{s(1,j')j'}}{\gamma_{s(1,j')j}} \pi(\mathbf{s}_{kjj'}) \quad (2)$$

We also observe that the product-form in (1) satisfies

$$\pi(\mathbf{s}_{kjj'}) = \frac{\gamma_{s(1,j')j'} \gamma_{s(k,j)j'}}{\gamma_{s(1,j')j} \gamma_{s(k,j)j}} \pi(\mathbf{s}) \quad (3)$$

Plugging (3) into (2) the expression reduces to an identity. This holds also if some of the innermost brackets in (1) are zero, since the terms are neglected from the balance equation.

E. Additional recurrence relations

We develop here additional recurrence relations for the normalizing constant $E(\mathbf{m})$ and for the quantities $\xi_l(\mathbf{m})$.

G. Casale is with the Department of Computing at Imperial College London, UK; e-mail: (g.casale@imperial.ac.uk).

N. Gast is with Inria Grenoble, France; e-mail: (nicolas.gast@inria.fr).

The authors would like to thank Benny van Houdt for helpful comments during preparation of this manuscript. The work of G.C. was supported by the EU H2020 grants DICE (644869) and RADON (825040) and by a UK EPSRC grant (EP/L00738X/1, iBids). The work of N.G. was supported by the ANR project REFINO (ANR-19-CE23-0015). Research data is available at <http://doi.org/10.5281/zenodo.1134886> (CC BY 4.0 license).

1) $O(n^m)$ recurrence relation for $E(\mathbf{m})$: Plugging (2) into the condition $\sum_i \pi_{il}(\mathbf{m}) = m_l$ yields the recurrence relation

$$E(\mathbf{m}) = \sum_{i=1}^n \gamma_{il} E_i(\mathbf{m} - 1_l) \quad (4)$$

for all $l = 1, \dots, h$. The right-hand side may be seen as summing over all the (un-normalized) conditional probabilities that item i is in position m_l in list l . Further, summing (3) over $k = 1, \dots, n$ we get

$$\sum_{k=1}^n E(\mathbf{m}) = \sum_{k=1}^n E_k(\mathbf{m}) + \sum_{j=1}^h m_j \sum_{k=1}^n \gamma_{kj} E_k(\mathbf{m} - 1_j)$$

Finally, using (4) on the last summation yields

$$(n - m)E(\mathbf{m}) = \sum_{i=1}^n E_i(\mathbf{m}) \quad (5)$$

2) Recurrence relation for $\xi_l(\mathbf{m})$: Let us observe that $\xi_l(\mathbf{m})$ admits the exact recurrence relation

$$\xi_l(\mathbf{m}) = \frac{1 + \sum_{j=1}^h \gamma_{kj} \xi_j^{-k}(\mathbf{m} - 1_l) + \gamma_{kl} \xi_l^{-k}(\mathbf{m} - 1_l)}{\frac{I\{n > m\}}{\xi_l^{-k}(\mathbf{m})} + \sum_{j=1}^h \gamma_{kj} \frac{\xi_j^{-k}(\mathbf{m} - 1_l)}{\xi_l^{-k}(\mathbf{m} - 1_j)} + \gamma_{kl}} \quad (6)$$

for arbitrary $k = 1, \dots, n$, where I is the indicator function. Relation (6) is a non-uniform access extension of known results for the function $F_l(\mathbf{m}) = m_l / \xi_l(\mathbf{m})$ [6, Eq. 2.7] and [15, Thm. 2] in uniform access models. The proof of (6) follows similarly to earlier work, that is plugging (3) into (5) we get

$$\xi_l(\mathbf{m}) = \frac{E_k(\mathbf{m} - 1_l) + \sum_{j=1}^h (m_j - \delta_{jl}) \gamma_{kj} E_k(\mathbf{m} - 1_j - 1_l)}{E_k(\mathbf{m}) + \sum_{j=1}^h m_j \gamma_{kj} E_k(\mathbf{m} - 1_j)} \quad (7)$$

The result is then obtained with simple passages by dividing both numerator and denominator by $E_k(\mathbf{m})$ and using that

$$\frac{T_l^{-k}(\mathbf{m})}{T_j^{-k}(\mathbf{m} - 1_l)} = \frac{E_k(\mathbf{m} - 1_j - 1_l)}{E_k(\mathbf{m} - 1_l)} \frac{E_k(\mathbf{m} - 1_l)}{E_k(\mathbf{m})}$$

Using (6), we can compute the miss rate in $\mathcal{O}(hn \prod_{j=1}^h (m_h + 1))$. The recursion has boundary conditions: i) $\xi_l(\mathbf{m}) = +\infty$ if $m > n$ or $\min_l m_l < 0$; ii) $\xi_l(\mathbf{0}) = 0$; iii) $\xi_l(1_l) = \gamma_{k,l}$ if $m_k = 1$ and $n = 1$; iv) $\xi_l(\mathbf{m}) = 0$ if $m_l = 0$.

F. Proof of Theorem 4

The proof relies on the following two lemmas, proven later in Appendix -G and Appendix -H.

Lemma 1. $\xi_j(\mathbf{m} - 1_l) \leq \xi_j(\mathbf{m}) \leq \xi_j^{-k}(\mathbf{m})$, $\forall j, l, k$.

Lemma 2. $\pi_{k0}(\mathbf{m}) \leq \pi_{k0}(\mathbf{m} - 1_l), \forall k, l$.

We obtain the bound by summing (7) over all l , and then by Lemma 2 replacing $\pi_i(\mathbf{m} - 1_l)$ with its upper bound $\pi_i(\mathbf{m})$. This yields $\pi_i(\mathbf{m}) \geq \sum_{l=1}^h \gamma_{il} \xi_l(\mathbf{m})(1 - \pi_i(\mathbf{m}))$. The proof follows after solving for $\pi_i(\mathbf{m})$ and using the result to compute $\pi_{i0}(\mathbf{m})$.

G. Proof of Lemma 1

The proof is by induction on the number of items n and depends on results in Appendix -E2, in particular on the recurrence relation

$$\xi_l(\mathbf{m}) = \frac{1 + \sum_{j=1}^h \gamma_{kj} \xi_j^{-k}(\mathbf{m} - 1_l) + \gamma_{kl} \xi_l^{-k}(\mathbf{m} - 1_l)}{\frac{I\{n > m\}}{\xi_l^{-k}(\mathbf{m})} + \sum_{j=1}^h \gamma_{kj} \frac{\xi_j^{-k}(\mathbf{m} - 1_l)}{\xi_l^{-k}(\mathbf{m} - 1_j)} + \gamma_{kl}} \quad (8)$$

where $I\{n > m\} = 1$ if and only if $n > m$.

We now prove the result.

Case $n = m+1$. The case holds since by (5) it is $\xi_j^{-i}(\mathbf{m} - 1_l) = 0$, being $E(\mathbf{m} - 1_l - 1_j) = 0$.

Induction step. Let the model with $n - 1$ items exclude item i , we focus on the hypothesis $\xi_j^{-i}(\mathbf{m} - 1_l) \leq \xi_j^{-i}(\mathbf{m}), \forall j, l$. We choose $k = i$ in (8) and divide both sides by $\xi_l^{-i}(\mathbf{m})$ so that

$$\frac{\xi_l(\mathbf{m})}{\xi_l^{-i}(\mathbf{m})} = \frac{1 + \sum_{j=1}^h \gamma_{ij} \xi_j^{-i}(\mathbf{m} - 1_l) + \gamma_{il} \xi_l^{-i}(\mathbf{m} - 1_l)}{1 + \sum_{j=1}^h \gamma_{ij} \xi_j^{-i}(\mathbf{m}) + \gamma_{il} \xi_l^{-i}(\mathbf{m})}$$

since $I\{n > m\} = 1$ and by (5) we note that $\xi_j^{-i}(\mathbf{m})/\xi_l^{-i}(\mathbf{m}) = E_j(\mathbf{m} - 1_l)/E_i(\mathbf{m} - 1_l) = \xi_j^{-i}(\mathbf{m} - 1_l)/\xi_l^{-i}(\mathbf{m} - 1_j)$. If we compare term-by-term numerator and denominator in the above expression, by the induction hypothesis we obtain the upper bound $\xi_l(\mathbf{m}) \leq \xi_l^{-i}(\mathbf{m})$. Plugging (5) in the last inequality, we get $E(\mathbf{m} - 1_l)E_i(\mathbf{m}) \leq E_i(\mathbf{m} - 1_l)E(\mathbf{m})$. Diving both sides by $E(\mathbf{m} - 1_l)E(\mathbf{m})$, by (36) we get $\pi_i(\mathbf{m}) \geq \pi_i(\mathbf{m} - 1_l)$. The last relationship implies in particular $\pi_i(\mathbf{m} - 1_j) \geq \pi_i(\mathbf{m} - 1_l - 1_j)$, thus by (6) we find the other bound $\xi_j(\mathbf{m} - 1_l) \leq m_j (\sum_{i=1}^n \gamma_{il} (1 - \pi_i(\mathbf{m} - 1_j)))^{-1} = \xi_j(\mathbf{m})$.

H. Proof of Lemma 2

We give the proof for $\pi_i(\mathbf{m}) = 1 - \pi_{i0}(\mathbf{m})$ by showing that $\pi_i(\mathbf{m}) \geq \pi_i(\mathbf{m} - 1_l)$, for any l . By (36), we need show that $E(\mathbf{m} - 1_l)E_i(\mathbf{m}) \leq E_i(\mathbf{m} - 1_l)E(\mathbf{m})$. Plugging (3) in $E(\mathbf{m} - 1_l)$ and $E(\mathbf{m})$, with the choice $k = i$, after simplifications it is sufficient to show that $E_i(\mathbf{m} - 1_l - 1_j)E_i(\mathbf{m}) \leq E_i(\mathbf{m} - 1_l)E_i(\mathbf{m} - 1_j)$, for all $j \neq l$. Dividing by $E_i(\mathbf{m})E_i(\mathbf{m} - 1_j)$ this may be seen as stating $\xi_j^{-i}(\mathbf{m} - 1_l) \leq \xi_j^{-i}(\mathbf{m})$, which holds after applying Lemma 1 to a model without item i .

I. Proof of Theorem 3

To prove the lower bound, we show that the assumption implies $\pi_i(\mathbf{m}) \geq \pi_k(\mathbf{m}), \forall \mathbf{m}$, so that the statement follows by applying Chebyshev's sum inequality [?] to the denominator of (6). Using (36), we need to show that $E_k(\mathbf{m}) \geq E_i(\mathbf{m})$ for all $i, k > i$. Note that $E_k(\mathbf{m})$ is obtained from $E_i(\mathbf{m})$ by

replacing γ_{kl} in $E_i(\mathbf{m})$ with $\gamma_{il} \geq \gamma_{kl}$. The statement then holds since the normalizing constant is by (35) non-decreasing with the access factors.

J. Derivation of the refined mean-field approximation (RMF)

In this section, we discuss in detail how to compute the refined mean-field approximation.

The refined mean field is obtained by considering a scaled system with N identical copies of each item and by multiplying all list sizes by N . Recall that $X_{il}^{(N)}(t)$ the fraction of class- i items that are in list l at time t , which implies that $NX_{il}^{(N)}(t)$ is the number of class- i items that are in list l at time t .

To simplify the notations, let $\rho_{ill'} = \sum_{v=1}^u \lambda_{v,k}(l)c_{vk}(l, l')$ be the rate at which a given item k is promoted from list l to list l' . At rate $\rho_{ill'}NX_{il}$, a class- i item will be promoted to list l' . It will be exchanged with an item from list l' . In such a list, there are $NX_{jl'}$ items of type j among the $Nm_{l'}$ items. Combining the two, this means that a class i items of state l is exchanged with a class j item of list l' at rate $N\rho_{ill'}NX_{il}X_{jl}/m_{l'}$. Such a transition will modify the state $X^{(N)}$ into

$$X^{(N)} + \frac{1}{N} (-\mathbf{e}_{il} + \mathbf{e}_{il'} + \mathbf{e}_{jl} - \mathbf{e}_{jl'}), \quad (9)$$

where $\mathbf{e}_{il}$ is a vector whose only non-zero coordinate is the (il) coordinate.

This shows that the process is a density dependent population process as defined in [36]. The case $N = 1$ corresponds to our original caching model. According to [?], there exists a deterministic dynamical system $x(t) = (x_{il}(t))$ such that for all item i and list l : $X_{il}(t) = x_{il}(t) + O(1/N)$, where x satisfies the differential equation $\dot{x} = f(x)$ with

$$f(x) = \sum_{il, j, l'} (-\mathbf{e}_{il} + \mathbf{e}_{il'} + \mathbf{e}_{jl} - \mathbf{e}_{jl'}) x_{il} x_{jl} / m_{l'}. \quad (10)$$

In particular, the (il) coordinate of f is

$$f_{il}(x) = \sum_{l'} \left[\rho_{ill'} x_{il'} - \rho_{ill'} x_{il} + \frac{x_{il'}}{m_{l'}} \sum_j \rho_{jll'} x_{jl} - \frac{x_{il}}{m_l} \sum_j \rho_{jl'l} x_{jl'} \right]. \quad (11)$$

In the above equation, the first line correspond to items that are promoted to list l from list l' (the term $\rho_{ill'} x_{il'}$) or promoted from list l to list l' (the term $\rho_{ill'} x_{il}$). The second line correspond to class i items that are demoted from list l' to list l (here $\sum_j \rho_{jll'} x_{jl}$ is the rate at which any item is promoted from list l to list l' and $\frac{x_{il'}}{m_{l'}}$ is the probability that a class i item is exchanged with it) or demoted from list l to list l' ($\frac{x_{il}}{m_l} \sum_j \rho_{jl'l} x_{jl'}$).

It is straightforward to verify that $\hat{\pi}^{FPI}$ satisfies $f(\hat{\pi}^{FPI}) = 0$, i.e., it is a fixed point of the differential equation $\dot{x} = f(x)$. Following [?], it can be shown that this fixed point is unique and is a global attractor of all the trajectories of the differential equation. It is shown in [21], that this implies that there exists a vector V such that, in steady-state,

$$\mathbb{E}[X_{il}] = \hat{\pi}_{il}^{FPI}(t) + \frac{1}{N} V_{il} + O(1/N^2).$$

Note that in the above equation $\mathbb{E}[X_{il}(t)]$ is equal to the steady-state probability that an item i is in list l .

The above vector V_{il} can be computed by solving the following linear system that depends on the first (Jacobian) and second derivative of the drift evaluated in $\hat{\pi}^{FPI}$:

$$\forall il, j'l' : \sum_{kl''} (A_{il,kl''} W_{kl'',jl'} + W_{il,kl''} A_{kl'',jl'}) + Q_{il,jl'} = 0.$$

$$\forall il : \sum_{jl'} A_{il,jl'} V_{jl'} + \frac{1}{2} \sum_{j'l', kl''} B_{il,jl',kl''} W_{jl',kl''} = 0,$$

In the above equation, $A_{il,jl'}$ is the derivative of $f_{il}(x)$ with respect to $x_{jl'}$, $B_{il,jl',kl''}$ the second derivative of $f_{il}(x)$ with respect to $x_{jl'}$ and $x_{kl''}$ and Q is a variance term whose expression is given later. The terms could be computed using symbolic analysis tools such a `rmf-tool`¹.

The above equations form a system of linear equations with $nh + (nh)^2$ unknowns (nh unknowns for V and $(nh)^2$ for W). This system can be solved by first computing W and then computing V . Note that the equation for W is a Lyapunov equation that can be solved by using appropriate algorithms such as the ones implemented in the function `scipy.linalg.solve_lyapunov` of the library `scipy`.

To compute the derivatives of f , the simplest numerical procedure is to use the definition of the drift of Equation (10). Denoting $e_{il,kl''}$ a matrix of size $(nh \times nh)$ whose only non-zero element is the (il, kl'') element, the derivative of f with respect to $x_{kl''}$ is the vector $A_{.,kl''}$:

$$A_{.,kl''} = \sum_{il,jl'} (-e_{il} + e_{il'} + e_{jl} - e_{jl'}) \frac{\partial(x_{il}x_{jl})}{\partial x_{kl''}} \frac{\rho_{ill'}}{m_{l'}}$$

$$= \sum_{il,jl'} (-e_{il} + e_{il'} + e_{jl} - e_{jl'}) (x_{jl'} \mathbf{1}_{il=kl''} + x_{il} \mathbf{1}_{jl=kl''}) \frac{\rho_{ill'}}{m_{l'}}$$

$$= \sum_{jl'} (-e_{kl''} + e_{kl'} + e_{jl''} - e_{jl'}) \frac{\rho_{kl''jl'}}{m_{l'}}$$

$$+ \sum_{il} (-e_{il} + e_{il''} + e_{kl} - e_{kl''}) \frac{\rho_{ill'} x_{il}}{m_{l''}}.$$

Note that $\frac{\partial(x_{il}x_{jl})}{\partial x_{kl''}} = (x_{jl'} \mathbf{1}_{il=kl''} + x_{il} \mathbf{1}_{jl=kl''})$ only when $il \neq jl'$. In the above sum, we can write this identity because when $l = l'$, the first factor of the sum is equal to 0.

Similarly, the second derivative of f with respect to $x_{kl''}$ and $x_{k'l''}$ is the vector $B_{.,kl'',k'l''}$ that is such that:

$$B_{.,kl'',k'l''} = \frac{\partial A_{.,kl''}}{\partial x_{k'l''}}$$

$$= \sum_{il,jl'} (-e_{il} + e_{il'} + e_{jl} - e_{jl'}) 2\mathbf{1}_{il=kl'',jl=k'l''} \frac{\rho_{ill'}}{m_{l'}}$$

$$= 2(-e_{kl''} + e_{kl'''} + e_{k'l''} - e_{k'l'''}) \frac{\rho_{kl''l''}}{m_{l''}}.$$

Finally following [21], the matrix Q is

$$Q = \sum_{il,jl'} \rho_{ill'} \frac{x_{il}x_{jl'}}{m_{l'}} \left(e_{il,il} + e_{il',il'} + e_{jl,jl} + e_{jl',jl'} \right. \\ \left. + 2(-e_{il,il'} - e_{il,jl} + e_{il,jl'} + e_{il',jl} - e_{il',jl'} - e_{jl,jl'}) \right).$$

¹https://github.com/ngast/rmf_tool/

K. Further validation: models with synchronous streams

1) *Parameters for models with synchronous streams:* The following table lists the model parameters for the experiments discussed in Section VIII-A.

n	Number of items	10, 50, 100
n/m	Items-to-capacity ratio	2, 4, 10
h	Number of lists	1, 2, 3
u	Number of streams	2
α_1	Zipf parameter, stream 1	0.6, 1.0, 1.4
α_2	Zipf parameter, stream 2	1.0
σ_1	Mean think time, stream 1	0.001s, 0.01s, 0.1s, 1.0s, 10s
σ_2	Mean think time, stream 2	$2\sigma_1$
θ_{v0}	Mean fetch (hit) latency	0.01s
θ_{v1}	Mean miss latency	0.10s
s_v	Maximum outstanding requests per stream	1
c	Normalized access prob.	1.0

2) *Sensitivity to number of outstanding requests:* We have also considered variants where for every stream v the maximum number of outstanding requests is varied as $s_v = 1, 4, 16$, and different mixes of think time values are used in the model. We keep the number of samples used in the simulations to 10^7 . In these experiments, we set $n = 10, 50, 100$, $h = 2$, $c = 0.5$, $\alpha_1 = 1.4$, and use both `FIFO-C(m)` and `RR-C(m)`, for a total of 162 experiments. The results indicate that average errors remain small for all mixes (σ_1, σ_2), with FPI ranging between 2.01% and 2.33% mean absolute percentage errors across the mixes and with SPA between 0.18% and 0.49%. For both methods, the largest error for FPI is observed for mix (16, 1), while for SPA is for mix (1, 16).

3) *Sensitivity to access probabilities and tree structure:* We now consider another variant where the cache has $h = 5$ lists arranged in a tree structure. As before, we consider $n = 10, 50, 100$, $n/m = 10, 50, 100$, $\alpha_1 = 0.6, 1.0, 1.4$, $\alpha_2 = 1.0$, and $s_1 = s_2 = 1$. Access probabilities for list $l = 0$, corresponding to cache misses, are set to 1.0, for all $v = 1, \dots, u$. Access probabilities for stream v requests to item k with a cache hit are given by the following matrix

$$C_{vk} = \begin{bmatrix} 0 & \frac{1}{1+k \cdot v} & \frac{1}{3(1+k \cdot v)} & \frac{2}{3(1+k \cdot v)} \\ \frac{k \cdot v}{1+k \cdot v} & \frac{v}{1+v} & \frac{1}{1+v} & 1 \\ & 1 & & \\ & & \frac{3}{10} & \frac{7}{10} \\ & & & 1 \end{bmatrix}$$

where element in row l and column j corresponds to $c_{vk}(l, j)$, for $l, j = 0, \dots, h$. Costs on the diagonal give the probability that the item upon a hit does not get promoted to another list.

With the above parameters we find that for `FIFO-C(m)` the average errors are 1.23% for FPI and 0.38% for SPA, while maximum errors are 4.34% for FPI and 1.41% for SPA. We also repeat the experiment with `RR-C(m)` finding average

errors of 1.21% for FPI and 0.32% for SPA, and maximum errors of 4.17% for FPI and 1.41% for SPA, which are close to the FIFO-C(m) results and compatible with simulation error.

L. Validation: hybrid caching-queueing models

1) *Parameters for models with queueing*: The following table lists the model parameters for the experiments on synchronous models with queueing stations. We index nodes as t (think time), c (cache), h (hit queue), and m (miss queue).

n	Number of items	10, 50, 100
n/m	Items-to-capacity ratio	2, 4, 10
h	Number of lists	2
u	Number of streams	2
α_{c1}	Zipf parameter at cache c , stream 1	1.4
α_{c2}	Zipf parameter at cache c , stream 2	1.0
θ_{t1}	Mean think time, stream 1	1.0s
θ_{t2}	Mean think time, stream 2	2.0s
θ_{h1}	Mean fetch (hit) latency, stream 1	0.01s
θ_{h2}	Mean fetch (hit) latency, stream 2	0.01s
θ_{m1}	Mean miss latency, stream 1	0.10s
θ_{m2}	Mean miss latency, stream 2	0.50s
s	Maximum outstanding requests in each stream	1, 4, 16, 32
c	Normalized access prob.	1.0

2) *Correlated item reads*: We consider a variant of the parallel topology in Figure 8 in which a request that completes in stream 1 deterministically switches class at the reference station (think time node) so to re-enter in the following cycle as a synchronous request of stream 2, and vice-versa for requests completing in stream 2. This is sufficient to introduce a degree of serial correlation in the item references at the cache, since each requests will deterministically alternate requests drawn for Zipf-like distributions with different parameters ($\alpha_{c1} \neq \alpha_{c2}$).

While our approximations cannot explicitly represent correlations, the results shown in Figure 11 show that both FPI and SPI remain fairly accurate on the considered model, especially with small and large number of outstanding requests. This suggests that the proposed methods can still provide reasonably accurate results in models that are reasonable close to the reference ones, but that strictly speaking do not meet the assumptions.




Fig. 11: FPI and SPA item miss ratios on the parallel topology in Figure 8 with correlated item reads.