Improved Rate-Based Pull and Push Strategies in Large Distributed Networks

Wouter Minnebo and Benny Van Houdt
Department of Mathematics and Computer Science
University of Antwerp - iMinds
Middelheimlaan 1, B-2020 Antwerp, Belgium
Email: {wouter.minnebo,benny.vanhoudt}@ua.ac.be

***Abstract*—Large distributed systems benefit from the ability to exchange jobs between nodes to share the overall workload. To exchange jobs, nodes rely on probe messages that are either generated by lightly-loaded or highly-loaded nodes, which corresponds to a so-called pull or push strategy. A key quantity of any pull or push strategy, that has often been neglected in prior studies, is the resulting overall probe rate. If one strategy outperforms another strategy in terms of the mean delay, but at the same time requires a higher overall probe rate, it is unclear whether it is truly more powerful.**

In this paper we introduce a new class of rate-based pull and push strategies that can match any predefined maximum allowed probe rate, which allows one to compare the pull and push strategy in a fair manner. We derive a closed form expression for the mean delay of this new class of strategies in a homogeneous network with Poisson arrivals and exponential job durations under the infinite system model. We further show that the infinite system model is the proper limit process over any finite time scale as the number of nodes in the system tends to infinity and that the convergence extends to the stationary regime.

Simulation experiments confirm that the infinite system model becomes more accurate as the number of nodes tends to infinity, while the observed error is already around 1% for systems with as few as 100 nodes.

I. Introduction

One of the key features of a contemporary distributed network is its ability to (re)distribute the workload among a large number of processing nodes. Jobs in such a network can either enter the network via one (or multiple) centralized job dispatchers [1], [2] or may enter via the processing nodes themselves [3]–[6]. In the latter case, the workload is redistributed by the exchange of jobs between the processing nodes. Two important families of strategies have been identified for redistributing jobs: the *pull* and *push* strategy. Under the pull strategy lightly-loaded nodes try to attract work from highly-loaded nodes, a strategy that is also known as *work stealing*. Under the push strategy the highly-loaded nodes take the initiative to transfer jobs to lightly-loaded nodes.

Many performance studies of pull and push strategies have been presented over the years. The performance of two traditional pull and push strategies in homogeneous networks with Poisson arrivals and exponential job durations was analyzed in [3] and extended to heterogeneous networks in [7]. Both studies showed that the pull strategy is superior under high load conditions, while the push strategy results in a lower mean delay under low to medium loads. More recent analytic studies of the performance of pull and push strategies include [5], [6].

All of these studies provided valuable insights with respect to the performance of pull and push strategies. However, they also paid hardly any attention to the probe rate, that is, the number of probe messages that the strategies under consideration generate per time unit. These probe messages are used to enable the exchange of jobs and thus to balance the load. Typically, when a node wants to pull or push a job it probes another node at random to see whether a job can be transferred. Some of the probes result in a job transfer, while others do not. Clearly, sending more probe messages tends to result in more job transfers and therefore in lower mean delays.

As different strategies tend to have different probe rates that depend to a large extent on the arrival rate λ, it is typically not possible to adapt the parameters of the strategies under consideration such that they generate the same overall probe rate (for arbitrary λ), making any comparison biased. Further, some so-called optimal parameter settings also result in the highest probe rate, which makes the optimality questionable (e.g., the hybrid pull/push strategy in [4]).

To mitigate this, we introduced a class of rate-based pull and push strategies in [8] that make use of a single parameter r. Under the pull strategy idle nodes generate probe messages at rate r, while under the push strategy probe messages at rate r are generated by the nodes with at least 2 queued jobs (including the one in service). As shown in [8], these rate-based strategies can match any predefined maximum allowed probe rate R by setting r in the appropriate manner and this for any arrival rate λ, allowing a fair comparison between the pull and push strategy. The main results in [8] showed that the rate-based push strategy results in a lower mean delay if and only if

$$\lambda < \frac{\sqrt{(R+1)^2 + 4(R+1)} - (R+1)}{2},$$

under the so-called infinite system model and that a hybrid pull/push strategy is always inferior to the pure pull or push strategy.

To evaluate and compare the different strategies considered in this paper we introduce an infinite system model, the evolution of which is described by a set of ordinary differential equations (ODEs) as in [5], [8]. To assess the mean delay and overall probe rate of a strategy, we define a set of ODEs, give an explicit expression for its unique fixed point and express the mean delay and probe rate using this fixed point. To guarantee all trajectories converge to the fixed point, we prove that the fixed point is a global attractor. We also show that the set of ODEs captures the evolution of the limit process of a family of density dependent Markov chains as introduced by Kurtz in [9], [10]. Simulation experiments confirm that the infinite system model becomes exact as the number of nodes in the system tends to infinity, while the error is about 1% for systems with as few as 100 nodes.

This paper makes the following contributions:

1) We introduce a more general class of rate-based pull and push strategies that rely on two parameters T and r and that coincide with the strategies introduced in [8] when $T = 1$. Closed form results for the mean delay of this new class of pull and push strategies are presented (under the infinite system model).
2) We show that setting $T > 1$ reduces the mean delay of the rate-based push strategy (for larger λ and smaller R values). This is in contrast to earlier findings for the traditional strategy [4], [7], for which smaller T values result in higher probe rates, making the comparison biased. For the rate-based pull strategy we show that setting $T = 1$ is optimal.
3) We introduce the so-called max-push strategy and derive a closed form expression for its mean delay (under the infinite system model). We show that the max-push strategies further reduce the mean delay of the best rate-based pull and push strategies with $T \geq 1$ for certain combinations of (λ, R).
4) Finally, we prove that the infinite system models introduced in this paper are the proper limit processes of the finite stochastic systems with N nodes as N tends to infinity over any finite time scale. In addition, we prove that the convergence extends to the stationary regime (i.e., the ODEs have a global attractor).

The paper is structured as follows. Section II introduces the rate-based pull and push strategies. For the rate-based strategies with $T \geq 1$ we present the infinite system model in Section III. In Section IV we validate this model using simulation results and present some numerical examples that compare the performance of the rate-based pull and push strategy. Section V introduces the max-push strategy and its infinite system model, while numerical results for the the max-push strategy are presented in Section VI.

II. Pull and Push Strategies

We consider a continuous-time system consisting of N queues, where each queue consists of a single server with an infinite buffer. As in [3], [5]–[7], jobs arrive locally according to a Poisson process with rate $\lambda < 1$, and have an exponentially distributed duration with mean 1. Servers process jobs in a first-come first-served order. Servers can send probe messages to each other to query for queue length information, and to transfer jobs. We assume that the time required to transfer probe messages and jobs is sufficiently small in comparison with the processing time of a job, i.e., transfer times are considered zero.

We consider the following load balancing strategies that all make use of two parameters: an integer $T \geq 1$ and a real number $r > 0$. We note that the rate-based strategies with $T = 1$ were initially introduced in [8].

1) *Rate-based Push:* As soon as the queue length exceeds T, a server starts to generate probe messages according to a Poisson process with rate r. Whenever the queue length drops below T, this process is interrupted until the queue length exceeds T again. The node that is probed is selected at random and is only allowed to accept a job if it is idle.
2) *Rate-based Pull:* Whenever a server is idle it generates probe messages according to a Poisson process with rate r. This process is interrupted whenever the server is busy. The node that is probed is selected at random and is only allowed to transfer one of its jobs if its queue length exceeds T.

For each of the above strategies transferred jobs are immediately served by the accepting node, hence any job transfer results in a reduction of the mean delay. To make the comparison fair the mean overall probe rate R should be identical. The rate R is defined as the mean number of probes that is sent by a server per time unit irrespective of its queue length, where R is clearly less than r. Further on we will show that r can be set in such a manner that it can match any predefined $R \geq 0$.

III. Performance of rate-based strategies with $T \geq 1$

In this section we introduce the infinite system model to assess the performance of the rate-based strategies with $T \geq 1$. This model, the evolution of which is captured by a set of ODEs, is validated by simulation in Section IV, while in Section VII it is argued to be the proper limit process of the stochastic finite system model with N nodes as N tends to infinity.

The evolution of both the rate-based pull and push strategy model is given by a set of ODEs denoted as $\frac{d}{dt}x(t) = F(x(t))$, where $x(t) = (x_1(t), x_2(t), \dots)$ and $x_i(t)$ represents the fraction of the number of nodes with at least i jobs at time t. As explained below, this set of ODEs can be written as

$$\frac{dx_1(t)}{dt} = (\lambda + r x_{T+1}(t))(1 - x_1(t)) - (x_1(t) - x_2(t)), \tag{1}$$

$$\frac{dx_i(t)}{dt} = \lambda(x_{i-1}(t) - x_i(t)) - (x_i(t) - x_{i+1}(t)), \tag{2}$$

for $2 \leq i \leq T$ and

$$\begin{aligned}\frac{dx_i(t)}{dt} &= \lambda(x_{i-1}(t) - x_i(t)) - (x_i(t) - x_{i+1}(t)), \\ &\quad - r(1 - x_1(t))(x_i(t) - x_{i+1}(t)),\end{aligned} \tag{3}$$

for $i > T$. The terms $\lambda(x_{i-1}(t) - x_i(t))$ and $x_i(t) - x_{i+1}(t)$, for $i \geq 1$, correspond to arrival and service completions, respectively. Under the pull strategy probes are sent at rate $r(1 - x_1(t))$ and a probe is successful with probability $x_{T+1}(t)$, while under the push strategy the probe rate equals $rx_{T+1}(t)$ and a probe is successful with probability $(1 - x_1(t))$. Hence, for both strategies queues of length 1 are created by job transfers at rate $rx_{T+1}(t)(1 - x_1(t))$. Similarly, job transfers reduce the number of queues with exactly i jobs, for $i > T$, at rate $r(1 - x_1(t))(x_i(t) - x_{i+1}(t))$ under both strategies.

Let $E = \{(x_1, x_2, \ldots) | x_i \in [0, 1], x_i \geq x_{i+1}, i \geq 1, \sum_{j\geq 1} x_j < \infty\}$. The next two theorems show that this set of ODEs is Lipschitz on E and it has a unique fixed point in E.

Theorem 1. *The function F is Lipschitz on E.*

Proof: F is Lipschitz provided that for all $x, y \in E$ there exists an $L > 0$ such that $|F(x) - F(y)| \leq L\,|x - y|$. By definition of $F(x)$ one finds

$$\begin{aligned}|F(x) - F(y)| \leq 2(\lambda + 1 + 2r)\,|x - y| + \\ 2r \sum_{i>T} |x_1(x_i - x_{i+1}) - y_1(y_i - y_{i+1})|\,.\end{aligned}$$

The above sum can be bounded by

$$\sum_{i>T} |(x_1 - y_1)(x_i - x_{i+1}) + y_1(x_i - x_{i+1} - y_i + y_{i+1})|\,,$$

which is bounded by $2\,|x - y|$ on E. Hence, F is Lipschitz by letting $L = 2\lambda + 2 + 8r$. ∎

As E is a Banach space the Lipschitz condition of F suffices to guarantee that the set of ODEs $\frac{d}{dt}x(t) = F(x(t))$, with $x(0) \in E$, has a unique solution[1] $\phi_t(x(0))$ [11, Section 1.1].

Theorem 2. *The set of ODEs given by* (1) *to* (3) *has a unique fixed point $\bar{\pi} = (\bar{\pi}_1, \bar{\pi}_2, \ldots)$ with $\sum_{i\geq 1} \bar{\pi}_i < \infty$. Let $\eta_i = \bar{\pi}_i - \bar{\pi}_{i+1}$ and $\eta_0 = 1 - \lambda$, then the fixed point can be expressed as*

$$\begin{aligned}
\eta_1 &= \frac{\lambda(1 + (1 - \lambda)r - \lambda)}{1 + (1 - \lambda^T)r}, \\
\eta_i &= \eta_1 \lambda^{i-1}, && 2 \leq i \leq T, \\
\eta_i &= \eta_T \left(\frac{\lambda}{1 + (1 - \lambda)r}\right)^{i-T}, && i > T.
\end{aligned}$$

Proof: Assume $\bar{\pi}$ is a fixed point with $\sum_{i\geq 1} \bar{\pi}_i < \infty$, meaning $F_i(\bar{\pi}) = 0$ for $i \geq 1$, where $F(x) = (F_1(x), F_2(x), \ldots)$. When $\sum_{i\geq 1} \pi_i < \infty$, we can simplify $\sum_{i\geq 1} F_i(\pi) = 0$ to $\lambda - \bar{\pi}_1 = 0$. Hence, $\bar{\pi}_1$ must equal λ. The expressions for η_i then readily follow from the conditions $F_i(\bar{\pi}) = 0$, for $i \geq 1$. ∎

This fixed point is also the unique solution of the Kolmogorov differential equation for a state dependent M/M/1 queue with $\lambda_0 = \lambda + r\bar{\pi}_{T+1}$, $\lambda_i = \lambda$, for $i \geq 1$, $\mu_i = 1$, for $i = 1, \ldots, T$, and $\mu_i = 1 + (1 - \lambda)r$, for $i \geq T + 1$. The arrival process of such an M/M/1 queue is Poisson with rate λ_i and the service is exponential with rate μ_i whenever the queue length equals i.

The set of ODEs in (1) to (3) describes the transient evolution of the infinite system, while we are in fact interested in its behavior as t goes to infinity. Thus, we are interested in the limit of all the trajectories of this set of ODEs. In Appendix A we prove the following theorem:

Theorem 3. *All the trajectories of the set of ODEs given by* (1) *to* (3)*, starting from $x \in E$ converge towards the unique fixed point π.*

Due to the above theorem, we can now express the main performance measures of the push and pull strategies with $T \geq 1$ via Theorem 2:

Corollary 1. *The mean delay D of a job under the push or pull strategy equals*

$$D_{both} = \frac{1}{1 - \lambda} - \frac{r\lambda^T \left(\frac{\lambda}{(1-\lambda)(1+r)} + T\right)}{1 + r(1 - \lambda^T)}.$$

Proof: Using Theorem 2, one can apply Little's law to express the mean response time D as $\sum_{i\geq 1} \pi_i / \lambda$. ∎

Corollary 2. *Given a predefined maximum allowed probe rate R, the rate r must be set as*

$$r_{pull} = \frac{R}{1 - \lambda}, \tag{4}$$

$$r_{push} = \frac{R}{\lambda^{T+1} - (1 - \lambda^T)R}, \tag{5}$$

with $r_{push} = \infty$ for $R > \lambda^{T+1}/(1 - \lambda^T)$. Hence, if the predefined value of R exceeds $\lambda^{T+1}/(1 - \lambda^T)$, the rate r_{push} can be set arbitrarily high.

Proof: From the relationships $R = (1 - \bar{\pi}_1)r_{pull}$ and $R = r_{push}\bar{\pi}_{T+1}$, we find

$$R = (1 - \lambda)r_{pull},$$

and

$$R = \frac{\lambda^{T+1}}{(1 - \lambda^T) + 1/r_{push}}.$$

∎

Theorem 4. *The mean response time D of a job under the push stategy equals*

$$D_{push} = \frac{1 - R}{1 - \lambda} - \frac{RT}{\lambda} + \frac{R^2}{\lambda^T(1 - \lambda)(\lambda + R)},$$

if $R \leq \lambda^{T+1}/(1 - \lambda^T)$, while for $R > \lambda^{T+1}/(1 - \lambda^T)$ the rate $r_{push} = \infty$, and the mean delay D_{push} is given by:

$$D_{push|r=\infty} = \frac{1}{1 - \lambda} - \frac{T\lambda^T}{1 - \lambda^T}.$$

Remark, $D_{push|r=\infty} = 1$ for $T = 1$.
Under the pull strategy the mean response time equals

$$D_{pull} = \frac{1}{1 - \lambda} - \frac{R\lambda^T(\frac{\lambda}{1-\lambda+R} + T)}{1 - \lambda + R(1 - \lambda^T)}.$$

[1]The solution $\phi_t(x)$ belongs to the class of continuously differentiable functions as in the finite dimensional case.

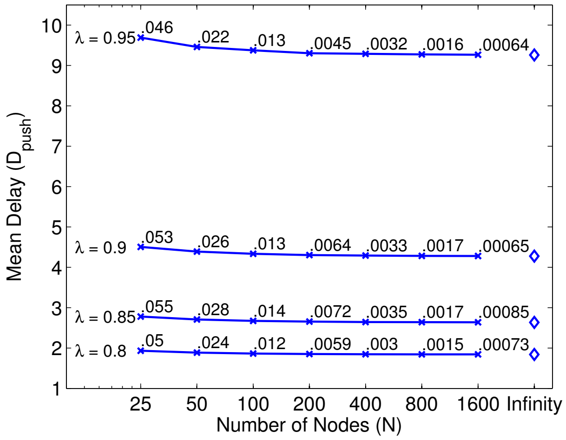

Fig. 1. Simulated mean delay for a finite system of varying size, using a rate-based push strategy with $T = 2$, matching an overall request rate of $R = 1$. The relative error, shown above a simulated point, becomes smaller when simulating larger systems. In addition, the infinite system model approximates systems of moderate size well.

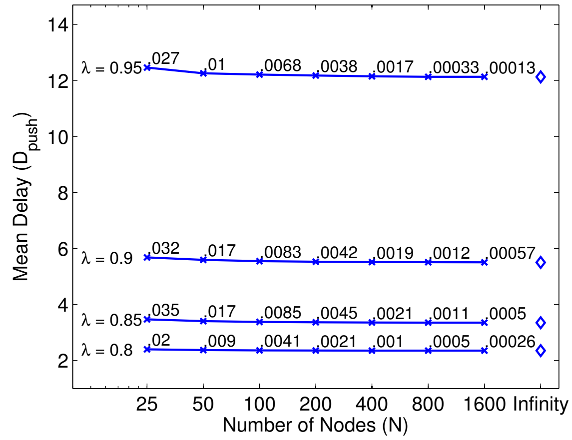

Fig. 2. Simulated mean delay for a finite system of varying size, using a rate-based push strategy with $T = 4$, matching an overall request rate of $R = 0.5$. The relative error, shown above a simulated point, becomes smaller when simulating larger systems.

Proof: The expressions for D_{push} and D_{pull} are found using Corollary 1, by plugging in the appropriate value for r, given by Corollary 2, in the expression for D_{both}. ■

Theorem 5. *The optimal choice for the rate-based pull strategy is* $T = 1$.

Proof: It can be verified that increasing T by one will increase D_{pull} if and only if

$$\frac{\lambda/(1-\lambda+R)+T}{1-\lambda+R(1-\lambda^T)} \geq \lambda\frac{\lambda/(1-\lambda+R)+T+1}{1-\lambda+R(1-\lambda^{T+1})},$$

which is equivalent to stating

$$T(1-\lambda+R)(1-\lambda) \geq \lambda R(1-\lambda^T).$$

This condition can be rewritten as

$$T(1-\lambda+R) \geq R\sum_{i=1}^{T}\lambda^i,$$

which holds as $\lambda^i < 1$, for $i = 1, \ldots, T$. Hence, increasing T by one always increases the mean delay of the rate-based pull strategy. ■

Although it may at first seem sensible to steal jobs from *long* queues only if the maximum allowed probe rate R is low, the pull strategy is in some sense blind as it also needs to send probes to locate these *long* queues. This is contrary to the push strategy, where nodes will only probe if their queue is *long*, as a result setting $T = 1$ is not always optimal for the rate-based push strategy (see numerical results in Section IV).

IV. Numerical results for rate-based strategies with $T \geq 1$

A. Validation

In this section we present validation results for the rate-based push strategy with $T \geq 2$ as the model for both rate-based strategies with $T = 1$ was already validated in [8] and the mean delay of the pull strategy is minimized for $T = 1$. The infinite system model and simulation setup only differ in the system size. The rate r_{push} in the simulation experiments is independent of N and was determined by λ and R using the expression for R in (5). Each simulated point in the figures represents the average value of 25 simulation runs. Each run has a length of 10^6 (where the service time is exponentially distributed with mean 1) and a warm-up period of length $10^6/3$.

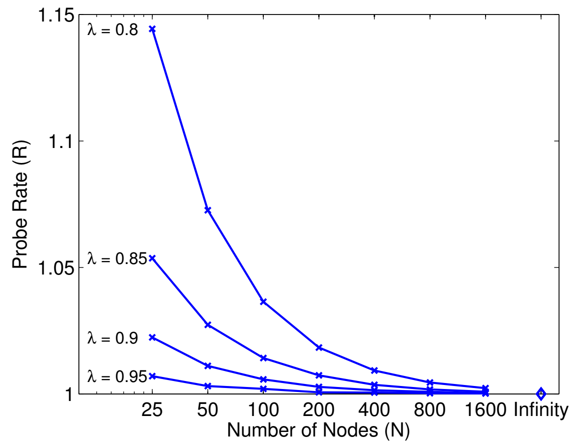

Fig. 3. Request rate for the finite system using a rate-based push strategy with $T = 2$.

Figure 1 depicts the mean delay as a function of N for $T = 2$, $R = 1$ and $\lambda = 0.8, 0.85, 0.9$ and 0.95, while Figure 2 depicts the same for $T = 4$, $R = 0.5$. In both cases the relative error shown above the simulation results decreases as N tends to infinity. The error for a system with as few as 100 nodes is only slightly above 1% when $T = 2$. We should note that the actual overall probe rate observed in the finite system exceeds R for smaller N values as shown in Figure 3 and 4. In other words, the relation between R and r_{push} given by (5) is not very accurate for small N values as the

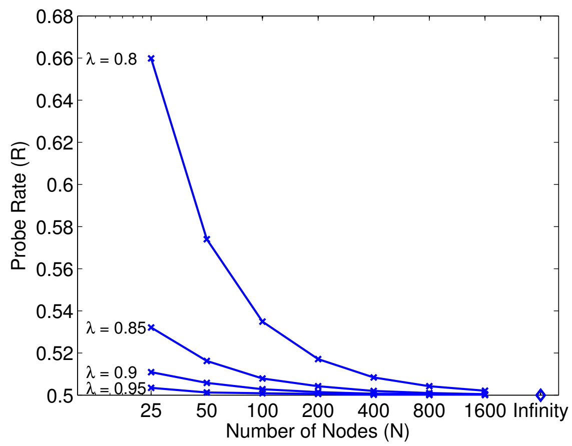

Fig. 4. Request rate for the finite system using a rate-based push strategy with $T = 4$.

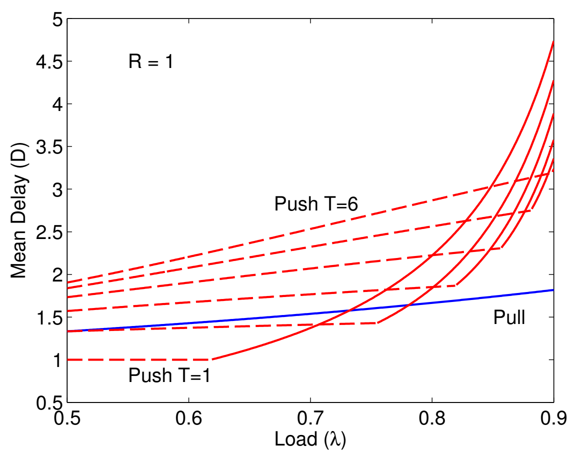

Fig. 5. The mean delay for the rate-based pull ($T = 1$) and push strategy ($T = 1, \ldots, 6$) with $R = 1$. For the push strategy the part of the curve with $r = \infty$ is dashed.

infinite model is optimistic with respect to the queue length distribution. However, as soon as the system consists of several hundred nodes, there is a fairly good agreement. Similar results were observed for other parameter settings.

B. Comparison of push and pull strategy

The mean delay of the rate-based push (for $T = 1, \ldots, 6$) and pull (for $T = 1$) strategy is shown in Figure 5 as a function of λ for a mean overall probe rate $R = 1$. The curves for the push strategy consist of two parts and $r_{push} = \infty$ for the dashed part of the curve. For these loads λ the rate r_{push} can be set arbitrarily high without violating the maximum allowed probe rate $R = 1$. The results indicate that the optimal T value for the push strategy increases as λ increases (while R remains fixed). This is as expected as small T values allow queues with a length below average to probe for idle servers, using part of the available probe rate. For the same reason smaller R values also give rise to larger optimal T values (for fixed λ). Figure 5 also indicates that setting $T > 1$ implies that the rate-based push strategy can outperform the pull strategy

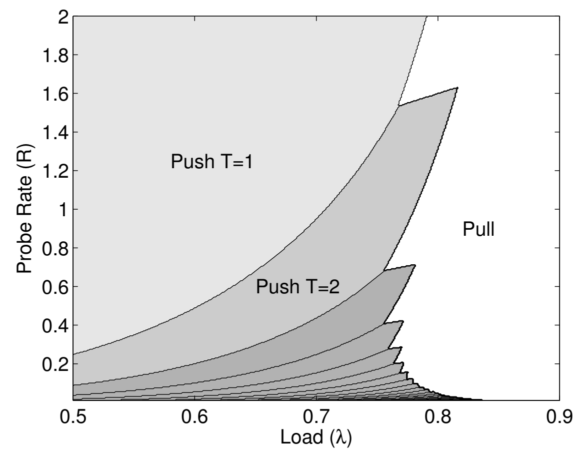

Fig. 6. The optimal rate-based strategy with variable T as a function of the load λ and the overall probe rate R. The pull strategy is optimal for high loads. The optimal T for the push strategy increases as R decreases.

for a larger range of loads λ.

The rate-based strategy (with variable T) that minimizes the mean delay for $\lambda \in [0.5, 0.9]$ and $R \in [0, 2]$ is depicted in Figure 6. The pull strategy is superior for loads above 75 to 80%. For lower loads the push strategy prevails and lower maximum allowed probe rates R give rise to larger optimal T values.

V. The Max-Push Strategy

The mean delay under the rate-based push strategy given in Theorem 4 can be further reduced as follows. Recall, whenever $R > \lambda^{T+1}/(1 - \lambda^T)$, the rate r_{push} can be chosen arbitrarily large (i.e., $r_{push} = \infty$). In other words, even if requests are sent at infinite rate when the queue length exceeds T, the overall probe rate remains below R. Hence, in order to use this remaining request rate, we introduce the max-push strategy when $T > 1$ and

$$\lambda^{T+1}/(1 - \lambda^T) < R < \lambda^T/(1 - \lambda^{T-1}). \tag{6}$$

Note, for any $R > 0$ and $0 < \lambda < 1$, there exists a single $T > 0$ such that the above relationship holds. Under the max-push strategy we send probes at an infinite rate whenever the queue length exceeds T and at rate $r < \infty$ if the queue length equals T. Note, under this strategy jobs are instantaneously transferred to another queue if the queue length equals T upon arrival (at the expense of a number of probe messages). The evolution of the infinite system model for this strategy is also readily formulated as a set of ODEs $\frac{d}{dt}x(t) = G(x(t))$, where $x(t) = (x_1(t), \ldots, x_T(t))$ and $x_i(t)$ represents the fraction of the number of nodes with at least i jobs at time t:

$$\begin{aligned}\frac{dx_1(t)}{dt} &= \lambda(1 - x_1(t) + x_T(t)) - (x_1(t) - x_2(t)) \\ &\qquad + r x_T(t)(1 - x_1(t))\end{aligned} \tag{7}$$

$$\frac{dx_i(t)}{dt} = \lambda(x_{i-1}(t) - x_i(t)) - (x_i(t) - x_{i+1}(t)), \tag{8}$$

for $2 \leq i \leq T-1$ and

$$\frac{dx_T(t)}{dt} = \lambda(x_{T-1}(t) - x_T(t)) - x_T(t)(1 + r(1 - x_1(t))). \tag{9}$$

The terms of the form $\lambda(x_{i-1}(t) - x_i(t))$ and $(x_i(t) - x_{i+1}(t))$, for $1 \leq i \leq T$, are again due to arrival and service completion events, respectively. Additionally, queues of length 1 are created at rate $\lambda x_T(t)$ due to the instantaneous job transfers and rate $r x_T(t)(1 - x_1(t))$ due to successful probes sent by a queue of length T, while the latter event also reduces the number of queues of length T by one.

Let $E_T = \{(x_1, \ldots, x_T) | 1 \geq x_1 \geq x_2 \geq \ldots x_T \geq 0\}$. The next two theorems show that this set of ODEs is Lipschitz on E_T and it has a unique fixed point in E_T.

Theorem 6. *The function G is Lipschitz on E_T.*

Proof: G is Lipschitz provided that for all $x, y \in E_T$ there exists an $L > 0$ such that $|G(x) - G(y)| \leq L\,|x - y|$. By definition of $G(x)$ one finds

$$|G(x) - G(y)| \leq 2(2\lambda + 1 + 2r)\,|x - y|\,.$$

Hence, G is Lipschitz by letting $L = 4\lambda + 2 + 4r$. ■

As E_T is a finite dimensional space the Lipschitz condition of G suffices to guarantee that the set of ODEs $\frac{d}{dt}x(t) = G(x(t))$, with $x(0) \in E_T$, has a unique solution (due to the Picard Lindelöf theorem).

Theorem 7. *The set of ODEs given by* (7)-(9) *has a unique fixed point $\dot{\pi} = (\dot{\pi}_1, \ldots, \dot{\pi}_T)$ in E_T that can be expressed as*

$$\dot{\pi}_i = \lambda^i \frac{1 + (\frac{\lambda}{1-\lambda} + r)(1 - \lambda^{T-i})}{1 + (\frac{\lambda}{1-\lambda} + r)(1 - \lambda^{T-1})},$$

for $1 \leq i \leq T$.

Proof: Assume $\dot{\pi}$ is a fixed point with $\sum_{i\geq 1} \dot{\pi}_i \leq T$, meaning $G_i(\dot{\pi}) = 0$ for $i \geq 1$, where $G(x) = (G_1(x), G_2(x), \ldots, G_T(x))$. When $\sum_{i\geq 1} \dot{\pi}_i \leq T$, we can simplify $\sum_{i\geq 1} G_i(\dot{\pi}) = 0$ to $\lambda - \dot{\pi}_1 = 0$. Hence, $\dot{\pi}_1$ must equal λ. The expressions for $\dot{\pi}_i$ then follow from the condition $G_i(\dot{\pi}) = 0$. ■

In Appendix B we prove the following theorem:

Theorem 8. *All the trajectories of the set of ODEs given by* (7)-(9)*, starting from $x \in E_T$ converge towards the unique fixed point $\dot{\pi}$.*

Due to the above theorem, we can now express the main performance measures of the max-push strategy via Theorem 7:

Corollary 3. *The mean delay D_{mp} of a job under the max-push strategy equals*

$$D_{mp} = \frac{1 - \lambda^T + (\frac{\lambda}{1-\lambda} + r)(1 - T\lambda^{T-1} + (T-1)\lambda^T)}{1 + r(1-\lambda)(1 - \lambda^{T-1}) - \lambda^T}.$$

A predefined overall probe rate R can be matched by setting

$$r_{mp} = \frac{R}{\lambda^{T-1}(R + \lambda) - R} - \frac{\lambda}{1 - \lambda}, \tag{10}$$

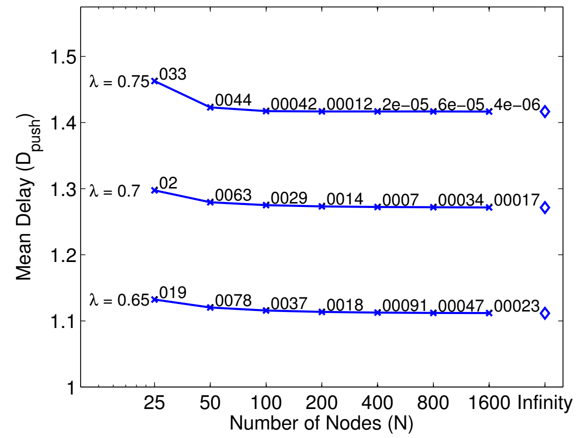

Fig. 7. Simulated mean delay for a finite system of varying size, using a max-push strategy with $T = 2$ and $R = 1$. The relative error, shown above a simulated point, becomes smaller when simulating larger systems. In addition, the infinite system model approximates systems of moderate size well.

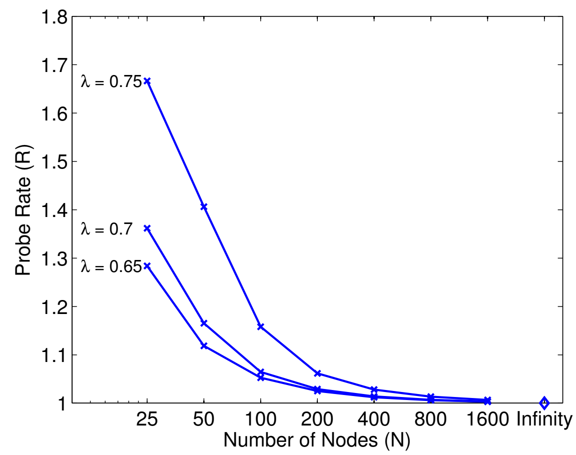

Fig. 8. Observed probe rate for the finite system using a max-push strategy with $T = 2$.

where $0 \leq r_{mp} < \infty$ for $\lambda^{T+1}/(1 - \lambda^T) \leq R < \lambda^T/(1 - \lambda^{T-1})$.

Proof: The mean response time D can be expressed as $\sum_{i\geq 1}^{T} \pi_i/\lambda$ by Little's law. For the max-push strategy the overall probe rate R equals

$$R = \dot{\pi}_T \left(\frac{\lambda}{1 - \lambda} + r_{mp} \right),$$

as the instantaneous transfer of an arrival to a queue with T jobs requires $1/(1 - \lambda)$ probe messages on average. ■

VI. Numerical results for the max-push strategy

A. Validation

In this section we validate the infinite system model for the max-push strategy using the same approach as in Section IV-A for the rate-based push strategy with $T > 1$. The rate r_{mp} in the simulation was determined using the relationship in (10).

The mean delay as a function of the number of nodes N and the relative error are shown in Figure 7 for $T = 2$,

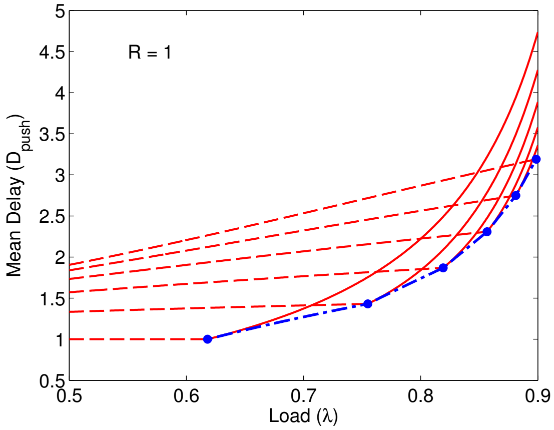

Fig. 9. The mean delay for the rate-based push strategy (full lines) for $T = 1, \ldots, 6$ with $R = 1$. For each strategy, the probe rate $r = \infty$ when λ is below the load marked by a dot. The max-push strategy interconnects the dots, as shown by dash-dotted curves.

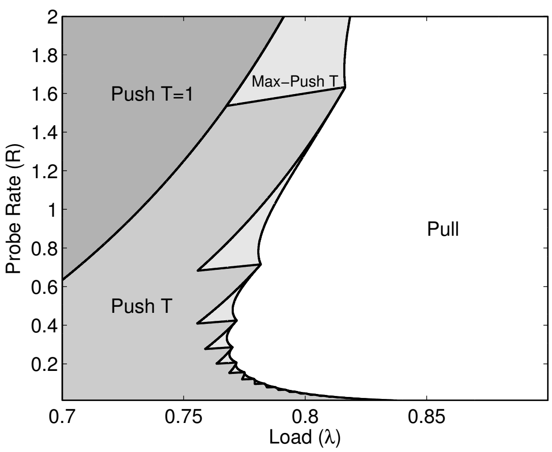

Fig. 10. The different areas identify the (λ, R) combinations for which the rate-based pull strategy (with $T = 1$) is outperformed by the rate-based push strategy with $T = 1$, by the rate-based push strategy with $T \geq 1$ and by the max-push strategy.

$R = 1$ and $\lambda = 0.65, 0.7$ and 0.75. Notice, the max-push strategy with $T = 2$ and $R = 1$ is only properly defined for $\lambda \in [0.6180, 0.7549]$ due to (6), larger λ values would result in the choice of a larger T value. The relative errors are small (below 1% for $N = 100$ nodes) and decrease as N increases. We should note that the probe rates observed during the simulation are well above 1 for small N as illustrated in Figure 8. Hence, the relationship in (10) for the max-push strategy is less accurate than (5) for the rate-based push strategy for small N. Nonetheless, the observed probe rate still seems to decrease to 1 as N tends to infinity.

B. Comparison of pull and max-push strategy

The mean delay of the max-push strategy is depicted in Figure 9. The dots represent the points where $r = \infty$ for the rate-based push strategy, i.e., the positive real roots of $\lambda^{T+1} + (\lambda^T - 1)R = 0$. The max-push strategy nicely interconnects these points as it utilizes the remaining probe rate.

The combination of (λ, R) values for which the pull strategy is outperformed by the rate-based push strategy with $T = 1$, by the rate-based push strategy with $T \geq 1$ and by the max-push strategy, respectively, is shown in Figure 10. The pull strategy is still the most effective for larger loads λ, however, for a large range of (λ, R) values the delay of the pull strategy can be reduced using a rate-based push strategy with $T > 1$ or a max-push strategy.

VII. Finite versus infinite system model

Similar to [8] for the rate-based strategies with $T = 1$, we can define a family of density dependent Markov chains [9] to describe the behavior of the stochastic finite systems with N nodes for both the rate-based pull/push and max-push strategy. In case of the max-push strategy convergence towards the infinite system model over finite time scales follows from Kurtz's well-known theorem [9] and the convergence extends to the stationary regime as we showed that the set of ODEs given by (7)-(9) has a unique global attractor in E_T, due to a result by Benaïm [12].

For the rate-based pull/push strategy with $T \geq 1$ we can rely on the following generalization of Kurtz's theorem [13, Theorem 3.13]:

Theorem 9 (Kurtz). *Consider a family of density dependent CTMCs, with F Lipschitz. Let $\lim_{N\to\infty} X^{(N)}(0) = \tilde{x}$ a.s. and let $\phi_t(\tilde{x})$ be the unique solution to the initial value problem $\frac{d}{dt}x(t) = F(x(t))$ with $x(0) = \tilde{x}$. Consider the path $\{\phi_t(\tilde{x}), t \leq T\}$ for some fixed $T \geq 0$ and assume that there exists a neighborhood K around this path satisfying*

$$\sum_{\ell \in L} |\ell| \sup_{x \in K} \beta_\ell(x) < \infty, \tag{11}$$

then

$$\lim_{N\to\infty} \sup_{t \leq T} \left| X^{(N)}(t) - \phi_t(\tilde{x}) \right| = 0 \quad a.s.$$

For the rate-based pull/push strategy, the above condition (11) corresponds to showing that there exists an environment K such that $\sum_{i \geq 2} \sup_{x \in K}(x_i - x_{i+1}) < \infty$. Such an environment can be shown to exist by repeating the argument for $T = 1$ from [14, Theorem 7]. To show that the convergence extends to the stationary regime, we can make use of a theorem by Benaïm and Le Boudec [15] as in the $T = 1$ case, where the required proof for the tightness of the measures can be proven as in [14].

VIII. Conclusion and future work

In this paper we introduced a new class of rate-based pull and push strategies that can match any predefined maximum allowed probe rate R. This class relied on a threshold parameter T such that jobs can only be exchanged between idle nodes and nodes with a queue length exceeding T, where the class of strategies introduced in [8] corresponds to setting $T = 1$. We derived a closed form expression for the mean delay of this new class of strategies in a homogeneous network with Poisson arrivals and exponential job durations under the so-called infinite system model.

We showed that setting $T = 1$ is optimal for the pull strategies considered, while for the push strategy setting $T > 1$ may reduce the mean delay for some values of (λ, R), i.e., for larger λ and smaller R values. We further introduced the max-push strategy, which utilizes the remaining probe rate capacity in case $R > \lambda^{T+1}/(1-\lambda^T)$, derived a close form expression for its mean delay and (numerically) identified the (λ, R) region where it outperforms the pull strategy.

We proved that the infinite system models of both the rate-based strategies with $T > 1$ and the max-push strategy, are the proper limit processes of the finite stochastic systems with N nodes as N tends to infinity over any finite timescale. Moreover, the convergence was shown to extend to the stationary regime by proving that the ODEs have a global attractor. We validated these theoretical results by simulation, and have shown that the infinite model is an accurate approximation for finite systems of moderate size.

The current results can be extended in a number of ways: networks with finite queues can be considered (this actually makes some of the technical issues less involved), the assumptions on the arrival and service time distribution could be relaxed (which makes the analysis more challenging) or heterogeneous networks could be studied. The class of strategies considered in the paper can also be generalized. For instance, it should be possible to incorporate another parameter B, such that any node with a queue length below B is allowed to accept push requests instead of only the idle nodes.

References

[1] M. Mitzenmacher, "The power of two choices in randomized load balancing," *IEEE Trans. Parallel Distrib. Syst.*, vol. 12, pp. 1094–1104, October 2001.

[2] Y. Lu, Q. Xie, G. Kliot, A. Geller, J. R. Larus, and A. Greenberg, "Join-idle-queue: A novel load balancing algorithm for dynamically scalable web services," *Perform. Eval.*, vol. 68, pp. 1056–1071, 2011.

[3] D. Eager, E. Lazowska, and J. Zahorjan, "A comparison of receiver-initiated and sender-initiated adaptive load sharing," *Perform. Eval.*, vol. 6, no. 1, pp. 53–68, 1986.

[4] R. Mirchandaney, D. Towsley, and J. Stankovic, "Analysis of the effects of delays on load sharing," *IEEE Trans. Comput.*, vol. 38, no. 11, pp. 1513–1525, 1989.

[5] M. Mitzenmacher, "Analyses of load stealing models based on families of differential equations," *Theory of Computing Systems*, vol. 34, pp. 77–98, 2001.

[6] N. Gast and B. Gaujal, "A mean field model of work stealing in large-scale systems," *SIGMETRICS Perform. Eval. Rev.*, vol. 38, no. 1, pp. 13–24, 2010.

[7] R. Mirchandaney, D. Towsley, and J. A. Stankovic, "Adaptive load sharing in heterogeneous distributed systems," *J. Parallel Distrib. Comput.*, vol. 9, no. 4, pp. 331–346, 1990.

[8] W. Minnebo and B. Van Houdt, "Pull versus push mechanism in large distributed networks: Closed form results," in *Proc. of the 24-th International Teletraffic Congress*, Krakau (Poland), 2012.

[9] T. Kurtz, *Approximation of population processes*. Society for Industrial and Applied Mathematics, 1981.

[10] S. Ethier and T. Kurtz, *Markov processes: characterization and convergence*. Wiley, 1986.

[11] K. Deimling, *Ordinary Differential Equations in Banach spaces*. Lect. Notes in Math. 596, 1977.

[12] M. Benaïm, "Recursive algorithms, urn processes and chaining number of chain recurrent sets," *Ergodic Theory and Dynamical Systems*, vol. 18, pp. 53–87, 1998.

[13] M. Mitzenmacher, "The power of two choices in randomized load balancing," Ph.D. dissertation, University of California, Berkeley, 1996.

[14] W. Minnebo and B. Van Houdt, "A fair comparison of pull and push strategies in large distributed networks," *to appear in IEEE/ACM Transactions on Networking*, 2013.

[15] M. Benaïm and J. Le Boudec, "On mean field convergence and stationary regime," *CoRR*, vol. abs/1111.5710, Nov 24 2011.

[16] J. Walker, *Dynamical Systems and Evolution Equations. Theory and Applications.* Plenum Press, New York, 1980.

Appendix A
Proof of Theorem 3

We start by proving the following Lemma:

Lemma 1. *Let $x(t)$ be the unique solution of the set of ODEs given by* (1) *to* (3) *with $x(0) \in E$. The L_1-distance to the unique fixed point $\sum_{i\geq 1} |x_i(t) - \bar{\pi}_i|$ does not increase as a function of t.*

Proof: Define $\epsilon_i(t) = x_i(t) - \bar{\pi}_i$, for $i \geq 1$, such that $\Phi(t) = \sum_{i\geq 1} |\epsilon_i(t)|$ represents the L_1-distance. As $\frac{d}{dt}x_i(t) = \frac{d}{dt}\epsilon_i(t)$ and $\bar{\pi}$ is a fixed point of (1) to (3), we find

$$\begin{aligned}\frac{d}{dt}\epsilon_1(t) &= -\epsilon_1(t)(1+\lambda) - r\epsilon_1(t)(\bar{\pi}_{T+1} + \epsilon_{T+1}(t)) \\ &\quad + r\epsilon_{T+1}(t)(1-\bar{\pi}_1) + \epsilon_2(t), \qquad (12)\\ \frac{d}{dt}\epsilon_i(t) &= \lambda\epsilon_{i-1}(t) - (1+\lambda)\epsilon_i(t) + \epsilon_{i+1}(t), \qquad (13)\end{aligned}$$

for $2 \leq i \leq T$, and

$$\begin{aligned}\frac{d}{dt}\epsilon_i(t) &= \lambda(\epsilon_{i-1}(t) - \epsilon_i(t)) + r\epsilon_1(t)(\bar{\pi}_i - \bar{\pi}_{i+1}) \\ &\quad - (\epsilon_i(t) - \epsilon_{i+1}(t))(1 - r\epsilon_1(t) + r(1-\bar{\pi}_1)), \qquad (14)\end{aligned}$$

for $i > T$. Assume for now that $\epsilon_i(t) \neq 0$ for all i such that $\frac{d}{dt}\Phi(t)$ is properly defined as

$$\frac{d}{dt}\Phi(t) = \sum_{i:\epsilon_i(t)>0} \frac{d}{dt}\epsilon_i(t) - \sum_{i:\epsilon_i(t)<0} \frac{d}{dt}\epsilon_i(t).$$

If $\epsilon_i(t)$ has the same sign for all i, one finds that $\frac{d}{dt}\Phi(t) = -|\epsilon_1(t)|$ by summing (12) to (14), we will show that this inequality also holds in general. Let $I = \{i_1, i_2, \ldots\}$, with $i_1 < i_2 < \ldots$, be the set of indices where $\epsilon_i(t)$ changes sign, that is, $\epsilon_{i-1}(t)$ and $\epsilon_i(t)$ have a different sign if and only if $i \in I$. Assume $\epsilon_1(t) < 0$ and let $I_k = \{i \in I : i \leq T+1\}$, $I_m = \{i \in I : i > T+1\}$,

By means of (12) to (14), we find that if $\epsilon_{i-1}(t)$ and $\epsilon_i(t)$ differ in sign, $\frac{d}{dt}\Phi(t)$ contains an extra term given by

$$sign(\epsilon_i(t))2(\lambda\epsilon_{i-1}(t) - \epsilon_i(t)),$$

for $i = 2, \ldots, T+1$,

$$sign(\epsilon_i(t))2[\lambda\epsilon_{i-1}(t) + r\epsilon_1(t)\bar{\pi}_i - \epsilon_i(t)(1 - r\epsilon_1(t) + r(1-\bar{\pi}_1))],$$

for $i > T+1$. Further, if $\epsilon_1(t)$ and $\epsilon_{T+1}(t)$ differ in sign, $\frac{d}{dt}\Phi(t)$ contains an extra term given by

$$sign(\epsilon_{T+1}(t))2r[\epsilon_1(t)(\bar{\pi}_{T+1} + \epsilon_{T+1}(t)) - \epsilon_{T+1}(t)(1-\bar{\pi}_1)].$$

This implies that for $\epsilon_1(t) \leq 0$

$$\begin{aligned}\frac{d}{dt}\Phi(t) &= \underbrace{\epsilon_1(t)}_{\leq 0} + \alpha \\ &+ 2\sum_{i\in I_m} \underbrace{sign(\epsilon_i(t))\{\lambda\epsilon_{i-1}(t) - \epsilon_i(t)(1 + r(1-\bar{\pi}_1))\}}_{\leq 0} \\ &+ 2\sum_{i\in I_m} sign(\epsilon_i(t))\{r\epsilon_1(t)(\epsilon_i(t) + \bar{\pi}_i)\} \\ &+ 2\sum_{i\in I_k} \underbrace{sign(\epsilon_i(t))\{\lambda\epsilon_{i-1}(t) - \epsilon_i(t)\}}_{\leq 0}\end{aligned}$$

where α is equal to

$$2r\epsilon_1(t)(\epsilon_{T+1}(t) + \bar{\pi}_{T+1}) - \underbrace{2\epsilon_{T+1}(t)r(1-\bar{\pi}_1)}_{\leq 0}.$$

if $\epsilon_{T+1}(t) > 0$ and zero otherwise.

Hence, $\frac{d}{dt}\Phi(t) \leq \epsilon_1(t)$ provided that

$$\sum_{i\in I_m} sign(\epsilon_i(t))(\epsilon_i(t) + \bar{\pi}_i) = \sum_{i\in I_m} sign(\epsilon_i(t))x_i(t) \geq 0,$$

if $\epsilon_{T+1}(t) \leq 0$ and

$$x_{T+1}(t) + \sum_{i\in I_m} sign(\epsilon_i(t))x_i(t) \geq 0,$$

if $\epsilon_{T+1}(t) > 0$.

Let $I_m = \{i_0, i_1, \ldots\}$. In case $\epsilon_{T+1}(t) \leq 0$, the $sign(\epsilon_{i_n}(t))$ is equal to 1 for n even and -1 for n odd. Hence, the condition reduces to

$$\sum_{k\geq 0}(x_{i_{2k}}(t) - x_{i_{2k+1}}(t)) \geq 0,$$

which holds as $x_i(t) \geq x_j(t)$ for $i < j$. Similarly, if $\epsilon_{T+1}(t) > 0$, the $sign(\epsilon_{i_n}(t))$ is equal to -1 for n even and 1 for n odd. Hence, the condition reduces to

$$(x_{T+1}(t) - x_{i_0}(t)) + \sum_{k\geq 0}(x_{i_{2k+1}} - x_{i_{2k+2}}) \geq 0,$$

which again holds as $x_i(t) \geq x_j(t)$ for $i < j$.

Hence, $\frac{d}{dt}\Phi(t) \leq -|\epsilon_1(t)|$ if $\epsilon_1(t) \leq 0$. A similar argument can be used for $\epsilon_1(t) \geq 0$.

Finally, the technical issue of defining $\frac{d}{dt}\Phi(t)$ in case $\epsilon_i(t) = 0$ for some i and $t = t_0$, can be resolved as in [14]. ∎

The above lemma shows that the L_1-distance to the fixed point does not increase along any trajectory $x(t)$ in E, and can only remain the same whenever $x_1(t) = \bar{\pi}_1$ and there are no sign changes in the $\epsilon_i(t)$'s.

Lemma 2. *The only trajectory $x(t)$ of the ODEs given by* (1) *to* (3) *with $x(0) \in E$ for which the L_1-distance does not decrease is given by $x(t) = \bar{\pi}$ for all t.*

Proof: From the proof of Lemma 1, we know that $x_1(t) = \bar{\pi}_1 = \lambda$ for all t, whenever the L_1-distance does not decrease. Equation (1) therefore implies that $x_{T+1}(t) = \frac{\lambda^2 - x_2(t)}{r(1-\lambda)}$ on such a trajectory. Hence, if $x_2(t) = \bar{\pi}_2 + c$, then

$$x_{T+1}(t) = \frac{\lambda^2 - \bar{\pi}_2(t)}{r(1-\lambda)} - \frac{c}{r(1-\lambda)} = \bar{\pi}_{T+1} - \frac{c}{r(1-\lambda)}.$$

Hence, $\epsilon_2(t) = x_2(t) - \bar{\pi}_2$ and $\epsilon_{T+1}(t) = x_{T+1}(t) - \bar{\pi}_{T+1}$ differ in sign unless $x_2(t) = \bar{\pi}_2$ and $x_{T+1}(t) = \bar{\pi}_{T+1}$. The fact that $x(t) = \bar{\pi}$ on such a trajectory now follows from (1) to (3). ∎

We now recall La Salle's invariance principle for Banach spaces, where a (positively) invariant subset of $K \subset E$ of an ODE defined on E is such that $x(t) \in K$ for all t provided that $x(t)$ is the unique solution of the ODE with $x(0) \in K$.

Theorem 10 ([16]). *Let $V(x)$ be a continuous real valued function from E to $\mathbb{R}$ with $\frac{d}{dt}V(x) = \limsup_{t\to 0^+} \frac{1}{t}(V(x(t)) - V(x)) \leq 0$, where $x(t)$ is the unique solution of an ODE with $x(0) = x$. Let $K = \{x \in E | \frac{d}{dt}V(x) = 0\}$ and let M be the largest (positively) invariant subset of K. If $x(t)$ is precompact (i.e., remains in a compact set) for $x(0) \in E$, then*

$$\lim_{t\to\infty} dist(x(t), M) = 0,$$

where $dist(x, M)$ represents the Banach distance between the point $x \in E$ and the set $M \subset E$.

Using La Salle's invariance principle, Theorem 3 can be proven analogously to [14, Theorem 3] with $V(x)$ equal to the L_1-distance.

Appendix B
Proof of Theorem 8

We start by proving the following Lemma:

Lemma 3. *Let $x(t)$ be the unique solution of the ODEs given by* (7) *to* (9) *with $x(0) \in E_T$. The L_1-distance to the unique fixed point $\sum_{i\geq 1}|x_i(t) - \dot{\pi}_i|$ does not increase as a function of t.*

Proof: Using the same definitions as in Appendix A, we find

$$\begin{aligned}\frac{d}{dt}\epsilon_1(t) &= \lambda(\epsilon_T(t) - \epsilon_1(t)) - (\epsilon_1(t) - \epsilon_2(t)) \\ &\quad + r\epsilon_T(t)(1 - (\epsilon_1(t) + \dot{\pi}_1)) - r\dot{\pi}_T\epsilon_1(t), \quad (15)\\ \frac{d}{dt}\epsilon_i(t) &= \lambda\epsilon_{i-1}(t) - (1+\lambda)\epsilon_i(t) + \epsilon_{i+1}(t), \quad (16)\end{aligned}$$

for $1 < i < T$, and

$$\begin{aligned}\frac{d}{dt}\epsilon_T(t) &= \lambda(\epsilon_{T-1}(t) - \epsilon_T(t)) + r\epsilon_1(t)\dot{\pi}_T \\ &\quad - \epsilon_T(t)(1 + r(1 - (\epsilon_1(t) + \dot{\pi}_1))). \quad (17)\end{aligned}$$

Assume for now that $\epsilon_i(t) \neq 0$ for all i such that $\frac{d}{dt}\Phi(t)$ is properly defined. If $\epsilon_i(t)$ has the same sign for all i, one finds that $\frac{d}{dt}\Phi(t) = -|\epsilon_1(t)|$ by summing (15) to (17), we will show that this inequality also holds in general. If $\epsilon_i(t)$ and $\epsilon_{i-1}(t)$ differ in sign, $\frac{d}{dt}\Phi(t)$ changes by

$$sign(\epsilon_i(t))2(\lambda\epsilon_{i-1}(t) - \epsilon_i(t)),$$

for $i = 2, \ldots, T$, while a difference in sign between the terms $\epsilon_1(t)$ and $\epsilon_T(t)$ creates a term of the form

$$sign(\epsilon_T(t))2\{\lambda\epsilon_T(t) + r\epsilon_1(t)\dot{\pi}_T - \epsilon_T(t)r(1 - (\epsilon_1(t) + \dot{\pi}_1))\}.$$

Assume $\epsilon_1(t) \leq 0$, then we find

$$\frac{d}{dt}\Phi(t) = \underbrace{\epsilon_1(t)}_{\leq 0} + \alpha + 2\sum_{i \in I}\{\underbrace{sign(\epsilon_i(t))(\lambda\epsilon_{i-1}(t) - \epsilon_i(t))}_{\leq 0}\} \tag{18}$$

where α is equal to

$$\alpha = \underbrace{-2\lambda\epsilon_T(t)}_{\leq 0} \underbrace{-2r\epsilon_T(t)(1 - (\epsilon_1(t) + \dot{\pi}_1))}_{\leq 0} + \underbrace{2r\dot{\pi}_T\epsilon_1(t)}_{\leq 0}. \tag{19}$$

if $\epsilon_T(t) > 0$ and $\alpha = 0$ otherwise. Hence, $\frac{d}{dt}\Phi(t) \leq \epsilon_1(t)$. A similar argument can be used for $\epsilon_1(t) > 0$ by reversing all the signs.

As in Appendix A, the technical issue of defining $\frac{d}{dt}\Phi(t)$ in case $\epsilon_i(t) = 0$ for some i and $t = t_0$ is resolved by relying on the upper right-hand derivative (as in [1, Theorem 3]). ∎

The above lemma shows that the L_1-distance to the fixed point does not increase along any trajectory $x(t)$ in E_T, and only remains the same whenever $x_1(t) = \pi_1$ (as $\epsilon_1(t) = 0$ in such a case).

Lemma 4. *The only trajectory $x(t)$ of the ODE given by* (7) *to* (9) *with $x(0) \in E_T$ for which the L_1-distance does not decrease is given by $x(t) = \dot{\pi}$ for all t.*

Proof: If $x_1(t) = \dot{\pi}_1 = \lambda$ for all t, then (7) implies that $x_T(t) = \frac{\lambda^2 - x2(t)}{\lambda + r(1-\lambda)}$ and the proof proceeds as in Lemma 2. ∎

Using La Salle's invariance principle for Banach spaces as given by Theorem 10, we can now prove theorem 8:

Proof of Theorem 8: We rely on La Salle's invariance principle for Banach spaces by setting $V(x)$ equal to the L_1-distance to the fixed point, i.e., $V(x) = \sum_{i=1}^{T} |x_i - \dot{\pi}_i|$. Lemma 3 implies that $\frac{d}{dt}V(x) \leq 0$, while Lemma 4 shows that $M = \{\dot{\pi}\}$ is a singleton. Hence, $\dot{\pi}$ is a global attractor since E_T itself is a compact set and all trajectory are contained within E_T by definition. ∎