

Reduction of a polling network to a single node

Paul Beekhuizen^{*†}

Dee Denteneer[‡]

Jacques Resing[§]

February 1, 2008

Abstract

We consider a discrete-time tree network of polling servers where all packets are routed to the same node (called node 0), from which they leave the network. All packets have unit size and arrive from the exterior according to independent batch Bernoulli arrival processes. The service discipline of each node is work-conserving and the service discipline of node 0 has to be *HoL-based*, which is an additional assumption that is satisfied by, a.o., m_i -limited service, exhaustive service, and priority disciplines.

Let a type i packet be a packet that visits queue i of node 0. We establish a distributional relation between the number of type i packets in the network and in a *single* station system, and we show equality of the mean end-to-end delay of type i packets in the two systems. Essentially this reduces an arbitrary tree network to a much simpler system of one node, while preserving the mean end-to-end delay of type i packets.

1 Introduction

In polling systems jobs requiring a certain amount of service arrive to multiple queues sharing a single server. The fact that all queues are served by the same server leads to all kinds of research topics, such as determining the mean waiting time in each queue, optimisation of the order in which queues are served, optimisation of the queue to which each job is routed, etc. Polling models have many applications, for example in telecommunications, transportation, healthcare, etc., and have been the subject of numerous studies (for a recent overview, see [13]).

Few attempts, however, have been made to analyse *networks* of polling servers; one of the rare examples is a heavy-traffic study [7]. We study a discrete-time polling network with unit service times and show that this network can be reduced to a single node. This reduction allows for an analysis of the network via a simpler analysis of the single-node system.

The main application that motivates this study is a Network-on-Chip (see [2]), where multiple processors share a single memory and data is transmitted via routers. The range of applications of our study is of course much broader; for instance a number of workstations connected to a single server via a number of routers, a manufacturing environment where all jobs require the same final operation after a number of intermediate operations, etc.

We consider a concentrating tree network with an arbitrary number of nodes, as depicted in Figure 1a. Packets of fixed size 1 arrive from the exterior of the network to each node, where they are stored and eventually transmitted to the next node. The network operates in discrete time and we assume that packets arrive in batches at slot boundaries.

All packets leave the network via the same node (the sink), which we call node 0. Node 0 is a single server polling system with N queues. Queue i may store packets arriving from some node upstream (which we call node i) or from the exterior directly (in which case node i does not exist). For the other nodes in the network it is unimportant which number identifies them. We refer to packets that pass through queue i of node 0 as ‘type i ’ packets, for $i = 1, \dots, N$.

^{*}Philips Research, Digital Signal Processing group, Eindhoven, The Netherlands

[†]EURANDOM, Eindhoven, e-mail: beekhuizen@eurandom.tue.nl

[‡]Philips Research, Connectivity Systems and Networks group, Eindhoven

[§]Eindhoven University of Technology, Department of Mathematics and Computer Science




Figure 1: A schematic representation of the reduction. We consider an arbitrary tree network for which Assumption 2.1 is satisfied and reduce it to a single node while preserving the mean end-to-end delay of type i packets. In the figure all queues of node 0 store packets coming from nodes upstream, but it is also possible that they store packets directly arriving from the exterior (but not both).

In the remainder, we think of each node as a polling system with an arbitrary unspecified number of queues. All nodes have work-conserving service disciplines (i.e., the server has to serve precisely one packet if there is at least one and it may not create or remove additional packets). Except for node 0, the service disciplines are arbitrary.

The service discipline of node 0 is not only work-conserving, but also *HoL-based*. This means that the decision which queue is served may only depend on whether or not queues are occupied, and not on the *number* of packets present in each queue. Examples of HoL-based service disciplines are 1-limited service (if queue i was served in the previous slot, serve queue $i + 1$ if it has a packet), exhaustive service (serve queue i again if there is another packet left), and priority disciplines (serve queue 1 if there is a packet in queue 1, serve queue 2 if there is one in queue 2 but not in queue 1, etc.). Service disciplines that are not HoL-based are disciplines such as gated service and longest or shortest queue first.

We construct a system consisting of one node (see Figure 1b), called the reduced system. The reduced system uses the same service discipline as node 0 and its arrival processes are given by superpositions of the arrivals to nodes upstream of node i . We establish a relation in distribution between the contents of the network and the reduced system and we show that the mean end-to-end delay of type i packets in the network and the reduced system are equal.

The reason why the service discipline of node 0 has to be HoL-based for our results to be applicable is the following: As the arrival processes to the reduced system are given by the superposition of arrival processes in the network, the number of packets in queue i of node 0 is typically smaller than the number of packets in queue i of the reduced system. If the server is allowed to make decisions based on the number of packets in the queues, the behaviour in both systems might be entirely different, which would hence lead to different mean end-to-end delays. Since HoL-based service disciplines may not use any information other than whether or not queues are empty, we can show that the behaviour in the reduced and original system is stochastically identical if a HoL-based service discipline is used. We do so by introducing a coupled system with one node, of which queue i is empty if and only if queue i of node 0 of the original system is empty. In Section 4 we give an example without a HoL-based service discipline for which the reduction fails.

There is a large amount of literature available on reductions of networks. The earliest results are those of Avi-Itzhak [1] and Friedman [3], where a tandem network of deterministic multi-server queues with an arbitrary arrival process is considered and reduced to a single node. Rubin [8, 9]

studies a tandem network that is based on packet-switched communication networks. Mandjes and Van Uitert [4] apply the results of Avi-Itzhak and Friedman to obtain results for the overflow probability in a two-queue tandem network.

These papers concern tandem networks without cross-traffic (i.e., with a single source of arrivals). Rubin [10] devises an approximation for a tandem network with deterministic service times and cross traffic. Shalmon and Kaplan [12] perform an exact analysis of a tandem network with Poisson arrivals and deterministic packet sizes with FIFO service order and Shalmon [11] extends this analysis to priority disciplines. Neely, Rohrs and Modiano [6] analyse a continuous-time concentrating tree network with general arrival processes at each of the sources and reduce this to a two-stage equivalent network.

The result that is most related to our work is that of Morrison [5]. He first considers a tree network where all packets pass through at most two nodes before leaving the network. He constructs a reduced system consisting of one node and establishes a sample path relation between the queue contents in both systems. Morrison argues that this reduction can be applied repeatedly to arbitrary tree networks but restricts his analysis to arbitrary tandem queues.

At the heart of our reduction result is a deterministic relation on sample paths as well. This relation is an extension of Morrison's since the latter is a relation between the *total* contents of the reduced system and the network, while we establish a relation between the number of type i packets in the reduced system and the network. For HoL-based service disciplines, Morrison's result follows from our sample path relation by summing over all i . In addition to this, we address the precise implications of this relation for a network with stochastic arrivals.

This paper is organised as follows: In Section 2 we describe the model and the reduced system in more detail and we state our main results. These results are then proved in Section 3. In Section 4 we give an example that shows why HoL-based service disciplines are required. A conclusion is drawn in Section 5.

2 Formalisation

In Section 2.1, we formalise the network model and the condition the service discipline of node 0 has to satisfy. In Section 2.2 we define the reduced system and present a more precise formulation of the main result.

2.1 The original system

We define $\mathcal{N}^{(m)}$ as the set of nodes whose output reaches node m at some point in time, combined with node m itself. Recall that if queue $i = 1, \dots, N$ of node 0 stores packets directly arriving from the exterior, node i does not exist. In this case, we define $\mathcal{N}^{(i)} = \emptyset$.

We denote the total number of packets (i.e., summed over all queues) arriving from the exterior to node m at time t by $X_t^{(m)}$. For all queues i of node 0 that store packets directly arriving from the exterior, we denote the number of arriving packets by $X_{i,t}^{(0)}$. We assume that the external arrivals are governed by independent batch Bernoulli processes, i.e., $X_t^{(m)}$ and $X_{i,t}^{(0)}$ are both i.i.d. with respect to t , with generic random variables $X^{(m)}$ and $X_i^{(0)}$, and expectation $\lambda^{(m)}$ and $\lambda_i^{(0)}$. The notation used here is indicative of the notation used throughout the paper: We add the node as a superscript and the queue (if applicable) as subscript.

The total arrival rate at node m , $\gamma^{(m)}$, comprises both external arrivals and arrivals from nodes upstream. Observe that every packet arriving at node m must have arrived from the exterior at some node upstream or at m itself:

$$\gamma^{(m)} = \sum_{l \in \mathcal{N}^{(m)}} \lambda^{(l)} \quad \text{for all } m. \quad (2.1)$$

We assume $\gamma^{(0)} < 1$ so all queues are stable.

We denote the number of packets in all queues of node $m \neq 0$ at time t by $Q_t^{(m)}$. Arrivals in time slot $(t-1, t]$ take place at time t and are not in the queue before time t . A packet that arrives at time t may be served in time slot $(t, t+1]$ and leave at time $t+1$. For convenience we assume that the network starts operating at time 0: If $t \leq 0$ then $X_t^{(m)} = 0$. Consequently, $Q_t^{(m)} = 0$ for all $t \leq 0$. We denote the steady state queue lengths by dropping the subscript t : $Q_i^{(0)}$ and $Q^{(m)}$.

We denote the queue the server of node 0 serves in $(t, t+1]$ (or the position the server moves to at time t) by $P_t^{(0)} \in \{0, \dots, N\}$, where $P_t^{(0)} = 0$ if the server is idle in $(t, t+1]$. Precisely one packet departs from queue i at time $t+1$ if and only if $P_t^{(0)} = i$, so

$$Q_{i,t+1}^{(0)} = \begin{cases} Q_{i,t}^{(0)} + \varepsilon(Q_t^{(i)}) - 1(P_t^{(0)} = i) & \text{if } \mathcal{N}^{(i)} \neq \emptyset \\ Q_{i,t}^{(0)} + X_{i,t+1}^{(0)} - 1(P_t^{(0)} = i) & \text{if } \mathcal{N}^{(i)} = \emptyset \end{cases} \quad (2.2)$$

where $Q_{i,t}^{(0)}$ is the length of queue i at time t and $\varepsilon(x) = 1(x > 0)$.

We define a vector $\mathbf{Q}_t^{(0)} = (Q_{1,t}^{(0)}, \dots, Q_{N,t}^{(0)})$ containing all the queue lengths, and $\varepsilon(\mathbf{Q}_t^{(0)}) = (\varepsilon(Q_{1,t}^{(0)}), \dots, \varepsilon(Q_{N,t}^{(0)}))$.

Assumption 2.1. Node 0 uses a HoL-based service discipline, which means that it satisfies

$$\begin{aligned} \mathbb{P}\left(P_t^{(0)} = j \mid P_{t-1}^{(0)}, \dots, P_1^{(0)}, \mathbf{Q}_t^{(0)}, \dots, \mathbf{Q}_1^{(0)}\right) \\ = \mathbb{P}\left(P_t^{(0)} = j \mid P_{t-1}^{(0)}, \dots, P_{t-M}^{(0)}, \varepsilon(\mathbf{Q}_t^{(0)}), \dots, \varepsilon(\mathbf{Q}_{t-M}^{(0)})\right), \end{aligned}$$

for some (finite) M . In other words, the server makes a (random) decision which queue it starts serving at time t based on whether there were packets in the queues and the queues it served during a history of M time slots.

Two main classes of service disciplines are HoL-based: The first is Bernoulli scheduling, where after a service completion at queue i , the server decides to serve queue i again with probability p_i and moves to one of the other queues with probability $1 - p_i$. This class also contains the 1-limited and exhaustive service disciplines as extreme cases ($p_i = 0$ and $p_i = 1$ respectively). The second main class is the class of m_i -limited service disciplines, where the server serves at most m_i packets from a queue before moving to one of the other queues. The history parameter M ensures that this is allowed; provided $m_i \leq M$, the number of packets served consecutively is contained in $P_{t-1}^{(0)}, \dots, P_{t-M}^{(0)}$.

If the server moves to one of the other queues it may choose among the non-empty queues according to some fixed order, such as round robin, a polling table, or a priority index, or according to a random order that is independent of the queue lengths.

Service disciplines that are not HoL-based for instance take arrival times or deadlines into account (e.g., earliest arrival time first, earliest deadline first), or queue lengths (e.g., gated service and shortest/longest queue first). In Section 4 we give an example that shows why our reduction fails if Assumption 2.1 is violated.

Remark 2.2. The history parameter M has to be finite due to a technicality. We will come back to this issue in Remark 3.6.

Remark 2.3. Assumption 2.1 also prevents the server of node 0 to take into account the size of the batch in which a packet arrived, since this information is not contained in $\varepsilon(\mathbf{Q}_{t-M}^{(0)}), \dots, \varepsilon(\mathbf{Q}_t^{(0)})$.

Remark 2.4. We have not specified in which order packets are served *within* each queue (for instance FIFO, LIFO, etc.). We only specify that a packet has to be served from a queue that is selected according to Assumption 2.1. The specific order in which packets are served within a queue is irrelevant for our results.

The mean sojourn time of packets in node m is denoted by $\mathbb{E}[S^{(m)}]$, for all m . We also define $\mathbb{E}[S_i^{(0)}]$ as the mean sojourn time of a packet in queue i of node 0. The corresponding mean waiting times (or delays) are denoted by a W : $\mathbb{E}[W^{(m)}] = \mathbb{E}[S^{(m)}] - 1$ and $\mathbb{E}[W_i^{(0)}] = \mathbb{E}[S_i^{(0)}] - 1$.

In order to describe the end-to-end delay of type i packets, we introduce the set $P(m, l)$ containing every node on the path from m to l , including m and l themselves. That is,

$$P(m, l) = \{k : k \in \mathcal{N}^{(l)} \text{ and } m \in \mathcal{N}^{(k)}\} \quad \text{for } m \in \mathcal{N}^{(l)}. \quad (2.3)$$

The distance from node m to l , with $m \in \mathcal{N}^{(l)}$, is given by

$$d(m, l) = |P(m, l)| - 1,$$

so that $d(m, m) = 0$, $d(m, l) = 1$ if the output of m goes to l directly, $d(m, l) = 2$ if there is one intermediate node, and so on.

The mean end-to-end delay of type i packets is now given by

$$\mathbb{E}[W_i] = \sum_{m \in \mathcal{N}^{(i)}} \frac{\lambda^{(m)}}{\gamma^{(i)}} \sum_{l \in P(m, i)} \mathbb{E}[W^{(l)}] + \mathbb{E}[W_i^{(0)}].$$

The reasoning behind this equation is the following: A fraction $\lambda^{(m)}/\gamma^{(i)}$ of the packets that arrive at queue i of node 0, arrived from the exterior at node m . The mean end-to-end delay of these packets is the sum of waiting times at nodes from m to i . Note that if $\mathcal{N}^{(i)} = \emptyset$, $\mathbb{E}[W_i] = \mathbb{E}[W_i^{(0)}]$, i.e., if packets arrive from the exterior to queue i of node 0 directly, their mean end-to-end delay is given only by their mean waiting time in queue i of node 0.

2.2 The reduced system

In this section we describe the construction of the reduced system and state our main result. All quantities related to the reduced system are denoted by primes next to the normal letters. We will refer to the network formalised in Section 2.1 as the original system.

The arrivals to the reduced system are constructed by aggregating all arrivals to the node i subtree of the original system ($i = 1, \dots, N$) into a single stream. We let $X'_{i,t}$ denote the number of arrivals to queue i of the reduced system at time t . The random variables $X'_{i,t}$ are i.i.d. with respect to t , with generic random variable X'_i given by

$$X'_i \stackrel{d}{=} \begin{cases} \sum_{m \in \mathcal{N}^{(i)}} X^{(m)}, & \text{if } \mathcal{N}^{(i)} \neq \emptyset \\ X_i^{(0)}, & \text{if } \mathcal{N}^{(i)} = \emptyset. \end{cases}$$

The evolution of the queues can now be described as follows:

$$Q'_{i,t+1} = Q'_{i,t} + X'_{i,t+1} - 1(P'_t = i),$$

where P'_t is the position of the server of the reduced system at time t , and $Q'_{i,t}$ is the length of queue i at time t . The steady state queue length is denoted by Q'_i .

The service discipline of the reduced system is the same HoL-based discipline as that of node 0 in the original system. In practice, this means that if node 0 uses m_i -limited, then so must the reduced system, if node 0 uses exhaustive service, then so must the reduced system, etc. In a more general setting, we mean the following: Given identical server positions and queue contents during the last M time slots, the probability that a certain queue is served is the same under both service disciplines:

$$\begin{aligned} \mathbb{P}\left(P'_t = p_t \mid \varepsilon(\mathbf{Q}'_t) = \varepsilon(\mathbf{q}_t), \dots, \varepsilon(\mathbf{Q}'_{t-M}) = \varepsilon(\mathbf{q}_{t-M}), P'_{t-1} = p_{t-1}, \dots, P'_{t-M} = p_{t-M}\right) = \\ \mathbb{P}\left(P_t^{(0)} = p_t \mid \varepsilon(\mathbf{Q}_t^{(0)}) = \varepsilon(\mathbf{q}_t), \dots, \varepsilon(\mathbf{Q}_{t-M}^{(0)}) = \varepsilon(\mathbf{q}_{t-M}), P_{t-1}^{(0)} = p_{t-1}, \dots, P_{t-M}^{(0)} = p_{t-M}\right). \end{aligned} \quad (2.4)$$

The mean waiting time of a type i packet in the reduced system is denoted by $\mathbb{E}[W'_i]$. Our two main results are the following:

Theorem 2.5 (Queue length reduction). *In steady-state, the length of queue i of the reduced system is in distribution equal to the number of type i packets in the network, minus external arrivals to node m in the last $d(m, 0)$ time slots, summed over all m :*

$$Q'_i \stackrel{d}{=} \lim_{t \rightarrow \infty} \left(\sum_{m \in \mathcal{N}^{(i)}} Q_t^{(m)} + Q_{i,t}^{(0)} - \sum_{m \in \mathcal{N}^{(i)}} \sum_{d=1}^{d(m,0)} X_{t+1-d}^{(m)} \right). \quad (2.5)$$

Theorem 2.6 (Waiting time reduction). *The mean end-to-end delay (total waiting time) of type i packets is the same in the original and the reduced system:*

$$\mathbb{E}[W'_i] = \mathbb{E}[W_i].$$

Remark 2.7. Theorem 2.5 also implies that the number of type i packets in the single-node system is stochastically smaller than that in the network: $Q'_i \leq_d \sum_{m \in \mathcal{N}^{(i)}} Q_t^{(m)} + Q_{i,t}^{(0)}$. We expect that this bound is tight in heavy traffic.

3 Proof of the main results

In this section, we prove Theorem 2.5 and 2.6. In order to do so, we use a coupling argument: In Section 3.1, we introduce another single station system, called the *coupled* system. The arrivals to this system are constructed in a deterministic way from the arrivals to the original system. This deterministic construction allows us to prove a sample path relation between the *original* system and the *coupled* system in Section 3.2. In Section 3.3 we show that the queue lengths of the *coupled* system and the *reduced* system are equal in distribution. These two relations combined provide us with a relation between the *original* and *reduced* system, namely Theorem 2.5. In Section 3.4 we prove Theorem 2.6 by applying Little's law to Theorem 2.5.

3.1 Coupling

In this section, we construct a single node from the original system, as defined in Section 2.1. All quantities related to this system are denoted by tildes above the normal letters.

We define

$$\tilde{X}_{i,t} = \begin{cases} \sum_{m \in \mathcal{N}^{(i)}} X_{t-d(m,0)}^{(m)}, & \text{if } \mathcal{N}^{(i)} \neq \emptyset, \\ X_{i,t}^{(0)}, & \text{if } \mathcal{N}^{(i)} = \emptyset. \end{cases} \quad (3.1)$$

The evolution of the queues can now be described as follows:

$$\tilde{Q}_{i,t+1} = \tilde{Q}_{i,t} + \tilde{X}_{i,t+1} - 1(\tilde{P}_t = i), \quad (3.2)$$

where \tilde{P}_t is the position of the server.

We couple the service discipline to that of node 0 in the original system in the following way: If both servers have the same history, their next position will also be the same, i.e., if $\varepsilon(\tilde{Q}_s) = \varepsilon(Q_s^{(0)})$, for all $t - M \leq s \leq t$, and $\tilde{P}_s = P_s^{(0)}$, for all $t - M \leq s < t$, then $\tilde{P}_t = P_t^{(0)}$. In Section 3.2 we prove that $\tilde{P}_t = P_t^{(0)}$ and $\varepsilon(\tilde{Q}_t) = \varepsilon(Q_t^{(0)})$ for all t .

The intuition behind this coupling is that a packet reaches node 0 in the original system after or at the same time it arrives to the coupled system; a packet requires *at least* $d(m, 0)$ time slots to reach node 0 from node m , which is precisely the number of time slots by which an arrival to the coupled system is delayed. Using induction we can show that queue i of the coupled system and node 0 of the original system are empty at precisely the same times. Together with the (inductive)

definition of \tilde{P}_t this implies a sample path relation between the *original* system and the *coupled* system for all t .

Furthermore, the steady state queue lengths of the *coupled* and *reduced* systems are stochastically the same due to the following argument: By the nature of batch Bernoulli arrival processes, imposing a time-delay does not stochastically change them. In addition to this, the service disciplines of the coupled and the reduced system are shown to be stochastically equal. As the arrival processes and service disciplines of the coupled and reduced system are stochastically identical, their queue lengths must be so too.

3.2 The original and the coupled system

We establish a sample path relation between the original and the coupled system for all t , which leads to a similar steady state relation. We first give Lemma 3.1 for further reference. Next, we obtain our sample path relation in Proposition 3.2. Taking limits then yields the steady state relation (Cor. 3.3).

We first describe the evolution of the *total* contents of the nodes. In order to do so, we define for all m , $\mathcal{N}_1^{(m)}$ as the set of nodes whose output enters node m directly, i.e.,

$$\mathcal{N}_1^{(m)} = \{l \in \mathcal{N}^{(m)} : d(l, m) = 1\}.$$

The total number of packets arriving to node m is given by the number of external arrivals plus the arrivals from nodes upstream of m , and one packet is served if there is at least one packet present, so for all m ,

$$Q_{t+1}^{(m)} = Q_t^{(m)} + X_{t+1}^{(m)} + \sum_{l \in \mathcal{N}_1^{(m)}} \varepsilon(Q_t^{(l)}) - \varepsilon(Q_t^{(m)}). \quad (3.3)$$

The following lemma is presented for further reference:

Lemma 3.1. *For all t and $i = 1, \dots, N$,*

$$\sum_{m \in \mathcal{N}^{(i)}} Q_t^{(m)} = \sum_{m \in \mathcal{N}^{(i)}} Q_{t-d(m,0)}^{(m)} + \sum_{m \in \mathcal{N}^{(i)}} \sum_{d=1}^{d(m,0)} X_{t+1-d}^{(m)} - \sum_{m \in \mathcal{N}^{(i)}} \varepsilon(Q_{t-d(m,0)}^{(m)}). \quad (3.4)$$

Proof. We first observe that, by (3.3),

$$Q_t^{(m)} - Q_{t-1}^{(m)} = X_t^{(m)} - \varepsilon(Q_{t-1}^{(m)}) + \sum_{l \in \mathcal{N}_1^{(m)}} \varepsilon(Q_{t-1}^{(l)}).$$

If we apply this argument $d(m, 0)$ times, we obtain

$$Q_t^{(m)} - Q_{t-d(m,0)}^{(m)} = \sum_{d=1}^{d(m,0)} \left(X_{t+1-d}^{(m)} - \varepsilon(Q_{t-d}^{(m)}) + \sum_{l \in \mathcal{N}_1^{(m)}} \varepsilon(Q_{t-d}^{(l)}) \right),$$

for all m . Summing over all $m \in \mathcal{N}^{(i)}$ yields

$$\begin{aligned} \sum_{m \in \mathcal{N}^{(i)}} Q_t^{(m)} &= \sum_{m \in \mathcal{N}^{(i)}} Q_{t-d(m,0)}^{(m)} + \sum_{m \in \mathcal{N}^{(i)}} \sum_{d=1}^{d(m,0)} X_{t+1-d}^{(m)} - \sum_{m \in \mathcal{N}^{(i)}} \sum_{d=1}^{d(m,0)} \varepsilon(Q_{t-d}^{(m)}) + \\ &\quad \sum_{m \in \mathcal{N}^{(i)}} \sum_{l \in \mathcal{N}_1^{(m)}} \sum_{d=1}^{d(m,0)} \varepsilon(Q_{t-d}^{(l)}). \end{aligned} \quad (3.5)$$

Observe that in the double summation over the nodes in the last term of (3.5), we include for each $m \in \mathcal{N}^{(i)}$ all l immediately upstream of m . We thus sum over all $l \in \mathcal{N}^{(i)}$, except $l = i$. As $d(m, 0) = d(l, 0) - 1$ for such m and l , we get

$$\sum_{m \in \mathcal{N}^{(i)}} \sum_{l \in \mathcal{N}_1^{(m)}} \sum_{d=1}^{d(m,0)} \varepsilon(Q_{t-d}^{(l)}) = \sum_{\substack{l \in \mathcal{N}^{(i)} \\ l \neq i}} \sum_{d=1}^{d(l,0)-1} \varepsilon(Q_{t-d}^{(l)}) = \sum_{l \in \mathcal{N}^{(i)}} \sum_{d=1}^{d(l,0)-1} \varepsilon(Q_{t-d}^{(l)}), \quad (3.6)$$

since $d(i, 0) = 1$. Substitution of the latter expression in (3.5) indeed yields (3.4). \square

Proposition 3.2. *For all t and $i = 1, \dots, N$,*

$$\tilde{Q}_{i,t} = \sum_{m \in \mathcal{N}^{(i)}} Q_t^{(m)} + Q_{i,t}^{(0)} - \sum_{m \in \mathcal{N}^{(i)}} \sum_{d=1}^{d(m,0)} X_{t+1-d}^{(m)}, \quad (3.7)$$

$$\varepsilon(\tilde{Q}_{i,t}) = \varepsilon(Q_{i,t}^{(0)}), \quad (3.8)$$

$$\tilde{P}_t = P_t^{(0)}. \quad (3.9)$$

Proof. We prove this lemma by induction. Note that (3.7)-(3.9) trivially hold for $t \leq 0$ due to the assumption that the systems are initially empty. We hypothesise that the proposition is true for $1, \dots, t$ and prove it for $t + 1$. First, assume $\mathcal{N}^{(i)} \neq \emptyset$. By (3.1), (3.2), and the induction hypothesis,

$$\begin{aligned} \tilde{Q}_{i,t+1} &= \tilde{Q}_{i,t} + \tilde{X}_{i,t+1} - 1(\tilde{P}_t = i) \\ &= \sum_{m \in \mathcal{N}^{(i)}} Q_t^{(m)} + Q_{i,t}^{(0)} - \sum_{m \in \mathcal{N}^{(i)}} \sum_{d=1}^{d(m,0)} X_{t+1-d}^{(m)} + \sum_{m \in \mathcal{N}^{(i)}} X_{t+1-d(m,0)}^{(m)} - 1(P_t^{(0)} = i) \\ &= \sum_{m \in \mathcal{N}^{(i)}} Q_t^{(m)} + Q_{i,t}^{(0)} - \sum_{m \in \mathcal{N}^{(i)}} \sum_{d=1}^{d(m,0)-1} X_{t+1-d}^{(m)} - 1(P_t^{(0)} = i). \end{aligned}$$

By applying (2.2) and (3.3) we obtain

$$\begin{aligned} \tilde{Q}_{i,t+1} &= \sum_{m \in \mathcal{N}^{(i)}} \left(Q_{t+1}^{(m)} + \varepsilon(Q_t^{(m)}) - X_{t+1}^{(m)} - \sum_{l \in \mathcal{N}_1^{(m)}} \varepsilon(Q_t^{(l)}) \right) + Q_{i,t+1}^{(0)} - \varepsilon(Q_t^{(i)}) \\ &\quad - \sum_{m \in \mathcal{N}^{(i)}} \sum_{d=1}^{d(m,0)-1} X_{t+1-d}^{(m)} \\ &= \sum_{m \in \mathcal{N}^{(i)}} Q_{t+1}^{(m)} + \sum_{m \in \mathcal{N}^{(i)}} \varepsilon(Q_t^{(m)}) - \sum_{m \in \mathcal{N}^{(i)}} \sum_{l \in \mathcal{N}_1^{(m)}} \varepsilon(Q_t^{(l)}) + Q_{i,t+1}^{(0)} - \varepsilon(Q_t^{(i)}) \\ &\quad - \sum_{m \in \mathcal{N}^{(i)}} \sum_{d=0}^{d(m,0)-1} X_{t+1-d}^{(m)}. \end{aligned}$$

We proceed by interchanging the order of summation in the first double sum. Similar to the proof of Lemma 3.1, we include for all $m \in \mathcal{N}^{(i)}$ all l immediately upstream of m . This implies that we sum over all $l \in \mathcal{N}^{(i)}$, except $l = i$. We thus get (cf. (3.6))

$$\sum_{m \in \mathcal{N}^{(i)}} \sum_{l \in \mathcal{N}_1^{(m)}} \varepsilon(Q_t^{(l)}) = \sum_{\substack{l \in \mathcal{N}^{(i)} \\ l \neq i}} \varepsilon(Q_t^{(l)}) = \sum_{l \in \mathcal{N}^{(i)}} \varepsilon(Q_t^{(l)}) - \varepsilon(Q_t^{(i)}),$$

which yields (3.7).

Now assume $\mathcal{N}^{(i)} = \emptyset$. Then,

$$\begin{aligned}
\tilde{Q}_{i,t+1} &= \tilde{Q}_{i,t} + \tilde{X}_{i,t+1} - 1(\tilde{P}_t = i), && \text{by (3.2),} \\
&= \tilde{Q}_{i,t} + X_{i,t+1}^{(0)} - 1(\tilde{P}_t = i), && \text{by (3.1),} \\
&= Q_{i,t}^{(0)} + X_{i,t+1}^{(0)} - 1(P_t^{(0)} = i), && \text{by the induction hypothesis,} \\
&= Q_{i,t+1}^{(0)}, && \text{by (2.2).}
\end{aligned}$$

To prove (3.8), we first prove that $\tilde{Q}_{i,t+1} = 0$ implies $Q_{i,t+1}^{(0)} = 0$ and then that $Q_{i,t+1}^{(0)} = 0$ implies $\tilde{Q}_{i,t+1} = 0$. Note that if $\mathcal{N}^{(i)} = \emptyset$, (3.8) immediately follows from (3.7), so we restrict our proof to i for which $\mathcal{N}^{(i)} \neq \emptyset$. From (3.4) and (3.7) with $t+1$ substituted for t , it follows that

$$\tilde{Q}_{i,t+1} = Q_{i,t+1}^{(0)} + \sum_{m \in \mathcal{N}^{(i)}} \left(Q_{t+1-d(m,0)}^{(m)} - \varepsilon(Q_{t+1-d(m,0)}^{(m)}) \right). \quad (3.10)$$

Hence $\tilde{Q}_{i,t+1} = 0$ implies that $Q_{t+1-d(m,i)}^{(m)} = \varepsilon(Q_{t+1-d(m,i)}^{(m)})$ for all $m \in \mathcal{N}^{(i)}$ and $Q_{i,t+1}^{(0)} = 0$.

Suppose now that $Q_{i,t+1}^{(0)} = 0$. This means that $0 = Q_{i,t+1}^{(0)} = Q_{i,t}^{(0)} - 1(P_t^{(0)} = i) + \varepsilon(Q_t^{(i)})$, which entails that $Q_t^{(i)} = 0$. We can then apply this argument to $Q_t^{(i)} = Q_{t-1}^{(i)} - \varepsilon(Q_{t-1}^{(i)}) + X_t^{(i)} + \sum_{m \in \mathcal{N}_1^{(i)}} \varepsilon(Q_{t-1}^{(m)})$ to obtain $Q_{t-1}^{(m)} = 0$ for $m \in \mathcal{N}_1^{(i)}$, and so on. This eventually results in $Q_{t+1-d(m,0)}^{(m)} = 0$ for all m . Together with (3.10) we conclude $\tilde{Q}_{i,t+1} = 0$.

In Section 3.1 we defined \tilde{P}_t such that if $\{\tilde{P}_s\}_{s=t-M}^{t-1} = \{P_s^{(0)}\}_{s=t-M}^{t-1}$ and $\{\varepsilon(\tilde{Q}_s)\}_{s=t-M}^t = \{\varepsilon(Q_s^{(0)})\}_{s=t-M}^t$, then $\tilde{P}_t = P_t^{(0)}$. This implies that $\tilde{P}_{t+1} = P_{t+1}^{(0)}$, which completes the proof. \square

Corollary 3.3. *The following relation holds:*

$$\tilde{Q}_i \stackrel{d}{=} \lim_{t \rightarrow \infty} \left(\sum_{m \in \mathcal{N}^{(i)}} Q_t^{(m)} + Q_{i,t}^{(0)} - \sum_{m \in \mathcal{N}^{(i)}} \sum_{d=1}^{d(m,0)} X_{t+1-d}^{(m)} \right). \quad (3.11)$$

Remark 3.4. Proposition 3.2 is actually much stronger than Corollary 3.3; Proposition 3.2 holds for all t , for any realisation of $X_t^{(m)}$, and hence regardless of the underlying arrival process. Proposition 3.2 thus gives a deterministic relation on sample paths. We use it here primarily to obtain Corollary 3.3, but in itself it is an extension of the reduction result of Morrison [5].

3.3 The reduced and the coupled system

In this section, we show that the steady state queue lengths of the coupled and the reduced system are the same (Prop. 3.5). By combining this with Cor. 3.3, we prove Theorem 2.5 at the end of the subsection.

Proposition 3.5. *The steady state queue lengths of the coupled and the reduced system are the same:*

$$Q'_i \stackrel{d}{=} \tilde{Q}_i. \quad (3.12)$$

Proof. We introduce the discrete time process

$$Z'_t = (Q'_t, \varepsilon(Q'_{t-1}), \dots, \varepsilon(Q'_{t-M}), P'_{t-1}, \dots, P'_{t-M}),$$

and \tilde{Z}_t , defined similarly.

The processes Z'_t and \tilde{Z}_t are aperiodic and irreducible Markov chains, and therefore have a unique stationary distribution. This implies that $\tilde{Q}_{i,t}$, and $Q'_{i,t}$, being components of these Markov

chains, have unique stationary distributions too. To prove $Q'_i \stackrel{d}{=} \tilde{Q}_i$ we show that the transition probabilities of Z'_t and \tilde{Z}_t are the same for $t > \max_m d(m, 0)$. After all, if these Markov chains have the same transition probabilities, they have the same equilibrium distribution and hence so do $Q'_{i,t}$ and $\tilde{Q}_{i,t}$.

Due to the time-delay of arrivals to the coupled system, and the assumption that $X_t^{(m)} = 0$ for $t \leq 0$, it can easily be checked that the transition probabilities of \tilde{Z}_t depend on t if $t \leq \max_m d(m, 0)$. Nevertheless, this does not affect the equilibrium distribution of \tilde{Z}_t since it only affects a finite initial period of time. For $t > \max_m d(m, 0)$ the process \tilde{Z}_t is time-homogeneous.

Let $\omega_t = \{\mathbf{q}_t, \varepsilon(\mathbf{q}_{t-1}), \dots, \varepsilon(\mathbf{q}_{t-M}), p_{t-1}, \dots, p_{t-M}\}$. Assume furthermore $t > \max_m d(m, 0)$. We will show

$$\mathbb{P}(Z'_{t+1} = \omega_{t+1} | Z'_t = \omega_t) = \mathbb{P}(\tilde{Z}_{t+1} = \omega_{t+1} | \tilde{Z}_t = \omega_t).$$

Let $q_{j,t+1} = q_{j,t} - 1(p_t = j) + x_{j,t+1}$ for all j and t . Then, since the transition probabilities are only determined by the arrival probabilities and the HoL-based service discipline,

$$\begin{aligned} \mathbb{P}(Z'_{t+1} = \omega_{t+1} | Z'_t = \omega_t) \\ = \mathbb{P}\left(P'_t = p_t \mid \varepsilon(\mathbf{Q}'_s) = \varepsilon(\mathbf{q}_s), P'_s = p_s, \varepsilon(\mathbf{Q}'_t) = \varepsilon(\mathbf{q}_t)\right) \prod_{j=1}^N \mathbb{P}(X'_{j,t+1} = x_{j,t+1}), \end{aligned}$$

where $s = t-1, \dots, t-M$. Using $X'_{j,t+1} \stackrel{d}{=} \tilde{X}_{j,t+1}$ (see Section 3.1) and the fact that the service discipline is equal to that of node 0 (see (2.4)), we obtain

$$\begin{aligned} \mathbb{P}(Z'_{t+1} = \omega_{t+1} | Z'_t = \omega_t) \\ = \mathbb{P}\left(P_t^{(0)} = p_t \mid \varepsilon(\mathbf{Q}_s^{(0)}) = \varepsilon(\mathbf{q}_s), P_s^{(0)} = p_s, \varepsilon(\mathbf{Q}_t^{(0)}) = \varepsilon(\mathbf{q}_t)\right) \prod_{j=1}^N \mathbb{P}(\tilde{X}_{j,t+1} = x_{j,t+1}). \end{aligned}$$

Finally, by $\varepsilon(\tilde{\mathbf{Q}}_t) = \varepsilon(\mathbf{Q}_t^{(0)})$ and $\tilde{P}_t = P_t^{(0)}$ for all t , we have

$$\begin{aligned} \mathbb{P}(Z'_{t+1} = \omega_{t+1} | Z'_t = \omega_t) \\ = \mathbb{P}\left(\tilde{P}_t = p_t \mid \varepsilon(\tilde{\mathbf{Q}}_s) = \varepsilon(\mathbf{q}_s), \tilde{P}_s = p_s, \varepsilon(\tilde{\mathbf{Q}}_t) = \varepsilon(\mathbf{q}_t)\right) \prod_{j=1}^N \mathbb{P}(\tilde{X}_{j,t+1} = x_{j,t+1}) \\ = \mathbb{P}(\tilde{Z}_{t+1} = \omega_{t+1} | \tilde{Z}_t = \omega_t). \quad \square \end{aligned}$$

Remark 3.6. In Section 2.1 we stated that $M < \infty$. This is necessary for the processes Z'_t and \tilde{Z}_t to have unique equilibrium distributions. If $M = \infty$ the queue that is served at time t might depend on the queue that was served at time 1 for all t . The transition of \tilde{Z}_t might thus depend on \tilde{Z}_1 and the equilibrium distribution does not have to be unique.

The proof of Theorem 2.5 is now elementary:

Proof of Theorem 2.5. Theorem 2.5 follows immediately from Corollary 3.3 and Proposition 3.5. \square

3.4 Waiting times

In this section we establish equality of the mean waiting times in the original and reduced system. The mean end-to-end sojourn time of type i packets passing through queue i of node 0 is given by

$$\mathbb{E}[S_i] = \sum_{m \in \mathcal{N}^{(i)}} \frac{\lambda^{(m)}}{\gamma^{(i)}} \sum_{l \in P(m,i)} \mathbb{E}[S^{(l)}] + \mathbb{E}[S_i^{(0)}].$$

The mean distance a type i packet must traverse before it reaches queue i of node 0 is given by

$$\mathbb{E}[D_i] = \sum_{m \in \mathcal{N}^{(i)}} \frac{\lambda^{(m)}}{\gamma^{(i)}} d(m, 0).$$

Note that $\mathbb{E}[D_i] + 1$ gives the mean total service time of type i packets.

We can now prove Theorem 2.6:

Proof of Theorem 2.6. First assume $\mathcal{N}^{(i)} \neq \emptyset$. In this case the total arrival rate to queue i of node 0 is given by $\gamma^{(i)}$. Using (2.5) and Little's law yields

$$\begin{aligned} \mathbb{E}[S'_i] &= \frac{1}{\gamma^{(i)}} \mathbb{E}[Q'_i] = \sum_{m \in \mathcal{N}^{(i)}} \frac{1}{\gamma^{(i)}} \mathbb{E}[Q^{(m)}] + \frac{1}{\gamma^{(i)}} \mathbb{E}[Q_i^{(0)}] - \frac{1}{\gamma^{(i)}} \sum_{m \in \mathcal{N}^{(i)}} \sum_{d=1}^{d(m,0)} \mathbb{E}[X_d^{(m)}] \\ &= \frac{1}{\gamma^{(i)}} \sum_{m \in \mathcal{N}^{(i)}} \gamma^{(m)} \mathbb{E}[S^{(m)}] + \mathbb{E}[S_i^{(0)}] - \frac{1}{\gamma^{(i)}} \sum_{m \in \mathcal{N}^{(i)}} d(m, 0) \lambda^{(m)} \\ &= \frac{1}{\gamma^{(i)}} \sum_{m \in \mathcal{N}^{(i)}} \sum_{l \in \mathcal{N}^{(m)}} \lambda^{(l)} \mathbb{E}[S^{(m)}] + \mathbb{E}[S_i^{(0)}] - \mathbb{E}[D_i], \end{aligned}$$

where the last equation is a consequence of the fact that any arrival to node m must be an external arrival to some node $l \in \mathcal{N}^{(m)}$ (see (2.1)).

We continue by interchanging the order of summation. First, observe that for every $m \in \mathcal{N}^{(i)}$ we include all nodes l upstream of m , so that every $l \in \mathcal{N}^{(i)}$ is eventually included in the summation at least once. Second, for every node l , we include all nodes m downstream of l and upstream of i , i.e., all m for which $m \in \mathcal{N}^{(i)}$ and $l \in \mathcal{N}^{(m)}$. This, however, corresponds precisely to the definition of the path from l to i (see (2.3)), so that

$$\mathbb{E}[S'_i] = \sum_{l \in \mathcal{N}^{(i)}} \frac{\lambda^{(l)}}{\gamma^{(i)}} \sum_{m \in P(l, i)} \mathbb{E}[S^{(m)}] + \mathbb{E}[S_i^{(0)}] - \mathbb{E}[D_i] = \mathbb{E}[S_i] - \mathbb{E}[D_i].$$

The equality of $\mathbb{E}[W'_i]$ and $\mathbb{E}[W_i]$ follows from the observation that $\mathbb{E}[D_i] + 1$ is the mean total service time.

For $\mathcal{N}^{(i)} = \emptyset$ the proof is similar, except that now the arrival rate to queue i of node 0 is given by $\lambda_i^{(0)}$. Taking expectations of (2.5), dividing by $\lambda_i^{(0)}$, and applying Little's law immediately yields $\mathbb{E}[S'_i] = \mathbb{E}[S_i]$, which implies $\mathbb{E}[W'_i] = \mathbb{E}[W_i]$. \square

4 Discussion

In this section, we illustrate why Assumption 2.1 is needed by means of an example. We consider a tree network consisting of two nodes (node 0 and 1), as depicted in Figure 2. Node 0 consists of two queues and node 1 of one. We assume that node 0 serves packets according to the shortest queue first policy (so the service discipline is *not* HoL-based). In case of equal queue lengths, it serves queue 1. The service policy of node 1 is arbitrary. We will show that the mean end-to-end delays of type 1 packets are different in the original and the reduced system.

We assume that batches of size $K > 1$ arrive to node 1, whereas batches of size 1 arrive to queue 2 of node 0:

$$\begin{aligned} X^{(1)} &= \begin{cases} 0 & \text{w.p. } 1 - p_1, \\ K & \text{w.p. } p_1, \end{cases} \\ X_2^{(0)} &= \begin{cases} 0 & \text{w.p. } 1 - p_2, \\ 1 & \text{w.p. } p_2. \end{cases} \end{aligned}$$

In the *original* system type 1 batches arrive to node 1 and are sent to node 0 packet-by-packet. Because node 0 serves queue 1 in case of equal queue lengths, type 1 packets never have to wait




Figure 2: The network we consider in order to show that Assumption 2.1 is needed. Node 0 uses shortest queue first and the service discipline of node 1 is arbitrary.

at node 0. The mean end-to-end delay of a type 1 packet is thus given by the mean waiting time in node 1, which is an ordinary $geo^{X^{(1)}}/D/1$ queue.

In the *reduced* system type 1 batches arrive to the queues in their entirety. The length of queue 1 is therefore typically larger than that of queue 2 and queue 2 is usually served first (except if both queues have one packet). The mean waiting time of type 1 packets is thus equal to that in a $geo^{X^{(1)}}/D/1$ queue plus some additional delay due to services of type 2 packets.

This example illustrates why a HoL-based service discipline is required; the number of packets in queue 1 of the reduced system is typically different from that of queue 1 of node 0, because batches arrive to the reduced system in their entirety, whereas they arrive to node 0 packet-by-packet. Even though the service discipline is shortest queue first in both systems, the server of the original system essentially gives priority to type 1 and the server of the reduced system to type 2.

5 Conclusion

We have proved a distributional relation between queue lengths in a tree network and a reduced system consisting of one node, and equality of the mean end-to-end delay in both systems. Our results are extensions of earlier work, in particular Morrison [5], in the sense that they hold conditioned on the type of the packet. By this we mean that where we prove a relation between the number of type i packets in both systems, earlier results relate the *total* number of packets in both systems, i.e., summed over all i . This extension comes at the price that the class of allowed service disciplines is slightly more restrictive; the service discipline of node 0 has to be HoL-based.

A HoL-based service discipline is required because the number of packets in queue i of the reduced system is typically different from that in queue i of node 0. If the server is allowed to choose a queue based on the number of packets in a queue rather than whether or not it is empty, the behaviour of the systems might be entirely different, which could in turn lead to different mean waiting times.

The reduction result of this paper facilitates the analysis of a polling network through a simpler analysis of a single node. Even if the per-queue mean waiting time is unknown (such as for m_i -limited systems) the reduction result entails that single-station approximations can also be applied to networks without additional loss of accuracy.

6 Acknowledgements

The authors would like to thank Onno Boxma and Lasse Leskelä for their helpful suggestions.

References

- [1] B. Avi-Itzhak, *A sequence of service stations with arbitrary input and regular service times*, Management Science **11** (1965), no. 5, pp. 565–571.
- [2] W.J. Dally and B. Towles, *Route packets, not wires: on-chip interconnection networks*, Proc. of the 38th Design Automation Conference, 2001, pp. 684–689.
- [3] H.D. Friedman, *Reduction methods for tandem queueing systems*, Operations Research **13** (1965), pp. 121–131.
- [4] M. Mandjes and M. van Uitert, *Sample-path large deviations for tandem and priority queues with gaussian inputs*, The Annals of Applied Probability **15** (2005), no. 2, pp. 1193–1226.
- [5] J.A. Morrison, *A combinatorial lemma and its application to concentrating trees of discrete-time queues*, The Bell Systems Technical Journal **57** (1978), no. 5, pp. 1645–1652.
- [6] M.J. Neely, C.E. Rohrs, and E. Modiano, *Equivalent models for queueing analysis of deterministic service time tree networks*, IEEE Transactions on Information Theory **51** (2005), no. 10, pp. 1–10.
- [7] M.I. Reiman and L.M. Wein, *Heavy traffic analysis of polling systems in tandem*, Operations Research **47** (1999), no. 4, pp. 524–534.
- [8] I. Rubin, *Communication networks: Message path delays*, IEEE Transactions on Information Theory **20** (1974), no. 6, pp. 738–745.
- [9] I. Rubin, *Message path delays in packet-switching communication networks*, IEEE Transactions on Communications **23** (1975), no. 2, pp. 186–192.
- [10] I. Rubin, *An approximate time-delay analysis for packet-switching communication networks*, IEEE Transactions on Communications **24** (1976), no. 2, pp. 210–222.
- [11] M.S. Shalmon, *Exact delay analysis of packet-switching concentrating networks*, IEEE Transactions on Communications **35** (1987), no. 12, pp. 1265–1271.
- [12] M.S. Shalmon and M.A. Kaplan, *A tandem network of queues with deterministic service and intermediate arrivals*, Operations Research **23** (1984), no. 4, pp. 753–773.
- [13] V.M. Vishnevskii and O.V. Semenova, *Mathematical methods to study the polling systems*, Automation and Remote Control **67** (2006), no. 2, pp. 173–220.