Companion

Performance analysis of parallel networks with collaborating multi-tasking resources

Itai Gurvich and Jan A. Van Mieghem

October 15, 2016 – *Management Science*, Forthcoming

Some preliminaries. Let $B(t)$ be the configuration used at time t, i.e., $B(t)$ is a column of the matrix C; see §4. Then,

$$T_c(t) = \int_0^t \mathbb{1}\{B(s) = c\} ds.$$

Define the state descriptor

$$\mathbb{X}(t) = (Q(t), B(t))$$

and denote by $\mathcal{X}$ its state space. We use $x = (q, b)$ for a point in this state space. A policy is said to be stationary if the process $\mathbb{X}(t)$ is a Continuous Time Markov Chain. This rules out for example, policies that use residual service or inter-arrival times (e.g. shortest processing time first) but is otherwise fairly general.

We use standard Markov chain notation throughout: for a function $f : \mathcal{X} \to \mathbb{R}$, $\mathbb{E}_x[f(\mathbb{X}(t))]$ is the expectation of $f(\mathbb{X}(t))$ conditional on the state at time $t = 0$ being $x \in \mathcal{X}$. We write $\mathbb{E}_\pi[f(\mathbb{X}(t))]$ to denote this expectation when the chain is initialized with a distribution π. Also, we follow the convention for positive sequences that $a_n = \mathcal{O}(b_n)$ means $\limsup_{n\to\infty} a_n/b_n < \infty$, $a_n = o(b_n)$ means $\lim_{n\to\infty} a_n/b_n = 0$, and $a_n = \Theta(b_n)$ means that $a_n = \mathcal{O}(b_n)$ but $a_n \neq o(b_n)$.

In our proofs we make frequent use of the relationship between a continuous time Markov chain and a discrete time chain sampled at hitting times of a subset $\mathcal{B} \subset \mathcal{X}$. The following known fact will be used: Let τ_l be the l^{th} hitting time of the set $\mathcal{B}$ and suppose that $Q^l = \mathbb{X}(\tau_l)$ is a discrete time chain with a steady-state distribution π such that $\sum_{x\in\mathcal{B}} \pi_x \mathbb{E}_x[\tau_1] < \infty$. The CTMC $\mathbb{X}(t)$ is then also positive recurrent with a steady-state distribution ν, and π and ν are related through

$$\mathbb{E}_\nu[f(Q(0))] = \frac{\sum_{x\in\mathcal{B}} \pi_x \mathbb{E}_x\left[\int_0^{\tau_1} f(Q(s))ds\right]}{\sum_{x\in\mathcal{B}} \pi_x \mathbb{E}_x\left[\tau_1\right]}; \tag{23}$$

see e.g. Asmussen (2003)[Proposition VII.5.2].

A.1. Proofs for Section 3

Proof of Proposition 1. Under preemptive priority to the individual tasks, the process $(Q_0(t), Q_1(t), Q_2(t))$ is a continuous time Markov chain. Let τ_l be the time of the l^{th} return of the resources to the collaborative task. At these moments the individual queues are empty so that $Q_0(\tau_l), l = 0, \ldots,$ is a discrete time Markov chain.

Because of preemption, queues 1 and 2 are two independent $M/M/1$ queues. In particular, $\mathbb{E}_{(q_0,0,0)}[\tau_1]$ is (independently of q_0) the expected time it takes two independent $M/M/1$ queues, starting at 0, until the first return to $(0,0)$: for any q_0

$$\mathbb{E}_{(q_0,0,0)}[\tau_1] = \frac{1}{(\lambda_1+\lambda_2)\pi^P(0,0)} = \frac{1}{(\lambda_1+\lambda_2)\pi^P(0,0)} = \frac{1}{(\lambda_1+\lambda_2)(1-\rho_1^a)(1-\rho_2^a)} < \infty,$$

where π^P is the steady state of two independent $M/M/1$ queues so that $\pi^P(0,0) = (1-\rho_1^a)(1-\rho_2^a)$. (recall that $\rho_i^a = \lambda_i/\mu_i$.)

We will show that the DTMC $Q_0^l = (Q_0(\tau_l), l = 1,2,\ldots)$ is positive recurrent if

$$(\rho_0^a - (1-\rho_1^a)(1-\rho_2^a)) < 0. \tag{24}$$

Let π be its stationary distribution. Since $\sup_{q_0} \mathbb{E}_{(q_0,0,0)}[\tau_1] < \infty$, $\mathbb{E}_\pi[\tau_1] < \infty$ so that by (23) the CTMC is also stable.

To prove that Q_0^l is stable let $D_0(t)$ be the number of service completions in the collaborative queue 0 by time t. Then, the server processes until the first arrival $\tilde{\tau}_1 < \tau_1$ to an individual queue and then stops. Thus, the expected number of completions by τ_1 is $\mathbb{E}_{(q_0,0,0)}[D_0(\tau_1)] = \mathbb{E}_{(q_0,0,0)}\left[S_0\left(\mu_0\int_0^{\tilde{\tau}_1}\mathbb{1}\{Q_0(s)>0\}ds\right)\right] = \mathbb{E}_{(q_0,0,0)}\left[\mu_0\int_0^{\tilde{\tau}_1}\mathbb{1}\{Q_0(s)>0\}ds\right]$. $\mathbb{E}[\tilde{\tau}_1] = (\lambda_1+\lambda_2)^{-1} < \infty$ (independently of q_0) and for any finite t, $\int_0^t \mathbb{1}\{Q_0(s)>0\}ds \to t$ as $q_0 \to \infty$ almost surely. It follows that $\mathbb{E}_{(q_0,0,0)}[\int_0^{\tilde{\tau}_1}\mathbb{1}\{Q_0(s)>0\}ds] - \mathbb{E}_{(q_0,0,0)}[\tilde{\tau}_1] = 0$ as $q_0\to\infty$ and, in particular, that $\limsup_{q_0\to\infty}\mathbb{E}_{(q_0,0,0)}[D_0(\tau_1)] = \frac{\mu_0}{\lambda_1+\lambda_2}$. Since $\mathbb{E}_{(q_0,0,0)}[A_0(\tau_1)] = \lambda_0\mathbb{E}_{(q_0,0,0)}[T_1]$ we have

$$\begin{aligned}\limsup_{q_0\to\infty}(\mathbb{E}_{q_0}[Q_0^1]-q_0) &= \limsup_{q_0\to\infty}\left(\mathbb{E}_{(q_0,0,0)}[A_0(\tau_1)] - \mathbb{E}_{(q_0,0,0)}[D_0(\tau_1)]\right)\\ &\le \frac{\lambda_0}{(\lambda_1+\lambda_2)(1-\rho_1^a)(1-\rho_2^a)} - \frac{\mu_0}{\lambda_1+\lambda_2}\\ &\le \frac{\mu_0}{(\lambda_1+\lambda_2)(1-\rho_1^a)(1-\rho_2^a)}\left(\rho_0^a - (1-\rho_1^a)(1-\rho_2^a)\right).\end{aligned}$$

The right hand side is negative if (24) holds which concludes the sufficiency argument; see e.g. (Robert, 2003, Theorem 8.6).

Necessity: Since task 0 idles whenever the individual queues have work, it follows from the strong law of large numbers that, almost surely,

$$\begin{aligned}\liminf_{t\to\infty}\frac{1}{\mu_0}\frac{1}{t}Q_0(t) &\ge (\lambda_0/\mu_0 - 1) + \lim_{t\to\infty}\frac{1}{t}\int_0^t \mathbb{1}\{Q_1(s)+Q_2(s)>0\}ds\\ &\ge (\rho_0^a - 1) + (1-(1-\rho_1^a)(1-\rho_2^a))\\ &\ge (\rho_0^a - (1-\rho_1^a)(1-\rho_2^a)) > \epsilon > 0.\end{aligned}$$

Thus, if the necessary condition holds $Q_0(t)/t$ diverges to infinity almost surely which implies the transience of $\mathbb{X}(t)$.

We turn to the non-preemptive policy. Under non-preemptive priority to the individual tasks the process $\mathbb{X}(t) = (Q(t), B(t))$ is a continuous time Markov chain. We study the collaborative queue at moments τ_l of return to that queue. Let $\bar{\tau}_0$ be the first time after $t = 0$ that resources move to the individual tasks. Then, for $x = ((q_0, 0, 0), (1, 0, 0))$ (to simplify notation, we write below $\mathbb{E}_{q_0}[\cdot]$ to mean $\mathbb{E}_{(q_0,0,0),(1,0,0)}[\cdot]$), $\mathbb{E}_{q_0}[\bar{\tau}_0] \leq \frac{1}{\lambda_1+\lambda_2} + m_0$ (the time until an arrival to one of the individual queues plus one collaborative service time). In expectation there are at most $\mu_0 \mathbb{E}_{q_0}[\bar{\tau}_0]$ service completions on $[0, \bar{\tau}_0)$ in the collaborative queue so that

$$\mathbb{E}_{q_0}[Q_0(\tau_1)] - q_0 \geq \lambda_0 \mathbb{E}_{q_0}[\tau_1] - \mu_0 \left(\frac{1}{\lambda_1 + \lambda_2} + m_0 \right).$$

Thus, the chain $Q_0^l = Q_0(\tau_l)$ is unstable if $\liminf_{q_0 \to \infty} (\mathbb{E}_{q_0}[Q_0(\tau_1)] - q_0) > 0$ which holds, in particular, if

$$\rho_0^a > \limsup_{q_0 \to \infty} \frac{\frac{1}{\lambda_1+\lambda_2} + m_0}{\mathbb{E}_{q_0}[\tau_1]}. \tag{25}$$

We will show that

$$\liminf_{q_0 \to \infty} \mathbb{E}_{q_0}[\tau_1] \geq \underline{T} := \frac{1}{\lambda_1 + \lambda_2} + m_0 + \frac{1 - (1 - \rho_1^a)(1 - \rho_2^a)}{(\lambda_1 + \lambda_2)(1 - \rho_1^a)(1 - \rho_2^a)},$$

so that the chain is unstable if

$$\rho_0^a > \frac{\frac{1}{\lambda_1+\lambda_2} + m_0}{\underline{T}},$$

which implies (25).

To lower bound $\mathbb{E}_{q_0}[\tau_1]$ denote by $\tilde{\tau}(q_1, q_2)$ the time it takes for two independent $M/M/1$ queues to reach $(0, 0)$ starting at (q_1, q_2). Then, $\tau_1 = \bar{\tau}_0 + \tilde{\tau}(Q_1(\bar{\tau}_0), Q_2(\bar{\tau}_0))$. The stopping time $\tilde{\tau}$ is monotone (in the sense of standard stochastic ordering) in the initial states of the queues, that is $\tilde{\tau}(q_1, q_2) \geq_{st} \tilde{\tau}(\tilde{q}_1, \tilde{q}_2)$ when $q_i \geq \tilde{q}_i,\ i = 1, 2$. Conditioning on whether the move is triggered by arrival to queue 1 or to queue 2, we have

$$\mathbb{E}_{q_0}[\tilde{\tau}(Q_1(\bar{\tau}_0), Q_2(\bar{\tau}_0))] \geq \frac{\lambda_1}{\lambda_1 + \lambda_2} \mathbb{E}_{q_0}[\tilde{\tau}(1, 0)] + \frac{\lambda_2}{\lambda_1 + \lambda_2} \mathbb{E}_{q_0}[\tilde{\tau}(0, 1)] =: Z.$$

The sum $Z + \frac{1}{\lambda_1+\lambda_2}$ is the time it takes, starting at 0, for two independent $M/M/1$ queues to return to 0 after having left it. From the standard relationship between return times and the stationary distributions of a Markov chain (e.g. (Asmussen, 2003, Chapter II.4)) applied to the two $M/M/1$ queues, it holds that

$$Z + \frac{1}{\lambda_1 + \lambda_2} = \frac{1}{(\lambda_1 + \lambda_2)\pi^P(0, 0)},$$

where $\pi^P(0,0) = (1-\rho_1^a)(1-\rho_2^a)$ is the steady-state probability of the two $M/M/1$ queues being empty. Further, since $\mathbb{E}_{q_0}[\bar{\tau}_0] = \frac{1}{\lambda_1+\lambda_2} + m_0 \mathbb{P}_{q_0}\{Q_0(\tau_1) > 0\}$ independently of q_0 we have, as $q_0 \to \infty$, that $\mathbb{E}_{q_0}[\bar{\tau}_0] \to \frac{1}{\lambda_1+\lambda_2} + m_0$. We conclude that

$$\liminf_{q_0\to\infty} \mathbb{E}_{q_0}[\tau_1] \geq \liminf_{q_0\to\infty} \mathbb{E}_{q_0}[\bar{\tau}_0] + Z = \frac{1}{\lambda_1+\lambda_2} + m_0 + \frac{1-\pi^P(0,0)}{(\lambda_1+\lambda_2)\pi^P(0,0)}.$$

This establishes (25) which implies the transience of the chain as well as the existence of a constant $\gamma > 0$ such that, almost surely,

$$\liminf_{l\to\infty} \frac{Q_0(\tau^l)}{l} \geq \gamma; \tag{26}$$

see e.g. (Robert, 2003, Theorem 8.11).

We finally translate the transience of the DTMC to that of the CTMC. First, each time the resources return to the individual tasks, the expected queue in task i is in expectation smaller than $\lambda_i\left(\frac{1}{\lambda_1+\lambda_2} + m_0\right)$, and it follows that $\limsup_{q_0\to\infty} \mathbb{E}_{q_0}[\tau_1] < \infty$. Let $N(t)$ the number of returns to the collaborative task by time t. It then follows from (26) that $\liminf_{t\to\infty} Q_0(\tau_{N(t)+1})/(N(t)+1) \geq \gamma$ almost surely. Further, by Wald's lemma $0 \leq \mathbb{E}_{q_0}[\tau_{N(t)+1} - t] \leq \sup_{q_0} \mathbb{E}_{q_0}[\tau_1] < \infty$ and, in particular, $\tau_{N(t)+1}/t \to 1$ in probability and $(A_0(\tau_{N(t)+1}) - A_0(t))/t \to 0$ in probability. Finally, since $Q(\tau_{N(t)+1}) \leq Q_0(t) + A_0(\tau_{N(t)+1}) - A_0(t)$, then

$$\frac{Q_0(t)}{t} \geq \frac{Q_0(\tau_{N(t)+1})}{\tau_{N(t)+1}} \frac{\tau_{N(t)+1}}{t} - \frac{A(\tau_{N(t)+1}) - A(t)}{t}.$$

Using (26) we conclude that $Q_0(t)/t \to \infty$ in probability as required. ■

Proof of Theorem 1. Notice that this proof does not rely on the indivisibility of resources.

We first prove that, regardless of whether the policy is preemptive or not, at least one of the individual queues must be non-negligible. We will then prove that there exists a preemptive policy under which one of the queues does not grow with ρ^{BN}. Finally, for the fact that, with non-preemption, both individual queues must be non-negligible; see the comment and the end of Theorem's 3.3 proof.

Let $T_i^a(t) = (CT(t))_i$ be the time allocated to task i by time t and $T_I(t) = t - T_0^a(t)$ be the time remaining after allocating $T_0^a(t)$ to the collaborative task. In particular, $T_I(t) \geq T_i^a(t)$ for all $t \geq 0$ and $i = 1,2$. Then,

$$\begin{aligned} Q_i(t) &= Q_i(0) + A_i(t) - S_i(T_i^a(t)) \\ &= Q_i(0) + \lambda_i t - \mu_i T_I(t) + \mu_i(T_I(t) - T_i^a(t)) + M_i(t), \end{aligned}$$

where $M_i(t) := A_i(t) - \lambda_i t - (S_i(T_i^a(t)) - \mu_i T_i^a(t))$. We add the superscript $\rho^{\rm BN}$ to make explicit the dependence on the bottleneck load. Towards contradiction suppose that, for each $\rho^{\rm BN}$, the network has a steady-state distribution $\pi^{\rho^{\rm BN}}$ and that

$$\limsup_{\rho^{\rm BN}\uparrow 1}(1-\rho^{\rm BN})\mathbb{E}_{\pi^{\rho^{\rm BN}}}[Q_i^{\rho^{\rm BN}}(0)] = 0 \text{ for } i=1,2. \tag{27}$$

That the same applies to the waiting time $W_i^{\rho^{\rm BN}}$ follows from $\mathbb{E}_{\pi^{\rho^{\rm BN}}}[Q_i^{\rho^{\rm BN}}(0)] = \lambda_i^{\rho^{\rm BN}}\mathbb{E}_{\pi^{\rho^{\rm BN}}}[W_i^{\rho^{\rm BN}}(0)]$ and the assumption that $\lambda_i^{\rho^{\rm BN}}$ is bounded below by a constant $\eta > 0$.

Since $m_i = \mu_i^{-1} > 0$ is fixed, the sequence $\lambda_i^{\rho^{\rm BN}}$ is bounded. We assume without loss of generality that the whole sequence converges since the argument below apply to all convergent subsequences.

Fix $\epsilon > 0$ and let $t^{\rho^{\rm BN}} = \epsilon(1-\rho^{\rm BN})^{-2}$. Initializing the network at $t=0$ with this distribution, it follows that

$$(1-\rho^{\rm BN})^{-2} T_i^{a,\rho^{\rm BN}}(t^{\rho^{\rm BN}}) \Rightarrow (\lambda_i/\mu_i)\epsilon, \ i=1,2, \tag{28}$$

where $\lambda_i = \lim_{\rho^{\rm BN}\uparrow 1}\lambda_i^{\rho^{\rm BN}}$. A standard random time change argument leads to

$$(1-\rho^{\rm BN})M_i^{\rho^{\rm BN}}(t^{\rho^{\rm BN}}) \Rightarrow \hat{M}_i, \tag{29}$$

where $\hat{M}_i$ is a zero mean normally distributed random variable with variance $2\epsilon\lambda_i$. Further, from the properties of the Poisson process it follows that $(1-\rho^{\rm BN})M_i^{\rho^{\rm BN}}(t^{\rho^{\rm BN}})$ is a uniformly integrable sequence and, consequently, that

$$(1-\rho^{\rm BN})\mathbb{E}[|M_i^{\rho^{\rm BN}}(t^{\rho^{\rm BN}})|] \to \mathbb{E}[|\hat{M}_i|]. \tag{30}$$

Since $\rho_1 = \rho_2 = \rho^{\rm BN}$ we have $\lambda_1^{\rho^{\rm BN}}/\mu_1 = \lambda_2^{\rho^{\rm BN}}/\mu_2$. (All the below holds under the relaxed requirement that $\rho_1 - \rho_2 = o(1-\rho^{\rm BN})$). Notice that $I_i^{\rho^{\rm BN}}(t) := T_I^{\rho^{\rm BN}}(t) - T_i^{a,\rho^{\rm BN}}(t) = t - T_0^{\rho^{\rm BN}}(t) - T_i^{\rho^{\rm BN}}(t)$ is the accumulated idleness of resource i by time t. Then, for all $t \geq 0$.

$$\frac{Q_1^{\rho^{\rm BN}}(t)}{\mu_1} - \frac{Q_2^{\rho^{\rm BN}}(t)}{\mu_2} = \frac{Q_1^{\rho^{\rm BN}}(0)}{\mu_1} - \frac{Q_2^{\rho^{\rm BN}}(0)}{\mu_2} + (I_1^{\rho^{\rm BN}}(t) - I_2^{\rho^{\rm BN}}(t)) + M_1^{\rho^{\rm BN}}(t) - M_2^{\rho^{\rm BN}}(t),$$

and, in turn,

$$\begin{aligned} \left|I_1^{\rho^{\rm BN}}(t) - I_2^{\rho^{\rm BN}}(t))\right| = & \\ & \left|\frac{Q_1^{\rho^{\rm BN}}(t)}{\mu_1} - \frac{Q_2^{\rho^{\rm BN}}(t)}{\mu_2} - \left(\frac{Q_1^{\rho^{\rm BN}}(0)}{\mu_1} - \frac{Q_2^{\rho^{\rm BN}}(0)}{\mu_2}\right) - (M_1^{\rho^{\rm BN}}(t) - M_2^{\rho^{\rm BN}}(t))\right|. \end{aligned} \tag{31}$$

Using (27) and the triangular inequality we have

$$\begin{aligned} (1-\rho^{\rm BN})\mathbb{E}_{\pi^{\rho^{\rm BN}}}\left[\left|\frac{Q_1^{\rho^{\rm BN}}(t^{\rho^{\rm BN}})}{\mu_1} - \frac{Q_2^{\rho^{\rm BN}}(t^{\rho^{\rm BN}})}{\mu_2} - \left(\frac{Q_1^{\rho^{\rm BN}}(0)}{\mu_1} - \frac{Q_2^{\rho^{\rm BN}}(0)}{\mu_2}\right) - (M_1^{\rho^{\rm BN}}(t^{\rho^{\rm BN}}) - M_2^{\rho^{\rm BN}}(t^{\rho^{\rm BN}})\right|\right] & \\ -(1-\rho^{\rm BN})\mathbb{E}_{\pi^{\rho^{\rm BN}}}\left[\left|(M_1^{\rho^{\rm BN}}(t^{\rho^{\rm BN}}) - M_2^{\rho^{\rm BN}}(t^{\rho^{\rm BN}}))\right|\right] \to 0, & \end{aligned}$$

as $\rho^{\mathrm{BN}} \uparrow 1$. Plugging this into (31) we get

$$\lim_{\rho^{\mathrm{BN}}\uparrow 1}(1-\rho^{\mathrm{BN}})\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[|I_1^{\rho^{\mathrm{BN}}}(t^{\rho^{\mathrm{BN}}}) - I_2^{\rho^{\mathrm{BN}}}(t^{\rho^{\mathrm{BN}}})|] = \lim_{\rho^{\mathrm{BN}}\uparrow 1}(1-\rho^{\mathrm{BN}})\mathbb{E}[|M_1^{\rho^{\mathrm{BN}}}(t^{\rho^{\mathrm{BN}}}) - M_2^{\rho^{\mathrm{BN}}}(t^{\rho^{\mathrm{BN}}})|].$$

By (29), $(1-\rho^{\mathrm{BN}})(M_1^{\rho^{\mathrm{BN}}}(t^{\rho^{\mathrm{BN}}}) - M_2^{\rho^{\mathrm{BN}}}(t^{\rho^{\mathrm{BN}}})) \Rightarrow \hat{M}_{1,\epsilon} - \hat{M}_{2,\epsilon}$ which is normally distributed with variance $2\epsilon(\lambda_1+\lambda_2)$ so that, for all ρ^{BN} sufficiently close to 1,

$$(1-\rho^{\mathrm{BN}})\left(\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[I_1^{\rho^{\mathrm{BN}}}(t^{\rho^{\mathrm{BN}}})] + \mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[I_2^{\rho^{\mathrm{BN}}}(t^{\rho^{\mathrm{BN}}})]\right) \geq c\sqrt{\epsilon} \tag{32}$$

for some constant $c > 0$. Denote the total idleness (summing over resources) by $I_+^{\rho^{\mathrm{BN}}}(t^{\rho^{\mathrm{BN}}})$. Then,

$$\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[I_+^{\rho^{\mathrm{BN}}}(t^{\rho^{\mathrm{BN}}})] = \mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[I_1^{\rho^{\mathrm{BN}}}(t^{\rho^{\mathrm{BN}}})] + \mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[I_2^{\rho^{\mathrm{BN}}}(t^{\rho^{\mathrm{BN}}})].$$

The inequality (32) implies that for ρ^{BN} sufficiently close to 1

$$\frac{\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[I_+^{\rho^{\mathrm{BN}}}(t^{\rho^{\mathrm{BN}}})]}{t^{\rho^{\mathrm{BN}}}} = \frac{\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[I_+^{\rho^{\mathrm{BN}}}(\epsilon(1-\rho^{\mathrm{BN}})^{-2}])}{\epsilon(1-\rho^{\mathrm{BN}})^{-2}} \geq \frac{c\sqrt{\epsilon}(1-\rho^{\mathrm{BN}})^{-1}}{\epsilon(1-\rho^{\mathrm{BN}})^{-2}} = \frac{c(1-\rho^{\mathrm{BN}})}{\sqrt{\epsilon}}.$$

Since ϵ was arbitrary we have, in particular, that

$$\frac{\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[I_+^{\rho^{\mathrm{BN}}}(t^{\rho^{\mathrm{BN}}})]}{t^{\rho^{\mathrm{BN}}}} \geq 4(1-\rho^{\mathrm{BN}}),$$

for all ρ^{BN} sufficiently close to 1. Stationarity, however, requires that

$$\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[I_+^{\rho^{\mathrm{BN}}}(t^{\rho^{\mathrm{BN}}})] \leq 2(1-\rho^{\mathrm{BN}})t^{\rho^{\mathrm{BN}}}. \tag{33}$$

(Recall that $\rho_1 = \rho_2 = \rho^{\mathrm{BN}}$) which is a contradiction. We conclude that there is no stationary stabilizing control that has $(1-\rho^{\mathrm{BN}})\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[Q_i^{\rho^{\mathrm{BN}}}(0)] \to 0$ for both individual queues.

We next construct a preemptive policy that keeps one individual queue from growing with ρ^{BN} while retaining network stability. We will prove that given $\epsilon > 0$, the proposed policy stabilizes $Q(t) = (Q_0(t), Q_1(t), Q_2(t))$ and that under the corresponding (sequence of) steady state distributions $\nu^{\rho^{\mathrm{BN}}}$ it holds that

$$\limsup_{\rho^{\mathrm{BN}}\uparrow 1}(1-\rho^{\mathrm{BN}})\mathbb{E}_{\nu^{\rho^{\mathrm{BN}}}}[W_1^{\rho^{\mathrm{BN}}}(0)] = 0, \text{ and } \limsup_{\rho^{\mathrm{BN}}\uparrow 1}(1-\rho^{\mathrm{BN}})^{1+\epsilon}\mathbb{E}_{\nu^{\rho^{\mathrm{BN}}}}[\mathcal{W}_+^{\rho^{\mathrm{BN}}}(0)] < \infty,$$

where $\mathcal{W}_+(t) = 2\rho_0 Q_0(t) + \rho_1 Q_1(t) + \rho_2 Q_2(t)$ is the total workload. Or similarly, with 1 replaced with 2.

Our policy is as follows: if, upon a service completion of resource 1, $\mathcal{W}_+(t) \leq (1-\rho^{\mathrm{BN}})^{-2}$, resource 1 picks a job from queue 1 (its individual queue) if it is non-empty. If queue 1 is empty at that point, both resources move to the collaborative task. In other words, when $\mathcal{W}_+(t) \leq (1-\rho^{\mathrm{BN}})^{-2}$, resource 1 prioritizes its individual task while resource 2 prioritizes the collaborative task but works in its individual task if resource 1 is working in activity 1. When $\mathcal{W}_+(t) > (1-\rho^{\text{BN}})^{-2}$ both resources prioritize activity 0 (the collaborative task). Mostly, the policy is one where resource 1 prioritizes his individual task and resource 2 follows. The difference is only in point in time where the total workload is very large.

Under our proposed policy, resource 1 is always busy as long as it has work in either queue 0 or queue 1. Hence, as in a work conserving single-server queue, resource 1's workload is of the order of $(1-\rho^{\text{BN}})^{-1}$. The workload $m_0 Q_0(t)$ in queue 0 (that is a part of resource 1's workload) will, in particular be of this order. If the workload of resource 2, $\mathcal{W}_2(t) = m_0Q_0(t) + m_2Q_2(t)$ is greater than $(1-\rho^{\text{BN}})^{-(1+\epsilon)} >> (1-\rho^{\text{BN}})^{-1}$, then $m_2Q_2(t) = \mathcal{W}_+(t) - m_0Q_0(t) > 0$ and resource 2 will be working and depleting his workload at a speed of $(1-\rho^{\text{BN}})$. In words, the "constrained drift" of $\mathcal{W}_2(t)$ (constrained on the workload of resource 1's being $\mathcal{O}(1-\rho^{\text{BN}})^{-1}$) will be suitably negative. A mechanism for deriving steady-state bounds based on such "constrained drift" was developed in Gurvich (2013) but the analysis below is self-contained.

Lemma A.1 *Suppose that Q is a J dimensional non-explosive and positive recurrent continuous time Markov chain on $\mathbb{Z}_+^J$ with generator $\mathcal{Q}$ and let π be its steady-state distribution. Suppose further that there exist functions V, U integrable with respect to π,* inclusion *set $\mathcal{E} \subseteq \mathbb{Z}_+^J$ and constants c_1, c_2, c_3 such that $|\mathcal{Q}U(x)| \le c_1(1+V(x))$ for all $x \in \mathbb{Z}_+^J$ and*

$$\mathcal{Q}U(x) \le -c_2V(x) + c_3, \ \ x \in \mathcal{E}.$$

Then,

$$\mathbb{E}_\pi[V(Q(0))] \le \frac{c_3}{c_2} + \frac{c_1}{c_2}\mathbb{E}_\pi[(1+V(Q(0)))\mathbb{1}\{Q^{\rho^{BN}}(0) \notin \mathcal{E}\}]. \tag{34}$$

All auxiliary lemmas are proved in Section A.4 at the end of this companion.

Given a state $x = (q_0, q_1, q_2)$, let $w_i(x) = m_0q_0 + m_iq_i$, $i = 1, 2$ and $w_+(x) = w_1(x) + w_2(x)$. Then, $\mathcal{W}_i(t) = w_i(Q(t))$ and $\mathcal{W}_+(t) = w_+(Q(t))$. Consider the Markov chain $Q^{\rho^{\text{BN}}}(t) = (Q_0^{\rho^{\text{BN}}}(t), Q_1^{\rho^{\text{BN}}}(t), Q_2^{\rho^{\text{BN}}}(t))$ and denote by $\mathcal{Q}^{\rho^{\text{BN}}}$ its generator. Define the function

$$V_{\rho^{\text{BN}}}(x) = e^{(1-\rho^{\text{BN}})\left[w_2(x)-(1-\rho^{\text{BN}})^{-(1+\epsilon)}\right]^+}.$$

The following allows us to apply Lemma A.1 to our setting. We will show that $U = V = V_{\rho^{\text{BN}}}$ exhibit the desired drift condition with the exclusion set

$$\mathcal{E}^{\rho^{\text{BN}}} := \{x : w_1(x) \le (1-\rho^{\text{BN}})^{-(1+\epsilon/2)}\}.$$

Lemma A.2 *Consider the scaled Markov chain* $\widehat{Q}^{\rho^{BN}}(t) = Q^{\rho^{BN}}((1-\rho^{BN})^{-2}t)$. *Let* $\widehat{\mathcal{Q}}^{\rho^{BN}}$ *be its generator. There exist constants* $\bar{c}_1, \bar{c}_2, \bar{c}_3$ *that do not depend on* ρ^{BN} *and such that*

$$\widehat{\mathcal{Q}}^{\rho^{BN}} V_{\rho^{BN}}(x) \leq -\bar{c}_2 V_{\rho^{BN}}(x) + \bar{c}_3 (1-\rho^{BN})^{-1},\ x \in \mathcal{E}^{\rho^{BN}},$$

and, $|\mathcal{Q}^{\rho^{BN}} V_{\rho^{BN}}(x)| \leq \bar{c}_1 (1-\rho^{BN})^{-1}(1+V_{\rho^{BN}}(x))$, *for all queue values* $x = (q_1, q_2, q_3)$.

The next lemma completes the requirements for Lemma A.2 and provides some initial crude bounds.

Lemma A.3 *Suppose that* $\rho^{BN} < 1$. *The process* $Q^{\rho^{BN}}(t)$ *is positive recurrent (with the steady-state distribution denoted by* $\pi^{\rho^{BN}}$*) and, for each* k*, the* k^{th} *moment satisfies*

$$\mathbb{E}_{\pi^{\rho^{BN}}}[\mathcal{W}_1^k(0)] = \mathcal{O}((1-\rho^{BN})^{-k}), \tag{35}$$

and

$$\mathbb{E}_{\pi^{\rho^{BN}}}[\mathcal{W}_+^k(0)] = \mathcal{O}((1-\rho^{BN})^{-2k}).$$

Using these we have that

$$\begin{aligned}\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[(1+V_{\rho^{\mathrm{BN}}}(Q^{\rho^{\mathrm{BN}}}(0)))\mathbb{1}\{Q^{\rho^{\mathrm{BN}}}(0)\notin\mathcal{E}^{\rho^{\mathrm{BN}}}\}] &\leq \widehat{c}_1\sqrt{1+\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[(\mathcal{W}_+^{\rho^{\mathrm{BN}}}(0))^2]}\sqrt{\mathbb{P}_{\pi^{\rho^{\mathrm{BN}}}}\{Q^{\rho^{\mathrm{BN}}}(0)\notin\mathcal{E}^{\rho^{\mathrm{BN}}}\}} \\ &\leq \widehat{c}_2(1-\rho^{\mathrm{BN}})^{-2}(1-\rho^{\mathrm{BN}})^{k\epsilon},\end{aligned}$$

for constants $\widehat{c}_1, \widehat{c}_2$, where the last inequality follows from lemma A.3 that guarantees, by Markov's inequality, the existence of $\widehat{c}_3$, such that

$$\mathbb{P}_{\pi^{\rho^{\mathrm{BN}}}}\{Q^{\rho^{\mathrm{BN}}}\notin\mathcal{E}^{\rho^{\mathrm{BN}}}\} = \mathbb{P}_{\pi^{\rho^{\mathrm{BN}}}}\{\mathcal{W}_1(0) > (1-\rho^{\mathrm{BN}})^{-(1+\epsilon)}\} \leq \widehat{c}_3(1-\rho^{\mathrm{BN}})^{k\epsilon},\ \text{for all } k=1,2,\ldots.$$

Choosing $k \geq 3/\epsilon$ we finally have a constant $\widehat{c}_4$ such that

$$\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[(1+V_{\rho^{\mathrm{BN}}}(Q^{\rho^{\mathrm{BN}}}(0)))\mathbb{1}\{Q^{\rho^{\mathrm{BN}}}(0)\notin\mathcal{E}^{\rho^{\mathrm{BN}}}\}] \leq \widehat{c}_4,$$

Replacing in Lemma A.1, $c_2 = \bar{c}_2$, $c_3 = \bar{c}_3(1-\rho^{\mathrm{BN}})^{-1}$ and $c_1 = \bar{c}_1(1-\rho^{\mathrm{BN}})^{-1}$ (with $\bar{c}_1, \bar{c}_2, \bar{c}_3$ from Lemma A.3) we then have

$$\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[V_{\rho^{\mathrm{BN}}}(Q^{\rho^{\mathrm{BN}}}(0))] \leq \frac{\bar{c}_3}{1-\rho^{\mathrm{BN}}} + \frac{\bar{c}_1}{1-\rho^{\mathrm{BN}}}\widehat{c}_4,$$

and we conclude that

$$\limsup_{\rho^{\mathrm{BN}}\uparrow 1}(1-\rho^{\mathrm{BN}})\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}\left[V_{\rho^{\mathrm{BN}}}(Q^{\rho^{\mathrm{BN}}}(0))\right] < \infty.$$

Recalling that $V_{\rho^{\mathrm{BN}}}(x)=e^{(1-\rho^{\mathrm{BN}})\left[w_2(x)-(1-\rho^{\mathrm{BN}})^{-(1+\epsilon)}\right]^+}$, this implies that

$$\limsup_{\rho^{\mathrm{BN}}\uparrow 1}(1-\rho^{\mathrm{BN}})^{l(1+\epsilon)}\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}\left[(W_2^{\rho^{\mathrm{BN}}})^l(0)\right]<\infty, \tag{36}$$

for each integer l.

To complete the proof it remains to show that

$$(1-\rho^{\mathrm{BN}})\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[Q_1^{\rho^{\mathrm{BN}}}(0)]\to 0,\ \text{ as } \rho^{\mathrm{BN}}\uparrow 1.$$

Here, we take $V(x)=m_1q_1$ and $U=V^2$ and for the exclusion set we take

$$B^{\rho^{\mathrm{BN}}}:=\{x: w_+(x)\le(1-\rho^{\mathrm{BN}})^{-(1+2\epsilon)}\}.$$

In contrast to Lemma A.2, no time or space scaling is used below.

Lemma A.4 *There exist constants b,c_1,c_2 that do not depend on ρ^{BN} and such that*

$$\mathcal{Q}^{\rho^{BN}}\widehat{V}^2(x)\le -2(1-\rho_1^{a,\rho^{BN}})\widehat{V}(x)+(1+\rho_1^{a,\rho^{BN}})m_1,\ x\in B^{\rho^{BN}},$$

and, $|\mathcal{Q}^{\rho^{BN}}V^2(x)|\le c_1(1+V(x))$, for all queue values $x=(q_1,q_2,q_3)$.

Since $\mathcal{W}_+((t)=\mathcal{W}_1(t)+\mathcal{W}_2(t)$, we have from (36) and (35) that $\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[(\mathcal{W}_+^{\rho^{\mathrm{BN}}})^k(0)]\le 2^{k-1}(\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[(\mathcal{W}_2^{\rho^{\mathrm{BN}}})^k(0)]+\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[(\mathcal{W}_1^{\rho^{\mathrm{BN}}})^k(0)])=\mathcal{O}((1-\rho^{\mathrm{BN}})^{k(1+\epsilon)})$. By Markov's inequality

$$\mathbb{P}_{\pi^{\rho^{\mathrm{BN}}}}\{\mathcal{W}_+^{\rho^{\mathrm{BN}}}(0)>(1-\rho^{\mathrm{BN}})^{-(1+2\epsilon)}\}=\mathcal{O}\left(\frac{(1-\rho^{\mathrm{BN}})^{-(k+k\epsilon)}}{(1-\rho^{\mathrm{BN}})^{-(k+2k\epsilon)}}\right)=\mathcal{O}((1-\rho^{\mathrm{BN}})^{k\epsilon})$$

Then,

$$\begin{aligned}\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[|\mathcal{Q}U(Q^{\rho^{\mathrm{BN}}(0)})|\mathbb{1}\{Q^{\rho^{\mathrm{BN}}}(0)\notin B^{\rho^{\mathrm{BN}}}\}]&\le\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[c(1+V(Q^{\rho^{\mathrm{BN}}}(0)))\mathbb{1}\{Q^{\rho^{\mathrm{BN}}}(0)\notin B^{\rho^{\mathrm{BN}}}]\\&\le\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[c(1+V(Q^{\rho^{\mathrm{BN}}}(0)))\mathbb{1}\{Q^{\rho^{\mathrm{BN}}}(0)\notin B^{\rho^{\mathrm{BN}}}\}]\\&\le c\sqrt{1+\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[(W_1^{\rho^{\mathrm{BN}}}(0))^2]}\sqrt{\mathbb{P}\{Q^{\rho^{\mathrm{BN}}}(0)\notin B^{\rho^{\mathrm{BN}}}\}}\\&\le c(1-\rho^{\mathrm{BN}})^{-1}(1-\rho^{\mathrm{BN}})^{k\epsilon}.\end{aligned}$$

Taking $k\epsilon>1$ we have, as required, that

$$\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[V(Q^{\rho^{\mathrm{BN}}}(0))]\le\frac{(1+\rho_1^{a,\rho^{\mathrm{BN}}})m_1+c(1-\rho^{\mathrm{BN}})^{k\epsilon-1}}{2(1-\rho_1^{a,\rho^{\mathrm{BN}}})},$$

so that for any $\delta>0$, and for all ρ^{BN} sufficiently large,

$$\limsup_{\rho^{\mathrm{BN}}\uparrow 1}\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[V(Q^{\rho^{\mathrm{BN}}}(0))]\le\frac{\rho_1^{a,\rho^{\mathrm{BN}}}m_1+m_1}{2(1-\rho_1^{a,\rho^{\mathrm{BN}}})}+\delta.$$

Finally, since $\rho_1^{a,\rho^{\mathrm{BN}}}<1$,

$$\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[W_1(0)]\le m_1(\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[Q_1(0)]+1)=\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[V(Q^{\rho^{\mathrm{BN}}}(0))]+m_1\le\frac{2m_1}{(1-\rho_1^{a,\rho^{\mathrm{BN}}})}+\delta.$$

■

Proof of Theorem 3.3. For each ρ^{BN}, we fix a stationary policy that induces a steady-state distribution $\pi^{\rho^{\text{BN}}}$. Define

$$\mathcal{B} := \{x \in \mathcal{X} : b \notin \{(1,0,0),(0,1,1)\}\}.$$

This is the set of states in which at least one of the two resources idles. Set $\tau_0 = 0$ and define recursively

$$\tau_j = \inf\{t \geq \tau_{j-1} : B(t-) \in \{(0,1,0),(0,0,1),(0,1,1)\} \text{ and } B(t) \in \{(1,0,0),(0,0,0)\}\}.$$

These are times where the two resources leave the individual tasks. Let $N(t)$ be the number of such switches by time t, i.e.,

$$N(t) = \sup\{m : \tau_m \leq t\}.$$

Let

$$\bar{\tau}_j = \inf\{t \geq \tau_{j-1} : B(t) \in \{(0,1,0),(0,0,1),(0,1,1)\}.$$

This is the first time after the j^{th} visit to the collaborative task that at least one of the resources 1 or 2 begins working on an individual activity. Due to non-preemption there must exist $t \in [\bar{\tau}_j, \tau_j)$ with $\mathbb{X}(t) \in \mathcal{B}$, i.e., at least one of the servers 1 or 2 must be idle before switching to activity 0. Here, we use also the fact that because of the exponential service times, the probability of simultaneous service completions in two tasks is 0.

Let X_j be the time that the process stays in $\mathcal{B}$ when visiting it for the first time after $\bar{\tau}_j$. The random variables $X_1, X_2, \ldots,$ are independent and X_j is at least as long as the time it takes until some arrival or service completion. In particular, $\mathbb{E}[X_j] > 1/(2\Lambda)$ where $\Lambda = \sum_i(\lambda_0 + \mu_i)$. Choose a constant c_X such that $\mathbb{E}[X_j \wedge c_X] \geq \frac{1}{2\Lambda}$. Let $I_i(t)$ be the cumulative idleness of resource i by time t. Let $I_i(t)$ be the accumulated idleness of resource i by time t. The total idleness $I_+(t) = I_1(t) + I_2(t)$ is bounded from below by the idleness accumulated during visits to $\mathcal{B}$, i.e.,

$$\mathbb{E}_{\pi^{\rho^{\text{BN}}}}\left[I_+(t)\right] \geq \mathbb{E}_{\pi^{\rho^{\text{BN}}}}\left[\sum_{j=1}^{N(t)} X_j\right],$$

and

$$\begin{aligned}\mathbb{E}_{\pi^{\rho^{\text{BN}}}}[I_+(t)] &\geq \mathbb{E}_{\pi^{\rho^{\text{BN}}}}\left[\sum_{j=1}^{N(t)} X_j \wedge c_X\right] \geq \mathbb{E}_{\pi^{\rho^{\text{BN}}}}\left[\sum_{j=1}^{N(t)+1} X_j \wedge c_X\right] - \mathbb{E}_{\pi^{\rho^{\text{BN}}}}\left[X_{N(t)+1} \wedge c_X\right] \\ &\geq \frac{1}{2\Lambda}\mathbb{E}_{\pi^{\rho^{\text{BN}}}}\left[N(t)+1\right] - c_X,\end{aligned}$$

where the last inequality follows from Wald's identity. Thus,

$$\mathbb{E}_{\pi^{\rho^{\text{BN}}}}[N(t)]\mathbb{E}[X_j \wedge c_X] \leq (1-\rho_1+1-\rho_2)t + c_X \leq 2(1-\rho^{\text{BN}})t + c_X,$$

so that

$$\mathbb{E}_{\pi^{\rho^{\rm BN}}}[N(t)] \leq 2\Lambda(1-\rho^{\rm BN})t + 2\Lambda c_X \tag{37}$$

Denote by $Z_j = \tau_{j+1} - \bar{\tau}_j$ the time allocated to individual tasks during the j^{th} cycle. The total time spent in individual task 1 must satisfy, in stationarity,

$$\begin{aligned}\mathbb{E}_{\pi^{\rho^{\rm BN}}}\left[\sum_{j=1}^{N(t)+1} Z_j\right] &\geq \mathbb{E}_{\pi^{\rho^{\rm BN}}}\left[\int_0^t \mathbb{1}\{B(s)\in\{(0,1,0),(0,1,1)\}\}ds\right] \\ &= t\mathbb{E}_{\pi^{\rho^{\rm BN}}}\left[B(0)\in\{(0,1,0),(0,1,1)\}\right] = \rho_1^{a,\rho^{\rm BN}} t, \ t\geq 0,\end{aligned} \tag{38}$$

where, recall, $\rho_i^{a,\rho^{\rm BN}} = \lambda_i^{\rho^{\rm BN}}/\mu_0$. The same applies to activity 2. Moreover, using (37), we have for any $\delta > 0$

$$\begin{aligned}\mathbb{E}_{\pi^{\rho^{\rm BN}}}\left[\sum_{j=1}^{N(t)+1} Z_j \mathbb{1}\{Z_j \geq \delta\}\right] &\geq \rho_i^{a,\rho^{\rm BN}} t - \mathbb{E}_{\pi^{\rho^{\rm BN}}}\left[\sum_{j=1}^{N(t)} Z_j \mathbb{1}\{Z_j\leq \delta\}\right] \\ &\geq \rho_i^a t - \delta\mathbb{E}_{\pi^{\rho^{\rm BN}}}[N(t)] \\ &\geq \gamma^{\rho^{\rm BN}} t - 2\delta(1-\rho^{\rm BN})\Lambda t - 2\Lambda c_X,\end{aligned}$$

where $\gamma^{\rho^{\rm BN}} = \min\{\rho_0^{a,\rho^{\rm BN}}, \rho_2^{a,\rho^{\rm BN}}, \rho_3^{a,\rho^{\rm BN}}\}$. Letting $\gamma := \liminf_{\rho^{\rm BN}\uparrow 1}\gamma^{\rho^{\rm BN}}$ and setting $\delta^{\rho^{\rm BN}} = \frac{1}{4}\frac{\gamma}{2(1-\rho^{\rm BN})\Lambda}$ we have

$$\mathbb{E}_{\pi^{\rho^{\rm BN}}}\left[\sum_{j=1}^{N(t)} Z_j \mathbb{1}\{Z_j \geq \delta^{\rho^{\rm BN}}\}\right] \geq \frac{3}{4}\gamma t - 2\Lambda c_X.$$

Taking $t^{\rho^{\rm BN}} = (1-\rho^{\rm BN})^{-1}$, we have that for all $\rho^{\rm BN}$ sufficiently close to one

$$\mathbb{E}_{\pi^{\rho^{\rm BN}}}\left[\sum_{j=1}^{N(t^{\rho^{\rm BN}})} Z_j \mathbb{1}\{Z_j \geq \delta^{\rho^{\rm BN}}\}\right] \geq \frac{1}{2}\gamma t^{\rho^{\rm BN}}. \tag{39}$$

Intuitively speaking, we expect to see in queue 0 the pattern in Figure A.1 with an accumulation of the order of $A_0(\bar{\tau}_j + Z_j) - A_0(\bar{\tau}_j) \approx \lambda_0^{\rho^{\rm BN}} Z_j$ during the j^{th} cycle through the individual tasks. The area of the j^{th} triangle should be of the order $\lambda_0^{\rho^{\rm BN}} Z_j^2$ so that, initializing the chain in stationarity, the average queue should be bounded below by $\frac{1}{t}\mathbb{E}_{\pi^{\rho^{\rm BN}}}\left[\sum_{j=1}^{N(t)} \lambda_0^{\rho^{\rm BN}} Z_j \delta^{\rho^{\rm BN}} \mathbb{1}\{Z_j \geq \delta^{\rho^{\rm BN}}\}\right]$ which can be further bounded using (39).

To formalize this intuition introduce the event

$$\mathcal{A}_{t^{\rho^{\rm BN}}} := \left\{\omega\in\Omega : \frac{|A_0(u) - A_0(s) - \lambda_0^{\rho^{\rm BN}}(u-s)|}{(\lambda_0^{\rho^{\rm BN}}(u-s))^{1/2+\epsilon}} \leq 1, \text{ for all } s,u\leq t^{\rho^{\rm BN}}, |u-s|\geq \delta^{\rho^{\rm BN}}\right\}.$$

On $\mathcal{A}_{t^{\rho^{\rm BN}}}$, we have for all $j\leq N(t)$ with $Z_j \geq \delta^{\rho^{\rm BN}}$ that

$$|A_0(\bar{\tau}_j + Z_j) - A_0(\bar{\tau}_j) - \lambda_0^{\rho^{\rm BN}} Z_j| \leq (\lambda_0^{\rho^{\rm BN}} Z_j)^{1/2+\epsilon}.$$

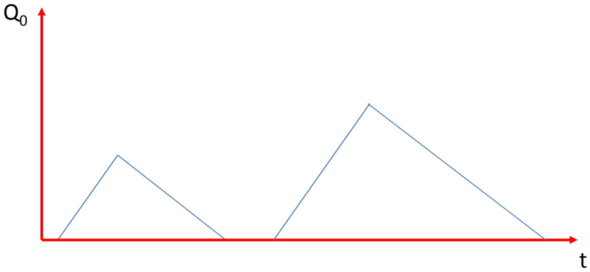

Figure 15 A buildup diagram for the collaborative queue

Noticing that $\lambda_0^{\rho^{\mathrm{BN}}}\delta^{\rho^{\mathrm{BN}}}\geq 1$ for all sufficiently large ρ^{BN}, we then get

$$\int_0^t Q_0(s)ds \geq \sum_{j=1}^{N(t)} \frac{\lambda_0^{\rho^{\mathrm{BN}}} Z_j^2}{2} 1\{Z_j \geq \delta^{\rho^{\mathrm{BN}}}\} - \sum_{j=1}^{N(t)} Z_j(\lambda_0^{\rho^{\mathrm{BN}}} Z_j)^{1/2+\epsilon} 1\{Z_j \geq \delta^{\rho^{\mathrm{BN}}}\}.$$

For all $z \geq \delta^{\rho^{\mathrm{BN}}} = \frac{1}{8}\frac{\gamma}{(1-\rho^{\mathrm{BN}})\Lambda}$ and all ρ^{BN} sufficiently close to 1, $\lambda_0^{\rho^{\mathrm{BN}}} z^2/2 - (\lambda_0^{\rho^{\mathrm{BN}}})^{1/2+\epsilon} z^{1/2+\epsilon} \geq \bar{c}z^2$ for some constant $\bar{c}$ that does not depend on ρ^{BN}. On $\mathcal{A}_{t^{\rho^{\mathrm{BN}}}}$,

$$\int_0^{t^{\rho^{\mathrm{BN}}}} Q_0(s)ds \geq \sum_{j=1}^{N(t^{\rho^{\mathrm{BN}}})} \bar{c} Z_j^2 1\{Z_j \geq \delta^{\rho^{\mathrm{BN}}}\} \geq \bar{c}\delta^{\rho^{\mathrm{BN}}} \sum_{j=1}^{N(t^{\rho^{\mathrm{BN}}})} Z_j 1\{Z_j \geq \delta^{\rho^{\mathrm{BN}}}\},$$

where in the last inequality we replaced Z_j^2 with $\delta^{\rho^{\mathrm{BN}}} Z_j$. Thus,

$$\begin{aligned}
\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}\left[\mathbb{1}\{\mathcal{A}_{t^{\rho^{\mathrm{BN}}}}\}\int_0^{t^{\rho^{\mathrm{BN}}}} Q_0(s)ds\right] &\geq \mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}\left[\mathbb{1}\{\mathcal{A}_{t^{\rho^{\mathrm{BN}}}}\}\bar{c}\delta^{\rho^{\mathrm{BN}}} \sum_{j=1}^{N(t^{\rho^{\mathrm{BN}}})} Z_j 1\{Z_j \geq \delta^{\rho^{\mathrm{BN}}}\}\right] \\
&= \mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}\left[\bar{c}\delta^{\rho^{\mathrm{BN}}} \sum_{j=1}^{N(t^{\rho^{\mathrm{BN}}})} Z_j 1\{Z_j \geq \delta^{\rho^{\mathrm{BN}}}\}\right] \\
&\qquad - \mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}\left[\mathbb{1}\{(\mathcal{A}_{t^{\rho^{\mathrm{BN}}}})^c\}\bar{c}\delta^{\rho^{\mathrm{BN}}} \sum_{j=1}^{N(t^{\rho^{\mathrm{BN}}})} Z_j 1\{Z_j \geq \delta^{\rho^{\mathrm{BN}}}\}\right] \\
&\geq \bar{c}t^{\rho^{\mathrm{BN}}}(\delta^{\rho^{\mathrm{BN}}} - \mathbb{P}_{\pi}\{\mathcal{A}_{t^{\rho^{\mathrm{BN}}}}\})
\end{aligned}$$

for a re-defined constant $\bar{c}$. In the last inequality we used (39) and $\sum_{j=1}^{N(t^{\rho^{\mathrm{BN}}})} Z_j \leq t^{\rho^{\mathrm{BN}}}$. In particular,

$$\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}\left[\int_0^{t^{\rho^{\mathrm{BN}}}} Q_0(s)ds\right] \geq \mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}\left[\mathbb{1}\{\mathcal{A}_{t^{\rho^{\mathrm{BN}}}}\}\int_0^{t^{\rho^{\mathrm{BN}}}} Q_0(s)ds\right] \geq \bar{c}\delta^{\rho^{\mathrm{BN}}} t^{\rho^{\mathrm{BN}}} - \bar{c}\mathbb{P}_{\pi^{\rho^{\mathrm{BN}}}}\{(\mathcal{A}_{t^{\rho^{\mathrm{BN}}}})^c\} t^{\rho^{\mathrm{BN}}};$$

Lemma A.5

$$t^{\rho^{BN}}\mathbb{P}_{\pi^{\rho^{BN}}}\{(\mathcal{A}_{t^{\rho^{BN}}})^c\} \to 1, \;\; as \;\rho^{BN}\uparrow 1. \tag{40}$$

Using (A.5) and stationarity we have

$$t^{\rho^{\rm BN}}\mathbb{E}_{\pi^{\rho^{\rm BN}}}[Q_0(0)] = \mathbb{E}_{\pi^{\rho^{\rm BN}}}\left[\int_0^{t^{\rho^{\rm BN}}} Q_0(s)ds\right] \geq \bar{c}\delta^{\rho^{\rm BN}}t^{\rho^{\rm BN}}$$

for a re-defined constant $\bar{c}$. We assume here without loss of generality that $\mathbb{E}_{\pi^{\rho^{\rm BN}}}[Q_0^{\rho^{\rm BN}}(0)] < \infty$. If it does not, our result holds trivially. Dividing by $t^{\rho^{\rm BN}}$ on both sides and using $\delta^{\rho^{\rm BN}} = \frac{1}{4}\gamma/(4(1-\rho^{\rm BN})\Lambda)$ we have then

$$\mathbb{E}_{\pi^{\rho^{\rm BN}}}[Q_0(0)] = \frac{1}{t^{\rho^{\rm BN}}}t^{\rho^{\rm BN}}\mathbb{E}_{\pi^{\rho^{\rm BN}}}[Q_0(0)] \geq \bar{c}\frac{1}{1-\rho^{\rm BN}},$$

for a re-defined constant $\bar{c}$ which, by Little's law and since $\lambda_0 m_0 \leq 1$, implies $\mathbb{E}_{\pi^{\rho^{\rm BN}}}[W_0(0)] \geq \frac{1}{\lambda_0}\frac{\bar{c}}{(1-\rho^{\rm BN})} \geq m_0\frac{\bar{c}}{(1-\rho^{\rm BN})}$.

Finally, we note that one could repeat the arguments above taking an individual task i as the focal activity to conclude that each of the individual tasks must grow with $\rho^{\rm BN}$. In that proof $N(t)$ would be the number of returns to individual tasks. ■

Proof of Proposition 2. Under polling, the process $\mathbb{X}(t) = (Q(t), B(t))$ is a continuous time Markov chain. We study the collaborative queue Q_0 sampled at times in which the resources move from the individual tasks to the collaborative task. These are return times to the family of states $\{((q_0,0,0),(1,0,0))\}$ with $q_0 > 0$. Formally, set $\tau_0 = 0$, $(Q(0), B(0)) = ((q_0,0,0),(1,0,0))$, and define

$$\tau_{l+1} = \inf\{t \geq \tau_l : Q_0(t) > 0, Q_1(t) + Q_2(t) = 0, B(t-) \in \{(0,1,0),(0,0,1)\}\};$$

(as the probability of simultaneous service completions is 0 it cannot be that $B(t-) = (0,1,1)$). Since the individual queues are empty at τ_l, the discrete time process $Q_0^l = Q_0(\tau_l)$ is a one-dimensional discrete time Markov chain. It is also aperiodic and irreducible.

For $l = 0, \ldots,$ let

$$\bar{\tau}_l = \inf\{t \geq \tau_l : Q_0(t) = 0, Q_1(t) + Q_2(t) > 0\}$$

be the first moment after τ_l that resources drain the collaborative activity 0 and return to the individual activities. Since $\rho_1 = \rho_2 = \rho^{\rm BN}$ $\bar{\rho}^a := \max\{\rho_1^a, \rho_2^a\} = \rho_1^a = \rho_2^a$. To simplify notation we write $\mathbb{E}_{q_0}[\cdot]$ to mean $\mathbb{E}_{(q_0,0,0),(1,0,0)}[\cdot]$

Lemma A.6 *Suppose that $x = ((q_0,0,0),(1,0,0))$ for some $q_0 > 0$. Then, $\mathbb{E}_x[(\bar{\tau}_0)^2] < \infty$,*

$$\mathbb{E}_{q_0}[\bar{\tau}_0] = \frac{q_0}{\mu_0 - \lambda_0}, \text{ and } \bar{q}_i := \mathbb{E}_{q_0}[Q_i(\bar{\tau}_0)] = \lambda_i\frac{q_0}{\mu_0 - \lambda_0} = q_0\mu_i\frac{\rho_i^a}{\mu_0(1-\rho_0^a)} = \mu_i\tilde{q}, \tag{41}$$

where $\tilde{q} := q_0\bar{\rho}^a/(\mu_0(1-\rho_0^a))$. *Also,*

$$\mathbb{E}_{q_0}[\tau_1 - \bar{\tau}_0] = \frac{\tilde{q}}{1-\bar{\rho}^a} + \Theta(\sqrt{q_0}) = q_0\frac{\bar{\rho}^a}{(1-\bar{\rho}^a)(\mu_0-\lambda_0)} + \Theta(\sqrt{q_0}), \tag{42}$$

and

$$\mathbb{E}_{q_0}[\tau_1] = \Theta(q_0). \tag{43}$$

Lemma A.6 shows that, as the initial collaborative queue length q_0 grows, the time until all resources deplete their individual queues and ready to move to the collaborative task takes $\sqrt{q_0}$ more than the *fluid* depletion time $\tilde{q}/(1-\bar{\rho}^a)$. The queue that hits zero first will oscillate and accumulate substantial (specifically order $\sqrt{q_0}$) idle time before the resources move. This is the mathematical manifestation of the simulation in Figure 8.

Since $Q_0(\bar{\tau}_0) = 0$ by definition, we have that

$$\begin{aligned}\mathbb{E}_{q_0}[Q_0(\tau_1)] &= \mathbb{E}_{q_0}[A_0(\tau_1) - A_0(\bar{\tau}_0)] = \lambda_0\mathbb{E}_{q_0}[\tau_1 - \bar{\tau}_0] \\ &= q_0\frac{\lambda_0\bar{\rho}^a}{(1-\bar{\rho}^a)(\mu_0-\lambda_0)} + \lambda_0\Theta(\sqrt{q_0}) = q_0\frac{\bar{\rho}^a\rho_0^a}{(1-\bar{\rho}^a)(1-\rho_0^a)} + \Theta(\sqrt{q_0}),\end{aligned}$$

and

$$\begin{aligned}\mathbb{E}_{q_0}[Q_0^1] - \mathbb{E}_{q_0}[Q_0^0] &= \mathbb{E}_x[Q_0(\tau_1)] - \mathbb{E}_{q_0}[Q_0(\tau_0)] = -q_0\left(1 - \frac{\bar{\rho}^a\rho_0^a}{(1-\bar{\rho}^a)(1-\rho_0^a)}\right) + \Theta(\sqrt{q_0}) \\ &= -q_0\frac{1}{(1-\bar{\rho}^a)(1-\rho_0^a)}(1-\rho^{\mathrm{BN}}) + \Theta(\sqrt{q_0}).\end{aligned}$$

There then exists constant c, η, b such that

$$\sup_{q_0 \geq c(1-\rho^{\mathrm{BN}})^{-2}} \left(\mathbb{E}_{q_0}[Q_0^1] - q_0\right) \leq -\eta q_0(1-\rho^{\mathrm{BN}}), \text{ and } \mathbb{E}_{q_0}[Q_0^1] - q_0 \leq b(1-\rho^{\mathrm{BN}})^{-1}.$$

In particular,

$$\left(\mathbb{E}_{q_0}[Q_0^1] - q_0\right) \leq -\eta q_0(1-\rho^{\mathrm{BN}}) + b(1-\rho^{\mathrm{BN}})^{-1},$$

This implies that the DTMC is positive recurrent (e.g. (Robert, 2003, Theorem 8.6)) and moreover, that under its steady-state distribution π

$$\mathbb{E}_\pi[Q_0^0] \leq \frac{b}{\eta}(1-\rho^{\mathrm{BN}})^{-2};$$

see (Glynn and Zeevi, 2008, Corollary 4). In turn, $\mathbb{E}_\pi[\tau_1] = \Theta(\mathbb{E}_\pi[Q_0^0]) < \infty$. By (23) this further guarantees that the CTMC is also positive recurrent. Applying expectations with respect to the stationary distribution, we get

$$0 = \mathbb{E}_\pi[Q_0^1] - \mathbb{E}_\pi[Q_0^0] = -\mathbb{E}_\pi[Q_0^0]\frac{1}{(1-\bar{\rho}^a)(1-\rho_0)}(1-\rho^{\mathrm{BN}}) + \Theta(\mathbb{E}_\pi[\sqrt{Q_0^0}]),$$

so that

$$\frac{\mathbb{E}_\pi[Q_0^0]}{\mathbb{E}_\pi[\sqrt{Q_0^0}]}=\Theta\left(\frac{1}{1-\rho^{\mathrm{BN}}}\right).$$

If

$$\mathbb{E}_\pi[\sqrt{Q_0^0}]=\Omega((1-\rho^{\mathrm{BN}})^{-1/2}), \tag{44}$$

then there exists a constant $\bar{c}>0$ for which

$$\frac{\mathbb{E}_\pi[Q_0^0]}{\bar{c}(1-\rho^{\mathrm{BN}})^{-1/2}}\geq\frac{\mathbb{E}_\pi[Q_0^0]}{\mathbb{E}_\pi[\sqrt{Q_0^0}]}=\Theta(1/(1-\rho^{\mathrm{BN}})),$$

and, in turn,

$$\mathbb{E}_\pi[Q_0^0]=\Omega((1-\rho^{\mathrm{BN}})^{-3/2}).$$

We postpone the argument for (44) and translate this DTMC bound to one for the continuous time chain. The following is standard.

Lemma A.7 *Fix x and let $\widehat{\tau}_1$ be a stopping time with $\mathbb{E}_x[(\widehat{\tau}_1)^2]<\infty$. Then,*

$$\mathbb{E}_x[Q_0^2(\widehat{\tau}_1)]=\mathbb{E}_x\left[Q_0^2(0)\right]+2\lambda_0\mathbb{E}_x\left[\int_0^{\widehat{\tau}_1}Q_0(s)ds\right]-2\mu_0\mathbb{E}_x\left[\int_0^{\widehat{\tau}_1}Q_0(s)\mathbb{1}\{Q_0(s)>0\}ds\right]. \tag{45}$$

Taking $\widehat{\tau}_1=\bar{\tau}_0$, we have by Lemma A.6 that $\mathbb{E}_x[(\bar{\tau}_0)^2]<\infty$. A standard argument then also show that $\mathbb{E}_x[Q_0^2(\bar{\tau}_0)]=0$. Since, by definition, $\mathbb{1}\{Q_0(s)>0\}=1$ for all $t\in[0,\bar{\tau}_0)$, Lemma A.7 then gives

$$\mathbb{E}_x\left[\int_0^{\bar{\tau}_0}Q_0(s)ds\right]=\frac{q_0^2}{2(\mu_0-\lambda_0)},$$

for all $x=((q_0,0),(1,0,0))$. Let ν be the steady-state distribution of the CTMC and suppose that $\mathbb{E}_\nu[Q_0(0)]<\infty$ (otherwise, our result holds trivially). Since $\mathbb{E}_\pi[\tau_1]<\infty$, we have by (23) that $\mathbb{E}_\pi\left[\int_0^{\tau_1}Q_0(s)ds\right]<\infty$ and, since $\tau_1\geq\bar{\tau}_0$, also that $\mathbb{E}_\pi\left[\int_0^{\bar{\tau}_0}Q_0(s)ds\right]<\infty$ and $\mathbb{E}_\pi[(Q_0^0)^2]<\infty$. We have that

$$\mathbb{E}_\nu[Q_0(0)]=\frac{\mathbb{E}_\pi[\int_0^{\tau_1}Q_0(s)ds]}{\mathbb{E}_\pi[\tau_1]}\geq\frac{\mathbb{E}_\pi[\int_0^{\bar{\tau}_0}Q_0(s)ds]}{\mathbb{E}_\pi[\tau_1]}\geq\frac{\mathbb{E}_\pi[(Q_0^0)^2]}{2(\mu_0-\lambda_0)\mathbb{E}_\pi[\tau_1]}\geq\frac{\mathbb{E}_\pi^2[Q_0^0]}{2(\mu_0-\lambda_0)\mathbb{E}_\pi[\tau_1]},$$

where the last step follows from Jensen's inequality. By (43), $\mathbb{E}_\pi[\tau_1]=\Theta(\mathbb{E}_\pi[Q_0^0])$ and we conclude that

$$\mathbb{E}_\nu[Q_0(0)]=\Omega(\mathbb{E}_\pi[Q_0^0])=\Omega((1-\rho^{\mathrm{BN}})^{-3/2}),$$

as in the statement of the theorem.

It only remains to argue (44). This follows from a simple bound. Consider a two class single server polling queue with (λ_0,m_0) and (λ_1,m_1) as the parameters for the two classes. Let $\tilde{Q}_0^l$ be the length of queue 0 in this queue upon returns to queue 0 and let $\tilde{\pi}$ be the stationary distribution of this discrete time chain. It is a simple argument that, in stationarity, $\tilde{Q}_0^0 \leq_{st} Q_0^0$. By (van der Mei, 2007, Theorem 2 and Remark 1), there exists $C, \epsilon > 0$ such that $\mathbb{P}_{\tilde{\pi}}\{\tilde{Q}_0^0 > C(1-\rho^{\text{BN}})^{-1}\} \geq \epsilon$. The stochastic ordering then implies that $\mathbb{E}_{\pi}[\sqrt{Q_0^0}] = \Omega((1-\rho^{\text{BN}})^{-1/2})$ as needed. ■

Proof of Proposition 3. Let τ_l be the time of the l^{th} return of the resources to the individual tasks after depleting the collaborative queue. The process $Q^l = (Q_0(\tau_l), Q_1(\tau_l), Q_2(\tau_l))$ is a discrete time Markov chain and has $Q_0(\tau_l) = 0$ for all $l = 1, \ldots$. We will prove that this discrete time chain is positive recurrent and has a steady-state distribution π^D. Since, for any initial state with an empty collaborative queue, $x = (0, q_1, q_2)$,

$$\mathbb{E}_x[\tau_1] \leq \frac{S}{\lambda_0} + \frac{S + \lambda_0 \mathbb{E}[T_s]}{\mu_0 - \lambda_0},$$

we have $\mathbb{E}_{\pi^D}[\tau_1] < \infty$. This subsequently guarantees, by (23), that the chain $Q(t) = (Q_0(t), Q_1(t), Q_2(t))$ is positive recurrent.

The following result will be useful in the study of the discrete chain.

Lemma A.8 *Let $X = ((X_1^l, \ldots, X_J^l); l = 0, 1, \ldots)$ be an irreducible and aperiodic Markov chain on $\mathbb{Z}_+^J$ with transition probability $\mathbb{P}_x\{\cdot\}$. Suppose that for each coordinate i there exists a one dimensional Markov chain $Y_i = (Y_i^l; l = 0, 1, \ldots)$ with transition probability $\mathbb{P}_y^i\{\cdot\}$ such that for each $x \in \mathbb{Z}_+^J$ and all $y \in \mathbb{Z}_+$*

$$\mathbb{P}_x\{X_i^1 \geq y\} \leq \mathbb{P}_{x_i}^i\{Y_i^1 \geq y\}, \text{ for all } i. \tag{46}$$

Suppose further that each Y_i is aperiodic, monotone ($\mathbb{P}_x^i\{Y_i^1 \geq y\} \leq \mathbb{P}_z^i\{Y_i^1 \geq y\}$ for all $x \leq z$ and all y) and positive recurrent. Then X is positive recurrent.

To generate the bounding chains Y_i, consider a two-class single-server queue with **two** queues labeled 0 and i and corresponding arrival and service time means λ_0, m_0 and λ_i, m_i. The server follows a threshold rule: When queue 0 reaches S jobs in the queue, the server moves to queue 0 as soon as its current processing is complete plus a switchover time T_s which is distributed as the maximum of 2 exponential random variables with means m_0 and m_i.

Let τ_l^i be the time of the l^{th} return of this server to queue i and let Y_i^l be the length of queue i upon this return. Then, Y_i^l is a (one-dimensional) discrete time Markov chain that satisfies the comparison (46). Indeed, in the original chain, the "switchover" time will be at most as the maximum of exponentials above. It is also monotone and aperiodic. We omit the simple formalization of these facts.

Tp apply Lemma A.8 it only remains to prove that, for each i, Y_i^l is positive recurrent. We fix i and omit the subscripts from Y. Set $\tau_0 = 0$ and formally define $l \geq 1$

$$\tau_l = \inf\{t \geq \tau_{l-1} : Q_0(t-) = 1,\ Q_0(t) = 0\}.$$

At τ_l the resource moves back to queue i. We will show that $\limsup_{q_i \to \infty} (\mathbb{E}_{q_i}[Y^1] - q_i) \leq -c$, which, in particular, implies the existence of $\bar{q}_i$ such that

$$\left(\mathbb{E}_{q_i}[Y^1] - q_i\right) \leq -c/2, \tag{47}$$

for all $q_i \geq \bar{q}_i$. This guarantees that the discrete chain Y^l is positive recurrent. It is also irreducible and aperiodic and, thus, has a steady-state distribution.

Let $D_i(t)$ be the number of service completions by time t in queue i. Then,

$$\mathbb{E}_{q_i}[D_i(\tau_1) - D_i(\tau_0)] = \mu_i \left(\frac{S}{\lambda_0} + m_i \mathbb{P}_{q_i}\{Y^1 > 0\} \right).$$

The expected time until the threshold is hit is S/λ_0. If Q_i is positive at that moment, there will be an additional service before the resource actually moves. Between the time S is reached and until the resource moves there are arrivals but no service completions. The time that the server works in queue 0 equals in expectation to $(S + \lambda_0 \mathbb{E}[T_s])/(\mu_0 - \lambda_0)$ so that the arrivals to queue i on $[0, \tau_1)$ satisfy

$$\mathbb{E}_{q_i}[A_i(\tau_1) - A_i(\tau_0)] \leq \lambda_i \left(\frac{S}{\lambda_0} + \mathbb{E}[T_s] + \frac{S + \lambda_0 \mathbb{E}[T_s]}{\mu_0 - \lambda_0} \right).$$

In sum,

$$\begin{aligned} \mathbb{E}_{q_i}[Y^1] - q_i \leq & - \left(\frac{S}{\lambda_0} + m_i \mathbb{P}_{q_i}\{Y^1 > 0\} \right) (\mu_i - \lambda_i) \\ & + \lambda_i (\mathbb{E}[T_s] - m_i \mathbb{P}_{q_i}\{Y^1 > 0\}) + \frac{\lambda_i (S + \lambda_0 \mathbb{E}[T_s])}{\mu_0 - \lambda_0} \end{aligned}$$

As $q_i \to \infty$, $\mathbb{P}_{q_i}\{Y^1 > 0\} \to 1$ (the probability of serving all q_i customers before the threshold is reached goes to 0 as the initial queue grows – we omit the simple argument) and we get

$$\begin{aligned} \limsup_{q_i \to \infty} \left(\mathbb{E}_{q_i}[Y^1] - q_i\right) \leq & - \left(\frac{S}{\lambda_0} + m_i \right) (\mu_i - \lambda_i) \\ & + \lambda_i (\mathbb{E}[T_s] - m_i) + \frac{\lambda_i (S + \lambda_0 \mathbb{E}[T_s])}{\mu_0 - \lambda_0}. \end{aligned}$$

Dividing both sides by $\eta := \lambda_0 \lambda_i$ and multiplying by $\delta := \rho_0^a \rho_i^a / (S/\lambda_0 + m_i)$ (recall $\rho_i^a = \lambda_i/\mu_i$) we have that

$$\begin{aligned} & \limsup_{q_i \to \infty} \frac{\delta}{\eta} \left(\mathbb{E}_{q_i}[Y^1] - q_i\right) \\ & \quad \leq -(1 - \rho_0^a - \rho_i^a) + \frac{\rho_i^a (1 - \rho_0^a)(\mathbb{E}[T_s] - m_i) + \left(\frac{S}{\lambda_0} + \mathbb{E}[T_s] \right) \rho_0^a \rho_i^a - \rho_0^a \rho_i^a \left(\frac{S}{\lambda_0} + m_i \right)}{\frac{S}{\lambda_0} + m_i} \\ & \quad \leq -(1 - \rho_0^a - \rho_i^a) + \frac{\rho_i^a (\mathbb{E}[T_s] - m_i)}{\frac{S}{\lambda_0} + m_i}. \end{aligned}$$

The right-hand side is strictly negative if

$$\rho_0^a + \rho_i^a + \frac{\rho_i^a(\mathbb{E}[T_s] - m_i)}{\frac{S}{\lambda_0} + m_i} < 1,$$

as in the statement of the Proposition. Finally, the fact that $2\rho_0^a \mathbb{E}W_0(\infty) + \rho_1^a \mathbb{E}W_1(\infty) + \rho_2^a \mathbb{E}W_2(\infty) \leq c_T(1-\rho^{\text{BN}})^{-1}$ is a special case of Theorem 5. ■

A.2. Proofs for Section 6

Proof of Theorem 6. The theorem is a direct corollary of the representation for W^{net} that precedes it. ■

Proof of Theorem 7. We start with an observation about hierarchical networks and the set $\mathcal{Y}$ of solutions to the SPPC dual

$$\begin{aligned} \max_{y \in \mathbb{R}_+^J} \ & y'(m * \lambda) \\ \text{s.t.} \ & y'C \leq e'. \end{aligned}$$

Take a bottleneck resource j. Let y^j be the vector given by $y_l^j = 1$ for all activities $l : j \in \mathcal{R}(l)$ and $y_k^j = 0$ otherwise. That y^j is an optimal solution to the dual follows then from $\sum_l y_l^j \lambda_l m_l = \rho^{\text{BN}} = \rho^{\text{net}}$ (where $\rho^{\text{BN}} = \rho^{\text{net}}$ follows from the hierarchy and the results of GVM). Further, as a feasible configuration contains at most one of a resource's activities, we have $(y^j)'C \leq e'$. Thus, for every input parameters λ, m there exists an integer optimal solution to the dual corresponding to a bottleneck resource. This guarantees that all extreme points of the polyhedron $y'C \leq e,\ y \geq 0$ are convex combinations of such (bottleneck-based) vectors.

A configuration k that uses a bottleneck resource j has $((y^j)'C)_k = 1$ for each bottleneck j and in particular $1 - (y^j)'C_k = 0$. Hence, only configurations that do not use bottleneck resource j can contribute to the network availability idleness under y^j. Further, by assumption, such a sub-optimal configuration is used only when j has no work in any of its queues. In turn, the network availability idleness under y^j is bounded above by the average amount of time that bottleneck j has no work. In *parallel networks* this is further bounded by $(1-\rho^{\text{BN}})$ as required. Thus, we have

$$\limsup_{t\to\infty} \frac{1}{t} \sum_{k:\pi_k^*=0} (1 - (y^j)'C)_k) I_k^a(t) \leq (1 - \rho^{\text{BN}}).$$

Let $\bar{y}_\alpha = \sum_j \alpha_j y^j$ be a convex combination of such bottleneck-based extreme points. Then,

$$\limsup_{t\to\infty} \frac{1}{t} \sum_{k:\pi_k^*=0} (1 - (\bar{y}_\alpha)'C)_k) I_k^a(t) = \sum_j \alpha_j \limsup_{t\to\infty} \frac{1}{t} \sum_{k:\pi_k^*=0} \alpha_j (1 - (y^j)'C)_k) I_k^a(t) \leq (1 - \rho^{\text{BN}}),$$

as required.

The second part of the theorem follows noting that, under the stated conditions, $\limsup_t \frac{1}{t} Q_i(t) = 0$ almost surely for all queue i. Since, $Q_i(t) = Q_i(0) + A_i(t) - S_i((CT)_i(t))$, the conclusion follows from strong law of large numbers applied to the underlying Poisson processes. ■

A.3. Proofs for Section 5

Proof of Theorem 4. The key observation is that under hierarchical preemptive priorities the queues served by any given resource evolve marginally like a multiclass single-server queue. Consider, for example, the network in Figure 12 and the activities $1,2,3$. When there is work in activity 1, resources 1 and 2 are working in this activity. When there is work in activity 2 but none in 1 these resources are working in activity 2. Preemption and hierarchy guarantee that resource 3 will also move to activity 1 when resources 1 and 2 do so. The process $(Q_1(t),Q_2(t),Q_3(t))$, has the distribution of a three-class $M/M/1$ queue with utilization $\lambda_1 m_1+\lambda_2 m_2+\lambda_3 m_3\le\rho^{\rm BN}<1$. The formulas in theorem then follow those of static priority queues. ■

Proof of Theorem 5. The non-preemptive case presents two challenges relative to Theorem 4: resources may idle waiting for the thresholds to be reached (coordination idleness) and, once thresholds are reached, waiting for other resources (switching idleness). For coordination idleness, hierarchy will guarantee that when a resource's queues are sufficiently long there is no coordination idlenesses – the thresholds coordinate the transition of resources. We will also show that switching idleness is kept sufficiently small by choosing sufficiently large threshold coefficients K_i (in $S=K_i(1-\rho^{\rm BN})^{-1}$). Combined, these will guarantee that when a resource's workload is large, it decreases at rate that is proportional to $(1-\rho^{\rm BN})$. Such a drift sets the ground for the application of standard Lyapunov-based bounds. What follows is the formalization of the above.

For simplicity of notation, we do not superscript all processes by $\rho^{\rm BN}$ but the reader should keep in mind that all statement are made *for ρ^{BN} sufficiently close to 1* and that the statement “there exists a constant c” points to the existence of a constant that does not depend on $\rho^{\rm BN}$.

It is useful to recall that, in networks with hierarchical architectures, activities on an (acyclic) path from the top level to the bottom level share a resource and, moreover, that all of a resource's activities lie on a single such path; in Figure 12 the activities on the path $a1->a2->a4$ share resource 2. Similarly $a1->a2->a3$ correspond to resource 1. We refer to these as *resource paths.*

For a resource path p and given a state $x=(q,b)$, let $w_k(x)=\sum_{i:k\in\mathcal{R}(i)}m_iq_i$ be the workload of resource k. Define $\mathcal{W}_k(t)=w_k(Q(t))$ and introduce the scaled versions

$$\widehat{\mathcal{W}}_k(t)=(1-\rho^{\rm BN})\mathcal{W}_k(t(1-\rho^{\rm BN})^{-2}),\quad \widehat{w}_k(x)=(1-\rho^{\rm BN})w_k(x),\quad \bar{w}_k=\sum_{i:k\in\mathcal{R}(i)}\kappa_i m_i,$$

where κ_i, recall, is the threshold coefficient in $S_i=\kappa_i(1-\rho^{\rm BN})^{-1}$.

We assume that all resources are bottlenecks (or asymptotic bottlenecks, i.e., that $(1-\rho_k)=(1-\rho^{\rm BN})+o(1-\rho^{\rm BN})$). The general case works similarly but different scaling is required for different resources.

Central to our argument is establishing the existence of threshold coefficients $\kappa_1,\dots,\kappa_J$ as well as constants $t_0,\bar{K},\delta$ such that for all ρ^{BN} sufficiently close to 1 and each resource k

$$\sup_{x\in\mathcal{X}:\widehat{w}_k(x)>\bar{K}}\left(\mathbb{E}_x[\widehat{\mathcal{W}}_k(t_0)]-\widehat{w}_k(x)\right)\leq-\delta t_0. \tag{48}$$

This linear drift is the key. Our drift argument looks at the individual resources separately as if they are independent. They are connected through the switching and coordination idleness bounds that appear in the proof of (48). Nevertheless, the hierarchical structure facilitates is key in what follows.

First, we deduce a geometric drift from the linear one.

Lemma A.9 *There exist thresholds $\kappa_1,\dots,\kappa_J$ and constants $\gamma<1$ and $b,t_0,\theta>0$ such that for each resource k and all $x\in\mathcal{X}$*

$$\mathbb{E}_x\left[e^{\theta\widehat{\mathcal{W}}_k(t_0)}\right]\leq\gamma e^{\theta\widehat{w}_k(x)}+b. \tag{49}$$

Notice that (49) implies the existence of $\gamma<1$ and b such that

$$\sum_k\mathbb{E}_x\left[e^{\theta\widehat{\mathcal{W}}_k(t_0)}\right]\leq\gamma\left(\sum_k e^{\theta\widehat{w}_k(x)})\right)+b.$$

In particular, since the chain $\mathbb{X}(t)$ is non-explosive and (under our policy) irreducible, this guarantees that the CTMC is positive recurrent and has a steady-state distribution ν; see (Robert, 2003, Theorem 8.13).

Taking expectations with respect to ν on both sides of (49) we have

$$\mathbb{E}_\nu[e^{\theta\widehat{\mathcal{W}}_k(0)}]\leq\frac{b}{1-\gamma}, \tag{50}$$

and by Jensen's inequality that

$$(1-\rho^{\text{BN}})\mathbb{E}_\nu[\mathcal{W}_k(0)]=\mathbb{E}_\nu[\widehat{\mathcal{W}}_k(0)]\leq c,$$

for a constant c that does not depend on ρ^{BN}. Since we can repeat this argument for all resources, this proves (10).

We turn to prove (48). Each of the queues satisfies

$$Q_i(t)=Q_i(0)+A_i(t)-S_i(T_i^a(t))=Q_i(0)+\lambda_i t-\mu_i T_i^a(t)+M_i(t),$$

where $T_i^a(t) = (CT)_i(t)$, $M_i(t) = A_i(t) - \lambda_i t + \mu_i T_i^a(t) - S_i(T_i^a(t))$. In turn, for all $t \geq 0$,

$$\mathcal{W}_k(t) = \sum_{i:k\in\mathcal{R}(i)} m_i Q_i(t) = \mathcal{W}_k(0) + t(\sum_{i:k\in\mathcal{R}(i)} \lambda_i m_i - 1) + (t - \sum_{i:k\in\mathcal{R}(i)} T_i^a(t)) + \sum_{i:k\in\mathcal{R}(i)} m_i M_i(t).$$

Let $T_k^R(t) = \sum_{i:k\in\mathcal{R}(i)} T_i^a(t)$. We decompose $t - T_k^R(t)$ into two components: (i) idleness incurred at times when the total work exceeds $\bar{w}_k$ (so that at least one queue is above its threshold) and (ii) idleness incurred when $\widehat{\mathcal{W}}_k$ is below $\bar{w}_k$. Specifically, $T_k^R(t) = T_{k,S}(t) + I_k(t)$ where

$$T_{k,S}(t) = \int_0^t \mathbb{1}\{\widehat{\mathcal{W}}_k(s) > \bar{w}_k, dT_k^R(s) = 0\}ds, \text{ and } I_k(t) = \int_0^t \mathbb{1}\{\widehat{\mathcal{W}}_k(s) \leq \bar{w}_k, dT_k^R(s) = 0\}ds.$$

Thus,

$$\begin{aligned} W_k(t) &= w_k(Q(0)) + (\sum_{i:k\in\mathcal{R}(i)} \lambda_i m_i - 1)t + T_{k,S}(t) + I_k(t) + \sum_{i:k\in\mathcal{R}(i)} m_i M_i(t) \\ &= w_k(x) - (1 - \rho^{\mathrm{BN}})t + T_{k,S}(t) + I_k(t) + \sum_{i:k\in\mathcal{R}(i)} m_i M_i(t). \end{aligned}$$

Since M_i is a zero mean martingale we have

$$\mathbb{E}_x[\widehat{\mathcal{W}}_k(t)] - \widehat{w}_k(x) = -t + \mathbb{E}_x[\widehat{T}_{k,S}(t)] + \mathbb{E}_x[\widehat{I}_k(t)], \tag{51}$$

where

$$\widehat{T}_{k,S}(t) = (1 - \rho^{\mathrm{BN}})T_{k,S}(t(1 - \rho^{\mathrm{BN}})^{-2}) \text{ and } \widehat{I}_k(t) = (1 - \rho^{\mathrm{BN}})I_k(t(1 - \rho^{\mathrm{BN}})^{-2}).$$

The following lemma is instrumental to the proof. It is here that collaboration hierarchy is used. Consider again the network in Figure 12. When $W_k(s) > \bar{w}_k$, there is at least one queue that requires this resource and exceeds its threshold. Suppose it is activity 1. If not already working there, all resources $1, 2$ and 3 will, by our policy, switch to activity 1 as soon as they complete their current processing. Idleness incurred at such times corresponds to switching delays and the cumulative effect of these can be made small by choosing the thresholds to be sufficiently large.

Lemma A.10 *Given $\epsilon > 0$, there exists a choice of threshold coefficients κ_i and constants $t_0, \bar{K}$ such that*

$$\sup_{x\in\mathcal{X}} \mathbb{E}_x[\widehat{T}_{k,S}(t)] \leq \epsilon t, \tag{52}$$

for all $t \geq 0$ and

$$\sup_{x\in\mathcal{X}:\widehat{w}_k(x)>\bar{K}} \mathbb{E}_x[\widehat{I}_k(t_0)] \leq \epsilon t_0. \tag{53}$$

Plugging (52) and (53) into (51) implies, in particular, the existence of constants $t_0, \bar{K}$ (that do not depend on t_0) such that

$$\sup_{x\in\mathcal{X}:\widehat{w}_k(x)>\bar{K}} \Big(\mathbb{E}_x[\widehat{W}_k(t_0)] - \widehat{w}_k(x)\Big) \leq -(1-2\epsilon)t_0,$$

which proves (48) and concludes the argument for (10) in the statement of the theorem.

We turn to prove (11). The argument is similar in spirit to the previous one but, for each resource, looks only at its queues at levels $1,\ldots,\mathcal{I}_k - 1$. It excludes resource k's lowest-level activity.

We re-define let $w_k(x) = \sum_{i:k\in\mathcal{R}(i)} m_i q_i - m_{i_{k,\mathcal{I}_k}} q_{k,\mathcal{I}_k}$, and update all other definition accordingly.

We will prove that (this time with no time or space scaling) there exists $t_0, \widetilde{K}, c$ such that

$$\sup_{x\in\mathcal{X}:w_k(x)\geq \bar{w}_k(1-\rho^{\mathrm{BN}})^{-1}+\widetilde{K}} \mathbb{E}_x[(\mathcal{W}_k(t_0) - \bar{w}_k(1-\rho^{\mathrm{BN}})^{-1})^+] - (w_k(x) - \bar{w}_k(1-\rho^{\mathrm{BN}})^{-1})^+ \leq -c. \quad (54)$$

This will, again, imply a geometric drift.

Lemma A.11 *Fix resource k. There then exist constants $\gamma < 1$ and $\theta, b, t_0, \widetilde{K} > 0$ such that*

$$\mathbb{E}_x\left[e^{\theta(\mathcal{W}_k(t_0)-\bar{w}_k(1-\rho^{BN})^{-1})^+}\right] \leq \gamma e^{\theta(w_k(x)-\bar{w}_k(1-\rho^{BN})^{-1})^+} + b.$$

Consequently, under the stationary distribution

$$\mathbb{E}_\nu[e^{\theta(\mathcal{W}_k(0)-\bar{w}_k(1-\rho^{\mathrm{BN}})^{-1})^+}] \leq \frac{b}{1-\gamma}$$

This guarantees that

$$\mathbb{E}_\nu[\mathcal{W}_k(0)] \leq \frac{\bar{w}_k}{1-\rho^{\mathrm{BN}}} + c,$$

for some constant c and, in particular, that (11) holds.

It remains to establish (54). As before

$$\mathbb{E}_x[\mathcal{W}_k(t)] - w_k(x) \leq -c_k t + \mathbb{E}_x[T_{k,S}(t)] + \mathbb{E}[I_k(t)],$$

where $c_k = \sum_{i:k\in\mathcal{R}(i)} \lambda_i m_i - 1$. Since we do not include resource k's activity at level $\mathcal{I}_k$ we have that $\sum_{i:k\in\mathcal{R}(i)} \lambda_i m_i < \rho^{\mathrm{BN}} - \min_l \lambda_l m_l$. In turn, as ρ^{BN} approaches 1, $c_k \geq \min_l \lambda_l m_l$ remains bounds away from 0. The following analogue of Lemma A.10 concludes the argument. Here, in contrast to that lemma, $T_{k,S}$ and I_k are not scaled.

Lemma A.12 *Fix $\epsilon > 0$. There exist $t_0, \widetilde{K} > 0$ such that*

$$\sup_{x \in \mathcal{X}: w_k(x) \geq \bar{w}_k(1-\rho^{BN})^{-1} + \widetilde{K}} \mathbb{E}_x[T_{k,S}(t)] \leq \epsilon t,$$

for all $t \geq 0$, and

$$\sup_{x \in \mathcal{X}: w_k(x) \geq \bar{w}_k(1-\rho^{BN})^{-1} + \widetilde{K}} \mathbb{E}_x[I_k(t_0)] \leq \epsilon t_0.$$

■ .

A.4. Proofs of Auxiliary Lemmas

Proof of Lemma A.1. Under the condition that Q is non-explosive and $\mathcal{Q}U(x) \leq (c(1+U(x)))$ Dynkin's formula holds: for each x, and all $t \geq 0$,

$$\mathbb{E}_x[U(Q(t))] = U(x) + \mathbb{E}_x\left[\int_0^t \mathcal{Q}U(Q(s))ds\right];$$

see e.g. (Klebaner, 2005, Theorem 9.19). Provided that $\mathbb{E}_\pi[U(Q(0))] < \infty$ it also holds

$$\mathbb{E}_\pi[U(Q(t))] = \mathbb{E}_\pi\left[U(Q(0))\right] + \mathbb{E}_\pi\left[\int_0^t \mathcal{Q}U(Q(s))ds\right];$$

By stationarity $\mathbb{E}_\pi\left[U(Q(t))\right] = \mathbb{E}_\pi\left[U(Q(0))\right]$ for all $t \geq 0$ so that $\mathbb{E}_\pi\left[\int_0^t \mathcal{Q}U(Q(s))ds\right] = 0$ and, by the conditions of the lemma,

$$\begin{aligned} 0 = \mathbb{E}_\pi\left[\int_0^t \mathcal{Q}U(Q(s))ds\right] &= \mathbb{E}_\pi\left[\int_0^t \mathcal{Q}U(Q(s))\mathbb{1}\{Q(s)\notin A\}ds\right] + \mathbb{E}_\pi\left[\int_0^t \mathcal{Q}U(Q(s))\mathbb{1}\{Q(s)\in A\}ds\right] \\ &\leq -c_2\mathbb{E}_\pi\left[\int_0^t V(Q(s))ds\right] + c_3 t + \mathbb{E}_\pi\left[\int_0^t \mathcal{Q}U(Q(s))\mathbb{1}\{Q(s)\notin A\}ds\right]. \end{aligned}$$

In particular,

$$c_2\mathbb{E}_\pi\left[\int_0^t V(Q(s))ds\right] \leq c_3 t + \mathbb{E}_\pi\left[\int_0^t \mathcal{Q}U(Q(s))\mathbb{1}\{Q(s)\notin A\}ds\right].$$

Since $|\mathcal{Q}U(x)| \leq c_1(1+V(x))$ and U, V are π integrable we can interchange expectation and integration to get

$$\mathbb{E}_\pi[V(Q(0))] \leq \frac{c_3}{c_2} + \frac{c_1}{c_2}\mathbb{E}_\pi[(1+V(Q(0)))\mathbb{1}\{Q(0)\notin A\}].$$

■

Proof of Lemma A.2. Recall $w_2(x) = m_0q_0 + m_2q_2$. Let $w_2^{\rho^{\mathrm{BN}}}(x) = w_2(x) - (1-\rho^{\mathrm{BN}})^{-(1+\epsilon)}$ and $\mathcal{B}_0 = \{(q,b) \in \mathcal{X} : q_1 > 0, b_0 = 1\}$. The set $\mathcal{B}_0$ contains the states in which resource 2 works in activity 0 and let $\mathcal{B}_2 = \{(q,b) \in \mathcal{X} : q_2 > 0, b_2 = 1\}$ be those on which it works in activity 2. Then,

$$\begin{aligned}
&\widehat{\mathcal{Q}}^{\rho^{\mathrm{BN}}} V_{\rho^{\mathrm{BN}}}(x) = \\
&\quad \lambda_0(1-\rho^{\mathrm{BN}})^{-2}\left[e^{(1-\rho^{\mathrm{BN}})[w_2^{\rho^{\mathrm{BN}}}(x)+m_0]^+} - e^{(1-\rho^{\mathrm{BN}})[w_2^{\rho^{\mathrm{BN}}}(x)]^+}\right] + \\
&\quad \lambda_2(1-\rho^{\mathrm{BN}})^{-2}\left[e^{2(1-\rho^{\mathrm{BN}})[w_2^{\rho^{\mathrm{BN}}}(x)+m_2]^+} - e^{(1-\rho^{\mathrm{BN}})[w_2^{\rho^{\mathrm{BN}}}(x)]^+}\right] + \\
&\mu_0(1-\rho^{\mathrm{BN}})^{-2}\mathbb{1}\{x \in \mathcal{B}_0\}\left[e^{(1-\rho^{\mathrm{BN}})[w_2^{\rho^{\mathrm{BN}}}(x)-m_0]^+} - e^{(1-\rho^{\mathrm{BN}})[w_2^{\rho^{\mathrm{BN}}}(x)]^+}\right] + \\
&\mu_2(1-\rho^{\mathrm{BN}})^{-2}\mathbb{1}\{x \in \mathcal{B}_2\}\left[e^{(1-\rho^{\mathrm{BN}})[w_2^{\rho^{\mathrm{BN}}}(x)-m_2]^+} - e^{(1-\rho^{\mathrm{BN}})[w_2^{\rho^{\mathrm{BN}}}(x)]^+}\right].
\end{aligned} \tag{55}$$

Fix a constant $K \geq 3\max\{m_0, m_2\}$. Notice that if x is such that $w_2(x) \leq (1-\rho^{\mathrm{BN}})^{-(1+\epsilon)} + K$, then for all ρ^{BN} sufficiently close to 1,

$$\widehat{\mathcal{Q}}^{\rho^{\mathrm{BN}}} V_{\rho^{\mathrm{BN}}}(x) \leq b(1-\rho^{\mathrm{BN}})^{-1},$$

for some constant $b > 0$. For instance, as $\rho^{\mathrm{BN}} \uparrow 1$,

$$(1-\rho^{\mathrm{BN}})^{-2}\left[e^{(1-\rho^{\mathrm{BN}})[w_2^{\rho^{\mathrm{BN}}}(x)+m_0]^+} - e^{[w_2^{\rho^{\mathrm{BN}}}(x)]^+}\right] \leq (1-\rho^{\mathrm{BN}})^{-2}\left[\left(e^{(1-\rho^{\mathrm{BN}})K} - 1\right)\right] = \mathcal{O}(1-\rho^{\mathrm{BN}})^{-1},$$

and the argument is identical for the remaining summands in (55).

Otherwise, for all ρ^{BN} sufficiently close to 1 and $x \in \mathcal{E}^{\rho^{\mathrm{BN}}} := \{x : w_1(x) \leq (1-\rho^{\mathrm{BN}})^{-(1+\epsilon/2)}\}$ with $w_2(x) > (1-\rho^{\mathrm{BN}})^{-(1+\epsilon)} + K$ we have both $m_0q_0 \leq (1-\rho^{\mathrm{BN}})^{-(1+\epsilon/2)} << (1-\rho^{\mathrm{BN}})^{-(1+\epsilon)}$ and $m_0q_0 + m_2q_2 \geq (1-\rho^{\mathrm{BN}})^{-(1+\epsilon)} + K$ and it must be the case that $q_2 > 0$. In particular, resource 2 is working in one of the queues 0 or 2. For the point $x \in \mathcal{E}^{\rho^{\mathrm{BN}}}$ for which $w_2(x) > (1-\rho^{\mathrm{BN}})^{-(1+\epsilon)} + K$, $w_2^{\rho^{\mathrm{BN}}}(x) \geq \max\{m_0, m_2\}$ (so that the terms inside $[]^+$ are positive) and it can be verified that

$$\mathcal{Q}^{\rho^{\mathrm{BN}}} V_{\rho^{\mathrm{BN}}}(x) \leq -V_{\rho^{\mathrm{BN}}}(x) + b(1-\rho^{\mathrm{BN}})^{-1},$$

for a re-defined constant b. Finally, the fact that $|\mathcal{Q}^{\rho^{\mathrm{BN}}} V^2_{\rho^{\mathrm{BN}}}(x)| \leq c(1-\rho^{\mathrm{BN}})^{-1}(1 + V^2_{\rho^{\mathrm{BN}}}(x))$ follows from (55) using the fact that $e^x = 1 + x + o(1)$ as $x \to 0$. For example, $e^{(1-\rho^{\mathrm{BN}})w_2^{\rho^{\mathrm{BN}}}(x)+m_0} - e^{(1-\rho^{\mathrm{BN}})w_2^{\rho^{\mathrm{BN}}}(x)} = e^{(1-\rho^{\mathrm{BN}})w_2^{\rho^{\mathrm{BN}}}(x)}(e^{(1-\rho^{\mathrm{BN}})m_0} - 1) \approx (1-\rho^{\mathrm{BN}})m_0e^{(1-\rho^{\mathrm{BN}})w_2^{\rho^{\mathrm{BN}}}(x)}$. ■

Proof of Lemma A.3. The bound on the workload for resource 1 is trivial as this resource never idles as long as it has work. Viewed marginally, queues 0 and 1, follow a two-class $M/M/1$ queue. For the total workload $W_+(t)$, we take the Lyapunov function $V_{\rho^{\mathrm{BN}}}(x) := [x - (1-\rho^{\mathrm{BN}})^{-2}]^+$. Since above $(1-\rho^{\mathrm{BN}})^{-2}$ both resources work, it is easy to show that $\mathcal{Q}^{\rho^{\mathrm{BN}}} V_{\rho^{\mathrm{BN}}}(x) \leq -(1-\rho^{\mathrm{BN}})V_{\rho^{\mathrm{BN}}}(x) + b$

which gives the desired result as this implies that the chain is positive recurrent (see e.g. (Robert, 2003, Proposition 8.14)) and also, by a standard argument, that

$$\begin{aligned}0 \leq \mathbb{E}_x\left[V^2_{\rho^{\mathrm{BN}}}(Q^{\rho^{\mathrm{BN}}}(t))\right] &= V^2_{\rho^{\mathrm{BN}}}(x) + \mathbb{E}_x\left[\int_0^t \mathcal{Q}^{\rho^{\mathrm{BN}}} V^2_{\rho^{\mathrm{BN}}}(Q^{\rho^{\mathrm{BN}}}(s))ds\right]\\ &\leq V^2_{\rho^{\mathrm{BN}}}(x) - (1-\rho^{\mathrm{BN}})\mathbb{E}_x\left[\int_0^t V_{\rho^{\mathrm{BN}}}(Q^{\rho^{\mathrm{BN}}}(s))ds\right] + bt.\end{aligned}$$

For any given $x \in \mathcal{X}$, $V(x)/t \to 0$ as $t \to \infty$ we have that

$$\limsup_{t\to\infty} \frac{1}{t}\mathbb{E}_x\left[\int_0^t V_{\rho^{\mathrm{BN}}}(Q^{\rho^{\mathrm{BN}}}(s))ds\right] \leq b.$$

For each M, ergodicity guarantees that

$$\begin{aligned}\mathbb{E}_{\pi^{\rho^{\mathrm{BN}}}}[V_{\rho^{\mathrm{BN}}}(Q^{\rho^{\mathrm{BN}}}(0)) \wedge M] &= \limsup_{t\to\infty}\frac{1}{t}\mathbb{E}_x\left[\int_0^t V_{\rho^{\mathrm{BN}}}(Q^{\rho^{\mathrm{BN}}}(s)) \wedge M ds\right]\\ &\leq \limsup_{t\to\infty}\frac{1}{t}\mathbb{E}_x\left[\int_0^t V_{\rho^{\mathrm{BN}}}(Q^{\rho^{\mathrm{BN}}}(s))ds\right] \leq b,\end{aligned}$$

Taking M to infinity and applying the monotone convergence theorem gives the result. This argument can be repeated with higher moments by taking V^k for arbitrary integer k. ■

Proof of Lemma A.4. Take $x \in B^{\rho^{\mathrm{BN}}} := \{w_+(x) \leq (1-\rho^{\mathrm{BN}})^{-(1+2\epsilon)}\}$. For all ρ^{BN} sufficiently close to 1 and all $x \in B^{\rho^{\mathrm{BN}}}$, $w_+(x) \leq (1-\rho^{\mathrm{BN}})^{-2}$. In particular, resource 1 prioritizes queue 1 in these states so that

$$\begin{aligned}\mathcal{Q}^{\rho^{\mathrm{BN}}}U(x) &= \lambda_1[(m_1q_1+m_1)^2-(m_1q_1)^2] + \mu_1\mathbb{1}\{q_1>0\}[(m_1q_1-m_1)^2-(m_1q_1)^2]\\ &\leq 2(\lambda_1m_1-1)m_1q_1 + \rho_1^a m_1 + m_1\\ &= -2(1-\rho_1^a)V(x) + \rho_1^a m_1 + m_1.\end{aligned}$$

Furthermore, regardless of x, $Q^{\rho^{\mathrm{BN}}}U(x) \leq \lambda_1[(m_1q_1+m_1)^2-(m_1q_1)^2] \leq 2\lambda_1m_1^2q_1 + \lambda_1m_1^2$ and $Q^{\rho^{\mathrm{BN}}}U(x) \geq -2\mu_1m_1^2q_1 - 2\lambda_1m_1^2$ so that $|Q^{\rho^{\mathrm{BN}}}U(x)| \leq 2(\lambda_1+\mu_1)(m_1\vee 1)^2(1+m_1q_1) = 2(\lambda_1+\mu_1)(m_1\vee 1)(1+U(x))$. ■

Proof of Lemma A.5. There exists a Brownian motion $\mathcal{W}$ such that

$$\mathbb{P}\left\{\sup_{0\leq t\leq t^{\rho^{\mathrm{BN}}}}\left|\frac{A_0(t)-\lambda_0 t}{\sqrt{\lambda}} - \mathcal{W}(t)\right| \geq C_{21}\log(t^{\rho^{\mathrm{BN}}}) + x\right\} \leq C_{22}e^{-C_{23}x}$$

for all $x>0$ and fixed constants C_{21}, C_{22} and C_{23}; see Csörgo and Horváth (1996)[Theorem 2.2.1]. Fixing $\bar{c}$ and taking $x = \bar{c}(1-\rho^{\mathrm{BN}})^{-1}$ we have (with possibly re-defined constants)

$$\mathbb{P}\left\{\sup_{0\leq t\leq t^{\rho^{\mathrm{BN}}}}|A_0(t)-\lambda_0t-\sqrt{\lambda}\mathcal{W}(t)| \geq \bar{c}(1-\rho^{\mathrm{BN}})^{-1}\right\} \leq C_{22}e^{-C_{23}(1-\rho^{\mathrm{BN}})^{-1}}.$$

For ρ^{BN} sufficiently close to 1 (and some $\kappa > 0$)

$$\begin{aligned}\mathbb{P}\{(\mathcal{A}_{t^{\rho^{\mathrm{BN}}}})^c\} &\le C_{22}e^{-C_{23}(1-\rho^{\mathrm{BN}})^{-1}} + \mathbb{P}\left\{\sup_{0\le s\le t\le t^{\rho^{\mathrm{BN}}}:|t-s|\ge\delta\rho^{\mathrm{BN}}} \frac{|\mathcal{W}(t)-\mathcal{W}(s)|}{|t-s|^{1/2+\epsilon}} \ge 1+\epsilon\right\} \\ &\le C_{22}e^{-C_{23}(1-\rho^{\mathrm{BN}})^{-1}} + \mathbb{P}\left\{\sup_{0\le t\le t^{\rho^{\mathrm{BN}}}} |\mathcal{W}(t)| \ge \kappa((1-\rho^{\mathrm{BN}})^{-1})^{1/2+\epsilon}\right\}\end{aligned}$$

Notice that $t^{\rho^{\mathrm{BN}}}C_{22}e^{-C_{23}(1-\rho^{\mathrm{BN}})^{-1}} \to 0$ and it remains to show that

$$t^{\rho^{\mathrm{BN}}}\mathbb{P}\left\{\sup_{0\le t\le t^{\rho^{\mathrm{BN}}}} |\mathcal{W}(t)| \ge \kappa((1-\rho^{\mathrm{BN}})^{-1})^{1/2+\epsilon}\right\} \to 0.$$

By known fact for Brownian motion $\mathbb{P}\{\sup_{0\le s\le T}|\mathcal{W}(s)| > K\} \le e^{-\frac{K^2}{2T}}$, so that

$$\mathbb{P}\left\{\sup_{0\le s\le t^{\rho^{\mathrm{BN}}}} |\mathcal{W}(s)| > \kappa\right\} \le e^{-\frac{\kappa^2((1-\rho^{\mathrm{BN}})^{-1})^{1+2\epsilon}}{2t^{\rho^{\mathrm{BN}}}}} \le e^{-\tilde{c}(1-\rho^{\mathrm{BN}})^{-2\epsilon}}$$

for a constant $\tilde{c}$. This concludes the proof since for any $\epsilon > 0$, $(1-\rho^{\mathrm{BN}})^{-1}e^{-\tilde{c}(1-\rho^{\mathrm{BN}})^{-2\epsilon}} \to 0$ as $\rho^{\mathrm{BN}} \uparrow 1$. ■

Proof of Lemma A.6. Let $x = ((q_0,0,0),(1,0,0))$, $\bar{\tau}_0$ is the hitting time of 0 in an $M/M/1$ queue with parameters λ_0, m_0. It is known that $\mathbb{E}_x[\bar{\tau}_0] = q_0/(\mu_0-\lambda_0)$ and that $\mathbb{E}_x[(\bar{\tau}_0)^2] < \infty$; see e.g. (Robert, 2003, Chapter 5.3). Since $q_1 = q_2 = 0$ and there are no service completions in the individual tasks on $[0,\bar{\tau}_0)$ $Q_i(\bar{\tau}_0)$ is, conditional on $\bar{\tau}_0$, a Poisson random variable with parameter $\lambda_i\bar{\tau}_0$. thus, $\mathbb{E}_{q_0}[Q_i(\bar{\tau}_0)] = \lambda_i\mathbb{E}_{q_0}[\bar{\tau}_0] = \lambda_i q_0/(\mu_0-\lambda_0) = \tilde{q}\mu_i$. This proves (41).

From Robert (2003)[Proposition 5.5] is follows that $\bar{\tau}_0(q)/q \to \frac{1}{\mu_0-\lambda_0}$ in probability and from the central limit theorem and an application of the random time change theorem that, jointly,

$$\frac{1}{\sqrt{q_0}}\left(Q_i(\bar{\tau}_0) - \lambda_i q_0/(\mu_0-\lambda_0)\right) = \frac{1}{\sqrt{q_0}}\left(Q_i(q_0(\bar{\tau}_0/q_0)) - \lambda_i q_0/(\mu_0-\lambda_0)\right) \Rightarrow X_i.$$

The independence of X_1, X_2 follows from the independence of the arrival processes. The proof of (42) then builds on the following:

Consider two independent $M/M/1$ queues with arrival rate $\lambda_i < \mu_i$ for the i^{th} queue and with $\rho_1 = \rho_2 = \rho < 1$. Fix the sequence of initial conditions $(Q_1^q(0), Q_2^q(0)) = (q\mu_1 + X_1^q, q\mu_2 + X_2^q)$ where X_1^q, X_2^q satisfy $(\frac{X_1^q}{\sqrt{q}}, \frac{X_2^q}{\sqrt{q}}) \Rightarrow (\hat{X}_1, \hat{X}_2)$ and $\hat{X}_i$ are independent zero mean normal random variables distributed with standard deviation σ_i^x. Let

$$\tau_q = \inf\{t \ge 0 : (Q_1^q(t), Q_2^{q_0}(t)) = (0,0)\}.$$

Then, as $q \to \infty$,

$$\frac{\tau_q - \frac{q}{1-\rho}}{\sqrt{q}} \Rightarrow \max\{Y_1, Y_2\}, \tag{56}$$

where Y_i $i=1,2$ are independent zero-mean random variables with

$$Var(Y_i)=\sqrt{\sigma_B^2+\frac{\sigma_x^2}{\mu^2(1-\rho)^2}},$$

and

$$\sigma_B^2=\frac{1+\rho}{\mu^2(1-\rho)^3}$$

is the variance of an $M/M/1$ busy period (starting at 1 until hitting 0). If the sequence X_i^q is uniformly integrable, so is the sequence $\frac{\tau_q-\frac{q}{1-\rho}}{\sqrt{q}}$ and consequently

$$\mathbb{E}[\tau_q]=\frac{q}{1-\rho}+\sqrt{q}\mathbb{E}[\max\{Y_1,Y_2\}]+o(\sqrt{q})$$

Notice that, since $\mathbb{E}[\max Y_1,Y_2]>0$, this last result is exactly (42). Equation (43) then follows by combining (41) and (42). To conclude the proof it remains to prove (56). To that end, consider the individual-queue's hitting time

$$\tau_i(q_i)=\inf\{t\geq 0:Q_i(t)=0\},$$

where the argument q_i captures the dependence on the initial condition.

Note that

$$\tau_q=\max\{\tau_1(q+X_1^q),\tau_2(q+X_2^q)\}=\frac{q}{\mu-\lambda}+\max\left\{\tau_1(q+X_1^q)-\frac{q}{\mu-\lambda},\tau_2(q+X_2^q)-\frac{q}{\mu-\lambda}\right\},$$

and so,

$$\frac{\tau_q-\frac{q}{\mu-\lambda}}{\sqrt{q}}=\max\left\{\frac{\tau_1(q+X_1^q)-\frac{q}{\mu-\lambda}}{\sqrt{q}},\frac{\tau_2(q+X_2^q)-\frac{q}{\mu-\lambda}}{\sqrt{q}}\right\}. \tag{57}$$

We next apply the central limit theorem for the hitting time of each of the individual queues and the continuous mapping theorem using the continuity of the max operation. The weak convergence for each individual queue is known; see e.g. Robert (2003)[Proposition 5.5]. We prove it here so that we can use the same infrastructure for uniform integrability.

The hitting time of queue i to 0 is a sum of $q+X_i^q$ $M/M/1$ busy periods (starting at 1 and reaching 0) – notice that the time going from q to $q-1$ is identically distributed as that from $q-1$ to $q-2$ and from 1 to 0. In other words

$$\tau_i(q_i)=\sum_{i=1}^{q}Z_i$$

where Z_i has the distribution of the $M/M/1$ busy period and, in particular,

$$\mathbb{E}[Z_i]=\frac{1}{\mu-\lambda}\text{ and }Var(Z_i)=\frac{1+\rho}{\mu^2(1-\rho)^3}=:\sigma_B^2.$$

The strong law of large numbers gives

$$\frac{\tau_i(q+X_i^q)}{q} \to \frac{1}{\mu-\lambda},$$

and by the central limit theorem

$$\frac{\sum_{i=1}^{q+X_i^q} Z_i - \frac{q+X_i^q}{\mu-\lambda}}{\sqrt{q}} \Rightarrow \hat{\tau}, \text{ as } q \to \infty,$$

where $\mathbb{E}[\hat{\tau}]=0$ and $Var(\hat{\tau})=\sigma_B^2$. Further, by the assumption on X_i^q

$$\frac{\sum_{i=1}^{q+X_i^q} Z_i - \frac{q}{\mu-\lambda}}{\sqrt{q}} \Rightarrow \hat{\tau} - \frac{\hat{X}}{\mu-\lambda} := Y.$$

In particular,

$$Var(Y) = \sigma_B^2 + \frac{\sigma_x^2}{\mu^2(1-\rho)^2}.$$

Plugging this into (57) we have that

$$\max\{\frac{\tau_1(q+X_1^q) - \frac{q}{\mu-\lambda}}{\sqrt{q}}, \frac{\tau_2(q+X_2^q) - \frac{q}{\mu-\lambda}}{q}\} \Rightarrow \max\{Y_1, Y_2\},$$

which concludes the convergence argument. It only remain to establish uniform integrability but this is immediate from the assumptions as then (by independence of the Z_i and their independence from X_i^q)

$$\begin{aligned}\mathbb{E}\left[\left(\frac{\sum_{i=1}^{q+X_i^q}(Z_i - \frac{1}{\mu-\lambda})}{\sqrt{q}}\right)^2\right] &\leq \mathbb{E}\left[\left(\frac{\sum_{i=1}^{q+|X_i^q|}(Z_i - \frac{1}{\mu-\lambda})}{\sqrt{q}}\right)^2\right] \\ &= \mathbb{E}\left[\frac{1}{q}\sum_{i=1}^{q+|X_i^q|}\left(Z_i - \frac{1}{\mu-\lambda}\right)^2\right] = \frac{q+\mathbb{E}|X_i^q|}{q}\mathbb{E}[Z^2].\end{aligned}$$

By assumption $\mathbb{E}[Z^2]=\sigma_B^2<\infty$ and since (again, by assumption) $\limsup_q \mathbb{E}[|X_i^q|]/\sqrt{q}<\infty$ we have that the right hand side is bounded uniformly in q and the uniform integrability of each of the sequences follows. Finally, we use the fact that

$$\left(\max\left\{\frac{\tau_1(q+X_1^q) - \frac{q}{\mu-\lambda}}{\sqrt{q}}, \frac{\tau_2(q+X_2^q) - \frac{q}{\mu-\lambda}}{\sqrt{q}}\right\}\right)^2 \leq \left(\frac{\tau_1(q+X_1^q) - \frac{q}{\mu-\lambda}}{\sqrt{q}}\right)^2 + \left(\frac{\tau_1(q+X_2^q) - \frac{q}{\mu-\lambda}}{\sqrt{q}}\right)^2.$$

■

Proof of Lemma A.7. By Dynkin's formula, for each $t>0$,

$$\mathbb{E}_x[Q_0^2(\widehat{\tau}_1 \wedge t)] = \mathbb{E}_x[Q_0^2(0)] + 2\mathbb{E}_x\left[\int_0^{\widehat{\tau}_1 \wedge t}(\lambda_0 - \mu_0 \mathbb{1}\{Q_0(s)>0\})Q_0(s)ds\right]$$

Using the fact that $\mathbb{E}_x[Q_0(\widehat{\tau}_1)] \leq \mathbb{E}_x[q_0 + A_0(\widehat{\tau}_1)] \leq q_0 + \lambda_0\mathbb{E}_x[\widehat{\tau}_1] < \infty$ and that $\mathbb{E}_x[\int_0^{\widehat{\tau}_1 \wedge t} Q_0(s)ds] \leq \mathbb{E}_x[q_0\widehat{\tau}_1 + A_0(\widehat{\tau}_1)\widehat{\tau}_1] \leq q_0\mathbb{E}_x[\widehat{\tau}_1] + \lambda_0\mathbb{E}_x[(\widehat{\tau}_1)^2] < \infty$ we can apply the dominated convergence theorem and take $t\to\infty$ to conclude that $\mathbb{E}_x[(Q_0(\widehat{\tau}_1))^2] = \mathbb{E}_x[Q_0^2(0)] - 2\mathbb{E}_x\left[\int_0^{\bar{\tau}_0}(\lambda_0 - \mu_0\mathbb{1}\{Q_0(s)>0\})Q_0(s)ds\right]$. ■

Proof of Lemma A.8. We first show that the conditions of the lemma guarantee that if one initializes X in state x_i and Y_i in state x_i, then for all $l=1,\dots,$

$$X_i^l \leq_{st} Y_i^l, \tag{58}$$

where $\leq_{st}$ is standard stochastic ordering. We can argue this by induction. For $l=1$, it holds by assumption that $\mathbb{P}_x\{X_i^1 \geq y\} \geq \mathbb{P}_{x_i}\{Y_i^1\}$. The following argument is an adaptation of Derman and Ignall (1975).

Suppose the result holds for all $l=1,\dots,n-1$.

$$\begin{aligned}
\mathbb{P}_x\{X_i^n \geq y\} - \mathbb{P}_{x_i}\{Y_i^n \geq y\} &\leq \sum_{z=0^\infty} \left(\mathbb{P}_x\{X_i^{n-1}=z\} \sup_{x:x_i=z} \mathbb{P}_x\{X_i^1 \geq y\} - \mathbb{P}^i_{x_i}\{Y_i^{n-1}=z\}\mathbb{P}^i_z\{Y_i^1 > y\}\right) \\
&\leq \sum_{z=0}^{\infty} \mathbb{P}^i_z\{Y_i^1 > y\} \left(\mathbb{P}_x\{X_i^{n-1}=z\} - \mathbb{P}^i_{x_i}\{Y_i^{n-1}=z\}\right) \\
&= \sum_{r=0}^{\infty}\sum_{z\geq r} \left(\mathbb{P}_x\{X_i^{n-1}=z\} - \mathbb{P}^i_{x_i}\{Y_i^{n-1}=z\}\right)\left(\mathbb{P}^i_r\{Y_i^1>y\} - \mathbb{P}^i_{r-1}\{Y_i^1>y\}\right) \\
&= \sum_{r=0}^{\infty} \left(\mathbb{P}_x\{X_i^{n-1}\geq r\} - \mathbb{P}^i_{x_i}\{Y_i^{n-1}\geq r\}\right)\left(\mathbb{P}^i_r\{Y_i^1>y\} - \mathbb{P}^i_{r-1}\{Y_i^1>y\}\right),
\end{aligned}$$

where the second inequality follows from the assumption on the one step transition probability, the first equality is based on telescoping sums (and we take $P^i_{-1}\equiv 0$). In the last row, $\mathbb{P}_x\{X_i^{n-1}\geq r\} - \mathbb{P}^i_{x_i}\{Y_i^{n-1}\geq r\} \geq 0$ by the induction assumption and $\mathbb{P}^i_r\{Y_i^1>y\} - \mathbb{P}^i_{r-1}\{Y_i^1>y\} \geq 0$ by the assumed monotonicity of Y_i. This concludes the induction argument and establishes (58).

Since Y_i^l is assumed to be positive recurrent and aperiodic we have that Y_i^l is a tight sequence (that converges to the steady-state distribution of Y_i). By the stochastic ordering X_i^l is also a tight sequence for each i and so is, in turn, the sum of the coordinates $X_+^l = \sum_{i=1}^J X_i^l$. We conclude that the chain X^l is tight. Since it is irreducible and aperiodic, it is also positive recurrent; see (Asmussen, 2003, Proposition I.4.1). ■

Proof of Lemma A.9. It suffices to argue that, given t_0 there exists $\theta,\bar{\theta}>0$ such that for all ρ^{BN} sufficiently close to 1.

$$\sup_{x\in\mathcal{X}} \mathbb{E}_x[(\widehat{W}_k(t_0)-\widehat{w}_k(x))^2 e^{\theta(\widehat{W}_k(t_0)-\widehat{w}_k(x))^+}] \leq \bar{\theta}, \tag{59}$$

and

$$\sup_{x\in\mathcal{X}} \mathbb{E}_x[e^{\theta(\widehat{W}_k(t_0)-\widehat{w}_k(x))}] \leq \bar{\theta}. \tag{60}$$

From (Gamarnik and Zeevi, 2006, Theorem 6) (see also the first display in the proof of Theorem 5 there) it follows that these bounds together with the linear drift (48), guarantee (49).

To establish the bounds (59) and (60), recall that

$$\widehat{W}_k(t) = \hat{w}_k(x) + (1-\rho^{\mathrm{BN}})\left(\sum_{i:k\in\mathcal{R}(i)} \rho_i^a - 1\right)(1-\rho^{\mathrm{BN}})^{-2}t + \widehat{T}_{k,S}(t) + \widehat{I}_k(t) + \widehat{M}_k(t)$$

where $\widehat{M}_k(t) = \sum_{i:k\in\mathcal{R}(i)} m_i \widehat{M}_i(t)$. Recall also that, by definition, $\widehat{I}_k(t)$ does not increase when $\widehat{W}_k > \bar{w}_k$. In particular, using (Ghamami and Ward, 2013, Lemma 8.3), we have the existence of a constant $\bar{c}$ such that

$$\begin{aligned}\sup_{0\le u\le s\le t} |\widehat{W}_k(s) - \widehat{W}_k(u)| &\le \bar{w}_k + c(1-\rho^{\mathrm{BN}})\left(\sum_{i:k\in\mathcal{R}(i)} \rho_i^a - 1\right)(1-\rho^{\mathrm{BN}})^{-2}t \\ &\quad + \bar{c}\left(\sup_{0\le u\le s\le t} |\widehat{T}_{k,S}(s) - \widehat{T}_{k,S}(u)| + \sup_{0\le u\le s\le t} |\widehat{M}_k(s) - \widehat{M}_k(u)|\right) \\ &\le \left(\sum_{i:k\in\mathcal{R}(i)} \rho_i^a - 1\right)(1-\rho^{\mathrm{BN}})^{-2}t + 2\bar{c}\widehat{T}_{k,S}(t) + 2\bar{c}\sup_{0\le s\le t} |\widehat{M}_k(s)|,\end{aligned}$$

for all $t \ge 0$. For the second inequality we used the fact that $\widehat{T}_{k,S}(t)$ is increasing in t. Exponential bounds for the Poissonian martingale are standard and those for the idleness terms follow from Lemma A.10 proved below. ■

Proof of Lemma A.10. We provide only the essential ingredients of this proof with some standard details being omitted.

Notice that, because of hierarchy, $dT_k^R(t) = 0$ and $\widehat{\mathcal{W}}_k(t) > \bar{w}_k = \sum_{i:k\in\mathcal{R}(i)} \kappa_i m_i$ can hold simultaneously only when resource k is waiting for the other resources to switch to an activity that exceeded its threshold. Hierarchy guarantees that these other resources switch as soon as they complete their current processing.

In particular, $T_{k,S}(t) \le \sum_{i=1}^{N_u(t)} X_i$ where $N_u(t)$ is the number of up-switches and X_i is the switching time. We say that an up switch occurs at t if there exists $i : k \in \mathcal{R}(i)$ such that $T_i(t-) = 1$, $T_i(t) = 0$, and there is an activity j with $k \in \mathcal{R}(j)$ at a higher collaboration level than i, that exceeds its threshold. At time t, resource k either starts processing at the highest level activity that exceeds its threshold or is idling (waiting for other resources). Each X_i is stochastically smaller than a maximum of J exponentials with means $m_1, \ldots, m_J$ and, in particular, $\mathbb{E}[X_i] \le \sum_j m_j$.

Let d be the number of levels in the collaboration graph. We say that a down switch occurs at time t if there exists $i : k \in \mathcal{R}(i)$ such that $T_i(t-) = 1$, $T_i(t) = 0, Q_i(t) = 0$ and all the queues at a higher level are below their threshold. At time t, resource k either starts processing at a lower level activity or idles. Notice that if t is such that $N_u(t) \ge 2d$ there must be at least one down-switch by time t. Thus, $N_u(t) \le 2dN_d(t)$ where $N_d(t)$ is the number of down-switches by time t. The policy dictates that down-switches occur only when a queue is drained starting at its threshold or above (how much above depends on the time it took the resources to switch). Let Z_i be the random variable for the time it takes to drain queue i starting at its threshold. Then,

$$\mathbb{E}[Z_i] = \frac{K_i(1-\rho^{\mathrm{BN}})^{-1}}{\mu_i - \lambda_i}.$$

For a sufficiently large constant b, we have that $\mathbb{E}[Z_i \wedge b] \geq \frac{K_i}{2}(1-\rho^{\mathrm{BN}})^{-1}/(\mu_i-\lambda_i)$, so that applying standard renewal argument we have that

$$\mathbb{E}_x[N_d(t)] \leq t \min_i \left\{ \frac{2}{K_i}(\mu_i - \lambda_i)(1-\rho^{\mathrm{BN}}) \right\}.$$

Taking K_i such that $\frac{2}{K_i}(\mu_i - \lambda_i) \leq \epsilon / \sum_j m_j$ we have that $\mathbb{E}_x[N_d(t)] \leq \frac{\epsilon}{\sum_j m_j}(1-\rho^{\mathrm{BN}})t$ so that

$$\mathbb{E}_x[\widehat{T}_{k,S}(t)] = (1-\rho^{\mathrm{BN}})\mathbb{E}_x[T_{k,S}(t(1-\rho^{\mathrm{BN}})^{-2})] \leq 2d(1-\rho^{\mathrm{BN}})\mathbb{E}[N_d(t(1-\rho^{\mathrm{BN}})^{-2})]\mathbb{E}[X_i] \leq \epsilon t.$$

A simple extension of the above gives a bound on the exponential moment of $\widehat{T}_{k,S}$, i.e., that given $t_0 > 0$ there exist $\theta, \bar{\theta} > 0$ such that

$$\sup_{x \in \mathcal{X}: \widehat{w}_k(x) > \bar{K}} \mathbb{E}_x[e^{\theta \widehat{T}_{k,S}(t_0)}] \leq \bar{\theta}.$$

Finally, notice that $\widehat{I}_k(t) \leq \int_0^t \mathbb{1}\{\widehat{W}_k(s) \leq \bar{w}_k\} ds$. The path workload $\widehat{W}_k(t)$ is bounded stochastically from below by the workload in a multiclass $M/M/1$ queue (with the single server being a focal resource of the path). For the latter, $\mathbb{E}_q[\int_0^t \mathbb{1}\{\widehat{W}_{MM1}(s) \leq \bar{w}_k\} ds] \leq t\mathbb{P}_q\{\tau_{\bar{w}} \leq t\}$ where $\tau_{\bar{w}}$ is the hitting time of the scaled $M/M/1$ workload $\widehat{W}_{MM1}$ to $\bar{w}$ starting at q. It is a standard argument that $\mathbb{P}_q\{\tau_{\bar{w}} \leq t\} \leq \epsilon t$ for $q = \bar{w} + \bar{K}$ with sufficiently large $\bar{K}$. ■

Proofs of Lemmas A.11 and A.12. These are simpler versions of the proofs of Lemmas A.9 and A.10. We omit the details. ■

A.5. Expanded Material

This last section is not necessary for the reading of the paper or for the support of its results. It contains elaborated versions of content that is included in the paper and that could help the reader interested in a more detailed exposition of some parts of the paper.

Remarks on Maximum Pressure policies Dai and Lin (2005) prove that resource-splitting, preemptive maximum pressure policies, introduced by Tassiulas and Ephremides (1992), maximize throughput in open networks (here we use the term to refer to networks where each activity has a single buffer and routing is probabilistic). Dai and Lin (2005, Theorems 5 and 6) yields that a non-resource-splitting preemptive maximum pressure policy achieves the theoretical throughput and stabilizes the network for each $\rho^{\mathrm{BN}} < 1$. In the parallel network that we study here it also achieves optimal scaling.

Theorem 8 (Maximum pressure in hierarchical parallel networks) *Consider a parallel network with J activities and a hierarchical collaboration architecture. Then, the preemptive maximum pressure policy stabilizes the network (denote the stationary distribution by π) and achieves the optimal scaling:*

$$\mathbb{E}_\pi\left[\sum_i Q_i(0)\right] \le \frac{J\sum_i(\lambda_i+\mu_i)}{(1-\rho^{BN})} \tag{61}$$

Proof of Theorem 8. We use a standard Lyapunov argument. We prove that when queues are sufficiently large there is a "down drift" of the order of $(1-\rho^{\mathrm{BN}})$. With preemptive maximum pressure, Poisson arrivals, and exponential service times, the process $Q(t)=(Q_1(t),\dots,Q_{\mathbf{J}}(t))$ is a continuous time Markov chain. Since the transition rates are bounded by $\sum_i\lambda_i+\sum_i\mu_i$, this chain is non-explosive. It is also easily verified to be irreducible.

Let $\mathcal{Q}$ be the generator of the chain $Q(t)$. Let $g(q)=\sum_i q_i^2$. We will show that

$$\mathcal{Q}g \le -\frac{1}{J}(1-\rho^{\mathrm{BN}})|q| + \sum_i(\lambda_i+\mu_i). \tag{62}$$

This implies that the chain is positive recurrent with a steady-state distribution π. Being non-explosive the chain also satisfies

$$\mathbb{E}_\pi[|Q(0)|] \le \frac{J\sum_i(\lambda_i+\mu_i)}{\gamma(1-\rho^{\mathrm{BN}})};$$

see (Glynn and Zeevi, 2008, Corollary 1). This is consistent with the grounding of maximum pressure in quadratic Lyapunov functions; see Dai and Lin (2005).

Let $x_i(q)$ be maximum pressure's allocation to activity i when $Q(t)=q$. Then,

$$\begin{aligned}\mathcal{Q}g(q) &= \sum_i \lambda_i\left((q_i+1)^2-q_i^2\right) + \sum_i \mu_i x_i(q)\mathbb{1}\{q_i>0\}\left((q_i-1)^2-q_i\right)\\ &= 2\sum_i\lambda_i q_i - 2\sum_i \mu_i q_i x_i(q) + \sum_i \mu_i x_i(q)\mathbb{1}\{q_i>0\} + \sum_i \lambda_i\\ &\le 2\sum_i\lambda_i q_i - 2\sum_i \mu_i q_i x_i(q) + \sum_i(\lambda_i+\mu_i).\end{aligned}$$

With $p(q,x(q))=2\sum_i\mu_i q_i x_i(q) - 2\sum_i\lambda_i q_i$ (this is the pressure at state q), the above is re-written as

$$\mathcal{Q}g(x) \le -p(q,x(q)) + \sum_i(\lambda_i+\mu_i).$$

To prove (62) it remains to establish that

$$p(q,(x(q))) \le -\frac{1}{J}(1-\rho^{\mathrm{BN}})|q|. \tag{63}$$

By definition

$$p(q,x(q)) \in \arg\max_{x\geq 0}\{p(q,x) : Ax \leq 1\}, \tag{64}$$

and define the "allocation"

$$\bar{x}_i = \frac{\lambda_i}{\mu_i} + \min_{k\in\mathcal{R}(\{i\})}(1-\rho_k)/J.$$

Since $(1-\rho^{\mathrm{BN}}) \leq (1-\rho_k)$ for any resource k, we have that

$$p(q,x(q)) \geq \sum_i q_i(\mu_i\bar{x}_i - \lambda_i) = \sum_i q_i \min_{k\in\mathcal{R}(\{i\})}(1-\rho_k)/J \geq \sum_i q_i(1-\rho^{\mathrm{BN}})/J = |q|\frac{1}{J}|(1-\rho^{\mathrm{BN}}).$$

Since

$$\sum_i A_{ki}\bar{x}_i \leq \sum_{i:k\in\mathcal{R}(\{i\})} \lambda_i/\mu_i + (1-\rho^{\mathrm{BN}}) \leq \rho_k + (1-\rho^{\mathrm{BN}}) \leq 1,$$

$\bar{x}$ is feasible for (64), and we conclude that (63) holds with $\gamma = 1/J$. ■

Dai and Lin (2005) cover non-preemptive policies for the case of reversed Leontief networks (see Definition 5 there), where each activity (task) requires a single resource. This, in particular, rules out collaborative networks. Section 8 there includes a nice illustration of what could go wrong when an activity requires multiple resources.

Figure 16 demonstrates the capacity loss due to increasing switching idleness as throughput increases. In the symmetric base network, the non-preemptive maximum pressure policy becomes unstable around $\rho \simeq 0.86$ because switching idleness then consumes the entire "idleness budget" $1-\rho$. The policy thus no longer maximizes throughput, let alone achieves optimal scaling. In fact, there exist no general results for non-preemptive maximum pressure policies in networks where resources process multiple activities.

Achievable region of two-class $M/M/1$ queue. Here we use Dacre et al. (1999) to develop the achievable region of R1's two-class benchmark queue shown in Figure 3(LEFT). A pair (w_0,w_1) is in the achievable region $\mathcal{A}$ if there exists a control policy under which the expected waiting time $\mathbb{E}[W_i]$ in queue i equals w_i. The minimal delay for queue $i\in\{0,1\}$ stems from giving that queue static, preemptive priority and is given by a simple single-class $M/M/1$ delay expression:

$$w_i \geq \frac{\rho_i^a m_i}{1-\rho_i^a}.$$

The third boundary of $\mathcal{A}$ stems from the fact that, in steady state, the expected workload in the system is invariant among work-conserving policies. Consider serving the two-class arrival process in FIFO in steady state. This is a single-class $M/G/1$ queue for which the expected workload equals the expected delay, which is given by:

$$\mathbb{E}[W] = \frac{\mathbb{E}[R]}{1-\rho},$$

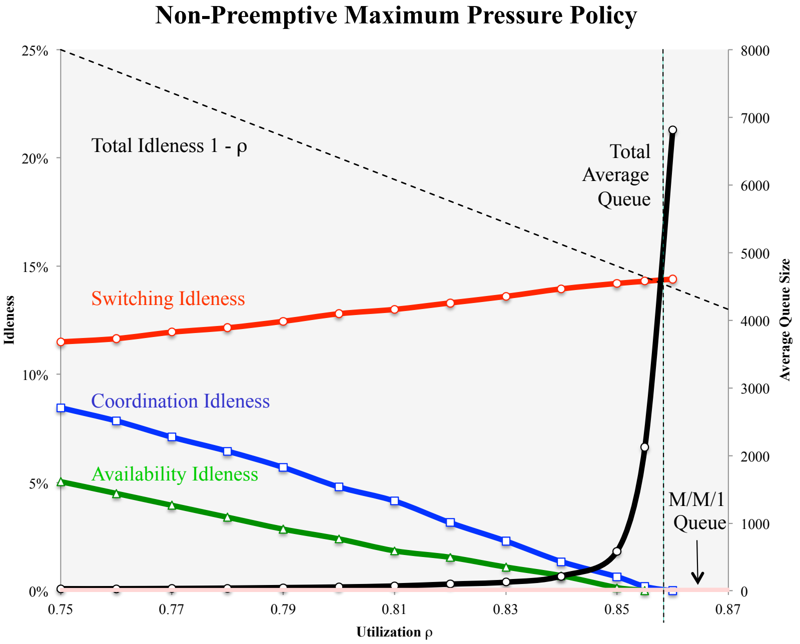

Figure 16 **The switching idleness of the non-preemptive maximum pressure policy in the symmetric base network leads to capacity losses: the maximal throughput (equal to ρ given $m_i = 0.5$) is about 0.86.**

where R is the remaining service time. In steady state, the job in service is class i with probability λ_i/μ_i and the system is idle otherwise. Given two independent exponential service times which are memoryless, we have that

$$\mathbb{E}[R] = (\lambda_0/\mu_0)\cdot(1/\mu_0) + (\lambda_1/\mu_1)\cdot(1/\mu_1) + (1-\lambda_0/\mu_0-\lambda_1/\mu_1)\cdot 0.$$

If $\mathbb{E}[W_i] = w_i$, Little's law yields that the expected queue-i count is $\lambda_i w_i$ with expected workload $\lambda_i w_i/\mu_i$. Summing over both classes yields total workload and the third boundary equation of $\mathcal{A}$.

Dynamics of network workload $W^{\mathbf{net}}$ A control rule specifies which configuration vector is active at each time t. Let $T(t)$ be the cumulative allocation (vector) process. Its component $T_c(t)$ is the cumulative amount of time a configuration vector c is active during $[0,t]$:

$$T_c(t) = \int_0^t \mathbb{1}\Big\{\text{configuration vector } c \text{ is active at time } \tau\Big\} d\tau,$$

where $\mathbb{1}\{\cdot\}$ denotes the indicator function. Given that only one configuration can be active at any point in time, the total time the network is processing during $[0,t]$ is $\sum_c T_c(t) = e'T(t) \le t$. The total time the network is idle (and no configuration is active) during $[0,t]$ is $I^{\text{net}}(t) = t - e'T(t)$.

To translate queue workload into network workload and activity load into network utilization we first consider a long-run average time allocation vector π and use the configuration matrix C. For a constant time-allocation vector π, the total amount of time activity i is processed equals $(C\pi)_i$.

We define the *network workload* $W^{\text{net}}(Q)$ as the minimal expected time needed to process the work embodied in the queue-length vector Q:

$$\begin{aligned} W^{\text{net}}(Q) = \min_{\pi \in \mathbb{R}_+^J} \; & e'\pi \\ \text{s.t.} \; & C\pi \geq m * Q. \end{aligned} \tag{65}$$

For the BC and BC+ networks of Figure 1 this yields the intuitive expressions:

$$W^{\text{net}}_{\text{BC}}(Q) = m_0 Q_0 + \max\{m_1 Q_1, m_2 Q_2\} \;\text{ and }\; W^{\text{net}}_{\text{BC+}}(Q) = m_0 Q_0 + m_1 Q_1 + m_2 Q_2. \tag{66}$$

It is instructive to consider the dual linear program to (65):

$$\begin{aligned} \max_{y \in \mathbb{R}_+^{\mathbf{I}}} \; & y'(m * Q) \\ \text{s.t.} \; & y'C \leq e'. \end{aligned} \tag{67}$$

Letting $y^\star(Q)$ denote the optimal solution, strong duality then gives us a simple decomposition of network workload directly in terms of queue workload:

$$W^{\text{net}}(Q) = \sum_j y_j^\star(Q) m_j Q_j. \tag{68}$$

There are two extreme deconstructions depending on A or C: When A is the identity matrix (networks without collaboration or multitasking), then $W^{\text{net}}(Q) = \max_j(m_j Q_j)$ and $y^\star$ is a piecewise linear function of Q. When C is the identity matrix (as in the BC+ networks), then $W^{\text{net}}(Q) = \sum_j(m_j Q_j)$ and $y^\star = e$ is the constant vector of ones, indicating that resources perform as a single resource. In general, the dual variables of the instantaneous network workload depend on Q.

A long run average view simplifies these dual variables. Consider a network with initial queue vector $Q(0)$ and arrival and service processes $A(t)$ and $S(t)$, which denote the number of arrivals and service completions, respectively, during $[0,t]$. Recall that $T_c(t)$ is the cumulative amount of time configuration vector c is active during $[0,t]$ so that:

$$Q(t) = Q(0) + A(t) - S \circ CT(t), \tag{69}$$

where $S \circ CT(t)$ is the vector with components $S_j((CT)_j(t))$. Dividing both sides by t, letting $t \to \infty$ and recalling the strong law that $A_j(t)/t \to \lambda_j$ and $S_j(t)/t \to 1/m_j$, then in a stable network, we must have

$$C\pi = m * \lambda \;\text{ and }\; \pi = \lim_{t\to\infty} \frac{T(t)}{t}. \tag{70}$$

As shown earlier, the optimal time-allocation rates $\bar{\pi}$ satisfy (**??**). If $\bar{\pi}_k > 0$, we say that configuration vector k is optimal; otherwise it is a suboptimal configuration. We label the configurations to allow the following block partitioning

$$\bar{\pi} = [\pi^\star, 0] \text{ and } C = [C^\star, C^0], \tag{71}$$

where $\pi^\star > 0$ and $C^\star$ contains all optimal configuration vectors.

Let $\bar{y} = y^\star(\lambda)$ denote the solution to (67) when $Q = \lambda$ (this is the dual of SPP). By strong duality $\sum_j \bar{y}_j m_j \lambda_j = \rho^{\text{net}}$ and complementary slackness further yields that $\bar{y}' C_{\cdot,k} = 1$ if the k^{th} column is an optimal configuration, meaning $\bar{\pi}_k > 0$. In matrix notation, recalling (71), we have

$$\bar{y}' C^\star = e'. \tag{72}$$

Let the 0 configuration mean that no resource is processing (the network idles). Then,

$$\bar{y}' C \bar{\pi} = \bar{y}' C^\star \pi^\star = e' \pi^\star = 1 - \bar{\pi}_0, \tag{73}$$

where $\bar{\pi}_0$ is the optimal idleness rate (the fraction of time the zero configuration is used). The dual variables $\bar{y}$ provide us with a simplification: it can be shown that in heavy traffic as $\rho^{\text{net}} \uparrow 1$, then

$$W^{\text{net}}(Q^{\rho^{\text{net}}}(t)) = \sum_j y_j^\star(Q^{\rho^{\text{net}}}(t)) m_j Q_j^{\rho^{\text{net}}}(t) \approx \sum_j \bar{y}_j m_j Q_j^{\rho^{\text{net}}}(t).$$

Henceforth, the network workload process refers to

$$W^{\text{net}}(t) = \sum_j \bar{y}_j m_j Q_j(t). \tag{74}$$

Essentially, $W^{\text{net}}(t)$ is the expected time the network needs to process the queue vector $Q(t)$ while using the steady-state optimal time-allocations $\bar{\pi}$. Combining (69) and (74) we have that

$$\begin{aligned} W^{\text{net}}(t) &= W^{\text{net}}(0) + \sum_j \bar{y}_j m_j A_j(t) - \sum_j \bar{y}_j m_j S_j((CT)_j(t)) \\ &= W^{\text{net}}(0) + \sum_j \bar{y}_j m_j \lambda_j t - \bar{y}' C T(t) + M(t), \end{aligned}$$

where M is the deviation from the fluid (first moment) processes:

$$M(t) = \sum_j \bar{y}_j m_j \Big(A_j(t) - \lambda_j t \Big) + \sum_j \bar{y}_j m_j \Big(\mu_j (CT)_j(t) - S_j((CT)_j(t)) \Big).$$

Similar to (71), block partition the control vector as

$$T(t) = [T^\star(t); T^0(t)] \tag{75}$$

where $T^\star$ denotes the allocation to all optimal configurations and T^0 to the suboptimal configurations. Complementary slackness then yields

$$\bar{y}' C T(t) = e' T^\star(t) + \bar{y}' C^0 T^0(t).$$

Recall that $e' T(t) = t - I^{\text{net}}(t)$ so that

$$\bar{y}' C T(t) = t - I^{\text{net}}(t) - (e' - \bar{y}' C^0) T^0(t).$$

Strong duality yields that $\sum_j \bar{y}_j m_j \lambda_j = \rho^{\text{net}}$ so that we arrive at

$$W^{\text{net}}(t) = W^{\text{net}}(0) - (1 - \rho^{\text{net}}) t + (e' - \bar{y}' C^0) T^0(t) + M(t) + I^{\text{net}}(t). \tag{76}$$

Prioritization in the BC network with underloaded resources. Our focus in the paper has been on the tension between maximizing capacity and prioritizing specific queues over others. A natural way to alleviate this tension is to give up on maximizing capacity. Here we provide an example to support for our claim that if one resource is kept underloaded some prioritization freedom is regained.

First, notice that Theorem 3.3 does not depend on whether one resource or both are bottlenecks: preemptive priority to the collaborative task maintains optimal scaling and keeps the collaborative queue from growing with ρ^{BN}.

Consider the case of non-preemptive prioritiztion of the collaborative task. In Theorem 3.3 we prove that maximizing capacity conflicts with making this queue short: there exists no policy that can keep queue 0 from growing with ρ^{BN}.

Suppose that we give up on maximizing capacity and have $\rho_1 = \alpha\rho^{\text{BN}} = \alpha\rho_2$ for some $\alpha < 1$. Then, we use the C-priority policy with an added threshold S as in the paper. We change, however, the *mechanism for switching*: when the collaborative queue hits the level S, resource 1 stops as soon he completes his current service; say this happens at time t. Resource 2 will does not stop working on his individual queue 2 until time t (assuming queue 2 is non-empty). Resource 2, stops upon his first service completion after time t.

This sequential switching is bad when both resources are bottlenecks but makes sense here. We want to load all the switching idleness on the underloaded resource 1. At each switch resource 1 incurs at most one service time of resource 2, m_2 so that $\mathbb{E}[T_s] \leq m_1 + m_2$. The network is stable if

$$\rho_1 + \frac{\rho_1^a m_2}{\frac{S}{\lambda_0} + m_1} < 1, \text{ and } \rho_2 = \rho^{\text{BN}} < 1. \tag{77}$$

Notice that (77) since we assume $\rho_1 = \alpha\rho^{\text{BN}} \leq \alpha < 1$. The network is stable if $\rho_1^a m_2/(S/\lambda_0 + m_1) < 1 - \alpha$ and, in turn, if

$$S > \lambda_0 \left(\frac{\rho_1^a m_2}{1-\rho_1} - m_1 \right) = \lambda_0 \left(\frac{\rho_1^a m_2}{1-\alpha\rho^{\text{BN}}} - m_1 \right)$$

If the right hand side is negative than no threshold is needed and the policy is non-preemptive static priority to the collaborative task with sequential switching. More generally, notice that since $\alpha < 1$, the right hand side is bounded in ρ^{BN}. The finite threshold $\bar{S} = \lambda_0 \left(\frac{m_2}{1-\alpha} - m_1 \right) + 1$ is sufficient. Thus, the collaborative queue *does not grow with* ρ^{BN}.

Companion References

Asmussen, S. 2003. *Applied Probability and Queues*. Springer Verlag.

Csörgo, M., L. Horváth. 1996. *Weighted approximations in probability and statistics*. John Wiley & Sons; Chicester.

Derman, C., E. Ignall. 1975. On the stochastic ordering of markov chains. *Operations Research* **23**(3) 574–576.

Gamarnik, D., A. Zeevi. 2006. Validity of heavy traffic steady-state approximations in generalized Jackson networks. *The Annals of Applied Probability* **16**(1) 56–90.

Ghamami, S., A.R. Ward. 2013. Dynamic scheduling of a two-server parallel server system with complete resource pooling and reneging in heavy traffic: Asymptotic optimality of a two-threshold policy. *Mathematics of Operations Research* **38**(4) 761–824.

Glynn, P.W., A. Zeevi. 2008. Bounding stationary expectations of Markov processes. *Markov processes and related topics: A Festschrift for Thomas G. Kurtz. Selected papers of the conference, Madison, WI, USA, July*, vol. 4. 195–214.

Gurvich, I. 2013. Validity of heavy-traffic steady-state approximations in multiclass queueing networks: The case of queue-ratio disciplines. *Mathematics of Operations Research* **39**(1) 121–162.

Klebaner, F.C. 2005. *Introduction to stochastic calculus with applications*. Imperial College Pr.

Robert, P. 2003. *Stochastic networks and queues*. Springer-Verlag, Berlin New York.

van der Mei, R. D. 2007. On a unifying theory on polling models in heavy traffic. *Managing Traffic Performance in Converged Networks*. Springer, 556–567.