

# Dynamic server assignment in an extended machine-repair model

JAN-PIETER L. DORSMAN<sup>1,2,\*</sup>, SANDJAI BHULAI<sup>3,2</sup> and MARIA VLASIOU<sup>1,2</sup>

<sup>1</sup>*Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands*

*E-mail: j.l.dorsman@tue.nl*

*E-mail: s.bhulai@vu.nl*

*E-mail: m.vlasiou@tue.nl*

<sup>2</sup>*Centrum Wiskunde & Informatica (CWI), P.O. Box 94079, 1090 GB Amsterdam, The Netherlands*

<sup>3</sup>*VU University Amsterdam, De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands*

Received November 2012 and accepted April 2014

---

This article considers an extension of the classic machine-repair problem. The machines, apart from receiving service from a single repairman, now also supply service themselves to queues of products. The extended model can be viewed as a two-layered queueing network, in which the queues of products in the first layer are generally correlated, due to the fact that the machines have to share the repairman's capacity in the second layer. Of particular interest is the dynamic control problem of how the repairman should allocate his/her capacity to the machines at any point in time so that the long-term average (weighted) sum of the queue lengths of the first-layer queues is minimized. Since the optimal policy for the repairman cannot be found analytically due to the correlations in the queue lengths, a near-optimal policy is proposed. This is obtained by combining intuition and results from queueing theory with techniques from Markov decision theory. Specifically, the relative value functions for several policies for which the model can be decomposed in less complicated subsystems are studied, and the results are combined with the classic one-step policy improvement algorithm. The resulting policy is easy to apply, is scalable in the number of machines, and is shown to be highly accurate over a wide range of parameter settings.

**Keywords:** Layered queueing networks, Markov decision processes, one-step policy improvement, near-optimal policies

## 1. Background and motivation

In this article, we study a queueing model that consists of two layers. The first layer of the model contains two queues of products; see Fig. 1. Each of these queues is served by its own machine. At any point in time, a machine is subject to failure, irrespective of the number of products in each of the queues. When a failure occurs, the service of a product in progress is interrupted. Upon failure, the machine temporarily stops fulfilling its server role in the first layer and becomes a customer in the second layer of the model. The second layer of the model consists of a single repairman capable of repairing failed machines. When the repair of a machine in the second layer has finished, the machine again assumes its server role in the first layer.

The queueing model we consider in this article has wide applicability. Evidently, it has immediate applications in

manufacturing systems. However, this model is also of interest in less-obvious application areas, such as telecommunication systems. For instance, this model occurs naturally in the modeling of middleware technology, where multi-threaded application servers compete for access to shared object code. Access of threads to the object code is typically handled by a portable object adapter that serializes the threads trying to access a shared object. During the complete duration of the serialization and execution time a thread typically remains blocked, and upon finalizing the execution, the thread is de-activated and ready to process pending requests (Harkema *et al.*, 2004). In this setting, the application servers and the shared object code are analogous to the machines and the repairman. Likewise, several service systems can be modeled in this way.

The two-layered model is an example of a *layered queueing network*. In such networks, there exist servers that when executing a service, may request a lower-layer service and wait for it to be completed. Layered queueing networks occur naturally in information and e-commerce systems, grid systems, and real-time systems such as telecom switches; see Franks *et al.* (2009) and references therein for an overview.

---

\*Corresponding author

Color versions of one or more of the figures in the article can be found online at [www.tandfonline.com/uiie](http://www.tandfonline.com/uiie).

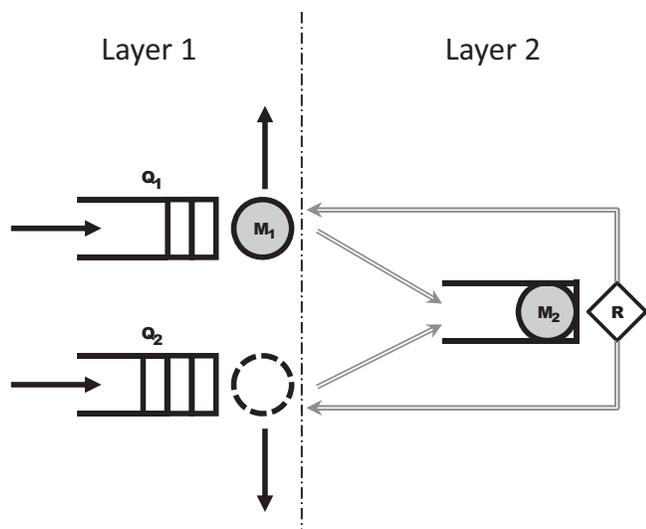


Fig. 1. The two-layered model under consideration.

A layered queueing network is characterized by simultaneous or separate phases where entities are no longer classified in the traditional role of “servers” or “customers,” but they may also have a dual role of being either a server to upper-layer entities or a customer to lower-layer entities. Think, for example, of a peer-to-peer network, where users are both customers when downloading a file but also servers to users who download their files. In our model, the machines are the entities fulfilling the dual role. They act as a server to the products in the first layer but may interrupt service to request a second-layer service from the repairman.

The model we consider is also motivated by an extension of the classic machine-repair problem. The machine-repair problem, also known as the *computer terminal model* (Bertsekas and Gallager, 1992) or as the *time sharing system* (Kleinrock (1976, Section 4.11)), is a well-studied problem in the literature. In the machine-repair problem, there is a number of machines working in parallel, and one repairman. As soon as a machine fails, it joins a repair queue in order to be repaired by the repairman. It is one of the key models to describe problems with a finite input population. An extensive survey of the machine-repair problem can be found in Haque and Armstrong (2007). We extend this model in two different directions. First, we allow different machines to have mutually different failure or repair rates. As observed in Gross and Ince (1981), this leads to technical complications. Second, we assume that each of the machines processes a stream of products, which leads to an additional layer in the model.

An important feature of the classic machine-repair problem and the two-layered model under consideration is the fact that the machines have to share the repair facilities offered by the repairman. When both of the machines are down, this may result in the repair of each of the machines occurring at a much slower rate or in one machine having to wait for the other machine to be repaired before its own

repair can be started. Because of this, significant positive correlations may arise in the downtimes of the machines. As a consequence of the dual roles of the machines, this leads to correlations between the lengths of the queues of products in the first layer. Due to these correlations, it is extremely hard to obtain expressions for the distribution or even the mean of these queue lengths. Even when machines are repaired in the same order as they break down, the queue lengths are notoriously difficult to analyze (see, for example, (Dorsman, Boxma, and Vasiou, 2013; Dorsman, van der Mei, and Vasiou, 2013; Dorsman, Vasiou, and Zwart, 2015).

For the two-layered model under consideration, we study the problem of how the repairman should distribute his capacity over the machines in order to minimize the lengths of the queues of products. To this end, we will use several techniques from Markov decision theory. In particular, we will apply the classic approximation method of one-step policy improvement. This method requires a relative value function of an initial policy, which can be obtained analytically by solving the Poisson equations known from standard theory on Markov decision processes. The result is then used in one step of the policy iteration algorithm from Markov decision theory to obtain an improved policy. Although the relative value function of the improved policy is usually hard to compute, the improved policy itself is known explicitly as a function of the state and the parameters of the model. Moreover, it is known that the improved policy performs better in terms of costs than the initial policy, so that it may be used as an approximation for the optimal policy. The intrinsic idea of one-step policy improvement goes back to Norman (1972). Since then, this method has been successfully applied to derive nearly optimal state-dependent policies in a variety of applications, such as the control of traffic lights (Hajjema and Van der Wal, 2008), production planning (Wijngaard, 1979), and the routing of telephone calls in a network or call center (Ott and Krishnan, 1992; Sassen, Tijms, and Nobel, 1997; Bhulai and Spieksma, 2003). In the next section, we explain in detail the contribution of this article to the problem at hand.

## 2. Contributions

In this article, we are concerned with the question of how the repairman should dynamically allocate his/her capacity to the machines at any point in time, given complete information on the lengths of the queues of products and the states of the machines at all times. We aim to formulate a policy that minimizes the long-term average (weighted) sum of the queue lengths of the first-layer queues as much as possible. Although there exists a considerable amount of literature on the machine-repair model, this is a problem that has not been studied so far. In particular, virtually all studies on machine-repair models investigate performance

measures such as the availability of the machines but do not consider any queues of products that accumulate for service at the machines. In other words, whereas the second layer as depicted in Fig. 1 has received lots of attention, hardly any of the studies in the literature considers the first layer. We, however, particularly focus on the first layer and take the interaction between the two layers into account. The few exceptions that study the first layer (cf. Wartenhorst, 1995; Dorsman, Boxma, and Vlasίου, 2013; Dorsman, van der Mei, and Vlasίου, 2013) only analyze the product queues in the first layer. We, however, aim to optimize the lengths of these queues, by correctly choosing the repairman's policy.

The identification of the repairman's optimal policy is hard, due to the general intractability of the analysis on the queue lengths because of their interdependence. When formulating this problem as a Markov decision problem, one may be able to numerically obtain the optimal policy for a specific set of parameter settings by truncating the state space. However, due to the multi-dimensionality of the model, the computation time needed to obtain reliable and accurate results may be infeasibly long. Moreover, these numerical methods are cumbersome to implement, do not scale well in the number of dimensions of the problem, lack transparency, and provide little insight into the effects of the model's parameters. To overcome these problems, we derive a near-optimal policy that can be explicitly expressed in terms of the model's parameters by use of the one-step policy improvement method. Here, our contribution lies in the more structured way we apply this method. It is tempting to construct an improved policy based on a policy that creates the simplest analytically tractable relative value function. However, this improved policy might not have the best performance, or may cause the lengths of the first-layer queues to grow without bound. In our case, we start with multiple policies that are complementary in different parameter regions in both of these aspects and combine them to create an improved policy that is more broadly applicable. In our objective to derive a near-optimal policy over a broad range of parameter values, we need to start with initial policies that do not allow for the derivation of closed-form expressions for the relative value function. In some cases, we use insights from queueing theory to provide an accurate approximation for the relative value function and the long-term averaged costs. We use these results to construct a near-optimal policy that requires no computation time, is easy to implement, and gives insights into the effects of the model's parameters.

The article is organized as follows. Section 3 gives a mathematical description of the control problem and introduces the notation required. Although the optimal policy for this control problem cannot be obtained explicitly, several of its structural properties can be derived. As we will see in Section 4, the optimal policy makes the repairman work at full capacity whenever there is at least one machine down, and behaves like a threshold policy. Subsequently, we focus on finding a policy that generally performs nearly as well as the optimal policy. As input for the one-step policy im-

provement algorithm, we study two policies in Section 5, for which the system decomposes into multiple subsystems, so that the system becomes easier to evaluate. The first of these policies, which we will call the static policy, is state-independent and always reserves a certain predetermined fraction of repair capacity to each machine, regardless of the state of the machines. Therefore, the machines behave independently of each other under this policy, which allows us to derive an exact expression for the relative value function. As the static policy cannot always be used as an input for the one-step policy improvement algorithm, we also study a second class of policies in Section 5. More specifically, we study the priority policy, in which the repairman always prioritizes the repair of a specific machine over the other when both machines are down. Under this policy, the repairman assigns his/her full capacity to the high-priority machine when it is down, irrespective of the state of the low-priority machine. This makes the system easier to analyze. Nevertheless, it is hard to exactly obtain the relative value function for this policy, but we are able to identify most of its behavior. Although analytic results on the relative value functions of these policies are of independent interest, we use these results in Section 6 in combination with the one-step policy improvement algorithm. This ultimately results in a well-performing, nearly optimal policy, which is given in terms of a few simple decision rules. The resulting policy turns out to be scalable in the number of machines and corresponding first-layer queues in the model, so that the policy can be readily extended to allow for a number of machines larger than two. Extensive numerical results in Section 7 show that the proposed policy is highly accurate over a wide range of parameter settings. We also identify the key factors that determine the performance of the near-optimal policy. Finally, Section 8 gives the main conclusions of this study and provides directions for future research.

### 3. Model description and problem formulation

The layered model consists of two machines  $M_1$  and  $M_2$  and a single repairman  $R$ , see Fig. 1. Each machine  $M_i$  serves its own first-layer queue  $Q_i$  of products. The products arrive at  $Q_i$  according to a Poisson process with rate  $\lambda_i$ . The service requirements of the products in  $Q_i$  are exponentially distributed with rate  $\mu_i$ . The service of these products may be interrupted, since the machines are prone to failure. After an exponentially ( $\sigma_i$ ) distributed uptime or lifetime, denoted by  $U_i$ , a machine  $M_i$  will break down, and the service of  $Q_i$  inevitably stops. When this happens, the service of a product in progress is aborted and will be resumed once the machine is operational again. To be returned to an operational state, the machine  $M_i$  needs to be repaired by the repairman. Whenever the repairman allocates his/her full repair capacity to  $M_i$ , its repair takes an exponential ( $\nu_i$ ) distributed time. However, the machines share the capacity of the repairman. At any moment in time, the repairman is able to decide how to divide his/her total repair

capacity over the machines. More specifically, he/she can choose the fractions of capacity  $q_1$  and  $q_2$  that are allocated to the repair of  $M_1$  and  $M_2$  respectively, so that the machines are getting repaired at rate  $q_1v_1$  and  $q_2v_2$ , respectively. We naturally have that  $0 \leq q_1 + q_2 \leq 1$  and that  $q_i = 0$  whenever  $M_i$  is operational. The objective of the repairman is to dynamically allocate his/her repair capacity in such a way that the average long-term weighted number of products in the system is minimized.

In order to mathematically describe this dynamic optimization problem, one needs to keep track of the queues of products as well as the conditions of the machines. To this end, we define the state space of the system as  $\mathcal{S} = \mathbb{N}^2 \times \{0, 1\}^2$ . Each possible state corresponds to an element  $s = (x_1, x_2, w_1, w_2)$  in  $\mathcal{S}$ , where  $x_1$  and  $x_2$  denote the number of products in  $Q_1$  and  $Q_2$ , respectively. The variables  $w_1$  and  $w_2$  denote whether  $M_1$  and  $M_2$  are in an operational (1) or in a failed state (0), respectively. The repairman bases his/her decision on the information  $s$  and, therefore, any time the state changes can be regarded as a decision epoch. At these epochs, the repairman takes an action  $\mathbf{a} = (q_1, q_2)$  out of the state-dependent action space  $\mathcal{A}_s = \{(q_1, q_2) : q_1 \in [0, 1 - w_1], q_2 \in [0, 1 - w_2], q_1 + q_2 \leq 1\}$ , where  $q_i$  denotes the fraction of capacity assigned to  $M_i$ ,  $i = 1, 2$ . The terms  $1 - w_1$  and  $1 - w_2$  included in the description of the action set enforce the fact that the repairman can only repair a machine if it is down. Now that the states and actions are defined, we introduce the cost structure of the model. The objective is modeled by the cost function  $c(s, \mathbf{a}) = c_1x_1 + c_2x_2$ , where  $c_1$  and  $c_2$  are non-negative real-valued weights. Thus, when the system is in state  $s$ , the weighted number of customers present in the system equals  $c(s, \cdot)$ , regardless of the action  $\mathbf{a}$  taken by the repairman.

With this description, the control problem can be fully described as a Markov decision problem. To this end, we uniformize the system (see, for example, Lippman (1975a)); i.e., we add dummy transitions (from a state to itself) such that the outgoing rate of every state equals a constant parameter  $\gamma$ , the uniformization parameter. We choose  $\gamma = \lambda_1 + \lambda_2 + \mu_1 + \mu_2 + \sigma_1 + \sigma_2 + v_1 + v_2$  and we assume that  $\gamma = 1$  without loss of generality, since we can always achieve this by scaling the model parameters. Note that this assumption has the intuitive benefit that rates can be considered to be transition probabilities, since the outgoing rates of each state sum up to one. Thus, for  $i = 1, 2$ , any action  $\mathbf{a} \in \mathcal{A}_s$  and any state  $s \in \mathcal{S}$ , the transition probabilities  $P$  are given by

$$\begin{aligned} P_a(s, s + e_i) &= \lambda_i, && \text{(product arrivals),} \\ P_a(s, s - e_i) &= \mu_i w_i \mathbb{1}_{\{x_i > 0\}}, && \text{(product services),} \\ P_a(s, s - e_{i+2}) &= \sigma_i w_i, && \text{(machine breakdowns),} \\ P_a(s, s + e_{i+2}) &= q_i v_i, && \text{(machine repairs),} \\ P_a(s, s) &= 1 - \lambda_i - w_i && \text{(uniformization),} \\ & \quad (\mu_i \mathbb{1}_{\{x_i > 0\}} + \sigma_i) \\ & \quad - q_i v_i, \end{aligned}$$

where  $\mathbb{1}_{\{E\}}$  denotes the indicator function on the event  $E$  and  $e_i$  represents the unit vector of which the  $i$ th entry equals one. All other transition probabilities are zero. The tuple  $(\mathcal{S}, \{\mathcal{A}_s : s \in \mathcal{S}\}, P, c)$  now fully defines the Markov decision problem at hand.

Define a deterministic policy  $\pi^*$  as a function from  $\mathcal{S}$  to  $\mathcal{A}$ ; i.e.,  $\pi^*(s) \in \mathcal{A}_s$  for all  $s \in \mathcal{S}$ , and let  $\{X^*(t) : t \geq 0\}$  be its corresponding Markov process taking values in  $\mathcal{S}$ , which describes the state of the system over time when the repairman adheres to policy  $\pi^*$ . Furthermore, let

$$u^*(s, t) = \mathbb{E} \left[ \int_{z=0}^t c(X^*(z), \pi^*(X^*(z))) dz \mid X^*(0) = s \right],$$

denote the total expected costs up to time  $t$  when the system starts in state  $s$  under policy  $\pi^*$ .

We call the policy  $\pi^*$  *stable*, when the average costs  $g^* = \lim_{t \rightarrow \infty} (u^*(s, t))/t$  per time unit that arise when the repairman adheres to this policy, remain finite. From this, it follows that the Markov process corresponding to the model under consideration in combination with a stable policy has a single positive recurrent class. As such, the number  $g^*$  is independent of the initial state  $s$ . Due to the definition of the cost function, the average expected costs may also be interpreted as the long-term average sum of queue lengths under policy  $\pi^*$ , weighted by the constants  $c_1$  and  $c_2$ . A stable policy thus coincides with a policy for which the average number of customers in each of the queues is finite. Observe that there does not necessarily exist a stable policy for every instance of this model. In fact, necessary (but not sufficient) conditions for the existence of a stable policy are given by

$$\lambda_1 < \mu_1 \frac{v_1}{\sigma_1 + v_1} \quad \text{and} \quad \lambda_2 < \mu_2 \frac{v_2}{\sigma_2 + v_2}. \quad (1)$$

This condition implies that for each queue  $Q_i$ , the rate  $\lambda_i$  at which products arrive is smaller than the rate at which the corresponding machine  $M_i$  is capable of processing products, given that  $M_i$  is always repaired instantly at full capacity when it breaks down. This assumption can in some sense be seen as the best-case scenario from the point of view of  $M_i$ . The latter processing rate is, of course, equal to the service rate  $\mu_i$  times the fraction  $(1/\sigma_i)/(1/\sigma_i + 1/v_i) = v_i/(\sigma_i + v_i)$  of time that  $M_i$  is operational under this best-case assumption. When this condition is not satisfied, it is sure that there is at least one queue where on average more products arrive per time unit than the machine can handle under any repair policy. Thus, the costs accrued will grow without bound over time for any policy in such case, eliminating the existence of a stable policy. Observe that the converse of the necessary conditions for existence of stable policy stated in Equation (1) constitutes two different sufficient conditions for non-existence. That is, if

$$\lambda_1 \geq \mu_1 \frac{v_1}{\sigma_1 + v_1} \quad \text{or} \quad \lambda_2 \geq \mu_2 \frac{v_2}{\sigma_2 + v_2},$$

then it is guaranteed that no stable policies exist.

Any policy  $\pi^*$  can be characterized through its relative value function  $V^*(s)$ . This function is a real-valued function defined on the state space  $\mathcal{S}$  given by

$$V^*(s) = \lim_{t \rightarrow \infty} (u^*(s, t) - u^*(s_{\text{ref}}, t))$$

and represents the asymptotic difference in expected total costs accrued when starting the process in state  $s$  instead of some reference state  $s_{\text{ref}}$ ; see, for example, Hernández-Lerma and Lasserre (1996, Equation (5.6.2)). Among all policies, the optimal policy  $\pi^{\text{opt}}$  with relative value function  $V^{\text{opt}}$  minimizes the average costs (i.e., the long-term average weighted sum of queue lengths); thus,  $g^{\text{opt}} = \min_{\pi^*} g^*$ . Its corresponding long-term optimal actions are a solution of the Bellman optimality equations  $g^{\text{opt}} + V^{\text{opt}}(s) = \min_{a \in \mathcal{A}_s} \{c(s, a) + \sum_{t \in \mathcal{S}} P_a(s, t) V^{\text{opt}}(t)\}$  for all  $s \in \mathcal{S}$ . For our problem, this constitutes

$$g^{\text{opt}} + V^{\text{opt}}(x_1, x_2, w_1, w_2) = H^{\text{opt}}(x_1, x_2, w_1, w_2) + K^{\text{opt}}(x_1, x_2, w_1, w_2),$$

for every  $(x_1, x_2, w_1, w_2) \in \mathcal{S}$ , where  $H^{\text{opt}}$  and  $K^{\text{opt}}$  are defined in the following way. For an arbitrary policy  $\pi^*$  with relative value function  $V^*$ , the function  $H^*$  is given by

$$\begin{aligned} H^*(x_1, x_2, w_1, w_2) &= c_1 x_1 + c_2 x_2 \\ &+ \lambda_1 V^*(x_1 + 1, x_2, w_1, w_2) + \lambda_2 V^*(x_1, x_2 + 1, w_1, w_2) \\ &+ \mu_1 w_1 V^*((x_1 - 1)^+, x_2, 1, w_2) \\ &+ \mu_2 w_2 V^*(x_1, (x_2 - 1)^+, w_1, 1) + \sigma_1 w_1 V^*(x_1, x_2, 0, w_2) \\ &+ \sigma_2 w_2 V^*(x_1, x_2, w_1, 0) \\ &+ \left(1 - \sum_{i=1}^2 (\lambda_i + w_i (\mu_i + \sigma_i))\right) V^*(x_1, x_2, w_1, w_2), \end{aligned} \quad (2)$$

and it models the costs and the action-independent events of product arrivals, product service completions, machine breakdowns, and dummy transitions, respectively. The function  $K^*$  given by

$$\begin{aligned} K^*(x_1, x_2, w_1, w_2) &= \min_{(q_1, q_2) \in \mathcal{A}(x_1, x_2, w_1, w_2)} \{q_1 v_1 (V^*(x_1, x_2, 1, w_2) \\ &- V^*(x_1, x_2, 0, w_2)) + q_2 v_2 (V^*(x_1, x_2, w_1, 1) \\ &- V^*(x_1, x_2, w_1, 0))\} \end{aligned} \quad (3)$$

models the optimal state-specific decisions of how to allocate the repair capacity over the machines and includes corrections for the uniformization term.

As already mentioned in Section 1, these equations are exceptionally hard to solve analytically. Alternatively, the optimal actions can also be obtained numerically by recursively defining  $V^{n+1}(s) = H^n(s) + K^n(s)$  for an arbitrary function  $V^0$ . For  $n \rightarrow \infty$ , the minimizing actions converge to the optimal ones (see Lippman (1975b) for conditions on existence and convergence). We use this procedure called *value iteration* or *successive approximation* for our numerical experiments in Section 7.

## 4. Structural properties of the optimal policy

As mentioned before, it is hard to give a complete, explicit characterization of the optimal policy for the problem sketched in Section 3. Therefore, we derive a near-optimal policy later in Section 6. Nevertheless, several important structural properties of the optimal policy can be obtained. It turns out that the optimal policy is a non-idling policy, always dictates the repairman to work on one machine only, and can be classified as a threshold policy. In this section, we inspect these properties more closely.

### 4.1. Non-idling property

We show in this section that the optimal policy is a non-idling policy, which means the repairman always repairs at full capacity whenever a machine is not operational; i.e.,  $q_1 + q_2 = 1 - w_1 w_2$ . Intuitively, this makes sense, as there are no costs involved for the repairman's service. On the other hand, having less repair capacity go unused has decreasing effects on the long-term weighted number of products in the system. There is no trade-off present, and therefore the repair capacity should be used exhaustively whenever there is a machine in need of repair.

This property can be proved mathematically. Note that the minimizers of the right-hand side of Equation (3) represent the optimal actions. From this, it follows that the optimal action satisfies  $q_1 + q_2 = 1 - w_1 w_2$  for every state  $s \in \mathcal{S}$  (i.e., the optimal policy satisfies the non-idling property), if both  $V^{\text{opt}}(x_1, x_2, 0, w_2) - V^{\text{opt}}(x_1, x_2, 1, w_2)$  and  $V^{\text{opt}}(x_1, x_2, w_1, 0) - V^{\text{opt}}(x_1, x_2, w_1, 1)$  are non-negative for all  $(x_1, x_2, w_1, w_2) \in \mathcal{S}$ . The following proposition proves the latter condition. For the sake of reduction of the proof's complexity, it also concerns the trivial fact that under the optimal policy, the system incurs higher costs whenever the number of products in the system is increasing (i.e.,  $V^{\text{opt}}(x_1 + 1, x_2, w_1, w_2) - V^{\text{opt}}(x_1, x_2, w_1, w_2)$  and  $V^{\text{opt}}(x_1, x_2 + 1, w_1, w_2) - V^{\text{opt}}(x_1, x_2, w_1, w_2)$  are non-negative).

**Proposition 1.** *The relative value function  $V^{\text{opt}}(s)$  corresponding to the optimal policy satisfies the following properties for all  $s \in \mathcal{S}$ :*

1.  $V^*(x_1, x_2, 0, w_2) - V^*(x_1, x_2, 1, w_2) \geq 0$  and  $V^*(x_1, x_2, w_1, 0) - V^*(x_1, x_2, w_1, 1) \geq 0$ ;
2.  $V^*(x_1 + 1, x_2, w_1, w_2) - V^*(x_1, x_2, w_1, w_2) \geq 0$  and  $V^*(x_1, x_2 + 1, w_1, w_2) - V^*(x_1, x_2, w_1, w_2) \geq 0$ .

**Proof.** The proof is based on induction and the guaranteed convergence of the value iteration algorithm. We arbitrarily pick a function  $V^0(s) = 0$  for all  $s \in \mathcal{S}$ . Obviously, this function satisfies properties 1 and 2. We show that these properties are preserved when performing one step of the value iteration algorithm. In mathematical terms, we show for any  $n \in \mathbb{N}$  that the function  $V^{n+1}$  defined by  $V^{n+1}(s) = H^n(s) + K^n(s)$  also satisfies the properties if  $V^n$

also satisfies these properties. Because of the guaranteed convergence,  $V^{opt}$  then satisfies properties 1 and 2 by induction. For an extensive discussion of this technique to prove structural properties of relative value functions, see Koole (2006).

The induction step is performed as follows. We assume that properties 1 and 2 hold for  $V^n$  (the induction assumption). We now show that property 1 holds for  $V^{n+1}$ . Observe that by interchanging the indices of the model parameters, one obtains another instance of the same model, since the structure of the model is symmetric. Therefore, the left-hand side of property 1 implies the right-hand side. To prove the first part of property 1, we expand  $V^{n+1}(x_1, x_2, 0, w_2) - V^{n+1}(x_1, x_2, 1, w_2)$  into  $V^n$ :

$$\begin{aligned} & V^{n+1}(x_1, x_2, 0, w_2) - V^{n+1}(x_1, x_2, 1, w_2) \\ &= H^n(x_1, x_2, 0, w_2) - H^n(x_1, x_2, 1, w_2) \\ & \quad + K^n(x_1, x_2, 0, w_2) - K^n(x_1, x_2, 1, w_2). \end{aligned} \quad (4)$$

By rearranging the terms arising from Equation (2) and applying the induction assumption, we have that

$$\begin{aligned} & H^n(x_1, x_2, 0, w_2) - H^n(x_1, x_2, 1, w_2) \\ & \geq (1 - \lambda_1 - \lambda_2 - \sigma_1 - w_2(\mu_2 + \sigma_2))(V^n(x_1, x_2, 0, w_2) \\ & \quad - V^n(x_1, x_2, 1, w_2)). \end{aligned} \quad (5)$$

Furthermore, since  $V^n(x_1, x_2, w_1, 1) - V^n(x_1, x_2, w_1, 0)$  and  $V^n(x_1, x_2, 1, w_2) - V^n(x_1, x_2, 0, w_2)$  are both non-positive numbers, we can limit the set of possible minimizing actions in  $K^n$  (see Equation (3)) to  $\{(q_1, q_2) : q_1 \in \{0, 1 - w_1\}, q_2 \in \{0, 1 - w_2\}, q_1 + q_2 = 1 - w_1 w_2\}$ . By this and Equation (3), we obtain

$$\begin{aligned} & K^n(x_1, x_2, 0, w_2) - K^n(x_1, x_2, 1, w_2) \\ &= \min\{v_1(V^n(x_1, x_2, 1, w_2) - V^n(x_1, x_2, 0, w_2)), \\ & \quad (1 - w_2)v_2(V^n(x_1, x_2, 0, 1) - V^n(x_1, x_2, 0, 0)) \\ & \quad - (1 - w_2)v_2(V^n(x_1, x_2, 1, 1) - V^n(x_1, x_2, 1, 0))\}, \end{aligned} \quad (6)$$

where the second equality holds because of the induction assumption. Let  $E_1$  denote the event that the last minimum is only minimized by its first argument and let  $E_2$  be its complementary event. As a conclusion, we find by combining Equations (4) to (6) that

$$\begin{aligned} & V^{n+1}(x_1, x_2, 0, w_2) - V^{n+1}(x_1, x_2, 1, w_2) \\ & \geq (1 - \lambda_1 - \lambda_2 - \sigma_1 - w_2(\mu_2 + \sigma_2) - \mathbb{1}_{\{E_1\}}v_1 - \mathbb{1}_{\{E_2\}} \\ & \quad \times (1 - w_2)v_2)(V^n(x_1, x_2, 0, w_2) - V^n(x_1, x_2, 1, w_2)) \\ & \quad + \mathbb{1}_{\{E_1\}}(1 - w_2)v_2(V^n(x_1, x_2, 1, 0) - V^n(x_1, x_2, 1, 1)) \\ & \quad + \mathbb{1}_{\{E_2\}}(1 - w_2)v_2(V^n(x_1, x_2, 0, 1) - V^n(x_1, x_2, 1, 1)) \\ & \geq 0. \end{aligned}$$

The last inequality holds by applying the induction assumption on each term of the expression in front of it and observing, for the first term, that  $(1 - \lambda_1 - \lambda_2 - \sigma_1 - w_2(\mu_2 + \sigma_2) - \mathbb{1}_{\{E_1\}}v_1 - \mathbb{1}_{\{E_2\}}(1 - w_2)v_2)$  is non-negative due to uniformization.

This proves that property 1 holds. By similar techniques of expanding  $V^{n+1}$  into  $V^n$  and termwise elimination, one

can also show that  $V^{n+1}$  satisfies property 2 under the induction assumption, which completes the proof. ■

By proving that  $V^{opt}$  satisfies property 1 as stated in Proposition 1, we have established that the optimal policy is a non-idling policy, implying that  $q_1 + q_2 = 1 - w_1 w_2$  at all times. We finish this section by pointing out that the optimal policy always dictates the repairman to focus all his/her attention on one machine. That is, at all times, the optimal action reads  $(q_1, q_2) = (1 - w_1, 0)$  or  $(q_1, q_2) = (0, 1 - w_2)$ . This is easily derived from Equation (3) in combination with property 1 in Proposition 1. Even when there are states for which  $w_1 w_2 = 0$  and  $v_1(V^n(x_1, x_2, 1, w_2) - V^n(x_1, x_2, 0, w_2)) = v_2(V^n(x_1, x_2, w_1, 1) - V^n(x_1, x_2, w_1, 0))$ , the actions  $(q_1, q_2) = (1 - w_1, 0)$  and  $(q_1, q_2) = (0, 1 - w_2)$  will be optimal (although they are not uniquely optimal), so that there are always optimal policies that concentrate all repair capacity on one machine. Therefore,  $K^{opt}$  can be simplified to

$$\begin{aligned} & K^{opt}(x_1, x_2, w_1, w_2) \\ &= \min_{(q_1, q_2) \in \{(1-w_1, 0), (0, 1-w_2)\}} \{q_1 v_1(V^{opt}(x_1, x_2, 1, w_2) \\ & \quad - V^{opt}(x_1, x_2, 0, w_2)) + q_2 v_2(V^{opt}(x_1, x_2, w_1, 1) \\ & \quad - V^{opt}(x_1, x_2, w_1, 0))\}. \end{aligned} \quad (7)$$

This is a welcome simplification when one wants to evaluate the optimal policy numerically, since now the minimum-operator only involves two arguments.

#### 4.2. Threshold policy

Now that we know that the optimal policy is a non-idling policy and always dictates the repairman to focus his/her attention on a single machine, the question arises regarding which machine is the one in question. In the event that both machines are down, this question is hard to answer explicitly, since the relative value function  $V^{opt}$  pertaining to the optimal policy defies an exact analysis. However, by inspection of numerical results, one can derive a partial answer.

To this end, we numerically examine the model with the settings  $c_1 = c_2 = \mu_2 = \sigma_1 = v_1 = 1.0$ ,  $\lambda_1 = 0.1$ ,  $\lambda_2 = 0.2$  and  $\mu_1 = \sigma_2 = v_2 = 0.5$ . By using the simplified version (7) of  $K^{opt}$  in the value iteration algorithm, we numerically obtain the optimal actions for the states  $(x_1, x_2, 0, 0)$ ,  $x_1 \in \{0, \dots, 50\}$ ,  $x_2 \in \{0, \dots, 100\}$ . Figure 2 shows the optimal actions in the form of a scatter plot. Given that both machines are down, a marked point  $(x_1, x_2)$  in the scatter plot indicates that it is optimal for the repairman to repair  $M_2$ . If a certain point  $(x_1, x_2)$  is not marked, then the optimal action is to repair  $M_1$  at full capacity.

Judging by Fig. 2, it seems that the optimal policy falls in the class of threshold policies. That is, if the optimal action for the state  $(x_1, x_2, 0, 0)$  is to repair  $M_1$  at full capacity, then this is also the optimal action for the states  $(x_1 + k, x_2, 0, 0)$ ,

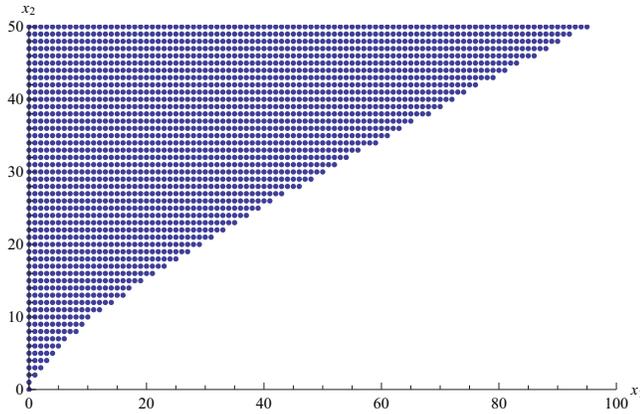


Fig. 2. The optimal actions for the model instance studied in Section 4.2.

$k \in \mathbb{N}$ . Meanwhile, if it would be optimal to repair  $M_2$  when the system is in the state  $(x_1, x_2, 0, 0)$ , then the optimal policy also prescribes to repair  $M_2$  if there are fewer products waiting in  $Q_1$ ; i.e., in the states  $(x_1 - k, x_2, 0, 0)$ ,  $k \in \{1, \dots, x_1\}$ . Thus, for any number  $x_2$ , the number of products in  $Q_1$  from which the optimal policy starts taking the decision to repair  $M_1$  can be seen as a threshold. Similar effects and definitions apply for varying number of products in  $Q_2$ . The figure clearly exposes a curve that marks the thresholds. At first glance, this threshold curve may seem linear. However, especially near the origin, this is not quite true.

One can reason intuitively that for any instance of the model, the optimal policy is a threshold policy. This is easily understood by the notion that an increasing number of products in  $Q_1$  makes it more attractive for the repairman to repair  $M_1$ . Then, if it was already optimal to repair  $M_1$ , this obviously will not change. Similar notions exist for a decreasing number of products in  $Q_1$  and varying numbers of products in  $Q_2$ . Although the threshold effects are easily understood, they are hard to prove mathematically. A possible approach to this would be to show that the difference between the arguments in Equation (7) is increasing in  $x_1$  using the same techniques as used in the proof of Proposition 1. However, this turns out to be very challenging.

**5. Relative value functions**

Recall that for any policy  $\pi^*$ , we defined  $V^*$  and  $g^*$  to be its corresponding relative value function and long-run expected weighted number of products in the system, respectively. The main reason why one cannot obtain the optimal policy  $\pi^{opt}$  other than through numerical means is because its corresponding relative value function  $V^{opt}$  does not allow for an exact analysis. As an intermediate step, we therefore study the relative value functions of two other policies, for which explicit expressions can be obtained. In Section 6,

these two policies and their relative value functions act as a basis for the one-step policy improvement method to obtain nearly optimal heuristic policies. We first examine the static policy in Section 5.1, where each machine is assigned a fixed part of the repair capacity regardless of the system’s state. However, there exist instances of the model for which no static policies are available that result in a finite average cost, while stable policies are available in general. Since a one-step policy improvement approach based on a static policy in that case, we will also study the priority policy in Section 5.2, which dictates the repairman to prioritize a specific machine (the high-priority machine), in the case where both machines are not operational; i.e., in such a case, all repair capacity is given to the high-priority machine.

**5.1. Static policy**

As the name of the static policy suggests, the actions taken under this policy do not depend on the state of the system. Under the static policy, the repairman always has a fraction  $p \in (0, 1)$  of his/her repair capacity reserved for the repair of  $M_1$ , regardless of whether  $M_1$  (or  $M_2$ ) is down or not. Likewise, the remaining fraction  $(1 - p)$  is reserved for  $M_2$ . Therefore, repair on  $M_1$  at rate  $p\nu_1$  starts instantly the moment it breaks down, and the same holds for  $M_2$  at rate  $(1 - p)\nu_2$ . Thus, under this policy, the repairman always takes the action  $(p(1 - w_1), (1 - p)(1 - w_2))$ . In the remainder of this article, we will refer to  $p$  as the splitting parameter. It is evident that this policy is not optimal, since the repairman does not use his/her repair capacity exhaustively when exactly one of the two machines is down; i.e., the static policy does not satisfy the non-idling property studied in Section 4.1. However, when the splitting parameter is chosen well, this policy is not totally unreasonable. When analyzing this policy, we assume that the system is stable when adhering to it. That is, for each queue the rate of arriving products is smaller than the rate at which the corresponding machine is capable of serving products

$$\lambda_1 < \mu_1 \frac{p\nu_1}{\sigma_1 + p\nu_1} \quad \text{and} \quad \lambda_2 < \mu_2 \frac{(1 - p)\nu_2}{\sigma_2 + (1 - p)\nu_2}, \quad (8)$$

where the two fractions denote the fractions of time  $M_1$  and  $M_2$  are operational, respectively.

Observe that the capacity that  $M_1$  receives from the repairman is now completely independent of that received by  $M_2$  at any given time and *vice versa*. Analysis of the relative value function of the static policy is tractable, since the machines do not compete for repair resources anymore under this policy, making the queue lengths in each of the queues uncorrelated. In a way, it is as if each machine now has its own repairman, repairing at rate  $p\nu_1$  and  $(1 - p)\nu_2$ , respectively. Therefore, the system can be decomposed into two components, which do not interact. Each of these components can be modelled as a single-server queue of M/M/1 type with server vacations occurring independently of the

amount of work present in the queue. Because of this decomposition, the relative value function  $V^{sta}(x_1, x_2, w_1, w_2)$  of the total system can be seen as the weighted sum of the relative value functions  $V_1^{com}(x_1, w_1)$  and  $V_2^{com}(x_2, w_2)$  corresponding to the two components. As a result, the long-term average cost  $g^{sta}$  is also a weighted sum of the average costs  $g_1^{com}$  and  $g_2^{com}$ :

$$g^{sta} = c_1 g_1^{com} + c_2 g_2^{com} \quad \text{and} \quad V^{sta}(x_1, x_2, w_1, w_2) = c_1 V_1^{com}(x_1, w_1) + c_2 V_2^{com}(x_2, w_2). \quad (9)$$

To derive  $g_1^{com}$ ,  $g_2^{com}$ ,  $V_1^{com}(x_1, w_1)$ , and  $V_2^{com}(x_2, w_2)$ , we focus on the relative value function corresponding to one component in Section 5.1.1. We then finalize the analysis on  $V^{sta}$  in Section 5.1.2.

### 5.1.1. Relative value function for the components

We now derive the relative value function of one component of the model under the static policy and omit all indices of the parameters. Thus, we regard a single-server queue of M/M/1 type, in which products arrive at rate  $\lambda$  and are processed at rate  $\mu$  if the machine is up. Independent of this process, the server takes a vacation after an exponentially ( $\sigma$ ) distributed amount of time, even when there is a product in service. The service of the product is then interrupted and resumed once the server ends its vacation. A vacation takes an exponentially ( $\nu$ ) distributed amount of time, after which the server will process products again until the next vacation. This system can be interpreted as a Markov reward chain with states  $(x, w) \in \mathcal{S}^{com}$  representing the number  $x$  of products present in the system and the state of the server being in a vacation ( $w = 0$ ) or not ( $w = 1$ ), where  $\mathcal{S}^{com} = \mathbb{N} \times \{0, 1\}$  is its state space. The system is said to accrue costs in a state-dependent manner at rate  $c(x, w) = x$  per time unit. After uniformization at rate one, the transition probabilities  $P^{com}(s, t)$  from a state  $s \in \mathcal{S}^{com}$  to a state  $t \in \mathcal{S}^{com}$  are given by

$$\begin{aligned} P^{com}((x, w), (x + 1, w)) &= \lambda, \\ P^{com}((x, w), (x - 1, w)) &= \mu w \mathbb{1}_{\{x > 0\}}, \\ P^{com}((x, 1), (x, 0)) &= \sigma, \\ P^{com}((x, 0), (x, 1)) &= \nu, \\ \text{and } P^{com}((x, w), (x, w)) &= (1 - \lambda - w(\mu \mathbb{1}_{\{x > 0\}} + \sigma) \\ &\quad + \nu(1 - w)), \end{aligned}$$

with all other transition probabilities being zero. By the above description, the Poisson equations for this Markov reward chain with long-term average costs per time unit  $g^{com}$  and relative value function  $V^{com}(x, w)$  are given by

$$\begin{aligned} g^{com} + V^{com}(x, w) &= x + \lambda V^{com}(x + 1, w) + \mu w V^{com}((x - 1)^+, w) \\ &\quad + \sigma w V^{com}(x, 0) + \nu(1 - w) V^{com}(x, 1) \\ &\quad + (1 - \lambda - w(\mu + \sigma) - \nu(1 - w)) V^{com}(x, w), \quad (10) \end{aligned}$$

for all  $(x, w) \in \mathbb{N} \times \{0, 1\}$ .

To solve these equations, we first observe that the completion time for a product from the moment its service is

started until it leaves the system consists of an exponentially ( $\mu$ ) distributed amount of actual service time and possibly some interruption time due to server vacations. When interruption takes place, the number of interruptions is geometrically ( $\mu/(\mu + \sigma)$ ) distributed, due to the Markovian nature of the model. Combined with the fact that every interruption takes an exponential ( $\nu$ ) amount of time, this means that the total interruption time, given that it is positive, is exponentially ( $\mu\nu/(\mu + \sigma)$ ) distributed. Thus, the completion time consists of an exponential ( $\mu$ ) service phase and also, with a probability  $\sigma/(\mu + \sigma)$  that there is at least one interruption, an exponential ( $\mu\nu/(\mu + \sigma)$ ) interruption phase. The above implies that the distribution of the completion time falls in the class of Coxian distributions with order two. Due to this observation, the average costs per time unit  $g^{com}$  incurred by a component can be calculated by the use of standard queueing theory, see Remark 2. However, we are also interested in the relative value function of the component. If the server would only start a vacation if there is at least one product in the queue, the component could in principle be modeled as an M/Cox(2)/1 queue, by incorporating the interruption times into the service times (that is, by replacing the service times with the completion times). For the M/Cox(2)/1 queue, it is known that the relative value function can be expressed as a second-order polynomial in the queue length (Bhulai, 2006). However, in our case, a server may also start a vacation during an idle period, so that products arriving at an empty system may not be served instantly. Nevertheless, it is reasonable to conjecture that the relative value function  $V^{com}$  is also of a second-order polynomial type.

If this conjecture holds, substituting  $V^{com}(x, 0) = \alpha_1 x^2 + \alpha_2 x + \alpha_3$  and  $V^{com}(x, 1) = \beta_1 x^2 + \beta_2 x + \beta_3$  in Equation (10) should lead to a consistent system of equations and give a solution for the coefficients. After substitution, we find the equations

$$\begin{aligned} g^{com} + \alpha_3 &= \lambda(\alpha_1 + \alpha_2) + (1 - \nu)\alpha_3 + \nu\beta_3, \\ g^{com} + \beta_3 &= \sigma\alpha_3 + \lambda(\beta_1 + \beta_2) + (1 - \sigma)\beta_3, \\ g^{com} + \alpha_1 x^2 + \alpha_2 x + \alpha_3 &= ((1 - \nu)\alpha_1 + \nu\beta_1)x^2 + (1 + 2\lambda\alpha_1 + (1 - \nu)\alpha_2 + \nu\beta_2)x \\ &\quad + \lambda(\alpha_1 + \alpha_2) + (1 - \nu)\alpha_3 + \nu\beta_3, \\ g^{com} + \beta_1 x^2 + \beta_2 x + \beta_3 &= (\sigma\alpha_1 + (1 - \sigma)\beta_1)x^2 + (1 + \sigma\alpha_2 + 2(\lambda - \mu)\beta_1 \\ &\quad + (1 - \sigma)\beta_2)x + \sigma\alpha_3 + (\lambda + \mu)\beta_1 + (\lambda - \mu)\beta_2 \\ &\quad + (1 - \sigma)\beta_3, \end{aligned}$$

for all  $x \in \mathbb{N}_+$ . One can easily verify that the system of equations is indeed consistent. By solving for the coefficients, a solution for  $g^{com}$  and  $V^{com}$  up to a constant can be found. The constant can be chosen arbitrarily—e.g., by assuming that  $V^{com}(0, 1) = 0$ —but is of no importance. In principle, there may exist other solutions to Equation (10) that do not behave like a second-order polynomial in  $x$ . In fact, when the state space is not finite, as is the case in our

model, it is known that there are many pairs of  $g$  and  $V$  that satisfy the Poisson equations (10) (see, for example, Bhulai and Spieksma (2003)). There is only one pair satisfying  $V(0, 1) = 0$  that is the correct stable solution, however, and we refer to this as the unique solution. Showing that a solution to Equation (10) is the unique solution involves the construction of a weighted norm so that the Markov chain is geometrically recurrent with respect to that norm. This weighted norm imposes extra conditions on the solution to the Poisson equations, so that the unique solution can be identified. Lemma 1 summarizes the solution resulting from the set of equations above, and states that this is also the unique solution.

**Lemma 1.** *For a stable component instance, the long-term average number of products  $g^{com}$  and the relative value function  $V^{com}$  are given by*

$$g^{com} = \frac{\lambda((\sigma + \nu)^2 + \mu\sigma)}{(\sigma + \nu)(\mu\nu - \lambda(\sigma + \nu))},$$

$$V^{com}(x, 0) = \alpha_1 x^2 + \alpha_2 x + \alpha_3 \quad \text{and}$$

$$V^{com}(x, 1) = \alpha_1 x^2 + \alpha_1 x, \tag{11}$$

for  $x \in \mathbb{N}$ , where

$$\alpha_1 = \frac{\sigma + \nu}{2(\mu\nu - \lambda(\sigma + \nu))}, \quad \alpha_2 = \frac{2\mu + \sigma + \nu}{2(\mu\nu - \lambda(\sigma + \nu))},$$

$$\text{and} \quad \alpha_3 = \frac{\lambda\mu}{(\mu\nu - \lambda(\sigma + \nu))(\sigma + \nu)},$$

when taking  $V^{com}(0, 1) = 0$  as a reference value.

**Proof.** One simply verifies by substitution that the solution given in Equation (11) satisfies  $V^{com}(0, 1) = 0$  and the Poisson equations in Equation (10). It is left to show that the above solution is the unique solution. To this end, we use Bhulai and Spieksma (2003, Theorem 6). Suppose that there exists a finite subset of states  $M$  and a weight function  $u : S^{com} \rightarrow \{0, 1\}$ , such that the Markov chain, which satisfies the stability and aperiodicity conditions needed for the theorem to hold, is  $u$ -geometrically recurrent; that is,

$$R_{M,u}(x, w) := \sum_{(x', w') \notin M} \frac{P^{com}((x, w), (x', w'))u(x', w')}{u(x, w)} < 1,$$

for all  $(x, w) \in \mathcal{S}$  and

$$\|c\|_u = \sup_{s \in S^{com}} \frac{|c(s)|}{u(s)} < \infty.$$

Then, this theorem implies that a pair  $(g, V)$  satisfying the Poisson equations (10) is the unique solution when

$$\|V\|_u = \sup_{s \in S^{com}} \frac{|V(s)|}{u(s)} < \infty.$$

To invoke this theorem, we set  $M = \{(0, 0), (0, 1)\}$  and  $u(x, w) = (1 + \delta)^x(1 - \epsilon)^w$ , with

$$\delta \in \left( 0, \frac{\mu + \nu + \sigma - \sqrt{(\lambda - \mu - \nu - \sigma)^2 + 4(\lambda\nu - \mu\nu + \lambda\sigma)}}{2\lambda} \right. \\ \left. - \frac{1}{2} \right)$$

and

$$\epsilon \in \left( \frac{\lambda}{\nu}\delta, \frac{\delta\mu - \lambda\delta(1 + \delta)}{\delta\mu - \lambda\delta(1 + \delta) + \sigma(1 + \delta)} \right).$$

Then, we have that

$$R_{M,u}(x, w) = \lambda(1 + \delta) + w \left( \mu \mathbb{1}_{\{x>1\}} \frac{1}{1 + \delta} + \sigma \frac{1}{1 - \epsilon} \right) \\ + \nu(1 - w)(1 - \epsilon) + (1 - \lambda - w(\mathbb{1}_{\{x>1\}}\mu + \sigma) - (1 - w)\nu).$$

For all  $x \in \mathbb{N}$ , the lower bound on  $\epsilon$  ensures that  $R_{M,u}(x, 0) < 1$ , whereas the upper bound guarantees that  $R_{M,u}(x, 1) < 1$ . The upper bound of  $\delta$  is derived by equating the two bounds of  $\epsilon$  and thus warrants that the lower bound of  $\epsilon$  does not exceed the upper bound of  $\epsilon$ . In its turn, the stability condition  $\lambda < \mu\nu/(\sigma + \nu)$  (see Equation (8)) guarantees that the upper bound of  $\delta$  is positive. Observe that for the assessment of the validity of the conditions  $\|c\|_u < \infty$  and  $\|V^{com}\|_u < \infty$ , the value of  $w$  does not play an essential role, as it can only influence the value of  $u(x, w)$  up to a finite factor  $(1 - \epsilon)$  for any  $x \in \mathbb{N}$ . We clearly have that the cost function  $c(x, w) = x$  satisfies  $\|c\|_u < \infty$ , since it is linear in  $x$ , and the weight function  $u$  is exponential in  $x$ . Likewise, the function  $V^{com}$  as given in Equation (11) satisfies  $\|V^{com}\|_u < \infty$ , since it behaves as a quadratic polynomial as opposed to exponential in  $x$ . Hence, by Bhulai and Spieksma (2003, Theorem 6) the solution given by Equation (11) is the unique solution to the Poisson equations. ■

This concludes the derivation of the relative value function for a component with parameters  $\lambda, \mu, \sigma$  and  $\nu$ .

*Remark 1.* For  $\sigma = 0$  and  $w = 1$ , the component model degenerates to a regular M/M/1 queue. As expected,  $g^{com}$  and  $V^{com}(x, 1)$  then simplify to the well-known expressions

$$g^{M/M/1} = \frac{\lambda}{\mu - \lambda} \quad \text{and} \quad V^{M/M/1}(x) = \frac{1}{2(\mu - \lambda)} x(x + 1).$$

For the general case, we may rewrite

$$V^{com}(x, 1) = \frac{1}{2(\mu\nu/(\sigma + \nu) - \lambda)} x(x + 1).$$

Observe that  $\mu\nu/(\sigma + \nu)$  is the maximum rate at which the server is able to process products in the long term. When interpreting this as an effective service rate, we may conclude that the structure of the relative value function  $V^{com}$  is similar to that of the regular M/M/1 queue.

*Remark 2.* As observed above, the component can be alternatively modeled as a single-server vacation queue with the Coxian completion time  $C$  of a product regarded as the service time and with server vacations occurring exclusively when the queue is empty. As a result, the average costs per time unit, or rather, the average queue length  $g^{com}$  (including any possible product in service) can also be obtained by applying the Fuhrmann–Cooper decomposition (Fuhrmann and Cooper, 1985):

$$g^{com} = \mathbb{E}[V] + \mathbb{E}[L_{M/Cox/1}],$$

where  $\mathbb{E}[V]$  represents the expected queue length when observed during a server vacation (at the start of which there are no products in the queue) and  $\mathbb{E}[L_{M/Cox/1}] = \lambda \mathbb{E}[C] + (\lambda \mathbb{E}[C^2]) / (2(1 - \lambda \mathbb{E}[C]))$  is the well-known expectation of the queue length  $L_{M/Cox/1}$  of a similar single-server queueing system with Poisson ( $\lambda$ ) arrivals and the completion times as service times but excluding any server vacations. The term  $\mathbb{E}[V]$  equals the probability  $\sigma / (\sigma + \nu)$  that a product arriving in an empty system finds the server in a vacation, times the mean number of Poisson ( $\lambda$ ) arrivals during a residual exponentially ( $\nu$ ) distributed vacation time, and thus amounts to  $\lambda \sigma / (\nu(\sigma + \nu))$ . The moments  $\mathbb{E}[C]$  and  $\mathbb{E}[C^2]$  can be determined by studying the relation between the completion time and the service requirement of a product. When substituting these expressions, we obtain  $g^{com}$  as given in Lemma 1.

### 5.1.2. Resulting expression for $V^{sta}$

We now turn back to the relative value function of the model as described in Section 3 under the static policy with parameter  $p$ . As mentioned before, this model consists of two components with rates  $\lambda_1, \mu_1, \sigma_1, p\nu_1$ , and  $\lambda_2, \mu_2, \sigma_2, (1-p)\nu_2$ , respectively. Now that we have found an expression for the relative value functions pertaining to one such component, we readily obtain an expression for the relative value function for the complete system. Combining Equation (9) with Lemma 1 results in the following theorem.

**Theorem 1.** *Given that the stability conditions in Equation (8) are satisfied, the long-term average costs  $g_p^{sta}$  and the relative value function  $V_p^{sta}(x_1, x_2, w_1, w_2)$  corresponding to the static policy with parameter  $p$  are given by*

$$g_p^{sta} = c_1 \frac{\lambda_1((\sigma_1 + p\nu_1)^2 + \mu_1\sigma_1)}{(\sigma + p\nu_1)(\mu_1 p\nu_1 - \lambda_1(\sigma_1 + p\nu_1))} + c_2 \frac{\lambda_2((\sigma_2 + (1-p)\nu_2)^2 + \mu_2\sigma_2)}{(\sigma_2 + (1-p)\nu_2)(\mu_2(1-p)\nu_2 - \lambda_2(\sigma_2 + (1-p)\nu_2))},$$

and

$$V_p^{sta}(x_1, x_2, w_1, w_2) = \alpha_{1,1}c_1x_1^2 + c_1(\alpha_{2,1}(1-w_1) + \alpha_{1,1}w_1)x_1 + \alpha_{3,1}c_1(1-w_1) + \alpha_{1,2}c_2x_2^2 + c_2(\alpha_{2,2}(1-w_2) + \alpha_{1,2}w_2)x_2 + \alpha_{3,2}c_2(1-w_2),$$

for all  $(x_1, x_2, w_1, w_2) \in \mathcal{S}$ , where

$$\begin{aligned} \alpha_{1,1} &= \frac{\sigma_1 + p\nu_1}{2\mu_1 p\nu_1 - \lambda_1(\sigma_1 + p\nu_1)}, \\ \alpha_{1,2} &= \frac{\sigma_2 + (1-p)\nu_2}{2\mu_2(1-p)\nu_2 - \lambda_2(\sigma_2 + (1-p)\nu_2)}, \\ \alpha_{2,1} &= \frac{2\mu_1 + \sigma_1 + p\nu_1}{2\mu_1 p\nu_1 - \lambda_1(\sigma_1 + p\nu_1)}, \\ \alpha_{2,2} &= \frac{2\mu_2 + \sigma_2 + (1-p)\nu_2}{2\mu_2(1-p)\nu_2 - \lambda_2(\sigma_2 + (1-p)\nu_2)}, \\ \alpha_{3,1} &= \frac{\lambda_1\mu_1}{(\mu_1 p\nu_1 - \lambda_1(\sigma_1 + p\nu_1))(\sigma_1 + p\nu_1)}, \quad \text{and} \\ \alpha_{3,2} &= \frac{\lambda_2\mu_2}{(\mu_2(1-p)\nu_2 - \lambda_2(\sigma_2 + (1-p)\nu_2))(\sigma_2 + (1-p)\nu_2)}. \end{aligned}$$

### 5.2. Priority policy

In the previous section, we have explicitly derived the relative value function for the static policy. In Section 6, this policy will act as an initial policy for the one-step policy improvement algorithm to obtain a well-performing heuristic policy. However, for certain instances of the model, there may be no static policy available for which stability holds, whereas the optimal policy does result in stable queues. Then, one-step policy improvement based on the static policy is not feasible, since the initial policy for this procedure must result in a stable system. In these cases, a priority policy may still result in stability and thus be suitable as an initial policy, so that a heuristic policy can still be obtained. For this reason, we study the relative value function of the priority policy in the current section.

Under priority policy  $\pi_i^{prio}$ , the repairman always prioritizes machine  $M_i$ , which we will refer to as the high-priority machine. This means that in case both machines are down, the repairman allocates his/her full capacity to  $M_i$  as a high-priority machine. If there is only one machine non-operational, the repairman dedicates his/her capacity to the broken machine, regardless of whether it is the high-priority machine. In case all machines are operational, the repairman obviously remains idle. Thus, the repairman always takes the action  $((1-w_1), w_1(1-w_2))$  if  $i=1$  or  $((1-w_1)w_2, (1-w_2))$  if  $i=2$ . The priority policy  $\pi_1^{prio}$ , where  $M_1$  acts as the high-priority machine, is stable if and only if for each queue the rate at which products arrive is smaller than the effective service rate of its machine:

$$\lambda_1 < \mu_1 \frac{\nu_1}{\sigma_1 + \nu_1} \quad \text{and} \quad \lambda_2 < \mu_2^{eff}, \quad (12)$$

where  $\mu_2^{eff}$  refers to the effective service rate of  $M_2$ . The right-hand side of the first inequality represents the effective service rate of the high-priority machine  $M_1$  and consists of the actual service rate  $\mu_1$  times the fraction of time  $M_1$  is operational under the priority policy. The effective service

rate of  $M_2$  analogously satisfies

$$\mu_2^{eff} = \mu_2 \frac{1/\sigma_2}{1/\sigma_2 + z/\nu_1 + \mathbb{E}[R_2]}. \quad (13)$$

The expression  $(z/\nu_1) + \mathbb{E}[R_2]$  in the right-hand side represents the expected downtime of  $M_2$ . The constant  $z$  refers to the probability that  $M_2$  observes the repairman busy on  $M_1$  when it breaks down, so that  $z/\nu_1$  represents the expected time  $M_2$  has to wait after its breakdown until the start of its repair as a result of an  $M_1$  failure. The probability  $z$  is computed by the fixed-point equation

$$z = \frac{\sigma_1}{\sigma_1 + \sigma_2} \left( \frac{\sigma_2}{\nu_1 + \sigma_2} + \frac{\nu_1}{\nu_1 + \sigma_2} z \right) \Rightarrow z = \frac{\sigma_1}{\sigma_1 + \sigma_2 + \nu_1}. \quad (14)$$

Likewise,  $\mathbb{E}[R_2]$  represents the expected time from the moment the repairman starts repair on  $M_2$  until its finish and is computed by the fixed-point equation:

$$\begin{aligned} \mathbb{E}[R_2] &= \frac{1}{\sigma_1 + \nu_2} + \frac{\sigma_1}{\sigma_1 + \nu_2} \left( \frac{1}{\nu_1} + \mathbb{E}[R_2] \right) \\ \Rightarrow \mathbb{E}[R_2] &= \frac{1}{\nu_2} + \frac{\sigma_1}{\nu_1 \nu_2}. \end{aligned}$$

By repeating the arguments above, it is easy to see that the priority policy  $\pi_2^{prio}$  is stable if and only if

$$\lambda_1 < \mu_1^{eff} \quad \text{and} \quad \lambda_2 < \mu_2 \frac{\nu_2}{\sigma_2 + \nu_2}, \quad (15)$$

where  $\mu_1^{eff}$  has an expression similar to  $\mu_2^{eff}$  but with indices interchanged.

In the remainder of this section, we study the relative function corresponding to the priority policy, under the assumption that this policy is stable. We will only study the priority policy  $\pi_1^{prio}$  where  $M_1$  acts as the high-priority machine in this section, since results for the other case will follow immediately by similar arguments or simply by interchanging indices. As such, we drop the machine-specific index in this section, so that  $V^{prio}$  actually refers to  $V_1^{prio}$ .

Deriving an expression for the relative value function  $V^{prio}$  of the priority policy is hard. Before, in the case of the static policy, the model could be decomposed into several components that exhibit no interdependence. This allowed us to obtain an explicit expression for  $V^{sta}$ . In contrast, a similar decomposition under the current policy does lead to *interacting* components. The first component, which contains the high-priority machine  $M_1$  and its corresponding queue, acts independently of any other component, since  $M_1$  is not affected by  $M_2$  when accessing repair resources. However,  $M_2$  is affected by  $M_1$ . This interference causes the second component, which contains the other machine and its queue of products, to become dependent on the events occurring in the first component. Therefore, there exist correlations, which make an explicit analysis of  $V^{prio}$  intractable. Nevertheless, we are still able to derive certain characteristics of the relative value function.

When decomposing the model in the same way as was done in Section 5.1, we have, similar to Equation (9), that the long-term average costs  $g^{prio}$  per time unit and the relative value function  $V^{prio}$  pertaining to the priority policy can be written as

$$\begin{aligned} g^{prio} &= c_1 g^{prc} + c_2 g^{nprc} \quad \text{and} \quad V^{prio}(x_1, x_2, w_1, w_2) \\ &= c_1 V^{prc}(x_1, w_1) + c_2 V^{nprc}(x_2, w_1, w_2), \end{aligned} \quad (16)$$

where  $g^{prc}$  and  $V^{prc}(x_1, w_1)$  are the long-term average costs and the relative value function pertaining to the first component, which we will also call the priority component. Similarly,  $g^{nprc}$  and  $V^{nprc}(x_2, w_1, w_2)$  denote the long-term average costs and the relative value function of the second component, which we will also refer to as the non-priority component. In both of these subsystems, the products present are each assumed to incur costs at rate one. Note that the function  $V^{nprc}(x_2, w_1, w_2)$  of the second component now includes  $w_1$  as an argument, since the costs accrued in the second component are now dependent on the state of  $M_1$  in the first component. In Section 5.2.1, we obtain an expression for  $V^{prc}$ . While  $V^{nprc}$  defies an explicit analysis due to the aforementioned dependence, we make several conjectures on its form in Section 5.2.2. In Section 6, it will turn out that these conjectures still allow us to use  $\pi^{prio}$  as an initial policy for the one-step policy improvement algorithm.

### 5.2.1. Relative value function for the priority component

In the priority component, the machine  $M_1$  faces no competition in accessing repair facilities. If  $M_1$  breaks down, the repairman immediately starts repairing  $M_1$  at rate  $\nu_1$ . Thus, from the point of view of  $M_1$ , it is as if  $M_1$  has its own dedicated repairman. Therefore, the priority component behaves completely similar to a component of the static policy studied in Section 5.1.1 but now with  $\lambda_1$ ,  $\mu_1$ ,  $\sigma_1$ , and  $\nu_1$  as product arrival, product service, machine breakdown, and machine repair rates. As a result, we obtain by Lemma 1 that, when products in the queue incur costs at rate one, the long-term average costs  $g^{prc}$  and the relative value function  $V^{prc}$  are given by

$$\begin{aligned} g^{prc} &= \frac{\lambda_1((\sigma_1 + \nu_1)^2 + \mu_1 \sigma_1)}{(\sigma_1 + \nu_1)(\mu_1 \nu_1 - \lambda_1(\sigma_1 + \nu_1))}, \\ V^{prc}(x_1, 0) &= \nu_1 x_1^2 + \nu_2 x_1 + \nu_3, \quad \text{and} \\ V^{prc}(x_1, 1) &= \nu_1 x_1^2 + \nu_1 x_1, \end{aligned} \quad (17)$$

for  $x_1 \in \mathbb{N}$ , where

$$\begin{aligned} \nu_1 &= \frac{\sigma_1 + \nu_1}{2(\mu_1 \nu_1 - \lambda_1(\sigma_1 + \nu_1))}, \\ \nu_2 &= \frac{2\mu_1 + \sigma_1 + \nu_1}{2(\mu_1 \nu_1 - \lambda_1(\sigma_1 + \nu_1))}, \quad \text{and} \\ \nu_3 &= \frac{\lambda_1 \mu_1}{(\mu_1 \nu_1 - \lambda_1(\sigma_1 + \nu_1))(\sigma_1 + \nu_1)}, \end{aligned}$$

when taking  $V^{prc}(0, 1) = 0$  as a reference value.

5.2.2. Heuristic for approximating the relative value function for the non-priority component

As mentioned earlier, the relative value function  $V^{nprc}$  of the non-priority component defies an explicit analysis, due to its dependence on the priority component. Explorative numerical experiments suggest that  $V^{nprc}$  asymptotically behaves like a second-order polynomial in  $x_2$  as  $x_2 \rightarrow \infty$ . We support this insight by arguments from queueing theory, which are given in Conjecture 1 below. Building on this, we also pose certain conjectures on the first-order and second-order coefficients of this polynomial. This leads to a heuristic for approximating the relative value function for the non-priority component, which we present in this section. Finally, we present an approximation for the long-term expected costs  $g^{nprc}$ .

In the non-priority component, products arrive at rate  $\lambda_2$  and are served at rate  $\mu_2$  by  $M_2$  when it is operational. Independently of this,  $M_2$  breaks down at rate  $\sigma_2$  when it is operational. In case  $M_2$  is down, it gets repaired at rate  $\nu_2$  if  $M_1$  is operational, and at rate zero otherwise. Obviously, if  $M_1$  is operational, it breaks down at rate  $\sigma_1$  and gets repaired at rate  $\nu_1$  if it is down. The resulting system can again be formulated as a Markov reward chain with states  $(x_2, w_1, w_2) \in \mathcal{S}^{nprc}$ , representing the number of products in the component ( $x_2$ ), and the indicator variables corresponding to each of the machine's operational states  $(w_1, w_2)$ , where  $\mathcal{S}^{nprc} \in \mathbb{N} \times \{0, 1\}^2$  is its state space. This chain is said to accrue costs in a state-dependent manner at rate  $c(x_2, w_1, w_2) = x_2$ . After uniformization at rate one, the transition probabilities  $P^{nprc}(s, t)$  from a state  $s \in \mathcal{S}^{nprc}$  to a state  $t \in \mathcal{S}^{nprc}$  are given by

$$\begin{aligned} P^{nprc}((x_2, w_1, w_2), (x_2 + 1, w_1, w_2)) &= \lambda_2, \\ P^{nprc}((x_2, w_1, w_2), (x_2 - 1, w_1, w_2)) &= \mu_2 w_2 \mathbb{1}_{\{x_2 > 0\}}, \\ P^{nprc}((x_2, 1, w_2), (x_2, 0, w_2)) &= \sigma_1, \\ P^{nprc}((x_2, w_1, 1), (x_2, w_1, 0)) &= \sigma_2, \\ P^{nprc}((x_2, 0, w_2), (x_2, 1, w_2)) &= \nu_1, \\ P^{nprc}((x_2, 1, 0), (x_2, 1, 1)) &= \nu_2, \quad \text{and} \\ P^{nprc}((x_2, w_1, w_2), (x_2, w_1, w_2)) &= (1 - \lambda_2 - \sigma_1 w_1 - w_2(\mu_2 \mathbb{1}_{\{x_2 > 0\}} + \sigma_2) \\ &\quad - \nu_1(1 - w_1) - \nu_2 w_1(1 - w_2)), \end{aligned}$$

while all other transition probabilities are zero. For this Markov reward chain, the Poisson equations are given by

$$\begin{aligned} g^{nprc} + V^{nprc}(x_2, w_1, w_2) &= x_2 + \lambda_2 V^{nprc}(x_2 + 1, w_1, w_1) \\ &\quad + \mu_2 w_2 V^{nprc}((x_2 - 1)^+, w_1, 1) \\ &\quad + \sigma_1 w_1 V^{nprc}(x_2, 0, w_2) + \sigma_2 w_2 V^{nprc}(x_2, w_1, 0) \\ &\quad + \nu_1(1 - w_1) V^{nprc}(x_2, 1, w_2) \\ &\quad + \nu_2 w_1(1 - w_2) V^{nprc}(x_2, 1, 1) \\ &\quad + (1 - \lambda_2 - \sigma_1 w_1 - w_2(\mu_2 + \sigma_2) - \nu_1(1 - w_1) \\ &\quad - \nu_2 w_1(1 - w_2)) V^{nprc}(x_2, w_1, w_2). \end{aligned} \tag{18}$$

**Conjecture 1.** Assume that the stability conditions in Equation (12) are satisfied. Then, the relative value function  $V^{nprc}(x_2, w_1, w_2)$  of the non-priority component asymptotically behaves as a second-order polynomial in  $x_2$  with second-order coefficient  $\phi_1 = (1/2)(\mu_2^{eff} - \lambda_2)^{-1}$  as  $x_2 \rightarrow \infty$  for each  $w_1, w_2 \in \{0, 1\}$ , where  $\mu_2^{eff}$  represents the effective service rate of  $M_2$  as given in Equation (13).

**Argument.** Recall that  $V^{nprc}(x_2 + 1, w_1, w_2) - V^{nprc}(x_2, w_1, w_2)$  represents the long-term difference in total expected costs accrued in the non-priority component when starting the system in state  $(x_2 + 1, w_1, w_2)$  instead of  $(x_2, w_1, w_2)$ . Since every customer generates costs at rate one per time unit, it is easily seen by a sample path comparison argument that this difference asymptotically (for  $x_2 \rightarrow \infty$ ) amounts to the expected time it takes for the queue to empty when the system is started in the state  $(x_2 + 1, w_1, w_2)$ . For small values of  $x_2$ , this difference may depend slightly on  $w_1$ , since the event of  $M_1$  being down at the start of the process may have a relative significant impact on the time to empty the queue, as the first repair on  $M_2$  is likely to take longer than usual. However, as  $x_2$  becomes larger, the time needed for the queue to empty becomes larger too, so that the process describing the conditions of the machines is more likely to have advanced towards an equilibrium in the meantime. As a result, the initial value of  $w_1$  does not have a relatively significant impact on the difference (i.e., the time for the queue to empty) for larger  $x_2$ -values. In fact, the extra delay in the time to empty imposed by an initial failure of  $M_1$  is expected to converge to a constant as  $x_2$  increases. Based on these observations, we expect that asymptotically, the value  $w_1$  will only appear in front of  $x_2$ -terms (and not higher powers) in  $V^{nprc}(x_2, w_1, w_2)$ . This asymptotic linear effect is studied in Conjecture 2. We also expect that  $V^{nprc}$  starts to exhibit this asymptotic behavior very quickly as  $x_2$  increases, since the process describing the conditions of the machines regenerates every time  $M_2$  is repaired and thus moves to an equilibrium rather quickly.

Now that we have identified the contribution of  $w_1$ , we study the behavior of  $V^{nprc}$  in the direction of  $x_2$  that is not explained by  $w_1$ . When ignoring the interaction with the priority queue (thus ignoring  $w_1$ ), the queue of products in the non-priority component may be interpreted as an M/PH/1 queue, by incorporating the service interruptions (consisting of  $M_1$  and  $M_2$  repairs) into the service times of the products. Thus, queueing-theoretic intuition suggests that the relative value function for our problem may behave similarly to that of the M/PH/1 queue, particularly if the degree of interdependence between the queue lengths of  $Q_1$  and  $Q_2$  is not very high. It is known that the relative value function of such a queue is a quadratic polynomial (see, for example, Bhulai (2006)), so that asymptotically,  $V^{nprc}$  is also likely to behave as a quadratic polynomial. The second-order coefficient of the relative value function

of the M/PH/1 queue satisfies the form  $(1/2)(\mu_2^{eff} - \lambda_2)^{-1}$ , where  $\lambda_2$  is the arrival rate and  $\mu_2^{eff}$  is the effective service rate; i.e., the maximum long-term rate at which the server can process the products. As observed in Remark 1, the second-order coefficient  $\alpha_1$  of the static component in Lemma 1 is also of this form, independent of the value of  $w_2$ . Therefore, it is reasonable to assume that the second-order coefficient of  $V^{nprc}$  also satisfies this form, independent of the values for  $w_1, w_2$ . The involved effective service rate of  $M_2, \mu_2^{eff}$ , is given in Equation (13). By combining all arguments above, the conjecture follows. ■

Note that the first-order coefficient of the polynomial, unlike the second-order coefficient, is expected to be dependent on  $w_1$  as mentioned in the argument of Conjecture 1 but also on  $w_2$ , in line with the results on the components of the static policy. The first-order coefficient is studied in the next conjecture.

**Conjecture 2.** *Suppose that Conjecture 1 holds true, so that asymptotically*

$$\begin{aligned} V^{nprc}(x_2, 0, 0) &= \phi_1 x_2^2 + \phi_2 x_2 + \phi_3, \\ V^{nprc}(x_2, 1, 0) &= \phi_1 x_2^2 + \psi_2 x_2 + \psi_3, \\ V^{nprc}(x_2, 0, 1) &= \phi_1 x_2^2 + \chi_2 x_2 + \chi_3, \quad \text{and} \\ V^{nprc}(x_2, 1, 1) &= \phi_1 x_2^2 + \omega_2 x_2 + \omega_3, \end{aligned} \quad (19)$$

as  $x_2 \rightarrow \infty$ . Then,  $\psi_2 = \phi_2 - \Delta_{1,0}$ ,  $\chi_2 = \phi_2 - \Delta_{0,1}$ ,  $\omega_2 = \phi_2 - \Delta_{1,1}$ , where

$$\begin{aligned} \Delta_{0,1} &= \frac{\mu_2(v_1 + \sigma_1)(v_1 + v_2 + \sigma_1 + \sigma_2)}{\mu_2 v_1 v_2 (v_1 + \sigma_1 + \sigma_2) - \lambda_2 (v_1 + \sigma_1)(v_1 (v_2 + \sigma_2) + \sigma_2 (v_2 + \sigma_1 + \sigma_2))}, \\ \Delta_{1,0} &= \frac{\mu_2 v_2 (v_1 + \sigma_1 + \sigma_2)}{\mu_2 v_1 v_2 (v_1 + \sigma_1 + \sigma_2) - \lambda_2 (v_1 + \sigma_1)(v_1 (v_2 + \sigma_2) + \sigma_2 (v_2 + \sigma_1 + \sigma_2))}, \quad \text{and} \\ \Delta_{1,1} &= \frac{\mu_2 (v_1 + v_2 + \sigma_1)(v_1 + \sigma_1 + \sigma_2)}{\mu_2 v_1 v_2 (v_1 + \sigma_1 + \sigma_2) - \lambda_2 (v_1 + \sigma_1)(v_1 (v_2 + \sigma_2) + \sigma_2 (v_2 + \sigma_1 + \sigma_2))}. \end{aligned}$$

**Argument.** The relative value function  $V^{nprc}$  is expected to satisfy the Poisson equations (18), also asymptotically for  $x_2 \rightarrow \infty$ . When substituting Equation (19) into Equation (18) for  $x_2 > 0$ , the constraints on  $\phi_2, \chi_2, \psi_2$ , and  $\omega_2$  mentioned above are necessary for the first-order terms in  $x_2$  on both sides of the equations to be equal. ■

**Remark 3.** As costs in the non-priority component are generated primarily by having customers in the queue, we expect the values of  $\phi_3, \chi_3, \psi_3$ , and  $\omega_3$  in Equation (19) to be of very moderate significance compared with the second-order and first-order coefficients. As mentioned before, we also expect that  $V^{nprc}$  starts to exhibit its asymptotic behavior very quickly as  $x_2$  increases. Although we have not found an explicit solution for the first-order coefficients  $\phi_2, \chi_2, \psi_2$ , and  $\omega_2$ , we can still obtain accurate approximations for expressions such as  $V^{prio}(x_1, x_2, 1, 0) - V^{prio}(x_1, x_2, 0, 0)$  and  $V^{prio}(x_1, x_2, 0, 1) - V^{prio}(x_1, x_2, 0, 0)$  based on the information we have obtained. In particular, we have by combining the results in Equations (16) and (17), Conjecture 1 and

Conjecture 2 that

$$\begin{aligned} V^{prio}(x_1, x_2, 1, 0) - V^{prio}(x_1, x_2, 0, 0) &\approx c_1((v_1 - v_2)x_1 - v_3) - c_2 \Delta_{1,0} x_2, \\ V^{prio}(x_1, x_2, 0, 1) - V^{prio}(x_1, x_2, 0, 0) &\approx -c_2 \Delta_{0,1} x_2, \end{aligned} \quad (20)$$

with the parameters  $v_1, v_2, v_3, \Delta_{1,0}$ , and  $\Delta_{0,1}$  as defined in Section 5.2.1 and Conjecture 2, respectively. These two accurate approximations allow us to apply the one-step policy improvement algorithm based on the priority policy in Section 6.1.2.

In the two conjectures above, we have not studied the long-term expected costs per time unit  $g^{nprc}$ . However, to predict which of the two possible priority policies  $\pi_1^{prio}$  and  $\pi_2^{prio}$  will lead to the best one-step improved policy, we will need an expression for the overall long-term average costs  $g^{prio}$ , which includes the costs  $g^{nprc}$  generated by the non-priority queue. Therefore, we end this section by deriving an approximation for  $g^{prio}$ , which is obtained by combining Equations (16) and (17) with an independence argument.

**Approximation 1.** An accurate approximation for the long-term expected costs per time unit  $g^{prio}$  is given by

$$g_{app}^{prio} \approx c_1 \frac{\lambda_1((\sigma_1 + v_1)^2 + \mu_1 \sigma_1)}{(\sigma_1 + v_1)(\mu_1 v_1 - \lambda_1(\sigma_1 + v_1))} + c_2 g_{app}^{nprc}, \quad (21)$$

where

$$g_{app}^{nprc} = \lambda_2 \mathbb{E}[\Gamma_{app}] + \frac{\lambda_2^2 \mathbb{E}[\Gamma_{app}^2]}{2(1 - \lambda_2 \mathbb{E}[\Gamma_{app}])} + \lambda_2 \frac{\sigma_2 \mathbb{E}[D^2]}{2(1 + \sigma_2 \mathbb{E}[D])}, \quad (22)$$

$$\mathbb{E}[\Gamma_{app}^i] = (-1)^i \frac{d^i}{ds^i} \left( \frac{\mu_2}{\mu_2 + s + \sigma_2(1 - \tilde{D}(s))} \right) \Big|_{s=0},$$

$$\mathbb{E}[D^i] = (-1)^i \frac{d^i}{ds^i} \tilde{D}(s) \Big|_{s=0} \quad \text{and}$$

$$\tilde{D}(s) = \left( (1 - z) + z \frac{v_1}{v_1 + s} \right) \frac{v_2}{v_2 + s + \sigma_1(1 - v_1/(v_1 + s))},$$

with  $z$  as defined in Equation (14).

**Justification.** The form of Equation (21) is a consequence of Equations (16) and (17). It thus remains to obtain an approximation for  $g^{nprc}$ . We do this by ignoring the interaction between the two components. Inspired by Remark 2, we approximate  $g^{nprc}$  by studying the queue length in an M/G/1 queue with server vacations, with completion times denoted by  $\Gamma$  instead of service times, which incorporate the time lost due to service interruptions as a result of a breakdown of  $M_2$  during service. The server vacations, which start each time the queue becomes empty, include the down-periods of  $M_2$  following a breakdown occurring when the queue is empty. Let  $\tilde{D}(s) = \mathbb{E}[e^{-sD}]$  be the Laplace–Stieltjes Transform (LST) of the duration  $D$  of an  $M_2$  down-period. This period  $D$  consists of an exponential ( $v_2$ ) repair time  $R_2$  with LST  $\tilde{R}_2(s) = v_2/(v_2 + s)$  and a Poisson ( $\sigma_1 R_2$ ) number of interruptions  $N$ , each caused by a breakdown of  $M_1$ . Since  $M_1$  has priority, these interruptions take an exponential

( $\nu_1$ ) repair time  $R_1$  with LST  $\tilde{R}_1(s) = \nu_1/(\nu_1 + s)$ . Finally, when  $M_2$  breaks down, it will have to wait for an  $M_1$ -repair to finish before repair on  $M_2$  can start with probability  $z$  as defined in Equation (14). Since the repair time of  $M_1$  is memoryless, this waiting time also has the LST  $\tilde{R}_1(s)$ . Thus, we have that

$$\begin{aligned} \tilde{D}(s) &= ((1 - z) + z\tilde{R}_1(s)) \int_{t=0}^{\infty} e^{-st} \\ &\quad \times \left( \sum_{n=0}^{\infty} \tilde{R}_1^n(s) \mathbb{P}(N = n) \right) \nu_2 e^{-\nu_2 t} dt \\ &= ((1 - z) + z\tilde{R}_1(s)) \int_{t=0}^{\infty} e^{-st} \\ &\quad \times \left( \sum_{n=0}^{\infty} e^{-\sigma_1 t} \frac{(\sigma_1 t \tilde{R}_1(s))^n}{n!} \right) \nu_2 e^{-\nu_2 t} dt \\ &= ((1 - z) + z\tilde{R}_1(s)) \tilde{R}_2(s + \sigma_1(1 - \tilde{R}_1(s))) \\ &= \left( (1 - z) + z \frac{\nu_1}{\nu_1 + s} \right) \frac{\nu_2}{\nu_2 + s + \sigma_1(1 - \nu_1/(\nu_1 + s))}. \end{aligned}$$

The completion time  $\Gamma$  of a product, of which the distribution is represented by its LST  $\tilde{\Gamma}(s)$ , consists of an exponentially ( $\mu_2$ ) distributed service time  $B_2$  with LST  $\tilde{B}_2(s) = \mu_2/(\mu_2 + s)$  and a Poisson ( $\sigma_2 B_2$ ) number of interruptions, each caused by a breakdown of  $M_2$ . Due to the interaction between the components, we know that at the start of the first completion time after a vacation period, both machines are operational, biasing the duration of the first breakdown of  $M_2$  after that time. When ignoring this effect of interaction, and assuming that each breakdown has a duration distributed according to  $D$ , we obtain similar to the computations above that

$$\tilde{\Gamma}(s) \approx \tilde{B}_2(s + \sigma_2(1 - \tilde{D}(s))). \tag{23}$$

An application of the Fuhrmann–Cooper decomposition similar to Remark 2 suggests that

$$\begin{aligned} g^{nprc} &\approx \mathbb{E}[L_{M/G/1}] + \mathbb{E}[L_{vac}], \\ g^{nprc} &\approx \mathbb{E}[L_{M/G/1}] + \mathbb{E}[L_{vac}], \end{aligned}$$

where

$$\mathbb{E}[L_{M/G/1}] = \lambda \mathbb{E}[\Gamma] + \frac{\lambda_2 \mathbb{E}[\Gamma^2]}{2(1 - \lambda_2 \mathbb{E}[\Gamma])}$$

is the mean queue length of the number of products in an M/G/1 queue with Poisson ( $\lambda_2$ ) arrivals and service times distributed according to the completion times  $\Gamma$ . Approximations for the moments of  $\Gamma$  follow by differentiation of Equation (23) with respect to  $s$ . The term  $\mathbb{E}[L_{vac}]$  represents the expected queue length observed when the server is on a vacation, which is initiated any time the queue empties. This vacation lasts until the time of the next product arrival in case  $M_2$  is operational at that point or otherwise, until the first time a repair of  $M_2$  completes after that time. When conditioning on the former case, the expected queue

length as observed by a vacation is obviously zero. When conditioning on the latter, this expectation resolves to the expected number of Poisson ( $\lambda_2$ ) arrivals during the past part of a down-period  $D$ . The duration of this past part has expectation  $\mathbb{E}[D^2]/2\mathbb{E}[D]$ , of which the moments of  $D$  can be computed by differentiation of  $\tilde{D}(s)$  with respect to  $s$ . The probability of the latter event occurring is hard to compute, as the durations of the vacation periods are weakly interdependent, due to the interaction between the components. By ignoring this dependence, we have by a renewal argument that this probability is well approximated by  $\mathbb{E}[D]/((1/\sigma_2) + \mathbb{E}[D])$ , where  $\frac{1}{\sigma_2}$  is the expected duration of an up-period of  $M_2$ . As a result

$$\mathbb{E}[L_{vac}] \approx \lambda_2 \frac{\mathbb{E}[D^2]}{2\mathbb{E}[D]} \frac{\mathbb{E}[D]}{1/\sigma_2 + \mathbb{E}[D]} = \lambda_2 \frac{\sigma_2 \mathbb{E}[D^2]}{2(1 + \sigma_2 \mathbb{E}[D])}.$$

By combining the results above, we obtain the approximation  $g_{app}^{nprc}$  as given in Equation (22). Note that the application of the Fuhrmann–Cooper decomposition requires that the completion times are mutually independent. However, in our case, this requirement is not met, again due to the interaction between the components. For example, a very long completion time may imply that the last actual service period of  $M_2$  has been longer than usual. This in its turn implies that  $M_2$  has been in operation for some time, so that if a  $M_2$ -breakdown occurs in the next completion time, it is likelier than usual that  $M_1$  is also down at that point. This obviously has effects on the completion time. Due to this interdependence, the application of the Fuhrmann–Cooper decomposition also results in a computation error. However, all computation errors made share the same source, namely, the interaction between the components, and in particular the role of  $M_1$ . As we already saw in Conjecture 1, the influence of  $M_1$  on the relative value function is limited, especially for states with a large number of products in the queue. Therefore, we expect this approximation to be accurate, especially for the purpose of deciding which of the two priority policies available performs best (see also Remark 9). ■

### 6. Derivation of a near-optimal policy

Based on the explicit expressions for the relative value functions of the static policy and the priority policy as obtained in the previous section, we derive a nearly optimal dynamic policy. We do so by applying the one-step policy improvement method on both the static policy and the priority policy in Section 6.1. The resulting improved policies  $\pi^{oss}$ ,  $\pi_1^{osp}$ , and  $\pi_2^{osp}$  can then be used to construct a nearly optimal policy, as discussed in Section 6.2. By construction, this near-optimal policy is applicable in a broader range of parameter settings than each of the improved policies separately.

### 6.1. One-step policy improvement

One-step policy improvement is an approximation method that is distilled from the policy iteration algorithm in Markov decision theory. In policy iteration, one starts with an arbitrary policy  $\pi^{init}$  for which the relative value function  $V^{init}$  is known. Next, using these values, an improved policy  $\pi^{imp}$  can be obtained by performing a policy improvement step:

$$\pi^{imp}(s) = \arg \min_{a \in \mathcal{A}_s} \left\{ \sum_{s' \in \mathcal{S}} P_a(s, s') V^{init}(s') \right\}; \quad (24)$$

i.e., the minimizing action of  $K^{init}(s)$  as defined in Equation (3). If  $\pi^{imp} = \pi^{init}$ , the optimal policy has been found. Otherwise, the procedure can be repeated with the improved policy by setting  $\pi^{init} := \pi^{imp}$ , generating a sequence converging to the optimal policy. However, as the relative value function of the improved policy may not be known explicitly, subsequent iterations may have to be executed numerically. To avoid this problem, the one-step policy improvement consists in executing the policy improvement step only once. In this case, the algorithm starts with a policy for which an expression for the relative value function is known. The resulting policy is then explicit and can act as a basis for approximation of the optimal policy. We now derive two one-step improved policies based on the results of the static policy and the priority policy as obtained in Section 5.

#### 6.1.1. One-step policy improvement based on the static policy

In Section 5.1, we found the relative value function  $V^{sta}$  for the class of static policies, in which each policy corresponds to a splitting parameter  $p \in (0, 1)$ . As an initial policy for the one-step policy improvement, we take the policy which already performs best within this class with respect to the weighted number of products in the system. Thus, we take as an initial policy the static policy with splitting parameter:

$$p^{oss} = \arg \min_p \{g_p^{sta} : p \in \mathcal{P}\}, \quad (25)$$

with  $g_p^{sta}$  as defined in Theorem 1 and where  $\mathcal{P} \subset (0, 1)$  is the set of splitting parameters that satisfy the stability conditions in Equation (8). Then, by performing one step of policy improvement as given in Equation (24), we obtain

$$\begin{aligned} \pi^{oss}(x_1, x_2, w_1, w_2) &= \arg \min_{(q_1, q_2) \in \mathcal{A}_{(x_1, x_2, w_1, w_2)}} \left\{ q_1 v_1 (V_{p^{oss}}^{sta}(x_1, x_2, 1, w_2) \right. \\ &\quad - V_{p^{oss}}^{sta}(x_1, x_2, 0, w_2)) + q_2 v_2 (V_{p^{oss}}^{sta}(x_1, x_2, w_1, 1) \\ &\quad \left. - V_{p^{oss}}^{sta}(x_1, x_2, w_1, 0)) \right\}. \end{aligned} \quad (26)$$

It is easily seen that  $V_{p^{oss}}^{sta}(x_1, x_2, 1, w_2) - V_{p^{oss}}^{sta}(x_1, x_2, 0, w_2)$  and  $V_{p^{oss}}^{sta}(x_1, x_2, w_1, 1) - V_{p^{oss}}^{sta}(x_1, x_2, w_1, 0)$  are non-positive for any  $(x_1, x_2, w_1, w_2) \in \mathcal{S}$  by observing that  $\alpha_{2,i} \geq \alpha_{1,i}$  and  $\alpha_{3,i} \geq 0$ ,  $i = 1, 2$ . This means that  $\pi^{oss}$  satisfies the

properties mentioned in Section 4.1. Therefore, we can simplify Equation (26) to

$$\begin{aligned} \pi^{oss}(x_1, x_2, w_1, w_2) &= \arg \min_{(q_1, q_2) \in \{(1-w_1, 0), (0, 1-w_2)\}} \left\{ q_1 v_1 (V_{p^{oss}}^{sta}(x_1, x_2, 1, w_2) \right. \\ &\quad - V_{p^{oss}}^{sta}(x_1, x_2, 0, w_2)) + q_2 v_2 (V_{p^{oss}}^{sta}(x_1, x_2, w_1, 1) \\ &\quad \left. - V_{p^{oss}}^{sta}(x_1, x_2, w_1, 0)) \right\}. \end{aligned}$$

Substituting  $V_{p^{oss}}^{sta}$  as obtained in Theorem 1 in this expression yields the following one-step improved policy:

$$\begin{aligned} \pi^{oss}(x_1, x_2, w_1, w_2) &= \begin{cases} (0, 0) & \text{if } w_1 = w_2 = 1, \\ (1, 0) & \text{if } w_1 = 1 - w_2 = 0, \text{ or if } w_1 w_2 = 0 \text{ and} \\ & c_1 v_1 ((\alpha_{1,1} - \alpha_{2,1})x_1 - \alpha_{3,1}) \\ & \leq c_2 v_2 ((\alpha_{1,2} - \alpha_{2,2})x_2 - \alpha_{3,2}), \\ (0, 1) & \text{otherwise,} \end{cases} \end{aligned} \quad (27)$$

for  $(x_1, x_2, w_1, w_2) \in \mathcal{S}$ , where expressions for the  $\alpha$ -coefficients are obtained by substituting the value for  $p$  in the expressions given in Theorem 1 by its optimized counterpart  $p^{oss}$ . Thus, whenever both machines are not operational, the repairman repairs the machine  $M_i$  for which  $c_i v_i ((\alpha_{1,i} - \alpha_{2,i})x_i - \alpha_{3,i})$  is smallest, when adhering to  $\pi^{oss}$ .

*Remark 4.* If  $\mathcal{P}$  is empty, there is no static policy available that results in a system with stable queues. In such circumstances, the static policy cannot be used as an initial policy for the one-step policy improvement approach. The priority policy as studied in Section 5.2 may, however, still result in a stable system. If this is the case, the priority policy may act as an initial policy for the one-step policy improvement method, as explained in the next section.

*Remark 5.* Whenever  $\mathcal{P}$  is not empty, the optimal splitting parameter  $p^{oss}$  is guaranteed to exist. As  $g_p^{sta}$  is a continuous function in  $p$  for  $p \in \mathcal{P}$ , the optimal splitting parameter  $p^{oss}$  is then a root of the derivative of  $g_p^{sta}$  with respect to  $p$  in the domain  $\mathcal{P}$ . This derivative, which forms a sixth-order polynomial in  $p$ , defies the possibility of deriving an explicit expression for  $p^{oss}$ . For implementation purposes, however, this poses no significant problems, as such roots can be found numerically up to arbitrary precision with virtually no computation time needed.

#### 6.1.2. One-step policy improvement based on the priority policy

Although an explicit expression for the relative value function  $V_i^{prio}$  is not tractable, we have identified enough of its characteristics in Section 5.2 to allow the use of a priority policy as an initial policy. We now show how to compute the one-step improved policy  $\pi_1^{osp}$  based on the priority policy  $\pi_1^{prio}$ ; i.e., the priority policy where  $M_1$  is the high-priority machine. Of course, the one-step improved policy

$\pi_2^{osp}$  based on the priority policy  $\pi_2^{prio}$  again follows by interchanging indices in the expressions below. The improvement step as given in Equation (24) implies, after performing the same simplification as in the case of the static policy, that

$$\begin{aligned} \pi_1^{osp}(x_1, x_2, w_1, w_2) &= \arg \min_{(q_1, q_2) \in \{(1-w_1, 0), (0, 1-w_2)\}} \{q_1 v_1 (V_1^{prio}(x_1, x_2, 1, w_2) \\ &\quad - V_1^{prio}(x_1, x_2, 0, w_2)) + q_2 v_2 (V_1^{prio}(x_1, x_2, w_1, 1) \\ &\quad - V_1^{prio}(x_1, x_2, w_1, 0))\}. \end{aligned} \quad (28)$$

The simplification is justified by the fact that  $V_1^{prio}(x_1, x_2, 1, w_2) - V_1^{prio}(x_1, x_2, 0, w_2)$  and  $V_1^{prio}(x_1, x_2, w_1, 1) - V_1^{prio}(x_1, x_2, w_1, 0)$  are obviously non-positive, since also under the priority policy it is always beneficial for the system to have a machine operational. Due to this, it is clear that  $\pi_1^{osp}(x_1, x_2, w_1, w_2)$  in Equation (28) resolves to  $((1-w_1), (1-w_2))$  in the case where  $w_1 = w_2 = 1$ ,  $w_1 = 1-w_2 = 1$ , or  $1-w_1 = w_2 = 1$ . For the case  $w_1 = w_2 = 0$ , expressions for  $V_1^{prio}(x_1, x_2, 1, 0) - V_1^{prio}(x_1, x_2, 0, 0)$  and  $V_1^{prio}(x_1, x_2, 0, 1) - V_1^{prio}(x_1, x_2, 0, 0)$  are, however, not available. Due to their general intractability, we use the approximations for these differences as derived in Equation (20) instead. By plugging these approximations into Equation (28) in case  $w_1 = w_2 = 0$ , we obtain, with a slight abuse of notation, that

$$\begin{aligned} \pi_1^{osp}(x_1, x_2, w_1, w_2) &= \begin{cases} (1, 0) & \text{if } w_1 = 1 - w_2 = 0, \text{ or if } w_1 w_2 = 0 \text{ and} \\ & v_1(c_1((v_1 - v_2)x_1 - v_3) - c_2 \Delta_{1,0} x_2) \\ & \leq -c_2 \Delta_{0,1} v_2 x_2, \\ (0, 1) & \text{otherwise,} \end{cases} \end{aligned} \quad (29)$$

where the parameters  $v_1, v_2, v_3, \Delta_{1,0}$ , and  $\Delta_{0,1}$  are as defined in Section 5.2.1 and Conjecture 2, respectively.

*Remark 6.* We have based  $\pi^{osp}$  on approximations for the relative value function  $V^{prio}$ , rather than exact expressions. Nevertheless, we have already argued in Section 5.2.2 that these approximations are accurate. Moreover, the argmin-operator in Equation (28) only checks which of the two arguments is smallest. Therefore, the improvement step is very robust against approximation errors, especially since both arguments share the same source of approximation error.

### 6.2. Resulting near-optimal policy

In the previous section, we have constructed the improved policies  $\pi^{oss}, \pi_1^{osp}$ , and  $\pi_2^{osp}$  based on the static policy and the priority policy. However, the question remains about which of these policies should be followed by the repairman, given a particular case of the model. In this section, we suggest a

near-optimal policy, which chooses one of the three policies based on the model parameters. To this end, we now inspect these improved policies as well as their initial policies.

First, we observe that each of the improved policies satisfy the structural properties of the optimal policy. The improved policies  $\pi^{oss}, \pi_1^{osp}$ , and  $\pi_2^{osp}$  each instruct the repairman to work at full capacity whenever at least one of the machines is down and therefore satisfy the non-idling property as derived in Section 4.1. Furthermore, when both of the machines are down, the improved policies base the action on threshold curves (or, in this case, threshold lines), so that they also satisfy the properties discussed in Section 4.2. As each of the improved policies satisfy the required properties, we base the decision on which of the improved policies to follow on their respective initial policies.

In terms of feasibility, the static policy and the priority policy complement each other. For any model instance, one can construct an improved static policy, if there exists a static policy that results in stable queues; i.e., there exists a value  $p \in (0, 1)$  such that Equation (8) holds. Similarly, an improved priority policy can be constructed if either Equation (12) or Equation (15) holds. There are cases of the model for which there is no stable static policy, whereas a stable priority policy exists. There are also cases for which the reverse holds true. In these cases, it is clear whether to use an improved static policy or an improved priority policy as a near-optimal policy. However, in case both of the approaches are feasible, other characteristics of the improved policies need to be taken into account.

A static policy would have been optimal if the repairman would have no information to base his/her decision on (i.e., he/she has no knowledge about the queue lengths and the state of the machines). It is not optimal in the current model, however, as it does not have the non-idling property; i.e., the server works only at partial capacity when exactly one of the machines is down. However, this problem does not arise in the improved version of the static policy. Furthermore, if the amount of load presented to the system would be such that the queues of products are never exhausted, it is easily seen that the optimal policy is in the class of priority policies. In such a case, the possibility of having a machine in an operational but idle state then disappears, so that the optimal policy always gives priority to one machine over the other, due to faster service of products, a slower breakdown, faster repair times, or a higher cost rate. We therefore expect the priority policy (and thus also its improved version) to work particularly well in our model when the model parameters are skewed in the favor of repair of a certain machine and when the queues of the products are particularly heavily loaded such that the machines are almost never idling. The performance of the improved static policy, however, is not expected to be as sensitive to the load of the system, since the static policy balances the repair fractions based on, among other things, the load offered to each of the queues.

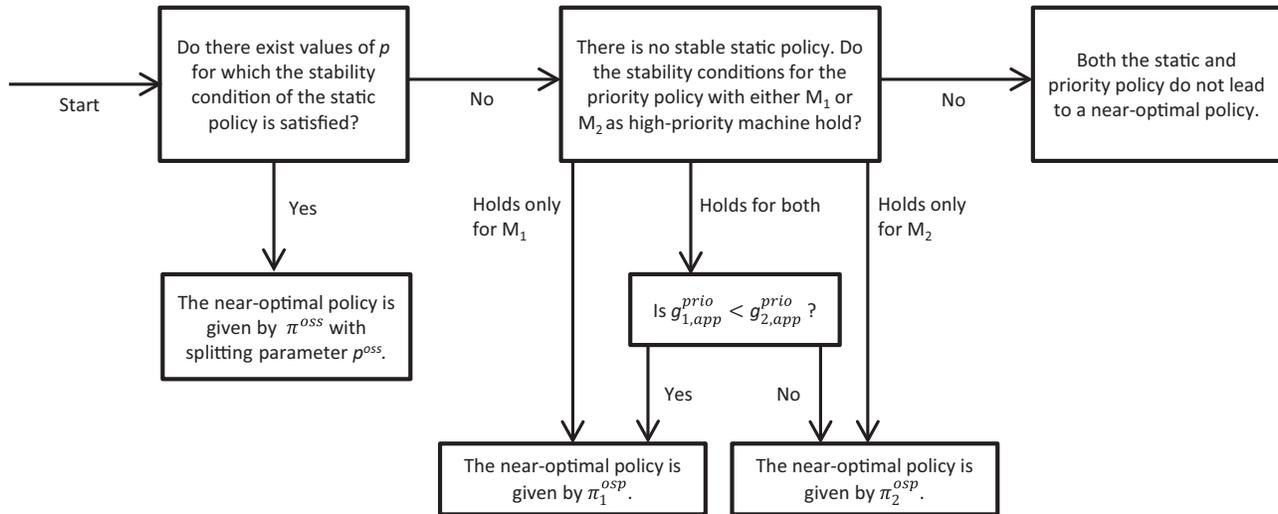


Fig. 3. Schematic representation of the near-optimal policy.

Based on the observations above, we suggest a near-optimal policy that is expressed in terms of a few simple decision rules. A schematic representation of this near-optimal policy is given in Fig. 3. This near-optimal policy prescribes to follow the improved static policy as derived in Section 6.1.1, if there is a static policy available that results in a stable system (i.e., when there exists a  $p$  for which Equation (8) holds). Otherwise, an improved priority policy should be followed, provided that a stable priority policy exists. In the case where only one of the priority policies is stable (i.e., either only Equation (12) or only Equation (15) holds), the choice of which improved priority policy to follow is easy. When both of them are stable, the choice is based on which of the two initial priority policies are expected to perform best. That is, the near-optimal policy will then select the improved policy corresponding to policy  $\pi_1^{prio}$ , if its approximated long-term average costs  $g_{1,app}^{prio}$  as given in Approximation 1 is smaller than its equivalent  $g_{2,app}^{prio}$  obtained by interchanging the indices in Equation (21). Observe that this near-optimal policy is applicable in a wider range of parameter settings than each of the improved policies  $\pi^{oss}$ ,  $\pi_1^{osp}$ , and  $\pi_2^{osp}$  separately.

We end this section with several remarks concerning the obtained nearly optimal policy.

*Remark 7.* As the nearly optimal policy requires a stable static policy or a stable priority policy as a basis for one-step improvement, the approach only works when either Equation (8) (for some value of  $p \in (0, 1)$ ), Equation (12), or Equation (15) holds. However, in theory, it is possible for some parameter settings that none of these conditions are satisfied, whereas stable policies do actually exist. Recall that necessary conditions for the existence of stable policies are given in Equation (1). One can, however, reason that the parameter region where this occurs is fairly

small. First, it is trivially seen that the stability condition (8) for the static policy only reasonably differs from Equation (1) when the breakdowns are large compared with the repair rates. In practice, however, breakdown rates are often much smaller than repair rates. Furthermore, for the priority policy, the conditions for  $\lambda_1$  in Equation (12) and  $\lambda_2$  in Equation (15) coincide with the requirements given in Equation (1) for  $\lambda_1$  and  $\lambda_2$ , respectively. Thus, the parameter region where our approach does not work only captures values for which both  $\lambda_1$  and  $\lambda_2$  are close to their boundary values of  $\mu_1 v_1 / (\sigma_1 + v_1)$  and  $\mu_2 v_2 / (\sigma_2 + v_2)$  at the same time. Finally, we observe that Equation (1) only presents necessary conditions for the existence of a stable policy but does not provide sufficient conditions. As such, the size of this parameter region is limited even further.

*Remark 8.* Many optimization approaches in Markov decision theory suffer from the curse of dimensionality. When dimensions are added to the state space—e.g., by adding more machines to the problem—the size of the state space increases considerably, so that numerical computation techniques break down due to time and resource constraints. Note, however, that the approach presented in this article generally scales well in the number of machines and the corresponding queues of products. The one-step improved policy based on the static policy is easily modified to allow for models with  $N > 2$  machines, since a decomposition of the system in the fashion of Equation (9) can then be done into  $N$  components. After finding a vector of splitting parameters  $(p_1^{oss}, p_2^{oss}, \dots, p_N^{oss})$ , the execution of the one-step policy improvement algorithm will then still result into a simple decision rule similar to Equation (27). Likewise, the priority policy may be used to derive near-optimal policies in a model with larger dimensions. The current approximation for the relative value function  $V^{prio}$  in the case of

$N = 2$  already accounts for the components containing the two most prioritized machines in a model with  $N > 2$  machines, as the repair capacity assigned to a machine is not affected by the breakdown of a machine with lower priority. When approximations for the relative value function pertaining to lower-prioritised components can be found, a nearly optimal policy follows similarly to the case of  $N = 2$ .

### 7. Numerical study

In this section, we numerically assess the performance of the near-optimal policy obtained in Section 6 with respect to the optimal policy. We do this by comparing the average costs per time unit of both policies applied to a large number of model instances. To ensure that there is heavy competition between the machines for the resources of the repairman, we study instances with breakdown rates that are roughly of the same order as the repair rates. In these cases, the event that both machines are in need of repair is not a rare one, which allows us to compare the performance of the near-optimal policy to that of the optimal policy. We will see that the near-optimal policy performs very well over a wide range of parameter settings. Moreover, we observe several parameter effects. Throughout, we also give results for the improved static and priority policies (insofar as they exist), in order to observe how the near-optimal policy compares to these policies in terms of performance.

The complete test bed of instances that are analysed contains all 2916 possible combinations of the parameter values listed in Table 1. This table lists multiple values for the cost weights of having products in  $Q_1$  and  $Q_2$  ( $c_1$  and  $c_2$ ), the service rates at which  $M_1$  and  $M_2$  serve products when operational ( $\mu_1$  and  $\mu_2$ , their breakdown rates ( $\sigma_1$  and  $\sigma_2$ ) as well as their repair rates ( $\nu_1$  and  $\nu_2$ ). Finally, the product arrival rates  $\lambda_1$  and  $\lambda_2$  are specified by the values of the parameters  $\rho_1^{FCFS}$  and  $\rho_2^{FCFS}$  given in the table, where  $\rho_i^{FCFS}$  represents the workload offered to  $M_i$ , if the repairman would repair the machines in a First-Come-First-Serve (FCFS) manner. More specifically, the arrival rates are taken such that the values of the workload

$$\rho_i^{FCFS} = \frac{\lambda_i}{\mu_i} \frac{1}{\theta_i^{FCFS}}$$

would coincide with those given in Table 1, if the repairman were to follow a FCFS policy. In this expression,  $\theta_i^{FCFS}$  represents the fraction of time that  $M_i$  is operational under a FCFS policy, and can be obtained by modeling the second layer of the model (cf. Fig. 1) as a continuous-time Markov chain with a finite state space. The values for  $\rho_i^{FCFS}$ ,  $\sigma_i$ , and  $\nu_i$  are varied in the order of magnitude through the values  $a_i^\rho$ ,  $a_i^\sigma$ , and  $a_i^\nu$  as specified in the table and in the imbalance through the values  $b_j^\rho$ ,  $b_j^\sigma$ , and  $b_j^\nu$ . For example, the workload values ( $\rho_1^{FCFS}$ ,  $\rho_2^{FCFS}$ ) run from  $(0.25 \times 2/3, 0.25 \times 4/3) = (1/6, 1/3)$ , being small

and putting the majority of the load on the second queue, to  $(0.75 \times 4/3, 0.75 \times 2/3) = (1, 0.5)$ , being large and putting the majority on the first queue. Observe that in the latter case,  $\rho_1^{FCFS}$  takes the value of one. Thus, we also consider cases where not all of the queues would be stable, if the repairman would repair the machines in a FCFS fashion.

For the systems corresponding to each of the parameter combinations in Table 1, it turns out that there is always at least one static policy or priority policy available as an initial policy, so that the near-optimal policy is feasible. We numerically compute the average costs  $g^{n-opt}$  incurred per time unit by the system if the repairman would follow the near-optimal policy as suggested in Section 6.2, as well as the average costs  $g^{opt}$  incurred per time unit if the repairman would follow the optimal policy. We do this by using the value iteration algorithm (see, for example, Puterman (1994)). Subsequently, we compute the relative difference  $\Delta^{n-opt}$  between these approximations; that is:

$$\Delta^{n-opt} = 100\% \times \frac{g^{n-opt} - g^{opt}}{g^{opt}}$$

For instances where the corresponding initial policy exists, we also compute the relative differences of the improved policies considered in this paper. That is, we compute similarly defined relative differences  $\Delta^{oss}$  and  $\Delta^{osp}$  for the improved static policy and the improved policy based on the priority policy with the smallest value for  $g_{app}^{prio}$  as computed in Approximation 1, respectively. Obviously,  $\Delta^{n-opt}$ ,  $\Delta^{oss}$ , and  $\Delta^{osp}$  cannot take negative values. Furthermore, the closer these values are to zero, the better the corresponding policy performs.

In Table 2 the computed relative differences are summarized.

**Table 1.** Parameter values of the test bed

Parameter	Considered parameter values
$c_1$	{0.25, 0.75}
$c_2$	{1}
$(\rho_1^{FCFS}, \rho_2^{FCFS})$	$a_i^\rho \times b_j^\rho \quad \forall i, j$ where $\mathbf{a}^\rho = \{0.25, 0.5, 0.75\}$ and $\mathbf{b}^\rho = \{(\frac{2}{3}, \frac{4}{3}), (1, 1), (\frac{4}{3}, \frac{2}{3})\}$
$(\mu_1, \mu_2)$	{(0.75, 1.25), (1.25, 0.75), (1, 1)}
$(\sigma_1, \sigma_2)$	$a_i^\sigma \times b_j^\sigma \quad \forall i, j$ where $\mathbf{a}^\sigma = \{0.1, 1\}$ and $\mathbf{b}^\sigma = \{(\frac{1}{2}, \frac{3}{2}), (1, 1), (\frac{3}{2}, \frac{1}{2})\}$
$(\nu_1, \nu_2)$	$a_i^\nu \times b_j^\nu \quad \forall i, j$ where $\mathbf{a}^\nu = \{0.025, 0.1, 1\}$ and $\mathbf{b}^\nu = \{(\frac{1}{2}, \frac{3}{2}), (1, 1), (\frac{3}{2}, \frac{1}{2})\}$

**Table 2.** Relative differences  $\Delta^{n-opt}$ ,  $\Delta^{oss}$ , and  $\Delta^{osp}$  categorized in bins

Percentages	0–0.1%	0.1–1%	1–10%	10–25%	25%+
of rel. differences $\Delta^{n-opt}$	36.73	30.39	32.55	0.34	0.00
of rel. differences $\Delta^{oss}$	32.29	32.44	35.01	0.26	0.00
of rel. differences $\Delta^{osp}$	57.95	18.55	16.43	5.27	1.80

We note that the vast majority of relative differences corresponding to the near-optimal policy do not exceed 10%, and more than half of the cases constitute a difference lower than 1%. These results show that the near-optimal policy works very well. The worst performance of the near-optimal policy encountered in the test bed is the exceptional case with parameters  $(c_1, c_2) = (0.25, 1)$ ,  $(\rho_1^{FCFS}, \rho_2^{FCFS}) = (1, 0.5)$ ,  $(\mu_1, \mu_2) = (0.75, 1.25)$ ,  $(\sigma_1, \sigma_2) = (0.15, 0.05)$ , and  $(\nu_1, \nu_2) = (0.05, 0.15)$ , as well as the results  $g^{opt} = 26.37$ ,  $g^{n-opt} = 32.70$ , and, consequently,  $\Delta^{n-opt} = 24.05\%$ . For this instance, any static policy, as well as the FCFS policy, re-

sults in unstable queues. Moreover, this instance is characterized by highly asymmetric model parameters but in such a way that neither of the machines would be a clear candidate for the role of the high-priority machine in the priority policy.

We also see in Table 2 that the improved static policy performs in a similar manner as to the near-optimal policy in terms of the calculated relative differences. This is not surprising, as by construction the near-optimal policy follows the improved static policy in case the latter exists. However, the gain of the near-optimal policy lies

**Table 3.** Mean relative differences categorized in each of the parameters as specified in Table 1.

	(a)		(b)				
	$c_1$		$a_i^\rho$				
	0.25	0.75	0.25	0.5	0.75		
Mean rel. diff. $\Delta^{n-opt}$	1.26%	1.06%	Mean rel. diff. $\Delta^{n-opt}$	0.66%	1.00%	1.83%	
Mean rel. diff. $\Delta^{oss}$	1.30%	1.14%	Mean rel. diff. $\Delta^{oss}$	0.66%	1.01%	2.22%	
Mean rel. diff. $\Delta^{osp}$	2.23%	2.50%	Mean rel. diff. $\Delta^{osp}$	2.63%	1.30%	3.24%	
	(c)			(d)			
	$b_j^\rho$			$(\mu_1, \mu_2)$			
	$(\frac{2}{3}, \frac{4}{3})$	$(1, 1)$	$(\frac{4}{3}, \frac{2}{3})$	$(0.75, 1.25)$	$(1, 1)$	$(1.25, 0.75)$	
Mean rel. diff. $\Delta^{n-opt}$	1.39%	0.98%	1.11%	Mean rel. diff. $\Delta^{n-opt}$	1.26%	1.15%	1.07%
Mean rel. diff. $\Delta^{oss}$	1.58%	0.99%	1.13%	Mean rel. diff. $\Delta^{oss}$	1.31%	1.22%	1.14%
Mean rel. diff. $\Delta^{osp}$	1.51%	3.99%	1.77%	Mean rel. diff. $\Delta^{osp}$	2.38%	2.38%	2.40%
	(e)			(f)			
	$a_i^\sigma$			$a_i^\nu$			
	0.025	0.1	1	0.1	1		
Mean rel. diff. $\Delta^{n-opt}$	0.42%	0.77%	2.29%	Mean rel. diff. $\Delta^{n-opt}$	1.84%	0.48%	
Mean rel. diff. $\Delta^{oss}$	0.48%	0.76%	2.32%	Mean rel. diff. $\Delta^{oss}$	1.89%	0.50%	
Mean rel. diff. $\Delta^{osp}$	0.07%	2.21%	6.58%	Mean rel. diff. $\Delta^{osp}$	4.58%	1.06%	
	(g)			(h)			
	$b_j^\sigma$			$b_j^\nu$			
	$(\frac{1}{2}, \frac{3}{2})$	$(1, 1)$	$(\frac{3}{2}, \frac{1}{2})$	$(\frac{1}{2}, \frac{3}{2})$	$(1, 1)$	$(\frac{3}{2}, \frac{1}{2})$	
Mean rel. diff. $\Delta^{n-opt}$	1.09%	1.07%	1.31%	Mean rel. diff. $\Delta^{n-opt}$	1.39%	1.20%	0.90%
Mean rel. diff. $\Delta^{oss}$	1.19%	1.13%	1.35%	Mean rel. diff. $\Delta^{oss}$	1.39%	1.28%	0.99%
Mean rel. diff. $\Delta^{osp}$	2.12%	1.90%	3.12%	Mean rel. diff. $\Delta^{osp}$	3.24%	1.70%	2.23%

primarily in the fact that the near-optimal policy can handle a far broader range of parameter settings than the static policy. For example, of all instances with  $a^\rho = 0.75$ , there are 268 instances for which the improved static policy is not available due to stability issues. The near-optimal policy, however, does result in an implementable policy for all 2916 instances considered in the test bed.

Judging by Table 2, the performance of the improved priority policy does differ from that of the near-optimal policy, as opposed to the improved static policy. In 57.95% of the cases where an improved priority policy is available, the performance of the improved priority policy is less than 0.1% away from that of the optimal policy. The relative difference, however, exceeds 10% in more than 7% of the cases. Thus, there is far more variation in the performance of the priority policy than in the performance of the near-optimal policy. Furthermore, for 414 of the instances considered in this section, there is no improved priority policy available. It is, however, important to note that the set of instances for which no priority policy exists is completely disjoint of the set consisting of instances with no available improved static policy. This illustrates the fact that the improved static policy and the improved priority policy are complementary. These complementary parameter regions are combined in the near-optimal policy.

To observe any further parameter effects, Table 3 displays tables that show the mean relative difference  $\Delta^{n-opt}$ ,  $\Delta^{oss}$ , and  $\Delta^{osp}$  categorized in some of the variables. Based on these results, we identify four factors determining the quality of the near-optimal policy.

1. Table 3(a) suggests that the closer the value of  $c_1$  is to the value of  $c_2$ , the better the quality of the near-optimal policy becomes. A similar effect can be observed in the table shows in Table 3(c) with the values  $\rho_1^{FCFS}$  and  $\rho_2^{FCFS}$ . These effects suggest that the level of asymmetry in the parameters plays a role in the effectiveness of the near-optimal policy. Intuitively, this makes sense, as the optimal policy gets easier to predict when the system becomes more symmetric. For example, in the case of a completely symmetric model (i.e.,  $\lambda_1 = \lambda_2$ ,  $\mu_1 = \mu_2$ , etc.), the threshold curve of the optimal policy is easily seen to be the line  $x_1 = x_2$  by a switching argument. In that case, the improved static policy also attains this curve, which suggests that the near-optimal policy is optimal in symmetric systems, provided that the initial static policy is stable.
2. Judging by Table 3(b), the performance of the near-optimal policy with respect to the optimal policy becomes worse when the workload of products offered to the queues increases. This can be explained by the fact that in the case of a smaller workload, products on average encounter less waiting products in their respective queue, and therefore are less influenced by the downtimes of their machines, which occurred before their

arrival. In its turn, this means that the sojourn time of products in the system is less sensitive to any sub-optimal decisions taken in the past, improving the accuracy of the near-optimal policy. In the extreme case where the workload offered to each queue equals zero (i.e., there are no products arriving), any policy is optimal, as the system does not incur any costs in that case.

3. From Tables 3(e) and 3(f) it is apparent that the quality of the near-optimal policy is influenced by the values of  $a_i^\sigma$  and  $a_i^\nu$ . This can be mainly explained by the fact that these values determine the level of competition between the machines for access to the repairman. When breakdowns do not occur often and repairs are done quickly, the event of having both machines down is exceptional, so that any sub-optimality of the policy used is expected to have a relatively little impact on the average costs.
4. Tables 3(d) and 3(h) seem to contradict with the first factor mentioned, as the near-optimal policy seems to perform better when  $M_2$  has more “favorable” characteristics with respect to  $M_1$ ; i.e., the fast product services and fast repairs make it lucrative to repair  $M_2$  at the expense of additional downtime for  $M_1$ . However, note that this effect occurs because the cost weights are already taken in favor of the repair of  $M_2$  in every instance of the test bed. When the workloads are such that the static policy becomes infeasible, a priority policy with  $M_2$  as the high-priority machine will then already be close to optimal. Therefore, its improved version also works particularly well. However, if, as opposed to the cost weights, the rates of product services, breakdowns and repairs are in favor of  $M_1$ , a priority policy works less well, since there is no clear candidate for the high-priority machine anymore. This leaves room for sub-optimality of the improved priority policy.

As for the other policies, Table 3 suggests that the performance of the improved static policy exhibits similar parameter effects as that of the near-optimal policy. Again, this is not surprising considering the way the near-optimal policy is constructed. The improved priority policy, however, behaves differently in a number of ways. First, Tables 4(a) and 4(c) show that the improved priority policy performs better in systems with skewed model parameters. For these systems, usually the operational state of one machine is evidently more important than the other, so that the initial priority policy already performs quite well. Unlike the near-optimal policy and the improved static policy, we see in Table 3(b) that the performance of the improved priority policy does not necessarily increase in the workload offered to the system. Finally, Table 3(d) suggests that the performance of the improved priority policy is highly insensitive to any difference in the service rates of the machines.

*Remark 9.* In Section 5.2.2, we introduced an approximation  $g_{app}^{prio}$  for the long-term average costs of the priority

policy with either machine as the high-priority machine. We did this for the purpose of predicting which of the two improved priority policies performs best, in case both of them exist. Of the 2916 instances considered in the test bed, there are 1782 instances for which both priority policies lead to an improved policy. For each of these instances, it turns out that the best performing improved priority policy corresponds to the initial priority policy with the smallest approximated costs. This suggests that the approximation for the long-term average costs fulfills its purpose well.

*Remark 10.* In this section, we have considered models with two machines, of which the breakdown rates and repair rates are of a comparable size. Interference between machines may in practice, however, also occur in systems with a larger number of machines that have breakdown rates that are much smaller than their repair rates. In that case, having two or more machines in need of repair is not a rare event, so that the question of how to allocate the repairman's resources is still an important one. For models with a larger number of machines (and thus a larger number of queues), we have already established in Remark 8 that the near-optimal policy scales well. Numerical computation techniques, however, break down, so that a numerical study, similar to the one in this section for  $N = 2$ , becomes infeasible. Observe, however, that if  $N$  increases, but the breakdown rates decrease at a similar intensity, the average number of machines that are in need of repair fluctuates around the same mean. The situation of  $N = 2$  and similar rates as considered in this section is thus comparable to the case of a large number of machines with dissimilar rates. Therefore, it is expected that the near-optimal policy also performs well for the case of  $N > 2$ .

## 8. Conclusions and future research

In this article, we have considered a novel queueing network with two interacting layers and with machines being servers of products in one layer and customers in the other. Since the machines have to share repair facilities offered by the repairman in the second layer, the queue lengths of products in the first layer are generally correlated, which makes analysis of this model complicated in general.

For this model, we have studied the question of how the repairman should distribute his/her capacity over the machines in order to minimize the queue lengths of the first-layer queues. It turns out that the optimal policy for the repairman is a non-idling policy and falls in the class of threshold policies. Analytical expressions for the optimal policy are hard to derive. Therefore, we have constructed a near-optimal policy that is expressed analytically in terms of the model's parameters by using the classic one-step policy improvement approach. We have used multiple initial policies that are complementary in nature to construct a near-optimal policy. Subsequently, we have examined the accuracy of the near-optimal policy by performing an ex-

tensive numerical study in Section 7. In general, we saw that the near-optimal policy performs extremely well for a wide range of parameter settings. Furthermore, it turned out that the near-optimal policy is applicable in a far broader range of parameter settings than the improved static policy and the improved priority policy separately.

Further avenues for research include relaxing the exponentiality assumptions in the model. In order to deal with more general distributions, one could model these as phase-type distributions to remain in the dynamic programming framework. This leads to numerically completely intractable models, for which the techniques we developed are promising approaches. Furthermore, it would be interesting to study a more general class of cost functions. In particular, many quality of service constraints are usually expressed in terms of tail probabilities instead of averages. Finally, the expressions we derived for the relative value functions for the class of static and priority policies may be applied to obtain heuristic policies in a wide array of other models.

## Acknowledgements

The authors wish to thank Onno Boxma for valuable comments on earlier drafts of the present article.

## Funding

This work was funded in the framework of the STAR-project "Multilayered Queueing Systems" by the Netherlands Organization for Scientific Research (NWO). The research of Maria Vlasiou is also partly supported by an NWO individual grant through project 632.003.002.

## References

- Bertsekas, D. and Gallager, R. (1992) *Data Networks*, Prentice-Hall, Engelwood Cliffs, NJ.
- Bhulai, S. (2006) On the value function of the  $M/\text{Cox}(r)/1$  queue. *Journal of Applied Probability*, **43**, 363–376.
- Bhulai, S. (2009) Dynamic routing policies for multiskill call centers. *Probability in the Engineering and Informational Sciences*, **23**, 101–119.
- Bhulai, S. and Spieksma, F.M. (2003) On the uniqueness of solutions to the Poisson equations for average cost Markov chains with unbounded cost functions. *Mathematical Methods of Operations Research*, **58**, 221–236.
- Dorsman, J.L., Boxma, O.J. and Vlasiou, M. (2013) Marginal queue length approximations for a two-layered network with correlated queues. *Queueing Systems*, **75**, 29–63.
- Dorsman, J.L., van der Mei, R.D. and Vlasiou, M. (2013) Analysis of a two-layered network by means of the power-series algorithm. *Performance Evaluation*, **70**, 1072–1089.
- Dorsman, J.L., Vlasiou, M. and Zwart, B. (2015) Heavy-traffic asymptotics for parallel queues with Markov-modulated service speeds. To appear in *Queueing Systems*.
- Franks, G., Al-Omari, T., Woodside, M., Das, O. and Derisavi, S. (2009) Enhanced modeling and solution of layered queueing networks. *IEEE Transactions on Software Engineering*, **35**, 148–161.

- Fuhrmann, S.W. and Cooper, R.B. (1985) Stochastic decompositions in the M/G/1 queue with generalized vacations. *Operations Research*, **33**, 1117–1129.
- Gross, D. and Ince, J.F. (1981) The machine repair problem with heterogeneous populations. *Operations Research*, **29**, 532–549.
- Haijema, R. and Van der Wal, J. (2008) An MDP decomposition approach for traffic control at isolated signalized intersections. *Probability in the Engineering and Informational Sciences*, **22**, 587–602.
- Haque, L. and Armstrong, M.J. (2007) A survey of the machine interference problem. *European Journal of Operational Research*, **179**, 469–482.
- Harkema, M., Gijzen, B.M.M., van der Mei, R.D. and Hoekstra, Y. (2004) Middleware performance: a quantitative modeling approach, in *Proceedings of the International Symposium on Performance Evaluation of Computer and Communication Systems (SPECTS)*, Society for Modeling and Simulation International, San Jose, California, USA, pp. 733–742.
- Hernández-Lerma, O. and Lasserre, J.B. (1996) *Discrete-Time Markov Control Processes: Basic Optimality Criteria*, Springer-Verlag, New York, NY.
- Kleinrock, L. (1976) *Queueing Systems, Volume II: Computer Applications*, Wiley, New York, NY.
- Koole, G.M. (2006) Monotonicity in Markov reward and decision chains: Theory and applications. *Foundations and Trends in Stochastic Systems*, **1**, 1–76.
- Lippman, S.A. (1975a) Applying a new device in the optimization of exponential queueing systems. *Operations Research*, **23**, 687–710.
- Lippman, S.A. (1975b) On dynamic programming with unbounded rewards. *Management Science*, **21**, 1225–1233.
- Norman, J.M. (1972) *Heuristic Procedures in Dynamic Programming*, Manchester University Press, Manchester, UK.
- Ott, T.J. and Krishnan, K.R. (1992) Separable routing: a scheme for state-dependent routing of circuit switched telephone traffic. *Annals of Operations Research*, **35**, 43–68.
- Puterman, M.L. (1994) *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, John Wiley & Sons, Inc., Chichester, UK.
- Sassen, S.A.E., Tijms, H.C. and Nobel, R.D. (1997) A heuristic rule for routing customers to parallel servers. *Statistica Neerlandica*, **51**, 107–121.
- Wartenhorst, P. (1995)  $N$  parallel queueing systems with server breakdown and repair. *European Journal of Operational Research*, **82**, 302–322.
- Wijngaard, J. (1979) Decomposition for dynamic programming in production and inventory control. *Engineering and Process Economics*, **4**, 385–388.

## Biographies

Jan-Pieter Dorsman received his master's degree (cum laude) in Business Mathematics and Informatics from the VU University Amsterdam in 2010. He did his master's thesis project on the performance analysis of polling systems in a traineeship at the Centrum Wiskunde & Informatica (CWI). He presently works as a Ph.D. student at the Eindhoven University of Technology and CWI under the supervision of O.J. Boxma, R.D. van der Mei, and M. Vlasiou. His Ph.D. research focuses on the performance analysis of layered queueing networks. He is the co-recipient of the best paper award in ICORES 2013.

Sandjai Bhulai (1976) received his M.Sc. degrees in Mathematics and in Business Mathematics and Informatics, both from the VU University Amsterdam, The Netherlands. He carried out his Ph.D. research on Markov decision processes: the control of high-dimensional systems at the same university for which he received his Ph.D. degree in 2002. After that he was a postdoctoral researcher at Lucent Technologies, Bell Laboratories as NWO Talent Stipend fellow. In 2003 he joined the Mathematics department at the VU University Amsterdam, where he is an associate professor in Applied Probability and Operations Research. His primary research interests are in the general area of stochastic modeling and optimization, in particular, the theory and applications of Markov decision processes. His favorite application areas include telecommunication networks and call centers. He is currently involved in the control of time-varying systems, partial information models, dynamic programming value functions, and reinforcement learning.

Maria Vlasiou is an Associate Professor in the Department of Mathematics and Computer Science at the Eindhoven University of Technology (TU/e), a research fellow of the European research institute EURANDOM, and scientific staff member of CWI. She received her B.Sc. (2002, Hons.) and Ph.D. (2006) from the Aristotle University of Thessaloniki and TU/e, respectively. Her research interests center around stochastic processes and stochastic operations research and their applications to production and computer networks. She received a personal grant in the Women in Science framework from TU/e (2008), the MEERVOUD individual grant (2011), and a TOP grant for the best emerging scientists in a field (2013) from the Netherlands Organisation for Scientific Research. She is the co-author of some 30 refereed papers and the co-recipient of the best paper award in ICORES 2013 and the Marcel Neuts student paper award in MAM8.