Proceedings of the 2019 Winter Simulation Conference
N. Mustafee, K.-H.G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, eds.

MULTIOBJECTIVE RANKING AND SELECTION WITH CORRELATION AND HETEROSCEDASTIC NOISE

Sebastian Rojas-Gonzalez

Decision Sciences and Information Management
KU Leuven
Naamsestraat 69, Leuven, 3000, BELGIUM

Juergen Branke

Warwick Business School
The University of Warwick
Coventry, CV4 7AL, UK

Inneke Van Nieuwehuyse

Research Group Logistics
Hasselt University
Agoralaan, Building D
Diepenbeek, 3590, BELGIUM

ABSTRACT

We consider multi-objective ranking and selection problems with heteroscedastic noise and correlation between the mean values of alternatives. From a Bayesian perspective, we propose a sequential sampling technique that uses a combination of screening, stochastic kriging metamodels, and hypervolume estimates to decide how to allocate samples. Empirical results show that the proposed method only requires a small fraction of samples compared to the standard EQUAL allocation method, with the exploitation of the correlation structure being the dominant contributor to the improvement.

1 INTRODUCTION

A multiobjective optimization problem can be formulated as follows (Branke et al. 2008): $\min[f_1(\mathbf{x}),...,f_m(\mathbf{x})]$ for m objectives where the *decision vectors* $\mathbf{x} = [x_1,...,x_n]^T$ (also referred to as *systems* or *design points*) are contained in the decision space D (usually $D \subset \mathbb{R}^n$), with $f : D \to \mathbb{R}^m$ the vector-valued function with coordinates $f_1,...,f_m$ in the objective space $\Theta \subset R^m$. As there are tradeoffs between the different objectives, the goal then is to find a set F of all vectors $\mathbf{x}^* = [x_1^*,...,x_n^*]^T$ where no objective can be improved without negatively affecting any other objective. The systems in this solution set are referred to as *non-dominated* or *Pareto-optimal*, and form the *Pareto set* in the decision space, and the *Pareto front* in the objective space.

In stochastic simulation optimization, the objectives are also perturbed by noise. In general, relying on the *observed* mean objective values to determine the non-dominated systems may lead to two possible errors due to sampling variability: designs that actually belong to the non-dominated set can be wrongly considered dominated, or designs that are truly dominated are considered Pareto-optimal. Chen and Lee (2010) refers to these errors as *Error Type 1* (ET1) and *Error Type 2* (ET2) respectively, whereas Hunter et al. (2019) refer to them as *misclassification by exclusion* (MCE) and *misclassification by inclusion* (MCI). Throughout this paper we adopt the former terminology.

The most commonly used method for reducing the impact of noise during optimization is to evaluate the same point a number of times and use the mean of these samples as the response value (referred to as *static resampling*). However, when the noise is high and/or strongly heterogeneous, this method may fail to provide accurate approximations with limited computational budget. It is thus necessary to use

more advanced procedures that aim to correctly identify the non-dominated systems, such as *probabilistic dominance* or *multiobjective ranking and selection* (MORS).

Some of the early works that discussed probabilistic dominance appear in Fieldsend and Everson (2005) and Basseur and Zitzler (2006), who propose to use the *expected* values of any deterministic indicator to compare the quality of different Pareto fronts with a certain confidence level, under the assumption that each solution is inherently associated with a probability distribution over the objective space. More recently, in Voß et al. (2010), probabilistic dominance is defined by comparing the volume in the objective space of the confidence intervals, and the center point of these volumes is used to determine the dominance relationship. Similarly, in Trautmann et al. (2009), the standard deviation is added to the mean such that dominance is determined with the quantile objective values.

A more advanced alternative is to use MORS methods; these, however, are relatively scarce in the literature (Hunter et al. 2019). Most MORS procedures aim to ensure a high probability of *correctly* selecting (PCS) a true non-dominated design, by smartly distributing the *finite* replication budget among critically competitive designs, in order to avoid unnecessary resampling (i.e., they avoid spending replications on those designs that are clearly dominated and thus are not interesting to the decision-maker). Hunter et al. (2019) refer to these methods as fixed-budget procedures, and they represent the largest group in the literature. Some of the most relevant works include MOCBA, a well-known Bayesian Multiobjective Optimal Computing Budget Allocation framework based on the single-objective OCBA (Chen and Lee 2010); MOCBA+ (Li et al. 2018), an extension of MOCBA where the authors approach the problem from a large deviations perspective, and a formulation that maximizes the lower bound of the rate of decay in the probability of false selection (i.e., the probability of not identifying the true Pareto set) is derived. Under certain conditions, both MOCBA and MOCBA+ are asymptotically optimal, but only MOCBA+ accounts for sampling correlations in the objectives.

Furthermore, the SCORE allocations for bi-objective ranking and selection framework presented in Feldman et al. (2015) and Feldman and Hunter (2018), accounts for correlations between the objectives, is asymptotically optimal and aims to allocate replications to maximize the rate of decay of the probability of misclassification. The recent M-MOBA (Branke and Zhang 2015) and M-MOBA-HV (Branke et al. 2016) frameworks propose, from a Bayesian perspective, to use the *expected value of information* (Chick et al. 2010) to determine the observed solution which is expected to change the current Pareto set the most if more replications are performed on it. In the M-MOBA algorithm, this system is the one with the highest probability of changing the current Pareto set, whereas in M-MOBA-HV, it is the one leading to the largest change in the observed hypervolume.

In this paper we propose a MORS method that builds upon the M-MOBA and M-MOBA-HV procedures. Its main advantage is that it uses *stochastic kriging* metamodels (Ankenman et al. 2010) to build the predictive distributions of each objective, thus it is able to exploit correlations in the decision space. By using stochastic kriging we are also able to take into account the intrinsic *heterogeneous* noise affecting the observed performance, as opposed to standard (deterministic) kriging, which assumes that the observations in the prior distributions are deterministic (Kleijnen 2015). We exploit the distance between the stochastic kriging predictions and the observed sample means (resp. the sample variance and the predictor uncertainty) in combination with the expected hypervolume difference (HVD) (Branke et al. 2016). This information is used to determine which points in the solution set should get more replications, in order to improve the identification of the Pareto-optimal solutions.

2 STOCHASTIC KRIGING

Stochastic kriging is a recently developed metamodeling technique for representing the response surface implied by a stochastic simulation (Ankenman et al. 2010). For a given objective and an arbitrary design point $\mathbf{x}^i$, the model represents the observed objective value $\tilde{f}_r(\mathbf{x}^i)$ in the r^{th} replication as:

$$\tilde{f}_r(\mathbf{x}^i) = \mathbf{f}(\mathbf{x}^i)^T \boldsymbol{\beta} + M(\mathbf{x}^i) + \varepsilon_r(\mathbf{x}^i), \tag{1}$$

where $\mathbf{f}(\mathbf{x}^i)$ is a vector of known trend functions of $\mathbf{x}^i$, $\boldsymbol{\beta}$ is a vector of unknown parameters of compatible dimension, and $M(\mathbf{x}^i)$ is a realization of a mean 0 covariance stationary Gaussian random field at the design point $\mathbf{x}^i$. It is assumed that this field exhibits spatial correlation: i.e., $M(\mathbf{x}^i)$ and $M(\mathbf{x}^h)$ will tend to be similar when $\mathbf{x}^i$ is close to $\mathbf{x}^h$ in the design space. This assumption is analogous to the assumption made in the standard deterministic kriging metamodel (see, e.g., Kleijnen 2015); essentially, this type of uncertainty is imposed on the problem to aid in developing the metamodel. Hence, it is referred to as *extrinsic uncertainty*. Different spatial correlation functions (also referred to as *kernels*) exist. The most popular ones in the kriging literature are the Gaussian and Matérn kernels (Rasmussen and Williams 2005); the former is used in the experiments presented in Section 4.2:

$$k_G(\mathbf{x}^i, \mathbf{x}^h) = \lambda^2 \exp\left[-\sum_{d=1}^{n} \left(\frac{|\mathbf{x}_d^i - \mathbf{x}_d^h|}{\sqrt{2} l_d} \right)^2 \right] \tag{2}$$

where λ^2 and l_d ($d = 1, ..., n$) are the *hyperparameters* to denote the process variance and the length-scale of the process along dimension d, respectively (see Rasmussen and Williams 2005 for further details on kernel functions).

The *intrinsic uncertainty* $\varepsilon_r(\mathbf{x}^i)$ is, naturally, independent and identically distributed across replications, having mean 0 and variance $\tau^2(\mathbf{x}^i)$ at any arbitrary point $\mathbf{x}^i$. Note that the model allows for heterogenous noise, implying $\tau^2(\mathbf{x}^i)$ need not be constant throughout the design space. The model also allows for $Corr[\varepsilon_r(\mathbf{x}^i), \varepsilon_r(\mathbf{x}^h)] > 0$, as tends to be the case with the use of common random numbers (CRN); yet, this is not desirable, as discussed in Chen et al. (2012). In what follows, we only consider the case where the first term in Equation (1) is a constant, β_0, representing the overall mean of the response surface, as this has been shown to be the most useful model in practice (Ankenman et al. 2010; Kleijnen 2015).

The **stochastic kriging prediction** $\hat{f}(\mathbf{x}^i)$ at *any* design point $\mathbf{x}^i$ (whether it has been sampled or not) is given by

$$\hat{f}(\mathbf{x}^i) = \beta_0 + \Sigma_M(\mathbf{x}^i, \cdot)^T [\Sigma_M + \Sigma_\varepsilon]^{-1} (\bar{\mathbf{f}} - \beta_0 \mathbf{1}_p) \tag{3}$$

This expression is analogous to the kriging predictor in the well-known deterministic kriging model, except for the impact of the intrinsic noise, present through Σ_ε. Σ_ε is the $p \times p$ covariance matrix with (i,h) element $cov\left[\sum_{j=1}^{r_i} \varepsilon_j(\mathbf{x}^i)/r^i, \sum_{j=1}^{r_h} \varepsilon_j(\mathbf{x}^h)/r^h\right]$ across all design points $\mathbf{x}^h$ and $\mathbf{x}^i$; when CRN are not used, this reduces to the diagonal matrix $\text{diag}[\tau^2(\mathbf{x}^1)/r^1, ..., \tau^2(\mathbf{x}^p)/r^p]$. The notation $\bar{\mathbf{f}}$ is the vector containing all the observed mean outcomes at the already sampled design points: $\bar{\mathbf{f}} = [\bar{f}(\mathbf{x}^1), ..., \bar{f}(\mathbf{x}^p)]^T$, with $\bar{f}(\mathbf{x}^i) = \sum_{k=1}^{r^i} \tilde{f}_k(\mathbf{x}^i)/r^i$, and $\mathbf{1}_p$ is a $p \times 1$ vector of ones. Analogous to the deterministic kriging model, Σ_M denotes the $p \times p$ matrix containing the covariances between each couple of already sampled points, as implied by the extrinsic spatial correlation model: $\Sigma_M(\mathbf{x}^i, \mathbf{x}^j) = Cov[M(\mathbf{x}^i), M(\mathbf{x}^j)]$. The notation $\Sigma_M(\mathbf{x}^i, \cdot)$ is the $p \times 1$ vector containing the covariances between the point under study, and the p already sampled points:

$$\Sigma_M(\mathbf{x}^i, \cdot) = \left[Cov[M(\mathbf{x}^i), M(\mathbf{x}^1)], Cov[M(\mathbf{x}^i), M(\mathbf{x}^2)], \ldots, Cov[M(\mathbf{x}^i), M(\mathbf{x}^p)] \right]^T.$$

As opposed to the deterministic kriging predictor, the stochastic kriging predictor is not an exact interpolator, due to the presence of the intrinsic noise. The MSE of the stochastic kriging predictor (i.e., the predictor uncertainty), denoted $s^2(\mathbf{x}^i)$, is given by (Ankenman et al. 2010):

$$s^2(\mathbf{x}^i) = \Sigma_M(\mathbf{x}^i, \mathbf{x}^i) - \Sigma_M(\mathbf{x}^i, \cdot)^T [\Sigma_M + \Sigma_\varepsilon]^{-1} \Sigma_M(\mathbf{x}^i, \cdot) + \frac{\gamma^T \gamma}{\mathbf{1}_p^T [\Sigma_M + \Sigma_\varepsilon]^{-1} \mathbf{1}_p} \tag{4}$$

$$\text{with } \gamma = 1 - \mathbf{1}_p^T [\Sigma_M + \Sigma_\varepsilon]^{-1} \Sigma_M(\mathbf{x}^i, \cdot)$$

Again, the difference with the deterministic kriging expressions lies in the impact of the intrinsic noise through Σ_ε, which allows for modeling heterogeneous noise variances (as opposed to the common assumption in the literature of homogeneous noise). In the absence of intrinsic noise, expressions (3) and (4) thus reduce to the deterministic kriging predictor and its variance. The presence of intrinsic noise *inflates* the MSE, as discussed in Ankenman et al. (2010). The above expressions suppose that β_0, Σ_ε and Σ_M are known; clearly, in a realistic application, they must be estimated. This is commonly done using maximum likelihood estimation (MLE); yielding $\hat{\beta}_0$, $\hat{\Sigma}_\varepsilon$ and $\hat{\Sigma}_M$. We refer to Ankenman et al. (2010) for the detailed derivation of these MLE estimators. These estimators are then used in expressions (2) and (3) to yield the estimated kriging predictor $\hat{f}$, and its variance $\hat{s}^2$.

3 HYPERVOLUME DIFFERENCE

Measuring the quality of a Pareto front approximation is difficult, as the so-called "true" Pareto front is usually unknown. Intuitively, a *good* Pareto front is characterized by being well-populated (i.e., richness), with the non-dominated points well-spread across the front with respect to all the objectives (i.e., *diversity*). Numerous quantitative performance indicators have been developed for assessing the quality of the Pareto front in *deterministic* problem settings (see Riquelme et al. (2015) for a recent review). One of the most widely used quality indicators is the *hypervolume* (Zitzler et al. 2007), as it is a *Pareto compliant* indicator: for a given point a in front A, if a point a' is found that dominates a, then the hypervolume of front A' is larger than for A. However, the runtime complexity of computing the hypervolume is exponential in the number of objectives (Zitzler et al. 2007).

The hypervolume dominated by a given Pareto front A with respect to a reference point $\mathbf{r}$ is defined as the Lebesgue measure, denoted Λ, of the set of objective vectors dominated by the solutions in A, but not by $\mathbf{r}$:

$$HV(A,\mathbf{r}) = \Lambda\left(\bigcup_{z\in A}\{z \prec z' \prec \mathbf{r}\}\right), \quad z \in \Theta \tag{5}$$

where the notation $z \prec z'$ denotes domination of vector z over z'. Thus, all the non-dominated vectors contribute equally to the indicator value, and the dominated vectors do not contribute. Then, for two fronts A and B, the hypervolume difference (HVD) is defined as (Branke et al. 2016):

$$HVD(A,B,\mathbf{r}) := HV(A,\mathbf{r}) + HV(B,\mathbf{r}) - 2\times\Lambda(HV(A,\mathbf{r}) \cap HV(B,\mathbf{r})) \tag{6}$$

As discussed in Branke et al. (2016) and Hunter et al. (2019), using the expected HVD is more likely to be useful to the decision-maker, as it invests more computational effort in estimating the performance of outstanding systems, rather than differentiating between systems that are marginally dominated or non-dominated. Figure 1 (right) shows the HVD between a non-dominated set A (black points) and another potential non-dominated set B (black squares).

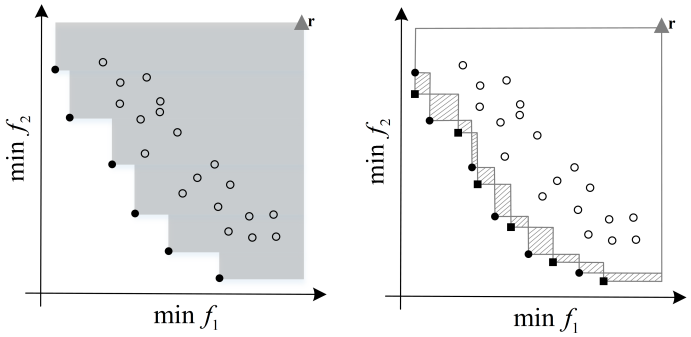

Figure 1: The figure on the left shows the hypervolume covered (shaded area) by a set of non-dominated points (black points) with respect to reference point $\mathbf{r}$ (triangle). Circles represent the dominated points.

4 PROPOSED METHOD: SK-MORS

Given that stochastic kriging is not an exact interpolator, the systems observed so far, used to build the stochastic kriging metamodels, can be re-evaluated using the predictor to yield a hopefully more accurate approximation of its response. We are then using both the sample means and variances (denoted $\bar{f}$ and σ^2), respectively the predictor values and uncertainties (denoted $\hat{f}$ and $\hat{s}$), to allocate replications to those points that are expected to change the current observed Pareto-optimal set the most. We thus rely on the stochastic kriging predictions to make decisions, as opposed to relying on the observed means; the observed means and variances are used to guide the allocation procedure.

Let S be the entire set of sampled points, PS the *observed* Pareto set of S, and $\widehat{PS}$ the *predicted* Pareto set of S. Then, the evaluation of the points in PS and $\widehat{PS}$ yields the fronts PF and $\widehat{PF}$, respectively. Furthermore, for a given point $\mathbf{x}^i$, let $\bar{\mathbf{f}}(\mathbf{x}^i)$ and $\hat{\mathbf{f}}(\mathbf{x}^i)$ be the observed mean objective vectors and predicted objective vectors, respectively. Similar to Branke et al. (2016) (see Equation (6)), we calculate the *expected hypervolume difference* (EHVD) for a single point $\mathbf{x}^i$ as:

$$EHVD^i = \mid HV(\widehat{PF}) - HV\left(\widehat{PF} \backslash \hat{\mathbf{f}}(\mathbf{x}^i) \cup \bar{\mathbf{f}}(\mathbf{x}^i)\right) \mid, \qquad \forall i \in S. \tag{7}$$

We use Equation (7) as indicator for how much we expect the hypervolume to change if we run more replications on point $\mathbf{x}^i$. There are three cases to be considered:

- Case 1: $\bar{\mathbf{f}}(\mathbf{x}^i) \prec \hat{\mathbf{f}}(\mathbf{x}^i)$, the HV will change (increase) as the observed vector dominates the predicted vector.
- Case 2: $\hat{\mathbf{f}}(\mathbf{x}^i) \prec \bar{\mathbf{f}}(\mathbf{x}^i)$, and there are no other systems observed to dominate $\bar{\mathbf{f}}(\mathbf{x}^i)$. The HV will change (decrease) as the predicted vector dominates the observed vector.
- Case 3: $\hat{\mathbf{f}}(\mathbf{x}^i) \prec \bar{\mathbf{f}}(\mathbf{x}^i)$, and there are one or more systems that are observed to dominate $\bar{\mathbf{f}}(\mathbf{x}^i)$. The HV will change (decrease) as the predicted vector dominates the observed vector, but the change is not equally significant as in the previous case, as some of the previously dominated systems become non-dominated.

When the noise is high and/or strongly heterogeneous, an extra replication on a given point might significantly change its position in the objective space. As the EHVD does not reward those points which have both their observed means and predictions (marginally) dominated, we propose to use the Euclidean distance, denoted ED, between these two in addition to the EHVD, such that if the mean and prediction are far away from each other, by running more replications on these points, we expect both the mean and prediction to come closer to each other, and aid in minimizing both ET1 and ET2.

Instead of using the predictions directly, we opt for using the *farthest* confidence bound of $\hat{f}_j(\mathbf{x}^i)$ with respect to $\bar{f}_j(\mathbf{x}^i)$ for objective j. The choice of the farthest confidence bound is because we intend to inflate the ED when the uncertainty of the predictor is high, such that points with high uncertainty will be rewarded more (see Figure 2(left) for an illustration of the ED). The farthest confidence bound is analogous to using the *lower confidence bound* ($\widehat{LCB}$) of the dominated points, and the *upper confidence bound* ($\widehat{UCB}$) of the non-dominated points of the predicted values for each objective. The $\widehat{LCB}$ and $\widehat{UCB}$ are defined as $\hat{f}_j(\mathbf{x}) \pm \omega \hat{s}_j(\mathbf{x}), j = 1, ..., m$, where ω is usually an integer in the interval $[1,3]$, yielding the 68%, 95% and 99% confidence intervals (CI) of the prediction respectively (Rasmussen and Williams 2005). These bounds are commonly used in kriging-based optimization as a way to take into account the predictor uncertainty during the search for solutions (Ponweiser et al. 2008; Emmerich et al. 2006), as it minimizes the exclusion of potential promising solutions at the cost of a higher number of evaluations. Analogously, the CI of the observed means (denoted LCB and UCB) are defined as $\bar{f}_j(\mathbf{x}) \pm \omega \sigma(\mathbf{x}), j = 1, ..., m$.

We then define the ED between the observed means and the predictor as:

$$ED^i = \sqrt{\sum_{j=1}^{m} \left[\mid \bar{f}_j(\mathbf{x}^i) - \hat{f}_j(\mathbf{x}^i) \mid + \hat{s}_j^i \right]^2}, \qquad \forall i \in S. \tag{8}$$

Consequently, the sampling criterion used to rank and select the candidates in S is calculated as:

$$\mathrm{RS}^i = EHVD^i + ED^i, \qquad \forall i \in S. \tag{9}$$

with both ED and $EHVD$ being normalized in the $[0,1]$ interval. As the algorithm we propose is sequential, the system or subset of systems with highest RS is selected to perform more replications in the current iteration. After the budget in the current iteration has been depleted, the kriging parameters are recomputed with the new (more accurate) observed means and variances to perform a new iteration. The proposed algorithm is described in the next section.

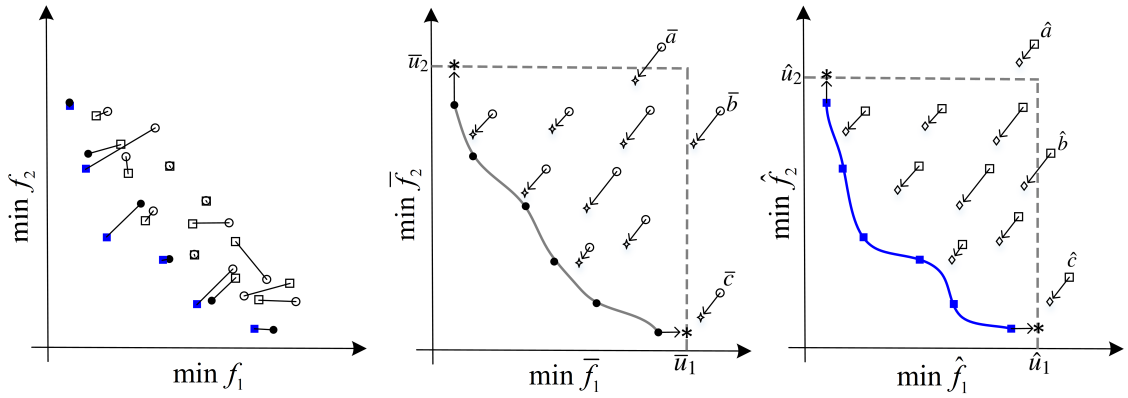

Figure 2: The left panel shows the ED between the observed means (circles) and predicted means (squares). The observed fronts are depicted as filled circles and squares respectively. The center and right figures denote the confidence regions delimited by $(\bar{u}_1, \bar{u}_2)$ and $(\hat{u}_1, \hat{u}_2)$ respectively, used in the screening procedure. The arrows on the dominated points denote the LCB (center panel) and $\widehat{LCB}$ (right panel).

4.1 Algorithm Outline

As is well-known in ranking and selection, in many cases some of the systems included in the solution set will be clearly inferior to other solutions in the set (Boesel et al. 2003). *Screening* heuristics are widely used prior to running the ranking and selection procedure in order to filter out these clearly inferior systems. We propose a subset-selection procedure (steps 2-4 in Algorithm 1), that will use both the observed and predicted Pareto fronts.

As shown in Figure 2, the worst 95% CI among all the non-dominated points for each objective j, denoted $\bar{u}_j$ for the observed means, and $\hat{u}_j$ for the predicted means, are used to enclose a confidence region. Then, the best 95% CI of each dominated system for each objective are used to improve the performance of each point. The improved performance is denoted with arrows in Figure 2, resulting in the improved positions for the sample means (denoted with stars) and the predicted values (denoted with diamonds). If the improved position does not enter the confidence region in *none* of the observed fronts (see e.g., point $\bar{c}$ (resp. $\hat{c}$) in Figure 2), then this point is filtered out from the candidate set for the current iteration (i.e., the point will be considered in further iterations). If the improved position of a point enters only one of the confidence regions (see e.g., points $\bar{a}$ (resp. $\hat{a}$) and $\bar{b}$ (resp. $\hat{b}$) in Figure 2), then it will still be considered in the candidate set of the current iteration. This is to ensure that, even with very heavy noise levels (as in the experiments in Section 4.2), with a high confidence we are not excluding truly non-dominated solutions. The steps of the entire proposed sequential procedure are given in Algorithm 1, and the different sets of points used are summarized in Table 1.

Table 1: Overview of the different sets of points used in the proposed algorithm.

Notation	Description
S	Entire set of sampled systems (candidates).
$PS(PF)$	Observed Pareto set (resp. Pareto front) based on sample means.
$\widehat{PS}(\widehat{PF})$	Observed Pareto set (resp. Pareto front) based on predicted means.
$UCB(\widehat{UCB})$	Upper 95% CI of all systems observed in PF (resp. $\widehat{PF}$).
$LCB(\widehat{LCB})$	Lower 95% CI of all systems observed in $S\backslash PF$ (resp. $S\backslash\widehat{PF}$).

Algorithm 1 SK-MORS algorithm

Input:
b $\leftarrow$ Initial number of replications per system for each objective.
$B = |S| \times b \leftarrow$ Total number of replications to be allocated in one iteration.
δ $\leftarrow$ Parameter to select the systems with highest RS.
Output:
$PS \rightarrow$ The observed Pareto set.
while stopping criterion not met **do**
  Step 1: Fit a stochastic kriging metamodel to each objective j: $\hat{f}_j(\mathbf{x})$.
  Step 2: Compute UCB, $\widehat{UCB}$, LCB and $\widehat{LCB}$.
  $\bar{u}_j$: Maximum UCB of objective j from all systems in PF, $j = 1,...,m$.
  $\hat{u}_j$: Maximum $\widehat{UCB}$ of objective j from all systems in $\widehat{PF}$, $j = 1,...,m$.
  Step 3: Initialize set of clearly inferior systems $\bar{S} = \emptyset$.
  $NP = S\backslash PS \cap S\backslash\widehat{PS}$.
  for $i \in NP$ **do**
    for $j = 1 : m$ **do**
      if $LCB_j^i > u_j$ **and** $\widehat{LCB}_j^i > \hat{u}_j$ **then**
        $\bar{S} \cup \{\mathbf{x}^i\}$.
      end if
    end for
  end for
  Step 4: Remove clearly inferior systems from the candidate set $\tilde{S} = S\backslash\bar{S}$.
  Step 5: Compute the $EHVD^i$ and ED^i values, $\forall i \in \tilde{S}$.
  Step 6: Normalize the $EHVD$ and ED values. Compute the RS values, $\forall i \in \tilde{S}$.
  Step 7: Rank the systems according to their RS. Select the Δ systems with RS $\geq \delta$.
  Step 8: $T = \sum_{i=1}^{|\Delta|} \mathrm{RS}^i$; $T^i = \frac{\mathrm{RS}^i}{T} \times B, i \in \Delta$.
  Step 9: Allocate T^i replications to each i in Δ.
end while
Step 10: Return the PS.

In **step 1**, the algorithm fits a stochastic kriging metamodel to each objective, based on the observed mean objective values and respective variances of the sampled systems in S. The metamodels are then used to approximate each response on these sampled systems. The current observed Pareto-optimal sets are obtained based on both the observed means (PS and PF) and predicted means ($\widehat{PS}$ and $\widehat{PF}$). In **step 2**, the upper and lower confidence bounds for each point in all non-dominated and dominated sets are calculated

based on the observed and predicted performances. Using these confidence bounds, in **steps 3 and 4** the algorithm filters out from the current iteration the clearly inferior systems using the screening procedure described above.

In **step 5**, the algorithm will calculate the $EHVD$ values using Equations (5) and (7), and the ED values using Equation (8). These are then normalized in **step 6**, as it is possible to have very high values of $EHVD$ relative to ED. Consequently, the RS is computed using Equation (9). Due to normalization in every iteration, the RS values are always in the $[0,2]$ interval. Thus, in **step 7**, the user-defined parameter δ will determine the subset of systems to which more replications will be allocated. In the experiments presented in Section 4.2, we opt for $\delta = 0.5$, as on average it selects 40% of the candidate points for resampling.

In **steps 8 and 9**, the allocation quantity is determined proportional to the individual RS values (using general rounding rules). The input parameter b determines how many replications could be performed per candidate system if these were distributed equally. Thus, the total number of replications B to be allocated per iteration, depends on the initial number of candidate points in S. As the algorithm will not allocate replications equally, a significantly higher number of replications $T^i \gg b$ will be allocated to the selected points in Δ. If the stopping criterion is met, in **step 10** the algorithm returns the observed Pareto set. In the experiments presented in the following section, we opt for stopping the algorithm in two ways: when the entire true Pareto set has been correctly identified (i.e., no ET1 or ET2 points), and after a fixed number of samples have been allocated.

4.2 Empirical Results

To assess the performance of the proposed procedure, we use the WFG4 function (concave Pareto front), as it is common in the literature (Huband et al. 2006). We present the results for 2 objectives and 5 decision variables, but the extension to more dimensions in both spaces is straightforward. We refer to the analytical expressions and detailed characteristics of this function in Huband et al. (2006). We set the parameter $b = 5$, and stop the algorithm after at most 25 iterations. The decision space D is discrete and set to $|S| = 100$, with a true Pareto set (denoted PS_t) of size 20. This set of candidate systems is obtained using the algorithm proposed in Rojas-Gonzalez et al. (2018).

We compare the performance of the proposed method, denoted SK-MORS, against EQUAL allocation (i.e., b replications are allocated to every point in S at each iteration). In EQUAL, the *observed means* are used instead of the stochastic kriging predictions. To show the benefit of using the stochastic kriging predictions instead of on the observed means, we compare SK-MORS against an EQUAL allocation algorithm that uses the stochastic kriging predictions instead of the observed means (denoted EQUAL-SK). The impact of the proposed sampling criterion, i.e., Equation (9), is analyzed by comparing the performance of an algorithm using only the $EHVD$ criterion (denoted SK-MORS-HV), and using only the ED criterion (denoted SK-MORS-ED). To further show the convergence rate of the algorithms, we propose to use a performance measure that evaluates the *accuracy of the true Pareto set* identified considering both error types, which we refer to as APS:

$$APS = 1 - \frac{ET1 + ET2}{S} \tag{10}$$

Thus, if the algorithm correctly identifies the entire PS_t and does not misclassify any truly dominated points, then $APS = 1$. Table 2 summarizes the performance metrics used.

The objectives are perturbed with heterogeneous Gaussian noise. Hence, we obtain noisy observations $\tilde{f}_r^j(\mathbf{x}^i) = f^j(\mathbf{x}^i) + \varepsilon_r(\mathbf{x}^i)$, with $\varepsilon_r(\mathbf{x}^i) \sim \mathcal{N}(0, \tau(\mathbf{x}^i))$ for $j = 1, ..., m$ objectives at the r^{th} replication. In practice, the noise can follow any type of structure. In our experiments we assume that the standard deviation of the noise, $\tau(\mathbf{x})$, varies linearly with respect to the objective values, as is common in the literature (Jalali et al. 2017; Rojas-Gonzalez et al. 2018). The maximum and minimum values of $\tau(\mathbf{x})$ are linked to the range of each objective value in the region of interest (i.e., $R_f^j = \max_{\mathbf{x} \in S} f^j(\mathbf{x}) - \min_{\mathbf{x} \in S} f^j(\mathbf{x}), j = 1, ..., m$). We

Table 2: Overview of performance measures used to evaluate the algorithm.

Notation	Description
Iterations	Number of iterations necessary to identify the current Pareto set.
$ET1$	Number of points in $S \backslash \widehat{PS}$ that are incorrectly identified as dominated (Type I error).
$ET2$	Number of points in $\widehat{PS}$ that are incorrectly identified as non-dominated (Type II error).
APS	Accuracy of the $\widehat{PS}$.

consider a heavy level of noise perturbing the responses, varying between $0.1R_f^j$ and $1.5R_f^j$. Moreover, we consider the case where the noise decreases linearly with the objective values; we have minimum noise at the global minimum of each individual objective. By considering the opposite case (i.e., maximum noise at the global minimum of each individual objective) we expect to need more samples and thus the impact of accounting for correlation becomes even more important. With increasing number of objectives, the tradeoff becomes more complex; for the bi-objective case, when the noise is maximal at the optimum of one objective, will be minimal at the optimum of the other objective.

Figure 3 compares the rate of minimization of both error types for all the algorithms, and Table 3 shows the computational budget used by the algorithms after convergence to $APS = 1$. The positive impact of the proposed screening procedure is conveyed in the metric $|S|$ at the end of the algorithm; the number of candidate points is reduced sequentially as the observed performance becomes more accurate. This reduction significantly benefits the proposed method, as the computational cost of doing the double hypervolume calculations in Equation (7) for every candidate point at each iteration, grows exponentially with the number of objectives (see also Equation (5)).

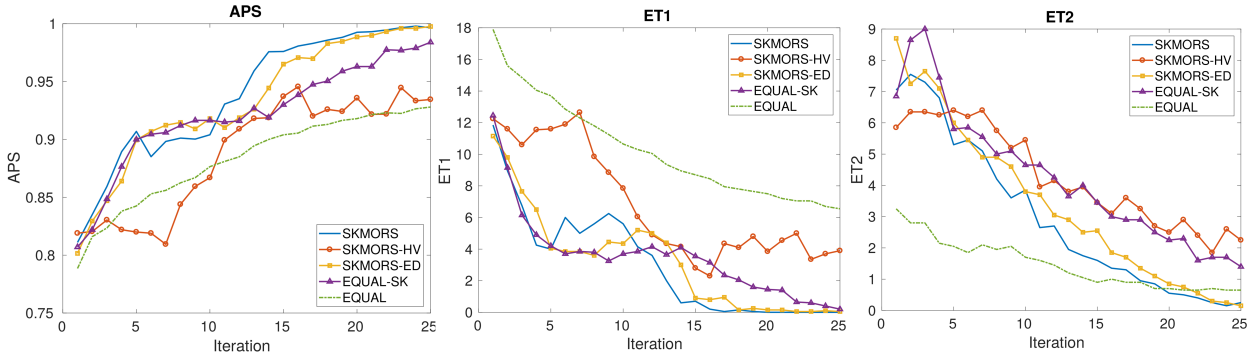

Figure 3: Performance of all the algorithms after 20 runs. At each iteration, the averages are shown.

It is clear that the speed of convergence of EQUAL allocation towards the correct identification of all the systems that are truly Pareto-optimal is drastically outperformed by the the proposed algorithm. As shown in Table 3, the number of allocated samples (i.e., iterations) required to converge to PS_t is much lower for the proposed method: slightly above 3.5% of EQUAL. Furthermore, the convergence to PS_t is observed to be more effective in SK-MORS than using the ED and $EHVD$ separately. As also reflected in the number of iterations in Table 3, SK-MORS and SKMORS-ED use significantly less budget than all the other algorithms to converge to PS_t; the performance of the ED criterion is observed to be competitive. As expected, SKMORS-HV is less likely to allocate replications to those points that are (truly) marginally dominated, and thus the algorithm requires a large budget to differentiate these. However, it is clear that the EHVD combined with the ED is the most effective sampling criterion. An interesting avenue for further work is to develop an indifference zone procedure for the EHVD.

The minimization of ET1 allows similar conclusions as to the maximization of APS; all the algorithms using stochastic kriging remain largely superior, while EQUAL requires a very large replication budget

to correctly identify the true non-dominated points (see also Table 3). With respect to the ET2, EQUAL allocation seems to reduce these faster at the beginning, but the rate of decay of the proposed algorithm surpasses EQUAL allocation after a few iterations.

Table 3: Summary of results (averages) for all algorithms after 20 runs of the same instance until convergence to $APS = 1$.

	SK-MORS	SK-MORS-HV	SK-MORS-ED	EQUAL-SK	EQUAL
$\|S\|$	80.85	81.85	83.40	100	100
Iterations	19.25	195.20	33.35	30.80	520.60

In general we observe that the identification of the points with the true best expected performance is largely facilitated by exploiting the stochastic kriging information (as opposed to relying on the observed means). A drawback of the proposed method is the need to specify the parameter δ. As this parameter determines the percentage of points to resample, it is dependent on the size of S and the specific problem at hand. In practical settings, ideally, this parameter (and the stopping criterion) should automatically adapt based on the progress of the algorithm after each iteration. This also presents an interesting opportunity for further work. The outcome of the proposed algorithm against EQUAL allocation is shown in Figure 4; the *start* of the algorithm refers to the first iteration, and the *end* refers to the moment when no ET1 or ET2 are observed.

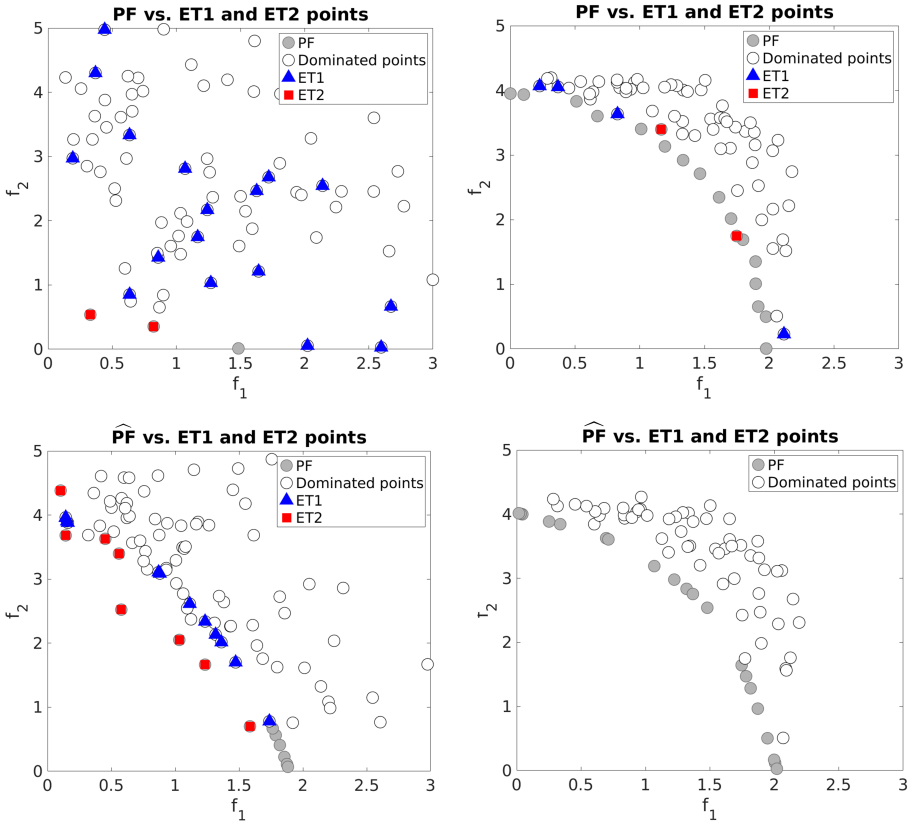

Figure 4: The upper panels show the results for EQUAL allocation, the lower panels the results for SK-MORS. The figures on the left side show the PF and $\widehat{PF}$ at the start the algorithm, while the two figures on the right side show the PF and $\widehat{PF}$ at the end the *proposed* algorithm.

5 CONCLUSION

We have proposed a multi-objective ranking and selection technique that can deal with heteroscedastic noise and still exploit correlation between the mean values of alternatives by using stochastic kriging metamodels.

The biggest reduction in the number of samples compared to EQUAL allocation came from using the values predicted by the stochastic kriging metamodel instead of the observed mean values for determining whether a solution is non-dominated. Thus, if it is possible to exploit a correlation structure between the mean values of alternatives, this should definitely be done as it is shown to bring great benefit. In previous papers, the expected hypervolume difference was proposed as a good criterion to allocate samples. However, estimating the hypervolume change is computationally expensive. We have proposed an approximation that simply looks at the effect of replacing the observed mean by the mean predicted by the stochastic kriging metamodel. The proposed measure on its own didn't perform particularly well as criterion to allocate samples to alternatives. Best results were obtained when combining it with a measure of the difference of an alternative's observed means and predicted values. An initial screening step further reduces the computational cost needed to compute the hypervolume differences.

The ideas presented on this paper need more exploration. Future work should include a comparison with other state-of-the-art MORS algorithms and a performance comparison based on probability of correct selection of hypervolume difference.

REFERENCES

Ankenman, B., B. L. Nelson, and J. Staum. 2010. "Stochastic Kriging for Simulation Metamodeling". *Operations Research* 58(2):371–382.

Basseur, M., and E. Zitzler. 2006. "A Preliminary Study on Handling Uncertainty in Indicator-Based Multiobjective Optimization". In *Applications of Evolutionary Computing*, edited by F. Rothlauf, J. Branke, S. Cagnoni, E. Costa, C. Cotta, R. Drechsler, E. Lutton, P. Machado, J. H. Moore, J. Romero, G. D. Smith, G. Squillero, and H. Takagi, 727–739. Berlin, Heidelberg: Springer.

Boesel, J., B. L. Nelson, and S.-H. Kim. 2003. "Using Ranking and Selection to "Clean Up" After Simulation Optimization". *Operations Research* 51(5):814–825.

Branke, J., K. Deb, K. Miettinen, and R. Slowinski. 2008. *Multiobjective Optimization: Interactive and Evolutionary Approaches*. Berlin, Heidelberg: Springer-Verlag.

Branke, J., and W. Zhang. 2015. "A New Myopic Sequential Sampling Algorithm for Multi-objective Problems". In *Proceedings of the 2015 Winter Simulation Conference*, edited by Y. Yilmaz, W. Chan, I. Moon, T. Roeder, C. Macal, and M. Rosseti, 3589–3598. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Branke, J., W. Zhang, and Y. Tao. 2016. "Multiobjective Ranking and Selection Based on Hypervolume". In *Proceedings of the 2016 Winter Simulation Conference*, edited by T. Roeder, P. Frazier, R. Szechtman, E. Zhou, T. Huschka, and S. Chick, 859–870. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Chen, C.-H., and L. H. Lee. 2010. *Stochastic Simulation Optimization*, Volume 1. Singapore: World Scientific.

Chen, X., B. E. Ankenman, and B. L. Nelson. 2012, March. "The Effects of Common Random Numbers on Stochastic Kriging Metamodels". *ACM Transactions on Modeling and Computer Simulation* 22(2):7:1–7:20.

Chick, S. E., J. Branke, and C. Schmidt. 2010. "Sequential Sampling to Myopically Maximize the Expected Value of Information". *INFORMS Journal on Computing* 22(1):71–80.

Emmerich, M. T., K. C. Giannakoglou, and B. Naujoks. 2006. "Single and Multiobjective Evolutionary Optimization Assisted by Gaussian Random Field Metamodels". *IEEE Transactions on Evolutionary Computation* 10(4):421–439.

Feldman, G., and S. R. Hunter. 2018, January. "SCORE Allocations for Bi-objective Ranking and Selection". *ACM Transactions on Modeling and Computer Simulation* 28(1):7:1–7:28.

Feldman, G., S. R. Hunter, and R. Pasupathy. 2015. "Multiobjective Simulation Optimization on Finite Sets: Optimal Allocation Via Scalarization". In *Proceedings of the 2015 Winter Simulation Conference (WSC)*, edited by Y. Yilmaz, W. Chan, I. Moon, T. Roeder, C. Macal, and M. Rosseti, 3610–3621. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Fieldsend, J. E., and R. M. Everson. 2005. "Multiobjective Optimisation in the Presence of Uncertainty". In *2005 IEEE Congress on Evolutionary Computation*, Volume 1, 243–250. Institute of Electrical and Electronics Engineers, Inc.

Huband, S., P. Hingston, L. Barone, and L. While. 2006. "A Review of Multiobjective Test Problems and a Scalable Test Problem Toolkit". *IEEE Transactions on Evolutionary Computation* 10(5):477–506.

Hunter, S. R., E. A. Applegate, V. Arora, B. Chong, K. Cooper, O. Rincón-Guevara, and C. Vivas-Valencia. 2019, January. "An Introduction to Multiobjective Simulation Optimization". *ACM Transactions on Modeling and Computer Simulation* 29(1):7:1–7:36.

Jalali, H., I. Van Nieuwenhuyse, and V. Picheny. 2017. "Comparison of Kriging-based Algorithms for Simulation Optimization with Heterogeneous Noise". *European Journal of Operational Research* 261(1):279 – 301.

Kleijnen, J. P. C. 2015. *Design and Analysis of Simulation Experiments*. 2^{nd} ed. New York: Springer.

Li, J., W. Liu, G. Pedrielli, L. H. Lee, and E. P. Chew. 2018. "Optimal Computing Budget Allocation to Select the Nondominated Systems: A Large Deviations Perspective". *IEEE Transactions on Automatic Control* 63(9):2913–2927.

Ponweiser, W., T. Wagner, D. Biermann, and M. Vincze. 2008. "Multiobjective Optimization on a Limited Budget of Evaluations Using Model-Assisted S-Metric Selection". In *Parallel Problem Solving from Nature - PPSN X*, edited by G. Rudolph, T. Jansen, N. Beume, S. Lucas, and C. Poloni, 784–794. Berlin, Heidelberg: Springer.

Rasmussen, C. E., and C. K. I. Williams. 2005. *Gaussian Processes for Machine Learning (Adaptive computation and machine learning)*. 1^{st} ed. Cambridge, Massachusetts: The MIT Press.

Riquelme, N., C. Von Lücken, and B. Baran. 2015. "Performance metrics in multi-objective optimization". In *Latin American Computing Conference (CLEI), 2015*, 1–11. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Rojas-Gonzalez, S., H. Jalali, and I. Van Nieuwenhuyse. 2018. "A Stochastic-kriging-based Multiobjective Simulation Optimization Algorithm". In *Proceedings of the 2018 Winter Simulation Conference*, edited by M. Rabe, A. Juan, N. Mustafee, A. Skoogh, S. Jain, and B. Johansson, 2155–2166. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Trautmann, H., J. Mehnen, and B. Naujoks. 2009. "Pareto-dominance in Noisy Environments". In *Proceedings of the Eleventh Conference on Congress on Evolutionary Computation*, 3119–3126. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Voß, T., H. Trautmann, and C. Igel. 2010. "New Uncertainty Handling Strategies in Multi-objective Evolutionary Optimization". In *Parallel Problem Solving from Nature, PPSN XI*, edited by R. Schaefer, C. Cotta, J. Kołodziej, and G. Rudolph, 260–269. Berlin, Heidelberg: Springer.

Zitzler, E., D. Brockhoff, and L. Thiele. 2007. "The Hypervolume Indicator Revisited: On the Design of Pareto-compliant Indicators Via Weighted Integration". In *Evolutionary Multi-Criterion Optimization*, edited by S. Obayashi, K. Deb, C. Poloni, T. Hiroyasu, and T. Murata, 862–876. Berlin, Heidelberg: Springer.

AUTHOR BIOGRAPHIES

SEBASTIAN ROJAS-GONZALEZ is a Ph.D. student in the Department of Decision Sciences and Information Management at KU Leuven, Belgium. His doctoral work is centered on methods for multi-criteria decision making in settings with uncertainty. His website is https://feb.kuleuven.be/Sebastian.RojasGonzalez/ and his email address is sebastian.rojasgonzalez@kuleuven.be.

JUERGEN BRANKE is Professor of Operational Research and Systems of Warwick Business School, University of Warwick, UK. He is Area Editor for the Journal of Heuristics and the Journal on Multi Criteria Decision Analysis, and Associate Editor for IEEE Transaction on Evolutionary Computation, and the Evolutionary Computation Journal. His research interests include metaheuristics and Bayesian optimization, multiobjective optimization and decision making, optimization in the presence of uncertainty, and simulation-based optimization. He has published over 180 peer-reviewed papers in international journals and conferences. His e-mail address is Juergen.Branke@wbs.ac.uk.

INNEKE VAN NIEUWENHUYSE is a Full Professor at the Research Group Logistics, Hasselt University (Belgium), and guest professor at KU Leuven (Belgium). She holds M.Sc. and PhD degrees in Applied Economics-Business Engineering from the University of Antwerp (Belgium). She teaches courses in Supply Chain Management, Facilities Design and Operations Strategy. Her research focuses on the design and analysis of stochastic manufacturing and supply chain systems. Her email address is inneke.vannieuwenhuyse@uhasselt.be.