A MULTI-OBJECTIVE COAL INVENTORY MANAGEMENT MODEL USING MONTE CARLO COMPUTER SIMULATION

R. Brits[1][#] & J. Bekker[1][*]

Contact details
*   Corresponding author
   jb2@sun.ac.za

Author affiliations
1   Department of Industrial
   Engineering, Stellenbosch
   University, South Africa

\#   Author was enrolled for a M Eng
   (Industrial) in the Department of
   Industrial Engineering,
   Stellenbosch University, South
   Africa

ABSTRACT

The ability of a coal-fired power station to meet its generation targets is influenced by periods of coal shortage. In this article, we propose a multi-objective inventory model to assist with the management of coal stockpiles. The model is applicable to power utilities with a network of two or more coal-fired power stations. The aim is to determine the near-optimal amount of coal inventory to stockpile at each station in the network. A Monte Carlo coal stockpile simulator is used to incorporate stochastic uncertainty into the stockpile levels, while a metaheuristic uses the simulator as an estimator of two objective functions. The metaheuristic finds good values for the coal stockpile level at each power station in the network. The algorithm for multi-objective optimisation using the cross-entropy method is proposed as a suitable metaheuristic. A hypothetical case study is used to validate the inventory model and to showcase the optimisation results.

OPSOMMING

Die vermoë van 'n steenkool-aangedrewe kragsentrale om sy opwekkingteikens te bereik word beïnvloed deur periodes van 'n steenkooltekort. In hierdie artikel stel ons 'n meerdoelige voorraadmodel voor om te help met die bestuur van steenkoolvoorraad. Die model is van toepassing op elektrisiteitsverskaffers met 'n netwerk van twee of meer steenkoolaangedrewe kragsentrales. Die doel is om die naby-optimale steenkoolvoorraadvlak te bepaal wat by elke kragsentrale in die netwerk gestoor moet word. 'n Monte Carlo steenkool voorraad-simulator word gebruik om stogastiese onsekerheid in die steenkoolverbruik te inkorporeer, terwyl 'n metaheuristiek die simulator gebruik as 'n beramer vir twee doelfunksies. Die metaheuristiek vind goeie waardes vir steenkoolvlakke by elke kragsentrale in die leweringsnetwerk. Die algoritme vir meerdoelige optimering met die kruis-entropie metode word voorgestel as 'n geskikte metaheuristiek. 'n Hipotetiese gevallestudie word gebruik om die voorraadmodel te valideer en optimeringsresultate te toon.

1 INTRODUCTION

The continued evolution of civilisations is highly dependent on a secure and accessible supply of energy and, as the human population continues to grow, global energy demand will continue to increase. Power utilities are at the forefront of global energy supply. Currently, more than 40 per cent of the world's electricity is generated at coal-fired power stations [1]. This figure does not seem that high, but many countries — such as South Africa (93%), Poland (92%), China (79%), India (69%), and the USA (49%) — rely primarily on coal for electricity generation [2]. The unpredictable demand for electricity, along with the uncertainties related to inventory replenishment and coal quality, places great emphasis on accurate and effective inventory planning at coal-fired power stations. Like most industries, the inventory problem encountered by power producers is typically of a multi-objective nature.

In this article, we propose a multi-objective inventory model to assist with the management of coal stockpiles. The model is applicable to power utilities with a network of two or more coal-fired power stations, and assumes that some of the stations in the network are capable of burning coal with similar properties. This assumption allows for inventory to be moved between the stations, or for coal deliveries to be redirected from the source in emergency situations. A Monte Carlo coal stockpile simulator coded in the open source R programming language is used to emulate stochastic depletion and replenishment of the stockpile levels. The simulator requires an electricity production plan as input. This production plan is determined using an independently-developed production planning module that allows for the input of different electricity demand profiles. The simulator allows for the input of such profiles that consider weather (cold, hot), sectors (agriculture, industry, mining, households), and different geographical areas. The basic simulator, called the energy flow simulator (EFS), was obtained from Enerweb [3].

The simulated coal stockpiles are used as a basis for the decision variables and objective functions that we propose. The algorithm for multi-objective optimisation using the cross-entropy method (MOO CEM) of Bekker & Aldrich [4] is proposed as a suitable metaheuristic to solve the inventory model approximately.

The remainder of this article is organised as follows: A short discussion on inventory management in the electricity supply sector is provided in Section 2, followed by important fundamentals of multi-objective optimisation in Section 3. The proposed coal stockpile simulator is described in Section 4. In Section 5, we explain how the output of the simulator is used to formulate the multi-objective coal inventory model. In Section 6, we use a hypothetical case study to illustrate the capability of the inventory model while also testing the effectiveness of the MOO CEM algorithm in finding approximate solutions for the model. Concluding remarks are provided in Section 7.

2   INVENTORY MANAGEMENT IN THE ELECTRICITY SUPPLY SECTOR

Logistics plays a significant role in the operation of thermal power stations. Several simulation and optimisation studies have been conducted on the logistics systems and supply chains of coal-fired power stations in particular. In their paper, Li and Li [5] propose a simulation and optimisation model for the logistics system of a coal-fired power station using WITNESS software. The model, which was applied to a power station in China, included the coal suppliers, the transportation system, and the power station itself. These researchers regarded these as the three major components of a power station's logistics system. The selection of suppliers is critical to guaranteeing the quantity and quality of the coal supply [6],[7]. The transportation system has many requirements and constraints that must be considered, while at the power station the coal storage systems, the coal conveying systems, and human resources must be managed.

In a typical coal supply chain, the coal suppliers and transportation companies are responsible for managing inventory levels at mines and throughout the transportation network, while the electricity suppliers are responsible for managing the coal stockpiles at their power stations. According to Zhanwu, Xiao, Zhonglei, and Xiongbiao [8], the survival and development of coal-fired power stations is seriously affected by the bottleneck that is created during periods of coal shortage. Compared with other enterprises, inventory management at coal-fired power stations has its own features and requirements due to the nature of the goods being stored. Zhanwu *et al.* [8] mention the following as the four main features of coal inventory management:

1. The uncertainty of coal demand: The uncertain demand for coal at power stations is triggered by the uncertain demand for electricity. Many factors contribute to variation in demand, such as temperature, humidity, and wind speed [9].
2. The uncertainty of coal suppliers: The price of coal varies over time, especially where short-term contracts with suppliers are in place. Similar to any other market, high demand means high prices, and *vice versa*.
3. The uncertainty of inventory replenishment: Due to weather conditions and other uncertainties within coal transportation networks, the inventory replenishment of coal is not constant.
4. The requirement of safety stock: Each of the above uncertainties emphasises the need to stockpile coal because coal shortages will result in an electricity shortage. Safety stock is an effective management tool for protecting electricity suppliers against uncertainty. However, there is a trade-off between having the ability always to provide customers with the promised service level and the costs involved in storing large amounts of coal at a power station.

These four features of coal inventory management must be considered when developing a coal stockpile inventory management model. It also follows from the four features mentioned above that at least two conflicting objectives prevail in coal stockpile inventory management: carrying sufficient inventory while trying to minimise the capital invested in inventory.

Next, multi-objective optimisation principles and solution methods are discussed.

3   MULTI-OBJECTIVE OPTIMISATION PRINCIPLES AND SOLUTION METHODS

Multi-objective optimisation (MOO) problems differ from single objective problems, in that they have at least two conflicting objectives. This means that the meaning of 'optimum' changes. There are many acceptable solutions for a given problem; the aim is to find good 'trade-offs' between the objective functions [10].

The MOO problem for K objectives, D decision variables, and M unequal and Q equal constraints, is defined by:

$$\text{Minimise } \mathbf{f}(\mathbf{x}) = [f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_K(\mathbf{x})]^T$$

$$\text{subject to } \mathbf{x} \in \Omega$$

$$\Omega = \{\mathbf{x} \,|\, g_i(\mathbf{x}) \leq 0, \qquad i = 1, 2, \dots, M;$$

$$h_j(\mathbf{x}) = 0, \qquad\qquad i = 1, 2, \dots, Q\}$$

where $\mathbf{x} = [x_1, x_2, \dots, x_D]^T$ is a D dimensional vector of decision variables. Each x_i ($i = 1, 2, \dots, D$) can be real-valued, integer-valued, or boolean-valued, and is constrained within some feasible set Ω, which consists of M equality constraints g_i and Q inequality constraints h_j. The degrees of freedom are given by $M - Q$ and, to avoid an over-constrained problem, it is required that $Q < M$.

Approaches to solving MOO problems can be broadly organised into two categories: *scalarisation approaches* and *Pareto approaches*. In scalarisation approaches, the MOO problem is solved by translating it back into a single-objective problem (or a series of single-objective problems). The preferences of the decision-maker must thus be incorporated into the optimisation model before solutions are found. A few examples are the weighted sum approach, lexicographic ordering, and goal programming. Pareto approaches differ from scalarisation approaches in that they do not admit a unique solution but a set of solutions, based on the concept of Pareto-optimality. In these approaches, the decision-maker's preferences are incorporated after the optimisation model has been solved. MOO metaheuristics belong to this category ([11], [12]).

Metaheuristics are global search algorithms that are widely-recognised for their ability to obtain near-optimal solutions within reasonable time periods. These algorithms are generally designed for combinatorial optimisation problems in the deterministic context, but have recently been successfully applied in simulation optimisation. A major advantage of metaheuristics lies in their ability to approximately solve problems for which no satisfactory problem-specific algorithm is available. Also, metaheuristics perform well in the case of a high-dimensional, discontinuous, and non-differentiable response surface. This means that they are capable of overcoming the trap of local optimality ([13], [14], [15]).

Metaheuristics used in MOO are mostly population-based. This means that the algorithm searches by iteratively improving a population of solutions. The aim of an MOO metaheuristic is to find a set of decision-variable vectors, known as the *Pareto optimal set*. (For a comprehensive discussion of these algorithms and some of their variants, see Chapter 2 of Coello, Lamont, and Van Veldhuizen [10].)
Multi-objective optimisation using the cross-entropy method (MOO CEM) was selected as a suitable approach to solving the MOO inventory model described in Section 5. The MOO CEM algorithm is a population-based metaheuristic inspired by statistical principles. The cross-entropy method (CEM) for optimisation is a versatile Monte Carlo method, developed by Rubinstein and Kroese [16]. The approach was motivated by the work of Rubinstein [17] on variance minimisation methods for rare-event probability estimation, and modified in Rubinstein [18] to solve continuous and combinatorial optimisation problems ([19], [20]). In contrast to a random search algorithm that searches for an

optimal solution by sampling decision variable values from the same probability density function (pdf), the CEM assigns a pdf to each decision variable. The search works on the principle that the parameters of each of these pdfs are iteratively adjusted to increase the probability of drawing decision variables that result in good objective function values. The aim of the CEM is to estimate the parameters of each decision variable's ideal pdf in order for them to converge to an optimal (or near-optimal) solution. The algorithm does this by using the cross-entropy (or Kullback-Leibler distance). As a measure of the distance between two pdfs, a minimum cross-entropy is desired between the sampling distribution associated with each decision variable and the optimal pdf from which to sample ([12], [16]).

The CEM has been proven to converge quickly when applied to optimisation problems with one objective [16]. It is thus an ideal approach for the computationally-expensive, time-dependent problems often encountered in simulation optimisation (SO). This was the motivational factor for expanding the CEM to solve multi-objective problems. The MOO CEM was introduced by Bekker and Aldrich [4] and Bekker [21], and has since been the topic of multiple research studies.

4   THE PROPOSED COAL STOCKPILE SIMULATOR

The Monte Carlo simulation model that we propose is an extension of the coal stockpile simulator of Micali and Heunis [22]. We have added a multi-objective optimisation capability to their coal stockpile simulator. For each replication (or sample path), the simulator computes the daily power production and the end-of-day coal stockpile level for each power station $s \in S$ over a user-specified study horizon. This is done by stochastically adding uncertainty (or noise) to the simulator's two main inputs: the planned power production for power station s on day t (g_{st}), and the planned coal deliveries to power station s on day t (d_{st}). Further inputs are the heat rate of power station s (H_s), the estimated calorific value (CV) of the coal burnt at power station s (CV_s), and the energy availability factor (EAF) of power station s on day t (EAF_{st}). Each power station has a specific heat rate, and is designed to burn coal within a specific CV range. The energy availability factor is a percentage value that represents a power station's generation availability with respect to its total capacity ([3]). It is denoted by

$$EAF_{st} = 100 - PCLF_{st} - OCLF_{st} - UCLF_{st} \tag{1}$$

where $PCLF_{st}$ refers to the power generation capacity lost due to planned maintenance at power station s on day t. $UCLF_{st}$ refers to the power generation capacity lost as a result of unplanned maintenance at power station s on day t, while $OCLF_{st}$ refers to the power generation capacity lost due to extraordinary events at power station s on day t, such as employee strikes or the theft of transmission cables. The PCLF percentages are extracted from the power utility's maintenance plan, while the UCLF and OCLF percentages are forecast based on historical events.

The following uncertainties are considered in the simulator:

1. A power station's actual UCLF on a given day may differ from the expected UCLF.
2. The actual CV of the coal burnt at a power station on a given day may differ from the estimated CV.
3. A power station's actual coal deliveries on a given day may differ from the planned coal deliveries.

The three uncertainties are incorporated daily for each power station, as shown by (2) to (4). Stochastic noise is added to the UCLF at each power station s on a given day t by adding a random value sampled from a normal probability density function (pdf) with a mean of zero and a standard deviation that is obtained from historical unplanned maintenance data. The actual power production at power station s on day t (G_{st}) is subsequently computed by multiplying g_{st} with the ratio of actual EAF over planned EAF ([3]). This is mathematically denoted by

$$G_{st} = g_{st} \times \left(\frac{100 - PCLF_{st} - OCLF_{st} - (UCLF_{st} + N(0, \sigma_s^{(UCLF)}))}{EAF_{st}} \right) \tag{2}$$

where σ_s^{UCLF} is the UCLF standard deviation for power station s.

The actual CV of the coal burnt at each station is computed in a similar manner. The CV is adjusted with a standard deviation that is obtained from historical coal quality data, and the actual coal burnt at power station s on day $t(B_{st})$ is subsequently computed by

$$B_{st} = G_{st} \times \left(\frac{H_s \times 10^{-6}}{CV_s + N(0, \sigma_s^{CV})} \right) \tag{3}$$

where σ_s^{CV} is the CV standard deviation for power station s. The original expressions in (1), (2), and (3) were based on monthly time periods, but the resolution has been adjusted to daily time periods in this study. This is to allow for potential daily spikes in demand that are otherwise averaged out when a monthly resolution is used.

A triangular distribution is assumed for the actual coal deliveries because its parameters can be specified by subject-matter experts. A triangular distribution's pdf is defined by three values: the minimum value, the most likely value, and the maximum value. We propose that the actual coal deliveries to power station s on day $t(D_{st})$ be sampled from a triangular pdf with the most likely value set as d_{st}. The minimum and maximum value of each is set as ad_{st} and bd_{st} respectively, where a $(0 < a < 1)$ and b $(1 < b < 2)$. The values of a and b are user-specified values that define the distribution's range. This is mathematically denoted by

$$D_{st} = Tri(ad_{st}, d_{st}, bd_{st}). \tag{4}$$

After incorporating the uncertainties for power s on day t, the coal stockpile level for power station s at the end of day t is computed by

$$S_{s,t} = S_{s,t-1} + D_{st} - B_{st}. \tag{5}$$

Negative stockpile levels are not allowed. Thus, if $S_{s,t}$ is negative, $S_{s,t}$ is set equal to zero and the actual coal burnt and actual power production are recalculated by

$$B_{st} = D_{st} - L_{s,t-1} \tag{6}$$

And

$$G_{st} = B_{st} \times \left(\frac{CV_s + N(0, \sigma_s^{CV})}{H_s \times 10^{-6}} \right) \tag{7}$$

respectively, where $N(0, \sigma_s^{CV})$ is the same as in (3). The expression in (7) has been adapted from Enerweb [3].

A power station's stockpile level at the end of a given day can be expressed in kilotonnes (ktonnes) or in stockpile days. To obtain the number of stockpile days on hand at a given station, the number of ktonnes is divided by the station's standard daily burn (SDB). The SDB of a power station, measured in ktonnes per day, is the amount of coal that the station could burn if it were to operate for one full day without any outages (i.e., EAF = 100%). The units of measurement for all the parameters in the coal stockpile simulator are provided in Table 1.

Table 1: Units of measurement for the parameters in the coal stockpile simulator

Symbol	Description	Unit of measurement
g_{st}, G_{st}	Power production (planned, actual)	MWh
d_{st}, D_{st}	Coal deliveries (planned, actual)	ktonnes
H_s	Heat rate	MJ/MWh
EAF_{st}	Energy availability factor	%
CV_s	Calorific value	MJ/kg
B_{st}	Coal burnt	ktonnes
$S_{s,t}$	Coal stockpile level	ktonnes (or number of stockpile days)

5   MULTI-OBJECTIVE MODEL FORMULATION

The multi-objective coal inventory model that we propose is applicable to power utilities with a network of two or more coal-fired power stations. A network such as this represents a continuous system: a certain load is constantly produced and fed into the transmission network. The model is dependent on the assumption that certain stations in the network are capable of burning coal with similar properties. This assumption allows for coal to be moved between certain stations in emergency situations. A power station with some reserve thus supplies coal to another coal-compatible power station that is in need, if it is practically possible.

The aim of the model is to determine the optimal amount of coal inventory to stockpile — referred to as the 'target stockpile levels' — at each station in the network. We denote the target stockpile level for power station s as T_s. The coal stockpile level at each power station s is simulated by the coal stockpile simulator described in Section 4. Recall that the coal stockpile simulator is a Monte Carlo simulation model that adds noise to the planned power production and planned coal deliveries. The noise is added in the form of uncertainty about unplanned maintenance (UCLF), coal quality (CV), and delivery reliability. The unplanned maintenance and the coal quality at a particular power station affect the amount of coal being burnt by the station, while the delivery reliability affects inventory replenishment. Ideally, coal delivery and burn should balance each other out so that the target stockpile level can be maintained. However, this is almost never the case.

Two hypothetical scenarios for extreme variation between coal delivery and burn at a power station are shown in Figure 1. The effective decrease in stockpile level shown in scenario 1 may be as a result of low delivery reliability, poor coal quality, or less unplanned maintenance than was expected. Scenario 2 illustrates the opposite, where the amount of coal delivered is significantly more than the coal burnt. A typical case would involve a combination of the two scenarios, as shown in Figure 2, where the inventory level fluctuates around the target level.

The inventory model we propose is to allow coal transfers between appropriate power stations in the network to try to maintain each station's coal inventory level as close as possible to its target level. To achieve this, two transfer levels, called the lower warning limit (L_s) and the upper warning limit (U_s), are required for each power station s to trigger the movement of coal.

At the end of every day t, each station's stockpile level ($S_{s,t}$) is evaluated to determine whether transfers are required. When a given station's stockpile is depleted below L_s, inventory is replenished by transferring coal from one or more of the other stations. Similarly, when a given station's stockpile becomes more than U_s, coal is transferred to some of the other stations to reduce the inventory level. Both L_s and U_s are decision variables for the inventory model, and since the main goal is to determine each station's T_s, it also has to be included as a decision variable. To trigger the target stockpile levels, we formulated the inventory model by specifying that a transfer is only allowed when, for at least one of the power stations involved in a transfer transaction, the stockpile level can be forced towards T_s. Furthermore, coal can only be moved between power stations capable of burning similar coal. All transfers are expressed in terms of the station from which the coal is moved. The coal transferred from power station s on day t is denoted by t. The basic constraints on the decision variables are:

$$L_s < T_s < U_s \qquad \forall s, \qquad\qquad (8)$$
$$L_s > 0 \qquad \forall s, \qquad\qquad (9)$$
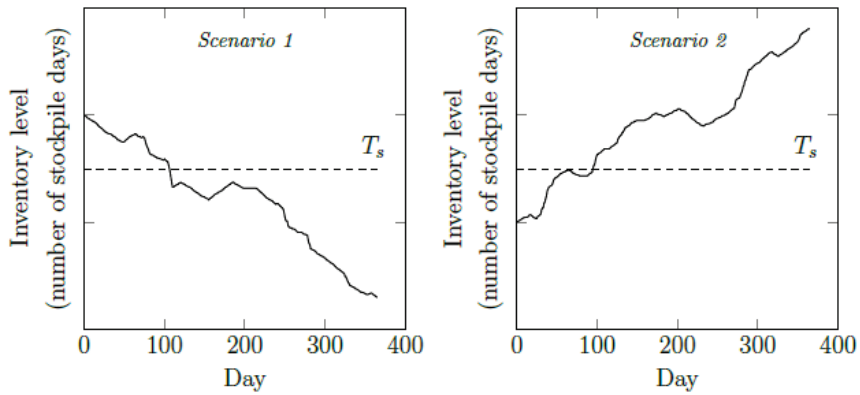$$L_s, T_s, U_s \in \mathbb{Z}^+ \qquad \forall s. \qquad\qquad (10)$$

Figure 1: Two hypothetical scenarios for extreme variation between coal delivery and burn
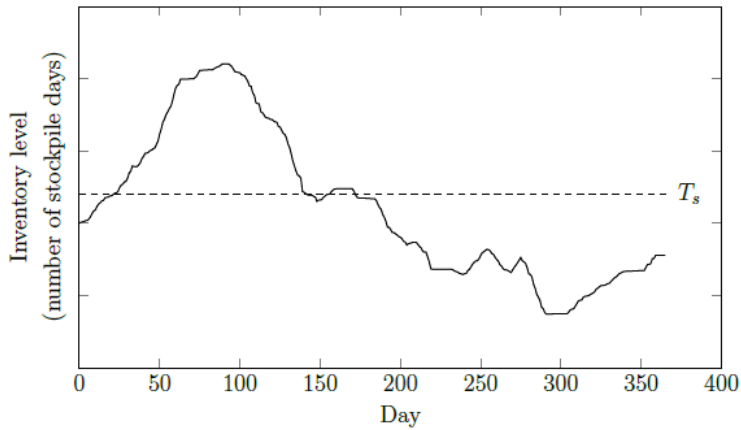

Figure 2: Hypothetical scenario for typical variation between coal delivery and burn

To simplify the complex inventory system of large power utilities, the following assumptions are made:

1. L_s, T_s, and U_s must be positive integers when expressed in terms of stockpile days.
2. There is no limit on the maximum stockpile level at any of the power stations.
3. By default, no constraint is placed on the amount of coal that can be moved between two power stations on a given day. Also, transfers are allowed on any day throughout the course of the study period, and there is no limit on the number of transfers that may be made.
4. All lead times are one day. Thus, if a coal transfer is triggered on a given day t, the changes on the stockpiles are visible on day $t + 1$.

The user has the option to discard assumption 3 by setting a cap for the amount of coal that can be moved between certain stations. Furthermore, to limit the total number of transfers that may be made, the user can specify that the stockpiles may only be evaluated on certain days. This allows the model to comply with practical considerations, as the amount of coal that can be moved between any two power stations on a given day is limited. The available infrastructure and throughput rates dictate transfer amounts; infrastructure includes trains and trucks for movement of coal by road.

For a better description of the coal transfer functions, consider the following: two coal-fired power stations A and B are capable of burning coal with similar properties. If $S_{A,t} < L_A$ becomes true on day t, a transfer is required from station B. However, coal may only be moved if $S_{B,t} > T_B$. If this condition is true, the amount of coal transferred from B to A on day $t + 1$ is calculated by

$$Y_{B,t+1} = min\{S_{B,t} - T_B; T_A - S_{A,t}; C\} \tag{11}$$

where C is the coal transfer cap between the two stations. By default, $C = \infty$.

After incorporating both the actual coal delivery and burn on day $t + 1$, $S_{B,t+1}$ and $S_{A,t+1}$ are adjusted according to

$$S_{B,t+1} \leftarrow S_{B,t+1} - Y_{B,t+1} \tag{12}$$

and

$$S_{A,t+1} \leftarrow S_{A,t+1} + Y_{B,t+1} \tag{13}$$

respectively.

The power stations at which $S_{s,d} < L_s$ receive priority (i.e., coal transfers for them are sought first); and the stations are served in descending order based on the difference between $S_{s,t}$ and L_s. Two policies are proposed to select the power stations involved in a transfer transaction:

1. **Closest first**: If a given power station requires coal ($S_{s,t} < L_s$), the stations from which the transfer will be made are selected by considering the closest one first. Similarly, if a given power station has too much coal ($S_{s,t} > U_s$), the stations to which the transfer is made are selected by considering the closest first.
2. **Most urgent:** If a given power station requires coal, the stations from which a transfer is made are selected by first considering the one at which $S_{s,t}$ is furthest above T_s. Similarly, if a given station has too much coal, the stations to which the transfer is made are selected by first considering the one at which $S_{s,t}$ is furthest below T_s.

The first objective proposed for the model is to minimise the total average coal stockpile level ($\bar{S}$), measured in ktonnes. This is a typical objective for inventory management models. Assuming a study period of $t = 365$ days, $\bar{S}$ is given by

$$\bar{S} = \sum_{s=1}^{n} \left(\frac{\sum_{t=1}^{365} S_{s,t}}{365} \right) \tag{14}$$

where n is the number of power stations included in the system.
As a second objective, it is proposed that the total coal transfers (Z) throughout the course of the study period be minimised. This is given by

$$Z = \sum_{s=1}^{n} \sum_{t=1}^{365} Y_{s,t} X_{st} \tag{15}$$

where X_{st} is the distance (in km) that coal is transferred from power station s on day t. The unit of measurement for Z is thus ktonnes.km. It is clear that (14) aims to obtain low values for all T_s, while (15) tries to space L_s, T_s, and U_s so that coal is transferred optimally, since there is a trade-off between large coal transfers that occur occasionally and frequent smaller transfers. The two objectives are conflicting, and the decision space has a combinatorial nature that requires an MOO metaheuristic. The reader is reminded that coal transfers are expensive and must be avoided as far as possible; in this study they are allowed as a way of penalising the optimisation metaheuristic for 'bad choices' of stockpile levels. In this way, the metaheuristic 'learns' and avoids bad values chosen previously.

6   HYPOTHETICAL CASE STUDY

In this section, we use a hypothetical case study to test the proposed inventory model.

We use a large hypothetical power utility with 14 coal-fired power stations. Each of these stations receives coal daily, and in quantities that vary from one delivery to the next. We denote the 14 stations 'A' to 'N'. The coal transfer matrix indicating which of the stations can burn coal with similar properties is given in Table 2. A '0' indicates that a transfer may not take place between two stations, while a '1' indicates that a transfer is allowed. The distance matrix showing the distances (in km) between the 14 stations is given in Table 3.

Table 2: Coal transfer matrix

	A	B	C	D	E	F	G	H	I	J	K	L	M	N
A	-	0	0	0	1	0	0	1	1	0	0	0	0	0
B	-	-	0	0	0	0	0	0	0	0	0	0	0	0
C	-	-	-	1	0	1	1	0	0	1	1	1	1	1
D	-	-	-	-	0	1	1	0	0	1	1	1	1	1
E	-	-	-	-	-	0	0	1	1	0	0	0	0	0
F	-	-	-	-	-	-	1	0	0	1	1	1	1	1
G	-	-	-	-	-	-	-	0	0	1	1	1	1	1
H	-	-	-	-	-	-	-	-	1	0	0	0	0	0
I	-	-	-	-	-	-	-	-	-	0	0	0	0	0
J	-	-	-	-	-	-	-	-	-	-	1	1	1	1
K	-	-	-	-	-	-	-	-	-	-	-	1	1	1
L	-	-	-	-	-	-	-	-	-	-	-	-	1	1
M	-	-	-	-	-	-	-	-	-	-	-	-	-	1
N	-	-	-	-	-	-	-	-	-	-	-	-	-	-

Table 3: Distance matrix (km)

	A	B	C	D	E	F	G	H	I	J	K	L	M	N
A	-	343.2	308	328.8	329.6	328.6	328.7	302	439	387.7	411.8	336	328.3	356
B	-	-	160.6	180	126.5	131.2	131.3	122.6	182.3	127.5	210.2	201.2	165	51.9
C	-	-	-	26.9	40.7	35.8	35.9	39.5	134	90.9	105.1	45.7	18.9	122.3
D	-	-	-	-	53.8	49	49.1	62.6	121.5	87.8	83.1	21.6	14.6	137.5
E	-	-	-	-	-	5	4.9	28.1	100.9	57.9	101.1	75.5	38.9	83.1
F	-	-	-	-	-	-	0.1	26.9	111.1	60.6	99.7	70.2	34.2	88.9
G	-	-	-	-	-	-	-	27	111.2	60.7	99.6	70.1	34.1	89
H	-	-	-	-	-	-	-	-	138	85.5	125.5	83.1	50.7	89.5
I	-	-	-	-	-	-	-	-	-	55.6	62.1	127.6	115.3	131.7
J	-	-	-	-	-	-	-	-	-	-	76.9	103.3	77.2	83.9
K	-	-	-	-	-	-	-	-	-	-	-	80.9	84.9	159.6
L	-	-	-	-	-	-	-	-	-	-	-	-	37.6	158.3
M	-	-	-	-	-	-	-	-	-	-	-	-	-	122.9
N	-	-	-	-	-	-	-	-	-	-	-	-	-	-

The necessary parameter values for the various power stations are shown in Table 4.

We analyse the MOO inventory model with each of the two coal transfer policies (i.e., closest first, and most urgent) and for the default case. Thus no constraint is placed on the amount of coal that may be moved between two power stations on a given day, and an unlimited number of transfers is allowed (recall assumption 3 in Section 5). A study period of 365 days is used. As output statistic, we examine the expected value, the 20th percentile value, and the 80th percentile value of the two objective functions, each computed using 10 Monte Carlo replications. Recall that the first objective (f_1) is to minimise the total average coal stockpile level (in ktonnes), while the second objective (f_2) is to minimise the total coal transfers throughout the study period (in ktonnes.km).

Since we are including 14 power stations in the model, there are 3 x 14 = 42 decision variables. During each iteration of the MOO CEM algorithm, truncated normal distributions are used to sample N decision variable vectors X_i ($1 \leq i \leq 42$) such that $x_1 = L_1$, $x_2 = T_1 - L_1$, $x_3 = U_1 - T_1$, $x_4 = L_2$, and so on. Each x_i is sampled from the range $[a_i, b_i]$. For our analysis, we set $a_i = 1$ and $b_i = 20$ for all X_i. This means that the maximum values allowed for L_s, T_s, and U_s are 20 stockpile days, 40 stockpile days, and 60 stockpile days respectively. The number of possible decision variable combinations is thus 42 x 20 x 20 x 20 = 336 000, because $L_s \leq T_s \leq U_s$.

The three decision variables associated with each coal-fired power station s are thus constrained by

$$1 \leq L_s \leq 20 \qquad \forall s, \qquad (16)$$
$$L_s + 1 \leq T_s \leq L_s + 20 \qquad \forall s, \qquad (17)$$
$$T_s + 1 \leq U_s \leq T_s + 20 \qquad \forall s. \qquad (18)$$

Table 4: Parameters for the coal stockpile simulator

	H_s (MJ/MWh)	CV_s (MJ/kg)	σ_s^{CV}	σ_s^{UCLF}	a	b
A	10 000	19.65	1.5	1.4	0.9	1.1
B	10 000	15.50	1.5	3.1	0.9	1.1
C	10 000	22.85	1.5	10.0	0.9	1.1
D	10 000	22.85	1.5	5.3	0.9	1.1
E	10 000	19.65	1.5	3.9	0.9	1.1
F	10 000	22.85	1.5	8.5	0.9	1.1
G	10 000	22.85	1.5	8.5	0.9	1.1
H	10 000	19.65	1.5	1.8	0.9	1.1
I	10 000	19.65	1.5	3.3	0.9	1.1
J	10 000	22.85	1.5	3.9	0.9	1.1
K	10 000	22.85	1.5	8.5	0.9	1.1
L	10 000	22.85	1.5	4.4	0.9	1.1
M	10 000	22.85	1.5	3.9	0.9	1.1
N	10 000	22.85	1.5	14.0	0.9	1.1

Also, recall that the decision variable values must be positive integers when expressed in terms of stockpile days. No maximum stockpile levels are specified, and we set each station's initial stockpile level equal to T_s.

The approximate Pareto set achieved for the closest-first transfer policy and for the most urgent transfer policy are shown graphically in Figure 3. The expected value is used as the output statistic. A good distribution of solutions was obtained for both transfer policies, and the non-dominated solutions cover a wide range of values for each objective function. One can clearly see that there is not much difference between the two transfer policies when only the non-dominated sets are considered.
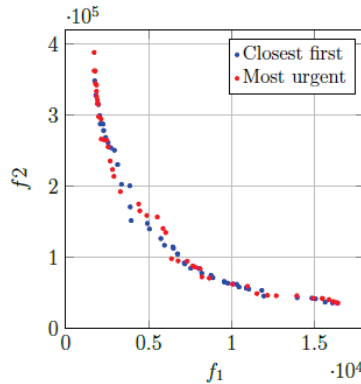

Figure 3: Approximate Pareto set achieved for the two coal transfer policies

To evaluate the quality of the solutions, we compare the approximate Pareto set achieved for each transfer policy with sets that were achieved with 50 000 objective function evaluations in the MOO CEM. These plots are given in Figure 4(a) for the closest-first transfer policy, and in Figure 4(b) for the most urgent transfer policy. One can see that slightly better solutions were found with 50 000 evaluations — i.e., the solutions are closer to the origins of the plots due to minimisation.

When solving an MOO problem, the process does not stop when the Pareto set (which is approximate) is achieved. A single best solution must still be selected by the decision-maker. This process is not trivial, because the members of the approximate Pareto set are the estimations of a stochastic process. A few methodologies for multi-objective decision-making under uncertainty have been developed over the years. The two-phase selection procedure by Chen and Lee [23] may be used, while the technique proposed by Lee, Chew, Teng and Goldsman [24] is fairly recent. We do not elaborate on these techniques here, as they are not part of the research scope.
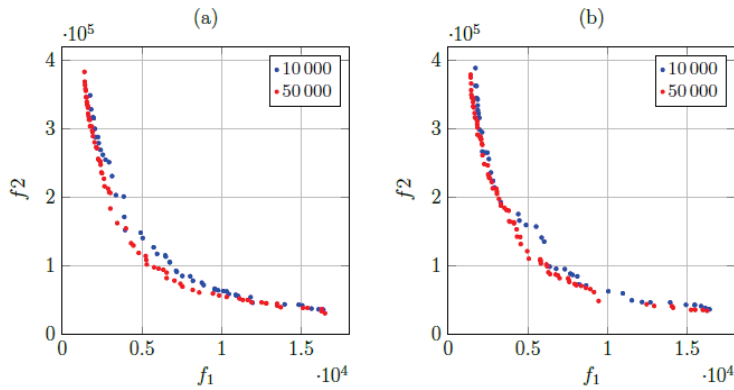
Figure 4: A comparison of the approximate Pareto sets achieved for 10 000 vs 50 000 MOO CEM objective function evaluations for (a) the closest-first transfer policy, and (b) the most urgent transfer policy
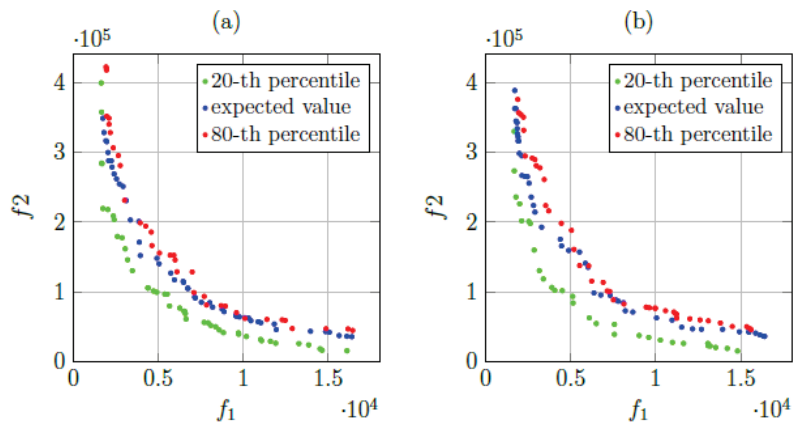

Figure 5: Approximate Pareto sets achieved for (a) the closest-first transfer policy, and (b) the most urgent transfer policy when different output statistics are used

A graphical comparison of the approximate Pareto sets for the statistics of the three outputs (i.e., the expected value, the 20th percentile value, and the 80th percentile value) is shown in Figure 5(a) for the closest-first transfer policy and Figure 5(b) for the most urgent transfer policy. We again used 10 000 objective function evaluations in the MOO CEM. One can see from Figure 5 that there is very little difference between the solutions achieved when using the expected value and the 80th percentile value as output statistic. The results produced by the model are thus slightly skewed towards the lower percentiles. We conclude that the expected value is sufficient, due to the small differences we observed between the approximate Pareto sets.

7   CONCLUSION

In this article, we proposed a multi-objective inventory model to assist with the management of coal stockpiles. The model is applicable to power utilities with a network of two or more coal-fired power stations, and is dependent on the assumption that some of the stations in the network are capable of burning coal with similar properties. This assumption allows for inventory to be moved between the stations in emergency situations. We used a Monte Carlo coal stockpile simulator to incorporate stochastic behaviour into the power generation process. The uncertainties considered are the unplanned outages at the power stations, the reliability of coal deliveries, and the quality of the coal.

The aim of the multi-objective inventory model is to determine the optimal (or near-optimal) amount of coal to stockpile at each station in the network. As objectives for the model, we proposed that both the total average coal stockpile level and total coal transfers be minimised. The algorithm

for multi-objective optimisation using the cross-entropy method (MOO CEM) was found to be a suitable metaheuristic for creating approximate solutions. In future work, the optimisation model will be evaluated with a real-world case.

ACKNOWLEDGEMENT

REFERENCES

[1]   **World Coal Association**, http://www.worldcoal.org/coal/uses-coal/coal-electricity, accessed on 24 October 2016.
[2]   **Sarker, M., Chowdhury, D. & Hossain, I.** 2012. Power generation from coal: A review, *Journal of Chemical Engineering*, 27(2), pp. 50-54.
[3]   **Enerweb**. 2014. www.enerweb.co.za, accessed on 11 April 2016.
[4]   **Bekker, J. & Aldrich, C.** 2010. The cross-entropy method in multi-objective optimisation: An assessment, *European Journal of Operational Research*, 211(1), pp. 112-121.
[5]   **Li, Y. & Li, R.** 2008. Simulation and optimization of the power station coal-fired logistics system based on WITNESS simulation software. In *2008 Workshop on Power Electronics and Intelligent Transportation System*, Guangzhou, China. Workshop Chair: Weitao Zheng, IEEE Computer Society, pp. 394-398.
[6]   **West, J.** 2011. Optimising coal stockpiles in a supply chain using a dynamic cost flow model. Technical report, Griffith University, Department of Accounting, Finance and Economics.
[7]   **Yucekaya, A.** 2013. Cost minimizing coal logistics for power plants considering transportation constraints, *Journal of Traffic and Logistics Engineering*, 1(2), pp. 122-127.
[8]   **Zhanwu, W., Xiao, Y., Zhonglei, F. & Xiongbiao, X.** 2011. Research on improving the inventory management of coal in thermal power plants: Take H Corporation as an example. In *Proceedings of the 2011 International Conference on Business Management and Electronic Information*, 1, pp. 321-324.
[9]   **Psiloglou, B.E., Giannakopoulos, C., Majithia, S. & Petrakis, M.** (2009). Factors affecting electricity demand in Athens, Greece and London, UK: A comparative assessment, *Energy*, 34(11), pp. 1855-1863.
[10]   **Coello, C.C., Lamont, G.B. & Van Veldhuizen, D.A.** 2007. *Evolutionary algorithms for solving multi-objective problems*. Springer Science & Business Media.
[11]   **De Weck, O.L.** 2004. Multiobjective optimization: History and promise. *Invited Keynote Paper, GL2-2, The Third China-Japan-Korea Joint Symposium on Optimization of Structural and Mechanical Systems*, 2, p. 34.
[12]   **Scholtz, E.** 2014. *A comparative study on the value of accounting for possible relationships between decision variables when solving multi-objective problems*. Master's thesis, Stellenbosch: Stellenbosch University.
[13]   **Tekin, E. & Sabuncuoglu, I.** 2004. Simulation optimization: A comprehensive review on theory and applications, *IIE Transactions*, pp. 1067-1081.
[14]   **Gendreau, M. & Potvin, J.Y.** 2005. Metaheuristics in combinatorial optimization. *Annals of Operations Research*, 140(1), pp. 189-213.
[15]   **Boussaïd, I., Lepagnot, J. & Siarry, P.** 2013. A survey of optimization metaheuristics. *Information Sciences*, 237, pp. 82-117.
[16]   **Rubinstein, R. & Kroese, D.** 2004. *The cross entropy method: A unified approach to combinatorial optimization, Monte-Carlo simulation, and machine learning*. Springer.
[17]   **Rubinstein, R.Y.** 1997. Optimization of computer simulation models with rare events, *European Journal of Operational Research*, 99(1), pp. 89-112.
[18]   **Rubinstein, R.** 1999. The cross-entropy method for combinatorial and continuous optimization, *Methodology and Computing in Applied Probability*, 1(2), pp. 127-190.
[19]   **Kroese, D.P., Porotsky, S. & Rubinstein, R.Y.** 2006. The cross-entropy method for continuous multi-extremal optimization, *Methodology and Computing in Applied Probability*, 8(3), pp. 383-407.
[20]   **Kroese, D.P., Rubinstein, R.Y., Cohen, I., Porotsky, S. & Taimre, T.** (2013). Cross-entropy method, *Encyclopedia of Operations Research and Management Science*. Springer, pp. 326-333.
[21]   **Bekker, J.** 2013. Multi-objective buffer allocation with the cross-entropy method, *International Journal of Simulation Modelling*, 12(1), pp. 50-61.
[22]   **Micali, V. & Heunis, S.** 2011. Coal stock pile simulation. *8th Conference on the Industrial and Commercial Use of Energy*, pp. 198-203.
[23]   **Chen, E.J. & Lee, L.H.** 2009. A multi-objective selection procedure of determining a Pareto set, *Computers & Operations Research*, 36(6), pp. 1872-1879.
[24]   **Lee, L.H., Chew, E.P., Teng, S. & Goldsman, D.** 2010b. Finding the nondominated Pareto set for multi-objective simulation models. *IIE Transactions*, 42(9), pp. 656-674.