

**Requirements specification for the
optimisation function of an electric utility's
energy flow simulator**



Marc Hatton

Department of Industrial Engineering

University of Stellenbosch

Study leader: James Bekker

Thesis presented in partial fulfilment of the requirements for the
degree of Master of Engineering in the Faculty of Engineering at
Stellenbosch University

M. Eng Industrial (Research)

March 2015

Declaration

By submitting this thesis electronically, I declare that the entirety of the work contained therein is my own, original work, that I am the sole author thereof (save to the extent explicitly otherwise stated), that reproduction and publication thereof by Stellenbosch University will not infringe any third party rights and that I have not previously in its entirety or in part submitted it for obtaining any qualification.

Date: 1 December 2014

Abstract

Efficient and reliable energy generation capability is vital to any country's economic growth. Many strategic, tactical and operational decisions take place along the energy supply chain. Shortcomings in South Africa's electricity production industry have led to the development of an energy flow simulator. The energy flow simulator is claimed to incorporate all significant factors involved in the energy flow process from primary energy to end-use consumption. The energy flow simulator thus provides a decision support system for electric utility planners.

The original aim of this study was to develop a global optimisation model and integrate it into the existing energy flow simulator. After gaining an understanding of the architecture of the energy flow simulator and scrutinising a large number of variables, it was concluded that global optimisation was infeasible. The energy flow simulator is made up of four modules and is operated on a module-by-module basis, with inputs and outputs flowing between modules. One of the modules, namely the primary energy module, lends itself well to optimisation. The primary energy module simulates coal stockpile levels through Monte Carlo simulation. Classic inventory management policies were adapted to fit the structure of the primary energy module, which is treated as a black box. The coal stockpile management policies that are introduced provide a prescriptive means to deal with the stochastic nature of the coal stockpiles.

As the planning horizon continuously changes and the entire energy flow simulator has to be re-run, an efficient algorithm is required to optimise stockpile management policies. Optimisation is achieved through the rapidly converging cross-entropy method. By integrating the

simulation and optimisation model, a prescriptive capability is added to the primary energy module. Furthermore, this study shows that coal stockpile management policies can be improved. An integrated solution is developed by nesting the primary energy module within the optimisation model. Scalability is incorporated into the optimisation model through a coding approach that automatically adjusts to an ever-changing planning horizon as well as the commission and decommission of power stations.

As this study is the first of several research projects to come, it paves the way for future research on the energy flow simulator by proposing future areas of investigation.

Opsomming

Effektiewe en betroubare energie-opwekkingsvermoë is van kardinale belang in enige land se ekonomiese groei. Baie strategiese, taktiese en operasionele besluite word deurgaans in die energie-verskaffingsketting geneem. Tekortkominge in Suid-Afrika se elektrisiteitsopwekkingsindustrie het tot die ontwikkeling van 'n energie-vloei-simuleerder gelei. Die energie-vloei-simuleerder vervat na bewering al die belangrike faktore wat op die energie-vloei-proses betrekking het van primêre energie-verbruik tot eindgebruik. Die energie-vloei-simuleerder verskaf dus 'n ondersteuningstelsel aan elektrisiteitsdiensbeplanners vir die neem van besluite.

Die oorspronklike doel van hierdie studie was om 'n globale optimeringsmodel te ontwikkel en te integreer in die bestaande energie-vloei-simuleerder. Na 'n begrip aangaande die argitektuur van die energie-vloei-simuleerder gevorm is en 'n groot aantal veranderlikes ondersoek is, is die slotsom bereik dat globale optimering nie lewensvatbaar is nie. Die energie-vloei-simuleerder bestaan uit vier eenhede en werk op 'n eenheid-tot-eenheid basis met insette en uitsette wat tussen eenhede vloei. Een van die eenhede, naamlik die primêre energiemodel, leen dit goed tot optimering. Die primêre energiemodel boots steenkoolreserwevlakke deur Monte Carlo-simulering na. Tradisionele voorraadbestuursbeleide is aangepas om die primêre energiemodel se struktuur wat as 'n swartboks hanteer word, te pas. Die steenkoolreserwebestuursbeleide wat ingestel is, verskaf 'n voorgeskrewe middel om met die stogastiese aard van die steenkoolreserwes te werk.

Aangesien die beplanningshorison deurgaans verander en die hele energie-vloei-simulering weer met die energie-vloei-simuleerder uitgevoer moet word, word 'n effektiewe algoritme benodig om die re-

serwebestuursbeleide te optimeer. Optimering word bereik deur die vinnige konvergerende kruis-entropie-metode. 'n Geïntegreerde oplossing is ontwikkel deur die primêre energiemodel en die optimering funksie saam te voeg. Skalering word ingesluit in die optimeringsmodel deur 'n koderingsbenadering wat outomaties aanpas tot 'n altyd-veranderende beplanningshorison asook die ingebruikneem en gebruikstel van kragstasies.

Aangesien hierdie studie die eerste van verskeie navorsingsprojekte is, baan dit die weg vir toekomstige navorsing oor die energie-vloei-simuleerder deur ondersoekareas vir die toekoms voor te stel.

Acknowledgements

I would like to express my sincere gratitude to the following people and organisations who supported me during the past two years:

- Foremost, Professor James Bekker for his sound guidance during the course of this study.
- The International Offices of Vrije Universiteit in Amsterdam and the University of Stellenbosch, for going out of their way to provide financial and logistic support which enabled me to complete a semester of study abroad, immerse myself in the Dutch culture and become a global citizen.
- The Oppenheimer Memorial Trust for providing further funding towards this exchange programme.
- The Department of Industrial Engineering at Stellenbosch University for providing the necessary resources, including bottomless coffee, to fuel this study.
- Eskom for providing partial financial assistance towards this study.
- Dr. van Harmelen and his team at Enerweb (the industry partner to this study) for their assistance.
- My loving father, mother, and sister for their patience and support, which was tested during the final stages of this study.
- Last but not least, special thanks goes to my esteemed colleagues from the *Unit for Systems Modelling and Analysis* for the wealth of review I received and the great friendships I am fortunate enough to have made.

Contents

Declaration	i
Abstract	ii
Opsomming	iv
Acknowledgements	vi
Nomenclature	xix
1 Introduction	1
1.1 Background	1
1.2 Problem statement	4
1.3 Aim and objectives	4
1.4 Structure of the thesis	5
2 The energy flow simulator	6
2.1 Background to Chapter 2	7
2.1.1 Literature review of electricity sector modelling techniques	7
2.1.1.1 Modelling problems within the sector	7
2.1.1.2 Modelling the sector as a whole	8
2.1.2 An overview of the energy supply chain	9
2.2 An introduction to the energy flow simulator	15
2.2.1 The need for the energy flow simulator	15
2.2.2 The aim of the energy flow simulator	15
2.2.3 An overview of the energy flow simulator	15

CONTENTS

2.3	Components of the energy flow simulator	16
2.4	Operating the energy flow simulator	19
2.5	Chosen area of optimisation: Primary energy module	19
2.5.1	Why the primary energy module was chosen	20
2.5.2	An overview of the primary energy module	22
2.5.3	Analysis of the primary energy module	24
2.5.4	The value of coal inventory on hand	27
2.5.5	Classification of the primary energy module as a simulation model	28
2.5.6	Limitations of the primary energy module	29
2.6	Concluding remarks on Chapter 2	30
3	Inventory and coal stockpile management systems	31
3.1	Coal stockpile management systems	31
3.1.1	Coal stockpiles as blending facilities	32
3.1.2	Coal stockpiles as buffers	33
3.2	Modelling of inventory management systems	34
3.2.1	Inventory management	34
3.2.2	Deterministic inventory problems	36
3.2.3	Stochastic inventory problems	36
3.2.3.1	News vendor problem	36
3.2.3.2	EOQ model with uncertain demand and lead time	37
3.3	Concluding remarks on Chapter 3	38
4	Simulation optimisation	39
4.1	An introduction to simulation optimisation	40
4.2	Considerations in simulation optimisation	41
4.3	Simulation optimisation techniques	43
4.3.1	Local optimisation	45
4.3.1.1	Discrete parameter set	45
4.3.1.2	Continuous parameter set	45
4.3.2	Global optimisation	47
4.3.3	The chosen simulation optimisation technique	48
4.4	Concluding remarks on Chapter 4	48

CONTENTS

5	Model formulation and solution approach	50
5.1	Conceptual model formulation	50
5.1.1	Decision variables	50
5.1.2	Components of the objective function	56
5.2	Mathematical problem formulation	57
5.2.1	Problem notation	57
5.2.2	Model	59
5.3	Solution approach	60
5.3.1	Literature: Cross-entropy method	60
5.3.1.1	Theoretical foundations	62
5.3.1.2	Applications of the cross-entropy for optimisation	65
5.3.2	Cross-entropy method applied to this study	65
5.3.2.1	Use of a normal distribution to generate solutions	65
5.3.2.2	Parameter settings	66
5.3.2.3	Nesting the primary energy module within the cross-entropy method	67
5.4	Concluding remarks on Chapter 5	69
6	Experimental design	70
6.1	Financial parameter estimation	70
6.1.1	Shortage cost	70
6.1.2	Holding cost	71
6.1.3	Emergency delivery cost	73
6.1.4	Delivery cancellation cost	75
6.2	Simulation settings for primary energy module	75
6.3	Experimental setup	76
6.3.1	Base case	76
6.3.2	Sensitivity analysis of the financial parameter estimates . .	77
6.3.3	Using different output statistics	79
6.3.4	Varying the baseline deliveries	80
6.4	Concluding remarks on Chapter 6	81

CONTENTS

7	Analysis of experimental results	82
7.1	Verification and validation	82
7.1.1	Validation and verification of the primary energy module	83
7.1.2	Validation and verification of the combined primary energy module and cross-entropy method	84
7.1.3	Running the primary energy module without the cross-entropy method	85
7.2	Factors that influence the results	88
7.3	Analysis of the results for the base case	88
7.4	Experimental results	101
7.4.1	Sensitivity analysis of the financial parameter estimates	101
7.4.2	Using different output statistics	102
7.4.3	Varying the baseline deliveries	109
7.5	Concluding remarks on Chapter 7	119
8	Conclusion	120
8.1	Project summary	120
8.2	Suggested future research	121
8.3	Skills acquired	124
8.4	Final remarks	124
	References	126
A	Parameters of the energy flow simulator	A-1
B	Analysis of the primary energy module	B-1

List of Figures

2.1	Distribution of power stations in South Africa.	11
2.2	Location of the coal-fired power stations in Mpumalanga.	12
2.3	Energy flow in the South African electricity industry.	13
2.4	High-level representation of the energy flow simulator.	17
2.5	Coal stockpile level policies for 14 coal-fired power stations.	21
2.6	An example of the simulated burn of four coal-fired power stations.	25
2.7	An example of the simulated delivery of four coal-fired power stations.	26
2.8	An example of the difference in simulated delivery and burn of four coal-fired power stations.	26
2.9	The monetary value of coal stock on hand.	27
3.1	Characteristics of the (s, S) inventory process.	37
4.1	Integration of simulation and optimisation algorithms.	40
4.2	Classification of simulation optimisation techniques.	44
5.1	Controllable inputs, uncontrollable inputs and outputs of the primary energy module.	51
5.2	Four hypothetical scenarios for variation between coal delivery and burn.	52
5.3	Four hypothetical scenarios for the upper and lower warning limits.	54
5.4	The workings of the optimiser in the case of the lower warning limit.	55
5.5	The integrated simulation model and optimisation algorithm.	56
5.6	A typical example of how the CEM progresses over time.	62

LIST OF FIGURES

6.1	Additional cost for the emergency purchasing of coal, for a minimum stockpile level of five days and a target stockpile level of 20 days.	73
6.2	Additional cost for the emergency purchasing of coal, for two arbitrary cases for both the minimum and target stockpile level.	74
7.1	Stockpile level variation on a month-by-month basis.	86
7.2	Cumulative monthly stockpile variation from the initial stockpile level.	87
7.3	Progression of the objective function value for the base case.	89
7.4	Progression of the μ values of L_p for the base case.	90
7.5	Progression of the μ values of T_p for the base case.	90
7.6	Progression of the μ values of U_p for the base case.	91
7.7	Progression of the σ values of L_p for the base case.	91
7.8	Progression of the σ values of T_p for the base case.	92
7.9	Progression of the σ values of U_p for the base case.	92
7.10	Progression of the μ values of all three decision variables for the base case.	93
7.11	Final approximation of the decision variable values for the base case.	94
7.12	Emergency deliveries for the final approximation of the base case.	96
7.13	Delivery cancellations for the final approximation of the base case.	96
7.14	Final approximation of the decision variable values for the output statistic experiment.	103
7.15	Emergency deliveries for the final approximation of the 5-th percentile scenario.	104
7.16	Stockpile policies and the changes in stockpile levels for the final approximation of the 5-th percentile scenario.	105
7.17	Emergency deliveries for the final approximation of the 95-th percentile scenario.	107
7.18	Delivery cancellations for the final approximation of the 95-th percentile scenario.	107
7.19	Stockpile policies and the changes in stockpile levels for the final approximation of the 95-th percentile scenario.	108

LIST OF FIGURES

7.20	Final approximation of the decision variable values for the baseline delivery experiment.	110
7.21	Emergency deliveries for the final approximation of the reduced baseline deliveries scenario.	111
7.22	Delivery cancellations for the final approximation of the reduced baseline deliveries scenario.	111
7.23	Stockpile policies and the changes in stockpile levels for the final approximation of the reduced baseline deliveries scenario.	112
7.24	Emergency deliveries for the final approximation of the increased baseline deliveries scenario.	114
7.25	Stockpile policies and the changes in stockpile levels for the final approximation of the increased baseline deliveries scenario.	115
7.26	Comparison of the cumulative stockpile variation between the original and randomly varied baseline delivery.	116
7.27	Emergency deliveries for the final approximation of the randomly varied baseline deliveries scenario.	117
7.28	Delivery cancellations for the final approximation of the randomly varied baseline deliveries scenario.	117
7.29	Stockpile policies and the changes in stockpile levels for the final approximation of the randomly varied baseline deliveries scenario.	118
8.1	Suggested future research for the energy flow simulator.	122
B.1	Simulated burn of all fourteen coal-fired power stations.	B-2
B.2	Simulated delivery of all fourteen coal-fired power stations.	B-3
B.3	Difference in simulated delivery and burn of all fourteen coal-fired power stations.	B-4

List of Tables

2.1	Production parameters of the primary energy module.	22
2.2	Coal supply parameters of the primary energy module.	23
5.1	Notation used in the mathematical problem formulation.	58
5.2	CEM parameter settings.	66
6.1	Assigned values of the parameters for the base case.	76
6.2	The experiment design matrix for the sensitivity analysis of the financial parameter estimates.	79
6.3	Values of each factor in the sensitivity analysis.	79
6.4	Randomly generated scale factor values used to alter the baseline deliveries.	80
7.1	Final approximated values of the objective function.	97
7.2	Final approximated values of L_p measured in stockpile days.	98
7.3	Final approximated values of T_p measured in stockpile days.	99
7.4	Final approximated values of U_p measured in stockpile days.	100
7.5	Final approximated objective function values for sensitivity analysis of the financial parameter estimates.	101
A.1	Parameters in the EFS.	A-2

Nomenclature

Acronyms

API	Application programming interface
BD	Baseline delivery
CC	Delivery cancellation cost
CEM	Cross-entropy method
CSPS	Coal stockpile simulator
CV	Calorific value
DCA	A delivery cancellation
EFS	Energy flow simulator
EMD	An emergency delivery
EOQ	Economic order quantity
GBT	Gradient-based techniques
GDP	Gross domestic product
GUI	Graphical user interface
HC	Holding cost
IPP	Independent power producer
IRP	Integrated resource planning
JSON	Java-script object notation
lb	Lower bound of a decision variable
MC	Monte Carlo
NP	Non-deterministic polynomial time
OCGT	Open-cycle gas turbine
symbol	description
OCLF	Other capability loss factor
PCLF	Planned capability loss factor
pdf	Probability density function

PEM	Primary energy module
R	R programming language
SC	Shortage cost
SO	Simulation optimisation
SQL	Standard query language
SSM	Statistical selection methods
ub	Upper bound of a decision variable
UCLF	Unplanned capability loss factor
V&V	Verification and validation
ZAR	South African Rands
Greek Symbols	
α	Smoothing parameter for the cross-entropy method
δ	The holding cost per year per kton
κ	Penalty fee percentage for delivery cancellations
μ	Vector of mean values of the decision variables
ϕ	Vertical shift of the negative exponential pdf for the cost of emergency coal
Ψ	Scaling and normalisation function of $EC_{p,t}$
ρ	User-specified rare-event threshold value for the cross-entropy method
σ	Vector of standard deviation values of the decision variables
τ	Scaling value of $EC_{p,t}$
γ	Rare-event threshold of the cross-entropy method
φ	Mean of the negative exponential pdf for the cost of emergency coal
ϱ_i	Effect of factor i in sensitivity analysis.
ξ_i	Elite sample of the cross-entropy method at iteration i
Roman Symbols	
$A_{p,t}$	Actual stockpile level at the end of month t at power station p
$AL_{p,t}$	Additional load for power coal-fired station p at month-end t
$B_{p,t}$	Coal burnt at power coal-fired station p during month t

Nomenclature

$c_{p,t}$	Cancellations of deliveries during month t at power station p
C_p	Generating capacity of power coal-fired station p
$CV_{p,t}$	Calorific value of stockpiled coal at power coal-fired station p during month t
$D_{p,t}^{(\max)}$	Maximum allowable delivery to power station p during month t
$D_{p,t}^{(\min)}$	Minimum allowable delivery to power station p during month t
$D_{p,t}$	Baseline coal deliveries to power coal-fired station p during month t
$e_{p,t}$	Emergency deliveries during month t at power station p
$EAF_{p,t}$	Energy availability factor for power coal-fired station p during month t
$EC_{p,t}$	Emergency delivery cost for month t at power station p
$G_{p,t}$	Electricity generation at coal-fired power station p during month t
H_p	Heat rate for coal-fired power station p
i	Iteration of the cross-entropy method
L_p	Decision variable for the lower warning limit of power station p
$l^{(c)}$	Lead time for delivery cancellations
$l^{(e)}$	Lead time for emergency deliveries
m	Number of months in the primary energy module Monte Carlo simulation range
N	Population size for cross-entropy method algorithm
n	Number of coal-fired power stations
N_m	Maximum number of loops for the cross-entropy algorithm
$OCLF_{p,t}$	Other capability loss factor for power coal-fired station p during month t
$O_{p,t}$	Outages at coal-fired power station p during month t
p	Index for coal-fired power station
$PCLF_{p,t}$	Planned capability loss factor for power coal-fired station p during month t

Nomenclature

r	Rare-event probability in important sampling for the cross-entropy method
$R_{p,t}$	Resultant coal delivery, after emergency deliveries and delivery cancellations, during month t at power station p
$s_{p,t}$	Shortages in coal reserves during month t at power station p
$SG_{p,t}$	Scheduled electricity generation at coal-fired power station p during month t
$S_p^{(\max)}$	Maximum stockpile level of power station p
$S_p^{(\min)}$	Minimum stockpile level of power station p
$S_{p,t}$	Nett stockpile volume at power coal-fired station p at month-end t
t	Index for monthly time periods
T_p	Decision variable for the target stockpile level of power station p
U_p	Decision variable for the upper warning limit of power station p
$UCLF_{p,t}$	Unplanned capability loss factor for power coal-fired station p during month t
$w_{p,t}$	Binary variable representing shortages in coal reserves occurring or not during month t at power station p
$x_{p,t}$	Binary variable representing emergency deliveries occurring or not during month t at power station p
$y_{p,t}$	Binary variable representing delivery cancellations occurring or not during month t at power station p

Other Symbols

\mathcal{D}	Kullback-Leibler cross-entropy
\mathcal{P}	Set of indices for coal-fired power stations
\mathcal{R}	Response value from the 2^k factorial design method
\mathcal{T}	Set of indices for monthly time periods

Terminology

Baseline delivery	The coal deliveries based on long-term contracts
Sample path	A single Monte Carlo simulation replication
Simulation optimisation	Simulation optimisation is an iterative process whereby the PEM is nested within an optimisation model and run until convergence or some other stopping criteria is reached

Nomenclature

Simulation run	A simulation run refers to 1 000 independent replications of the PEM's Monte Carlo simulation model
Stockpile day	On average, the quantity of coal required to fuel power station p for one day
Stockpile level	The quantity of coal on hand at power station p

Chapter 1

Introduction

This chapter provides an overview of the problem addressed in the study. The problem statement, research aim and scope are defined. The methodology is put forth and a summary of the document structure is provided.

1.1 Background

Throughout the world, electric utilities form an essential foundation for nations' economies. The operations planning thereof is a highly complex process as it is characterised by decentralized decision-making and involves many processes structured in an intricate hierarchy (Hobbs, 1995; Liu *et al.*, 2005). In developing countries, electric utility operations planning is even more challenging. With the rapid economic growth of emerging economies, electricity demand is rapidly increasing (Asif & Muneer, 2007). Commission of power systems to meet such demand places huge pressure on often already constrained capital reserves. Adding to the problem, developing countries have to keep local energy policies in-line with the worldwide push to cleaner energy; which comes at a far greater cost than traditional methods, such as coal-fired power stations (D'Sa, 2005). Therefore, electric utility operations planners are assigned a daunting decision-making task.

In South Africa, the state-owned electricity utility Eskom provides for approximately 96% of the country's electricity needs (Eskom Holdings Limited, 2014a). Eskom is a vertically integrated monopoly which generates, transmits and distributes energy to customers in all sectors of society.

1.1 Background

In 1923, the Electricity Supply Commission (Escom) was established by the government of South Africa. ESCOM was also known by its Afrikaans name *Elektrisiteitsvoorsieningskommissie* (Eskom), and the two acronyms were combined as Eskom in 1986. During the 1970s and 1980s, the government embarked on a massive electricity capacity expansion programme. The excess supply resulted in Eskom not having to meet the ambitious targets set forth, which resulted in surplus capacity for two decades (Kessides *et al.*, 2007). Due to the excess supply, investment was not made in generation expansion. However, ironically after having excess capacity for many years, South Africa now faces a highly constrained demand and supply balance due to lack of investment in generation expansion. Poor planning on Eskom's behalf and a failed privatisation scheme on the government's behalf are mostly to blame.

Between 1999 and 2004, the government attempted to create a competitive electricity industry through the introduction of independent power producers (IPPs). In order to encourage IPP's to invest in the South African electricity industry, Eskom was prohibited from investing in capacity expansion programmes (Kessides *et al.*, 2007). However, due to insufficient political buy-in and various forms of resistance, the privatisation scheme failed and resulted in Eskom falling five years behind in future expansion investment (Eberhard, 2011). In addition, numerous delays in the integration of Medupi and Kusile coal-fired power stations into the national grid have further exacerbated Eskom's problems. Less efficient coal-fired power stations — decommissioned in the 1980's — were returned to service in order to provide more (much-needed) production capacity.

Recent demand-side management programmes have helped to reduce the peak load. An example is the recent *49 million* electricity savings campaign which urges the nation as a whole to reduce electricity usage during peak periods. Furthermore, various demand-side management technologies have been implemented, such as geyser timers and solar panels. Although improvements due to demand-side management have helped reduce peak load, the power system remains highly constrained and load shedding is a possibility when unforeseen events occur, particularly during peak periods.

In a further attempt to reduce consumption, Eskom introduced buy-back contracts. Eskom “buys back” electricity from large consumers during critical

1.1 Background

periods. Eskom reduces electricity supply to large consumers (in-line with the pre-determined buy-back contracts) and in turn those large consumers are compensated for production losses. The large consumers are typically companies within the manufacturing and mining sectors. The buy-back contracts are an innovative short-term solution, but in the long run stunt economic growth.

Regardless of demand reduction and supply expansion measures put in place by Eskom, the desired reserve margin of 15% is currently only 7.5% (Crowley & Janse van Vuuren, 2014). In 2008 — in addition to the system being heavily constrained — higher than expected electricity demand, unplanned generation unit outages, and unseasonal rain resulted in a national energy crisis. Blackouts began in the Western Cape and spread nationwide. In the subsequent years, the system has remained heavily constrained. Moreover, the system is constrained throughout the year — not only in the winter months — because maintenance work during the winter months is deferred to the summer months, which reduces the summer months' production capacity.

Running such a constrained system results in a limited ability to mitigate the effects of unexpected events and disturbances therein can have catastrophic effects. Planning is of utmost importance to ensure that the effects of those unforeseen events are mitigated. A holistic view of the electricity supply chain is needed to provide assistance to electric utility planners. The aforementioned shortcomings in the South African energy production industry have led to the development of an energy flow simulator (EFS) by an industry partner of Eskom. The said industry partner of Eskom is also the industry partner of this study.

The EFS models the energy supply chain from “fuel to fridge”. As electricity is predominantly produced by coal-fired power stations, one can also think of the EFS modelling “coal to consumption”. The energy supply chain is characterised by a few distinct functional areas, namely: primary energy supply, production planning, transmission, distribution, and end-use consumption. In addition, weather and system losses are incorporated. The effects of weather are ubiquitous throughout the simulation model, as exhibited in real-life.

The EFS is a stochastic, dynamic simulator aimed at mitigating the effects of uncertainties in the foreseeable future. Most importantly for decision makers, the EFS allows for *what-if* analysis of various scenarios. Scenarios include different

weather patterns, various gross domestic product levels, and the commission and decommission of power stations.

1.2 Problem statement

The values of many variables chosen by management within the EFS are known to be sub-optimal (van Harmelen, 2014b). Thus, there exists a need to optimise decision-making within the EFS. Tweaking certain parameters has the potential to increase efficiency and enable more informed decisions to be made. Adding more prescriptive capability will enable utility planners to further optimise decision making. Therefore, it is of national interest to optimise the EFS.

1.3 Aim and objectives

The *primary aim* is to develop an optimiser for the energy flow simulator. As this study is the first of a potential seven research projects for post-graduate students on the EFS, the *secondary aim* is to perform pioneering work for further research.

In order to achieve the desired research output, the aim is decomposed into the following objectives:

- Study the architecture of the existing EFS and relevant literature (**Chapters 2 and 3**).
- Study literature on simulation optimisation techniques (**Chapter 4**).
- Study further literature on the chosen optimisation technique (**Chapter 5**).
- Designate decision variables and determine an objective function to optimise (**Chapter 5**).
- Perform experiments on the optimiser (**Chapters 6 and 7**).
- Validate the combined simulation and optimisation model (**Chapter 7**).
- Provide recommendations (**Chapter 7**).
- Master the use of programming languages R and SQL, and the document preparation system \LaTeX .

1.4 Structure of the thesis

The document structure is based on the defined objectives and is as follows. **Chapter 2** first provides, as background, literature on the energy supply chain and modelling thereof. Subsequently, the EFS is decomposed and it is concluded that only one of the modules can be optimised as a whole, namely the primary energy energy module (PEM). The PEM is then further investigated and sets the platform for the rest of the study. The PEM module is essentially a coal stockpile simulator.

As optimising coal stockpile level policies is of interest, literature of coal and inventory management techniques are covered in **Chapter 3**. Simulation optimisation techniques are investigated in **Chapter 4**, in the search of finding an efficient algorithm that can be applied to the black box simulation of the PEM.

The problem and solution approach are formulated in **Chapter 5**. The conceptual model formulation is first described and then the mathematical model formulation is put forth. Subsequently, the chosen solution approach is detailed.

Experimental design is documented in **Chapter 6** and in **Chapter 7** analysis of experimental results is presented. Verification and validation is also covered in **Chapter 7**. Finally, the thesis is concluded in **Chapter 8**.

Additionally, two appendices are included. **Appendix A** documents the 260 types of parameters that are included in the existing energy flow simulator. Regarding the primary energy module, variation in the mass of coal burnt and delivered, at each power station and for each month, is illustrated in **Appendix B**.

Chapter 2

The energy flow simulator

The main aim of this chapter is to describe the existing energy flow simulator (EFS). This chapter builds upon Eskom's internal reports by [Eskom Holdings Limited \(2013b\)](#); [Ramjith \(2014\)](#); [van Harmelen \(2014a\)](#).

Background is provided on electricity sector modelling techniques used worldwide and an overview of the South African electricity sector is presented. Thereafter, the EFS is briefly discussed. After discussing the components of the EFS, the chosen area of optimisation is presented.

Before commencing with Chapter 2, this paragraph serves as a precursor. Strictly speaking, electricity falls under the greater field of energy. However, in this study electricity and energy will be used interchangeably. Coal-fired power stations provide for the majority of electricity generation in South Africa. Hence, the focus will be upon coal-fired power stations. Nuclear power stations, open-cycle gas turbines (OCGTs), hydroelectric power stations, and pumped storage schemes are included but are not a prominent feature. It is also important to note that Eskom is a vertically integrated monopoly. This has an important effect on the modelling process. For instance, cost minimization could be the objective because Eskom is a monopoly, as opposed to the standard approach of maximising a utility's profit.

2.1 Background to Chapter 2

In this section, modelling techniques applicable to electricity production are first briefly covered. Thereafter, an overview of the South African electricity sector is provided.

2.1.1 Literature review of electricity sector modelling techniques

This subsection is grouped into two parts. At first, modelling techniques applied to specific problems within the energy sector are analysed. Subsequently, the holistic modelling of the energy sector is investigated, treating the energy sector in a like manner to an economic supply and demand market.

The field of energy modelling and associated problems is rich in literature. Specifically, much high-quality research dates to the end of the 20th century (see [Heron \(1985\)](#); [Hobbs \(1995\)](#); [Schramm \(1993\)](#); [Scott & Read \(1996\)](#)).

2.1.1.1 Modelling problems within the sector

In energy industry literature, there are many specific optimisation problems. Some examples of such problems are economic dispatch ([Wood *et al.*, 2013](#)), generator maintenance scheduling ([Dahal & Chakpitak, 2007](#); [Schlünz, 2011](#); [Wood *et al.*, 2013](#)), coal handling process scheduling ([Conradie, 2011](#)), transmission losses ([Wood *et al.*, 2013](#)), hydrothermal coordination ([Wood *et al.*, 2013](#)), coal stockpile simulation ([Micali & Heunis, 2011](#)), optimal power flow ([Wood *et al.*, 2013](#)), generation expansion planning ([Sirikum *et al.*, 2007](#); [Tekiner *et al.*, 2010](#)), and state estimation problems ([Wood *et al.*, 2013](#)).

Additionally, there have been countless studies on load forecasting and demand reduction, namely: modelling household electricity-saving in South Africa by using a modified stochastic multi-criteria acceptability analysis approach ([Durbach & Davis, 2013](#)) and by using system dynamics ([Davis & Durbach, 2010](#)), including variants of autoregressive integrated moving average modelling for electricity demand ([Chikobvu & Sigauke, 2012](#); [Ediger & Akar, 2007](#)), and a seasonal demand forecasting hybrid procedure ([Zhu *et al.*, 2011](#)). [Bhattacharyya & Timilsina \(2009\)](#) and [Suganthi & Samuel \(2012\)](#) provide surveys of energy demand models.

Initially, power generation studies focused upon very specific problems, such as optimal power flow ([Wood *et al.*, 2013](#)). Nowadays, models are being developed

2.1 Background to Chapter 2

for integrated resource planning (IRP). IRP aims to incorporate all facets of the energy supply chain, within a specific region (Loulou & Labriet, 2008; Seebregts *et al.*, 2002; Ventosa *et al.*, 2005). The future trend is an increased focus towards modelling the entire global energy supply chain, with an emphasis on emissions (Loulou & Labriet, 2008; Seebregts *et al.*, 2002).

2.1.1.2 Modelling the sector as a whole

Given the importance of efficient and reliable energy planning, many decades of research have culminated into countless energy flow planning models (Ventosa *et al.*, 2005). Essentially, the aim of electric power utility planning is to provide reliable electricity at an acceptable economic and environmental cost (Hobbs, 1995). Ventosa *et al.* (2005) provide a survey of three types of electricity modelling techniques: simulation, equilibrium and optimisation. Equilibrium models incorporate multiple utilities, and are thus irrelevant to this single utility study. Usually, only one of the three modelling techniques is used. However, that is where this study differs because simulation and optimisation models are used.

Advances in electricity market planning methods can be separated into five main groups (Ge *et al.*, 2011; Zhou *et al.*, 2007):

1. traditional programming
2. mathematical programming
3. metaheuristics
4. integrated resource planning
5. agent-based modelling.

Firstly, in the pre-1960s, *traditional programming* was used to determine when and where to locate generation units based on available capacity, with no focus upon consumption.

Secondly, in the post-1960s, developments in operations research and increased computing power allowed for the application of *mathematical programming* techniques.

2.1 Background to Chapter 2

Thirdly, although not originally termed “metaheuristics” (it is still debatable as to whether metaheuristics is the all encompassing term to describe the field), *metaheuristics* have been applied since Holland’s genetic algorithm was formulated in 1975 and they are still being applied in energy market modelling to this day.

Fourthly, *IRP* places greater emphasis on incorporating environmental cost by including emissions and greener technologies in modelling.

Finally, recent studies apply *agent-based modelling* techniques to energy markets. Agent-based modelling is a fairly recent simulation approach to modelling systems comprised of dynamically interacting, autonomous agents with the aim of analysing their effects on the system as a whole (Macal & North, 2007).

In this study, the EFS is a holistic representation of the energy flow supply chain. Thus, the EFS can be viewed as an *IRP* method. *Mathematical programming* techniques are used in the simulator. For instance, linear programming is used to minimise cost in the scheduling of electricity generation per coal-fired power station per time period. The optimiser makes use of a *metaheuristic* technique, which will be discussed in Section 5.3.

2.1.2 An overview of the energy supply chain

The energy sector is comprised of many complex and interlinked systems and processes which are structured in an intricate hierarchy. For the purposes of this study, only medium-to-high level key characteristics will be analysed.

For this utility, the bulk of electricity is generated through coal-fired power stations. In 2012, 83% of the 41 933 MW generation capacity consisted of coal-fired power stations (Eskom Holdings Limited, 2012b). Currently 23 power stations are operational, with four new power stations being commissioned. Eskom’s power stations are classified as follows (Eskom Holdings Limited, 2014d):

- base load stations
- return-to-service
- peak demand stations
- renewable energy

2.1 Background to Chapter 2

- new build
- distribution.

Base load stations are comprised of one nuclear power station (Koeberg), and 10 coal-fired power stations with generating capacities of between 2 000 MW and 4 110 MW. Koeberg has a generating capacity of 1 940 MW.

There are three return-to-service coal-fired power stations that were mothballed in 1990 — namely Camden, Grootvlei, and Komati — with generating capacities of between 940 MW and 1 510 MW.

Peak demand stations include hydroelectric, pumped storage scheme and OCGT power systems. Generation capacities of peak demand stations vary between 171 MW and 1 338 MW.

The Klipheuwel Wind Facility is a 3 MW capacity pilot project for the Sere Wind Facility — with a planned generation capacity of 100 MW — which is part of the new build power stations.

Also part of the new build programme is a solar power station (100 MW), Ingula pumped storage scheme (1 332 MW), Kusile coal-fired power station (4 800 MW), and Medupi coal-fired power station (4 788 MW). There have been numerous delays in the construction of both Medupi and Kusile power stations.

Distribution is comprised of four hydroelectric power stations ranging from 2 MW to 42 MW. They are used to stabilise the distribution network in the Eastern Cape.

Figure 2.1 depicts 22 of the 23 power stations in South Africa. The Klipheuwel Wind Facility is excluded because it has a generating capacity of only 3 MW and is thus deemed insignificant. Three power stations are located near Cape Town: Koeberg nuclear power station, Acacia OCGT and Ankerlig OCGT. The majority of electricity is produced by coal-fired power stations and thus 11 of the 13 coal-fired power stations are located in Mpumalanga, adjacent to the Witbank, Highveld, and Ermelo coalfields. Figure 2.2 depicts the power stations in Mpumalanga.

The two remaining coal-fired power stations are Matimba and Lethabo. Matimba is located adjacent to the Waterberg coalfield in Limpopo. Lethabo is situated adjacent to the Sasolburg-Vereeniging coalfield in the Free State. Jeffrey

2.1 Background to Chapter 2

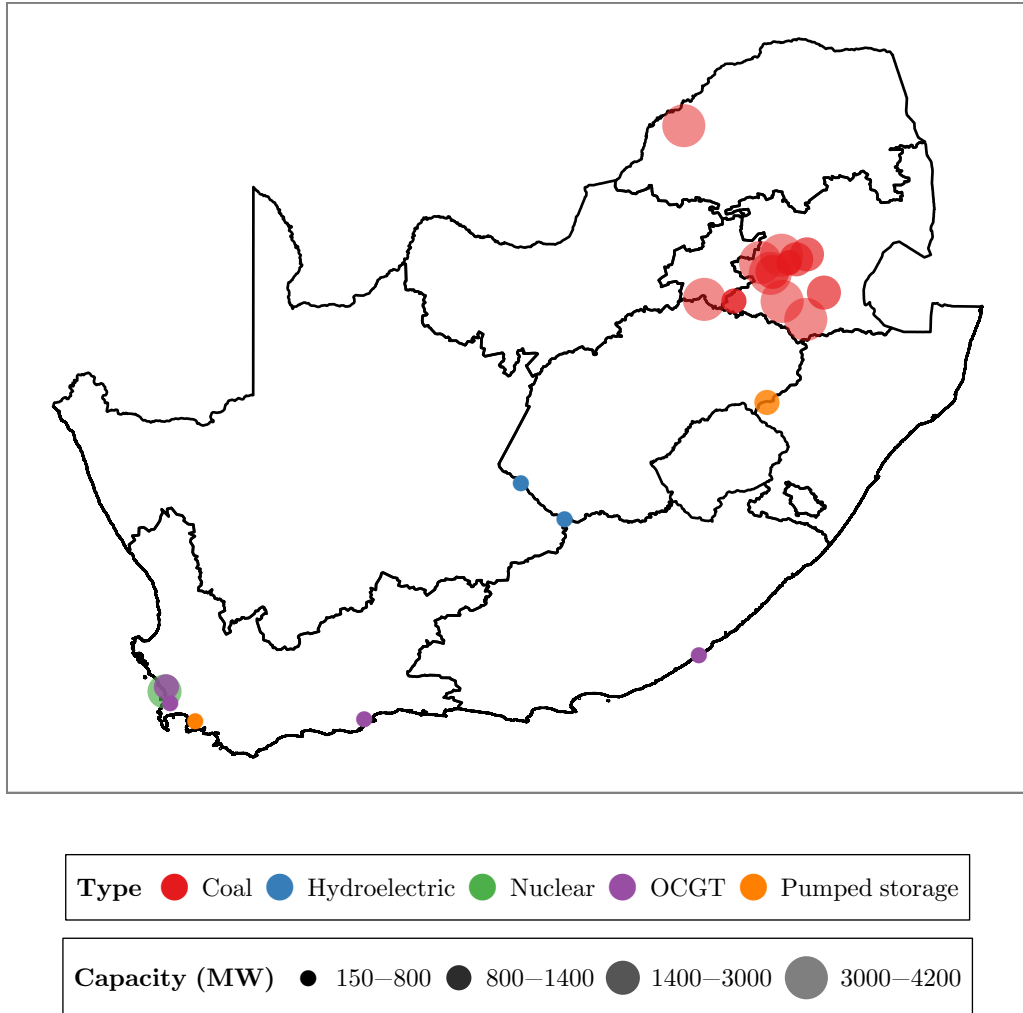


Figure 2.1: Distribution of power stations in South Africa.

(2005) provides an analysis of the characterisation of coal resources in South Africa. The nett cost of coal for each coal-fired power station is significantly different due to varying coal quality, different proximities from the mines (related to transportation cost), and the cost associated with different modes of transport.

Figure 2.3 provides a decomposition of the energy flow in South Africa. An overview of the four main functional areas represented in Figure 2.3, namely: generation, transmission, distribution, and consumption — along with primary energy — is provided. The effect of weather is also mentioned.

2.1 Background to Chapter 2

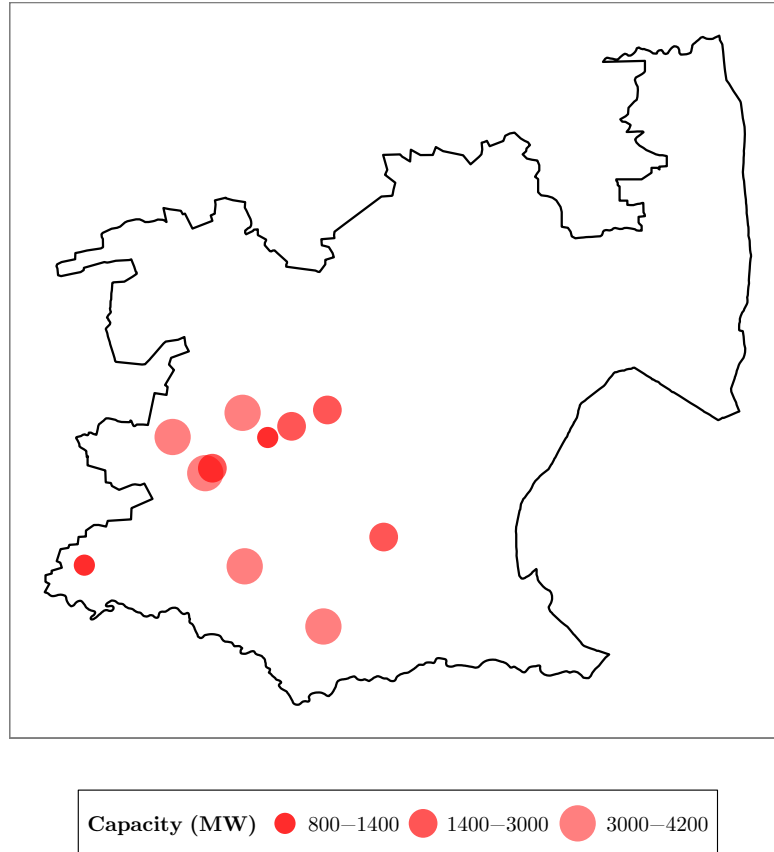


Figure 2.2: Location of the coal-fired power stations in Mpumalanga.

Primary energy

The majority of electricity produced is through coal-fired power stations and thus primary energy is mostly focussed on coal operations. The quality of coal received from the majority of suppliers is of high ash content and thus classified as low-grade (Eberhard, 2011). High ash content reduces the calorific value (CV), resulting in lower thermal efficiency during combustion. In addition, high ash content leads to increased coal cleaning costs. CV refers to the amount of stored chemical energy converted to thermal energy upon combustion.

2.1 Background to Chapter 2

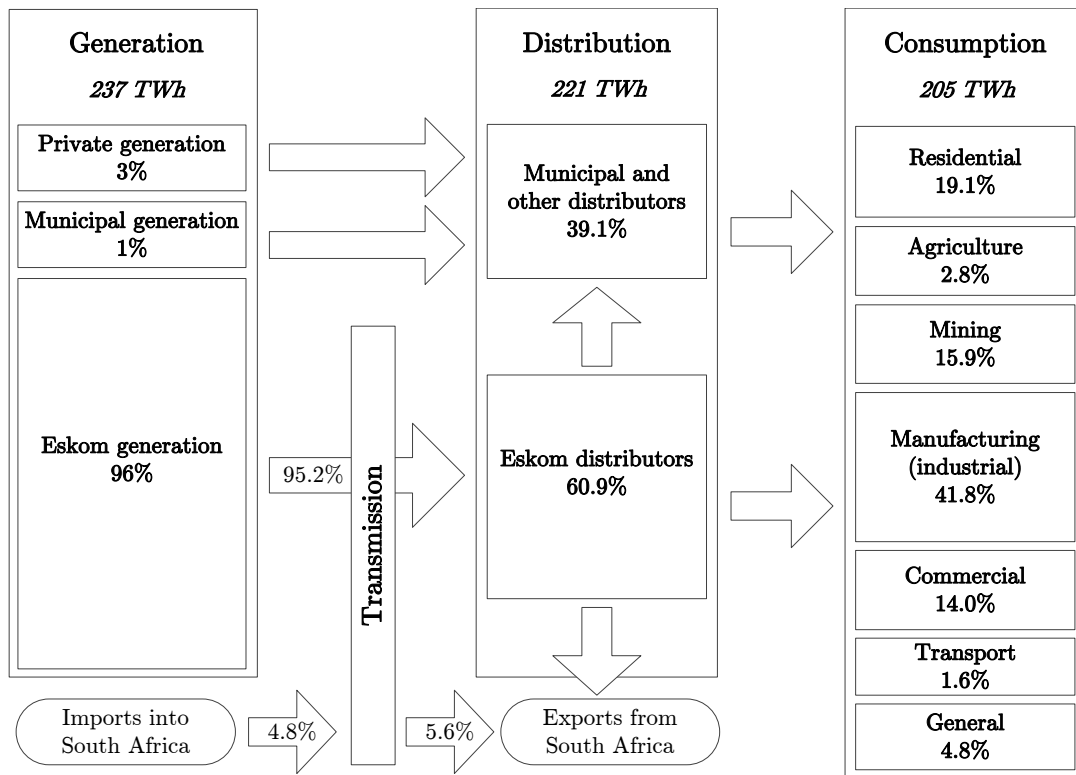


Figure 2.3: Energy flow in the South African electricity industry (National Energy Regulator of South Africa, 2006).

Generation

As seen in Figure 2.3, independent power producers (IPPs) and municipalities account for approximately only four percent of electricity generation. IPPs produce their own electricity and sell surplus electricity to municipalities. Municipal producers provide all their electricity to municipal distributors. Considering generation of electricity by Eskom, generation costs differ significantly between coal-fired power stations. Furthermore, generation costs differ vastly between nuclear, hydroelectric, OCGT, pumped storage, and coal-fired power systems — with OCGT's being extremely expensive to operate.

Transmission

Output from Eskom's generation is provided to Eskom distributors. This is done through the transmission system. Electricity is also imported and exported within the Southern African Power Pool through the transmission system.

2.1 Background to Chapter 2

Distribution

As seen in Figure 2.3, Eskom is responsible for the distribution of about 60% of South Africa's electricity. Municipalities and small distributors are responsible for the remainder of the electricity distribution.

Consumption

Consumption is segregated into seven types of consumers, as illustrated in Figure 2.3. Listed in descending order, the largest consumers are: manufacturing, residential, mining, and commercial. South Africa has a unique demand curve when compared to other countries. One of the contributing factors is the use of electricity for cooking — many European countries use gas for cooking which reduces the evening peak demand.

Weather

Weather affects both the supply-side and demand-side. Cold and very warm weather both increase consumption, for example the operation of heating units during winter and air-conditioning systems during summer. Also, large quantities of rainfall can have a negative impact on reliability, particularly for open-cast mines and coal stockpiles as they are unsheltered. Eskom uses third grade coal of low CV to produce electricity. The low quality coal is comprised of a high ash content. When mixed with water, it forms a slurry which is troublesome to transport along conveyor belts and can block inlet pipes leading to the combustion chambers.

The future of South African electricity generation is governed by a two-decade IRP. The IRP outlines South Africa's current energy situation, documents energy policies, and projects future foreseeable scenarios. In 2010, the South African Department of Energy released the IRP for the years 2010 to 2030. Refer to [The Department of Energy \(2013\)](#) for the most recent version of the document at the time of writing. The IRP provides a mandate for a cleaner energy mix, which includes the commission of nuclear and renewable energy resources by 2030.

2.2 An introduction to the energy flow simulator

2.2 An introduction to the energy flow simulator

“Simulation” is a word used in many contexts, such as flight simulation package, simulation training, and computer games. For the purposes of this study, *simulation* is defined as the imitation of the energy flow from primary energy to end-use consumption.

Simulation is a powerful and versatile tool to model different types of systems and environments. The *simulation model* is an abstracted and simplified representation of Eskom’s energy flow process.

2.2.1 The need for the energy flow simulator

When uncertainty exists so does an element of risk. The simulator aims to quantify the risk so that it can be suitably managed.

Furthermore, Eskom has carried out many projects over the years, and there is a need to integrate successful research and projects into a single decision support system ([Eskom Holdings Limited, 2013b](#)).

2.2.2 The aim of the energy flow simulator

The EFS serves to simulate the impact of uncertainties in strategic, tactical and operational planning of energy flow from source to end-use. By incorporating the randomness of and the interdependence between parameters, the EFS models — predominantly through simulation — the complexity of energy flow which would otherwise be too difficult to model analytically. However, given the scale and complexity of energy sector systems, it would still be impossible to incorporate all the detailed aspects of the system. Therefore, the simulator has been developed with the intention of modelling only the key, medium-to-high level characteristics.

2.2.3 An overview of the energy flow simulator

As a supply and demand equilibrium must be met, the electric utility’s system models primary energy, production planning and consumption; which culminate in the final electricity generation plan. Weather effect on the system and system losses are incorporated. Planning forms an integral part of the simulator and thus the model is designed to allow for *what-if* scenarios. Typical what-if scenarios are

2.3 Components of the energy flow simulator

economic growth and decline, unforeseen weather conditions, and the commission and decommission of power stations.

The simulator has 260 types of parameters. Each type of parameter has tens, hundreds or thousands of instances, all of which are stored in a large database. Many parameters are defined by a probability function, whilst other parameters are assigned a set value. A few examples of parameters are the *average ash content of coal delivered*, *residential area demand*, *time-of-use tariff for the mining sector*, and *sulphur dioxide factor*. The reader is referred to Table A.1 (in Appendix A) for a listing of all the parameters used in the EFS.

The modules of the EFS are not run simultaneously, but rather on a module by module basis. Provided that input and output data already exist for each module, the modules can be run one-by-one in any order. Referring to Figure 2.4, typically modules are run in an anti-clockwise manner.

First, the *consumption* module is run to forecast approximately 720 hourly demand profiles (24 hours per day for 30 days in a month). Secondly, in the *Production planning* module, a linear programming model schedules production to minimise production cost subject to various constraints. Thirdly, the *Primary Energy* module determines the expected effect of production and delivery variation on coal stockpile levels — on a monthly basis. Fourthly, the outputs from the previous three modules culminate in the *Generation plan* module which quantifies emissions and time-of-use tariffs.

2.3 Components of the energy flow simulator

Figure 2.4 represents an overview of the interaction between the modules of the EFS, namely:

1. consumption
2. production planning
3. primary energy
4. generation.

2.3 Components of the energy flow simulator

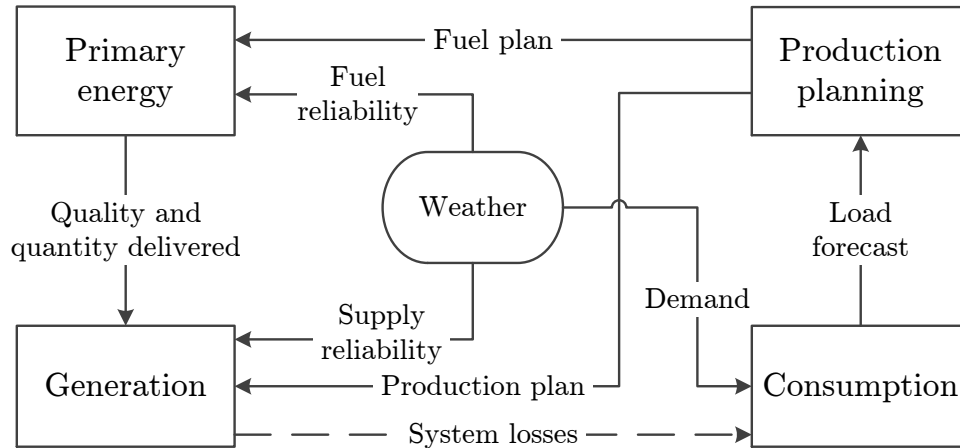


Figure 2.4: High-level representation of the energy flow simulator (van Harmelen, 2014a).

The first module is *consumption*. Simply put, consumption (demand) is forecast per region and customer type. The country is split into three main regions: Central, Eastern, and Southern. Customers are segmented into four types: manufacturing, mining, residential, and other. A culmination of previous research showed that each of the four customer types exhibits a different demand curve (van Harmelen, 2014a). Segmenting customers allows for more accurate demand forecasting. In addition to region and customer type, demand is forecast based on the selected *gross domestic product* (GDP) scenario and weather scenario. GDP scenarios are defined as high, medium, and low. The EFS assumes a correlation between GDP and electricity demand — the higher the GDP, the more electricity is required to provide for greater economic growth. Weather scenarios are created by analysing historic weather data and determining profiles for hot, normal, and cold years. From 20 years of historical data, Monte Carlo (MC) simulation is used to forecast weather scenarios. Based on the chosen GDP scenario (high, medium, or low) and weather scenario (hot, normal, or cold), hourly profiles per region, customer type, and month are created. Developing hourly profiles, amongst other benefits, helps to test the effect of time-of-use tariffs and demand-side management technologies.

2.3 Components of the energy flow simulator

Apart from reducing overall demand, both schemes aim to specifically reduce peak demand by levelling demand.

The second module is *production planning*. The load forecast is aggregated to monthly intervals. Smaller time intervals would be preferable, although in reality energy utility planners tend to plan on a month-to-month basis. The production planning module schedules the planned energy production per coal-fired power station to minimise production cost. Coal-fired power stations that have cheaper production costs are scheduled first. Demand must be met whilst taking into account production capacity. Linear programming is used to schedule electricity production to minimise total production cost. Although the forecast demand inputs are on an hourly basis, the electricity production schedule sent out is aggregated to a monthly basis. To allow for more realistic production schedules, the outputs should have a finer resolution, such as daily or better yet hourly.

The third module is *primary energy*. The primary energy module (PEM) deals exclusively with coal, because coal-fired power plants produce most of the utility's electricity. Through MC simulation, the PEM simulates coal supply and demand unreliability. The PEM incorporates and builds upon the work of [Micali & Heunis \(2011\)](#) who developed a coal stockpile simulator. They describe the coal stockpile simulator as a dynamic model that enables *what-if* analysis to evaluate different plans and scenarios. Examples of stochastic inputs include unplanned coal-fired power station maintenance, variation in the CV of coal, and variation in the quantity of coal delivered. Weather plays an important role in the reliability of coal, because open cast mines and stockpiles at the coal-fired power stations are not sheltered from the rain.

The fourth module is *generation*. The production plan, supply reliability, and the quality and quantity of coal to be delivered are fed into the generation module. The generation module then quantifies emissions, such as sulphur and nitrogen oxides. The generation module also provides a time-of-use tariff report. In addition, system losses are incorporated as the estimated percentage of losses through electricity supply. Correctly estimating system losses is important to reflect the true supply and demand balance.

2.4 Operating the energy flow simulator

2.4 Operating the energy flow simulator

The EFS can be operated in two ways, through the graphical user interface (GUI) or the application programming interface (API)¹.

The GUI provides an easy-to-use dashboard-like environment for electric utility planners, many of whom do not have programming experience. Parameters can easily be adjusted and *what-if* analyses can be performed.

A limitation in using the GUI is that each time a simulation run is performed, it is initiated by the user through the click of a button. In simulation optimisation this is impractical and thus the API is favoured as it allows for remote triggering of the PEM.

At the time of writing, an API has been created for the R programming language (R). R was developed by [R Core Team \(2014\)](#). R is a freely available programming language designed for statistical and data analysis. As an API was created for R, the optimisation algorithm was developed in R. At a later stage, the API will be expanded to Java-script Object Notation (JSON), by the industry partner of this study. JSON is a standardised, lightweight format used to transmit data across different software platforms.

The EFS makes use of the H2 database to store inputs and outputs of the simulation to the database. H2 is a Java-based SQL database. It is fast, open source, can run in embedded or server mode, and has a small memory footprint ([Mueller, 2014](#)). Another advantage of the H2 database is that it is simple to integrate into Java applications, such as the Java application used to provide the GUI of the EFS.

2.5 Chosen area of optimisation: Primary energy module

Many values for controllable inputs that have been chosen by management are known to be sub-optimal ([van Harmelen, 2014b](#)). Originally, it was thought that global optimisation of the EFS could be performed. However, after gaining an understanding of the internal workings of the EFS, it was seen that global

¹An API is a set of rules, protocols, and tools detailing how software programs should communicate. APIs aid programmers in developing software by providing the interface for one software application to communicate with another software application.

2.5 Chosen area of optimisation: Primary energy module

optimisation was not possible. Therefore, optimisation of only one of the four modules was investigated. [Hatton & Bekker \(2014\)](#) step through the challenges and approach followed in developing an optimiser for the EFS.

After scrutinising the 260 types of parameters in the EFS (see [Table A.1](#)) and through personal communication ([van Harmelen, 2014b](#)), it was decided that the EFS module with the most potential for optimisation was primary energy. Not because the PEM was poorly modelled, but rather the contrary. The PEM was built around the Coal Stockpile Simulator (CSPS) developed by [Micali & Heunis \(2011\)](#). The CSPS was developed with the intention that an optimisation capability be added at a later stage ([Micali & Heunis, 2011](#)).

2.5.1 Why the primary energy module was chosen

The traditional means of determining coal stockpile size, for this electric utility, was to set the minimum level at 20 stockpile days and the maximum level at the capacity of the stockpile yard. A coal stockpile day is equivalent to the average daily quantity of coal burnt. As expected, the amount of coal required for one stockpile day differs from power station to power station. Hence, 20 stockpile days will, on average, provide a coal-fired power station with 20 days of fuel. The stockpiles act as buffers to mitigate the effect of unforeseen events.

With growth in electricity demand and lack of generation capacity expansion both reducing the reserve margin (system's overall capacity), the load factor at each coal-fired power station has subsequently been increased. The higher the load factor required, the more coal to be burnt. Many coal contracts could not provide for the increased burn requirements and hence additional suppliers were contracted. Many of the additional suppliers were contracted on a short-term basis, placing more risk in the system and calling for a sophisticated coal stockpile management system. Short to medium term contracts account for about one quarter of Eskom's coal supply ([Eberhard, 2011](#)). The increase in risk of supply lead to the development of the CSPS by [Micali & Heunis \(2011\)](#). In short, the CSPS simulates uncertainty in the delivery rate and burn rate of coal.

In 2008, increasing demand and a lack of additional capacity coupled with unexpected events caused nationwide blackouts. Some of the problems were attributed to depleted stockpiles at a few of the coal-fired power stations. One of

2.5 Chosen area of optimisation: Primary energy module

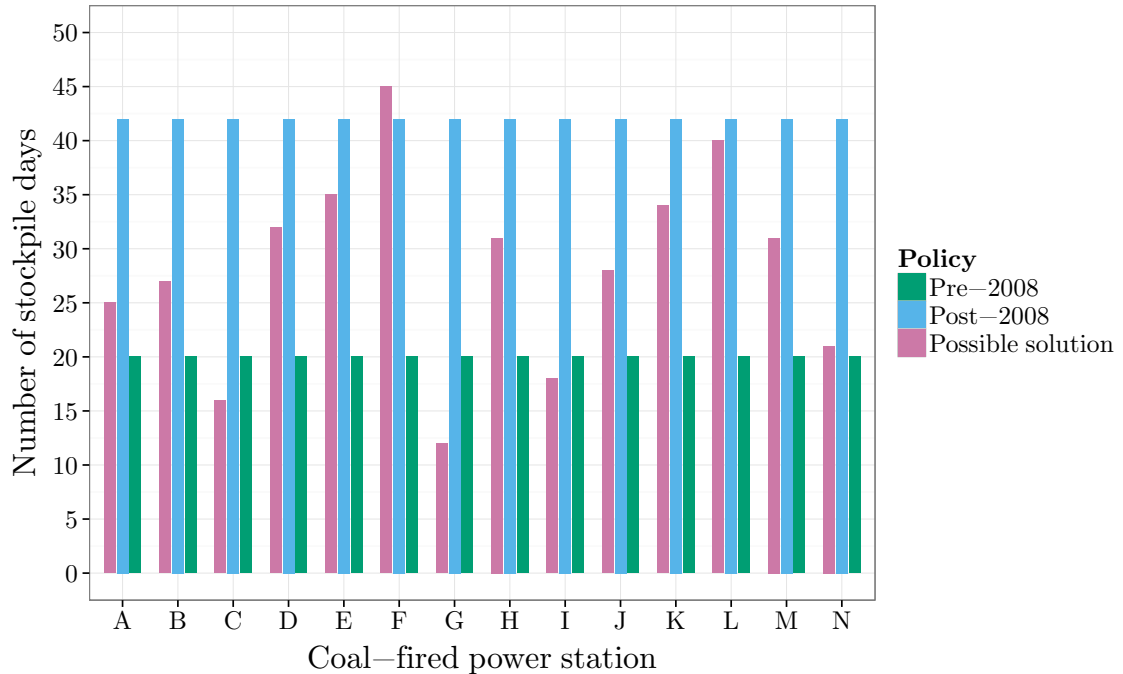


Figure 2.5: Coal stockpile level policies for 14 coal-fired power stations.

the actions taken by senior management was to increase the minimum stockpile level from 20 stockpile days to 42 stockpile days, for all powerstations. Still the traditional approach, but with a much higher lower limit on stockpiles levels.

The problem with the traditional approach is that each coal-fired power station exhibits different behaviour due to variation (Micali & Heunis, 2011; van Harmelen, 2014a,b). Some of the contributing factors to variation in the system are unplanned power station outages and the different degrees of coal reliability. Some stockpiles may require larger buffers due to high levels of variation, whilst other stockpiles may exhibit low variation and thus require much smaller stockpiles. Thus, the purpose of this research is to optimise stockpile level policies. Figure 2.5 depicts such a possible solution, along with the management policy before and after 2008. For purposes of confidentiality, coal-fired power stations are not referred to by name and are labelled alphabetically.

2.5 Chosen area of optimisation: Primary energy module

2.5.2 An overview of the primary energy module

The variables in the PEM stem from two sets of inputs, namely the:

1. production plan
2. coal supply plan.

Table 2.1: Production parameters of the primary energy module.

Parameter	Unit of measurement
Capacity	MW
Maximum generator load factor	%
Energy sent out	GWh
Coal burnt	ktons
Heat rate	MJ/kWh
Planned capability loss factor (PCLF)	%
Unplanned capability loss factor (UCLF)	%
Other capability loss factor (OCLF)	%
PCLF – standard deviation	%
UCLF – distribution type	pdf
UCLF – location	%
UCLF – shape	%
UCLF – shift	%

Table 2.1 specifies the production parameters used as inputs to the PEM. Capacity is the maximum energy a power station can generate in perfect conditions. The maximum generator load factor is the maximum percentage per month that a power station can be operated, usually approximately 96%. The 4% drop is attributed to downtime such as mandatory maintenance in monthly windows. The coal burnt is proportional to the energy sent out, but varies according to the CV. Heat rate is a means of measuring power station efficiency ([The Engineering Toolbox, 2014](#)). Downtime for power stations is modelled as a capability loss factor and is grouped into three categories: planned, unplanned and other. *Planned* refers to scheduled maintenance. *Unplanned* refers to breakdowns, often as a result of lack of planned maintenance. *Other* refers to extraordinary events such as employee strikes and theft of transmission cables. Variation in planned

2.5 Chosen area of optimisation: Primary energy module

maintenance is incorporated through a standard deviation. UCLF is disaggregated randomly throughout the month through the use of probability density functions (pdfs).

Table 2.2: Coal supply parameters of the primary energy module.

Parameter	Unit of measurement
Nett coal delivery reliability – mean	ktons
Nett coal delivery reliability – standard deviation	ktons
Cost of supply – mean	ZAR/MWh
Cost of supply – standard deviation	ZAR/MWh
Calorific value	MJ/kg
Maximum mine production	ktons
Maximum stockyard volume	ktons
Initial stockpile	ktons

Table 2.2 specifies the coal supply parameters used as inputs to the PEM. Nett delivery and cost of supply are subject to variation from suppliers. The higher the CV, the more energy is produced during combustion and less coal is required to meet electricity demand. Stockyards and coal production from mines are subject to capacity constraints. The stockpiles levels at the beginning of the specified simulation range need to be defined.

The equations (2.1) to (2.5) are used to calculate the resultant stockpile volumes at each power station $p \in \mathcal{P} = \{1, \dots, n\}$ for each month $t \in \mathcal{T} = \{1, \dots, m\}$.

In (2.1), the energy availability factor is calculated based on the three capability loss factors. As a rule of thumb, the monthly percentages of $\text{PCLF}_{p,t}$, $\text{UCLF}_{p,t}$ and $\text{OCLF}_{p,t}$ are approximately 10%, 9% and 1% respectively. Before running the PEM, the capability loss factor monthly percentages for each power station are defined. $\text{PCLF}_{p,t}$ is varied according to a standard deviation, as defined in Table 2.1. A point estimate for $\text{UCLF}_{p,t}$ is randomly sampled from a pdf, also as defined in Table 2.1. $\text{OCLF}_{p,t}$ is assigned an exact value, usually 1%. The energy availability factor ($\text{EAF}_{p,t}$) is approximately 80% per power station, but varies according to the point estimates of capability loss factors, and is calculated using

$$\text{EAF}_{p,t} = 1 - (\text{PCLF}_{p,t} + \text{UCLF}_{p,t} + \text{OCLF}_{p,t}). \quad (2.1)$$

2.5 Chosen area of optimisation: Primary energy module

Point estimates for outages

$$O_{p,t} = (1 - \text{EAF}_{p,t}) \cdot C_p \quad (2.2)$$

are calculated from the energy availability factor percentage $\text{EAF}_{p,t}$ and the maximum capacity C_p (in MW). The point estimate for generation

$$G_{p,t} = \text{SG}_{p,t} - O_{p,t} + \text{AL}_{p,t} \quad (2.3)$$

is determined by the scheduled generation $\text{SG}_{p,t}$, outages $O_{p,t}$ and the additional load $\text{AL}_{p,t}$. Operational rescheduling occurs due to $O_{p,t}$ and results in $\text{AL}_{p,t}$ at certain power stations in order to meet electricity demand. The scheduled generation $\text{SG}_{p,t}$ is an output from the linear programming solver discussed in Section 2.3.

The point estimates for coal burnt

$$B_{p,t} = G_{p,t} \cdot \frac{H_p}{\text{CV}_{p,t}} \quad (2.4)$$

are determined by the electricity generation $G_{p,t}$, heat rate H_p , and the calorific value of coal supplied $\text{CV}_{p,t}$. H_p differs per power station. Point estimates for $\text{CV}_{p,t}$ are determined from a pdf.

Finally, the resultant nett stockpile levels, at month-end,

$$A_{p,t} = A_{p,t-1} + D_{p,t} - B_{p,t} \quad (2.5)$$

are determined from the nett stockpile levels at the end of the previous month $A_{p,t-1}$, the estimated delivery $D_{p,t}$, and the $B_{p,t}$. Delivery is defined by a pdf. Burn is as a result of the relationship of the variables in (2.1) to (2.4). The PEM was developed in a generic manner to allow for any pdf. Needless to say, the initial stockpile level at the beginning of the planning horizon is $A_{p,0}$.

2.5.3 Analysis of the primary energy module

The stochastic nature of the PEM is investigated in this subsection, as illustrated by Figures 2.6, 2.7, and 2.8. The box-and-whisker plots represent the mean output values of 500 simulation runs. Four power stations, out of a total of 14, were

2.5 Chosen area of optimisation: Primary energy module

arbitrarily chosen for analysis. In this instance, the planning horizon was eight months. In Appendix B, Figures B.1, B.2, and B.3, depict the stochastic nature of all 14 power stations.

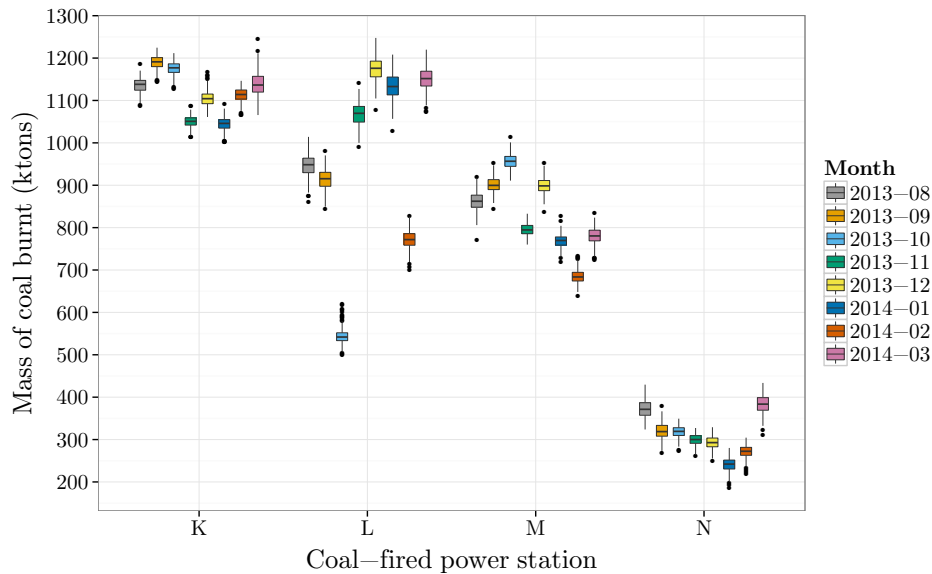


Figure 2.6: An example of the simulated burn of four coal-fired power stations.

An example of the variation in coal burnt per power station per month is shown in Figure 2.6. Power station L exhibits slightly more variation than Power station K, although Power station L burns a slightly lower mass of coal. The lower mass of coal burnt at Power station L during month “2013-10” is as a result of a high level of maintenance during the said month. The mass of coal burnt during each month for all the power stations differs significantly.

An example of the variation in the mass of coal delivered per power station per month is shown in Figure 2.7. Interestingly, Power station L exhibits a far greater variation than power station K. Power station N has a smaller number of deliveries and variation in those deliveries. The mass of coal delivered differs significantly per month for Power station K.

The difference in variation between coal delivered and burnt per power station per month is depicted in Figure 2.8. On average, coal is under delivered for Power stations L and M. Coal is over delivered for Power station N, while Power station L exhibits large variation.

2.5 Chosen area of optimisation: Primary energy module

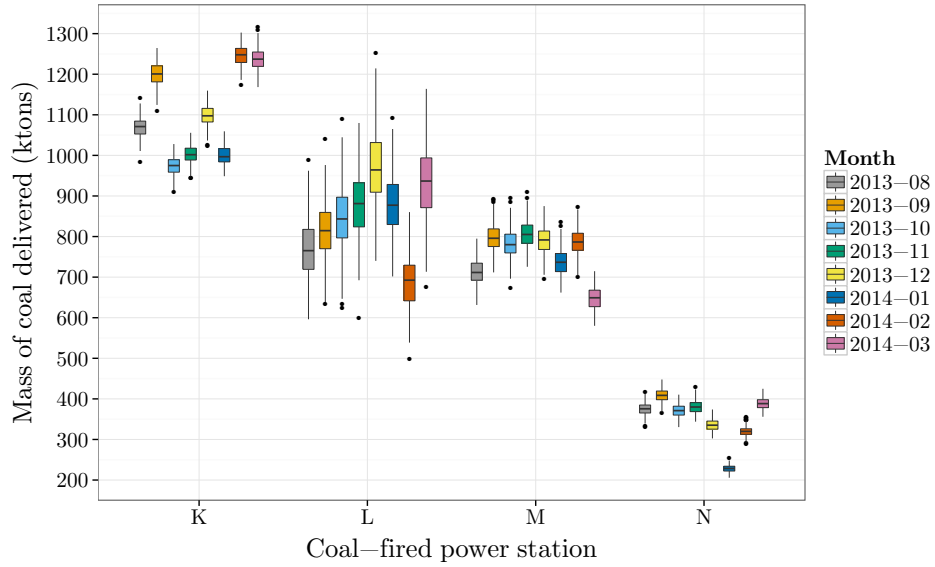


Figure 2.7: An example of the simulated delivery of four coal-fired power stations.




Figure 2.8: An example of the difference in simulated delivery and burn of four coal-fired power stations.

2.5 Chosen area of optimisation: Primary energy module

2.5.4 The value of coal inventory on hand

The value of coal inventory on hand for the national electric utility from 2002 to 2014 is depicted in Figure 2.9. The values are taken from Eskom's financial statements (see [Eskom Holdings Limited \(2003, 2005, 2007, 2008, 2009, 2010, 2011a, 2012a, 2013a, 2014b\)](#)). Inventory levels have risen sharply since 2008. Between 2006 and 2008, stockpiles levels dipped dangerously low and in 2008 nationwide load-shedding was partially attributed to the lack of coal reserves at certain power stations.

In this already highly constrained system, disturbances in coal supply will have severe impacts. In 2008, to mitigate against coal supply disturbances and ensure sufficient fuel for each power station, management increased stockpiles levels at all the power stations to 42 stockpile days. The value of the stockpiles' levels increased from below ZAR 0.7 billion in 2008 to over ZAR 5.3 billion in 2014 — a 655% increase in the monetary value of stockpile levels in six years. Such a tremendous increase may in fact be too much and the hypothesis is that opportunity costs have arisen. Perhaps, the capital tied up in coal stockpiles could rather be used to finance projects in other business domains in the electric utility.




Figure 2.9: The monetary value of coal stock on hand.

2.5 Chosen area of optimisation: Primary energy module

It is important to note that Eskom's financial year-end changed from 31 December 2003 to 31 March 2005, resulting in a 15 month reporting period. Therefore, there is no value for financial year-end 2004.

2.5.5 Classification of the primary energy module as a simulation model

In the field of Operations Research, simulation is defined as the imitation of a real-world system as it evolves probabilistically over time (Hillier & Lieberman, 2001; Winston, 2004). Operations Research practitioners often use simulation to model systems that are too complex to model analytically (Hillier & Lieberman, 2001).

Simulation is a descriptive modelling approach that involves generating an artificial archive of a system, based on modelling assumptions, and observing the artificial archive to draw inferences about the characteristics of the real-life system that is modelled (Banks *et al.*, 1998). In a simulation modelling mindset, systems can be classified by three dimensions (Law *et al.*, 2000; Winston, 2004):

Discrete or continuous. In continuous systems, such as a chemical process, the state of the system is continuously changing. In discrete systems, one is only interested in events occurring at distinct points in time that result in change in the state of the system.

Static or dynamic. Static models represent the system at a specific point in time, whereas dynamic models represent the time varying behaviour of a system. "Dynamic" implies the *time-dependency* of one or more model variables (Bekker, 2012).

Deterministic or stochastic. Deterministic systems do not contain random variables, whereas stochastic systems contain at least one random variable. "Stochastic" implies the *distribution-dependency* of one or more model variables (Bekker, 2012).

The PEM is modelled on a month-to-month basis and is thus classified as discrete. Typically, MC simulation models are static. However, a dynamic element

2.5 Chosen area of optimisation: Primary energy module

is incorporated in the PEM through the inclusion of different monthly values for each parameter. The PEM is stochastic because it contains probabilistic inputs.

MC simulation is a term generally used to describe any simulation involving random sampling and is used to realise the value of one or more random variable(s) (Henderson & Nelson, 2006).

2.5.6 Limitations of the primary energy module

The author came across the following shortcomings while operating the PEM:

Simulation resolution. The PEM is aggregated to a monthly level. Although management may make monthly decisions, aggregation to a monthly level flattens the variation in the process. To more effectively account for the inherent stochastic nature of the real-world system, the PEM should rather be developed on a daily resolution.

Computationally expensive. Running the PEM once takes approximately 15 seconds. To optimise target stockpile levels, the PEM module will need to be evaluated (run) thousands of times, leading to a lengthy optimisation.

PCLF estimates. The planned capability loss factor represents planned maintenance and is estimated as a monthly downtime percentage. However, in the real-world situation, power station generator units are either in operation or not. Therefore, a daily maintenance schedule would provide more accurate results.

Complicated database architecture. This problem is relevant to the entire EFS. The input and output values of the simulator are stored in a complicated hierarchy of relational tables which makes it hard to navigate through and read and write values to and from the database (Eskom Holdings Limited, 2014c).

The monthly aggregation affects the modelling approach which is detailed in Chapter 5. Due to the computationally expensive nature of performing simulation runs, an efficient solution approach is required. The planned capability loss factor is proposed as a future study and is discussed in the conclusion to this study

2.6 Concluding remarks on Chapter 2

(Chapter 8). Due to the complicated nature of the EFS, a function is coded in R to simplify reading from and writing to the database.

2.6 Concluding remarks on Chapter 2

Some modelling techniques applied to the energy sector were investigated and an overview of Eskom's energy supply chain was presented, providing background to the EFS. Thereafter, the EFS was discussed and it was concluded that the EFS could not be optimised in a global manner. The PEM showed the most potential for optimisation and thus it is simulation optimisation of the PEM that forms the basis for the rest of this study. Only coal-fired power stations are considered from this point onwards, because optimising coal stockpile policies is the renewed aim of this study. Coal-fired are sometimes referred to as power stations.

As optimising coal reserves has its roots in inventory management, literature on classic inventory management systems is investigated in the following chapter. Coal stockpile management systems will also be discussed.

Chapter 3

Inventory and coal stockpile management systems

Adding optimisation capability to the primary energy module(PEM) presents itself as a fuel-scheduling problem and coal management problem. The focus of this study is not on the former — determining the optimal fuel-schedule for Eskom — as there is already a division that does so. The focus of this study is on the latter — developing an integrated optimisation model that can optimise coal stockpile policies in the presence of stochastic elements.

After presenting a review of coal stockpile management literature, inventory models are discussed because coal stockpile management has its roots in classic inventory models. Inventory control is an important supporting activity in many businesses. It has been studied extensively and there are countless inventory models. This chapter merely aims to provide a review of literature relevant to this study and is by no means an exhaustive overview.

3.1 Coal stockpile management systems

The two most common uses of coal are to generate electricity or produce fuel. In either case, stockpiles serve two main purposes ([Carpenter, 1999](#)):

1. To enable coal blending to satisfy quality requirements.
2. To provide a buffer between delivery and burn to mitigate against perturbations.

3.1 Coal stockpile management systems

This study is focussed on the latter and reducing costs thereof. Nonetheless, coal blending is addressed in this chapter due to some important concepts that relate to coal stockpiles as buffers.

3.1.1 Coal stockpiles as blending facilities

In the case of coal-fired power stations, coal blending problems are associated with determining the optimal coal mix to be fed into boilers, given certain quality requirements. Generation of electricity is sensitive to the quality of the coal and the variation thereof (Carpenter, 1999). Blending of coal needs to take place to homogenise coal of differing quality. Some coal is pre-blended at collieries, but nonetheless all coal needs to be blended at the stockyards.

Stackers and reclaimers are machinery used for bulk material handling, such as coal blending. Stacking is the process of depositing coal delivered onto stockpile heaps. Reclaiming is the process of retrieving stacked coal and sending it to the combustion chamber. Through methodical stacking and reclaiming strategies, the efficiency of coal blending operations can be improved.

Bituminous coal is used to produce electricity in South Africa. The varying volatility, moisture, and abrasiveness of the bituminous coal makes it difficult to estimate the calorific value (CV) of coal. Thus, estimates for CV are subject to a high degree of variation. Washing of coal is required when the calorific value does not meet quality requirements. De-stoning is also often required.

Carpenter (1999) emphasises the importance of representative coal sampling techniques. In this study, it is assumed that satisfactory coal sampling techniques are used to determine the average CV and the variation thereof. The values used in the PEM — and the energy flow simulator (EFS) in general — were obtained from various departments within Eskom, such as the Primary Energy Division and the Generation Division.

For the supply of coal, Eskom relies predominantly on long-term contracts with mines adjacent to coal-powered power stations. However, of recent some coal has been acquired through short-term agreements. Majuba is the only power station not located in the vicinity of a colliery. The coalfields at Majuba could not be economically mined and therefore coal is sourced elsewhere and transported by trucks instead of trains and conveyors. Due to lack of rail infrastructure, coal is

3.1 Coal stockpile management systems

also transported via truck to three other power stations, which comes at a greater cost (Eskom Holdings Limited, 2011b).

In a South African case study, Conradie (2011) optimised Sasol's coal handling facility's blending process through the application of metaheuristic approaches. Although the blending of coal is an important process, it is not modelled by the PEM. The PEM models coal stockpiles as buffers.

3.1.2 Coal stockpiles as buffers

Coal stockpiles act as buffers to mitigate against perturbations in the system. However, over-sized stockpiles lead to capital being unnecessarily tied up in coal inventory. Furthermore, reducing stockpile levels is important to shorten the turnover period of coal inventory. A shorter inventory turnover period reduces the weathering and atmospheric oxidation of coal, which means coal of higher CV loaded into boilers (Carpenter, 1999). This is especially important in the case of low grade coal, as in the case of this electric utility.

Additional costs associated with maintaining coal on hand are as a result of operating stackers and reclaimers. Coal needs to be continuously blended and the greater the stockpile, the more blending needs to be performed. In addition, coal has to continuously be stacked and reclaimed, and in some cases irrigated, to mitigate against spontaneous combustion.

A combination of different types of stacking and reclaiming technologies can be used in the coal stockyard. The coal blending process begins with stacking. Stacking the coal in a strategic way makes reclaiming for blending easier. Stacking methods commonly applied are chevron, strata, window, and cone shell (Muller, 2010). Typically used reclaimers are the portal scraper reclaimer, boom type bucket wheel reclaimer, bridge scraper reclaimer, bridge type bucket wheel reclaimer, and drum reclaimer — each with their own advantages and disadvantages (Muller, 2010).

Enterprise resource planning systems are commonly applied in inventory management problems. An example of an enterprise resource planning system applied to coal stockpiles is a coal management system developed by Sinha *et al.* (2008). They designed a sophisticated system to deal with the operational, commercial, and administrative burdens encountered when estimating coal stockpile requirements.

3.2 Modelling of inventory management systems

3.2 Modelling of inventory management systems

Inventory management principles, problems, and techniques are investigated in this section.

3.2.1 Inventory management

Mercado (2007) defines the objective of inventory management as “meeting *customer needs* while keeping inventory costs at a reasonable level to *produce a profit* for the firm”. Firstly, the term *customer needs* relates to the total electricity demand of South Africans. Sufficient electricity production should be viewed as a requirement rather than a demand (**D’Sa, 2005**), and thus providing sufficient coal to fuel the coal-fired power stations is essential. Secondly, cost minimisation is chosen instead of *producing a profit*, because Eskom is a vertically integrated monopoly. **Winston (2004)** defines the aim of inventory management as determining the policies to reduce the costs related to meeting customer demand.

In short, the aim of inventory management is to reduce the costs of ordering and handling, whilst having enough safety stock to meet customer demand within a target confidence level. Safety stock acts as a buffer to reduce the effect of perturbations in a system. Costs associated with inventory management are holding costs, ordering costs, and shortage costs. **Zipkin (2000)** specifies four eternal issues associated with inventory management:

- What is the optimal inventory quantity?
- What is the current inventory level?
- How much inventory should be ordered?
- When should inventory be ordered?

Chapter 5 describes the modelling approach used in this study, providing policies to address three of the four eternal issues. Determining the *current inventory level* is not incorporated as it is assumed that the stockpile levels calculated by the PEM are sufficiently accurate.

Sherbrooke (2004) describes two approaches to inventory management: the system approach and the item approach. The system approach is generally

3.2 Modelling of inventory management systems

concerned with meeting a service level whilst reducing costs. Questions are asked such as: how can a company ensure that 95% of orders are fulfilled? How can the process of fulfilling 95% of orders be made more efficient? On the other hand, the traditional item approach is concerned with minimising costs by finding the balance between ordering, stockout and holding inventory. The item approach is more straightforward because decisions are made per item, independent of other items.

The system approach essentially consists of many item approaches combined into one. In this study, only a single item is considered, namely coal, which entails an item approach. However, due to numerous locations of stockpiles (a stockpile for each power station), a system approach is also incorporated.

Systems can be classified according to three dimensions (Sherbrooke, 2004):

Single-item vs. multi-item. The system consists only of coal and thus is a single-item system.

Single-echelon vs. multi-echelon. Multi-echelon systems involve supporting depots for distribution centres. In this case, distribution centres are coal stockpiles for power stations (first echelon). The coal stockpiles at mines could be seen as the supporting depots (second echelon). However, as the mines are not owned by the electric utility, they are not in the scope of this study. Hence, the system is classified as single-echelon.

Single-indenture vs. multi-indenture. By modelling the quality of batches of coal through different probability distributions, coal can be seen as consisting of sub-items and therefore being multi-indenture. However, this is not modelled in the PEM. Only a single probability density function is designated per power station. This is a potential area of investigation for future studies.

The next two subsections will discuss deterministic and stochastic inventory systems.

3.2 Modelling of inventory management systems

3.2.2 Deterministic inventory problems

The most common approach to modelling deterministic inventory problems is the classic economic order quantity (EOQ) which was first formulated in 1915 by F.W Harris (Winston & Venkataramanan, 2002). Variations of EOQ models include quantity discounts and backorders. EOQ models generally have four key assumptions (Winston, 2004):

1. repetitive ordering
2. constant demand
3. constant lead time
4. continuous ordering.

Repetitive ordering implies that orders are made for multiple periods. Constant demand coupled with constant lead time allows for determining the optimal ordering policy pre-emptively. The assumption that continuous review of the quantity of inventory on-hand is possible, is referred to as continuous ordering. Classic EOQ models cannot be applied to this study because, amongst others, the assumptions of constant demand and constant lead time are not applicable to the stochastic nature of the PEM.

3.2.3 Stochastic inventory problems

In this subsection, the news vendor problem and the adaptation of the deterministic EOQ to the EOQ with uncertain demand are discussed.

3.2.3.1 News vendor problem

The news vendor problem is commonly found in literature and is characterised by the following sequence of events (Winston, 2004):

1. An entity orders a certain number of goods, say a .
2. A certain demand b occurs, with probability $p(b)$.
3. A cost is incurred, dependant on a and b .

3.2 Modelling of inventory management systems

Such problems are named news vendor problems because they are analogous to the challenges news vendors face. If a news vendor orders more newspapers than the demand, he will be left with excess papers at the end of the day. On the other hand, if too few newspapers are ordered, the news vendor will lose out on sales.

News vendor problems are classified as single-period, perishable models (Hillier & Lieberman, 2001). Winston (2004) defines single-period decision models as models that involve making a decision only once. Therefore, news vendor problems are not used in this study, because they do not model inventory carried over from the end of one period to the beginning of another period.

3.2.3.2 EOQ model with uncertain demand and lead time

EOQ models are multi-period decision models. As the coal demand per power station is uncertain, the coal stockpile levels need to be continuously reviewed. Two examples of such models are (r, q) and (s, S) continuous review policies.

(r, q) continuous review policies are classified by an order of q units being

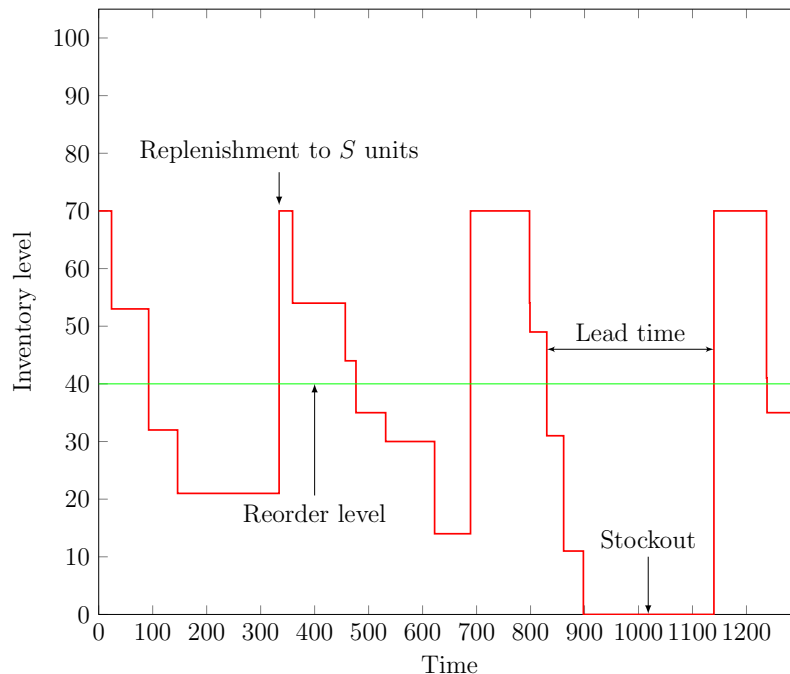


Figure 3.1: Characteristics of the (s, S) inventory process.

3.3 Concluding remarks on Chapter 3

placed when inventory levels fall below re-order point r . The (r, q) policy has been shown to be optimal under the assumption that only single unit orders are received, which allows for orders to be placed as soon as inventory levels drop below r (Winston, 2004). However, when multiple unit orders are received, (r, q) policies are no longer optimal. In that circumstance, (s, S) policies have been shown to be optimal (Winston, 2004). For the (s, S) continuous review policy, an order is placed when the inventory level is equal to or less than re-order point s and the re-order size is the amount required to return the inventory level to S . Figure 3.1 illustrates typical characteristics of a (s, S) continuous review policy.

In continuous review, the inventory level status is automatically updated each time a transaction occurs. Due to system constraints, sometimes only periodic review is possible. As the PEM runs on a monthly interval, the most appropriate review policy is to review the inventory level status on a monthly basis. This rules out the application of traditional (r, q) and (s, S) continuous review policies. Principles from continuous review policies can, however, be adapted and then used to optimise stockpile level policies. As introduced in Chapter 5: the lower warning limit; upper warning limit; and target stockpile level are adapted from the (s, S) continuous review policy to fit the structure of the primary energy module.

3.3 Concluding remarks on Chapter 3

As optimisation of the PEM is of interest, this chapter investigated literature on coal stockpiles as blending facilities and buffers. Coal stockpile management has its roots in inventory problems, and therefore inventory models were discussed, with a focus on stochastic problems. The next chapter reviews literature on simulation optimisation techniques.

Chapter 4

Simulation optimisation

Operations Research, otherwise known as management science and decision science, involves the application of advanced analytical methods to help people and organisations make better decisions. Mathematical representations of real-life situations are formulated to understand systems better (Winston & Venkataramanan, 2002). Countless optimisation methods form part of the operations research practitioner's toolbox. Common approaches include: linear programming, integer programming, mixed integer programming, dynamic programming, game theory, Markov chains, queueing theory, non-linear programming, and metaheuristics.

Of interest to this study are optimisation techniques that can be applied to stochastic systems, specifically simulation optimisation. Many traditional deterministic optimisation methods have been adapted to the stochastic problem setting, for instance stochastic dynamic programming, otherwise known as Markov decision theory (see Puterman (2009)).

The terms *simulation optimisation* (SO) and *stochastic optimisation* are closely related because they both involve the optimisation of stochastic systems. However, SO involves optimising the performance measures estimated as the outputs of simulation models, whereas stochastic optimisation is a general term used to describe the optimisation of a stochastic system.

In this chapter, an introduction to SO is provided, followed by the detailing of important considerations in SO. Thereafter, the bulk of the chapter investigates SO techniques in search of an applicable SO technique for this study.

4.1 An introduction to simulation optimisation

4.1 An introduction to simulation optimisation

Rosen *et al.* (2008) define the general SO problem setting, for a single-objective, as

$$\text{Minimise } f(\theta) \quad (4.1)$$

$$\text{subject to } \theta \in \Theta, \quad (4.2)$$

where the expected system performance $f(\theta) = \mathbb{E}[\psi(\theta, \omega)]$ is estimated by $\hat{f}(\theta)$ from numerous samples of a simulation model with instances of feasible input parameters θ . Input parameters can be discrete or continuous, and are subject to constraint set(s) $\Theta \subset \mathbb{R}^d$, while ω represents the stochastic effects of the system.

Since the objective function $f(\theta)$ cannot be explicitly defined in closed form, computer simulation is used to estimate system performance. Simulation models are often of a complex nature and thus $f(\theta)$ is treated as a black box, with inputs and outputs (Azadivar, 1999; Bekker, 2012).

Figure 4.1 illustrates the combination of a simulation model and an optimisation algorithm. Based on the response value(s), the optimisation algorithm adjusts the values for the decision variables, which are then in turn evaluated by the simulation model. The optimisation algorithm guides the process towards an optimal combination of decision variables. The process is repeated until no more improvement is shown, which means that an optimal solution may have been

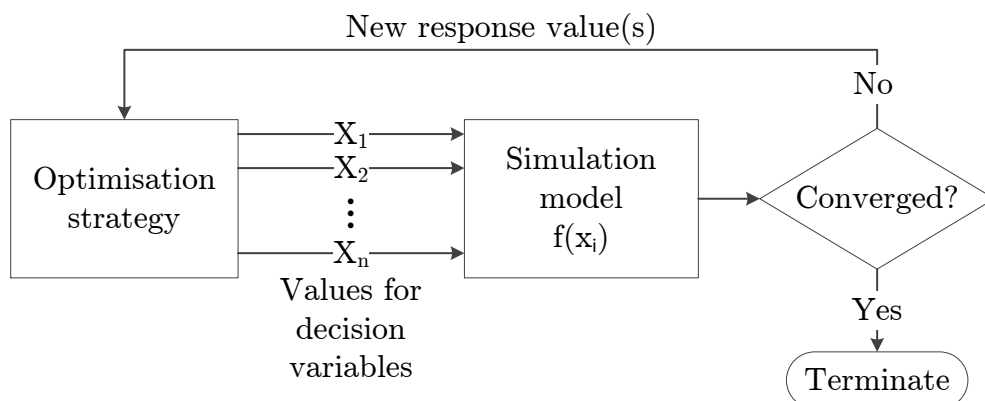


Figure 4.1: Integration of simulation and optimisation algorithms (Bekker, 2013).

4.2 Considerations in simulation optimisation

reached. For the purpose of this study, two definitions are put forth:

A simulation run refers to 1 000 independent replications of the primary energy module (PEM)'s Monte Carlo simulation model.

Simulation optimisation is the iterative process whereby the PEM is nested within an optimisation model and run until convergence or some other stopping criteria is reached.

Fu (2001) provides examples of typical applications of SO:

- Integrating an optimisation technique with a discrete-event simulation model of a manufacturing system to determine the optimal system design.
- Optimising an existing simulation model of a manufacturer's supply chain.
- Given a complicated call-centre model, minimising staffing costs and maximising customer service levels.
- Maximising the expected return of a portfolio of financial instruments.
- Minimising the mean waiting time of single-server queueing models when increased server speed comes at a greater cost.
- Determining the optimal policies in an (s, S) inventory control system.

4.2 Considerations in simulation optimisation

Optimisation for deterministic problems can be challenging, if the structure of the performance function is for the most part unknown and the number of decision variables is large. SO further complicates the optimisation process because the performance of a function cannot be determined exactly. As the performance is estimated, optimisation algorithms sometimes experience difficulty in conclusively determining which candidate solution outperforms the other. This problem can be addressed by running a simulation model many times, but due to the computationally expensive nature of simulation this is usually infeasible (**Banks et al., 1998**).

Three key challenges exist in SO (**Azadivar, 1999; Fu, 2002**):

4.2 Considerations in simulation optimisation

1. There is no closed form expression for the objective function. This excludes many of the most commonly used optimisation methods, such as linear programming, integer programming, and non-linear programming.
2. The objective function is a stochastic function of the decision variables, which makes determining point estimates challenging.
3. Running computer simulations to evaluate performance is much more expensive than analytical functions. Therefore, it is important to use efficient optimisation algorithms that require a relatively small number of evaluations to reach a “good enough” solution.

Banks et al. (1998) classifies SO approaches according to techniques that:

- i. Guarantee asymptotic convergence to the optimal solution.
- ii. Guarantee optimality for the problem’s deterministic counterpart.
- iii. Guarantee a probability of correct selection.
- iv. Are robust.

The four approaches are specifically ordered. There is a trade-off between the confidence in optimality and the practicality of solutions. The higher up on the list, the greater the confidence in optimality. On the contrary, the lower the position, the more workable solutions are. A robust solution is favoured because the SO in this study is a highly repetitive problem.

Eiben & Smith (2003) discuss design problems and repetitive problems. A trade-off exists between algorithm accuracy and speed. Design problems require a once-off solution with a high level of accuracy. In such a case, algorithm speed is less important. However with repetitive problems execution speed is much more important, and sometimes comes at the expense of accuracy. The PEM will be re-run frequently by electric utility planners to plan for the upcoming planning horizon.

Moving onto the next section, an approach is not to be confused with a technique. A technique can be thought of as the tool used and the approach as the general mindset when tackling a problem.

4.3 Simulation optimisation techniques

4.3 Simulation optimisation techniques

Carson & Maria (1997) group SO techniques into six categories:

1. gradient based search methods
2. stochastic optimisation
3. response surface methodology
4. heuristic methods
5. asynchronous teams (hybrids)
6. statistical methods.

A drawback of this categorisation is that it is not based on the problem's underlying characteristics. The answer to “Is the problem highly dimensional, non-differentiable or discontinuous?” determines whether a local optimisation or global optimisation technique is applicable. Additionally, in the case of local optimisation, the answer to “Are the parameters discrete or continuous?” prescribes the types of local optimisation techniques that are applicable (Tekin & Sabuncuoglu, 2004).

The classification by Tekin & Sabuncuoglu (2004) is favoured because it is based on the underlying characteristics of the problem. Hachicha *et al.* (2010) improved upon the work of Tekin & Sabuncuoglu (2004) and developed a hierarchical diagram. This is illustrated in Figure 4.2 and forms the structure of the search process for a SO technique applicable to this study.

This section is merely an investigation into commonly used SO techniques. The interested reader is referred to the following sources for additional material on SO. Carson & Maria (1997) provide a critical review of SO techniques. Issues relating specifically to SO are detailed in Azadivar (1999). Andradóttir (1998a) review SO methods, focussing on gradient-based and random search methods. An entire chapter, written by Andradóttir (1998b), is dedicated to SO in the *Handbook of Simulation* by Banks *et al.* (1998).

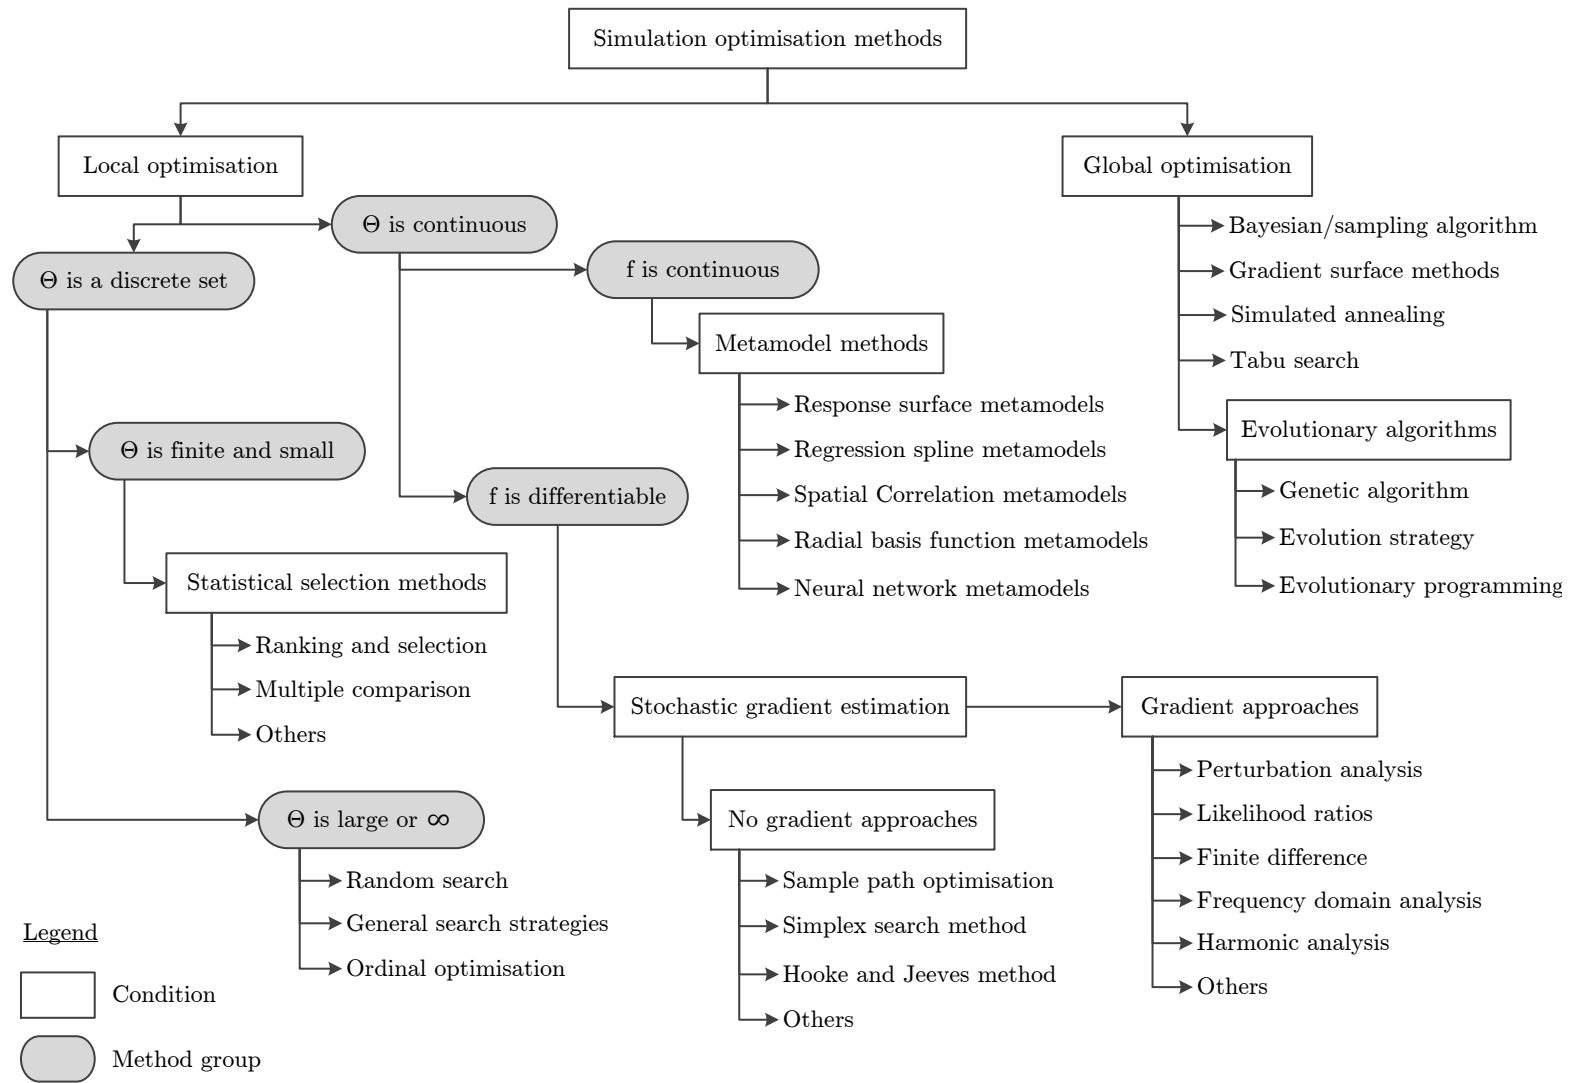


Figure 4.2: Classification of simulation optimisation techniques (Hachicha *et al.*, 2010).

4.3 Simulation optimisation techniques

4.3.1 Local optimisation

Local optimisation techniques are arranged according to discrete and continuous decision spaces, as illustrated in Figure 4.2.

4.3.1.1 Discrete parameter set

As optimisation complexity increases when the number of input parameters increases, discrete parameter techniques are differentiated upon based on the number of input parameters (Tekin & Sabuncuoglu, 2004).

First, when the number of input parameters is small, *statistical selection methods* (SSM) are best-suited. Examples of SSM include ranking and selection, importance sampling, and multiple comparison. *Ranking and selection* analyses a fixed number of alternatives and then ranks them. If the decision space is too large, this brute force approach becomes infeasible (Fu *et al.*, 2005). *Importance sampling* has shown success in speeding up simulation involving rare-events, by sampling from distributions that favour rare-events (Rubinstein, 1999). A limitation is the challenge in choosing an appropriate change of measure to increase the probability of rare-events being sampled (Carson & Maria, 1997). The last discussed SSM is *multiple comparison methods* which use certain pairwise comparisons to quantify the differences between systems' performance (Swisher *et al.*, 2003).

Secondly, when the number of input parameters is large, the decision space may be too large for SSM to perform adequately. To that end, *ordinal optimisation* aims to obtain “good enough” solutions by using a subset method whereby a large set of solutions is sampled from, resulting in a smaller decision space which is subsequently solved (Tekin & Sabuncuoglu, 2004). Ho (1999) describes ordinal optimisation as soft computing for hard problems. *Random search methods* move in a “greedy” manner from the current best solution to the next best solution in their neighbourhood. The disadvantage of random search is that if the neighbourhood structure is poorly defined or the initial solution is bad, values will converge to a poor local optimal (Fu *et al.*, 2005).

4.3.1.2 Continuous parameter set

A great amount of research has been carried out on problems that have a continuous decision space (Azadivar, 1999; Tekin & Sabuncuoglu, 2004). This section is split

4.3 Simulation optimisation techniques

into three parts: *Gradient-based* (GBTs), *stochastic approximation*, and *meta-model* techniques.

First, GBTs adapt their standard deterministic methods to account for randomness, usually through the use of stochastic approximation algorithms (Carson & Maria, 1997). They have been shown to converge towards a local minimum (Andradóttir, 1998a). Benefits include the fact that relatively few simulation runs are required and that the simulation model can be treated as a black-box. Examples of a few gradient-based procedures are finite difference estimates, infinitesimal perturbation analysis, likelihood ratios, and frequency domain analysis.

Finite difference estimation is the crudest GBT. It requires at least one model run more than the number of decision variables. As response estimates are noisy, finite difference estimation is particularly susceptible to being misled into searching in the wrong “direction” (Andradóttir, 1998b).

Infinitesimal perturbation analysis has the ability to estimate all the gradients of the objective function in only one simulation run. However, this is subject to certain conditions. Most importantly, thorough knowledge of the simulation model’s internal workings is required (Azadivar, 1999). The main principle of this technique is to perturb decision variable values by an infinitesimal amount, and track the effects thereof as a simulation run progresses over time. A pitfall is that tracking capabilities are often challenging to implement (Andradóttir, 1998a).

Likelihood ratios, otherwise known as score functions, determine the gradient of the expected value of outputs for given inputs (Carson & Maria, 1997). As with any SO technique: the more simulation runs, the more accurate the estimates.

Frequency domain analysis was born out of the concept of *Fourier analysis*. Inputs are oscillated sinusoidally at irregular frequencies during a lengthy simulation run (Carson & Maria, 1997). Fourier analysis is then used to measure the sensitivity of input parameters. It has the same drawback as infinitesimal perturbation analysis, because of the difficulty posed in incorporating the technique into an independently built simulation model (Azadivar, 1999).

Secondly, continuous parameter local optimisation techniques that are not focussed on gradient estimates are investigated. They are typically referred to as stochastic approximation techniques (Andradóttir, 1998a). Stochastic optimisation techniques include sample path optimisation, and the simplex search method.

4.3 Simulation optimisation techniques

Sample path optimisation saves estimates from many simulation runs, building an archive so as to transform the problem from a stochastic to a deterministic setting. The strong law of large numbers allows for this transformation, granted that there are a sufficient amount of simulated estimates (Fu *et al.*, 2005). The major advantage of this method is that standard deterministic optimisation methods can be used to find a near optimal solution.

Nelder and Mead's *simplex search* technique assumes a convex decision space. The heuristic begins with points in a simplex. At each iteration, the worst point is excluded and a new point replaces it. The new point is chosen to be the reflection of the worst point about the centroid. Application of this technique is problematic due to handling feasibility constraints (Carson & Maria, 1997).

Thirdly, meta-models draw a relationship between a response of interest and its corresponding input variables (Hachicha *et al.*, 2010). The most popular technique is *response surface methodology*. Regression series, artificial neural networks, or hybrid models are fitted to the output variable. The simulation model is re-run to optimise the chosen method. Response surface methodology is advantageous due to its straightforward approach, but it lacks diversity in solutions (Fu *et al.*, 2005).

4.3.2 Global optimisation

Global optimisation techniques are preferred to local optimisation techniques if the decision space is highly dimensional, discontinuous or non-differentiable (Tekin & Sabuncuoglu, 2004). Most of the techniques used in global optimisation are *metaheuristics*. Metaheuristics are efficient high-level approaches that obtain near-optimal solutions in a relatively short time period, without limiting the search space (Boussaïd *et al.*, 2013; Gendreau & Potvin, 2005).

Commonly applied metaheuristics that have shown success in SO, include (Fu *et al.*, 2005; Gendreau & Potvin, 2005):

- simulated annealing
- tabu search
- scatter search

4.4 Concluding remarks on Chapter 4

- evolutionary algorithms:
 - genetic algorithm
 - evolutionary strategies
 - evolutionary programming.

[Gendreau & Potvin \(2005\)](#) describe metaheuristics as “solution methods that orchestrate an interaction between local improvement procedures and higher level strategies to create a process capable of escaping from local optima and performing a robust search of a solution space.” Metaheuristics are based on the principle of balancing exploration and exploitation. Exploration maintains diversity and therefore reduces the chance of premature convergence to a local optima, whereas exploitation focuses on refinement of current solutions.

4.3.3 The chosen simulation optimisation technique

Global optimisation is favoured over local optimisation due to the underlying characteristics of the problem in this study. As detailed in Chapter 5, there are three decision variables per power station. At the time of writing, 14 coal-fired power stations are defined as commissioned, resulting in 42 decision variables.

Metaheuristic techniques are chosen because they lend themselves well to this highly dimensional SO problem. More specifically, [Fu *et al.* \(2005\)](#) proposed using a less commonly applied metaheuristic, namely the cross-entropy method (CEM). The CEM fits a probability distribution on the space of solutions ([Rubinstein & Kroese, 2004](#)), thus making it more versatile than other metaheuristics. [Fu *et al.* \(2005\)](#) state that the CEM shows great promise in the field of SO, because it is not dependant explicitly on the current set of simulated values. Such versatility is beneficial in a stochastic environment where much simulation noise exists. Therefore, the CEM is the proposed algorithm for the optimiser of the PEM.

4.4 Concluding remarks on Chapter 4

The literature presented in this chapter is by no means an exhaustive overview. The goal of the chapter was to investigate simulation optimisation techniques,

4.4 Concluding remarks on Chapter 4

and find a technique that is applicable to use an optimisation strategy for the primary energy module. The CEM was chosen.

The following chapter is dedicated to presenting the optimisation model and explaining the reasoning behind the model formulation. The CEM is also described in more detail.

Chapter 5

Model formulation and solution approach

This chapter divides the modelling approach into three sections: conceptual model formulation, mathematical model formulation, and the application of the cross-entropy method (CEM) as the solution approach. Objectives, decision variables and constraints are discussed. This chapter also presents a review of literature on the CEM, focussed by its application to continuous stochastic parameter optimisation.

5.1 Conceptual model formulation

In this section, the reasoning behind the modelling approach is explained. The decision variables are designated and then the objective function is formulated.

5.1.1 Decision variables

In simulation, inputs are divided into two groups: controllable inputs and inputs that represent uncontrollable factors. Controllable inputs influence the performance of the system and are known as decision variables (Winston, 2004). Uncontrollable factors are usually defined by a probability distribution and generated randomly. Uncontrollable factors add noise to the system. “Controllable” refers to whether actions taken by management result in changes to inputs (Law *et al.*, 2000).

5.1 Conceptual model formulation

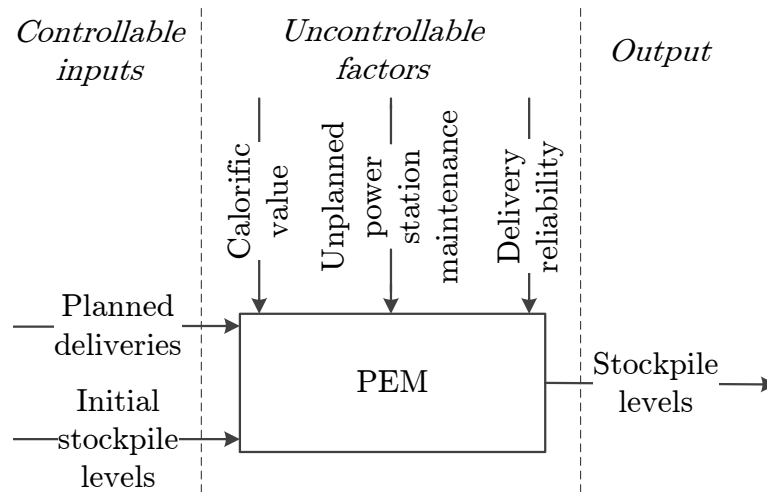


Figure 5.1: Controllable inputs, uncontrollable inputs and outputs of the primary energy module.

The inputs and outputs of interest for the primary energy module (PEM) are shown in Figure 5.1. First, the *planned deliveries* and *initial stockpile levels* are controllable inputs. The initial stockpile level is the stockpile level at the beginning of the planning horizon, for each power station. The initial stockpile level can be thought of as the target or desired stockpile level.

Secondly, the *calorific value*, *unplanned power station maintenance*, and *delivery reliability* are uncontrollable inputs. The three uncontrollable inputs illustrated are of interest, because they are defined by probability functions, and are thus responsible for the stochastic nature of the PEM. For each of the uncontrollable inputs, different probability density function (pdfs) are defined in the PEM per month per power station.

Thirdly, the combined controllable inputs and uncontrollable factors shown in Figure 5.1 — along with other inputs in the PEM — result in the *stockpile levels* at the end of each month, for each power station.

The controllable inputs, namely deliveries and initial stockpile levels, are potential decision variables for the optimiser. However, due to the fact that all the values in the PEM are aggregated to a monthly basis, it does not make sense to model the system as a classic inventory model. Neither deliveries nor the review of stockpile levels can occur more frequently, such as daily. Furthermore, the

5.1 Conceptual model formulation

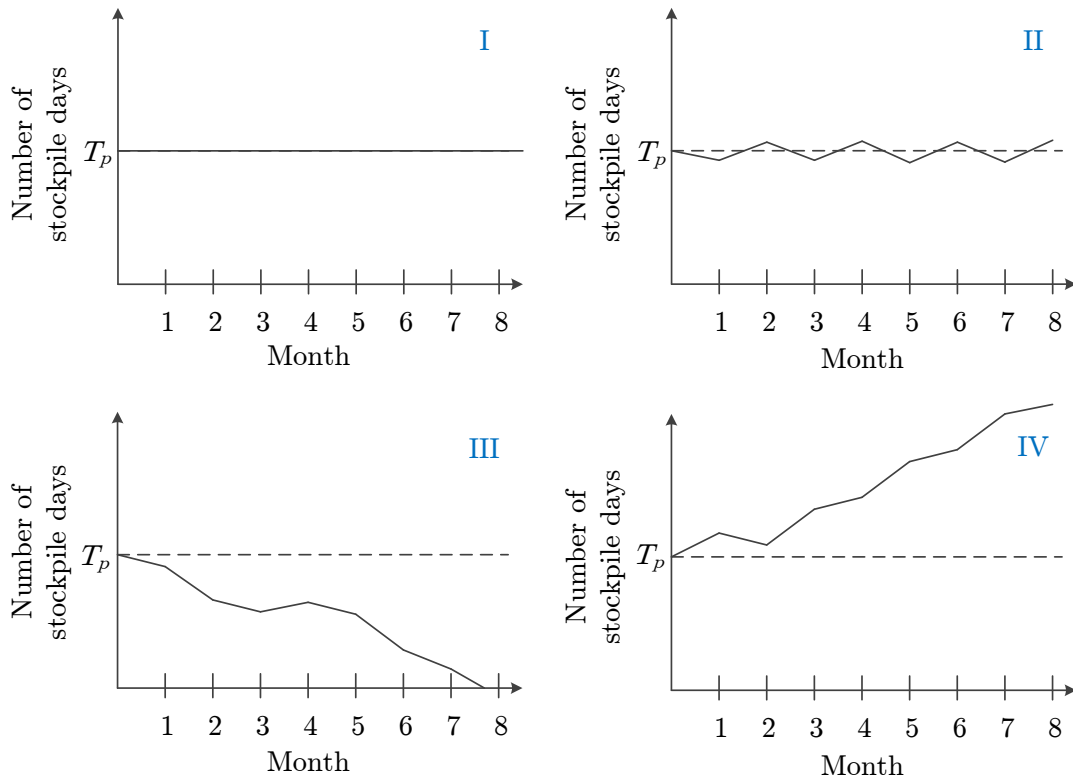


Figure 5.2: Four hypothetical scenarios for variation between coal delivery and burn.

uncertainty in the process cannot be forecast. Hence, a different approach needs to be followed.

Simply put, the approach followed is to determine optimal coal stockpile management policies to cater for the variation between deliveries and burn. Figure 5.2 depicts four hypothetical cases of the variation between deliveries and burn for a power station, from the target stockpile level (T_p), on a month-to-month basis. Stockpile days is used as a common measure to compare the stockpile levels at different power stations. As expected, on average one stockpile day will provide a specific power station with enough fuel for one day.

Case I represents a deterministic process, whereas Case II represents a stochastic process with variation around the mean. Both Case I and Case II do not represent reality — Case I is almost impossible and Case II is improbable and fortunate. In Case III, management would need need to respond to the decreasing

5.1 Conceptual model formulation

stockpile level by placing additional orders for coal. Additional orders are referred to as emergency orders. In Case IV, management should try reduce increasing stockpile levels to cut back on carrying costs by cancelling baseline deliveries.

The transfer of coal between power stations, as opposed to the cancellation of coal deliveries, was initially considered. However it was concluded that coal is too expensive to be transferred from one coal-fired power station to another, and thus transferral of coal is not modelled.

To cater for the variation between deliveries and burn, the proposed decision variables are the upper warning limit (U_p), the lower warning limit (L_p), and T_p — per coal-fired power station. The reasoning behind the proposed decision variables is illustrated in Figure 5.3. However, before discussing the proposed decision variables, the trivial method used to equate deliveries and burn is explained.

Coal delivered is equated to the average coal burnt, for each power station, by running the PEM simulator many times (say 500) and calculating the average coal burnt. The more simulation runs, the smaller the statistical estimation error. The baseline deliveries ($D_{p,t}$) are then equated to the average coal burnt. Although coal delivery is equated to the average coal burnt, there is variation in both coal delivered and burnt. This variation is what the optimisation model aims to cater for.

When running the PEM to estimate the mass of coal burnt at each power station, it is important to set arbitrarily high initial (target) stockpile levels. This is done to ensure that there are no coal shortages at any of the power stations. Coal shortages at one power station will result in additional load uptake at other power stations to meet electricity demand. In turn, this will result in additional burn at certain power stations, leading to misrepresented coal burn estimates. By the same token, coal stockpile policies at each power station cannot be optimised in an isolated manner.

Resuming the discussion of decision variables, hypothetical situations for the warning limits and target stockpile level are illustrated in Figure 5.3, for a single power station. T_p also represents the initial stockpile level. The two warning limits — L_p and U_p — serve as alarm sirens, each representing a policy that defines an action:

5.1 Conceptual model formulation

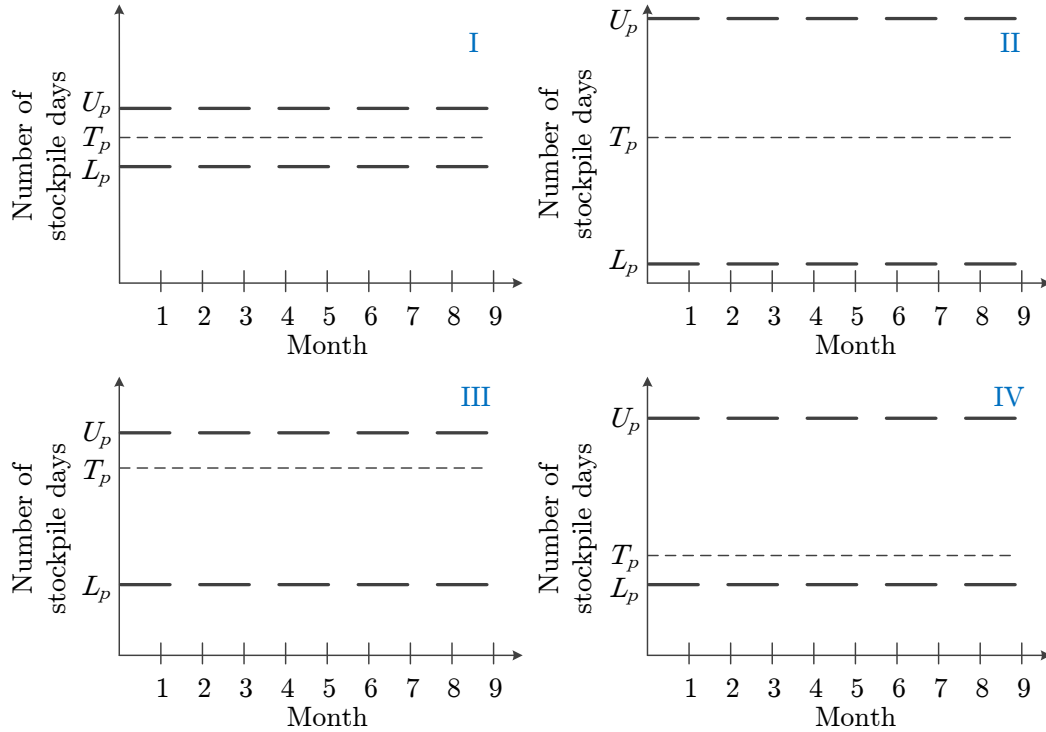


Figure 5.3: Four hypothetical scenarios for the upper and lower warning limits.

1. L_p . If stockpile levels are too low, order emergency coal to replenish coal stockpile levels to T_p .
2. U_p . If stockpiles are too high, cancel coal deliveries to reduce stockpile levels to T_p .

If U_p is too high, not enough cancellations will occur and the stockpile levels have the potential to become excessively large. Likewise, if L_p is too low, the stockpiles may stoop drastically low and be at risk of a stockpile shortage. On the other hand, if either of the warning limits are fitted too tightly around T_p , the optimiser will penalise even the smallest variation.

The reader is reminded that the inspiration from this approach is drawn from the (s, S) continuous review policy discussed in Section 3.2.3.2. The situation when L_p comes into play is depicted in Figure 5.4. In this example, L_p can be thought of as the re-order point, and T_p as the quantity to be replenished to when

5.1 Conceptual model formulation




Figure 5.4: The workings of the optimiser in the case of the lower warning limit.

the emergency delivery arrives.

As illustrated in Figure 5.4, at Point “a” the stockpile drops below L_p , which results in an emergency order of coal. The quantity of the emergency order is equal to the difference in the target level and current level. However, the order is subject to a random lead time (in this case three months). Between the time of order (Point “a”) and delivery (Point “b”), variation continues to take place in the system. When the order arrives (Point “b”), the stockpile level is closer to the target level, but is unlikely to be equal to the target level.

A general approach was chosen for the lead time of emergency deliveries and delivery cancellations. The lead time for emergency deliveries ($l^{(e)}$) is uniformly distributed on the range (1,3) months. Lead time for delivery cancellations ($l^{(c)}$) is also uniformly distributed on the range (1,3) months.

Figure 5.5 illustrates the combined simulation and optimisation model. The values of decision variables L_p , T_p , and U_p are adjusted. T_p is directly input into the model as the initial stockpile level. The warning limits, L_p and U_p , alter the deliveries that are fed into the PEM.

Based on the response values which are the actual stockpile levels ($A_{p,t}$), $D_{p,t}$

5.1 Conceptual model formulation

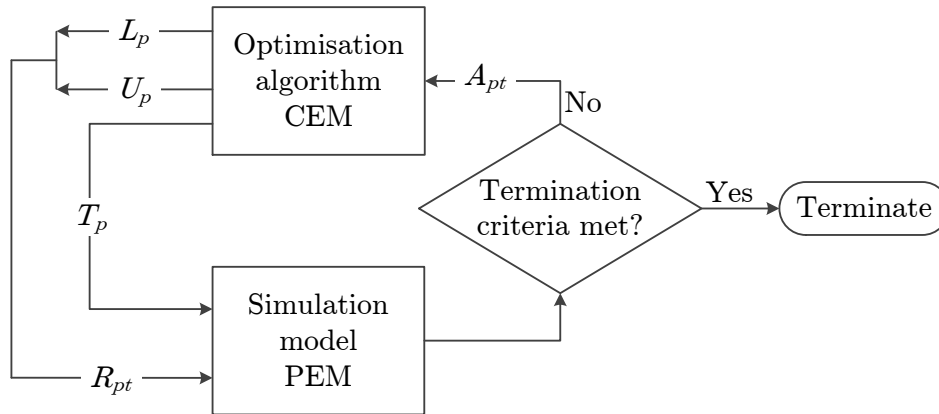


Figure 5.5: The integrated simulation model and optimisation algorithm.

is altered through L_p and U_p . In other words, emergency deliveries are incurred through L_p , and delivery cancellations are incurred through U_p . Emergency deliveries ($e_{p,t}$) and delivery cancellations ($c_{p,t}$) give rise to the resultant delivery ($R_{p,t}$). The process is repeated until convergence or some other termination criteria occurs.

It is important to note that when the PEM is run by itself, all the months defined in the simulation range are simulated at once. However, when the PEM is run in unison with the CEM, simulation is performed on a month-to-month basis so as to enable actions — based on policies — to take place at the end of each month. In other words, emergency deliveries ($e_{p,t}$) and delivery cancellations ($c_{p,t}$) occur based on the lower and upper warning limits, respectively, in relation to the actual stockpile level at month-end.

5.1.2 Components of the objective function

A classic problem in electric utility planning is to match supply and demand at minimum cost (Wood *et al.*, 2013). In saying that, the proposed objectives of the optimiser is to minimise coal on hand, coal shortages, emergency deliveries, and delivery cancellations.

The objectives could be modelled in a multi-objective manner. However, as all the objectives can be expressed in monetary value, a combined single-objective function is defined. The costs are estimated in Section 6.1 and are provided as

5.2 Mathematical problem formulation

ZAR per kton. Based on the response value (actual stockpile levels $A_{p,t}$), measured in ktons, the performance measure is estimated through the objective function.

5.2 Mathematical problem formulation

In this section, some points regarding the PEM are first discussed, along with providing notation for the mathematical model. The bulk of the section is dedicated to documenting the mathematical model.

5.2.1 Problem notation

Based on the straightforward notation structure introduced by [Schlünz & Van Vuuren \(2012\)](#), the notation for the mathematical problem formulation is shown in [Table 5.1](#).

The PEM is treated as a black-box simulation model, and thus only inputs and outputs relevant to the optimiser are of interest. In other words, constraints in the PEM that do not directly influence the optimiser are ignored. However, there are two exceptions: the resultant delivery and its effect on stockpile level.

As deliveries are adjusted by actions taken due to the stockpile warning limit policies (refer back to the description of [Figure 5.5](#) on [page 56](#)), the following two equations are defined for the optimisation model. That is, emergency deliveries ($e_{p,t}$) increase baseline deliveries $D_{p,t}$, and delivery cancellations ($c_{p,t}$) reduce baseline deliveries ($D_{p,t}$). The resultant delivery ($R_{p,t}$) is defined as

$$R_{p,t} = D_{p,t} + e_{p,t} - c_{p,t} \quad \forall p \in \mathcal{P}, \forall t \in \mathcal{T}. \quad (5.1)$$

Therefore, actual stockpile levels ($A_{p,t}$) are calculated with resultant ($R_{p,t}$) instead of baseline ($D_{p,t}$) deliveries, as defined by

$$A_{p,t} = A_{p,t-1} + R_{p,t} - B_{p,t} \quad \forall p \in \mathcal{P}, \forall t \in \mathcal{T}. \quad (5.2)$$

5.2 Mathematical problem formulation

Table 5.1: Notation used in the mathematical problem formulation.

Indices	
p	Index for coal-fired power stations
t	Index for monthly time periods
Sets	
\mathcal{P}	Set of indices for coal-fired power stations
\mathcal{T}	Set of indices for monthly time periods
Parameters	
n	Number of coal-fired power stations in the PEM, therefore $p \in \mathcal{P} = \{1, \dots, n\}$
m	Number of months in the planning horizon of the PEM, therefore $t \in \mathcal{T} = \{1, \dots, m\}$
$l^{(e)}$	Lead time of emergency deliveries
$l^{(c)}$	Lead time of delivery cancellations
$D_{p,t}$	Baseline coal delivery during month t at power station p
$R_{p,t}$	Resultant coal delivery, after emergency deliveries and delivery cancellations, during month t at power station p
$B_{p,t}$	Coal burnt during month t at power station p
$A_{p,t}$	Actual stockpile level at the end of month t at power station p
$S_p^{(\min)}$	Minimum allowed stockpile level of power station p
$S_p^{(\max)}$	Maximum allowed stockpile level of power station p
$D_{p,t}^{(\min)}$	Minimum allowed delivery to power station p during month t
$D_{p,t}^{(\max)}$	Maximum allowed delivery to power station p during month t
HC	Holding cost
SC	Shortage cost
$EC_{p,t}$	Emergency delivery cost for month t at power station p
CC	Delivery cancellation cost
Variables	
L_p	Decision variable for the lower warning limit of power station p
T_p	Decision variable for the target stockpile level of power station p
U_p	Decision variable for the upper warning limit of power station p
$w_{p,t}$	A binary auxiliary variable of value 1, if shortages in coal reserves occur during month t at power station p , or zero otherwise
$x_{p,t}$	A binary auxiliary variable of value 1, if emergency deliveries occur during month t at power station p , or zero otherwise
$y_{p,t}$	A binary auxiliary variable of value 1, if delivery cancellations occur during month t at power station p , or zero otherwise
$s_{p,t}$	Shortages in coal reserves during month t at power station p
$e_{p,t}$	Emergency deliveries during month t at power station p
$c_{p,t}$	Cancellations of deliveries during month t at power station p

5.2 Mathematical problem formulation

5.2.2 Model

The optimisation model is represented analytically in this section. This is done to show the logic incorporated in the optimisation model. Each of the four objective function components are based on the stochastic output of the simulator, namely $A_{p,t}$. It is important to note that the problem cannot be solved analytically, even though it is represented analytically.

The objective is:

$$\text{Minimise } \sum_{p \in \mathcal{P}} \sum_{t \in \mathcal{T}} (\text{HC} \cdot A_{p,t} + \text{SC} \cdot s_{p,t} + \text{EC}_{p,t} \cdot e_{p,t} + \text{CC} \cdot c_{p,t}) \quad (5.3)$$

subject to

$$w_{p,t} = \begin{cases} 1 & \text{if } A_{p,t-1} + R_{p,t} - B_{p,t} < S_p^{(\min)} \\ 0 & \text{Otherwise} \end{cases} \quad \forall p \in \mathcal{P}, \forall t \in \mathcal{T}, \quad (5.4)$$

$$x_{p,t} = \begin{cases} 1 & \text{if } A_{p,t} + \sum_{t \leq m} e_{p,t} < L_p \\ 0 & \text{Otherwise} \end{cases} \quad \forall p \in \mathcal{P}, \forall t \in \mathcal{T}, \quad (5.5)$$

$$y_{p,t} = \begin{cases} 1 & \text{if } A_{p,t} - \sum_{t \leq m} c_{p,t} > U_p \\ 0 & \text{Otherwise} \end{cases} \quad \forall p \in \mathcal{P}, \forall t \in \mathcal{T}, \quad (5.6)$$

$$0 < L_p < T_p < U_p < S_p^{(\max)} \quad \forall p \in \mathcal{P}, \quad (5.7)$$

$$D_{p,t}^{(\min)} \leq R_{p,t} \leq D_{p,t}^{(\max)} \quad \forall p \in \mathcal{P}, \forall t \in \mathcal{T}, \quad (5.8)$$

$$S_p^{(\min)} < A_{p,t} < S_p^{(\max)} \quad \forall p \in \mathcal{P}, \forall t \in \mathcal{T}, \quad (5.9)$$

$$s_{p,t} = w_{p,t} \cdot (A_{p,t-1} + R_{p,t} - B_{p,t} - S_p^{(\min)}) \quad \forall p \in \mathcal{P}, \forall t \in \mathcal{T}, \quad (5.10)$$

$$e_{p,t+l(e)} = x_{p,t} \cdot (T_{p,t} - A_{p,t} - \sum_{t \leq m} e_{p,t}) \quad \forall p \in \mathcal{P}, \forall t \in \mathcal{T}, \quad (5.11)$$

$$c_{p,t+l(c)} = y_{p,t} \cdot (A_{p,t} - T_{p,t} - \sum_{t \leq m} c_{p,t}) \quad \forall p \in \mathcal{P}, \forall t \in \mathcal{T}, \quad (5.12)$$

$$s_{p,t}, e_{p,t}, c_{p,t} \geq 0 \quad \forall p \in \mathcal{P}, \forall t \in \mathcal{T}, \quad (5.13)$$

$$w_{p,t}, x_{p,t}, y_{p,t} \in \{0, 1\} \quad \forall p \in \mathcal{P}, \forall t \in \mathcal{T}. \quad (5.14)$$

The mathematical problem formulation is defined by (5.3) to (5.14). The objective function (5.3) is the sum of the holding, shortage, emergency delivery,

5.3 Solution approach

and delivery cancellation costs incurred. The objective function is subject to constraint sets (5.4) to (5.14).

The three auxiliary variables are assigned binary values in (5.4), (5.5) and (5.6). Whether shortages occurred or not is determined by (5.4). (5.5) and (5.6) determine if actual stockpile levels are not within the warning limits. (5.5) takes into account future deliveries and (5.6) takes into account future cancellations. This is done to prevent unnecessary duplication of emergency orders or delivery cancellations.

The allowable ranges for the warning limits and target stockpile level are ensured by (5.7). The resultant delivery may not exceed the maximum and minimum bounds defined in (5.8). The actual stockpile levels are bounded by (5.9).

(5.10), (5.11), and (5.12) make use of the auxiliary binary variables in (5.4), (5.5), and (5.6), respectively, to ensure non-negativity. (5.10) determines the quantity of shortages. The emergency delivery constraint is specified in (5.11), taking into account future emergency deliveries. Likewise, the delivery cancellation constraint is specified in (5.12). Emergency deliveries and delivery cancellations are subject to random lead times $l^{(e)}$ and $l^{(c)}$, respectively.

Finally, (5.13) and (5.14) specify the ranges of the variables.

5.3 Solution approach

First, this section reviews literature and provides a summary of the CEM. Subsequently, the application of the CEM to this study is detailed.

5.3.1 Literature: Cross-entropy method

The CEM is a relatively new, significant development providing a simple, efficient and general method for solving complicated optimisation problems; specifically combinatorial, stochastic, or continuous multi-extremal problems (De Boer *et al.*, 2005; Rubinstein & Kroese, 2004). This section touches upon the most important concepts of the CEM, relevant to this study. For more information, the reader is referred to the book by Rubinstein & Kroese (2004) documenting the CEM. De Boer *et al.* (2005) provide a tutorial on the CEM. Another useful source is Rubinstein & Kroese (2011).

5.3 Solution approach

Attributed to Reuven Rubinstein, the CEM was originally developed as an efficient algorithm to determine the optimal sampling measure for simulation in rare event networks in order to minimise variance (Rubinstein, 1997). It was soon realised that the CEM could be used for optimisation by setting optimal rare-events as the events of interest (Rubinstein, 1999). Chan & Kroese (2012) state that the CEM is an adaptive and versatile Monte Carlo (MC) method. The CEM has been predominantly applied to three domains:

1. rare-event simulation
2. optimisation
3. machine learning.

The domain of interest for this study is optimisation. In the field of optimisation, the CEM is a metaheuristic. Suppose one has a stochastic performance function $f(\mathbf{x})$ and would like to find the optimal values associated with $f^*(\mathbf{x})$. The expected value of $f(\mathbf{x})$ is defined as

$$\mathbb{E}[f(\mathbf{X})] = \int f(\mathbf{x})g(\mathbf{x})d\mathbf{x}, \quad (5.15)$$

where $f(\mathbf{X})$ is the sample performance and the probability density of \mathbf{X} is $g(\mathbf{x})$. To estimate $f^*(\mathbf{x})$, the CEM assigns a pdf to each decision variable.

In short — samples are generated from the pdfs; the samples are evaluated through the objective function; and the parameters of the pdfs are updated based on the best sample values of the current generation. This is continued iteratively. As the CEM progresses in the search of the optimal combination of decision variable values, the pdfs become more refined and typically converge towards an optimal or near-optimal solution.

This progression is represented in Figure 5.6 for three decision variables. By way of example, normally distributed pdfs were used. The CEM allows for the use of any uni-modal pdf. At “Time i”, the standard deviation is initialised to be extremely large, resulting in pdfs that are flat. The CEM begins favouring pdfs that generate samples with better performance, as seen at “Time ii”. At “Time iii”, the classic bell-shape curve has taken form. The pdfs continually become

5.3 Solution approach




Figure 5.6: A typical example of how the CEM progresses over time.

more refined as illustrated at “Time iv”. Eventually, the CEM aims to draw the optimal values with probability of one.

The theoretical foundations discussed in the following subsection are in most part general to the CEM. However, due to the stochastic nature of the system under study and the variables of interest being continuous, the focus is on continuous, stochastic optimisation.

5.3.1.1 Theoretical foundations

The CEM has mathematical foundations in *importance sampling* and the *Kullback-Leibler distance*. The CEM derives its name from the Kullback-Leibler distance $\mathcal{D}(g, h)$ in information theory, otherwise known as cross-entropy. A measure of the distance between g and h is

$$\mathcal{D}(g, h) = \mathbb{E}_g \left[\ln \frac{g(\mathbf{X})}{h(\mathbf{X})} \right] \quad (5.16)$$

$$= \int g(\mathbf{x}) \ln g(\mathbf{x}) d\mathbf{x} - \int g(\mathbf{x}) \ln h(\mathbf{x}) d\mathbf{x}. \quad (5.17)$$

Strictly speaking, cross-entropy does not represent distance because it is asymmetric: $\mathcal{D}(g, h) \neq \mathcal{D}(h, g)$. However, it is referred to as distance because $\mathcal{D}(g, h) \geq 0$

5.3 Solution approach

and it represents a relative measure (Rubinstein & Kroese, 2004). If $g(\mathbf{x})$ and $h(\mathbf{x})$ are identical, then $\mathcal{D}(g, h) = 0$. In saying that, the CEM aims to minimise \mathcal{D} .

The sampling pdf $h(\mathbf{x})$ and the optimal pdf $g^*(\mathbf{x})$ are assumed to come from the same family of distributions. This assumption simplifies the procedure so that only the optimal parameters of $h(\mathbf{x})$ need to be found to minimise cross-entropy between the importance sampling pdf $h(\mathbf{x}, \mathbf{v})$ and the optimal pdf $g^*(\mathbf{x})$, with parameter vector \mathbf{v} (Rubinstein, 1999). This is equivalent to solving the maximisation problem (Rubinstein & Kroese, 2011), represented here as a program,

$$\max_{\mathbf{v}} [D(\mathbf{v}) = \mathbb{E}_{\mathbf{u}} I_{\{f(\mathbf{X}) \leq \gamma\}} \ln h(\mathbf{X}; \mathbf{v})], \quad (5.18)$$

where the indicator function $I_{\{f(\mathbf{X}) \leq \gamma\}}$ is defined as

$$I_{\{f(\mathbf{X}) \leq \gamma\}} = \begin{cases} 1 & \text{if } f(\mathbf{X}) \leq \gamma \\ 0 & \text{if } f(\mathbf{X}) > \gamma. \end{cases} \quad (5.19)$$

The random vector \mathbf{X} is defined by pdf $h(\mathbf{x}, \mathbf{u})$ with parameter vector \mathbf{u} . Solving (5.18) instead of its variance minimisation counterpart is advantageous because the maximisation program can be solved analytically, whereas the said minimisation program can only be solved numerically (Rubinstein & Kroese, 2004).

Without the loss of generality, let us consider the minimisation of $f(\mathbf{x})$ with a minimum value of γ^* . This is not to be confused with solving the maximisation problem in (5.18). Instead of solving the deterministic problem

$$\gamma^* = \min_{\mathbf{x}} f(\mathbf{x}), \quad (5.20)$$

this study is focussed on the associated stochastic problem

$$r(\gamma) = \mathbb{P}_{\mathbf{u}}(f(\mathbf{X}) \leq \gamma) = \mathbb{E}_{\mathbf{u}} I_{\{f(\mathbf{X}) \leq \gamma\}}. \quad (5.21)$$

By setting optimal events as rare-events of interest r , the optimal solution can be estimated through adaptive changes to the pdf using cross-entropy \mathcal{D} . To guide the algorithm towards the optimal solution, a sequence of pdfs $h(\mathbf{x}, \mathbf{v}_0), h(\mathbf{x}, \mathbf{v}_1), h(\mathbf{x}, \mathbf{v}_2), \dots$ is created, with $\mathbf{v}_0 = \mathbf{u}$. A sequence of tuples

5.3 Solution approach

$\{(\hat{\gamma}_i, \hat{\mathbf{v}}_i)\}$ is generated with the aim of convergence to the optimal tuple $\{\gamma^*, \mathbf{v}^*\}$, through an iterative updating procedure (Rubinstein & Kroese, 2011):

1. **Adaptive updating of γ_i .** For a fixed \mathbf{v}_{i-1} , let γ_i be the ρ -quantile of $f(\mathbf{x})$ under \mathbf{v}_{i-1} . Namely, $\mathbb{P}_{\mathbf{v}_{i-1}}(f(\mathbf{X}) \leq \gamma_i) \leq \rho$, with a small value for ρ such as $\rho = 10^{-2}$. γ_i is estimated by drawing a random sample $\mathbf{X}_1, \dots, \mathbf{X}_n$ from $h(\mathbf{x}, \mathbf{v}_{i-1})$ and then evaluating the sample ρ -quantile of performances

$$\hat{\gamma}_i = f_{[\rho N]}.$$
 (5.22)

2. **Adaptive updating of \mathbf{v}_i .** Based on (5.18), for a fixed γ_i and \mathbf{v}_{i-1} , derive \mathbf{v}_i from the solution of the program

$$\max_{\mathbf{v}} [D(\mathbf{v}) = \mathbb{E}_{\mathbf{v}_{i-1}} I_{\{f(\mathbf{X}) \leq \gamma_i\}} \ln h(\mathbf{X}; \mathbf{v})].$$
 (5.23)

In the stochastic problem setting, the value of $\max_{\mathbf{v}} D(\mathbf{v})$ in (5.23) can be estimated by means of the stochastic program

$$\max_{\mathbf{v}} \left[\hat{D}(\mathbf{v}) = \frac{1}{N} \sum_{a=1}^N I_{\{f(\mathbf{x}_a) \leq \hat{\gamma}_i\}} \ln h(\mathbf{X}_a; \mathbf{v}) \right].$$
 (5.24)

A smoothing function can be employed to update \mathbf{v} so as to avoid premature convergence

$$\hat{\mathbf{v}}_i = \alpha \tilde{\mathbf{v}}_i + (1 - \alpha) \hat{\mathbf{v}}_{i-1},$$
 (5.25)

Algorithm 1 Main CEM algorithm: Continuous optimisation

- 1: Choose an initial parameter vector for the pdf $h(\mathbf{x}, \mathbf{v})$.
 - 2: Set $i \leftarrow 1$.
 - 3: **while** termination criteria not met **do**,
 - 4: Generate a random sample $\mathbf{X}_1, \dots, \mathbf{X}_n$ from the pdf $h(\mathbf{x}, \mathbf{v}_{i-1})$.
 - 5: Evaluate the sample.
 - 6: Compute the ρ -quantile $\hat{\gamma}_i$ of the sample performance.
 - 7: Solve the stochastic program in (5.23).
 - 8: Smooth the parameter vector \mathbf{v}_i using (5.25).
 - 9: **end while**
-

5.3 Solution approach

with smoothing constant α usually in the range of 0.6 to 0.9. The parameter vector $\tilde{\mathbf{v}}_i$ is equivalent to \mathbf{v} obtained from (5.24).

Based on the aforementioned procedure, the basic CEM algorithm for continuous optimisation is shown in Algorithm 1 (Rubinstein, 1997; Rubinstein & Kroese, 2004, 2011).

5.3.1.2 Applications of the cross-entropy for optimisation

The CEM is widely applicable to many optimisation problems (Rubinstein & Kroese, 2004). The CEM can be applied to static or noisy optimisation problems. For instance, the CEM has performed well in optimising the noisy buffer allocation problem where the performance measure cannot be evaluated analytically and has to be estimated through simulation (Rubinstein & Kroese, 2011).

It has also shown promise in the field of combinatorial optimisation (De Boer *et al.*, 2005; Rubinstein & Kroese, 2004). The CEM has shown success in, amongst others, optimising the quadratic assignment problem, the travelling salesman problem, and the max-cut problem, which are NP (non-deterministic polynomial-time) hard problems and thus are computationally hard to solve, even to near optimality (Rubinstein & Kroese, 2004).

The CEM has shown success by rapidly converging to the global optima for continuous optimisation of well-known test functions with many local extrema, such as the Rastrigin function and Rosenbrock function (Rubinstein & Kroese, 2011).

In the electricity industry, Kothari & Kroese (2009) applied the CEM to the generation expansion problem. The problem is concerned with minimising the cost of meeting future electricity demand through the commission of electric power systems.

5.3.2 Cross-entropy method applied to this study

In this subsection, the chosen probability density function of the CEM method is put forth and the parameters used are discussed.

5.3.2.1 Use of a normal distribution to generate solutions

In the case of continuous optimisation, a simple yet effective choice for a pdf is the normal distribution (Rubinstein & Kroese, 2004). Assuming elements of \mathbf{X}

5.3 Solution approach

are independent, the adaptive updating procedure in (5.22) and (5.24) becomes more straightforward and is defined as

$$\tilde{\mu}_{ij} = \frac{\sum_{\mathbf{X} \in \xi_i} \mathbf{X}_j}{|\xi_i|} \quad j = 1, \dots, N, \quad (5.26)$$

and

$$\tilde{\sigma}_{ij} = \sqrt{\frac{\sum_{\mathbf{X} \in \xi_i} (\mathbf{X}_j - \tilde{\mu}_{ij})^2}{|\xi_i| - 1}} \quad j = 1, \dots, N, \quad (5.27)$$

with the elite set $\xi_i = \{\mathbf{X}_j : f(\mathbf{X}_j) \leq \hat{\gamma}_i\}$, where \mathbf{X}_i is the random sample X_1, \dots, X_n , for n decision variables at iteration i , drawn from a normal distribution with $\tilde{\mu}_{i-1}$ and $\tilde{\sigma}_{i-1}$ (Rubinstein & Kroese, 2011). That is, based on the maximum likelihood estimators for the best samples of each iteration, the means and standard deviations are updated.

The smoothing function in (5.25) is applied to $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ such that

$$\hat{\boldsymbol{\mu}}_i = \alpha \tilde{\boldsymbol{\mu}} + (1 - \alpha) \hat{\boldsymbol{\mu}}_{i-1} \quad \text{and} \quad (5.28)$$

$$\hat{\boldsymbol{\sigma}}_i = \alpha \tilde{\boldsymbol{\sigma}} + (1 - \alpha) \hat{\boldsymbol{\sigma}}_{i-1}. \quad (5.29)$$

5.3.2.2 Parameter settings

Table 5.2: CEM parameter settings.

Description	Symbol	Value
Smoothing parameter	α	0.75
Percentage of samples to include in the elite set	ρ	20
Population size	N	50
Maximum number of iterations	N_m	100

The chosen CEM parameters for this study are shown in Table 5.2. As the aim of the research is to determine if the CEM can be used to optimise the PEM, the fine-tuning of parameters is not investigated. The values chosen for α and ρ are default values. As the combined optimisation and simulation model need to be re-run by electric utility planners, reducing the run-time duration is important.

5.3 Solution approach

Therefore, population size (N) and the maximum number of iterations (N_m) are chosen to be small.

5.3.2.3 Nesting the primary energy module within the cross-entropy method

The three decision variables per power station are grouped together for practical purposes. Therefore, $\boldsymbol{\mu} = \{\mu_1, \dots, \mu_{3n}\}$ and $\boldsymbol{\sigma} = \{\sigma_1, \dots, \sigma_{3n}\}$. Elements $\{1, \dots, n\}$ represent T_p , elements $\{n+1, \dots, 2n\}$ represent L_p , and elements $\{2n+1, \dots, 3n\}$ represent U_p .

As defined in (5.7), the generated values for the decision variables are subject to constraints:

1. T_p must fall between the minimum $S_p^{(\min)}$ and maximum $S_p^{(\max)}$ allowable stockpile levels.
2. L_p must fall between between 0 and T_p .
3. U_p must fall between T_p and $S_p^{(\max)}$.

To ensure that generated solutions fall between the lower bounds (lb)s and upper bounds (ub)s of the decision variables, truncated normal distributions are used. Truncated distributions provide a simple means of containing the search space of an algorithm.

As with any metaheuristic, decision variables need to be initialised. In other words, $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ require values at iteration $i = 1$. In the case of the CEM, the initialised values of $\boldsymbol{\mu}$ are not of as much importance relative to other metaheuristics. This is because the CEM commences with extremely “flat” normal distributions (see “Time i” in Figure 5.6). To generate the said “flat” normal distributions, the values of $\boldsymbol{\sigma}$ are initially assigned arbitrarily high values, using $\boldsymbol{\sigma} = 10 \cdot (\mathbf{ub} - \mathbf{lb})$.

As mentioned in Points 2 and 3, L_p and U_p are partly bound by the generated solutions of T_p . Therefore, solutions are first generated for T_p . In other words, solutions are first generated using $\{\mu_1, \dots, \mu_n\}$ and $\{\sigma_1, \dots, \sigma_n\}$. Thereafter, the remainder of the solutions are generated from $\{\mu_{n+1}, \dots, \mu_{3n}\}$ and $\{\sigma_{n+1}, \dots, \sigma_{3n}\}$.

The procedure for the combined CEM and PEM is illustrated in Algorithm 2. N_m and N can be chosen by the analyst. In this study, the values defined in Table

5.3 Solution approach

Algorithm 2 Combined primary energy module and cross-entropy method.

- 1: Let N_m be the maximum number of iterations, $\boldsymbol{\mu}$ be the set of mean values in the elite set, and $\boldsymbol{\sigma}$ be the set of standard deviations in the elite set.
 - 2: $i \leftarrow 0$
 - 3: Initialise $\boldsymbol{\mu}_i$ and $\boldsymbol{\sigma}_i$.
 - 4: Set PEM's MC simulation replications.
 - 5: **while** termination criteria not met **do**
 - 6: $i \leftarrow i + 1$
 - 7: Generate a population of N solutions from a truncated normal distribution with $\boldsymbol{\mu}_{i-1}$ and $\boldsymbol{\sigma}_{i-1}$.
 - 8: **for** each solution j of the population **do**
 - 9: Calculate emergency deliveries and delivery cancellations through the PEM using Algorithm 3.
 - 10: Evaluate each solution.
 - 11: **end for**
 - 12: Rank the solutions.
 - 13: Determine the elite set.
 - 14: Update $\boldsymbol{\mu}_i$ and $\boldsymbol{\sigma}_i$.
 - 15: Smooth $\boldsymbol{\mu}_i$ and $\boldsymbol{\sigma}_i$ according to (5.28) and (5.29), respectively.
 - 16: **end while**
-

5.2 are used. In Step 4, the Monte Carlo simulation settings range from 1 to 1 000. After performing tests, it was found that increasing the number of simulation replications in the MC simulator does not significantly increase computational time. The majority of the computational time is spent reading input files and writing output files. Thus, the number of MC replications is set to the maximum value of 1 000. This minimises the statistical estimation error. The loop of Steps 5 to 16 runs until the termination criteria is reached. Solutions are drawn from the pdfs and evaluated, and then used to update the pdf parameters. $e_{p,t}$ and $c_{p,t}$ are calculated as shown in Algorithm 3.

The procedure in Algorithm 3 runs the simulator on a month-to-month basis, as indicated in Step 3. This is contrary to how the PEM is run when not combined with the CEM. Running the PEM on a month-by-month allows for actions — based on policies — to take place. In other words, month-end stockpile levels are reviewed and then if actual stockpile levels are below or above the warning limits, then emergency deliveries and delivery cancellations are made, respectively.

5.4 Concluding remarks on Chapter 5

Algorithm 3 Calculation of emergency deliveries and delivery cancellations

```

1: Set matrix  $e \leftarrow \emptyset$  and matrix  $c \leftarrow \emptyset$ .
2: for each solution do
3:   for each month do
4:     Run the PEM MC simulator.
5:     if stockpile level + future emergency deliveries < lower warning limit
6:       then order emergency coal of quantity  $e$  that will result in
           replenishment to the target stockpile level.
7:     if stockpile level - future delivery cancellations > upper warning limit
8:       then cancel coal deliveries by quantity  $c$  that will result in
           reduction to the target stockpile level.
9:   end for
10: end for

```

All the power stations are optimised simultaneously by assigning three decision variables to each power station, namely: L_p , T_p , and U_p . When running the PEM, values for all the power stations and months are read from and written to the database.

This chapter is concluded in the following section.

5.4 Concluding remarks on Chapter 5

This chapter stepped through the conceptual and mathematical model formulation and thereafter discussed the chosen solution approach, in the form of the CEM. Integration of the PEM and CEM was also discussed, taking into account the nature of the Monte Carlo simulator. A robust solution approach is proposed through the use of default parameters in the CEM.

In the following chapter, the experimental design is put forth. A large portion of the following chapter is dedicated to the financial parameter estimation defined by the objective function in Section 5.2.2.

Chapter 6

Experimental design

In this chapter, financial parameters used for the objective function are first estimated. Thereafter, the simulation settings of the primary energy module are briefly discussed. Finally, the experimental setup is described.

6.1 Financial parameter estimation

Financial parameters are required to measure the performance of the combined optimisation and simulation model. However, exact data for the shortage, handling, emergency, and cancellation costs could not be obtained. Thus, it was necessary to estimate these costs. To be more clear, upon implementation of the optimiser, the estimated costs can be replaced by improved values. The estimated values serve as a proof of concept and in Section 7.4.1 of the following chapter, a sensitivity analysis is performed on the estimated values. The four components of the objective function (as defined in Section 5.2) are estimated in this section: shortage, holding, emergency delivery, and delivery cancellation costs.

6.1.1 Shortage cost

Shortage cost (SC) can be estimated in two ways:

1. The cost to the economy in the case of load-shedding.
2. The cost to operate open cycle gas turbines (OCGTs).

First, the latest version of the Integrated Resource Plan (IRP) for 2010–2030, as set out by [The Department of Energy \(2013\)](#), quantifies the cost to the economy

6.1 Financial parameter estimation

in the case of load-shedding. The cost to the economy is known as unserved energy. Unserved energy is the opportunity cost for the South African economy from electricity supply interruptions (The Department of Energy, 2013). The cost of unserved energy is estimated by The Department of Energy (2013) to be ZAR 75 000/MWh, aggregated over all consumers.

Secondly, OCGTs have high operating costs and are operated as a last resort. Estimating shortage costs as the cost of operating OCGTs is problematic because there is a limited OCGT capacity. OCGTs make up only 9.4% of Eskom’s total generating capacity of 41 933 MW (Eskom Holdings Limited, 2012b). The concern is “Will the mathematical model accurately represent shortage cost, if coal shortages are greater than 9.4%?”.

Although shortages in electricity production (stemming from coal stockpile shortages) being greater than the OCGTs generation capacity is highly unlikely, it is important to correctly represent the possible cost of shortages for the mathematical model. Furthermore, the OCGT generation capacity may in fact be reduced from 9.4% due to OCGTs that are already in use.

Given these points, the shortage penalty cost is chosen to be estimated as the unserved energy cost. Unserved energy is estimated in terms of ZAR per MWh, but the holding, emergency delivery, and cancellation of delivery costs are estimated as ZAR per kton. A single unit of measure (ZAR) is required to group the objectives into a single objective function. To do so, the response values from the actual stockpile levels need to be converted from ktons to expected MWhs. (2.4) and (5.10) are combined to form

$$s_{p,t} = w_{p,t} \cdot (A_{p,t-1} + R_{p,t} - B_{p,t} - S_p^{(\min)}) \cdot \frac{CV_{p,t}}{H_p}. \quad (6.1)$$

Due to the fact that calorific value is defined by a probability density function, the mean calorific value per power station per month is used as an estimate.

6.1.2 Holding cost

Holding cost (HC) is the cost of carrying inventory per unit per time period (Winston, 2004). Typically, HCs consist of opportunity, storage, insurance, and

6.1 Financial parameter estimation

risk costs (Hillier & Lieberman, 2001). Risk costs are associated with theft and spoilage of coal. As coal is of low monetary value, risk costs and insurance costs play an insignificant role in this study. Likewise, storage costs such as rental fee are not relevant because the land at each coal stockyard is owned by Eskom. The most important cost associated with inventory on hand is opportunity cost (Winston, 2004).

Opportunity cost is the cost incurred due to capital being tied up in inventory. It is a percentage of the unit cost per time period per unit. The percentage is typically estimated as the interest rate attainable. The interest rate can be the rate of return on possible investments or the lending rate for loan repayment.

Due to Eskom's current state of affairs, the opportunity cost percentage is estimated as the probable attainable interest rate on a loan repayment. The lending rates considered are for a short to medium-term time period.

In summary, the lending rate represents the opportunity cost percentage which represents the holding cost percentage. The holding cost percentage is approximated by recent lending rates from three sources:

1. The South African Reserve Bank daily value.
2. The World Bank annual value for South Africa.
3. The South African government bond daily value.

First, The South African Reserve Bank (2014) set the prime lending rate (benchmark rate for private banks) at 9.25% as at October 16, 2014. Secondly, The World Bank (2014) calculated an average lending rate of 8.5% for 2013. Thirdly, government bonds are used by national governments to raise capital. The South African government bond lending rate was 7.77% as at October 16, 2014 (Trading Economics, 2014).

These three values were used to estimate the lending rate. It is also assumed that Eskom will be charged a sub-prime lending rate because it is a large company. Given these points, the estimated lending rate for Eskom is 8.25% (prime less 1%). The holding cost per year per kton of coal δ is thus estimated as 8.25% of the value of the coal.

6.1 Financial parameter estimation

6.1.3 Emergency delivery cost

Coal managers are willing to pay a higher unit price for coal as coal reserves near the minimum allowable level. A negative exponential distribution pdf was chosen to model this behaviour. Figure 6.1 illustrates the emergency cost ($EC_{p,t}$) function. $EC_{p,t}$ is normalised and scaled by the minimum stockpile level ($S^{(\min)}$) and the target stockpile level (T_p). $S^{(\min)}$ is a constant input and T_p is a decision variable. Only if the actual stockpile level ($A_{p,t}$) dips below the lower warning limit (L_p) decision variable, an emergency order is placed.

The closer L_p is to T_p , more emergency orders will be placed, but those orders will come at a cheaper unit cost. Conversely, the further L_p is from T_p , fewer emergency orders will be placed, but those orders will come at a more expensive unit cost.

The effect of scaling and normalising the $EC_{p,t}$ is represented in Figure 6.2. The further apart $S^{(\min)}$ and T_p , the flatter the emergency cost curve is (as seen in Case 1). On the contrary, the closer together $S^{(\min)}$ and T_p , the steeper the emergency cost curve is (as seen in Case 2).

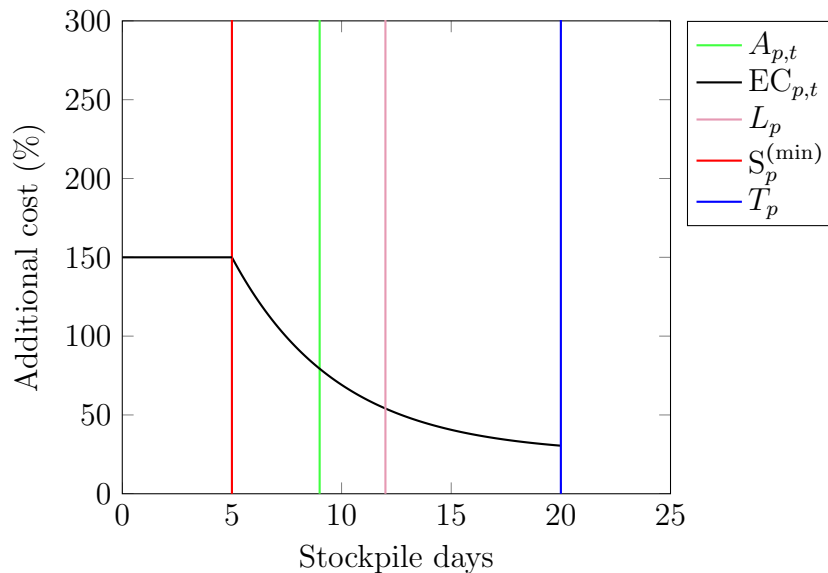


Figure 6.1: Additional cost for the emergency purchasing of coal, for a minimum stockpile level of five days and a target stockpile level of 20 days.

6.1 Financial parameter estimation

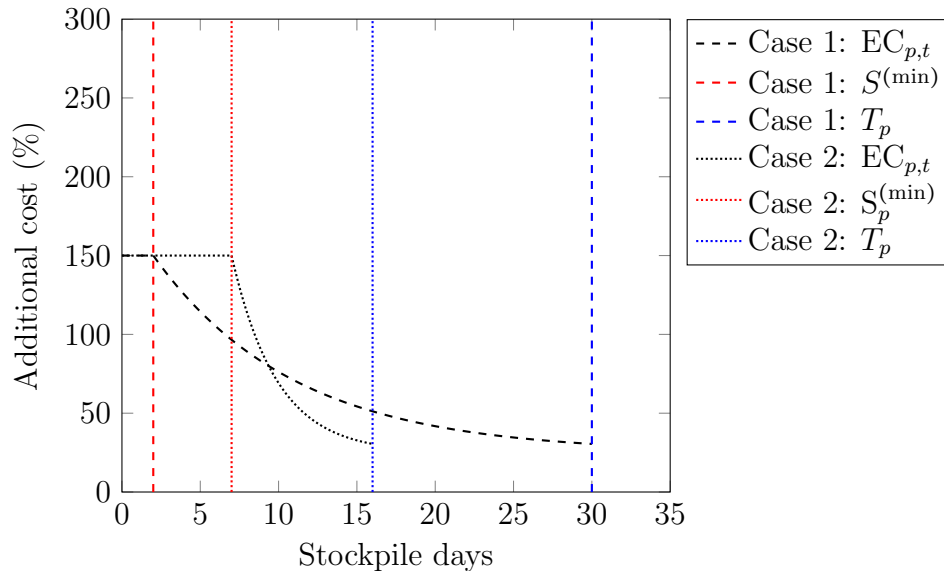


Figure 6.2: Additional cost for the emergency purchasing of coal, for two arbitrary cases for both the minimum and target stockpile level.

Emergency delivery cost $EC_{p,t}$ is a pdf and is defined as

$$EC_{p,t} = \begin{cases} \varphi + \phi & \text{if } A_{p,t} < S_p^{(\min)} \\ \varphi \cdot e^{-\varphi \cdot \Psi \cdot A_{p,t}} + \phi & \text{if } S^{(\min)} \leq A_{p,t} < L_p \\ 0 & \text{otherwise} \end{cases} \quad (6.2)$$

with the mean of the pdf $\varphi = 1.25$, and the vertical shift $\phi = 0.25$, and with the scaling and normalisation function

$$\Psi = \tau \cdot \frac{A_{p,t} - S^{(\min)}}{A_{p,t} \cdot (T_p - S^{(\min)})}. \quad (6.3)$$

Normalisation of $EC_{p,t}$ is done by

$$\frac{A_{p,t} - S^{(\min)}}{A_{p,t} \cdot (T_p - S^{(\min)})}$$

and is required to anchor the upper limit and lower limit of $EC_{p,t}$ to $S^{(\min)}$ and T_p respectively, allowing the “stretching” and “compressing” of $EC_{p,t}$. The scaling of $EC_{p,t}$ is done by $\tau = 2\varphi$, allowing for a good shape of the negative exponential function — one that is not too steep, nor too flat.

6.2 Simulation settings for primary energy module

6.1.4 Delivery cancellation cost

Delivery cancellation cost (CC) is modelled as the action taken when $A_{p,t}$ becomes too large and exceeds the upper warning limit (U_p). As no cost data could be acquired to estimate the cancellation penalty, it is simply modelled as a scale factor of the cost of purchasing coal ordinarily. Let κ be the penalty fee percentage for CC. κ is estimated as 0.5, which is half of the cost of purchasing coal for each power station.

6.2 Simulation settings for primary energy module

As experiments involve adjusting the settings of the PEM, these are briefly discussed. The PEM module provides the settings for the Monte Carlo (MC) simulation, including:

- start date
- end date
- seed
- number of replications
- output statistic.

The start and end dates define the range. The seed is used to generate different sequences of pseudo-random numbers. The value of the seed can be changed, but is not necessary for the purposes of this study. As the seed is only of interest if multiple algorithms are applied to this problem and compared against one another. Controlling the seed allows for the generation of common random numbers, which in turn provides a fair and unbiased means of comparing algorithm performance. The number of replications refers to the number of MC simulations to perform. The number of replications ranges from 1 to 1 000. The output statistic relates to any of the stochastic output variables in the EFS. In this study, the stochastic output variable of interest is the stockpile level. The minimum, 5-th percentile, mean, 95-th percentile, or maximum can be chosen as the output statistic. The output statistics are varied in this study to examine various risk approaches.

6.3 Experimental setup

6.3 Experimental setup

In this study, the combined simulation and optimisation model is applied to a base case scenario and three types of experiments, namely:

1. Sensitivity analysis of the financial parameter estimates.
2. Use of different output statistics.
3. Varying the baseline deliveries.

The parameters of the CEM, as detailed in Section 5.3.2.2, are not varied in the experiments. The focus is rather on analysing the application of the CEM to different instances of the PEM.

6.3.1 Base case

Parameters that are varied in the experiments are represented in Table 6.1 along with the assigned values for the base case. Further insight into why these values were chosen is now discussed. A minimum stockpile level is required at each power station, and this was set to five stockpile days. After performing multiple test runs by varying the maximum stockpile level, it was concluded that a reasonable limit was 50 stockpile days.

Table 6.1: Assigned values of the parameters for the base case.

Description	Symbol	Value
Minimum stockpile level	$S_p^{(\min)}$	5 stockpile days
Maximum stockpile level	$S_p^{(\max)}$	50 stockpile days
Initialised values of the target stockpile level	μ_1, \dots, μ_n	20 stockpile days
Initialised values of the lower warning limit	$\mu_{n+1}, \dots, \mu_{2n}$	12.5 stockpile days
Initialised values of the upper warning limit	$\mu_{2n+1}, \dots, \mu_{3n}$	35 stockpile days
Initialised values of the standard deviations	σ	$10 \cdot (\mathbf{ub} - \mathbf{lb})$
Mean of the emergency cost function	φ	1.25
Cancellation penalty percentage per kton	κ	50
Holding cost percentage per year per kton	δ	8.25
Baseline deliveries	$D_{p,t}$	Average $B_{p,t}$
Output statistic	–	Mean

6.3 Experimental setup

The values of T_p were initialised to the pre-2008 stockpile policy, as illustrated by Figure 2.5. The values of L_p were initialised as the midpoint between $S_p^{(\min)}$ and T_p . Likewise, the values of U_p were initialised as the midpoint between T_p and $S_p^{(\max)}$. To ensure that a “flat” normal distribution was sampled from initially, the standard deviations were initialised with arbitrarily high values, chosen as 10 times the difference of the upper bounds (ubs) and lower bounds (lbs) of the decision variables.

A sensitivity analysis is performed on three of the four financial parameters estimates, namely: φ , κ , and δ . The reasoning behind the choice of financial parameter estimates incorporated in the sensitivity analysis and the experimental setup thereof is detailed in Section 7.4.1. The baseline deliveries (BDs) are assigned the average quantity of coal burnt, at each power station for each month, as discussed in Section 5.1.1. Lastly, the chosen output statistic is the mean of the 1 000 generated MC sample paths.

Although data for the optimisation component of the model could not be attained and were therefore estimated, data for the energy flow simulator (EFS) were provided by the industry partner of this study. Specifically for the PEM, eight months of populated data were provided. All the results presented in this study are based on that instance of data. That said, the optimiser is designed to be versatile and work on any instance of data, granted that the data structure is consistent with the structure of the database of the EFS. Currently 14 power stations are defined in the EFS. Some of the information used in this thesis is not in the public domain. For reasons of confidentiality, some information is not disclosed in this thesis. For instance, power stations are referred to anonymously.

Coal stockpiles levels are represented as stockpile days, so as to have a common measure when comparing between power stations. The CEM was run for 100 iterations, because it was observed that decision variables values sufficiently converged after 100 iterations.

6.3.2 Sensitivity analysis of the financial parameter estimates

Sensitivity analysis is the systematic approach of varying input parameters and observing their effect on the model (Banks *et al.*, 1998), and is typically applied in two cases (Law *et al.*, 2000):

6.3 Experimental setup

1. Determining which parameters have the greatest effect on performance measures.
2. Determining which set of model specifications leads to optimal performance.

Considering the former, a robust approach was chosen for the parameters of the CEM. Hence, it is not important to analyse which parameters of the CEM lead to optimal performance. Considering the latter, as cost data for the optimiser's objective function is estimated, it is natural to ask the question, "how does the accuracy of the estimates affect solution quality?". Sensitivity analysis provides a means to quantify this effect.

The objective function of the optimisation model, as presented in Section 5.2.2, is comprised of four costs: holding, shortages, emergency deliveries and delivery cancellations. After analysing many test runs of the combined simulation and optimisation model, it was seen that shortage costs were never incurred. It makes sense that the optimisation model aims to avoid shortages, because shortages are essentially a soft (penalty) constraint. It is assumed that small changes in the estimate of shortage costs is unlikely to influence the results. Shortages were thus not included in the sensitivity analysis. This reduced the number of experiments required from 16 to eight.

Input parameters are defined as factors. Output performance measures are defined as responses. A standard approach of conducting sensitivity analysis is to keep the values of $k - 1$ factors the same, whilst adjusting the value of the remaining factor. This process is repeated for each factor.

A more computationally economical approach is to use the 2^k factorial design method which is detailed in Law *et al.* (2000). In the 2^k factorial design, each factor is defined by two values, resulting in 2^k responses. Law *et al.* (2000) state that the values should be "opposite" from one another, but not very far apart. If the values are too far apart, important elements of the response may possibly be hidden.

Each unique combination of factors, referred to as experiments, are represented in Table 6.2. The "+" and "-" symbols represent the greater and smaller value, respectively, for each decision value as shown in Table 6.3.

6.3 Experimental setup

Table 6.2: The experiment design matrix for the sensitivity analysis of the financial parameter estimates.

Experiment	Factors			Response
	Factor 1: φ	Factor 2: κ	Factor 3: δ	
1	-	-	-	\mathcal{R}_1
2	+	-	-	\mathcal{R}_2
3	-	+	-	\mathcal{R}_3
4	+	+	-	\mathcal{R}_4
5	-	-	+	\mathcal{R}_5
6	+	-	+	\mathcal{R}_6
7	-	+	+	\mathcal{R}_7
8	+	+	+	\mathcal{R}_8

Table 6.3: Values of each factor in the sensitivity analysis.

	Normal value	+	-
Factor 1: φ	1.25	1.75	0.75
Factor 2: κ	0.5	0.8	0.2
Factor 3: δ	8.25%	8.75%	7.75%

The effect of each factor is taken as the difference between responses. Hence all the greater responses are assigned the “+” value and, likewise, all the smaller responses are assigned the “-” value. The responses \mathcal{R}_j , with $j = \{1, \dots, 8\}$ and $i = \{1, \dots, 3\}$ in this case, determine the effect of the factor

$$\varrho_i = \frac{\sum_{j=1}^8 (-1)^{\lceil \frac{j}{2^{i-1}} \rceil} \mathcal{R}_j}{8}. \quad (6.4)$$

$\frac{j}{2^{i-1}}$ is rounded up to $\lceil \frac{j}{2^{i-1}} \rceil$. This is used to generate the same sequence of “+” and “-” signs as shown in the columns of Table 6.2.

6.3.3 Using different output statistics

The chosen output statistic to measure performance is usually the mean value. This represents a risk neutral approach. A best-case (risk taking) or worst-case (risk averse) approach can also be followed with the aim of analysing the effect this

6.3 Experimental setup

has on the final decision variable values. Therefore, the chosen output statistic for a risk averse approach is the 95-th percentile of the stockpile level. On the other hand, a risk taking approach would make use of the 5-th percentile. The 5-th percentile and 95-th percentile scenarios are compared against the base case.

6.3.4 Varying the baseline deliveries

The final experiment involves varying the BDs. In the base case, BDs are assigned the average quantity of coal burnt, at each power station and for each month. Varying BDs is done with the aim of testing L_p and U_p . The three experiments are:

1. Increase BDs by 20% for all power stations and months.
2. Decrease BDs by 20% for all power stations and months.
3. Randomly vary BDs for some power stations and months.

Table 6.4: Randomly generated scale factor values used to alter the baseline deliveries.

Power station	2013					2014		
	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar
A	1.00	0.91	0.73	1.02	0.85	1.00	1.25	0.82
B	1.20	1.00	0.84	1.00	1.22	1.00	0.77	1.00
C	1.00	1.00	1.00	0.97	1.00	1.11	0.83	1.00
D	1.17	0.78	1.00	1.16	0.94	0.94	0.73	1.00
E	1.22	0.95	0.94	0.90	1.04	1.07	0.77	0.77
F	1.00	1.00	1.03	1.00	0.91	1.00	1.00	0.98
G	0.84	1.21	1.00	0.89	1.16	0.70	1.24	0.70
H	1.00	0.82	1.00	1.00	1.00	1.00	1.00	1.00
I	1.27	1.21	1.13	0.96	1.01	1.02	1.01	1.11
J	1.00	1.00	0.91	0.94	1.00	1.00	1.00	1.08
K	0.86	1.23	1.00	1.00	0.74	1.00	0.83	0.76
L	1.09	1.12	1.00	1.16	0.83	1.00	0.88	1.00
M	1.29	0.80	0.96	1.13	1.28	1.00	1.08	1.27
N	1.04	0.95	0.70	1.00	1.00	1.00	0.77	1.00

6.4 Concluding remarks on Chapter 6

In the third case, BDs for a random number of months (1–8) per power station are varied. BDs for the designated months are varied by scale factors that are randomly drawn from $U(0.7, 1.3)$, where “1” represents the original BDs. The randomly drawn scale factor values are shown in Table 6.4.

The increased, decreased, and randomly varied baseline delivery scenarios are compared against the base case.

6.4 Concluding remarks on Chapter 6

This chapter first presented the cost estimates for holding inventory, inventory shortages, emergency deliveries, and delivery cancellations. Simulation settings for the primary energy module were then discussed. Lastly, experimental setup of the base case, financial parameter sensitivity analysis, “using different output statistics”, and “varying the baseline deliveries” was detailed. The results for the base case and three experiments are presented in Chapter 7, along with verification and validation of the proposed solution.

Chapter 7

Analysis of experimental results

The previous chapter discussed the details of the experimental design.

In this chapter, validation and verification of the primary energy module (PEM) and cross-entropy method (CEM) is first covered. The bulk of the chapter is dedicated to analysing and discussing the experimental results, according to the structure defined in the experimental setup of the previous chapter. Results for the base case are first analysed. Thereafter, results for three experiments are analysed, namely sensitivity analysis of the financial parameters; using different output statistics; and varying baseline deliveries.

7.1 Verification and validation

Verification and validation (V&V) of the combined simulation and optimisation model ensure that it was correctly developed. Verification involves ensuring the correctness of the computer model which includes correcting syntax errors, examining logic, fixing compiler and runtime errors, and debugging. Validation is focussed on determining whether the model is a sufficient representation of the real-world process.

V&V overlap in many cases and are therefore discussed jointly in this chapter. V&V are performed on both the PEM and the CEM. The work by [Banks *et al.* \(1998\)](#) and [Law *et al.* \(2000\)](#) serves as the basis for V&V issues addressed in this section.

V&V are not limited to this section, since observations from experiments

7.1 Verification and validation

also serve as a means of verifying and validating the combined simulation and optimisation model.

7.1.1 Validation and verification of the primary energy module

First, to have confidence in the PEM, input data was varied substantially in an attempt to crash the model. The model only crashed when values that were out of range were entered. For example, when the minimum and maximum delivery limits were reached.

Secondly, model reasonableness of the PEM was investigated. Changes were made to the model and the resultant model behaviour and outputs were analysed by the following factors:

Continuity. The PEM exhibited good continuity. A high initial stockpile level resulted in high stockpiles at the end of each month. Increasing the calorific value (CV) resulted in less coal being burnt to meet the same electricity demand requirements. As expected, reducing deliveries leads to reduced stockpile levels. Increased planned maintenance resulted in a lesser mass of coal burnt.

Consistency. The PEM exhibited consistent outputs for similar model runs.

Absurd conditions. The PEM was run for extreme delivery and initial stockpile conditions. The combination of low initial stockpiles and low deliveries resulted in significant reduction of the coal burnt at each power station, so much so that electricity demand could not be met. The three remaining combinations of extreme conditions for deliveries and stockpile levels did not effect the mass of coal burnt at each power station. Either stockpile levels were high enough to cater for eight months of shortfall, or there was an excess delivery. Setting the deliveries and initial stockpile levels to zero at a single power station resulted in additional load uptake at other power stations.

Thirdly, limitations of the PEM are revisited. These limitations could possibly influence the results of the optimisation model. The three main limitations are:

7.1 Verification and validation

Simulation resolution. The PEM is aggregated to a monthly level, which “flattens” variation in the system.

Calorific value. The CV of coal at each power station is estimated through a single probability distribution. A more accurate approach would be to model batches of coal each with their own distribution. In other words, defining multiple distributions for CV at each power station.

Limited data. Only eight months of data were attained, which means that the simulator only runs for eight time intervals. This problem is linked to the simulation resolution.

7.1.2 Validation and verification of the combined primary energy module and cross-entropy method

A basic understanding of the workings of the CEM was required before optimisation of the PEM could take place. This was achieved by applying the CEM to two continuous parameter test functions, namely: the *Rastrigin* and *Rosenbrock* functions. The CEM converged to the optimal solutions, as it should. Applying the CEM to the two test functions not only allowed the researcher to become well-acquainted with the CEM, it also provided a means for the researcher to learn the R programming language.

Thereafter, the PEM was integrated into the CEM. A skeleton model approach was used. First, the CEM was applied to optimise only one of the decision variables — the target stockpile level (T_p) — for a single power station. Secondly, it was then scaled up to optimise T_p for all 14 power stations. Thirdly, all three decision variables were optimised for a single power station. Finally, all three decision variables — including the lower warning limit (L_p) and upper warning limit (U_p) — were optimised for all the power stations.

Code was written in a modular manner, with each module representing a certain task. The main module was first developed. Thereafter, as each subsequent module was added, the code was scrutinised and, if necessary, debugged. An electronic copy of the code is provided on the compact disc (CD) that accompanies this submitted study.

7.1 Verification and validation

Multiple meetings were held with the industry partner, responsible for development of the energy flow simulator, to ensure that the modelling approach followed was suitable.

7.1.3 Running the primary energy module without the cross-entropy method

To begin with, the PEM is run without the optimiser in order to further illustrate the need for the optimiser. As discussed in Section 5.1.1, baseline deliveries are assigned the average mass of coal burnt, at each power station for each month, through 1 000 independent simulation runs. To ensure that shortages do not occur so that they do not have an influence on the coal burn estimates at any of the power stations, initial stockpile levels were set arbitrarily high.

Differences in monthly stockpile levels are represented in Figure 7.1. 500 Monte Carlo (MC) sample paths are shown, with the mean as the chosen output statistic. 14 power stations are illustrated as A–N, for a planning horizon of eight months. A thicker band represents greater variation between delivery and burn. For some power stations, such as Power station B, there is continuous over-delivery. On the contrary, Power stations L and N exhibit large quantities of under-delivery. The cumulative monthly difference in stockpile levels, from the initial stockpile level, is represented in Figure 7.2. 500 MC sample paths are again represented.

7.1 Verification and validation




Figure 7.1: Stockpile level variation on a month-by-month basis.

7.1 Verification and validation




Figure 7.2: Cumulative monthly stockpile variation from the initial stockpile level.

7.2 Factors that influence the results

7.2 Factors that influence the results

Before presenting the experimental results, likely results and some factors that influence the results are discussed.:

- When a large amount of variation in stockpile levels is exhibited at a power station, a large buffer exists in the form of T_p .
- Likewise, when there is little variation at a power station, a small buffer size exists in the form of T_p .
- In the case of continual over-delivery, only U_p comes into effect and L_p slowly converges.
- By the same token, in the case of continual under-delivery, only L_p comes into effect and U_p slowly converges.
- Holding costs and shortage costs always drive T_p . Thus, T_p should converge to a sensible value.
- Shortages are essentially a soft constraint. Solutions should therefore tend to avoid incurring shortage costs.
- The combined PEM and CEM is run for a range of only eight months (due to limited data availability). It may take a few months for stockpiles to significantly vary from the the T_p and, in turn, emergency deliveries (EMDs) and delivery cancellations (DCAs) may only occur towards the end of the planning horizon, or they may not even occur.

7.3 Analysis of the results for the base case

As shown in Figure 7.3, the objective function value for the base case approaches and then varies around ZAR 600 million, due to the stochastic nature of the PEM. Progression of the approximated decision variable values is shown in Figures 7.4, 7.5, and 7.6. As a personal touch by the researcher, the colours used to represent the results were chosen from the colour-blind-friendly palette developed by [Okabe & Ito \(2014\)](#).

7.3 Analysis of the results for the base case

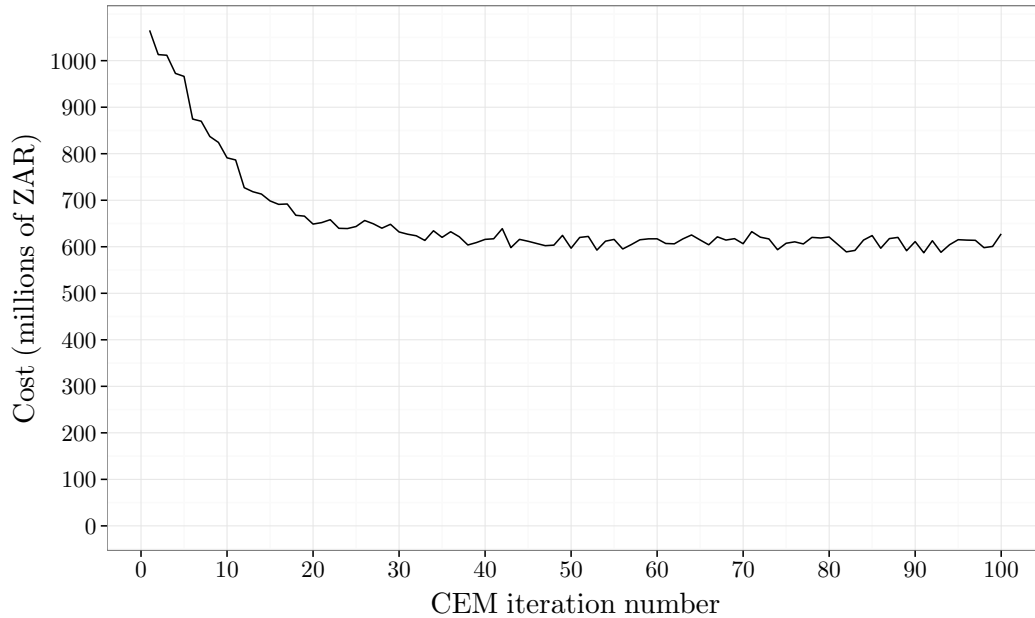


Figure 7.3: Progression of the objective function value for the base case.

As can be seen in Figures 7.4, 7.5, and 7.6, T_p converges before L_p and U_p , due to the objective function always driving T_p . In the case of continual over-delivery, only DCAs come into effect, through U_p . On the contrary, when continual under-delivery occurs, only EMDs come into effect, through U_p . In other words, the objective function does not always drive L_p and U_p , leading to slow convergence of some of the values for L_p and U_p .

The progress of the standard deviation of the decision variable values is represented in Figures 7.7, 7.8, and 7.9. The σ values decrease quickly over-time and almost approach zero. This behaviour is as a result of the nature of the CEM, because it originates from a variance reduction method.

As Figures 7.4, 7.5, and 7.6 are slightly cluttered, Figure 7.10 is introduced. The progression of all three decision variable values for all the power stations is illustrated. The shaded area represents the region between the warning limits L_p and U_p , and the solid line represents T_p . For the majority of power stations, T_p decreases from the initialised values to minimise holding costs while avoiding shortages.

7.3 Analysis of the results for the base case

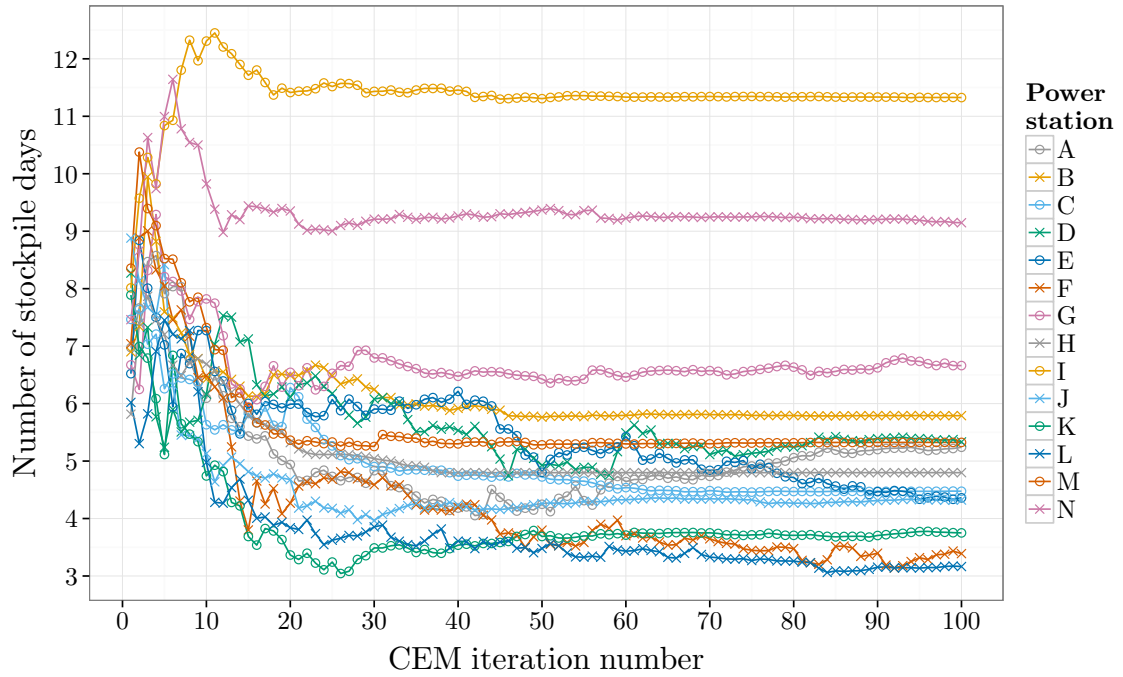


Figure 7.4: Progression of the μ values of L_p for the base case.

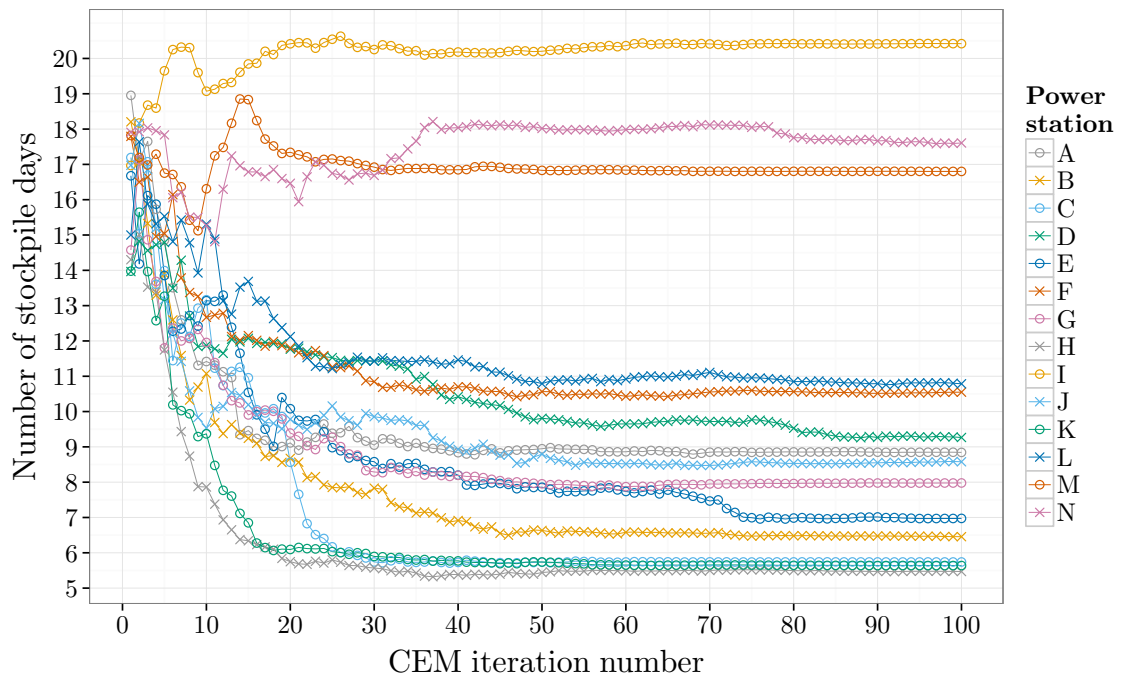


Figure 7.5: Progression of the μ values of T_p for the base case.

7.3 Analysis of the results for the base case




Figure 7.6: Progression of the μ values of U_p for the base case.

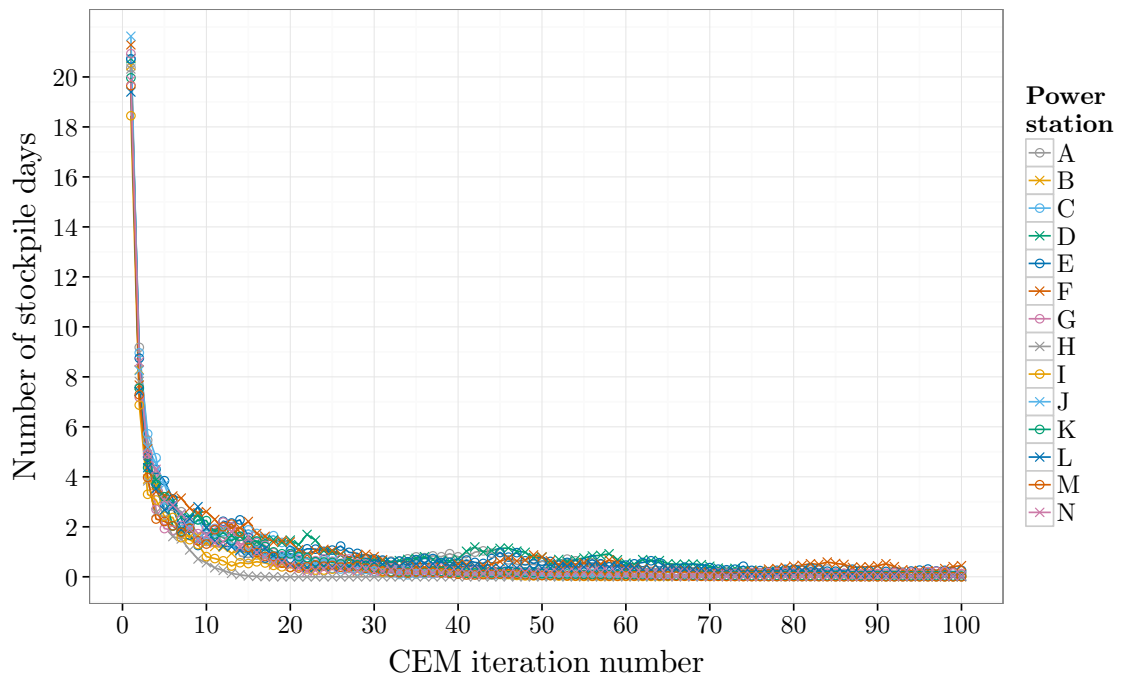


Figure 7.7: Progression of the σ values of L_p for the base case.

7.3 Analysis of the results for the base case

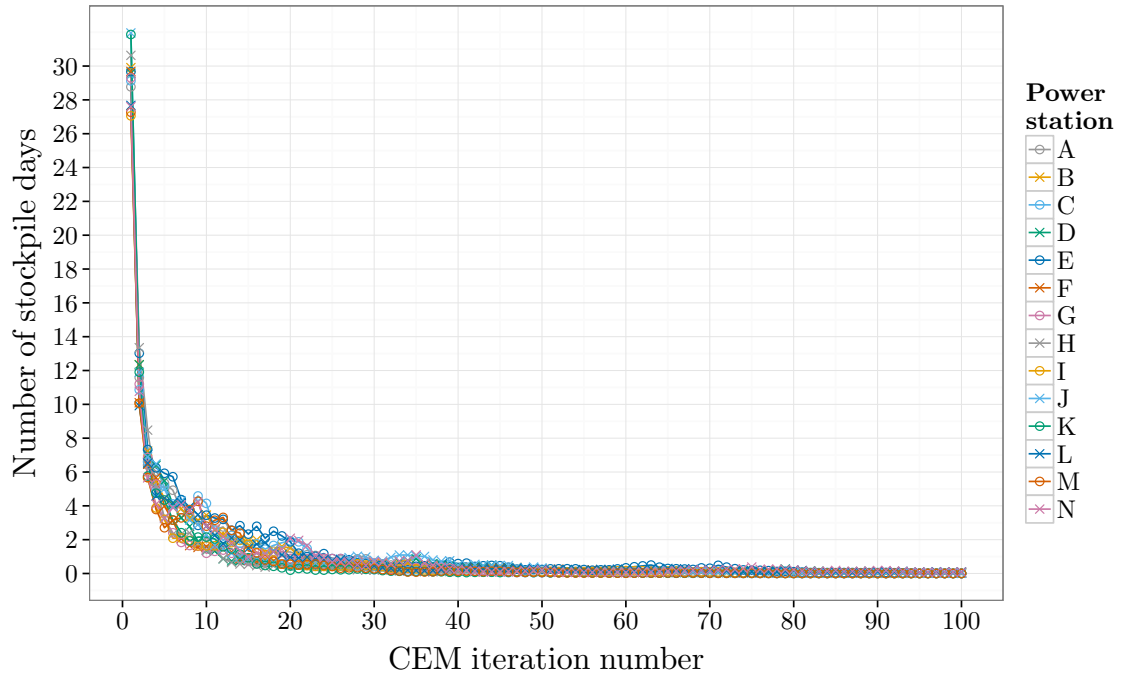


Figure 7.8: Progression of the σ values of T_p for the base case.

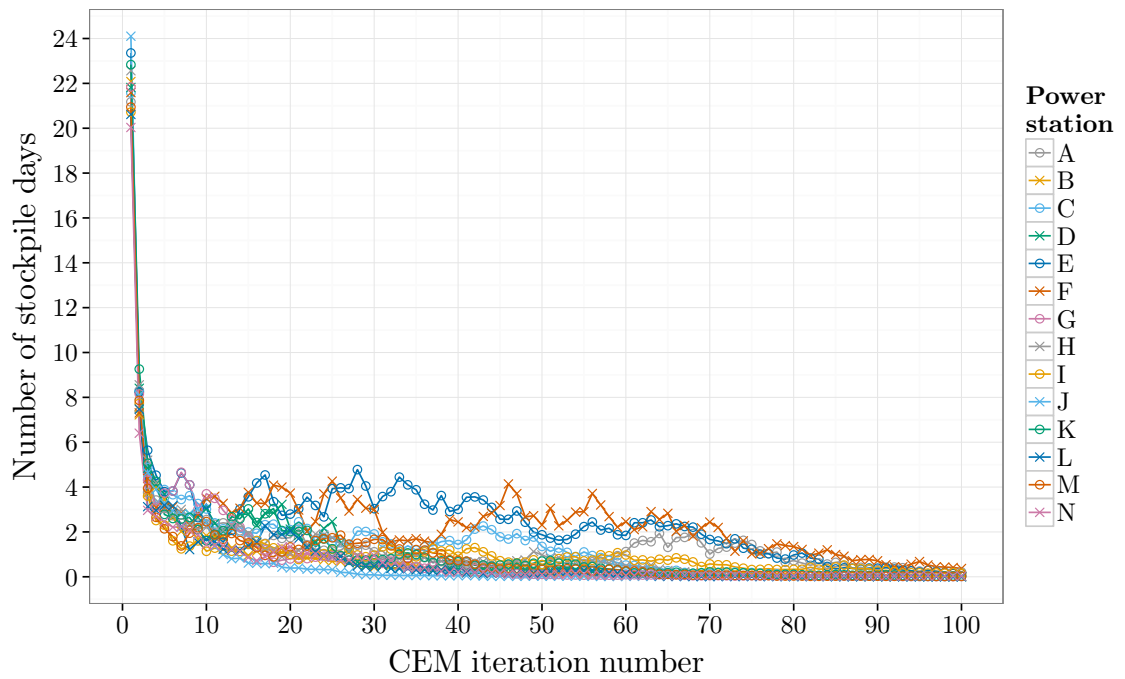


Figure 7.9: Progression of the σ values of U_p for the base case.

7.3 Analysis of the results for the base case

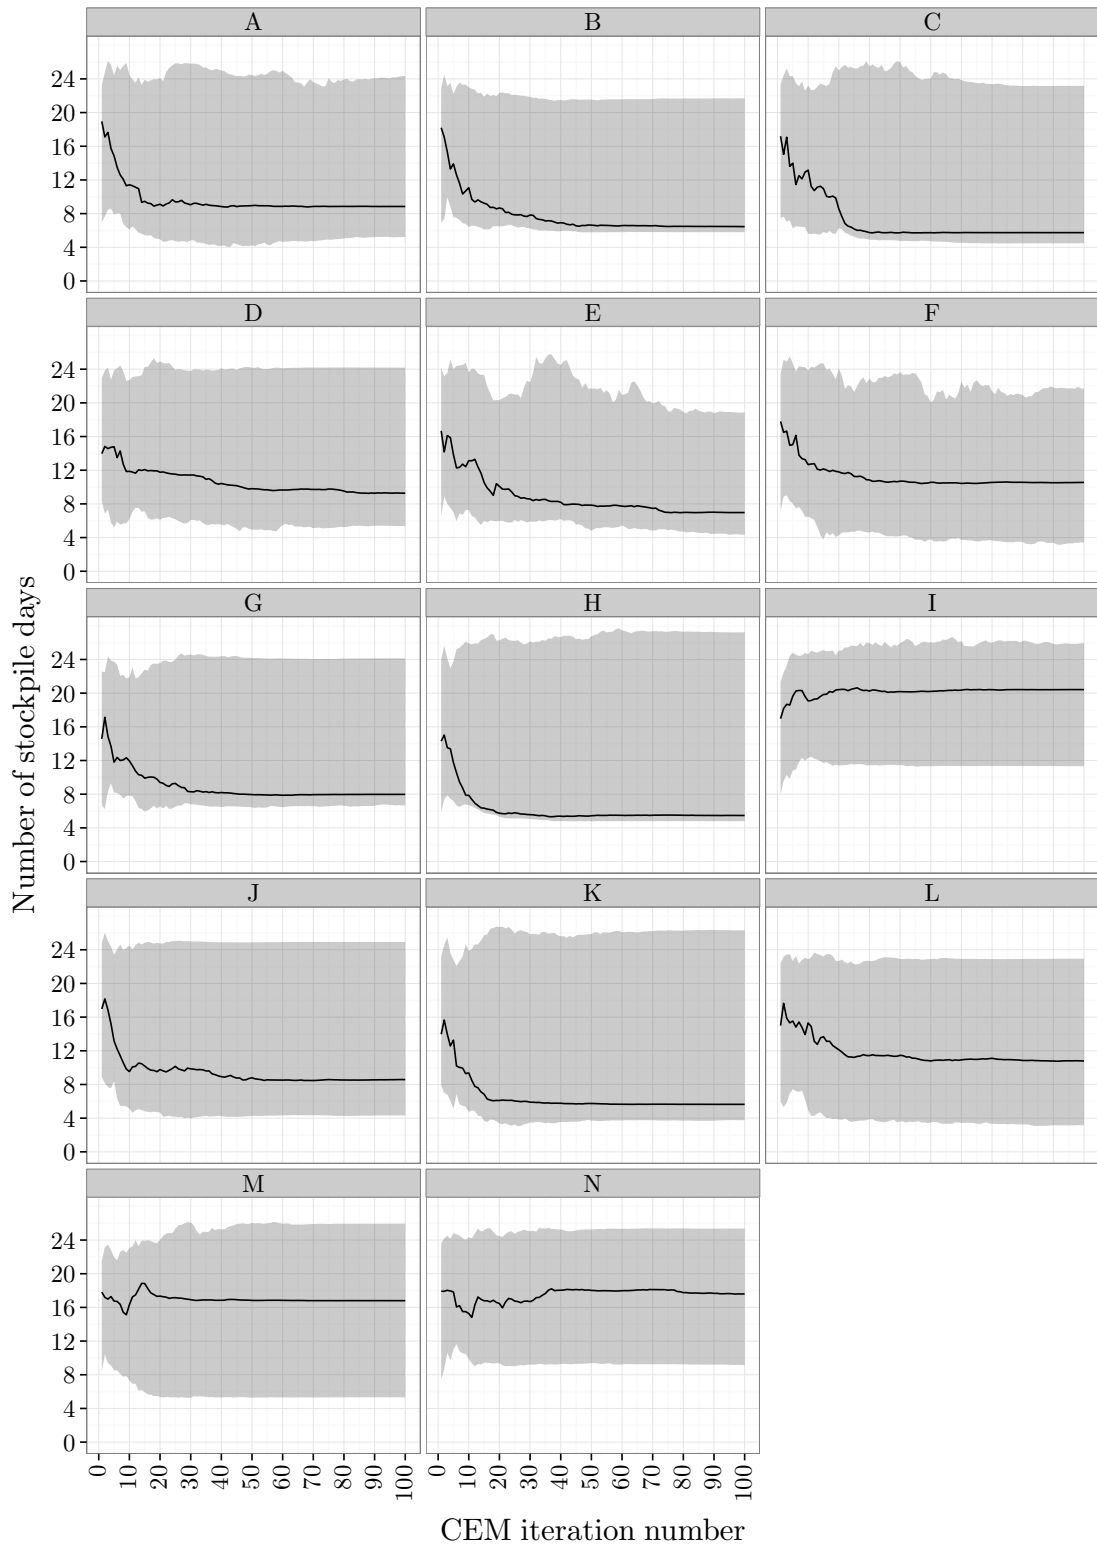


Figure 7.10: Progression of the μ values of all three decision variables for the base case.

7.3 Analysis of the results for the base case

It is important to note that a simulation run of 100 iterations takes approximately 10 days to complete when running on an Intel i7-3632QM 2.2GHZ processor with 8GB of RAM and a hard drive that operates at 5400RPM. This computational burden is predominantly attributed to the slow I/O disk rate. Each time the PEM is triggered, it reads from and writes to the database. With this process being repeated thousands of times, simulation optimisation quickly becomes computationally expensive. One means of reducing the computational burden is to launch the database in RAM. Although the lengthy computational time is undesirable, decision-makers are still able to optimise coal stockpile policies through the combined simulation and optimisation model.

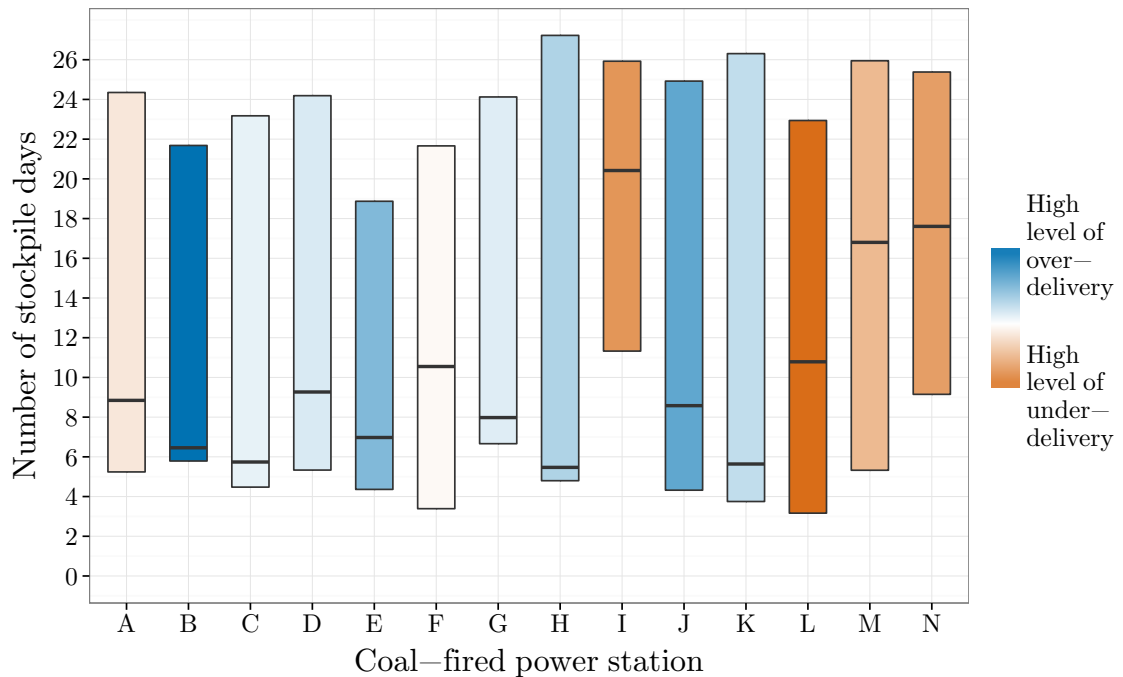


Figure 7.11: Final approximation of the decision variable values for the base case.

In Figure 7.11, each box represents the final approximation of all three decision variable values for a power station. The upper horizontal line represents U_p , the lower horizontal line represents L_p , and the horizontal line in between L_p and U_p represents T_p . A gradient colour scale is introduced to represent over-delivery, balanced delivery, and under-delivery. The results in this study differ from the

7.3 Analysis of the results for the base case

electric utility's pre-2008 and post-2008 stockpile policies of 20 and 42 stockpile days (refer to Figure 2.5). Insight into the lower than expected values of T_p is gained through analysis of the results.

When continual over-delivery occurs, L_p does not come into effect and slowly converges. Likewise, when continual under-delivery occurs, U_p slowly converges. The values of T_p for over-delivery and balanced delivery are generally lower than the values of T_p for under-delivery, as expected. T_p — for many of the power stations with over-delivery and balanced delivery — thus has a value that is only slightly larger than the minimum stockpile level of five days.

EMDs and DCAs for the final approximation are depicted in Figures 7.12 and 7.13, respectively. It makes sense that Power station I and Power station N incur EMDs, because they are subject to continual under-delivery (refer to Figure 7.2). A DCA occurs for Power station B, since it is subject to continual over-delivery. EMDs and DCAs occur from the middle to the end of the simulation range, since the initial stockpile level is equivalent to T_p . Running the simulator for a greater number of time intervals would result in more EMDs and DCAs occurring at the aforementioned power stations and also at many of the remaining power stations.

Revisiting Figure 7.11, consider Power stations I, L, M, and N. Power stations L and M have lesser values for L_p , since they do not incur EMDs. On the other hand, Power station I and N incur EMDs and thus require a buffer, between L_p and the minimum stockpile level of five days, to cater for random delivery lead time.

7.3 Analysis of the results for the base case




Figure 7.12: Emergency deliveries for the final approximation of the base case.

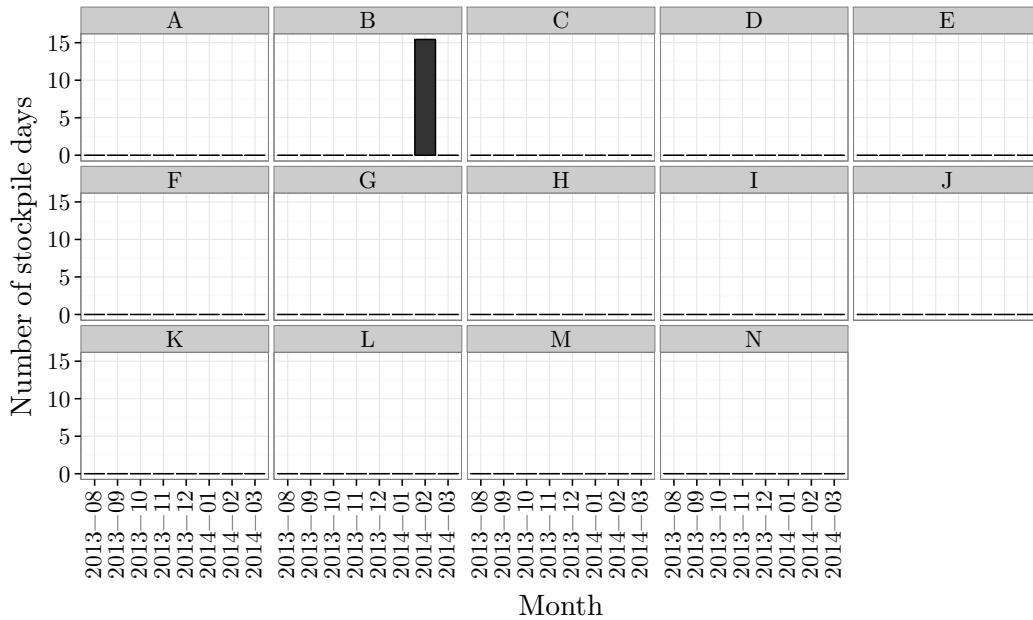


Figure 7.13: Delivery cancellations for the final approximation of the base case.

7.3 Analysis of the results for the base case

Final approximation values of the objective function are shown in Table 7.1. The “%” symbol represents the difference in the objective function value, relative to the base case. Final approximations of L_p , T_p , and U_p are shown in Tables 7.2, 7.3, and 7.4, respectively. The tabulated results are for the base case and three experiments. The sensitivity analysis is made up of eight scenarios, referred to as design points. The experiment titled “using different output statistics” is comprised of two scenarios, namely: the 5-th percentile and the 95-th percentile. Lastly, the experiment whereby BDs are randomly varied is made up of three scenarios: reduced, increased, and randomly varied BDs. This results in a total of 13 scenarios, along with the base case.

Table 7.1: Final approximated values of the objective function.

Experiments		Cost	
		(millions of ZAR)	(% difference)
Base case		627.7	0.0
Sensitivity analysis of the financial parameter estimates	Design point 1	888.4	41.5
	Design point 2	687.5	9.5
	Design point 3	613.1	-2.3
	Design point 4	640.0	2.0
	Design point 5	788.0	25.5
	Design point 6	625.6	-0.3
	Design point 7	604.7	-3.7
	Design point 8	534.8	-14.8
Using different output statistics	5-th percentile	1683.0	168.1
	95-th percentile	1145.0	82.4
Varying the baseline deliveries	$0.8 \cdot D_{p,t}$	592.9	-5.5
	$1.2 \cdot D_{p,t}$	644.4	2.7
	$U(0.7, 1.3) \cdot D_{p,t}$	612.3	-2.5

Significant findings of the base case have already been discussed, and significant findings of the three experiments will be explained in the following section.

7.3 Analysis of the results for the base case

Table 7.2: Final approximated values of L_p measured in stockpile days.

Experiments	L_A	L_B	L_C	L_D	L_E	L_F	L_F	L_H	L_I	L_J	L_K	L_L	L_M	L_N
Base case	5.2	5.8	4.5	5.3	4.4	3.4	6.7	4.8	11.3	4.3	3.8	3.2	5.3	9.1
Design point 1	5.0	5.2	3.9	5.5	5.3	7.1	3.8	4.4	7.8	6.4	5.2	2.6	4.2	4.5
Design point 2	4.6	6.1	5.7	5.1	5.6	4.8	6.6	5.3	11.2	5.9	4.8	4.6	5.1	8.3
Design point 3	4.9	5.6	6.9	3.5	4.1	2.7	5.1	4.7	11.0	3.6	3.7	2.4	3.9	10.6
Design point 4	5.3	5.5	4.6	5.5	5.2	5.8	6.5	5.4	13.1	6.8	4.2	4.1	5.8	7.5
Design point 5	6.6	5.1	3.6	5.4	6.0	4.4	4.7	4.8	9.7	4.8	4.7	3.4	5.2	4.1
Design point 6	6.0	6.6	5.2	3.4	5.8	2.5	5.1	4.9	12.2	7.9	4.8	5.1	5.7	10.0
Design point 7	5.1	6.0	7.3	5.6	5.5	3.9	4.8	5.2	10.6	6.9	4.9	3.7	4.1	7.8
Design point 8	5.0	4.9	8.4	3.9	5.7	5.6	5.2	4.8	11.7	4.5	4.8	5.0	5.5	8.8
5-th percentile	9.8	11.2	9.6	7.4	6.8	10.8	9.9	7.6	13.2	6.7	9.8	13.1	11.4	10.6
95-th percentile	3.3	6.1	7.5	5.3	7.2	5.9	5.4	5.5	10.6	9.3	5.6	5.7	5.5	5.5
$0.8 \cdot D_{p,t}$	3.8	5.2	5.3	5.2	5.7	6.4	3.0	3.9	10.8	4.2	3.5	2.7	4.0	8.0
$1.2 \cdot D_{p,t}$	6.1	5.9	5.6	7.0	6.6	5.9	5.1	5.9	11.2	7.3	3.6	5.2	5.9	7.0
$U(0.7, 1.3) \cdot D_{p,t}$	3.1	3.2	5.7	7.6	5.3	6.4	4.1	4.5	12.3	6.2	3.5	2.8	5.3	8.6

7.3 Analysis of the results for the base case

Table 7.3: Final approximated values of T_p measured in stockpile days.

Experiments		T_A	T_B	T_C	T_D	T_E	T_F	T_F	T_H	T_I	T_J	T_K	T_L	T_M	T_N
Base case		8.8	6.5	5.7	9.3	7.0	10.5	8.0	5.5	20.4	8.6	5.6	10.8	16.8	17.6
Design point 1		5.6	6.5	7.0	6.2	6.1	9.8	5.7	5.6	18.1	8.4	5.8	8.0	13.3	7.9
Design point 2		8.3	8.7	6.3	6.6	7.1	10.4	7.3	6.0	21.1	7.5	5.4	14.2	16.8	16.7
Design point 3		8.4	8.2	8.2	8.3	6.8	10.4	5.7	5.2	20.4	6.5	5.4	10.8	17.0	16.1
Design point 4		8.3	7.0	5.6	6.1	6.4	11.6	7.3	6.4	18.2	8.9	5.8	11.2	17.3	17.5
Design point 5		7.7	10.3	10.0	6.2	12.3	5.9	5.2	5.3	18.1	9.2	5.3	8.8	14.1	9.1
Design point 6		10.2	8.6	9.4	7.2	6.5	10.6	5.6	6.1	18.4	12.5	5.6	11.6	16.9	14.8
Design point 7		8.2	6.9	8.1	6.2	7.5	10.7	5.5	7.1	20.0	11.6	5.5	10.7	16.8	17.4
Design point 8		8.6	5.5	9.4	5.5	6.5	11.3	5.9	5.3	20.1	6.4	5.3	12.4	17.4	17.2
5-th percentile		18.5	18.6	18.4	17.9	15.5	18.2	18.9	17.0	19.9	19.2	16.2	23.2	20.3	19.2
95-th percentile		6.7	7.3	12.0	8.8	12.8	6.6	6.9	7.2	14.0	15.8	7.4	6.6	6.3	7.5
Varying the baseline		9.2	10.6	8.9	5.7	6.1	10.7	6.1	4.6	17.0	6.7	5.3	9.0	15.1	14.0
deliveries		9.7	6.6	7.0	7.8	7.2	12.0	8.4	6.5	21.4	9.2	6.5	18.0	18.7	27.2
$U(0.7, 1.3) \cdot D_{p,t}$		7.2	6.8	8.3	9.2	9.2	11.2	6.7	5.3	22.8	11.0	6.2	12.9	18.0	17.0

7.3 Analysis of the results for the base case

Table 7.4: Final approximated values of U_p measured in stockpile days.

Experiments		U_A	U_B	U_C	U_D	U_E	U_F	U_F	U_H	U_I	U_J	U_K	U_L	U_M	U_N
Base case	Design point 1	24.3	21.7	23.2	24.2	18.9	21.7	24.1	27.2	25.9	24.9	26.3	22.9	25.9	25.4
	Design point 2	25.7	21.0	24.7	22.2	26.4	23.3	23.2	19.5	26.5	23.0	21.7	18.5	20.4	18.6
	Design point 3	21.5	23.4	23.3	22.9	21.8	26.8	20.1	23.8	26.2	26.4	29.0	23.0	23.7	26.6
	Design point 4	17.2	23.1	23.2	25.5	24.9	26.6	26.4	25.5	25.6	26.1	20.1	20.4	23.0	23.7
	Design point 5	27.0	22.3	20.6	26.1	21.2	19.7	19.7	23.1	25.1	27.5	21.6	21.2	21.7	23.9
	Design point 6	22.8	21.4	26.2	22.2	21.4	24.1	21.9	26.2	25.1	22.4	23.5	19.0	25.6	27.2
	Design point 7	23.1	20.2	25.1	23.7	22.9	23.3	24.1	23.1	24.5	27.3	26.8	21.5	23.4	23.0
	Design point 8	20.7	17.9	24.3	22.3	23.5	23.6	21.2	19.0	26.4	27.9	19.1	20.6	23.1	27.1
Using different output statistics	5-th percentile	24.4	18.3	28.2	20.1	27.6	27.0	21.9	25.4	24.3	27.0	18.0	25.7	24.9	24.9
	95-th percentile	27.6	30.1	30.8	33.8	25.2	23.3	29.6	27.9	29.8	30.0	29.1	30.9	33.3	30.6
Varying the baseline deliveries	$0.8 \cdot D_{p,t}$	28.0	30.5	25.1	22.1	25.2	25.8	29.8	29.4	32.3	29.9	30.4	25.2	31.6	29.9
	$1.2 \cdot D_{p,t}$	21.3	22.5	25.6	22.5	20.6	22.3	27.4	19.4	23.3	27.4	21.1	24.6	26.3	21.0
	$U(0.7, 1.3) \cdot D_{p,t}$	35.2	37.4	37.0	26.7	36.6	38.4	38.9	36.1	30.4	28.4	42.1	37.6	30.7	38.7
		29.5	36.6	26.5	28.2	29.0	33.1	27.4	29.8	37.1	30.0	29.1	30.4	31.2	28.6

7.4 Experimental results

7.4 Experimental results

Experiments were conducted as detailed in Section 6.3. Results of the experiments are reviewed in this section, beginning with sensitivity analysis of the financial parameters. Secondly, experimental results are presented for the use of different outputs statistics, namely: the 5-th and 95-th percentiles. Lastly, experimental results of the varied BDs are analysed.

7.4.1 Sensitivity analysis of the financial parameter estimates

Results for the sensitivity analysis of financial parameters are shown in Table 7.5. The response values are identical to the costs of Design points 1–8 in Table 7.1. The effect measures the average change in the response, as a result of a change in an individual factor.

Table 7.5: Final approximated objective function values for sensitivity analysis of the financial parameter estimates.

Design point	Factor 1 φ	Factor 2 κ	Factor 3 δ	Response (millions of ZAR)
1	-	-	-	888.4
2	+	-	-	687.5
3	-	+	-	613.1
4	+	+	-	640.0
5	-	-	+	788.0
6	+	-	+	625.6
7	-	+	+	604.7
8	+	+	+	534.8
Effect (millions of ZAR)	-50.8	-74.6	-34.5	

The effect of raising the mean of the emergency delivery function (φ) from 0.75 to 1.75 was ZAR 50.8 million, for the eight month period. The effect of raising the delivery cancellation penalty (δ) from 20% to 80% increased costs by ZAR 74.6 million. Raising the holding cost percentage (δ) from 7.75% to 8.75% increased costs by ZAR 34.5 million. Negotiating a reduced delivery cancellation fee has the potential to greatly reduce costs. Costs can also be reduced by negotiating a lower interest rate and a reduced fee for emergency coal delivered.

7.4 Experimental results

7.4.2 Using different output statistics

The aim of this experiment is to determine the effect of varying the choice of output statistic on the final approximation of decision variable values. To begin with, the final approximation of objective function values for the 5-th and 95-th percentile scenario are discussed, as shown in Table 7.1. The 5-th percentile scenario incurs a high level of shortage and EMD costs and the total cost is therefore 168.1% greater than the base case. The 95-th percentile scenario incurs a high level of holding and DCA costs, which results in a total cost 82.4% greater than the base case.

Final approximations for the decision variable values of the 5-th and 95-th percentile scenario are compared against the base case, as depicted in Figure 7.14. As expected, T_p has greater values for the 5-th percentile scenario than the base case. This results in high initial stockpile values so as to avoid incurring EMDs and shortages. For the 95-th percentile scenario, T_p does not however have lesser values for all the power stations. This phenomenon is explained later.

DCAs do not occur for the 5-th percentile scenario. EMDs for the 5-th percentile scenario are represented in Figure 7.15. For the 5-th percentile scenario, it makes sense that when stockpile levels are lower than expected, many EMDs come into effect. Moreover, the only three power stations not to incur deliveries are subject to over-delivery, namely Power stations B, H, and K. Stockpile levels for the final approximated values of the 5-th percentile scenario are represented in Figure 7.16. Power stations C, D, I, J, M, and N drop below the minimum stockpile level of five days and therefore incur shortages. The 5-th percentile is the only scenario, out of a total of 13, that incurs shortages.

There is a problem with having only eight intervals in the planning horizon: the stockpile level at the end of Month 8 might be undesirable. In Figure 7.16, for instance, the final stockpile level of Power station A is only slightly above L_p . On an operational planning level, planning eight months into the future is sufficient, granted that the simulation resolution is daily. However, this is not the case for the PEM. On a strategic planning level, simulating eight months into the future is plausible. But once again, the pitfall is that aggregating deliveries and burn to a monthly basis “flattens” variation in the system.

7.4 Experimental results




Figure 7.14: Final approximation of the decision variable values for the output statistic experiment.

7.4 Experimental results




Figure 7.15: Emergency deliveries for the final approximation of the 5-th percentile scenario.

7.4 Experimental results

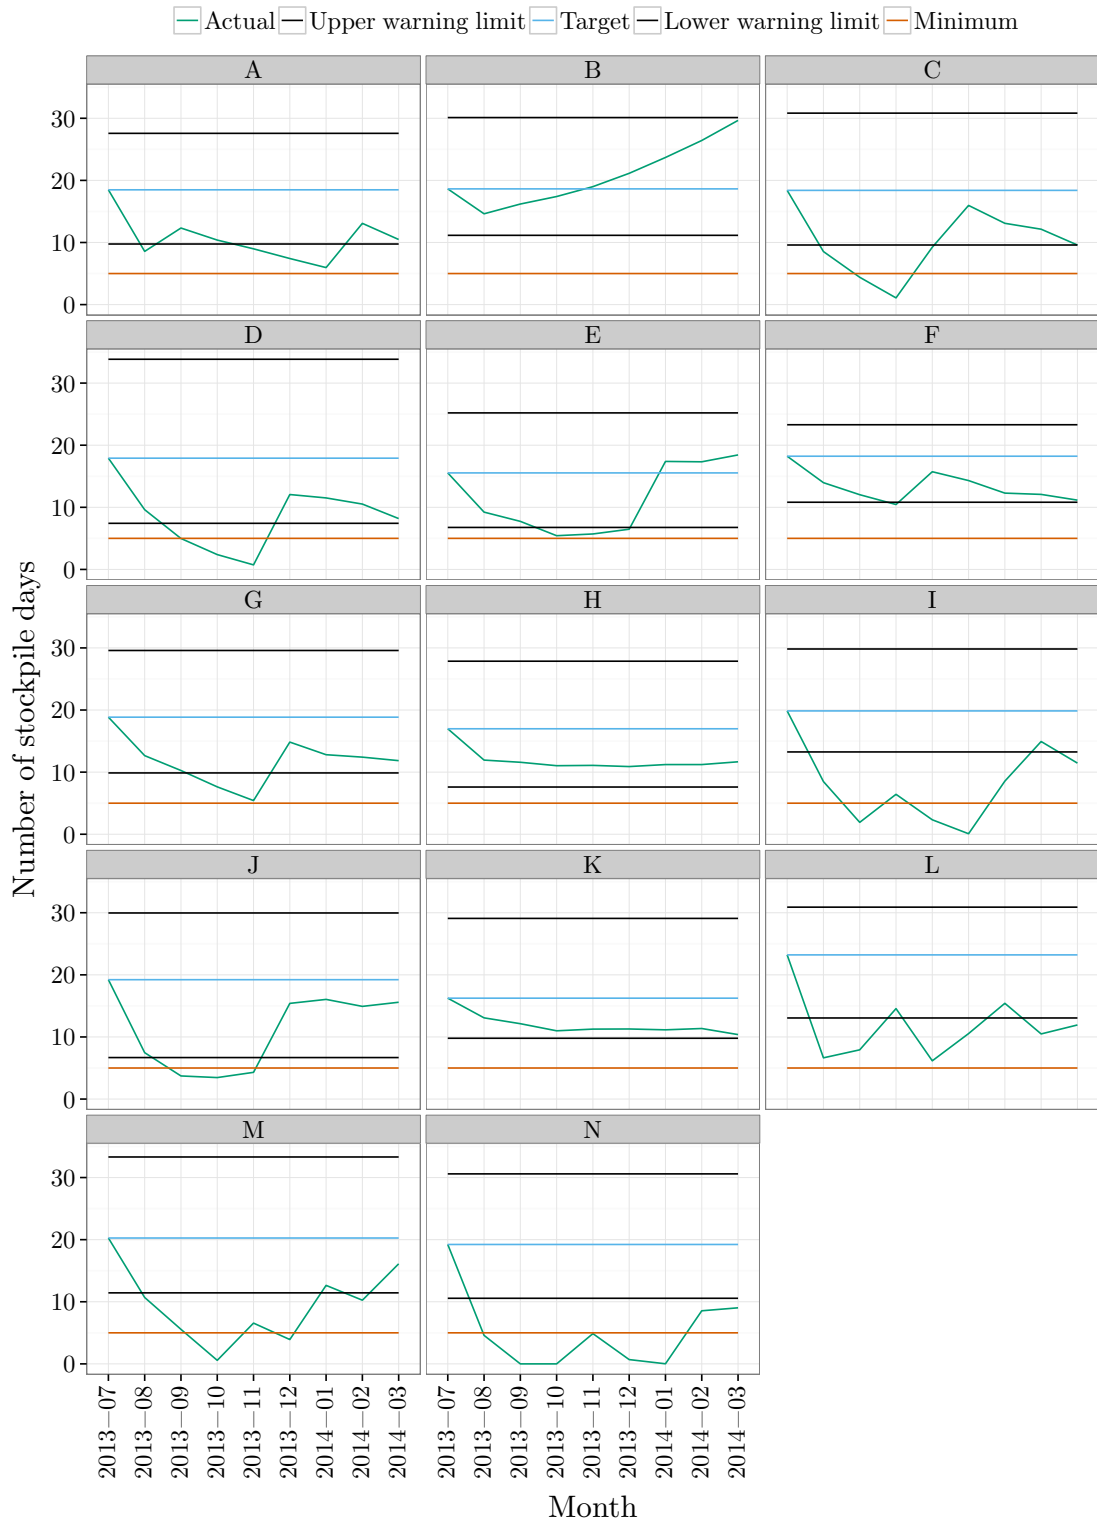


Figure 7.16: Stockpile policies and the changes in stockpile levels for the final approximation of the 5-th percentile scenario.

7.4 Experimental results

Still referring to Figure 7.16, coal stockpile replenishment through EMDs is discussed by way of example. At Power station D, the actual stockpile level ($A_{p,t}$) falls below L_p at month-end “2013-09”. An order of quantity equal to the difference between T_D and $A_{D,2}$ is made. The order arrives during month “2013-12”. Stockpile levels are not returned to the exact value of T_D , because variation occurred between order and delivery.

EMDs and DCAs for the 95-th percentile scenario are represented by Figures 7.17 and 7.18, respectively. For the 95-th percentile scenario, it makes sense that when stockpile levels are higher than estimated, only an EMD occurs for Power station I, since Power station I exhibits a high level of under-delivery. That said, the EMD is of a much smaller quantity compared to the base case (see Figure 7.12). Interestingly, Power station E incurs an EMD despite the fact it is subject to over-delivery. Upon analysis of Figure 7.19, it can be seen that a large DCA occurs at month-end “2013-09”. During the following month, the stockpile level then drops slightly below L_p and an EMD order is made at month-end.

As depicted in Figure 7.18, DCAs occur for the 95-th percentile scenario at Power stations B, C, D, E, and J. The discussion of “Why does T_p not have a lesser value for the 95-th percentile than the base case, for all the power stations?” is now resumed. Usually, T_p is driven by shortage and holding costs. In this case, EMDs drive T_p because DCAs reduce the stockpile levels by the difference between $A_{p,t}$ and T_p . In turn, this influences the cost of DCAs, since the greater the difference, the greater the cost. This results in the DCAs driving T_p “upwards”, against the holding costs driving T_p “downwards”.

Month-end stockpile levels for the final approximated values of the 95-th percentile scenario are represented in Figure 7.19. Decreasing coal stockpiles through DCAs is illustrated by way of example. At Power station B, the actual stockpile level ($A_{p,t}$) rises above L_p at month-end “2013-10”. An order of quantity equal to the difference between $A_{B,3}$ and T_B is made. The order arrives during month “2013-11”. Stockpile levels are not returned to the exact value of T_B , because variation occurred between order and delivery.

The results of the third and final experiment are discussed next.

7.4 Experimental results




Figure 7.17: Emergency deliveries for the final approximation of the 95-th percentile scenario.




Figure 7.18: Delivery cancellations for the final approximation of the 95-th percentile scenario.

7.4 Experimental results

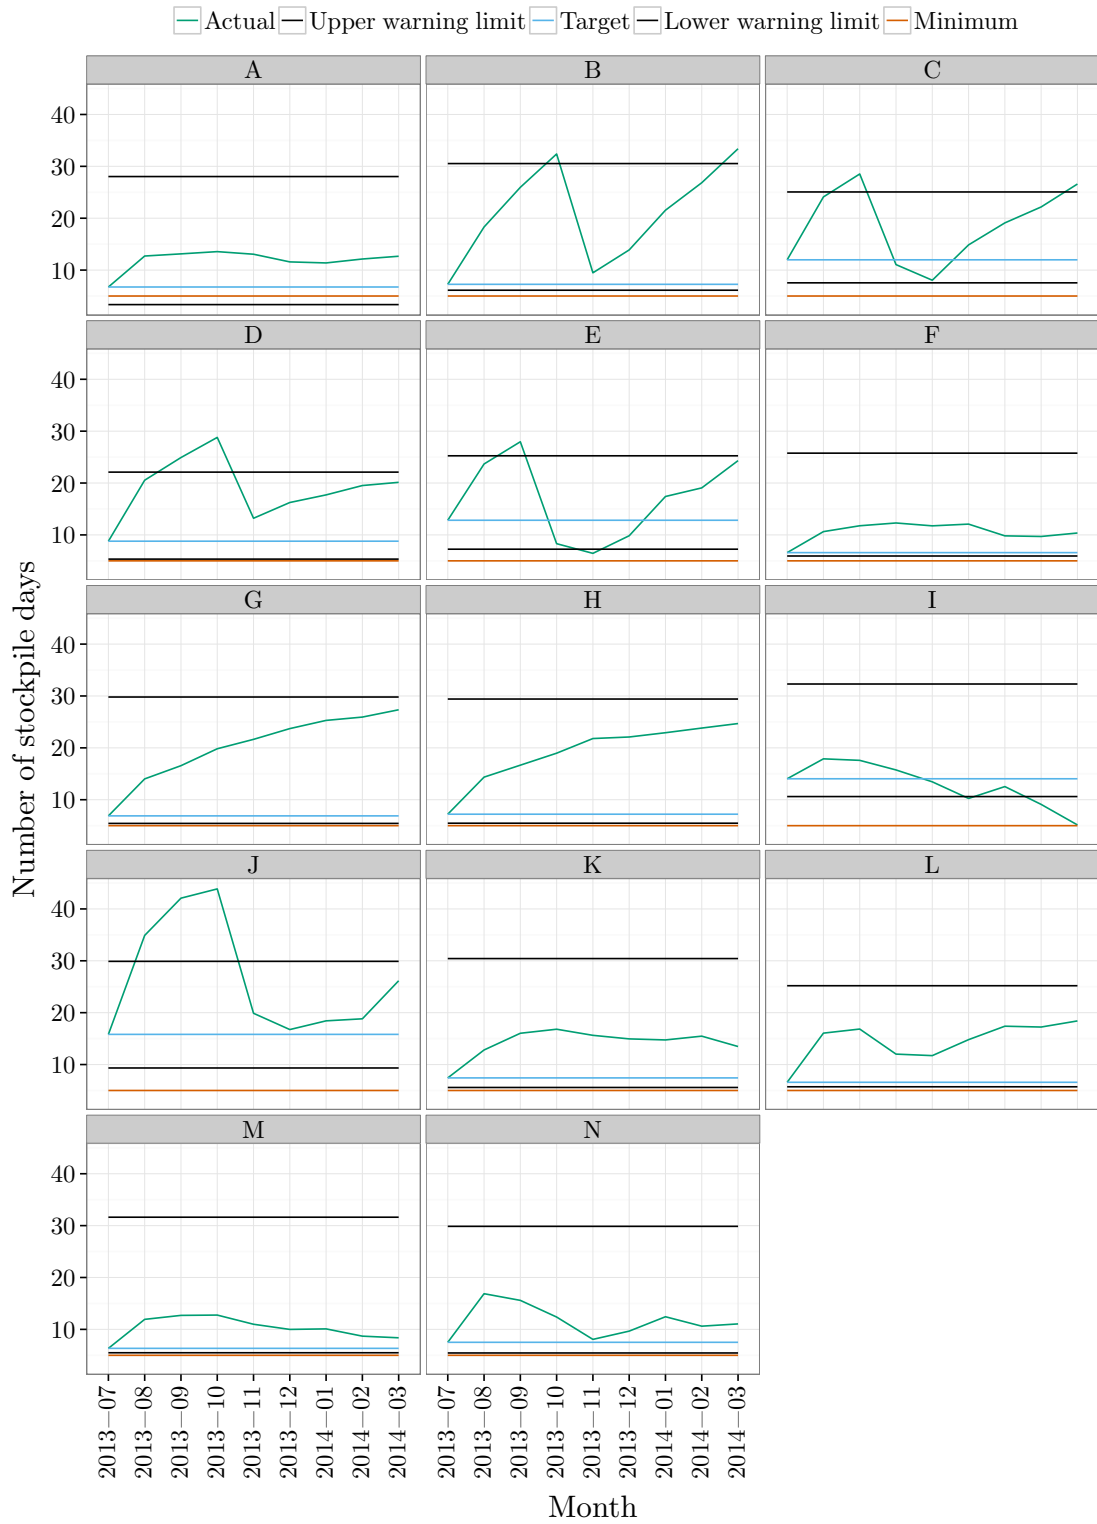


Figure 7.19: Stockpile policies and the changes in stockpile levels for the final approximation of the 95-th percentile scenario.

7.4 Experimental results

7.4.3 Varying the baseline deliveries

As detailed in Section 6.3, the experiment whereby BDs are varied is comprised of three scenarios:

1. Increasing BDs by 20% $\forall p \in \mathcal{P}, t \in \mathcal{T}$.
2. Decreasing BDs by 20% $\forall p \in \mathcal{P}, t \in \mathcal{T}$.
3. Randomly varying BDs, as detailed in Section 6.3.4.

The focus of this experiment is to analyse the effect of varying BDs on the final approximated decision variable values. To begin with, the final approximation of objective function values for the reduced, increased, and randomly varied baseline delivery scenarios are discussed, as shown in Table 7.1. These three scenarios had a small effect on costs: reduced BDs reduced the costs by 5.5%, increased BDs increased the costs by 2.7%, and randomly varying the BDs reduced the cost by 2.5%. It is important to note that in the case of increased BDs, the costs associated with an increased overall delivery quantity are not included in the model. Likewise, in the case of reduced BDs, the savings associated with a reduced overall delivery quantity are not included.

Final approximations for the decision variable values of all three scenarios, compared against the base case, are depicted in Figure 7.20. The cost values of the optimisation function cause the algorithm to avoid incurring EMDs and DCAs, at the expense of greater holding costs. In other words, inventory is rather carried than rushing EMDs and DCAs.

The reduced BD scenario is first considered. EMDs and DCAs for the reduced BD scenario are represented in Figures 7.21 and 7.22, respectively. EMDs and DCAs do not differ much from the base case (refer to Figures 7.12 and 7.13). In fact, exactly the same number of EMDs and DCAs occur, and moreover, they occur at the same power stations. The triggering of the EMDs through actual stockpile levels falling below L_p can be seen in Figure 7.23. Likewise, the single DCA is ordered when the actual stockpile level of Power station B rises above U_p .

7.4 Experimental results




Figure 7.20: Final approximation of the decision variable values for the baseline delivery experiment.

7.4 Experimental results




Figure 7.21: Emergency deliveries for the final approximation of the reduced baseline deliveries scenario.




Figure 7.22: Delivery cancellations for the final approximation of the reduced baseline deliveries scenario.

7.4 Experimental results

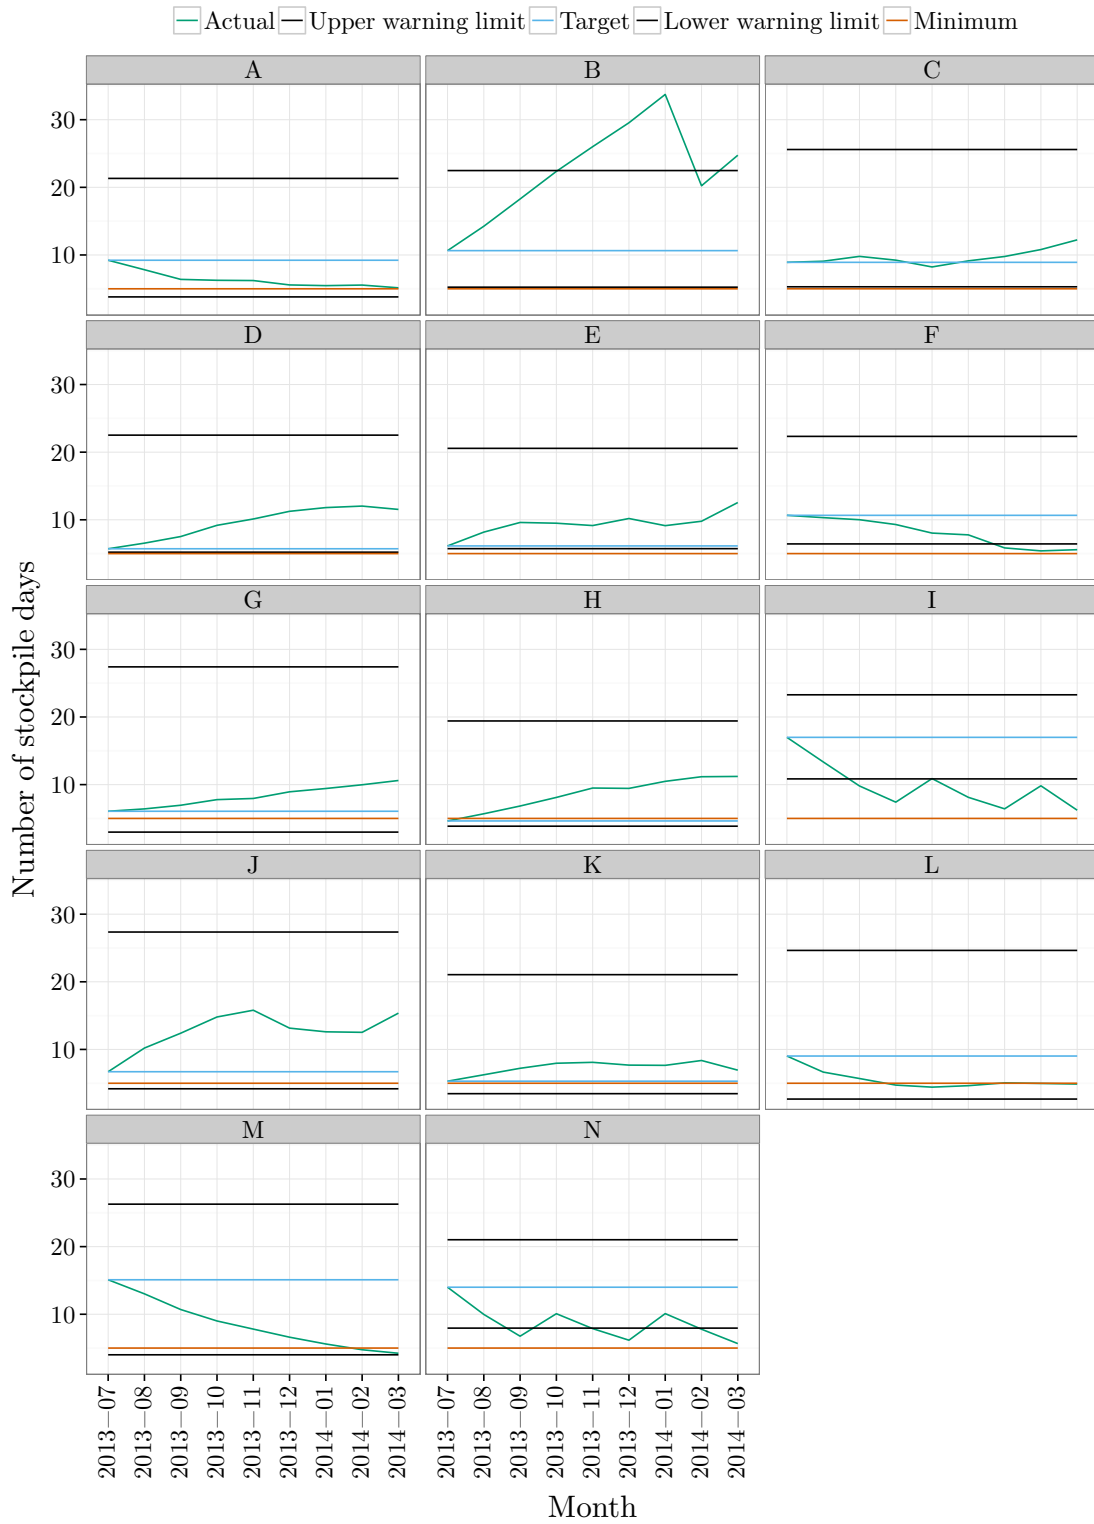


Figure 7.23: Stockpile policies and the changes in stockpile levels for the final approximation of the reduced baseline deliveries scenario.

7.4 Experimental results

Moving onto the scenario for increased BDs, no DCAs occur for the planning horizon. EMDs are illustrated in Figure 7.24. Only one EMD occurs. It seems to be that the model is attempting to avoid EMDs by setting a high value for T_p , which comes at the expense of increased holding costs. In the same manner, U_p is assigned large values so as not to incur DCAs, which also comes at the expense of increased holding costs. It is assumed that if a greater number of months are included in the planning horizon, EMDs and DCAs will occur more frequently in order to reduce holding costs, while not incurring shortage costs. This behaviour can be seen in Figure 7.25.

Finally, the scenario for randomly varied BDs is analysed. In saying that, cumulative variation in the stockpile levels, as a result of the modified BDs schedule, is investigated. The difference between simulated values of the randomly varied BDs and the coal burnt is illustrated in Figure 7.26. The randomly varied BDs introduce a greater amount of month-on-month variation into the system. 500 MC sample paths are represented.

An EMD occurs at Power stations I and N, as illustrated in Figure 7.27. Once again, only one DCA occurs and it occurs at Power station B, as depicted in Figure 7.28. As seen in Figure 7.29, stockpile levels seem to be “flatter”, when compared to the previous scenarios presented. With the exception of Power stations B, I, and N — which incur EMDs or DCAs — the randomly varied delivery scenario seems to reduced the effect of over-delivery and under-delivery, resulting in fewer EMDs and DCAs than the base case (refer to Figures 7.12 and 7.13).

This chapter is concluded in the following section.

7.4 Experimental results




Figure 7.24: Emergency deliveries for the final approximation of the increased baseline deliveries scenario.

7.4 Experimental results

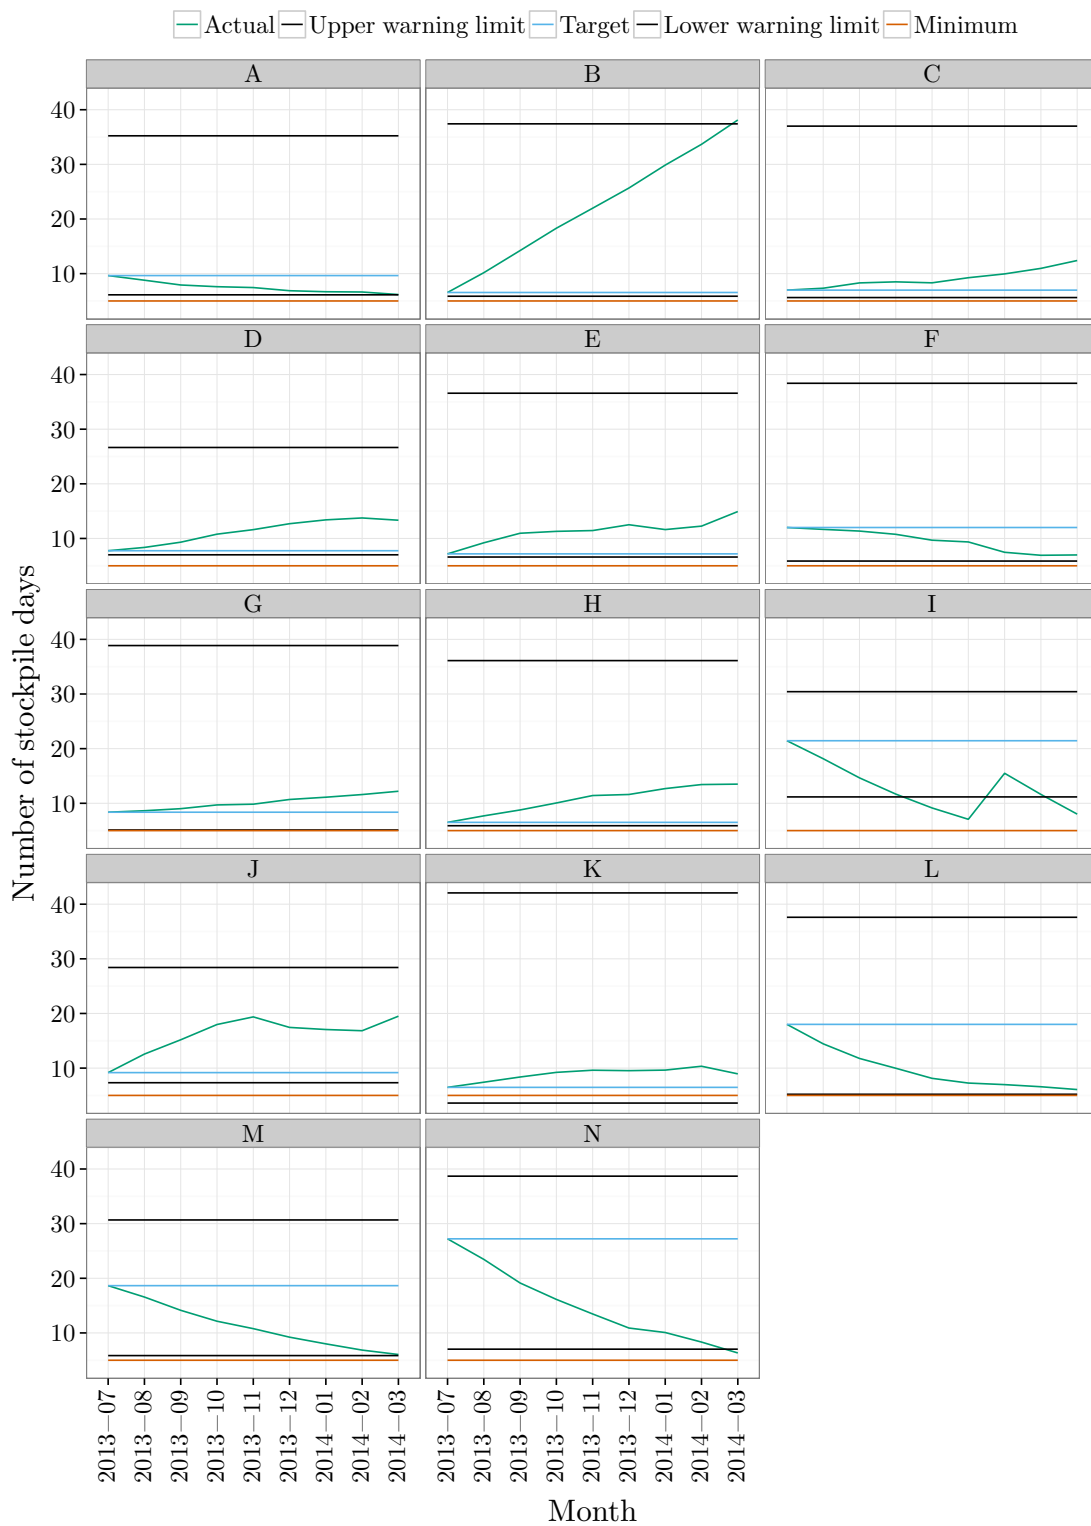


Figure 7.25: Stockpile policies and the changes in stockpile levels for the final approximation of the increased baseline deliveries scenario.

7.4 Experimental results

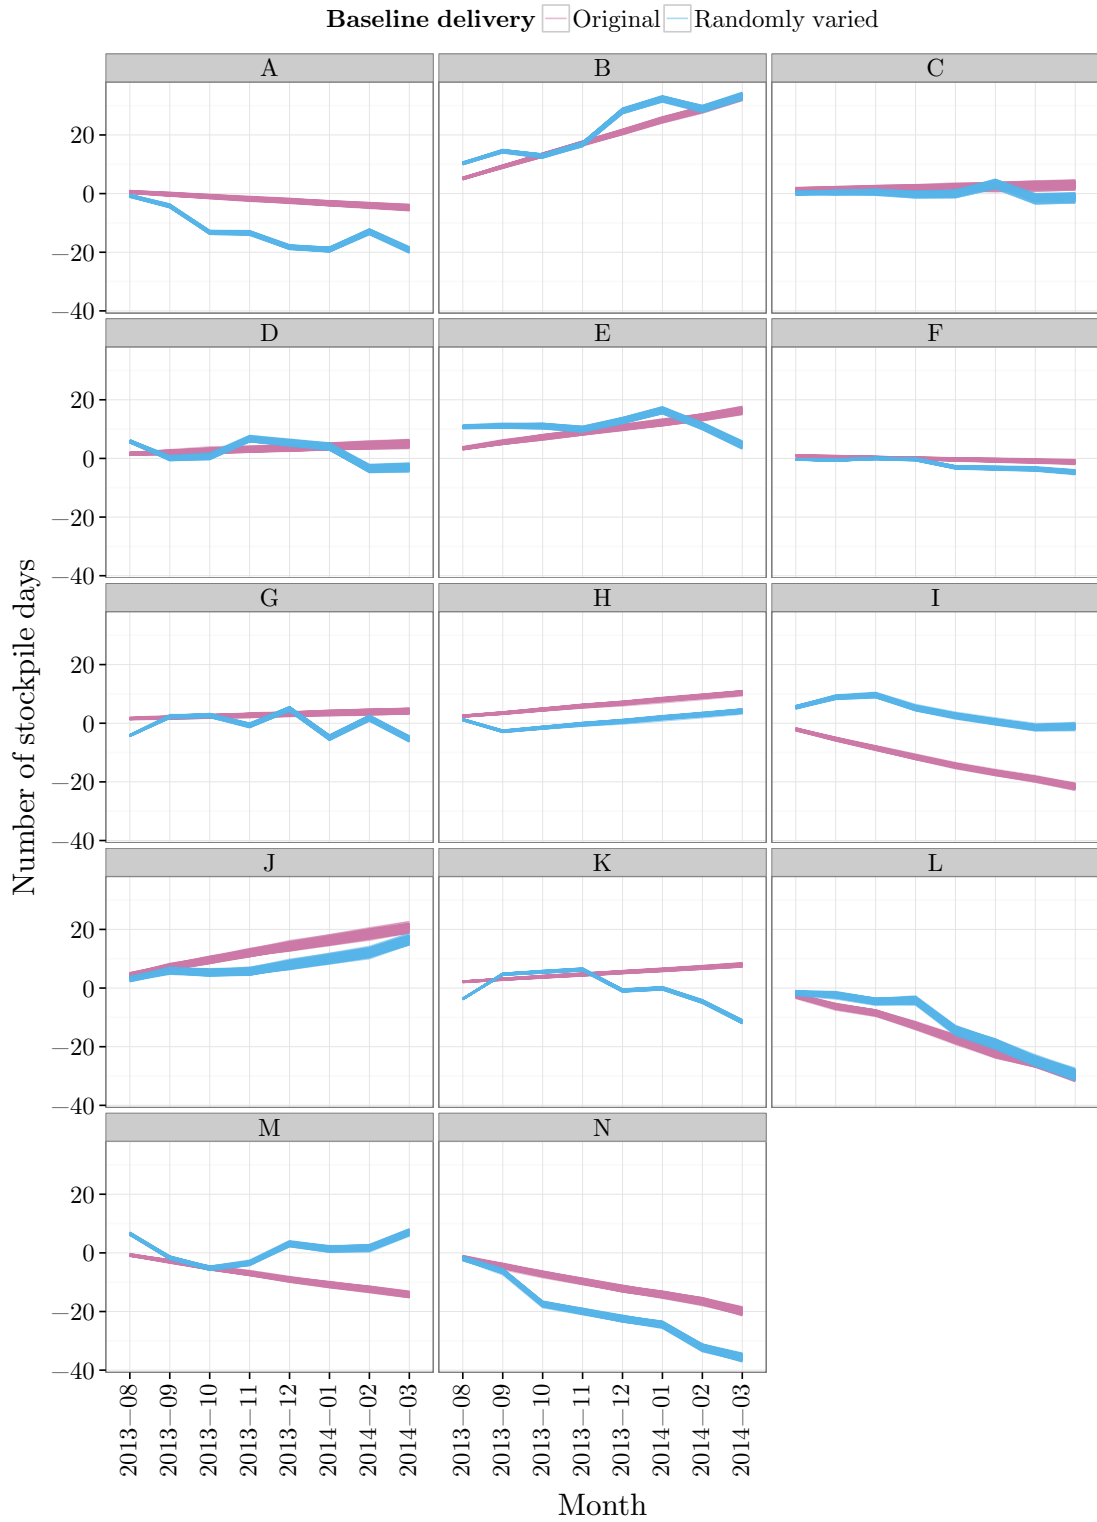


Figure 7.26: Comparison of the cumulative stockpile variation between the original and randomly varied baseline delivery.

7.4 Experimental results




Figure 7.27: Emergency deliveries for the final approximation of the randomly varied baseline deliveries scenario.




Figure 7.28: Delivery cancellations for the final approximation of the randomly varied baseline deliveries scenario.

7.4 Experimental results

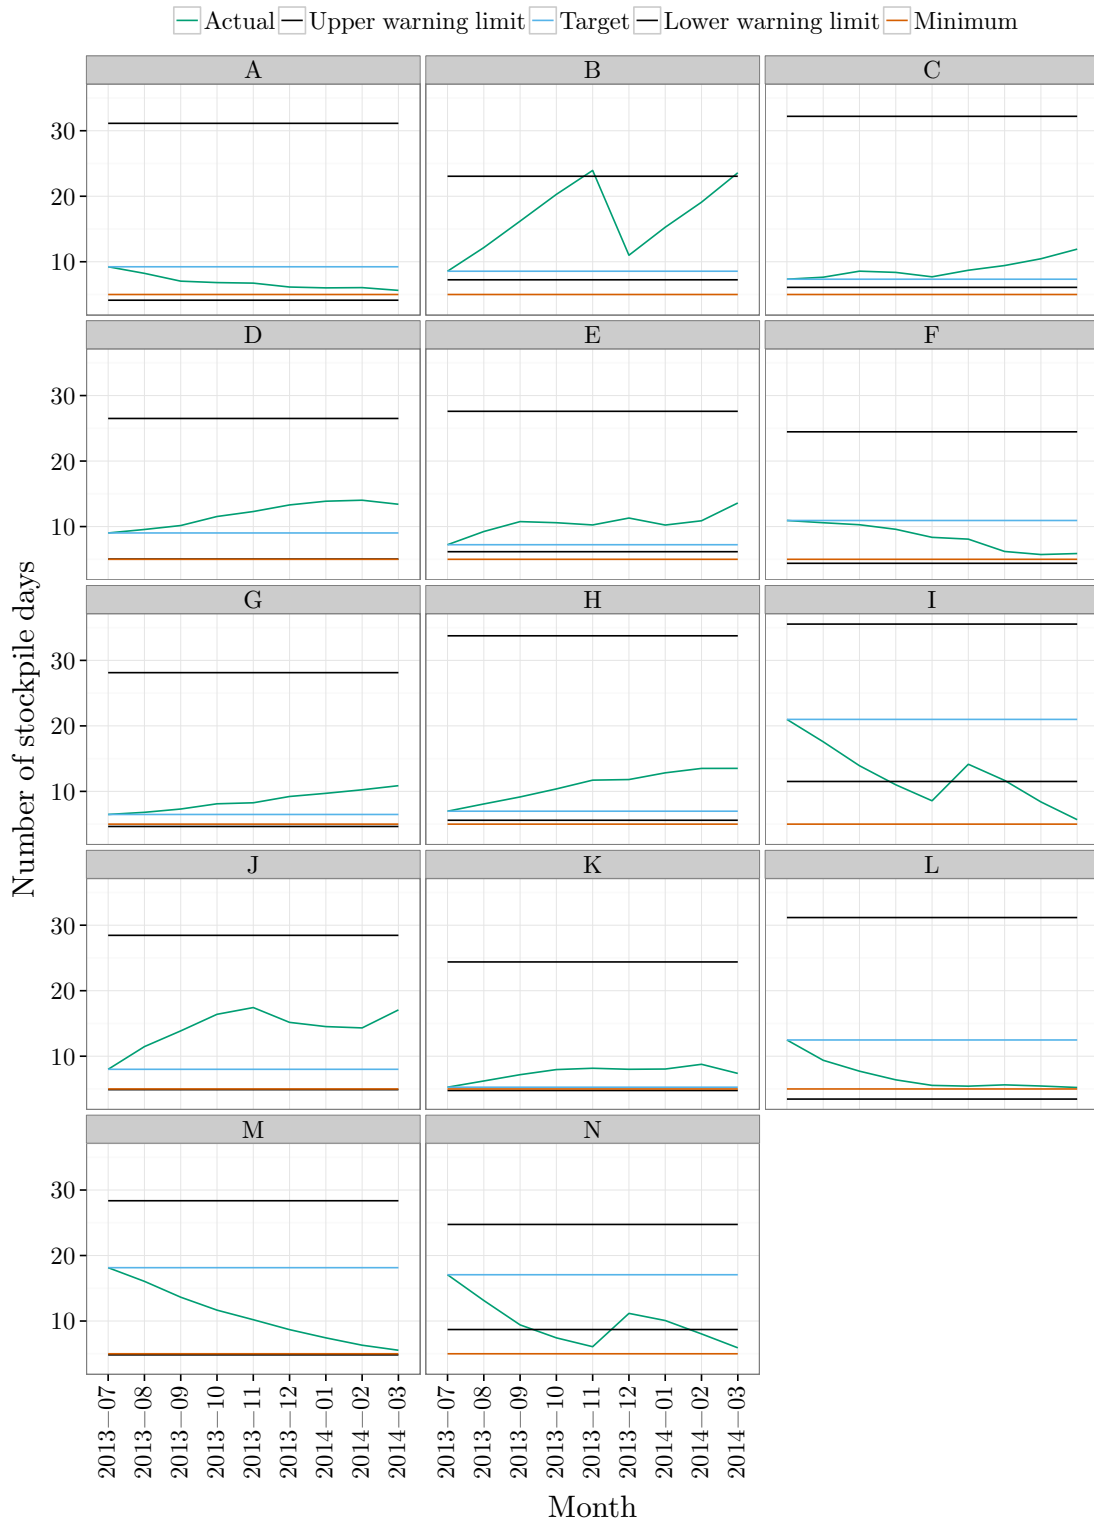


Figure 7.29: Stockpile policies and the changes in stockpile levels for the final approximation of the randomly varied baseline deliveries scenario.

7.5 Concluding remarks on Chapter 7

7.5 Concluding remarks on Chapter 7

In this chapter, the PEM and CEM were verified and validated. Limitations of the PEM were brought to light. It was shown that there is merit in the approach proposed in this study. The combined simulation and optimisation model provides prescriptive capability to coal stockpile management in the energy flow simulator, where none previously existed.

More specifically, it was shown that when large variation in stockpile levels occurs, EMDs and DCAs occur frequently. However, when a small amount of variation exists in the stockpile levels, EMDs and DCAs are avoided at the expense of greater holding costs.

The study is concluded in the following chapter.

Chapter 8

Conclusion

This chapter presents a summary of the research conducted and the findings of the study. Recommendations for future research are also provided.

8.1 Project summary

The primary and secondary aims of this research were introduced in **Chapter 1**. This study forms part of a greater research project for the energy flow simulator (EFS). Potentially, seven post-graduate students will work on the greater EFS project over the course of the next few years. Therefore, the secondary aim was to perform pioneering work so as to investigate potential areas for future research. This played an important role in what can be thought of as a “base” study.

The original principal aim of the study was to develop an optimiser for the EFS. To achieve this, **Chapter 2** was dedicated to detailing the decomposition of the database structure of the EFS. Parameters of the EFS were then investigated and scrutinised in search of potential decision variables for an optimisation model. It was concluded that global optimisation of the EFS was not possible due to its modular nature. The primary energy module (PEM) did, however, show potential for optimisation. The PEM was subsequently investigated. The renewed aim of the study was to provide the PEM with an optimiser. Limitations of the PEM — some of which apply in general to the EFS — were documented. These limitations influenced the modelling approach and the choice of a solution approach.

As the PEM is essentially a coal stockpile simulator, inventory and coal stockpile management techniques were studied in **Chapter 3**. The most important

8.2 Suggested future research

finding of this chapter was to propose using continuous review inventory policies in the optimiser. Due to the monthly resolution limitation of the PEM, standard continuous review policies for inventory problems were adapted and then used to model the lower and upper warning limits, as discussed in **Chapter 5**.

Simulation optimisation techniques were investigated in **Chapter 4**. The cross-entropy method (CEM) was chosen as the solution approach and studied further in **Chapter 5**. **Chapter 5** also detailed the approach followed to integrate the CEM into the PEM.

In **Chapter 6**, experimental design was outlined for a base case and three experiments. Results were analysed and conclusions drawn in **Chapter 7**. The application of the CEM to the PEM was achieved with reasonable success. A prescriptive capability was added to the PEM, which did not previously exist. By changing the resolution of the PEM to a daily basis, the combined simulation and optimisation model would have the potential to further optimise decision making. This forms part of the suggested future research presented in the following section.

The application of the CEM provided a robust optimisation model. In addition, the optimisation algorithm was coded to fit the size of the decision space and the planning horizon. In other words, the combined simulation and optimisation model automatically scales itself to the number of commissioned power stations and the specified start and end dates.

8.2 Suggested future research

As the secondary aim of this research was to provide pioneering work for the EFS project, this section plays an important role by providing suggestions for future research. Potential areas of research for the EFS, along with possible improvements to this work are presented.

It is important to note that the EFS was originally built in H2 which is a Java-based database engine. The original EFS is described as being “cumbersome, intricate and complex to operate, and essentially not usable” ([Eskom Holdings Limited, 2014c](#)). A simpler, more maintainable solution was required and this led to redevelopment of the EFS. At the time of writing, the new version of the simulator was under development solely in the R programming language (R),

8.2 Suggested future research

incorporating the same logic and concepts from the previous version. The new version allows code to be edited and rewritten, which will enable future research to be integrated in a straightforward manner. That said, the EFS will not have to be treated as a black box. The main areas for research are illustrated in Figure 8.1 and discussed in the following list:

This study. Other metaheuristics or simulation optimisation techniques can also be applied to the PEM, and compared against the CEM. However, the greatest improvement can be achieved by redeveloping the PEM to a daily resolution. This is where “Proposed research area 2” comes into play.

Proposed research area 2. The simulation resolution should be reduced to a daily basis. Hourly electricity forecasts are fed into the production planning module. However, the planned generation is aggregated to a monthly basis. In order to reduce the PEM to a daily resolution, the production planning

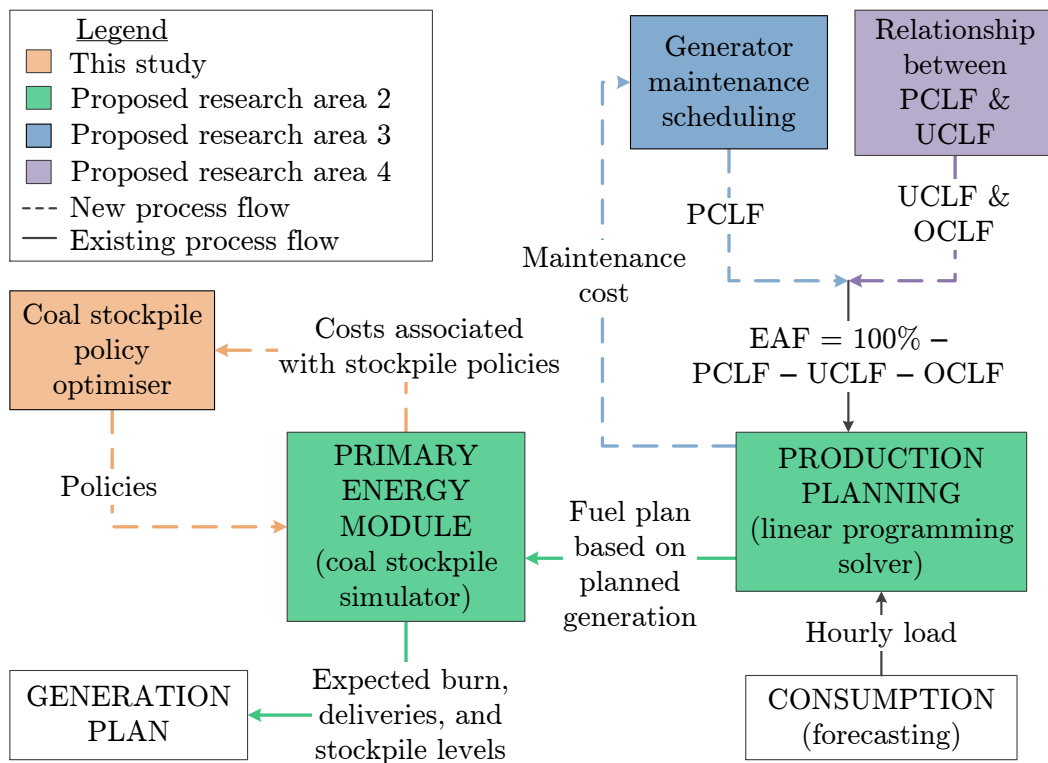


Figure 8.1: Suggested future research for the energy flow simulator.

8.2 Suggested future research

module has to be reformulated. Thereafter, the PEM will also need to be rebuilt. An approach similar to what was used in this study could then be applied to optimise stockpile policies.

Proposed research area 3. Development of a generator maintenance scheduling optimisation model. Planned maintenance is currently incorporated as a monthly percentage. By means of example, this unrealistic approach is explained. A generating unit of a power station does not operate at 80% capacity for the entire duration of a month. When undergoing maintenance, generating units “operate” at 0% capacity for the duration of the maintenance period. Otherwise, generating units can operate at full capacity. The generator maintenance scheduling problem is common to literature. It is essentially a highly constrained timetabling problem, making it highly combinatorial. Metaheuristics are popular solution approaches for the generator maintenance scheduling problem.

Proposed research area 4. Determining the relationship between planned and unplanned power station maintenance. Instead of assigning an estimated monthly downtime percentage for unplanned maintenance, formulate failure rate distributions for unplanned maintenance. This could be used to quantify the risk of not performing sufficient maintenance. A hypothesis is that a lag relationship exists between planned maintenance and unplanned maintenance.

Apart from the aforementioned suggested research areas, there are few more suggestions for future research:

- The modeller can develop a coal stockpile policy solution that incorporates season dependant policies. For example, a model that differentiates between Winter and non-Winter months.
- Coal stockpiles can be modelled more accurately by treating them as a coal blending problem. In a South African case study, [Conradie \(2011\)](#) optimised Sasol’s coal blending process. The use of such an approach to model Eskom’s coal blending process has the potential to provide further descriptive capability to the PEM.

8.3 Skills acquired

- Incorporating a start-up cost for the power stations in the linear program. Currently, the power stations are only assigned an operational cost.

8.3 Skills acquired

The researcher learnt how to successfully apply the CEM to an existing Monte Carlo simulation model, subject to its limitations. In doing so, the researcher learnt how to decompose a complex system into individual parts that can be optimised. Additionally, the researcher learnt how to apply an optimisation technique in the field of Operations Research to a real-world problem. An understanding of the field of simulation optimisation was also gained. The researcher gained an understanding of the South African electricity supply chain, and in particular coal stockpile management. Use of the document preparation system L^AT_EX was mastered. Finally, the researcher also learnt how to code in three programming languages: R, Python, and Structured Query Language (SQL). R was the main programming language used in this study. With the use of an R library, the database was read from and written to with the SQL. However, the researcher originally applied the CEM to test problems in Python. As time progressed, it was clear that a means of remotely triggering the PEM would be required. The industry partner to this study then developed an application programming interface in R that enabled remote triggering of the PEM. Thus, the researcher could no longer use Python and R was used instead.

8.4 Final remarks

Operations research can be defined as the art of providing bad answers to problems to which otherwise worse answers are given (Saaty, 2004).

The researcher believes that the optimisation model developed is able to optimise stockpile policies for the existing PEM, providing a prescriptive capability where one previously did not exist. The researcher believes that, through the proposed model and suggested future research areas for the energy flow simulator, the path has been paved for adding further prescriptive capability to the EFS in general.

8.4 Final remarks

Essentially, all models are wrong, but some are useful (George E. P. Box).

The researcher claims that the modelling approach used in this study adds to the existing decision support system, providing a simplification of reality that is useful to decision makers.

References

- ANDRADÓTTIR, S. (1998a). A review of simulation optimization techniques. In *Proceedings of the 1998 Winter simulation Conference*, 151–158. [43](#), [46](#)
- ANDRADÓTTIR, S. (1998b). Simulation optimization. In J. Banks *et al.*, eds., *Handbook of simulation*, 307–333, Wiley. [43](#), [46](#)
- ASIF, M. & MUNEER, T. (2007). Energy supply, its demand and security issues for developed and emerging economies. *Renewable and Sustainable Energy Reviews*, **11**, 1388–1413. [1](#)
- AZADIVAR, F. (1999). Simulation optimization methodologies. In *Proceedings of the 1999 Winter Simulation Conference*, 93–100. [40](#), [41](#), [43](#), [45](#), [46](#)
- BANKS, J. *et al.* (1998). *Handbook of simulation*. Wiley Online Library. [28](#), [41](#), [42](#), [43](#), [77](#), [82](#)
- BEKKER, J. (2012). *Applying the cross-entropy method in multi-objective optimisation of dynamic stochastic systems*. Ph.D. thesis, Stellenbosch University. [28](#), [40](#)
- BEKKER, J. (2013). Simulation of discrete-event stochastic processes, university lecture notes. [40](#)
- BHATTACHARYYA, S.C. & TIMILSINA, G.R. (2009). Energy demand models for policy formulation: A comparative study of energy demand models. Tech. rep., World Bank. [7](#)
- BOUSSAÏD, I., LEPAGNOT, J. & SIARRY, P. (2013). A survey on optimization meta-heuristics. *Information Sciences*, **237**, 82–117. [47](#)
- CARPENTER, A.M. (1999). Management of coal stockpiles. Tech. rep., IEA Coal Research. [31](#), [32](#), [33](#)

REFERENCES

- CARSON, Y. & MARIA, A. (1997). Simulation optimization: Methods and applications. In *Proceedings of the 1997 Winter Simulation Conference*, 118–126. [43](#), [45](#), [46](#), [47](#)
- CHAN, J.C. & KROESE, D.P. (2012). Improved cross-entropy method for estimation. *Statistics and Computing*, **22**, 1031–1040. [61](#)
- CHIKOBVU, D. & SIGAUKE, C. (2012). Regression-SARIMA modeling of daily peak electricity demand in South Africa. *Journal of Energy in South Africa*, **23**, 23–30. [7](#)
- CONRADIE, D.G. (2011). *Scheduling coal handling processes using metaheuristics*. Master's thesis, University of Pretoria. [7](#), [33](#), [123](#)
- CROWLEY, K. & JANSE VAN VUUREN, A. (2014). Eskom declares power emergency and asks South Africa to cut use, [Online], [Accessed 25 August 2014]. Available from www.bloomberg.com/news/2014-02-20/eskom-declares-power-emergency-and-asks-south-africa-to-cut-use.html. [3](#)
- DAHAL, K.P. & CHAKPITAK, N. (2007). Generator maintenance scheduling in power systems using metaheuristic-based hybrid approaches. *Electric Power Systems Research*, **77**, 771–779. [7](#)
- DAVIS, S. & DURBACH, I. (2010). Modelling household responses to energy efficiency interventions via system dynamics and survey data. *ORiON: The Journal of ORSSA*, **26**. [7](#)
- DE BOER, P.T., KROESE, D.P., MANNOR, S. & RUBINSTEIN, R.Y. (2005). A tutorial on the cross-entropy method. *Annals of operations research*, **134**, 19–67. [60](#), [65](#)
- D'SA, A. (2005). Integrated resource planning (IRP) and power sector reform in developing countries. *Energy Policy*, **33**, 1271–1285. [1](#), [34](#)
- DURBACH, I. & DAVIS, S. (2013). Decision support for selecting a shortlist of electricity-saving options: A modied SMAA approach. *ORiON: The Journal of ORSSA*, **28**, 99–116. [7](#)
- EBERHARD, A. (2011). The future of south african coal: market, investment, and policy challenges. Tech. rep., PESD Working Paper 100, Stanford. [2](#), [12](#), [20](#)
- EDIGER, V.Ş. & AKAR, S. (2007). ARIMA forecasting of primary energy demand by fuel in Turkey. *Energy Policy*, **35**, 1701–1708. [7](#)

REFERENCES

- EIBEN, A.E. & SMITH, J.E. (2003). *Introduction to evolutionary computing*. Springer. 42
- ESKOM HOLDINGS LIMITED (2003). Annual financial statement: Year-end 31 December 2003, [Online], [Accessed 25 July 2014]. Available from http://www.eskom.co.za/OurCompany/Investors/IntegratedReports/Documents/ESKOM_ANNUAL_REPORT_2003.pdf. 27
- ESKOM HOLDINGS LIMITED (2005). Annual financial statement: Year-end 31 March 2005, [Online], [Accessed 25 July 2014]. Available from http://www.eskom.co.za/OurCompany/Investors/IntegratedReports/Documents/ESKOM_ANNUAL_REPORT_2005.pdf. 27
- ESKOM HOLDINGS LIMITED (2007). Annual financial statement: Year-end 31 March 2007, [Online], [Accessed 25 July 2014]. Available from http://www.eskom.co.za/OurCompany/Investors/IntegratedReports/Documents/ESKOM_ANNUAL_REPORT_2005.pdf. 27
- ESKOM HOLDINGS LIMITED (2008). Annual financial statement: Year-end 31 March 2008, [Online], [Accessed 25 July 2014]. Available from http://financialresults.co.za/eskom-ar2008/ar_2008/downloads/eskom-ar2008.pdf. 27
- ESKOM HOLDINGS LIMITED (2009). Annual financial statement: Year-end 31 March 2009, [Online], [Accessed 25 July 2014]. Available from http://www.financialresults.co.za/eskom_ar2009/ar_2009/downloads/05-consolidated_fin_statements.pdf. 27
- ESKOM HOLDINGS LIMITED (2010). Annual financial statement: Year-end 31 March 2010, [Online], [Accessed 25 July 2014]. Available from http://financialresults.co.za/2010/eskom_ar2010/note_18.html. 27
- ESKOM HOLDINGS LIMITED (2011a). Annual financial statement: Year-end 31 March 2011, [Online], [Accessed 25 July 2014]. Available from http://financialresults.co.za/2011/eskom_ar2011/downloads/eskom-ar2011.pdf. 27
- ESKOM HOLDINGS LIMITED (2011b). Coal transport by road, [Online], [Accessed 9 September 2014]. Available from http://www.eskom.co.za/aboutelectricity/factsfigures/documents/co_0010coal-truckingrev41.pdf. 33

REFERENCES

- ESKOM HOLDINGS LIMITED (2012a). Annual financial statement: Year-end 31 March 2012, [Online], [Accessed 25 July 2014]. Available from http://financialresults.co.za/2012/eskom_ar2012/integrated-report/downloads/full-downloads/02_eskom_annual_financial_statements2012.pdf. 27
- ESKOM HOLDINGS LIMITED (2012b). Electricity Generation Options considered by Eskom, [Online], [Accessed 9 September 2014]. Available from <http://www.energy.gov.za/files/IEP/presentations/ElectricityGenerationOptionsConsidered-30March2012.pdf>. 9, 71
- ESKOM HOLDINGS LIMITED (2013a). Annual financial statement: Year-end 31 March 2013, [Online], [Accessed 25 July 2014]. Available from http://integratedreport.eskom.co.za/iir2014/integrated_report_2013/financials/fin-position.php. 27
- ESKOM HOLDINGS LIMITED (2013b). Future impact investigation with the energy flow simulator, confidential report. 6, 15
- ESKOM HOLDINGS LIMITED (2014a). About electricity and electricity technologies: Coal power, [Online], [Accessed 8 June 2014]. Available from <http://www.eskom.co.za/AboutElectricity/ElectricityTechnologies/Pages/Coal-Power.aspx>. 1
- ESKOM HOLDINGS LIMITED (2014b). Annual financial statement: Year-end 31 March 2014, [Online], [Accessed 25 July 2014]. Available from <http://integratedreport.eskom.co.za/financials/fin-pos.php>. 27
- ESKOM HOLDINGS LIMITED (2014c). Energy flow simulator: Visually interactive financial strategic decision support, confidential report. 29, 121
- ESKOM HOLDINGS LIMITED (2014d). Eskom power stations, [Online], [Accessed 9 September 2014]. Available from <http://www.eskom.co.za/Whatweredoing/ElectricityGeneration/PowerStations/Documents/EskomGenerationDivMapREV81.pdf>. 9
- FU, M.C. (2001). Simulation optimization. In *Proceedings of the 2001 Winter Simulation Conference*, 53–61. 41
- FU, M.C. (2002). Optimization for simulation: Theory vs. practice. *INFORMS Journal on Computing*, 14, 192–215. 41

REFERENCES

- FU, M.C., GLOVER, F. & APRIL, J. (2005). Simulation optimization: A review, new developments, and applications. In *Proceedings of the 2005 Winter Simulation Conference*, 83–95. [45](#), [47](#), [48](#)
- GE, J., ZHANG, C. & DU, M. (2011). Research on modeling and optimization methods of integrated resource planning. In *Proceedings of the 4th International Conference on Electric Utility Deregulation and Restructuring and Power Technologies (DRPT)*, 639–644, IEEE. [8](#)
- GENDREAU, M. & POTVIN, J.Y. (2005). Metaheuristics in combinatorial optimization. *Annals of Operations Research*, **140**, 189–213. [47](#), [48](#)
- HACHICHA, W., AMMERI, A., MASMOUDI, F. & CHACHOUB, H. (2010). A comprehensive literature classification of simulation optimisation methods. In *Proceedings of the 2010 Multi-objective Optimization and Goal Programming Conference*, 1–13. [43](#), [44](#), [47](#)
- HATTON, M. & BEKKER, J. (2014). Development of an optimiser for a simulator of an electric utility: Challenges and approach. In *Proceedings of the 43rd Annual Conference of the Operations Research Society of South Africa*, 18–26. [20](#)
- HENDERSON, S.G. & NELSON, B.L. (2006). *Handbooks in Operations Research and Management Science: Simulation*, vol. 13. Elsevier. [29](#)
- HERON, A. (1985). Financing electric power in developing countries. *International Atomic Energy Agency Bulletin*, **27**, 44–49. [7](#)
- HILLIER, F. & LIEBERMAN, G. (2001). *Introduction to Operations Research*. McGraw Hill, New York, ninth edn. [28](#), [37](#), [72](#)
- HO, Y.C. (1999). An explanation of ordinal optimization: Soft computing for hard problems. *information Sciences*, **113**, 169–192. [45](#)
- HOBBS, B.F. (1995). Optimization methods for electric utility resource planning. *European Journal of Operational Research*, **83**, 1–20. [1](#), [7](#), [8](#)
- JEFFREY, L. (2005). Characterization of the coal resources of South Africa. *Journal of the South African Institute of Mining and Metallurgy*, **105**, 95–102. [10](#)
- KESSIDES, I.N., BOGETIC, Z. & MAURER, L. (2007). Current and forthcoming issues in the South African electricity sector. Tech. rep., World Bank. [2](#)

REFERENCES

- KOTHARI, R.P. & KROESE, D.P. (2009). Optimal generation expansion planning via the cross-entropy method. In *Proceedings of the 2009 Winter Simulation Conference*, 1482–1491. [65](#)
- LAW, A.M., KELTON, W.D. & KELTON, W.D. (2000). *Simulation modeling and analysis*, vol. 3. McGraw-Hill, New York. [28](#), [50](#), [77](#), [78](#), [82](#)
- LIU, M., YANG, L. & GAN, D. (2005). A survey on agent based electricity market simulation. *Power System Technology*, **29**, 76–80. [1](#)
- LOULOU, R. & LABRIET, M. (2008). ETSAP-TIAM: the TIMES integrated assessment model Part I: Model structure. *Computational Management Science*, **5**, 7–40. [8](#)
- MACAL, C.M. & NORTH, M.J. (2007). Agent-based modeling and simulation: Desktop ABMS. In *Proceedings of the 2007 Winter Simulation Conference*, 95–106. [9](#)
- MERCADO, E.C. (2007). *Hands-on Inventory Management*. CRC Press. [34](#)
- MICALI, V. & HEUNIS, S. (2011). Coal stock pile simulation. In *Proceedings of the 8th Conference on the Industrial and Commercial Use of Energy*, 198–203, IEEE. [7](#), [18](#), [20](#), [21](#)
- MUELLER, T. (2014). The H2 database engine, [Online], [Accessed 13 February 2014]. Available from <http://www.h2database.com/html/main.html>. [19](#)
- MULLER, K. (2010). Stacking, reclaiming and blending effects, [Online], [Accessed 20 March 2014]. Available from http://www.thyssenkrupp-materialshandling.co.za/Images/332_StackingReclaiming_MechanicalTechnology_NovemberIssue_13122010.pdf. [33](#)
- NATIONAL ENERGY REGULATOR OF SOUTH AFRICA (2006). Electricity supply statistics for South Africa, [Online], [Accessed 14 April 2014]. Available from <http://www.nersa.org.za/Admin/Document/Editor/file/News> [13](#)
- OKABE, M. & ITO, K. (2014). Color Universal Design: How to make figures and presentations that are friendly to color-blind people, [Online], [Accessed 22 September 2014]. Available from <http://jfly.iam.u-tokyo.ac.jp/color/>. [88](#)
- PUTERMAN, M.L. (2009). *Markov decision processes: discrete stochastic dynamic programming*, vol. 414. John Wiley & Sons. [39](#)

REFERENCES

- R CORE TEAM (2014). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. [19](#)
- RAMJITH, P. (2014). Coal Stockpile Simulator, confidential report: Enerweb. [6](#)
- ROSEN, S.L., HARMONOSKY, C.M. & TRABAND, M.T. (2008). Optimization of systems with multiple performance measures via simulation: Survey and recommendations. *Computers & Industrial Engineering*, **54**, 327–339. [40](#)
- RUBINSTEIN, R. (1999). The cross-entropy method for combinatorial and continuous optimization. *Methodology and computing in applied probability*, **1**, 127–190. [45](#), [61](#), [63](#)
- RUBINSTEIN, R.Y. (1997). Optimization of computer simulation models with rare events. *European Journal of Operational Research*, **99**, 89–112. [61](#), [65](#)
- RUBINSTEIN, R.Y. & KROESE, D.P. (2004). *The cross-entropy method: a unified approach to combinatorial optimization, Monte-Carlo simulation and machine learning*. Springer. [48](#), [60](#), [63](#), [65](#)
- RUBINSTEIN, R.Y. & KROESE, D.P. (2011). *Simulation and the Monte Carlo method*. John Wiley & Sons. [60](#), [63](#), [64](#), [65](#), [66](#)
- SAATY, T.L. (2004). *Mathematical methods of operations research*. Courier Dover Publications. [124](#)
- SCHLÜNZ, E. (2011). *Decision support for generator maintenance scheduling in the energy sector*. Master's thesis, Stellenbosch University. [7](#)
- SCHLÜNZ, E. & VAN VUUREN, J. (2012). The application of a computerised decision support system for generator maintenance scheduling: a South African case study. *South African Journal of Industrial Engineering*, **23**, 169–179. [57](#)
- SCHRAMM, G. (1993). Issues and problems in the power sectors of developing countries. *Energy Policy*, **21**, 735–747. [7](#)
- SCOTT, T.J. & READ, E.G. (1996). Modelling hydro reservoir operation in a deregulated electricity market. *International Transactions in Operational Research*, **3**, 243–253. [7](#)
- SEEBREGTS, A.J., GOLDSTEIN, G.A. & SMEKENS, K. (2002). Energy/environmental modeling with the MARKAL family of models. In *Operations Research Proceedings 2001*, 75–82, Springer. [8](#)

REFERENCES

- SHERBROOKE, C.C. (2004). *Optimal inventory modeling of systems: multi-echelon techniques*. Springer. 34, 35
- SINHA, A., LAHIRI, R., BYABORTTA, S., CHOWDHURY, S., CHOWDHURY, S. & CRUSLEY, P. (2008). Coal Management Module (CMM) for power plants. In *Proceedings of the 2008 International Universities Power Engineering Conference*, 1–7, IEEE. 33
- SIRIKUM, J., TECHANITISAWAD, A. & KACHITVICHYANUKUL, V. (2007). A new efficient GA-benders' decomposition method: For power generation expansion planning with emission controls. *Power Systems, IEEE Transactions on*, **22**, 1092–1100. 7
- SUGANTHI, L. & SAMUEL, A.A. (2012). Energy models for demand forecasting: A review. *Renewable and Sustainable Energy Reviews*, **16**, 1223–1240. 7
- SWISHER, J.R., JACOBSON, S.H. & YÜCESAN, E. (2003). Discrete-event simulation optimization using ranking, selection, and multiple comparison procedures: A survey. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, **13**, 134–154. 45
- TEKIN, E. & SABUNCUOGLU, I. (2004). Simulation optimization: A comprehensive review on theory and applications. *IIE Transactions*, **36**, 1067–1081. 43, 45, 47
- TEKINER, H., COIT, D.W. & FELDER, F.A. (2010). Multi-period multi-objective electricity generation expansion planning problem with Monte-Carlo simulation. *Electric Power Systems Research*, **80**, 1394–1405. 7
- THE DEPARTMENT OF ENERGY (2013). Integrated resource plan for electricity (IRP) 2010-2030, [Online], [Accessed 21 February 2014]. Available from http://www.doe-irp.co.za/content/IRP2010_updatea.pdf. 14, 70, 71
- THE ENGINEERING TOOLBOX (2014). Power Plant Performance Factors, [Online], [Accessed 21 May 2014]. Available from http://www.engineeringtoolbox.com/power-plant-efficiency-d_960.html. 22
- THE SOUTH AFRICAN RESERVE BANK (2014). Current market rates, [Online], [Accessed 16 October 2014]. Available from <https://www.resbank.co.za/Research/Rates/Pages/CurrentMarketRates.aspx>. 72
- THE WORLD BANK (2014). Lending interest rate percentage, [Online], [Accessed 16 October 2014]. Available from <http://data.worldbank.org/indicator/FR.INR.LEND/countries>. 72

REFERENCES

- TRADING ECONOMICS (2014). Africa Government Bond 10Y, [Online], [Accessed 16 October 2014]. Available from <http://www.tradingeconomics.com/south-africa/government-bond-yield>. 72
- VAN HARMELEN, G. (2014a). Introduction to the energy flow simulator, confidential report: Enerweb. 6, 17, 21
- VAN HARMELEN, G. (2014b). Utility Analytics Business Area Manager at Enerweb. Personal communication, contactable at gerard.van.harmelen@gmail.com. 4, 19, 20, 21
- VENTOSA, M., BAILLO, A., RAMOS, A. & RIVIER, M. (2005). Electricity market modeling trends. *Energy policy*, **33**, 897–913. 8
- WINSTON, W. (2004). *Operations Research: Applications and Algorithms*. Thomson Brooks/Cole. 28, 34, 36, 37, 38, 50, 71, 72
- WINSTON, W. & VENKATARAMANAN, M. (2002). *Introduction to Mathematical Programming: Applications and Algorithms*. Thomson Brooks/Cole. 36, 39
- WOOD, A., WOLLENBERG, B. & SHEBLÉ, G. (2013). *Power Generation, Operation and Control*. Wiley. 7, 56
- ZHOU, Z., CHAN, W.K.V. & CHOW, J.H. (2007). Agent-based simulation of electricity markets: A survey of tools. *Artificial Intelligence Review*, **28**, 305–342. 8
- ZHU, S., WANG, J. & ZHAO, W. (2011). A seasonal hybrid procedure for electricity demand forecasting in China. *Applied Energy*, **88**, 3807–3815. 7
- ZIPKIN, P.H. (2000). *Foundations of inventory management*. McGraw-Hill. 34

Appendix A

Parameters of the energy flow simulator

The types of parameters are listed on the following page in Table [A.1](#). In total, 260 types of parameters are tabulated. Parameters are grouped to form the hierarchical structure of the energy flow simulator. As is common to databases, each parameter is assigned a unit type. Furthermore, each parameter is either designated as an input to the energy flow simulator or as an output.

Table A.1: Parameters in the EFS.

	Parameter group	Parameter	Data type	Unit	Parameter type
1	Weather Stations	Rainfall from SAWS	Number	MM	Input
2	Weather Stations	Heating Degree Days - base 18	Number	DEGC	Input
3	Weather Stations	Cooling Degree Days - base 18	Number	DEGC	Input
4	Weather Stations	Average Temperature from SAWS	Number	DEGC	Input
5	Weather Stations	Rainfall from SAWS	Num_Array	MM	Denorm
6	Weather Stations	Heating Degree Days - base 18	Num_Array	DEGC	Denorm
7	Weather Stations	Cooling Degree Days - base 18	Num_Array	DEGC	Denorm
8	Weather Stations	Average Temperature from SAWS	Num_Array	DEGC	Denorm
9	Weather Stations	Solar Radiation - Energy	Number	KWH	Input
10	Weather Stations	Solar Radiation - Energy	Num_Array	KWH	Denorm
11	Weather Stations	Geographical area of node	GIS_Location_Array	DD	Input
12	Network Nodes	Lat / long coordinates (X)	Number	DD	Input
13	Network Nodes	Lat / long coordinates (Y)	Number	DD	Input
14	Network Nodes	Deficit	Number	GWH	Input
15	Area Sector Loads	Load in Sector and Area	Number	MW	Input
16	Area Sector Loads	High level economic sector	SECTOR	PU	Input
17	Powerstations - General	Geographical area of node	GIS_Location_Array	DD	Input
18	Powerstations - General	UCLF Scale	Number	PU	Input
19	Powerstations - General	Commision date	Text	PU	Input
20	Powerstations - General	Min Load Factor	Number	PU	Input
21	Powerstations - General	DeCommision date	Text	PU	Input
22	Powerstations - General	UCLF Planned	Number	PU	Input
23	Powerstations - General	Hourly PCLF Planned	Number	PU	Input
24	Powerstations - General	Hourly UCLF Planned	Number	PU	Input
25	Powerstations - General	Reserve Price Step	Number	ZAR/MW	Input
26	Powerstations - General	Hourly Energy Scheduled	Number	GWH	Input
27	Powerstations - General	Reserve Margin	Number	MW	Input

Continued on the next page

	Parameter group	Parameter	Data type	Unit	Parameter type
28	Powerstations - General	Reserve Cost	Number	ZAR/MW	Input
29	Powerstations - General	Additional Load TakeUp	Number	PU	Input
30	Powerstations - General	Cost of Supply	Number	ZAR/MW	Input
31	Powerstations - General	Energy Sent Out Actual	Number	GWH	Input
32	Powerstations - General	Station Capacity	Number	MW	Input
33	Powerstations - General	Station Capacity	Number	MW	Input
34	Powerstations - General	Energy Available	Number	GWH	Result
35	Powerstations - General	Energy Pickup	Number	GWH	Result
36	Powerstations - General	Energy Scheduled	Number	GWH	Result
37	Powerstations - General	UCLF	Number	MW	Result
38	Powerstations - General	PCLF	Number	MW	Result
39	Powerstations - General	PCLF	Number	PU	Result
40	Powerstations - General	OCLE	Number	PU	Result
41	Powerstations - General	EAF	Number	PU	Result
42	Powerstations - General	Cycle Efficiency	Number	PU	Input
43	Powerstations - General	PCLF Actual	Number	PU	Input
44	Powerstations - General	UCLF Actual	Number	PU	Input
45	Powerstations - General	UCLF - Shift	Number	PU	Input
46	Powerstations - General	UCLF - Shape	Number	PU	Input
47	Powerstations - General	UCLF - Location	Number	PU	Input
48	Powerstations - General	PCLF	Number	PU	Input
49	Powerstations - General	PCLF - Percent Stdev	Number	PU	Input
50	Powerstations - General	OCLE	Number	PU	Input
51	Powerstations - General	Max Generator LF	Number	PU	Input
52	Powerstations - General	Energy Sent Out Budget	Number	GWH	Input
53	Powerstations - General	UCLF - Family	Number	PU	Input
54	Powerstations - General	Powerstation Type	Text	PU	Input
55	Powerstations - General	Heat Rate	Number	PU	Input
56	Powerstations - General	Initial Stock	Number	PU	Input

Continued on the next page

	Parameter group	Parameter	Data type	Unit	Parameter type
57	Powerstations - General	SDB - Availability	Number	PU	Input
58	Powerstations - General	SDB Load factor	Number	PU	Input
59	Powerstations - General	Stockpile Volume Actuals	Number	PU	Input
60	Powerstations - General	Stockpile Days Actuals	Number	PU	Input
61	Powerstations - General	Calorific Value Actuals	Number	PU	Input
62	Powerstations - General	Nett Delivery Actual	Number	PU	Input
63	Powerstations - General	Coal Burn	Number	PU	Input
64	Powerstations - General	Standard Daily Burn	Number	PU	Input
65	Powerstations - General	Min Alarm Level	Number	PU	Input
66	Powerstations - General	Max Alarm Level	Number	PU	Input
67	Powerstations - General	Cost Std Dev	Number	PU	Input
68	Powerstations - General	Calorific Value	Number	PU	Input
69	Powerstations - General	Calorific Value Stdev	Number	PU	Input
70	Powerstations - General	Nett Coal Reliability	Number	PU	Input
71	Powerstations - General	Nett Coal Reliability Std Dev	Number	PU	Input
72	Powerstations - General	Nett Coal Delivery	Number	PU	Input
73	Powerstations - General	Max Mine Production	Number	PU	Input
74	Powerstations - General	Stockpile Volume	Number	PU	Result
75	Powerstations - General	Stock Pile Days	Number	PU	Result
76	Powerstations - General	Coal Delivered	Number	PU	Result
77	Powerstations - General	Calorific Value	Number	PU	Result
78	Powerstations - General	Coal Burn	Number	PU	Result
79	Powerstations - Coal	Geographical area of node	GIS_Location_Array	DD	Input
80	Powerstations - Coal	Load Takeup	Number	PU	Input
81	Powerstations - Coal	PCLF Std Dev	Number	PU	Input
82	Powerstations - Coal	UCLF Family	Number	PU	Input
83	Powerstations - Coal	UCLF Location	Number	PU	Input
84	Powerstations - Coal	UCLF Shape	Number	PU	Input
85	Powerstations - Coal	UCLF Shift	Number	PU	Input

Continued on the next page

	Parameter group	Parameter	Data type	Unit	Parameter type
86	Powerstations - Coal	UCLF Planned	Number	PU	Input
87	Powerstations - Coal	Stockpile Days Actuals - Days	Number	PU	Input
88	Powerstations - Coal	PCLF Actual	Number	PU	Input
89	Powerstations - Coal	Cost of Supply - mean	Number	PU	Input
90	Powerstations - Coal	Cost of Supply Stdev	Number	PU	Input
91	Powerstations - Coal	Min Stockpile Warning - Days	Number	PU	Input
92	Powerstations - Coal	Max Stockpile Warning - Days	Number	PU	Input
93	Powerstations - Coal	UCLF Scale	Number	PU	Input
94	Powerstations - Coal	Commision date	Text	PU	Input
95	Powerstations - Coal	Min Load Factor	Number	PU	Input
96	Powerstations - Coal	DeCommision date	Text	PU	Input
97	Powerstations - Coal	Hourly UCLF Planned	Number	PU	Input
98	Powerstations - Coal	Hourly PCLF Planned	Number	PU	Input
99	Powerstations - Coal	Reserve Price Step	Number	ZAR/MW	Input
100	Powerstations - Coal	Hourly Energy Scheduled	Number	GWH	Input
101	Powerstations - Coal	Reserve Margin	Number	MW	Input
102	Powerstations - Coal	Reserve Cost	Number	ZAR/MW	Input
103	Powerstations - Coal	CV	Number	MJ/KG	Input
104	Powerstations - Coal	CV Stdev	Number	PU	Input
105	Powerstations - Coal	Nett Coal Delivery	Number	KTONS	Input
106	Powerstations - Coal	Nett Coal Reliability	Number	KTONS	Input
107	Powerstations - Coal	CV Actual	Number	MJ/KG	Input
108	Powerstations - Coal	Nett Coal Reliability Std Dev	Number	KTONS	Input
109	Powerstations - Coal	Heat Rate	Number	MJ/KG	Input
110	Powerstations - Coal	Cost of Supply	Number	ZAR/MW	Input
111	Powerstations - Coal	Energy Sent Out	Number	GWH	Input
112	Powerstations - Coal	Capacity	Number	MW	Input
113	Powerstations - Coal	Capacity	Number	MW	Input
114	Powerstations - Coal	Initial Stock	Number	KTONS	Input

Continued on the next page

	Parameter group	Parameter	Data type	Unit	Parameter type
115	Powerstations - Coal	Stockpile Volume Actuals	Number	KTONS	Input
116	Powerstations - Coal	Nett Delivery Actual	Number	KTONS	Input
117	Powerstations - Coal	Max Stockpile Volume	Number	KTONS	Input
118	Powerstations - Coal	Stock Pile Days - Days	Number	PU	Result
119	Powerstations - Coal	Calorific Value	Number	MJ/KG	Result
120	Powerstations - Coal	Stockpile Volume	Number	KTONS	Result
121	Powerstations - Coal	Coal Burn	Number	KTONS	Input
122	Powerstations - Coal	Coal Delivered	Number	KTONS	Result
123	Powerstations - Coal	Coal Burn	Number	KTONS	Result
124	Powerstations - Coal	Nett Coal Delivery	Number	KTONS	Input
125	Powerstations - Coal	Coal Delivered	Number	KTONS	Result
126	Powerstations - Coal	Standard Daily Burn	Number	KTONS	Input
127	Powerstations - Coal	Min Stockpile Level	Number	KTONS	Input
128	Powerstations - Coal	Energy Sent Out Actual	Number	GWH	Input
129	Powerstations - Coal	Energy Available	Number	GWH	Result
130	Powerstations - Coal	Energy Pickup	Number	GWH	Result
131	Powerstations - Coal	Energy Scheduled	Number	GWH	Result
132	Powerstations - Coal	Max Load Factor	Number	PU	Input
133	Powerstations - Coal	OCLF Mean	Number	PU	Input
134	Powerstations - Coal	PCLF Mean	Number	PU	Input
135	Powerstations - Coal	SDB - Availability	Number	PU	Input
136	Powerstations - Coal	SDB Load factor	Number	PU	Input
137	Powerstations - Coal	UCLF Actual	Number	PU	Input
138	Powerstations - Coal	Max Mine Production	Number	PU	Input
139	Powerstations - Coal	UCLF	Number	PU	Result
140	Powerstations - Coal	PCLF	Number	PU	Result
141	Powerstations - Coal	PCLF	Number	PU	Result
142	Powerstations - Coal	OCLF	Number	PU	Result
143	Powerstations - Coal	EAF	Number	PU	Result

Continued on the next page

	Parameter group	Parameter	Data type	Unit	Parameter type
144	Powerstations - Coal	Coal Station - Additional Load TakeUp	Number	PU	Input
145	Powerstations - Coal	Powerstation Type	Text	PU	Input
146	Coal Station Supplies	Sensitivity of rain on coal reliability	Number	PU	Input
147	Coal Station Supplies	Delivery uncertainty	Number	PU	Input
148	Coal Station Supplies	Base tied colliery ratio	Number	PU	Input
149	Coal Station Supplies	Sensitivity of tied colliery ratio	Number	PU	Input
150	Coal Station Supplies	Second moment around the mean (stdev) of Act/Bud	Number	PU	Input
151	Coal Station Supplies	Average Moisture of Coal Delivered	Num_Array	PU	Denorm
152	Coal Station Supplies	Average Abrasiveness Index of Coal Delivered	Num_Array	PU	Denorm
153	Coal Station Supplies	Average Ash content of Coal Delivered	Num_Array	PU	Denorm
154	Coal Station Supplies	Average Calorific Value of Coal Delivered	Num_Array	PU	Denorm
155	Coal Station Supplies	Cost of coal supply	Num_Array	PU	Denorm
156	Coal Station Supplies	Average Calorific Value of Coal Delivered	Num_Array	PU	Denorm
157	Coal Station Supplies	Coal Supplier - Cost of Supply	Number	ZAR/TON	Input
158	Coal Station Supplies	Coal Supplier - CV	Number	MJ/KG	Input
159	Coal Station Supplies	Mean of Actual/Budget	Number	PU	Result
160	Loads - System Level	Projections	Number	PU	Input
161	Loads - System Level	Hourly Noad Load	Number	MW	Input
162	Loads - System Level	Residential Node Load	Number	MW	Input
163	Loads - System Level	Manufacturing Node Load	Number	MW	Input
164	Loads - System Level	Mining Node Load	Number	MW	Input
165	Loads - System Level	Other Node Load	Number	MW	Input
166	Loads - System Level	Residential Losses	Number	GWH	Input
167	Loads - System Level	Manufacturing Losses	Number	GWH	Input
168	Loads - System Level	Mining Losses	Number	GWH	Input
169	Loads - System Level	Other Losses	Number	GWH	Input
170	Loads - System Level	Reserve Cost	Number	ZAR/MW	Input
171	Loads - System Level	Reserve Margin	Number	MW	Input
172	Loads - System Level	Load	Number	GWH	Input

Continued on the next page

	Parameter group	Parameter	Data type	Unit	Parameter type
173	Loads - System Level	Expected Input Load	Number	GWH	Input
174	Loads - System Level	Output of Load Module - iterations of load	Num_Array	MW	Result
175	Loads - System Level	Unservd Load	Num_Array	GWH	Result
176	Loads - System Level	Station Load	Num_Array	GWH	Result
177	Solar Heating Technologies	Expected impact in kW per household	Number	KW	Input
178	Solar Heating Technologies	Number of potential households for solar	Number	PU	Input
179	Solar Heating Technologies	Percentage of households implemented	Number	PU	Input
180	Efficiency impact - Industrial Loads	Percentage implementation of the impact	Number	PU	Input
181	Efficiency impact - Industrial Loads	Upper limit for the impact	Number	PU	Input
182	Efficiency impact - Commercial Loads	Percentage implementation of the impact	Number	PU	Input
183	Efficiency impact - Commercial Loads	Upper limit for the impact	Number	PU	Input
184	Heat pumps	Expected impact in kW per household	Number	KW	Input
185	Heat pumps	Number of potential households for solar	Number	PU	Input
186	Heat pumps	Percentage of households implemented	Number	PU	Input
187	Coal Contracts	Transport Type	Text	PU	Input
188	Coal Contracts	Ash	Number	PU	Input
189	Coal Contracts	Reliability	Number	PU	Input
190	Coal Contracts	CV Standard Deviation	Number	PU	Input
191	Coal Contracts	Reliability Standard Deviation	Number	PU	Input
192	Coal Contracts	Nett Coal Delivery	Number	PU	Input
193	Coal Contracts	Nett Coal Delivery Actual	Number	PU	Input
194	Coal Contracts	Reliability Actual	Number	PU	Input
195	Coal Contracts	Cost of Supply	Number	PU	Input
196	Coal Contracts	Cost of Supply Stdev	Number	PU	Input
197	Coal Contracts	CV	Number	MJ/KG	Input
198	Coal Contracts	CV Actual	Number	MJ/KG	Input
199	Aggregated Coal Contracts	Transport Type	Text	PU	Input
200	Aggregated Coal Contracts	Volume per month	Number	PU	Input
201	Aggregated Coal Contracts	Ash	Number	PU	Input

Continued on the next page

Parameter group	Parameter	Data type	Unit	Parameter type	
202	Aggregated Coal Contracts	Reliability	Number	PU	Input
203	Aggregated Coal Contracts	Reliability Standard Deviation	Number	PU	Input
204	Aggregated Coal Contracts	CV Standard Deviation	Number	PU	Input
205	Aggregated Coal Contracts	Cost of Supply	Number	PU	Input
206	Aggregated Coal Contracts	CV	Number	MJ/KG	Input
207	Aggregated Coal Contracts	Cost of Supply Stdev	Number	PU	Input
208	Loads - Node Level	Geographical area of node	GIS_Location_Array	DD	Input
209	Loads - Node Level	Residential Node Load	Number	MW	Input
210	Loads - Node Level	Manufacturing Node Load	Number	MW	Input
211	Loads - Node Level	Mining Node Load	Number	MW	Input
212	Loads - Node Level	Other Node Load	Number	MW	Input
213	Loads - Node Level	DSM Load Reduction	Number	MW	Input
214	Loads - Node Level	Base Node Load	Number	MW	Input
215	Loads - Node Level	CO2 factor	Number	PU	Input
216	Loads - Node Level	CO factor	Number	PU	Input
217	Loads - Node Level	SO2 factor	Number	PU	Input
218	Loads - Node Level	NOX factor	Number	PU	Input
219	Loads - Node Level	CO2 price	Number	PU	Input
220	Loads - Node Level	PM10 price	Number	PU	Input
221	Loads - Node Level	SOx price	Number	PU	Input
222	Loads - Node Level	NOx price	Number	PU	Input
223	Loads - Node Level	Residential Losses	Number	GWH	Input
224	Loads - Node Level	Manufacturing Losses	Number	GWH	Input
225	Loads - Node Level	Mining Losses	Number	GWH	Input
226	Loads - Node Level	Other Losses	Number	GWH	Input
227	Loads - Node Level	Mining Loss Multipliers	Number	GWH	Input
228	Loads - Node Level	Residential Loss Multipliers	Number	GWH	Input
229	Loads - Node Level	Other Loss Multipliers	Number	GWH	Input
230	Loads - Node Level	Manufacturing Loss Multipliers	Number	GWH	Input

Continued on the next page

	Parameter group	Parameter	Data type	Unit	Parameter type
231	Loads - Node Level	Reserve Cost	Number	ZAR/MW	Input
232	Loads - Node Level	Reserve Margin	Number	MW	Input
233	Loads - Node Level	Expected Node Load	Number	GWH	Input
234	Info Document Types	Info Document Types	Resource_URL	PU	Input
235	Info Document Types	Info Document Types	Text	PU	Input
236	Info Sql Types	Info Sql Types	Text	PU	Input
237	Info Sql Types	Info Sql Types	Text	PU	Input
238	Info Sql Types	Info Sql Types	Boolean	PU	Input
239	Network Lines	Line Load Factor	Number	PU	Input
240	Network Lines	Transmission Flow	Number	GWH	Input
241	Network Lines	Line Maximum Capacity	Number	GWH	Input
242	Network Lines	Transmission Losses	Number	GWH	Input
243	DSM Technologies	DSM technology amount	Number	PU	Input
244	TOU Definitions	Tariff Time of use	Text	PU	Input
245	TOU Definitions	Tariff TOU Item	Number	PU	Input
246	TOU Definitions	Tariff TOU Color	Text	PU	Input
247	Dashboards	Serialized Diagram	Text	PU	Input
248	Simulated Connections	Simulated Connection Load	Number	GWH	Input
249	Simulated Connections	Connection Sector	Text	PU	Input
250	Simulated Connections	Connection Start Date	Text	PU	Input
251	Simulated Connections	Connection Energy	Number	GWH	Input
252	Simulated Connections	Connection Month Capacity	Number	GW	Input
253	Renewable Generator	Renewable Generator Type	Text	PU	Input
254	Renewable Generator	Renewable Generator Start Date	Text	PU	Input
255	Renewable Generator	Renewable Generator Energy	Number	MW	Input
256	Renewable Generator	Renewable Generation	Number	GWH	Input
257	Renewable Generator	Commission date	DateTime	PU	Input
258	Renewable Generator	Commission date	DateTime	PU	Input
259	Renewable Generator	Powerstation Type	Text	PU	Input

Continued on the next page

	Parameter group	Parameter	Data type	Unit	Parameter type
260	System Dashboard	Serialized Diagram	Text	PU	Input

Appendix B

Analysis of the primary energy module

On the following page, 1 000 Monte Carlo sample paths are represented in Figures [B.1](#), [B.2](#), and [B.3](#). All fourteen commissioned power stations are represented for the planning horizon of eight months. The simulated values illustrated by box-plots are: burn, delivery, and the difference in delivery and burn, respectively.




Figure B.1: Simulated burn of all fourteen coal-fired power stations.




Figure B.2: Simulated delivery of all fourteen coal-fired power stations.




Figure B.3: Difference in simulated delivery and burn of all fourteen coal-fired power stations.

