The Stochastic Optimization of Long and Short-term Mine Production Schedules Incorporating Uncertainty in Geology and Equipment Performance

Matthew Quigley

Department of Mining and Materials Engineering

McGill University, Montréal, Quebec, Canada

December 2016

McGill COSMO
Stochastic Mine Planning Laboratory

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of Master of Engineering

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my supervisor, Professor Roussos Dimitrakopoulos, for believing in me and encouraging me to enroll in this program. His mentorship and guidance has helped me develop my critical thinking skills and has taught me to never fear challenging the norms. I am very grateful to now be a part of such a distinguished group of alumni. I would also like to thank Deborah Frankland, Rosalie Nardelli, and Barbara Hanley for their selfless support in all administrative manners.

I extend my thanks to Dr. Luis Montiel, Dr. Ryan Goodfellow, and Dr. Ilnur Minniakhmetov for always taking the time to assist me in my studies. I would like to give a very special thank you to Dr. Amina Lamghari for her input on my second research topic as well as her tireless efforts reviewing this thesis. I had the pleasure of sharing this experience at COSMO with a special group of individuals, including Cosmin, Ray, Ashish, Fernanda, Adrien, Lingqing, Renaud, Maria, Joao, and my good friends Iain, Rein, and James. I truly hope we can all stay in close contact throughout our careers.

Of course, this would not have been possible without the support from our sponsors. This work was funded by the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant 239019, and the COSMO consortium of mining companies (AngloGold Ashanti, Barrick Gold, BHP Billiton, De Beers, Kinross Gold, Newmont Mining, and Vale).

And last, but not least, I would like to thank my family, Cathy, Tony, and Sarah, and my girlfriend Ali for their unconditional support.

Contribution of Authors

The author of this thesis is also the first author for both manuscripts contained within. All work was completed under the supervision and guidance of Professor Roussos Dimitrakopoulos, who is the co-author of each manuscript.

Chapter 3 – Quigley, M. and Dimitrakopoulos, R. (2016) Risk-Resilient Mine Production Schedules Yielding Favourable Product Blends for Rare Earth Element Deposits. (Submitted) Mining Technology.

Chapter 4 – Quigley, M. and Dimitrakopoulos, R. (2016) Optimizing Short-term Mine Production Schedules Incorporating Geological and Equipment Performance Uncertainty. (In preparation)

Abstract

Mine production scheduling consists of defining the extraction sequence and process allocation of mineralized material over some length of time. These decisions can be made at different time steps, which will entail varying objectives subject to different technical and operational constraints. Long-term mine production scheduling usually takes place at an annual scale for the entire life of mine and aims to maximize the net present value of the project while satisfying the mining and processing capacities. Short-term mine production scheduling consists of developing an extraction sequence on a shorter time scale, either months, weeks, or days. The goal is typically to maximize compliance with the production targets imposed by the long-term plan while considering more detailed operational constraints. Historically, these optimization frameworks have relied on the assumption of perfect knowledge of highly uncertain inputs. Developments in the field of stochastic mine planning have shown that incorporating uncertainty into the optimization of mine production schedules can add significant economic value while also minimizing the risk of deviating from production targets. This thesis will explore the benefits that stochastic mine planning can offer when applied to both long and short-term production scheduling problems.

For the first exercise, the long-term mine production schedule of a rare earth element (REE) project is generated under geological uncertainty using a stochastic optimization framework. The uncertainty in REE grades is modelled using an efficient joint-simulation technique to preserve the strong cross-element relationships. The proposed approach avoids the use of the conventional total rare earth oxide grade. The stochastic long-term schedule is benchmarked against a deterministic schedule generated using an industry standard optimizer. The stochastic solution generates a 20% increase in expected NPV, ensures better utilization of the processing plant, and delivers a superior ore feed in terms of satisfying mineral and REE blending targets.

For the second exercise, a formulation is proposed that simultaneously optimizes the short-term equipment plan and production schedule under both geological and equipment performance uncertainty. The proposed approach rectifies certain limitations of previous work in stochastic short-term planning by: incorporating a location-dependant shovel movement optimization; generating more realistic equipment performance scenarios; developing a new approach to facilitate more practical mine designs; and proposing model improvements to allow for a more efficient optimization of very large problem instances. The model is applied to a large copper mining complex and is compared to a more traditional approach, where the same formulation is implemented using averaged inputs for geology and equipment performance. The stochastic solution is more effective in mitigating the risk of deviating from tonnage targets at each processing destination, and the integration of equipment performance variability allows the stochastic optimizer to generate a block extraction sequence that is far more likely to be achieved.

Résumé

La cédulation de la production minière consiste à définir la séquence d'extraction et le mode de traitement du minerai, discrétisé en un modèle de blocs, au sein de différentes fenêtres et échelles de temps. Ces dernières répondent à des objectifs distincts et sont soumises à des contraintes techniques et opérationnelles différentes. La planification minière à long terme vise habituellement à définir la production annuelle de la mine jusqu'à la fin de l'exploitation et pour cela cherche à maximiser la valeur actuelle nette du projet tout en satisfaisant les capacités d'extraction et de traitement. La planification à court terme cherche quant à elle à déterminer une séquence d'extraction à une échelle de temps plus courte puisqu'elle se concentre sur quelques mois, semaines ou même jours de production uniquement. En général, l'objectif est alors de maximiser la satisfaction des cibles de production définies par la planification à long terme tout en considérant des contraintes opérationnelles plus détaillées. Historiquement, ce type d'optimisation considère l'hypothèse d'une connaissance parfaite des paramètres incertains. Cependant, les études sur la planification stochastique montrent qu'intégrer l'incertitude dans l'optimisation de la production des mines est en mesure d'augmenter significativement la valeur économique d'un projet, tout en minimisant le risque de déviation par rapport aux cibles de production. Ce mémoire explorera les gains que permettent d'obtenir des considérations stochastiques dans la planification des mines lors d'une optimisation à long terme et à court terme.

On présente d'abord une étude portant sur la planification minière à long terme d'un projet d'extraction d'éléments de terres rares (ETR) prenant en compte l'incertitude géologique dans le cadre d'une optimisation stochastique. L'incertitude concernant les teneurs des terres rares est modélisée en utilisant une technique de simulation-conjointe efficace pour préserver les relations inter-éléments. L'approche proposée évite l'utilisation conventionnelle de la teneur en oxydes de terres rares totales. La planification stochastique obtenue est comparée à une planification dite déterministe générée avec un optimiseur conventionnel. Les gains de la solution stochastique sont divers: une augmentation de 20% de la valeur actuelle nette espérée; une meilleure utilisation de la capacité de l'usine de traitement; et un concentré de minerai de qualité supérieure en termes de contraintes de mélanges des ETR.

La deuxième étude propose une formulation stochastique qui vise à simultanément optimiser l'allocation de l'équipement et la planification de la production à court terme en prenant de nouveau en compte l'incertitude géologique mais également l'incertitude sur la performance de l'équipement. L'approche proposée complète certaines limitations de précédents travaux en planification stochastique à court terme, notamment en: incorporant l'optimisation du mouvement et de l'emplacement des pelleteuses; en générant des scénarios de performance de l'équipement plus réalistes; en présentant une nouvelle approche pour définir des contours d'extraction plus pratiques; et en proposant des améliorations de stratégies d'optimisation afin de résoudre plus efficacement des problèmes de grandes dimensions. Le modèle est appliqué à un complexe minier d'extraction de cuivre et est comparé à une approche plus traditionnelle, déterministe, où la même formulation est résolue en utilisant uniquement les paramètres moyens pour la géologie et la performance des équipements au lieu d'un ensemble de simulations. La solution stochastique s'avère être plus efficace pour contrôler le risque de déviation des cibles de tonnages de chaque destination de traitement. De plus, l'intégration de la variabilité de la performance de l'équipement permet à l'optimiseur stochastique de générer une séquence d'extraction pour laquelle les objectifs de production sont beaucoup plus susceptibles d'être atteints.

Contents

List of Figures

List of Tables

List of Terms

AMS	Automated manufacturing systems
B&B	Branch and bound
CAD	Dollars Canadian
CCDF	Condition cumulative distribution function
CDF	Cumulative distribution function
Ce	Cerium
CoG	Cut-off grade
CSSR	Conditional simulation by successive residuals
DBMAFSIM	Direct block min/max autocorrelation factors simulation
DBSIM	Direct block simulation
DS	Direct sampling
Dy	Dysprosium
Er	Erbium
Eu	Europium
Gd	Gadolinium
Ho	Holmium
HOSIM	High-order spatial simulation
HREE	Heavy rare earth element
KPI	Key performance indicator
La	Lanthanum
LOM	Life of mine
LP	Linear programming
LREE	Light rare earth element
LTMPS	Long-term mine production schedule
LU	Lower and upper (decomposition)
Lu	Lutetium
MAF	Min/max autocorrelation factors
MIP	Mixed-integer programming

MPS	Multiple point statistics
Nd	Neodymium
PB	Pushbacks
PCA	Principle component analysis
Pr	Praseodymium
REE	Rare earth element
ROM	Run of mine
SGS	Sequential Gaussian simulation
SIP	Stochastic integer programming
SIS	Sequential indicator simulation
SNESIM	Single normal equation simulation
Sm	Samarium
SMU	Selective mining unit
SSTMPS	Stochastic short-term mine production schedule
STMPS	Short-term mine production schedule
Tb	Terbium
TI	Training image
Tm	Thulium
TREO	Total rare earth oxide
UPL	Ultimate pit limit
VND	Variable neighborhood descent
Yb	Ytterbium

Chapter 1

Introduction

Discovering, financing and developing new mining projects is a unique and challenging task. Since the peak of the commodity super cycle in 2011, the incumbents of the mining industry have witnessed the adverse effects of shrinking commodity prices, tougher environmental regulations, difficulty in procuring capital, and rising capital and operating costs. These challenges, coupled with the future projections of even deeper deposits with declining metal grades, means it is now more important than ever for miners to establish and implement stronger engineering and operational decision making frameworks in order to succeed. Effective mine planning practices can generate realistic designs that guide the progression of mining activity while maximizing the return on investment. More specifically, optimal mine production scheduling can be used to maximize the generation of cash flows, ensure efficient use of capital spending over time, and shield the operation from risk.

Mine production scheduling consists of defining the extraction sequence and process allocation of mineralized material over some length of time. These decisions can be made at different time scales, which will entail varying objectives subject to different technical and operational constraints. Long-term mine production scheduling usually takes place at an annual scale for the entire life of mine. A long-term plan aims to maximize the net present value of the project, while abiding by the mining and processing capacities. The resulting schedule defines the annual ore, waste, and metal tonnages, and the annual cash flows which are crucial in valuing the asset in order to attract investors. Short-term mine production scheduling consists of developing an extraction sequence on a shorter time scale, either months, weeks, or days. The goal in short-term scheduling is often to maximize compliance with the production targets and restrictions imposed by the long-term plan, while considering more detailed operational constraints, such as available equipment, equipment access, material haulage time, and the ability to deliver consistent ore quantity and quality to each processing destination. These goals can be achieved by optimizing the location and movement of mobile mining equipment, the extraction sequence, and the material destinations.

Historically, most optimization frameworks in the field of open pit mine production scheduling have relied on estimated orebody models to represent the geological characteristics. However, it has been understood for years that estimated models, which are primarily used for reserve assessment, are not fit to use for planning purposes since they cannot adequately represent the spatial variability or extreme values of deposits (David et al., 1974). These shortcomings have motivated researchers to develop risk-based mine planning frameworks that are able to incorporate and manage geological uncertainty by representing orebodies using multiple geostatistical simulations (Menabde et al., 2007; Boland et al., 2008; Ramazan and Dimitrakopoulos, 2013). Furthermore, when considering more detailed aspects of short-term planning, the failure to account for uncertainty in equipment availability and productivity can have adverse effects as well. Equipment performance is subject to random fluctuations in road conditions, operator habits, and maintenance practices that should be measured and managed in order to generate short-term schedules that have stronger likelihoods of being realized in practice.

Despite significant contributions in the field of stochastic mine planning in recent years, the practical implementation of these risk-based approaches is still limited in the mining industry today. The primary deterrent is the additional complexity that is introduced by incorporating multiple scenarios to problems that are already computationally demanding. This thesis will explore methods to generate quality realizations of geological deposits and equipment performance scenarios as well as develop stochastic formulations to efficiently optimize long and short-term mine production schedules for different commodities.

1.1 Goals and Objectives

The goal of this thesis is to investigate and test frameworks that optimize long and short-term mine production schedules considering uncertainty in geology and equipment performance. In order to achieve this goal, the following objectives must be fulfilled.

1. Review the literature pertaining to the stochastic optimization of open-pit mine production scheduling. More specifically, the literature on past and recent development in long and short-term production scheduling and the stochastic simulation of mineral deposits.
2. Optimize the long-term production schedule considering uncertain geology for a rare earth element project, a commodity that has yet to be explored using stochastic mine planning techniques:
    a. Model the volumetric and grade uncertainty of the deposit;
    b. Efficiently optimize the long-term production schedule.
3. Develop a stochastic formulation for the simultaneous optimization of the short-term equipment plan and production schedule for a very large copper mining complex:
    a. Propose a novel approach to co-simulate equipment performance scenarios;
    b. Optimize the short-term production schedule using heuristic techniques and a general purpose solver.
4. Benchmark these stochastic optimization frameworks against the industry's best practices to demonstrate the benefits of incorporating uncertainty in mine production scheduling.

1.2 Thesis Outline

1. This chapter introduces the motivation, the goals, the objectives and the outline of the thesis.
2. Chapter 2 reviews the literature pertinent to the stochastic optimization of mine production schedules. This entails covering advances in the fields of long-term production scheduling, short-term production scheduling, and stochastic orebody simulation.
3. Chapter 3 investigates the procedures for generating geostatistical simulations and optimizing long-term production schedules under uncertain geology for rare earth element deposits.
4. Chapter 4 introduces a new approach to co-simulate equipment performance scenarios and develops a formulation to simultaneously optimize the short-term equipment plan and production schedule under both geological and equipment uncertainty.
5. Chapter 5 summarizes the finding of this thesis, addresses the limitations, and proposes suggestions for future work.

Chapter 2

Literature Review

2.1 Outline

This chapter outlines the literature pertinent to the topics of stochastic optimization of long and short-term mine production schedules. This review is organized in three sections.

1. Section 2.2 reviews previous work in the field of long-term mine production scheduling. It covers the progression of conventional optimization techniques, extensions to stochastic optimization frameworks, and new concepts developed in the field of global optimization of mining complexes under uncertainty.
2. Section 2.3 covers the literature focusing on short-term mine production scheduling. A wide variety of operational objectives can be optimized in the short-term which has led to many different decision making models. The progression from earlier works to more modern short-term optimization methods is discussed. Then, short-term production scheduling models that incorporate operational and geological uncertainty are reviewed, followed by techniques that facilitate practical mining patterns.
3. Section 2.4 reviews the development of geostatistical simulation techniques, from more traditional methods to algorithms that offer significant computational efficiency improvements. Following this, advancements in multiple point and high-order statistical simulation techniques are covered.

2.2 Long-term mine production scheduling

The optimization of a long-term mine production schedule (LTMPS) is the process of defining the timing of extraction and process allocation of mineralized material. This production schedule is designed in a way that maximizes the total value of the products sold while obeying certain operational and technical constraints. Long-term optimization is most often executed at an annual scale and spans the entire life of mine (LOM). This task is essential in determining the annual production tonnages, cash flows, and in turn, the economic value of the project to secure capital investment. All of the material in the orebody model is discretized into selective mining units (SMU), or blocks, each having a fixed position in space, a uniform volume, and a set of geological properties which define its quality. The following sub-sections will review some of the progress in the field of long-term mine production scheduling.

2.2.1 Conventional optimization models

To develop optimal long-term mine production schedules, many earlier works aim to decompose the problem into disjoined stages. The first stage is to eliminate all of the blocks that lie below the boundary of profitable extraction; the remaining blocks are known as the ultimate pit limit (UPL). The UPL is most often solved using some version of the original Lerchs-Grossmann algorithm (Lerchs and Grossmann, 1965), which maximizes the total value of enclosed blocks considering their precedence relationships and economic values. A popular extension of this algorithm involves parameterizing the economic values of blocks in order to obtain smaller pits within the UPL, referred to as nested pit shells. In the second stage of this step-wise optimization, these nested pit shells are grouped together to form larger units within the UPL, known as phases or pushbacks (Whittle, 1988). Pushbacks are used to guide the progression of mining activity throughout the LOM, but most importantly, they provide the means to define the fundamental scheduling units in the third and final stage of this optimization framework. Mathematical programming techniques are used to schedule the ore, waste, and metal production of these units on an annual basis to optimize the discounted cash flows while satisfying the operational capacities (Johnson, 1968; Dagdelen and Johnson, 1986; Tolwinski and Underwood, 1996). This step-wise optimization of the production schedule is the foundation of most current commercial mine scheduling packages since it reduces the complexity and total solving time of the optimization. However, this simplicity comes at the cost of the quality of the final solution. There is no guarantee that the generation of the final scheduling units are optimal since they are created without accurate cash flow discounting and without the final problem's key components, such as grade blending, processing capacities and stockpiling constraints.

More recent works understand this major drawback and focus on optimizing the block-level mine production schedule using mixed integer programming (MIP) techniques (Gershon, 1983). Caccetta and Hill (2003) propose a branch-and-cut method to solve their MIP formulation of the mine production scheduling problem. The authors test the method on different problem sizes, the largest of which consists of 210 thousand blocks and is solved to an 8% optimality gap in 20 hours. Cullenbine et al. (2011) combine Lagrangian relaxation techniques with a sliding time window heuristic (STWH) (Pochet and Wolsey, 2006). The algorithm uses a fix-and-optimize scheme and sequentially solves sub-problems where the integrality constraints are relaxed for smaller subsets of periods. The authors are able to solve a problem instance of 25 thousand blocks and 15 periods but highlight the fact that the algorithm offers no guarantee of finding a feasible solution. Moreno et al. (2010) generate near-optimal mine production schedules by first solving linear relaxations of the problem and then creating feasible integer solutions using a heuristic based on topological sorting. The method is applied to a very large problem containing 3.5 million blocks and 15 periods. However, the topological sorting heuristic requires that the problem only has one resource knapsack constraint per period.

2.2.2 Stochastic optimization models

2.2.2.1 Motivation for stochastic techniques

Historically, all optimization frameworks in the field of open pit mine production scheduling have relied on one basic assumption: that the geological characteristics of the deposit are fully understood. More specifically, the metal grades that define economic values of blocks, which drive the optimization techniques, are assumed to fully characterize the quality of the true orebody. The most common convention is to represent these metal grades using an estimated orebody model (David, 1988; Isaaks and Srivastava, 1989). These geostatistical estimation techniques assign the expected grade value to each location in the orebody using some weighted average of all surrounding points. By construction, these techniques do not adequately represent the deposit since they cannot preserve the connectivity of high grade values and they fail to reproduce the univariate distributions or the spatial correlations of the input data.

Stochastic geostatistical simulation methods (Journel, 1974; Goovaerts, 1997) address the shortcomings of estimation techniques by randomly sampling a conditional probability distribution for each location rather than taking its expectation. These methods can generate multiple, equally probable representations of the deposit that better reproduce the extreme values, the univariate distributions, and the spatial correlations of the input data. Together, these simulations can also be used to quantify the geological uncertainty of the deposit. As these techniques gain popularity, researchers have begun to investigate how these simulated models might impact reserve estimations, ultimate pit limits, and production schedules. It is known that the existing optimizers that rely on estimated orebody models will always underperform in the presence of uncertainty, motivating the need for new stochastic optimization techniques (Ravenscroft, 1992; Vallée, 2000). Dimitrakopoulos et al. (2002) support this conclusion by quantifying the geological risk of a conventional mine production schedule using simulated orebodies. The risk analysis shows that the forecasted net present value (NPV) only has a 5% chance of being realized, and that the expected NPV is actually 25% lower than originally predicted. The authors attribute these misleading results to the inability of deterministic optimizers to effectively prioritize high grade ore material, stemming from the deficiencies of estimated orebody models in reproducing the spatial variability and connectivity of extreme values.

$$Value_i = \begin{cases} +Net\ Profit_i & \text{if } g_i \geq CoG \\ -Mining\ Cost_i & \text{if } g_i < CoG \end{cases} \qquad 2.1$$

Using an average grade model has detrimental impacts on the mine production schedule since the economic value of a block is calculated using a non-linear function (Equation 2.1). From this non-linearity, it follows that mine production schedules optimized with these averaged values will not necessarily perform well on average (Birge and Louveaux, 1997) when exposed to geological simulations, which are shown to better represent the true deposit. Depending on the value of the cut-off grade (CoG) relative to the deposit's grade distribution, an estimated orebody model will either over-predict or under-predict the total quantity of ore material. Both of these situations will cause the mine plan to perform sub-optimally. The next frontier of mine production optimizers should exploit multiple simulated orebody scenarios in order to capitalize on the improved understanding of connected of high-grade blocks, and the ability to shield the production schedule from geological uncertainty. Sections 2.2.2.2 to 2.2.2.6 will describe different approaches that have been proposed to integrate geological uncertainty when generating optimal long-term mine production schedules.

2.2.2.2 Evaluation of production schedules using risk analyses

The earliest attempt to develop a long-term mine production schedule while integrating geological uncertainty is based on the idea of maximum upside and minimum downside (Dimitrakopoulos et al., 2007). The authors generate multiple conventional production schedules, each using a simulated orebody model as input. Each schedule is then evaluated given its probability of being above certain key performance indicator (KPI) thresholds. This probability is quantified through a risk analysis using all simulated geological scenarios. A subset of schedules that pass the preliminary threshold requirements are retained, and the best schedule is chosen if it maximizes the upside potential and minimizes the downside risk, meaning that it has the highest likelihood of performing better than expected. Although this approach can be implemented with existing tools, the resulting schedules are only optimal for one particular scenario. This approach is an improvement over deterministic methods that ignore uncertainty, yet, it does not attempt to directly optimize the production schedule in the presence of geological uncertainty and hence, will always yield sub-optimal results.

2.2.2.3 Probability driven production scheduling

The next generation of stochastic optimizers aim to integrate geological uncertainty using probabilistic attributes at the block level. Ramazan and Dimitrakopoulos (2004b) derive a measure of probability that a block should be mined in a certain period. The authors begin by generating a unique production schedule for each simulated orebody using a standard MIP formulation. A probability distribution is then calculated for each block, characterizing its likelihood of being extracted in a given period. This probability measure is then included in the objective function of a new MIP formulation that encourages the extraction of blocks in the period in which they were most likely to be mined.

Dimitrakopoulos and Ramazan (2004) develop a similar concept, whereby the optimizer penalizes blocks that do not have strong likelihoods of portraying desirable geological properties. There are two major advantages to this approach: (1) it can be easily extended to multi-element operations; (2) it does not require solving multiple, time-consuming MIPs to derive the probability measure; it can be pre-processed for each block very quickly given the range of geological scenarios. One of the more important contributions of this method is the concept of geological risk discounting. The penalty costs associated with deviating from the desirable geological properties are reduced over time. Similar to the concept of discounting cash flows, geological risk discounting encourages the optimizer to prioritize risk management earlier in time and defer more uncertain production to later periods. This is an attractive feature since it is imperative for operating companies to achieve production targets early in the LOM to pay back investors. Furthermore, deferring risk to later in the project life is preferable since more information will become available over time, allowing for corrections to be made. Grieco and Dimitrakopoulos (2007) extend this probabilistic approach to optimize the stope design of an underground gold mine. Each stope is given a probability of being above a pre-defined metal CoG, quantified using orebody simulations. The optimizer minimizes the risk of this underground stope design by only mining stopes that satisfy a minimum probability threshold. This allows an underground mine planner to explicitly control the relationship between ore tonnage and grade certainty.

All of the methods listed above suffer from one major drawback: they rely on a probabilistic measure for each block independently, ignoring the uncertainty of its surrounding neighbours. This joint-local uncertainty is better represented by simultaneously considering multiple orebody simulations in the optimization process.

2.2.2.4 Production risk management with simulated annealing

More recent developments in stochastic mine production scheduling aim to directly manage joint-local uncertainty by incorporating multiple geological scenarios in the optimization simultaneously. Godoy and Dimitrakopoulos (2004) develop an approach to minimize the ore and waste target deviations across multiple geological scenarios. Similar to the work of Ramazan and Dimitrakopoulos (2004b), the first step of this method is to generate a production schedule for each orebody simulation using some deterministic technique in order to derive the probability of each block belonging to a certain period. Then, a single optimal production schedule is generated using a metaheuristic based on simulated annealing (SA) (Kirkpatrick et al., 1983; Geman and Geman, 1984). The algorithm is initialized with a starting schedule, which is then gradually improved by changing the period when blocks are extracted. These block extraction perturbations (swaps) are generated by randomly sampling the block transition probability derived in the previous step. These changes are either rejected or accepted depending on how much they improve the objective of minimizing ore and waste target deviations across all scenarios. A useful feature of this approach is the ability to control the rate at which perturbations are accepted over time. The algorithm can thoroughly explore the solution space earlier in the run by accepting less desirable changes and can also narrow the focus on improving the objective value later on by only accepting the highest quality swaps. Godoy and Dimitrakopoulos (2004) apply this methodology to a gold case study, and their results demonstrate that the stochastic design only poses a 3% chance of failing to meet annual ore production targets as opposed to 13% for the deterministic design. The authors also show this risk-based approach yields a 28% improvement in expected NPV when compared to the deterministic solution. Leite and Dimitrakopoulos (2007) apply the same stochastic optimizer to a copper deposit. Their stochastic solution shows a 15% increase in expected NPV when compared to the deterministic design. Furthermore, the stochastic schedule ceases production one year earlier than the deterministic schedule. This is a result of the estimated orebody model indicating low-grade ore at the bottom of the deposit, where the simulations depict waste.

Although this stochastic optimization does not directly incorporate NPV in the objective function, the value is always increased since the block transition probabilities reflect the motivation to extract high-grade blocks earlier, and defer lower grades to later periods. This approach shows substantial improvement over the probabilistic methods and is also attractive since it does not require a general purpose solver to generate a mine design. However, this technique is not without limitations: the maximization of NPV is not directly optimized, complex grade blending requirements are not considered, and there is no guarantee of generating a proven optimal solution.

2.2.2.5 Production scheduling with stochastic integer programming

Ramazan and Dimitrakopoulos (2007) develop a two-stage stochastic integer programming (SIP) approach to solve the long-term mine production scheduling problem with uncertain geology. This two-stage approach integrates both anticipative and adaptive mechanisms of dealing with uncertain input. First stage (anticipative) decisions are designed to be robust to uncertainty, whereas second stage recourse (adaptive) decisions are used to facilitate corrections once the uncertainty is revealed. In the context of the production scheduling problem, the authors formulate the block extraction variables as first stage decisions to be robust to geological uncertainty. The recourse decisions are the annual tonnage and grade deviations, which vary by scenario. The optimizer generates a solution that maximizes the expected NPV while also minimizing the scenario-dependant production deviations, resulting in a schedule that is more valuable and risk-resilient. This SIP can be solved using commercial optimization packages, which is useful as it provides an indication as to how far the solution is from the theoretical optimum.

Ramazan and Dimitrakopoulos (2013) apply this two-stage SIP to a small gold deposit and show a 10% increase in expected NPV when compared to the deterministic design. The authors demonstrate how the stochastic optimizer is more effective in minimizing and distributing production risk throughout the LOM as opposed to the random fluctuations seen in the traditional schedule. Leite and Dimitrakopoulos (2014) apply the same SIP formulation to a copper deposit and highlight a 29% expected NPV increase over traditional methods. The authors comment on the optimizer's ability to better control the risk profiles and suggest that it can generate better balances between blocks with high but uncertain grades, and lower valued blocks with higher degrees of certainty. Benndorf and Dimitrakopoulos (2013) extend this approach to stochastically optimize a multi-element mining operation with more complex blending requirements. By minimizing the grade deviations of multiple elements in an iron ore deposit, the authors confirm that the production schedule will generate higher quality material with more certainty when compared to the deterministic design. The authors investigate the impact that the magnitudes of the deviation penalties have on the quality of the stochastic solution. They conclude that increasing these penalty values will minimize the annual production deviations to a certain limit; the production risk cannot be eliminated in full, and at some point the marginal increase in penalty magnitude no longer provides any benefit.

By increasing the expected NPV and minimizing the deviations from production targets, stochastic optimizers produce counter-intuitive results that value can be increased while also reducing risk, two entities that are traditionally understood as being inversely related. Integrating multiple simulated orebodies when generating long-term mine production schedules allows the optimizer to exploit the connectivity of high grade areas of the deposit that deterministic optimizers cannot identify. Furthermore, in the presence of uncertainty, stochastic optimizers are capable of explicitly managing the production risk overtime, specifically by blending the uncertainties of many blocks together and by deferring uncertainty to later time periods through geological risk discounting. However, these SIP optimizers carry certain limitations. Only production scheduling problems that can be modelled as linear programming formulations can be solved and downstream processing decisions, stockpiling, and more complex aspects of metal recovery cannot not be incorporated in this SIP model. Furthermore, the size of the problems that can be solved is very limited due to approach's reliance on general purpose solvers. Realistic mining operations containing hundreds of thousands to millions of blocks that span half-decade LOMs are too large to be solved efficiently with conventional tools. For these reasons, researchers have investigated non-exact solving techniques for these SIP formulations.

2.2.2.6 Solving SIP formulations using metaheuristics

Metaheuristics offer alternative means to solve very large, complicated integer programs by intelligently exploring a solution space to find a high quality solution quickly rather than searching for the theoretical optimum in an intractable amount of time. Lamghari and Dimitrakopoulos (2012) develop an efficient solution method for the production scheduling SIP using a Tabu Search (TS) metaheuristic. The authors investigate two separate diversification strategies and test them using ten different problem instances, ranging from a few thousand blocks in three periods to over 40 thousand blocks spanning 11 years. The authors use a commercial solver to generate an upper-bound by solving the linear relaxation of each problem instance and show that the TS algorithm is able to consistently find solutions within 2 – 3% of these upper-bounds. The most attractive feature of this algorithm is the running time. For the largest problem set, the Tabu Search solver generates a solution with a 2.4% gap in 104 minutes as opposed to over 15,000 minutes required by the commercial solver.

Lamghari et al. (2014) propose a variable neighborhood descent (VND) metaheuristic to solve the production scheduling problem with geological uncertainty. The authors investigate the VND algorithm's solution quality and running sensitivity to two different types of initial solutions: an exact solution and a heuristic solution to a sequential, period-separated series of sub-problems. The VND algorithm is tested on three large-scale problem instances and efficiently generates solutions within a 3% optimality gap when compared to results of the commercial solver. The authors show that the quality of the VND approach using the exact initial solution is better than when using the greedy heuristic, however, the greedy heuristic approach provides a final solution in significantly less time.

More complicated mine optimization formulations have been developed recently which rely heavily on the use of metaheuristic solvers due to their non-linearity and very large problem sizes. Some of the work in the unified modelling and optimization of mining complexes is presented next.

2.2.3 Global optimization of mining complexes under uncertainty

A mining complex is an integrated system comprised of multiple mines, multiple processing destinations, and complex downstream transportation arrangements that, together, convert raw material into saleable products. These complex operations necessitate the holistic, simultaneous optimization of all of the extraction, processing, routing, and transportation decisions in order to exploit the value that is lost when traditionally optimized using step-wise approaches. One of the major developments in this field is the paradigm shift that the value generated from a mining operation should be derived from the final products sold as opposed to the traditional convention, where value is assumed at the block-level. This reflects the reality that blocks are neither extracted nor processed independently, but rather, are combined together at various stages throughout the mineral value chain. These interactions have an impact on how material is blended and recovered within the mining complex, therefore, these material interactions should be explicitly modelled and managed throughout the optimization to ensure a more accurate valuation. The presence of geological uncertainty further complicates the flow of material through a mining complex but also presents the opportunity to generate production plans that are more valuable with less risk. Similar to the conventional SIP models, the extraction decisions in a mining complex are formulated to be robust to uncertainty. However, destination, stockpiling, and transportation decisions can be treated as corrective recourse variables.

Goodfellow and Dimitrakopoulos (2016) develop a unified model for the simultaneous optimization of mining complexes considering geological uncertainty. The authors develop a multi-element material destination policy that is optimal under uncertain geology but also robust and operationally implementable. First, the block values are clustered together in a high-dimensional material attribute space. The destination of each of these block clusters is then identified in the simultaneous optimization. The resulting destination policy is robust in that any block that is observed during production can be allocated to its optimal destination since it must belong to an existing cluster. The model is solved using three imbedded metaheuristics: simulated annealing, particle swarm optimization, and differential evolution. The authors apply the methodology to a very large copper-gold deposit and show that the proposed model is able to outperform a deterministic solver by increasing the expected NPV by 22.6% and decreasing the production risk.

Montiel and Dimitrakopoulos (2015) develop a modelling and solving approach to optimize large mining complexes under geological uncertainty with the integration of variable processor operating modes and flexible transportation options. Traditionally, these stages of a mining complex would be optimized independently but this particular model demonstrates how the simultaneous optimization of these downstream decisions can be used to fully exploit the operating capacities. This stochastic optimizer uses a simulated annealing solution method that explores perturbations on three different operational levels: the block extraction level, the operating mode level, and the transportation option level. This framework is applied to a large copper operation, and the results show how altering the operating grind-size of the crushing circuit in certain periods can optimize the throughput to better align with the production schedule. The authors show a 5% expected NPV increase compared to the operating company's current production plan. Montiel et al. (2016) also show the model's ability to adapt to different types of mining operations through the simultaneous optimization of an integrated open-pit and underground gold mining complex.

2.3 Short-term mine production scheduling

Similar to long-term planning frameworks, the optimization of short-term mine production schedules (STMPS) defines the best extraction sequence and destination allocation of mineralized material over some time horizon. The distinguishing factors are the time scale at which the optimization is performed and the specific objectives and constraints that are considered. Since shorter time scales usually negate the need for risk-discounted cash flows, short-term optimizers are motivated by a different set of objectives. Where long-term schedulers optimize the annual tonnages and discounted cash flows given fixed mining and processing capacities, short-term production schedules tend to prioritize tonnage and grade consistency at time steps of months, weeks, or days, while considering more granular operational decisions and constraints. Some of these operational additions might be the precise planning of equipment movement, the availability of drill and blast equipment, the productivity and matching of loading and hauling units, truck routing options, and more refined equipment access to material. The following sub-sections will review the literature pertaining to short-term mine production scheduling.

2.3.1 Conventional short-term optimization models

Although not as well-studied as LTMPS, STMPS optimization problems have gained considerable interest, given the substantial value that can be added by ensuring consistent production and by minimizing operating costs. Every mining project is unique in the sense that a wide variety of different operational aspects can be considered critical at one mine and neglected at the next. Since short-term planning often directly precedes operations, the primary goals in short-term optimization approaches may vary as well. In general, the extent of the material considered in short-term optimization is derived from the preceding long-term plan, and this material is also represented using mining blocks.

Wilke and Reimer (1977) develop a linear programming (LP) model for a STMPS that maximizes profit while delivering material that satisfies mining and milling capacities and grade blending requirements. Fytas (1985) also develops an LP formulation to optimize the STMPS of an iron ore mine that is constrained by the amount of ore produced, the stripping ratio, and multi-element blending requirements in each period. These past LP applications are relatively small, whereby sub-problems consisting of one or two periods are solved sequentially and blocks are grouped into bench polygons to limit the number of decision variables. More importantly, these methods rely on the continuous representation of extraction decisions that the authors state is not feasible in terms of mining activity. Both of these drawbacks are a result of the state of computing power at the time of these studies. Mixed integer programming (MIP) has since been deemed more appropriate to solve mine production scheduling problems since the resulting extraction precedence of blocks is feasible and implementable. Smith (1998) develops a short-term MIP formulation that maximizes ore production complying with grade blending requirements. The author applies the formulation to a small, single period, single bench study using a commercial solving package. Suggestions are made as to how to reduce the size of the problem to reduce the solving time. The author concludes that the major drawback of the approach is the computational burden of solving for many binary variables given large, multi-period applications.

Kumral and Dowd (2002) develop a multi-objective simulated annealing (MSSO) metaheuristic to optimize a STMPS problem for an iron ore deposit. In this approach, a sub-optimal solution is derived from Lagrangian parametrization and improved using MSSO. The different objectives are to: satisfy both total tonnage and ore tonnage mined, minimize the grade target deviations for two elements, and minimize the grade variance between blocks for the same two elements. In this MSSO approach, the perturbation acceptance criteria is not a simple binary decision since a candidate solution might improve one objective and deteriorate another. The prioritization of these objectives is subject to bias, and the authors generate over 250 solutions relatively quickly and select the best 5 manually. This particular approach relies on a significant amount of subjectivity and is unable to provide one optimized production schedule. Moreover, this approach does not consider more pertinent short-term constraints, such as equipment access or movement.

Eivazy and Askari-Nasab (2012) propose an MIP formulation to optimize a STMPS. The authors incorporate more complex operational aspects, such as: variable block destinations to avoid the use of a pre-defined cut-off grade, the ability to reclaim material from a fixed grade stockpile in order to satisfy blending requirements, and ramp routing options to minimize the haulage cost. In order to facilitate these added decision variables, the authors reduce the size of the problem by aggregating blocks into much larger scheduling units with averaged grades. The optimal solution is generated using a commercial solver, and the results show that the optimizer is able to deliver consistent monthly tonnages, within the grade bounds, to each of the processing destinations. In this particular formulation, the optimal ramp routing decisions for each block could have been set to the shortest-distance option *a priori* since the ramps are not capacitated in any time period.

The methods described above are what Alarie and Gamache (2002) define as the upper-stage of the common multi-stage STMPS optimization. In this framework, the upper-stage generates the production schedule by satisfying blending and capacity constraints, which then imposes the targets for the lower-stage optimization, which optimizes the dispatching plan for the mobile fleet. More recent methods attempt to simultaneously optimize these scheduling and fleet decisions to generate more robust solutions. L'Heureux et al. (2013) incorporate the sequencing of production activities such as drilling, blasting, and extraction into their MIP formulation along with a loader movement optimization. The authors define the different spatial units of the model: areas are regions that qualify long equipment moves, faces are small grouping of blocks that act as the scheduling units, and clusters are groupings of faces that can be made available through one production blast. The formulation ensures that a shovel must re-locate to a face before the material can be extracted, and the shovels' moves are penalized in the objective function. The proposed approach is tested through an application consisting of 200 blocks over 90 daily periods. The authors are constrained by the computational limitations of commercial solvers and make modifications to the formulation in order to reduce the solving time. Substantial running time reductions are achieved, but the authors conclude that heuristic solvers offer a more attractive option for solving applications of realistic sizes.

Mousavi et al. (2016) develop a short-term scheduler that simultaneously optimizes the block extraction sequence and the destination decisions, while incorporating fixed-grade stockpile reclamation decisions and more comprehensive block access constraints. The solutions are generated using a hybrid heuristic consisting of simulated annealing (SA), branch and bound (B&B) and a large neighbourhood search (LNS). The authors give three reasons that stress the importance of scheduling at the block-level as opposed to using larger aggregated units: (1) using aggregating blocks reduces the granularity of the solution and will lead to sub-optimal results since pre-processed grouping algorithms often lack sophistication; (2) averaging block grades through aggregation intensifies the grade smoothing effect and suppresses the presence of grade variability; (3) when schedules optimized using aggregated units are evaluated at the block level, the solutions are often infeasible in the original context of the problem. The authors test the proposed method on a variety of problem instances, the largest of which consists of 2,500 blocks, 6 destinations, and 12 periods. The heuristic method consistency rivals the performance of the commercial solver in terms of quality. Furthermore, solutions are generated for instances that are too large for the commercial solver. The authors are successful in developing a comprehensive model that efficiently optimizes short-term mine production schedules at the block-scale. However, this formulation does not directly incorporate an optimal loader movement plan nor the availability or productivity of the mobile equipment fleet.

2.3.2 Stochastic short-term optimization models

The poor performance of deterministic LTMPS techniques is caused by the failure of estimated orebody models to adequately represent the spatial variability and presence of extreme values in mineral deposits, coupled with the non-linearity of production scheduling optimization, as discussed in Section 2.2.2.1. These adverse effects of deterministic optimization are amplified in short-term production scheduling for two major reasons: (1) more operational complexity is usually incorporated in short-term scheduling, and therefore, more uncertain systems are introduced; (2) more drilling information is available in the short-term, which tends to expose even more geological variability than previously observed with exploration data. For these reasons, it is critical to incorporate stochasticity into short-term production scheduling models. Very few works have combined uncertainty in short-term production scheduling. The following sections will review the literature on uncertain geology and uncertain equipment performance in short-term planning.

2.3.2.1 Geological uncertainty

Dimitrakopoulos and Jewbali (2013) incorporate long and short-term metal grade uncertainty in a short-term production scheduling model. The authors propose a multi-stage approach to update the orebody models and optimize the production schedule. In the first stage, high-density future grade control information is simulated using exploration data and existing grade control data from previously mined areas. The second stage updates the pre-existing simulated orebody models with the grade control simulations from stage one using a technique known as conditional simulation by successive residuals (CSSR) (Vargas-Guzmán and Dimitrakopoulos, 2002). The third and final stage of the approach is the stochastic optimization of the short-term production schedule using the updated orebody models from the second stage. The SIP model from Section 2.2.2.5, which maximizes the expected NPV and minimizes deviations from production targets, is used. The authors apply this multi-stage approach to a large gold mine where two stochastic production schedules are generated on a quarter-annual scale; one using the pre-existing orebody simulations conditioned on exploration data, and the other using the updated orebody models from the CSSR technique. The results show that the latter stochastic approach adds 36% more ore tonnes and a 70% increase in expected NPV when compared to the mine's existing schedule. Furthermore, the stochastic approach using the updated models yields a 1.6% NPV increase over the stochastic approach using the original simulated orebodies. These benefits arise from the fact that the dense grade control data indicates more high-grade values than the exploration information drilling that is traditionally used. This work provides a novel approach to incorporate useful short-term information when optimizing mine production schedules. However, the optimization model is simply a long-term scheduling formulation applied at a quarter-annual time scale; the formulation does not simultaneously optimize some of the other aspects considered short-term planning, such as fleet capacity and equipment access.

2.3.2.2 Equipment uncertainty

Random fluctuations in weather, operator habits, road conditions, and overall maintenance trends are all factors that prompt variability in mobile equipment performance. In turn, these factors influence uncertainty in the extraction capacity of a mining operation, which has a substantial impact on how short-term decisions are made. Deviations from the production schedule caused by the mobile fleet extracting more, or less, than planned will inevitably result in tonnage and grade deviations that may cause rippling effects throughout the entire LOM plan. There is very little literature on how the uncertain performance of mobile equipment can be incorporated into mine production scheduling.

In the field of automated manufacturing systems (AMS), Jain and Foley (2016) review three types of strategies available to optimize machine-job scheduling tasks to cope with unexpected machine breakdowns, lower yields, and hot-jobs. The strategies are: (1) do nothing and absorb the compound effects of deviating from the plan, (2) slightly modify the pre-existing schedule to accommodate these events as they occur, and (3) switch to a flexible dispatching approach. The authors evaluate these types of strategies based on their responsiveness and their ease of implementation. Panagiotidou and Tagaras (2007) optimize the preventative maintenance schedule of a series of units in an AMS in order to maximize the productivity at minimal cost. The authors define the operating level of each machine belonging to one of two states. The first state entails regular performance with a low likelihood of machine failure, and the second state has lower productivity, higher operating cost, and higher likelihood of failure. The authors state that aggressive maintenance schedules may reduce occurrences of failure but at the expense of overages in downtime and maintenance costs. They generate an optimal preventive maintenance schedule given the initial states of the machines, the operating state transition probabilities, and the preventive and failure maintenance costs. Unfortunately, neither of these approaches translate well to the complex nature of block-wise mine production scheduling, given the strict grade requirements, block precedence, and equipment access constraints.

Topal and Ramazan (2010) propose an MIP formulation to minimize the total maintenance cost of operating a mobile equipment fleet in an open-pit mine, assuming that these maintenance costs vary with equipment age. The authors represent the age of each piece of equipment as belonging to a discrete age bin with an associated operating cost per hour. Given the starting age of equipment, the authors optimize the scheduled trucking hours and the timing of major engine rebuilds to minimize the total maintenance costs while ensuring that the desired mine production is achieved. The methodology is tested at a gold mine, and the results show a maintenance cost decrease of 16% when compared to the mine's current practices. Topal and Ramazan (2012) extend the previous formulation to incorporate stochastic maintenance costs, and Fu et al. (2014) include a new truck purchase option to the deterministic formulation. Burt et al. (2016) adopt a similar concept to the problem of purchasing and salvaging equipment for long-term mine production, given that equipment availability and maintenance costs vary with age. Unfortunately, these works in optimal maintenance scheduling for mining equipment consider equipment performance on a long-term scale and therefore cannot be extended to short-term mine production scheduling.

2.3.2.3 Combining geological and equipment uncertainty

Matamoros and Dimitrakopoulos (2016) introduce a method that incorporates both geological and equipment uncertainty into an STMPS optimization. The authors use historical data to characterize three Gaussian equipment performance distributions: (1) hourly shovel production rates, (2) shovel and truck availabilities, and (3) truck haulage cycle times. The orebody models are generated by updating pre-existing simulations with simulated future grade-control data using CSSR (Dimitrakopoulos and Jewbali, 2013). The optimization model is formulated as an SIP that simultaneously optimizes the block-wise production schedule and the shovel movement plan under geological and equipment uncertainty. The SIP aims to maximize compliance with the long-term plan by minimizing the tonnage and grade target deviations at each destination. The optimizer also aims to minimize the deviation between available equipment times and planned production times in order to induce better utilization. The approach is used to develop a monthly schedule considering one year's worth of production at a small iron ore mine. This methodology has three major drawbacks: (1) the model only allows the monthly block-wise production to deviate below the uncertain shovel capacity. The production schedule is then ultimately bounded by the worst-case shovel scenario, negating the value of uncertain information; (2) the utilization factor in this work, measured as the ratio between scheduled production time and the available equipment time, is unrealistic. It is assumed that once a piece of equipment has completed its planned production, it will remain idle until the end of the period, artificially decreasing utilization; and (3) when comparing the equipment utilities between the stochastic and deterministic solutions, the risk profile of the stochastic solution is only compared to the expected value of the deterministic solution meaning that there is no way to induce that utilization has been decreased or rendered less variable.

2.3.3 Operationally implementable schedules

It is imperative that short-term production schedules be very practical since they usually directly precede operational planning. Shovel movement, shovel proximity, equipment access, and bench progression should all be incorporated in STMPS optimization in order to make this planning transition as seamless as possible. Aside from traditional vertical precedence constraints, encouraging blocks that are spatially connected in a bench to be mined in the same period will yield extraction patterns that are easily minable and will limit unnecessary shovel moves.

Dimitrakopoulos and Ramazan (2004) propose a formulation to encourage smooth schedules when generating LTMPS. In this approach, production schedules that contain neighboring blocks that are not extracted in the same period are discouraged by incurring penalty costs in the objective function. This approach requires introducing an additional integer decision variable for each block, whose value is influenced by the number of surrounding blocks that are not extracted in the same period. These smoothness deviation variables are then penalized in the objective function. The authors apply this formulation to a small, two-dimensional deposit and show that the resulting schedules appear smooth and practical. Benndorf and Dimitrakopoulos (2013) apply this methodology to a larger, three-dimensional iron ore mine and perform a sensitivity analysis on the smoothness penalty values. Matamoros and Dimitrakopoulos (2016) then use this approach to develop an implementable short-term production schedule at an iron ore deposit. The authors also use a fixed horizontal direction of block precedence to ensure that the progression of mining follows a pre-defined path. The additional integer variables and constraints that are required to implement this approach significantly increase the computational burden of an already cumbersome problem. This approach is not feasible when using commercial solvers for realistic problem sizes and the largest application above consists of only 3,000 blocks and 5 periods.

Mousavi et al. (2016) introduce a more mining operations-focused approach to generate practical mine designs. The authors distinguish block access into drop-cuts and side-cuts. The former being access from a bench directly above and the latter being access from within the same bench from any direction. The authors develop a formulation that allows the optimizer to decide which access path to take, given that drop-cuts are significantly more expensive. The authors are able to accommodate this added complexity since these access decision are embedded in the perturbation mechanism of the hybrid heuristic solver, described in Section 2.3.1.

2.4 Stochastic orebody modelling

For both long and short-term stochastic optimization models to be effective, it is imperative to have high quality realizations of deposits in order to adequately represent the spatial variability and uncertainty of geological phenomena. The stochastic representation of both continuous and discrete geological entities is essential in driving the optimization of the ultimate pit and the mine production schedule. Attributes that vary over a continuum, such as primary and deleterious metal grades, and densities and porosities, are critical in defining the quality of material. Discrete characteristics can be used to model mineralogical, geometallurgical, or lithological categories that may be instrumental in deciding how material should be processed. The arrangements of these attributes that comprise mineral deposits are understood to be random spatial fields. These fields can be modelled by sampling conditional distribution functions at each location. Geostatistical simulation methods are used to create multiple, equally probable realizations of a spatial random field that accurately reproduce the univariate and spatial statistics of the known information. This information is derived from location-dependant geological assays measured from drill holes. This section will cover the developments in traditional geostatistical simulation methods as well as more recent progress in the fields of multiple point and high-order techniques.

2.4.1 Sequential Gaussian simulation

Sequential geostatistical simulation algorithms rely on the decomposition of the complex, multi-variate cumulative distribution function (CDF) of an entire random field, into the product of a random sequence of marginal conditional distributions (Equation 2.2) (Rosenblatt, 1952). In other words, the simulation is generated sequentially, one point at a time. The conditional marginal distribution of one location is sampled, and this value is retained as conditioning information for the remainder of the points on the simulation grid.

$$f(z(\mathbf{x}_1), \dots, z(\mathbf{x}_N); \mathbf{x}_1, \dots, \mathbf{x}_N | \mathbf{\Lambda}_0) = \prod_{i=1}^{N} f(z(\mathbf{x}_i); \mathbf{x}_i | \mathbf{\Lambda}_{i-1}) \qquad 2.2$$

Logically, the final quality of a simulation is highly sensitive to the way in which these marginal distributions are created. Sequential Gaussian simulation (SGS) (Isaaks, 1990) is a widely popular simulation technique that assumes these conditional distributions are Gaussian and can be fully characterized by a mean and a variance. This mean and variance is computed for each location using the traditional kriging mean and variance (David, 1988). Although relatively simple to implement, SGS suffers from substantial computational limitations since a matrix, whose size is the entire simulation grid squared, is inverted to calculate the kriging weights for each point.

Dimitrakopoulos and Luo (2004) propose generalized sequential Gaussian simulation (GSGS) which modifies the SGS algorithm by simulating multiple neighboring points using one single matrix inversion. The technique relies on two major underlying observations: (1) the complete conditioning set of a point can be approximated by a very small subset of the conditioning nodes without deteriorating the quality of the realization, and (2) a local group of v adjacent nodes all share the same neighborhood of conditioning points and can be simulated simultaneously using the lower-upper (LU) decomposition method (Davis, 1987). Using a measure of relative screen effect approximation loss, the authors show that there is little to no reduction in quality of the realization when tested on various grid sizes. The authors also show that the simulation running time of GSGS is 50 times less than SGS. The major drawback of GSGS and SGS alike is the significant amount of memory that is required to store all of the simulated points. This is particularly problematic once the simulation task approaches a realistic problem size, containing millions to hundreds of millions of nodes.

Godoy (2003) improves on this group simulation approach by developing the direct block simulation (DBSIM) method. Similar to GSGS, a group of nodes, that all share the same conditioning neighborhood, is simulated simultaneously. More specifically, this group is all of the nodes that discretize one mining block. Once these discretizing nodes are simulated, they are averaged together to generate a single simulated value for the entire block. The individual nodes are discarded and the block value is retained for future conditioning instead. Not only does this directly relieve the burden on memory requirement, but the total number of conditioning points is also significantly reduced, which accelerates the search sub-routine for each subsequent node. Benndorf and Dimitrakopoulos (2005) implement this DBSIM methodology to a large deposit and show a 1000% memory requirement reduction and a 7% running time decrease when compared to GSGS. This simulation technique also tends to better preserve the frequency and connectivity of high grades at the block level due to the direct block conditioning. This is a critical aspect in driving mine production scheduling algorithms.

2.4.1.2 Joint simulation of correlated variables

It is very common to have mineral deposits consisting of multiple elements of interest. In order to accurately evaluate co-products or manage the impact of deleterious elements, it is critical in some cases to generate high quality realizations of many spatially cross-correlated attributes within one orebody model. The sequential geostatistical simulation methods above only have the capacity to model random fields consisting of one attribute. Joint conditional co-simulation of multi-variate random fields has been developed in the past (Verly, 1993). However, when dealing with real deposit sizes, containing three or more correlated variables, these methods become overly complex to model and computationally intractable to implement. An alternative approach is based on the orthogonalization of spatially correlated data to facilitate the use of efficient simulation tools that already exist. Desbarats and Dimitrakopoulos (2000) borrow the min/max autocorrelation factors (MAF) technique from signal processing (Switzer and Green, 1984) for applications in geostatistics.

This approach aims to generate spatially un-correlated service factors that are linear combinations of the original variables. These decorrelated factors are then independently simulated using any available method, and the resulting realizations are back-transformed to the original variable space, preserving the original spatial cross-correlations. Boucher and Dimitrakopoulos (2009, 2012) develop the direct block min/max autocorrelation factor simulation (DBMAFSIM) method to model large, multi-variate deposits. This method exploits the decorrelation transformation of MAF and the computational improvements of DBSIM to efficiently model industrial sized multi-variate deposits.

2.4.2 Multiple-point simulation methods

Traditional geostatistical simulation methods rely on two-point statistics, through the standard covariance model, to infer correlations of attributes through space. Journel (2007) states that these conventional two-point methods are incapable of adequately reproducing the connectivity of more complex patterns in naturally occurring deposits. A new class of multiple-point statistic (MPS) algorithms have been developed to overcome these limitations. The major differentiating factor is that the conditional distribution for a location is derived by considering multiple points simultaneously, rather than multiple pair-wise relationships, as seen in traditional simulation techniques.

Guardiano and Srivastava (1993) propose ENESIM, a sequential MPS algorithm to simulate discrete random fields. The primary challenge for multi-point techniques is the inability to collect a sufficient number of MPS from extremely sparse data sets. It is proposed that spatial continuity can be communicated through explicit, analog images, known as training images (TI). A TI is an exhaustive set of data that is assumed *a priori* and used to derive the high order spatial statistics used for simulation. For each node, the ENESIM algorithm populates a fixed template with surrounding conditioning data from drill hole information and previously simulated points. It then searches the entire TI to find replicates of this data event to derive a conditional distribution that is used to simulate the value of the node. However, the repetitive TI search is computationally expensive, and the proposed method is too inefficient to be used in real applications.

Strebelle (2002) proposes the single normal equation (SNESIM) algorithm to overcome the computational drawbacks of previous MPS simulation techniques. Rather than scanning the entire TI for each node to be simulated, the algorithm only scans the TI once at the beginning, and arranges the frequencies of all data events in a tree data structure. This allows for the rapid retrieval of each conditional distribution required to simulate the entire grid. Liu (2006) expands the SNESIM algorithm so that the univariate distributions can be explicitly controlled by skewing the local conditional distributions using a servo-system factor. The author also introduces the multi-grid simulation extension, whereby, the complete simulation is carried out sequentially at increasingly finer template sizes in order to better reproduce both large and small scale patterns. Strebelle and Zhang (2005) alter the SNESIM algorithm to accommodate for the non-stationarity of the geological domains. This is done through zonation, rotation, and stretching. The zonation approach partitions the simulation grid so that each node within a zone draws its conditional distribution from the appropriate TI. Rotation and stretching both alter the configuration of the template prior to mapping the TI in order to induce more complex geological patterns in the simulation. The primary drawback of the SNESIM algorithm is the substantial amount of memory that is required to store the pattern database in a tree structure. These memory constraints are amplified when using the multi-grid and non-stationarity extensions, which all require additional pattern trees to be retained in memory.

Straubhaar et al. (2011) propose the IMPALA simulation technique, which, unlike SNESIM, arranges the pattern database in a list structure rather than a tree. Although the database search component is slightly slower than SNESIM, there are two major advantages to this methodology: (1) the memory requirement of a list is substantially less than a tree; a list will increase linearly with template size whereas a tree will grow exponentially, and (2) the expensive task of scanning the list structure during the simulation can be parallelized very easily. This added efficiency compensates for the longer pattern database search times.

Mariethoz et al. (2010) develop the direct sampling (DS) MPS simulation technique. Unlike SNESIM or IMPALA, this particular algorithm does not retain a pattern database in memory. Instead, it searches the TI directly for each grid node. However, it does not scan the entire image, nor does it explicitly compute the conditional distribution. Instead, it randomly scans a very small subset of the TI and uses a similarity measure to directly sample a data event that is found within a certain threshold. The authors demonstrate the mathematical equivalence to SNESIM and report comparable quality and running times without any of the memory limitations. The DS algorithm is easily extended to use the multi-grid, servo-system, non-stationarity, and parallelized search methodologies. Furthermore, DS has the capacity to accommodate continuous attributes as well as multi-variate simulation. The drawback of the DS algorithm is the substantial number of input parameters that must be tuned to optimize the balance between running time and pattern reproduction.

One of the major criticisms of these MPS simulation techniques is the strong reliance on training images. The representation of complex geological phenomena is highly subjective, and the means to generate high quality TIs *a priori* is not yet fully understood. Furthermore, since the spatial relationships are derived solely from the TI, any conflicting patterns observed in the drilling data will not necessarily be reproduced in the final simulation. This is particularly problematic in the mining industry, where known information is generally far more abundant than in the oil and gas industry. Dimitrakopoulos et al. (2009); Mustapha and Dimitrakopoulos (2010, 2011) introduce high-order spatial simulation (HOSIM), a mathematically complete methodology to simulate continuous variables using high-order spatial cummulants and Legendre polynomials. This method infers information from a TI sparingly, and instead, maximizes the use of high-order statistics derived from hard data. The authors show strong potential, but the current computational requirements are too demanding for real application sizes.

Chapter 3

Risk-Resilient Mine Production Schedules Yielding Favourable Product Blends for Rare Earth Element Deposits

3.1 Introduction

The term "rare earth elements" (REE) is used to describe 17 chemically similar elements of the periodic table, including scandium, yttrium and the 15 elements of the lanthanide series. Their unique chemical and physical properties are critical in driving strategic industries such as renewable energy, communications, defence, healthcare, advanced optics and many other high tech sectors that have gained popularity in the 21st century. For the last decade, China has dominated the global production and consumption of REE products accounting for approximately 83% and 70%, respectively (Roskill, 2014). China's tremendous demand for REE is only expected to increase over the course of the next decade as they move towards high-tech industry and clean energy initiatives. In 2011, China drastically tightened their REE export quotas, causing prices to skyrocket as the risk of an impending global supply shortage became apparent. This monopoly, paired with the expected increase in REE demand, brought global attention to the imminent supply risk of these metals. To secure and maintain competitive rare earth market share, nations such as Australia, Japan, South Korea, the USA and Canada have initiated government and industry funded programs to secure resource supply, facilitate financing for domestic production and advance all downstream technologies of rare earth products (Ernst and Young, 2010).

Contrary to their name, REEs are relatively common in the earth's crust and have been known to exist in over 200 minerals. The primary challenge lies in locating them in high enough concentrations with favourable mineralogical characteristics and elemental distributions to allow for profitable extraction. Bringing a potentially viable project into production can be particularly difficult when considering the immense amount of risk that must be considered to locate, finance and develop a REE operation. One of the primary sources of risk is the geological uncertainty of REE deposits, or more specifically, the uncertainty in material grades, tonnages, mineralogy, and geology. The failure to properly address these geological risks may cause a project to deviate from expectations (Dimitrakopoulos et al., 2002), which can be catastrophic for a mining business (Vallée, 2000). The focus of this paper will be the effective assessment and integration of these risks when developing a mine production schedule that will result in more realistic forecasts and allow for more informed engineering and business decisions.

The goal when developing a long-term, mine production schedule is to identify the most profitable extraction sequence of material in a deposit throughout the life of mine (LOM). This schedule dictates the annual metal outputs and cash flows and that are critical in accurately valuing a project to attract investors. Production scheduling is typically formulated as a mixed integer program (MIP) where the sequence and destination of extracted mining units, known as blocks, are modelled as binary decision variables contributing to some objective function while subject to a series of operational constraints. Early works use exact methods to find optimal solutions to mine production scheduling problems (Gershon, 1983; Caccetta and Hill, 2003; Ramazan and Dimitrakopoulos, 2004a; Bley et al., 2010). However, these applications are limited twofold: they can be computationally intractable when scaled to realistic problem sizes; and most importantly, they rely on a single deterministic representation of the orebody. The grade values in these conventional deterministic orebody models are often estimated using some type of weighted average of all surrounding sampled points (Isaaks and Srivastava, 1989). These estimation approaches, by construction, assign values to unknown locations that minimize the global variance of the deposit. The failure to represent the in-situ grade variability inevitably leads to the production target deviations mentioned above. The geostatistical approach known as stochastic spatial simulation, offers the capacity to model these uncertainties through the generation of multiple, equally probable realizations of the deposit (Journel, 1979; Goovaerts, 1997). Unfortunately, there is no commercial software that can incorporate multiple simulated orebodies when generating mine production schedules; Ravenscroft (1992) highlights the need for new stochastic optimization techniques. Dimitrakopoulos et al. (2007) use the potential upside and downside of multiple schedules, generated using independent simulated orebodies, to choose the best mine design. Godoy and Dimitrakopoulos (2004) use multiple mining schedules to develop one risk-resilient design using an approach based on simulated annealing. Dimitrakopoulos and Ramazan (2004); Grieco and Dimitrakopoulos (2007) use a probabilistic measure to discourage the mining of blocks with a high likelihood of being below a cut-off grade. These sequential and probabilistic approaches have since been deemed ineffective in finding an optimal solution when considering geological uncertainty. Ramazan and Dimitrakopoulos (2013) propose an approach, based on two-stage stochastic integer programming (Birge and Louveaux, 1997), to develop a mine production schedule robust to uncertainty. Results show significant improvement over conventional optimization techniques by increasing the expected net present value (NPV) while decreasing the risk of not meeting key production targets. One major shortfall of this approach, and many before it, is that all material destinations must be determined prior to optimization. Menabde et al. (2007) devise an approach that allots material into discrete grade "bins". The proposed optimizer decides the associated destination allocation of these bins which results in a cut-off grade policy that can vary over time. The effectiveness of these approaches that rely on exact mathematical solving techniques is still limited by the size of the problem. This can be overcome by using a class of methods to solve integer programs known as metaheuristics. By intelligently exploring a large, complex solution space, these algorithms have shown to find high-quality solutions to the mine production scheduling problem in reasonable amounts of time (Lamghari and Dimitrakopoulos, 2012; Montiel and Dimitrakopoulos, 2015; Goodfellow and Dimitrakopoulos, 2016). In addition, these methods can efficiently handle non-linear constraints that must be added when stockpiling and non-linear recovery functions are introduced.

For the stochastic mine planning approaches to be effective, it is imperative to generate high quality realizations of the deposit. Like most other commodities, the common practice to model REE deposits begins by assaying core samples for each of the metals of interest. This information is used to inform the remainder of the deposit which is bounded by some geological domain. Logically, the final representation of the deposit is very sensitive to the methods that describe the final extents of the domain, as well as the approach used to interpolate the material grades within. Traditionally, orebody boundaries are created using some attribute from the sampled data to trace the outlines of the material domains. This can capture the large-scale trends, but will be locally inaccurate, and thus will result in a subjective, overly-smoothed interpretation of the mineralized zones (Journel, 2007). Simulating the boundaries of geological domains allows for the quantification of the volumetric uncertainty. However, modelling algorithms that only consider two-point spatial statistics, fail to reproduce complex curvilinear patterns present in a naturally occurring deposit (Guardiano and Srivastava, 1993). Simulation algorithms that simultaneously consider multiple points can offer a more realistic reproduction of geological boundaries. The technique known as single normal equation simulation (SNESIM) (Strebelle, 2002) is developed to efficiently simulate discrete geological entities using multiple point statistics (MPS). Due to the strong correlations between REEs of similar atomic size, it is imperative to respect and reproduce these naturally occurring cross-elemental relationships. Joint conditional simulation methods can be used to quantify the uncertainty surrounding the spatial relationships of correlated variables. However, when dealing with large deposits containing three or more correlated variables, these methods become overly complex to model and computationally intractable to implement. An alternative approach that aims to decompose a set of correlated variables into uncorrelated "factors" can be used instead. The geostatistical application of this approach, based on principle component analysis (PCA) (David, 1988; Goovaerts, 1993; Wackernagel, 2003), is limited as it only guarantees decorrelation at a lag distance of zero. Alternatively, min/max autocorrelation factors (MAF) (Switzer and Green, 1984; Desbarats and Dimitrakopoulos, 2000; Boucher and Dimitrakopoulos, 2009) transforms the original variables into a set of orthogonal factors that are decorrelated at all lag distances. These factors are then simulated independently and back-transformed to the initial variable space, preserving the original spatial relationships.

This work explores the process of simulating REE deposits and optimizing the annual mine production schedule while integrating geological uncertainty. A two-step simulation approach is used to model both the volumetric and grade uncertainty (Goodfellow et al., 2012). First, the geological boundaries of the deposit are simulated using SNESIM, then MAF is used to jointly simulate all 15 of the REEs contained within the simulated domains. The LOM production schedule is generated using a multi-element SIP formulation with variable destination decisions that is solved using a Multi-Neighbourhood Tabu Search (MNTS) metaheuristic (Lamghari and Dimitrakopoulos, 2012; Senécal and Dimitrakopoulos, 2014). The remainder of this chapter is organized as follows. First, the conventional and stochastic approaches to model REE deposits are reviewed and then the formulation of the SIP and the details of the optimization method are described. The stochastic modelling and optimization approach is then applied to a case study at an open pit, monazite-bastnäsite REE operation. The stochastic design is benchmarked against the best-known conventional practices to demonstrate the advantages of designing a production schedule that integrates geological uncertainty.

3.2 Modelling rare earth element deposits

3.2.1 Conventional approach

There has been very little published on how to effectively model REE deposits and even less in terms of mine optimization. This is in part due to the fact that there are only a handful of REE focused projects currently operating around the world. Like most other commodities, the common practice is to assume fixed boundaries of geological domains to act as the extents of some estimation technique for the metal grades within. However, since every REE deposit contains all of the REEs in varying concentrations, it can be tedious to model each of them, given that they are strongly correlated (Figure 3.1). One particular convention widely used in the REE industry is the use of a cumulative representation of ore quality known as the total rare earth oxide (TREO) grade. This simplified grade is derived from a weighted summation of the individual REEs in a sample as seen in Equation 3.1. The material quality is often estimated using this TREO grade which has major ramifications throughout the later processes of mine optimization. The primary drawback of using this convention is that the locations and magnitudes of the less concentrated, and more valuable REEs are lost due to their smaller contribution to the TREO grade. This loss of resolution leads to a misguided understanding of the eventual ultimate pit limit, cut-off grade, and production schedule, indicating that this convention should be avoided all together. For REE producers that aim to mine and sell their product as a bulk ore concentrate at a fixed TREO basket price, this approach may suffice. However, for producers that require more granularity to accurately declare their reserves or forecast the production of individual REE to achieve a strict blend, it is necessary to improve how the ore body is modelled.

$$TREO^{grade} = \sum_{e \in REE} REE_e^{grade} \cdot Oxide_e^{weight} \qquad 3.1$$

Extensions of this cumulative grade representation might be the separation of rare earths into various classes to group the elemental distribution according to similar properties, similar abundancy, and subsequently, similar market value. One of the more common groupings involves splitting the light rare earth elements (LREE), consisting of lanthanum to gadolinium, from the remaining heavy rare earth elements (HREE), terbium to lutetium. Whether the TREO grade (Commerce Resources Corp, 2012; IAMGOLD, 2012; Molycorp Inc., 2012) or these groupings (Matamec, 2013) are used to estimate the quality of the orebody, both approaches will inevitably require some linear scaling to eventually assess individual REE content to forecast reserves. Unfortunately, scaling naturally infers perfect correlation which is unrealistic (Figure 3.1). Another common approach to interpolate REE grades is to independently estimate each of the elements (Avalon Rare Metals Inc., 2013; Quest Rare Minerals Ltd., 2014; Great Western Minerals Group Ltd., 2014). Since all REEs in the sample data are highly correlated, by construction, an estimation technique will be sufficient in reproducing the cross-element relationships assuming that the each elements spatial statistics do not vary drastically. Although it is an improvement over the grouping and scaling methods, any estimation technique will ignore the in-situ variability of grades throughout the deposit. Because of this, it should be sought to perform a joint conditional simulation to model the grades of all REEs in a deposit. The following sections will cover the algorithms necessary to implement the two step stochastic simulation approach to model the volumetric and grade uncertainties pertaining to any typical REE deposit.

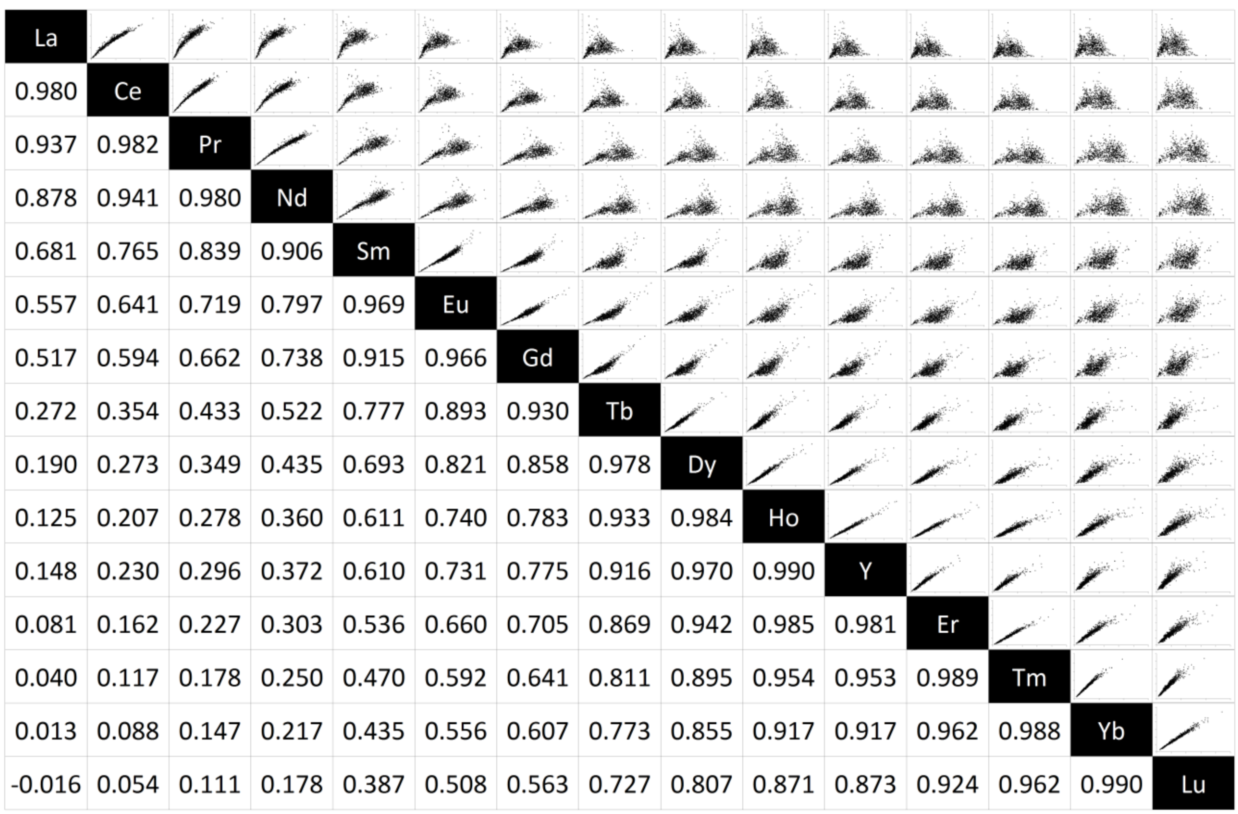

Figure 3.1 - Cross plots (upper-diagonal) and Pearson correlation coefficients (lower-diagonal) of sampled REE grades

3.2.2 Multiple point simulation with SNESIM

3.2.2.1 Outline

Multiple point statistics (MPS) consider the joint neighbourhood of n points, centered around a location, $\mathbf{x}_0$. The arrangement of these points is known as a template, $\boldsymbol{\tau}_n$, which is defined by a set of vectors, $\mathbf{h}_\alpha$ $\alpha = 1 \ldots, n$. The realization of the n values within the template is known as a data event, $\mathbf{d}_n$, and these values are represented by an attribute s taking on one of K possible discrete states. These states will likely represent types of lithologies or metallurgical ore types. An example of a data event on a template of size, $n = 4$, can be seen in Figure 3.2. A brief outline of the SNESIM algorithm can be found below and a more detailed explanation can be found in (Strebelle, 2002).

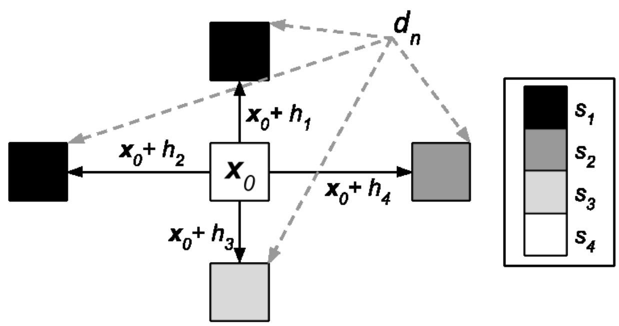

Figure 3.2 - Configuration of a data event. Conditioning points centered on location, $\mathbf{x}_0$

3.2.2.2 General algorithm steps

1. Scan the training image and store all occurrences of all data events in a tree structure.
2. Define a random path to visit each node in the simulation grid.
3. For each node i in the random path:
    a. Retrieve the data event of node i, populated with conditioning data and previously simulated data, and search for occurrences of this data event in the tree.
    b. Derive the local probability distribution of central values. The probability of finding a certain categorical state given the surrounding data event is given by Bayes' relation for conditional probability.
    c. Randomly sample the conditional cumulative distribution function and add the simulated node to the grid.
4. Repeat all previous steps for the desired number of realizations.

3.2.3 Joint conditional simulation using MAF

3.2.3.1 Outline

Let $\mathbf{Z}(\mathbf{x}) = [z_1(\mathbf{x}) \dots, z_p(\mathbf{x})]^T$ denote a multivariate, stationary and ergodic random field over a region, representing p correlated, continuous variables. Now consider its point-wise standard Gaussian transformation, $y_i(x) = \phi(z_i(\mathbf{x}))$, which has zero mean and unit variance:

$$\mathbf{Y}(\mathbf{x}) = [y_1(\mathbf{x}) \dots, y_1(\mathbf{x})]^T \qquad 3.2$$

The decorrelation at lag zero is achieved by generating the PCA factors, $\mathbf{F}_{PCA}(\mathbf{x})$:

$$\mathbf{F}_{PCA}(\mathbf{x}) = \boldsymbol{\Lambda}_1^{-\frac{1}{2}} \mathbf{Q}_1 \mathbf{Y}(\mathbf{x}) \qquad 3.3$$

where $\boldsymbol{\Lambda}_1$ and $\mathbf{Q}_1$ are from the spectral decomposition of the covariance matrix, $\mathbf{B}$, of $\mathbf{Y}(\mathbf{x})$.

$$\mathbf{B} = \mathbf{Q}_1^T \boldsymbol{\Lambda}_1 \mathbf{Q}_1 \qquad 3.4$$

To ensure the spatial decorrelation at all lag distances, a second data transformation is applied:

$$\mathbf{F}_{MAF}(\mathbf{x}) = \mathbf{Q}_2 \mathbf{F}_{PCA}(\mathbf{x}) \qquad 3.5$$

where $\mathbf{Q}_2$ is derived from the spectral decomposition of the experimental cross-covariance matrix, $\boldsymbol{\Gamma}(\Delta)$, of $\mathbf{F}_{PCA}(\mathbf{x})$, given some experimental lag distance, Δ.

$$\boldsymbol{\Gamma}(\Delta) = \mathbf{Q}_2^T \boldsymbol{\Lambda}_2 \mathbf{Q}_2 \qquad 3.6$$

In this data-driven decorrelation approach (Rondon, 2012), it is common to test many values of Δ and select the one that yields the best spatial decorrelation. The factors from Equation 3.5 are, by definition, spatially uncorrelated at all lag distances (Desbarats and Dimitrakopoulos, 2000) meaning that each one can be simulated independently. Furthermore, since the factors are normally distributed, any efficient sequential Gaussian simulation method can be utilized (Isaaks, 1990; Dimitrakopoulos and Luo, 2004). In this work, a computationally efficient method known as direct block simulation (DBSIM) (Godoy, 2003) is used. A more detailed description of the decorrelation procedure and the integration with DBSIM can be found in Boucher and Dimitrakopoulos (2009).

3.2.3.2 General algorithm steps

1. Transform the data, $\mathbf{Z}(\mathbf{x})$, to normal scores, $\mathbf{Y}(\mathbf{x})$.
2. Orthogonalize $\mathbf{Y}(\mathbf{x})$ using the MAF transformation to obtain $\mathbf{F}_{MAF}(\mathbf{x})$.
3. Define a random path visiting each block to be simulated.
4. For each block in the random path:
   a. Define a random path visiting each point in the block.
   b. Simulate each of the factors at each point.
   c. For each factor, take the average of all points within the block and retain for further conditioning in the MAF space.
   d. Back-transform the simulated factors to the original data space.
5. Repeat steps 3 and 4 for the desired number of realizations.

3.3 Stochastic long-term mine production scheduling

The following two-stage stochastic recourse model offers a means to solve the long-term mine production scheduling (LTMPS) problem with uncertain geology. In general, this two-stage approach is a type of SIP model that integrates both anticipative and adaptive mechanisms to deal with uncertain input. First stage (anticipative) decisions involve making decisions prior to the unveiling of uncertainty, whereas second stage (adaptive) decisions, are made once uncertain parameters have been observed. In mining terms, this can be thought of as making extraction sequence decisions independent of scenario, whereas the risk of deviating from production targets can be controlled on a per-scenario basis acting as a corrective strategy (Dimitrakopoulos and Ramazan, 2008).

The first stage decisions in this formulation are when blocks should be extracted, and where they should be sent. The latter decisions are particularly useful when considering a complex, multi-element deposit. All REEs have unique concentrations, market values and associated operational constraints. This added complexity makes it substantially more difficult to derive a multi-element cut-off grade policy. Informing the optimizer with: the uncertainty about the grades of each element in a block, the various blending requirements, and the state of the mining and processing capacities, allows for the derivation of a more robust destination policy. The following SIP formulation is based on the model introduced in Ramazan and Dimitrakopoulos (2013) with modifications necessary to expand the problem to a multi-element operation (Benndorf and Dimitrakopoulos, 2013).

3.3.1 Model formulation

3.3.1.1 Indices

t is a time period, $t = 1 \dots, T$

i is a block in the mine, $i = 1 \dots, N$

p is a processing destination, $p = 1 \dots, P$

ε is a metal element, $\varepsilon = 1 \dots, E$

s is a geological scenario, $s = 1 \dots, S$

3.3.1.2 Model parameters

Ton_i is the tonnage of block i

$G_{i\varepsilon s}$ is the grade of element ε in block i in scenario s

M_t^{max} is the maximum capacity for tonnage mined in period t

D_{tp}^{tar} is the target for tonnage processed at destination p in period t

$G_{tp\varepsilon}^{tar}$ is the target for the metal grade of element ε at destination in period t

dr^f is the financial discount rate, $dr^f \in [0,1]$

dr^g is the geological discount rate, $dr^g \in [0,1]$

MC is the mining cost per tonne

PC_p is the processing cost per tonne at destination p

$R_{p\varepsilon}$ is the recovery of element ε at destination p, $R_{p\varepsilon} \in [0,1]$

π_ε is the selling price of element ε

α_{tips} is the discounted profit from processing block i at destination p in period t for scenario s

$$\alpha_{tips} = \frac{1}{(1 + dr^f)^t} \cdot \left(\left(\sum_{\varepsilon=1}^{E} Ton_i \cdot G_{i\varepsilon s} \cdot R_{p\varepsilon} \cdot \pi_\varepsilon \right) - PC_p \cdot Ton_i \right) \qquad 3.7$$

β_{ti} is the discounted cost of mining block i in period t

$$\beta_{ti} = \frac{1}{(1 + dr^f)^t} \cdot (MC \cdot Ton_i) \qquad 3.8$$

C_p^{D-}, C_p^{D+} are the lower/upper target deviation costs for tonnage at destination p

C_{tp}^{D-}, C_{tp}^{D+} are the lower/upper target deviation costs for tonnage at destination p in period t following the application of the geological discount rate dr^g:

$$C_{tp}^{D\mp} = \frac{1}{(1+dr^g)^t} \cdot \left(C_p^{D\mp}\right) \qquad 3.9$$

$C_{p\varepsilon}^{G-}$, $C_{p\varepsilon}^{G+}$ are the lower/upper grade target deviation costs for element ε at destination p

$C_{tp\varepsilon}^{G-}$, $C_{tp\varepsilon}^{G+}$ are the lower/upper grade target deviation costs for element ε at destination p in period t following the application of the geological discount rate dr^g:

$$C_{tp\varepsilon}^{G\mp} = \frac{1}{(1+dr^g)^t} \cdot \left(C_{p\varepsilon}^{G\mp}\right) \qquad 3.10$$

$\rho(i)$ is the set of blocks that precede block i

3.3.1.3 Decision variables

$x_{ti} = \begin{cases} 1 & \text{if block } i \text{ is extracted during period } t \\ 0 & \text{otherwise} \end{cases}$

$y_{tip} = \begin{cases} 1 & \text{if block } i \text{ is processed at destination } p \text{ during period } t \\ 0 & \text{otherwise} \end{cases}$

d_{tp}^{D-}, d_{tp}^{D+} are the shortage and surplus tonnage deviations at destination p in period t

$d_{tp\varepsilon s}^{G-}$, $d_{tp\varepsilon s}^{D+}$ are the grade deviations for element ε at destination p in period t in scenario s

3.3.1.4 Objective function – Maximize:

$$+\frac{1}{S}\sum_{t=1}^{T}\sum_{i=1}^{N}\sum_{p=1}^{P}\sum_{s=1}^{S}\alpha_{tips}\cdot y_{tip} - \sum_{t=1}^{T}\sum_{i=1}^{N}\beta_{ti}\cdot x_{ti} \qquad \text{Part 1}$$

$$-\sum_{t=1}^{T}\sum_{p=1}^{P} C_{tp}^{D-}\cdot d_{tp}^{D-} + C_{tp}^{D+}\cdot d_{tp}^{D+} \qquad \text{Part 2} \qquad 3.11$$

$$-\frac{1}{S}\sum_{t=1}^{T}\sum_{p=1}^{P}\sum_{\varepsilon=1}^{E}\sum_{s=1}^{S} C_{tp\varepsilon}^{G-}\cdot d_{tp\varepsilon s}^{G-} + C_{tp\varepsilon}^{G+}\cdot d_{tp\varepsilon s}^{G+} \qquad \text{Part 3}$$

The objective function (Equation 3.11) of this SIP model aims to maximize the expected NPV and minimize the costs of deviating from planned production targets. Part 1 refers to the total expected discounted cash flows from mining and processing material. Part 2 accounts for the total tonnage target deviations at each destination. Part 3 accounts for the total expected grade deviations of each element at each destination. Parts 2 and 3 incorporate the geological risk management parameters; a mechanism known as geological discounting (Dimitrakopoulos and Ramazan, 2004) seen in Equation 3.9 and Equation 3.10. Notice that these deviation cost parameters are indexed by time. By discounting these costs across the time horizon, the optimizer is encouraged to defer riskier production later into the project life where it is assumed that corrections can be made more easily.

3.3.1.5 Constraints

Constraints 3.12 ensure that a block can be sent to exactly one destination, but only if it has been extracted in the same period.

$$\sum_{p=1}^{P} y_{tip} = x_{ti} \quad \forall\, t,\; i \tag{3.12}$$

Constraints 3.13 ensure that a block can only be extracted once throughout the LOM.

$$\sum_{t=1}^{T} x_{ti} \leq 1 \quad \forall\, i \tag{3.13}$$

Constraints 3.14 inform the model that in order to extract any block, all of its overlying blocks must be extracted first.

$$x_{ti} \leq \sum_{\tau=1}^{t} x_{\tau j} \quad \forall\, t,\; i,\; j \in \rho(i) \tag{3.14}$$

Constraints 3.15 ensures that the mining capacity is not exceeded in any period.

$$\sum_{i=1}^{N} Ton_i \cdot x_{ti} \leq M_t^{max} \quad \forall\, t \tag{3.15}$$

Constraints 3.16 are used to determine the shortage and surplus tonnage deviations from the production targets at each destination in each period.

$$\sum_{i=1}^{N} Ton_i \cdot y_{tip} + d_{tp}^{D-} - d_{tp}^{D+} = D_{tp}^{tar} \quad \forall\, t,\ p \tag{3.16}$$

Constraints 3.17 are used to determine the deviations above or below the grade production targets at each destination in each period across all geological scenarios.

$$\sum_{i=1}^{N} Ton_i \cdot (G_{i\varepsilon s} - G_{tp\varepsilon}^{tar}) \cdot y_{tip} + d_{tp\varepsilon s}^{G-} - d_{tp\varepsilon s}^{G+} = 0 \quad \forall\, t,\ p,\ \varepsilon,\ s \tag{3.17}$$

Constraints 3.18 to 3.21 enforce the variable bounds.

$$x_{ti} \in 0 \text{ or } 1 \quad \forall\, t,\ i \tag{3.18}$$

$$y_{tip} \in 0 \text{ or } 1 \quad \forall t, i, p \tag{3.19}$$

$$d_{tp}^{D-}, d_{tp}^{D+} \geq 0 \quad \forall t, p \tag{3.20}$$

$$d_{tp\varepsilon s}^{G-}, d_{tp\varepsilon s}^{G+} \geq 0 \quad \forall t, p, \varepsilon, s \tag{3.21}$$

3.3.2 Solution approach: Multi-neighbourhood Tabu Search

When the LTMPS problem approaches realistic problem sizes, conventional mathematical programming solvers become less viable options due to their prolonged solving times. For this reason, approximate methods known as metaheuristics are used. The Multi-neighbourhood Tabu Search (MNTS) method is used in this work (Glover and Laguna, 1997; Lamghari and Dimitrakopoulos, 2012; Senécal and Dimitrakopoulos, 2014). The MNTS method uses basic move operators to explore the solution space to gradually improve the objective function of the proposed model. Pre-defined combinations of these moves are defined as neighbourhoods of the current solution vector. There are four different neighbourhoods considered in this case: (1) changing a block's extraction period without violating precedence, (2) changing a block's destination, (3) changing a block's destination and extraction period without violating precedence, and (4) swapping the destinations of two blocks extracted in the same period. At each iteration, all four neighbourhoods are explored and the neighbour that results in the best change in objective value is accepted, and the solution vector is updated. The most important feature of this solving method is what is referred to as the Tabu list. A list of previously accepted neighbours is kept in memory which are rendered "tabu" and are removed from contention in the succeeding iterations. These tabu moves allow the solver to avoid cycling and to escape local optima. More detailed information pertaining to the efficient parallelization, the diversification strategy and other mechanisms of this solution method can be found in Senécal and Dimitrakopoulos (2014).

3.4 Application – Jupiter REE project

To demonstrate the effectiveness of methods described herein, an application is carried out at a fictitious REE operation known as the Jupiter project. The following section will cover the two-step simulation of the Jupiter deposit followed by the stochastic optimization of the long-term production schedule.

3.4.1 Simulation of the Jupiter deposit

The features of the Jupiter orebody are based on a real, REE deposit which will remain undisclosed for confidentiality purposes. The deposit is predominantly abundant in LREE with an area of higher HREE concentration hosted near the surface, covered by a very thin layer of overburden. Almost all of the deposit's REEs are hosted by monazite and bastnäsite, which are two of the most common and historically exploitable REE bearing minerals. The data used for this study is comprised of 21 exploration drill holes, spaced roughly 50 m apart, containing approximately 8,000 core samples at an average length of 1 m. The samples are assayed for yttrium and 14 of the 15 lanthanide elements: lanthanum (La), cerium (Ce), praseodymium (Pr), neodymium (Nd), samarium (Sm), europium (Eu), gadolinium (Gd), terbium (Tb), dysprosium (Dy), holmium (Ho), erbium (Er), thulium (Tm), ytterbium (Yb) and lutetium (Lu). In addition, the ratio of monazite to bastnäsite is also assessed throughout the deposit. 20 geological domain simulations are generated using SNESIM (Section 3.2.2) and each one acts as the unique extents of one joint conditional simulation of REE grades using MAF (Section 3.2.3). The simulation of the monazite-bastnäsite ratio is done independently using sequential indicator simulation (SIS).

3.4.1.1 Simulation of geological boundaries

A traditional wireframe serves as a training image (TI) to encompass the large scale geometric shapes and features of the deposit. This TI is generated by first coding the borehole samples as inside or outside the interpreted domain. The envelope is traced and converted to a block model on a regular grid where all blocks lying within the wireframe are considered to be mineralized, and blocks lying outside are considered to be barren rock. The SNESIM algorithm is implemented, given the coded borehole samples as conditioning data on point support and the block model as the TI. The simulation grid measures 630, 850 and 840 m in the East-West, North-South and vertical directions respectively, and the grid is discretized by blocks measuring 10 x 10 x 10 m. The massive nature of the deposit allows for relatively consistent and simple patterns about the boundary of the domain, thus, the reproduction of the simulations are relatively similar in terms of their total volume. The distribution of the simulated volumes can be seen in Figure 3.3. The results show that approximately 40% of the simulations are larger than the TI and 60% are smaller in terms of total mineralized volume. Figure 3.4 compares a cross section of the TI to the smallest, median, and largest domain simulations after using a post-processing algorithm to remove erroneous mineralized blocks from the outer corners of the simulation grid. The large scale features of the training image are preserved, however, the simulated domains show the presence of spatial variability at the local scale. After quantifying the volumetric uncertainty of the mineralized domain, the uncertainty pertaining to spatial continuity of REE grades is assessed through a joint conditional simulation within each of these realizations using DBMAFSIM.

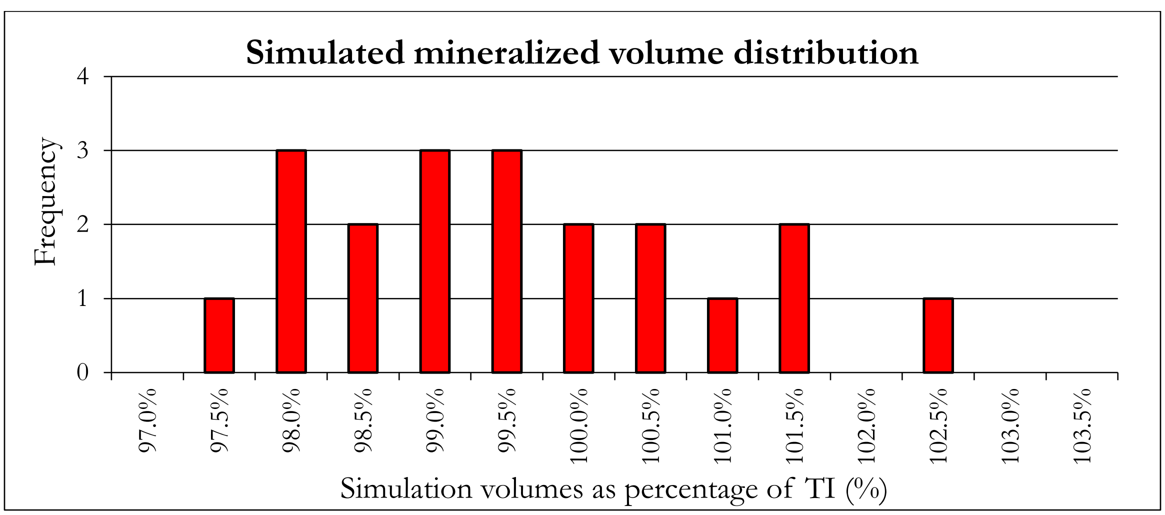

Figure 3.3 - Frequency distribution of simulated mineralized volumes expressed as percentage of the TI

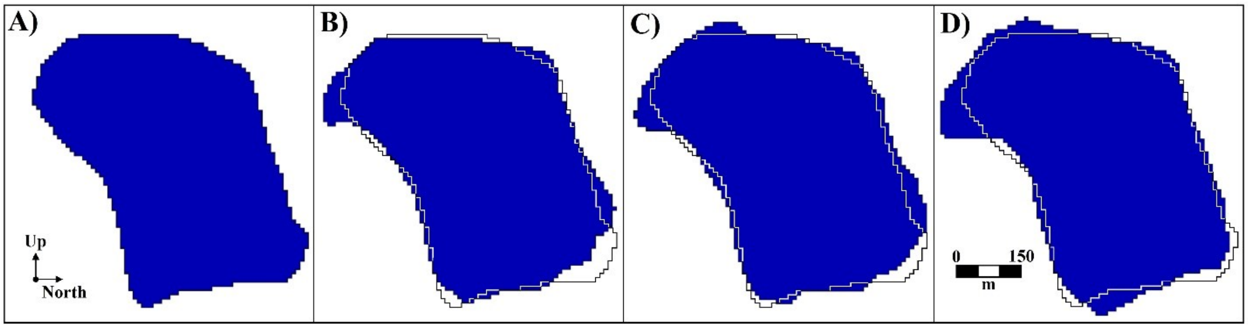

Figure 3.4 - Vertical cross-sections of the TI (A) and the minimum (B), median (C), and maximum (D) mineralized domain simulations. Outline of TI is included for reference

3.4.1.2 Joint simulation of REE grades

The samples are declustered to remove any bias introduced by concentrating drilling in areas of higher grades, which is not uncommon in most exploration programs. The conditioning points are independently filtered through each of the domain simulations resulting in 20 unique data sets. In this case, the REE grades are partitioned into three segments. Groups A (La to Nd), B (Sm to Gd), and C (Tb to Lu) are treated independently to simplify the decorrelation process and the eventual validation steps without compromising the most relevant relationships. This simulation approach is not limited by the number of variables considered, as long as the decorrelation is deemed satisfactory. It is encouraged to jointly simulate all of the REEs together for the most accurate representation of the orebody, but this segmented approach is sufficient for this application. Following a normal score transformation of the input data, an experimental lag distance, Δ, of 75 $\pm$ 35 m is used to ensure adequate decorrelation. An example cross-variogram of two of the group A service factors can be seen in Figure 3.5. The proximity to zero suggests that these factors are well-decorrelated at all lag distances.

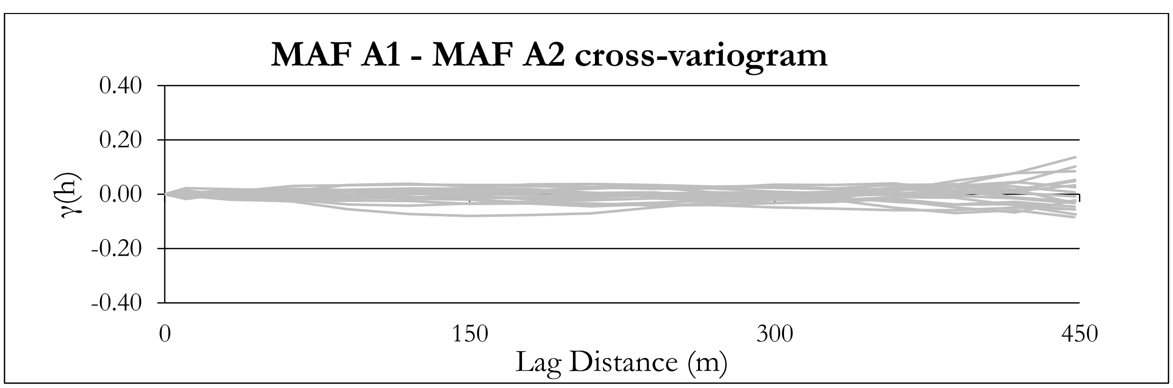

Figure 3.5 - Cross-variogram of transformed input data factor A1 and factor A2

DBSIM is performed on each of the service factors at a block discretization of 5 x 5 x 5 m. Due to the varying extents of the mineralized zones from Section 3.4.1.1, the number of simulated nodes varies between 635 and 660 thousand. The appropriate back-transformation is applied to the simulated factors to yield the final simulated points in the original data space. The simulated values from groups A, B and C are combined and re-blocked to produce the final simulated orebodies.

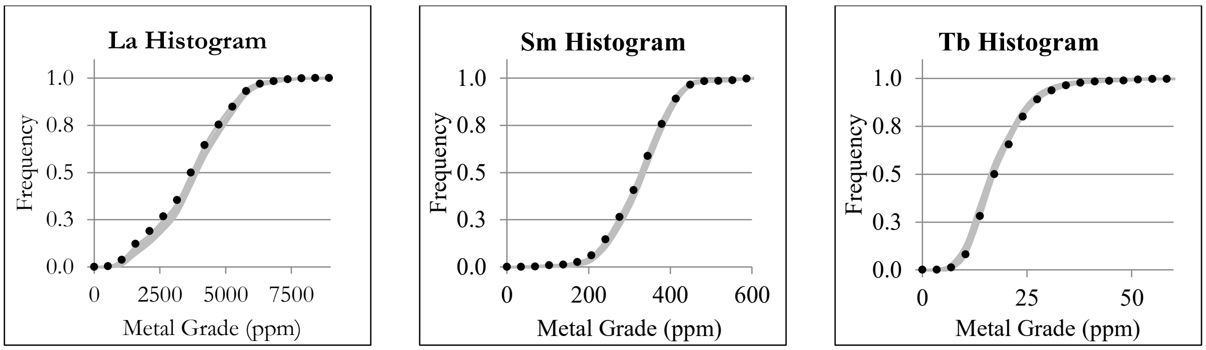

Figure 3.6 - Histograms for lanthanum, samarium and terbium grades

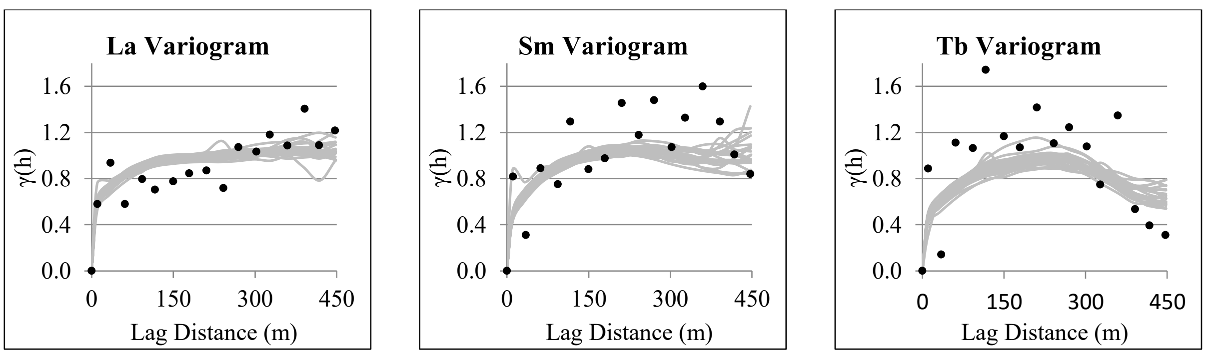

Figure 3.7 - Standardized variograms for lanthanum, samarium and terbium grades

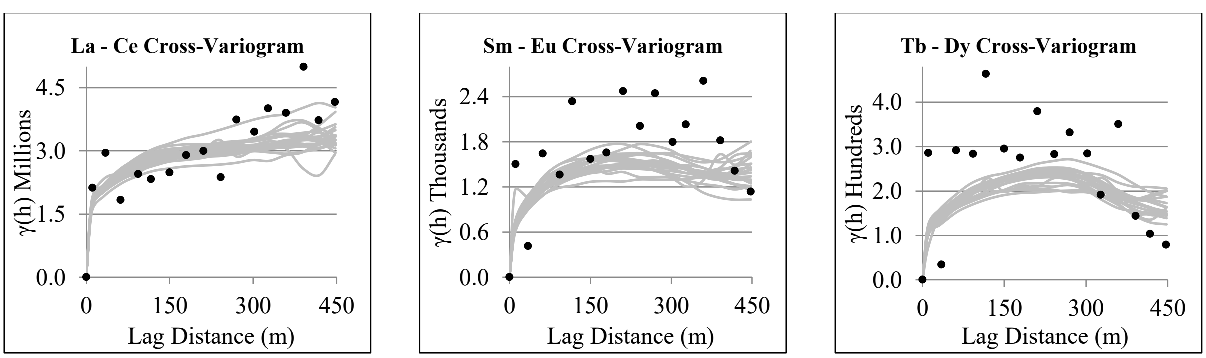

Figure 3.8 - Cross-variograms for lanthanum and cerium, samarium and europium and terbium and dysprosium

Figure 3.6, Figure 3.7 and Figure 3.8 demonstrate the reproduction of the spatial continuity of the conditional simulations with histograms, variograms and cross-variograms respectively. The black dots represent the original data, and the light grey lines represent the simulations on point support. The statistical validation demonstrates that the simulation is successful in reproducing the univariate and bivariate spatial relationships of the conditioning information. It should be noted that the conditioning data is relatively sparse in terms of pairings found, which results in its sporadic nature. Figure 3.9 demonstrates the varying spatial continuity of REE grades between different simulations. The warmer colours represent higher grades for dysprosium across three example simulations. Figure 3.10 also validates that the simulation approach is successful in reproducing the spatial cross-element relationships within one realization. The areas of concentrated high grades are similar amongst lanthanum and cerium (LREE), however, once compared to dysprosium (HREE), the locations of the concentrated zones become dissimilar (Figure 3.1).

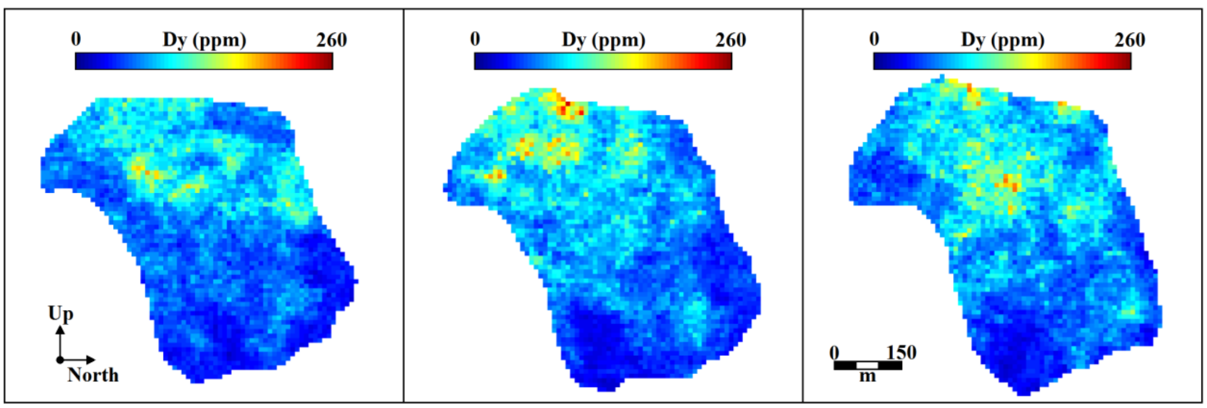

Figure 3.9 - Dysprosium grades for three different simulations: SIM11 (left), SIM08 (middle), SIM17 (right)

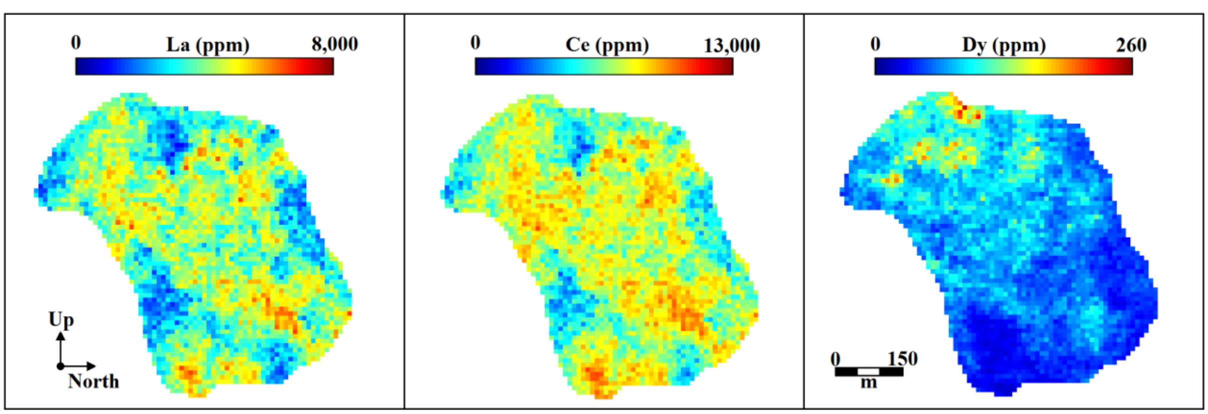

Figure 3.10 - Simulated lanthanum (left), cerium (middle) and dysprosium (right) grades for one simulation

3.4.2 Long-term mine production schedule of the Jupiter project

3.4.2.1 Outline

This open-pit operation consists of one mine, and two destinations: a waste dump and a processing facility that produces four distinct, REE products. The annual mining and processing rates are gradually increased throughout the first three years of production to emulate fleet and plant expansions. The ore material delivered from the mine is crushed prior to flowing through a series of flotation cells to recover both monazite and bastnäsite. The flotation process has been designed to handle specific concentrations of each of these minerals. To ensure optimal recovery, the run-of-mine (ROM) ore feed should fall within these designed grade ranges. The concentrate from this stage then moves to an acid cracking step, where a bulk rare earth oxide solution is purified. The solution then flows through four sequential separation stages to produce the four Jupiter products as seen in Figure 3.11: (1) a heavy rare earth concentrate, (2) a praseodymium and neodymium (didymium) mixture, (3) a cerium product, and (4) a lanthanum product. Each product stream has its own unique metal recovery. Once separated, the HREE product is sampled for its individual REE content and is sold accordingly. It is assumed that all of this product, along with the La and Ce products, can be sold in full at fixed prices. However, there is a very stringent offtake agreement with a customer purchasing the didymium product. The agreement states that there are very tight requirements on the quantity of product that is purchased annually, and that the quality of the blend in terms of the ratio of Pr to Nd is to be heavily controlled. Breaching the terms of this agreement will result in various penalty fees.

The goals of the optimization are to: maximize the discounted value of the mine, achieve the ore production targets, and mitigate the risk of the mill feed deviating from the desired mineral concentrations and REE blend (Equation 3.11). The geological risk is quantified using a set of 20 orebody simulations from Section 3.4.1. This stochastic approach is benchmarked against the best known conventional practices, namely Geovia's Whittle mine optimization package using an estimated ore body model as input. The original TI of the Jupiter deposit is used to bound the mineralized domain of this deterministic model, and each of the REE grades and the monazite-bastnäsite ratio are independently interpolated using ordinary kriging. The schedules are compared using annual key performance indicators (KPI), such as: NPV, annual tonnages mined and processed, and ROM mineral contents, and REE grades.

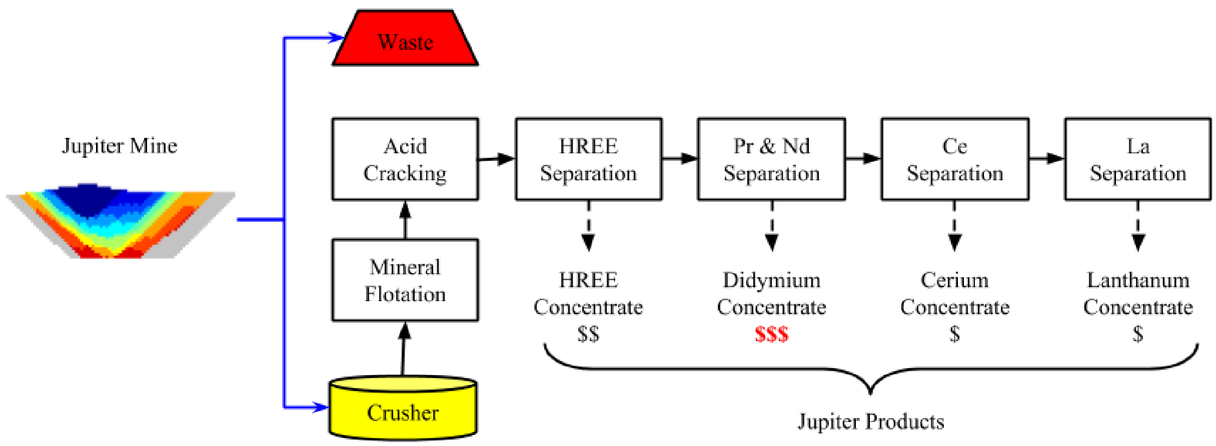

Figure 3.11 - Material destinations and the procedure for creating the Jupiter products

3.4.2.2 Jupiter operational parameters

The annual grade targets for Pr and Nd are derived from the annual didymium production agreement of 24 thousand tonnes per year, the required Pr to Nd ratio of 78% Nd, and the ore production target of 7 million tonnes per year. Capital costs are not included in the NPV assessment of the Jupiter project. The capital requirements would be identical for both the stochastic and deterministic designs, and therefore, would not affect the comparison. Since the capital costs are omitting, the cost of capital, (discount rate) is held constant throughout the entire time horizon.

Table 3.1 - Operational and financial parameters of the Jupiter project

Parameter	Value	Parameter	Value
Periods	11 years	Monazite range	2.1 – 2.3 %
Slope angle	45°	Bastnäsite range	0.365 – 0.395 %
Blocks	106,487	Praseodymium target	750 ppm
Density	2.98 t/m³	Neodymium target	2,700 ppm
Mining capacity	20 – 35 Mt/year	Flotation recovery	70 %
Processing target	4.5 – 7.0 Mt/year	Cracking recovery	95 %
Financial discount	10 %	HREE recovery	92 %
Geological discount	20 %	Didymium recovery	95 %
Mining cost	6.50 $/t	Lanthanum recovery	93 %
Processing cost	112.00 $/t	Cerium recovery	91 %

Table 3.2 - REE metal prices in dollars (Canadian) per pound

Element	Price $CAD/lb	Element	Price $CAD/lb	Element	Price $CAD/lb
Yttrium	16.20	Samarium	6.20	Holmium	320.00
Lanthanum	6.24	Europium	450.00	Erbium	275.00
Cerium	6.20	Gadolinium	25.00	Thulium	500.00
Praseodymium	60.60	Terbium	385.00	Ytterbium	360.00
Neodymium	72.80	Dysprosium	216.00	Lutetium	650.00

3.4.2.3 Stochastic optimization results

The stochastic solution to the SIP formulation is found using the parallelized MNTS method described in Section 3.3.2 on an Intel Xeon 5650 (2.67 GHz) with a total of 24 threads and 24 GB of RAM. A crude conventional schedule from Whittle is used as an initial solution. The final solution was found in approximately eight hours. Figure 3.12 shows a cross-section displaying the general progression of pit development. The schedule obeys all slope constraints and yields relatively mineable shapes that would require minor corrections to be operationally implementable.

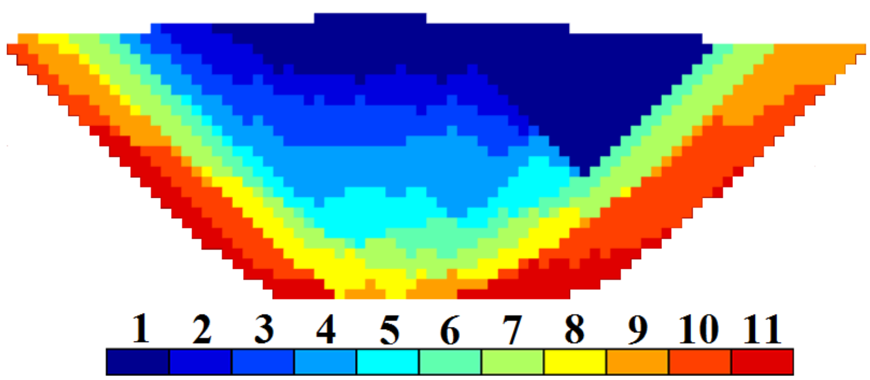

Figure 3.12 - East-West cross-section of the stochastic schedule. Colours represent the period of extraction

Block destination decisions are modelled to be robust to uncertainty in this formulation. As a result, the annual tonnages mined and processed are consistent across all geological scenarios, as shown in Figure 3.13. Alternatively, the mineral and REE blend will vary, and is shown using the individual lines in Figure 3.14. Figure 3.13 demonstrates that the stochastic schedule utilizes the processing capacity efficiently across all periods with the exception of year nine, due to the large concentration of waste encountered in this area of the pit.

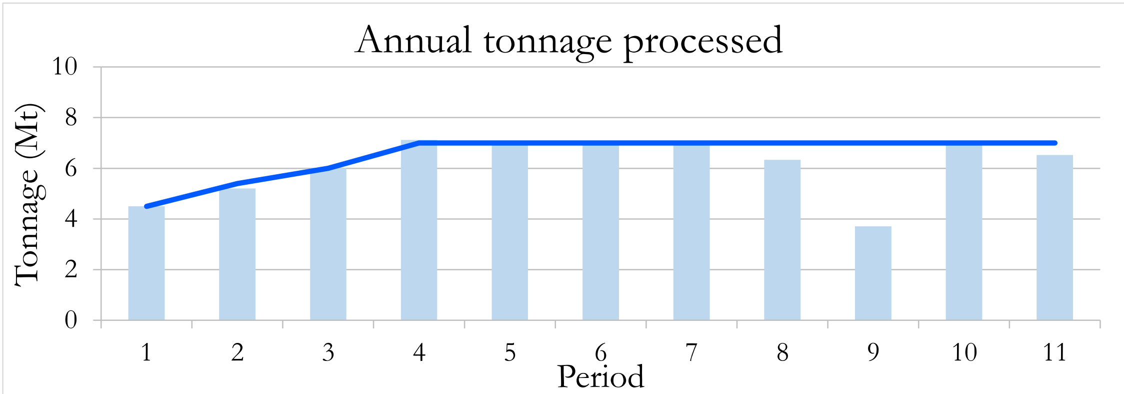

Figure 3.13 - Annual tonnage processed in the stochastic schedule and processing capacity (dark blue)

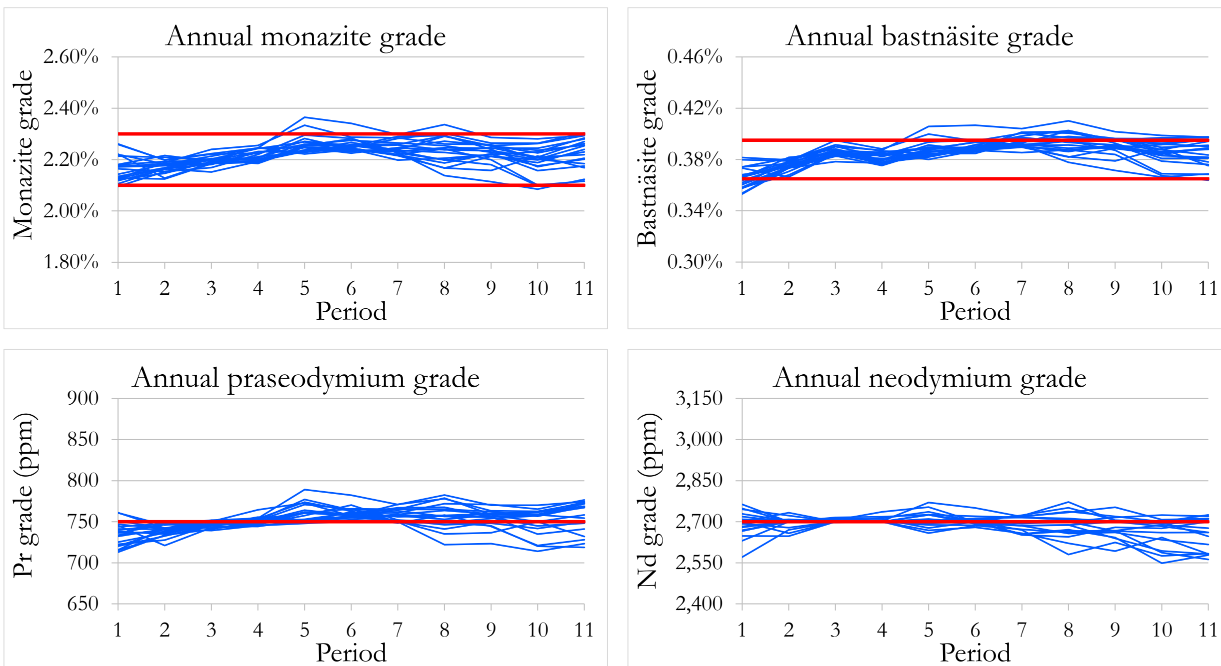

Figure 3.14 - Annual mineral and REE targets (red) and grades (blue) of the stochastic schedule. Monazite (top-left), bastnäsite (top-right), praseodymium (bottom-left), and neodymium (bottom-right)

In terms of the annual mineral ratio targets shown in Figure 3.14, the stochastic schedule is successful in controlling the ore feed to be within the optimized ranges across all geological scenarios. This ensures that the planned recovery at the flotation stage will have a strong likelihood of being realized, eventually leading to more metal and more revenue. Similarly, the Pr and Nd grades in the ore feed are expected to be within acceptable bounds of the targets throughout the mine life. This indicates that the didymium blend is suitable for the customer and that the correct quantity of product will be sold each year, resulting in little, or no penalty costs. These results are useful in demonstrating the effects of geological discounting. The risk profiles are relatively tight in the earlier years of production, indicating effective management of risk, whereas, the later years show a larger variation in REE grades. It is assumed that over time, more information will become available and corrective measures can be applied to reduce this risk that has been deferred to later years. Next, this stochastic schedule is bench marked against the conventional mine production schedule and comparisons are drawn.

3.4.2.4 Conventional optimization results

A deterministic mine production schedule is generated in Whittle using the same operational and financial parameters in Section 3.4.2.2. A risk analysis of the conventional schedule is performed by running each geological simulation through the design.

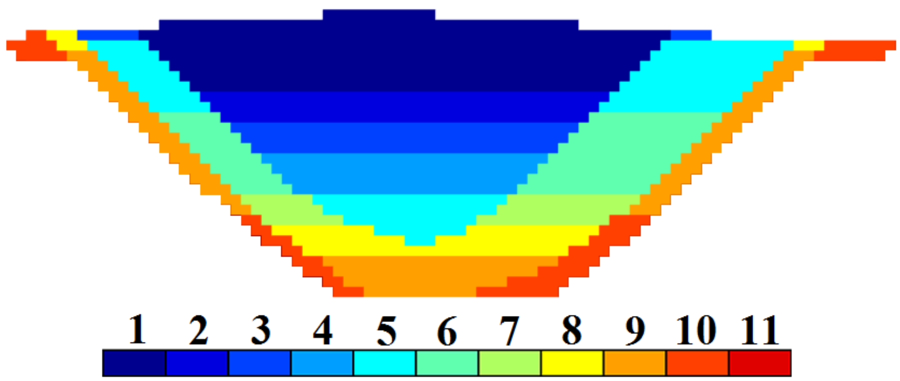

Figure 3.15 - East-West cross-section of the conventional schedule. Colors represent the period of extraction

The LOM in this case is one year shorter than the stochastic approach. The final mining extents cannot be bounded by time directly using this particular solver. An ultimate pit limit is decided earlier in the optimization process by discretizing the space and using approximated discounting, which is considered to be a suboptimal practice. The final conventional schedule is found to be 12 periods, however, since this study does not aim to determine the ultimate pit, the conventional schedule is trimmed to 11 years for a balanced comparison. The eleventh year of production yields negative cash flows and is therefore truncated from the solution to preserve optimism. Figure 3.16 indicates that the processing capacity of the conventional schedule is not utilized as efficiently when compared to the stochastic solution. The limit is exceeded in two of the periods and the planned rate falls short of the target in the remaining eight.

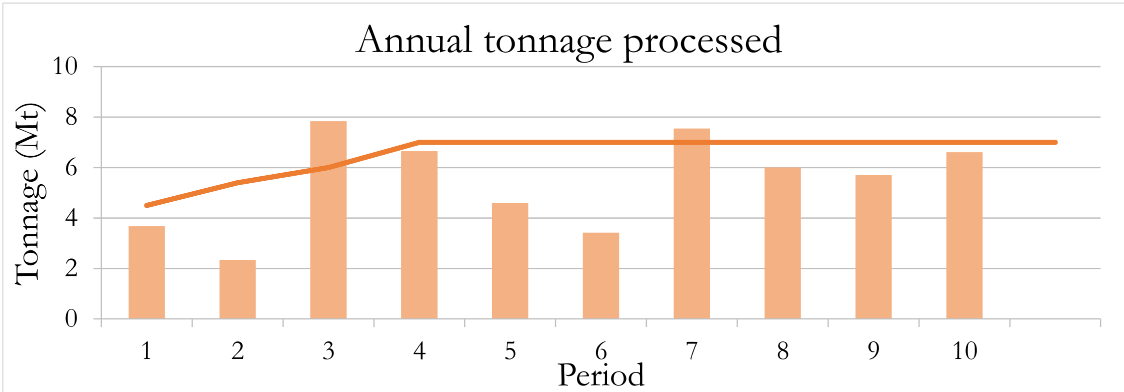

Figure 3.16 - Annual tonnage processed in the conventional schedule and processing capacity (dark orange)

The stochastic schedule also outperforms the conventional schedule in term of mineral targets (Figure 3.17; top); not only in meeting these targets with higher accuracy, but also with far more certainty, evidenced by the tighter risk profiles in Figure 3.14. It is unlikely that the mineral ratio of the ore feed delivered in the conventional schedule will fall within the optimal ranges which will cause a deterioration of the flotation recovery. It is also apparent that there is a very strong likelihood that penalty costs will be incurred due to the fluctuating quality of the didymium blend, as well as inventory costs of stockpiling unsold product in periods five, eight and nine (Figure 3.17; bottom).

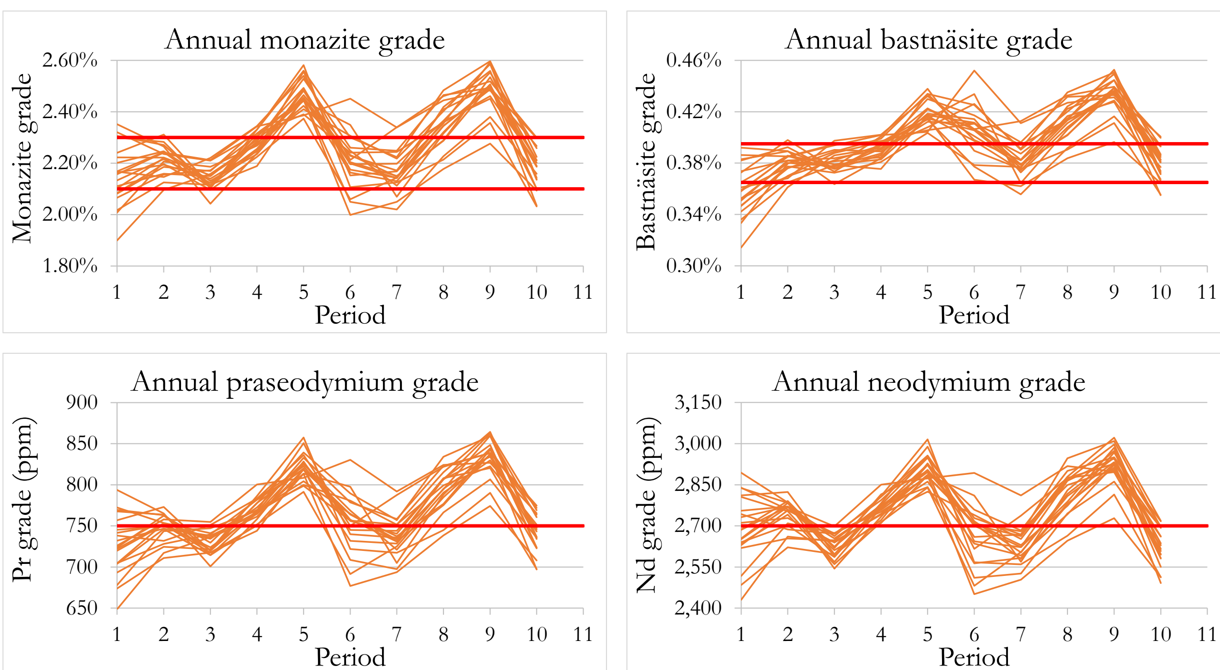

Figure 3.17 - Annual mineralization and REE grades of the conventional schedule. Monazite (top-left), bastnäsite (top-right), praseodymium (bottom-left), and neodymium (bottom-right)

The final comparison between the two schedules is the projects LOM NPV in Figure 3.18. The NPV is described using the 10th, 50th and 90th percentiles of all geological scenarios, rather than each one individually. The stochastic schedule yields an expected NPV increase of 20% over the conventional design. It should also be noted that the final NPV of the conventional schedule does not incorporate the detrimental aspects described above. Due to the strong likelihood of decreased recoveries and the penalty fees for violating the didymium contract, the deterministic cash flows in Figure 3.18 are misleading and will likely to be even lower than demonstrated.

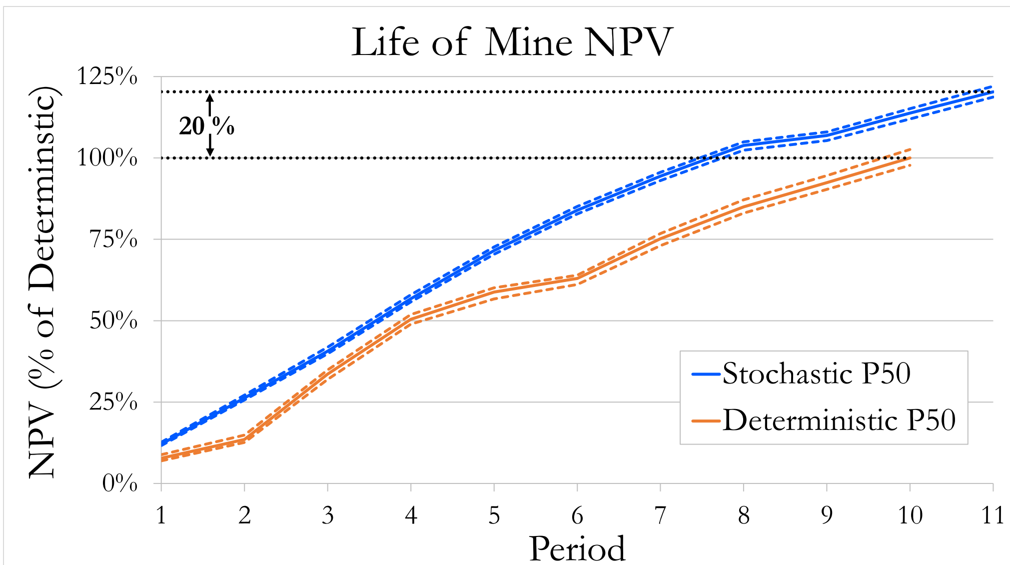

Figure 3.18 - Cumulative NPV of the stochastic schedule (blue) and the conventional schedule (orange)

The results presented herein demonstrate the advantages that stochastic mine optimization can offer over the more conventional methods. By incorporating geological uncertainty into the optimization, logically, the risk of not meeting production targets can be explicitly controlled. In addition, it has also been shown that the expected NPV can also be significantly improved by adopting this risk-based scheduling approach. These results are particularly counter-intuitive given that value is traditionally considered inversely proportional to risk. These advantages can be attributed to the input data that the optimizer receives. By informing a solver with a smooth, estimated orebody model, it becomes far more difficult to discern ore from waste and this can render the task of targeting well-connected, high grade zones far more difficult. Geostatistical simulation not only better reproduces the statistical properties of the conditioning data, but also allows for the quantification of risk, both of which a stochastic solver is able to exploit in full.

3.5 Conclusions and future work

This work presents an approach to assess the geological uncertainty of REE deposits and integrate these uncertainties into the optimization of the long-term mine production schedule. This approach is applied to a factious REE operation and the results show significant improvement both in terms of maximizing the value of the asset as well as mitigating the risk of deviating from annual production targets. SNESIM is used to quantify the volumetric uncertainty of the mineralized domain and MAF, integrated with DBSIM, allows for the efficient joint conditional simulation of all the REE grades. The latter can provide an engineering team with means to: assess the in-situ REE grade variability, reproduce the strong cross-element spatial relationships, and avoid the convention of using the TREO grade representation.

A set of 20 geological simulations are used within the SIP formulation, which is solved using an efficient implementation of the MNTS metaheuristic in an amount of time that would be considered feasible in industry. The optimized stochastic mine production schedule is benchmarked against the industry's best known practices and is shown to increase the expected LOM NPV by 20%. Moreover, the stochastic solution derives a schedule that yields a superior blend of REE bearing minerals at the head of the plant and a didymium product that is far more likely to meet the buyer's strict quality constraints. Future work would aim to test this approach on a REE operation with more complex REE blending requirements.

Chapter 4

Optimizing Short-term Mine Production Schedules Incorporating Geological and Equipment Performance Uncertainty

4.1 Introduction

Production scheduling of open pit mines consists of defining the extraction sequence and process allocation of mineralized material over some length of time. These decisions can be made at different time horizons which will consider different objectives, subject to different technical and operational constraints. Long-term mine production scheduling (LTMPS) usually aims to maximize the net present value (NPV) of the project, while abiding by the mining and milling capacities, which are often pre-defined. The resulting schedule defines the annual ore, waste, and metal tonnages, as well as the annual cash flows which are critical in valuing the project in order to attract investors. Short-term mine production scheduling (STMPS) entails developing an extraction sequence on a shorter time-scale, either months, weeks or days. The goal in short-term scheduling is often to maximize compliance with the production targets and restrictions imposed by the long-term production schedule while considering more detailed operational constraints, such as: available equipment, equipment access, transportation routing, and the ability to deliver consistent ore quantity and quality to each processing destination. These goals can be achieved by optimizing the locations and movements of mobile mining equipment, and the extraction sequence and destination decisions of individual units of material, usually referred to as blocks.

The literature is extensive for methods that optimize LTMPS problems using mathematical programming techniques. Some of the more notable contributions are: Gershon (1983); Caccetta and Hill (2003); Ramazan and Dimitrakopoulos (2004a); Bley et al. (2010). A more comprehensive review of LTMPS optimization methods can be found in Newman et al. (2010). Although not as well-studied as LTMPS, STMPS optimization problems have gained considerable interest, given the substantial savings that can be achieved by minimizing the operating costs. Wilke and Reimer (1977) develop a linear programming (LP) model for a STMPS that maximizes profit while delivering material that satisfies mining and milling capacities and blending requirements. Fytas (1985) also develops an LP formulation to optimize an STMPS constrained by the amount of ore and waste that can be produced each period, along with multi-element blending requirements. As LP approaches may lead to infeasible mine designs, integer programming has since been deemed more appropriate for optimizing mine production schedules. Smith (1998) develops a mixed-integer programming (MIP) formulation to maximize ore production complying with grade blending requirements. The author concludes that the major drawback of the approach is the computation time required to solve for many binary variables given large, multi-period applications. Alternatively, Kumral and Dowd (2002) develop a multi-objective simulated annealing (MSSO) heuristic to optimize the STMPS problem. In this approach, sub-optimal solutions, derived from Lagrangian parametrization, are improved using MSSO considering ore quantity and ore quality objectives. Eivazy and Askari-Nasab (2012) formulate an MIP that incorporates fixed-grade stockpile reclamation, variable destinations and variable haulage routes. The model is solved using CPLEX by aggregating blocks into clusters to reduce the size of the problem.

The LP and MIP methods described above are what Alarie and Gamache (2002) define as the upper-stage of the common multi-stage STMPS optimization approach. In this framework, the upper-stage generates the production schedule to satisfy blending and capacity constraints. This schedule then imposes the targets for the lower-stage, which entails optimizing the dispatching of mobile equipment throughout the mine. L'Heureux et al. (2013) incorporate the sequencing of production activities such as drilling, blasting and extraction into their MIP formulation, along with a shovel movement optimization. A relatively small example where production is scheduled on block aggregates, referred to as faces, is studied to determine how modifications to the formulation can improve the solving time. Moving away from block aggregation approaches, Mousavi et al. (2016) stress the importance of scheduling on at the block level to achieve better grade accuracy throughout the blending optimization.

The authors include fixed-grade stockpile reclamation and more detailed block access requirements. The problem is solved using a hybrid heuristic consisting of simulated annealing (SA), branch and bound (B&B) and large neighbourhood search (LNS).

Although these models are solved considering varying degrees of operational complexity, they all rely on the simple assumption that the geological characteristics and the operational capacities are fully known prior to the optimization. More specifically, the metal grades that define the quality of the orebody are generated using some geostatistical estimation technique (Journel, 1979; Goovaerts, 1997). The shortcomings of conventional grade estimation techniques, and the detrimental impacts they impose on reserve forecasting and LTMPS, have been studied extensively in the past (Ravenscroft, 1992; Vallée, 2000; Godoy, 2003; Ramazan and Dimitrakopoulos, 2013). It has also been shown that incorporating geological uncertainty into optimization models can lead to better decision making and substantial economic benefits (Dimitrakopoulos and Godoy, 2014). Dimitrakopoulos and Jewbali (2013) incorporate grade uncertainty in the STMPS optimization of a large gold mine. The short-term grade uncertainty is modelled by updating conditional simulations with simulated future grade control data, using conditional simulation by successive residuals (CSSR). These models are then used in a stochastic integer programming (SIP) formulation (Dimitrakopoulos and Ramazan, 2008) to maximize NPV and minimize ore and grade target deviations at a quarter-annual scale.

Just as it has been shown that ignoring grade uncertainty inherently leads to unforeseen complications in achieving production targets, uncertainty pertaining to equipment performance can also have substantial negative ramifications on short-term production. The failure to actually extract the scheduled material will always result in deviations from the plan that might cause rippling effects throughout the long-term plan. There has been very little investigation as to how equipment performance uncertainty can be incorporated into STMPS considering that all mining operations observe these effects in practice. In the field of automated manufacturing systems (AMS), Jain and Foley (2016) review multiple strategies available for machine-job scheduling to handle unexpected machine breakdowns, lower yields, and hot-jobs. Panagiotidou and Tagaras (2007) optimize the preventative maintenance schedule of a series of units in an AMS in order to maximize the productivity at minimal cost. Unfortunately, these strategies do not translate well to the complex nature of block-wise mine production scheduling, due to strict grade requirements, block precedence, and equipment access constraints. Topal and Ramazan (2012) propose an MIP formulation to minimize the maintenance cost of operating mining equipment under the assumption that the cost varies with equipment age. This formulation is also extended to incorporate stochastic maintenance costs (Topal and Ramazan, 2012). Burt et al. (2016) adopt a similar concept for the problem of purchasing and salvaging equipment for long-term mine production, given that equipment availability and maintenance costs vary with age. Unfortunately, these approaches consider equipment performance on a long-term scale and, therefore, cannot be extended to STMPS. Matamoros and Dimitrakopoulos (2016) introduce a model that integrates both geological, and equipment uncertainty into a STMPS optimization. Conditional simulations, updated with future grade control data, are coupled with historical equipment performance data to expose the optimization model to both of these sources of uncertainty. The optimizer aims to meet the monthly grade blending requirements while generating minable shapes that respect a pre-defined direction of mining.

The new formulation to optimize STMPS proposed in this chapter builds on the method introduced by Matamoros and Dimitrakopoulos (2016). The model optimizes short-term schedules by minimizing the production deviations from targets imposed by the long term plan, and by generating extraction patterns that can be achieved with far greater confidence given the fleet's uncertain productivity. These objectives will be optimized while accounting for uncertainty in both material grades and equipment performance, in order to shield the production plan from risk. Below are the novel aspects and model improvements that will be incorporated in this stochastic short-term mine production scheduling (SSTMPS) optimization approach.

1. Adapt the short-term scheduling problem to coincide with the concept of mining complexes which are large, integrated mining systems comprised of multiple mines, multiple material types, and multiple processing destinations.
2. Develop a new co-simulation approach to generate more realistic equipment performance scenarios and route dependant truck cycle time scenarios.
3. Introduce a more comprehensive concept of block access to facilitate more practical mine designs.
4. Include more realistic, location-dependant shovel movement costs in the formulation.

The optimization model is formulated as an SIP (Birge and Louveaux, 1997) and solved with CPLEX v.12.6.1.0 in a Visual Studio 13 (C++) environment. The following section outlines the main concepts of the proposed SSTMPS optimization model in more detail. Then, the full formulation of the optimization model is presented. The approach is then applied to a very large copper mining complex to assess the effectiveness of the stochastic optimizer when compared to a more conventional approach. Finally, conclusions and venues of future work are discussed.

4.2 Concepts of short-term optimization

4.2.1 Short-term optimization objectives

The proposed stochastic optimization model is designed to adapt the STMPS problem to fit the notion of mining complexes (Montiel and Dimitrakopoulos, 2015; Goodfellow and Dimitrakopoulos, 2016), which entails the simultaneous optimization of mining operations that contain multiple mines, multiple processors, and complex multi-element ores. The model yields two primary deliverables: (1) a complete mobile equipment plan, consisting of shovel placements and truck allocation; and (2) a block-level extraction sequence that will guide the progression of mining activity. The main objectives in STMPS are significantly different than those in long-term scheduling since the optimizer is no longer motivated by time-discounted cash flows. Since the long-term plan has already established that all of the value will be recovered within one fiscal year, the motivation to target high grade ore in earlier periods is no longer present. Instead, the optimizer prioritizes the following three objectives.

1. Maximize compliance with the long-term plan. This is achieved by delivering consistent tonnage and grades to each destination in each period. This will minimize re-handling costs and maximize processor utilization and recovery.
2. Minimize shovel movement costs. This can be done by keeping shovels in the same location longer or by encouraging them to move smaller distances to extract new material.
3. Minimize the non-productive time of the fleet. By ensuring that adequate trucking resource is allocated to each shovel, the production from both trucks and shovels can be synchronised to minimize idle times. This is particularly difficult considering that material destinations, equipment performance, and truck cycle times are uncertain.

Previous STMPS optimization methods attempt to directly minimize the mining cost in the objective function using a per-block cost; a concept borrowed from long-term scheduling. Since the long-term plan already states that all blocks being considered in the short-term must be extracted by the end of the time horizon, these costs are incurred regardless of the schedule. The three objectives listed above better encapsulate the challenges pertaining to short-term planning.

4.2.2 Short-term concepts of space and time

The time-step for STMPS problems can vary from months, weeks, days or shifts. However, the final extent of the time horizon should coincide with the amount of available material taken from the long-term plan. The concept of space is partitioned into two distinct domains, as can be seen in Figure 4.1.

1. Selective mining units (SMU), or blocks, are the smallest discretization of space in mine production scheduling. Each block has a known position, a metal grade(s), a tonnage, and can be extracted and sent to a destination in one period.
2. An area represents a pre-defined region of the mine comprised of a large grouping of connected blocks. In the optimization model, areas are used to describe equipment positions, truck cycle times, shovel movements, and the associated movement cost.

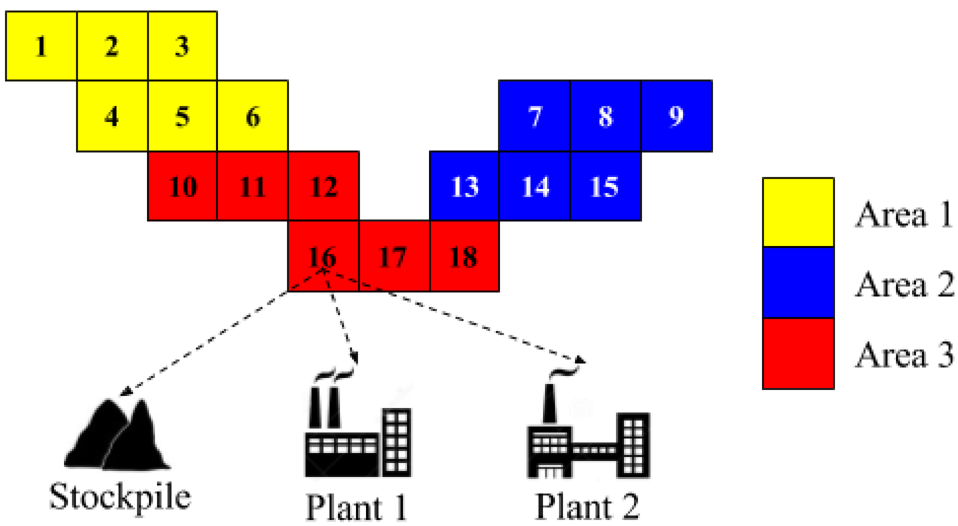

Figure 4.1 - Configuration of blocks within areas and possible destinations of a block

It is worth noting that in the proposed model, all scheduling decisions are made at the block level. Most short-term modules in commercial mine planning packages, and recent works in academia (Eivazy and Askari-Nasab, 2012) use a block aggregation approach to simplify the problem. An aggregate is a pre-processed grouping of spatially connected blocks that form larger units with cumulative tonnages and averaged grades. Although this reduces the number of decision variables, precedence constraints, and subsequent solving times, Mousavi et al. (2016) highlight three reasons why this approximation should be avoided: (1) it reduces the granularity of the solution, and since the pre-processing algorithms often lack sophistication, it will lead to sub-optimal results; (2) averaging block grades through aggregation intensifies the grade smoothing effect and the negative impacts of this phenomenon and how it deteriorates the quality of mine production schedules is a well-studied topic; and (3) when schedules that are optimized using aggregated units are evaluated at the block level, the solutions are often infeasible in the original context of the problem due to the re-exposure of grade variability.

4.2.3 Practical schedules using horizontal precedence

Solutions from STMPS optimizations must be easily implementable since they usually directly precede operational planning. Shovel movement, shovel proximities, equipment access and face progression should all be considered to make this planning transition as seamless as possible. Aside from traditional vertical precedence constraints, encouraging spatially connected blocks in a bench to be extracted in the same period will yield extraction patterns that are more practical and will limit unnecessary shovel moves. However, a trade-off exists between objective function value and practical schedules achieved by heavily-constraining the problem.

One approach to encourage smooth schedules is to apply a penalty cost when blocks are not extracted in the same period as their adjacent neighbor. This approach requires a significant amount of additional constraints and integer decision variables (Dimitrakopoulos and Ramazan, 2004; Matamoros and Dimitrakopoulos, 2016). The additional variables and complex constraints required to implement this approach render it ineffective, especially when using commercial solvers, since it cannot be scaled to realistic problem sizes. Mousavi et al. (2016) uses a more mining operations focused approach. The authors distinguish block access into drop-cuts and side-cuts. They propose a heuristic that selects an access path for each block and encourages smooth schedules by discouraging drop-cuts with high penalty costs.

The direction of mining in an implementable schedule is also important. This is usually enforced by adding some of the immediate in-bench neighbours to each block's precedence set. This precedence direction is consistent for every block in the mine and must be pre-defined as one of four directions: access from the north, south, east, or west. There are two major drawbacks to this approach: (1) it tends to produce square, un-natural looking schedules; and (2) the direction of mining should be allowed to vary by mine and by regions of a mine, all of which will have different access points.

A new approach is proposed in this work to ensure reasonable progression of mining. More specifically, a modified pre-processing technique is proposed that allows for each block in the mine to have a unique horizontal precedence set along a continuum of access directions. The underlying assumption is that a significant amount of planning has already been established prior to the STMPS. For each bench, in each area being considered in the short-term time horizon, a permanent or temporary ramp should already be designed. Knowing the locations of these ramp entrance points for each bench, a unique direction of access for each block can be established. Figure 4.2 shows an illustration of how a ramps positioning (blue) is used to establish the horizontal precedence set (cyan) of certain blocks (red).

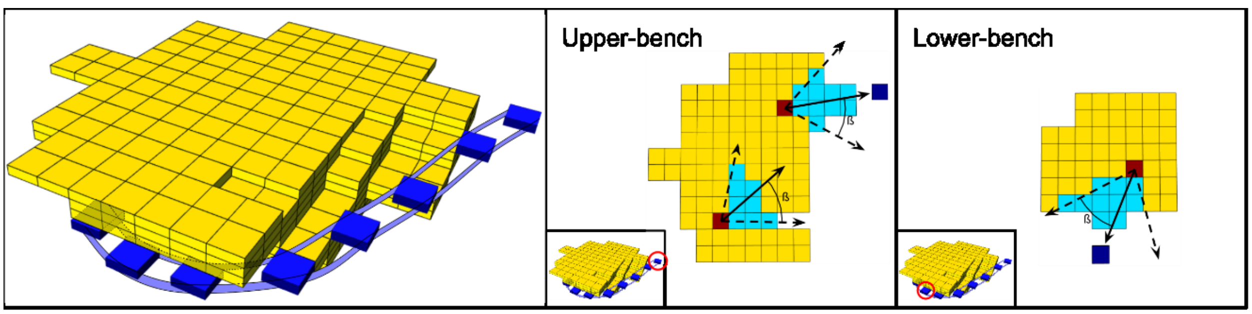

Figure 4.2 - An example area of a mine and the associated ramp design (left), horizontal precedence sets on the upper-bench (middle), and a horizontal precedence set on the lower-bench (right)

The pre-processing technique to determine the horizontal precedence of each block is summarized as follows.

1. Identify the bench and area that a block belongs to.
2. Establish the direction to the appropriate ramp access point.
3. Allow for a tolerance angle, β, about the access direction. The magnitude of this angle can be varied to impose equipment access requirements.
4. Add each block within the tolerance range to the horizontal precedence set. Since the problem is solved using mathematical programming techniques, the transitivity property of the precedence constraints is exploited; only the minimal set of preceding blocks is needed.

4.2.4 Uncertain model parameters

In this work, the optimization is performed in a stochastic framework considering three sources of uncertainty: (1) metal grades and associated material types, (2) equipment performance, and (3) truck cycle times. Each of these types of uncertainty are included in the model using simulated scenarios. Although more geological information has typically been collected from dense blast-hole data, the characteristics of the orebody are still not fully understood when executing short-term scheduling. First, if the extent of the planning time horizon exceeds 3-6 months, much of the material of interest will be too far from the available blast-hole information to have significant improvement in grade value confidence. Second, gathering more densely-spaced information tends to expose even more grade variability than previously understood. For these reasons, it is still very important to model the grade variability through simulated orebodies to optimize the short-term production plan. The mapping of unknown correlated metal grades at the block level is generated through a joint conditional simulation approach known as DBMAFSIM (Boucher and Dimitrakopoulos, 2009). It can be particularly cumbersome to perform large joint conditional simulations with conventional methods. Instead, DBMAFSIM transforms a set of correlated variables into uncorrelated "factors" that can be simulated independently. The simulations are back-transformed to the initial variable space to preserve the original spatial relationships.

Random fluctuations in weather, operator habits, road conditions, and overall maintenance trends are all factors that lead to mobile equipment performance variability. Since equipment performance can be highly unpredictable and can have a substantial impact on short-term decision making, it should be incorporated stochastically in the optimization model. To quantify the uncertain equipment performance, three different parameters are simulated: (1) the availability, and (2) the utilization of all mobile equipment, and (3) the hourly production of the shovels. The overall tonnage production for each unit of equipment can be described using these parameters. Matamoros and Dimitrakopoulos (2016) include uncertain equipment availabilities in a short-term model by randomly sampling independent Gaussian distributions derived from historical data. However, since the factors that cause this variability impact the entire fleet simultaneously, the performance measurements tend to be strongly correlated; the simulations must reproduce these relationships. Borrowing the decorrelation idea from DBMAFSIM, the joint simulation of equipment performance parameters is done following a projection into some orthogonal space using principle component analysis (PCA) (Hotelling, 1933).

Following the independent simulation of the uncorrelated components, the final simulation is generated using a back-rotation to the original space. Finally, since truck cycle times are also highly variable, although not necessarily directly related to availability and utilization, including them in a stochastic framework might allow for more robust equipment matching decisions in short-term mine planning. Lacking reliable data to develop a thorough investigation as to how these cycle times vary in a mining system, the approach used by Matamoros and Dimitrakopoulos (2016) is replicated. For each area in each mine and each potential destination, a Gaussian cycle time distribution is sampled using a mean and variance derived from historical data.

4.3 Stochastic short-term mine production scheduling

The SSTMPS model presented in this section is formulated as an SIP that optimizes the mobile equipment plan and the block extraction sequence under geological and equipment performance uncertainty. The binary extraction decisions are modelled in a way that reduces the size of the linear relaxation and facilitates simpler branching (Caccetta and Hill, 2003). The objectives are to minimize tonnage and grade target deviations at each destination under geological uncertainty, minimize equipment idle times, and generate a schedule that can be achieved with greater confidence under equipment uncertainty. The formulation is introduced first, followed by the solution approach.

4.3.1 Model formulation

4.3.1.1 Indices

t is a time period, $t = 1, \dots, T$

l is an individual shovel (loader), $l = 1, \dots, L$

m is a model of haulage truck, $m = 1, \dots, M$

j is an area of the mine, $j = 1, \dots, J$

i is a block within area j, $i \in i(j)$

p is a processing destination (crusher, waste dump leach-pad, etc.), $p = 1, \dots, P$

k is a material type, $k = 1, \dots, K$

ε is a metal element, $\varepsilon = 1, \dots, E$

s is a geological scenario, $s = 1, \dots, S$

α is an equipment performance scenario, $\alpha = 1, \dots, A$

4.3.1.2 Model parameters

4.3.1.2.1 Equipment parameters

H_t is the total number of hours in period t

$A^L_{tlj\alpha}$ is the availability of shovel l in period t in area j for equipment scenario α

$U^L_{tlj\alpha}$ is the utilization of shovel l in period t in area j for equipment scenario α

$Q^L_{tlj\alpha}$ is the production rate (t/h) of shovel l in period t in area j for equipment scenario α

$A^M_{tmj\alpha}$ is the availability of truck type m in period t in area j for equipment scenario α

$U^M_{tmj\alpha}$ is the utilization of truck type m in period t in area j for equipment scenario α

Q^M_m is the tonnage capacity of truck type m

$R_{tmjp\alpha}$ is the cycle time for truck type m delivering material from area j to destination p in period t for equipment scenario α

Λ^{start}_{0l} is the area location of shovel l at time $t = 0$

Λ^{max}_j is the maximum allowable loaders in area j

M^{max}_m is the maximum number of trucks of type m available in the fleet

4.3.1.2.2 Operational parameters

Ton_{ji} is the tonnage of block i in area j

θ_{jis} is the material type of block i in area j for geological scenario s

$G_{ji\varepsilon s}$ is the grade of element ε of block i in area j for geological scenario s

$D^{min}_{tp}, D^{max}_{tp}$ are the minimum/maximum tonnage targets at destination p in period t

$G^{min}_{tp\varepsilon}, G^{max}_{tp\varepsilon}$ are the minimum/maximum grade targets for element ε at destination p in period t

$C^{Lmove}_{lj'j}$ is the cost to move shovel l from area j' to area j

C^{L-}_j, C^{L+}_j are the lower/upper deviation costs for shovel tonnage production in area j

C^{M-}_j, C^{M+}_j are the lower/upper deviation costs for truck hours in area j

C^{D-}_p, C^{D+}_p are the lower/upper tonnage target deviation costs at destination p

$C^{G-}_{p\varepsilon}, C^{G+}_{p\varepsilon}$ are the lower/upper grade target deviations costs for element ε at destination p

F_{jips} dictates whether block i in area j can be sent to destination p in geological scenario s

$$F_{jips} = \begin{cases} 1 & \text{If the material type of block } i \text{ in geological scenario } s \text{ can be sent to destination } p \\ 0 & \text{otherwise} \end{cases}$$

$\rho_v(i(j))$ is the set of blocks that precede block i in the vertical direction

$\rho_h(i(j))$ is the set of blocks that precede block i in the horizontal direction (Figure 4.2)

$\rho(i(j))$ is the complete set of blocks that precede block i, $\rho(i(j)) = \rho_v(i(j)) \cup \rho_h(i(j))$

4.3.1.3 Decision variables

4.3.1.3.1 Integer decision variables

$$x_{tji} = \begin{cases} 1 & \text{if block } i \text{ in area } j \text{ is extracted by period } t \\ 0 & \text{otherwise} \end{cases}$$

$$\lambda_{tlj} = \begin{cases} 1 & \text{if shovel } l \text{ is located in area } j \text{ in period } t \\ 0 & \text{otherwise} \end{cases}$$

$$\Omega_{tlj'j} = \begin{cases} 1 & \text{if shovel } l \text{ moves from area } j' \text{ to area } j \text{ in period } t \\ 0 & \text{otherwise} \end{cases}$$

τ_{tmj} is the number of trucks of model m allocated to area j in period t

4.3.1.3.2 Continuous decision variables

n_{tmjps} is the number of trips that trucks of model m allocated to area j make to destination p in period t for geological scenario s

$h_{tmjps\alpha}$ is the hours that trucks of model m allocated to area j require to deliver material to destination p in period t for geological scenario s and equipment performance scenario α

$d_{tps}^{D-}, d_{tps}^{D+}$ are the lower/upper tonnage deviations at destination p in period t for scenario s

$d_{tp\varepsilon s}^{G-}, d_{tp\varepsilon s}^{G+}$ are the lower/upper grade deviations of element ε at destination p in period t for geological scenario s

$d_{tj\alpha}^{L-}, d_{tj\alpha}^{L+}$ are the lower/upper shovel production deviations in area j in period t for equipment performance scenario α

$d_{tjs\alpha}^{M-}, d_{tjs\alpha}^{M+}$ are the lower/upper truck hour deviations in area j in period t for geological scenario s and equipment performance scenario α

4.3.1.4 Objective function – Minimize:

$$\sum_{t=1}^{T}\sum_{l=1}^{L}\sum_{j'=1}^{J}\sum_{j=1}^{J} C_{lj'j}^{Lmove} \cdot \Omega_{tlj'j} \qquad \text{Part 1}$$

$$+\frac{1}{S}\sum_{t=1}^{T}\sum_{p=1}^{P}\sum_{s=1}^{S} C_{p}^{D-} \cdot d_{tps}^{D-} + C_{p}^{D+} \cdot d_{tps}^{D+} \qquad \text{Part 2}$$

$$+\frac{1}{S}\sum_{t=1}^{T}\sum_{p=1}^{P}\sum_{\varepsilon=1}^{E}\sum_{s=1}^{S} C_{p\varepsilon}^{G-} \cdot d_{tp\varepsilon s}^{G-} + C_{p\varepsilon}^{G+} \cdot d_{tp\varepsilon s}^{G+} \qquad \text{Part 3} \qquad 4.1$$

$$+\frac{1}{A}\sum_{t=1}^{T}\sum_{j=1}^{J}\sum_{\alpha=1}^{A} C_{j}^{L-} \cdot d_{tj\alpha}^{L-} + C_{j}^{L+} \cdot d_{tj\alpha}^{L+} \qquad \text{Part 4}$$

$$+\frac{1}{S \cdot A}\sum_{t=1}^{T}\sum_{j=1}^{J}\sum_{s=1}^{S}\sum_{\alpha=1}^{A} C_{j}^{M-} \cdot d_{tjs\alpha}^{M-} + C_{j}^{M+} \cdot d_{tjs\alpha}^{M+} \qquad \text{Part 5}$$

Part 1 of the objective function (Equation 4.1) minimizes the total shovel movement costs. The cost of moving a shovel between two areas is a function of the lateral and vertical distance between each area. Parts 2 and 3 minimize the expected costs of deviating from tonnage and grade targets at each destination respectively. For each area in the mine, Part 4 minimizes the costs associated to the discrepancy between the scheduled extraction tonnage and the production provided from the shovels across all equipment scenarios. Lastly, Part 5 minimizes the deviations between truck hours required to deliver the scheduled material in an area, and the total truck hours assigned to that area. If trucking hours deviate above the required hours, trucks will be idle and underutilized, and if the hours deviate below, shovels will endure waiting times instead. The optimal number of trucks is assigned to an area in order to improve the equipment productivity in the face of performance and cycle time uncertainty.

4.3.1.5 Model constraints

4.3.1.5.1 Scheduling constraints

$$x_{tji} \geq x_{(t-1)ji} \quad \forall\, t,\ j,\ i(j) \qquad 4.2$$

$$x_{tji} \leq x_{tn} \quad \forall\, t,\ j,\ i \in i(j),\ n \in p(i(j)) \qquad 4.3$$

$$\sum_{j=1}^{J}\sum_{i=1}^{i(j)} Ton_{ji} \cdot F_{jips} \cdot (x_{tji} - x_{(t-1)ji}) + d_{tps}^{D-} \geq D_{tp}^{min} \quad \forall\, t,\ p,\ s$$

$$\sum_{j=1}^{J}\sum_{i=1}^{i(j)} Ton_{ji} \cdot F_{jips} \cdot (x_{tji} - x_{(t-1)ji}) - d_{tps}^{D+} \leq D_{tp}^{max} \quad \forall\, t,\ p,\ s \qquad 4.4$$

$$\sum_{j=1}^{J}\sum_{i=1}^{i(j)} Ton_{ji} \cdot \left(G_{ji\varepsilon s} - G_{tp\varepsilon}^{min}\right) \cdot F_{jips} \cdot (x_{tji} - x_{(t-1)ji}) + d_{tp\varepsilon s}^{G-} \geq 0 \quad \forall\, t,\ p,\ \varepsilon,\ s$$

$$\sum_{j=1}^{J}\sum_{i=1}^{i(j)} Ton_{ji} \cdot \left(G_{ji\varepsilon s} - G_{tp\varepsilon}^{max}\right) \cdot F_{jips} \cdot (x_{tji} - x_{(t-1)ji}) - d_{tp\varepsilon s}^{G+} \leq 0 \quad \forall\, t,\ p,\ \varepsilon,\ s \qquad 4.5$$

Constraints 4.2 ensure that each block cannot be extracted more than once, and block precedence is enforced by Constraints 4.3. Constraints 4.4 and 4.5 set the tonnage and average grade deviations at each destination for each geological scenario respectively.

4.3.1.5.2 Shovel constraints

$$\sum_{l=1}^{L} \lambda_{tlj} \leq L_j^{max} \quad \forall\, t,\ j \qquad 4.6$$

$$\sum_{j=1}^{J} \lambda_{tlj} = 1 \quad \forall\, t,\ l \qquad 4.7$$

Constraints 4.6 limit the number of shovels that can operate in a given area in one period. The size of the area will dictate the maximum number of shovels permitted to ensure all units have enough space to operate safely. Constraints 4.7 enforce the rule that each shovel must be assigned to exactly one area in each period.

$$\lambda_{tlj} + \lambda_{(t-1)lj'} - \Omega_{tlj'j} \leq 1 \quad \forall\, t,\ l,\ j,\ j' \neq j \qquad 4.8$$

$$-\lambda_{tlj} + \Omega_{tlj'j} \leq 0 \quad \forall\, t,\ l,\ j,\ j' \neq j \qquad 4.9$$

$$-\lambda_{(t-1)lj'} + \Omega_{tlj'j} \leq 0 \quad \forall\, t,\ l,\ j,\ j' \neq j \qquad 4.10$$

Constraints 4.8 to 4.10 are used to set the shovel movement variables. All three are necessary to ensure that if a shovel is present in two different areas in consecutive periods, the movement variable, $\Omega_{tlj'j}$, is activated and the associated movement cost is accounted for in the objective function.

$$x_{tji} - x_{(t-1)ji} \leq \sum_{l=1}^{L} \lambda_{tlj} \quad \forall\, t,\ j,\ i(j) \tag{4.11}$$

$$\sum_{i=1}^{i(j)} Ton_{ji} \cdot \left(x_{tji} - x_{(t-1)ji}\right) - \sum_{l=1}^{L} \left(Q_{tlj\alpha}^{L} \cdot H_{t}^{L} \cdot A_{tlj\alpha}^{L} \cdot U_{tlj\alpha}^{L}\right) \cdot \lambda_{tlj} + d_{tj\alpha}^{L-} - d_{tj\alpha}^{L+} = 0 \quad \forall\, t,\ j,\ \alpha \tag{4.12}$$

Constraints 4.11 ensure that a block can only be extracted if at least one shovel is present in the area where this block is located. Constraints 4.12 link the tonnage extracted in each area to the production provided by the shovels present in that area. Since shovel production is uncertain, violations are permitted but are penalized in the objective function.

4.3.1.5.3 Truck constraints

$$\sum_{j=1}^{J} \tau_{tmj} \leq M_{m}^{max} \quad \forall\, t,\ m \tag{4.13}$$

$$\sum_{i=1}^{i(j)} Ton_{ji} \cdot F_{jips} \cdot (x_{tji} - x_{(t-1)ji}) - \sum_{m=1}^{M} \left(Q_{m}^{M} \cdot n_{tmjps}\right) = 0 \quad \forall\, t,\ j,\ p,\ s \tag{4.14}$$

$$h_{tmjps\alpha} - n_{tmjps} \cdot R_{mjp\alpha} = 0 \quad \forall\, t,\ m,\ j,\ p,\ s,\ \alpha \tag{4.15}$$

$$\sum_{m=1}^{M} \tau_{tmj} \cdot \left(H_{tm}^{M} \cdot A_{tmj\alpha}^{M} \cdot U_{tmj\alpha}^{M}\right) - \sum_{m=1}^{M} \sum_{p=1}^{P} h_{tmjps\alpha} + d_{tjs\alpha}^{M-} - d_{tjs\alpha}^{M+} = 0 \quad \forall\, t,\ j,\ s,\ \alpha \tag{4.16}$$

Constraints 4.13 limit the number of trucks assigned across all areas in each period by the fleet size of each truck model. Constraints 4.14 set the number of trips that each truck model in each area requires to deliver the scheduled tonnage to each destination. Unlike shovel production, truck production does not depend on the equipment performance scenario since truck production only consists of constant truck payloads and the number of trips. However, the number of trips to each destination does vary by geological scenario since a block is sent to different destinations depending on its uncertain material type. Furthermore, the time required to deliver material to its destination and the productive trucking time that will be provided by assigning a truck to an area both depend on the equipment performance scenario. Constraints 4.15 set the trucking hours required to deliver material in each geological and equipment performance scenario. Constraints 4.16 link the number of trucks that are assigned to an area in order to achieve the planned production schedule with as much certainty as possible. The hour deviations set in Constraints 4.16 are minimized in the objective function to maximize equipment productivity. Constraint 4.17 through 4.26 are the variable bounds.

$$x_{tji} \in 0 \text{ or } 1 \quad \forall\, t,\ j,\ i \tag{4.17}$$

$$\lambda_{tlj} \in 0 \text{ or } 1 \quad \forall\, t,\ l,\ j \tag{4.18}$$

$$\Omega_{tlj'j} \in 0 \text{ or } 1 \quad \forall\, t,\ l,\ j,\ j' \neq j \tag{4.19}$$

$$\tau_{tmj} \in 0 \text{ or } 1 \quad \forall\, t,\ m,\ j \tag{4.20}$$

$$n_{tmjps} \geq 0 \quad \forall\, t,\ m,\ j,\ p,\ s \tag{4.21}$$

$$h_{tmjps\alpha} \geq 0 \quad \forall\, t,\ m,\ j,\ p,\ s,\ \alpha \tag{4.22}$$

$$d_{tps}^{D-}, d_{tps}^{D+} \geq 0 \quad \forall\, t,\ p,\ s \tag{4.23}$$

$$d_{tp\varepsilon s}^{G-}, d_{tp\varepsilon s}^{G+} \geq 0 \quad \forall\, t,\ p,\ \varepsilon,\ s \tag{4.24}$$

$$d_{tj\alpha}^{L-}, d_{tj\alpha}^{L+} \geq 0 \quad \forall\, t,\ j,\ \alpha \tag{4.25}$$

$$d_{tjs\alpha}^{L-}, d_{tjs\alpha}^{L+} \geq 0 \quad \forall\, t,\ j,\ s,\ \alpha \tag{4.26}$$

4.3.2 Solution approach

The SSTMPS model described in Section 4.3.1 is solved using the branch-and-cut algorithm implemented in IBM's CPLEX package. In contrast to long-term planning, there are fewer binary decisions variables in short-term planning since there are fewer blocks. However, this reduction in extraction variables is negated by the additional equipment plan decision variables, which substantially increase the complexity of the problem. As a result, the SSTMPS model cannot be solved efficiently, and two heuristic algorithms are proposed to reduce the size of the problem.

4.3.2.1 Sliding time window heuristic

The basic goal of the sliding time window heuristic (STWH) (Pochet and Wolsey, 2006) is to reduce the solving time of large integer programs. This is done by iteratively solving relaxed sub-problems with fewer binary variables. Cullenbine et al. (2011) extend the STWH to the mine production scheduling problem. The block extraction variables within a pre-defined number of periods, or window, are kept as binary, whereas, those associated with the remaining periods are relaxed to be continuous. The sub-problem is solved, the binary variables are fixed to the solved values, the window is shifted forward by one period, and the process is repeated until all binary variables in the original problem have been fixed. This heuristic approach has shown to substantially reduce the solving time for realistic mine scheduling problem sizes and still achieve an optimality gap of less than 2.5% (Cullenbine et al., 2011). The STWH applied to the SSTMPS is outlined as follows:

1. A time window of one period is selected.
2. Model all variables and constraints in period t consistent with the formulation described in Section 4.3.1. Model all binary decisions as continuous in periods t', where $t' > t$.
3. Solve the relaxed sub-problem and fix the binary variables associated with period t.
4. Move the window to period $t + 1$.
5. Repeat until all periods have been considered.

4.3.2.2 Earliest start heuristic

The solving time of each STWH iteration can be reduced even further by fixing a number of binary variables using the earliest start heuristic (ESH) (Topal, 2008). This heuristic eliminates binary block extraction decisions based on the logic that it is impossible to mine some blocks in certain periods. It exploits the nature of block precedence and the capacitated mining constraints in production scheduling models. In short, the binary variable associated with extracting a block in a given period can be fixed to zero if that block is preceded by an amount of material greater than the extraction capacity of that period. The earliest period that a block can be extracted can be defined, and all earlier extraction decisions can be eliminated. Three variants of this ESH are used at the beginning of each sliding window iteration where the extraction variables are only binary in period t:

1. **Shovel production**. For each block i that has yet to be mined by period t, a recursive algorithm is used to determine the complete set of blocks that precede block i that also remain to be extracted by period t. The total tonnage contained within this complete precedence set is compared to the total shovel production capacity of period t. Since shovel production is uncertain, the performance scenarios with the highest production rates for each shovel are combined to form an upper-bound. If the total tonnage of blocks requiring extraction in order to access block i exceeds the best case shovel performance, block i is deemed unreachable and the decision variable x_{tji}, is fixed to zero.
2. **Shovel presence**. This variant exploits the structure of the SSTMPS model. Blocks in area j can only be extracted in period t if area j is active, meaning at least one shovel is present in area j during period t (Constraints 4.11) and shovels must be present in exactly one area for the entire duration of one period (Constraints 4.7). Therefore, given the complete precedence set of block i and knowing which area j each of these preceding blocks belongs to, the total number of active areas required to access block i can be determined. If the number of active areas required exceeds the total number of shovels in the fleet, the block is deemed unreachable and the decision variable x_{tji}, is fixed to zero. Figure 4.3 helps demonstrate this variant of the EHS. In a mine with two shovels, the extraction of block 2 in period t could be eliminated since three shovels would have to be present in areas 1, 2, and 3 in period t. On the other hand, block 1 could be extracted in period t since only two shovels would need to be present in areas 2 and 3.
3. **Area elimination**. Following the previous two EHS variants, another pass is made to determine if mining activity in an entire area of the mine can be eliminated. If the extraction variables of all of the remaining blocks contained in area j have been fixed to zero in period t, then all of the shovel presence variables λ_{tlj}, and truck allocation variables τ_{tmj}, are fixed to zero in period t. These variable eliminations are far less likely to occur, but have a significant impact on the solving time since the equipment presence variables tend to complicate the problem the most.

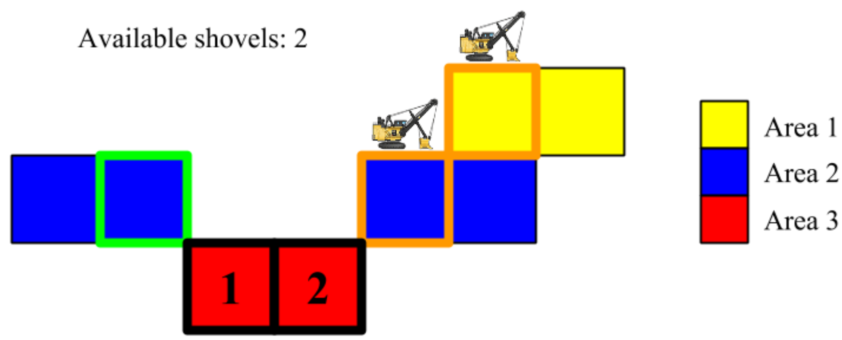

Figure 4.3 - The precedence set of block 1 (green) and block 2 (orange)

4.4 Application – Copper mining complex

In this section, the proposed SSTMPS formulation is applied to a very large copper mining complex to demonstrate its effectiveness is generating near-optimal short-term production schedules and equipment plans. The outline of the operation and the details of the input parameters are presented below. Some of the specific operational parameters, such as processing tonnages, copper grades and equipment production rates have been withheld for confidentiality purposes.

4.4.1 Copper operation outline

This copper mining complex uses conventional truck and shovel methods to extract material from two different mines; the Large Pit and the Small Pit. Material is extracted from each of these sources and is then trucked to a processing destination depending on the material's type. A block's material type is defined as sulphide, oxide, or waste material depending on its total copper and soluble copper contents. The various processing streams of this mining complex convert raw material into one of two saleable products, copper concentrate or copper cathodes. However, the only components in the system that are relevant to this short-term optimization are those where mobile equipment is directly involved; the points of shovel extraction and the truck dumping locations. Figure 4.4 outlines the processing destinations within this mining complex and the types of material each one accepts. Waste and oxide blocks from both mines are sent to the waste dump and oxide leach pad respectively. Sulphides, on the other hand, are allocated depending on the block's location and its total copper grade. Sulphide blocks from either pit that are below a threshold total copper value are sent to the bio-leach pad, whereas, blocks above this threshold are send to each mines specific sulphide crusher.

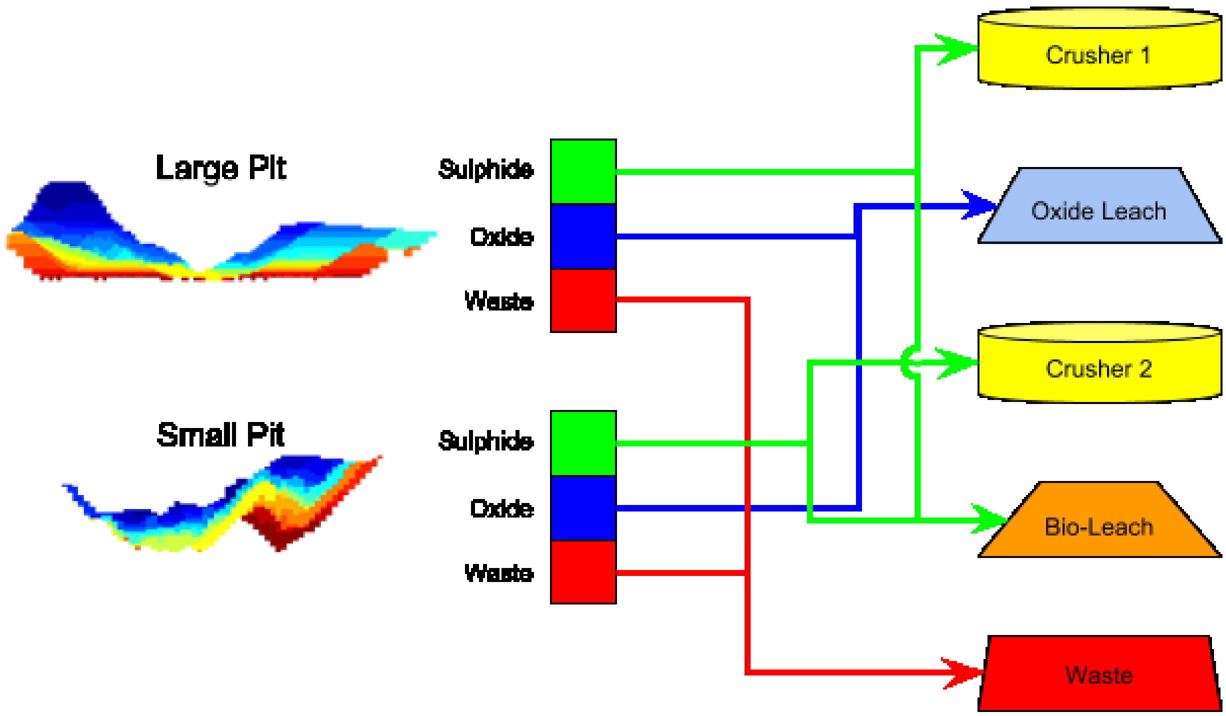

Figure 4.4 - Flow of material types from each mine to their processing destinations

The chosen period length is one month, and the extent of the material to be scheduled is one years' worth of production taken from period nine of the long-term plan (Figure 4.5). This particular year of production can be easily separated into seven mining areas; two areas in the Small Pit (0 and 1) and the remaining five in the Large Pit (2, 3, 4, 5, and 6). The long-term plan also defines the number of mobile equipment units that can be allocated, and the desired tonnage targets at each destination. The annual processing targets are equally divided into 12 parts (months) to ensure that the destinations receive uniform tonnages throughout the year. Metal grade targets are not relevant in this particular study since material type definitions already reflect minimum copper grade thresholds; grade requirements at each destination will always be respected as a result. The vertical slope angle is 45° and the horizontal precedence set for each block is determined using the method outlined in Section 4.2.3. The black blocks in Figure 4.5 represent the ramp positions for each bench in each area, which are used to define the horizontal precedence direction with a tolerance angle, β, of 46°.

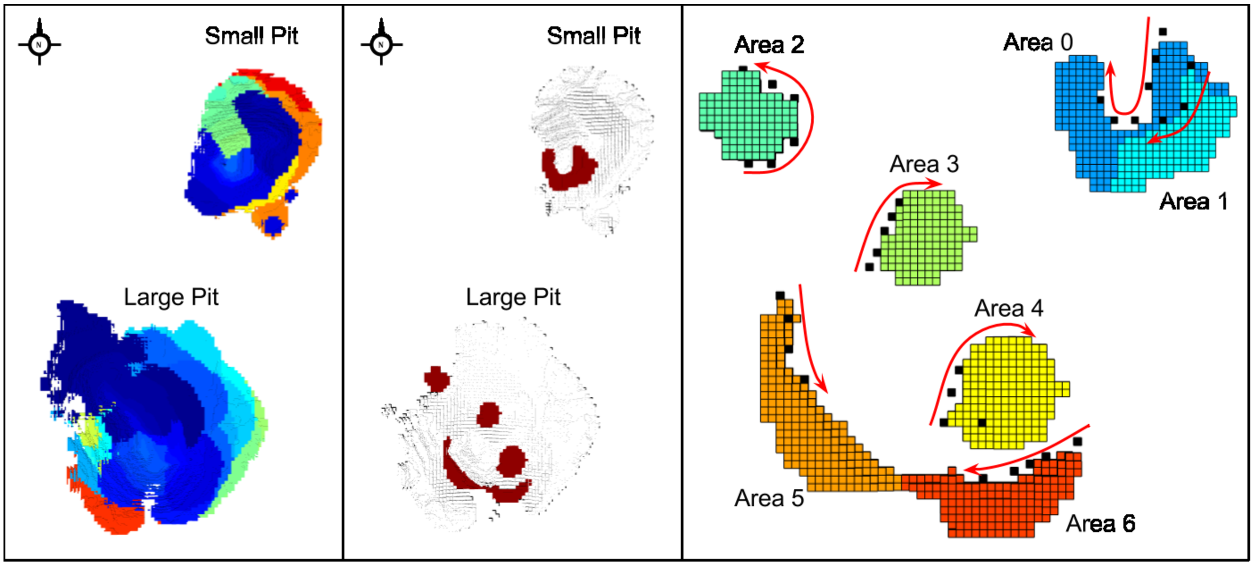

Figure 4.5 - Long-term schedule (left). Year nine of the long-term schedule (middle). The seven areas within year nine, and the descending (arrows) direction of the ramps (right)

The tonnage deviation costs at each destination can be seen in Table 4.1. The sulphide crushers have the highest priority and waste production is penalized the least. Amongst most destinations, shortages are penalized more heavily than surpluses to reflect this operation's emphasis on satisfying minimum tonnage requirements to maximize processor utilization. The complete mobile equipment fleet consists of 215 hauling trucks and 16 individual shovels (Table 4.2). The shovel movement costs between each pair of areas can be seen in Table 4.3. It should be noted that the cost of moving a shovel from one mine to the other is substantially higher than movements within the same mine. This is to heavily discourage unnecessarily long shovel moves to keep the equipment plan realistic.

Table 4.1 - Tonnage deviation penalty costs for each processing destination

Processing Destination	Shortage deviation costs (units/tonne)	Surplus deviation costs (units/tonne)
Large Pit Crusher	3.15	2.95
Small Pit Crusher	2.85	2.65
Bio-Leach	1.3	1.0
Oxide Leach	3	2
Waste Dump	2	2

Table 4.2 - Mobile equipment fleet size

Trucks			Shovels	
Model type	Payload (tonnes/trip)	Number of units	Model type	Number of units
1	350	142	1	8
2	320	54	2	3
3	220	19	3	5
		215		**16**

Table 4.3 - Shovel movement and production deviation costs

Area	Shovel movement cost (000' units/move)							Shovel deviation cost (units/tonne)		Truck deviation cost (units/hour)	
	0	1	2	3	4	5	6				
0	-	34	2,816	2,938	3,382	3,653	3,725	3	2	60	60
1	34	-	2,830	2,915	3,337	3,636	3,679	3	2	60	60
2	2,816	2,830	-	144	254	212	304	3	2	60	60
3	2,938	2,915	144	-	118	146	180	3	2	60	60
4	3,382	3,337	254	118	-	130	69	3	2	60	60
5	3,653	3,636	212	146	130	-	131	3	2	60	60
6	3,725	3,679	304	180	69	131	-	3	2	60	60

A total of 20 geological scenarios are used to map the copper grade variability for this stochastic optimization. The joint conditional simulations of total copper and soluble copper for the entire deposit are generated using DBMAFSIM. The simulated grades are then filtered through the 4,437 blocks in year nine for this short-term optimization study. The material type of each block, in each geological scenario, is then computed given the simulated copper values. The Small Pit tends to display higher copper grade variability when compared to the Large Pit, and this phenomena will become apparent in the results of the optimization. To model the equipment performance uncertainty, the PCA approach from Section 4.2.4 is used to jointly simulate 10 equipment performance scenarios consisting of the correlated availabilities, utilizations for all equipment, and hourly productions for shovels. A brief description and validation of these equipment scenarios and cycle time simulations are presented in Section 4.4.2.

4.4.2 Equipment simulation

4.4.2.1 Equipment performance

Table 4.4 - Mean and standard deviations for both original and simulated data for the Large Pit

		Mean		Standard Deviation	
		Original	Simulated	Original	Simulated
Shovel 1	Availability (%)	0.836	0.836	0.0496	0.0449
	Utilization (%)	0.679	0.685	0.0610	0.0562
	Production (t/h)	4521.0	4528.8	313.44	298.04
Shovel 2	Availability (%)	0.906	0.907	0.0246	0.0223
	Utilization (%)	0.714	0.716	0.0434	0.0399
	Production (t/h)	4810.9	4823.6	468.68	479.96
Shovel 3	Availability (%)	0.806	0.810	0.0517	0.0481
	Utilization (%)	0.760	0.765	0.0484	0.0452
	Production (t/h)	4767.0	4831.1	811.23	784.73
Truck 1	Availability (%)	0.751	0.749	0.0503	0.0460
	Utilization (%)	0.830	0.830	0.0123	0.0107
Truck 2	Availability (%)	0.812	0.811	0.0621	0.0538
	Utilization (%)	0.830	0.832	0.0232	0.0206
Truck 3	Availability (%)	0.811	0.809	0.0584	0.0515
	Utilization (%)	0.830	0.831	0.0224	0.0201

These simulations are generated using historical measurements for the monthly availability and utilizations for shovels and trucks, and hourly shovel production at the mine. This data was collected and aggregated on an equipment model basis (3 truck types and 3 shovel types), for each pit. However, since shovels are modelled independently in this SSTMPS formulation, the CDFs that are constructed in the PCA space are sampled enough times to populate the simulations for each shovel individually. Table 4.4 and Figure 4.6 show how the simulated Large Pit equipment performance scenarios reproduce the univariate statistics and bivariate correlations respectively. There are very clear relationships shown in Figure 4.6, specifically, the strong utilization correlations between the different shovel models, and similarly, between the different truck models. This is logical since equipment will tend to share similar levels of productivity when affected by the same unplanned events, such as mine shutdowns, inclement weather, etc. Most of the availabilities share these strong relationships as well since maintenance performance will tend to impact the entire fleet in a similar fashion.

		Loader 1			Loader 2			Loader 3			Truck 1		Truck 2		Truck 3	
		A	U	P	A	U	P	A	U	P	A	U	A	U	A	U
Loader 1	A	-	0.226	0.294	0.120	0.241	0.111	0.141	0.277	0.066	0.155	0.055	0.224	0.244	0.078	0.269
	U	0.226	-	0.182	0.058	0.643	0.155	0.544	0.837	0.497	0.149	0.003	0.272	0.143	0.398	0.373
	P	0.294	0.182	-	0.030	0.305	0.615	0.425	0.068	0.083	0.254	0.122	0.022	0.304	0.323	0.082
Loader 2	A	0.120	0.058	0.030	-	0.291	0.073	0.152	0.175	0.224	0.062	0.275	0.435	0.389	0.187	0.167
	U	0.241	0.643	0.305	0.291	-	0.181	0.170	0.677	0.222	0.288	0.054	0.099	0.079	0.465	0.041
	P	0.111	0.155	0.615	0.073	0.181	-	0.172	0.200	0.448	0.211	0.010	0.058	0.218	0.195	0.119
Loader 3	A	0.141	0.544	0.425	0.152	0.170	0.172	-	0.393	0.271	0.167	0.114	0.022	0.029	0.084	0.399
	U	0.277	0.837	0.068	0.175	0.677	0.200	0.393	-	0.649	0.319	0.168	0.065	0.310	0.474	0.452
	P	0.066	0.497	0.083	0.224	0.222	0.448	0.271	0.649	-	0.081	0.144	0.062	0.260	0.015	0.379
Truck 1	A	0.155	0.149	0.254	0.062	0.288	0.211	0.167	0.319	0.081	-	0.156	0.175	0.044	0.592	0.246
	U	0.055	0.003	0.122	0.275	0.054	0.010	0.114	0.168	0.144	0.156	-	0.185	0.888	0.024	0.691
Truck 2	A	0.224	0.272	0.022	0.435	0.099	0.058	0.022	0.065	0.062	0.175	0.185	-	0.150	0.047	0.137
	U	0.244	0.143	0.304	0.389	0.079	0.218	0.029	0.310	0.260	0.044	0.888	0.150	-	0.057	0.723
Truck 3	A	0.078	0.398	0.323	0.187	0.465	0.195	0.084	0.474	0.015	0.592	0.024	0.047	0.057	-	0.155
	U	0.269	0.373	0.082	0.167	0.041	0.119	0.399	0.452	0.379	0.246	0.691	0.137	0.723	0.155	-

		Loader 1			Loader 2			Loader 3			Truck 1		Truck 2		Truck 3	
		A	U	P	A	U	P	A	U	P	A	U	A	U	A	U
Loader 1	A	-	0.248	0.303	0.189	0.311	0.143	0.174	0.295	0.102	0.089	0.028	0.203	0.268	0.118	0.233
	U	0.248	-	0.133	0.034	0.650	0.217	0.595	0.864	0.498	0.191	0.035	0.212	0.217	0.340	0.435
	P	0.303	0.133	-	0.024	0.265	0.641	0.326	0.004	0.167	0.289	0.141	0.051	0.391	0.343	0.066
Loader 2	A	0.189	0.034	0.024	-	0.303	0.044	0.145	0.175	0.209	0.079	0.234	0.428	0.335	0.172	0.133
	U	0.311	0.650	0.265	0.303	-	0.136	0.224	0.682	0.199	0.307	0.045	0.077	0.108	0.459	0.086
	P	0.143	0.217	0.641	0.044	0.136	-	0.285	0.277	0.501	0.191	0.012	0.084	0.302	0.226	0.162
Loader 3	A	0.174	0.595	0.326	0.145	0.224	0.285	-	0.432	0.267	0.135	0.162	0.039	0.004	0.033	0.393
	U	0.295	0.864	0.004	0.175	0.682	0.277	0.432	-	0.641	0.337	0.174	0.056	0.346	0.429	0.468
	P	0.102	0.498	0.167	0.209	0.199	0.501	0.267	0.641	-	0.056	0.152	0.112	0.271	0.020	0.391
Truck 1	A	0.089	0.191	0.289	0.079	0.307	0.191	0.135	0.337	0.056	-	0.093	0.166	0.028	0.636	0.211
	U	0.028	0.035	0.141	0.234	0.045	0.012	0.162	0.174	0.152	0.093	-	0.190	0.847	0.046	0.656
Truck 2	A	0.203	0.212	0.051	0.428	0.077	0.084	0.039	0.056	0.112	0.166	0.190	-	0.144	0.050	0.137
	U	0.268	0.217	0.391	0.335	0.108	0.302	0.004	0.346	0.271	0.028	0.847	0.144	-	0.105	0.691
Truck 3	A	0.118	0.340	0.343	0.172	0.459	0.226	0.033	0.429	0.020	0.636	0.046	0.050	0.105	-	0.115
	U	0.233	0.435	0.066	0.133	0.086	0.162	0.393	0.468	0.391	0.211	0.656	0.137	0.691	0.115	-

Figure 4.6 - Absolute coefficient of correlation for each equipment performance parameter. Original historical data (top) and simulated data (bottom). Stronger relationships are highlighted in darker green

4.4.2.2 Truck cycle times

Table 4.5 - Mean and variances for route cycle times in minutes and minutes squared respectively

Processing destinations		From area 0	1	2	3	4	5	6
Crusher 1	Mean	-	-	62.10	69.79	55.96	41.91	41.71
	Var.	-	-	6.21	6.98	5.60	4.19	4.17
Crusher 2	Mean	45.52	42.20	-	-	-	-	-
	Var.	4.55	4.22	-	-	-	-	-
Bio-Leach	Mean	89.08	84.90	90.49	97.08	85.22	73.16	72.99
	Var.	8.91	8.49	9.05	9.71	8.52	7.32	7.30
Oxide Leach	Mean	119.52	116.47	79.67	90.10	71.34	52.27	51.99
	Var.	11.95	11.65	7.97	9.01	7.13	5.23	5.20
Waste Dump	Mean	62.73	59.10	74.59	81.77	68.86	55.72	55.54
	Var.	6.27	5.91	7.46	8.18	6.89	5.57	5.55

For this particular application, the only information available pertaining to truck cycle times is the average length of time it takes to reach each destination from each bench in each mine, invariant of the truck model. A route cycle time in this formulation is defined as the amount of time it takes a truck to complete one cycle from a given area to some destination. The route times are calculated by taking the average cycle time to each destination across all of the blocks contained in an area. A Gaussian distribution is then created, considering a variance of 10% of the mean route time. The mean and variances for each of these route times can be seen in Table 4.5. Each of these route time distributions are independently sampled 10 times and randomly linked to an equipment performance scenario to yield 10 complete equipment performance simulations.

4.4.3 Optimization results

The SSTMPS formulation in Section 4.3.1 is solved using the approach described in Section 4.3.2. CPLEX v.12.6.1.0 in a Visual Studio 13 (C++) environment is used to solve the sub-problems of each STWH iteration. The final stochastic solution is generated in approximately 9 hours on an Intel Xeon E5-2697 (2.60 GHz, 128 GB RAM). The physical short-term production schedule is displayed in Figure 4.7. The monthly mining patterns tend to be well connected and obey the directions of mining imposed by the ramp locations in Figure 4.5.

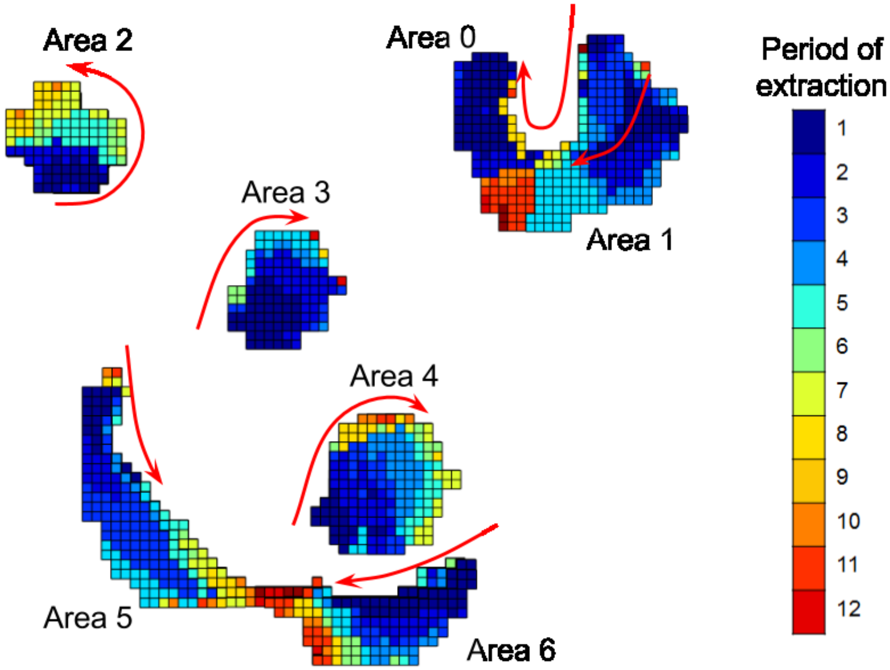

Figure 4.7 - Stochastic short-term production schedule with descending ramp directions

Figure 4.9 demonstrates how the optimizer assigns equipment to certain areas of the mine in order to move material. Allocating more shovels to area 4 near the end of the year allows for an increased extraction rate. More trucks are also required in these later months in order to deliver the extracted material. However, there is some variation as to how many trucking units are assigned in each of these months. Since different truck models have different payloads and cycle times depend on the destinations of blocks, similarly sized groupings of blocks in the same area may require significantly different trucking times. Figure 4.8 shows a best-case and a worst-case example of the shovel movement schedule. Shovel 11 only makes one very short move in the Small Pit throughout the entire year, whereas Shovel 15 makes four moves within the Large Pit. Only three shovel moves between the two pits occurred in this optimized equipment plan. These inter-pit moves could be discouraged even further using higher penalty costs, but this sensitivity was not explored in this particular study.

Shovel #11	Period	Start	1	2	3	4	5	6	7	8	9	10	11	12
	Area	1	1	1	0	0	0	0	0	0	0	0	0	0

Shovel #15	Period	Start	1	2	3	4	5	6	7	8	9	10	11	12
	Area	6	6	4	4	4	6	3	3	3	4	4	4	4

Figure 4.8 - Movement schedules for two different shovels. Areas are colour-coded

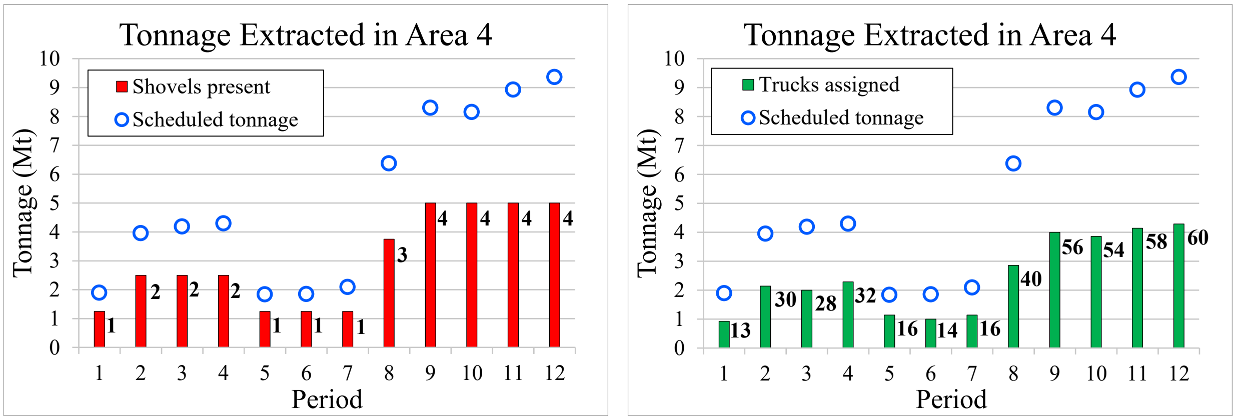

Figure 4.9 - Example of tonnage extracted from area 4. Shovels present (left) and trucks allocated (right) to extract the scheduled material

In order to demonstrate the benefits that incorporating grade and equipment uncertainty have in short-term production scheduling, the stochastic solution is benchmarked against a more conventional approach. A deterministic version of the proposed model is solved, and comparisons are drawn in terms of satisfying production targets at each destination, and minimizing equipment performance deviations. This deterministic version of the model uses an estimated orebody model and average equipment performances. More specifically, the orebody model is generated using Ordinary Kriging and the identical material type definitions are applied. The availabilities and utilizations are considered to be the average across all simulated scenarios, and the mean values for all route cycle times are used.

4.4.3.1 Impact of geological uncertainty

The monthly tonnages sent to each destination are shown in Figure 4.10. Since block destinations depend on the blocks uncertain material type, a distribution of possible tonnages sent is to be expected. The 10th, 50th and 90th percentiles of this distribution are shown in the graphs below. The blue risk profiles in the left column are the tonnages from the stochastic solution, and the orange profiles on the right are from the deterministic solution. All tonnages presented below are expressed as a percentage of the destination target for confidentiality purposes. A copper grade comparison at each destination is irrelevant in this study since destination decisions are based solely on material types; grade targets are always respected in both the stochastic and deterministic solutions. This comparison shows how the stochastic solution is able to manage the geological risk more effectively through near-optimal production scheduling. At every destination in this mining complex, the expected tonnages in the stochastic solution are within the targets more consistently, and the tonnage profile is far less variable when compared to the deterministic results. It can be noted that the satisfaction of these tonnage targets tends to deteriorate more in the later period for both the stochastic and deterministic solutions. This is an artifact of the sequential nature of the STWH heuristic (Section 4.3.2.1).

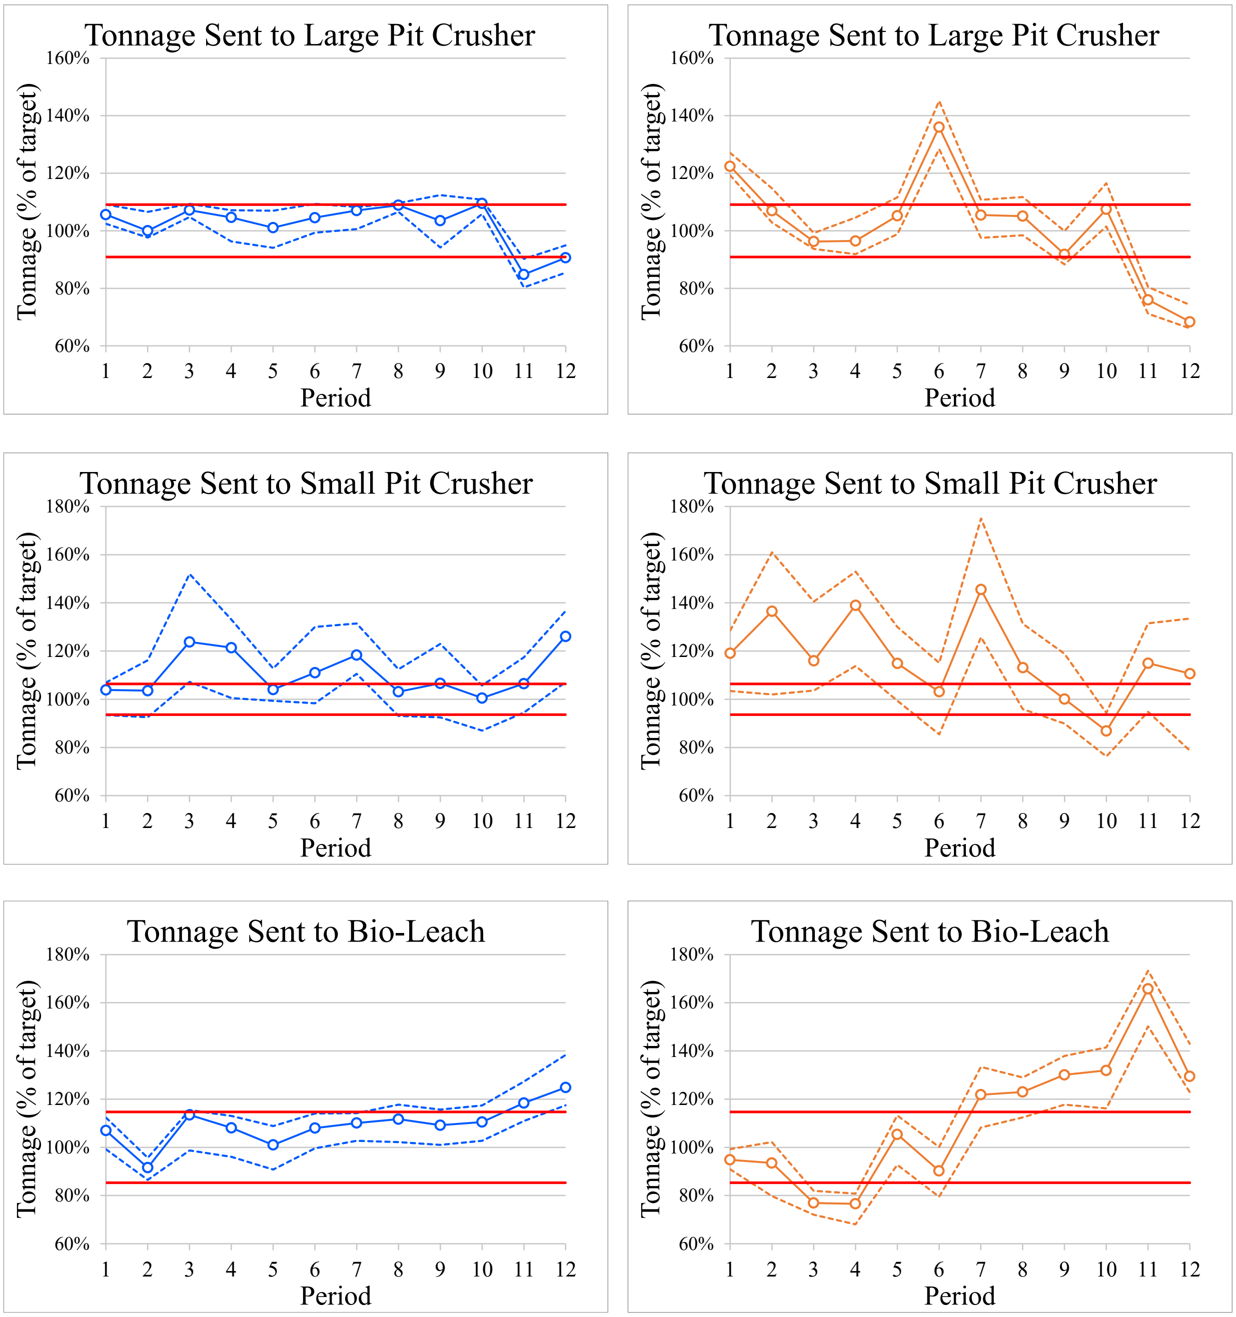

Figure 4.10 - Tonnage risk profiles for the stochastic (left) and deterministic (right) production schedules. The dotted lines are the 10th and 90th percentiles, the solid lines are the 50th percentiles, and the red lines are the destination's upper and lower targets (continued on Page 81)

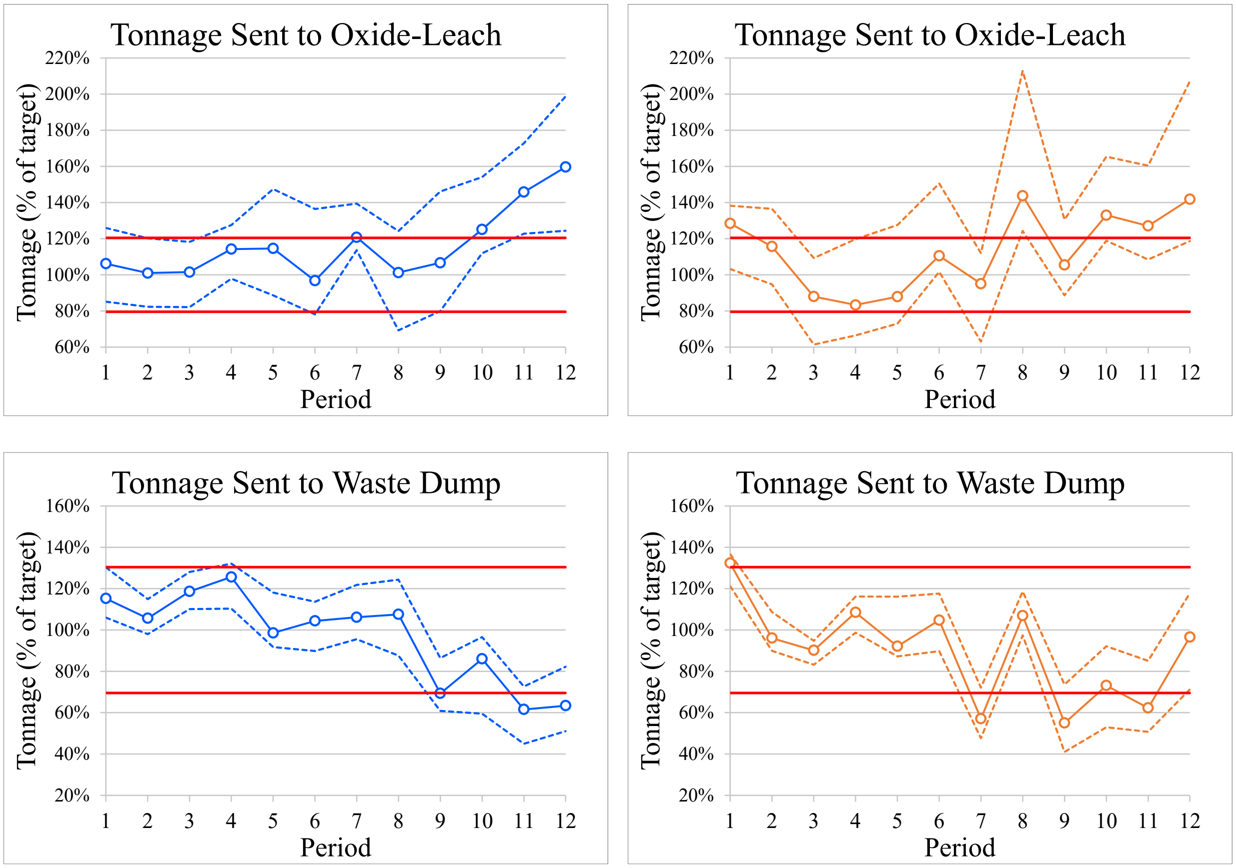

Figure 4.10 - Tonnage risk profiles for the stochastic (left) and deterministic (right) production schedules. The dotted lines are the 10th and 90th percentiles, the solid lines are the 50th percentiles, and the red lines are the destination's upper and lower targets (continued)

Through an improved understanding of material type variability, the stochastic optimizer can blend the geological risk amongst different periods better, facilitating a mine plan that can be implemented with far greater confidence. The deterministic plan is expected to frequently violate the tonnage targets which will inevitably lead to increased re-handling costs as well as decreased utilization of the processors. The greater copper grade variability present in the Small Pit can be observed from the results in Figure 4.10. This pit provides the majority of the oxide-leach material and all of the Small Pit crusher material; both of which display the highest degree of tonnage variability in the stochastic and deterministic solutions alike. Even still, the stochastic solution tends to better manage this variability and deliver more certain monthly tonnages to these destinations.

4.4.3.2 Impact of equipment uncertainty

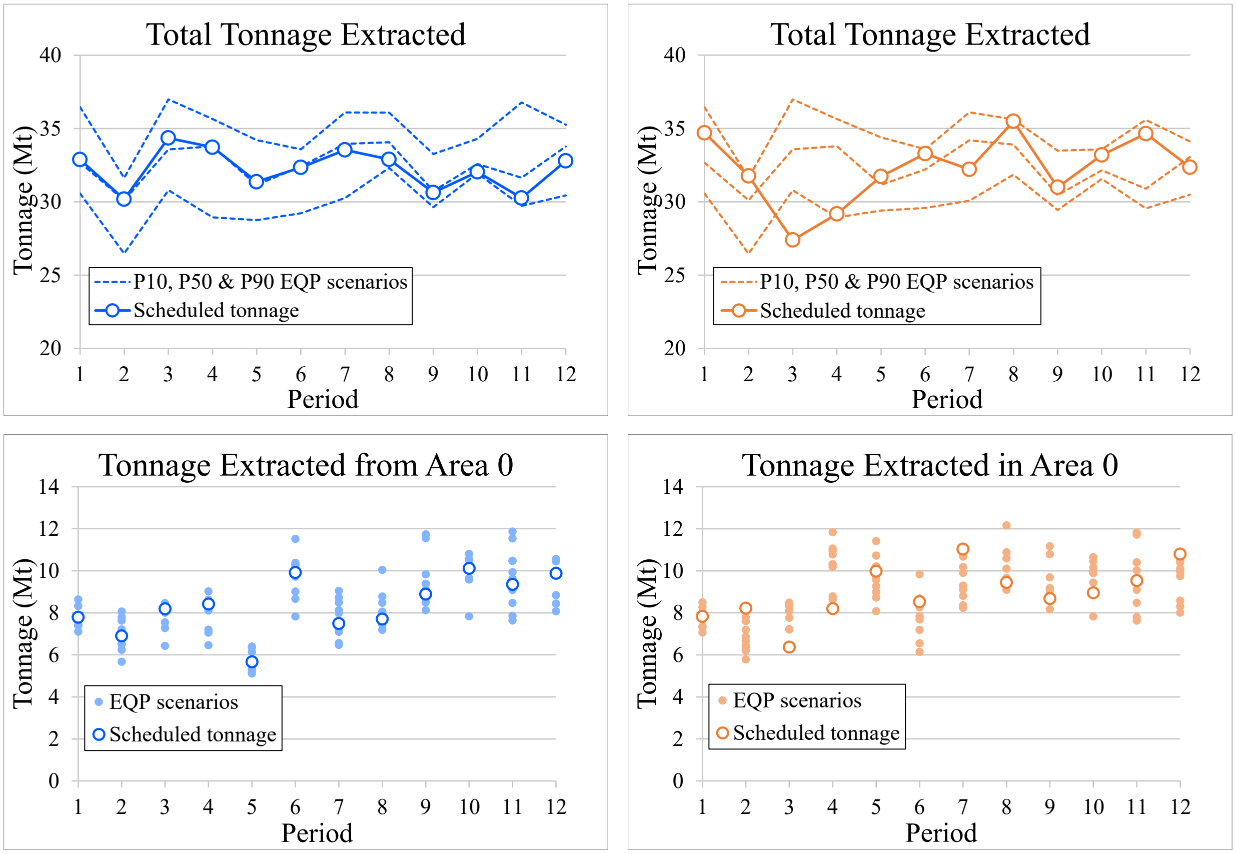

Figure 4.11 - Scheduled production each month compared to the uncertain equipment production for the stochastic solution (left) and deterministic solution (right). Total tonnage extracted (top) and an example of tonnage extracted from one area (bottom)

Figure 4.11 shows the production risk profiles of both solutions with respect to the equipment uncertainty. The integration of the equipment variability tends to allow the stochastic optimizer to yield a production schedule that is far more likely to be achieved. The stochastic production is comfortably within the expected operating ranges throughout the year (Figure 4.11; left). In contrast, the deterministic solution often chooses to extract groupings of blocks that are highly unlikely to be produced by the mobile fleet (Figure 4.11; right). The overall risk profile of the stochastic schedule is only slightly tighter than the deterministic. This is because the historical hourly productions are very similar amongst all loading units in this study. There is little to be gained by grouping different shovels together in the same areas, whose production profiles would complement each other's to minimize these variabilities by significant amounts.

Also, the total extraction graphs can be misleading since lower and upper equipment production deviations from different areas may appear to balance out. The production in Area 0 (Figure 4.11; bottom) and the absolute shovel deviations in Figure 4.12 are more reliable in demonstrating how the stochastic solution yields production schedules that are far more attainable.

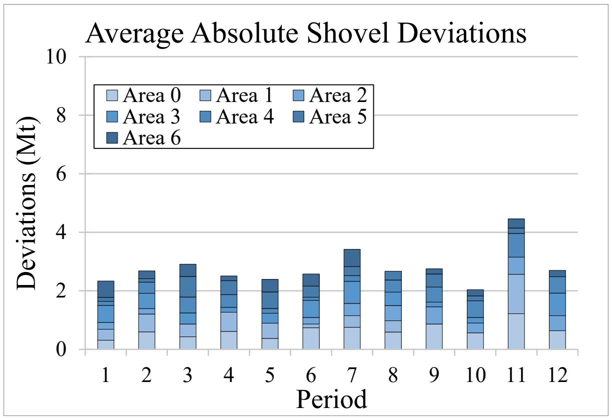

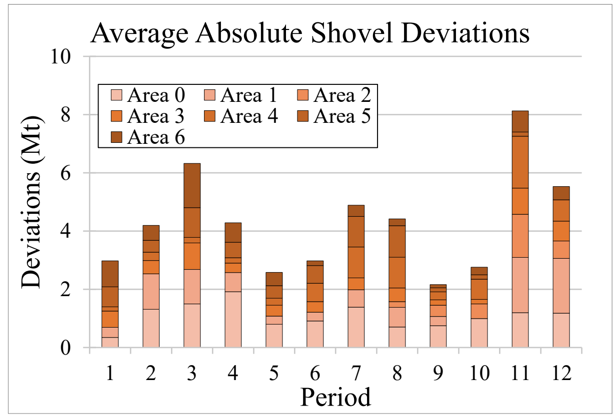

Figure 4.12 - Absolute shovel production deviations in each area averaged across all equipment performance scenarios for the stochastic (left) and deterministic (right) solutions

Figure 4.13 shows the cumulative absolute truck hour deviations averaged across all cycle time scenarios. This graph is effective in demonstrating the overall magnitudes of truck hour deviations between both solutions. Although more consistent, the overall magnitude of deviations in the deterministic case is equal to, if not less than, that of the stochastic solution. From this, it can be inferred that there is little benefit of including cycle time uncertainty in the manner proposed herein. There are two major factors that may contribute to this conclusion. (1) By some incitement of the estimated resource model, the deterministic optimizer decides to extract less total tonnage than in the stochastic solution, specifically in waste and oxide material (Figure 4.10). Having relatively long route cycle times, under-producing these material types places less strain on the total truck production which results in less truck hour shortages. This can be further supported by the fact that 70% of the total truck hour deviations are shortages in the stochastic solution compared to 60% in the deterministic case. (2) More importantly, the chosen cycle time simulation approach adds very little value to the stochastic solution. The symmetric nature of Gaussian distributions, coupled with equivalent shortage and surplus truck hour deviation penalties (Table 4.3), allows for cycle time uncertainty to be fully characterized by its mean value; just as proposed in the deterministic approach. In order for there to be value gained from optimal truck assignments through incorporating cycle time uncertainty, more insights would need to be drawn from a better understanding of the factors that result in cycle time variability. Some potential factors could be: truck payload densities for different material types, road quality variations by season, ramp grade, etc.

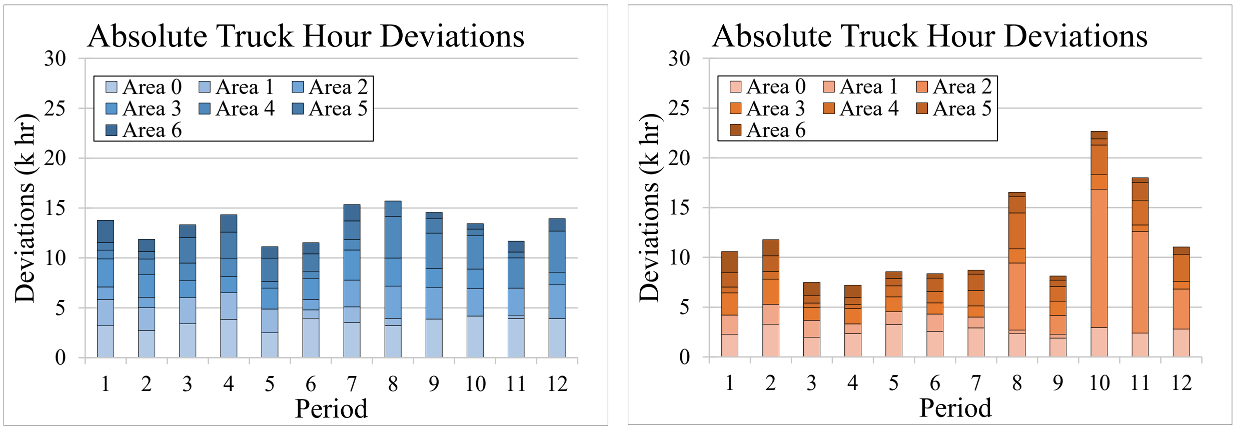

Figure 4.13 - Absolute trucking hour deviations in each area averaged across all cycle time scenarios for the stochastic (left) and deterministic (right) solutions

4.5 Conclusions and future work

This work proposes a new stochastic short-term model to simultaneously optimize the equipment plan and production schedule of mining complexes while incorporating both geological and equipment performance uncertainty. This formulation introduces: a new approach to generate more realistic equipment performance scenarios, a more accurate shovel movement formulation, and an efficient means to encourage connected extraction patterns with the notion of horizontal precedence. The SSTMPS formulation is applied to a very large copper mining complex and the results are compared to conventional methods to show the benefits of including geological and equipment performance uncertainties in the problem. The stochastic formulation generates practical equipment plans and extraction patterns while successfully controlling the operational risk in terms of delivering more reliable tonnages to the processors, and defining an extraction sequence that is far more likely to be achieved.

There are three main areas that may be improved and should be the focus for future work. (1) The complexity and size of the case study presented in this chapter approaches the limits that commercial solvers can handle efficiently. In order to accommodate more blocks, more mining areas and more detailed equipment decisions, future work on short-term problems should be developed using customized metaheuristic solvers. (2) More comprehensive cycle time distributions should be investigated. There is significant value to be gained through allowing new stochastic solvers to make well-informed decisions considering complex equipment variability factors in short-term production scheduling. (3) In order to expand into different areas of mine plan optimization, different short-term aspects could be optimized in parallel to the production schedule, such as: optimal equipment maintenance scheduling, optimal temporary ramp placement, and optimal equipment usage in terms of loading factors and trucking speeds.

Chapter 5

Conclusions and Future Work

5.1 Conclusions and objectives met

This thesis has shown that integrating uncertainty into mine production scheduling can increase the expected value of a mining project, significantly reduce the risk of deviating from key production targets, and minimize the risk of deviating from the physical schedule itself. These results are attributed to a stochastic optimizer's ability to explicitly manage uncertainty through the incorporation of multiple simulated scenarios. The optimizer then understands the spatial variability of extreme material grades throughout the deposit and the variability in how the mobile mining fleet will perform. This thesis explores both long and short-term stochastic mine production scheduling techniques and demonstrates how the inclusion of different sources of uncertainty offers substantial benefit when compared to traditional, deterministic methods. Below, the findings of this thesis are expressed in the context of the four original objectives in Section 1.1.

1. **Review the literature pertaining to the stochastic optimization of open-pit mine production scheduling**.

A critical literature review is performed on the topics pertinent to the developments in long and short-term mine production scheduling and the stochastic simulation of mineral deposits. One of the primary findings in Section 2.2 is the inherent limitations of traditional deterministic optimization techniques. Estimated orebody models are unable to represent the true spatial variability nor the presence of extreme grades in mineral deposits. As a result, deterministic optimizers cannot understand this variability and will not necessarily perform well under uncertainty due to the non-linearity of mine production scheduling optimization. Stochastic optimizers are introduced to overcome these drawbacks, and over time, have shown to improve the expected economic value of mining projects while also reducing the risk of deviating from key production targets. The review of short-term scheduling practices in Section 2.3 uncovers a lack of techniques that simultaneously optimize the mobile equipment plan with the production schedule at the block-level. Furthermore, very few works have addressed how uncertainty can be incorporated into short-term production scheduling; both in terms of uncertain geology and uncertain equipment performance. One key emphasis of future short-term production scheduling techniques should be the generation of more practical mining patterns in order to make the transition from planning to operations as seamless as possible. Lastly, the literature in Section 2.4 reveals that significant efficiency improvements have been developed for sequential Gaussian simulation techniques. However, future developments in the field of geostatistical simulation will require deviations from Gaussian frameworks through the use of high-order statistical methods in order to reproduce more complex geological patterns.

2. **Perform a stochastic optimization of the long-term production schedule of a rare earth element project considering uncertain geology.**

The principles of stochastic mine planning are applied to a commodity that has received little attention from academics in the field of mine optimization. First, the spatial variability of a REE deposit is quantified using a two-stage simulation approach (Goodfellow et al., 2012). The boundaries of the mineralized domain are simulated using SNESIM, and the REE metal grades are jointly simulated using DBMAFSIM to preserve the strong cross-elemental relationships. The primary contribution is the avoidance of the TREO grade convention, common in REE orebody modelling, which tends to misconstrue the locations of the valuable, but scarce, HREEs. A stochastic mine production scheduling formulation is proposed that maximizes the expected NPV while minimizing the risk of deviating from mineral and REE blending targets. The stochastic formulation is solved using and efficient, parallelized MNTS metaheuristic method.

3. **Develop a stochastic formulation for the simultaneous optimization of the short-term equipment plan and production schedule of a very large copper mining complex.**

A formulation is proposed to simultaneously optimize the short-term equipment plan and production schedule under both geological and equipment performance uncertainty. The proposed approach rectifies certain limitations of previous work in stochastic short-term planning by: incorporating a location-dependant shovel movement optimization into the formulation; generating more realistic equipment performance scenarios using a decomposition technique; developing a new approach to facilitate practical designs using existing ramp locations; and lastly, making improvements to the formulation to allow for the efficient optimization of very large problem instances. The model is implemented at a large copper mining complex and the problem is solved in approximately 9 hours using a general purpose solver (CPLEX). In order to facilitate the optimization of such large applications, STWH and ES heuristic methods are proposed to generate near-optimal solutions more quickly.

4. **Benchmark these stochastic optimization frameworks against the industry's best practices**

Both stochastic mine production scheduling exercises in Chapters 3 and 4 are applied to realistic mining case studies so they can be evaluated and benchmarked against the industry's current practices to demonstrate the benefits of incorporating uncertainty during the optimization process.

The long-term stochastic mine production schedule of the REE project is benchmarked against a traditional deterministic mine optimization package (Whittle). This comparison shows that the stochastic optimizer is able to generate a mine production schedule that: delivers a more appropriate amount of ore each period that will result in less rehandling costs and increased utilization of the processing facility; delivers an ore feed that displays superior mineral properties which will lead to higher recoveries at the flotation stage; delivers a didymium product that is far more likely to meet the buyer's strict quality constraints which will minimize the incurrence of penalty fees; and lastly, increases the expected NPV by 20%. These results are attributed to the stochastic optimizer's ability to understand the presence of extreme metal grades, and explicitly manage geological risk.

The stochastic short-term optimizer is applied to one of the largest copper mining complexes in the world. Due to the use of horizontal block precedence, this optimizer is able to efficiently generate production schedules that are practical and will limit unnecessary equipment movement. The stochastic formulation is compared to a more traditional approach, where the same formulation is implemented using averaged inputs for geology and equipment performance. The stochastic model tends to be more effective in mitigating the risk of deviating from tonnage targets at each processing destination. This will ensure less material rehandling, better utilization, and fewer unexpected changes to the design of the waste dumps and leach pads. Furthermore, the integration of equipment performance variability allows the stochastic optimizer to better understand the uncertain extraction capacities throughout the mine, and tends to generate a block extraction sequence that is far more likely to be achieved.

5.2 Future work

The mine production scheduling formulations presented in this thesis are successful in adding value to mining operations while minimizing the deviations from various production goals. However, the long and short-term approaches proposed are not without limitations.

The stochastic long-term production scheduling model for REE operations, proposed in Chapter 3, has three major venues of future work. First, the physical production schedules generated using this MNTS method should be made much more practical. Both minimum mining widths and maximum bench sink rates should be integrated into the formulation in order to generate near-optimal production schedules that can actually be implemented at a mine without significant re-design. Second, the destination policy in the proposed SIP formulation should be improved. Although the idea of solving for the destination of a block regardless of geological scenario is simple and robust to uncertainty, it is also overly pessimistic. Some manner of corrective recourse should be included without creating an independent destination policy for each geological scenario. The robust destination policy proposed by Goodfellow and Dimitrakopoulos (2016) would work well with the multi-element nature of REE projects. Lastly, future work should aim to test this stochastic methodology on REE operations with more complicated grade blending requirements and downstream processing decisions. Global optimization techniques, briefly outlined in Section 2.2.3, could unlock substantial value in the presence of more complex downstream aspects of REE projects.

The formulation for the simultaneous optimization of the short-term equipment plan and extraction sequence, proposed in Chapter 4, has four major areas for future research. First, future work should aim to improve the link between the uncertain performance of equipment and the physical production schedule. More specifically, the implications this type of uncertainty has on the production targets. In the proposed formulation, discrepancies between the production schedule and the equipment production are only quantified through tonnage deviations. Instead, the optimizer should understand the spatial (block-level) context of these discrepancies. The optimizer should be aware of precisely which blocks fail to be mined or are unexpectedly mined in a given period in order to accurately quantify the grade and tonnage production deviations caused by equipment uncertainty. Second, the size of the case study presented in Section 4.4 approaches the complexity limits that current general purpose solvers can handle efficiently. In order to accommodate more blocks, more mining areas, and more detailed equipment decisions, these short-term models should be developed using customized metaheuristic solvers in the future. Third, more comprehensive cycle time distributions should be investigated. There is significant value to be gained through allowing stochastic solvers to make well-informed decisions in matching trucks to shovels within the framework of production planning to minimize idle times and maximize productivity. Lastly, In order to expand into different areas of mine plan optimization, different short-term aspects should be optimized parallel to the production schedule, such as optimal equipment maintenance scheduling, optimal temporary ramp placement, and optimal equipment usage in terms of loading factors and trucking speeds, etc.

References

Alarie, S., and Gamache, M. (2002). Overview of Solution Strategies Used in Truck Dispatching Systems for Open Pit Mines. International Journal of Surface Mining, Reclamation and Environment, 16(1), 59-76.

Avalon Rare Metals Inc. (2013). Nechalacho Rare Earth Elements Project - Feasibility Study. NI 43-101 Technical Report, MICON International Ltd.

Benndorf, J., and Dimitrakopoulos, R. (2005). New efficient methods for conditional simulation of large orebodies. Orebody Modelling and Strategic Mine Planning, 14, 7.

Benndorf, J., and Dimitrakopoulos, R. (2013). Stochastic long-term production scheduling of iron ore deposits: Integrating joint multi-element geological uncertainty. Journal of Mining Science, 49(1), 68-81.

Birge, J. R., and Louveaux, F. (1997). Introduction to stochastic programming (2 Ed.): Springer.

Bley, A., Boland, N., Fricke, C., and Froyland, G. (2010). A strengthened formulation and cutting planes for the open pit mine production scheduling problem. Computers and Operations Research, 37, 1641-1647.

Boland, N., Dumitrescu, I., and Froyland, G. (2008). A multistage stochastic programming approach to open pit mine production scheduling with uncertain geology. Optimization Online.

Boucher, A., and Dimitrakopoulos, R. (2009). Block Simulation of Multiple Correlated Variables. Mathematical Geosciences, 41(2), 215-237.

Boucher, A., and Dimitrakopoulos, R. (2012). Multivariate Block-Support Simulation of the Yandi Iron Ore Deposit, Western Australia. Mathematical Geosciences, 44(4), 449-468.

Burt, C., Caccetta, L., Fouché, L., and Welgama, P. (2016). An MILP approach to Multi-location, Multi-Period Equipment Selection for Surface Mining with Case Studies. Journal of industrial and management optimization, 12(2), 403-430.

Caccetta, L., and Hill, S. P. (2003). An application of branch and cut to open pit mine scheduling. Journal of Global Optimization, 27(2-3), 349-365.

Commerce Resources Corp. (2012). Ashram Project - Preliminary Economic Assessment. NI 43-101 Technical Report, SGS Canada Inc.

Cullenbine, C., Wood, R. K., and Newman, A. (2011). A sliding time window heuristic for open pit mine block sequencing. Optimization Letters, 5(3), 365-377.

Dagdelen, K., and Johnson, T. (1986). Optimum multi period open pit mine production scheduling by Lagrangian Parameterization. Paper presented at the 19th APCOM Symposium of the Society of Mining Engineers, 127-142.

David, M. (1988). Handbook of applied advanced geostatistical ore reserve estimation. Amsterdam; New York: Elsevier: Distributors for the United States and Canada Elsevier Science Pub. Co.

David, M., Dowd, P., and Korobov, S. (1974). Forecasting departure from planning in open pit design and grade control. Paper presented at the 12th International Applications of Computers and Operations Research in the Minerals Industry (APCOM), Golden, Colorado.

Davis, M. (1987). Production of conditional simulations via the LU triangular decomposition of the covariance matrix. Mathematical Geology, 19(2), 91-98.

Desbarats, A. J., and Dimitrakopoulos, R. (2000). Geostatistical Simulation of Regionalized Pore-Size Distributions Using Min/Max Autocorrelation Factors. Mathematical Geology, 32(8), 919-942.

Dimitrakopoulos, R., Farrelly, C. T., and Godoy, M. (2002). Moving forward from traditional optimization: grade uncertainty and risk effects in open-pit design. Mining Technology, 111(1), 82-88.

Dimitrakopoulos, R., and Godoy, M. (2014). Grade control based on economic ore/waste classification functions and stochastic simulations: examples, comparisons and applications. Mining Technology, 123(2), 17.

Dimitrakopoulos, R., and Jewbali, A. (2013). Joint stochastic optimisation of short and long term mine production planning: method and application in a large operating gold mine. Mining Technology, 122(2), 110-123.

Dimitrakopoulos, R., and Luo, X. (2004). Generalized sequential Gaussian simulation on group size ν and screen-effect approximations for large field simulations. Mathematical Geology, 36(5), 567-591.

Dimitrakopoulos, R., Martinez, L., and Ramazan, S. (2007). A maximum upside / minimum downside approach to the traditional optimization of open pit mine design. Journal of Mining Science, 43(1), 73-82.

Dimitrakopoulos, R., Mustapha, H., and Gloaguen, E. (2009). High-order Statistics of Spatial Random Fields: Exploring Spatial Cumulants for Modeling Complex Non-Gaussian and Non-linear Phenomena. Mathematical Geosciences, 42(1), 65.

Dimitrakopoulos, R., and Ramazan, S. (2004). Uncertainty-based production scheduling in open pit mining. SME Transactions, 316, 106-112.

Dimitrakopoulos, R., and Ramazan, S. (2008). Stochastic integer programming for optimising long term production schedules of open pit mines: methods, application and value of stochastic solutions. Mining Technology, 117(4), 155-160.

Eivazy, H., and Askari-Nasab, H. (2012). A mixed integer linear programming model for short-term open pit mine production scheduling. Mining Technology, 121(2), 97-108.

Ernst, and Young. (2010). Material risk - Access to technology minerals.

Fu, Z., Topal, E., and Erten, O. (2014). Optimisation of a mixed truck fleet schedule through a mathematical model considering a new truck-purchase option. Mining Technology, 123, 30-35.

Fytas, K. G. (1985). A computerized model of open pit long and short term production scheduling. (Ph.D. thesis), Queen's University, Kingston, Ontario, Canada.

Geman, S., and Geman, D. (1984). Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. IEEE Transactions on Pattern Analysis and Machine Intelligence, 6(6), 721-741.

Gershon, M. E. (1983). Optimal mine production scheduling: evaluation of large scale mathematical programming approaches. International Journal of Mining Engineering, 1(4), 315-329.

Glover, F., and Laguna, M. (1997). Tabu Search. Kluwer Academic Publishers.

Godoy, M. (2003). The effective management of geological risk in long-term production scheduling of open pit mines. (Ph.D. thesis), University of Queensland, Brisbane, Australia.

Godoy, M., and Dimitrakopoulos, R. (2004). Managing risk and waste mining in long-term production scheduling of open-pit mines. SME Transactions, 316, 43-50.

Goodfellow, R., Albor Consuegra, F., Dimitrakopoulos, R., and Lloyd, T. (2012). Quantifying multi-element and volumetric uncertainty, Coleman McCreedy deposit, Ontario, Canada. Computers and Geosciences, 42(0), 71-78.

Goodfellow, R., and Dimitrakopoulos, R. (2016). Global optimization of open pit mining complexes with uncertainty. Applied Soft Computing, 40, 292-304.

Goovaerts, P. (1993). Spatial orthogonality of the principal components computed from coregionalized variables. Mathematical Geology, 25(3), 281-302.

Goovaerts, P. (1997). Geostatistics for natural resources evaluation. New York: Oxford University Press.

Great Western Minerals Group Ltd. (2014). Steenkampskraal Rare Earth Element Project - Feasibility Study. NI 43-101 - Technical Report, Venmyn Deloitte.

Grieco, N., and Dimitrakopoulos, R. (2007). Managing grade risk in stope design optimisation: probabilistic mathematical programming model and application in sublevel stoping. Mining Technology, 116(2), 49-57.

Guardiano, F. B., and Srivastava, R. M. (1993). Multivariate Geostatistics: Beyond Bivariate Moments. In A. Soares (Ed.), Geostatistics Tróia '92 (Vol. 5, pp. 133-144): Springer Netherlands.

Hotelling, H. (1933). Analysis of a complex of statistical variables into principal components. Journal of Educational Psychology, 24(6), 417-441.

IAMGOLD. (2012). Niobec Mine Property - Preliminary Economic Assessment NI 43-101 - Technical Report, IAMGOLD Corporation.

Isaaks, E. H. (1990). The application of Monte Carlo methods to the analysis of spatially correlated data. (Ph.D. thesis), Stanford University, Stanford California.

Isaaks, E. H., and Srivastava, R. M. (1989). Applied Geostatistics. New York: Oxford University Press.

Jain, S., and Foley, W. J. (2016). Dispatching strategies for managing uncertainties in automated manufacturing systems. European Journal of Operational Research, 248(1), 328-341.

Johnson, T. (1968). Optimum open pit mine production scheduling: University of California: Berkley.

Journel, A. G. (1974). Geostatistics for conditional simulation of ore bodies. Economic Geology, 69.

Journel, A. G. (1979). Geostatistical simulation - Ch. 6 Methods for exploration and mine planning. Engineering and Mining Journal, 180(12), 6.

Journel, A. G. (2007). Roadblocks to the evaluation of ore reserves - the simulation overpass and putting more geology into numerical models of deposits. Parkville Victoria: Australasian Institute of Mining and Metallurgy.

Kirkpatrick, S., Gelatt, C. D., and Vecchi, M. P. (1983). Optimization by Simulated Annealing. Science, 220(4598), 671.

Kumral, M., and Dowd, P. A. (2002). Short-term mine production scheduling for industrial minerals using multi-objective simulated annealing. Paper presented at the Proceedings 30th APCOM Symposium, Fairbanks, Alaska, USA.

L'Heureux, G., Gamache, M., and Soumis, F. (2013). Mixed integer programming model for short term planning in open-pit mines. Mining Technology, 122(2), 101-109.

Lamghari, A., and Dimitrakopoulos, R. (2012). A diversified Tabu search approach for the open-pit mine production scheduling problem with metal uncertainty. European Journal of Operational Research, 222(3), 642-652.

Lamghari, A., Dimitrakopoulos, R., and Ferland, J. A. (2014). A variable neighbourhood descent algorithm for the open-pit mine production scheduling problem with metal uncertainty. Journal of the Operational Research Society, 65(9), 1305-1314.

Leite, A., and Dimitrakopoulos, R. (2007). Stochastic optimisation model for open pit mine planning: application and risk analysis at copper deposit. Mining Technology, 116(3), 109-118.

Leite, A., and Dimitrakopoulos, R. (2014). Stochastic optimization of mine production scheduling with uncertain ore/metal/waste supply. International Journal of Mining Science and Technology, 24(6), 755-762.

Lerchs, H., and Grossmann, I. F. (1965). Optimum design of open-pit mines. CIM Bulletin, 58, 8.

Liu, Y. (2006). Using the Snesim program for multiple-point statistical simulation. Computers and Geosciences, 32(10), 1544-1563.

Mariethoz, G., Renard, P., and Straubhaar, J. (2010). The Direct Sampling method to perform multiple-point geostatistical simulations. Water Resources Research, 46(11).

Matamc Explorations Inc. (2013). Kipawa Project - Feasibility Study. NI 43-101 - Technical Report, Roche Ltd.

Matamoros, M. E. V., and Dimitrakopoulos, R. (2016). Stochastic short-term mine production schedule accounting for fleet allocation, operational considerations and blending restrictions. European Journal of Operational Research, 255(3), 911-921.

Menabde, M., Froyland, G., Stone, P., and Yeates, G. (2007). Mining Schedule Optimisation for Conditionally Simulated Orebodies. AusIMM Spectrum Series, 14, 379-384.

Molycorp Inc. (2012). Mountain Pass Rare Earth Project - Feasibility Study. NI 43-101 - Technical Report, SRK Consulting.

Montiel, L., and Dimitrakopoulos, R. (2015). Optimizing mining complexes with multiple processing and transportation alternatives: An uncertainty-based approach. European Journal of Operational Research, 247(1), 166-178.

Montiel, L., Dimitrakopoulos, R., and Kawahata, K. (2016). Globally optimising open-pit and underground mining operations under geological uncertainty. Mining Technology, 125(1), 2-14.

Moreno, E., Espinoza, D., and Goycoolea, M. (2010). Large-scale multi-period precedence constrained knapsack problem: A mining application. Electronic Notes in Discrete Mathematics, 36, 407-414.

Mousavi, A., Kozan, E., and Liu, S. Q. (2016). Open-pit block sequencing optimization: A mathematical model and solution technique. Engineering Optimization, 1-19.

Mustapha, H., and Dimitrakopoulos, R. (2010). High-order Stochastic Simulation of Complex Spatially Distributed Natural Phenomena. Mathematical Geosciences, 42(5), 457-485.

Mustapha, H., and Dimitrakopoulos, R. (2011). HOSIM: A high-order stochastic simulation algorithm for generating three-dimensional complex geological patterns. Computers and Geosciences, 37(9), 1242-1253.

Newman, A. M., Rubio, E., Caro, R., Weintraub, A., and Eurek, K. (2010). A review of operations research in mine planning. Interfaces, 40(3), 222-245.

Panagiotidou, S., and Tagaras, G. (2007). Optimal preventive maintenance for equipment with two quality states and general failure time distributions. European Journal of Operational Research, 180(1), 329-353.

Pochet, Y., and Wolsey, L. A. (2006). Production planning by mixed integer programming: Springer Science and Business Media.

Quest Rare Minerals Ltd. (2014). Strange Lake Property - Preliminary Economic Assessment. NI 43-101 Technical Report, MICON International Ltd.

Ramazan, S., and Dimitrakopoulos, R. (2004a). Recent applications of operations research and efficient MIP formulations in open pit mining. SME Transactions, 315, 73-78.

Ramazan, S., and Dimitrakopoulos, R. (2004b). Traditional and new MIP models for production scheduling with in-situ grade variability. International Journal of Surface Mining, Reclamation and Environment, 18(2), 85-98.

Ramazan, S., and Dimitrakopoulos, R. (2007). Stochastic optimization of long-term production scheduling for open pit mines with a new integer programming formulation. Paper presented at the Orebody Modelling and Strategic Mine Planning.

Ramazan, S., and Dimitrakopoulos, R. (2013). Production scheduling with uncertain supply: a new solution to the open pit mining problem. Optimization and Engineering, 14(2), 361-380.

Ravenscroft, P. J. (1992). Risk analysis for mine scheduling by conditional simulation. Transactions Australian Intitute of Mining and Metallurgy, pp 104-108.

Rondon, O. (2012). Teaching Aid: Minimum/Maximum Autocorrelation Factors for Joint Simulation of Attributes. Mathematical Geosciences, 44(4), 469-504.

Rosenblatt, M. (1952). Remarks on a multivariate transformation. Annals of Mathematical Statistics, 23, 3.

Roskill. (2014). What does the WTO ruling mean for the Chinese rare earth industry? Roskill briefing paper.

Senécal, R., and Dimitrakopoulos, R. (2014). Long-term mine production scheduling with multiple processing destinations: a new and efficient multi-neighborhood tabu search metaheuristic implementation. (M.Sc. thesis). McGill University.

Smith, M. L. (1998). Optimizing short-term production schedules in surface mining: Integrating mine modeling software with AMPL/CPLEX. International Journal of Surface Mining, Reclamation and Environment, 12(4), 149-155.

Straubhaar, J., Renard, P., Mariethoz, G., Froidevaux, R., and Besson, O. (2011). An Improved Parallel Multiple-point Algorithm Using a List Approach. Mathematical Geosciences, 43(3), 305-328.

Strebelle, S. (2002). Conditional Simulation of Complex Geological Structures Using Multiple-Point Statistics. Mathematical Geology, 34(1), 1-21.

Strebelle, S., and Zhang, T. (2005). Non-Stationary Multiple-point Geostatistical Models. In O. Leuangthong and C. V. Deutsch (Eds.), Geostatistics Banff 2004 (pp. 235-244). Dordrecht: Springer Netherlands.

Switzer, P., Green, A. A. (1984). Min/max autocorrelation factors for multivariate spatial imagery. Standford University, Department of Statistics(6). Stanford, California.

Tolwinski, B., and Underwood, R. (1996). A scheduling algorithm for open pit mines. IMA Journal of Management Mathematics, 7(3), 247-270.

Topal, E. (2008). Early start and late start algorithms to improve the solution time for long-term underground mine production scheduling. Journal of the Southern African Institute of Mining and Metallurgy, 108, 99-107.

Topal, E., and Ramazan, S. (2010). A new MIP model for mine equipment scheduling by minimizing maintenance cost. European Journal of Operational Research, 207(2), 1065-1071.

Topal, E., and Ramazan, S. (2012). Mining truck scheduling with stochastic maintenance cost. Journal of Coal Science and Engineering (China), 18(3), 313-319.

Vallée, M. (2000). Mineral resource + engineering, economic and legal feasibility = ore reserve. CIM Bulletin, 93(1038), 53-61.

Vargas-Guzmán, J. A., and Dimitrakopoulos, R. (2002). Conditional Simulation of Random Fields by Successive Residuals. Mathematical Geology, 34(5), 597-611.

Verly, G. W. (1993). Sequential Gaussian Cosimulation: A Simulation Method Integrating Several Types of Information. In A. Soares (Ed.), Geostatistics Tróia '92: Volume 1 (pp. 543-554). Dordrecht: Springer Netherlands.

Wackernagel, H. (2003). Multivariate geostatistics an introduction with applications (3 ed.). Berlin: Springer.

Whittle, J. (1988). Beyong optimization in open pit design. Paper presented at the 1st Canadian Conference on Computer Applications in the Mineral Industries.

Wilke, F. L., and Reimer, T. H. (1977). Optimizing the short-term production schedule for an open-pit iron ore mining operation. Paper presented at the Proceedings 15th International APCOM Symposium, Brisbane, Australia.