Data assimilation in the minerals industry

Real-time updating of spatial models using online production data

Wambeke, Tom

DOI
10.4233/uuid:3acfe30a-1c01-4851-b491-ca20b3b459ce

Publication date
2018

Document Version
Final published version

Citation (APA)
Wambeke, T. (2018). *Data assimilation in the minerals industry: Real-time updating of spatial models using online production data*. [Dissertation (TU Delft), Delft University of Technology].
https://doi.org/10.4233/uuid:3acfe30a-1c01-4851-b491-ca20b3b459ce

DATA ASSIMILATION IN THE MINERALS INDUSTRY

REAL-TIME UPDATING OF SPATIAL MODELS USING ONLINE PRODUCTION DATA

DATA ASSIMILATION IN THE MINERALS INDUSTRY

REAL-TIME UPDATING OF SPATIAL MODELS USING ONLINE PRODUCTION DATA

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology,
by the authority of the Rector Magnificus, prof. dr. ir. T. H. J. J. van der Hagen,
chair of the Board for Doctorates,
to be defended publicly on
monday 19 March 2018 at 15:00 o'clock

by

Tom WAMBEKE

Master of Science in Applied Earth Sciences,
Delft University of Technology, Delft, the Netherlands,
born in Leuven, Belgium.

This dissertation has been approved by the promotors.

Composition of the doctoral committee:

Rector Magnificus,	chairperson
Prof. dr. ir. J.D. Jansen	Delft University of Technology, promotor
Prof. dr. Dipl.-Ing. J. Benndorf	Freiberg University of Mining and Technology, promotor

Independent members:

Prof. dr. ir. A. Vervoort	KU Leuven, University of Leuven, Belgium
Prof. dr. ir. M. Verlaan	Delft University of Technology
Prof. dr. M.A. Hicks	Delft University of Technology
Dr. J. Ortiz	Queens University, Canada

Other members:

Dr. M.W.N. Buxton	Delft University of Technology

This research was funded by Royal IHC.

Cover design by T. Wambeke

The front cover shows an aerial photograph of the Tropicana Gold Mine, Australia.
Photo credit: Damon Elder, Anglo Gold Ashanti.

Printed by Gildeprint, The Netherlands
ISBN 978-94-6186-904-3

An electronic version of this dissertation is available at

http://repository.tudelft.nl/.

To my mother Dominique Pierards,
forever loved, never forgotten

CONTENTS

Summary ix

Samenvatting xiii

Symbols and Notation xvii

1 Introduction 1
- 1.1 Closed-Loop Resource Management . . . 3
- 1.2 Research Objectives . . . 5
- 1.3 Thesis Outline . . . 8
- References . . . 10

2 Literature Review 11
- References . . . 16

3 Algorithm 21
- 3.1 Gaussian Anamorphosis . . . 22
- 3.2 Solving the Updating Equations . . . 24
- 3.3 Neighbourhood . . . 26
- 3.4 Covariance Error Correction . . . 26
- 3.5 Pseudocode . . . 28
- 3.6 Conclusions . . . 31
- References . . . 31

4 Synthetic Experiment: 2D Case Study 35
- 4.1 Methodoloy . . . 36
  - 4.1.1 Reference Field . . . 36
  - 4.1.2 Prior Set of Realizations . . . 36
  - 4.1.3 Experimental Scenario . . . 42
- 4.2 Results . . . 44
  - 4.2.1 Visual Inspection . . . 45
  - 4.2.2 Global Assessment Statistics . . . 49
  - 4.2.3 Local Assessment Statistics . . . 52
- 4.3 Conclusions . . . 58
- References . . . 60

5 Sensitivity Analysis 61
- 5.1 Methodology . . . 62
  - 5.1.1 General Setup . . . 62
  - 5.1.2 Geology . . . 64
  - 5.1.3 Experimental Scenarios . . . 66
  - 5.1.4 Forward simulator and sensor response . . . 67

5.2 Assessment Statistics . . . 73
5.2.1 Field statistics . . . 73
5.2.2 Production statistics . . . 74
5.3 Results - Single Parameter Variation. . . . 82
5.3.1 Measurement Volume . . . 82
5.3.2 blending ratio . . . 84
5.3.3 Measurement Error . . . 88
5.4 Results - Double Parameter Variation . . . 89
5.4.1 Measurement volume and measurement error. . . . 89
5.4.2 Measurement volume and blending ratio . . . 93
5.4.3 Blending ratio and measurement error. . . . 95
5.5 Discussion . . . 96
5.5.1 Objective 1 - maximize error reduction in GC model. . . . 97
5.5.2 Objective 2 - optimize reconciliation of production data. . . . 98
5.5.3 Objective 3 - maximize error reduction in future predictions. . . . 98
5.6 Practical implications . . . 100
5.7 Conclusions. . . . 100
References . . . 102

6 Pilot Study **103**
6.1 Introduction . . . 104
6.2 Updating Algorithm. . . . 107
6.3 Background Information and Data Sources . . . 109
6.3.1 Geology . . . 109
6.3.2 Mining. . . . 110
6.3.3 Comminution . . . 111
6.4 Forward Simulator . . . 111
6.4.1 From Pit to Crusher . . . 111
6.4.2 From Crusher to Mill. . . . 113
6.5 Results case study. . . . 116
6.5.1 GeoMet model . . . 116
6.5.2 Mill Feed Estimates . . . 116
6.6 Conclusions. . . . 117
References . . . 125

7 Conclusions and Recommendations **127**
7.1 Synopsis and Conclusion . . . 129
7.2 Recommendations for Future Work. . . . 133
7.2.1 The Updating Approach . . . 133
7.2.2 Additional Experiments . . . 135
7.2.3 Practical Application. . . . 136
References . . . 137

Acknowledgements **139**

Curriculum Vitæ **143**

List of Publications **145**

SUMMARY

Declining ore grades, extraction at greater depths and longer hauling distances put pressure on maturing mines. Not enough new mines will be commissioned on time to compensate for the resulting shortages. Ore-body replacement rates are relatively low due to a reduced appetite for exploration. Development times are generally increasing and most new projects are remote, possibly pushing costs further upwards.

To reverse these trends, the industry must collect, analyse and act on information to extract and process material more productively (i.e. maximize resource efficiency). This paradigm shift, driven by digital innovations, aims to (partly) eliminate the external variability that has made mining unique. The external variability results from the nature of the resource being mined. This type of variability can only be controlled if the resource base is sufficiently characterized and understood.

Recent developments in sensor technology enable the online characterization of raw material characteristics and equipment performance. To date, such measurements are mainly utilized in forward loops for downstream process control. A backward integration of sensor information into the resource model does not yet occur. Obviously, such a backward integration would significantly contribute to the progressive characterization of the resource base.

This dissertation presents a practical updating algorithm to continuously assimilate recently acquired data into an already existing resource model. The updating algorithm addresses the following practical considerations. (a) At each point in time, the latest solution implicitly accounts for all previously integrated data (sequential approach). During the next update, the already existing resource model is further adjusted to honour the newly obtained observations as well. (b) Due to the nature of a mining operation, it is nearly impossible to formulate closed-form analytical expressions describing the relationship between observations and resource blocks. Rather, the relevant relationships are merely inferred from the inputs (the resource model realizations) and outputs (distribution of predicted observations) of a forward simulator. (c) The updating algorithm is able to assimilate noisy observations made on a blend of material originating from multiple sources and locations. Differences in scale of support are dealt with automatically.

The developed algorithm integrates concepts from several existing (geo)statistical techniques. Co-Kriging approaches for example are designed to integrate both direct and indirect measurements and are well capable to handle differences in accuracy and sampling volume. However, they do fail to extract information from blended measurements and can not sequentially incorporate new observations into an already existing resource model. To overcome the latter issue, the co-Kriging equations are merged into a sequential linear estimator. Existing resource models can now be improved using a weighted sum of differences between observations and model-based predictions (forward simulator output). The covariances, necessary to compute the

weights, are empirically derived from two sets of Monte Carlo samples (another statistical technique); the resource model realizations (input forward simulator) and the observation realizations (output forward simulator). This approach removes the need to formulate analytical functions modelling spatial correlations, blending and difference in scale of support.

The resulting mathematical framework bears some resemblances to that of a dynamic filter (Ensemble Kalman filter), used in other research areas, althoughthe underlying philosophy differs significantly. Weather forecasting and reservoir modelling, for example, consider dynamic systems repetitively sampled at the same locations. Each observation characterizes a volume surrounding the sample locations. Mineral resource modelling, on the other hand, focuses on static systems gradually sampled at different locations. Each observation is characteristic for a blend of material originating from multiple sources and locations. Each part of the material stream is sampled only once, the moment it passes the sensor.

Various options are implemented around the mathematical framework to either reduce computation time, memory requirements or numerical inaccuracies. (a) A Gaussian anamorphosis is included to deal with suboptimal conditions related to non-Gaussian distributions. The algorithm structure ensures that the sensor precision (measurement error) can be defined on its original units and does not need to be translated into a normal score equivalent. (b) An interconnected parallel updating sequence (double helix) can be configured to avoid a covariance collapse (filter inbreeding). This occurs as degrees of freedom are lost over time due to the empirical calculation of the covariances. (c) A neighbourhood option is implemented to constrain computation time and memory requirements. Different neighborhoods need to be considered simultaneously as material streams are blended. (d) Two covariance correction options are implemented to further inhibit the propagation of statistical sampling errors originating from the empirical computation of covariances.

A case specific forward simulator is built and run parallel to the more generally applicable updating code. The forward simulator is used to translate resource model realizations (input) into observation realizations (output). Empirical covariances are subsequently lifted from both realization sets and mathematically describe the link between sensor observations and individual blocks in the model. This numerical inference avoids the cumbersome task of formulating, linearising and inverting an analytical forward observation model. The application of a forward simulator further ensures that the distribution of the Monte Carlo samples already reflect the support of the concerned random values. As a result, the necessary covariances, derived from these Monte Carlo samples, inherently account for differences in scale of support.

A synthetic experiment is conducted to showcase that the algorithm is capable of assimilating inaccurate observations, made on blended material streams, into an already existing resource model. The experiment is executed in an artificial environment, representing a mining environment with two extraction points of unequal production rate. A visual inspection of cross-sections shows that the model converges towards the ”true but unknown reality”. Global assessment statistics quantitatively confirm this observation. Local assessment statistics further indicate that the global improvements mainly result from correcting local estimation biases.

Another 125 artificial experiments are conducted to study the effects of variations in measurement volume, blending ratio and sensor precision. The experiments investigate whether and how the resource model and the predicted observations improve over time. Based on the outcome, recommendations are formulated to optimally design and operate a monitoring system.

This work further describes the pilot testing of the updating algorithm at the Tropicana Gold Mine (Australia). The pilot aims to evaluate whether the updating algorithm can automatically reconcile ball mill performance data against the spatial Work Index estimates of the GeoMet model. The focus here lies on the ball mill since it usually is the single largest energy consumer at the mine site. The spatial Work Index estimates are used to predict a ball mill's throughput. In order to maximize mill throughput and optimize energy utilization, it is important to get the Work Index estimates right. At the Tropicana Gold Mine, Work Index estimates, derived from X-Ray Fluorescence and Hyperspectral scanning of grade control samples, are used to construct spatial GeoMetallurgical models (GeoMet). Inaccuracies in the block estimates exist due to limited calibration between grade control derived and laboratory Work Index values. To improve the calibration, the updating algorithm was tested at the mine during a pilot study. Deviations between predicted and actual mill performance are monitored and used to locally improve the Work Index estimates in the GeoMet model. While assimilating about a week of mill performance data, the spatial GeoMet model converged towards a previously unknown reality. The updating algorithm improved the spatial Work Index estimates, resulting in a real-time reconciliation of already extracted blocks and a recalibration of future scheduled blocks. The case study shows that historic and future production estimates improve on average by about 72% and 26%.

SAMENVATTING

Bestaande mijnen kampen met dalende ertsgraden, extractie op grotere diepten and langere transport afstanden. Niet genoeg nieuwe mijnen zullen tijdig worden geopend om de daaruit voortvloeiende tekorten te kunnen compenseren. Er worden te weinig nieuwe ertslichamen ontdekt door een afname in de zin voor exploratie. De ontwikkelingstijden nemen in het algemeen toe en de meeste nieuwe projecten zijn afgelegen, hetgeen de kosten verder te hoogte in zal jagen.

Om deze trends om te keren moet de industrie informatie verzamelen, analyseren en ernaar handelen met als doel het materiaal efficiënter te ontgraven en te verwerken. Deze paradigmaverschuiving, gedreven door digitale innovaties, is gericht op het (gedeeltelijk) elimineren van de externe variabiliteit die zo kenmerkend is in de mijnbouw. De externe variabiliteit is inherent verbonden met de aard van de gemijnde natuurlijke hulpbron. Dit soort variabiliteit kan enkel onder controle worden gehouden wanneer het ertslichaam voldoende is gekarakteriseerd en begrepen.

Nieuwe ontwikkelingen in de sensortechnologie maken het nu mogelijk om grondstofeigenschappen en machineprestaties online te meten. Tot op heden worden dergelijke metingen voornamelijk gebruikt in downstream procescontrole. Het integreren van sensorinformatie in bestaande resource modellen komt nog niet voor. Het is duidelijk dat zo een integratie significant zal bijdragen tot het steeds beter karakteriseren van het ertslichaam.

Dit proefschrift presenteert een praktisch updating algoritme om recent verworven data continu te assimileren in een reeds bestaand resource model. Het updating algoritme biedt het hoofd aan volgende praktische overwegingen. (a) Op elk tijdstip dient de laatste oplossing impliciet rekening te houden met alle voordien geïntegreerde data (sequentiële aanpak). Tijdens de volgende update wordt het reeds bestaande model verder aangepast zodat het een verklaring biedt voor nieuw verkregen waarnemingen. (b) Vanwege de complexiteit van een mijnbouwoperatie is het nagenoeg onmogelijk om een analytische uitdrukking in gesloten vorm te formuleren die de relaties tussen blokken en waarnemingen beschrijft. In plaats daarvan worden de relevante relaties afgeleid op basis van slechts de input (resource model realisaties) de output (de verdeling van voorspelde waarnemingen) van een simulator. (c) Het updating algoritme is in staat om zeer onnauwkeurige waarnemingen te assimileren die bovendien gemaakt werden op gemengde materiaalstromen afkomstig van verschillende locaties. Er wordt automatisch rekening gehouden met verschillen in meetvolumes.

Het ontwikkelde algoritme combineert enkele ideëen van reeds bestaande (geo)statische technieken. Co-Kriging technieken zijn eerder ontworpen om zowel directe als indirecte metingen te integreren. Ze zijn bovendien goed in staat om om te gaan met verschillende meetvolumes en meetnauwkeurigheden. Deze technieken zijn echter ontoereikend als het gaat om informatie te verwerken, afkomstig van gemengde materiaalstromen. Het sequentieel aanpassen van resource modellen is ook niet mogelijk.

Dit laatste probleem kan opgelost worden door de co-Kriging vergelijkingen samen te voegen met een sequentiële lineaire schatter. Bestaande resource modellen kunnen dan worden verbeterd door gebruik te maken van een gewogen som van de verschillen tussen waarnemingen en voorspellingen (de output van de simulator). De covarianties, die nodig zijn om de waarden van de gewichten te berekenen, kunnen empirisch worden afgeleid op basis van de uitkomst van Monte Carlo steekproeven (nog een andere statistische techniek); de resource model realisaties (input simulator) en de realisaties van waarnemingen (output simulator). Deze aanpak verwijdert de noodzaak om analytische functies te formuleren die de ruimtelijke correlaties, het mengen van materiaal en de verschillen in meetvolumes beschrijven.

Het resulterende wiskundige kader vertoont een aantal gelijkenissen met dat van een dynamische filter (de Ensemble Kalman filter), die vaak gebruikt wordt in andere onderzoeksgebieden. Let wel, de onderliggende filosofie verschilt dusdanig. Het voorspellen van het weer en de exploitatie van reservoirs betreft dynamische systemen waar herhaaldelijk metingen worden uitgevoerd op dezelfde locaties. Elke meting is kenmerkend voor een bepaald volume dat het meetpunt omringt. In de mijnbouw echter gaat het om statische systemen die geleidelijk aan worden bemonsterd op steeds weer verschillende locaties. Elke waarneming is kenmerkend voor gemende materiaalstromen afkomstig van meerdere gebieden in de mijn. Elk deel van de materiaalstroom wordt slechts één enkele keer gemeten op het moment dat het onder de sensor door beweegt.

Verschillende opties werden tijdens de implementatie toegevoegd aan het wiskundig kader om de rekentijd, het benodigde geheugen en de numerieke onnauwkeurigheden te verminderen. (a) Een Gaussische anamorfose is opgenomen om te kunnen omgaan met de suboptimale omstandigheden die ontstaan door de aanwezigheid van niet-normale verdelingen. Het ontwerp van het algoritme is verder aangepast zodat de waarde van de meetfout gedefinieerd kan worden in zijn oorspronkelijke eenheden en niet hoeft vertaald te worden naar een equivalent in de getransformeerde ruimte. (b) Twee verbonden updating reeksen kunnen worden geconfigureerd om een “instorting” van de covariantie te vermijden. Dit fenomeen doet zich voor wanneer covarianties empirisch worden berekend en is een gevolg van een afname van het aantal vrijheidsgraden met de tijd. (c) De berekening kan beperkt worden tot enkele lokale gebieden, dit om de rekentijd en het benodigde geheugen te beperken. De gebieden vanwaar de materiaalstromen afkomstig zijn moeten tegelijkertijd worden behandeld in dezelfde berekening. (d) Verder bestaan er nog twee opties om de numeriek berekende covarianties te corrigeren en zo de verspreiding van numerieke onnauwkeurigheden tegen te gaan. Deze onnauwkeurigheden ontstaan tijdens de empirische berekening van de covarianties.

Een toepassings-specifieke simulator moet worden gebouwd en gedraaid parallel aan het meer algemeen toepasbare updating algoritme. Deze simulator wordt gebruikt om resource model realisaties (input) te vertalen naar voorspelde waarnemingen (output). covarianties worden vervolgens empirisch berekend aan de hand van de beide groepen realisaties. Deze covarianties beschrijven wiskundig de relatie tussen de sensor waarnemingen enerzijds en individuele blokken in het model anderzijds. Deze numerieke benadering vermijdt de noodzaak om analytische functies in gesloten vorm

neer te schrijven en ze vervolgens te lineariseren en inverteren. Het toepassen van een simulator zorgt er verder voor dat de verdeling van de Monte Carlo waarden kenmerkend zijn voor het volume van de desbetreffende component. Als gevolg hiervan houden de covarianties, die empirisch afgeleid worden van de Monte Carlo waarden, impliciet rekening houden met de verschillen in meetvolumes.

Een synthetisch experiment toont aan dat het algoritme in staat is om onnauwkeurige waarnemingen van gemengde materiaalstromen te assimileren in een reeds bestaand resource model. Het experiment is uitgvoerd in een artificiële omgeving die representatief is voor een mijnbouwoperatie met twee extractiepunten waarvan de productiehoeveelheden kunnen verschillen. De drawsdoorsneden geven visueel weer hoe het model convergeert de "werkelijke maar tot dan toe nog onbekende realiteit". De berekende globale statistieken bevestigen deze observatie kwantitatief. De lokale statistieken tonen verder aan dat de globale verbeteringen worden veroorzaakt door het verbeteren van lokaal schattingsfouten.

Bijkomend worden er 125 experimenten uitgevoerd om het effect van variaties in meetvolumes, mengverhouding en meetnauwkeurigheid te onderzoeken. De experimenten geven inzicht in hoe het resource model en de voorspelde waarnemingen verbeteren met de tijd. Op basis van deze resultaten worden aanbevelingen geformuleerd met als doel het meetnetwerk optimaal te ontwerpen en benutten.

Dit werk beschrijft verder de pilot test die werd uitgevoerd in de Tropicana goudmijn in Australië. Deze pilot had als doel om na te gaan in welke mate het updating algoritme in staat is om ruimtelijke hardheidsgetallen automatisch te verbeteren op basis van de prestatiewaarden van een kogelmolen. De nadruk ligt hier op de kogelmolen omdat dit in het algemeen de grootste energieverbruiker is in een mijn. De ruimtelijke hardheidsgetallen worden gebruikt om de doorstroom in de kogelmodel te voorspellen. Om deze doorstroom te maximaliseren alsook het energieverbruik te optimaliseren is het gelangrijk dat de hardheidsgetallen correct geschat zijn. De hardheidsgetallen, afgeleid op basis van scans van controlemonsters (röntgenfluorescentie en hyperspectraal), dienen als input voor het maken van ruimtelijke modellen (met geschatte hardheidsgetallen). De geschatte blokwaarden zijn onbetrouwbaar omwille van de beperkte kalibratie tussen hardheidsgetallen die enerzijds bepaald zijn op basis van de scans en anderzijds afkomstig zijn van proeven in een laboratorium. Er werd nagegaan of het updating algoritme in staat was om de geschatte blokwaarden te verbeteren. Verschillen tussen voorspelde en daadwerkelijke waarnemingen in de kogelmolen werden nauwlettend bijgehouden en gebruikt om de hardheidsgetallen lokaal aan te passen. De resultaten tonen aan dat het ruimtelijk model convergeert naar een nog eerder ongekende realiteit naarmate een week aan prestatiewaarden van de kogelmolen worden geassimileerd. Het updating algoritme pas niet enkel de alreeds gemijnde blokken aan maar verbetert ook de blokken waarvan de extractie ingepland staat. De studie toont aan dat zowel historische als toekomstige productievoorspellingen gemiddeld verbeteren met respectievelijk 72% en 26%.

SYMBOLS AND NOTATION

:	operator, includes all available rows or columns
$:\alpha$	the first α columns or rows
$\forall$	for all
$\in$	in, belonging to the set of
AE	absolute error field
$\mathscr{A}_t$	forward simulator, simulation step from $t-1$ to t
$\mathbf{A}_t$	N x 1 matrix - first order linear approximation of forward simulation step
$\mathbf{B}_t$	K x I matrix - predicted observations
$\mathbf{B}_t(:,i)$	i^{th} realization of all K observations collected at time t - vector (column)
$\mathbf{B}_t(k,:)$	predicted observation k - random vector (row)
$\mathbf{C}_{t,dd}$	K x K covariance matrix - covariance between individual observations
$\mathbf{C}_{t,qq}$	K x K covariance matrix - covariance between individual transformed perturbed predicted observations
$\mathbf{C}_{t,zz}$	N x N covariance matrix - covariance between individual grid nodes
$\mathbf{C}_{t,uy}$	N x K covariance matrix - covariance between individual grid nodes and transformed predicted observations
$\mathbf{d}_t$	vector with k real observations
$E[...]$	best estimate of ...
EM	estimated mean field
E	measurement error (chapter 5)
$\mathbf{E}_t$	K x I matrix - random noise
$\mathbf{E}_t(k,:)$	vector with random noise for observation k - I elements
F_{80}	80% passing size of the mill feed
F	cumulative probability
G	standard normal cumulative distribution function
i	refers to certain Monte Carlo realization
i'	rank of corresponding element in sorted Monte Carlo sample
I	total number of Monte Carlo realizations
$\mathbf{I}$	N x N identity matrix
k	refers to specific observation
K	total number of observations
L	the number of extraction points
n	refers to specific block
N	total number of blocks in the grid
N'	subset of blocks, the neighborhood set

$\Phi_{t,\mathbf{x}_n}$	local Gaussian anamorphosis function at block n
$\Phi_{t,k}$	Gaussian anamorphosis function to transform k^{th} observation
P	power draw
P_{80}	80% passing size of the mill product
$\mathbf{Q}_t$	K x I matrix - transformed perturbed predicted observations
$\mathbf{Q}_t(:,i)$	i^{th} realization of all transformed perturbed observations at time t - vector (column)
$\mathbf{Q}_t(k,:)$	transformed perturbed predicted observation k - random vector (row)
R	blending ratio (chapter 5)
R	thoughput (chapter 6)
$RMSE$	root mean square error
$\mathbf{R}$	K x K matrix - diagional matrix with sensor precision
ρ	correlation
$\hat{\rho}$	estimated correlation
SD	standard deviation field
$\mathbf{s}_t$	vector with k transformed real observation
σ	standard deviation
σ^2	variance
t	current time step
$t-1$	previous time step
$\mathbf{U}_t$	N x I matrix - transformed field realizations
$\mathbf{U}_t(:,i)$	i^{th} transformed field realization at time t - vector (column)
$\mathbf{U}_t(n,:)$	block n - transformed random vector (row)
V	measurement volume (chapter 5)
$\mathbf{W}_t$	N x K matrix with Kriging weights
$\mathbf{Wi}^s$	vector of length M, representing one metre long GC samples
$\mathbf{Wi}_t^b$	N x I matrix, GeoMet realizations
$\mathbf{Wi}_t^f$	K x I matrix, mill feed estimates
$\mathbf{wi}_t^f$	K x 1 matrix, real mill feed observations
$\mathbf{Y}_t$	K x I matrix - transformed predicted observations
$\mathbf{Y}_t(:,i)$	i^{th} realization of all transformed predicted observations at time t - vector (column)
$\mathbf{Y}_t(k,:)$	transformed predicted observation k - random vector (row)
$\mathbf{Z}^*$	N x 1 matrix - true but unknown field
$\mathbf{Z}_t$	N x I matrix - field realizations
$\mathbf{Z}_t(:,i)$	i^{th} field realization at time t - vector (column)
$\mathbf{Z}_t(n,:)$	block n - random vector (row)

1

INTRODUCTION

The mining industry must try to avoid its short sighted behaviour of ramping up production capacity during times of unsustainably high commodity prices. As a growth strategy, these tactics only go so far. Instead, the industry is urged to develop a more sustainable alternative to deliver long-term growth. The necessary paradigm shift will most likely be driven by digital technology innovations bearing promise of a higher resource efficiency. New technologies can be adopted to collect, analyse and act on information to optimally extract and process material. In order to make the best possible decisions, it is vital that the resource model is always entirely up to date. This requires a practical updating algorithm which continuously assimilates new information into an already existing resource model.

1

During the preceding mining boom (2001 to 2011), the sky seemed to be the limit. Rampant Chinese demand and increasing commodity prices (e.g. aluminum, coal, copper, gold, iron ore, lead, nickel, platinum, silver, tin, zinc) resulted in short sighted business decisions in favour of a fast buck. While believing China's demand would continue to grow, mining companies entered a fierce race to increase production capacity at any cost. Capital was plunged into new projects, mergers and high-priced acquisitions, all in preparation to take on the escalating demand (Price Waterhouse Coopers, 2016; Arndt *et al.*, 2017).

In 2011, the inevitable happened. The demand for bulk commodities started to decline, at a time of huge spendings to expand supply. Poor capital discipline exposed many companies to significant write-downs. During the subsequent years, growing liquidity concerns piled up and investors lost confidence. In 2015, after a pulverizing downturn, the mining industry grounded to a virtual halt. The Top 40 mining companies accumulated a first ever recorded aggregate net loss of about $28bn. Over the course of four years, their combined market capitalization dropped by about 60%, from $1,200bn. in 2011 to $494bn. in 2015 (Price Waterhouse Coopers, 2016).

In response, operational expenditures have been cut relentlessly in order to match the depressed price levels. Assets have been up for sale, marginal projects have been shelved and companies have been holding back on exploration. These restructuring efforts all aimed to improve debt-burdened balance sheets while avoiding any imminent bankruptcies. In 2016, a turning point was reached as industry fundamentals began to improve. Commodity prices rebounded, though not reaching pre-downturn levels, and the global demand for raw materials started to awake. Credit ratings have been restored, while investors appreciated the stronger balance sheets resulting in better valuations. The market capitalization of the top 40 mining companies went up to $748bn. (April 2017), an increase of 50% compared to the rock bottom situation of 2015, but still $452bn. short of previous highs (Price Waterhouse Coopers, 2017).

Cautious optimism prevails as price levels are bottoming out. With the lessons of the past in mind, the industry must try to avoid its old habits of investing in production capacity during a boom followed by a wave of write-offs during the bust that inevitably follows. Instead, the industry is urged to develop a long-term growth strategy. Failing to do so, mining companies will remain at the mercy of commodity speculators and shareholders will probably start reallocating their capital to other more stable and profitable industries.

Whilst the mining industry is getting back on its feet, other challenges are already appearing on the horizon. Declining ore grades, extraction at greater depths and longer hauling distances are putting pressure on maturing mines. Not enough new mines will be commissioned on time to compensate for the resulting shortages. Ore-body replacement rates are relatively low due to a reduced appetite for exploration. Development times are generally increasing and most new projects are remote, possibly pushing costs further upwards.

To address these upcoming challenges while providing a long-term strategy, the industry's focus should shift from "simply" extracting and processing material towards collecting, analysing and acting on information to extract and process material more productively (Durrant-Whyte *et al.*, 2015). This paradigm shift, driven by digital inno-

vations, aims to (partly) eliminate the internal and external variability that has made mining unique. Please note that the following definitions of internal and external variability deviate substantially from those commonly used in the field of Geostatistics.

Internal variability, created by miners themselves, results from operating outside of the plan. The adoption of sensor and communication technology enables a live recording of equipment state and location. In an operations control room, all live feeds are analysed simultaneously to ensure compliance to plan. When operational problems arise, limited deviations will be allowed. One central control room further ensures that all decisions regarding scheduling, dispatching, blending, and process control are aimed to improve the operation as a whole (instead of improving the performance indicators of local silos). A more mechanized operation (e.g. autonomous drilling, loading and hauling) can further reduce variability in execution (Simonite, 2016; Logan, 2016; Gershgorn, 2016).

External variability results from the nature of the resource being mined. This type of variability can only be controlled if the resource base is sufficiently characterized and understood. A first set of algorithms can be applied to turn the vast amounts of data collected by embedded sensors into new insights. The knowledge of the resource base should be refined continuously as sensor measurements (equipment state and mineralogical/chemical characterization) are integrated with already available geological information. A second set of algorithms is used to continuously re-evaluate the relevant operational decisions using the most up-to-date understanding of the resource. Optimal executable mine plans and process control parameters are generated avoiding variability and quality issues at the source.

At the time of writing, companies are already actively mitigating internal variability as part of their continuous efforts to reduce operational costs and improve productivity. A comprehensive management of external variability during operations is currently more of a research theme. An integrated practical application has yet to be observed.

The following section discusses how the "closed-loop resource management" framework, adopted from the oil and gas industry, can be applied to understand and control external variability.

1.1. CLOSED-LOOP RESOURCE MANAGEMENT

The "Closed-Loop Resource Management" framework aims to exploit mineral deposits more efficiently, by making the best possible operational decisions over and over again, given all available information (Fig. 1.1). The proposed framework is an adapted version of the "Closed-Loop Reservoir Management" framework originally developed for the oil and gas industry by Jansen *et al.* (2009).

The framework essentially consists of an actual mining operation (the system) and a virtual representation thereof (the system models). The so-called system models are constructed during the exploration campaign and subsequent development phases. A first set of system models describes the current understanding of the resource and is based on geological, geotechnical and metallurgical data collected during exploration drilling. Due to economic constraints, the number of exploration holes and corresponding laboratory analyses are rather limited. The sparse data impede a conclusive characterization of the elements in the three dimensional resource model. Despite

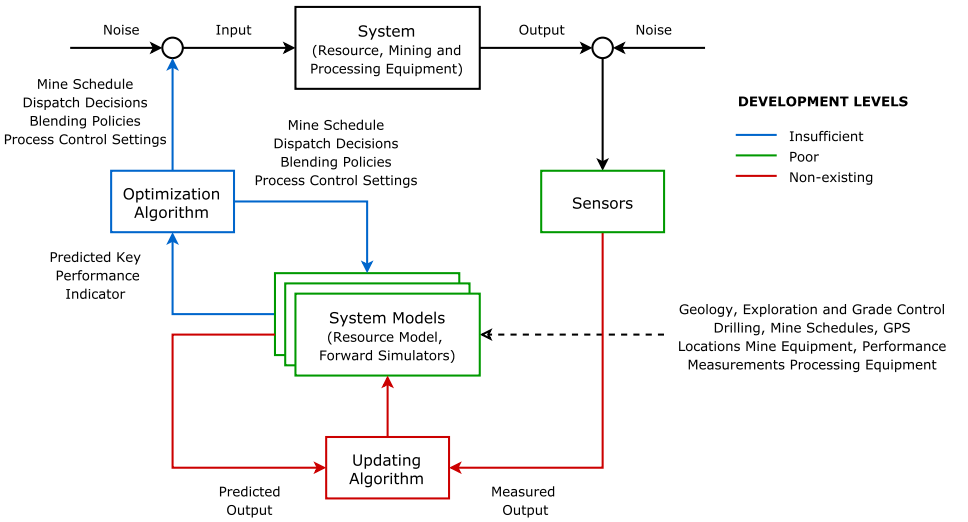

Figure 1.1: Closed-Loop Resource Management. Adapted from Jansen *et al.* (2009).

all exploration efforts, many of the elements remain rather poorly defined. Once the operation commences, data collected during grade control drilling are used to further refine the resource model and reduce its inherent uncertainty (grade control drilling is optional and not universally applied).

The resource model is one of two system models in the framework. The second model is a collection of forward simulators (Fig. 1.1). These simulators are but a virtual representation of the actual mining and processing operation. Simulators are used to study the effects of different development (long-term) and operational (short-term) scenarios on the overall performance of the mine. Short-term simulator behaviour can potentially be improved as additional operational data are taken into account (e.g. equipment state and performance measurements, GPS locations, executed schedules,...).

Once the system models are available, development and operational scenarios are translated into key performance indicators (Net Present Value, total metal or minerals produced, throughput, average head grade,...). A development scenario is mainly defined through its corresponding long-term mine schedule. An operational scenario concerns a combination of a short-term mine schedule, dispatch decisions, blending policies and process control setting. Algorithms are subsequently used to optimize either the development or operational scenario (blue optimization loop, Fig. 1.1).

The optimized scenario, i.e. the collection of related decisions, constitutes the input to control the actual “mining system”. The “mining system” comprises a resource, mining equipment and a processing facility (Fig. 1.1). The final system outputs are the same key performance indicators which were simulated before to select optimal controllable inputs.

The measured performance indicators commonly deviate from their target values

(resulting from optimized short- or long-term scenarios) due to internal and external variability. Assuming decisions and plans are executed correctly (mitigation of internal variability), most of the deviation directly results from our incomplete understanding of the resource base. All decisions taken thus far are merely optimal if the resource model exactly represents reality. The more the resource model deviates from reality, the worse are all congruent decisions and schedules. This is an important reason to keep the resource model up-to-date using all available sources of information.

Conventionally, the resource model is only reconciled quarterly against the measured performance indicators. The reconciliation techniques are fairly ad-hoc and limited to applying some overall correction factors. A balanced and careful integration of the intermediate system output does not generally occur. This intermediate system output is classified into two groups. The first group consists of equipment state and performance measurements. The second group consists of rapid non-destructive characterizations of material streams. The framework proposes to use algorithms to continuously extract hidden information from the resulting data streams and update the resource model where necessary (red updating or data assimilation loop, Fig. 1.1).

At the time of writing, the framework in its entirety has not been applied in the minerals industry. The colours in Fig. 1.1 refer to different development levels. Obviously, the "mining system" itself is considered to be given (at least with respect to this framework). Sufficient geostatistical and simulation-based techniques are available to construct reasonable models of the "mining system". However, the existing system models do not sufficiently account for the impact of geology and mineralogy on downstream processing steps. Further developments within the area of geometallurgy are necessary to establish such a link. Sensor technologies to monitor equipment state and performance are rather mature. The sensors for material characterization still require a lot of development and algorithms need to be designed to interpret the corresponding data streams. The optimization methodologies with respect to long-term scheduling are well established. Algorithms for the joint-optimization of the various short-term operational decisions can be significantly improved. Finally, practical algorithms, for real-time updating of resource models based on online sensor measurements, do not exist. The lack of a real-time updating algorithm is a major impediment to the application of the proposed "Closed-Loop Resource Management" framework.

1.2. RESEARCH OBJECTIVES

The aim of this research is to initiate the development of the previously outlined data assimilation loop (red loop, Fig. 1.1). During operation, this loop continuously improves local predictions in the resource model based on detected deviations between predicted and measured system output. The following four components need to be considered jointly:

- **Resource Model (System Model A)**: The resource model is a combination of different spatial stochastic models, all describing the spatial variability and geological uncertainty of their concerned attributes (e.g. grade, hardness, recovery, ...). Suitable simulation techniques are selected to generate the equally probable realizations of the in-situ orebody variability. All realization sets are conditioned on

scarce but accurate exploration data. Contrary to a set of single estimated models, the realization sets inherently describe uncertainty and correlations (empirical calculation of covariances, to be discussed later on).

- **Forward Simulator (System Model B)**: For each unique application, a case specific forward simulator has to be built. The forward simulator is a virtual representation of the actual mining and processing operation, describing which resource blocks are extracted, processed and measured within a specific time interval. Once built, the simulator computes intermediate and final system outputs based on the corresponding inserted resource realization. Eventually the entire set of resource realizations is propagated through the forward simulator and converted into distributions of predicted system outputs. These computed distributions are essential in determining how individual resource blocks are to be updated based on online sensor measurements.

- **Updating Algorithm**: In principle, the updating algorithm needs to solve an inverse problem. That is, the algorithm is tasked with inferring attributes of individual blocks based on time-averaged sensor observations. Initially, the mathematical problem is ill-posed, meaning that individual block attributes can be adjusted in a multitude of different ways to account for the sensor observations. Fortunately, the simulator output and the notion of in-situ correlation structures provide sufficient information to obtain a unique solution. To this end, an updating algorithm will be developed utilizing a Kalman filter-based approach to link forward propagated realizations (predicted system output) with real sensor observations to locally improve the resource model.

- **Online Sensor Measurements**: A variety of sensors embedded in the mine value chain continuously monitor intermediate system outputs such as equipment performance or material composition. The selected sensor data differs significantly from previously considered exploration data. (a) Whereas exploration data mainly result from a more accurate laboratory analysis, sensor data are significantly more noisy and recorded online in an operational environment. (b) An exploration data point can often be attributed to a single specific localized volume with a relatively high degree of accuracy (e.g. a core sample, one-to-one relationship). Sensor observations on the other hand potentially characterize blended material originating from multiple blocks, benches and even pits (a sensor observation characterizes the average property of multiple blended resource blocks, one-to-many relationship). Obviously, connecting a sensor observation to its constituent source material is more challenging. (c) The support of the exploration samples (i.e. the volume of the physical sample) generally remains constant and is rather small compared to the support of blocks in the spatial model. In contrast, a sensor observation characterizes blended material with a total volume that may well exceed the support of a single block. The support of a sensor observation depends on the selected interval to compute a time-averaged reading.

The research, presented in this thesis, is mainly focused on the development of an updating algorithm, capable of sequentially integrating sensor observations into an

already existing resource model. The resource model, the forward simulator and the online sensor observations are briefly discussed in the context of some case studies.

To structure and guide the development of the required updating algorithm, the following objectives were formulated:

1. Evaluate available geostatistical and data assimilation techniques and assess their applicability in the context of real-time updating. Select a method that meets the following requirements.

   (a) A sequential updating approach is preferred over a geostatistical resimulation technique. The former incorporates recently available sensor observations into the resource model without explicitly requiring the full data collection history. Before every update, the resource model should already implicitly account for all previously integrated data (both exploration and production). Each update should then be designed to adjust the resource model to take the newly obtained observations into account as well. Instead, a resimulation technique remodels the resource base using all data ever recorded. This approach would stretch the memory requirements as the amount of available data keeps growing.

   (b) Due to the nature of the mining operation, it is nearly impossible to formulate a closed-form analytical expression describing the link between the resource model and the observations. The absence of a closed-form expression excludes all techniques requiring the inversion of an (approximate) forward prediction model. Instead, the selected mathematical technique has to infer such a relationship based on merely the input (the resource model) and output (predicted observations) of the forward simulator. Such an approach would introduce great flexibility regarding the use of case specific simulators.

   (c) The selected mathematical framework must be able to handle the unique characteristics of the sensor observations. (i) A sensor observation is no longer attributable to a unique spatial volume but instead represents a blend of material originating from multiple sources and locations. (ii) Collected observations are less accurate due to imprecise sensors. (iii) Differences in scale of support will have to be dealt with, preferably without resorting to complex co-regionalization models. The support of the observations themselves might vary across time. Furthermore, the support of each observation might differ significantly with the support of the blocks in the resource model.

2. Based on the selected mathematical framework, design a computationally efficient algorithm specifically tailored to the requirements of the mining industry. That is, the algorithm should additionally cope with indirect observations, non-Gaussian distributions and large grids.

3. Once the algorithm is designed and implemented, conduct an extensive test in a synthetic environment to evaluate its performance. Monitor several assessment statistics while updating to ensure that the algorithm is operating as designed.

4. Study the relation between the monitoring system control parameters and algorithm performance. Investigate how differences in measurement volume, blending ratios and sensor precision affect the overall updating behaviour.

5. Apply the updating algorithm during a pilot study using data from an actual mining operation. Construct the necessary forward simulator and run the updating algorithm to assimilate online production data into the resource model. Evaluate how historic and future production estimates improve.

1.3. THESIS OUTLINE

The outline of this dissertation reflects the previously formulated research objectives.

Chapter 1 illustrates the relevance of digital innovation in proving a long-term growth strategy. The chapter presents the concept of a 'Closed-Loop Resource Management" framework and explains why a lacking updating algorithm is a major impediment to its application. The chapter further presents some research objectives to guide the development of the necessary algorithm. The chapter finally concludes with a thesis outline.

Chapter 2 provides a brief theoretical background on topics related to the updating or conditioning of spatial models based on additional production data. Several geostatistical techniques are reviewed, individually capable of solving at least one but not all of the practical problems identified (ref. research objective 1). The chapter further proves why the equations of a dynamic filter (the Ensemble Kalman Filter) would be suitable to solve the problem at hand. The chapter frames the underlying equations of the dynamic filter into a new static context and discusses the link with existing geostatistical concepts. The chapter further discusses how an application in the minerals industry differs significantly from those in various other research areas, particularly the fields of numerical weather forecasting, oceanography, physical geography and more recently in reservoir engineering.

Chapter 3 elaborates on the design of the updating algorithm, while presenting an overall mathematical formulation. Several additional features and options are highlighted. One of the first topics discussed is the rationale behind the Gaussian anamorphosis option. This option is implemented to deal with non-Gaussian distributions and indirect observations. The specific implementation, allowing for an intuitive treatment of measurement error, is explained. Subsequently, a computationally efficient strategy for solving the updating equations is provided. The presentation is generic, in that it includes an option to configure a parallel updating sequence (helix) reducing the effects of inbreeding (covariance collapse). The third subsection discusses the neighbourhood option which can be activated to further reduce computation times and memory requirements. Thereafter, two covariance correction techniques are introduced to contain the propagation of statistical sampling errors originating from their empirical computation. The fifth and last subsection presents the pseudocode which illustrates how the individual functional components are integrated.

The synthetic experiment, presented in **Chapter 4**, demonstrates that the algorithm is capable of improving the resource model based on inaccurate observations made on blended material streams. In order to conduct the experiment, an artificial environment is created, representing a mining operation with two extraction points of unequal production rate. The performance of the algorithm during the experiment is evaluated using three sets of criteria. (i) Horizontal and vertical cross-sections are visually inspected and compared with the "true but unknown" reality. (ii) Global assessment statistics are computed to study how the overall quality of the resource model evolves over time as production data are being assimilated. (iii) Local assessment statistics are computed to ensure that the observed global improvements also result in the correction of the local estimation biases.

Chapter 5 studies how the design and operational control of the monitoring system influences the overall performance of the updating algorithm. A total of 125 experiments are conducted to quantify the effects of variations in measurement volume, blending ratio and sensor precision. The following questions are addressed to compare behaviour across experiments. (i) Does the resource model improve over time? (ii) Does the accuracy of the predicted system output improve over time? (iii) Does the predicted system output in the future $24h$ improve after updating? Based on the outcome, recommendations are formulated to optimally design and operate the monitoring system, guaranteeing the best possible algorithm performance.

Chapter 6 describes the pilot testing of the entire updating concept (including the updating algorithm) at the Tropicana Gold Mine in Australia. The aim of the study is to evaluate whether the developed updating algorithm can be used to update spatial Work Index estimates based on actual ball mill performance data. The chapter starts with a detailed explanation of the practical problem. Thereafter, the updating equations are briefly reviewed and reformulated taking account of the problem specific terminology. The chapter further provides some background information regarding the geology at Tropicana, the operation and the available data. Subsequently, some insights are given regarding the construction of the application specific forward simulator. Then, results are presented illustrating improvements in both historic and future production estimates. The chapter finally concludes with an extensive discussion on modelling assumptions and potential improvements.

Chapter 7 is the last chapter of the dissertation and presents a brief overview of the general conclusions from this research. Thereafter, the chapter lists future recommendations and potential research areas to further improve and develop the "updating loop" within the overall "Closed-Loop Resource Management" framework. Note that specific conclusion and recommendations are provided at the end of most chapters.

1

REFERENCES

Price Waterhouse Coopers, *Mine 2016 - Slower, Lower, Weaker ... but not Defeated*, Tech. Rep. (Price Waterhouse Coopers, 2016).

N. Arndt, L. Fontbote, J. Hedenquist, S. Kesler, J. Thompson, and D. Wood, *Future global mineral resources,* Geochemical Perspectives **6**, 1–171 (2017).

Price Waterhouse Coopers, *Mine 2017 - Stop. Think ... Act.*, Tech. Rep. (Price Waterhouse Coopers, 2017).

H. Durrant-Whyte, R. Geraghty, F. Pujol, and R. Sellshop, *How digital innovation can improve mining productivity,* McKinsey&Company Metals and Mining (2015).

T. Simonite, *Mining 24 hours a day with robots,* MIT Technology Review (2016).

S. Logan, *Rio tinto: rolling out the world's first fully driverless mines,* Mining-Technology (2016).

D. Gershgorn, *Self-driving, 416-ton trucks are hailing raw materials around australia,* Quartz (2016).

J. Jansen, S. Douma, D. Brouwer, P. van den Hof, O. Bosgra, and A. Heemink, *Closed-loop reservoir management,* in *Proceedings SPE reservoir simulation symposium, Woodlands* (Society of petroleum engineers, 2009).

2

LITERATURE REVIEW

This chapter provides a brief theoretical background on topics related to the updating or conditioning of spatial models based on additional production data. Several geostatistical techniques are reviewed, individually capable of solving at least one but not all of the practical problems identified (ref. research objective 1). The chapter further proves why the equations of a dynamic filter (the Ensemble Kalman Filter) would be suitable to solve the problem at hand. The chapter frames the underlying equations of the dynamic filter into a new static context and discusses the link with existing geostatistical concepts. The chapter further discusses how an application in the minerals industry differs significantly from those in various other research areas, particularly the fields of numerical weather forecasting, oceanography, physical geography and more recently in reservoir engineering.

Parts of this chapter have been published in Mathematical Geosciences **49**, 1 (Wambeke and Benndorf, 2017) and in Proceedings of the 17th annual conference of the international association for mathematical geosciences (Wambeke and Benndorf, 2015).

Traditionally the mining industry has had mixed successes in achieving the production targets it has set out. Several projects have been identified where mineral grades are not as expected, schedules and plans are not met and recovery is lower than forecasted (Ward and McCarthy, 1999; Vallee, 2000; Tatman, 2001; McCarthy, 2003). The deviations of produced tonnages and grades from model-based expectations result from a mismatch between the scale of the exploration data and the short-term production targets (Benndorf, 2013). In other words, it is challenging to accurately define the characteristics of e.g. a few truck loads, designated to be transported to the processing plant, based on data gathered at relatively wide grids. For certain commodities, it is common to perform grade control (GC) drilling to further reduce the uncertainty (Peattie and Dimitrakopoulos, 2013; Dimitrakopoulos and Godoy, 2014). However, GC drilling is expensive and almost exclusively focused on sampling grades. Metallurgical properties are often ignored.

The mineral industry is increasingly looking for effective methods for monitoring and reconciling estimates and actual observations at different stages of the resource extraction process (Morley, 2014). Recent developments in sensor technology enable the on-line characterization of production performance and raw material characteristics. To date, sensor measurements are mainly utilized in forward loops for downstream process control and material handling (Zimmer, 2012; Lessard *et al.*, 2014; Nienhaus *et al.*, 2014). A backward integration of sensor information into the GC model to continuously improve the production forecasts and dispatch decisions does not yet occur.

The application of sensors carries a large potential regarding process improvements. Sensor responses could be used to progressively increase the knowledge about the in-situ material characteristics. This has two main consequences. First, the frequency of misallocation could decrease, i.e. a smaller amount of actual ore is incorrectly allocated to the waste dump and a smaller amount of actual waste enters the processing plant. Second, an improved characterization of metallurgical properties could lead to a better selection of process parameters. For example, the throughput of the comminution circuit can be reduced upfront when harder ore is expected to ensure that the resulting grain sizes stay within acceptable limits. A proactive selection of process parameters in combination with the elimination of low value material from the processing plant will further result in a reduction of dilution, an increase of concentrator recovery and a larger annual metal production.

The potential of real-time updating is obvious. In order to apply it in practice, algorithms need to be developed capable of assimilating direct and indirect measurements into the GC model. Thus, at any point in time when new observations become available, the following inverse problem needs to be solved (Tarantola, 2005; Oliver *et al.*, 2008):

$$\mathbf{z} = \mathscr{A}^{-1}(\mathbf{d}), \tag{2.1}$$

where $\mathscr{A}$ is a forward observation model (linear or non-linear) that maps the spatial attributes $\mathbf{z}$ of the GC model onto sensor observations $\mathbf{d}$. The observations result from either direct or indirect measurements. The following challenges are identified; (i) the latest solution should account for previously integrated data (sequential approach); (ii) due to the nature of a mining operation, it is nearly impossible to formulate an analytical approximation of the forward observation model, let alone compute its inverse and (iii)

observations are made on blended material streams originating from multiple extraction points. The objective of this chapter is to present a new algorithm to assimilate sensor observations into the grade control model, specifically tailored to the requirements of the mining industry.

Kitanidis and Vomvoris (1983) introduced a geostatistical approach to the inverse problem in groundwater modelling. Both scarce direct (local log conductivity) and more abundant indirect (hydraulic head and arrival time) measurements are used to estimate the hydraulic conductivity in geological media through a linear estimation procedure known in the geostatistical literature as co-Kriging (Deutsch and Journel, 1998; Journel and Huijbregts, 2003). In the Bayesian literature, the same procedure is referred to as updating or conditioning (Schweppe, 1973; Wilson *et al.*, 1978; Dagan, 1985). The geostatistical approach received considerable attention (Hoeksema and Kitanidis, 1984; Rubin and Dagan, 1987; Yates and A.W. Warrick, 1987; Sun and Yeh, 1992; Harter and Yeh, 1996; Tong, 1996).

Later, several simultaneous and independent developments resulted in a method to recursively incorporate subsets of data one at a time (Evensen, 1992; Harvey and Gorelick, 1995; Yeh and Zhang, 1996). The proposed sequential estimator improves previous subsurface models by using linearly weighted sums of differences between observations and model-based predictions

$$\mathbf{z}_t = \mathbf{z}_{t-1} + \mathbf{W}_t(\mathbf{d}_t - \mathscr{A}_t(\mathbf{z}_{t-1})), \tag{2.2}$$

where the vector $\mathbf{z}_t$ contains estimates of the spatial attributes after t updates; the vectors $\mathbf{d}_t$ and $\mathscr{A}_t(\mathbf{z}_{t-1})$ respectively hold actual observations and model-based predictions at time t and $\mathbf{W}_t$ is a matrix with Kriging weights defining the contribution of the detected deviations to the updated subsurface model. If the vector $\mathbf{d}_t$ only contains direct local measurements, then the linear estimator corresponds to simple Kriging, that is Kriging with a known mean. On the other hand, if also indirect measurements are included, then a single update results from solving a system of co-Kriging equations (Goovaerts, 1997; Chiles and Delfiner, 2012). The sequential linear estimator bears a remarkable resemblance to Kalman filter techniques (Evensen, 1992; Kalman, 1960; Bertino *et al.*, 2002). The Kriging weights are computed from the forecast and observation error covariance matrices, $\mathbf{C}_{t-1,zd}$ and $\mathbf{C}_{t-1,dd}$:

$$\mathbf{W}_t = \mathbf{C}_{t-1,zd}\mathbf{C}_{t-1,dd}^{-1} \tag{2.3a}$$

$$= \mathbf{C}_{t-1,zz}\mathbf{A}_t^T(\mathbf{A}_t\mathbf{C}_{t-1,zz}^{-1}\mathbf{A}_t^T + \mathbf{R})^{-1}, \tag{2.3b}$$

where $\mathbf{C}_{t-1,zz}$ is the prior error covariance matrix of the attribute field; $\mathbf{R}$ is a diagonal matrix which specifies the sensor precision (a large sensor precision corresponds to a low value on the diagonal) and $\mathbf{A}_t$ is a first order approximation of the non-linear observation model $\mathscr{A}_t$ (Evensen, 1992; Yeh and Zhang, 1996). If both the prior error covariance matrix and the sensor precision tend to be large then the Kriging weights tend to increase, indicating that a significant portion of the detected deviations are taken into account to update the attribute field. For completeness, the posterior error covariance matrix of the attribute field after one assimilation step is also given

$$\mathbf{C}_{t,zz} = (\mathbf{I} - \mathbf{W}_t\mathbf{A}_t)\mathbf{C}_{t-1,zz}. \tag{2.4}$$

Vargas-Guzmán and Yeh (1999) provided the theoretical evidence that under a linear observation model sequential Kriging and co-Kriging are equivalent to their traditional counterpart which includes all data simultaneously. In the case of a nonlinear observation model, a sequential incorporation of data increases the accuracy of the first-order (linear) approximations $\mathbf{A}_t$ since they are calculated around a progressively improving attribute field (Evensen, 1992; Harvey and Gorelick, 1995). The linear estimator thus propagates the conditional mean and covariances from one update cycle to the next.

At time zero, a global covariance model suffices to describe the degree and scale of variability in the attribute field. The covariance matrix $\mathbf{C}_{(0,zz)}$ is stationary. At any other time, the updated covariances reflect the assimilation history and indirectly depend on the location of the material sources (Harvey and Gorelick, 1995). This nonstationarity of the updated covariances results in perhaps the greatest limitation of the method, i.e. the necessity of storing large covariance matrices. Despite the promising results, the above mentioned techniques were yet not considered for resource modelling and reconciliation.

Thus far the discussion has focused on the sequential updating of a single best estimate. In geostatistics, it is common to simulate a set of realizations to assess uncertainty (Dowd, 1994; Dimitrakopoulos, 1998; Rendu, 2002). The propagated conditional mean and covariance provide an intuitive description of statistics required to perform geostatistical simulations. For example, a single realization can be generated through the combination of the propagated mean with the product of a decomposed covariance matrix (LU decomposition) and a vector filled with white noise (Davis, 1987; Alabert, 1987). Gomez-Hernandez and Cassiraga (2000), Hansen *et al.* (2006) and Hansen and Mosegaard (2008) opt for a different approach. They propose to use the entire collection of measurements simultaneously in combination with a co-Kriging based version of sequential Gaussian simulation. As time progresses, such an approach results in significant memory usage due to the substantial growth of available production data. The simulation approaches discussed thus far all require that the simulation algorithm is completely rerun after each timestep.

Vargas-Guzmán and Dimitrakopoulos (2002) presented an approach that can facilitate fast-updating of generated realizations based on new data, without repeating the full simulation process. The approach is termed conditional simulation of successive residuals and was designed to overcome the size-limitations of LU decomposition. The lower triangular matrix is obtained through a novel column partitioning, expressed in terms of successive conditional covariance matrices. The partitioning requires the specification of a sequence of (future) data locations. A stored L matrix can then facilitate the conditional updating of existing realizations if and only if the sequence of visited subsets and used production data are the same as the one used for generating the initial realizations (Jewbali and Dimitrakopoulos, 2011; Dimitrakopoulos and Jewbali, 2013).

A major limitation of the techniques discussed thus far results from the necessity to store and propagate the conditional nonstationary covariances. A significant portion of memory needs to be allocated to hold a collection of elements, the size of the square of the number of grid nodes. The very large grids commonly encountered in the mining industry make such approaches infeasible. To circumvent these limitations, the previously discussed sequential (linear) estimator (Eq. 2.2) can be integrated into a

Monte Carlo framework (Evensen, 1994). At each time step, a finite set of realizations is updated on the basis of sensor observations collected at time t:

$$\mathbf{Z}_t(:,i) = \mathbf{Z}_{t-1}(:,i) + \mathbf{W}_t(\mathbf{d}_t - \mathscr{A}_t(\mathbf{Z}_{t-1}(:,i)) + \mathbf{E}_t(:,i)) \quad \forall\, i \in I \tag{2.5}$$

where the observation errors $\mathbf{E}_t(:,i)$ are randomly drawn from a normal distribution with a zero mean vector and a diagonal covariance matrix $\mathbf{R}$. The conditional forecast and observation error covariances, $\mathbf{C}_{t-1,zd}$ and $\mathbf{C}_{t-1,dd}$, are computed empirically from both sets of model-based predictions $\mathscr{A}_t(\mathbf{Z}_{t-1}(:,i))$ and field realizations $\mathbf{Z}_{t-1}(:,i)$. Due to the applied Monte Carlo concept, the first-order approximation of the forward observation model can be avoided (Evensen, 1997; Burgers *et al.*, 1998). The initial set of realizations can be generated using techniques of conditional simulation. Aside from the empirical calculation of the covariances, the Monte Carlo based sequential conditioning approach bears some resemblance to the equations for conditioning simulations as presented in Journel and Huijbregts (2003). Applications of the sequential conditioning approach in a geoscientific context can be found in many documented studies (Bertino *et al.*, 2002; Jansen *et al.*, 2009; Aanonsen *et al.*, 2009; Hou *et al.*, 2015; Benndorf and Jansen, 2017; Heidari *et al.*, 2011; Hendricks Franssen *et al.*, 2011; Hu *et al.*, 2013). These other research disciplines refer to this technique as the Ensemble Kalman-Filter.

Before proceeding, it is important to comprehend the subtle differences to previous applications of the (Ensemble) Kalman-Filter. Weather forecasting and reservoir modelling (oil, gas and water) consider dynamic systems repetitively sampled at the same locations. Generally, each observation characterizes a volume surrounding a sample location. These local volumes are sampled repetitively in time. Mineral resource modelling on the other hand focuses on static systems gradually sampled at different locations. Each observation is characteristic for a blend of material originating from one or more extraction points. Each part of the material stream is sampled only once, at the moment it passes the sensor in the production chain. Although virtually the same, the terminology of sequential updating is preferred over (Ensemble) Kalman-Filtering. This is to stress the absence of dynamic components and to highlight its link to the field of geostatistics.

Benndorf (2015) was the first to recognize the potential of sequential updating in a context of mineral resource extraction. His preliminary work provides a mathematical description linked to the unique configurations of the mining industry (as discussed above). However the practicality of the description is limited due to the need for an explicit formulation of the $\mathbf{A}_t$ matrix. The presented description is based on 2.3b, which requires an empirical computation of the full $\mathbf{C}_{t-1,zz}$ matrix together with its inverse. Hence the description results in high computational costs and memory requirements. Moreover the diagonal R matrix characterizing the sensor precision needs to be formulated in the computation domain, which is straightforward when all distributions involved are Gaussian, but is much less intuitive when transformations are applied.

This contribution presents an algorithm based on Eq. 2.3a. The forward observation model $\mathscr{A}$ is excluded from the computer code (Fig 2.1). Instead, a forward simulator is used to convert field realizations into model-based predictions ($\mathscr{A}$ is applied to fields individually rather than to the covariance matrix). The use of a forward simulator overcomes the challenges of formulating an analytical approximation linking each

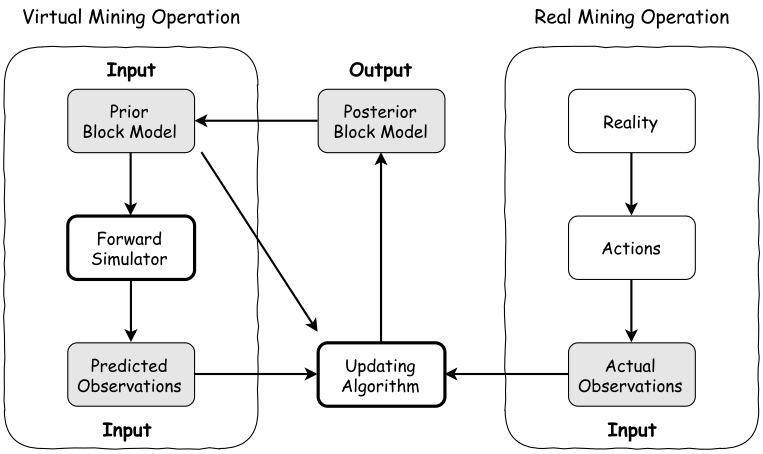

Figure 2.1: Overview of the interaction between the updating algorithm and the forward simulator

sensor observation to the grade control model. The necessary forecast and observation error covariances, $\mathbf{C}_{t-1,dz}$ and $\mathbf{C}_{t-1,dd}$, can now be computed empirically. Differences in scale of support are inherently accounted for.

The following concepts are integrated into the algorithm and will be discussed in detail in Chapter 3:

- A Gaussian anamorphosis option is implemented to deal with suboptimal conditions related to indirect observations and non-Gaussian distributions.
- A specific algorithm structure ensures that the measurement error can be defined on its original units and does not need to be translated into a normal score equivalent.
- An interconnected parallel updating sequence (double helix) can be configured to reduce the effects of filter inbreeding (covariance collapse).
- A neighbourhood option can be activated to further reduce computation times and memory requirements.
- Two covariance correction strategies are implemented to contain the propagation of statistical sampling errors originating from the empirical computation of covariances.

REFERENCES

T. Wambeke and J. Benndorf, *A simulation-based geostatistical approach to real-time reconciliation of the grade control model,* Mathematical Geosciences **49**, 1–37 (2017).

T. Wambeke and J. Benndorf, *Data assimilation of sensor measurements to improve production forecasts in resource extraction,* in *Proceedings of the 17th annual*

conference of the international association for mathematical geosciences (International association for mathematical geosciences, 2015).

D. Ward and P. McCarthy, *Start-up performance of new base metal projects,* in *Adding value to the Carpentaria mineral province, Mt Isa, Qld* (Australian Journal of Mining, 1999).

M. Vallee, *Mineral resources + engineering, economic and legal feasibility = ore reserve,* Bulletin of the Canadian Society of Mining Engineers **93**, 53–61 (2000).

C. Tatman, *Production rate selection for steeply dipping tabular deposits,* Mining Engineering **53**, 34–36 (2001).

P. McCarthy, *Managing technical risk for mine feasibility studies,* in *Mining Risk Management conference, Sydney* (The Australian Institute for mining and metallurgy, 2003).

J. Benndorf, *Application of efficient methods of conditional simulation for optimising coal blending strategies in large continuous open pit mining operations,* International Journal of Coal Geology **112**, 141 – 153 (2013).

R. Peattie and R. Dimitrakopoulos, *Forecasting recoverable ore reserves and their uncertainty at morila gold deposit, mali: An efficient simulation approach and future grade control drilling,* Mathematical Geosciences **45**, 1005–1020 (2013).

R. Dimitrakopoulos and M. Godoy, *Grade control based on economic ore/waste classification functions and stochastic simulations: examples, comparisons and applications,* Mining Technology **123**, 90–106 (2014).

C. Morley, *Monitoring and exploiting the reserve,* in *Mineral resource and reserve estimation - the AusIMM guide to good practice, monograph 30, 2nd edition* (The Australasian Institute of Mining and Metallurgy, 2014) pp. 647–657.

B. Zimmer, *Continuous mining equipment vs. complex geology challenges in mine planning,* in *11th Inter Symposium of Continuous surface mining, Miskolc* (The university of Miskolc, 2012).

J. Lessard, J. de Bakker, and L. McHugh, *Development of ore sorting and its impact on mineral processing economics,* Minerals Engineering **65**, 88 – 97 (2014).

K. Nienhaus, T. Pretz, and H. Wortruba, *Sensor technologies: impulses for the raw materials industry* (RWTH Aachen, Aachen, 2014).

A. Tarantola, *Inverse problem theory and methods for model parameter estimation* (SIAM, Philadelphia, 2005).

D. Oliver, A. Reynolds, and N. Liu, *Inverse theory for petroleum reservoir characterization and history matching* (Cambridge University Press, Cambridge, 2008).

P. Kitanidis and E. Vomvoris, *A geostatistical approach to the inverse problem in groundwater modeling (steady state) and one-dimensional simulations,* Water Resources Research **19**, 677–690 (1983).

C. Deutsch and A. Journel, *GSLIB: Geostatistical Software Library and User's Guide,* Applied geostatistics series (Oxford University Press, Oxford, 1998).

A. Journel and C. Huijbregts, *Mining Geostatistics* (Blackburn Press, Caldwell, New Jersey, 2003).

F. Schweppe, *Uncertain dynamic systems: modelling, estimation, hypothesis testing, identification and control* (Prentice-Hall, Englewood Cliffs, 1973).

J. Wilson, P. Kitanidis, and M. Dettinger, *State and parameter estimation in groundwater models,* in *Applications of Kalman filter to hydrology, hydraulics and water resources* (University of Pittsburg, 1978).

G. Dagan, *Stochastic modeling of groundwater flow by unconditional and conditional probabilities: The inverse problem,* Water Resources Research **21**, 65–72 (1985).

R. Hoeksema and P. Kitanidis, *An application of the geostatistical approach to the inverse problem in two-dimensional groundwater modeling,* Water Resources Research **20**, 1003–1020 (1984).

Y. Rubin and G. Dagan, *Stochastic identification of transmissivity and effective recharge in steady groundwater flow: 1. theory,* Water Resources Research **23**, 1185–1192 (1987).

S. Yates and A. A.W. Warrick, *Estimating soil water content using cokriging,* Soil Science Society America Journal **51**, 23–30 (1987).

N.-Z. Sun and W.-G. Yeh, *A stochastic inverse solution for transient groundwater flow: Parameter identification and reliability analysis,* Water Resources Research **28**, 3269–3280 (1992).

T. Harter and T.-C. Yeh, *Conditional stochastic analysis of solute transport in heterogeneous, variably saturated soils,* Water Resources Research **32**, 1597–1609 (1996).

A. Tong, *Stochastic parameter estimation, reliability analysis, and experimental design in ground water modeling,* phdthesis, University of California, Los Angeles (1996).

G. Evensen, *Using the extended kalman filter with a multilayer quasi-geostrophic ocean model,* Journal of Geophysical Research: Oceans **97**, 17905–17924 (1992).

C. Harvey and S. Gorelick, *Mapping hydraulic conductivity: Sequential conditioning with measurements of solute arrival time, hydraulic head, and local conductivity,* Water Resources Research **31**, 1615–1626 (1995).

T.-C. Yeh and J. Zhang, *A geostatistical inverse method for variably saturated flow in the vadose zone,* Water Resources Research **32**, 2757–2766 (1996).

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

J.-P. Chiles and P. Delfiner, *Geostatistics modeling spatial uncertainty* (John Wiley and Sons, Hoboken, New Jersey, 2012).

R. Kalman, *A new approach to linear filtering and prediction problems,* Transactions of the ASME, Journal of Basic Engineering **82**, 35–45 (1960).

L. Bertino, G. Evensen, and H. Wackernagel, *Combining geostatistics and kalman filtering for data assimilation in an estuarine system,* Inverse Problems **18**, 1 (2002).

J. A. Vargas-Guzmán and T.-C. Yeh, *Sequential kriging and cokriging: Two powerful geostatistical approaches,* Stochastic Environmental Research and Risk Assessment **13**, 416–435 (1999).

P. Dowd, *Risk assessment in reserve estimation and open-pit planning,* Transactions of the institution of mining and metallurgy, section A: mining technology **103**, 148–154 (1994).

R. Dimitrakopoulos, *Conditional simulation algorithms for modelling orebody uncertainty in open pit optimisation,* International Journal of Surface Mining, Reclamation and Environment **12**, 173–179 (1998).

J. Rendu, *Geostatistical simulations for risk assessment and decision making: The mining industry perspective,* International Journal of Surface Mining, Reclamation and Environment **16**, 122–133 (2002).

M. Davis, *Production of conditional simulations via the lu triangular decomposition of the covariance matrix,* Mathematical Geology **19**, 91–98 (1987).

F. Alabert, *The practice of fast conditional simulations through the lu decomposition of the covariance matrix,* Mathematical Geology **19**, 369–386 (1987).

J. Gomez-Hernandez and E. Cassiraga, *Sequential conditional simulations with linear constraints,* in *Geostatistics 2000, Cape Town, Monestiez P., Allard., D. and Froideveaux, R. (eds)* (Geostatistical association of Southern Africa, 2000).

T. Hansen, A. Journel, A. Tarantola, and K. Mosegaard, *Linear inverse gaussian theory and geostatistics,* Geophysics **71**, R101–R111 (2006).

T. Hansen and K. Mosegaard, *Visim: Sequential simulation for linear inverse problems,* Computers & Geosciences **34**, 53–76 (2008).

J. A. Vargas-Guzmán and R. Dimitrakopoulos, *Conditional simulation of random fields by successive residuals,* Mathematical Geology **34**, 597–611 (2002).

A. Jewbali and R. Dimitrakopoulos, *Implementation of conditional simulation by successive residuals,* Computers & Geosciences **37**, 129 – 142 (2011).

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

G. Evensen, *Sequential data assimilation with a nonlinear quasi-geostrophic model using monte carlo methods to forecast error statistics,* Journal of Geophysical Research: Oceans **99**, 10143–10162 (1994).

G. Evensen, *Advanced data assimilation for strongly nonlinear dynamics,* Monthly Weather Review **125**, 1342–1354 (1997).

G. Burgers, P. van Leeuwen, and G. Evensen, *Analysis scheme in the ensemble kalman filter,* Monthly Weather Review **126**, 1719–1724 (1998).

J. Jansen, S. Douma, D. Brouwer, P. van den Hof, O. Bosgra, and A. Heemink, *Closed-loop reservoir management,* in *Proceedings SPE reservoir simulation symposium, Woodlands* (Society of petroleum engineers, 2009).

S. Aanonsen, G. Naevdal, D. Oliver, A. Reynolds, and B. Valles, *The ensemble kalman filter in reservoir engineering - a review,* SPE Journal **14**, 393–412 (2009).

J. Hou, K. Zhou, X.-S. Zhang, X.-D. Kang, and H. Xie, *A review of closed-loop reservoir management,* Petroleum Science **12**, 114–128 (2015).

J. Benndorf and J. Jansen, *Recent developments in closed-loop approaches for real-time mining and petroleum extraction,* Mathematical Geosciences **49**, 277–306 (2017).

L. Heidari, V. Gervais, M. Le Ravalec, and H. Wackernagel, *History matching of reservoir models by ensemble kalman filtering: the state of the art and a sensitivity study,* in *Uncertainty analysis and reservoir modeling, AAPG memoir 96* (American Association of Petroleum Geologists, 2011) pp. 249–264.

H. J. Hendricks Franssen, H. P. Kaiser, U. Kuhlmann, G. Bauser, F. Stauffer, R. Müller, and W. Kinzelbach, *Operational real-time modeling with ensemble kalman filter of variably saturated subsurface flow including stream-aquifer interaction and parameter updating,* Water Resources Research **47**, 1–20 (2011).

L. Hu, Y. Zhao, Y. Liu, C. Scheepens, and A. Bouchard, *Updating multipoint simulations using the ensemble kalman filter,* Computers & Geosciences **51**, 7 – 15 (2013).

J. Benndorf, *Making use of online production data: Sequential updating of mineral resource models,* Mathematical Geosciences **47**, 547–563 (2015).

3

ALGORITHM

This chapter further elaborates on the design of the algorithm and its specific options. The first section discusses the rationale behind the Gaussian anamorphosis. The implementation, allowing for an intuitive treatment of measurement error, is explained. The second section provides a computationally efficient strategy for solving the updating equations. The presentation is generic, in that it includes the option to configure a parallel updating sequence (double helix). The third section discusses the neighbourhood option. The fourth section provides some insight into the covariance correction techniques. The fifth section presents the pseudocode of the updating algorithm. The pseudocode illustrates how the individual functional components are integrated.

Parts of this chapter have been published in Mathematical Geosciences **49**, 1 (Wambeke and Benndorf, 2017) and in Proceedings of the 17th annual conference of the international association for mathematical geosciences (Wambeke and Benndorf, 2015).

3.1. GAUSSIAN ANAMORPHOSIS

The updating approach presented in Benndorf (2015) is only optimal if all involved variables are multivariate Gaussian and if the forward observation model $\mathscr{A}$ is linear. Commonly, global normal score transformations are used to mitigate the effects of multivariate distributions which are not Gaussian (Goovaerts, 1997; Deutsch and Journel, 1998). However due to the developing non-stationarity, the underlying assumptions get violated. That is, after a few updates, the global cumulative field distribution does not anymore represent the local updated conditions. Instead, grid nodes are individually transformed according to node specific Gaussian anamorphosis functions (Béal *et al.*, 2010; Simon and Bertino, 2009). This approximation just renders marginal distributions as univariate Gaussian. Multivariate distributions are not explicitly corrected. Yet, the random vectors are closer to being multivariate Gaussian than prior to the univariate transformations (Zhou *et al.*, 2011). First the univariate transformation of grid nodes is discussed. The transformation of observations and the treatment of measurement error is addressed later on.

At each grid node $\mathbf{x}_n$, a local Gaussian anamorphosis function $\Phi_{\mathbf{x}_n}$ is built and used to forward transform the Monte Carlo representation of the random vector $\mathbf{z}(\mathbf{x}_n)$:

$$\mathbf{U}_t(n,i) = \Phi_{t,\mathbf{x}_n}(\mathbf{Z}_t(n,i)) \tag{3.1a}$$

$$= G^{-1}(F(\mathbf{Z}_t(n,i))) \quad \forall i \in I, \tag{3.1b}$$

where I refers to the total number elements in the Monte Carlo sample. A two step procedure is used for the implementation of the forward transformation. (i) The cumulative probability $F(\mathbf{Z}_t(n,i))$ is estimated based on the rank i' of the corresponding element in the full Monte Carlo sample ($\frac{i'-0.5}{I}$). (ii) The resulting cumulative probability is mapped onto its corresponding standard normal score value. In other words, the inverse of the standard normal cumulative distribution function G is computed. The left hand side of Figure 3.1 shows the procedure for the univariate transformation of grid nodes.

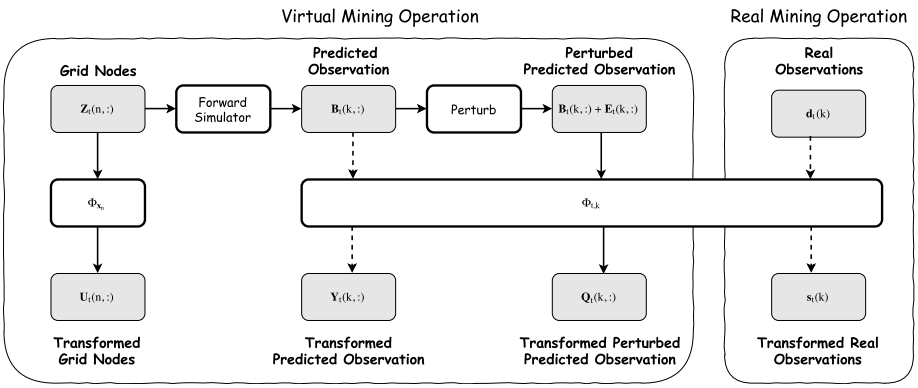

Figure 3.1: Univariate transformation of grid nodes, perturbed predicted observations, predicted observations and actual observations

The necessity and added value of a transformation of observations depends partially on the considered updating scheme (Eq. 2.3a versus Eq. 2.3b). To further clarify this statement, a situation with transformed grid nodes is assumed and the impact of non-Gaussian observations on both updating schemes is analysed.

A study of Eq. 2.3b learns that the Kriging weights $\mathbf{W}_t$ can be computed without explicitly considering the predicted observations. The transformed grid nodes $\mathbf{U}(n,:)$ are used to compute the empirical field error covariances $\mathbf{C}_{t-1,uu}$. Both forecast error and observation error covariances are then written as linear combinations hereof. The resulting matrices, $\mathbf{C}_{t-1,uu}\mathbf{A}_t^T$ and $\mathbf{A}_t\mathbf{C}_{t-1,uu}^{-1}\mathbf{A}_t^T + \mathbf{R}$ are fully described in a normal score space. As a consequence, sensor precisions need to be characterized in the normal score domain as well (R matrix). In conclusion, non-Gaussian observations thus have a limited impact on the computation of the weights $\mathbf{W}_t$ when using updating scheme 2.3b.

In contrast, due to the specific formulation of updating scheme 2.3a, predicted observations and their distributions do have a direct and significant impact on the computation of the weights $\mathbf{W}_t$. The forecast error and observation error covariances are after all empirically calculated based on untransformed predicted observations. They are no longer expressed as linear combinations of field error covariances, computed from transformed grid nodes. The use of a flexible forward simulator thus results in a less far-reaching effect of transformed grid nodes when computing the weights $\mathbf{W}_t$.

Since updating scheme 2.3a is less robust against non-Gaussian observations, a transformation is warranted. This will further increase the performance of the sequential updating technique inasmuch as all involved variables, including the observations, better approximate a multivariate Gaussian distribution. Moreover, a Gaussian anamorphosis of the observations leads to an apparent pseudolinearization of the forward observation model in the normal score domain (Schöniger *et al.*, 2012). This pseudolinearization can be more accurately exploited by the linear updating equations.

The K predicted observations at time step t are defined as

$$\mathbf{B}_t(:,i) = \mathscr{A}_t(\mathbf{Z}_{t-1}(:,i)) \quad \forall i \in I, \tag{3.2}$$

where K refers to the number of measurements taken over a period $[t-1,t]$ and $\mathscr{A}_t$ represents a forward simulator translating single field realizations $\mathbf{Z}_t(:,i)$ into vectors $\mathbf{B}_t(:,i)$ with K predicted observations. In order to account for measurement error, the predicted observations are randomly perturbed, that is I vectors $\mathbf{E}_t(:,i)$ drawn from $N(0,\mathbf{R})$ are added to the columns $\mathbf{B}_t(:,i)$. Note that the diagonal R matrix directly describes the measurement error in the original domain, so that a translation to a normal score space is not needed. The rows of the resulting perturbed predicted observation matrix are transformed individually by evaluating K Gaussian anamorphosis functions (Fig. 3.1)

$$\mathbf{Q}_t(k,i) = \Phi_{t,k}(\mathbf{B}_t(k,i) + \mathbf{E}_t(k,i)) \quad \forall i \in I. \tag{3.3}$$

The addition of random noise $\mathbf{E}_t(:,i)$ results in a kernel smoothing effect making the estimation of the empirical CDF more robust, even for small sample sizes (Cheng and Parzen, 1997). Each of the K actual inaccurate observations are assumed to originate from a distribution described by its corresponding Monte Carlo sample of perturbed predicted observations ($\mathbf{B}_t(k,:) + \mathbf{E}_t(k,:)$). Hence, the actual observations are

transformed according to the previously constructed anamorphosis functions (Fig. 3.1)

$$\mathbf{s}_t(k) = \Phi_{t,k}(\mathbf{d}_t(k)). \tag{3.4}$$

The integration of measurement noise into the transformations allows for a smooth computation of covariances and deviations, without the need to further consider measurement error in the normal score domain. The normal score equivalent **R** matrix is implicitly accounted for when computing observation error covariances from the rows of the $\mathbf{Q}_t$ matrix. Previously, a noise vector had to be explicitly added to the deviations between predicted and actual observations (Eq. 2.5). Now, a normal score equivalent is automatically inserted into the differences between $\mathbf{s}_t$ and the columns $\mathbf{Q}_t(:,i)$.

The currently available transformed data would not allow for a correct empirical computation of the forecast error covariances. A pairwise consideration of rows from the **U** and **Q** matrix would result in overestimated forecast error covariances. This is because the forecast error covariances should be computed without the consideration of measurement noise. Therefore, the previously constructed anamorphosis functions $\Phi_{t,k}$ are used once more to transform unperturbed predicted observations (Fig. 3.1)

$$\mathbf{Y}_t(k,i) = \Phi_{t,k}(\mathbf{B}_t(k,i)) \quad \text{for } i = 1,2,\ldots,I. \tag{3.5}$$

After the updating step, the grid nodes are back transformed using the inverse of the grid node specific anamorphosis functions $\Phi^{-1}_{t,\mathbf{x}_n}$. Back transformation of the observations is not required. Instead updated predicted observations are obtained from a subsequent run of the forward simulator.

3.2. SOLVING THE UPDATING EQUATIONS

The following provides an efficient strategy for solving the updating equations. The presentation is generic in that it includes an option to configure a pair of connected updating cycles (double helix). This option can be activated to avoid problems of inbreeding. Inbreeding is essentially caused by the empirical estimation of covariances from Monte Carlo samples of limited size (Zhou *et al.*, 2011). When occurring, the spread in the Monte Carlo sample becomes unrealistically small, in particular the spread would be smaller than the difference between the sample mean and the unknown truth (Zhou *et al.*, 2011). The problem generally arises when a group of grid nodes is being updated frequently during a certain period of time.

The idea of configuring a pair of sequential updating cycles originates from Houtekamer and Mitchell (1998), who argued that the computation of the weights $\mathbf{W}_t$ and the evaluation of their quality (diagonal elements in $\mathbf{C}_{t+1,uu}$) would be based on exactly the same information when using a single Monte Carlo sample. As the assimilation cycle proceeds, the propagation of this dependent quality test might eventually result in a collapse of the Monte Carlo sample. In order to maintain sufficient spread, two connected updating cycles can be configured. The weights $\mathbf{W}_t$ computed from one cycle are used to assimilate data into the other cycle. In this way, each of the two cycles uses different Monte Carlo samples to estimate the weights and to evaluate the quality of an update. As such, the concept bears some structural resemblance to the double helical structure of nucleic acids such as DNA. The two updating cycles represent the double

stranded molecules typically found in DNA. The cross-over of the Kriging weights ($\mathbf{W}_t$) are then represented by the base pairs holding together the double helical structure.

To configure two connected updating cycles, the previously transformed Monte Carlo samples need to be split into two subsets. The constructed matrices with I columns ($\mathbf{U}$, $\mathbf{Q}$ and $\mathbf{Y}$) are subdivided into two new matrices with respectively α and $I-\alpha$ columns. Considering an original matrix $\mathbf{X}$, the new matrices are denoted by $\mathbf{X}(:,:\alpha)$ and $\mathbf{X}(:,\alpha:)$. The operators $:\alpha$ and $\alpha:$ refer respectively to the first α and the last $I-\alpha$ columns. The default value of the integer parameter α is set to $I/2$. Note that the two subsets do not necessarily have to be of equal size.

Once the Monte Carlo samples are split, all necessary covariances are empirically calculated. First two K x K matrices with observation error covariances are constructed:

$$\mathbf{C}^1_{t,qq}(k,k) = E[(\mathbf{Q}_t(k,:\alpha) - E[\mathbf{Q}_t(k,:\alpha)])(\mathbf{Q}_t(k,:\alpha) - E[\mathbf{Q}_t(k,:\alpha)])] \tag{3.6a}$$

$$\mathbf{C}^2_{t,qq}(k,k) = E[(\mathbf{Q}_t(k,\alpha:) - E[\mathbf{Q}_t(k,\alpha:)])(\mathbf{Q}_t(k,\alpha:) - E[\mathbf{Q}_t(k,\alpha:)])] \tag{3.6b}$$

The observation error covariances are computed from transformed perturbed predicted observations and as such implicitly account for a normal score equivalent of the diagonal R matrix. Likewise, two N x K matrices with forecast error covariances are constructed:

$$\mathbf{C}^1_{t,uy}(n,k) = E[(\mathbf{U}_t(n,:\alpha) - E[\mathbf{U}_t(n,:\alpha)])(\mathbf{Y}_t(k,:\alpha) - E[\mathbf{Y}_t(k,:\alpha)])] \tag{3.7a}$$

$$\mathbf{C}^2_{t,uy}(n,k) = E[(\mathbf{U}_t(n,\alpha:) - E[\mathbf{U}_t(n,\alpha:)])(\mathbf{Y}_t(k,\alpha:) - E[\mathbf{Y}_t(k,\alpha:)])] \tag{3.7b}$$

It is important to stress that the transformation of the unperturbed predicted observations (Eq. 3.5) is only performed to calculate the forecast error covariances.

A significant speedup of the updating operation can be achieved by solving the equations first in the K dimensional 'observation space' and then projecting the solutions into the N dimensional 'field space' (Evensen and van Leeuwen, 1996; Cohn *et al.*, 1998). A Cholesky decomposition of the two observation error covariance matrices results in considerable computation savings when solving following equations:

$$\mathbf{C}^2_{t,qq}\mathbf{T}(:,i') = \mathbf{s}_t - \mathbf{Q}_t(:,i') \qquad \forall i' \in [1,\alpha], \tag{3.8a}$$

$$\mathbf{C}^1_{t,qq}\mathbf{T}(:,i'') = \mathbf{s}_t - \mathbf{Q}_t(:,i'') \qquad \forall i'' \in [\alpha,I]. \tag{3.8b}$$

The two resulting lower triangular matrices are combined with deviations between transformed actual and predicted observations of the complementary subset to calculate a series of solution vectors $\mathbf{T}(:,i)$. Both series of solution vectors are subsequently combined with previously calculated forecast error covariances to update existing realizations:

$$\mathbf{U}_{t+1}(:,i') = \mathbf{U}_t(:,i') + \mathbf{C}^2_{t,uy}\mathbf{T}(:,i') \qquad \forall i \in [1,\alpha], \tag{3.9a}$$

$$\mathbf{U}_{t+1}(:,i'') = \mathbf{U}_t(:,i'') + \mathbf{C}^1_{t,uy}\mathbf{T}(:,i'') \qquad \forall i'' \in [\alpha,I]. \tag{3.9b}$$

Finally the Monte Carlo samples of grid nodes are recombined and transformed back using the grid node specific anamorphosis functions which were stored in memory:

$$\mathbf{Z}_{t+1}(n,i) = \Phi^{-1}_{t,\mathbf{x}_n}(\mathbf{U}(n,i)) \quad \forall i \in I. \tag{3.10}$$

3.3. NEIGHBOURHOOD

Often it is impractical (especially for large grids) to consider all N grid nodes simultaneously for updating, when K observations become available. Since the spatial correlation generally decreases with distance, it is sufficient to only include a node when it is located in the near vicinity of one of the L extraction points. A search algorithm based on a space-partitioning tree structure is implemented to allocate grid nodes to a neighbourhood set N' replacing the full grid N in all subsequent computations. A local neighbourhood approach is not new. However, the current implementation requires that neighbourhoods around active extraction points are considered simultaneously. For example, due to measurements on blended material streams, an observation can be correlated to multiple nodes from different neighbourhoods.

The subset N' is constructed prior to performing the transformations. As a result, the order of the $\mathbf{U}$ matrices in Eq. 3.1, 3.7, 3.9 and 3.10 greatly reduces permitting a substantial computational economy to be realized. A local neighbourhood strategy also reduces the effects of spurious correlations which may arise between observations and remote grid nodes. Eventually, spurious correlations might lead to an update of several remote nodes based on unrelevant information. To avoid such a fallacy, a neighbourhood strategy can be used to exclude remote nodes from the analysis.

3.4. COVARIANCE ERROR CORRECTION

The previous sections already touched upon some of the issues related to the empirical estimation of covariances from finite Monte Carlo samples. Statistical theory states that the correlation between two normal distributions can be estimated with the following accuracy when using I Monte Carlo pairs:

$$E[(\rho - \hat{\rho})^2] \approx \frac{(1-\hat{\rho})^2}{I}, \tag{3.11}$$

where ρ and $\hat{\rho}$ are respectively the unknown and estimated correlation. The term $E[(\rho - \hat{\rho})^2]$ expresses the variance around the estimate. It follows that an accurate estimation of weak correlations (or covariances) would in any case require very large samples with thousands of realizations. When due to practical considerations the Monte Carlo samples are rather limited, numerical inaccuracies may arise. To mitigate possible numerical inaccuracies, two covariance correction techniques are implemented; an error correction and a localization correction. Both techniques apply a differential correction to the individual elements of the forecast error covariance matrix $\mathbf{C}_{uy}$. The magnitude of the correction (indirectly) depends on the estimated value of each covariance element.

The localization correction is based on the notion that spatial covariance generally decreases with distance. Equation 3.11 indicates that the accuracy of a covariance estimate reduces accordingly. To account for distant dependent numerical inaccuracies, the forecast error covariance matrix $\mathbf{C}_{t,uy}$ is localized. The concept of localization requires observations which can be attributed to a certain spatial location. As a result, a physically sensible distance between an observation k and a grid block n can be calculated (Sakov (2011)). Given a taper function ρ, a correction factor $\rho(n,k)$

is computed and used to correct the corresponding covariance ($\mathbf{C}_{t,uy}(n,k) * \rho(n,k)$). Since in mining applications observations are made on blended material streams, the conventional implementation had to be adjusted. The following assumes that the locations of the L simultaneous extraction points are known. The localized element of the $\mathbf{C}_{uy}$ matrix can then be written as a sum of the product between grid realization anomalies ($\mathbf{U}_t(n,:) - E[\mathbf{U}_t(n,:)]$) tapered around extraction location l and observation anomalies ($\mathbf{Y}_t(k,:) - E[\mathbf{Y}_t(k,:)]$):

$$\mathbf{C}_{t,uy}(n,k) = \sum_{l=1}^{L} \Big(E\Big[\rho(n,l) * (\mathbf{U}_t(n,:) - E[\mathbf{U}_t(n,:)]) * (\mathbf{Y}_t(k,:) - E[\mathbf{Y}_t(k,:)]) \Big] \Big). \tag{3.12}$$

The Gaspari-Cohn correlation function is used for tapering the grid realization anomalies (Gaspari and Cohn, 1999). Its use requires the definition of three range parameters (x, y, z direction). A good choice for the range parameters is based on the number of realizations and the prior covariance (Chen and Oliver, 2010). As the number of realizations increase, the accuracy of covariance estimates increase and thus larger range parameters can be selected.

The error correction technique uses an offline Monte Carlo simulation to characterize the potential variations between unknown and estimated correlations (Eq. 3.11). A lookup table is constructed linking estimated correlations with correction factors between 0 and 1 (Anderson (2012)). The use of lookup table requires a fixed value range ($[-1,1]$), hence covariances are replaced by correlations without any further consequences. The procedure for constructing the lookup table goes as follows. For every discrete correlation ρ_m in the range $[-1,1]$, a Monte Carlo experiment is performed to compute the corresponding correction factor. Each Monte Carlo experiment consists of five steps. (i) Define a bivariate normal distribution with zero mean and covariance $[[1,\rho_m],[\rho_m,1]]$. (ii) Draw I times two values from this distribution. Compute the correlation coefficient based on the I pairs. (iii) Repeat (ii) P times. (iv) Compute the mean $\hat{\rho}_m$ and standard deviation σ_m of the P estimated correlation values. (v) Define α as $(\hat{\rho}_m/\sigma_m)^2$. The correction factor (CF) for an estimated correlation of ρ_m is then computed as

$$CF(\rho_m) = \frac{\alpha}{1+\alpha} * \frac{\hat{\rho}_m}{\rho_m}. \tag{3.13}$$

The described Monte Carlo experiment needs to be repeated M times in order to construct the lookup table. The parameters M and P are respectively set to 201 and 10^6. I should equal the number of realizations used to estimate the empirical correlations.

The selected M value results in 201 discretized correlation values ranging from -1 to 1, each 0.005 apart (a value of 201, compared to 200, results in a nicer spread of discretization points: $[-1.0,-0.995,-0.99,...-0.005,0,0.005,...0.990,0.995,1.0]$). The 201 discretization points were deemed to be sufficient to map the relationship between the estimated correlation coefficient ($\hat{\rho}_m$) and the accuracy of that estimate (σ_m). The P value (10^6), on the other hand, refers to the number of experiments conducted to measure the accuracy of the estimated correlation (σ_m). That is, the million experiment

values are used to compute the mean ($\hat{\rho}_m$) and standard deviation (σ_m) of an underlying normal distribution.

Once a lookup table is constructed or read from file, the error correction can be applied to the $\mathbf{C}_{t,uy}$ matrix

$$\mathbf{C}_{t,uy}(n,k) = CF\left(\frac{\mathbf{C}_{t,uy}(n,k)}{\mathbf{S}_{t,uu}(n,n)\mathbf{S}_{t,yy}(k,k)}\right) * \mathbf{C}_{t,uy}(n,k). \tag{3.14}$$

This requires the additional computation of the empirical standard deviation at N grid nodes ($\mathbf{S}_{t,uu}(n,n)$) and K observations ($\mathbf{S}_{t,yy}(k,k)$).

3.5. PSEUDOCODE

The following pseudocode illustrates how the previously discussed functional components are integrated in one updating procedure (Algm. 1, p. 29). The option of using two interconnected updating cycles has been omitted to avoid unnecessary complexity. In order accommodate for this option, lines 26 to 44 (Algm. 1, p. 29) have to be adjusted based on Eq. 3.6 to 3.9.

Algorithm 1 Updating Algorithm

procedure UPDATE($\mathbf{Z}_t, \mathbf{d}_t, \mathbf{B}_t, \Sigma$) ▷ Σ holds L extraction locations $\mathbf{x}_l$

  $tree \leftarrow$ CONSTRUCTTREE($\mathbf{Z}_t$)
  for $\mathbf{x}_l$ in Σ **do** ▷ collect nodes around $\mathbf{x}_l$
    $\mathbf{Z}_t^l \leftarrow$ SEARCHTREE($tree, \mathbf{x}_l$) ▷ $\mathbf{Z}_t^l$ is a subset of $\mathbf{Z}_t$
    $\mathbf{Z}_t^\Sigma \leftarrow$ UNION($\mathbf{Z}_t^\Sigma, \mathbf{Z}_t^l$) ▷ combine subsets

  $\mathbf{U}_t \leftarrow$ INITIALIZE(N', I) ▷ N' nodes in combined subsets
  for n in N' **do**
    $\Phi_{t,\mathbf{x}_n} \leftarrow$ CONSTRUCT($\mathbf{Z}_t^\Sigma(n,:)$)
    for i in I **do**
      $\mathbf{U}_t(n,i) \leftarrow \Phi_{t,\mathbf{x}_n}(\mathbf{Z}_t^\Sigma(n,i))$

  $\mathbf{Q}_t \leftarrow$ INITIALIZE(K, I) ▷ K observations
  $\mathbf{Y}_t \leftarrow$ INITIALIZE(K, I)
  $\mathbf{s}_t \leftarrow$ INITIALIZE(K, 1)
  $\mathbf{E}_t \leftarrow$ GENERATE(sensor accuracy) ▷ K x I matrix with white noise
  $\mathbf{F}_t \leftarrow \mathbf{B}_t + \mathbf{E}_t$
  for k in K **do**
    $\Phi_{t,k} \leftarrow$ CONSTRUCT($\mathbf{F}_t(k,:)$)
    for i in I **do**
      $\mathbf{Q}_t(k,i) \leftarrow \Phi_{t,k}(\mathbf{F}(k,i))$ ▷ transform perturbed predictions
      $\mathbf{Y}_t(k,i) \leftarrow \Phi_{t,k}(\mathbf{B}(k,i))$ ▷ transform predictions
      $\mathbf{s}_t(k) \leftarrow \Phi_{t,k}(\mathbf{d}_t(k))$ ▷ transform real observations

  $\mathbf{C}_{t,qq} \leftarrow$ INITIALIZE(K, K)
  for k_i in K **do**
    for k_j in K **do**
      $\mathbf{C}_{t,qq}(k_i,k_j) \leftarrow$ COV($\mathbf{Q}_{t,qq}(k_i,:)$,$\mathbf{Q}_{t,qq}(k_j,:)$) ▷ Eq. 3.6

Algorithm 1 Updating Algorithm (continued)

$\mathbf{C}_{t,uy} \leftarrow$ INITIALIZE(N', K)
for n in N' **do**
  for k in K **do**
    if CovarianceLocalization is FALSE **then**
      $\mathbf{C}_{t,uy}(n,k) \leftarrow$ COV$(\mathbf{U}_t(n,:), \mathbf{Y}_t(k,:))$ ▷ Eq. 3.7
    else if CovarianceLocalization is TRUE **then**
      $\mathbf{C}_{t,uy}(n,k) \leftarrow$ COVLOC$(\mathbf{U}_t(n,:), \mathbf{Y}_t(k,:), \Sigma)$ ▷ Eq. 3.12

if ErrorCorrection is TRUE **then**
  for n in N' **do**
    $\mathbf{S}_{t,uu}(n,n) \leftarrow$ STANDARDDEVIATION$(\mathbf{U}_t(n,:))$
    for k in K **do**
      $\mathbf{S}_{t,yy}(k,k) \leftarrow$ STANDARDDEVIATION$(\mathbf{Y}_t(k,:))$
      $\rho \leftarrow \mathbf{C}_{t,uy}(n,k) / (\mathbf{S}_{t,uu}(n,n) * \mathbf{S}_{t,yy}(k,k))$
      $CF \leftarrow$ LOOPUP(ρ) ▷ from table
      $\mathbf{C}_{t,uy}(n,k) \leftarrow CF * \mathbf{C}_{t,uy}(n,k)$ ▷ correct & replace, Eq. 3.14

$\mathbf{T} \leftarrow$ INITIALIZE(K, I)
for i in I **do**
  $\mathbf{T}(:,i) \leftarrow$ SOLVE$(\mathbf{C}_{t,qq}\mathbf{T}(:,i) = \mathbf{s}_t - \mathbf{Q}_t(:,i))$
  $\mathbf{U}_{t+1}(:,i) \leftarrow \mathbf{U}_t(:,i) + \mathbf{C}_{t,uy}\mathbf{T}(:,i)$

for n in N' **do**
  for i in I **do**
    $\mathbf{Z}^{\Sigma}_{t+1}(n,i) \leftarrow \Phi^{-1}_{t,\mathbf{x}_n}(\mathbf{U}_{t+1}(n,i))$
$\mathbf{Z}_{t+1} \leftarrow$ OVERWRITE$(\mathbf{Z}^{\Sigma}_{t+1}, \mathbf{Z}_t)$ ▷ Overwrite the N' selected nodes

Return $\mathbf{Z}_{t+1}$

3.6. CONCLUSIONS

This thesis presents a new algorithm to automatically assimilate online data from a production monitoring network into the grade control model. The practicality of the algorithm mainly results from the design decision to exclude the forward observation model from the computer code. Instead, for each unique application, a case specific forward simulator is built and used independently from the existing code. This results in great flexibility, allows for a better integration of expert knowledge and facilitates interdisciplinary cooperation.

Once built, the forward simulator is run to translate grade control realizations into observation realizations. The resulting realization sets are subsequently used to compute empirical covariances. These covariances mathematically describe the link between sensor observations and blocks from the grade control model. This eliminates the need to formulate and linearise an analytical forward observation model, let alone to compute its inverse. The forward simulator further ensures that the distribution of the Monte Carlo samples already reflects the support of the concerned random values. As a result, the necessary covariance, derived from these Monte Carlo samples, inherently accounts for differences in scale of support.

A Gaussian anamorphosis option is implemented to deal with suboptimal conditions related to non-Gaussian distributions. A specific algorithm structure ensures that the sensor precision (measurement error) can be defined on its original units and does not need to be translated into a normal score equivalent.

When, due to practical considerations, the number of Monte Carlo samples is rather limited, numerical inaccuracies may arise. An interconnected parallel updating sequence (double helix) can be configured to reduce the effects of filter inbreeding. Due to the empirical computation of the covariance, degrees of freedom are lost over time. Eventually this might result in a collapse of covariances. A neighbourhood option is implemented to constrain computation time and memory requirements. Since observations are collected on blended material streams originating from multiple extraction zones, different neighbourhoods need to be considered simultaneously. A neighbourhood strategy also partially reduces the effect of spurious correlations. Two covariance correction options are implemented to contain the propagation of statistical sampling errors originating from the empirical computation of covariances. The localization correction is based on the assumption that the accuracy of a covariance estimate decreases with distance from an extraction point. Hence, grid anomalies are tapered around central extraction locations. The error correction uses an offline Monte Carlo simulation to precompute the accuracy of a certain correlation estimate. Corresponding correction factors are stored in lookup table and used to adjust the elements of the forecast error covariance matrix.

REFERENCES

T. Wambeke and J. Benndorf, *A simulation-based geostatistical approach to real-time reconciliation of the grade control model,* Mathematical Geosciences **49**, 1–37 (2017).

T. Wambeke and J. Benndorf, *Data assimilation of sensor measurements to improve production forecasts in resource extraction,* in *Proceedings of the 17th annual*

conference of the international association for mathematical geosciences (International association for mathematical geosciences, 2015).

J. Benndorf, *Making use of online production data: Sequential updating of mineral resource models,* Mathematical Geosciences **47**, 547–563 (2015).

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

C. Deutsch and A. Journel, *GSLIB: Geostatistical Software Library and User's Guide,* Applied geostatistics series (Oxford University Press, Oxford, 1998).

D. Béal, P. Brasseur, J.-M. Brankart, Y. Ourmières, and J. Verron, *Characterization of mixing errors in a coupled physical biogeochemical model of the north atlantic: implications for nonlinear estimation using gaussian anamorphosis,* Ocean Science **6**, 247–262 (2010).

E. Simon and L. Bertino, *Application of the gaussian anamorphosis to assimilation in a 3-d coupled physical-ecosystem model of the north atlantic with the enkf: a twin experiment,* Ocean Science **5**, 495–510 (2009).

H. Zhou, J. Gómez-Hernández, H. Hendricks Franssen, and L. Li, *An approach to handling non-gaussianity of parameters and state variables in ensemble kalman filtering,* Advances in Water Resources **34**, 844 – 864 (2011).

A. Schöniger, W. Nowak, and H.-J. Hendricks Franssen, *Parameter estimation by ensemble kalman filters with transformed data: Approach and application to hydraulic tomography,* Water Resources Research **48**, 1–18 (2012).

C. Cheng and E. Parzen, *Unified estimators of smooth quantile and quantile density functions,* Journal of Statistical Planning and Inference **59**, 291 – 307 (1997).

P. Houtekamer and H. Mitchell, *Data assimilation using an ensemble kalman filter technique,* Monthly Weather Review **126**, 796–811 (1998).

G. Evensen and P. van Leeuwen, *Assimilation of geosat altimeter data for the agulhas current using the ensemble kalman filter with a quasigeostrophic model,* Monthly Weather Review **124**, 85–96 (1996).

S. Cohn, A. da Silva, J. Guo, M. Sienkiewicz, and D. Lamich, *Assessing the effects of data selection with the dao physical-space statistical analysis system,* Monthly Weather Review **126**, 2913–2926 (1998).

L. Sakov, P. Bertino, *Relation between two common localisation methods for the enkf,* Computational Geosciences **15**, 225–237 (2011).

G. Gaspari and S. E. Cohn, *Construction of correlation functions in two and three dimensions,* Quarterly Journal of the Royal Meteorological Society **125**, 723–757 (1999).

Y. Chen and D. Oliver, *Cross-covariances and localization for enkf in multiphase flow data assimilation,* Computational Geosciences **14**, 579–601 (2010).

J. Anderson, *Localization and sampling error correction in ensemble kalman filter data assimilation,* Monthly Weather Review **140**, 2359–2371 (2012).

4

SYNTHETIC EXPERIMENT: 2D CASE STUDY

The synthetic experiment, presented in this chapter, demonstrates that the algorithm is capable of improving the resource model based on inaccurate observations made on blended material streams. In order to conduct the experiment, an artificial environment is created, representing a mining operation with two extraction points of unequal production rate. The performance of the algorithm during the experiment is evaluated using three sets of criteria. (i) Horizontal and vertical cross-sections are visually inspected and compared with the "true but unknown" reality. (ii) Global assessment statistics are computed to study how the overall quality of the resource model evolves over time as production data are being assimilated. (iii) Local assessment statistics are computed to ensure that the observed global improvements also result in the correction of the local estimation biases.

Parts of this chapter have been published in Mathematical Geosciences **49**, 1 (Wambeke and Benndorf, 2017).

4.1. METHODOLOY

A synthetic experimental environment is created to demonstrate the performance of the developed algorithm. The synthetic environment represents a mining operation with two extraction points of unequal production rate. The resulting material streams are blended and intentionally inaccurate observations are made. The goal of the experiment is to test whether the algorithm is capable of improving the Grade Control (GC) model based on unequally blended inaccurate observations. In order to stress the synthetic nature of the experiment, units for the modelled attribute are omitted.

4.1.1. REFERENCE FIELD

During the course of the artificial experiment a true but unknown state is sampled twice. First, small point samples are collected on a regular grid prior to the construction of the GC model (exploration phase). Second, observations are made on a blend of multiple mining blocks (operational phase). To correctly account for both scales of support, two representations of the true but unknown field are constructed. Both the high resolution point and the lower resolution block reference field characterize the true state over an area of $300m$ x $300m$ using a different level of discretization.

The high resolution point reference field is constructed on a discretized grid with 300 x 300 cells of size $1m$ x $1m$. This reference field is generated using a random field simulation algorithm (sequential Gaussian simulation). Attribute values are drawn from a standard normal distribution (centered around 0). The reference field is simulated based on an isotropic exponential covariance function with a variance of 1.0 and a range of $100m$ (units of the modelled attributes are omitted to stress their synthetic nature). Due to field dimensions which are relatively small compared to the range of the covariance model ($300m$ versus $100m$), a four times larger simulation grid is used (600 x 600 cells of size $1m$ x $1m$). Eventually only the South-West quarter of the grid is retained (300 x 300 cells of size $1m$ x $1m$). The resulting field has an average of $-0.251(-0.052)$ and a variance of 0.862(1.006). The 5% and 95% quantiles are respectively $-1.733(-1.733)$ and 1.312(1.518). The values between brackets refer to those computed from the larger simulation grid (the desired field statistics, e.g. the variance and mean, are reproduced within the larger simulation grid but not in the smaller one - ref. ergodicity assumption in geostatistics).

The low resolution block reference field is constructed on a discretized grid with 60 x 60 cells of size $5m$ x $5m$. This second reference field is obtained from reblocking the previous point field. Values of 5 x 5 $1m$ x $1m$ cells are averaged and assigned to larger $5m$ x $5m$ blocks. The resulting block reference field has an average of -0.251 and a variance of 0.792. The 5% and 95% quantiles are respectively -1.696 and 1.246. As expected, the variability reduces due to reblocking. The final block reference field is displayed in Fig. 4.1a.

4.1.2. PRIOR SET OF REALIZATIONS

A set of field realizations is used throughout to characterize the spatial variability and to quantify the geological uncertainty of the study environment. The geological uncertainty originates from the limited amount of data available. Prior to the operation, the study prescribes the availability of exploration data collected on a $50m$ x $50m$

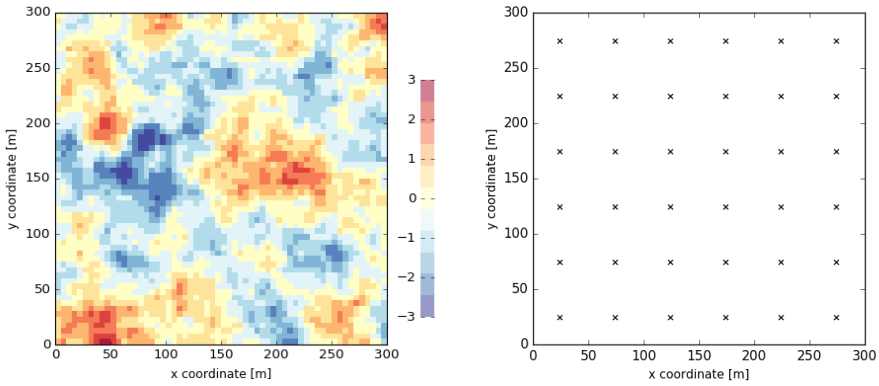

(a) (b)

Figure 4.1: True but unknown field (a) block representation, 60 x 60 blocks of size $5m$ x $5m$ colour coded according to their average block value (b) location of point samples on $50m$ x $50m$ exploration grid - point samples have a support of $1m$ x $1m$

sampling grid. This medium dense sampling grid (spacing equals half the range of the covariance model) results in a prior set of realizations with a realistic level of uncertainty and error.

A prior set of 200 realizations is simulated on a discretized grid with 600 x 600 $1m$ x $1m$ cells. Again, the simulation grid is four times larger than the studied area for reasons of small field dimensions compared to the range of the covariance model. A total of 144 sampling points are provided as conditioning data (sampled on a $50m$ x $50m$ exploration grid from the previous point simulation on the large 600 x 600 grid with $1m$ x $1m$ cells). From the 144, only 36 sampling points are actually located inside the study area (Fig. 4.1b). The covariance model of the true field is assumed to be known and was not inferred from the sampling points (isotropic exponential function with a variance of 1.0 and a range of $100m$). The effect of an incorrectly inferred covariance model is discussed in Sections 5.7 (last paragraph) and 7.2.2 (second paragraph).

Once simulated, only the South-West quarter of the grid is retained (300 x 300 cells of size $1m$ x $1m$). Subsequently, the 200 point simulations are reblocked onto a grid with 60 x 60 blocks of size $5m$ x $5m$. Figure 4.2 displays four arbitrarily chosen block realizations.

Only the block realizations on a discretized grid with 60 x 60 blocks of size $5m$ x $5m$ are further considered (one reference field $\mathbf{Z}^*$ and 200 prior block realizations $\mathbf{Z}_0(:,i)$). All previous point simulations were just an aid to this end.

Figure 4.3a and 4.3b display a mean and standard deviation (SD) field computed from the 200 prior realizations. The mean field, representing the current best block estimate, is considerably smoother than the reference field (compare Fig. 4.1a and 4.3a). The general larger scale patterns are nevertheless reasonably reproduced. This is due to the availability of sufficient conditioning data (on a point support). The approximate locations of the sampling points are marked by a lower block standard deviation (compare Figs. 4.3b and 4.1b). Due to effects of differences in scale of support,

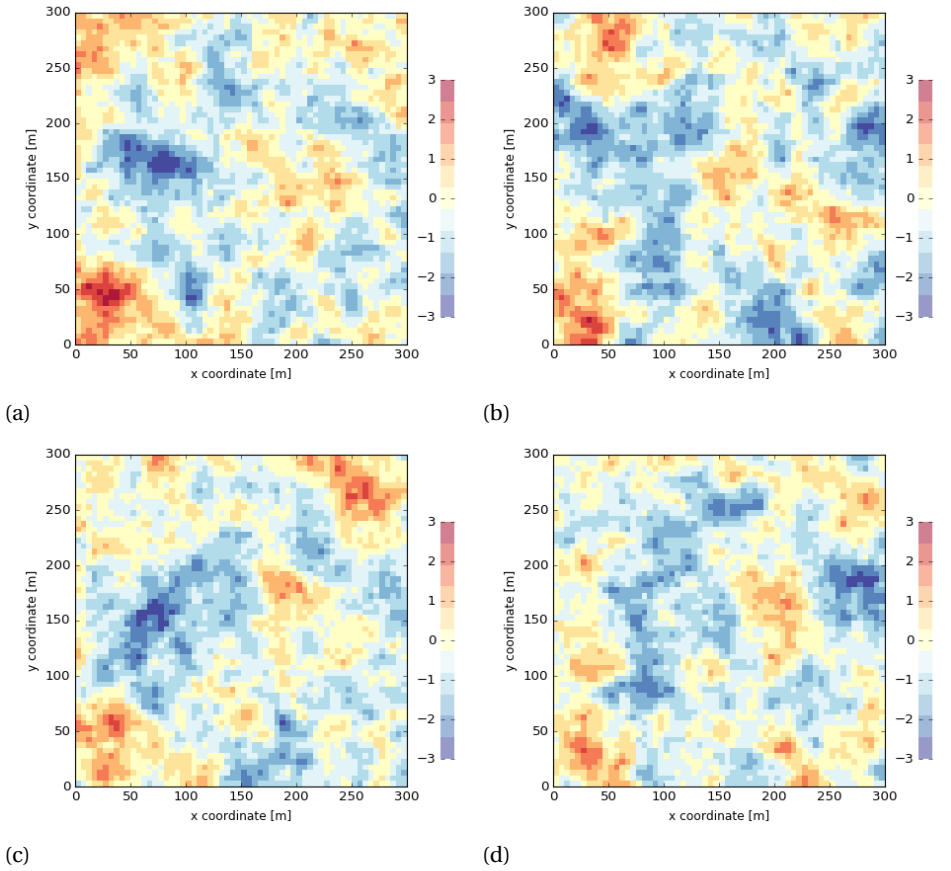

Figure 4.2: Four arbitrarily chosen block realizations (no. 39, 56, 125 and 183), 60x60 blocks of size $5m$ x $5m$ colour coded according to their average block value

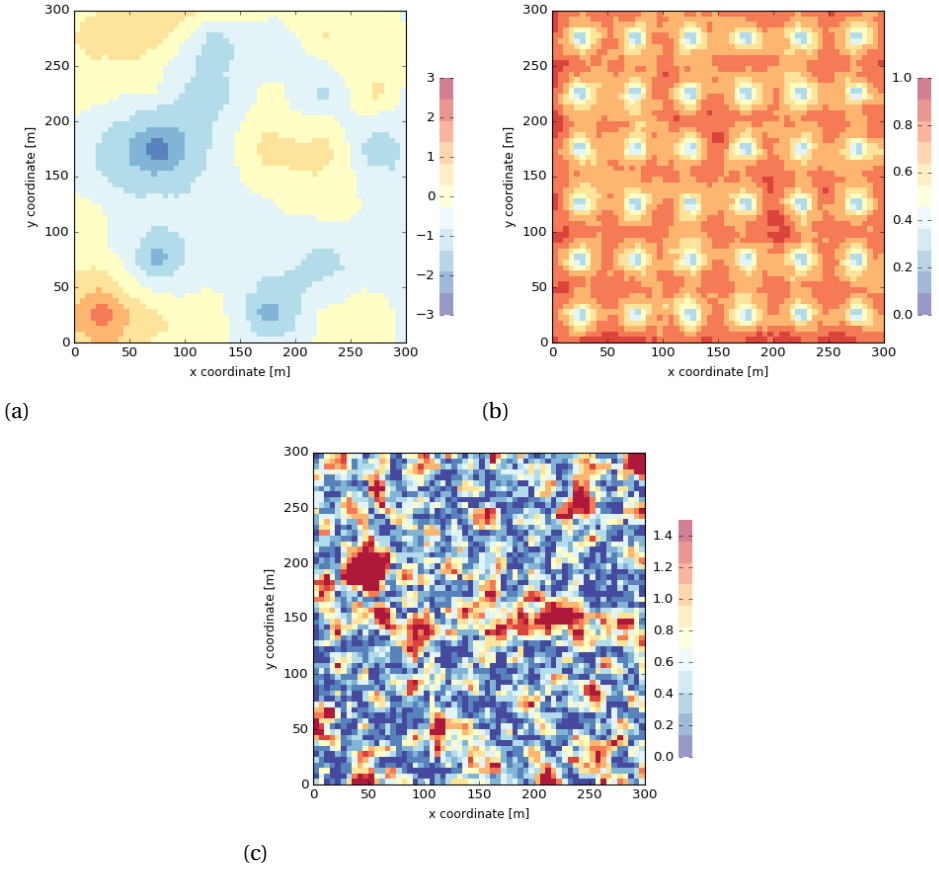

Figure 4.3: Grid with 60 x 60 blocks of size $5m$ x $5m$ colour coded according to the a computed statistic (a) mean (b) standard deviation (c) absolute error

the SD of $5m$ x $5m$ blocks does not entirely reduce to zero when conditioning based on cells of size $1m$ x $1m$. The minimum and maximum observed SD amount to 0.283 and 0.907. Figure 4.3c displays the absolute difference between true but unknown and estimated block values. The absolute error (AE) field is nothing more than the absolute difference between Fig. 4.1a and Fig. 4.3a. The figure illustrates that large deviations do occur locally, despite the considerable exploration effort. For example, the absolute error is consistently larger than 1.5 in an area bounded by $[25m, 75m]$ x $[175m, 225m]$ (in this particular area, the absolute error exceeds 2.0 in the majority of blocks, Fig. 4.5a).

Figure 4.4a and 4.4b depict respectively a North-South ($x = 57.5m$) and East-West section ($y = 167.5m$) across the prior realizations and the reference field. The information contained in the prior realizations is visualized as a combination of a most expected mean field (blue line) and a 90% confidence interval (blue area). The width of the 90% confidence interval is nearly uniform across both section lines (i.e. the

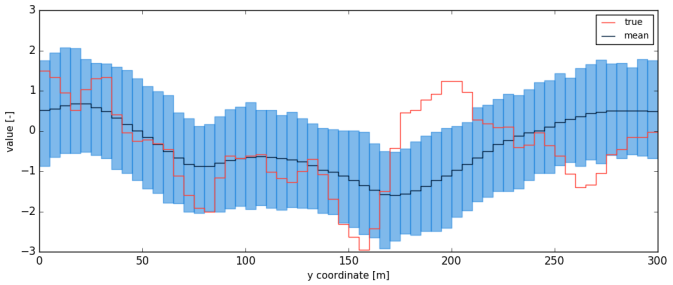

(a)

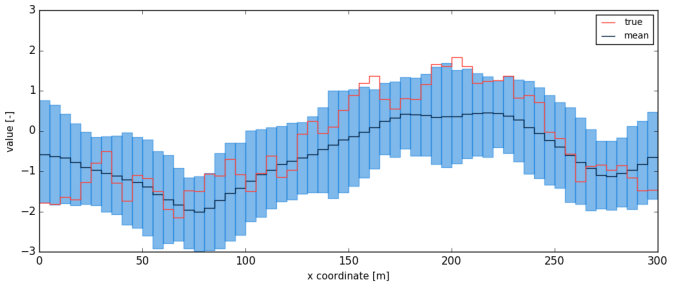

(b)

Figure 4.4: Sections across prior realizations and reference field (a) North-South section at $x = 57.5m$ (b) East-West section at $y = 167.5m$

influence of nearby sampling points is not immediately visible). The confidence interval of the North-South and East-West section encloses the reference field at respectively 44 (73.33%) and 53 (88.33%) blocks (Fig 4.4). Considering the entire grid, 3,225 out of 3,600 blocks (89.58%)are characterized by a confidence interval enclosing the reference value.

Figure 4.5a and 4.5b display respectively a North-South ($x = 57.5m$) and East-West section ($y = 167.5m$) across the AE (purple solid) and SD (orange dashed) field. A comparison between Fig. 4.4 and 4.5 shows that a severe misalignment between reference field and confidence interval directly results in an AE which is consistently larger than 1.0. The three peaks of large AE in the North-South section (Fig. 4.5a) result from a severe but local estimation bias (overestimation around $y = 150m$ and $y > 250m$, underestimation at $175m < y < 225m$, Fig. 4.4a). The two smaller peaks of AE in the East-West section (Fig. 4.5b) originate from a relatively large region with a smaller estimation bias (consistent underestimation between $125m < x < 250m$). The SD, which provides an alternative representation of remaining uncertainty, does not change significantly across both selected sections (slight decrease in the proximity of sampling points, Fig. 4.5b).

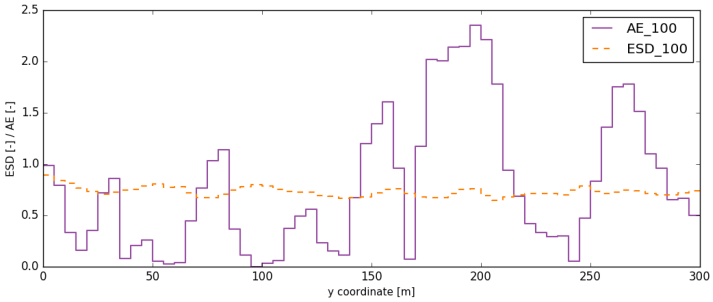

(a)

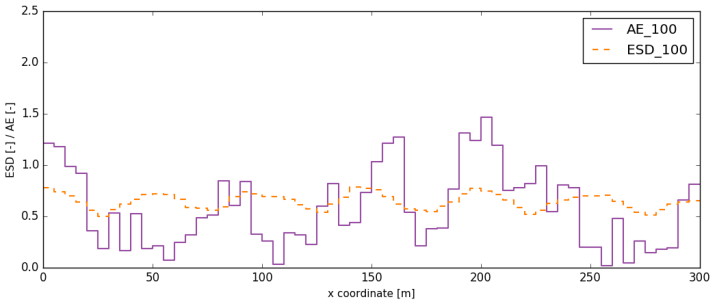

(b)

Figure 4.5: Sections across absolute error (AE, purple solid) and standard deviation (SD, orange dashed) field (a) North-South section at $x = 57.5m$ (b) East-West section at $y = 167.5m$

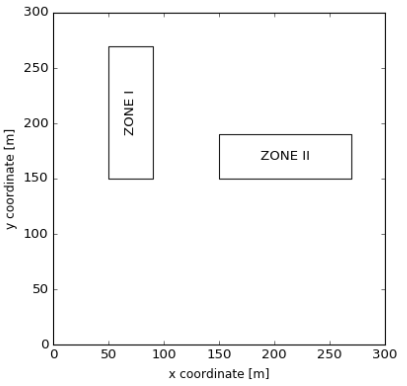

Figure 4.6: Outline of extraction zone I ([$50m, 150m$] x [$90m, 270m$]) and II ([$150m, 150m$] x [$270m, 190m$])

4.1.3. EXPERIMENTAL SCENARIO

A mining operation with two extraction points of unequal production rate is simulated and monitored. Figure 4.6 displays the outline of extraction zone I ($40m$ x $120m$) and zone II ($120m$ x $40m$). Each extraction zone has an area of $4,800m^2$ and contains 192 blocks. From this point forward, it will be assumed that the previous constructed block realizations have a thickness of $3m$. Considering a rock density of $2.8t/m^3$, it is not unreasonable to assume that each $75m^3$ block is transported by a truck with a payload capacity of $210t$. Hence the terms block and truckload are used interchangeably. The duration of the experiment equals the time T needed to excavate and process all 384 blocks from both extraction zones. This total time is subdivided into 12 discrete timesteps dt of equal length ($T = 12 * dt$). Since the total production rate is kept constant, 32 truckloads are processed during a single timestep.

Figure 4.7 provides a schematic overview of the extraction sequence and monitoring setup (forward simulator). The extraction sequence is visually prescribed through the numbers inside the demarcated digging blocks. During the first half time step ([$0, dt/2$]), a small digging block is extracted from mining zone I ($1a$, [$30m, 150m$] x [$40m, 160m$], 4 truckloads). Simultaneously, a large digging block is removed from mining zone II ($1a$, [$150m, 130m$] x [$165mx170m$], 12 truckloads). During the second half timestep ([$dt/2, dt$]), a small ($1b$, [$40m, 150m$] x [$50m, 160m$]) and large ($1b$, [$165m, 130m$] x [$180m, 150m$]) digging block are respectively extracted from mining zone I and II. A similar generic approach applies for the remaining timesteps. Generally, during each half timestep, 4 and 12 truckloads are extracted always from complementary mining zones (4 from one, 12 from the other). As a result, the total production rate remains constant; 16 truckloads per half timestep. Figure 4.7 however illustrates that the local production rates (4 or 12 truckloads per half timestep) reverse after every 3^{th} timestep.

The monitoring setup is strongly related to the extraction sequence. Every timestep yields two observations, each an average of 16 blended truckloads. The composition of the blend is prescribed by the defined extraction sequence. This is at time t, two observations become available, characterizing the blend resulting from the extraction during respectively periods $[t-1, t-1+dt/2]$ and $[t-1+dt/2, t]$. A forward simulator

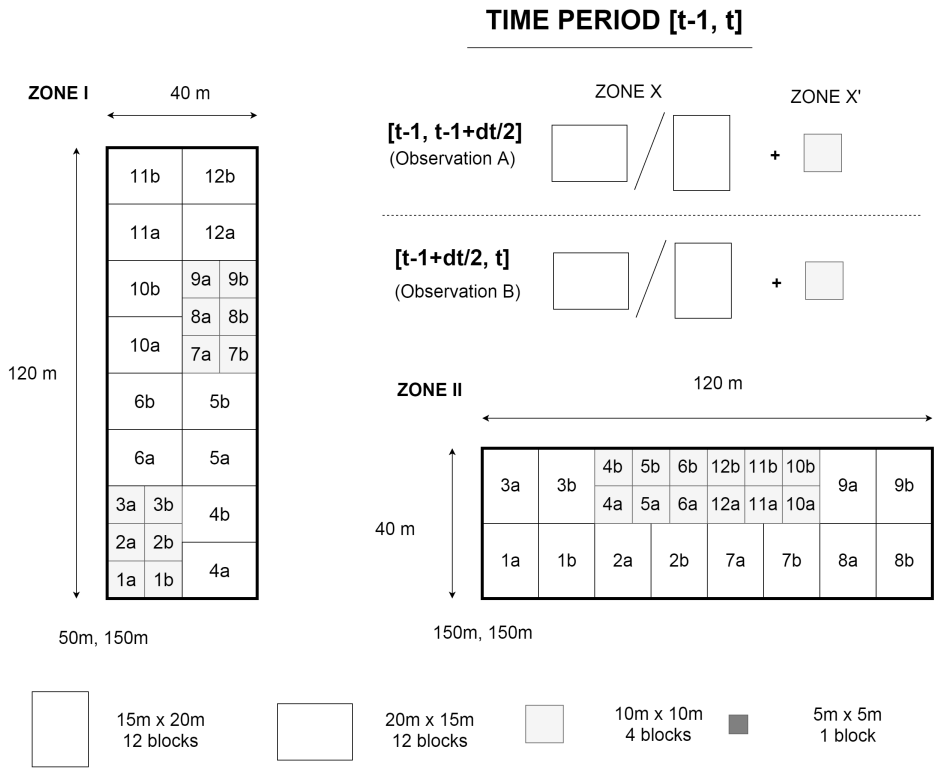

Figure 4.7: Schematic overview of extraction sequence and monitoring setup - during each half timestep dt, 4 and 12 blocks are extracted from complementary zones (X and X') and blended

is implemented based on the described extraction sequence and measurement setup. The forward simulator $\mathscr{A}_t$ translates each (prior) realization $\mathbf{Z}_{t-1}(:,i)$ into a vector $\mathbf{B}_t(:,i)$ with two observations (Fig. 2.1).

Obviously in practice, the real observation vector $\mathbf{d}_t$ results from actual measurements with inaccurate sensors. However, due to the synthetic nature of the experiment, a second forward simulator $\mathscr{A}_t^*$ needs to be built to mimic the behaviour of a real monitoring network. This simulator generates inaccurate but true observations $\mathbf{d}_t$ based on the reference field $\mathbf{Z}^*$ ($\mathbf{d}_t = \mathscr{A}_t^*(\mathbf{Z}^*)$). The previous constructed simulator $\mathscr{A}_t$ is adjusted to suit this purpose. Two main modifications are made. First, the reference field $\mathbf{Z}^*$ is propagated through the simulator and not one of the realizations $\mathbf{Z}_{t-1}(:,i)$. Second, random noise is added onto the resulting observations to mimic the behaviour of an inaccurate sensor. The random noise is drawn from a normal distribution with zero mean and a standard deviation of 0.0625 ($0.25/\sqrt{16}$). The reported standard deviation defines the measurement accuracy over a volume of 16 truckloads ($1{,}200m^3$). The applied sensor has a rather low accuracy. The corresponding accuracy of one truckload measurement would be 0.25, merely a quarter of the accuracy of a block estimate (block standard deviation varies between 0.283 and 0.907). In real applications, the observation and related accuracy would result from a time average of a continuous sensor response.

The described experimental scenario is executed in 29 seconds on a standard laptop (Intel Core i7-3520M CPU @ 2.90Ghz, 8Gb RAM memory). A single update is completed in less than 2.5 seconds. The computation time is expected to only slightly increase when applying this technique to ongoing operations. It is important to understand that the time needed to solve the updating equations is only dependent on the number of grid nodes in the neighbourhood subset. The size of the neighbourhood subset is more determined by the geological scale of correlation and the number of extraction locations than by e.g. the total number of nodes in a typical application problem (can be up to a billion). Larger grids though will impact the time needed to construct the neighbourhood subset. Considering L extraction locations and a space-partitioning search tree, the time to construct a neighbourhood subset will be in the order of $L\log_2(N)$. Therefore, it is to be expected that the time needed to complete a single update in real applications will be an order of magnitude smaller than the time needed to collect time-averaged measurements (a few minutes versus a few hours).

4.2. RESULTS

The performance of the algorithm is evaluated using three sets of criteria. (i) Fields and cross sections are visually inspected. Mean and reference fields are compared, the reliability of the 90% confidence interval is assessed and the evolution of the AE and SD is studied. (ii) Global assessment statistics are computed. The overall quality of a set of realizations is evaluated by the magnitude of two single value measures; the Root Mean Square Error (RMSE) and the spread. (iii) Local assessment statistics are computed for 14 different areas. RMSE and spread are calculated per local area to aid a quantitative comparison. Scatterplots, boxplots and empirical probability plots are drawn to visually evaluate local improvements.

4.2.1. VISUAL INSPECTION

Figure 4.8 displays how the mean field changes through time when assimilating sensor observations from the production chain. The two large rectangular outlines demarcate the extent of both extractions zone (I and II). The four smaller rectangles delineate the digging blocks which were extracted, blended and measured during the considered time interval. A comparison of Fig. 4.8a to 4.8l with the prior mean field (Fig. 4.3a) and the reference field (4.1a) indicates that the mean field (the current best estimate) continuously improves in and around both extraction zones.

Figure 4.9a and 4.9b depict respectively a North-South ($x = 67.5m$) and East-West ($y = 167.5m$) section across the reference field and final realizations (after 12 updates). Both figures support the previous conclusion. The mean field in the neighbourhood of both extraction zones ($150m < y < 270m$ and $150m < x < 270m$) considerably improved compared to the initial model (Fig. 4.4). Two out of the three occurrences of severe local estimation bias in the North-South section are corrected ($175m < y < 225m$ and $y > 250m$, Fig. 4.4a and 4.9a). The third occurrence, an overestimation bias around $y = 150m$, could not be corrected. This local deviation did go undetected because large positive deviations of small digging blocks were blended with smaller negative deviations of larger digging blocks ($1a$ and $1b$ from zone I mixed with $1a$ and $1b$ from zone II, Fig 4.7). The two local corrections along the North-South section can also be observed in Fig. 4.8f and 4.8j (increase mean field, $175m < y < 225m$) and 4.8k (decrease mean field, $y > 250m$). The algorithm further corrected the underestimation bias in the East-West section (Fig. 4.4b - 4.9b, $125m < x < 250m$). The observed correction in the East-West section corresponds to a higher mean field in mining zone II (e.g. compare Fig. 4.8l and 4.8a).

Figure 4.10a and 4.10b display a North-South ($x = 57.5m$) and an East-West ($y = 167.5m$) section across the initial and final AE field. The 12 updates result in a significant reduction of the absolute error around the extraction zones. Two of the three peaks of large AE in the North-South section have disappeared (Fig. 4.10a, $175m < y < 225m$ and $y > 250m$). In the corresponding areas, a reduction in AE from 2.0 to 0.5 is observed. The height of the third peak however slightly increases from about 1.6 to 2.3 (Fig. 4.10a $y = 150m$). A plausible explanation for this behaviour has been provided previously. The correction of the underestimation bias around mining zone II causes a reduction of AE in the East part of the East-West section (Fig. 4.10b, $125m < x < 250m$). The increase of AE in the West part of this section is directly related to the discussed local underperformance of the algorithm at the Southern border of mining zone I (Fig 4.8a $y = 167.5m$, Fig. 4.8b $x = 57.5m$).

The assimilation of sensor observations also results in an adjustment of the uncertainty model. The width of the 90% confidence interval is approximately halved within both extraction zones (compare Fig. 4.4 and 4.9, $150m < y < 270m$ and $< 150m < x < 270m$). The remaining uncertainty inside the excavated areas is attributed to a combination of measurement error and scale of support effects. In other words, average inaccurate measurements of blended truckloads are not sufficient to uniquely characterize individual blocks without any remaining uncertainty. The evolution of the SD field through time illustrates that the remaining uncertainty is not uniformly adjusted across the excavated areas (Fig. 4.11). The updating algorithm seems to assign a lower

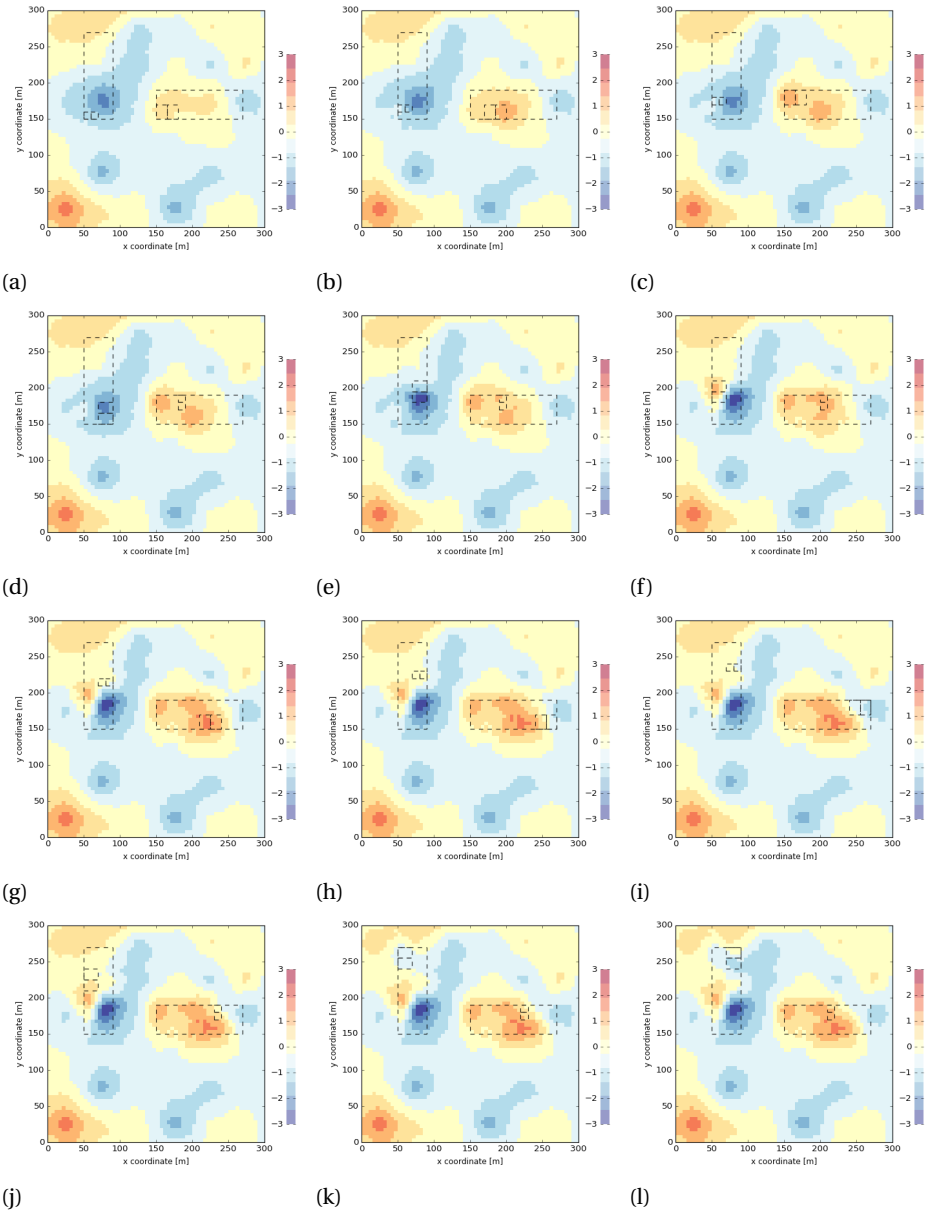

Figure 4.8: Evolution of mean field through time, from update 1 (a) to 12 (l) - grid with 60 x 60 blocks of size 5m x 5m colour coded according to the computed average block value

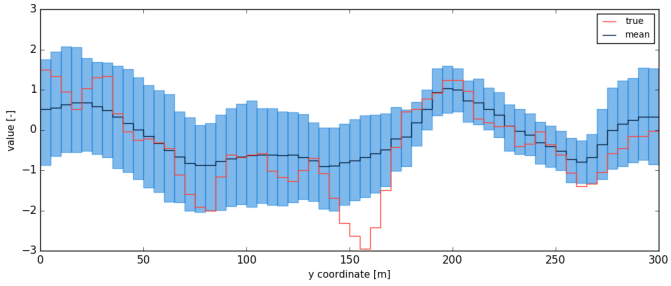

(a)

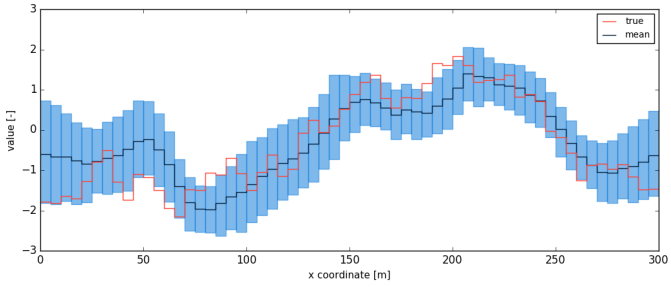

(b)

Figure 4.9: Sections across updated realizations (after update 12) and reference field (a) North-South section at $x = 57.5m$ (b) East-West section at $y = 167.5m$

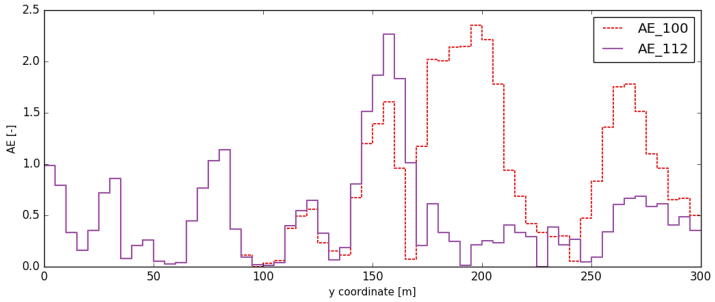

(a)

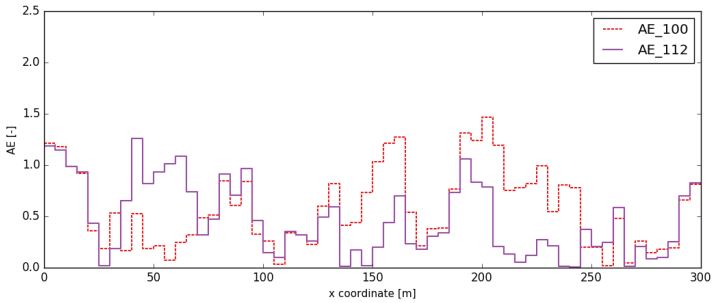

(b)

Figure 4.10: Sections across absolute error before (dashed red) and after 12 updates (solid purple) (a) North-South section at $x = 57.5m$ (b) East-West section at $y = 167.5m$

uncertainty (lower SD) to the larger digging blocks. The SD of updated blocks in large digging areas varies between 0.25 and 0.45 (Fig. 4.11). The SD of updated blocks in small digging areas is significantly higher, values between 0.55 and 0.75 are observed. A comparison of the SD fields across Fig. 4.11 indicates that the adjustment of the uncertainty model is not only limited to already excavated areas. For example, Figs. 4.11a and 4.11b show that the uncertainty near the middle of mining zone II reduces from somewhere around 0.7 to 0.5. This reduction occurs during the second update when none of the material near the middle of mining zone II has already been excavated. Fig. 4.11 and 4.3b illustrate that a reduction in uncertainty is observed up to a distance of $30m$ from the border of a digging block.

Figures 4.12a and 4.12b depict respectively a North-South ($x = 67.5m$) and East-West ($y = 167.5m$) section across the prior (Fig. 4.3b) and final (4.11l) SD field. The North-South section shows that the reduction of the SD is lower in the smaller digging areas ($150m < y < 180m$) compared to the larger ones ($180m < y < 270m$). North of the Northern border of extraction zone I ($y > 270m$), the SD increases from 0.25 (inside, $180m < y < 270m$) towards 0.75 over a distance of about $30m$. The uncertainty as well reduces up to $30m$ away from the Southern border ($120m < y < 150m$), though this reduction is less remarkable due to the larger SD values in the smaller digging areas ($150m < y < 180m$). The East-West section primarily intersects large digging areas (except for $50m < x < 70m$), hence the SD inside both extraction zones generally reduces to values between 0.25 and 0.35 (Fig. 4.12a). The SD inside the smaller digging areas ($50m < x < 70m$) reaches values near 0.65. The effects of a reduction in uncertainty are again observed up to $30m$ away from the boundaries of both extraction zones (zone I - $x < 50m$ and $x > 90m$; zone II - $x < 150m$ and $x > 270m$). The differences between the prior and final SD field displayed in the East-West section (Fig. 4.12b) are less striking compared to those in the North-South section (Fig. 4.12a). This is due to a lower and more variable prior SD field in the East-West section (larger effect of nearby sampling points) and has nothing to do with the updating results which are similar for both cross sections.

4.2.2. GLOBAL ASSESSMENT STATISTICS

The RMSE and spread are single value measures expressing the global quality of a set of realizations. As such, they allow for a more convenient comparison of realization sets across different timesteps. The RMSE measures the overall deviation between the reference field ($\mathbf{Z}^*$) and the mean field ($\hat{\mathbf{Z}}_t$) at a given timestep. The RMSE is defined as the root of the average squared differences between all available true and estimated block values:

$$RMSE_t = \sqrt{\frac{1}{N}\sum_{n=1}^{N}[\mathbf{Z}^*(n) - \hat{\mathbf{Z}}_t(n)]^2}. \tag{4.1}$$

The spread summarizes the overall estimated uncertainty and is calculated as the square root of the average block variance:

$$SPREAD_t = \sqrt{\frac{1}{N}\sum_{n=1}^{N} VAR(\mathbf{Z}_t(n,:))}. \tag{4.2}$$

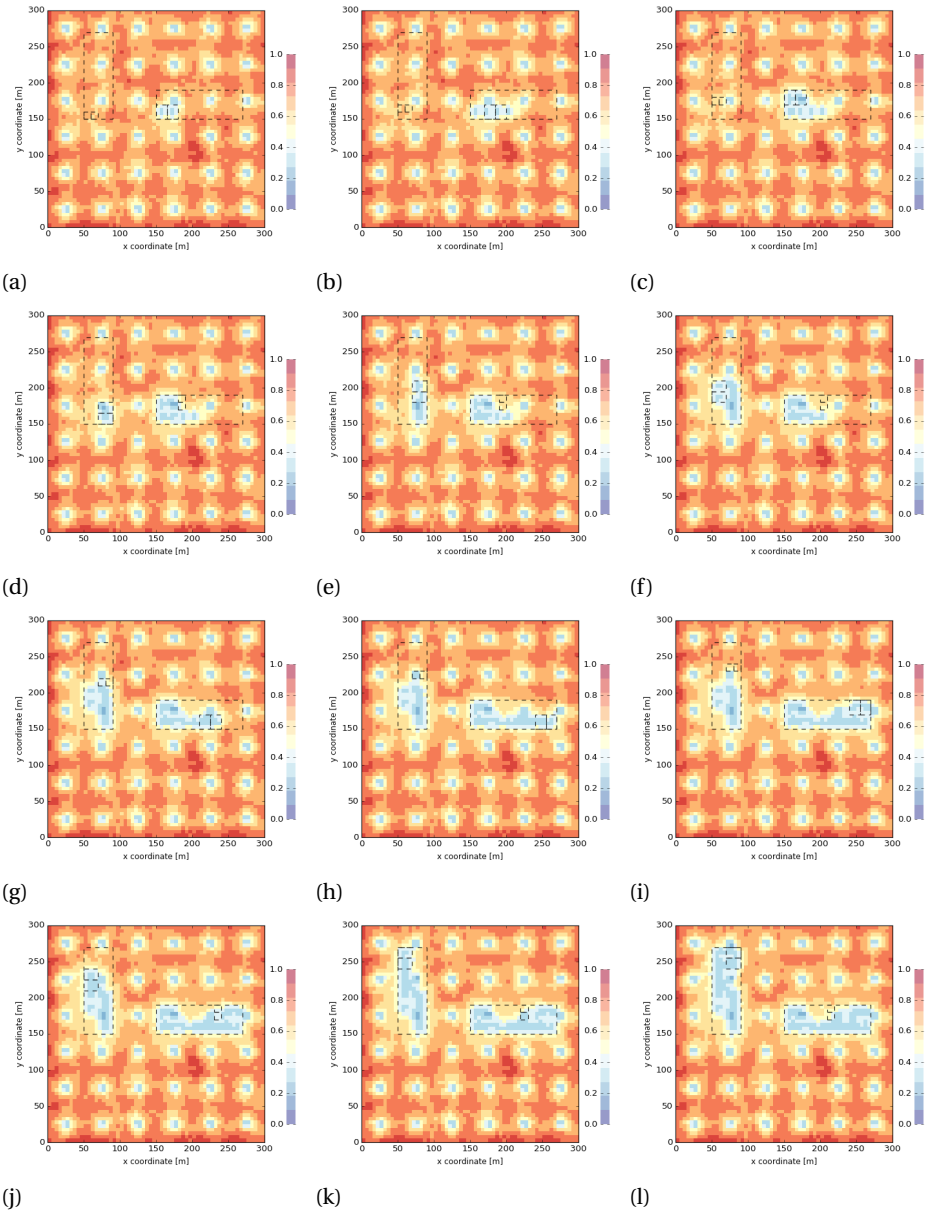

Figure 4.11: Evolution of SD field through time, from update 1 (a) to 12 (l) - grid with 60 x 60 blocks of size 5m x 5m colour coded according to the computed block standard deviation

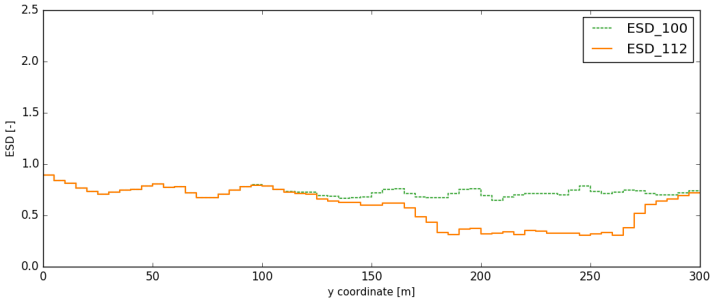

(a)

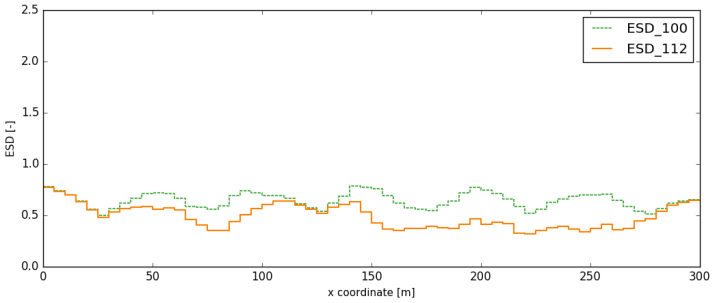

(b)

Figure 4.12: Sections across standard deviation before (dashed green) and after 12 updates (solid orange) (a) North-South section at $x = 57.5m$ (b) East-West section at $y = 167.5m$

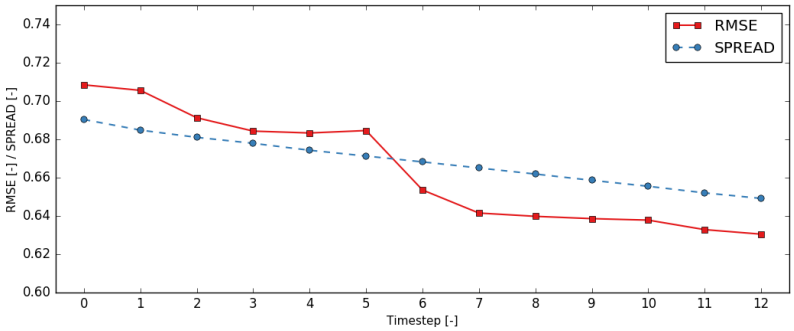

Figure 4.13: Evolution of RMSE and spread through time

In both formulations (Eq. 4.1 and 4.2) n iterates over all 3600 blocks in the grid. Ideally, the RMSE reduces as a function of the number of incorporated measurements. During the course of the experiment the uncertainty model is adjusted as well to account for the remaining error. Hence the need for a spread which approximates the RMSE.

Figure 4.13 shows the evolution of the global RMSE and spread through time. The RMSE of the prior realizations amounts to 0.708, while the prior spread is 0.690. Over the course of the experiment, the global RMSE drops by about 11% (from 0.708 to 0.630). The spread decreases accordingly, a reduction of 5% is observed (from 0.690 to 0.649). Figure 4.13 indicates that the spread decreases at a constant rate. This was to be expected since a constant material volume is measured during each timestep. There is no reason to assume that the information value of similar observations would be different over time. Figure 4.13 illustrates that the RMSE tends to reduce at the same rate. The more volatile behaviour is possibly caused by its large sensitivity to local improvements. For example, the rate of reduction in RMSE from timestep 5 to 6 is exceptionally large and inconsistent with the general pattern. This exceptionally large reduction originates from a local correction of an anomalous area. Figure 4.3c already indicated that the prior AE was consistently larger than 1.5 in an area bounded by $[25m, 75m]$ x $[175m, 225m]$. During the 6^{th} timestep a significant part of this area is excavated and measured. The algorithm detects this local anomaly and performs a correction. Figure 4.8e and 4.8f clearly show that the mean field in this local area is lifted during the 6^{th} timestep.

4.2.3. LOCAL ASSESSMENT STATISTICS

The previously reported global statistics are computed over all available 3600 blocks. During an update, only a fraction of this number is adjusted. The effects of local improvements are thus smeared out when computing a measure of global quality. The following presents a more local assessment to better comprehend the effects of an updating sequence.

A total of 14 local areas are defined and displayed in Fig. 4.14. Two local areas (I0 and II0) coincide with the previous demarcated extraction zones (zone I and II, Fig. 4.6). Antoher six frames, all $5m$ wide, are laid out around mine zone I and II(I1 to I6 and II1 to

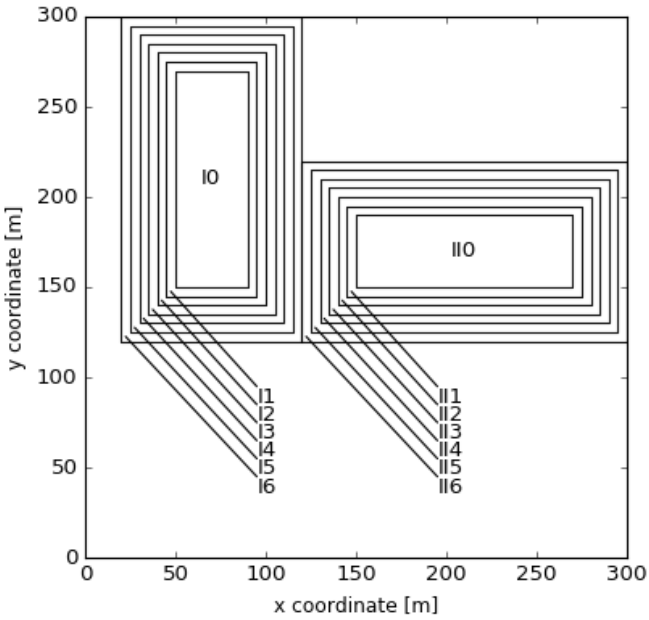

Figure 4.14: Outline of 14 local areas, 2 extraction zone of $4,800m^2$, 12 surrounding frames with a width of $5m$

II6).

For each local area, a true value is computed by averaging blocks of the reference field located inside the corresponding boundary. Similarly, each realization is converted into 14 local area values. Upon completion, 200 simulated local area values are collected and further used to compute a best estimate (i.e. mean) and quantile values. Figure 4.15 displays the true local area values versus their best estimates before and after 12 updates. A zero distance between a point and the bisector refers to a perfect alignment of a true and estimated value. The scatterplot indicates that the quality of the prior model (blue dots) is location dependent. In general, local area estimates deviate more from their reference values in and around mining zone II (open markers) than they do in and around mining zone I (filled markers). The 12 updates alter the realizations such that the estimated area values in and around zone II are shifted towards the bisector (red squares, Fig. 4.15). This corresponds to a general increase in mean field in the vicinity of zone II (compare Fig. 4.8a and 4.8l). Notice the large change in area value II0 and II1, both change respectively from 0.062 to 0.547 (0.673) and from −0.199 to 0.236 (0.427). Their corresponding reference values were indicated between brackets. The best estimates in mining zone I do not experience such drastic changes.

Figure 4.16 depicts boxplots constructed from the computed local area quantiles. The boxplots indicate that the prior uncertainty model partially accounts for the previously observed initial deviations (Fig 4.16a). For about 10 of the 14 local areas, the reference field is located inside the 90% confidence interval. The mismatch between reference value and confidence interval is remarkable in areas II0 to II4. After updating

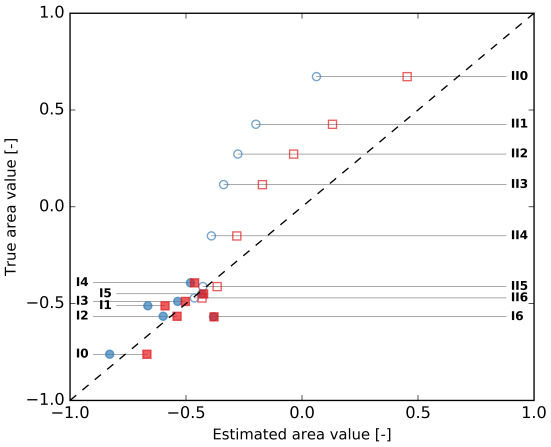

Figure 4.15: Scatterplot of true versus estimated area values before (blue dots) and after 12 updates (red squares) - filled markers refer to areas in mining zone I, open markers refer to areas in mining zone II

12 times, the boxplots characterizing local areas in mining zone II are shifted upwards and become shorter. Those characterizing mining I undergo a similar transition, although their upward shift is much more subtle. A uniform shift refers to a correction of bias, while shorter boxplots indicate a reduction of uncertainty.

The previously defined RMSE and spread can also be computed locally. In Eq. 4.1 and 4.2, n then refers to blocks confined to one of the specified local areas. Figure 4.17 displays the prior and posterior RMSE and spread per local area. Apparently, the prior RMSE of the local areas in mining zone I are slightly larger than those in mining zone II (Fig. 4.17a). This might seem contradictory, since previously the larger deviations between local area estimates and reference values were observed in mining zone II, and not in zone I (Fig. 4.15). The contradiction directly results from the effect of different bias mechanisms on the computed statistics. In mining zone I, three occurrences of a severe local under- or overestimation of the prior model are observed (zone I, Fig. 4.1a, 4.3a and 4.8l). When computing an area estimate, the local deviations are averaged out. However due to the squared difference in 4.1, these local large deviations do have a significant impact on the resulting RMSE value. The prior model in zone II is affected by a medium single bias extending over a larger area (zone II, Fig. 4.1a, 4.3a and 4.8l). This regional bias is not averaged out, hence the larger deviation between an area estimate and the reference value. Since the regional bias is of medium magnitude, the resulting RMSEs are not as large as those of mining zone I.

After 12 updates, the RMSEs of local areas I0 and II0 drop by about 38% and 45% (I0: from 0.984 to 0.608, II0: from 0.811 to 0.447, Fig. 4.17a). The first three local areas around mining zone I (I1, I2, I3, Fig. 4.14) observe a reduction in error of about 24%, 17% and 13%. The posterior RMSE values are respectively 0.833 (I1), 0.839 (I2) and 0.790 (I3). In the three frames surrounding mining zone II (II1, II2, II3, Fig. 4.14), the RMSE drops by about 39%, 27% and 19%, reaching posterior RMSE values of 0.547 (II1), 0.569 (II2)

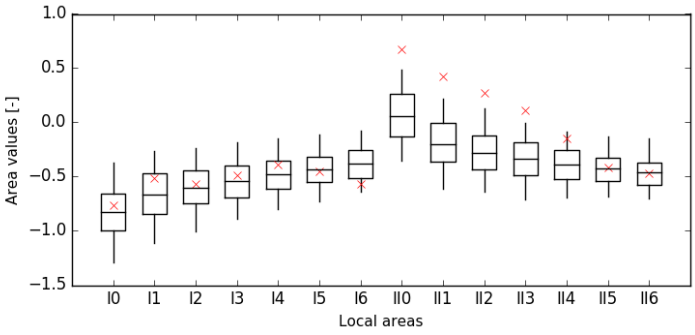

(a)

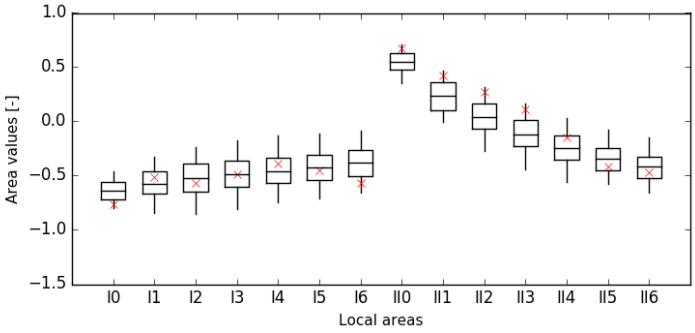

(b)

Figure 4.16: Boxplots constructed from the computed local area quantiles at timestep 0 (a) and timestep 12 (b) - reference values are indicated as red crosses, box and whiskers refer respectively to the 50% and 90% confidence interval

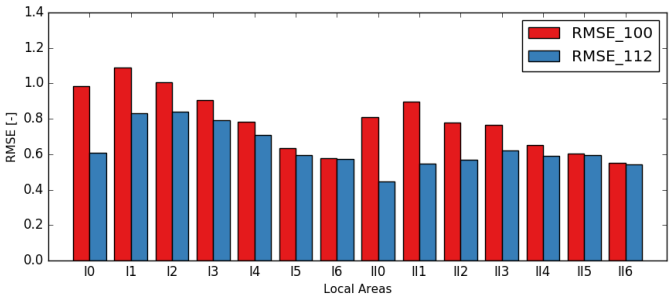

(a)

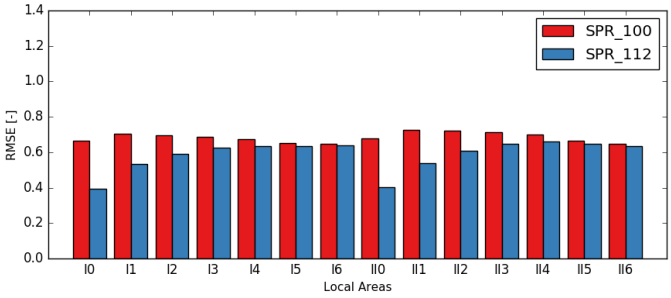

(b)

Figure 4.17: Barcharts displaying the local area RMSE (a) and spread (b) of the prior (red) and updated realizations (blue)

and 0.620 (II3). Around mining zone II, a significant improvement is observed and small posterior RMSE values are attained (0.55 to 0.60). The updates are still successful around mining zone, although medium RMSE values are obtained (0.80).

Figure 4.17b displays the prior and posterior spread per local area. The prior spread is more or less uniform across the areas and varies between 0.647 and 0.724. After 12 updates, the spread of local areas I0 and II0 drop by about 40% and 41% (I0: from 0.663 to 0.396, II0: from 0.677 to 0.0.401). The first local areas surrounding both extraction zones (I1 and II1) observe a reduction in spread of about 24% and 26% (I1: from 0.704 to 0.533, II1: from 0.724 to 0.539). The spread in local areas I2 and II2 reduces by about 15% and 16% (I2: from 0.696 to 0.591, II2: from 0.720 to 0.607). A reduction of 9% and 10% is observed in local areas I3 and II3 (I3: from 0.687 to 0.624, II3: from 0.714 to 0.646). The reported spread values indicate that the uncertainty gradually increases as a function of the distance from the extraction zones. The initial bias does not seem to affect the posterior spread values. Further notice the similarities between the reduction of RMSE in mining zone I and the overall spread reduction.

In order to conclude the discussion of the experimental results, global and local empirical probability curves are presented and analysed. For each block included in the analysis, empirical quantiles are computed according to a predefined set of theoretical probabilities (e.g. 0.025 to 0.975). For a given theoretical probability, corresponding block quantiles are retrieved and compared against the reference values. The proportion of reference values not exceeding the block quantiles is referred to as the empirical probability. The procedure is repeated for all predefined theoretical probabilities. The interested reader is referred to Olea (2012). A good uncertainty model would results in probability pairs located along the bisector. This would for example imply that for 10% of the investigated blocks (empirical probability) the reference value is lower than the corresponding 10% quantile (the 10% here refers to the theoretical probability).

Figure 4.18a displays the global empirical probability curves computed from prior and updates realizations. The curve confirms a global good agreement between deviations (error) and uncertainty (spread) in the prior model (due to the medium dense sampling grid). Updating results in an even better alignment of the global probability curve with the bisector.

Figure 4.18b, 4.18c and 4.18d all display two prior and two posterior local probability curves computed from blocks in respectively local areas I0 and II0, local areas I2 and II2 and local areas I4 and II4 (Fig. 4.14). The figures illustrate that the regional deviations in mining zone II are not accounted for in the prior uncertainty model (purple triangles). The occurrence of low values is severely overestimated. For example in only 6% of the blocks in area I0 , is the reference value below the computed 40% quantile value. The prior probability curves for areas II0, II2 and II4 indicate that the block distributions in and around extraction zone II are centred around values which are too low (purple triangles, Fig. 4.18b, 4.18c and 4.18d). Updating causes a significant upward shift of the corresponding distributions (orange pentagons, Fig. 4.18b, 4.18c and 4.18d). However, the posterior empirical probability curves indicate that the upward shift could have been slightly larger. Figure 4.16b confirms this conclusion. A slightly higher position of the boxplots, representing local areas II0 to II4, would result in reference values located inside the 50% confidence interval. The prior empirical probability curves of areas I0,

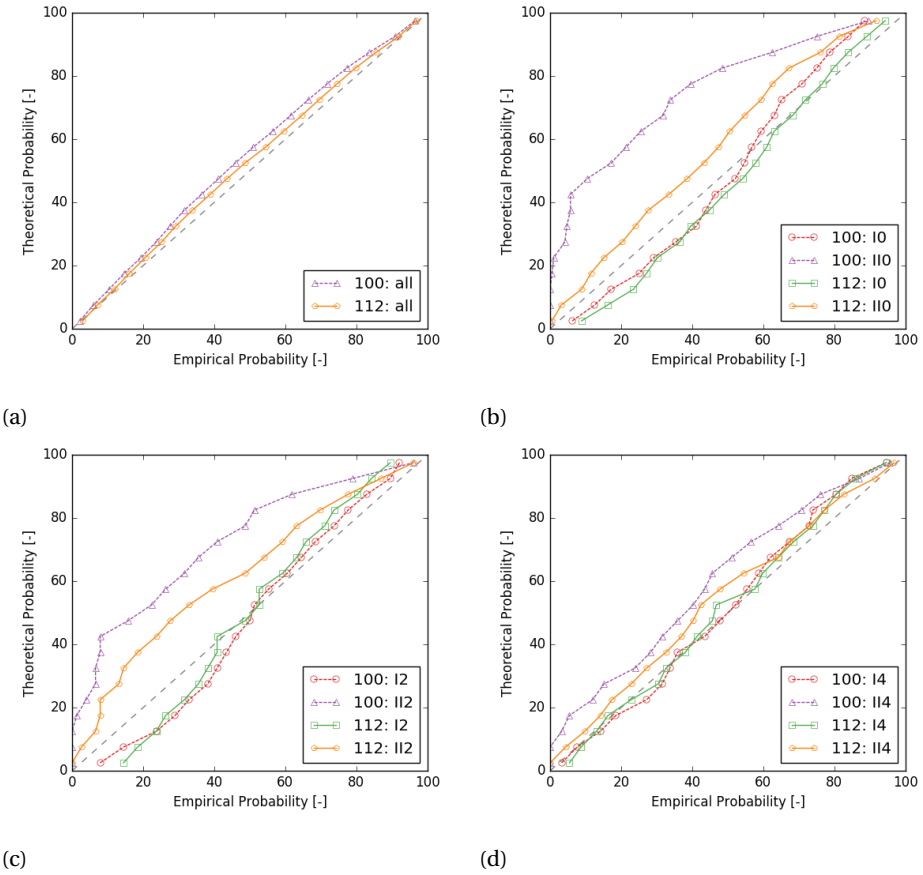

Figure 4.18: Empirical versus theoretical probabilities computed over prior and updates realizations - (a) entire grid (b) local area I0 and II0 (c) local areas I2 and II2 (d) local areas I4 and II4

I2 and I4 approve the overall quality of the prior uncertainty model around mining zone I (red circles, Fig. 4.18b, 4.18c and 4.18d). A more detailed analysis reveals that the tails of the distributions are a bit too long. Such a conclusion can be drawn from the fact that the occurrence of low values and high values is slightly overestimated ('S' shape instead of straight line). Updating generally results in a better alignment of the posterior curves with the bisector (green squares, Fig. 4.18b, 4.18c and 4.18d). This corresponds to shorter confidence interval, as observed in Fig. 4.16b.

4.3. CONCLUSIONS

An artificial experiment is conducted to demonstrate the capabilities of the developed algorithm. A mining operation with two extraction points of unequal production rate is simulated. The resulting material streams are blended and intentionally inaccurate observations are made. A total of 12 updates is performed, each one based on 2

inaccurate observations of 16 blended blocks. The performance of the algorithm is evaluated using various criteria.

(1) A visual inspection indicates that the mean field continuously improves in and around both inspection zones. Occurrences of severe local estimation bias are generally corrected. Only when individual block deviations were averaged out in the blend (large positive deviations associated with small digging blocks are cancelled against smaller negative deviations of larger digging blocks), did the local deviation go undetected. The assimilation of sensor observations also resulted in an adjustment of the uncertainty model. The confidence interval in the cross sections and the standard deviation fields show that the remaining uncertainty is not uniformly distributed across the excavated areas. Regions who contribute more material to the blend are estimated with a higher confidence (lower SD, smaller confidence interval). The adjustment of the uncertainty model extends beyond the already excavated areas.

(2) The global statistics demonstrate how the overall quality of the model evolves as production data are being assimilated. Over the course of the experiment, the global $RMSE$ drops by about 11%. The displayed evolution is somewhat volatile, reflecting the relative importance of a few local improvements. The spread, on the other hand, decreases with a constant rate. After all, the information value of similar observations do not change over time.

(3) An assessment of local statistics provides a better comprehension of the effects of an updating sequence. Prior to updating, the quality of the model is largely location dependent. Local biases do exist and the confidence interval only partially accounts for the observed deviations. As a result, certain areas are more drastically updated than others. The modifications in mining zone I are limited to a correction of the uncertainty model (to account for the assimilated observations). Mining zone II, on the other hand, experiences a significant alteration to correct the local estimation bias. The $RMSE$ in zone I and II drops with 38% and 45% respectively. Improvements of 24%, 17% and 13% are observed within the first three local areas surrounding zone I. The corresponding values for zone II amount to 39%, 27% and 19%.

The next chapter focusses on the influence of the monitoring network on the performance of the algorithm. A series of experiments is conducted to measure the effects of different sensor precisions, measurement volumes, update intervals and blending ratios. The results the study is used to formulate recommendations for designing an optimal monitoring network, guaranteeing optimal algorithm performance. The author expects that for a given blending ratio the order of extraction does not significantly impact the final results. This assumption has not been proven and should be verified.

It is also important to understand the limitations of the algorithm regarding the geological environment and the prior model hereof. The question still remains whether good performance can be achieved when the prior model and the reality differ substantially. Also the effects of different spatial correlation structures is not yet understood (different covariance models, ranges, nugget, ...). Another possible point of attention could be the further improvement of the localization error correction option. The current implementation (Gaspari-Cohn correlation function) results in a differential covariance correction of blocks inside the same digging area. Blocks near the extraction point receive a lower correction factor than blocks slightly further away. Especially when

considering large measurement volumes (digging blocks), a tapering function with a plateau might be more suitable. The function should be defined such that all blocks inside the digging area are not corrected (correction factor of 1) and that tapering starts from the border of the digging areas.

REFERENCES

T. Wambeke and J. Benndorf, *A simulation-based geostatistical approach to real-time reconciliation of the grade control model,* Mathematical Geosciences **49**, 1–37 (2017).

R. Olea, *Building on crossvalidation for increasing the quality of geostatistical modeling,* Stochastic Environmental Research and Risk Assessment **26**, 73–82 (2012).

5

SENSITIVITY ANALYSIS

This chapter studies how the design and operational control of the monitoring system influences the overall performance of the updating algorithm. A total of 125 experiments are conducted to quantify the effects of variations in measurement volume, blending ratio and sensor precision. The following questions are addressed to compare behaviour across experiments. (i) Does the resource model improve over time? (ii) Does the accuracy of the predicted system output improve over time? (iii) Does the predicted system output in the next $24h$ *improve after updating? Based on the outcome, recommendations are formulated to optimally design and operate the monitoring system, guaranteeing the best possible algorithm performance.*

Parts of this chapter have been submitted for publication in Mathematical Geosciences **accepted with revisions** (Wambeke and Benndorf, 2017a).

The previous chapter illustrates the potential of the developed real-time updating algorithm. In order to apply the concept in practice, a set of estimates $\mathscr{A}_t(\mathbf{Z}_{t-1}(:,i))$ needs to be generated and the corresponding measurements $\mathbf{d}_t$ are to be collected. At the beginning of each time interval $[t-1,t]$, a set of estimates $\mathscr{A}_t(\mathbf{Z}_{t-1}(:,i))$ is obtained by individually propagating the GC model realizations through a forward simulator $\mathscr{A}_t$ (a simulation based approach to characterize uncertainty). The forward simulator is a virtual model of the actual operation and describes which GC blocks are extracted, processed and measured between $t-1$ and t. Based on the provided GC model realizations, the forward simulator provides a set of estimates for the upcoming actual observation $\mathbf{d}_t$. This forward prediction step is essential in linking a specific observation $\mathbf{d}_t$ to its constituent GC blocks.

Figure 5.1 illustrates how actual observations $\mathbf{d}_t$ are collected. The figure shows an example of a monitoring system in an open pit mining operation. In this example, material is extracted from two benches with different local production rates. After hauling, material is tipped into a comminution circuit. Inside the comminution circuit, a sensor continuously records the properties of the crushed material. As soon as a time period ends, an average measurement $\mathbf{d}_t$ is computed based on the recorded sensor response. This average measurement characterizes all material that passed the sensor between $t-1$ and t (Fig. 5.1). The algorithm subsequently updates the GC model based on the differences between the estimates and the actual measurement (Eq. 2.5).

The performance of the algorithm directly depends on following three system parameters. (i) The measurement volume V_t equals the amount of material that has been characterized between $t-1$ and t. The measurement volume is controlled by adjusting the interval duration ($[t-1,t]$) and the local production rates (Fig. 5.1). (ii) The blending ratio $R_{t,X}$, formulated as a percentage, defines how much material in the measurement volume originates from Bench X. The blending ratios are determined by the local production rates. (iii) The measurement error E_t represents the accuracy of the applied sensor. Once installed, the sensor accuracy can hardly be adjusted. The objective of this paper is to study the influence of these three system parameters on the overall performance of the algorithm. A total of 125 experiments are conducted to derive some general recommendations to optimally operate the monitoring system.

5.1. METHODOLOGY

5.1.1. GENERAL SETUP

The 125 experiments represent a mining operation with two extraction zones located in different benches. Material streams originating from both benches are blended and inaccurate observations are made (Fig. 5.1).

The experiments are conducted using the following boundary conditions; (a) the true but unknown geological reality does not change across the experiments, (b) each experiment starts with the same prior model of this geological reality, (c) the global production rate is assumed constant across all timesteps and experiments, (d) the duration of each experiment is the same.

Albeit exactly the same material is excavated, processed and measured during the course of each experiment, the configuration of the monitoring system significantly

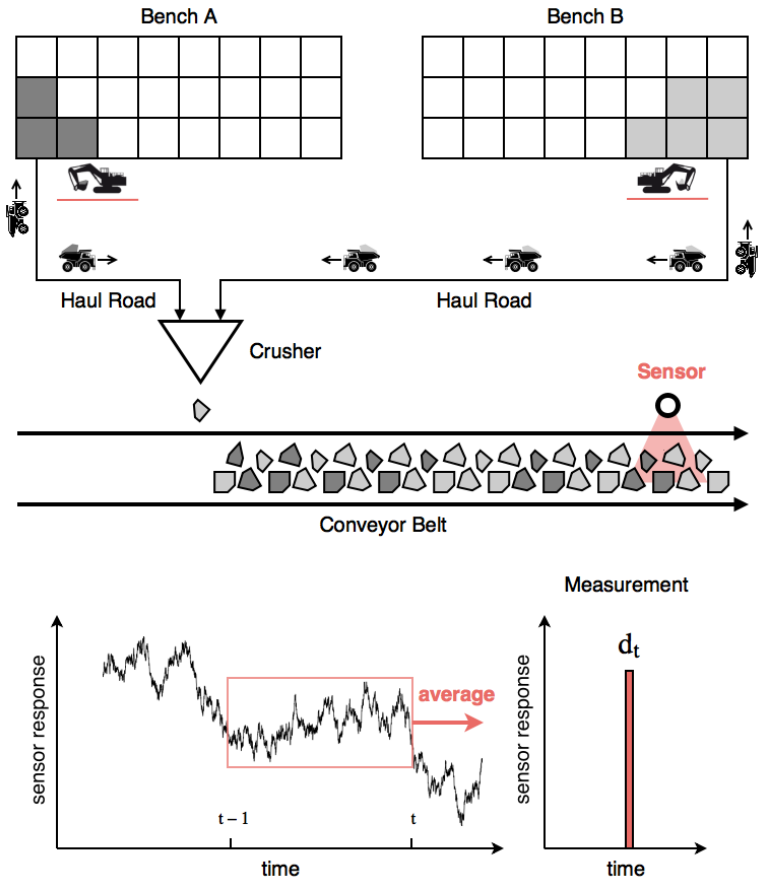

Figure 5.1: Monitoring setup in an open pit mining operation - the coloured blocks refer to material that has been extracted during the studied time interval ($t-1$ to t)

influences the end results. From this point forward, the measurement volume V_t is expressed in number of GC blocks. Measurement volumes of 8 (1), 16 (2), 32 (4), 48 (6) and 64 (8) GC blocks are considered. The values between brackets refer to corresponding time units. Note, the production rate remains constant at 8 GC blocks per time unit.

The blending ratio $R_{t,A}$ defines how many of the blocks in the measurement volume originate from bench A. Consequently, the remaining blocks originate from bench B. For example, a blending ratio of 75% and a measurement volume of 48 blocks prescribe that each single blended measurement is composed of 36 blocks from bench A (N_A) and 12 blocks from bench B (N_B). Blending ratios of 100%, 87.5%, 75%, 62.5% and 50% are considered.

The measurement error E_t, defined as the standard deviation of a zero mean normal distribution, represents the precision of the applied sensor. Measurement errors of 0.05, 0.25, 0.5, 0.75 and 1.00 are considered. The reported standard deviations define the measurement error on the scale of a single GC block (discussed later on).

The following presents more details regarding the construction and execution of the experiments. Units of the modelled attributes are omitted to stress their synthetic nature.

5.1.2. GEOLOGY

TRUE BUT UNKNOWN FIELD

During the course of an experiment, a true but unknown geological state is sampled twice. First, small point samples are collected on both benches prior to the construction of the GC model (exploration phase). Second, observations are made on a blend of multiple GC blocks, which do not necessarily originate from the same bench (operational phase). To correctly account for both scales of support, two representations of the true but unknown state are constructed. Both the high-resolution point and lower-resolution block reference fields characterize the true but unknown geological state in two benches of $300m$ x $300m$ using a different level of discretization.

The high resolution point representation is constructed on two discretized grids, each with 300 x 300 cells of size $1m$ x $1m$. The origins of bench A and B are located in ($0m$, $0m$) and ($1000m$, $1000m$) respectively. The two reference fields are generated using a random field simulation algorithm (sequential Gaussian simulation). Attribute values are drawn from a standard normal distribution (centred around 0). The correlation between attribute locations is described using an isotropic exponential function with a variance of 1.0 and a range of $100m$. Each attribute value thus reflects the direction (-/+) and magnitude of a deviation from a global mean.

Due to field dimensions which are relatively small compared to the range of the covariance model ($300m$ versus $100m$), four times larger simulation grids are used (600 x 600 cells of size $1m$ x $1m$). Eventually only the south-west quarter of the grids are retained. The resulting field in bench A has an average of -0.145 (-0.038) and a variance of 0.728 (1.025). The 5% and 95% quantiles are -1.500 (-1.709) and 1.271 (1.618). The values between brackets refer to those computed from the larger simulation grid. The resulting field in bench B has an average of -0432 (-0.079) and a variance of 1.089 (0.987). The 5% and 95% quantiles are -2.130 (-1.703) and 1.266 (1.569).

The low-resolution block representation is constructed on two discretized grids, each

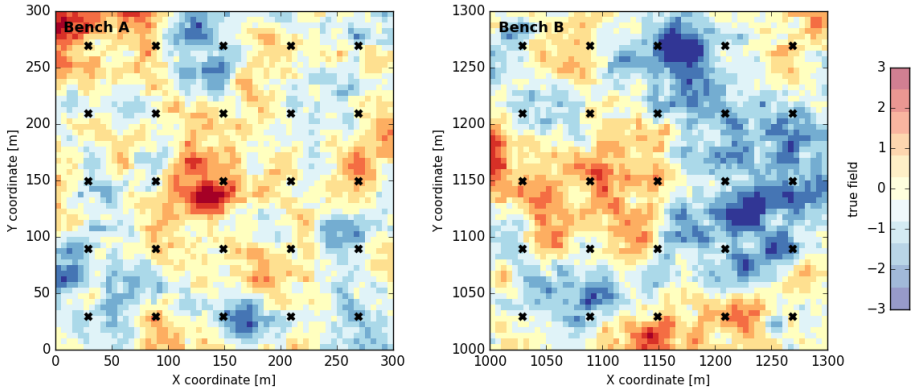

Figure 5.2: Block representation of true but unknown geological state in bench A (left) and bench B (right) - 60 x 60 blocks of size $5m$ x $5m$ colour coded according to their average block value - black crosses indicate the locations of the sampling points collected on a $60m$ x $60m$ exploration grid - point samples have a support of $1m$ x $1m$

with 60 x 60 cells of size $5m$ x $5m$. This second set of reference fields is obtained from reblocking the previous point fields. Values of 5 x 5 $1m$ x $1m$ cells are averaged and assigned to larger $5m$ x $5m$ blocks. The resulting block reference field in bench A has an average of -0.145 (unchanged) and a variance of 0.659. The 5% and 95% quantiles are -1.444 and 1.181. The resulting block reference field in bench B has an average of -0.432 (unchanged) and a variance of 1.012. The 5% and 95% quantiles are -2.068 and 1.209. As expected, the variability reduces due to reblocking. The final block reference fields of bench A and B are displayed in Fig. 5.2.

PRIOR MODEL

Two sets of field realizations are used throughout to characterize the spatial variability and quantify the geological uncertainty of the study environment in bench A and B. Prior to the operation, the study prescribes the availability of exploration data collected on a $60m$ x $60m$ sampling pattern. This sampling density results in two prior sets of realizations with a realistic level of uncertainty and error.

Two prior sets of 100 realizations are generated on discretized grids with 600 x 600 cells of size $1m$ x $1m$ using a sequential Gaussian simulation algorithm (same as before). Again, the simulation grids are four times larger than the studied areas for reasons of small field dimensions compared to the range of the covariance model. A total of 200 sampling points are provided as conditioning data (100 per grid, sampled on a $60m$ x $60m$ pattern from the previous point simulations on the large 600 x 600 grids with $1m$ x $1m$ cells). From the 200, only 50 sampling points are actually located inside one of both study areas (Fig. 5.2). The covariance model of the true state is assumed to be known and was not inferred from the sampling points (isotropic exponential function with a variance of 1.0 and a range of $100m$). The effect of an incorrectly inferred covariance model is discussed in Sections 5.7 (last paragraph) and 7.2.2 (second paragraph). Once simulated only the south-west quarter of grid A and grid B are retained (300 x 300 cells

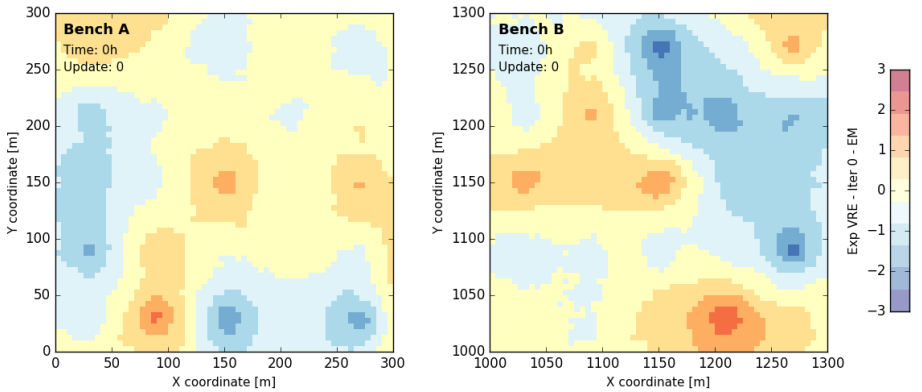

Figure 5.3: Bench A (left) and bench B (right) with 60 x 60 blocks of $5m$ x $5m$ colour coded according to computed mean field (EM)

of size $1m$ x $1m$). Subsequently, the 100 point simulation are reblocked onto a grid with 60 x 60 blocks of size $5m$ x $5m$.

Only the block realizations are further considered (one reference field $\mathbf{Z}^*$ and 100 prior block realizations $\mathbf{Z}_0(:,i)$ per bench). All previous point simulations were just an aid to this end.

Figure 5.3 displays the mean field (EM) computed from the 100 prior realizations. The mean fields of bench A and B are considerably smoother than the true state (compare Figs. 5.2 and 5.3). The general larger-scale patterns are nevertheless reproduced.

5.1.3. EXPERIMENTAL SCENARIOS

From this point forward, it will be assumed that the previously constructed block realizations have a thickness of $5m$. Considering a rock density of $2.8g/t$, it is not unreasonable to assume that each $125m^3$ GC block is transported by a truck with a payload capacity of $350t$. Hence the terms GC block and truck are used interchangeably.

During each experiment, a total of $403200t$ is blasted, excavated, processed and monitored. The corresponding 1152 GC blocks originate from two possible extraction zones in bench A or B (Fig. 5.4). Each experiment terminates after $114h$. This is the time needed to remove all GC blocks from the designated extraction zones assuming a constant global production rate of $2800t/h$ (8 GC blocks per time unit). For the sake of simplicity, blending and movement due to blasting are not being considered in this chapter (blast movements are accounted for in Chapter 6). Consequently, the position of a certain excavated volume corresponds directly to the original voume in the block model (no corrections have to be applied). The blending referred to here, in the this chapter, solely results from mixing material from two production areas.

COMPOSITION MEASUREMENT VOLUME

The extraction zones are further subdivided into 36 $20m$ x $20m$ digging blocks. The numbers in Fig. 5.4 indicate the order in which digging blocks are extracted from a bench. The same sequential pattern is used to extract 16 $5m$ x $5m$ x $5m$ GC blocks from

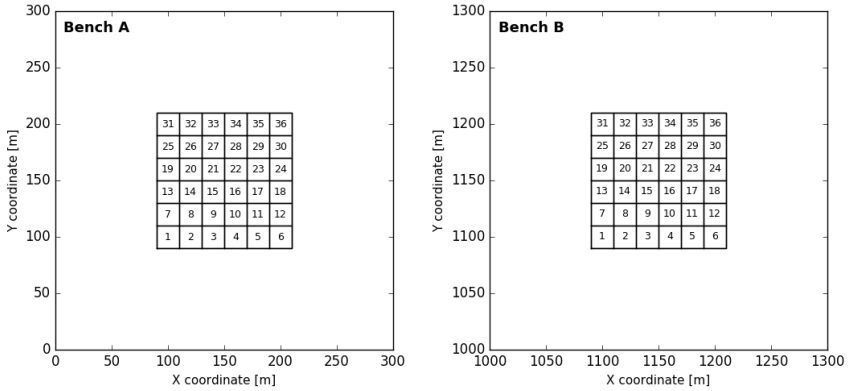

Figure 5.4: Extraction zones (120m x 120m x 5m) and digging blocks (20m x 20m x 5m) in bench A (left) and bench B (right) - the numbers indicate the order in which digging blocks are extracted from a bench

each 20m x 20m x 5m digging block (move a row from bottom to top, move within a row from left to right). Both extraction patterns are combined to construct two nested queues (one for each bench). Each nested queue uniquely defines the order in which GC blocks are removed from a particular bench. The order of extracting GC blocks from each bench does not change across the 125 experiments.

The specific time when a GC block is extracted, processed and measured does change across experiments. Algorithm 2 (p. 68) fully describes the composition of the blend at discrete timesteps given a measurement volume V_t and a blending ratio $R_{t,A}$. A blend refers to a collection of GC blocks passing the sensor between $t-1$ and t. Note that the blending ratio $R_{t,A}$ is changed to $100 - R_{t,A}$ when 576 GC blocks are extracted (lines 19-20, Algm. 2, p. 68). This to ensure that both extraction zones are fully depleted at the end of each experiment.

Figures 5.5, 5.6 and 5.7 display the composition of the blend at 36, 42 and 48 hours given three different blending ratios (75%, 62.5% and 50%) and a measurement volume of 48 GC blocks. The complexity of the corresponding mining schedules differs significantly across the various blending scenarios (compare Fig. 5.5 and 5.7). The digging outlines defined by blending scenario 43E (Fig. 5.5) do not immediately resemble an obvious pattern. Occasionally, scheduled GC blocks are even dispersed across the bench (Fig. 5.5a). In contrast, blending scenario 54E consistently yields two L-shaped digging outlines.

Table 5.1 provides an overview of the 25 different blending scenarios. Note that the measurement volume determines the time between updates and therefore also the total number of possible updates within each experiment.

5.1.4. FORWARD SIMULATOR AND SENSOR RESPONSE

Based on a defined blending scenario, a forward simulator $\mathscr{A}_t$ is implemented. The forward simulator is constructed based on the assumption that an actual measurement $\mathbf{d}_t$ would result from a time-averaged sensor response characterizing the composition of

Algorithm 2 Blending algorithm

```
1:  procedure BLEND(V_t, R_{t,A})
2:      Queue_A; Queue_B ← NESTEDQUEUES
3:      Dig_A; Dig_B ← POP(Queue_A); POP(Queue_B)
4:      N_A; N_B ← V_t * R_{t,A}; V_t * (100 − R_{t,A})
5:      N_tot ← 0
6:      while N_tot ≠ 1152 do
7:          Blend ← EMPTYBLEND
8:          for i ← 1, N_A do
9:              if LEN(Dig_A) = 0 then
10:                 Dig_A ← POP(Queue_A)
11:             Block ← POP(Dig_A)
12:             Blend ← APPEND(Blend, Block)
13:         for j ← 1, N_B do
14:             if LEN(Dig_B) = 0 then
15:                 Dig_B ← POP(Queue_B)
16:             Block ← POP(Dig_B)
17:             Blend ← APPEND(Blend, Block)
18:         N_tot ← N_tot + V_t
19:         if N_tot = 1152/2 then
20:             N_A, N_B ← N_B, N_A
21:         WRITEBLOCKIDS(Blend)
```

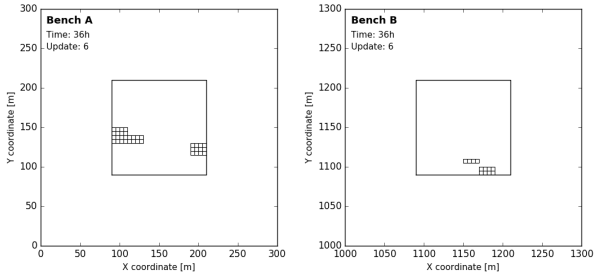

(a) 36h

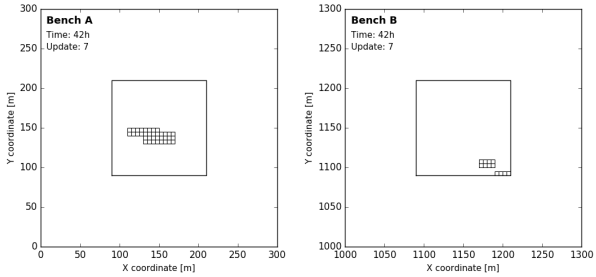

(b) 42h

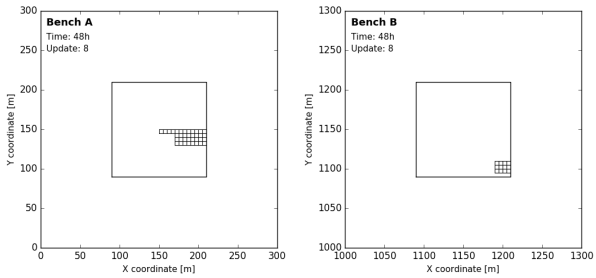

(c) 48h

Figure 5.5: GC blocks from bench A (left) and bench B (right) in blend at 36 (a), 42 (b) and 48 (c) hours - measurement volume of 48 blocks and blending ratio of 75% - blending scenario 43E

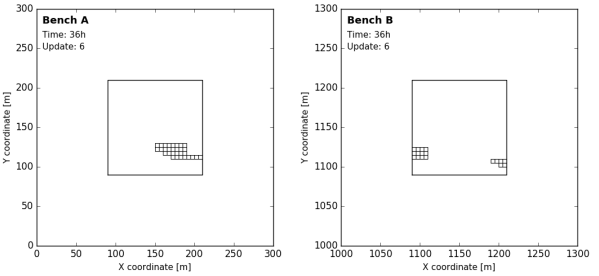

(a) 36h

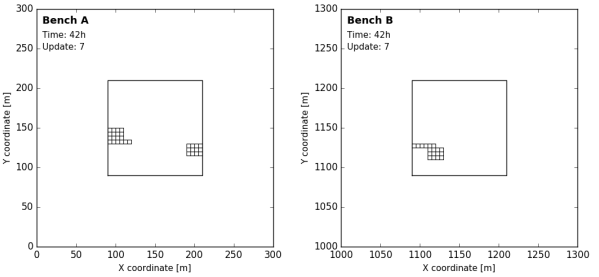

(b) 42h

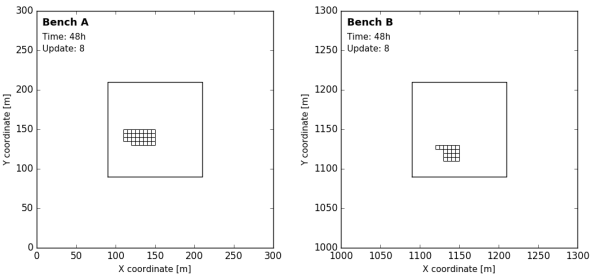

(c) 48h

Figure 5.6: GC blocks from bench A (left) and bench B (right) in blend at 36 (a), 42 (b) and 48 (c) hours - measurement volume of 48 blocks and blending ratio of 62.5% - blending scenario 44E

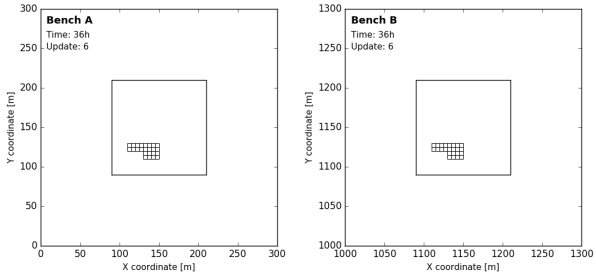

(a) 36h

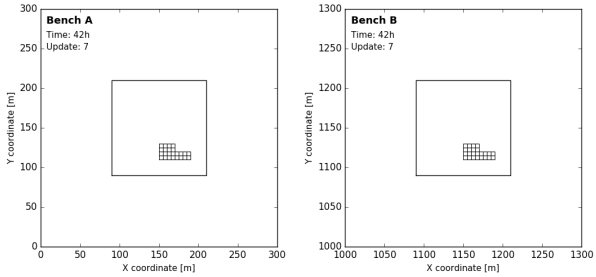

(b) 42h

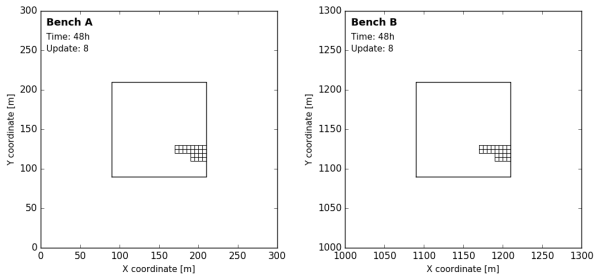

(c) 48h

Figure 5.7: GC blocks from bench A (left) and bench B (right) in blend at 36 (a), 42 (b) and 48 (c) hours - measurement volume of 48 blocks and blending ratio of 50% - blending scenario 45E

Table 5.1: Blending scenarios defined by measurement volumes V_t and blending ratios $R_{t,A}$ - $Upd.$ is number of updates within an experiment, dt is time between updates - N_a and N_b are the number of blocks from bench A and bench B constituting a blended measurement - the E in the experiment code refers to an unassigned measurement error

Exp.	V[#]	$Upd.$[#]	dt[h]	R[%]	N_A[#]	N_B[#]
11E	8	144	1	100	8	0
21E	16	72	2	100	16	0
31E	32	36	4	100	32	0
41E	48	24	6	100	48	0
51E	64	18	8	100	64	0
12E	8	144	1	87.5	7	1
22E	16	72	2	87.5	14	2
32E	32	36	4	87.5	28	4
42E	48	24	6	87.5	42	6
52E	64	18	8	87.5	56	8
13E	8	144	1	75	6	2
23E	16	72	2	75	12	4
33E	32	36	4	75	24	8
43E	48	24	6	75	36	12
53E	64	18	8	75	48	16
14E	8	144	1	62.5	5	3
24E	16	72	2	62.5	10	6
34E	32	36	4	62.5	20	12
44E	48	24	6	62.5	30	18
54E	64	18	8	62.5	40	24
15E	8	144	1	50	4	4
25E	16	72	2	50	8	8
35E	32	36	4	50	16	16
45E	48	24	6	50	24	24
55E	64	18	8	50	32	32

the blend. To mimic this behaviour, the forward simulator translates a set of realizations $\mathbf{Z}_{t-1}(:,i)$ into a set of estimated observations $\mathscr{A}_t(\mathbf{Z}_{t-1}(:,i))$ by averaging a selection of GC block values. The relevant blocks are selected with the aid of Algm. 2 (p. 68). The values to be averaged are obtained from the relevant GC model realization. The time between consecutive runs (time between $t-1$ and t) is determined by the measurement volume (Tab. 5.1).

Due to the synthetic nature of the experiment, a second forward simulator $\mathscr{A}_t^*$ is built to mimic the behaviour of a real monitoring network. This simulator generates inaccurate but true observations $\mathbf{d}_t$ based on the reference field $\mathbf{Z}^*$ ($\mathbf{d}_t = \mathscr{A}_t^*(\mathbf{Z}^*)$). The previous constructed simulator $\mathscr{A}_t$ is adjusted to suit this purpose. Two main modifications are made. First, the reference field $\mathbf{Z}^*$ is propagated through the simulator and not one of the realizations $\mathbf{Z}_{t-1}(:,i)$. Second, random noise is added onto the single GC block responses (individual GC block values). The random noise is drawn from a zero mean normal distribution with a standard deviation E_t, previously referred to as the measurement error (in order to compare results across experiments, the same random seed is used throughout). Once the noise is added onto the sensor response (GC block value), a time-averaged measurement $\mathbf{d}_t$ is computed. For further reference, the accuracy of the time-averaged measurement amounts to $E_t/\sqrt{V_t}$ (V_t, expressed in number of GC blocks).

Given the reference field, the prior set of realizations and both forward simulators, a total of 125 experiments are conducted based on different V-R-E combinations (5x5x5). Each experiment is given a unique three digit code. The first two digits refer to the blending scenario as indicated in Tab. 5.1. The third digit indicates the precision of the applied sensor (1 for an E of 0.05, 5 for an E of 1.00).

5.2. ASSESSMENT STATISTICS

The raw data of the 125 experiments is further processed. An assessment statistic, the root mean square error (RMSE), is computed on several subsets of data to address the following questions. (1) Does the GC model improve over time? (2) Does the accuracy of all predicted measurements improve over time? (3) Do predicted measurements in the next $24h$ improve after updating? Based on these questions, the behaviour across experiments is compared (next section). This section illustrates, based on results from a single experiment, how the necessary assessment statistics are calculated.

The computed assessment statistics are subdivided into two groups. (a) Field statistics are derived from spatial GC models, and are used to answer the first question. (b) Production statistics are computed based on predicted measurements, and used to address the second and third question. Experiment 223 is used as an example. As a reminder, experiment 223 is based on a measurement volume of 16 GC blocks ($5600t$), a blending ratio of 87.5% and a measurement error of 0.50.

5.2.1. FIELD STATISTICS

Figure 5.8 displays how the mean fields of bench A and B change through time when assimilating time-averaged sensor responses. The 16 smaller rectangles demarcate the GC blocks which were extracted, blended and measured during the considered time

interval. A comparison of Figs. 5.8a to 5.8h with the prior mean field (Fig. 5.3) and the true state (Fig. 5.2) indicates that the mean field (the current best estimate) continuously improves in and around both extraction zones.

It is hardly feasible to visually compare the evolution of spatial fields across a multitude of experiments. Instead a more quantitative measure such as the RMSE is computed. The $RMSE$ describes the overall deviation at time t between true ($\mathbf{Z}^*(n)$) and estimated block values ($E[\mathbf{Z}_t(n,:)]$):

$$RMSE_t = \sqrt{\frac{1}{N}\sum_{n=1}^{N}\left(\mathbf{Z}^*(n) - E[\mathbf{Z}_t(n,:)]\right)^2}. \tag{5.1}$$

Two time series of $RMSE$ are computed, one for each extraction zone. In this case, n only iterates over the 576 blocks located inside one of both studied extraction zones (Fig. 5.4). Both times series are further processed. The change in $RMSE$ quantifies obtained improvements (positive) or deteriorations (negative) relative to time 0:

$$\Delta RMSE_t = RMSE_0 - RMSE_t \tag{5.2}$$

5

Ideally, the $\Delta RMSE$ increases as more measurements are incorporated. Figure 5.9 displays the evolution of the $\Delta RMSE$ through time. The $RMSE$ of the prior model inside the extraction zones of bench A and B amounts to 0.718 and 0.781 respectively. Over the course of the experiment, the RMSE inside both extraction zones drops by about 45% (45.34% and 44.19%).

Figure 5.9 indicates that the $RMSE$ decreases approximately at a semi-constant rate during the first ($0h$ to $72h$) and second half ($72h$ to $144h$) of the experiment. The inflection point of both $\Delta RMSE$ curves lies at $72h$, i.e. when the blending ratio is reversed (lines 19-20, Algm. 2, p. 68). The blending scenario clearly influences the rate of reduction of the $RMSE$ in a bench. The period of steep incline (faster reduction in $RMSE$) coincides with a larger local production rate ($2450t/h$, 14 GC blocks per $2h$). A lower local production rate ($350t/h$, 2 GC blocks per $2h$) results in a more gentle incline. Note that the global production rate remains constant across the experiment ($2800t/h$, 8 GC blocks per h).

The somewhat volatile behavior results from a larger impact of a few local improvements (Fig. 5.9). For example, the exceptionally large reduction between $28h$ and $34h$ at bench A originates from a correction of the anomalous central area (compare Figs. 5.8a and 5.8b). A similar exceptionally large reduction in $RMSE$ occurs in Bench B between $72h$ and $74h$ (Fig. 5.9, decreasing mean field near $x = 1175m$ and $y = 1125m$, compare Figs. 5.8d and 5.8e).

5.2.2. PRODUCTION STATISTICS

Figs. 5.10a to 5.10i display predicted measurements $\mathscr{A}_t(\mathbf{Z}_{t-1}(:,i))$ (blue boxplots) computed from the most recent GC model at nine different timesteps (refer to time label in top left corner). The figures further show time-averaged sensor responses ($\mathbf{d}_t$, green cross) and true but unknown averages (red dots). Each plot displays the average composition of 72 $2h$ blends (16 GC blocks, $5600t$ per blend). Some blends are already processed and measured (grey area), one is currently being characterized (yellow bar)

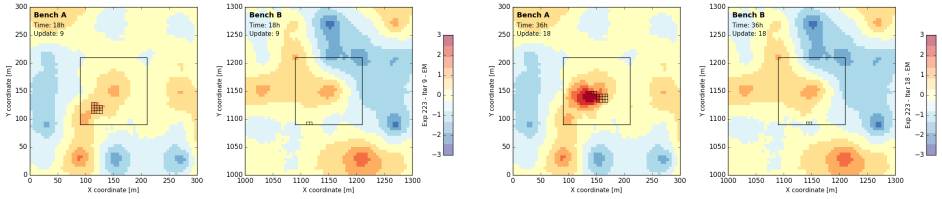

(a) $18h$ (b) $36h$

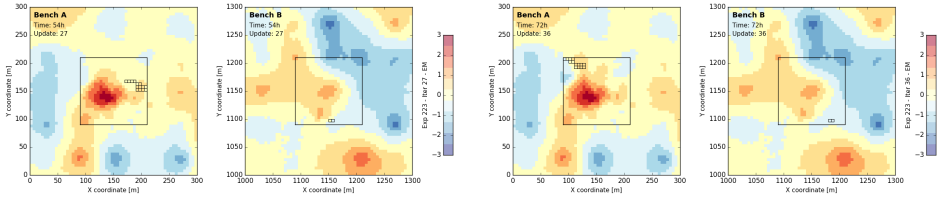

(c) $54h$ (d) $72h$

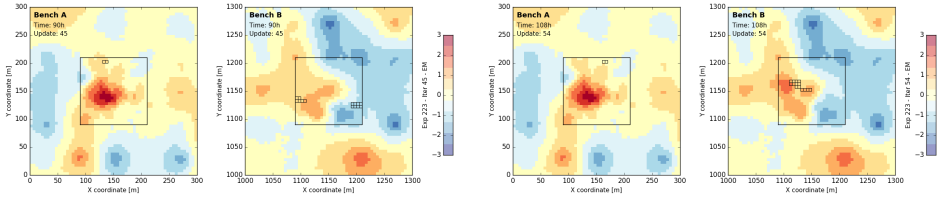

(e) $90h$ (f) $108h$

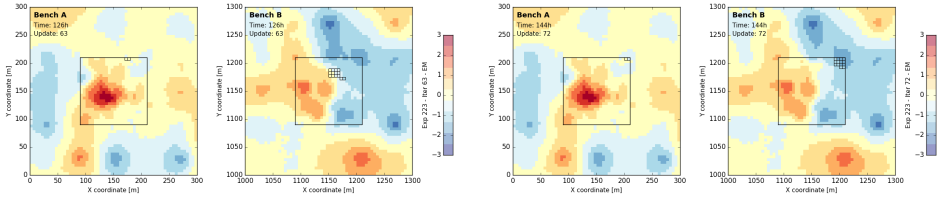

(g) $126h$ (h) $144h$

Figure 5.8: Mean field in bench A (left) and bench B (right) at $18h$ (a), $36h$ (b), $54h$ (c), $72h$ (d), $90h$ (e), $108h$ (f), $126h$ (g) and $144h$ (h) - experiment 223, measurement volume of 16 GC blocks ($5600t$), blending ratio of 87.5% and measurement error of 0.5

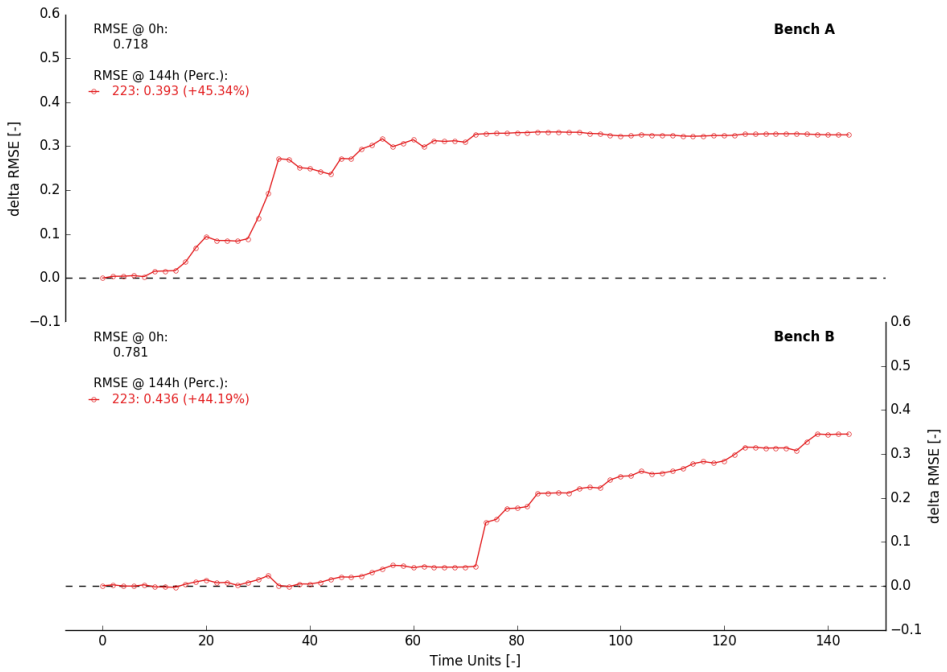

Figure 5.9: Changes in $RMSE$ as a function of time, computed over the extraction zones in bench A (top) and bench B (bottom) - results from experiment 223 are displayed

and others still have to be extracted (right of yellow bar). Once a new time interval ends, a new inaccurate measurement becomes available (green cross in the yellow bar) and an update is performed. Before moving to the next time interval, the forward simulator readjusts all predicted measurements.

In practice, inaccurate measurements are not yet available ahead of the current time interval (no green crosses right of the yellow bar). Moreover, due to sensor inaccuracies, the true averages of the blends (red dots) are never known.

Due to the artificial nature of the experiments, the true averages of the blends are known and production assessment statistics can be computed. The bottom axes in Figs. 5.10a to 5.10i display the differences in absolute error (DAE) between best estimates (blue horizontal lines in boxplots) and true averages (red dots) relative to time 0. For example, a DAE value of 0.5 indicates that the best estimate moved 0.5 units towards the true average. The bars are coloured green in case of improvements and red otherwise. Figures 5.10a to 5.10i show that deteriorations (red bars) are relatively rare and that their magnitudes stay within acceptable limits. The figures further illustrate that GC model updates do not only result in reconciled historic production data (grey area), but also improve future predictions (right of yellow bar).

Two sets of $\Delta RMSE$ curves are computed to describe the evolution of predicted measurements. The first set is based on all available data within a fixed $144h$ time window. That is, both historic and future predictions are jointly considered. The second set describes adjustments in error during the next $24h$ (moving window, purple area right of yellow bar). The calculation proceeds as follows: (a) at time t, construct time interval - either $[0h, 144h]$ or $[t, t+24h]$, (b) select predicted and real measurements within this time interval (c) compute $RMSE$ based on selected data (similar to Eq. 5.1), (d) repeat b-c, but use prior predicted measurements (derived from GC model at time $0h$), (e) compute $\Delta RMSE_t$, (f) go to (a) and increase t to $t+dt$.

The two resulting curves are displayed in Fig. 5.11. The $RMSE$ of all prior predicted measurements ($144h$ window) amounts to 0.546 (Fig. 5.11, top). Over the course of the experiment, the $RMSE$ drops by about 74%. The top graph in Fig. 5.11 indicates that the $RMSE$ of all predicted measurements tends to reduce at a nearly constant rate. The two exceptionally large reductions between $28h-34h$ and $72h-74h$ originate from significant GC model corrections in bench A and B respectively (previously discussed). The cumulative improvement in error mainly results from the increasing number of reconciled historic measurements in the static $144h$ window (Figs. 5.10a to 5.10i, larger grey area). Consequently, as time progresses fewer future predictions influence the assessment statistics.

Advancing a $24h$ time interval while updating yields a different kind of behaviour (Fig. 5.11, bottom). The $\Delta RMSE$ fluctuates around a constant value. The absence of a cumulative increase can be explained by a lack of historic reconciled measurements. At every point in time, the sample set is composed of an equal number of future predicted measurements, all derived from unsampled scheduled GC blocks. The $RMSE$ in future $24h$ windows drops on average by about 12% ($E[\frac{\Delta RMSE_t}{RMSE_0}]$). Improvements in error of over 0.1 do not seem to be uncommon. Between $44h$ and $48h$, the predicted future outlook severely deteriorates ($RMSE$ increases with 0.066, 0.078 and 0.037). Towards the end of the experiment, between $112h$ and $118h$, the $RMSE$ slightly increases between 0.008

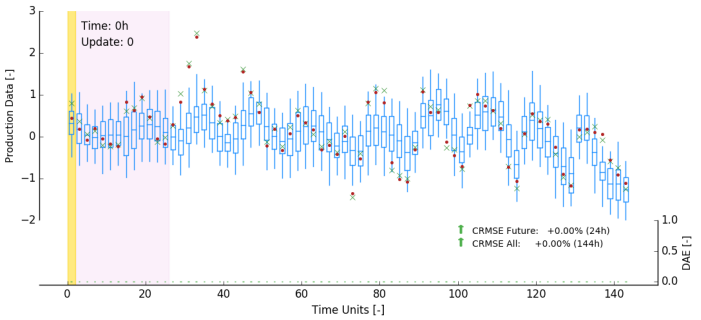

(a) $0h$

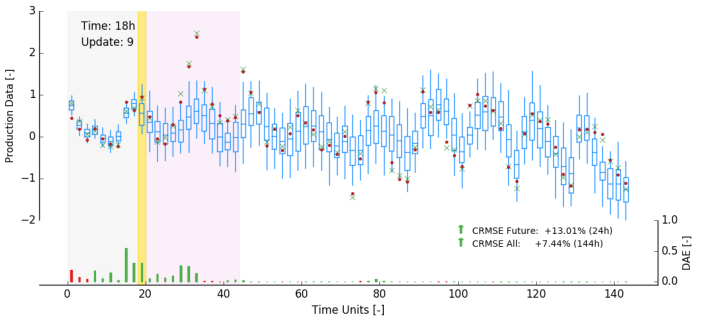

(b) $18h$

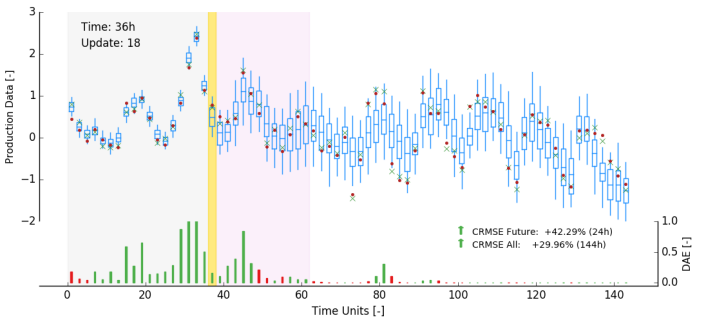

(c) $36h$

Figure 5.10: Predicted measurements (blue), time-averaged inaccurate sensor responses (green cross) and true but unknown averages (red dot) at $0h$ (a), $18h$ (b), $36h$ (c), $54h$ (d), $72h$ (e), $90h$ (f), $108h$ (g), $126h$(h) and $144h$ (i)

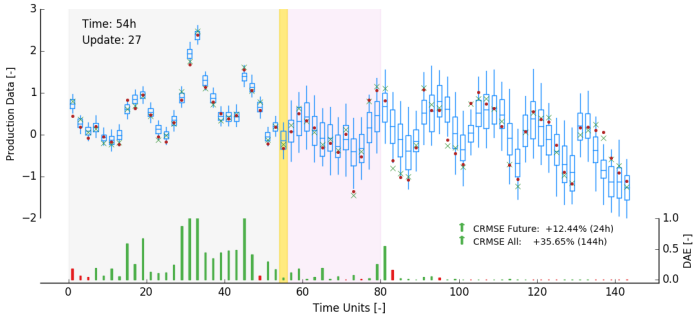

(d) 54h

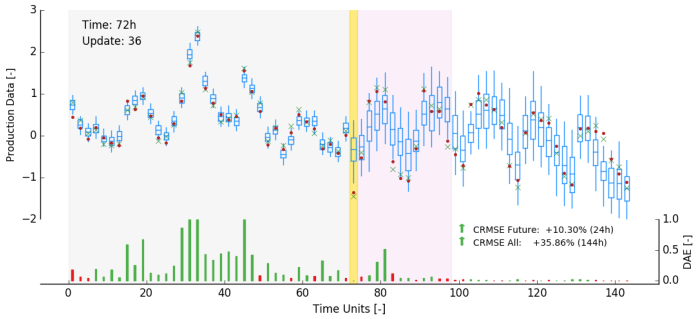

(e) 72h

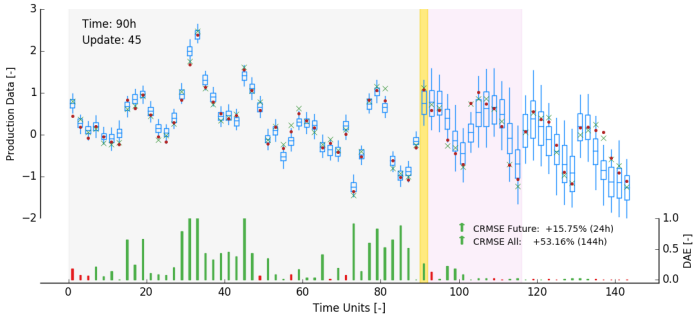

(f) 90h

Figure 5.10: CONTINUED - Predicted measurements (blue), time-averaged inaccurate sensor responses (green cross) and true but unknown averages (red dot) at 0h (a), 18h (b), 36h (c), 54h (d), 72h (e), 90h (f), 108h (g), 126h(h) and 144h (i)

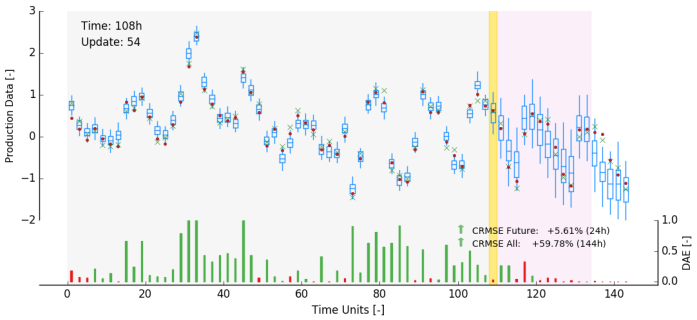

(g) 108h

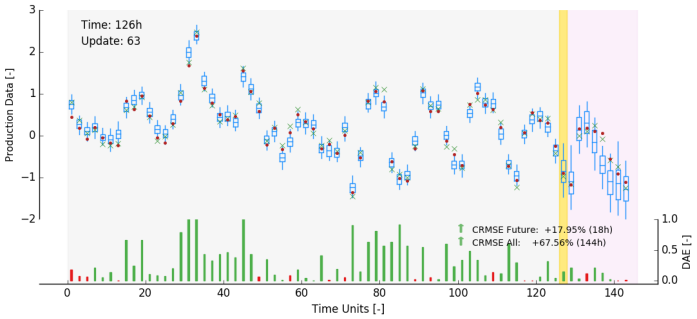

(h) 126h

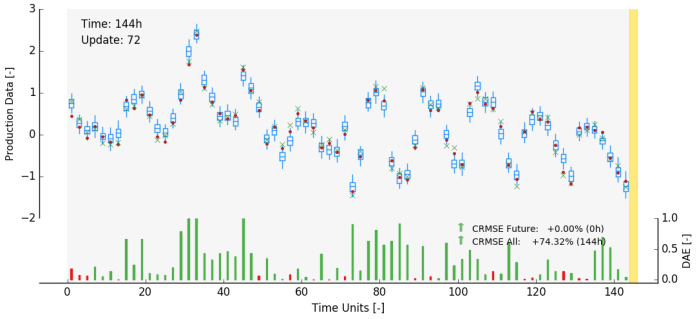

(i) 144h

Figure 5.10: CONTINUED (2) - Predicted measurements (blue), time-averaged inaccurate sensor responses (green cross) and true but unknown averages (red dot) at 0h (a), 18h (b), 36h (c), 54h (d), 72h (e), 90h (f), 108h (g), 126h(h) and 144h (i)

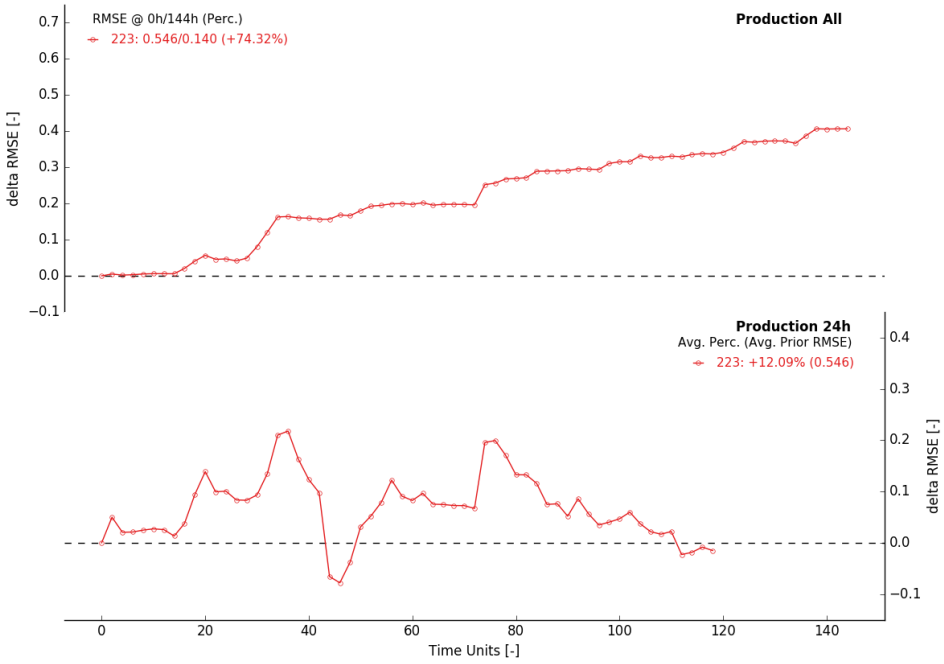

Figure 5.11: Changes in RMSE as a function of time, computed from predicted measurements within fixed 144h time window (top) and a moving 24h future outlook (bottom) - results from experiment 223 are displayed

and 0.019. No results are displayed from $120h$ onwards due to incomplete $24h$ windows. For example, Fig. 5.10h illustrates that at $126h$, only $18h$ of future measurements remain to be predicted.

5.3. RESULTS - SINGLE PARAMETER VARIATION

This section compares behaviour across experiments based on the previously introduced assessment statistics. The first subsection elaborates on the effects of variations in either the measurement volume, blending ratio or measurement error (single parameter variations). The second subsection investigates the impact of a simultaneous variation of two system parameters.

5.3.1. MEASUREMENT VOLUME

Figures 5.12 and 5.13 display results from five selected experiments, all conducted using a blending ratio of 100% and a measurement error of 0.05 (ref. sec. 5.1.1 for definition blending ratio). The measurement volume increases from 8 to 16, 32, 48 and 64 GC blocks (from $2800t$ to $22400t$).

Figure 5.12 displays how the $RMSE$ changes as a function of time inside the extraction zones of bench A and B. An 8 times larger measurement volume results in approximately half of the initial $RMSE$ reduction. The 144 updates in experiment 111 reduce the $RMSE$ in bench A and B to 0.464 (64.56%) and 0.475 (60.78%). In comparison, the 18 updates in experiment 511 result in a $RMSE$ reduction of 0.237 (32.96%) and 0.224 (28.26%).

The parabolic reduction in $RMSE$ reflects the exceptionally large significance of a few local improvements. These significant improvements occur when the algorithm detects and corrects local anomalies. When excavating a new area, local anomalies are generally detected and corrected early on. The other updates merely result in smaller incremental improvements. Hence the parabolic behaviour.

Larger measurement volumes apparently lead to fewer but more significant updating events. In experiment 111, five large updating events occur improving the model of bench A (Fig. 5.12 top, $t = 16h, 18h, 39h, 40h, 47h$). On average, the five updates yield individual improvements of about 0.053 (ranges from 0.030 to 0.098). Combined, they account for more than half of the total obtained improvement (0.265 versus 0.464). Contrarily, the 4^{th} update ($32h$) in experiment 511 is solely responsible for a reduction in $RMSE$ of 0.233 (Fig. 5.12, top). This single update is almost entirely responsible for the final obtained improvement in bench A. Similar observations can be made in bench B.

The top graph of Fig. 5.13 displays the $\Delta RMSE$ curves derived from predicted measurements and true averages (Ref. Fig. 5.10). Historic and future predicted measurements are jointly considered within a fixed $144h$ time window (entire experiment duration). The overall reduction in $RMSE$ decreases with an increase in measurement volume ($\Delta RMSE$, from 0.651 in exp. 111 to 0.340 in exp 511). However, some caution is needed in interpreting production related results. As measurement volumes grow larger, the $RMSE$ of predicted measurements is lower to start with ($RMSE_0$, 0.673 in exp. 111, 0.407 in exp. 511). As a result, the relative improvements barely vary across the

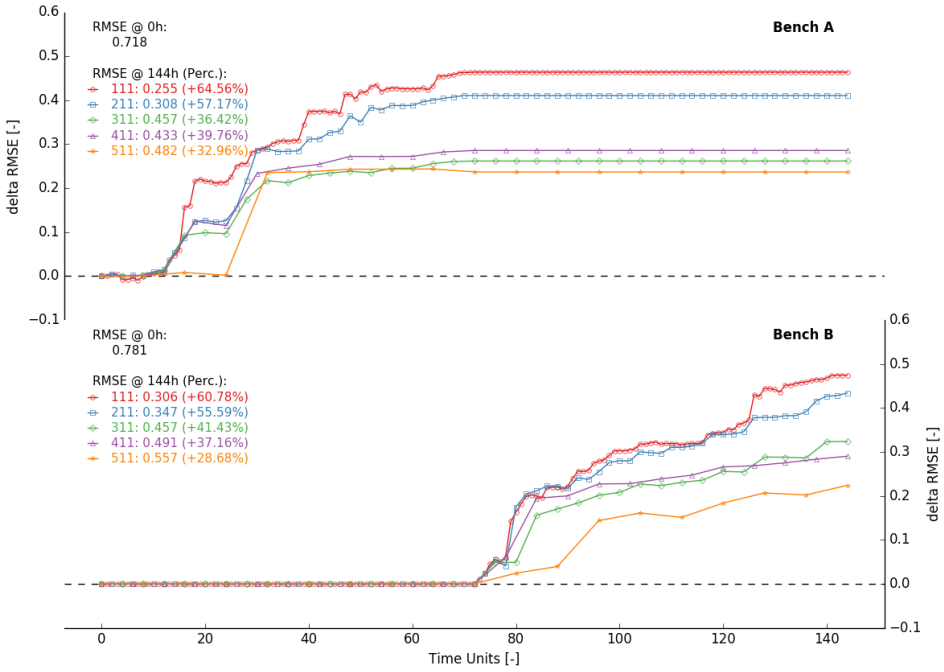

Figure 5.12: Changes in *RMSE* as a function of time, computed over the extraction zones in bench *A* (top) and bench *B* (bottom) - measurement volumes amount to 8 (111), 16 (211), 32 (311), 48 (411) and 64 (511) GC blocks - blending ratio and measurement error are fixed (100% and 0.05)

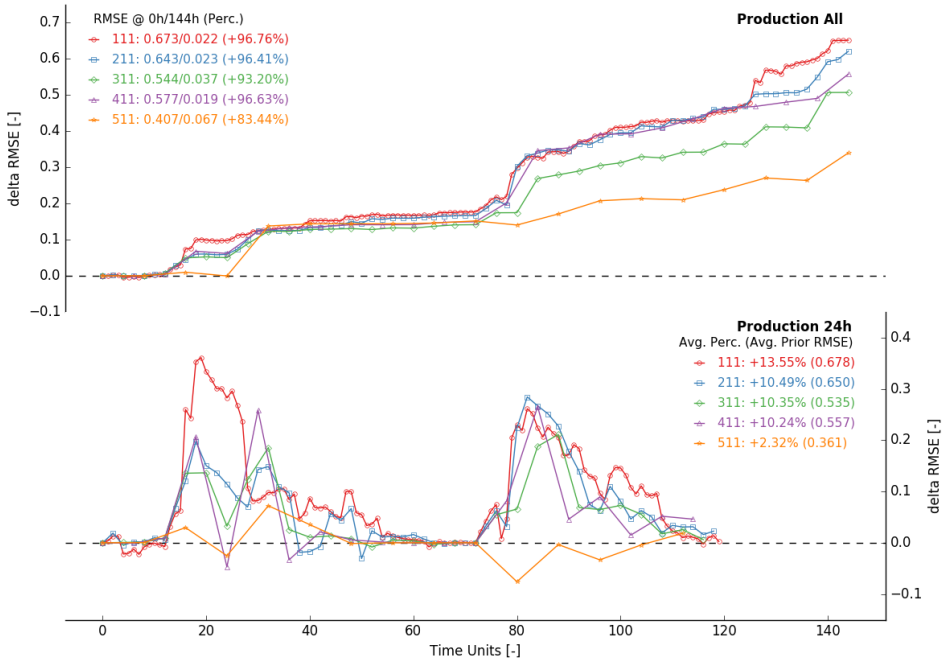

Figure 5.13: Changes in $RMSE$ as a function of time, computed from predicted measurements within a static $144h$ time window (top) and a moving $24h$ future outlook (bottom) - measurement volumes amount to 8 (111), 16 (211), 32 (311), 48 (411) and 64 (511) GC blocks - blending ratio and measurement error are fixed (100% and 0.05)

experiments. During most experiments, the $RMSE$ drops by about 95%.

The bottom graph of Fig. 5.13 depicts improvements in future predicted measurements (next $24h$). The $\Delta RMSE$ curves describe at time t, the reduction in error of upcoming predicted measurements in the time interval $[t, t+24h]$. Figure 5.13 illustrates that improvements diminish with larger measurement volumes. In experiment 111, the $RMSE$ decreases on average by 13.55%. These relative improvements drop to 2.32% when the measurement volume increases from 8 to 64 GC blocks (exp. 511). Note that the average prior $RMSE$ also reduces with a larger measurement volume (avg. $RMSE_0$, 0.678 in exp. 111, 0.361 in exp. 511).

5.3.2. BLENDING RATIO

Figures 5.14 and 5.15 display results from five selected experiments, all conducted using a measurement volume of 8 GC blocks and a measurement error of 0.05. The blending ratio is gradually reduced from 100% to 50% with steps of 12.5% (ref. sec. 5.1.1 for definition blending ratio).

A lower blending ratio results in a lower reduction of the $RMSE$ inside both extraction zones (Fig. 5.14). During all five experiments, 144 measurements are assimilated.

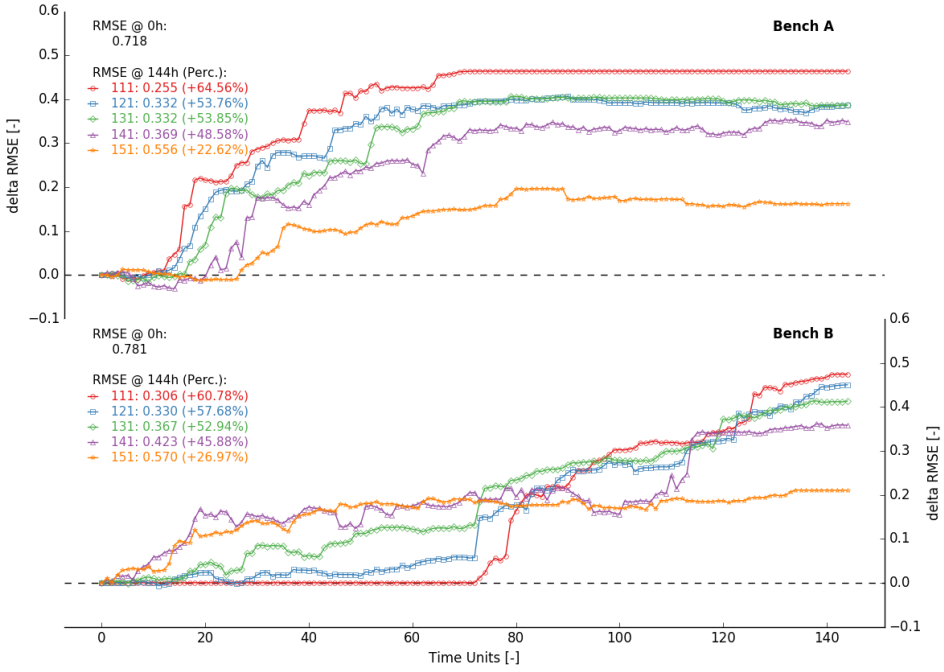

Figure 5.14: Changes in $RMSE$ as a function of time, computed over the extraction zones in bench A (top) and bench B (bottom) - blending ratios amount to 100% (111), 87.5% (121), 75% (131), 62.5% (141) and 50% (151) - measurement volume and measurement error are fixed (8 GC blocks and 0.05)

In experiment 111, the 8 GC blocks in the blend always originate from the same bench. This particular configuration yields improvements of 0.464 (64.56%, bench A) and 0.475 (60.78%, bench B). Improvements drops to about 0.162 (22.62%, bench A) and 0.211 (26.97%, bench B) when half of the blocks originate from the other bench (experiment 151).

When both areas contribute an equal number of blocks to the blend, it becomes harder to pinpoint the source of the detected deviation. The algorithm updates more cautiously leading to a longer sequence of smaller incremental improvements. Local anomalies are easier to detect if the blended material originates from a single area. The observed difference between predicted and recorded measurements can be easily attributed to a single area with a higher level of confidence. Hence, the algorithm corrects more aggressively.

The corresponding larger blending ratio further result in abrupt changes in local production rates (lines 19-20, Algm. 2, p. 68). These abrupt changes significantly influence the overall shape of the $\Delta RMSE$ curves. Periods of steep incline coincide with larger local production rates (e.g. exp. 121 - bench A, $2450t/h$ between $0h$ and $72h$). Lower local production rates result in more gentle slopes (e.g. exp. 121 - bench A, $350t/h$ between $72h$ to $144h$).

Figure 5.15 displays the $\Delta RMSE$ curves derived from all predicted measurements within a static $144h$ time window (top graph). As somewhat expected, the prior $RMSE$ generally decreases with a lower blending ratio ($RMSE_0$, 0.673 in exp. 111, 0.477 in exp. 141). Lower ratios often result in a lesser degree of correlation between the blended blocks (blocks from uncorrelated areas). More extreme block values (difficult to estimate) are likely to be averaged out, as are their larger associated errors. If the GC blocks in the blend originate from a single area (large blending ratio), their corresponding values are going to be correlated. Extreme values are no longer averaged out, nor are their associated larger errors.

The overall reduction in $RMSE$ also tends to be larger in experiments with larger blending ratios ($\Delta RMSE$ of 0.651 in exp. 111, 0.396 in exp. 141). Generally the performance improves with a lesser amount of blending. The differences expressed in terms of relative improvements are less significant. The $RMSE$ drops between 82% and 97%.

Contrary to what has been previously assumed, the blending ratio in itself does not always provide a suitable explanation of an observed phenomenon. In order to understand improvements in future predicted measurements (next $24h$), the more detailed underlying mining schedules need to be analysed instead. Two specific features are of importance. (1) A lower degree of correlation between blended blocks leads to a less informative time-averaged sensor response. The algorithm updates more cautiously. The errors in the GC model remain relatively high. Errors in subsequently derived future predictions do not significantly change (negative effect associated with lower blending ratios). (2) A smaller distance between extracted and scheduled blocks results in better informed future predictions (positive effect associated with lower blending ratios).

Two mining schedules are compared to demonstrate the interaction between both features. All 8 GC blocks, extracted at $24h$ according to schedule $11E$, originate from a single $20m$ x $10m$ rectangular area (Fig. 5.4, top half digging block 12). The 8 blended blocks will be correlated. The 192 GC blocks, scheduled to be extracted within the next $24h$, are located in digging blocks 13 to 24 of bench A (Fig. 5.4). Only half of the GC blocks (digging blocks 13 to 18) lie within a distance of $20m$ of the nearest mining face (horizontal line at $y = 130m$). In contrast, schedule $15E$ describes a simultaneous extraction from two mining faces; one at $y = 130m$ in bench A, the other at $y = 1130m$ in bench B (Fig. 5.4). The next 192 scheduled GC blocks all fit within the next row of digging blocks in either bench A or B. Consequently, all 192 blocks lie within a distance of $20m$ from the nearest mining face. However, schedule $15E$ results in a lower degree of correlation between the blended blocks. The 4 blocks extracted from bench A will hardly be correlated with the 4 blocks from bench B.

Figure 5.15 illustrates how both opposing features impact performance. During experiment 111, the $RMSE$ drops on average by 13.55% (Fig. 5.15, bottom). The average reduction in $RMSE$ in experiment 151 more than doubles (30.75%). Despite the additional extraction point, the performance improves due to a reduced average distance between extracted and scheduled blocks (feature 2 is dominant). This kind of performance gain is untenable when average distances increase. The degree of correlation between the blended blocks becomes the impetus behind the observed

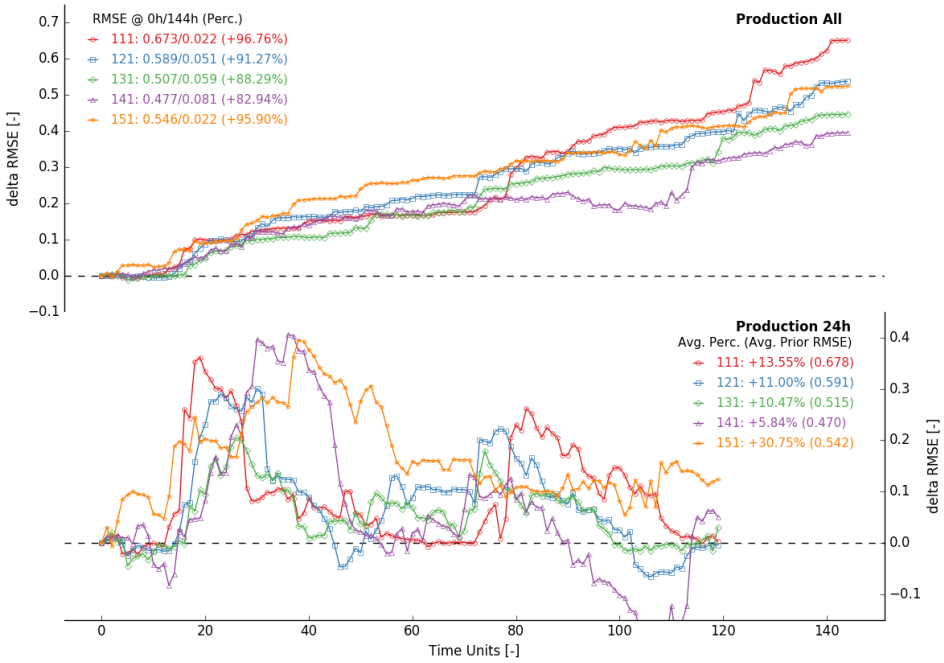

Figure 5.15: Changes in $RMSE$ as a function of time, computed from predicted measurements within a static $144h$ time window (top) and a moving $24h$ future outlook (bottom) - blending ratios amount to 100% (111), 87.5% (121), 75% (131), 62.5% (141) and 50% (151) - measurement volume and measurement error are fixed (8 GC blocks and 0.05)

$RMSE$ reduction (feature 1). During experiment 121, 131 and 141, the $RMSE$ drops by 11.00%, 10.47% and 5.84%.

Experiment 111 further demonstrates that large errors in future predicted measurements are not necessarily corrected on time (Fig. 5.15, bottom, between $60h$ and $72h$). During the last $12h$ of excavation in bench A, the majority of predicted measurements in the next $24h$ are based on block estimates from bench B. Until the processing of material from bench B actually commences, there is no way in knowing that the GC model needs to be locally corrected. Hence very little improvements are observed between 60 and $72h$ ($\Delta RMSE \approx 0$). As soon as data become available, the GC model is updated and the corresponding future predictions are corrected (eventually predictions are nearly always reconciled against measurements, ref. grey areas in Fig. 5.10). No such unpleasant surprises occur during experiment 151. Future predicted measurements are always corrected on time. Since both benches are extracted simultaneously, scheduled GC blocks are continuously better informed due to the proximity of localized production data (measurement attributed to one or multiple source areas). Hence, the $RMSE$ drops consistently around 30% within all $24h$ production windows (Fig. 5.15).

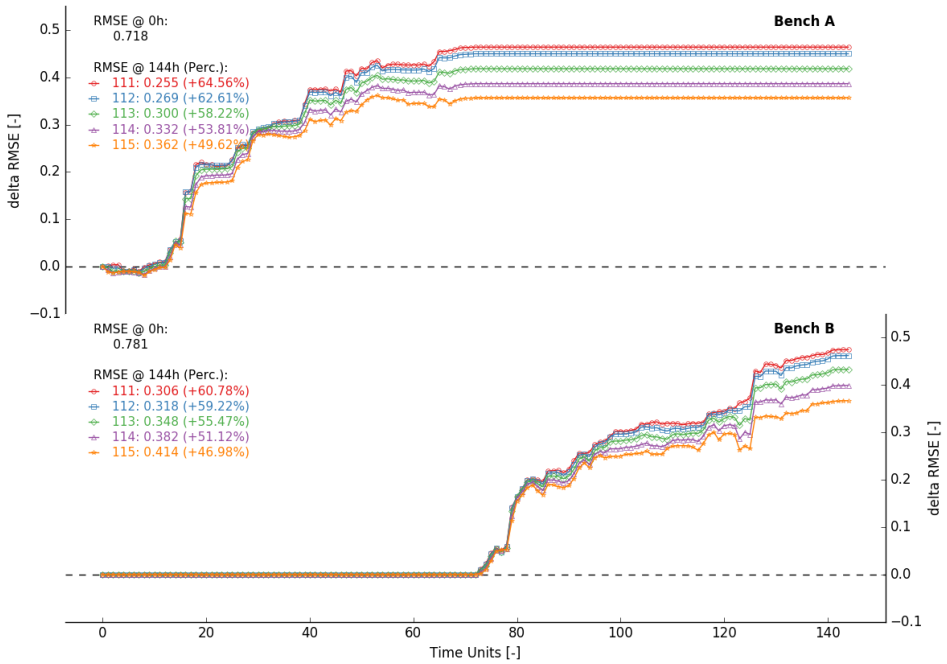

Figure 5.16: Changes in $RMSE$ as a function of time, computed over the extraction zones in bench A (top) and bench B (bottom) - measurement errors amount to 0.05 (111), 0.25 (112), 0.50 (113), 0.75 (114), 1.00 (115) - measurement volume and blending ratio are fixed (8 GC blocks and 100%)

5.3.3. MEASUREMENT ERROR

Figures 5.16 and 5.17 display results from five selected experiments, all conducted using a measurement volume of 8 GC blocks and a blending ratio of 100%. Sensors with measurement errors of 0.05, 0.25, 0.50, 0.75 and 1.00 are used during the experiments. The selected standard deviations define the measurement error E_t on a scale of a single GC block. To put the given numbers in perspective, the standard deviations of all prior GC block estimates vary between 0.284 (near exploration holes) and 1.033. The accuracy of the time-averaged sensor response characterizing 8 GC blocks then amounts to $E_t/\sqrt{8}$.

Figure 5.16 displays how the $RMSE$ changes as a function of time inside the extraction zones of bench A and B. The algorithm updates more cautiously when the measurement error is larger. Individual smaller updates accumulate into lower final $\Delta RMSE$ values (Fig. 5.16). For example, the $RMSE$ reduction observed over the course of experiment 111 amounts to 0.464 (64.56%, bench A) and 0.475 (60.78%, bench B). During experiment 115, the $RMSE$ in bench A and B drops by 0.356 (49.62%) and 0.367 (46.98%) respectively. Despite the 20 times larger measurement error, the performance loss remains manageable.

Figure 5.17 displays the $\Delta RMSE$ curves derived from all predicted measurements

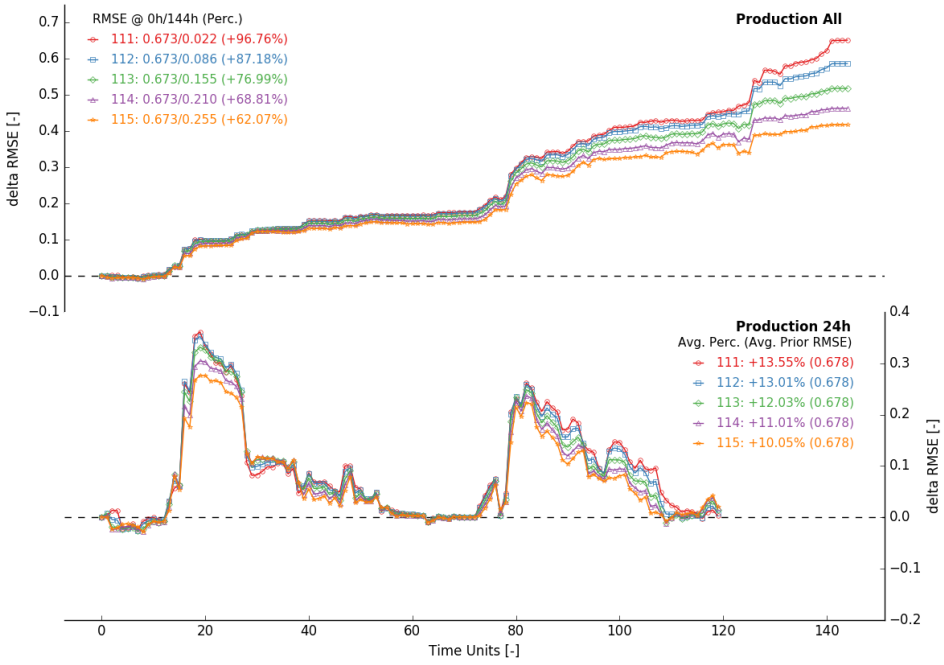

Figure 5.17: Changes in $RMSE$ as a function of time, computed from predicted measurements within a static $144h$ time window (top) and a moving $24h$ future outlook (bottom) - measurement errors amount to 0.05 (111), 0.25 (112), 0.50 (113), 0.75 (114), 1.00 (115) - measurement volume and blending ratio are fixed (8 GC blocks and 100%)

within a static $144h$ time window (top graph). The blending ratio and measurement volume do not change across the five experiments, hence the constant prior $RMSE$ (0.673). The performance drops considerably when increasing the measurement error. Conducting an experiment with a measurement error of 0.05 results in a 96.76% $RMSE$ reduction (exp. 111). The observed improvement drops to 62.07% when employing a sensor with a measurement error of 1.00 (exp. 115).

The bottom graph of Fig. 5.17 illustrates that relative improvements in future predicted measurements (next $24h$) are not significantly influenced by the magnitude of the measurement error. In experiment 111, the $RMSE$ decreases on average by 13.55%. Despite a 20 times larger measurement error, the average observed improvement during experiment 115 still amounts to 10.05%.

5.4. RESULTS - DOUBLE PARAMETER VARIATION

5.4.1. MEASUREMENT VOLUME AND MEASUREMENT ERROR

Figures 5.18 and 5.19 display results from six selected experiments, all conducted using a blending ratio of 75%. Measurement errors are either very low (0.05, accurate sensor,

$V31$) or very high (1.00, noisy sensor, $V35$). The measurement volume increases from 8 ($13E$) to 32 ($33E$) and 64 ($53E$) GC blocks. As a reminder, the accuracy of the time-averaged sensor response can be computed as $E_t/\sqrt{V_t}$.

Figure 5.18 displays the evolution of the $\Delta RMSE$ inside the extraction zones of bench A and B. An increase in measurement volume while applying an accurate sensor (131-331-531) leads to a severe deterioration in performance. Experiment 131 yields $RMSE$ reductions of 0.418 (53.85%, bench A) and 0.433 (52.94%, bench B). The $RMSE$ reductions observed over the course of experiment 531 only amount to 0.214 (29.76%, bench A) and 0.234 (30.01%, bench B). An increase in measurement volume while applying a noisy sensor results in a very different behavior (135-335-535). Initially a larger measurement volume seems to result in a stagnating performance (135-335). The final $RMSE$ reduction barely differs. The $RMSE$ reduction changes from 0.217 (30.18%) to 0.223 (30.97%) in bench A and from 0.239 (30.58%) to 0.256 (32.77%) in bench B. This was somewhat expected since the overall accuracy of recorded measurements ($E_t/\sqrt{V_t}$) increases as noisy sensor responses are averaged over larger volumes or periods of time (larger V_t). Further increasing the measurement volume from 32 to 64 GC blocks (335-535) eventually causes the expected lower performance. The $RMSE$ reduction in bench A and B drop respectively from 0.223 (30.97%) to 0.177 (24.69%) and from 0.256 (32.77%) to 0.207 (26.52%). Apparently, more accurate measurements ($1.00/\sqrt{64}$ versus $1.00/\sqrt{32}$) are not sufficient to counteract the adverse effects associated with larger volumes.

When applying an inaccurate sensor, an optimal measurement volume seems to exist. A volume too low increases the influence of individual errors on a time-averaged sensor response. A volume too large makes it harder to pinpoint the exact source of the detected deviation. An optimal volume is just large enough to average out individual errors, while still being small enough to yield sufficiently local updates.

The top graph of Fig. 5.19 displays the $\Delta RMSE$ curves derived from all predicted measurements within a static $144h$ window. Averaging accurate sensor measurements over larger volumes (131-331-531) leads to lower $RMSE$ reductions. Measurement volumes of 8, 32 and 64 GC blocks yield reductions of 0.448, 0.403 and 0.384. The relative improvements over the course of these three selected experiments are all approximate by 88.50%. An increase in measurement volume while applying a noisy sensor (135-335-535) does seem to have the opposite effect. The reduction in $RMSE$ gets larger as the measurement volume increases (0.237-0.281-0.285). An 8 times larger measurement volume (from 8 to 64 GC blocks) yields a 1.5 times larger relative improvement (from 46.48% to 65.54%).

The bottom graph of Fig. 5.19 depicts improvements in future predicted measurements (next $24h$). An increase in measurement volume, in combination with a low measurement error, leads to a loss of performance (131-331-531, from 10.47% to 5.51%). The performance barely reduces when increasing the measurement volume in combination with a high measurement error (135-335-535, from 7.83% to 7.48%). In certain circumstances, a larger measurement error seems to allow for more accurate future predictions (compare exp. 331-335 and 531-535).

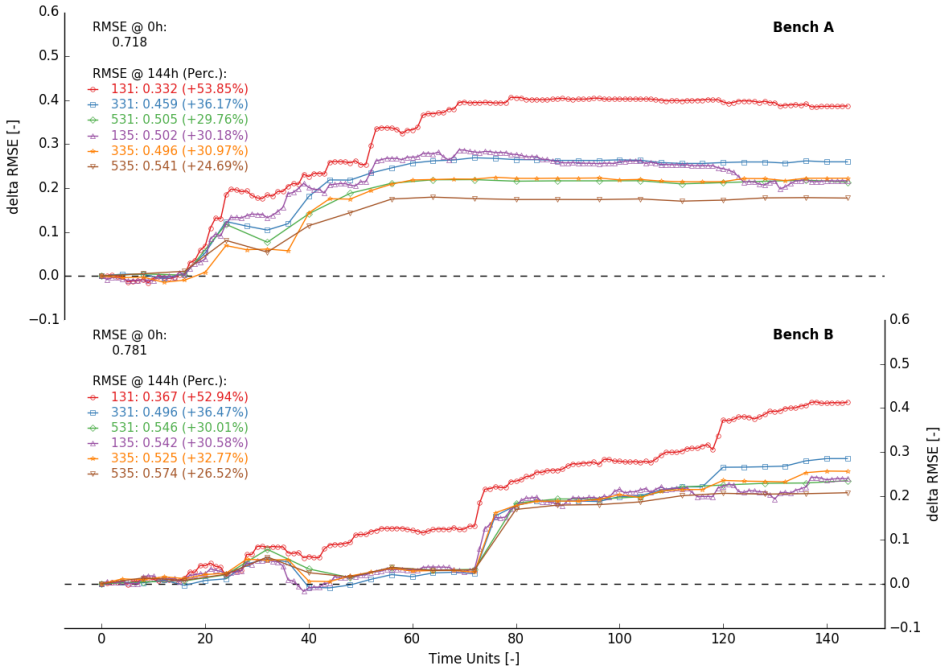

Figure 5.18: Changes in $RMSE$ as a function of time, computed over extraction zones in bench A (top) and bench B (bottom) - blending ratio fixed at 75% ($V3E$) - measurement volumes amount to 8 ($13E$), 32 ($33E$) and 64 ($53E$) GC blocks - measurement error is either very low (0.05, $V31$) or very high (1.00, $V35$)

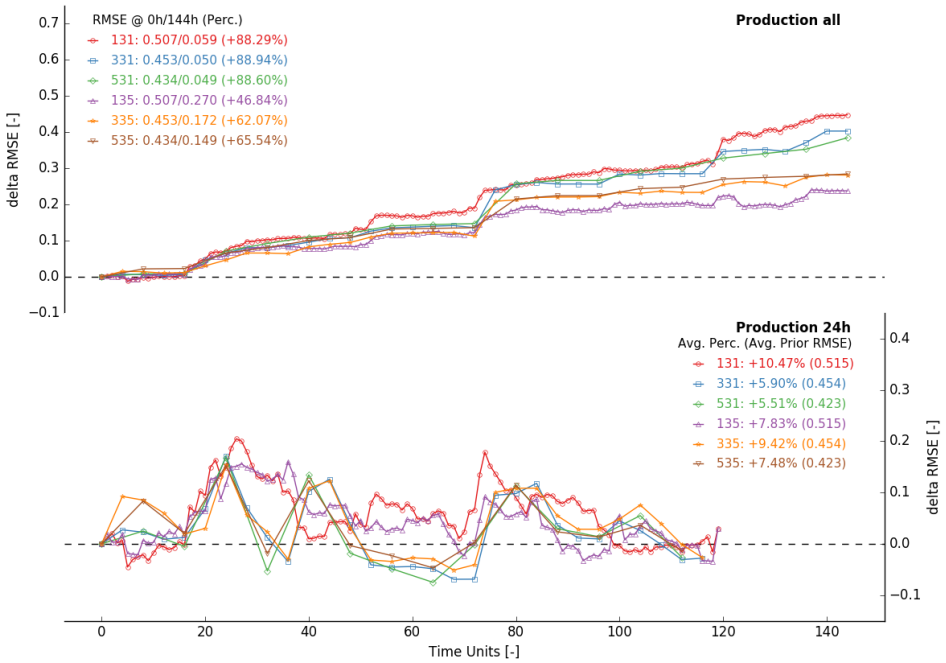

Figure 5.19: Changes in $RMSE$ as a function of time, computed from predicted measurements within a static $144h$ time window (top) and a moving $24h$ future outlook (bottom) - blending ratio fixed at 75% ($V3E$) - measurement volumes amount to 8 ($13E$), 32 ($33E$) and 64 ($53E$) GC blocks - measurement error is either very low (0.05, $V31$) or very high (1.00, $V35$)

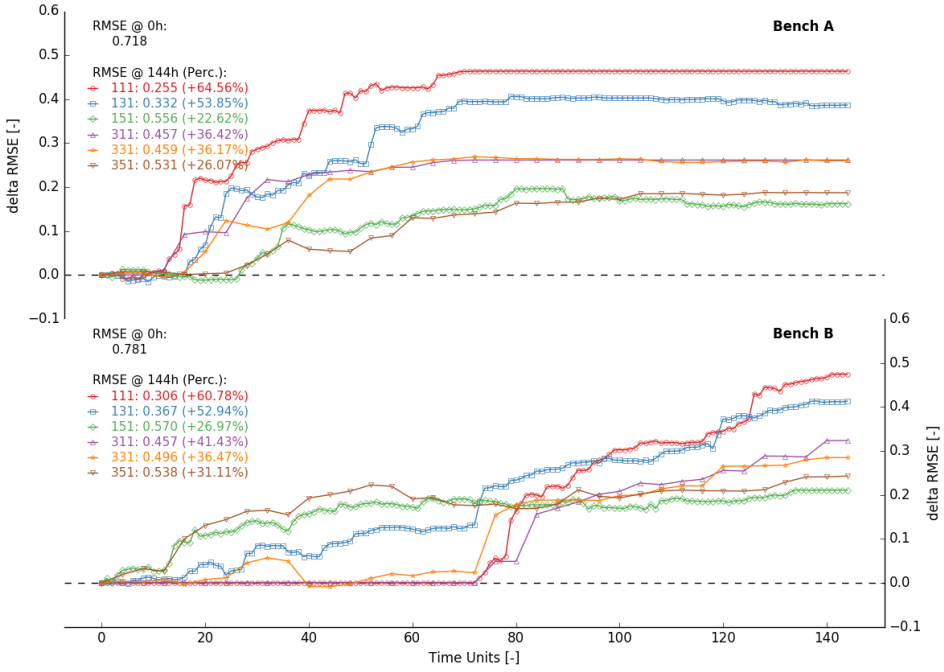

Figure 5.20: Changes in $RMSE$ as a function of time, computed over the extraction zones in bench A (top) and bench B (bottom) - measurement error fixed at 0.05 - measurement volumes amount to 8 ($1R1$) and 32 ($3R1$) GC blocks - blending ratios amount to 100% ($V11$), 75% ($V31$) and 50% ($V51$)

5.4.2. MEASUREMENT VOLUME AND BLENDING RATIO

Figures 5.20 and 5.21 display results from six selected experiments all conducted using an accurate sensor (0.05). Measurement volumes are either small (8 GC blocks, $1R1$) or medium (32 GC blocks, $3R1$). The blending ratio is gradually reduced from 100% to 50% with steps of 25% ($V11$-$V31$-$V51$).

Figure 5.20 displays how the $RMSE$ changes as a function of time inside the extraction zones of bench A and B. A significant performance loss occurs as local production rates converge, while averaging sensor responses over a relatively short period of time (111-131-151). During experiment 111, when blocks in the blend originate from a single area, the reduction in $RMSE$ amounts to 0.464 (64.65%, bench A) and 0.475 (60.78%, bench B). When both areas contribute an equal number of blocks to the blend, the $RMSE$ only drops 0.162 (22.62%, bench A) and 0.211 (26.97%, bench B). The effects of an increased amount of blending in combination with a larger measurement volume is less profound (311-331-351). The improvements in bench A and B reduce from 0.262 (36.42%) to 0.187 (26.07%) and from 0.323 (41.43%) to 0.243 (31.11%) respectively.

Figure 5.21 displays the $\Delta RMSE$ curves derived from all predicted measurements within a static $144h$ time window (top graph). Blending more, while considering a small measurement volume (111-131-151) does not necessarily result in a lower reduction

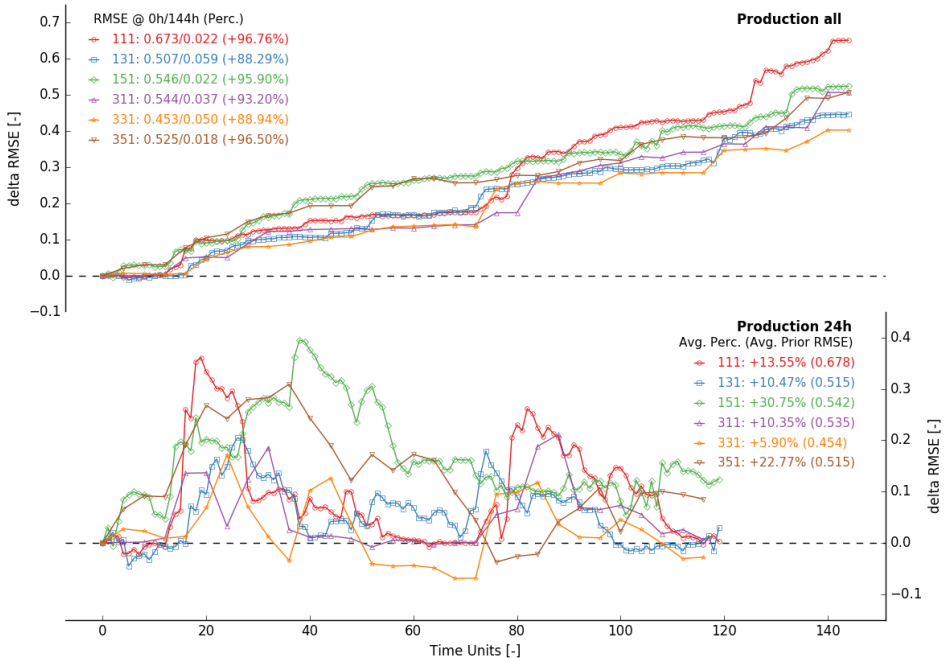

Figure 5.21: Changes in $RMSE$ as a function of time, computed from predicted measurements within a static $144h$ time window (top) and a moving $24h$ future outlook (bottom) - measurement error fixed at 0.05 - measurement volumes amount to 8 ($1R1$) and 32 ($3R1$) GC blocks - blending ratios amount to 100% ($V11$), 75% ($V31$) and 50% ($V51$)

of the $RMSE$. The performance loss in experiment 131 is for example larger than in experiment 151 ($\Delta RMSE$ of 0.448 versus 0.524). Blending within larger measurement volumes (311-331-351) does not seem to yield any different behaviour. Once again, the experiment with a blending ratio of 50% performs better than one with a ratio of 75%. Obviously, the magnitude of the $RMSE$ reductions are lower due to the larger measurement volume.

The bottom graph of Fig. 5.21 depicts improvements in future predicted measurements (next $24h$). As explained earlier, improvements in future predicted measurements rather result from an underlying mining schedule than from the actual values of the blending ratio. Figure 5.21 (bottom graphs) once again confirms that the mining schedules resulting from equal local production rates (blending ratio of 50%) are the ost optimal (with respect to future predictions). During experiment 151 and 351, the $RMSE$ reduces on average by 30.75% and 22.77%. In comparison, during experiment 111 and 311, the $RMSE$ reduces on average by 13.55% anf 10.35%. The results further indicate that an increase in measurement volume leads to a reduction in performance (compare 111-311 and 151-351).

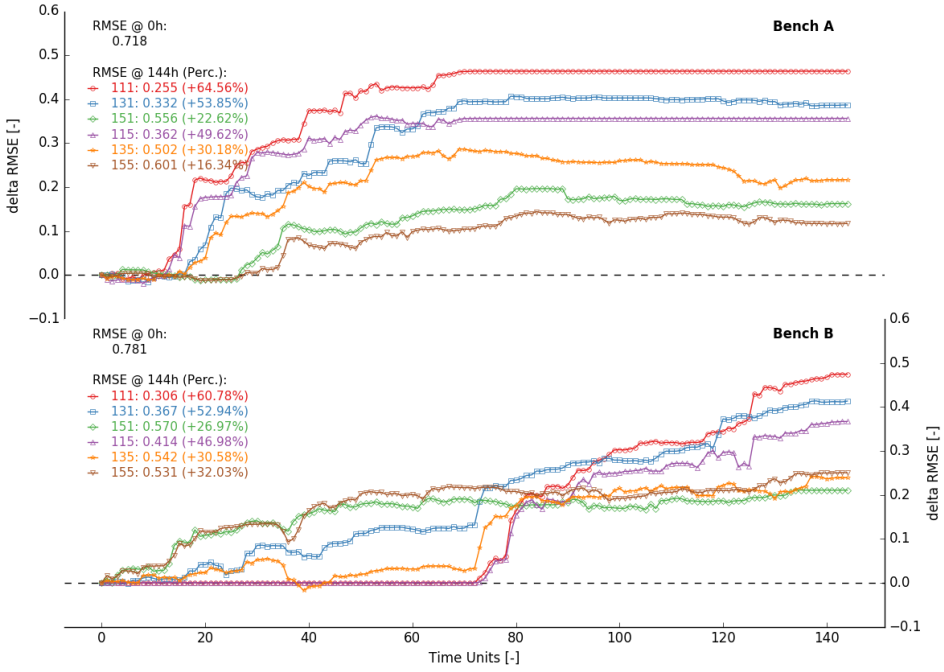

Figure 5.22: Changes in $RMSE$ as a function of time computed over the extraction zones in bench A (top) and bench B (bottom) - measurement volume fixed at 8 GC blocks - blending ratios amount to 100% ($11E$), 75% ($13E$) and 50% ($15E$) - measurement error is either very low (0.05, $1R1$) or very high (1.00, $1R5$)

5.4.3. BLENDING RATIO AND MEASUREMENT ERROR

Figures 5.22 and 5.23 display six selected experiments, all conducted using a measurement volume of 8 GC blocks ($1RE$). Measurement errors are either very low (0.05, $1R1$) or very high (1.00, $1R5$). The blending ratio gradually reduces from 100% to 50% with steps of 25% ($11E$-$13E$-$15E$).

Figure 5.22 displays how the $RMSE$ changes as a function of time inside the extraction zones of bench A and B. A decrease in blending ratio, while applying an accurate sensor (111-131-151) leads to a severe loss of performance. The $RMSE$ reductions in bench A and B drop from 0.464 (64.65%) to 0.162 (22.62%) and from 0.475 (60.78%) to 0.211 (26.97%). The use of a noisy sensor does not change this general pattern (115-135-155). The reduction in $RMSE$ in bench A and B decreases from 0.356 (49.62%) to 0.117 (16.34%) and from 0.367 (46.98%) to 0.250 (32.03%).

The top graph in Fig. 5.23 displays the $\Delta RMSE$ curves derived from all predicted measurements within a static $144h$ time window. Increasing the amount of blending while applying an accurate sensor (111-131-151) does not necessarily result in a lower reduction of the $RMSE$. The performance loss in experiment 131 is for example larger than in experiment 151 ($\Delta RMSE$ of 0.448 versus 0.524). Blending highly inaccurate sensor observations does not directly result in any different behaviour. Again, the

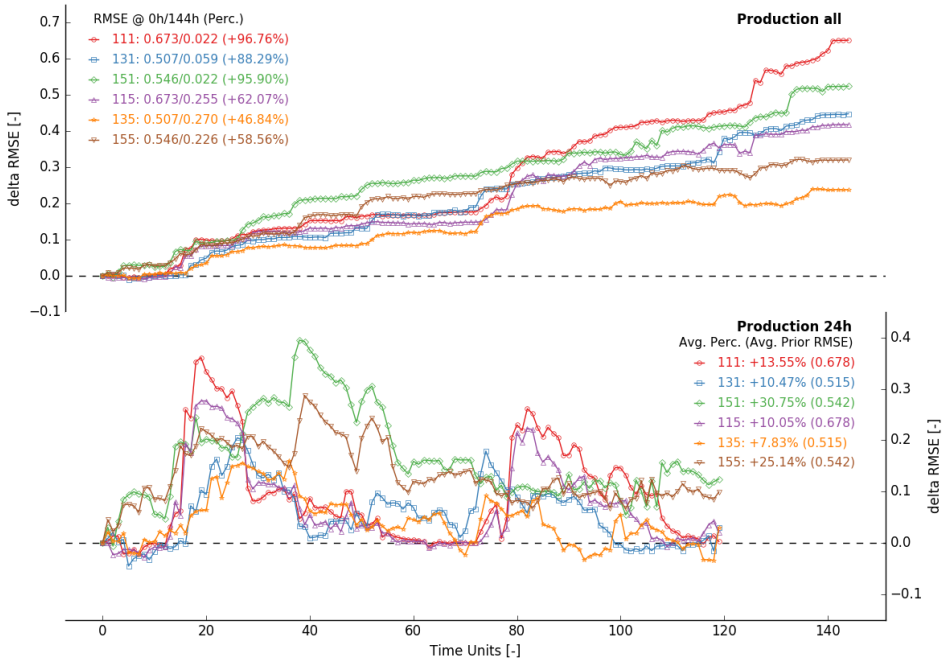

Figure 5.23: Changes in $RMSE$ computed from predicted measurements within a static $144h$ time window (top) and a moving $24h$ future outlook (bottom) - measurement volume fixed at 8 GC blocks - blending ratios amount to 100% ($11E$), 75% ($13E$) and 50% ($15E$) - measurement error is either very low (0.05, $1R1$) or very high (1.00, $1R5$)

experiment with a blending ratio of 50% performs better than one with a ratio of 75%. Obviously, the magnitude of the $RMSE$ reductions are lower due to the larger measurement error.

The bottom graph in Fig. 5.23 depicts relative improvements in future predicted measurements (next $24h$). Figure 5.23 illustrates that a 20 times larger measurement error still results in manageable performance reductions (compare 111-115, 131-135 and 151-155). The mining schedules, defined by underlying blending ratios (and measurement volumes), do have a more significant impact on the obtained $RMSE$ reductions. Experiments with blending ratios of 50% perform consistently better. During experiments 151 and 155, the $RMSE$ drops on average by 30.75% and 25.14%. In contrast, over the course of experiments 111 and 115, the $RMSE$ reduces on average by 13.55% and 10.05%.

5.5. DISCUSSION

The aim of this section is to translate previous observations into three sets of practical recommendations. The first set describes how to maximize the $RMSE$ reduction in the

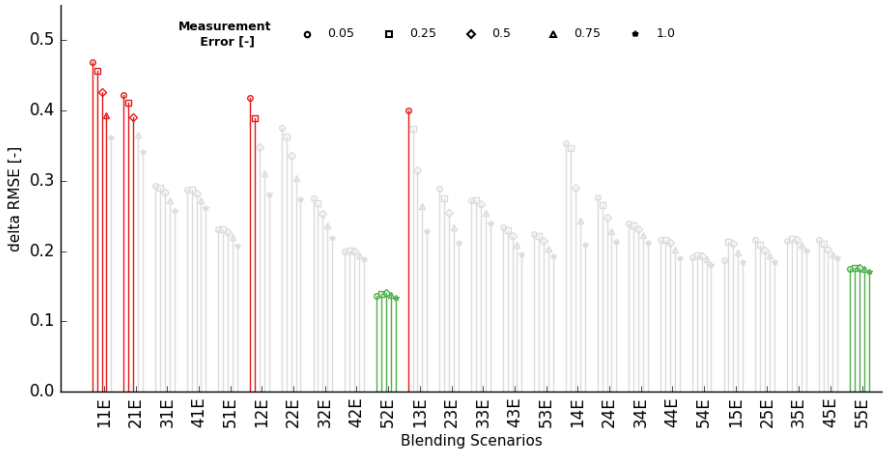

Figure 5.24: Aggregate change in $RMSE$ computed for the extraction zones in bench A and B - blending scenarios are indicated on the x-axis (Tab. 5.1) - markers refer to the measurement error - 10 best (red) and 10 worst (green) performing experiments are coloured

GC model. The second set provides some guidelines to optimize the reconciliation of production data (match predicted measurements with actuals). The third set directs actions towards obtaining the most accurate future predictions. The section concludes with some practical implications.

5.5.1. OBJECTIVE 1 - MAXIMIZE ERROR REDUCTION IN GC MODEL

System parameters are considered most optimal when leading to a maximum error reduction in the GC model. Instead of analysing two $\Delta RMSE$ values (one per bench), an aggregate measure is computed ($(\Delta RMSE_A + \Delta RMSE_B)/2$). Figure 5.24 provides a visual overview of this performance measure recorded at the end of each experiment. The experiments are grouped according to their blending scenario (Tab. 5.1). Within each group, a marker refers to the precision of the applied sensor. The 10 best and 10 worst performing experiments are coloured red and green respectively.

Controlling system parameters to improve algorithm performance is only worth the effort if local production rates are allowed to be unequal. Figure 5.24 illustrates that variations in interval duration (time between $t-1$ and t) and/or sensor precision have little to no impact if both areas contribute an equal number of blocks to the blend (Fig. 5.24, 25 rightmost experiments). As soon as local production rates are unequal, one of the most effective measures to increase performance is to reduce the interval duration. If possible, local production rates should be adjusted such that the majority of the blended material originates from a single area. Furthermore, investing in accurate sensors only seems to pay off when the interval durations are kept small.

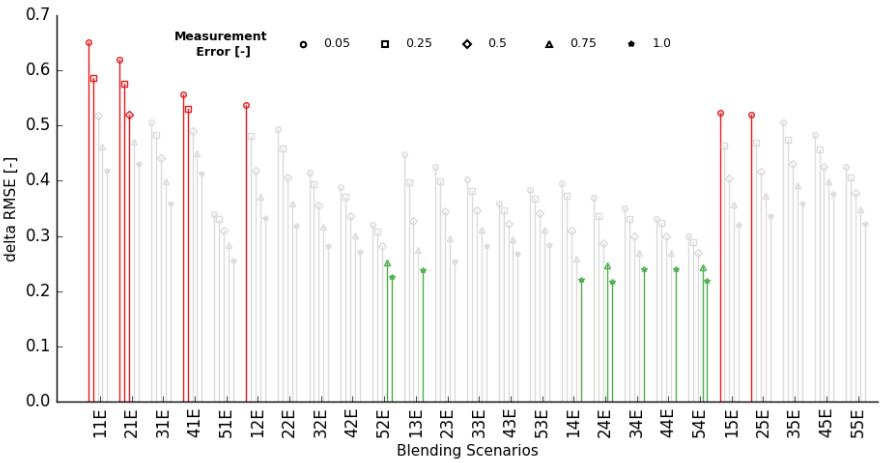

Figure 5.25: Change in $RMSE$ computed at the end of each experiment from all available predicted measurements - blending scenarios are indicated on the x-axis (Tab. 5.1) - markers refer to the measurement error - 10 best (red) and 10 worst (green) performing experiments are coloured

5

5.5.2. OBJECTIVE 2 - OPTIMIZE RECONCILIATION OF PRODUCTION DATA

The first objective aims to minimize the error at the input side of the forward simulator (the GC model, $\mathbf{Z}_{t-1}(:,i)$). The second objective however is directed towards maximizing the error reduction at its output side. In other words deviations between actual and predicted measurements are to be reduced as much as possible ($\mathbf{d}_t - \mathscr{A}_t(\mathbf{Z}_{t-1}(:,i))$). Figure 5.25 displays, per experiment, the last recorded $\Delta RMSE$ values computed from predicted measurements within static $144h$ windows (ref. 5.10i).

A lower interval duration generally results in a larger $RMSE$ reduction. Caution is advised as under certain circumstances the intended effect is limited or even reversed (converging local production rates - R of 75%, 62.5% or 50% and highly inaccurate sensors - E of 0.75 and 1.00). Opting for accurate sensors is generally more effective (larger $\Delta RMSE$ when E is lower). Adjusting the local production rates further affects the $RMSE$ reductions. Results are better when production rates are either very alike or very different (blending ratio of 100%, 82.5% and 50%).

Interestingly, the system parameters considered optimal in terms of reconciled production data do not necessarily lead to the best GC model improvements (compare Figs. 5.25 and 5.24). In other words, minimizing the error at the output side of the forward simulator does not necessarily result in a minimum amount of error at the input side.

5.5.3. OBJECTIVE 3 - MAXIMIZE ERROR REDUCTION IN FUTURE PREDICTIONS

Thus far, system parameters have been selected to either optimize GC model improvements or production reconciliations. The related assessment statistics, quantifying one of both objectives, nearly always improved over time. As a result, the last recorded values

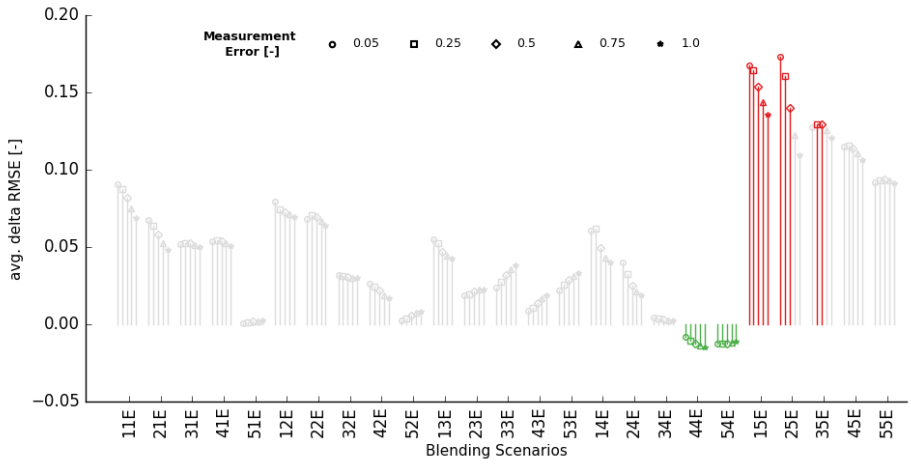

Figure 5.26: Average change in $RMSE$ computed from predicted measurements within moving $24h$ windows - blending scenarios are indicated on the x-axis (Tab. 5.1) - markers refer to the measurement error - 10 best (red) and 10 worst (green) performing experiments are coloured

could be used to describe and assess the overall performance during the experiments. A last recorded statistic would not suffice to evaluate the overall quality of future predictions (next $24h$). Instead, an average $\Delta RMSE$ is computed, summarizing the performance across all available $24h$ windows.

The average reduction in $RMSE$ (avg. $\Delta RMSE$) needs to be maximized to optimally improve future predictions. Equal local production rates seem to lead to exceptional performances (Fig. 5.26, blending ratio of 50%). As local production rates converge, distances between scheduled and extracted GC blocks reduce. Reduced distances however can only partially account for the observed phenomenon. After all, experiments with blending ratios of 82.5%, 75% and 62.5% do perform worse and not better than their counterparts with a blending ratio of 100%. A reasonable explanation follows from a detailed study of the various mining schedules. Presumably, converging local production rates only result in a performance gain if the average distance is sufficiently low (feature 2). As long as this threshold is not reached, the performance drops due to the ever lower correlation between the blended blocks (feature 1). Essentially, in order to optimize future predictions, an optimal schedule has to be found.

Figure 5.26 further illustrates that a lower interval duration generally results in larger $RMSE$ reductions. When interval durations are small, more accurate sensors definitely add value. The outcome of installing better sensors is less predictable as soon as interval durations start to increase. Within a minority of blending scenarios, more accurate sensors surprisingly result in lower $RMSE$ reductions. During 10 out of 125 experiments the error in future predictions even slightly increases.

5.6. PRACTICAL IMPLICATIONS

Nearly all results have shown that shorter interval durations and smaller measurement volumes improve algorithm performance, regardless of the considered metric. Due to the complexity of the material handling process, it might be challenging to construct a forward simulator accurate enough to track material on a scale as small as 8 GC blocks ($2800t$). A larger time interval might have to be selected to ensure that a forward predicted measurement and a time-averaged sensor response are actually characterizing roughly the same material. A representative forward prediction is essential in linking specific observations to their constituent GC blocks. As measurement volumes grow too large (64 GC blocks), the opportunity to improve future predictions might be lost.

It further would be advisable to install by default the most accurate sensor within the available budget (note the selected sensor also has to meet certain reliability requirements). First of all, a more accurate sensor significantly improves the reconciliation of production data (objective 2). Secondly, even in combination with medium measurement volumes (16 and 36 GC blocks), GC model improvements increase when more accurate sensors are applied.

In practice, it will be nearly impossible to control local production rates for the sake of improving the algorithm performance. Instead, local production rates are to be considered as external input and will very likely change over time. The previous results nevertheless can be used to predict how the algorithm performance will change as local production rates are adjusted. Diverging local production rates will always result in higher GC model improvements. Due to a better GC model, future predictions improve as well. Converging production rates generally harm algorithm performance. Errors in GC block estimates and future predicted measurements are less well corrected. However, at some point, the average distance between extracted and scheduled GC blocks reaches a critical threshold. If this happens, future predictions undergo exceptional improvements.

5.7. CONCLUSIONS

Recently, a new algorithm has been developed to repeatedly update the grade control model based on online data from a production monitoring system. The added value of the presented algorithm results from its ability to handle inaccurate observations made on blended material from two or more extraction points. A total of 125 artificial experiments are conducted to evaluate the influence of system parameters on the overall algorithm performance.

Each experiment mimics a virtual mining operation with two extraction points. Over the course of a $144h$ long experiment, a total of $403200t$ is excavated, processed and measured. Despite their mutual differences, all mining schedules ensure that at the end of each experiment the same 1152 GC blocks are extracted. Local production rates and interval durations do however vary significantly across the experiments. Moreover various sensors with a different precision are applied.

Results regarding GC model improvements are promising. Under fairly optimal conditions, the $RMSE$ in the GC model drops by around 60%. Optimal conditions occur when small operational volumes are extracted from a single location and characterized

with a relatively accurate sensor (8 GC blocks - 2800t, blending ratio of 100%, sensor precision of 0.05). Even when conditions are far from optimal, the $RMSE$ in the GC models can still be reduced by about 20%. This is remarkable, especially considering the large measurement volumes, the amount of blending and the extremely inaccurate sensor (64 GC blocks - 22400t, blending ratio of 50%, sensor precision of 1.00). The results indicate that even global inaccurate blended averages have some information potential to further improve the GC model.

The experiments further illustrated the outstanding reconciliation capabilities of the updating algorithm. The $RMSE$ between historic predicted and actual measurements decreased between 46% and 97% over the course of an experiment. More accurate sensors improved the reconciliation behavior. Larger $RMSE$ reductions were further observed when local production rates were sitting at either end of the spectrum (blending ratio of 100% or 50%). Interval duration did not significantly impact results.

Improvements regarding future predicted measurements (next 24h) turned out to be less robust against variations in system parameters. Average $RMSE$ reductions vary between -5% and $+30\%$. Two very different regimes were observed. As long as the average distance between scheduled and extracted GC blocks did not reach its lower threshold, performance drops as interval duration increase and local production rates start to converge. However, once an average distance threshold is reached, the $RMSE$ reduction becomes exceptionally high.

A clear distinction needs to be made between the GC model and production related improvements. On one hand, previous results clearly show that a single-value measure, such as the blending ratio, is sufficient to predict the behaviour of GC model updates. The blending ratio in itself can act as a proxy for the information content of a time-averaged sensor response. A high value refers to large correlations between all blended blocks and indicates that GC model updates will be more successful.

On the other hand, the value of the blending ratio does not explain variations in production related performance. The best performing experiments, related to reconciliation and future predicted measurements, are sitting on opposing sides of the blending ratio spectrum (either 50% or 100%). A plausible explanation is provided. The explanation is however not conclusive. Future research should be conducted to determine the influence of the correlation between extracted and scheduled blocks. A single-value statistic needs to be devised describing the updating potential of a mining schedule at a discrete timestep. Three types of correlations need to be considered; cross-correlations between recently excavated blocks, cross-correlations between near-term scheduled blocks and correlations between recently extracted and near-term scheduled blocks. All this information could possibly be captured in some sort of signal-to-noise ratio. A large ratio would predict a good amount of improvement in future predictions.

Future research should further focus on how the prior realizations impact local anomaly corrections later on. The empirical covariances lifted from the prior realizations largely determine the extent of updates in neighbouring block values. When the covariances between measurements and neighbouring blocks are overestimated, attribute values are corrected too aggressively, causing erroneous updates. A corrective localization technique is integrated into the algorithm to prevent such behaviour (Wambeke and Benndorf, 2017b). Underestimated covariances on the other hand yield

updates which are too modest. Certain local anomalies remain partly uncorrected. In reality, the correlation structure of the true field is obviously unknown. Several parallel updating tracks could be ran simultaneously, all initiated with a specific set of prior realizations (each set could be generated using a different variogram model and/or simulation technique). Error statistics, collected during an initial training period, could eventually be compared to select the most 'correct' realization set, best describing the true spatial correlation.

REFERENCES

T. Wambeke and J. Benndorf, *A study of the influence of measurement volume, blending ratios and sensor precision on real-time reconciliation of grade control models.* Mathematical Geosciences (accepted with revisions) , 1–45 (2017a).

T. Wambeke and J. Benndorf, *A simulation-based geostatistical approach to real-time reconciliation of the grade control model,* Mathematical Geosciences **49**, 1–37 (2017b).

6

PILOT STUDY

This chapter describes the pilot testing of the entire updating concept (including the updating algorithm) at the Tropicana Gold Mine in Australia. The aim of the study is to evaluate whether the developed updating algorithm can be used to update spatial Work Index estimates based on actual ball mill performance data. The chapter commences with a detailed explanation of the practical problem. Thereafter, the updating equations are briefly reviewed and reformulated taking account of the problem specific terminology. The chapter further provides some background information regarding the geology at Tropicana, the operation and the available data. Subsequently, some insights are given regarding the construction of the application specific forward simulator. Then, results are presented illustrating improvements in both historic and future production estimates. The chapter finally concludes with an extensive discussion on modelling assumptions and potential improvements.

Parts of this chapter have been submitted for publication in Transactions of the Institutions of Mining and Metallurgy, SectionA - Mining Technology **accepted, in press** (Wambeke *et al.*, 2017).

6.1. INTRODUCTION

Traditionally, the mining industry has had mixed success in achieving the production targets it has set out. Produced tonnages (and grades) nearly always deviate from model-based expectations due to ever present geological uncertainties. Even when numerous exploration samples are collected, it remains challenging to accurately characterize short-term production units equivalent to a few truckloads (Benndorf, 2013). In certain commodities, Grade Control (GC) drilling is performed to further reduce uncertainties (Peattie and Dimitrakopoulos, 2013; Dimitrakopoulos and Godoy, 2014). GC drilling is expensive and almost exclusively focused on sampling grades.

At the Tropicana Gold Mine, GC samples are collected at one meter intervals during Reverse Circulation drilling. Once collected, the samples are sent to an on-site laboratory for a semi-automated analysis. An autonomous system crushes, splits and pulverises the sample material prior to X-Ray Fluorescence (XRF) and Hyper-Spectral (HS) scanning. Conventional fire assaying techniques are used to determine the gold grade in a final prepared pulp. Calibrated relationships are subsequently applied to translate the obtained proxy measurements (XRF and HS) into geometallurgical estimates (e.g. work index, hardness or recovery). At this stage, the geometallurgical estimates describe the properties of one metre long cylindrical volumes, virtually located at the original down-hole positions of the GC samples. Geostatistical techniques are used at a later stage to convert the scattered interval data into metallurgical estimates for contiguous block volumes in the GeoMet model (Catto, 2015).

The calibrated relationships, vital in obtaining the metallurgical interval data, are largely untested. A larger number of metallurgical tests to improve the calibration is simply economically infeasible. Hence, despite all efforts, the derived geometallurgical block estimates remain (largely) inaccurate.

The Bond Ball Mill Work Index (Wi) is one such spatial estimate which remains difficult to infer correctly. The Wi defines the specific energy (kWh/t) required in grinding a ton of ore in the ball mill from a very large size (infinite) to 100 μm (Lynch *et al.*, 2015). At the time of writing, this variable is of particular interest for the following two reasons. (1) Collecting large datasets of Wi values is very expensive due to labour-intensive and time consuming laboratory work. It is worthwhile to investigate alternative options for improving the calibration of the concerned relationship. (2) Ball mill throughput could potentially be optimized by improving the Wi estimates of the mill feed. Both reasons justify the exclusive focus on improving Wi estimates during the remainder of the text. Before proceeding, a clear distinction is made between various types of Wi estimates.

1. $\mathbf{Wi}^s(m)$: Wi estimates describing one metre long cylindrical GC samples. $\mathbf{Wi}^s$ is a column vector with M Wi estimates. Each estimate m is centred at the original down-hole position of its corresponding GC sample.

2. $\mathbf{Wi}_t^b(n,i)$: Wi estimates describing block volumes in the GeoMet model. The N rows in the $\mathbf{Wi}_t^b$ matrix each refer to a unique block in the GeoMet model. The I columns contain different spatial realizations characterizing geological uncertainty (Monte Carlo approach, each realization represents an equally plausible scenario). The block estimates do change in time as new information is

assimilated into the GeoMet model, hence the subscript t.

3. $\mathbf{Wi}_t^f(1,i)$: Wi estimates describing the mill feed between $t-1$ and t. $\mathbf{Wi}_t^f(1,i)$ is a row vector containing I realizations. The I realizations approximate a distribution describing a best estimate (mean) and its related uncertainty (spread). The mill feed estimates are no longer attributable to a single spatial coordinate (the mill feed typically represents a blend of ore from multiple sources and locations).

The mill feed estimates $\mathbf{Wi}_t^f(1,i)$ are substituted in the following formula to compute the energy required in grinding a ton of ore in the mill from a known feed size to a required product size (Lynch *et al.*, 2015):

$$\frac{P}{R} = \mathbf{Wi}_t^f(1,i)\left(\frac{10}{\sqrt{P_{80}}} - \frac{10}{\sqrt{F_{80}}}\right) \quad \forall i \in I. \tag{6.1}$$

Assuming a constant power draw (P in kW), the energy delivered per ton is controlled by adjusting the mill throughput (R in t/h). The F_{80} and P_{80} represent the 80% passing sizes of the feed and product respectively (F_{80} and P_{80} in μm). To maximize mill throughput and optimize energy utilization, it is important to get the $\mathbf{Wi}_t^f$ estimates right.

When the ore is softer than expected ($\mathbf{Wi}_t^f$ is overestimated), the amount of energy transferred into each ton of material is too large and the resulting product will be too fine. This situation does not harm downstream recovery but rather results in an amount of wasted energy (the larger recovery due to a smaller product size does not outweigh the additional milling costs). An increase in throughput R would lead to a better distribution of energy per ton of material.

When the ore is harder than expected ($\mathbf{Wi}_t^f$ is underestimated), not enough energy is transferred into each ton of ore. The resulting product will be too coarse. A larger proportion of the product stream will be separated by a hydrocyclone (based on particle size) and recirculated as mill feed. A lower throughput R would reduce the amount of recirculated material, which in turn would result in an increased effective mill throughput (note the paradox).

Installed sensors continuously monitor throughput, power draw, feed and product sizes in the ball mill (Fig. 6.1). The sensor responses have the potential to be used in real-time to derive an actual Operating Work Index value $\mathbf{wi}_t^f$ of the material residing in the mill (Eq. 6.1). The lower case notation refers to an actual measurement (single value), whereas the upper case notation, previously used, indicates an estimate (multiple values in a row vector to characterize uncertainty). The observation $\mathbf{wi}_t^f$ characterizes all material that went through the ball mill between $t-1$ and t. Typically this material will represent a blend of ore from multiple sources and locations.

The online computation of $\mathbf{wi}_t^f$ carries a large potential as demonstrated in the pilot study. Mill observations of $\mathbf{wi}_t^f$ are used to progressively improve the block estimates $\mathbf{Wi}_t^b(n,i)$ in the GeoMet model. This backward integration has only been made possible through material tracking initiatives. Data, from both the fleet management and process monitoring system, are used to link mill observations with their constituent GeoMet blocks. A developed algorithm subsequently updates selected blocks based on the noisy

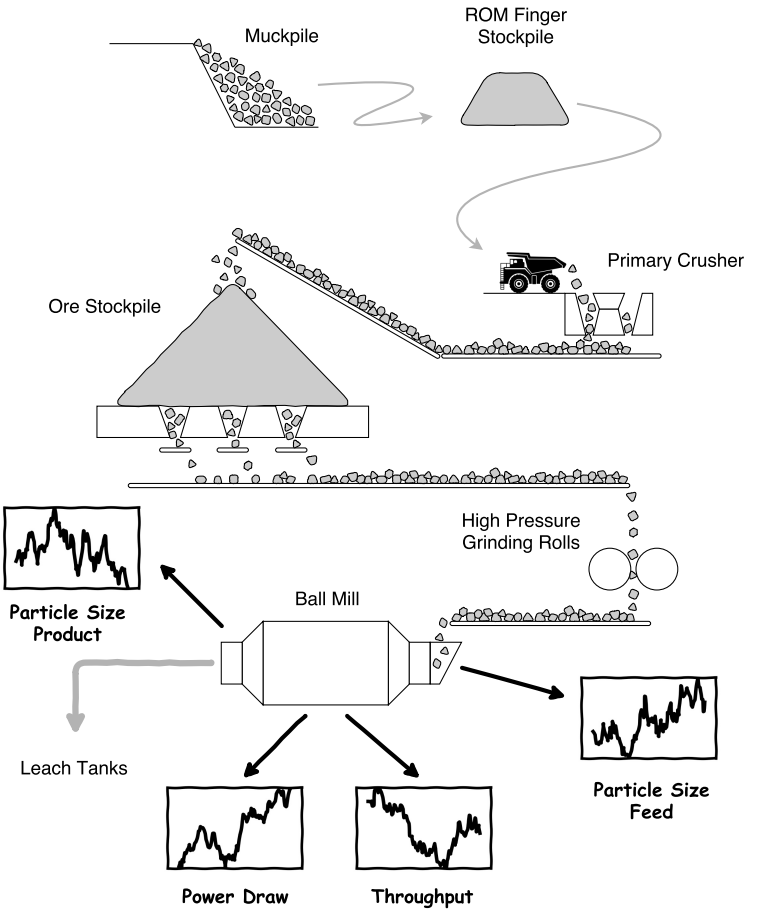

Figure 6.1: Simplified representation of the monitoring setup at the Tropicana Gold Mine. Research question: is it possible to use the ball mill performance measurements to better inform spatial Wi estimates?

time-averaged mill observation. The strength of the algorithm lies in its capability to differentiate between the more and less accurate block estimates constituting the mill feed. Updates aim to more aggressively correct the less accurate block estimates and their surroundings.

In the future, the updating algorithm can be expanded to integrate additional performance data back into the GeoMet model (e.g. recovery and reagent consumption). The methodology eventually could lead to a better and automated selection of ores for blending whilst providing advanced information for process control. For example, the throughput of the comminution circuit can be reduced/increased upfront when harder/softer ore is expected to ensure the most optimal energy utilization, while achieving a grind required for maximising gold recovery in the leach circuit.

This paper demonstrates an algorithm for continuous reconciliation of mill derived observations ($\mathbf{wi}_t^f$) against block estimates in the GeoMet model ($\mathbf{Wi}_t^b$). First, the updating algorithm, as presented in Wambeke and Benndorf (2017), is briefly reviewed. Then, background information is provided regarding the geology at Tropicana, the operation and the available data. Thereafter, a forward simulation model is constructed to convert (updated) block estimates ($\mathbf{Wi}_t^b$) into mill feed estimates ($\mathbf{Wi}_t^f$). The forward simulation step is essential in linking a specific observation ($\mathbf{wi}_t^f$) back to its constituent GeoMet blocks, while providing for a flexible approach to overcome a number of mathematical challenges and material tracking limitations (implementation details, to be discussed later on). Using the forward simulator, the spatial GeoMet models are updated every $4h$ over the course of one week. It is shown that the updates do not only result in a real-time reconciliation of extracted blocks but also significantly improve estimates of scheduled blocks. The paper concludes with an extensive discussion on modelling assumptions and potential improvements.

6.2. UPDATING ALGORITHM

At any point in time, when a new mill observation $\mathbf{wi}_t^f$ becomes available, the updating algorithm needs to solve the following inverse problem:

$$\mathbf{Wi}_t^b = \mathscr{E}_t^{-1}(\mathbf{wi}_t^f), \tag{6.2}$$

where $\mathscr{E}$ is a forward observation model (linear or non-linear) that maps block estimates $\mathbf{Wi}_t^b$ onto mill feed estimates $\mathbf{Wi}_t^f$. In other words, the algorithm is tasked with inferring attributes of individual blocks based on time-averaged mill observations.

The algorithm essentially solves the previous inverse problem using a sequential estimator within a Monte Carlo framework. At time zero, an initial set of I Monte Carlo realizations $\mathbf{Wi}_0^b(:,i)$ is generated using techniques of conditional simulation. All exploration information is inherently accounted for within these initial realizations; (a) sample values are approximated at their respective locations (no exact reproduction in order to account for measurement error), (b) the degree and scale of variability in the realizations follows a pattern described by a covariance model derived from the GC data ($\mathbf{Wi}^s$).

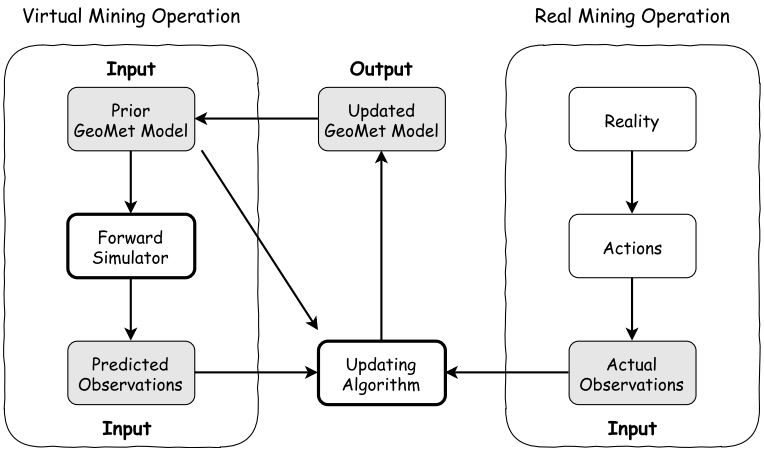

Figure 6.2: Closed-loop reconciliation framework to integrate ball mill performance measurements into the GeoMet model

Once collected, a mill observation $\mathbf{wi}_t^f$ is assimilated into the GeoMet model:

$$\mathbf{Wi}_t^b(:,i) = \mathbf{Wi}_{t-1}^b(:,i) + \mathbf{K}_t(\mathbf{wi}_t^f - \mathbf{Wi}_t^f(1,i)) \quad \forall\, i \in I. \tag{6.3}$$

Each realization $\mathbf{Wi}_{t-1}^b(:,i)$ is updated based on a weighted difference between a mill observation $\mathbf{wi}_t^f$ and a mill feed estimate $\mathbf{Wi}_t^f(1,i) = \mathscr{F}_t(\mathbf{Wi}_{t-1}^b(:,i))$. A mill feed estimate results from running a forward simulator $\mathscr{F}_t$ based on a most recent GeoMet realization $\mathbf{Wi}_{t-1}^b(:,i)$ (to be discussed later on). The latest solution set $\mathbf{Wi}_t^b(:,i)$ $(\forall\, i \in I)$ accounts for all previously collected exploration and production data ($\mathbf{Wi}^s$ and $\mathbf{wi}_t^f, \forall\, t \in [0,t]$).

The simulation-based approach (Monte Carlo framework) avoids the near impossible task of formulating an analytical approximation of the forward observation model $\mathscr{E}_t$ (and calculating its inverse, Eq. 6.2). Due to the complexity of the material handling process, it would indeed be very challenging to describe the link between individual blocks and a blended measurement as a single equation. Instead, for each unique operation, a case specific forward simulator $\mathscr{F}$ is built and run parallel to the more generally applicable updating code (Fig 6.2). The simulator is but a virtual model describing which blocks are extracted, processed and measured between $t-1$ and t (the complexity of the model should match the relevant problem specific features). The separation of the forward simulator from the updating code allows for a flexible integration with existing systems already installed at the mine site.

The forward simulator is thus used to propagate GeoMet realizations $\mathbf{Wi}_{t-1}^b(:,i)$ into mill feed estimates $\mathbf{Wi}_t^f(1,i) = \mathscr{F}_t(\mathbf{Wi}_{t-1}^b(:,i))$ (predicted mill performance). Both realization sets combined contain enough information to link observed deviations back to their constituent GeoMet blocks. This connection is established through the calculation of the Kriging weights $\mathbf{K}_t$:

$$\mathbf{K}_t = \mathbf{C}_{t-1,bf}\mathbf{C}_{t-1,ff}^{-1}, \tag{6.4}$$

where $\mathbf{C}_{t-1,bf}$ (a column vector of size N) and $\mathbf{C}^{-1}_{t-1,ff}$ (a single value) hold the conditional forecast and observation error covariances. Kriging weights tend to be larger when the observation is accurate (i.e. $\mathbf{C}_{t-1,ff}(1,1)$ is low) and strongly correlated to particular blocks (i.e. $\mathbf{C}_{t-1,bf}(n,1)$ is large). The Kriging weights eventually determine how individual block values have to be adjusted in order to shrink the detected deviations (Eq. 6.3).

The necessary covariances are in turn computed empirically from the available realization sets. Each entry $\mathbf{C}_{t-1,bf}(n,1)$ of the covariance vector describes the correlation between the observation and the n^{th} block in the GeoMet model. The covariance $\mathbf{C}_{t-1,ff}$ on the other hand describes the accuracy of the mill observation.

Several technical and practical challenges are solved by computing covariances empirically. (1) As time progresses, conditional forecast error covariances become non-stationary. Were it not that covariances are computed empirically, a large non-stationary field covariance matrix $\mathbf{C}_{t-1,bb}$ would have to propagated from one update cycle to the next (number of entries equal to the square of the number of grid nodes). Computing covariances empirically reduces computation costs and memory requirements (Wambeke and Benndorf, 2017). (2) Differences in scale of support are automatically dealt with. There is no need to perform a support correction on a non-stationary covariance model. (3) Empirical covariances are convenient to handle measurements on blended material streams originating from multiple extraction points. Based on the magnitude of the forecast error covariances, it is possible to pinpoint multiple blocks in the GeoMet model that are responsible for a single detected deviation. Furthermore, the forecast error covariances are of paramount importance in updating neighbouring correlated blocks.

The interested reader is referred to Wambeke and Benndorf (2017) for a detailed literature review and an elaborate presentation of the algorithm. The suggested paper discusses several other aspects of the algorithm which were omitted here.

6.3. BACKGROUND INFORMATION AND DATA SOURCES

This section briefly presents background information about the Tropicana Gold Mine. Aspects related to geology, mining and processing are discussed and relevant data sources are highlighted.

6.3.1. GEOLOGY

The Tropicana Gold Mine is located in Western Australia, approximately $330km$ East-North-East of Kalgoorlie. The mine is situated near the edge of the Great Victoria Desert along an ancient collision zone between the Yilgarn Craton and the Albany Fraser Origen. The regional geology is dominated by granitoid rocks, felsic to mafic paragneiss and orthogneiss, and felsic to ultramafic intrusive and volcano-sedimentary rocks. The area is characterized by extreme weathering that resulted in the formation of a $100m$ thick regolith. Mineralization is found within Archean-aged high grade quartzo-feldspathic gneisses and is associated with late biotite and pyrite alteration. The mineralization occurs as one or two laterally extensive planar lenses with a moderate dip. Post mineralization faulting resulted in four distinct structural domains offsetting the initial ore body.

6.3.2. MINING

The ore body is mined from four contiguous pits extending six kilometres in strike length (from North to South: Boston Shaker, Tropicana, Havana and Havana South). The mine is operated as a typical drill and blast, truck and shovel open pit mine.

Prior to extraction, GC drilling (Reverse Circulation) is completed on relatively dense $10m$ East x $12m$ North drill patterns to map out the ore reserves. The GC holes are drilled to intersect multiple benches at once and are drilled weeks ahead of extraction. The resulting $1m$ samples are sent to an on-site lab for analysis. Conventional fire assaying techniques are applied to determine the gold grade. During the sample preparation stage, samples are processed in an automated sample preparation system, which crushes, splits and pulverises the material prior to XRF and HS scanning. The resulting multivariate interval data are translated into geometallurgical properties using previously calibrated relationships. The inferred geometallurgical properties are subsequently modelled to populate the $3m$ x $3m$ x $3.33m$ blocks of the GeoMet model (used for ore design and short-term planning). Once populated, the mine geologist delineates ore polygons to group adjacent spatial blocks into semi-homogeneous digging volumes (known as dig blocks).

Subsequently during blasting operations, a $10m$ high bench is blasted. Transmitters are installed in blastholes and their locations are logged prior to and after the blast. Three dimensional displacement vectors are computed and applied to a lattice structure. The original undisturbed lattice structure consists of edges and nodes (corners) outlining the $3m$ x $3m$ x $3.33m$ blocks in the GeoMet model. The distorted lattice structure (after the displacement vectors have been applied) accounts for a repositioning of the GeoMet blocks due to blasting.

The fragmented material is then excavated in three passes (based on design flitches with a height of $3.33m$). When a truck is loaded, the position of the excavated material is being recorded (one GPS location per load). Following steps describe how the material characteristics of a truck load are defined:

1. Match the recorded position to a block in the distorted lattice (look-up the block which contains the recorded position).

2. Match the block in the distorted lattice to the corresponding block in the original non-distorted GeoMet model (a unique block identifier can be used to setup a one-to-one relationship between the blocks in both models).

3. Obtain the material characteristics of the selected block in the GeoMet model and assign them to the truck.

Before proceeding, it is important to stress the limitations of the applied correction technique. While the technique allows for differential distortion and movement, it assumes that the material of a block remains within its assigned distorted circumference (relative positioning is maintained). In other words, material cannot be exchanged between two neighboring distorted blocks. In certain circumstances (particular blast patterns and close proximity to the mining face), significant mixing could occur during blasting as lower situated material moves forward, while higher situated material falls

down. In such a situation, the previous assumption might be violated. A more advanced correction technique would need to devised.

Once the truck is loaded, ore is hauled directly to the primary crusher or to one of the reclaim stockpiles situated at the ROM pad (Run-Of-Mine). Direct crusher feed (material directly coming from the mine) is supplemented with ore reclaimed from ROM stockpiles. The fleet management system records each individual truck cycle in a central database. The recorded spatial coordinates are used for material tracking purposes, linking mill observations to their constituent blocks in the GeoMet model.

6.3.3. COMMINUTION

The comminution circuit comprises of a primary crusher, secondary crusher, High Pressure Grinding Rolls (HPGR) and a ball mill. The HPGR screen has a top size of 2.75 mm, resulting in a typical ball mill feed of 500 to 600 μm (F_{80}). The upper part of Fig. 6.4 displays a simplified version of the plant flowsheet (to be discussed later in detail). Conveyor belt speeds, throughput values, recirculating loads, flow velocities and mill performance are continuously monitored. The related sensor readings are written to a database at five minute intervals.

6.4. FORWARD SIMULATOR

A forward simulator is built to generate mill feed estimates $\mathbf{Wi}_t^f(1,i) = \mathscr{F}_t(\mathbf{Wi}_{t-1}^b(:,i))$. The realization set of mill feed estimates is used to compute the empirical covariances, which are essential in linking an observation $\mathbf{wi}_t^f$ to its constituent GeoMet blocks. The forward simulator is subdivided into two connected modules. The first module describes the material handling process in the mine. The second module tracks material flow in the comminution circuit.

6.4.1. FROM PIT TO CRUSHER

The material handling process in the mine can be replicated in great detail using truck cycle data stored in a fleet management database (Fig. 6.3). Four types of truck cycles are defined: (1) ore is hauled from the pit and dumped directly into the primary crusher (direct tip); (2) ore is hauled from the pit and stockpiled on one of the ROM stockpiles; (3) ore is reclaimed from ROM stockpiles and dumped into the primary crusher; (4) material is hauled from the pit to a waste dump (not shown or further discussed).

A type 1/2 mine cycle starts the moment a truck is being loaded in the pit (Fig. 6.3). A *DepartingTruck* is recorded in the **POD** (Pit Out Database). The *BlockID*, referring to a specific block in the GeoMet model, is determined using a combination of GPS data and a distorted lattice which accounts for differential movement due to blasting (ref. sec. 6.3.2). A type 1 cycle ends when a truck finishes tipping its load into the crusher. An *ArrivingTruck* is stored in the **CID** (Crusher In Database). Its *DataBaseLabel* refers to the **POD**.

A type 2 cycle ends the moment a truckload is stockpiled on one of the six ROM stockpiles (Fig. 6.3). An *ArrivingTruck* instance is written to the **RFID** (ROM Finger In Database). The assigned *CellID* refers to a $10m$ x $10m$ subdomain within a larger stockpile and is computed from the GPS location of a tipping truck (each of the six

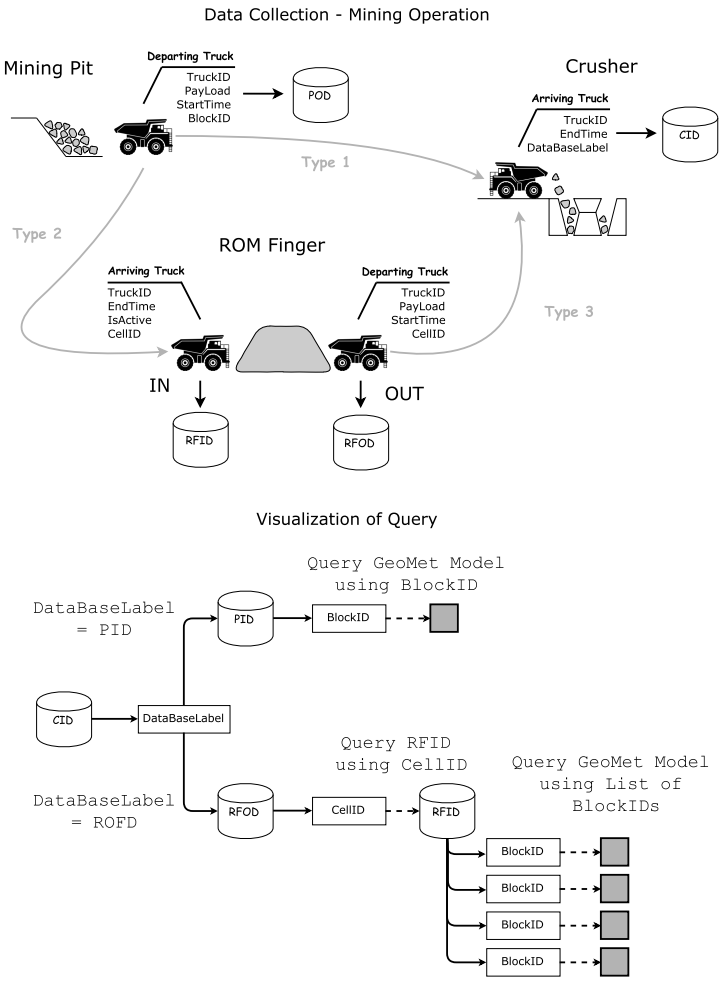

Figure 6.3: Schematic representation of the material handling process in the mine

stockpiles is modelled as a single layer of cells, hence each cell can be represented by its horizontal footpint). All 'active' *ArrivingTruck* instances within a particular ROM finger are deactivated the moment the entire stockpile is depleted (the *IsActive* boolean is set to False).

A type 3 mine cycle is initiated when reclaimed ROM finger material is loaded into a truck. A *DepartingTruck* instance is stored in the **RFOD** (ROM Finger Out Database). The cycle ends when loaded material is tipped into the primary crusher. An *ArrivingTruck* is written to the **CID** (Crusher In Database). Its *DataBaseLabel* points to the **RFOD**.

Querying the databases allows for a live characterization of the crusher feed and stockpile domains. The PayLoad of the trucks arriving at the crusher is obtained from the database referenced by the *DataBaseLabel*. A set of characteristic W_i values (multiple realization to characterize uncertainty) is obtained in one of two possible ways, depending on the assigned *DataBaselabel* (Fig. 6.3).

If the label points to the **POD**, the GeoMet model is queried using the *BlockID* of the corresponding *DepartingTruck* record (one-to-one relationship). The query returns a single set of simulated W_i values which are assigned to a specific truck.

If the label points to the **RFOD**, a more elaborate 'query' needs to be conducted.

1. The *CellID* is obtained from the corresponding DepartingTruck record in the **RFOD** (one-to-one relationship based on *TruckID*).

2. BlockIDs and *PayLoad*s are collected from 'active' *ArrivingTruck* records in the **RFID** (one-to-many relationship based on *CellID*).

3. For each *BlockID*, a set of simulated W_i values is extracted from the GeoMet model (one-to-one relationship).

4. Relative weighting factors are computed based on the obtained *PayLoad* values (each truckload receives a weight proportional to the amount of material it contributed to the stockpile subdomain). The weights are used to calculate a single set of W_i values characterizing the material within a stockpile subdomain.

5. This single set of W_i values can be connected to a truck departing from the ROM stockpiles and in extension thus also to a truck arriving at the crusher.

The module discussed thus far allows for the characterization of individual truckloads arriving at the crusher, independent of whether the material originates from one of the ROM stockpiles or directly from the pit.

6.4.2. FROM CRUSHER TO MILL

The second module of the forward simulator is designed to describe the material flow in the comminution circuit. Four sequential circuits have been identified and modelled: (a) primary crushing, (b) secondary crushing, (c) High Pressure Grinding Rolls (HPGR) and (d) ball mill. The material flow through each circuit is modelled using following the model components (Fig. 6.4); a 'merge' unit (circle) combines material streams, a 'queue' (hourglass) describes delays and a 'split' unit (trapezium) subdivides material streams into two substreams. The behaviour of each virtual unit is driven by information derived from the central processing database (five minute interval readings).

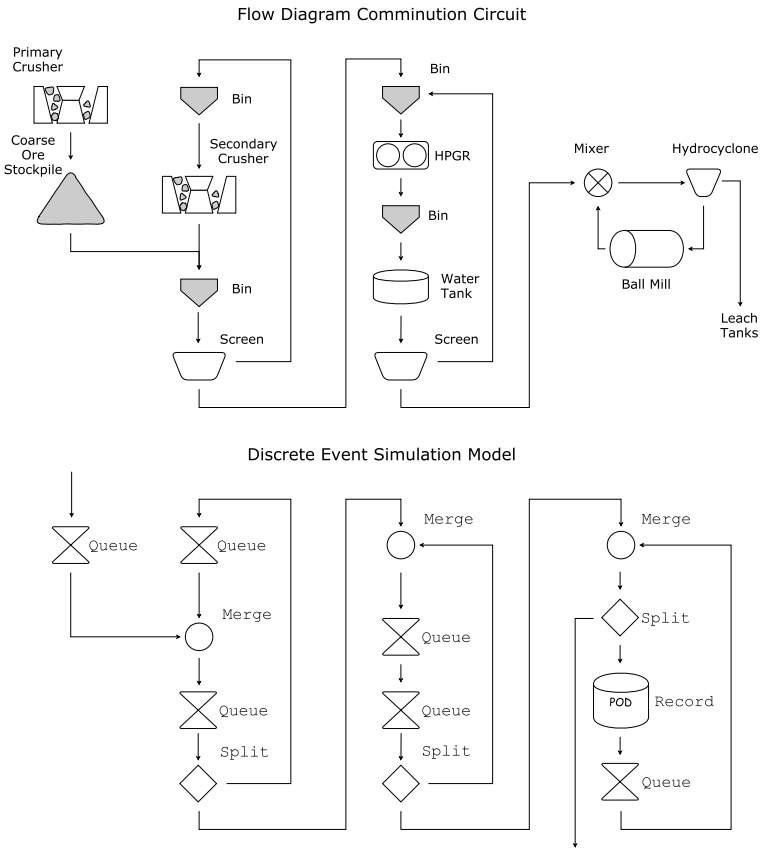

Figure 6.4: Schematic representation of the material flow in the comminution circuit

The moment a truck tips its load into the crusher, its virtual representation is subdivided into a large number of smaller *Parcel*s. Each *Parcel* has a payload and is linked to the *ArrivingTruck*. The sum of all *Parcel* payloads equals the *PayLoad* of the *ArrivingTruck*.

When material enters the comminution plant, it passes through a gyratory crusher and ends up on a coarse ore stockpile (COS). The behaviour of this circuit is modelled as a queue (first in, first out). The delay time of the queue corresponds to the residence time of a *Parcel* on the stockpile (the time to pass the crusher lies in the order of seconds). The popping rate of the queue matches the stockpile drawn down rate readings.

Drawn stockpile material is subsequently blended with the product of the secondary crusher (merge unit). Once blended, material resides in a bin (queue) before being dropped onto a screen (split unit). Oversized material is circulated back to the bin of the secondary crusher (queue), while the undersize is directed towards the HPGR circuit. The split unit randomly selects virtual *Parcel*s and circulates them into the secondary crushing circuit. The recirculating loads match the ones recorded in the processing database.

Arriving at the HPGR circuit, material is blended with screen oversize and stored in a bin, before being dropped into the HPGR (merge unit and queue). The HPGR grinds a loose collection of material into a conglomerated cake product. The cake product resides in a bin awaiting wet screening (queue). Prior to screening (split), the HPGR cake is deagglomerated using water jets and vibration. The screen oversize is circulated back to the bin installed above the HPGR, the undersize enters the milling circuit. Diverter gates to extract tramp metal and to construct emergency stockpiles are not accounted for (future work).

Material arriving from the HPGR circuit is mixed with the ball mill product and inserted into a hydrocyclone (a merge and split unit). The cyclone underflow is circulated back into the ball mill (queue), its overflow is transported to the carbon-in-leach tanks. The moment a virtual *Parcel* arrives in the mill it is recorded in the **MFD** (Mill Feed Database).

The comminution model is by far not accurate enough to track individual truckloads as they move through the plant. Consequently, the mill feed is characterized over $4h$ intervals to filter out possible inaccuracies. A mill feed estimate $\mathbf{Wi}_t^f(1,i) = \mathscr{F}_t(\mathbf{Wi}_{t-1}^b(:,i))$ is obtained as follows:

1. Query the **MFD**, collect all *Parcel* objects arriving at the mill between $t-1$ and t.
2. Group *Parcel* records according to their associated *TruckID*s (truck arriving at the crusher). Compute the total weight of the *Parcel*s within each group.
3. Compute weighting factors based on the group weights.
4. Connect each group to a truck which already arrived at the crusher at an earlier time. Assign the truck W_i values to the group.
5. Compute a weighted averaged set characterizing the mill feed within a $4h$ interval.

In comparison, an actual observation $\mathbf{wi}_t^f$ is obtained by processing $4h$ long time series of measured mill performance (previously discussed).

6.5. RESULTS CASE STUDY

In this pilot study, about 120h of mill performance data are used to conduct 30 updates of the GeoMet model (for a milling period between 21-Aug-2015, 06:00 - 26-Aug-2015, 06:00). A historic dataset is used to enable the validation of mill feed estimates $\mathbf{Wi}_t^f$ against already available mill observations $\mathbf{wi}_t^f$.

The material fed to the mill in the pilot study period, was sourced from ROM stockpiles and direct crusher feed. Due to intermediate stockpiling, the mill feed represents mining activity that occurred over a one month period. Hence, a month of truck cycle data needs to be analysed. At 21:20 on 02-Aug, the first truckload is tipped on a previously zeroed ROM stockpile. During the subsequent three weeks, the building of this finger is carefully tracked. As a result, ROM stockpile material, reclaimed to the crusher during week four, can be characterized. In summary, a month of mining data is required to connect a week of plant performance measurements back to their source locations.

6.5.1. GEOMET MODEL

The material milled between 21-Aug and 26-Aug mainly originates from two distinct benches; bench 2260 in the Tropicana pit and bench 2280 in the Havana pit (Fig. 6.5). Prior to extraction, two GeoMet models are constructed describing the spatial variation of the W_i values in both benches. Each model contains 100 GeoMet realisations $\mathbf{Wi}_t^b(:,i)$ on a block support of 3m x 3m x 3.33m. The field realization are generated using a sequential Gaussian simulation algorithm. All realizations are conditioned on $\mathbf{Wi}^s(m)$ estimates (Fig. 6.5), derived from XRF and HS proxies collected on GC samples. Figure 6.6a shows a horizontal section across the middle flitch of both benches. The figure displays the mean field computed over the 100 prior realisations.

Figures 6.6b to 6.6k illustrate how the mean field changes through time when assimilating mill observations $\mathbf{wi}_t^f$. A total of 30 updates are conducted. The time between updates amounts to 4h. Only results obtained at the end of each shift (i.e every 12h) are shown. The markers on each figure refer to the source locations of the material milled during the indicated time interval (last 12h). As the realizations are updated, the level of detail in the resulting mean field increases. Overall, the algorithm seems to correct for the globally occurring overestimation bias (i.e. in most blocks the W_i is lowered). However at specific locations, the algorithm learns that the ore is harder than initially expected (W_i values are increased).

6.5.2. MILL FEED ESTIMATES

Updated GeoMet models $\mathbf{Wi}_{t-1}^b$ are continuously propagated through the forward simulator $\mathscr{F}_t$ to adjust mill feed estimates $\mathbf{Wi}_t^f$ accordingly. The top axes of Figs. 6.7a to 6.7k display groups of mill feed estimates at 11 different timesteps (ref. time label in the top left corner). Each group of mill feed estimates (single plot) characterizes 30 distinct 4h long mill feeds (blue boxplots). Certain mill feeds are already processed (grey area); one is currently being milled and measured (yellow bar); others still have to be fed to mill (right of yellow bar). Once a 4h interval ends, a new observation $\mathbf{wi}_t^f$ becomes available

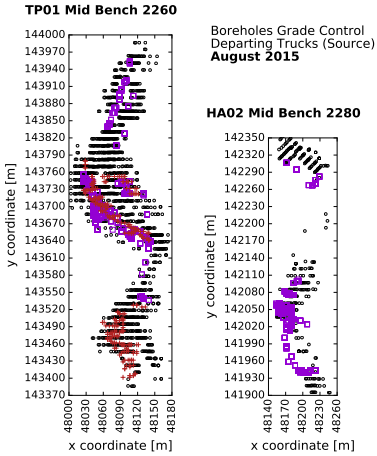

Figure 6.5: truck pit source locations from Tropicana and Havana pit during the entire month of August (2015). The displayed trucks are either dispatched to the crusher (red crosses - direct tip) or to ROM stockpile 4 (orange squares). The black dots display GC holes intersecting the bench.

(red bar) and an update is performed. Before switching to the next time interval, the forward simulator readjusts all predicted measurements.

Since the pilot study is based on a historic dataset, all future observations are known ahead of time. As a result assessment statistics can be computed not only describing improvements in historic predictions but also in future ones.

The bottom axis in Figs. 6.7b to 6.7k displays the Differences in Absolute Error (DAE) between best estimates (blue horizontal line in boxplot) and true measurements relative to time 0 (21-Aug 06:00). For example a DAE of 2 indicates that the best estimate moved 2 units towards the actual value. The bars are coloured green in case of improvements, red otherwise. Deteriorations do occasionally occur but their magnitude stays within acceptable limits.

The change in Root Mean Square Error ($CRMSE$) between predicted and actual measurements is computed within different dynamic time windows. The first and second window focuses on the upcoming $12h$ and $24h$ ($12h$ and $24h$ ahead of the yellow bar). The third window encompasses all historic $4h$ intervals (grey area). Future predictions improve on average by about 26% (next $12h$) and 22% (next $24h$). The error in historic estimates reduces on average by about 72%. A correction of 100% is not desired since time-averaged noisy observations do not contain enough information to fully eliminate all remaining inaccuracies in the blended blocks.

6.6. CONCLUSIONS

This chapter describes the pilot testing of a novel updating algorithm at the Tropicana Gold Mine. During the pilot, on-line mill observations are automatically reconciled against the spatial work index estimates of the GeoMet model. Deviations between

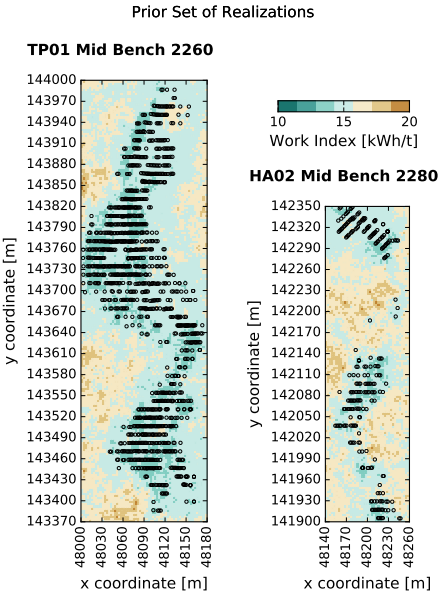

(a) $0h$

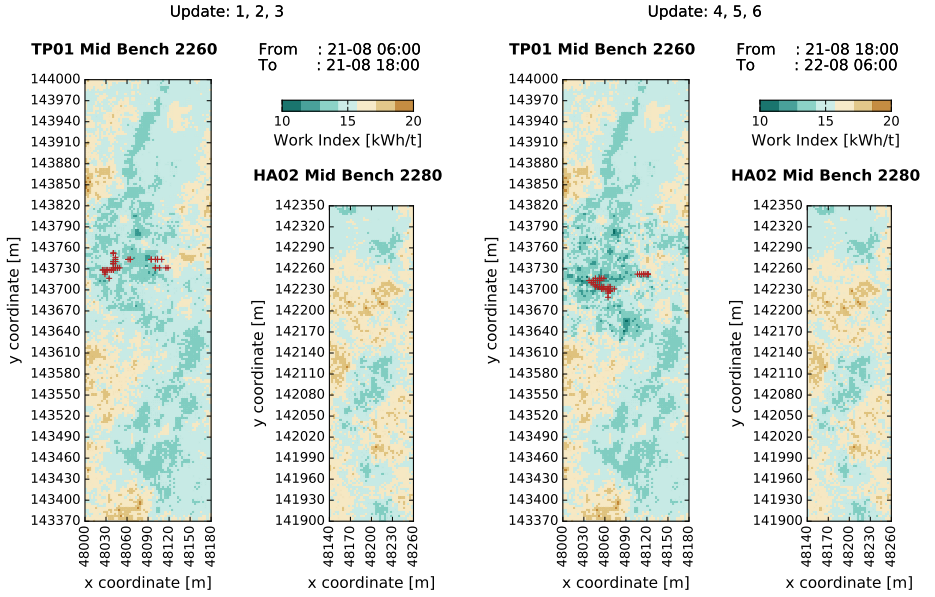

(b) $12h$ (c) $24h$

Figure 6.6: Mean field across the middle flitch of two benches in the Tropicana and Havana pit. Blocks are colour coded according to their best estimate at the indicated time. The source locations of the material milled during the indicated time interval are displayed using either red crosses (direct tip) or orange squares (material has resided on a stockpile finger)

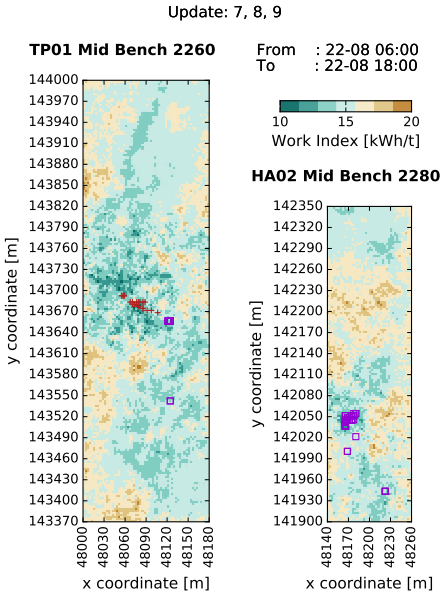

(d) 36h

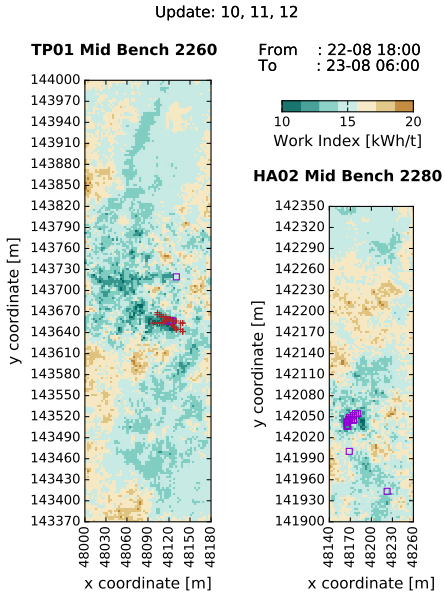

(e) 48h

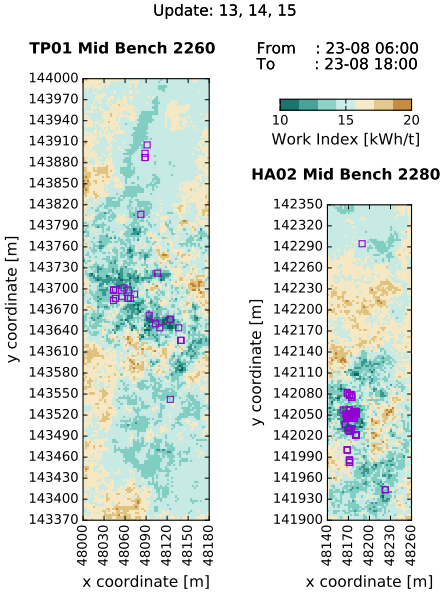

(f) 60h

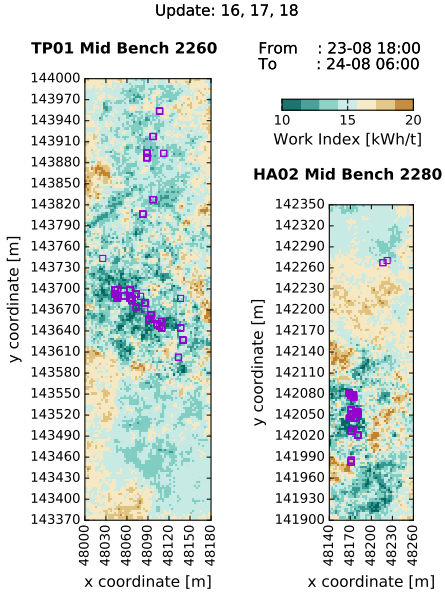

(g) 72h

Figure 6.6: CONTINUED (2) - Mean field across the middle flitch of two benches in the Tropicana and Havana pit. Blocks are colour coded according to their best estimate at the indicated time. The source locations of the material milled during the indicated time interval are displayed using either red crosses (direct tip) or orange squares (material has resided on a stockpile finger)

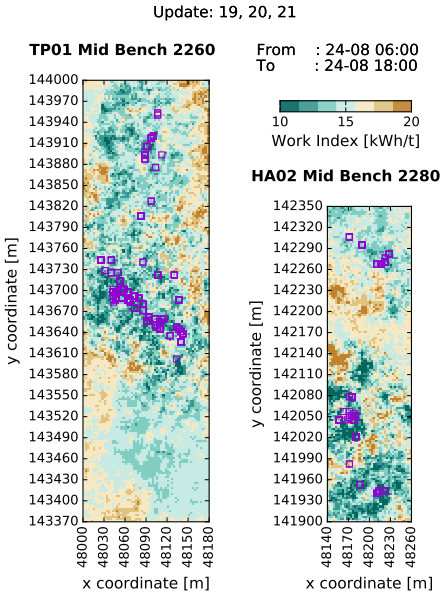

(h) 84h

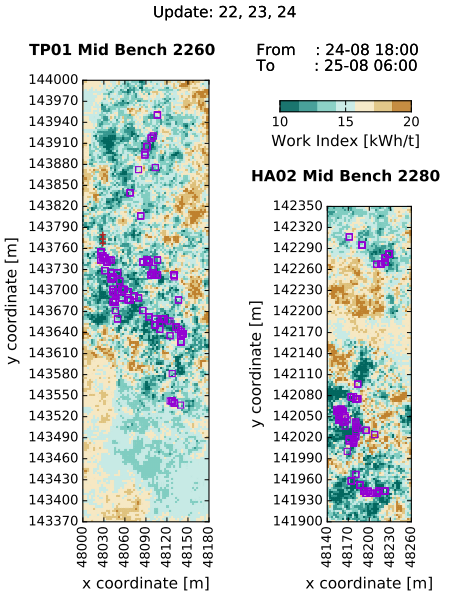

(i) 96h

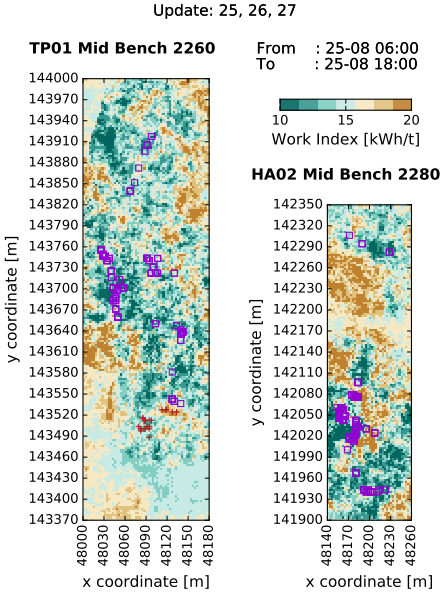

(j) 108h

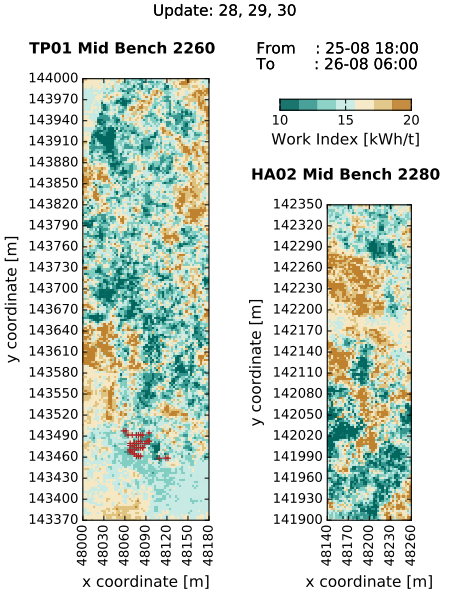

(k) 120h

Figure 6.6: CONTINUED (3) - Mean field across the middle flitch of two benches in the Tropicana and Havana pit. Blocks are colour coded according to their best estimate at the indicated time. The source locations of the material milled during the indicated time interval are displayed using either red crosses (direct tip) or orange squares (material has resided on a stockpile finger)

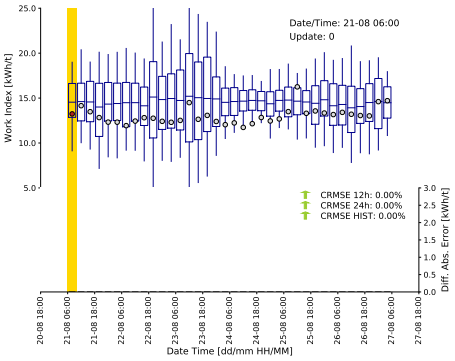

(a) $0h$

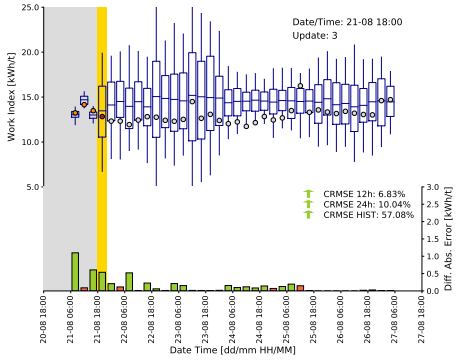

(b) $12h$

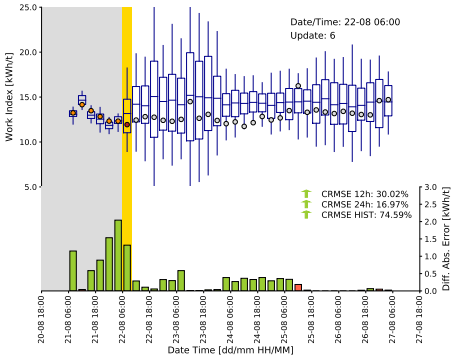

(c) $24h$

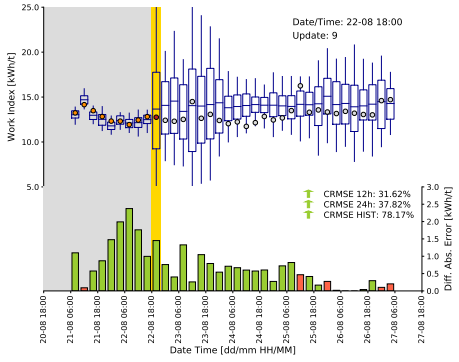

(d) $36h$

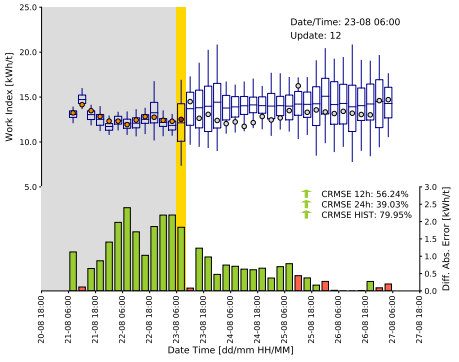

(e) $48h$

Figure 6.7: Predicted (blue) and actual time-averaged measurements (red). The lower axis displays difference in absolute error (DAE) relative to time 0. Changes in $RMSE$ between predicted and actual measurements, computed for three dynamic time windows, are shown in the bottom right corner

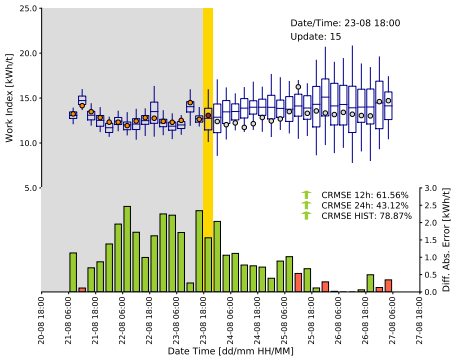

(f) 60h

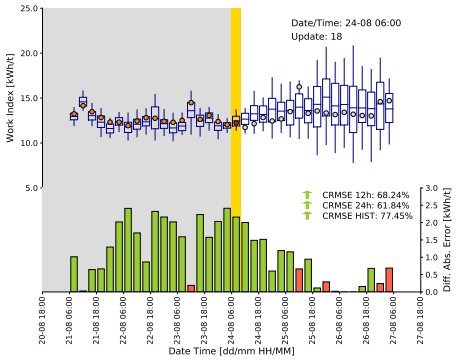

(g) 72h

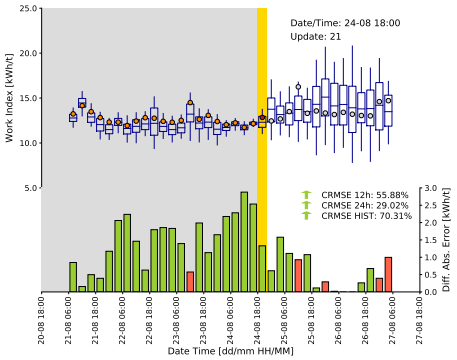

(h) 84h

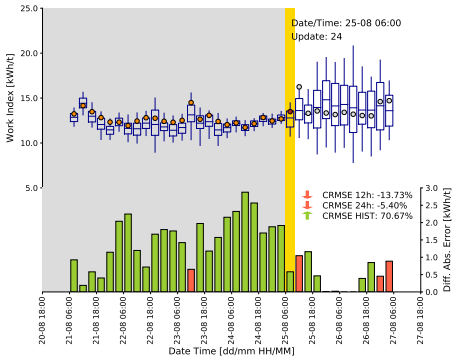

(i) 96h

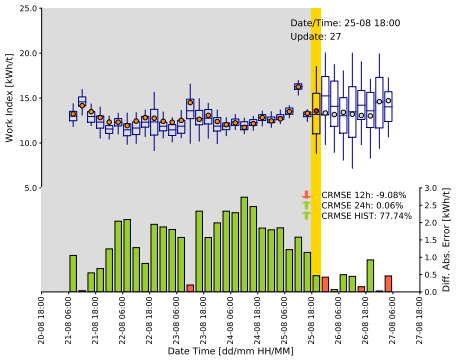

(j) 108h

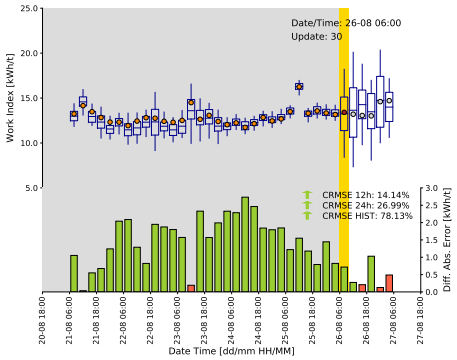

(k) 120h

Figure 6.7: CONTINUED (2) - Predicted (blue) and actual time-averaged measurements (red). The lower axis displays difference in absolute error (DAE) relative to time 0. Changes in $RMSE$ between predicted and actual measurements, computed for three dynamic time windows, are shown in the bottom right corner

predicted and actual mill performance are monitored and used to locally improve the GeoMet model. The novelty of the approach resides in its ability to trace detected deviations back to the predominant source. The algorithm automatically handles differences in scale of support, sensor inaccuracies and observations made on blended material originating from two or more extraction points.

In order to operate the updating algorithm, actual observations are to be compared against model-based expectations (the mill feed estimates). The model-based expectations result from the propagation of GeoMet realizations through a forward simulator. The resulting realization sets (block and mill feed estimates)are subsequently used to compute empirical covariances. The covariances describe the link between mill derived observations and blocks from the GeoMet model. There is no need to formulate and linearise an analytical forward observation model, let alone compute its inverse.

A total of 30 updates were performed to assimilate a week of mill performance data into the GeoMet model. The level of detail in the mean field increases significantly as the GeoMet realizations are updated. The algorithm corrects various local estimation biases resulting from ill-calibrated relationships between sensor data (XRF and NIR measurements) and work index values.

The obtained results are validated against readily available production measurements. Since the pilot was run off-line, future mill observations are known ahead of time (although this information is ignored during updating). Consequently, validation statistics are computed to evaluate whether updated GeoMet models result in more accurate mill feed estimates. Over the course of a week, the $RMSE$ between predicted and measured Word Index values drops by about 72%. The results further indicate that updating causes on average a reduction in error of about 26% in performance forecasts for the next shift (upcoming $12h$). Improvement in future forecasts of up to 68% have been observed.

Although the current implementation of the forward simulator does adequately describe the relevant operational features, further work needs to be done.

- The entire updating approach applied during the pilot study relies on the assumption that material can be tracked with sufficient accuracy during $4h$ long time intervals. The tracking assumptions underpinning the forward simulator are to be validated using e.g. RFID tags (radio frequency identification). These tags have to be placed within the source material (e.g. blastholes, muckpiles or stockpiles) and their initial location needs to be logged. Several antennae have to be installed at critical positions in the comminution circuit. Subsequently, events of tags passing an antenna are recorded. Finally these recorded events will have to be compared against the tracking behaviour of the forward simulator. A sufficient agreement will validate the tracking assumptions and the consequential behaviour of the forward simulator.

- The material source locations can be more accurately defined. Currently only the GPS locations of the trucks are recorded. Aerial photographs are used to derive a set of correction vectors (length and orientation) linking truck positions to actual digging locations. Very recently, the high-precision GPS locations of the loaders and excavators were made available as well. An algorithm can be written

to accurately determine the origin of each bucket loaded into the truck.

- A support correction algorithm needs to be designed and implemented to adjust the distribution of W_i values assigned to a truck when material is reclaimed from a stockpile subdomain. Some random noise proportional to the amount of volume reduction should be added to the characterizing set. As a result, the uncertainty in the truck would be larger than the one in the stockpile subdomain. The statistical consequences of this volume reduction are currently ignored.

- Equipment performance measurements might start to drift as critical components are wearing out (e.g. liners in the ball mill). Machine learning techniques could be applied to automatically correct for the occurring drift.

- Since the case study is based on an off-line execution of the algorithm, future predictions were generated using a 'tracking-based' forward simulator. When the algorithm would run on-line, the tracking data would only allow to generate model-based equivalents to current and historic measurements. A second 'schedule-based' forward simulator needs to be built to generate production forecasts. The production forecasts are by no means necessary to run the updating algorithm.

- The application of a 'schedule-based' forward simulator allows for a continuous re-evaluation of operational decisions based on the most up to date information. Production forecasts should be recorded and validated against performance measurements as soon as they become available. As such, the performance of the algorithm with respect to generating accurate forecasts is continuously monitored. The performance regarding reconciliation of historic measurements should obviously be monitored as well.

Future development should further focus on extending the capabilities of the updating algorithm. The current implementation is designed to update a single continuous attribute (spatial work index) based on a single continuous measurement variable (mill derived work index). This measurement variable is either directly or indirectly related to the attribute of interest. The algorithm needs to be extended to handle multivariate attributes, both in the block model as on the measurement side (updating multiple GeoMet estimates simultaneously). That is, correlated measurement variables need to be jointly considered to update (other) correlated attributes. Neglecting to do so will result in a loss of information. Additionally, it would be interesting to update categorical variables such as ore types, lithologies or weathering zones, based on equipment performance or other sensor measurements.

The application of machine learning algorithms to regularly recalibrate the relationships between sensor data and work index values opens up another avenue of research. A large database of reconciled work index values is built up as the updating algorithm is operated over a longer period of time. Each resulting block value is then to be linked with sensor data from neighbouring grade control samples. Both datasets are then to be used to regularly retrain the relationship between sensor data and work index estimates. The retrained relationship will eventually result in more accurate and reliable GeoMet

models. Finally, the gained knowledge, e.g. a better understanding of the relation between ore types and hardness (assuming sensor data can be used to differentiate between ore types), should be transferred into the longer-term resource model on a regular basis.

REFERENCES

T. Wambeke, D. Elder, A. Miller, J. Benndorf, and R. Peattie, *Real-time reconciliation of the geomet model based on ball mill performance measurements - a pilot study at the tropicana gold mine*. Transactions of the Institutions of Mining and Metallurgy - Section A, Mining Technogloy (under review) , 1–21 (2017).

J. Benndorf, *Application of efficient methods of conditional simulation for optimising coal blending strategies in large continuous open pit mining operations,* International Journal of Coal Geology **112**, 141 – 153 (2013).

R. Peattie and R. Dimitrakopoulos, *Forecasting recoverable ore reserves and their uncertainty at morila gold deposit, mali: An efficient simulation approach and future grade control drilling,* Mathematical Geosciences **45**, 1005–1020 (2013).

R. Dimitrakopoulos and M. Godoy, *Grade control based on economic ore/waste classification functions and stochastic simulations: examples, comparisons and applications,* Mining Technology **123**, 90–106 (2014).

B. Catto, *Tropicana Geometallurgy Programme - Integrated Planning*, Tech. Rep. (Anglo Gold Ashanti Australia, Tropicana Operation, 2015).

A. Lynch, A. Mainza, and S. Morell, *Comminution handbook, spectrum series 21,* (The Australian Institute of Mining and Metallurgy, 2015) Chap. Ore Comminution Measurement techniques, pp. 43–60.

T. Wambeke and J. Benndorf, *A simulation-based geostatistical approach to real-time reconciliation of the grade control model,* Mathematical Geosciences **49**, 1–37 (2017).

7

CONCLUSIONS AND RECOMMENDATIONS

The last chapter of the dissertation presents a brief overview of the general conclusions from this research. Thereafter, the chapter lists recommendations and potential research areas to further improve and develop the "updating loop" within the overall "Closed-Loop Resource Management" framework. Note that specific conclusions and recommendations are provided at the end chapters 3 (sec. 3.6), 4 (sec. 4.3), 5 (sec. 5.7) and 6 (sec. 6.6).

Parts of this chapter have been published in Mathematical Geosciences **49**, 1 (Wambeke and Benndorf, 2017a) or have been submitted for publication in Mathematical Geosciences **accepted with revisions** (Wambeke and Benndorf, 2017b) or in Transactions of the Institutions of Mining and Metallurgy, SectionA - Mining Technology **accepted, in press** (Wambeke *et al.*, 2017).

The goal of this research is to initiate the development of an updating framework capable of assimilating online sensor observations into an already existing resource model. The objectives, formulated in Chapter 1 (p. 7-8), are addressed as follows:

- **Objective 1**: A literature review compares the current state of the art in updating or conditioning spatial models based on additional production data. The Ensemble Kalman Filter is selected as the most suitable technique since it is capable of executing sequential updates (obj. 1a), does not require a closed-form analytical expression of the forward observation model (obj. 1b) and is able to handle noisy blended measurements on various scales of support (obj. 1c).

- **Objective 2**: The Ensemble Kalman Filter is subsequently integrated into an overall mathematical formulation and algorithm design. Several additional features and options are added to cope with the requirements of the mining industry: (a) a Gaussian Anamorphosis deals with non-Gaussian distributions and indirect observations, (b) a computation confined within the boundaries of a neighbourhood is faster and requires less memory, (c) a parallel updating sequence (double helix) and two covariance correction techniques contain the propagation of numerical errors, which do originate from the use of a limited number of spatial realizations (a larger number would simply be impractical), (d) the specific implementation of the overall framework allows for an intuitive treatment of measurement error.

- **Objective 3**: The performance of the algorithm is studied extensively in an artificial environment, representing a mining operation with two extraction points of unequal production rate. A visual inspection of cross sections teaches that the best estimates is significantly altered when production data is being assimilated into the model. The global assessment statistics further illustrate that the overall quality of the resource model improves over time. The local assessment statistics show that the observed global improvements result from correcting local estimation biases.

- **Objective 4**: The performance of the updating algorithm as a function of the design and operational control of the monitoring system is examined. A detailed analysis of the results and the subsequent conclusions are presented in great detail in Chapter 5.

- **Objective 5**: The developed updating concepts is tested during a pilot study in an operational mine. The pilot test demonstrates that it is possible to use online mill observations to automatically reconcile the spatial Work Index estimates in the GeoMet model. The error in historic and future production estimates drop on average by respectively 72% and 26%.

The following section presents a general synopsis and conclusion. Thereafter, future recommendations regarding the implementation, performance and applicability of the developed updating framework are provided.

7.1. SYNOPSIS AND CONCLUSION

In order to make the best possible decisions, it is vital that the resource model is always entirely up to date. This requires a practical updating algorithm which continuously assimilates new information into an already existing resource model.

The updating algorithm must address following practical considerations. (a) At each point in time, the latest solution has to implicitly account for all previously integrated data (sequential approach). During the next update, the already existing resource model has to be further adjusted to honour the newly obtained observations as well. (b) Due to the nature of a mining operation, it is nearly impossible to formulate closed-form analytical expressions describing the relationship between observations and resource blocks. The relevant relationships will rather have to be inferred merely from input (the resource model realizations) and the output (distribution of predicted observations) of a forward simulator. (c) The updating algorithm must be able to assimilate noisy observations made on a blend of material originating from multiple sources and locations. Differences in scale of support will have to be dealt with automatically.

Several existing (geo)statistical techniques are individually capable of solving at least one but not all of the practical requirements identified. Co-Kriging approaches for example can be used to jointly model the subsurface using both scarce direct and more abundant indirect measurements. The underlying equations are designed to handle differences in accuracy and measurement volume. They do however fail to take advantage of measurements made on blended material streams. Furthermore, the conventional co-Kriging approaches are not able to sequentially incorporate new information into an already existing resource model.

The latter issue is resolved by merging the co-Kriging equations into a sequential linear estimator. The resulting technique improves previous resource models by using linearly weighted sums of differences between observations and model-based predictions (forward simulator output). The weights, necessary to conduct each update, result from solving a system of co-Kriging equations. These kriging weights eventually determine how individual block values have to be adjusted in order to shrink the detected deviations.

The blending issue is elegantly addressed through the application of Monte Carlo principles. The covariances necessary to compute the Kriging weights are empirically derived from two sets of Monte Carlo samples; the resource model realizations (input forward simulator) and the observation realizations (output forward simulator). There is no need to provide analytical formulae modelling spatial correlations, blending and differences in scale of support.

Combined, the selected (geo)statistical techniques are capable of updating existing resource models based on noisy measurements of blended material streams. The governing equations bear some remarkable resemblance to those of a dynamic filter (the Ensemble Kalman Filter). Their underlying philosophies however do differ significantly. An application in the minerals industry is by no means comparable to those in various other research areas. Weather forecasting and reservoir modelling (oil, gas and water), for example, consider dynamic systems repetitively sampled at the same locations. Generally, each observation characterizes a volume surrounding a sample location.

Mineral resource modelling on the other hand focuses on static systems gradually sampled at different locations. Each observation is characteristic for a blend of material originating from one or more locations and sources. Each part of the material stream is sampled only once, at the moment it passes the sensor in the production chain. Although the equations are virtually the same, the terminology of sequential updating or conditioning is preferred over (Ensemble) Kalman filtering. This is to stress the absence of dynamic components and to highlight its link to the field of geostatistics.

The designed geostatistical approach is implemented into an efficient algorithm capable of automatically assimilating online production data. The practicality of the algorithm mainly results from the design decision to exclude the forward simulation model from the computer code. A case specific forward simulator is built and run parallel to the more generally applicable updating code. This results results in great flexibility, allows for a better integration of expert knowledge and facilitates interdisciplinary cooperation.

The forward simulator is run to translate resource model realizations into observation realizations. The resulting realization sets are subsequently used to compute empirical covariances. These covariances mathematically describe the link between sensor observations and individual blocks from the resource model. This numerical inference avoids the cumbersome task of formulating, linearising and inverting an analytical forward observation model. The application of a forward simulator further ensures that the distribution of the Monte Carlo samples already reflect the support of the concerned random values. As a result, the necessary covariances, derived from these Monte Carlo samples, inherently account for differences in scale of support.

Various options are integrated into the updating algorithm to either reduce computation time, memory requirements or numerical inaccuracies. (a) A Gaussian anamorphosis option is included into the algorithm to deal with suboptimal conditions related to non-Gaussian distributions. A specific algorithm structure ensures that the sensor precision (measurement error) can be defined in its original units and does not need to be translated into a normal score equivalent. (b) Numerical inaccuracies may arise when the number of Monte Carlo samples is kept rather small for practical reasons. Since covariances are computed empirically, degrees of freedom are lost over time. Eventually this might result in a collapse of covariances (filter inbreeding). To prevent this from happening, an interconnected parallel updating sequence (helix) can be configured. (c) A neighbourhood option is implemented to constrain computation time and memory requirements. Since observations are collected on blended material streams originating from multiple extraction zone, different neighbourhoods need to be considered simultaneously. A neighbourhood strategy also partially reduces the effect of spurious correlations. (d) Two covariance correction options are implemented to inhibit the propagation of statistical sampling errors originating from the empirical computation of covariances. The localization correction is based on the assumption that the accuracy of a covariance estimate decreases with distance from an extraction point. Hence, grid anomalies are tapered around central extraction locations. The error correction uses an offline Monte Carlo simulation to precompute the accuracy of a certain correlation estimate. Corresponding correction factors are stored in a lookup table and used to adjust the elements of the forecast error covariance matrix.

An artificial experiment is conducted to extensively test the developed algorithm and to showcase its capabilities. A mining operation with two extraction points of unequal production rate is simulated. The resulting material streams are blended and inaccurate observations are made. A total of 12 updates are performed, each one based on 2 inaccurate observations of 16 blended blocks. The $RMSE$ inside both extraction zones drops by 38% and 45%. Improvements ranging from 39% to 13% were observed in a close neighbourhood. The results are very promising, especially considering the large measurement volume and low sensor accuracy.

Another 125 artificial experiments are conducted to study how variations in measurement volume, blending ratio and sensor precision affect the overall performance of the updating algorithm. Each experiment mimics a virtual mining operation with two extraction points. Over the course of each $144h$ long experiment, a total of $403200t$ is excavated, processed and measured. Despite their mutual differences, all mining schedules ensure that at the end of each experiment the same 1152 blocks are extracted. Local production rates (blending ratios) and interval durations (measurement volumes) do however vary significantly across the experiments. Moreover various sensors with a different precision are applied.

Results regarding model improvements are promising. Under fairly optimal conditions, the $RMSE$ in the resource model drops by around 60%. Optimal conditions occur when small operational volumes are extracted from a single location and characterized with a relatively accurate sensor (8 GC blocks - $2800t$, blending ratio of 100%, sensor precision of 0.05). Even when conditions are far from optimal, the $RMSE$ in the resource models can still be reduced by about 20%. This is remarkable, especially considering the large measurement volumes, the amount of blending and the extremely inaccurate sensor (64 GC blocks - $22400t$, blending ratio of 50%, sensor precision of 1.00). The results indicate that even global inaccurate blended averages have some information potential to further improve the resource model.

The experiments further illustrate the outstanding reconciliation capabilities of the updating algorithm. The $RMSE$ between historic (re)predicted and actual measurements decreases between 46% and 97% over the course of an experiment. More accurate sensors and shorter measurement intervals significantly improve the reconciliation behaviour. Larger $RMSE$ reductions are further observed when local production rates are sitting at either end of the spectrum (blending ratio of 100% or 50%).

Improvements regarding future predicted measurements (next $24h$) turn out to be less robust against variations in system parameters. Average $RMSE$ reductions vary between -5% and $+30\%$. Longer interval durations (measurement volume) harm performance. The precision of the applied sensor does not seem to significantly impact the results. The overall effect of variations in local production rates (blending ratio) is harder to describe. Converging local production rates generally result in lower $RMSE$ reductions. This statement however only seems to apply as long as the local production rates are noticeably different. The moment they are equal (blending ratio of 50%), the $RMSE$ reduction becomes exceptionally large.

Nearly all results have shown that shorter interval durations have a positive impact on the algorithm performance, no matter the considered metric. Due to the complexity of the material handling process, it might be challenging to construct a forward

simulator accurate enough to track the material on a scale as small as a few blocks. A larger time interval might have to be selected to ensure that a forward predicted measurement and a time-averaged sensor response are actually characterizing roughly the same material (ensure a correct match between a virtual and real measurement). A representative forward prediction is essential in establishing an accurate link between specific observations and the right constituent blocks. As measurement volumes grow too large, the opportunity to improve future predictions might be lost.

It further would be wise to install by default the most accurate sensor available. First of all, a more accurate sensor significantly improves the reconciliation of production data. Secondly, even in combination with medium measurement volumes, inaccuracies in the resource model diminish when more accurate sensors are applied.

In practice, it will be nearly impossible to control local production rates or blending ratios for the sake of improving the algorithm performance. These parameters rather follow from a defined mining schedule and are to be considered as external input. Furthermore, local production rates and the corresponding blending regime will very likely change over time. The previous results nevertheless form a basis for understanding and anticipating changes in algorithm performance associated with adjustments in local production rates. Before proceeding, it important to stress that the algorithm performance can not be evaluated solely based on the blending ratio in itself. Differences in performance can only be explained by comparing the more inclusive mining schedules. It turns out that it is not the blending per se that matters but rather its impact on consequential covariances, necessary to compute the Kriging weights.

The mining schedule (in combination with the sensor precision) directly affects the magnitudes of following covariances; (1) Cross-covariances between extracted blocks; (2) Cross-covariances between scheduled blocks; (3) Cross-covariances between neighbouring blocks (not yet scheduled); (4) Covariances between extracted and scheduled blocks; (5) Covariances between extracted and neighbouring blocks. The interplay between these covariances eventually determines the success of resource model updates and their subsequent impact on future predicted measurements (via scheduled blocks).

When purely evaluating potential resource model improvements, covariances of type 1, 3 and 5 have to be considered. Mine schedules with converging local production rates (i.e. an increase of blending from two different sources) generally harm the quality of resource model improvements. This is mainly because the number of zero-correlated pairs of blocks increases as more material is being blended from two distant areas. In other words, more covariances of type 1, 3 and 5 amount to zero. The non-zero covariances do change as well with different mining schedules. Yet, their impact is believed to be of lesser importance.

On the other hand, when the aim is to anticipate improvements in future predicted measurements, covariances of type 2 and 4 have to be studied as well. Type 2 covariances increasingly become zero as more blocks are being blended from two different areas. Non-zero type 2 covariances increase or decrease depending on the specifics of the corresponding mining schedule. The results however indicate that the behaviour, with respect to improvements in future predictions, is mainly driven by type 4 covariances. These covariances entirely depend on the distance (and the associated geological correlation) between extracted and scheduled blocks. The schedules, describing the blending

of equal material amounts (blending ratio of 50%), apparently resulted in the lowest average distance. Hence, the associated experiments outperformed their counterparts with a lower degree of blending.

This dissertation further presents the results of a pilot test conducted at the Tropicana Gold Mine in Australia. The aim of the study is to investigate whether spatial Work Index estimates of a GeoMet model can be updated using actual mill performance data.

A total of 30 updates were performed to assimilate a week of mill performance data into the GeoMet model. The level of detail in the mean field increases significantly as the GeoMet realizations are updated. The algorithm corrects various local estimation biases resulting from ill-calibrated relationships between sensor data (XRF and NIR measurements) and Work Index values.

The obtained results are validated against readily available production measurements. Since the pilot was ran off-line, future mill observations are known ahead of time (although this information is ignored during updating). Consequently, validation statistics are computed to evaluate whether updated GeoMet models result in more accurate mill feed estimates. Over the course of a week, the $RMSE$ between predicted and measured Word Index values drops by about 72%. The results further indicate that updating causes on average a reduction in error of about 26% in performance forecasts for the next shift (upcoming $12h$). Improvements in future forecasts of up to 68% have been observed.

7.2. RECOMMENDATIONS FOR FUTURE WORK

This dissertation presents the results of a first attempt to develop an approach to real-time updating of resource models using online production data. The initial results are very promising. This work opens up a potential novel research avenue within the field of mining geotatistics. Hence, further research needs to be conducted to bring the developed methodology to a required mature level. The remainder of this chapter suggests some interesting new research directions which focus on (a) extending the existing updating algorithm or implementing new ones (b) conducting additional experiments to understand the effect of the prior model on subsequent updates (c) addressing the modelling assumptions behind the forward simulator and suggesting potential improvements.

7.2.1. THE UPDATING APPROACH

There is some potential to further enhance the localization error correction option (Gaspari-Cohn function, ref. Chapter 3) of the existing updating algorithm. The current implementation results in a differential covariance correction of blocks located in the same digging area. A block near the centre of the extraction area (specified as a point coordinate $\mathbf{x}_l$, ref. Algm. 1, p. 29) receives a lower correction factor than a block slightly further away. Especially, when considering larger measurement volumes (larger digging blocks), a tapering function with a plateau might be more suitable. The function should be defined such that all blocks inside the digging area are not corrected (correction factor of 1). Tapering then starts from the border onwards into the surrounding neighbourhood. Consequentially, some inputs will have to adjusted.

The central coordinates of the extraction area will have to be replaced with polygons delineating the digging volumes.

Another point of attention could be the numerical stochastic representation of a Gaussian anamorphosis function. During most forward transformations (those of grid nodes and model-based predictions, Eq. 3.1, 3.3), a non-linear non-parametric operator is used to map Monte Carlo samples onto their normal score equivalents. The forward transformation of actual observations (Eq. 3.4) and the backward-transformation of updated grid nodes (Eq. 3.10) requires a bit more care. Since these values of interest are not all included in the numerical representation of the corresponding Gaussian anamorphosis function, there is a need to establish rules of interpolation. If the value is located inside the limits of the Monte Carlo sample, a straightforward linear interpolation suffices. Outside the limits, the upper and lower tails are extrapolated using a power model. These extrapolations require tedious manual input. Before each update, the user has to specify a reasonable minimum and maximum value for the considered resource attributes and model-based predictions (a reasonable range might change when moving into an area with substantially higher or lower values). Polynomial curve fitting techniques can be applied to provide a continuous representation of the Gaussian anamorphosis function. Such an approach would circumvent the need to define this manual input.

It further would be interesting to develop some statistics to monitor the 'health' of a sequentially updated resource model. Once developed, these statistics could be used to identify malicious updates, which can then be rolled back. The resulting monitoring tools should detect a covariance collapse, a disappearance of spatial correlation structure and unusually large or small blocks values. Furthermore, predicted observations should be recorded and validated against actual measurements as soon as they become available. As such, the performance with respect to generating more accurate forecasts is continually monitored. The performance regarding reconciliation of historic measurements should obviously be monitored as well.

Besides monitoring the performance and 'health' of a sequence of updates, it can be helpful to display the computed empirical covariances. A visual representation of these covariances will be useful to learn and study the differences between seemingly similar observations. The observation error covariances (Eq. 3.6) could represent the joint-information content of a collection of observations, while the forecast error covariances (Eq. 3.7) could identify the blocks which need to be updated.

Another type of statistic can be developed to anticipate changes in future predicted measurements. A new statistic, preferably a single-value measure, needs to be devised describing the updating potential of a mining schedule at discrete timesteps. Three types of covariances need to be considered; cross-covariances between extracted blocks; cross-covariances between scheduled blocks and covariances between extracted and scheduled blocks. All this information could potentially be captured in some sort of signal-to-noise ratio. Such a ratio could characterize upfront the reliability of upcoming adjustments in predicted measurements. The ratio could as well indicate whether or not it is wise to change operational decisions based on the updated predicted measurements.

New updating algorithms have to be implemented to extend the capabilities of

the developed updating framework. The current algorithm is designed to update univariate continuous attributes based on single continuous measurement variables. These measurement variables are either directly or indirectly related to the attribute of interest. The algorithm needs to be extended to handle multivariate attributes, both in the block model as on the measurement side. That is correlated measured variables need to be jointly considered to update (other) correlated attributes. Neglecting to do so will result in a loss of information. Additionally, it would be interesting to update categorical variables such as ore types, lithologies, weathering zones based on equipment performance or other sensor measurements.

7.2.2. ADDITIONAL EXPERIMENTS

Additional experiments have to be conducted to understand the limitations of the algorithm regarding the geological environment and the prior model thereof. The question still remains whether good performance can be achieved when the prior model and the reality differ substantially. For the sake of simplicity, three detrimental scenarios are considered; (1) an incorrect representation of the spatial correlation structure, (2) an incorrect positioning of patterns of extreme attribute values (zones of clustered high and low values), (3) unrealistic characterization of the geological uncertainty.

A fist set of experiments focusses on studying how an incorrect spatial correlation structure impacts local anomaly corrections later on. The modelled correlation structure can be lifted from the prior realizations using empirical covariances. The resulting covariances largely determine the extent of updates in neighbouring block values. When the covariances between measurements and neighbouring blocks are overestimated, attribute values are corrected too aggressively, causing erroneous updates. A corrective localization technique is integrated into the algorithm to prevent such behaviour (ref. Chapter 3). Underestimated covariances on the other hand yield updates which are too modest. Certain local anomalies remain partly uncorrected. In reality, the correlation structure of the true field is obviously unknown. Several parallel updating tracks could be run simultaneously, all initiated with a specific set of prior realizations (each set could be generated using a different variogram model and/or simulation technique). Error statistics, collected during an initial training period, could eventually be compared to select the most 'correct realization set, best describing the true spatial correlation.

A second set of experiments evaluates whether the algorithm still performs if the prior realizations set is insufficiently conditioned. The inherent correlation structure is assumed to be approximated correctly (it could for example be derived from another area with similar geology). Under these conditions, it is assumed that the algorithm (partly) corrects for the local estimation biases and (approximately) identifies clusters of extreme values.

A third set of experiments investigates the necessity of realistically characterizing geological uncertainty (partly related to a correct representation of the spatial correlation structure). This geological uncertainty is propagated through the forward simulator and results in a quantification of the uncertainty around predicted measurements. The (geological) uncertainty is reflected by the spread in the Monte Carlo samples associated to block and measurement volumes. When the spread is too wide, insufficient confidence is placed in the initial prior model. As a result, the algorithm will update too

aggressively. The opposite happens when the spread is too narrow. The third experiment series studies whether the algorithm is capable of correcting under- and overconfident models. It is furthermore interesting to quantify the effect of an uncertainty bias on the overall algorithm performance.

The previously discussed experiment series focussed on analysing how closely prior realizations have to match reality in order for the algorithm to function properly. Some more experiments can be conducted varying the extraction sequence and the update interval (note that the update interval is not the same as the measurement interval). It is assumed that the order of extraction does not significantly impact the final results. This assumption has not been proven and should be verified. Furthermore, it is worthwhile to investigate whether longer update intervals, consisting of multiple measurement intervals, would produce worse results. In other words, a study needs to be made to differentiate between the assimilation of multiple observations simultaneously (update interval is a multiple of the measurement interval) and the assimilation of one observations at a time (update interval equals measurement interval).

7.2.3. PRACTICAL APPLICATION

A practical application requires two pieces of software: an updating algorithm and a forward simulator. A discussion of the former will be omitted here. Several suggestions have already been made to further improve/extend the updating algorithm. The following text discusses some of the modelling assumptions underpinning the case specific forward simulator and suggest potential improvements.

- The entire updating approach applied during the pilot study relies on the assumption that material can be tracked with sufficient accuracy during $4h$ long time intervals. The tracking assumptions underpinning the forward simulator could be validated using e.g. RFID tags (radio frequency identification). These tags have to be placed within the source material (blastholes, muckpiles, ...) and their initial location needs to be logged. Several antennas have to be installed at critical positions in the comminution circuit. Subsequently, events of tags passing an antenna are recorded. Finally these recorded events will have to be compared against the tracking behaviour of the forward simulator. A sufficient agreement will validate the tracking assumptions and the consequential behaviour of the forward simulator. This validation step might be challenging as it is yet unclear how the RFID points have to be related to corresponding material volumes.

- The material source locations can be more accurately defined. Currently only the GPS location of the trucks are recorded. Aerial photographs are used to derive a set of correction vectors (length and orientation) linking truck positions to actual digging locations. Very recently, the high-precision GPS locations of the loaders and excavators are made available as well. An algorithm can be written to accurately determine the origin of each bucket loaded into a truck.

- A support correction algorithm needs to be designed and implemented to adjust the distribution of W_i values assigned to a truck when material is reclaimed from a stockpile subdomain. Some random noise proportional to the amount of volume

reduction should be added to the characterizing set. As a result, the uncertainty in the truck would be larger than the one in the stockpile subdomain. The statistical consequences of this volume reduction are currently ignored.

- Equipment performance measurements might start to drift as critical components are wearing out (e.g. liners in the ball mill). Machine learning techniques could be applied to automatically correct for the occurring drift.
- Since the case study is based on an off-line execution of the algorithm, future predictions were generated using a 'tracking-based' forward simulator. When the algorithm would run on-line, the tracking data would only allow to generate model-based equivalents to current and historic measurements. A second 'schedule-based' forward simulator needs to be build to generate production forecasts. The production forecasts are by no means necessary to run the updating algorithm.
- The application of a 'schedule-based' forward simulator allows for a continuous re-evaluation of operational decisions based on the most up to date information. Production forecasts should be recorded and validated against performance measurements as soon as they become available. As such, the performance of the algorithm with respect to generating accurate forecasts is continuously monitored. The performance regarding reconciliation of historic measurements should obviously be monitored as well.

The application of machine learning algorithms to regularly recalibrate the relationships between sensor data and word index values opens up another avenue of research. A large database of reconciled work index values is built up as the updating algorithm is operated over a longer period of time. Each resulting block value is then to be linked with sensor data from neighbouring grade control samples. Both datasets are then to be used to regularly retrain the relationship between sensor data and work index estimates. The retrained relationship will eventually result in more accurate and reliable GeoMet models. Finally, the gained knowledge, e.g. a better understanding of the relation between ore types and hardness (assuming sensor data can be used to differentiate between ore types), should be transferred into the longer-term resource model on a regular basis.

REFERENCES

T. Wambeke and J. Benndorf, *A simulation-based geostatistical approach to real-time reconciliation of the grade control model,* Mathematical Geosciences **49**, 1–37 (2017a).

T. Wambeke and J. Benndorf, *A study of the influence of measurement volume, blending ratios and sensor precision on real-time reconciliation of grade control models.* Mathematical Geosciences (accepted with revisions) , 1–45 (2017b).

T. Wambeke, D. Elder, A. Miller, J. Benndorf, and R. Peattie, *Real-time reconciliation of the geomet model based on ball mill performance measurements - a pilot study at the tropicana gold mine.* Transactions of the Institutions of Mining and Metallurgy - Section A, Mining Technogloy (under review) , 1–21 (2017).

ACKNOWLEDGEMENTS

At the time of writing, I nearly completed the dissertation to illustrate my capabilities as an 'independent' researcher. To be honest, I do not believe research is meant to be truly independent and in my case it definitely wasn't. Numerous individuals have contributed directly or indirectly to this work or were indispensable in shaping the person I am today. I am and will always be indebted to you all! Without your support, assistance and guidance this dissertation would have never been written.

First and foremost, I would like to gratefully acknowledge **Prof. Jörg Benndorf**. You convinced me to pursue a PhD, while I always thought (when I was younger and more ignorant) that I would not be the kind of person who would enjoy the academic pursuit. I couldn't have been more wrong. Over the years, you provided me with valuable guidance, support, motivation and feedback. I truly appreciate your constant but patient efforts to improve my writing skills. I further would like to thank you for encouraging me to publish my work, speak at international conferences and expand my network. I really enjoyed the trips we made together and the numerous inspiring and invaluable discussions we had. You were abundantly helpful and have assisted me in numerous ways.

I would like to express my sincere gratitude to **Prof. Jan-Dirk Jansen**. Despite your packed agenda, you always found the time to provide me with advice and assistance. I really came to appreciate the critical questions you always asked. Your keen eye for detail continuously challenged me to improve the quality of my work. I highly respect you as a person. Your academic integrity and indefinite patience is truly remarkable.

I am grateful that I had the opportunity to conduct my research within the team led by **Dr. Mike Buxton**. Mike, your years of experience within the mining industry are invaluable to our section. Thank you for constantly sharing your knowledge with the rest of us. I am grateful for all the books, papers and reports you recommended me to read. I as well would like to thank you for your unconditional support and for believing in my capabilities.

I would like to thank **Prof. André Vervoort** who advised me to pursue a master's degree in Delft. You as well introduced me into the field of Geostatistics and provided me with some good introductory textbooks while I was working on my master thesis. Without your guidance and advice I would never have ended up where I am now.

I would like to thank as well the other committee members **Prof. Martin Verlaan**, **Prof. Michael Hicks** and **Dr. Julian Ortiz** for carefully reviewing my work. Currently, I am looking forward to discuss all your remaining questions and comments during the defense (I might and probably will think differently as the defense date will start to approach). Mike, thank you for spotting the last remaining grammar and spelling mistakes.

I gratefully acknowledge the support and generosity of **IHC MTI** and **Royal IHC**. I would like to especially thank **Henk van Muijen** for providing me with the opportunity to pursue a PhD while working for a company. Henk, without your support, this research would have never been possible. Thank you for keeping our project alive while IHC was facing difficult times. Not many people have the courage to invest in long-term research projects while the rest of the company resorted to short-sighted cost cutting.

I would like to thank **Robert van de Ketterij** who fiercely fought to protect the employees in the research department during the reorganization. Robert, I thank you for keeping me out of the unnecessary turmoil so I could focus on doing my job. Without your support and guidance, I would never have been able to finish within the respectable time-frame. Thank you for all your advice.

I would like to thank **Rick Lotman** for helping me to find a good balance between the time spend on my PhD and on other IHC projects. Your assistance and advice greatly helped me to reduce stress. I really enjoyed your friendly, correct and respectful management style. I learned from you that personal well-being always should come first.

I would like to thank **Léon Seijbel** for guiding me through the last stages of my PhD. Your optimism inspired me to always look on the bright side of life. Thank you for listening patiently to all our complaints. Considering what you have been through, it must not have always been easy.

I as well would like to thank all my other colleagues at MTI for creating such a nice working atmosphere. Most notably, I would like to thank **Mario Grima Alvarez** with whom I have been closely working together for the last six years. I really enjoyed all our technical and non-technical conversations. You are probably the most relaxed person I have ever met. I as well would like to thank **Joep Goeree**. Especially near the end, you were a great support. You truly understood the challenges of combining a PhD with an industry job. Our conversations often provided a welcome break from writing the last chapters of the dissertation. **Maarten in't Veld**, thank you for helping Mario and me in making FOCAL a success. It was a pleasure working with you. **Jan Los**, thank you for the many times we cycled together on the race bike to and from Kinderdijk. To all other colleagues who I did not yet mentioned, thank you all for the nice time!

I am greatly indebted to a large number of people at **Anglo Gold Ashanti**. I would like to thank you all for giving me the opportunity to pilot test our developed approach. I am very impressed with the ever present mentoring and open-minded culture. It was a real pleasure to collaborate with you all.

I would like to especially thank **Richard Peattie**, who internally advocated our proposal and introduced me to the right persons at the **Tropicana Mine**. You were always available on a very short notice to discuss the progress of the project and to help me resolve some of the technical challenges. Without your support and assistance, I would never have had a change to illustrate that what I developed could also work in practice.

I would like to thank **Peter Ketelaar** for acting as our onsite champion. I couldn't have wished for a better one. Peter, you were really invaluable during the initial stages of the project. Where it not for your retirement, you would have played a critical role until this day. I never quite understood how you did it, but you always seemed to get the right people involved at the right time.

I further would like to thank **Damon Elder** to assist us during the development of the material tracking model. Your technical excellence had a great impact on me. You always immediately understood our ideas and took them two steps further. The discussion I had with you were extremely instructive. I was then also very pleased to hear that you would become our new onsite champion. I as well have to thank you for the excellent tour we got when visiting the mine site and for providing the aerial photograph displayed on the front cover.

I would like to thank **Andrew Miller** for providing us access to all necessary data. If you would not have lead Tropicana through the transition of storing all information in a central data lake, this exercise would have been a complete nightmare. You learned me how to tackle a larger software project in a more professional way (although there is still room for improvement). Our shared passion of Python led to numerous interesting and inspiring discussions.

Numerous other persons from **Anglo Gold Ashanti** have been directly and indirectly involved in this project. **Nick Clarke** and **Anthony Anyimadu**, thank you for explaining how the processing plant at the Tropicana was designed. Your assistance was extremely valuable during the formulation of a prototype tracking model while you always ensured that I properly understood its limitations. **Jennifer Birch**, thank you for collecting and guiding us through the available Citect data. **Michael Chambers**, thank you for assisting us in correctly interpreting the MineStar data. **Cameron Thicker**, thank you for the interesting discussions regarding material flow in the primary ore stockpile, the primary and secondary crushing circuit. **Shae New**, thank you for arranging both our mine visits and making them so enjoyable. I do apologize to those who helped me out but who I forgot to acknowledge here.

I further would like to express my profound gratitude to my dear friends from the PhD room. As Cansin eloquently put it: "together we formed a band of brothers". I am really grateful for all your support and friendship. I am happy that we became much more than just colleagues. **Marinus**, you were there when I started and were a tremendous help to figuring out how things had to be arranged at the university. You were as well the instigator of some of the nicest evenings we had together as a group. **Adriana**, it was truly remarkable to learn that despite the large distance between our home countries, we still shared quite a bit of common views on life due to our Catholic uprising. I am not sure if you realized it but you were indispensable in diffusing the heated arguments we sometimes had in the group, especially in the very beginning. **Massoud**, I do miss our daily bicycle rides and really enjoyed the various trips we made together. It was nice to develop a joint passion for programming. I have to say, it is a pity that you stopped trying to literally translate Iranian sayings. You always made the group laugh when the donkey suddenly sneaked into the conversation. Thank you for being my paranymph. **Cansin**, you always ensured that we would not forget to enjoy life in the midst of all the PhD commotion. You as well taught me how to properly respond to Massoud in his own language. **Feven**, thank you for introducing us into the delicious Ethiopian cuisine. Your hospitality is heart-warming. **Volker**, thank you for all the advice and support. It really helped to discuss with someone who viewed things more from the outside. It is a petty that you were not always around. **Angel**, thank you for showing us all the nicest places in Valencia. Your enthusiasm and drive is really contagious. **Jeroen**, you are a brave

individual to start a PhD while you see most of us suffering from the home stretch. Thank you for joining the team. I as well would like to thank the other members of the resource engineering section. Thank you **Marco** and **Jack** for the nice conversations during lunch time.

I further would like to thank **my friends and family** for believing in me. You have certainly made this journey more pleasant and always provided a necessary and welcome distraction. I am grateful for all your unconditional support and encouragements. I would like to thank my **parents** for allowing me to grow up in a stimulating and loving environment. You were indispensable in shaping the person I am today. I dedicate this dissertation to my **mother** who was suddenly taken from our lives. Mom, I still miss you to this day. I am grateful to my brother **Jan** and sister **Kim** who (unknowingly) helped me through these dark days in my life (together with Marijke). I am proud to see what my little sister and brother have become. Jan, Thank you for being my paranymph. I further would like to thank the family of my girlfriend for embracing me like one of their own.

As we all know, the last paragraph is devoted to the most important person in your life, in my case **Marijke**. Marijke, you deserve way more than just a paragraph. Without your everlasting care, unconditional support and enormous love, this would have never been possible. Thank you for all the sacrifices you made without any complaints or hesitation! Thank you for always being there when I needed you the most. Our long foreign adventures were a most welcome distraction from the daily hassle of life and always made me realize what is truly important in life. You are my wonderwall! [1]

Tom Wambeke
Delft, January 2018

[1] I should not forget to acknowledge our cat **Linux** for contributing to the many typos and spelling mistakes as he occasionally walked over the keyboard when I was writing the manuscript.

CURRICULUM VITÆ

Tom WAMBEKE

08-06-1989 Born in Leuven, Belgium.

EDUCATION

2001–2007 Grammar School
Sint-Tarcisiusinstituut, Zoutleeuw
Specialisation: Mathematics and Science

2007–2010 Bachelor in Geotechnical and Mining Engineering - Cum Laude
University of Leuven

2010–2011 European Geotechnical and Environmental Course (EGEC) - Avg. grade of A
Federation of European Mineral Programs (FEMP)

2011–2013 Master in Resource Engineering - Cum Laude
Delft University of Technology

2013–2017 PhD. Applied Earth Sciences
Delft Univeristy of Technology
Thesis: Data Assimilation in the Minerals Industry
Real-Time Updating of Spatial Models using Online Production Data
Promotor: Prof. dr. ir. J.D. Jansen
Copromotor: Prof. dr. Dipl.-Ing. J. Benndorf

WORK EXPERIENCE

2013–2017 Research Engineer, IHC MTI (Royal IHC)

LIST OF PUBLICATIONS

JOURNAL PAPERS

7. **T. Wambeke**, D. Elder, A. Miller, J. Benndorf and R. Peattie, *Real-time reconciliation of the geomet model based on ball mill performance measurements - a pilot study at the Tropicana gold mine*,Transactions of the Institutions of Mining and Metallurgy - Section A, Mining Technogloy (under review), (2017).

6. **T. Wambeke** and J. Benndorf, *A study on the influence of measurement volume, blending ratios and sensor predicion on real-time updating of grade control models*, Mathematical Geosciences (accepted with revision), (2017).

5. **T. Wambeke** and J. Benndorf, *A simulation-based geostatistical approach to real-time reconciliation of the grade control model*, Mathematical Geosciences **49**, 1 (2017).

4. C. Yuksel, T. Thielemann, **T. Wambeke** and J. Benndorf, *Real-time resource model updating for improved coal quality control using online data*, International Journal of Coal Geology **162**, 61 (2016).

3. **T. Wambeke**, and J. Benndorf, *An integrated approach to simulate and validate orebody realizations with complex trends: a case study in heavy mineral sands*, Mathematical Geosciences, **48**, 7 (2016).

2. **T. Wambeke**, M. Alvarez Grima, G.A. Fenton, J. Benndorf and A. Vervoort, *Use of local average subdivision to characterize marine mineral reserves - a conceptual framework*, Journal of Canadian Institute of Mining **6**, 3 (2015).

1. J. Menz, J. Benndorf and **T. Wambeke**, *Signal-controlled least square collocation: A new quality in geostatistical estimation and simulation?*, Spatial Statistics **14**, Part C (2015)

CONFERENCE PROCEEDINGS

6. **T. Wambeke** and J. Benndorf, *Data-assimilation of sensor measurements to improve production forecasts in resource extraction*, Conference of the international association of mathematical geosciences, Freiberg (2015).

5. **T. Wambeke** and J. Benndorf, *Optimal inference of modelling parameters to simulate complex trends across soft boundaries - a case study in heavy mineral sands*, Ore body modelling and strategic mine planning, Perth (2014).

4. **T. Wambeke** and J. Benndorf, *Integrated geometallurgical modelling of heavy mineral sands accouting for profitability, extractability and processability under uncertainty*, Geokinematischer tag, Freiberg (2014).

3. **T. Wambeke** and J. Benndorf, *Embracing geological uncertainty in marine resource exploration and exploitation*, Minerals of the ocean and deep sea minerals and mining conference, Sint petersburg (2014).

2. **T. Wambeke**, M. Alvarez Grima and R. van Gelder. *Decision making on the ocean floor - a framework to assess variability and uncertainty in deep-sea mining*, Offshore Technology Conference, Houston (2013).

1. **T. Wambeke**, M. Alvarez Grima, G. Fenton, R. van Gelder, *Mining the ocean-floor - managing geological uncertainty*, World Mining Congress, Montreal (2013).