MODELOS E ALGORITMOS PARA PLANEJAMENTO INTEGRADO NA INDÚSTRIA DA MINERAÇÃO

BRUNO SANTOS PIMENTEL

MODELOS E ALGORITMOS PARA PLANEJAMENTO INTEGRADO NA INDÚSTRIA DA MINERAÇÃO

Tese apresentada ao Programa de Pós-Graduação em Ciência da Computação do Instituto de Ciências Exatas da Universidade Federal de Minas Gerais como requisito parcial para a obtenção do grau de Doutor em Ciência da Computação.

Orientador: Geraldo Robson Mateus

Belo Horizonte

Março de 2011

BRUNO SANTOS PIMENTEL

MODELS AND ALGORITHMS FOR INTEGRATED PLANNING IN THE MINING INDUSTRY

Thesis presented to the Graduate Program in Computer Science of the Federal University of Minas Gerais in partial fulfillment of the requirements for the degree of Doctor in Computer Science.

Advisor: Geraldo Robson Mateus

Belo Horizonte
March 2011

Pimentel, Bruno Santos.
P644m Modelos e algoritmos para planejamento integrado na indústria da mineração/ Bruno Santos Pimentel. — Belo Horizonte, 2011.
xxxv, 132f. : il. ; 29cm

Tese (doutorado) — Universidade Federal de Minas Gerais. Departamento de Ciência da Computação.
Orientador: Geraldo Robson Mateus.

1. Computação - Teses. 2. Sistemas Operacionais - Teses. 3. I. Orientador. II. Título.

CDU 519.6*34 (043)

UNIVERSIDADE FEDERAL DE MINAS GERAIS
INSTITUTO DE CIÊNCIAS EXATAS
PROGRAMA DE PÓS-GRADUAÇÃO EM CIÊNCIA DA COMPUTAÇÃO

FOLHA DE APROVAÇÃO

Modelos e algoritmos para planejamento integrado na indústria da mineração

BRUNO SANTOS PIMENTEL

Tese defendida e aprovada pela banca examinadora constituída pelos Senhores:

PROF. GERALDO ROBSON MATEUS - Orientador
Departamento de Ciência da Computação - UFMG

PROF. ABÍLIO PEREIRA DE LUCENA FILHO
Faculdade de Economia e Administração - UFRJ

DR. DARCTON POLICARPO DAMIÃO
Usinas Siderúrgicas de Minas Gerais S.A. - Usiminas

PROF. MAURÍCIO CARDOSO DE SOUZA
Departamento de Engenharia de Produção - EE- UFMG

PROF. REINALDO MORABITO NETO
Departamento de Engenharia de Produção - UFSCar

PROF. ALEXANDRE SALLES DA CUNHA
Departamento de Ciência da Computação - UFMG

Belo Horizonte, 18 de março de 2011.

À minha Grande Família.

Agradecimentos

Começar a escrever estes agradecimentos foi um dos momentos mais aguardados por mim nestes últimos meses. Eles são, de fato, o último pedaço de texto a ser incluído nesta Tese. E com certeza é um momento para ser saboreado. Como estou fazendo agora: bom vinho, ótima música e a segunda melhor vista do mundo: o janelão da minha sala.

("It's times like these...")

Não foram somente as longas horas investidas. Nem foram somente os tantos lugares em que me vi trabalhando: na Universidade, em casa, em Curvelo, nos saguões dos aeroportos deste País (e nos aviões, quando em vôo de cruzeiro), nas férias na praia, nos hotéis a trabalho, na fila do banco e até no carro, nos engarrafamentos. E também não foram poucos os finais de semana, as terças-feiras de Luizinho, os feriados, ou quaisquer outros dias em que a ideia era simplesmente ficar "de boa", que foram trocados pelos artigos e pelo notebook. Fato é que preciso compartilhar o momento com algumas pessoas que são muito importantes pra mim.

E como a lista é grande, vou começar primeiro pela turma que me apoiou diretamente no doutorado.

Este trabalho começou já há algum tempo e passou por diversos acertos de rumo ao longo do caminho. Pelo primeiro desses acertos eu agradeço ao amigo e professor Mario Campos, a quem tive a honra de conhecer ainda nos tempos do mestrado e cujo papel de orientador, à época, raramente se restringia, para a minha sorte, às coisas acadêmicas. Por todas as boas conversas e sábios conselhos, obrigado.

Outros ótimos acertos são fruto da orientação precisa do amigo e professor Geraldo Robson, que me guiou por este doutorado e que sempre esteve disposto a encarar minhas propostas, ansiedades, percalços e sumiços com compreensão, pragmatismo e bom humor, especialmente considerando minha dedicação (menos que) parcial. Trabalhar com o Robson no doutorado e no Synergia foram experiências valiosas. Melhor ainda, posso hoje dizer que me diverti um bocado nesse tempo todo; e espero que você também, Robson. Obrigado!

Meus sinceros agradecimentos às inestimáveis contribuições dos professores Abílio Lucena, Alexandre Cunha, Maurício Cardoso, Reinaldo Morabito e do Dr. Darcton Damião para este trabalho. Ao amigo Darcton, em especial, meu muito obrigado pelas oportunidades, pelos projetos e por todo o aprendizado, incentivo e companheirismo durante o tempo em que trabalhei sob sua direção na Pesquisa e Inovação da Usiminas — levo na bagagem muito mais que simplesmente experiência profissional.

E claro, os meus obviamente insuficientes agradecimentos ao Franklin, por todo o apoio e dedicação ao longo destes anos, e também pela notável habilidade com que foi capaz de transformar ideias complicadas em códigos simples. Só espero, Franklin, que você tenha ganhado tanto ou mais que eu nessas nossas interações. Valeu!

E aos amigos da Usiminas, da Accenture, do Synergia e do LaPO, muito obrigado pelo apoio, pela compreensão e pela torcida.

Agora, à parte emoção.

("There goes my hero... watch him as he goes.")

Meu avô José Mendes deixou este mundo muito antes que eu pudesse conhecê-lo — o que é uma pena, por tudo que já pude e procurei saber dele. Acontece que ele deixou muitos filhos (muitos mesmo) e mesmo com toda a adversidade conseguiu criar uma família unida, divertida e inspiradora. A amizade, a presença, a confiança, o apoio e até mesmo a importância de nos reunirmos para celebrar conquistas de qualquer magnitude são apenas alguns dos valores que aprendemos a carregar desde cedo.

O ambiente é tão acolhedor que meu pai resolveu deixar seu saudoso Recife pra vir se juntar à minha mãe cá nas Minas Gerais. Aos dois, por sinal, meu eterno obrigado: pela vida, pela amizade, pela educação, pela minha irmã e pelos belos valores.

E à confraria dos meus queridos tios e tias (todos os dez), dos meus grandes tios avôs e dos incontáveis primos de todos os graus, dos meus saudosos avós mineiros e pernambucanos, obrigado por me proporcionarem este lugar tão rico em união, alegria, lealdade e carinho. Tenham a certeza de que nunca conseguirei buscar, tampouco aceitar, nada menos que tudo isso para a minha vida. Todos vocês me deixaram muito mal acostumado.

Mas com isso tudo também vem um sério efeito colateral. Infelizmente, como disse antes, meus pais só me deram uma irmã. Verdade que eu queria pelo menos mais um rapa de tacho pra se unir a mim contra as forças do mal desde cedo, mas não deu. O jeito foi tentar expandir a minha irmandade com os melhores amigos que alguém pode querer. Só espero que eles tenham tido paciência de ler estas linhas até aqui — sei que precisaram de muita paciência comigo ao longo destes últimos anos. Então, aí vai: aos meus bons amigos Rodrigo, João, Menudo, Calaza, Daniela, Helbert, Chokito, Luís, Rodrigão, Vinícius, Xandoca, Fabiano, Rogerão, M., Sandrão, Bayão, Muzzi, Flávio, Martins, João Chequer, Tim, Cota — meus irmãos brancos —, às suas respectivas, à Segunda Geração da Moçada, à Ju, ao Mazzini e ao Caixeta, obrigado simplesmente pela amizade verdadeira e gratuita. Se chegamos até aqui depois de 2, 20 ou mesmo 30 anos, não foi por acaso; nem pouco provável será comemorarmos outras décadas juntos nos próximos anos.

("Someone like you...")

E a você, Júnia, a melhor vista do mundo, por todo o bem que você me faz, pela paixão, amor e carinho, pelo frio na barriga, pela inspiração, por tirar meu fôlego e por ser a pessoa por quem eu mais me esforço para merecer, obrigado. Quero tudo com você, sempre.

"Someday you will have to commit completely to something."
(Richter 10, Arthur C. Clarke)

Resumo

Neste trabalho desenvolvemos modelos e algoritmos para problemas de planejamento integrado de produção e logística na indústria da mineração. A partir de uma ampla revisão bibliográfica, aborda-se o conceito de Cadeia Global de Suprimentos da Mineração e discutem-se os principais problemas de Pesquisa Operacional sob a ótica da integração de decisões de planejamento.

Decisões estratégicas são contempladas em um modelo de programação estocástica multiestágio para planejamento de capacidade da Cadeia Global de Suprimentos da Mineração. O modelo integra decisões de localização de instalações e projeto de redes considerando as economias de escala inerentes aos investimentos em capacidade. Um primeiro estudo indica direções para o desenvolvimento de heurísticas e algoritmos específicos para o problema por meio de uma avaliação empírica de diferentes parâmetros do pacote CPLEX. Em um segundo estudo, uma Heurística Lagrangiana é proposta para determinar, para instâncias maiores, boas soluções viáveis para o problema em um tempo razoável quando comparado ao CPLEX. Adicionalmente, a habilidade de se determinar boas soluções viáveis nos momentos iniciais do processo computacional é tratada em uma abordagem de busca local com fixação flexível de variáveis. Tal algoritmo também é avaliado comparativamente aos demais procedimentos de solução discutidos no texto.

Decisões táticas, por sua vez, são abordadas em um modelo matemático para planejamento tático de vendas e operações. O modelo possui características de dimensionamento de lotes em um ambiente de rede, mas apresenta fluxos necessariamente inteiros (capacidades de trens e navios, por exemplo), níveis discretos de produção em minas e usinas de beneficiamento, além do desbalanceamento causado por perdas de massa inerentes ao beneficiamento de minérios e ao transporte de cargas em granel. Um conjunto de heurísticas Relax&Fix é desenvolvido para tentar resolver instâncias de tamanhos realistas. As estratégias propostas são capazes de produzir resultados melhores que o CPLEX para a maioria das instâncias, em especial para aquelas com horizontes de planejamento mais longos. O algoritmo de busca local com fixação flexível de variáveis também é avaliado em relação à sua habilidade de produzir boas soluções viáveis logo no início do processo computacional.

Decisões operacionais, por fim, são brevemente abordadas por meio de um modelo programação inteira por metas que aborda o problema de programação integrada de curto prazo de operações de lavra, beneficiamento, estocagem e despacho de produtos de minério de ferro em um complexo minerador. O conceito de Valor da Solução Integrada é proposto como forma de quantificar os ganhos obtidos por meio do investimento na solução de um problema integrado e (potencialmente) mais complexo.

Palavras-chave: Pesquisa Operacional, Cadeia de Suprimentos, Mineração.

Abstract

In this Thesis, we develop models and algorithms applied to integrated production and logistics problems in the mining industry. Based on an extensive literature review, we address the Global Mining Supply Chain concept and discuss the main related Operations Research problems under a integrated planning perspective.

Strategic decisions are evaluated in a novel multistage stochastic integer programming model to address the capacity planning problem in a Global Mining Supply Chain. The model integrates capacitated facility location and network design decisions with economies of scale on the capacity costs. We analyze the characteristics of the problem by means of an empirical study of different settings for the parameters of the CPLEX solver. Such analysis provides pointers to the development of specific algorithms and solution approaches. We then develop a Lagrangian Heuristic as a means to determine, for large problem instances, good feasible solutions in a reasonable amount of time when compared to CPLEX. Furthermore, the ability of determining good feasible solutions in the early stages of the computation is addressed in a soft-fixing local search framework, which is evaluated against the other solution approaches discussed.

Tactical decisions are tackled in a mixed-integer programming approach to the integrated sales and operations tactical planning problem in a Global Mining Supply Chain. The model has characteristics of a lot sizing problem in a network environment, but with challenging aspects related to integer flows, discrete production levels and mass losses in concentration and transportation processes. We develop a series of Relax&Fix strategies in order to address realistic sized problem instances. Those strategies are able to outperform CPLEX for most of the several problem instances considered, and with greater success in longer planning horizons. The soft-fixing local search is also evaluated for its ability of determining good feasible solutions in the early stages of the computation.

Operational decisions are briefly discussed in a mixed-integer goal programming model to address the integrated short-term programming of iron ore open pits, processing plants, stockyards and shipping operations. We propose the concept of the Value of the Integrated Solution, which determines how valuable is solving a more complex integrated decision problem given the potential losses incurred when individual decisions are undertaken.

Keywords: Operations Research, Supply Chain, Mining.

Resumo Estendido

Em um cenário econômico de alta competitividade, internacionalização e concentração, a indústria da mineração apresenta-se como um ambiente extremamente fértil à aplicação de técnicas de Pesquisa Operacional. A complexidade das operações, os desafios da integração entre produção e logística em toda a cadeia produtiva e os grandes investimentos envolvidos motivam estudos voltados para o aumento da produtividade dos ativos, para a redução dos custos de operação e para o aumento da geração de valor dos empreendimentos. Dessa forma, processos de tomada decisão nos níveis operacional, tático e estratégico, em especial aqueles fundamentados em Pesquisa Operacional, encontram oportunidades significativas de aplicação e captura de valor.

Dadas a complexidade e a escala dos ativos produtivos e logísticos envolvidos, a indústria da mineração deve ser tratada como uma grande cadeia de suprimentos, onde minas, usinas de beneficiamento, ferrovias, portos e entrepostos sejam geridos de forma a garantir o máximo de integração e efetividade na execução das operações. Nesse sentido, no desenvolvimento de sistemas de suporte à decisão em mineração, a integração de diferentes estágios da cadeia e de diferentes níveis de decisão deve ser cuidadosamente considerada. Existem ao mesmo tempo oportunidades e riscos significativos nas interfaces existentes, seja na dimensão funcional, seja nos diferentes níveis de decisão.

Neste trabalho, o termo planejamento é utilizado de forma ampla, abrangendo decisões estratégicas, táticas e operacionais. O termo integração, por outro lado, restringe-se ao aspecto funcional, ou seja, como operações em diferentes estágios da cadeia produtiva interagem e interferem com aquelas à jusante e à montante. Desenvolvemos e analisamos um conjunto de modelos matemáticos e algoritmos para planejamento integrado de produção e logística — considerando mina, ferrovia e porto — na indústria da mineração — sob as óticas de decisões estratégicas, táticas e operacionais.

Decisões estratégicas cobrem horizontes de longo prazo e envolvem investimentos em capacidade de produção e distribuição, seja na abertura de novas minas, usinas, entrepostos ou portos, seja na abertura de novas ferrovias ou extensão de ferrovias exis-

tentes, ou mesmo na aquisição de outros ativos relacionados à produção e transporte do minério de ferro desde as minas até os clientes finais. Como principal tipo de investimento, a expansão de capacidade encontra-se associada à previsão de crescimento da demanda por minério de ferro. No entanto, variações negativas na demanda também podem induzir decisões quanto ao fechamento, temporário ou definitivo, de instalações produtivas ou canais logísticos. Somadas a isso, as incertezas associadas às variações de demanda e preços de commodities sugerem uma abordagem estocástica para a tomada de decisões estratégicas. No Capítulo 2, esses conceitos são simultaneamente contemplados em um modelo estocástico multi-estágio de planejamento de capacidade sobre toda a cadeia produtiva global da mineração. De particular interesse é a habilidade de se desenvolverem políticas para decidir sobre a abertura, fechamento e reabertura de instalações produtivas e canais logísticos ao longo de um horizonte de planejamento e considerando diferentes cenários de variação de demanda. O modelo matemático proposto é original e apresenta características de dimensionamento de lotes, projeto de redes, localização de instalações e expansão de capacidade, demonstrando grande complexidade computacional e exigindo métodos específicos de solução. Num primeiro estudo, avaliamos as principais características do problema e diferentes abordagens de solução com auxílio do pacote de otimização CPLEX. Essa avaliação aponta para as direções mais promissoras de desenvolvimento de algoritmos e heurísticas para solução do problema. Uma das abordagens possíveis, a Relaxação Lagrangiana, é então explorada com o objetivo de se fortalecer o limite dual e derivar, por meio de uma Heurística Lagrangiana, bons limites primais para o problema. O desempenho da heurística é comparado com os resultados produzidos pelo CPLEX, indicando sua melhor aplicabilidade para instâncias de maior porte, com maior número de períodos no horizonte de planejamento e de nós na árvore de cenários. Uma abordagem de busca local com fixação flexível de variáveis também é proposta como alternativa para se determinar boas soluções viáveis logo no início do processo computacional. Ambas as abordagens são comparadas aos resultados obtidos pelo CPLEX, mostrando bons resultados para instâncias maiores e mais complexas.

Decisões táticas, por sua vez, cobrem horizontes de médio prazo e envolvem a alocação dos recursos de produção, estocagem e distribuição para atender as demandas do mercado de minério de ferro. Um maior grau de detalhamento dos elementos da cadeia produtiva da mineração torna-se necessário, uma vez que as decisões devem ser integradas. O Capítulo 3 apresenta um modelo matemático para planejamento tático de vendas e operações com o objetivo de integrar decisões de médio prazo ao longo de toda a cadeia. Grandes operações de mineração normalmente dispõem de diversas minas com diferentes capacidades, além de usinas de beneficiamento e pelotização, pátios de estocagem, ferrovias, minerodutos, terminais portuários e entrepostos. Dependendo das características do minério ofertado em cada mina e da demanda colocada em cada terminal portuário, as decisões de extração, beneficiamento, blendagem e manuseio podem utilizar diferentes instalações em diferentes estágios da cadeia produtiva. Cada uma dessas decisões deve considerar os custos e as capacidades de produção, estocagem e transporte de todo o sistema. Nesse sentido, um modelo de programação inteira mista aborda o problema de planejamento tático de vendas e operações tendo os cenários reais da indústria mineral brasileira como inspiração. O modelo tem características de um problema de dimensionamento de lotes em uma estrutura de redes, porém com aspectos complicadores relacionados à integralidade de determinados fluxos, ao desbalanceamento de massa causado por perdas inerentes ao transporte e à eficiência dos processos de beneficiamento, e às capacidades discretas de produção. Uma série de estratégias de solução baseadas na heurística Relax&Fix são desenvolvidas como forma de contornar a elevada complexidade computacional do problema. Apesar dos algoritmos não terem garantia de otimalidade, tampouco de viabilidade, algumas estratégias são capazes de gerar resultados razoáveis em instâncias para as quais o CPLEX alcança limites significativamente mais fracos. O algoritmo de busca local com fixação flexível de variáveis também é avaliado em relação à sua habilidade de produzir boas soluções viáveis logo no início do processo computacional, mostrando novamente um ganho modesto quando comparado ao CPLEX.

Decisões operacionais cobrem horizontes de curto prazo e envolvem a programação e o sequenciamento de operações em grau de detalhamento suficiente para orientar a execução de atividades do dia-a-dia de uma mina, usina de beneficiamento, pátio ou porto. A integração entre tais decisões é importante e pode produzir importantes resultados no que diz respeito à coordenação entre diferentes áreas operacionais de uma mesma instalação produtiva, por exemplo. Esse é o assunto do Apêndice A, que apresenta um modelo matemático de programação por metas inteira mista para tratar o problema de programação de curto prazo de operações de lavra, beneficiamento, estocagem e despacho de produtos de minério de ferro. Em grandes empresas mineradoras, a rígida estrutura organizacional comumente divide a gestão dessas operações entre departamentos bastante compartimentados. Essa estrutura e as pressões do dia-a-dia do ambiente produtivo promovem o estabelecimento de programas de produção independentes que podem gerar resultados potencialmente subótimos para as operações de um complexo minerador. Dessa forma, o modelo busca integrar os processos de programação de operações de curto prazo de um complexo minerador, utilizando restrições específicas para tratar os contratos — formais ou não — que surgem nas interfaces entre departamentos e unidades produtivas. O valor da solução integrada é

estimado por meio da comparação entre programas gerados por decisões compartilhadas e decisões individualizadas, evidenciando as perdas potenciais quando as restrições de interface são ignoradas. O modelo tem características bastante específicas e exige, assim, abordagens específicas de solução para instâncias de dimensões significativas.

O capítulo final desta Tese apresenta uma discussão sobre as principais contribuições, os resultados e as inúmeras oportunidades de extensão do trabalho. É importante ressaltar que os modelos desenvolvidos incorporaram importantes características observadas em situações reais da indústria da mineração. A experiência prática do autor e do orientador em projetos realizados na área foi fundamental nesse sentido.

Palavras-chave: Pesquisa Operacional, Cadeia de Suprimentos, Mineração.

List of Figures

List of Tables

Contents

Chapter 1

Operations Research applied to integrated iron ore production and logistics

> "He forced himself to take a single step forward and once he had done that the second was a little less difficult, and the third was almost easy."
>
> *The Pillars of the Earth, Ken Follet*

In today's highly competitive, globalized and concentrated economy, the application of Operations Research techniques in the mining industry is a promising approach. Given the size and complexity of its production and logistics assets, the mining industry must be addressed as a large-scale supply chain in which mines, processing plants, railways, ports and supply stations must be managed in a way to ensure integrated and effective operations. In the first chapter of this Thesis we present the fundamentals of the Global Mining Supply Chain, inspired by the Brazilian scenario, and according to an integrated planning approach. We use the term planning in a broad sense, which covers strategic, tactical and operational decisions. The term integration, on the other hand, is restricted to the functional interaction and coordination between decisions in different stages of the supply chain. The results given in this Chapter are also presented in the following publications:

> Pimentel, B. S., Mateus, G. R., and Almeida, F. A. (2010). Mathematical models for optimizing the global mining supply chain. In Nag, B., editor, *Intelligent Systems in Operations: Models, Methods and Applications*, pages 133–163. IGI Global, Hershey, Pennsylvania.

Pimentel, B. S., Mateus, G. R., and Almeida, F. A. (2010b). Operations research applied to integrated iron ore production and logistics. In *Operations Research in Mining Seminar*, Santiago, Chile.

1.1 Introduction

Brazil figures among the most important iron ore exporters, facing at the same time great opportunities and overwhelming challenges. On the one hand, the country has one of the largest and richest iron ore reserves in the world as well as a technically mature, traditional mining workforce. Also, given the elevated iron ore prices and demand, Brazilian mining companies have shown increasingly stronger revenues and, due to the higher economic feasibility of green field projects, unprecedented investment levels. On the other hand, there is a series of structural and economical issues that pose significant restrictions to the effective operation of the Brazilian mining industry. Firstly, although Brazilian mines and processing plants exhibit relatively high productivity levels, there are several constraints imposed by a limited and costly logistics infrastructure. Brazil still lacks important investments in roads, railways and ports which, coupled with excessive public taxation and an over appreciated currency, impose additional — and rather non-competitive — costs to the iron ore products destined to exportation. Secondly, the scale and complexity of Brazilian mining operations require a solid operations management expertise, which may not be as strong as the available mining and metallurgical technical expertise. Thirdly, given the continuous consolidation trends of the global mining markets, large mining companies often have to deal with the challenge of integrating several different assets located at potentially disperse sites and with specific capacities and managerial processes. All those issues motivate a thorough quantitative analysis of a Global Mining Supply Chain at the strategic, tactical and operational decision levels. An Operations Research approach to that call is therefore a natural choice.

Although there has been significant effort in developing intelligent decision support systems for the mining industry, those are usually focused in addressing isolated tasks such as mine development, hauling equipment dispatch, railway scheduling operations and ship queue management. An approach to the Global Mining Supply Chain, on the other hand, requires an integrated perspective that must take into account mine, railway and port operations, as well as domestic and international customers and supply stations served by the appropriate logistics channels.

1.2 An overview of mining operations

Mining is an extremely capital intensive industry and, as such, demands high efficiency with regard to operation management decisions. There are today two important trends that have been influencing the decision-making process in large-scale mining operations: the continuous ore grade decrease in mines throughout the world, which impacts product quality and processing costs; and the ever increasing pressure over operational costs. Managerial decisions must then be quantitatively evaluated in order to account for those issues and still attain the expected results.

The first major and critical decision in mining operations is the establishment of a new mining venture, since it often requires the highest levels of capital investment. Traditional evaluation of mining projects includes drilling and sampling, generating a representative orebody model, deciding mining and processing methods, assessing capital and operating costs and developing a technical and financial life-of-mine plan. After establishing a feasible mining operation, another major decision involves determining the mine design and development plan. Proper mine design must account for access opening, orebody detailing, open pit (or underground tunnel) preparation, among other strategic decisions. In the mine development plan, detailed orebody data is used to build a discrete block model that allows scheduling ore removal according to the net present value, expected grade, processing plant requirements and the terrain's structural constraints.

Ore extraction is performed by a fleet of heavy-duty, off-road trucks and shovels, which remove the run-of-mine (ROM) from the pits and take it to appropriate locations — valuable ore to crushing and screening facilities and valueless material to waste stockpiles. Since the operation of hauling equipment is rather costly, the appropriate allocation and dispatching of trucks and shovels to load and discharge points is of utmost importance.

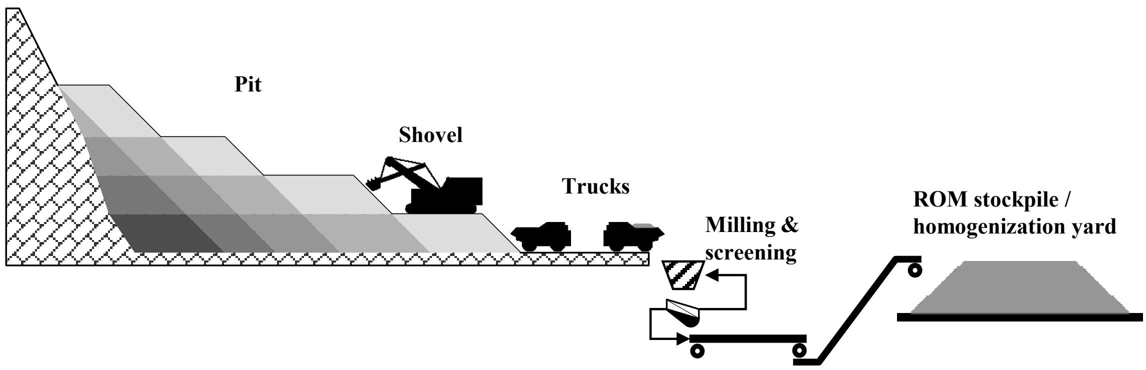

Figure 1.1. The mine pit and crushing subsystem.

Milling and screening facilities are employed to reduce ROM size according to the requirements of the processing plants. Besides size, those requirements may also include grade considerations, in which case the ROM may need to be blended and homogenized. Blending may be achieved by appropriately feeding the milling and screening facilities according to the characteristics of each mining work bench and thus allocating hauling equipment accordingly. Another way of achieving proper blending requires a homogenizing yard, where different ROM are properly mixed in stockpiles so that physical and chemical properties can be homogenized. Conveyor belts are used to transport ROM from the milling facilities to the homogenizing yard and from there to the processing plant feed.

Processing plants enhance ore product quality by means of classification (grinding and screening) and concentration (flotation and magnetic separation, for iron ore) processes. Processing plant operations are aimed at maximizing mass efficiency, which depends not only on the characteristics of the ROM offered by the mine, but also on adequate planning and control. In the iron ore industry, pelletizing plants can also be employed to process ultra-fine materials generated at the processing plants. Although such ultra-fine materials are also sold as end products, they cannot be used to feed a blast furnace directly and must, therefore, be agglomerated into pellets. Pelletizing plants can be located at the mine complex or at the port (or even at both). The pellets produced from pellet-feed ore and other agglomerating materials is also transferred to the mine or port product stockyards and from there to steelmaking customers.

In the product stockyard, ore products are stored according to commercial programs and logistics (railways, roads, rivers or sea) availability. A set of stackers form stockpiles and a set of reclaimers retrieve those stockpiles to compose specific ore product shipments. Again, a set of conveyor belts, or even trucks and shovels of smaller size, can be used to accomplish lot loading.

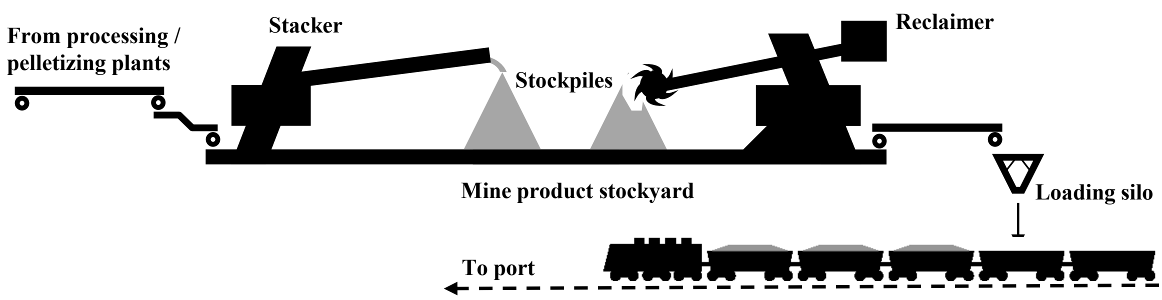

Figure 1.2. The mine's product stockyard and railroad interface.

Railroads are often used to transport ore products to the port, where exportation takes place, or to domestic supply stations, where additional processing — usually screening of lumpy ore products degraded during transport — may occur before serving local demand. Trains loaded with ore products may have to share the railway with trains used to convey different bulk cargoes, or even with passenger trains. Railroad capacity may thus be significantly constrained and impose important restrictions to the overall system operations.

The port receives ore products from a set of mine complexes and, according to a short-term program — which includes train and ship arrival, as well as stock levels and quality —, transfers those ore products to the product stockyard, pelletizing plants, screening plants, or even directly to awaiting ships. Additional blending may take place during stockpile formation or ship loading. Domestic demand may also be served by shipments originated at the port, often by rail or road.

Although some mining companies may also own and operate seaborne carriers, the majority of commercial contracts rely on third-party freight operators to transport ore products to offshore customers. Ports can ship ore products directly to customers located at, or close to, other offshore ports; alternatively, shipments may be directed to offshore supply stations, where additional blending and fine separation may take place, or where other pelletizing plants may be located. From offshore ports and supply stations, international demand can be served by train or road.

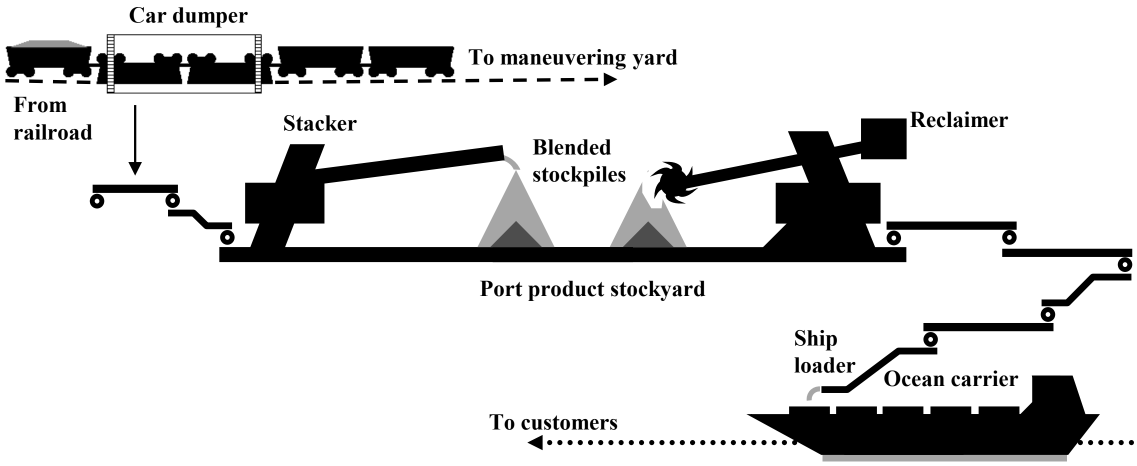

Figure 1.3. The port subsystem, with railroad and ocean interfaces.

Some characteristics of the commodity metals markets are essential to understanding the decision-making process in mining operations, especially when integrated approaches are considered. In the base metals market, such as iron ore, transportation costs amount to a significant fraction of the commodity's value. Usually, steel companies establish contracts with iron ore producers, in which prices are settled in accordance to chemical and physical characteristics of the products. Negotiated prices generally hold for approximately a quarter of a year, and iron ore producers are obliged to supply a certain amount of iron ore to each contracted buyer.

1.3 An integrated perspective and the Global Mining Supply Chain

Traditionally, a greater amount of effort has been directed towards the optimization of mine development and its long-term production plan. However, in a continuously consolidating industry, it is common to observe complex integrated operations involving not only open-pit and underground mining, but also all logistics concerned with hauling equipment dispatch, cargo composition, railroad transportation planning and traffic control, and port handling, blending and loading. Operations Research approaches to mining industry problems can provide management support along that functional dimension as well as in strategic, tactical and operational decision levels. Depending on the scenario and on the problem size, different solution approaches may be required, with an unavoidable tradeoff between computational complexity, time and solution quality.

The literature devoted to optimization of mining operations can be categorized according to decision-level and functional dimensions. At the decision-level dimension, Management Science categorizes the decision-making process according to the following structure [Bradley et al., 1977]:

1. **Strategic planning** deals with long-term decisions involving managerial policies and resource development. Decision analysis is large in scope and requires very aggregate information. In the mining context, strategic decisions include, for instance, the establishment of new mining ventures and capacity expansion of existing facilities.

2. **Tactical planning** seeks effective resource allocation to satisfy demand requirements and operation constraints on a given time horizon. Decision analysis still requires significant information aggregation. In the mining context, tactical decisions include, for instance, developing mine or port production plans, or railroad transportation plans.

3. **Operations control** is concerned with short-term decisions, usually involving low-level programming and scheduling. Information generated at higher decision levels must be disaggregated to a useful degree. In the mining context, oper-

ational decisions involve, for instance, allocating and dispatching mine haulage equipment, and scheduling train traffic in constrained railroads.

At the functional dimension, mining operations typically involve the following:

- Open-pit and underground mine scheduling, which involves determining an optimal production schedule over the life of the mineral deposit, from the feasibility study to the termination phase;
- Mine load and haulage equipment allocation and dispatching, which involves effectively deploying trucks and shovels according to some performance criteria, while maintaining a steady, reliable ore feed to the processing plants;
- Processing, which employs physical and chemical processes devoted to enhance ore quality by performing classification, concentration and agglomeration operations;
- Blending, which involves mixing ore by appropriate stockpiling and reclaiming, both at mine and port stockyards, in order to deliver ore at the required quality specifications;
- Railway scheduling and dispatch, which involves effectively routing and controlling train movement over a line, as well as planning meeting and passing of trains on single-line sections, while also attending to operational constraints;
- Port planning and scheduling, which involves determining an optimal plan to satisfy shipment demand, while minimizing delays in serving the ship arrival queue, and determining effective routes to convey ore products from the stockyard to the ship's holds.

A thorough analysis of Operations Research literature applied to mining operations is given in Pimentel et al. [2010a].

Large mining operations are often composed of several mines and processing plants, as well as stockyards, railroads, ports and supply stations, each with different capacities. According to the demand imposed by domestic and international customers, and depending on the characteristics of the ore products supplied by each mine, decisions such as mining, processing, handling, blending and shipping may involve different facilities at different stages of the supply chain. A Global Mining Supply Chain can thus be defined as *network of integrated facilities designed to process, using a variety of production techniques, and distribute, using a variety of transportation modals, bulk ore products from mines to customers, which can be (and usually are) at significant* *geographic distances*. This definition is consistent with that of a global production and supply network, as proposed by Tsiakis and Papageorgiou [2008].

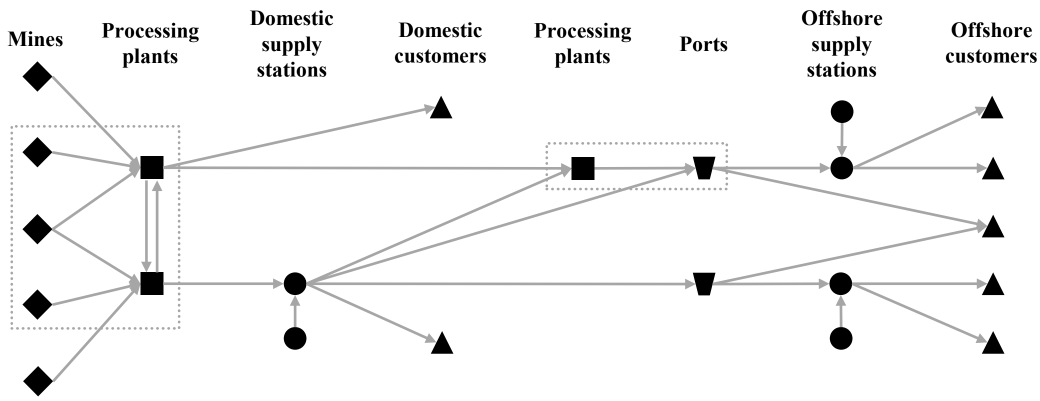

Figure 1.4. A Global Mining Supply Chain.

It is important to notice that end ore products in each mine may be considered intermediate products to the supply chain, as additional blending may occur either at the ports or at the supply stations. Supply stations may also operate as transshipment hubs, receiving ore products from the railway and dispatching them by road, or the other way around. Planning decisions are usually guided by expected revenue and variable and fixed costs modeled as linear or nonlinear functions on material flows and market conditions. In such a complex scenario, planning must be as detailed as lower is the decision level. However, the main opportunity, and also the main challenge in integrated planning of the mining supply chain is to explore the flexibility of the system to yield lower costs and increased product quality, especially in times of lower demand.

When inter-temporal integration comes into play, it is common to have strategic decisions defining high-level goals and constraints to tactical decisions, and the equivalent happens between tactical and operational decisions. If considered in isolation, decisions may turn out to be infeasible at lower levels, as higher-level decisions may consider averaged or aggregated data. For instance, a monthly production plan may generate infeasible daily production schedules if detailed operational time restrictions are larger than the averaged daily capacity. An integrated approach to optimal planning in the Global Mining Supply Chain should consider the framework [Pimentel et al., 2010b] depicted in Figure 1.5.

A final word must be said about uncertainty issues. All operations depicted above possess some degree of uncertainty and can significantly influence other downstream or upstream operations in the supply chain. For instance, unexpected downtime can decrease the nominal capacity and availability of all production and distribution elements

	Mine	Logistics	Port
Strategic	Strategic Mine Planning Capacity Investment	Capacity Investment	Capacity Investment
Tactical	Ultimate Pit Limit Production Planning	Loading & Shipping Planning Transportation Planning	Product Blending Loading & Shipping Planning
Operational	Truck & Shovel Allocation Routing Dispatching	Train Routing & Scheduling Yard Maneuvering Dispatching	Handling Scheduling Ship Queue Management

Figure 1.5. Mining decision problems according to functional and inter-temporal dimensions.

[Lin, 2009]. Uncertainties inherent to the orebody model can influence the processing plant throughput, both in quantity and quality, which in turn affects the expected production portfolio and makes additional blending required at stockyards and loading points. In the more general context of the supply chain, demand, market prices, operational costs and financial risks — varying taxes, duties, exchange rates and transfer prices — may vary unpredictably and with different global effects [Germain et al., 2008]. That would significantly affect the optimal production and distribution policies determined by solely deterministic approaches. Accounting for such uncertainties clearly brings additional complexity to an already difficult problem.

1.4 Scope and contributions

This Thesis presents a broad discussion on the most relevant planning and scheduling problems found in the mining industry, the corresponding mathematical programming approaches and the complex integrations that often arise in practice, but are seldom addressed in the literature. It is important to notice that we restrict the integration effort to the functional dimension of the Global Mining Supply Chain — that is, the integration between production and distribution stages on the same decision level. The integration between strategic, tactical and operational decisions — for instance, planning and scheduling decisions for a given stage — is left as future work.

The main contributions can be described as follows:

- In this Chapter, we have brought the supply chain concept to the mining industry scenario and discussed its inherent integrated planning problems. References to the corresponding mathematical programming approaches in strategic, tactical and operational decision levels can be found in [Pimentel et al., 2010a].

- In Chapter 2, we develop a novel multistage stochastic integer programming model to address the strategic capacity planning problem in a Global Mining Supply Chain. The model integrates capacitated facility location and network design decisions with economies of scale on the capacity costs. Also, the model is sensitive to both positive and negative variations in iron ore demand and allows the establishment of new (or capacity expansion of existing) production facilities and logistics channels, as well as the deactivation of specific facilities and channels in a temporary or permanent basis. We analyze the characteristics of the problem by means of an empirical study of different settings for the parameters of the CPLEX solver. Such analysis motivates the development of two different solution approaches: a Lagrangian heuristic, which attempts to improving the dual bounds and deriving quality primal bounds for the problem; and a Soft-Fixing Local Search approach, which attempts to determine good feasible solutions early in the solution process.

- In Chapter 3, we propose a mixed-integer programming approach to the integrated sales and operations tactical planning problem in a Global Mining Supply Chain. The model has characteristics of a lot sizing problem in a network environment, but presents challenging aspects related to integer flows, discrete production levels and flow conservation violations due to mass losses in processing plants and bulk transportation. We develop and evaluate a series of Relax&Fix strategies in order to tackle realistic sized problem instances and provide good primal bounds for the problem. The structure of the solution process also suggests the evaluation of the Soft Fixing Local Search approach as an alternative heuristic to the problem.

- A brief discussion on operational-level problems is given in Appendix A, where we present a mixed-integer goal programming model to address the integrated short-term programming of iron ore open pits, processing plants, stockyards and shipping operations. Additionally, we propose the concept of the Value of the Integrated Solution, which determines how valuable is solving a more complex integrated decision problem given the potential losses incurred when individual decisions are undertaken.

We conclude the text by discussing the main results presented in this Thesis and pointing out possible extensions and further research directions.

Chapter 2

Stochastic capacity planning in a Global Mining Supply Chain

> "'It will end well; almost certainly so for the project; and with reasonable probability for you.'
> 'What are the figures?' demanded Gaal.
> 'For the project, over 99.9%.'
> 'And for myself?'"
>
> *Foundation, Isaac Asimov*

At the strategic level, decisions must be made regarding major capital investments on the establishment of new (or on the capacity expansion of existing) mines, processing facilities, supply stations and logistics channels in order to satisfy increasing demand. However, in times of economic distress, strategic decisions should also consider whether or not to shutdown specific facilities in a temporary or permanent basis. In addition, the inherent risk associated to commodities' demand and price levels suggest a stochastic approach to strategic decision making. Hence, in this Chapter, we discuss those issues in a multistage stochastic mixed-integer programming model which addresses the capacity planning problem throughout the Global Mining Supply Chain. In a first study, we evaluate the main characteristics of the problem and broadly discuss possible solution approaches in order to drive the development of specific algorithms and heuristics. We explore one of such approaches, the Lagrangian Relaxation, which is employed as a means to determine, for large problem instances, good feasible solutions in a reasonable amount of time when compared to CPLEX. In yet another study, the ability of determining good feasible solutions in the early stages of the computation is addressed in a soft-fixing local search framework, which is evaluated against the other solution approaches discussed here.

The results given in this Chapter are also presented in the following publication:

Pimentel, B. S., Mateus, G. R., and Almeida, F. A. (2011). Stochastic capacity planning in a global mining supply chain. In *IEEE Workshop on Computational Intelligence in Production and Logistics Systems*, April 2011. Paris, France.

2.1 Introduction

The present decade has shown extreme variations in ore prices and demand. Such variations affect not only mining operations, but also all downstream production chain up to the ore consumers. Strategic, tactical and operational decisions must then be in accordance to market conditions, since it is the ultimate goal of any commodity enterprise to minimize its operational cost and maximize its throughput.

The recent financial crisis has proven that developing accurate predictions of commodities' prices and demand can be extremely difficult. The complexity of today's economical dynamics, the increasing levels of uncertainty and information asymmetry and the intricate relationships between productive and financial markets can be overwhelming [Kimura, 2002]. Such a scenario drives the motivation of contemplating stochastic aspects on analytics studies, especially those involving strategic decisions on high capital expenses. In a growing economy environment, when ore prices and demand are high, mining companies tend to invest in establishing new ventures, expanding capacity, or even consolidating the market. High selling prices improve the feasibility of lower grade operations and of ore producers that are geographically distant from consumer markets, especially if prices are significantly higher than the associated freight costs. Customers tend to prioritize quantity over quality and there is a higher pressure on the availability of mining production and distribution systems. However, in an unstable or decaying economic environment, the situation is quite the opposite. Demand and prices drop, and customers tend to prioritize quality over quantity, with a higher pressure on reducing mining production and distribution costs. Such dynamic scenario has driven large scale base metals mining companies to pursue higher efficiency levels in their global operations. Operations Research provides the basis to achieve such goals by adequately modeling and planning operations in an integrated framework.

In this Chapter, we are interested in developing a mathematical programming model for a stochastic capacity planning problem applied to a Global Mining Supply Chain. In this context, strategic planning should be devoted to evaluate major capital investments on the establishment of new (or on capacity expansion of existing) mines,

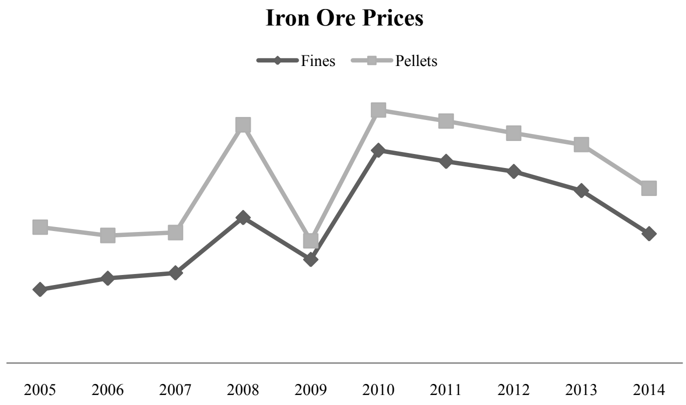

Figure 2.1. Qualitative variation of iron ore prices, historical and forecast.

processing facilities, supply stations and logistics channels in order to satisfy increasing demand. However, in times of decreasing demand, strategic decisions should also include whether or not to shutdown specific facilities in a temporary or permanent basis [Dias et al., 2006]. Those decisions may have to be made according to budget limitations for each time period, which can also affect other strategic supply chain decisions [Dogan and Goetschalckx, 1999]. Fixed costs, expected revenues and return on investment are thus the main drivers to strategic decisions in the mining supply chain. It is important to emphasize that capacity expansion and network design problems are known to be hard to solve. When risk is considered, computational complexity is further increased and special care must be taken when developing the mathematical formulation and, more importantly, the solution approach.

2.2 Capacity planning

We assume the term *capacity planning* to encompass capacitated facility location, network design and capacity expansion decisions. The first two are concerned with selecting the time schedule for installing resources at different locations in order to minimize, for instance, the total discounted costs for meeting customer demands specified over a planning horizon [Erlenkotter, 1981]. The third, on the other hand, aims at determining the additional capacity and the associated times at which it should be added so that the net present investment is minimized [Luss, 1982].

Facility location decisions play a critical role in the strategic design of supply chain networks. In a discrete facility location problem, the selection of the sites where new facilities are established is restricted to a finite set of candidate locations [Melo et al., 2009]. In many practical situations, however, the optimal topology of the underlying network must be determined together with the facility location decision, thus configuring a capacitated facility location/network design problem [Melkote and Daskin, 2001]. Solution strategies for those problems range from Lagrangian Relaxation [Shulman, 1991] and decomposition [Dogan and Goetschalckx, 1999; Singh et al., 2009] to a number of metaheuristics [Keskin and Üster, 2007].

It is important to notice that investments in capacity planning — in the form of new facilities/channels or discrete capacity increments — commonly exhibit substantial economies of scale. This means that any strategic investment policy should consider the tradeoff between the economies-of-scale savings of large expansion sizes versus the cost of installing capacity before it is actually needed. The capacity cost function is typically concave, strictly increasing, non-negative, and often specified as a piecewise, fixed charge linear approximation. Decisions of opening, closing and reopening facilities and logistics channels may then share specific costs.

The mathematical formulation for the capacity planning problem in the Global Mining Supply Chain is a novel approach that integrates lot sizing, capacity expansion, facility location and network design decisions such as:

- ROM production levels and its corresponding transformation to sinter-feed and pellet-feed at the processing plants, as well as the transformation of pellet-feed in pellets at the pelletizing plants;

- Iron ore flow between mines, plants, ports, supply stations and customers using trucks, railways, ducts or vessels;

- Dynamic capacitated facility location and network design, namely opening, shutting down and reopening facilities and logistics channels according to demand;

- Incremental capacity expansion on facilities and logistics channels, original or newly established.

It is important to notice that, since any facility should always be economically viable, a minimum level of production must be allocated at each period in order to cover the fixed operational costs. However, the combination of fixed charge and variable costs adequately represents that feature. Also, depletion of iron ore reserves are assumed to be independent on exploitation, at least for the duration of the planning horizon. The mathematical notation follows.

2.2.1 Mathematical notation

Sets

D : Customers (demand centers);

F_O : Existing facilities, which are already opened at the beginning of the planning horizon;

F_C : Candidate facilities, which can be opened from the beginning of the planning horizon;

F : All facilities, $F = \{F_O \cup F_C\}$;

L_O : Existing logistics channels, $L_O : \{(i,j)|i,j \in F_O\}$, which are already opened at the beginning of the planning horizon;

L_C : Candidate logistics channels, $L_C : \{(i,j)|i,j \in F\}$, which can be opened from the beginning of the planning horizon;

L : All logistics channels, $L = \{L_O \cup L_C\}$;

Q : Ore products, raw or processed;

$\hat{Y}$: Capacity increment magnitude levels;

$\bar{Y}$: Initial capacity magnitude levels.

Parameters

T : Number of time periods;

d_i^{qt} : Demand for product q at customer i and node n;

θ_i^n : Capital expense of opening $i \in F_C$ at time t;

θ_{ij}^t : Capital expense of establishing $(i,j) \in L_C$ at time t;

ι_i^t : Capital expense of closing facility $i \in F$ at time t;

ι_{ij}^t : Capital expense of closing channel $(i,j) \in L$ at time t;

π_i^t : Capital expense of reopening facility $i \in F$ at n;

π_{ij}^t : Capital expense of reopening channel $(i,j) \in L$ at t;

α_i^t : Fixed operational cost for facility $i \in F$ at time t;

α_{ij}^{t} : Fixed operational cost for channel $(i, j) \in L$ at time t;

β^{qt} : Penalties for unmet demand of product q at t;

δ_{i}^{qt} : Fixed cost of established facility capacity;

δ_{ij}^{qt} : Fixed cost of established channel capacity;

$\bar{\kappa}_{i}^{q}$: Initial capacity for production of q at $i \in F_O$;

$\bar{\kappa}_{ij}$: Initial capacity of channel $(i, j) \in L_O$;

$\hat{\mu}_{ik}^{qt}$: Cost of incrementing capacity at facility i, on level $k \in \hat{Y}$;

$\hat{\mu}_{ijk}^{t}$: Cost of incrementing capacity at channel (i, j), on level $k \in \hat{Y}$;

$\hat{\upsilon}_{ik}^{qt}$: Magnitude of capacity increment $k \in \hat{Y}$ at facility $i \in F$;

$\hat{\upsilon}_{ijk}^{t}$: Magnitude of capacity increment $k \in \hat{Y}$ at channel $(i, j) \in L$;

ϕ_{i} : Limit on capacity increment at facility $i \in F$;

ϕ_{ij} : Limit on capacity increment at channel $(i, j) \in L$;

$\bar{\mu}_{ik}^{qt}$: Cost of initial capacity at $i \in F_C$, level $k \in \bar{Y}$;

$\bar{\mu}_{ijk}^{t}$: Cost of initial capacity at $(i, j) \in L_C$, level $k \in \bar{Y}$;

$\bar{\upsilon}_{ik}^{qt}$: Magnitude of initial capacity $k \in \bar{Y}$ at $i \in F_C$;

$\bar{\upsilon}_{ijk}^{t}$: Magnitude of initial capacity $k \in \bar{Y}$ at $(i, j) \in L_C$;

σ_{i}^{t} : Production split of sinter and pellet-feed, $\sigma_{i}^{t} \in [0, 1]$;

ρ_{i}^{qt} : Unit production cost of q at $i \in F$;

χ_{ij}^{qt} : Unit transportation cost of q through $(i, j) \in L$;

η_{i}^{qt} : Process efficiency for q at $i \in F$, $\eta_{i}^{qt} \in [0, 1]$.

Variables

a_{i}^{t} : Determines whether facility $i \in F$ should be operational at time t, $a_{i}^{t} \in \{0, 1\}$;

a_{ij}^{t} : Determines whether channel $(i, j) \in L$ should be operational at time t, $a_{ij}^{t} \in \{0, 1\}$;

b^{qt} : Determines the amount of unmet demand for product $q \in Q$ at time t;

c_i^t : Determines whether facility $i \in F$ should be closed at time t, $c_i^t \in \{0,1\}$;

c_{ij}^t : Determines whether channel $(i,j) \in L$ should be closed at time t, $c_{ij}^t \in \{0,1\}$;

o_i^t : Determines whether facility $i \in F_C$ should be opened at time t, $o_i^t \in \{0,1\}$;

o_{ij}^t : Determines whether channel $(i,j) \in L_C$ should be established at time t, $o_{ij}^t \in \{0,1\}$;

p_i^{qt} : Determines the amount of product $q \in Q$ produced at facility $i \in F$ at time t;

r_i^t : Determines whether facility $i \in F$ should be reopened at time t, $r_i^t \in \{0,1\}$;

r_{ij}^t : Determines whether channel $(i,j) \in L$ should be reestablished at time t, $r_{ij}^t \in \{0,1\}$;

$\hat{u}_i^{qt}$: Determines the capacity increment at facility $i \in F$ at time t;

$\hat{u}_{ij}^t$: Determines the capacity increment of channel $(i,j) \in L$ at time t;

$\bar{u}_i^{qt}$: Determines the initial capacity at facility $i \in F_C$ at time t;

$\bar{u}_{ij}^t$: Determines the initial capacity of channel $(i,j) \in L_C$ at time t;

x_{ij}^{qt} : Determines the amount of product $q \in Q$ transported though channel $(i,j) \in L$ at t;

$\hat{y}_{ik}^{qt}$: Determines the selection of technology/magnitude $k \in \hat{Y}$ for the capacity increment decision of producing $q \in Q$ at facility $i \in F$ at time t, $\hat{y}_{ik}^{qt} \in \{0,1\}$;

$\hat{y}_{ijk}^t$: Determines the selection of technology/magnitude $k \in Y$ for the capacity increment decision at channel $(i,j) \in L$ at time t, $\hat{y}_{ijk}^t \in \{0,1\}$;

$\bar{y}_{ik}^{qt}$: Determines the selection of technology/magnitude $k \in Y$ for the initial capacity of producing $q \in Q$ at facility $i \in F_C$ at time t, $\bar{y}_{ik}^{qt} \in \{0,1\}$;

$\bar{y}_{ijk}^t$: Determines the selection of technology/magnitude $k \in Y$ for the initial capacity of channel $(i,j) \in L_C$ at time t, $\bar{y}_{ijk}^t \in \{0,1\}$.

2.2.2 Mathematical formulation

The objective function seeks minimizing investment and operational costs throughout the planning horizon:

$$\begin{aligned}
&\text{Minimize} \\
&\sum_t \sum_q \sum_i \rho_i^{qt} p_i^{qt} + \sum_t \sum_i \alpha_i^t a_i^t + \\
&\sum_t \sum_q \sum_i \sum_j \chi_{ij}^{qt} x_{ij}^{qt} + \sum_t \sum_i \sum_j \alpha_{ij}^t a_{ij}^t + \\
&\sum_t \sum_q \sum_k \sum_i \hat{\mu}_{ik}^{qt} \hat{y}_{ik}^{qt} + \sum_t \sum_q \sum_k \sum_i \bar{\mu}_{ik}^{qt} \bar{y}_{ik}^{qt} + \\
&\sum_t \sum_k \sum_i \sum_j \hat{\mu}_{ijk}^{t} \hat{y}_{ijk}^{t} + \sum_t \sum_k \sum_i \sum_j \bar{\mu}_{ijk}^{t} \bar{y}_{ijk}^{t} + \\
&\sum_t \sum_q \sum_i \delta_i^{qt} u_i^{qt} + \\
&\sum_t \sum_i \sum_j \delta_{ij}^{t} u_{ij}^{t} + \\
&\sum_t \sum_i \theta_i^t o_i^t + \sum_t \sum_i \iota_i^t c_i^t + \sum_t \sum_i \pi_i^t r_i^t + \\
&\sum_t \sum_i \sum_j \theta_{ij}^t o_{ij}^t + \sum_t \sum_i \sum_j \iota_{ij}^t c_{ij}^t + \sum_t \sum_i \sum_j \pi_{ij}^t r_{ij}^t
\end{aligned} \tag{2.1}$$

where several constraints apply, as described in the following sections.

2.2.2.1 The underlying production and transportation planning problem

For any active facility, the amount of ore produced, increased by the ore received must be greater than the amount of ore shipped to other facilities or customers:

$$p_i^{qt} + \sum_k x_{ki}^{qt} - \sum_j x_{ij}^{qt} \geq 0 \qquad \forall i \in F, q \in Q, t. \tag{2.2}$$

For any active facility where ore processing occurs, production must be limited by the amount of input material available, coming from the mine proper or transferred to other facilities:

$$p_i^{qt} \leq p_i^{rt} + \sum_k x_{ki}^{rt} - \sum_j x_{ij}^{rt} \qquad \forall i \in F, q \in Q, r \in Q, t. \tag{2.3}$$

It is important to notice that in some facilities, there may be coproduction of sinter-feed and pellet-feed, what imposes additional constraints on the production of the latter:

$$p_i^{\mathrm{PF},t} = \left(\frac{\sigma_i^t}{1-\sigma_i^t}\right) p_i^{\mathrm{SF},t} \qquad \forall i \in F, t, \tag{2.4}$$

including the special case $\sigma_i^n = 0$, where facility i has no concentration equipment and thus no pellet-feed is produced. But for every product, demand must be met with ore produced or acquired from third-parties:

$$\sum_k x_{ki}^{qt} \geq d_i^{qt} \qquad \forall i \in D, q \in Q, t. \tag{2.5}$$

Also, production and distribution capacities are finite and limited by physical constraints:

$$0 \leq p_i^{qt} \leq u_i^{qt} \qquad \forall i \in F, q \in Q, t, \tag{2.6}$$

$$0 \leq \sum_q x_{ij}^{qt} \leq u_{ij}^t \qquad \forall i \in F, j \in F, t. \tag{2.7}$$

And any production or transportation effort can only take place at active facilities and logistics channels:

$$\sum_q p_i^{qt} \leq u_i^{q*} a_i^t \qquad \forall i \in F, t, \tag{2.8}$$

$$\sum_q x_{ij}^{qt} \leq u_{ij}^* a_{ij}^t \qquad \forall (i,j) \in L, t \tag{2.9}$$

$$a_{ij}^t \leq a_i^t \qquad \forall (i,j) \in L, i \in F, t. \tag{2.10}$$

where u_i^{q*} and u_{ij}^* are known maximum theoretical capacities for facilities and logistics channels, respectively. Also, the following nonnegativity and integrality restrictions apply:

$$p_i^{qt}, x_{ij}^{qt} \geq 0 \qquad \forall i \in F, (i,j) \in L, q \in Q, t, \tag{2.11}$$

$$a_i^t, a_{ij}^t \in \{0, 1\} \qquad \forall i \in F, (i,j) \in L, t. \tag{2.12}$$

2.2.2.2 The dynamic facility location/network design problem

When a new facility is opened, it can be done from the beginning of the planning horizon, according to:

$$a_i^t \leq \sum_{\tau=1}^{t} o_i^\tau \qquad \forall i \in F_C, t, \tag{2.13}$$

$$\sum_{t=1}^{T} o_i^t \leq 1 \qquad \forall i \in F_C. \tag{2.14}$$

The establishment of new logistics channels, however, must be restricted to connections between facilities that are or will be opened at some point within the planning horizon:

$$o_{ij}^t \leq \sum_{\tau=1}^{t} o_i^\tau \qquad \forall (i,j) \in L_C, i \in F_C, t, \tag{2.15}$$

$$o_{ij}^t \leq 1 \qquad \forall (i,j) \in L_C, i \in F_C, j \in F_O, t. \tag{2.16}$$

Whenever an original facility is closed, it can be done from the beginning of the planning horizon. On the other hand, when a candidate facility is closed, it can only be done from $t > 1$, since it has to be opened at least a period before that:

$$c_i^t - a_i^{t-1} + a_i^t \geq 0 \qquad \forall t \text{ if } i \in F_O \text{ and } \forall t > 1 \text{ if } i \in F_C, \tag{2.17}$$

with $a_i^{t-1} = 1$ if $t = 1$. But for both original and candidate sets, a facility can only be closed if it has been opened before:

$$c_i^t \leq 1 \qquad \forall i \in F_O, t, \tag{2.18}$$

$$c_i^t \leq \sum_{\tau=1}^{t-1} o_i^\tau \qquad \forall i \in F_C, t > 1, \tag{2.19}$$

and if it was active at the end of the previous period:

$$c_i^t \leq a_i^{t-1} \qquad \forall i \in F, t > 1. \tag{2.20}$$

If a closed facility must be reopened, it can be done from $t > 1$ for original facilities, and from $t > 2$ for candidate ones, since they must have been opened and closed before that:

$$r_i^t - a_i^t + a_i^{t-1} \geq 0 \qquad \forall i \in F_O, t > 1, \tag{2.21}$$

$$r_i^t - a_i^t + a_i^{t-1} - o_i^t \geq 0 \qquad \forall i \in F_C, t > 2. \tag{2.22}$$

In the latter inequality, the term o_i^t is used to make the constraint inactive before facility $i \in F_C$ is opened. And again, any facility can only be reopened if it has been closed before:

$$r_i^t \leq \sum_{\tau=1}^{t-1} c_i^\tau \qquad \forall i \in F_O, t > 1, \tag{2.23}$$

$$r_i^t \leq \sum_{\tau=1}^{t-1} c_i^\tau \qquad \forall i \in F_C, t > 2, \tag{2.24}$$

and if it was inactive at the end of the previous period:

$$r_i^t \leq 1 - a_i^{t-1} \qquad \forall i \in F_O, t > 1, \tag{2.25}$$

$$r_i^t \leq 1 - a_i^{t-1} \qquad \forall i \in F_C, t > 2. \tag{2.26}$$

All facility location and network design variables are binary:

$$o_i^t, c_i^t, r_i^t \in \{0, 1\} \qquad \forall i \in F, t. \tag{2.27}$$

Analogous constraints are defined to the logistics channels whenever there is a need to open, close or reopen them to meet the system's requirements of material flows:

$$a_{ij}^t \leq \sum_{\tau=1}^{t} o_{ij}^\tau \qquad \forall (i,j) \in L_C, t, \tag{2.28}$$

$$\sum_{t=1}^{T} o_{ij}^t \leq 1 \qquad \forall (i,j) \in L_C, \tag{2.29}$$

$$c_{ij}^t - a_{ij}^{t-1} + a_{ij}^t \geq 0 \qquad \forall t \text{ if } (i,j) \in L_O \text{ and } \forall t > 1 \text{ if } (i,j) \in L_C \tag{2.30}$$

$$c_{ij}^t \leq 1 \qquad \forall (i,j) \in L_O, t, \tag{2.31}$$

$$c_{ij}^t \leq \sum_{\tau=1}^{t-1} o_{ij}^\tau \qquad \forall (i,j) \in L_C, t > 1, \tag{2.32}$$

$$c_{ij}^t \leq a_{ij}^{t-1} \qquad \forall (i,j) \in L, t > 1, \tag{2.33}$$

$$r_{ij}^t - a_{ij}^t + a_{ij}^{t-1} \geq 0 \qquad \forall (i,j) \in L_O, t > 1, \tag{2.34}$$

$$r_{ij}^t - a_{ij}^t + a_{ij}^{t-1} - o_{ij}^t \geq 0 \qquad \forall (i,j) \in L_C, t > 2, \tag{2.35}$$

$$r_{ij}^t \leq \sum_{\tau=1}^{t-1} c_{ij}^\tau \qquad \forall (i,j) \in L_O, t > 1, \tag{2.36}$$

$$r_{ij}^t \leq \sum_{\tau=1}^{t-1} c_{ij}^{\tau} \qquad \forall (i,j) \in L_C, t > 2, \tag{2.37}$$

$$r_{ij}^t \leq 1 - a_{ij}^{t-1} \qquad \forall (i,j) \in L_O, t > 1, \tag{2.38}$$

$$r_{ij}^t \leq 1 - a_{ij}^{t-1} \qquad \forall (i,j) \in L_C, t > 2, \tag{2.39}$$

$$o_{ij}^t, c_{ij}^t, r_{ij}^t \in \{0,1\} \qquad \forall (i,j) \in L, t. \tag{2.40}$$

2.2.2.3 Capacity expansion

Throughout the planning horizon, investments can be made in order to expand capacity from initial levels according to discrete increments:

$$u_i^{qt} = \sum_{\tau=1}^{t} \hat{u}_i^{q\tau} + \sum_{\tau=1}^{t} \bar{u}_i^{q\tau} \qquad \forall i \in F, q \in Q, t, \tag{2.41}$$

$$u_{ij}^t = \sum_{\tau=1}^{t} \hat{u}_{ij}^{\tau} + \sum_{\tau=1}^{t} \bar{u}_{ij}^{\tau} \qquad \forall i \in F, j, t, \tag{2.42}$$

where $\bar{u}_i^{q\tau}$ and $\bar{u}_{ij}^{\tau}$ are known for original facilities and logistics channels at the initial period. Capacity increments at existing facilities and channels are selected from a choice of technology and magnitude options with economies of scale:

$$\hat{u}_i^{qt} = \sum_{k \in \hat{Y}} \hat{\upsilon}_{ik}^{qt} \hat{y}_{ik}^{qt} \qquad \forall i \in F, q \in Q, t, \tag{2.43}$$

$$\hat{u}_{ij}^t = \sum_{k \in \hat{Y}} \hat{\upsilon}_{ijk}^t \hat{y}_{ijk}^t \qquad \forall (i,j) \in L, t, \tag{2.44}$$

and the same rationale applies to the initial capacity decision:

$$\bar{u}_i^{qt} = \sum_{k \in \bar{Y}} \bar{\upsilon}_{ik}^{qt} \bar{y}_{ik}^{qt} \qquad \forall i \in F_C, q \in Q, t, \tag{2.45}$$

$$\bar{u}_{ij}^t = \sum_{k \in \bar{Y}} \bar{\upsilon}_{ijk}^t \bar{y}_{ijk}^t \qquad \forall (i,j) \in L_C, t, \tag{2.46}$$

but in both cases, only one choice at each period is possible, and restricted to assets that are or will be opened within the planning horizon:

$$\sum_{k \in \hat{Y}} \hat{y}_{ik}^{qt} \leq \sum_{\tau=1}^{t-1} o_i^{\tau} \qquad \forall i \in F, q \in Q, t > 1 \tag{2.47}$$

$$\sum_{k \in \hat{Y}} \hat{y}_{ijk}^t \leq \sum_{\tau=1}^{t-1} o_{ij}^{\tau} \qquad \forall (i,j) \in L, t > 1 \tag{2.48}$$

$$\sum_{k \in \bar{Y}} \bar{y}_{ik}^{qt} \leq o_i^t \qquad \forall i \in F_C, q \in Q, t > 1 \tag{2.49}$$

$$\sum_{k \in \bar{Y}} \bar{y}_{ijk}^{t} \leq o_{ij}^t \qquad \forall (i,j) \in L_C, t > 1 \tag{2.50}$$

$$\tag{2.51}$$

Capacity increments must, however, be limited by practical constraints. Beyond that limit, any capacity expansion would have to be accomplished by establishing new assets:

$$\sum_{t=1}^{T} \hat{u}_i^{qt} \leq \phi \sum_{t=1}^{T} \bar{u}_i^{qt} \qquad \forall i \in F, q, \in Q, t, \tag{2.52}$$

$$\sum_{t=1}^{T} \hat{u}_{ij}^{t} \leq \phi \sum_{t=1}^{T} \bar{u}_{ij}^{t} \qquad \forall i \in F, t. \tag{2.53}$$

And again, the following nonnegativity and integrality restrictions apply:

$$\bar{u}_i^{qt}, \bar{u}_{ij}^t \in \mathbb{Z}_+ \qquad \forall i \in F_C, (i,j) \in L_C, q \in Q, t, \tag{2.54}$$

$$\hat{u}_i^{qt}, \hat{u}_{ij}^t \in \mathbb{Z}_+ \qquad \forall i \in F, (i,j) \in L, q \in Q, t, \tag{2.55}$$

$$u_i^{qt}, u_{ij}^t \geq 0 \qquad \forall i \in F, (i,j) \in L, q \in Q, t, \tag{2.56}$$

$$\hat{y}_{ik}^{qt}, \hat{y}_{ijk}^t, \bar{y}_{ik}^{qt}, \bar{y}_{ijk}^t \in \{0,1\} \qquad \forall i \in F_C, (i,j) \in L_C, k \in Y, q \in Q, t. \tag{2.57}$$

2.3 The stochastic approach

Due to the inherent uncertain behavior of the commodity market, a stochastic approach is a natural extension to the previous formulation. We adopt the concept of *decision making under risk*, in which every action leads to one of a set of possible specific outcomes, each occurring with a probability known to the decision maker [Luce and Raiffa, 1957]. A stochastic supply chain network design problem can then be tackled statically or dynamically. In the static approach, the first-stage decides the network configuration, while the recourse second-stage distributes products in an optimal fashion based on the configuration decided at the first stage and on the realization of the risky parameters [Santoso et al., 2005]. In the dynamic approach, an integer multi-stage stochastic formulation employs a scenario tree to discretize the evolution of the risky parameters, with fixed-charge cost functions accounting for economies of scale in expansion costs. Solution approaches include reformulation, decomposition, scenario-tree reduction [de Oliveira et al., 2010] and metaheuristics [Ahmed et al., 2003].

In the proposed stochastic approach, we assume the random parameters to evolve as a discrete time stochastic process with a finite probability space. Such an information structure can be represented by a tree $\mathcal{T}$ where each node $n \in \mathcal{T}$ represents a distinguishable state of world and each path from the root node to a leaf node constitutes a scenario $s \in \mathcal{S}$. That is a rather reasonable assumption since scenario planning [Schoemaker, 1995] has been used by organizations to design flexible strategic long-term plans.

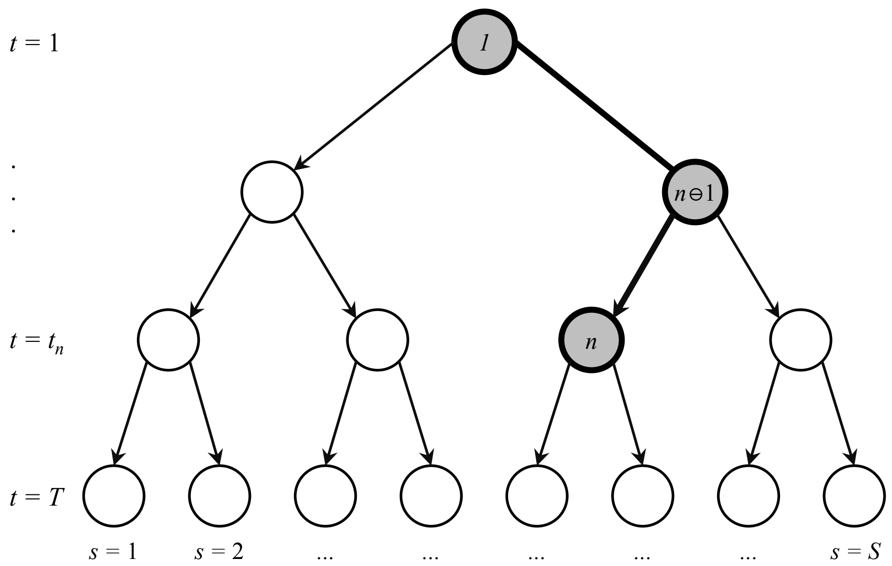

Figure 2.2. A general binary scenario tree.

Figure 2.2 illustrates a binary scenario tree. Nodes at any given level are all related to the same time period t of the planning horizon. Each scenario s is thus represented by the path from node 1 to the corresponding leaf node. Also, for each node n, there can be only one parent node $n \ominus 1$. Let $\mathbf{p}(.)$ be a function describing the probability of realizing a given event or state and let (i,j) represent a transition between any consecutive nodes i and j in the tree. The probability of the realization of any node $\mathbf{p(n)}$ is then given by:

$$\mathbf{p(n)} = \prod_{(i,j)\in\mathcal{Q}(n)} \mathbf{p}(i,j),$$

where $\mathcal{Q}(n)$ represents the set of (i,j) transitions in the scenario tree starting from the root node and reaching node n. That means that any node can be realized with the joint probability of every previous transitions within the path from the root of the scenario tree. For node n in Figure 2.2, we have $\mathbf{p(n)} = \mathbf{p}\,(1, n\ominus 1) \times \mathbf{p}\,(n\ominus 1, n)$.

The mathematical model then takes the general form:

$$
\begin{aligned}
z = \min \sum_{n \in \mathcal{T}} \mathbf{p}(n) c^n x^n \\
A x^n \geq b^n \\
x \geq 0,
\end{aligned}
$$

The formulation presented in the following section is based on the form above. One should notice that the sum of the probabilities of all nodes at a same level of the scenario tree (the same time period) must equal 1, since at any time period, only one possible state of the world can be realized. However, the same does not apply to the sum of all $\mathbf{p}(n), n \in \mathcal{T}$ in the scenario tree. Alternatively, one could also consider the probabilities of each scenario s, that is, the probability that a given path starting at the root node would reach a particular leaf node n. That would require the model to take the form:

$$
\begin{aligned}
z = \min \sum_{s \in \mathcal{S}} \sum_{n \in \vec{\mathcal{Q}}(s)} \mathbf{p}(s) c^n x^n \\
A x^n \geq b \\
x \geq 0,
\end{aligned}
$$

where, $\vec{\mathcal{Q}}(s)$ is the sequence of transitions in scenario $s \in \mathcal{S}$. This form, nevertheless, requires that $\sum_{s \in \mathcal{S}} \mathbf{p}(s) = 1$. These two general forms are, of course, equivalent.

The above discussion is also related to two distinct approaches to formulate a stochastic capacity planning problem. A first option is to define decision variables for every scenario at every time period within the planning horizon. In this case, a subscript s would have to be added to every parameter and every decision variable, which in turn would demand an explicit formulation of the nonanticipativity constraints. A second option, on the other hand, is to define decision variables for each node of the scenario tree, addressing the realizations of the random variables up to that node, and to implicitly consider the nonanticipativity constraints. In this work, we use this approach since it has the advantage of eliminating duplicate constraints and avoiding the nonanticipativity constraints which, although not representing significant computational gains, provides a more concise formulation. In both options, however, there is the challenge of determining a consistent scenario tree, which can become overwhelmingly complex for longer planning horizons. Thus, with the explicit inclusion of the nonanticipativity constraints, a scenario-based, multi-period stochastic program can be recast in the form of an equivalent large-scale deterministic program [Dupacová, 1999].

It is important to notice that any combination of risky parameters can be considered in the scenario tree. Of course, the larger that number, the harder it is to model the scenario tree, the more complex becomes its inherent structure and the larger is the problem size [Heitsch and Römisch, 2009].

2.3.1 Mathematical formulation

In the proposed stochastic mathematical formulation, we assume, without loss of generality, ore demand to be the only risky parameter — in general supply chains, several other parameters may also be considered risky, such as supplies, processing, transportation, shortage and capacity expansion costs [Azaron et al., 2008]. That means that each node at the scenario tree describes a possible realization of an expected demand level occurring with a known probability. In order to recast the formulation in a stochastic setting, the following must be considered in the model:

$\mathcal{S}$: Set of scenarios;

$\mathcal{T}$: A scenario tree and all its nodes;

n : A node in the scenario tree;

$n \ominus 1$: (Unique) Predecessor of node n;

$\mathcal{P}(n)$: Path from the root node of the scenario tree to node n; if n is a leaf node, then $\mathcal{P}(n)$ is a *scenario*;

$\mathbf{p}(n)$: Probability of node n.

Hence, instead of time periods, each parameter and variable must be indexed according to the corresponding node of the scenario tree. Any sequence of time periods must be recast as a path within the tree. For instance, the relationship:

$$c_i^t \leq \sum_{\tau=1}^{t-1} o_i^\tau \qquad \forall i \in F_C, t > 1,$$

must be recast in the stochastic settings as:

$$c_i^n \leq \sum_{m \in \mathcal{P}(n \ominus 1)} o_i^m \qquad \forall i \in F_C, n > 1.$$

In order to make the formulation more concise, we assume that variables and parameters of the form $*_i^{qn}$ are associated to facility $i \in F$, while variables and parameters of the form $*_{ijk}^{qn}$ are associated to logistics channel $(i,j) \in L$. The stochastic mixed-integer program then follows:

Minimize

$$\sum_{n \in \mathcal{T}} \mathbf{p}(n) \Bigg(\sum_q \sum_i \rho_i^{qn} p_i^{qn} + \sum_i \alpha_i^n a_i^n + \\ \sum_q \sum_i \sum_j \chi_{ij}^{qn} x_{ij}^{qn} + \sum_i \sum_j \alpha_{ij}^n a_{ij}^n + \\ \sum_q \sum_k \sum_i \hat{\mu}_{ik}^{qn} \hat{y}_{ik}^{qn} + \sum_q \sum_k \sum_i \bar{\mu}_{ik}^{qn} \bar{y}_{ik}^{qn} + \\ \sum_k \sum_i \sum_j \hat{\mu}_{ijk}^n \hat{y}_{ijk}^n + \sum_k \sum_i \sum_j \bar{\mu}_{ijk}^n \bar{y}_{ijk}^n + \\ \sum_q \sum_i \delta_i^{qn} u_i^{qn} + \sum_i \sum_j \delta_{ij}^n u_{ij}^n + \\ \sum_i \theta_i^n o_i^n + \sum_i \iota_i^n c_i^n + \sum_i \pi_i^n r_i^n + \\ \sum_i \sum_j \theta_{ij}^n o_{ij}^n + \sum_i \sum_j \iota_{ij}^n c_{ij}^n + \sum_i \sum_j \pi_{ij}^n r_{ij}^n \Bigg), \tag{2.58}$$

Subject to

$$p_i^{qn} + \sum_k x_{ki}^{qn} - \sum_j x_{ij}^{qn} \geq 0 \qquad \forall i \in F, q \in Q, n \in \mathcal{T}, \tag{2.59}$$

$$p_i^{qn} \leq p_i^{rn} + \sum_k x_{ki}^{rn} - \sum_j x_{ij}^{rn} \qquad \forall i \in F, q, r \in Q, n, \tag{2.60}$$

$$p_i^{\mathrm{PF},n} = \left(\frac{\sigma_i^n}{1 - \sigma_i^n} \right) p_i^{\mathrm{SF},n} \qquad \forall i \in F, n, \tag{2.61}$$

$$\sum_k x_{ki}^{qn} + b_i^{qn} \geq d_i^{qn} \qquad \forall i \in D, q \in Q, n, \tag{2.62}$$

$$0 \leq p_i^{qn} \leq u_i^{qn} \qquad \forall i \in F, q \in Q, n, \tag{2.63}$$

$$0 \leq \sum_q x_{ij}^{qn} \leq u_{ij}^n \qquad \forall (i,j) \in L, n, \tag{2.64}$$

$$\sum_q p_i^{qn} \leq u_i^{q*} a_i^n \qquad \forall i \in F, n, \tag{2.65}$$

$$\sum_q x_{ij}^{qn} \leq u_{ij}^* a_{ij}^n \qquad \forall (i,j) \in L, n, \tag{2.66}$$

$$a_{ij}^n \leq a_i^n \qquad \forall (i,j) \in L, i \in F, n, \tag{2.67}$$

$$p_i^{qn}, x_{ij}^{qn} \geq 0 \qquad \forall i \in F, (i,j) \in L, q \in Q, n, \tag{2.68}$$

$$a_i^n, a_{ij}^n \in \{0,1\} \qquad \forall i \in F, (i,j) \in L, n, \tag{2.69}$$

$$a_i^n \leq \sum_{m \in \mathcal{P}(n)} o_i^m \qquad \forall i \in F_C, n, \tag{2.70}$$

$$\sum_{n \in s} o_i^n \leq 1 \qquad \forall i \in F_C, s \in \mathcal{S}, \tag{2.71}$$

$$o_{ij}^n \leq \sum_{m \in \mathcal{P}(n)} o_i^m \qquad \forall (i,j) \in L_C, i \in F_C, n, \tag{2.72}$$

$$o_{ij}^n \leq 1 \qquad \forall (i,j) \in L_C, i \in F_C, j \in F_O, \tag{2.73}$$

$$c_i^n - a_i^{n\ominus 1} + a_i^n \geq 0 \qquad \forall n \text{ if } i \in F_O \text{ and } \forall n > 1 \text{ if } i \in F_C, \tag{2.74}$$

$$c_i^n \leq 1 \qquad \forall i \in F_O, n, \tag{2.75}$$

$$c_i^n \leq \sum_{m \in \mathcal{P}(n\ominus 1)} o_i^m \qquad \forall i \in F_C, n > 1, \tag{2.76}$$

$$c_i^n \leq a_i^{n\ominus 1} \qquad \forall i \in F, n, \tag{2.77}$$

$$r_i^n - a_i^n + a_i^{n\ominus 1} \geq 0 \qquad \forall i \in F_O, n > 1, \tag{2.78}$$

$$r_i^n - a_i^n + a_i^{n\ominus 1} - o_i^n \geq 0 \qquad \forall i \in F_C, n > 3, \tag{2.79}$$

$$r_i^n \leq \sum_{m \in \mathcal{P}(n\ominus 1)} c_i^m \qquad \forall i \in F_O, n > 1, \tag{2.80}$$

$$r_i^n \leq \sum_{m \in \mathcal{P}(n\ominus 1)} c_i^m \qquad \forall i \in F_C, n > 3, \tag{2.81}$$

$$r_i^n \leq 1 - a_i^{n\ominus 1} \qquad \forall i \in F_O, n > 1, \tag{2.82}$$

$$r_i^n \leq 1 - a_i^{n\ominus 1} \qquad \forall i \in F_C, n > 3, \tag{2.83}$$

$$o_i^n, c_i^n, r_i^n \in \{0,1\} \qquad \forall i \in F, n, \tag{2.84}$$

$$u_i^{qn} = \sum_{m \in \mathcal{P}(n)} \hat{u}_i^{qm} + \sum_{m \in \mathcal{P}(n)} \bar{u}_i^{qm} \qquad \forall i \in F, q \in Q, n, \tag{2.85}$$

$$\hat{u}_i^{qn} = \sum_{k \in \hat{Y}} \hat{v}_{ik}^{qn} \hat{y}_{ik}^{qn} \qquad \forall i \in F, q \in Q, n, \tag{2.86}$$

$$\bar{u}_i^{qn} = \sum_{k \in \bar{Y}} \bar{v}_{ik}^{qn} \bar{y}_{ik}^{qn} \qquad \forall i \in F_C, q \in Q, n, \tag{2.87}$$

$$\sum_{k \in \hat{Y}} \hat{y}_{ik}^{qn} \leq \sum_{m \in \mathcal{P}(n\ominus 1)} o_i^m \qquad \forall i \in F, q \in Q, n, \tag{2.88}$$

$$\sum_{k \in \bar{Y}} \bar{y}_{ik}^{qn} \leq o_i^n \qquad \forall i \in F_C, q \in Q, n, \tag{2.89}$$

$$\sum_{n \in \mathcal{P}(s)} \hat{u}_i^{qn} \leq \phi_i \sum_{n \in \mathcal{P}(s)} \bar{u}_i^{qn} \qquad \forall i \in F, q \in Q, n, \tag{2.90}$$

$$\hat{u}_i^{qn}, \bar{u}_i^{qn} \geq 0 \qquad \forall i \in F, (i,j) \in L, q \in Q, n, \tag{2.91}$$

$$u_i^{qn}, u_{ij}^{n} \geq 0 \qquad \forall i \in F, (i,j) \in L, q \in Q, n, \tag{2.92}$$

$$\hat{y}_{ik}^{qn}, \bar{y}_{ik}^{qn} \in \{0,1\} \qquad \forall i \in F_C, (i,j) \in L_C, q \in Q, n. \tag{2.93}$$

Analogous constraints are defined to the logistics channels but omitted for the sake of clarity.

2.3.2 A simple example

In order to illustrate the functionalities of our Stochastic Capacity Planning mathematical model, we present a simple example with one original facility (1), one candidate facility (2) and a single customer (d) demanding a single product. In Figure 2.3 below, Facility 1 (the shaded one) is already opened at the beginning of the planning horizon, while Facility 2 (the white, dotted one) is a candidate facility. Facility 1 and Channel

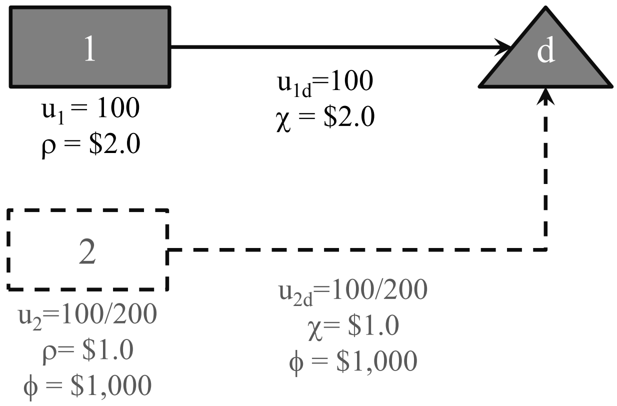

Figure 2.3. A simple example with two facilities: original and candidate.

1-d both have initial capacities of 100 units. Increments on capacity can be of 20 or 40 units, costing \$10 and \$15, respectively, but limited to a total 50% increase on initial capacity for facilities and logistics channels alike. This means that if the demand increases up to 50% of the original capacity of Facility 1, there would be no need to open the candidate Facility 2, nor the candidate Channel 2-d. On the other hand, if more than a 50% increase on Facilty 1's original capacity is needed, then Facility 2 and Channel 2-d would be required. Variable costs are \$2.0 for Facility 1 and Channel 1-d, and \$1.0 for Facility 2 and Channel 2. All fixed costs are \$1,000; opening, closing and reopening costs are \$1,000, \$10 and \$50, respectively.

The behavior of the expected demand during a three-period planning horizon is represented by the scenario tree shown in Figure 2.4. Each node in the set

$\mathcal{T} = \{1, 2, 3, 4, 5, 6, 7\}$ represents one possible demand configuration and its respective probability of realization. From the root node, there are two possible outcomes, represented by nodes 2 and 3, which can be realized according to probabilities 0.7 and 0.3, respectively. From node 2, two other transitions with probabilities 0.2 and 0.8 are possible. Each transition yields a new demand configuration with joint probability from the root: node 4, with $0.7 \times 0.2 = 0.14$ and node 5, with $0.7 \times 0.8 = 0.56$ probabilities. The same rationale applies to node 3 and its children: node 6 with probability $0.3 \times 0.3 = 0.09$ and node 7 with probability $0.3 \times 0.7 = 0.21$. The joint probabilities are shown in parentheses, below the demand levels on each node in Figure 2.4. The

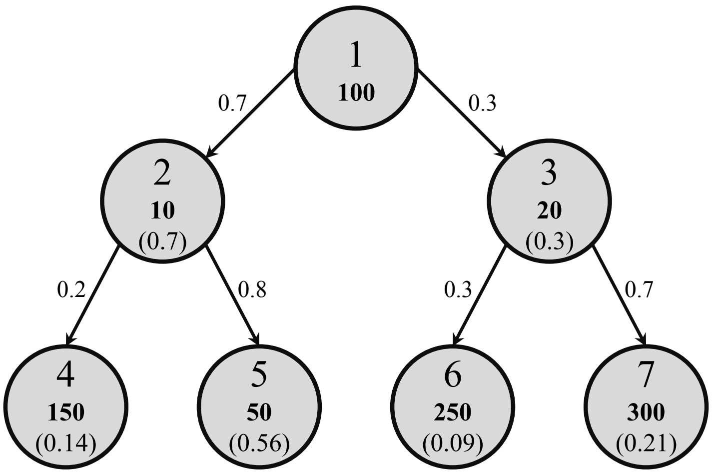

Figure 2.4. A scenario tree representing the risky demand on three periods and four scenarios, hence seven nodes. Demand levels are positioned below the node labels, with their respective probability in parentheses.

high variances in demand are such that different investment policies may be devised for each scenario, resulting in adequate demand/capacity ratios, or in idle capacity, or even in exceeding demand, which would imply in contractual penalties or third-party ore purchasing at unfavorable prices. Those undesirable costs are represented by the β parameter, which is initially set to a relatively high value in order to induce the network design dynamics.

In a stochastic program, decisions must be hedged against all scenarios in order to build a robust solution, as given in Table 2.1. At $t = 1$, since demand is higher than the available capacity in Facility 1 (including the 50% increment limit), Facility 2 must be activated. However, at $t = 2$ two very different demand levels may take place, which derive rather different policies for scenarios 1-2-**4** and 1-2-**5** and scenarios 1-3-**6** and 1-3-**7**. In the former set, when demand drops, as in node 2, the more expensive Facility 1 must be closed and shall remain like that until it is needed again at node 4.

In the latter, demand increases consistently and the corresponding capacity increments are required at nodes 6 and 7.

Table 2.1. Stochastic solution for the example problem, $\beta = 3000$.

n=1, d=160, b=0	a	o	c	r	$p; x$	$\bar{u}$	$\hat{u}$
1	1	—	0	0	60	—	0
2	1	1	0	0	100	100	0
1-d	1	—	0	0	60	—	0
2-d	1	1	0	0	100	100	0
n=2, d=10, b=0	a	o	c	r	$p; x$	$\bar{u}$	$\hat{u}$
1	0	—	1	0	0	—	0
2	1	0	0	0	10	0	0
1-d	0	—	1	0	0	—	0
2-d	1	0	0	0	10	0	0
n=3, d=200, b=0	a	o	c	r	$p; x$	$\bar{u}$	$\hat{u}$
1	1	—	0	0	100	—	0
2	1	0	0	0	100	0	0
1-d	1	—	0	0	100	—	0
2-d	1	0	0	0	100	0	0
n=4, d=150, b=0	a	o	c	r	$p; x$	$\bar{u}$	$\hat{u}$
1	1	—	0	1	50	—	0
2	1	0	0	0	100	0	0
1-d	1	—	0	1	50	—	0
2-d	1	0	0	0	100	0	0
n=5, d=50, b=0	a	o	c	r	$p; x$	$\bar{u}$	$\hat{u}$
1	0	—	0	0	0	—	0
2	1	0	0	0	50	0	0
1-d	0	—	0	0	0	—	0
2-d	1	0	0	0	50	0	0
n=6, d=250, b=10	a	o	c	r	$p; x$	$\bar{u}$	$\hat{u}$
1	1	—	0	0	140	—	40
2	1	0	0	0	100	0	0
1-d	1	—	0	0	140	—	40
2-d	1	0	0	0	100	0	0
n=7, d=300, b=60	a	o	c	r	$p; x$	$\bar{u}$	$\hat{u}$
1	1	—	0	0	140	—	40
2	1	0	0	0	100	0	0
1-d	1	—	0	0	140	—	40
2-d	1	0	0	0	100	0	0

It can be seen that it is impossible, under risk, to devise a policy that is ideal under all circumstances. Again, the stochastic solution must be well-hedged, building some flexibility to meet the risky demand at all scenarios. It is also important to emphasize that the β parameter has a strong influence on how much the model allows demand shortages. As β increases, however, the model becomes more intolerant to shortage penalties and hence investments on capacity become more economically attractive. If

$\beta = 5000$, for instance, all required capacity investments are made and no shortages would be allowed.

2.4 Model analysis

Having presented the mathematical formulation for the problem and having illustrated its functionality with a simple example, we evaluate the computational issues associated to solving a more realistic test instance (Figure 2.5) which is inspired in Brazilian large-scale iron ore operations. Our Global Mining Supply Chain has three systems, each comprised of pits, processing plants, railroads, ports and supply stations with some degree of interoperability. We are interested in evaluating how the quality of

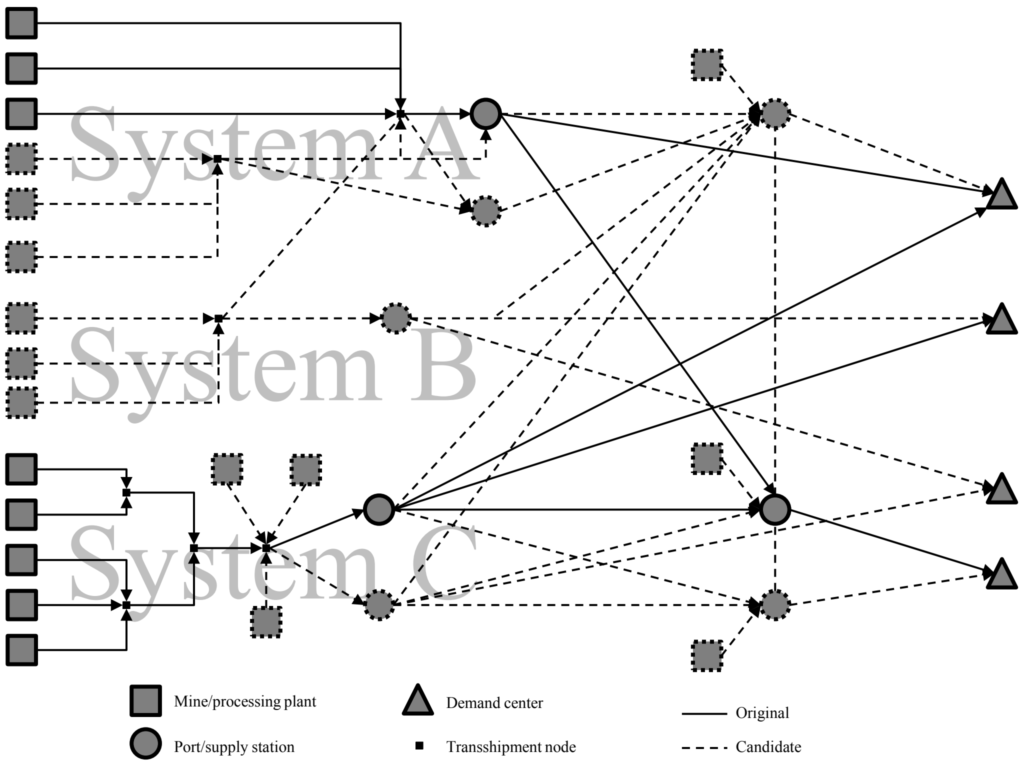

Figure 2.5. The Global Mining Supply Chain used in the tests.

the solution is influenced by the demand/capacity ratio and the size of the scenario tree. Also, we compare different solution approaches — from standard Branch&Cut to evolutionary heuristics. All runs use CPLEX 12.2 in an Intel Xeon E5520 workstation with four cores (up to eight threads), 2.27GHz and 16GB RAM.

2.4.1 Demand/capacity ratio

Since most of the integer variables are dedicated to incrementing or establishing new capacity at specific locales of the Global Mining Supply Chain, it is fair to expect that solving the model should become harder as the need to increment capacity or open new facilities increases. This assumption is evaluated for the test instance with the binary scenario tree of Figure 2.4, a time limit of one hour and a large enough penalty for demand shortages, so that the model would always try to use all available capacity expansion options. The demand/capacity ratio considers not only the initially established capacity, but also the additional capacity available from incrementing or establishing new facilities and channels. Data from Table 2.2 indicates that as the demand/capacity ratio increases up to 100%, so does the time needed to achieve optimality. As demand exceeds the overall capacity, however, the shortage variables b_i^n assume all unmet demand and solution time drops back to values of lower demand/capacity ratios. Different

Table 2.2. Influence of the demand/capacity ratio on solution quality.

Ratio	Time (s)	Gap
5%	47	0.00%
30%	187	0.00%
45%	1,540	0.00%
90%	3,600	0.01%
120%	66	0.00%
150%	16	0.00%

demand configurations and variance, even for the same demand/capacity ratios, can of course have a different effect on solution quality. In Section 2.5.3 we further explore those issues.

2.4.2 Size of the scenario tree

It is clear that the size of the problem increases exponentially with the size of the scenario tree. Here, we extend the binary tree shown in Figure 2.4 32 scenarios (6 time periods in the planning horizon). Although this is not a typically large number — one could be interested in planning for 10 or more time periods, resulting in 512 or more scenarios —, it is sufficient for the analytical purposes of this Section. Demand and probabilities are varied in such a way as to exercise all functionalities of the model[1].

It can be seen from Table 2.3 that the larger the scenario tree, the higher the solution gap. When CPLEX is used with its default settings and a one-hour time limit,

[1]All data and source codes in OPL format can be provided by the author.

Table 2.3. Influence of the scenario tree size on solution quality.

Periods	Scenarios	Time (s)	Gap	Binaries
2	2	87	0.00%	2,940
3	4	3,600	0.01%	7,776
4	8	3,600	0.77%	17,532
5	16	3,600	3.25%	36,948
6	32	3,600	26.16%	76,164

a 26.16% solution gap is achieved for the test instance. After 30 hours of processing, the resulting gap would still be around 2%. This suggests that specific solution approaches should be evaluated, especially for large-scale instances. That is the subject of next section.

2.4.3 Solution approaches

The Stochastic Capacity Planning Problem has a structure that shares characteristics of facility location, network design, capacity expansion and lot sizing. Such a structure may suggest specific solution procedures. A first step on this investigation is to evaluate the influence of specific CPLEX settings on the quality of the solution [Pimentel et al., 2011]. All runs are performed over the test instance in a 32 scenario planning horizon, again with a time limit of one hour. Table 2.4 shows the results for different CPLEX settings, each applied individually and, in bold, the resultant best effort for a selection of parameters that achieved the most interesting results. With the default settings, CPLEX usually decides whether or not to apply several additional algorithms during the solution process, often increasing the intensity of each algorithm if that seems to improve the solution. It is important to emphasize that each test over the default settings varies only the parameter indicated.

A few insights may be derived here. First, we notice that increasing the pre-processing efforts with symmetry breaking [Sherali and Smith, 2001] or probing — checking for logical implications of pre-fixing binary variables before processing the root node — provide no significant improvement and, in some cases, may even yield very poor results.

Regarding the application of non-default settings for the CPLEX's built-in search heuristics, results are also not significantly improved. Although local branching [Fischetti and Lodi, 2003] would be expected to produce quality solutions at early stages of the computation, most of the one-hour time limit is consumed by that procedure with no significant improvement on the first incumbent. The node heuristic, on the other hand, tries to construct a feasible solution from the current (fractional) Branch&Cut

node. Increasing the frequency of the node heuristic also consumes a significant amount of the time available for the search without improving the objective. Eliminating the node heuristic completely, however, ultimately prevents a feasible solution to be found. A similar behavior is achieved with the Relaxation Induced Neighborhood Search (RINS) [Danna et al., 2005]. RINS explores a neighborhood of the current incumbent by formulating the search as a mixed integer subproblem and truncating the number of nodes explored in the search tree. Although not a particularly expensive procedure, applying the RINS heuristic at each new incumbent does not significantly improve the solution. On the other hand, not applying it altogether yields very poor results.

Table 2.4. Influence of CPLEX Parameters on Solution Quality

	Configuration	**CPLEX Gap**	**CPLEX UB**
	Defaults	*26.16%*	*30,332,805*
	MIP Probing Level		
	None	feasible, but too high a gap	
	Higher than default	feasible, but too high a gap	
	Symmetry Breaking		
	None	42.44%	38,911,562
	Higher than default	no improvement	
	RINS Frequency		
	None	feasible, but too high a gap	
	Higher than default	no improvement	
	Local Branching (on)	feasible, but too high a gap	
	Node Heuristic Frequency		
	None	no feasible solution	
	Higher than default	no improvement	
	Highest (every node)	feasible, but too high a gap	
a	MIP Priority (Decreasing Cost)	23.55%	29,295,957
b	MIP Emphasis (Feasibility)	22.77%	29,051,366
c	Cutting Plane Algorithms (None)	15.12%	26,230,505
d	Solution Polishing (5/6)	12.76%	25,674,320
	Combined Settings I (a, b, c, d)	2.53%	22,866,143
	Combined Settings II (c, d)	7.54%	24,098,179
	Combined Settings III (a, c)	**2.30%**	**22,871,765**
	Combined Settings IV (b, c)	2.67%	22,871,068
	Combined Settings V (a, b, c)	2.54%	22,864,688

Better performance can be obtained by tackling CPLEX's MIP strategy parameters. When the *MIP Emphasis* parameter is set to *feasibility*, CPLEX attempts to generate more feasible solutions early in the solution process, at some sacrifice in speed of proving optimality. However, better solution gaps are achieved in the process. Other settings of *MIP Emphasis* do not further improve the solution. On another approach, setting the priority to *decreasing cost coefficients* when deciding which variable to branch on at a node, shows some improvement when compared to the no-priority default settings.

The default settings include the application of standard cutting-plane algorithms in an attempt to improve the efficiency of the solution procedure. Results show that when cutting planes are not applied, solution quality can be significantly improved. A closer look into the solution logs shows that a large number of cuts is generated at the initial node without improving the incumbent. The default Branch&Cut procedure then reaches the one-hour time limit having explored a few 500 nodes. This suggests a more detailed investigation of CPLEX's cutting plane algorithms. Without

Table 2.5. Analysis of cutting plane algorithms applied to the stochastic capacity planning problem.

Algorithm	# Applied	Solutions	Nodes	Objective	CPLEX Gap
MCF	51	117	6,163	24,476,179	8.97%
Clique	–	161	9,596	26,230,505	15.12%
GUB Covers	–	158	9,506	26,230,630	15.12%
Zero-Half	–	158	9,506	26,230,630	15.12%
Implied	3,231	31	9,430	26,341,425	15.40%
Gomory	459	52	8,039	26,603,175	15.92%
Disjunctive	–	65	10,156	26,863,809	17.12%
Covers	34	162	7,897	27,022,724	17.59%
Path Cuts	210	147	6,629	28,617,590	22.17%
Flow Covers	2,467	52	746	30,871,086	27.45%
MIR	2,630	41	388	feasible, but too high gap	

cutting planes, the simpler Branch&Bound algorithm is able to explore more than 17,000 nodes within the time limit, thus yielding a much better performance. It is important to notice that CPLEX ultimately addresses the problem as a generic integer program. Hence, besides trying to develop problem-specific valid inequalities[Melkote and Daskin, 2001], it would be interesting to apply techniques specially developed for multi-stage stochastic problems, such as the combination of inequalities that are valid among individual scenarios [Guan et al., 2009].

Metaheuristics can also be used in an attempt to improve solution quality. *Solution polishing* is a general purpose, coarse-grained evolutionary approach built within a large-neighborhood search framework [Rothberg, 2007]. We devote 50 minutes of the one-hour time limit to the traditional Branch&Cut procedure and the last 10 minutes to solution polishing. One can observe a significant improvement on solution quality within a short computation time. This indicates that further investigating specific metaheuristic approaches to such a problem structure may also be an interesting research direction. Alternatively, local search methods [Bai et al., 2010] could be used as an attempt to improve feasible solutions. Since local search focuses on the combinatorial aspects of the problem and polishing works with a population of solutions. The first could provide better results when facing tighter primal bounds.

One obvious approach, however, is to try combining the settings which produced the best results. We provide five of those combinations — all with no cutting plane algorithms — also in shown Table 2.4. Combination I applies Branch&Bound, decreasing cost coefficients, feasibility emphasis and one-sixth time-limit polishing. One can observe a significant improvement in the solution quality, with a 2.53% gap. However, an examination of the solution logs shows a very small contribution of polishing to the result. We then put together Branch&Bound and one-sixth time-limit polishing in Combination II in order to show that polishing does work better with worse primal bounds. This indicates that exact methods should be carefully considered when developing specific solution procedures. Combinations III and IV use plain Branch&Bound with decreasing cost priority for variables and feasibility emphasis, respectively. Results are relatively similar but slightly favour the usage of the decreasing cost priority. The benefits of providing an adequate ordering for the variables to be branched greatly compensate the computational cost involved in determining that ordering. It is also clear that in such a complex problem, it is better to concentrate efforts in determining quality feasible solutions early in the process, as driven by the feasibility MIP emphasis. Finally, in Combination V, decreasing cost priority, feasibility MIP emphasis and plain Branch&Bound are used together but with no significant improvement compared to the previous combinations. The additional computational cost does not compensate the benefits over the solution gap achieved by Combination III.

2.5 A Lagrangian Relaxation approach

We propose a Lagrangian Relaxation [Fisher, 1981] approach to strengthen the dual bounds — and possibly derive good approximate primal bounds — for the stochastic capacity planning problem. The basic idea is to relax and dualize a subset of constraints in order to derive a Lagrangian problem which should be easier to solve and whose optimal value is a dual bound on the optimal value of the original problem. Lagrangian Relaxation has been invoked in decomposition approaches for general multistage stochastic problems [Parpas and Rustem, 2007] as well as network design problems [Tanonkou et al., 2008]. The main advantages of this method are its simplicity and the ability to, when coupled with a primal heuristic, provide a feasible solution and a measure of the optimality gap for the original problem.

In this work, we explore the structure of the proposed stochastic capacity planning problem and dualize constraints which provide a good tradeoff between the quality of the Lagrangian dual and the complexity of the relaxed problem. The proposed Lagrangian Heuristic then attempts to build a primal feasible solution to the Lagrangian subproblem. A number of experiments are made in order to assess the heuristic's behavior against CPLEX.

2.5.1 Constraint dualization

Given a general optimization problem:

$$\begin{aligned} z = \min cx& \\ Ax &\geq b \\ Dx &\geq e \\ x &\in \mathbb{Z}_+^n, \end{aligned}$$

the basic idea of the Lagrangian Relaxation is to relax and dualize specific sets of constraints in order to derive a relaxed problem of the form:

$$\begin{aligned} z(\lambda) = \min cx + \lambda\,(e - Dx)& \\ Ax &\geq b \\ x &\in \mathbb{Z}_+^n \\ \lambda &\geq 0, \end{aligned}$$

where λ is the Lagrangian multiplier associated to constraints $Dx \geq e$. Since $z(\lambda)$ provides a lower bound to the original optimization problem, we are interested in determining the best lower bound over the possible values of λ — which is achieved by solving the Lagrangian dual problem [Wolsey, 1998]. One common approach to do so is the subgradient algorithm, to be detailed later on in this Chapter.

The mathematical formulation presented in Section 2.3.1 indicates that some specific restrictions are responsible for coupling the lot-sizing, network-design and capacity expansion subproblems, being therefore good candidates for dualization. Nevertheless, we perform a thorough preliminary analysis of several possible relaxations in order to determine the one which provides the best tradeoff between the quality of the dual bounds and the corresponding computational complexity required to obtain those bounds. The procedure involves solving different relaxed problems, each defined by the simple elimination of a given subset of constraints (which is equivalent to setting the corresponding Lagrangian multipliers to zero) and comparing the dual bounds obtained. According to this procedure, the best tradeoff is given by the relaxation of the following constraints:

$$\sum_q p_i^{qn} \leq u_i^{q*} a_i^n \qquad \forall i \in F, n, \tag{2.94}$$

$$\sum_q x_{ij}^{qn} \leq u_{ij}^* a_{ij}^n \qquad \forall (i,j) \in L, n, \tag{2.95}$$

$$\hat{u}_i^{qn} = \sum_{k \in \hat{Y}} \hat{\upsilon}_{ik}^{qn} \hat{y}_{ik}^{qn} \qquad \forall i \in F, q \in Q, n, \tag{2.96}$$

$$\hat{u}_{ij}^n = \sum_{k \in \hat{Y}} \hat{\upsilon}_{ijk}^n \hat{y}_{ijk}^n \qquad \forall (i,j) \in L, n, \tag{2.97}$$

$$\bar{u}_i^{qn} = \sum_{k \in \bar{Y}} \bar{\upsilon}_{ik}^{qn} \bar{y}_{ik}^{qn} \qquad \forall i \in F_C, q \in Q, n, \tag{2.98}$$

$$\bar{u}_{ij}^n = \sum_{k \in \bar{Y}} \bar{\upsilon}_{ijk}^n \bar{y}_{ijk}^n \qquad \forall i \in L_C, n. \tag{2.99}$$

The effect of relaxing the above constraints on the Lagrangian subproblem is that the activation variables and the dynamic network design variables become disconnected and thus may be set to incoherent values. Also, the capacity variables become unbounded and reflect the value of the corresponding production and transportation variables. That requires the following inequations to be added to the relaxed problem:

$$u_i^{qn} \leq (1 + \phi_i)\, \bar{\kappa}_i^q \qquad \forall i \in F_O, q \in Q, n, \tag{2.100}$$

$$u_{ij}^n \leq (1 + \phi_{ij})\, \bar{\kappa}_{ij} \qquad \forall (i,j) \in L_O, n, \tag{2.101}$$

$$u_i^{qn} \leq (1 + \phi_i) \max_k \bar{\upsilon}_{ik}^{qn} \qquad \forall i \in F_C, q \in Q, k \in \bar{Y}, n, \tag{2.102}$$

$$u_{ij}^n \leq (1 + \phi_i) \max_k \bar{\upsilon}_{ijk}^n \qquad \forall (i,j) \in L_C, k \in \bar{Y}, n. \tag{2.103}$$

Those inequations are redundant in the original formulation, but provide a theoretical upper bound on the capacities of the relaxed problem.

2.5.2 A Lagrangian Heuristic

The proposed Lagrangian Relaxation approach is applied as part of a heuristic procedure that iteratively approximates the Lagrangian dual solution z_{LD}^* using the subgradient method and, whenever that dual bound is improved, a greedy algorithm is applied to determine a feasible solution z^* for the problem. Both bounds are then used to update the Lagrangian multipliers according to a step size given by:

$$\mu^k = \epsilon^k \frac{(z^* - z^*_{LD})}{\|e - Dx\|^2},$$

where the parameter ϵ^k is a scalar satisfying $0 \leq \epsilon^k \leq 2$. The subgradient algorithm is terminated if one of the following conditions is achieved:

1. If $z^* = z^*_{LD}$ (or $z^* - z^*_{LD} < \xi$): an optimal solution has been determined by the Lagrangian Heuristic.

2. If the subgradient tends to zero (or $\|e - Dx\| < \xi$), the step size is undefined and the procedure stops with the currently best Lagrangian dual and primal feasible solutions.

3. As the parameter ϵ^k tends to zero, so does the step size μ^k. Below a certain threshold $\epsilon^k < \xi$, the algorithm is terminated since no further improvement on the Lagrangian dual should be expected. A limit on the number of iterations may yield that same effect.

Algorithm 1 below formalizes the proposed approach.

Roughly speaking, the greedy feasibility algorithm used to determine z^* can be described as follows:

1. Search for facilities and channels in which production and transportation occur and determine the corresponding activity variables a_i^n and a_{ij}^n.

2. According to the variation on the activity variables, determine the opening (o_i^n and o_{ij}^n), closure (c_i^n and c_{ij}^n) and reopening (r_i^n and r_{ij}^n) variables for every facility and channel, again at every node of the scenario tree.

3. Whenever new facilities or channels are opened, determine their initial capacities according to the largest activity level observed among all applicable scenarios; take into account the limit on the maximum initial capacity and the possible increments — $(1 + \phi_i)$ and $(1 + \phi_{ij})$ — that can be performed in subsequent nodes.

4. Whenever the activity level exceeds the capacity available on facilities and channels, determine the minimum necessary capacity increment in order to match that activity level, again observing the maximum capacity limits.

Algorithm 1: Lagrangian Heuristic

Input: Problem $z = \min\{\imath(n)c^n x^n : A^n x^n \geq b^n, D^n x^n \geq e^n\}$, the original stochastic capacity planning problem.

1: Let λ be the vector of Lagrangian multipliers.
2: Let $z_{LD} = \min\{p(n)\left[c^n x^n + \lambda\left(e^n - D^n x^n\right)\right] : A^n x^n \geq b^n\}$ be the relaxed problem
3: obtained by dualizing constraints $Dx \geq e$ represented by equations (2.94)-(2.99).
4: Let $z^*_{LD} \leftarrow 0$.
5: Let $z^* \leftarrow \infty$.
6: Let $k \leftarrow 1$.
7: Solve $z_{LD}(\lambda^k)$.
8: **if** $z_{LD} > z^*_{LD}$ **then**
9: $\quad z^*_{LD} \leftarrow z_{LD}$.
10: $\quad$ Execute greedy feasibility algorithm and determine z^*.
11: **end if**
12: **if** z^*_{LD} fails to increase after 30 iterations **then**
13: $\quad \epsilon^{k+1} = \epsilon^k/2$.
14: **end if**
15: Calculate the step size $\mu^k = \epsilon^k \frac{(z^* - z^*_{LD})}{\|e - Dx\|}$.
16: Update Lagrangian multipliers as $\lambda^{k+1} = \lambda^k + \mu^k\left(e - Dx(\lambda)\right)$.
17: **if** $(z^* - z^*_{LD}) < \xi$ **then**
18: $\quad z^*$ optimal, go to step 25 and terminate.
19: **else if** $\|e - Dx\| < \xi$ **then**
20: $\quad$ subgradient is null, go to step 25 and terminate.
21: **else if** $\epsilon^k < \xi$ **then**
22: $\quad$ go to step 25 and terminate.
23: **end if**
24: Increment k and go back to step 7.
25: Calculate optimality gap as $\left(1 - \frac{z^*_{LD}}{z^*}\right)$.

Output: z^*_{LD}, z^* and optimality gap.

5. Finally, having established the network infrastructure, recalculate all production, transportation and demand shortage variables (now as a linear problem) and determine a feasible solution z^*.

Algorithm 2 formalizes the greedy feasibility heuristic.

2.5.3 Numerical evaluation

In this section we evaluate how the proposed Lagrangian Heuristic behaves over a set of different problem instances. The implementation is based on IBM ILOG CPLEX Optimization Studio 12.2 and all tests use the same workstations described in Section 2.4. For these tests, however, we propose a different set of problem instances than the ones used in Section 2.4. Ternary scenario trees are more consistent to the expert evaluation commonly performed in scenario planning, since a decision maker may be interested in three possible outcomes of any course of action: a realistic one, a pessimist

Algorithm 2: Greedy Feasibility Algorithm

Input: Solution to z_{LD} at iteration k.

Let $a_i^n, a_{ij}^n, o_i^n, o_{ij}^n, c_i^n, c_{ij}^n, r_i^n, r_{ij}^n, \bar{u}_{ik}^{qn}, \bar{u}_{ijk}^n, \hat{u}_{ik}^{qn}, \hat{u}_{ijk}^n$ initially set to zero.
Given p_i^{qn} and x_{ij}^{qn} from z_{LD}, determine new values for the above variables as follows:
Determine activity variables:
$\forall i \in F, q \in Q, n \in N$:
    $p_i^{qn} > 0 \longrightarrow a_i^n = 1;$
$\forall (i,j) \in L, q \in Q, n \in N$
    $x_{ij}^{qn} > 0 \longrightarrow a_{ij}^n = 1, a_i^n = 1, a_j^n = 1;$
Determine opening variables:
$\forall i \in F_C, n \in N$:
    $a_i^n = 1$ **and** $\sum_{m \in P(n)} o_i^m = 0 \longrightarrow o_i^n = 1;$
$\forall (i,j) \in L_C, n \in N$
    $a_{ij}^n = 1$ **and** $\sum_{m \in P(n)} o_{ij}^m = 0 \longrightarrow o_{ij}^n = 1;$
Determine closure variables:
$\forall i \in F, n \in N$:
    $a_i^n < a_i^{n\ominus 1} \longrightarrow c_i^n = 1;$
$\forall (i,j) \in L, n \in N$
    $a_{ij}^n < a_{ij}^{n\ominus 1} \longrightarrow c_{ij}^n = 1;$
Determine reopening variables:
$\forall i \in F, n \in N$:
    $a_i^n > a_i^{n\ominus 1}$ **and** $\sum_{m \in P(n)} c_i^{\mathcal{A}(m)} \geq 1 \longrightarrow r_i^n = 1;$
$\forall (i,j) \in L, n \in N$:
    $a_{ij}^n > a_{ij}^{n\ominus 1}$ **and** $\sum_{m \in P(n)} c_{ij}^{\mathcal{A}(m)} \geq 1 \longrightarrow r_{ij}^n = 1;$
Determine initial capacity:
$\forall i \in F_C, n \in N, k \in \bar{Y}$:
    $o_i^n = 1 \longrightarrow \bar{y}_{ik^*}^{qn} = 1$ and $\bar{u}_i^{qn} = \bar{v}_{ik^*}^{qn}$, where $k^* = \arg_k \min \left\{ \left(\bar{v}_{ik}^{qn} - \frac{p^*}{1+\phi_i} \right) \geq 0 \right\}$ and $p^* = \max \{ p_i^{qm} : m \in (s \in \mathcal{S} : s \supset n) \}$,
$\forall (i,j) \in L_C, n \in N, k \in \bar{Y}$:
    $o_{ij}^n = 1 \longrightarrow \bar{y}_{ijk^*}^n = 1$ and $\bar{u}_{ij}^n = \bar{v}_{ijk^*}^n$, where $k^* = \arg_k \min \left\{ \left(\bar{v}_{ijk}^n - \frac{x^*}{1+\phi_{ij}} \right) \geq 0 \right\}$ and $x^* = \max \left\{ \sum_q x_{ij}^{qm} : m \in (s \in \mathcal{S} : s \supset n) \right\};$
Determine capacity increment in original facilities and channels:
$\forall i \in F_O, n \in N, k \in \hat{Y}, u_i^{qn} = \left(\bar{\kappa}_i^q + \sum_{m \in \mathcal{P}(n)} \hat{u}_i^{qm} \right), \delta = (p_i^{qn} - u_i^{qn})$:
    $\delta > 0 \longrightarrow \hat{y}_{ik^*}^{qn} = 1$, where $k^* = \arg_k \min \{ 0 \leq (\hat{v}_{ik}^{qn} - \delta) \leq u_i^{q*} \};$
$\forall (i,j) \in L_O, n \in N, k \in \hat{Y}, u_{ij}^{qn} = \left(\bar{\kappa}_{ij} + \sum_{m \in \mathcal{P}(n)} \hat{u}_{ij}^m \right), \delta = \left(\sum_q x_{ij}^{qn} - u_{ij}^n \right)$:
    $\delta > 0 \longrightarrow \hat{y}_{ijk^*}^n = 1$, where $k^* = \arg_k \min \left\{ 0 \leq \left(v_{ijk}^n - \delta \right) \leq u_{ij}^* \right\};$
Determine capacity increment in new facilities and channels:
$\forall i \in F_C, n \in N, k \in \hat{Y}, u_i^{qn} = \left(\sum_{m \in \mathcal{P}(n\ominus 1)} \bar{u}_i^{qm} + \sum_{min \mathcal{P}(n)} \hat{u}_i^{qn} \right), \delta = (p_i^{qn} - u_i^{qn})$:
    $\delta > 0 \longrightarrow \hat{y}_{ik^*}^{qn} = 1$, where $k^* = \arg_k \min \{ 0 \leq (\hat{v}_{ik}^{qn} - \delta) \leq u_i^{q*} \};$
$\forall i \in L, n \in N, k \in \hat{Y}, u_{ij}^n = \left(\sum_{m \in \mathcal{P}(n\ominus 1)} \bar{u}_{ij}^m + \sum_{min \mathcal{P}(n)} \hat{u}_{ij}^n \right), \delta = \left(\sum_q x_{ij}^{qn} - u_{ij}^n \right)$
    $\delta > 0 \longrightarrow \hat{y}_{ijk^*}^n = 1$, where $k^* = \arg_k \min \left\{ 0 \leq \left(v_{ijk}^n - \delta \right) \leq u_{ij}^* \right\};$
Fix above variables and re-solve to determine corrected values for p_i^{qn}, x_{ij}^{qn} and b_i^{qn};

Output: z^*, a feasible solution to z.

and an optimistic one [Schoemaker, 1995], each with specific realization probabilities. The ternary scenario trees are built as follows:

1. We assume every node in the scenario tree to have three children with individual probabilities 0.1, 0.7 and 0.2.

2. For any node n other than the root node, we associate a probability which can be determined as the joint probability of every node within the path from the

root node, or $\prod_{m \in \mathcal{P}(n)} p(m)$.

3. Regarding the probabilistic behavior of the demand, we wish to be able to control the magnitude of its variance throughout the scenario tree. Given a parameter $\varphi = [0, 1]$ and starting from an initial, known demand value at the root node, we decrease by φ the demand of the child node with 0.1 probability, we keep constant the demand of the child node with 0.7 probability and we increase by φ the demand of the child node of 0.2 probability.

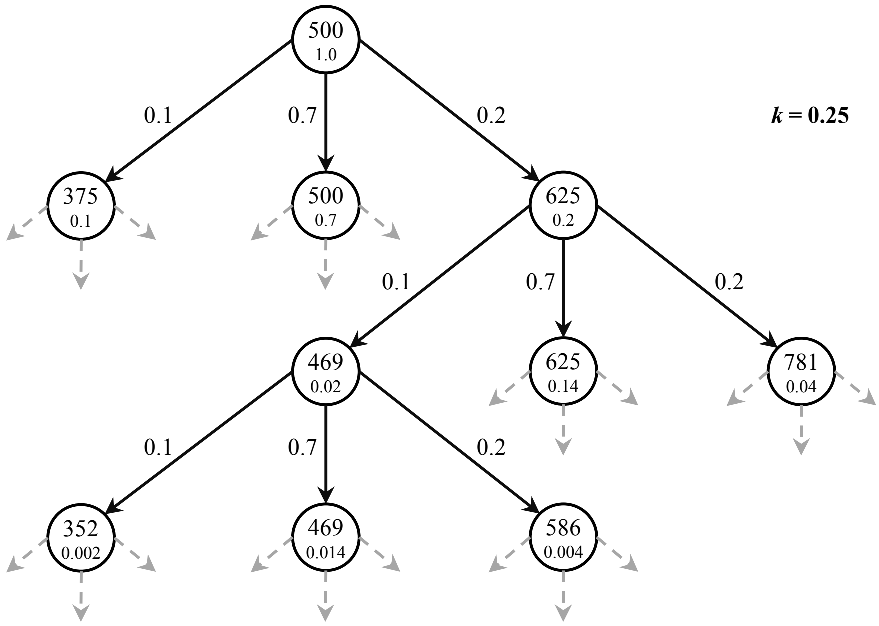

Figure 2.6. An example of building a ternary tree with probabilities 0.1, 0.7 and 0.2 and $\varphi = 0.25$.

Figure 2.6 illustrates the procedure described above. Demand is assumed to be low at the beginning of the planning horizon, tending to increase up to the overall system's capacity (including all possible expansions and extensions provided by the model) until the end of planning horizon[2].

The numerical experiments designed to evaluate the Lagrangian Heuristic use two of those ternary scenario trees with different values for φ: 0.25 and 0.50. The idea behind that choice is to provide us with some indication regarding the heuristic's robustness to increasing variance and magnitude of the demand figures within the scenario tree. We consider scenario trees of increasing size, from 3 to 7 periods in the planning horizon (respectively from 4 to 1,093 nodes and from 3 to 729 scenarios).

[2] All data and source codes in OPL format can be provided by the author.

Tables 2.6 e 2.7 below show the results of the Lagrangian Heuristic for the scenario trees constructed with $\varphi = 0.25$ and $\varphi = 0.50$ respectively. Each line presents the length of the planning horizon T, the corresponding number of scenarios $\mathcal{S}$ in the scenario tree, the linear relaxation z_{LR} for the problem instance, the Lagrangian dual z^*_{LD} and the primal bound z^* obtained by the Lagrangian Heuristic, the corresponding optimality gap and the computing time. All runs are limited to 1,000 iterations, except for

Table 2.6. Results of the Lagrangian Heuristic for scenario tree with $\varphi = 0.25$.

T	$\mathcal{S}$	z_{LR}	z^*_{LD}	z^*	**Gap**	**Iter.**	**Time (s)**
3	9	7,430,984	7,430,518	7,974,851	6.83%	1,000	3,111
4	27	9,944,929	9,942,471	10,714,650	7.21%	1,000	14,490
5	81	12,478,440	12,472,272	13,533,509	7.84%	1,000	45,446
6	243	15,030,290	15,017,426	16,381,346	8.33%	1,000	221,870
7	729	17,602,610	17,566,939	19,281,646	8.89%	300	735,769

Table 2.7. Results of the Lagrangian Heuristic for scenario tree with $\varphi = 0.50$.

T	$\mathcal{S}$	z_{LR}	z^*_{LD}	z^*	**Gap**	**Iter.**	**Time (s)**
3	9	7,968,985	7,503,145	8,224,179	8.77%	1,000	3,106
4	27	10,101,670	10,093,714	11,164,690	9.59%	1,000	14,644
5	81	12,753,100	12,735,990	14,223,416	10.46%	1,000	45,684
6	243	15,472,720	15,435,332	17,425,503	11.42%	1,000	229,808
7	729	18,273,850	18,142,638	20,779,377	12.68%	300	749,191

those with 729 scenarios, which were terminated after 300 iterations due to the much longer computing times[3]. The tables show that the proposed Lagrangian Heuristic can at most approximate the linear relaxation bounds for the problem, but not once in our tests does it provide better lower bounds than the linear relaxation. Also, by comparing the corresponding results with $\varphi = 0.25$ and $\varphi = 0.50$, one can notice that increased variance and magnitude of demand in the scenario trees tends to increase the problem complexity, as it can be inferred from the larger solution gaps of the latter.

Tables 2.8 and 2.9 show the corresponding results of the optimization using CPLEX. Here, we set the solver to run with approximately the same time limits allowed for the heuristic procedure. It can be seen that, for the smaller problem sizes, CPLEX easily outperforms the Lagrangian Heuristic. That can be explained, in part, by a closer examination of the solutions provided by the greedy feasibility algorithm, which show much higher costs due to demand shortage penalties — a consequence, for

[3]In fact, one could choose to run the algorithm for virtually any number of iterations. The limit is actually in the amount of time the workstation must be exclusively dedicated to solving the problem.

Table 2.8. Performance of CPLEX for scenario tree with $\varphi = 0.25$.

T	$\mathcal{S}$	CPLEX$_{LB}$	CPLEX$_{UB}$	**Gap**	**Time (s)**
3	9	7,883,339	7,937,825	0.69%	3,110
4	27	10,519,354	10,606,686	0.82%	14,489
5	81	13,063,345	13,367,450	2.27%	45,498
6	243	15,576,668	16,188,634	3.78%	222,895
7	729	18,182,600	20,415,600	6.78%	$663{,}456^{+}$

Table 2.9. Performance of CPLEX for scenario tree with $\varphi = 0.50$.

T	$\mathcal{S}$	CPLEX$_{LB}$	CPLEX$_{UB}$	**Gap**	**Time (s)**
3	9	7,980,966	8,042,528	0.77%	3,110
4	27	10,709,384	10,828,741	1.10%	14,643
5	81	13,398,237	13,772,149	2.71%	45,798
6	243	16,210,565	16,805,754	3.54%	229,802
7	729	18,348,900	21,670,100	15.33%	$27{,}642^{++}$

instance, of enforcing limited initial capacities on new facilities and channels that operated under much less strict restrictions in the Lagrangian dual solution — and higher long-term established capacity costs — since capacity decisions are always based on the immediate needs. Nevertheless, the exponential increase in both memory and time required by CPLEX to find good feasible solutions cannot be taken for granted — both problem instances with 729 scenarios proved extremely difficult to compute. For the $\varphi = 0.25$ tree, CPLEX stalled for a long time (at least three days) before printing out any evolution on the process. Then, the high levels of swapping observed in the workstation draw us to an early termination. For the $\varphi = 0.25$ tree, results are even more troublesome. The swapping caused the working station's operating system to stop responding altogether before reaching eight hours of processing.

A more detailed view of the approximation provided the Lagrangian Heuristic is illustrated in Figure 2.7. There we show the evolution of the lower and upper bounds obtained by the Lagrangian Heuristic for each problem instance of the scenario tree with $\varphi = 0.25$. Though not shown here, a similar behavior is perceived with $\varphi = 0.50$. It is interesting to notice that very early in the computation do the bounds display more significant improvement. That behavior can be explained by the steeper decrease on the evolution of the upper bounds rather than the smoother increase on evolution of the lower bounds. The graph also suggests that, though being outperformed by CPLEX on the long run, the Lagrangian Heuristic may yield more interesting results in the early stages of computation.

In order to further explore the behavior depicted above, we perform an assessment of the computational effort required by CPLEX in determining a solution when

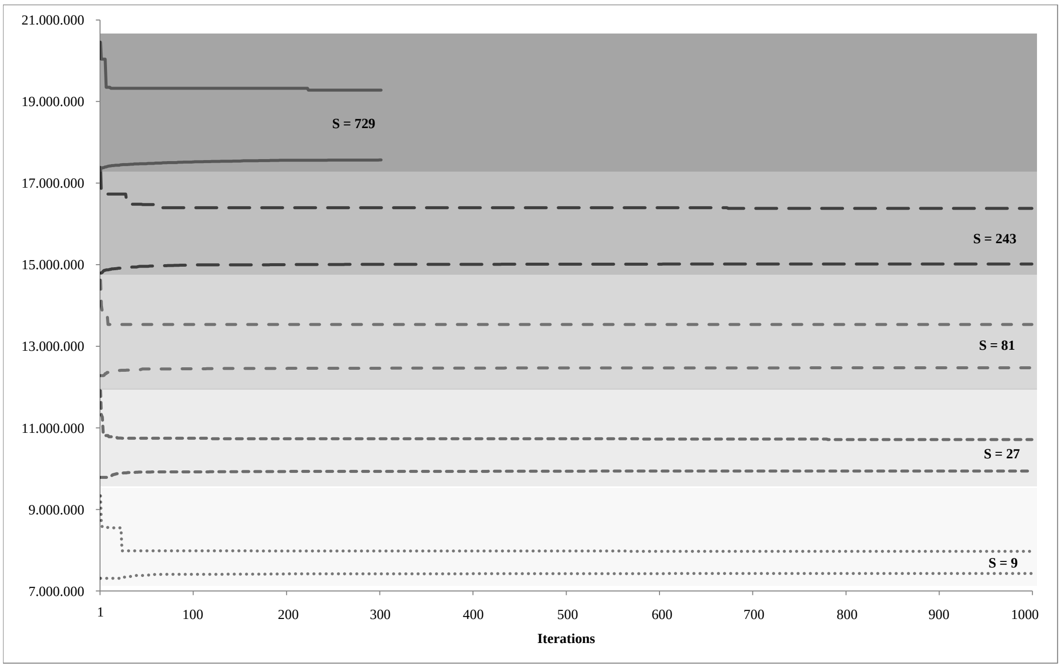

Figure 2.7. Evolution of the lower and upper bounds obtained by the Lagrangian Heuristic for the scenario tree with $\varphi = 0.25$. Each instance is identified by the number of scenarios and by a different shaded area on the graph.

limited to the time required for the Lagrangian Heuristic to reach a reasonable level of stabilization, say, after running 100 iterations. Tables 2.10 and 2.11 hence show that

Table 2.10. Results obtained with the Lagrangian Heuristic limited to 100 iterations and with CPLEX set to the corresponding time limit; tests performed on the scenario tree with $\varphi = 0.25$.

T	$\mathcal{S}$	z^*_{LD}	z^*	$\text{LH}_{\textbf{Gap}}$	CPLEX_{LB}	CPLEX_{UB}	$\text{CPLEX}_{\textbf{Gap}}$	**Time (s)**
3	9	7,411,269	7,984,816	7.18%	7,854,822	7,938,724	1.06%	298
4	27	9,922,241	10,745,678	7.66%	10,310,186	10,696,386	3.61%	1,305
5	81	12,448,900	13,533,509	8.01%	12,910,259	13,529,545	4.58%	4,887
6	243	14,990,446	**16,397,505**	**8.58%**	15,213,177	17,995,252	15.46%	22,188
7	749	17,519,461	**19,322,989**	**9.33%**	17,809,500	21,313,900	16.44%	126,295

in the early stages of computation, the Lagrangian Heuristic may yield better upper bounds than CPLEX when working significantly larger problem instances, which is emphasized by the figures in bold. In fact, although we have considered a 100-iteration threshold, our logs show that optimality gaps around 9% are already available for the 7-period problem instance at the tenth iteration (approximately 10,200 seconds of pro-

Table 2.11. Results obtained with the Lagrangian Heuristic limited to 100 iterations and with CPLEX set to the corresponding time limit; tests performed on the scenario tree with $\varphi = 0.50$.

T	$\mathcal{S}$	z^*_{LD}	z^*	LH$_{\mathbf{Gap}}$	CPLEX$_{LB}$	CPLEX$_{UB}$	CPLEX$_{\mathbf{Gap}}$	**Time (s)**
3	9	7,482,267	8,288,998	9.73%	7,968,389	8,042,528	0.92%	300
4	27	10,070,059	11,254,532	10.52%	10,538,063	10,880,329	3.15%	1,278
5	81	12,708,200	14,229,253	10.69%	13,319,878	14,202,873	6.22%	4,743
6	243	15,403,625	**17,438,025**	**11.67%**	15,551,457	19,127,146	18.69%	22,972
7	729	18,130,482	**20,779,377**	**12.75%**	18,348,900	21,670,100	15.33%	27,642+

cessing time) of the Lagrangian Heuristic. That fact suggests the use of the Lagrangian Heuristic as a means to obtaining reasonable feasible solutions in less CPU time than CPLEX. Those initial feasible solutions could then be used as a starting solution for the solver's Branch&Cut procedure, or even for refining metaheuristics.

2.6 Soft-Fixing Local Search

The previous section has shown that alternative solution approaches, such as the Lagrangian Heuristic, may yield interesting results when compared to CPLEX, specially in the early stages of the computation of large problem instances. Since we are still interested in using CPLEX as an inherent part of the solution process, in this Section we study yet another alternative solution approach inspired on well-known local search metaheuristics. The main idea behind the proposed method is to use CPLEX "as a black-box tool to explore solution subspaces defined by the introduction of (invalid) linear inequalities" and then, starting from an (poor, but easily obtained) incumbent, trying to improve it in a short amount of time before feeding it to the solver's exact Branch&Cut algorithm. The approach is inspired on the local branching framework [Fischetti and Lodi, 2003] in the sense of the local search effect provided by the local branching constraint, but without considering the strategic external branching framework required to configure an exact solution method.

2.6.1 The algorithm

Several heuristic methods are based on the *(hard) variable fixing* approach, which (iteratively) uses a black-box solver to quickly obtain (possibly infeasible) solutions and then applies some procedure to fix (some of) its binary variables to cleverly chosen values. One advantage of such methods is that the problem size is reduced after each fixing, therefore allowing the solver to handle smaller problems which, ideally, could be solved to optimality with less effort.

A critical issue, however, is the choice of which variables to fix. Since bad fixings may compromise the entire search procedure by establishing less promising solution subspaces, we are interested in fixing a relevant number of variables without losing the possibility of finding good feasible solutions. Here, we borrow the definitions of Fischetti and Lodi [2003] in their local branching paper by considering the mixed-integer program:

$$
\begin{aligned}
z = \min\, & cx \\
& Ax \geq b \\
& x_j \in \{0,1\} \qquad \forall j \in \mathcal{B} \neq \emptyset, \\
& x_j \in \mathbb{Z}_+^n \qquad \forall j \in \mathcal{G}, \\
& x_j \geq 0 \qquad \forall j \in \mathcal{C},
\end{aligned}
$$

where the variable index set is partitioned into $(\mathcal{B}, \mathcal{G}, \mathcal{C})$ being, respectively, the index set of binary variables and the possibly empty sets of general integer and continuous variables. Also, given a feasible reference solution $\bar{x}$ of z, let $\bar{S} = \{j \in \mathcal{B} : \bar{x}_j = 1\}$ be the binary support of $\bar{x}$, namely, the set of binary variables valued one in the reference solution. For a given positive integer parameter k, we define a neighborhood $\mathcal{N}(\bar{x}, k)$ of $\bar{x}$ as the set of feasible solutions of z satisfying the additional *soft fixing* constraint:

$$
\sum_{j \in \bar{S}} (1 - x_j) + \sum_{j \in \mathcal{B} \setminus \bar{S}} x_j \leq k, \tag{2.104}
$$

where the first term in the left-hand side accounts for the number of binary variables changing their values to one, while the second term accounts for the number of binary variables changing their values to zero, both with respect to $\bar{x}$.

We then develop a procedure to, starting from an initial incumbent $\bar{x}$, iteratively search for improved feasible solutions within the neighborhood $\mathcal{N}(\bar{x}, k)$. In this procedure, whenever a better feasible solution is found, the search restarts from that new improved solution, again considering a neighborhood of size k. On the other hand, if the procedure fails to determine an improved solution, the neighborhood size is expanded by a factor of, say, $\lceil k/2 \rceil$ and the previously searched, unfruitful region is eliminated by the addition of the constraint:

$$
\sum_{j \in \bar{S}} (1 - x_j) + \sum_{j \in \mathcal{B} \setminus \bar{S}} x_j \geq k + 1. \tag{2.105}
$$

Figure 2.8 illustrates the proposed approach. Starting from an initial feasible solution

$\bar{x}_0$, we add constraint (2.104) to the model and solve it again, if possible to optimality, obtaining an improved solution $\bar{x}_1$. The procedure is then repeated, providing improved solutions $\bar{x}_2$ and $\bar{x}_3$. However, while searching neighborhood $\mathcal{N}(\bar{x}_3, k)$, the procedure fails to determine an improved solution. We then increase k by $\lceil k/2 \rceil$ and add constraint (2.105) to the model and define an expanded neighborhood $\mathcal{N}(\bar{x}_3, k+1, \lceil k + k/2 \rceil)$, which is represented by the shaded area around $\bar{x}_3$. Fortunately, an improved solution $\bar{x}_4$ is found, followed by $\bar{x}_5$ and $\bar{x}_6$. And again, while searching neighborhood $\mathcal{N}(\bar{x}_6, k)$, the procedure fails to determine an improved solution. Then, after expanding the neighborhood around $\bar{x}_6$ for another three trials without finding an improved solution, the procedure is terminated — either because it has reached a given threshold on the number of neighborhood expansions, or because it has reached a given time limit. According to the characteristics described above, we name this procedure *soft-fixing*

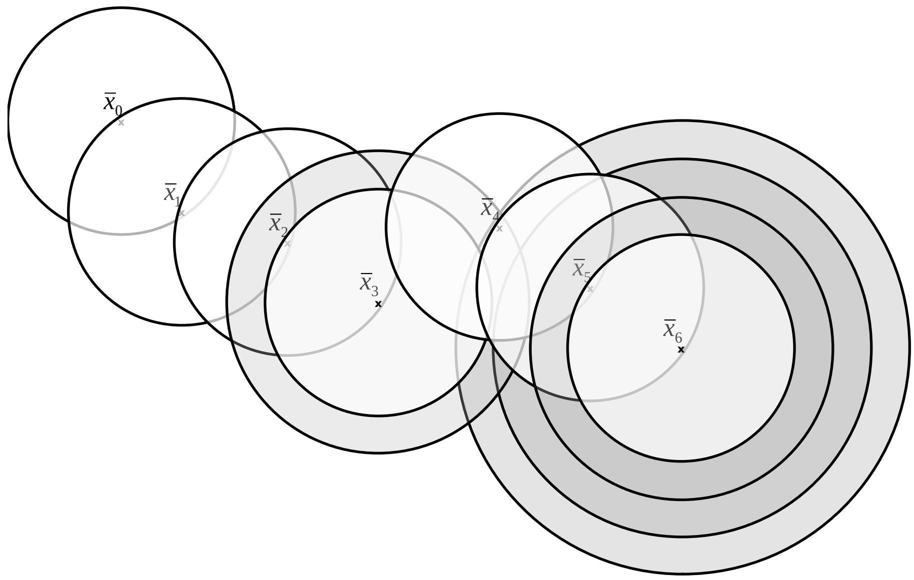

Figure 2.8. An illustrative diagram of the soft fixing local search approach.

local search and give its formal representation in Algorithm 3.

2.6.2 Numerical evaluation

In order to assess the behavior of the proposed soft-fixing local search algorithm, we devise a set of experiments based on the premise that interesting results are expected in the early stages of the computation. According to the description given in the previous section, the following parameters should be observed when running the algorithm:

- the contents of the binary variable subset $B \subseteq \mathcal{B}$ considered in the procedure;

Algorithm 3: Soft-Fixing Local Search

Input: Problem $z = \{\min cx : Ax \geq b, x \in \mathcal{B}\}, B \in \mathcal{B}$.

Let $k = k_0$ be the initial neighborhood size.
Let $\gamma \in (0, 1)$ be an acceptable optimality gap for the solution of each subproblem.
Let $v_0 = (-\infty, k)$.
Let $f = (x_i, v_0, \gamma)$ be a function returning the solution to z, starting from solution x_i and with added constraints (2.105) with lower bound $-\infty$ and (2.104) with upper bound k_0.
Let $\bar{x}_0 = f(*, (-\infty, +\infty), 0.9)$ be an initial feasible solution to z obtained by setting a relative MIP gap of 90% to the original z problem.
Let $e \leftarrow 0$ be the number of neighborhood expansions and E the limit on that parameter.
Let $t \leftarrow 0$ be the elapsed time and T be the amount of time allowed for computation.
Let $i \leftarrow 0$ be the number of iterations.
Let x^* be the best solution found to z.
while $t < T$ and $e < E$ **do**
    $\bar{x}_{i+1} = f(\bar{x}_i, v, \gamma)$.
    if $\bar{x}_{i+1} < \bar{x}_i$ **then**
        $x^* = x_i$.
        $i \leftarrow i + 1$.
        $k = k_0$.
        $v = (-\infty, k)$.
    else
        $v = (k + 1, k + \lceil k/2 \rceil)$.
        $k \leftarrow k + \lceil k/2 \rceil$.
    end if
end while
If there is still any time left for computation, use it to prove $\bar{x}_i$'s optimality.
if $(T - t) > 0$ **then**
    $v = (-\infty, +\infty)$.
    $x^* = f(\bar{x}_i, v, 0.0001)$.
end if

Output: x^*.

- the quality of the initial incumbent k_0, which is obtained by setting a high relative MIP gap on the original problem z;
- the initial neighborhood size, k_0, which influences the both the complexity and the quality of the solution of each subproblem;
- the relative MIP gap γ allowed for the solution of each subproblem;
- the total time T allowed for computation.

All those parameters — and the possible combinations between them — have a strong influence of the results. We have performed a number of preliminary experiments which indicated that the binary subset $B = \{a_i^n : i \in (F \cup S), n \in \mathcal{T}\}$ (please refer to the mathematical notation in Section 2.3.1) provides better results than larger subsets or even than the entire set $\mathcal{B}$. Also, we limit the number of consecutive neighborhood expansions $E = 5$.

Tables 2.12 and 2.13 show the behavior of the soft-fixing local search compared to CPLEX in the early stages of the computation. Note that only the three largest problem instances were considered since, for the smaller ones, CPLEX consistently yields much better results. Though we have tested many different configurations of k_0 and γ, we show only the best results found from a limited number of experiments. The soft-fixing local search algorithm does provide better results than CPLEX in the very early stages of the computation, at least for the 6-period instances of both scenario trees and the 7-period instances for $\varphi = 0.25$. The result is attributed to the cutting effect that the soft-fixing constraints produce on the solution space of each subproblem. Even with a limited number of changes in the binary variables, improved solutions can be determined in a shorter amount of CPU time than the corresponding results obtained by CPLEX. Also, the $\varphi = 0.50$ instances do pose more significant challenges to all three solution procedures. However, the soft-fixing local search results could not produce better results than the ones obtained by the Lagrangian Heuristic.

Table 2.12. Results obtained with the soft-fixing local search approach on the scenario tree with $\varphi = 0.25$.

T	$\mathcal{S}$	k_0	γ	**Iter.**	x^*	z^*_{LH}	CPLEX$_{UB}$	**Time (s)**
5	81	30	0.1	10	16,759,747	13,582,784	16,252,452	1,257
6	243	20	0.2	4	31,248,990	16,730,562	62,208,500	3,668
7	749	30	0.2	2	50,662,591	19,348,580	62,531,602	10,800

Table 2.13. Results obtained with the soft-fixing local search approach on the scenario tree with $\varphi = 0.50$.

T	$\mathcal{S}$	k_0	γ	**Iter.**	x^*	z^*_{LH}	CPLEX$_{UB}$	**Time (s)**
5	81	30	0.1	4	22,212,475	14,407,994	17,182,532	1,200
6	243	20	0.2	4	39,514,669	17,784,559	54,486,925	3,887
7	749	30	0.2	2	59,897,077	20,779,377	48,507,800	10,800

A more profound analysis must be performed in this procedure in order to determine if it is possible to achieve even more improved solutions. Another indication, and perhaps a more relevant one, is that the soft-fixing local search algorithm is actually able to improve a given feasible solution, what suggests its application as a refining procedure for the solutions obtained by Lagrangian Relaxation. Both initiatives, however, are left to future work.

2.7 Discussion and future work

In this Chapter, we have proposed and evaluated the Stochastic Capacity Planning Problem applied to a Global Mining Supply Chain. The model integrates capacity expansion, capacitated facility location and network design decisions within a supply chain planning framework. The size and complexity of the model poses significant challenges to powerful general-purpose optimization solvers and hence demands the development of specific solution procedures.

We have performed a number of analyses on the application of different CPLEX 12.2 settings to the optimization problem; a specific combination of parameters yielded good results for one-hour limited runs and was applied on all subsequent experiments. Those results, however, were still associated to a network and a scenario tree of restricted size. Larger networks, longer planning horizons and additional random parameters — iron ore long-term contract prices, for instance — are bound to be addressed in realistic instances and will certainly require more efficient solution procedures. The insights presented here may drive further developments on those procedures.

The main result presented in this Chapter is the Lagrangian Heuristic designed to derive good dual bounds and feasible solutions to the capacity planning problem. The Lagrangian relaxation considered a dualization of a subset of constraints which yielded a good tradeoff between the quality of the Lagrangian duals and the computation time required to determine them. Though not stronger than CPLEX for smaller problems, the Lagrangian Heuristic takes comparatively less time to run and may provide good approximate solutions when a larger number of scenarios is considered. That is specially interesting as a method to determine good feasible solutions early in the solution process.

Many improvements can be made in the Lagrangian approach. Firstly, one should notice that reducing the initial erratic behavior of the subgradient algorithm may result in faster convergence, even if there is no significant improvement in the Lagrangian dual. An approach based on the use surrogate constraint relaxations within the Lagrangian framework could be exploited to achieve that [Lorena and Senne, 1999]. Secondly, the greedy feasibility heuristic could be enhanced with local search methods in order to explore different capacity decisions or even to anticipate the establishment of new facilities and channels for a given feasible solution. That could avoid unnecessary demand shortage penalties and improve the primal bounds. Also, the initial capacity directives of Algorithm 2 (steps 23-27) establish that initial capacity must approximate the highest activity level among all scenarios derived from the subtree rooted in the node where a facility or channel is opened. However, other different criteria for deciding the initial capacity could be evaluated as, for instance, the scenario with the highest probability, or the scenario with maximum expected activity level. That might yield less conservative capacity configurations and provide interesting results for different demand behaviors. Assessing the dual information provided by the Lagrangian multipliers could provide us with some indication so as to decide initial capacities, as discussed above, or even other supply chain structure decisions.

In this Chapter, we have also proposed a soft-fixing local search procedure motivated by the need to derive good feasible solutions in the early stages of the computation. Though it has outperformed CPLEX in large problem instances, the procedure did not produce better results than the Lagrangian Heuristic. Hence, as indicated by the results, one interesting research direction would be the use of the soft-fixing local search as a refining procedure for the solutions obtained by the Lagrangian Heuristic.

Future work must also scope the development of strong valid inequations to be added to the proposed capacity planning problem. For that matter, specific formulations should be analyzed, as well as the relation of that problem to other similar subproblems [Balakrishnan, 1984]. Additionally, one must consider the development of specific cutting-plane algorithms which take into account the multi-stage, stochastic nature of the problem and might yield better results when compared to standard cutting-plane approaches [Guan et al., 2009].

Decomposition methods could also be used to explore the dynamically capacitated network problem structure as well as the discretized stochastic structure [Ruszczynski, 1997]. Dantzig-Wolfe decomposition, for instance, has been successfully applied on multistage stochastic capacity planning problems [Singh et al., 2009] and could be considered as an approach to the problem presented in this Chapter.

Although not addressed in this Thesis, the development of mathematical approaches to build coherent scenario trees and, more importantly, to reduce it to a minimum yet valuable size could also be studied. That, however, is a different problem altogether, involving the application of specific statistical modeling tools within the mathematical programming framework. Interesting results, however, may be achieved with Scenario Optimal Reduction techniques [de Oliveira et al., 2010].

It is also important to notice that explicit measures of risk could be used as an alternative to the traditional minimum cost, maximum value objective functions. Value-at-Risk measures, for instance, have been employed in facility location problems where the objective is to maximize the lower limit of future earnings based on a given confidence level [Wagner et al., 2009]. Those measures have also been receiving increasingly attention in infrastructure planning problems [Marzano, 2004].

Also, a more realistic, business oriented approach would require including valuation features in the objective function of the stochastic capacity planning problem. Since its main decisions involve investments in production and transportation capacity, and having in mind that it is the ultimate goal of an organization to add value to stockholders' capital expenses, the net present value would be a nature choice for the problem's objective function [Bagajewicz, 2008].

Finally, we point out that several additional analyses could also be made as to the characteristics of the stochastic capacity planning problem under different network structures, as well as different demand behaviors. Those might provide interesting insights to the establishment of conceptual managerial policies regarding investments on supply chain capacity. Also, as the planned decisions are confronted with the actual outcomes of the scenarios, methods to efficiently restart the solution process from a known state of the world may be of great interest.

Chapter 3

Integrated Sales and Operations Tactical Planning in a Global Mining Supply Chain

"However thirsty he was for knowledge, he hated to display ignorance."

When Nietzsche Wept, Irvin D. Yalom

In the context of this Thesis, tactical decisions focus on medium-term planning horizons and involve the effective allocation of production, storage and distribution resources in order to meet iron ore market demand. Large mining operations are often composed of several mines and processing plants, as well as stockyards, railroads, ports and supply stations, each with different capacities and characteristics. According to the demand imposed by domestic and international customers, and depending on the specifications of the ore products supplied by each mine, decisions such as mining, processing, handling, procurement, blending and shipping may involve different facilities at different stages of the supply chain. This Chapter presents a mixed-integer programming model to deal with the tactical sales and operations planning problem in the Global Mining Supply Chain. The model has characteristics of a lot sizing problem in a network environment, but with challenging aspects related to integer flows, discrete production levels and mass losses in concentration and transportation processes. We develop a series of Relax&Fix strategies in order to address realistic sized problem instances. Those strategies are able to outperform CPLEX for most of the several problem instances considered, and with greater success in longer planning horizons. The soft-fixing local search is also evaluated for its ability of determining good feasible solutions in the early stages of the computation. The results given in this Chapter are also presented in the following publications:

Pimentel, B. S., Almeida, F. A., and Mateus, G. R. (2010). Integrated Sales and Operations Planning in a Global Mining Supply Chain. *Submitted to the European Journal of Operational Research.*

Almeida, F. A., Pimentel, B. S., and Mateus, G. R. (2009). Algoritmos para planejamento integrado de produção e transporte de minérios. In *Anais do XLI Simpósio Brasileiro de Pesquisa Operacional*, Porto Seguro, Brazil. Sociedade Brasileira de Pesquisa Operacional.

3.1 Introduction

In the present Chapter, we develop a mathematical model applied to the integrated sales and operations tactical planning in the Global Mining Supply Chain. Again, tactical planning involves multistage decisions such as which ore and where to mine, process and store ore products, as well as which handling equipment to use and which logistic channels to employ in delivering those products to the end customers. As different mines operate at different grades and under different cost and capacity characteristics, the larger the system, the more flexible it may be in dealing with varying operational conditions - nevertheless, the more complex it becomes to determine optimal or even feasible solutions for practical-sized problem instances. Today's large scale mining operations may involve several of those elements organized in a rather intricate network, and decisions may span long planning horizons. Mathematical programming approaches to deal with such a scenario are thus challenged by high computational complexities, preventing the use of conventional, packaged optimization software.

We propose a set of Relax&Fix strategies to deal with such challenges, in an attempt to solve smaller sub-problems by relaxing only a part of the integer variable set and fixing the remaining integer variables to the values given by each sub-problem solution. The heuristic does not guarantee optimality, nor does it ensure that feasible solutions are found for any given partition strategy. However, good approximate solutions may be achieved with a clever selection of the partition criteria. We propose a set of strategies that consider partitioning the integer variable set according to reduced costs, geographic relations, time periods and even randomly. Each strategy is evaluated according to its ability to find feasible solutions.

We consider a problem scenario inspired in large-scale mining operations existing in Brazil, which is composed of three mining systems, three railroads and four ports, all of which operating with some degree of interoperability. Where exact solution procedures fail to achieve feasible solutions for realistic scenarios, we were able to develop specific Relax&Fix strategies which perform rather efficiently, achieving good approximate solutions.

3.2 Tactical Sales and Operations Planning

Sales & Operations Planning (S&OP) is a methodology for coordinating supply chain and demand management decisions [Shapiro, 2010]. Given accurate data inputs, a properly implemented optimization model can unravel the complex interactions and ripple effects across sourcing, production, distribution, inventory and demand management decisions that make S&OP difficult and important.

In the context of the Global Mining Supply Chain, tactical S&OP generally involve determining product portfolio for each mine, production amounts, utilization levels and material flows according to production, storage and distribution capacity constraints. Ore product portfolio for mines and processing plants is somewhat limited in scope, but may show some variation in the medium-term, since the characteristics of the ROM feed may change according to the mine development plan. Ideally, however, that plan should consider the integrated mining supply chain plan, including the overall demand, stock levels, and blending opportunities, as well as individual costs and capacities. Different ROM characteristics may drive different production campaigns and yield different costs and productivity, thus impacting on the ore product mix. It is important to notice that ore processing commonly involves the generation of coproducts, thus imposing a production split constraint on the processing plants. At the design phase, the processing plant features are defined according to the characteristics of the expected ROM. This fact limits the production split and makes setup and changeover times and costs relatively irrelevant on the short-term plan. When more than one production circuit is available — such as the sinter-feed and pellet-feed circuits in some iron ore processing plants — there may be some split flexibility. For instance, if the mine supplies a coarser ROM, the processing plant can still employ additional grinding to generate finer ore products. However, if the mine supplies a finer ROM, the processing plant will certainly yield a higher proportion of fine products, since there is no way of compensating for that input. This asymmetry between supply and demand can lead to difficulties in satisfying quality range specifications.

In such cases, the only way to satisfy demand in both quality and quantity is to explore the flexibility of the supply chain by blending ore products at mines and ports. Also, although production facilities may be relatively multi-purpose, it can be beneficial for the whole supply chain to operate specific plants in a dedicated mode, one that matches a typical ROM supply. Regarding utilization levels, processing plants usually operate at a specific nominal capacity level, where productivity is the highest. Hence, accommodating production to demand level is often achieved by determining discrete idle time, when production stops and management seizes the opportunity to perform maintenance and refurbishment activities. Discrete production levels, which provide an acceptable relation between cost and productivity, can also be considered. Additionally, material flows can be viewed as discrete quantities when transportation is made by trains, cars or trucks. Fixed and variable costs and expected revenues are the main drivers to those tactical decisions.

In general terms, the mathematical model aims at **minimizing production, transportation, procurement and storage variable costs, as well as operational and idle fixed costs throughout the planning horizon**, while observing the following restrictions:

- ROM supply limits at mine pits, as well as procurement limits at third-party supplies;
- Mass losses due to bulk handling processes and mass efficiency at the processing plants, which render unbalanced the flows at the respective nodes;
- Finite, discrete integer production levels at mines and processing plants, which impose minimum lot restrictions on production planning;
- Discrete, integer transportation levels for trains and ships, which impose minimum lot restrictions on transportation planning;
- Capacity limits for production, storage and transportation systems;
- Material balances at the stockyards and the corresponding blending restrictions when composing specific product qualities, and;
- The overall demand for iron ore by different customers throughout the Global Mining Supply Chain.

The following sections detail the mathematical formulation for the proposed Sales and Operation Planning problem.

3.2.1 Mathematical notation

The proposed mathematical model uses the following notation:

Sets :

B : Blending nodes ($B \subset X$);

D : Demand centers;

E_x : Links for (continuous) product flow;

E_y : Links for processing flows;

E_w : Links for third-party product procurement flows;

E_z^k : Links for discrete product flow using mode $k \in Q, E_z = \bigcup_{k \in Q} E_z^k$;

O : Supply centers (mine pits);

P : Ore products, where $P = (P_F \cup P_O \cup P_I)$ includes end, original and intermediate products, respectively;

Q : Discrete transportation modes used in distribution channels (trains, ships, etc.);

R : Third-party procurement supply centers;

S : Ore stockyards;

T : Time periods in the planning horizon;

V : All facilities, where $V = \{B \cup O \cup R \cup S \cup X \cup Y\}$;

X : Transshipment nodes (simple transfer or loading);

Y : Processing plants;

Parameters :

β_{bp} : Fraction of product $p \in P_F$ in the blending of product $b \in P_F$, such that $\beta_{bp} \in (0, 1)$, $\sum_{p \in P_F} \beta_{bp} = 1$, $\forall b \in P_F$;

c_*^{pt} : Unit transportation, storage, production and procurement costs for product $p \in P$ at $t \in T$;

$\hat{f}_i^t$: Fixed operational cost of facility $i \in V$ at $t \in T$;

$\check{f}_i^t$: Fixed idle cost of facility $i \in V$ at $t \in T$;

d_i^{pt} : Demand for product $p \in P_F$ at center $i \in D$ at $t \in T$, with $d_i^{pt} \in \mathbb{R}_+$;

δ_i^{pt} : Penalty applied over backlog of product $p \in P_F$ for customer $i \in D$ at $t \in T$;

ϵ_i^{pt} : Mass loss ratio due to handling, storage and transfer of product $p \in P$ at facility $i \in (S \cup X)$ at $t \in T$, where $\epsilon_i^{pt} \in (0, 1]$;

η_{ij}^{pt} : Mass efficiency of process $(i, j) \in E_y$ for $p \in P$ at $t \in T$, where $\eta_{ij}^{pt} \in (0, 1]$;

l_*^* : Lower bound on transportation and procurement flows and stock levels;

u_*^* : Upper bound on transportation and procurement flows and stock levels;

λ_{ij}^{pt} : Lower bound on process $(i, j) \in E_y$ for generating product $p \in P$ at $t \in T$;

υ_{ij}^{pt} : Upper bound on process $(i, j) \in E_y$ for generating product $p \in P$ at $t \in T$;

o_i^{pt} : Upper bound on the supply of $p \in P_O$ at pit $i \in O$ at $t \in T$;

r_i^{pt} : Upper bound on the supply of $p \in (P_I \cup P_F)$ by the supplier $i \in R$ at $t \in T$;

$\mathcal{C}_{ij}^p$: Production capacity step of process $(i, j) \in E_y$, such that $\mathcal{C}_{ij}^p = \upsilon_{ij}^{pt} / \mathcal{C}_{ij}^L$;

$\mathcal{C}_{ij}^L$: Number of feasible discrete production levels (including null) for process $(i, j) \in E_y$, where $\mathcal{C}_{ij}^L \in \mathbb{N}^*$;

$\mathcal{C}_Q^q$: Average capacity of a product lot transferred by mode $q \in Q$;

$\mathcal{C}_R^{ip}$: Basic lot of purchasing product $p \in (P_I \cup P_F)$ from supplier $i \in R$;

Γ : Demand/capacity ratio, representing how much of the overall system capacity is compromised to the demand imposed by customers;

Decision variables :

b_i^{pt} : Backlog of product $p \in (P_I \cup P_F)$ for demand center $i \in D$ at $t \in T$;

s_i^{pt} : Stock level of product $p \in (P_I \cup P_F)$ in stockyard $i \in S$ at $t \in T$;

x_{ij}^{pt} : Flow of product $p \in P$ through channel $(i, j) \in E_x$ at $t \in T$;

y_{ij}^{pt} : Production level of process $(i, j) \in E_y$ in generating product $p \in P$ at $t \in T$, such that $y_{ij}^{pt} \in \mathbb{Z}_+$;

w_{ij}^{pt} : Number of lots of product $p \in (P_I \cup P_F)$ purchased from supplier $i \in R$, through channel $(i,j) \in E_w$ at $t \in T$, such that $w_{ij}^{pt} \in \mathbb{Z}_+$;

z_{ijq}^{pt} : Number of transportation units employed on the flow of product $p \in P$ through channel $(i,j) \in E_z$, using mode $q \in Q$ at $t \in T$, such that $z_{ijq}^{pt} \in \mathbb{Z}_+$;

a_i^t : Allocation decision for facility $i \in V$, such that $a_i^t \in \{0,1\}$; if $a_i^t = 1$, the facility is operational, and if $a_i^t = 0$, the facility is inactive at $t \in T$.

We assume all capacity parameters to be averaged. It is important to notice that, although integer in the problem instances, those parameters need not be integer themselves — in fact, they define uniform capacity levels that can be accessed by the range of values of the integer devision variables. In order to simplify the model description, we define the following auxiliary terms:

E_χ : $E_\chi = \{E_x \cup E_z\}$, which represents the set of all transportation channels, both continuous and discrete in nature;

γ_{ij}^{pt} : $\gamma_{ij}^{pt} = \mathcal{C}_{ij}^{p} y_{ij}^{pt} \geq 0$, a decision variable that represents the number of units of $p \in P$ produced by process $(i,j) \in E_y$;

ζ_{ijq}^{pt} : $\zeta_{ijq}^{pt} = \mathcal{C}_{Q}^{q} z_{ijq}^{pt} \geq 0$, a decision variable that represents the number of units of $p \in P$ transported by a discrete modal through channel $(i,j) \in E_z$;

χ_{ij}^{pt} : $\chi_{ij}^{pt} = x_{ij}^{pt} + \sum_{q \in Q} \zeta_{ijq}^{pt} \geq 0$, a decision variable that represents the number of units of $p \in P$ transported by both discrete and continuous modals through channel $(i,j) \in E_\chi$;

ω_{ij}^{pt} : $\omega_{ij}^{pt} = \mathcal{C}_{R}^{ip} w_{ij}^{pt} \geq 0$, a decision variable that represents the number of units of $p \in (P_I \cup P_F)$ acquired from supplier $i \in R$.

3.2.2 Mathematical formulation

The objective function seeks minimizing production, transportation, procurement and storage variable costs, as well as operational and idle fixed costs throughout the planning horizon:

$$\begin{aligned}
&\text{Minimize} \\
&\sum_{t\in T}\sum_{p\in P}\left(\sum_{(i,j)\in E_\chi} c^{pt}_{ijq}\chi^{pt}_{ij} + \sum_{(i,j)\in E_y} c^{pt}_{ij}\gamma^{pt}_{ij} + \right. \\
&\left. \sum_{(i,j)\in E_w} c^{pt}_{ij}\omega^{pt}_{ij} + \sum_{i\in S} c^{pt}_{i} s^{pt}_{i} + \sum_{i\in D} \delta^{pt}_{i} b^{pt}_{i}\right) + \\
&+\sum_{t\in T}\left[\sum_{i\in V} \hat{f}^t_i a^t_i + \check{f}^t_i\left(1-a^t_i\right)\right],
\end{aligned} \tag{3.1}$$

where several constraints apply. Firstly, ROM supply limits at mine pits, as well as procurement supply limits must be observed:

$$\sum_j \chi^{pt}_{ij} \leq o^{pt}_i \qquad \forall i\in O, p\in P_O, q\in Q, t\in T, \tag{3.2}$$

$$\sum_j \omega^{pt}_{ij} \leq r^{pt}_i \qquad \forall i\in R, p\in \left(P_I\cup P_F\right), t\in T. \tag{3.3}$$

Also, in all transshipment nodes, mass losses due to bulk handling processes must be taken into account:

$$\sum_j \chi^{pt}_{ij} - \epsilon^{pt}_i \sum_k \chi^{pt}_{ki} = 0 \qquad \forall i\in X, p\in P, t\in T. \tag{3.4}$$

And at all input and output processing plant nodes, mass efficiency must be observed:

$$\gamma^{pt}_{ij} - \eta^{pt}_{ij}\sum_k \chi^{p't}_{ki} = 0 \qquad \forall (i,j)\in E_y, p'\in P, p\in\left(P_I\cup P_F\right), t\in T, \tag{3.5}$$

$$\sum_j \chi^{pt}_{ij} - \gamma^{pt}_{ki} = 0 \qquad \forall (k,i)\in E_y, \forall p\in\left(P_I\cup P_F\right), \forall t\in T. \tag{3.6}$$

At the stockyards, the difference between incoming and outgoing material flows determines the difference in the stock levels for each product at the end of each time period:

$$\sum_j \chi^{pt}_{ij} - \epsilon^{pt}_i\sum_k \chi^{pt}_{ki} - \epsilon^{pt}_i\sum_k \omega^{pt}_{ki} = s^{pt-1}_i - s^{pt}_i \qquad \forall i\in S, p\in P, t\in T. \tag{3.7}$$

Blending is achieved by adding together input products at specific proportions:

$$\sum_b \beta_{bp} \sum_j \chi_{ij}^{bt} - \sum_k \chi_{ki}^{pt} = 0 \qquad \forall i \in B, p \in P_F, t \in T. \tag{3.8}$$

Also, only active facilities can have positive production or transportation flows:

$$\sum_{p \in P} \left(\sum_j \chi_{ij}^{pt} + \sum_k \chi_{ki}^{pt} \right) \leq \mathcal{M} a_i^t \qquad \forall (i,j) \in E_\chi, i \in V, t \in T, \tag{3.9}$$

$$\sum_{p \in P} \gamma_{ij}^{pt} \leq \mathcal{M} a_i^t \qquad \forall (i,j) \in E_y, i \in Y, p \in P, t \in T, \tag{3.10}$$

$$\sum_{p \in P} \omega_{ij}^{pt} \leq \mathcal{M} a_i^t \qquad \forall (i,j) \in E_w, i \in R, p \in (P_I \cup P_F), \tag{3.11}$$

where $\mathcal{M}$ should be at least greater than the sum of the total capacity of the incoming and outgoing links.

Customer demand must be completely served at every time period; penalties will apply to any unmet demand. This is a fair assumption, given the overall system's capability and the constraints imposed by market and logistics dynamics:

$$\sum_k \chi_{ki}^{pt} + b_i^{pt} = d_i^{pt} \qquad \forall i \in D, p \in P, t \in T. \tag{3.12}$$

Alternatively, one can consider customer demand as being stated in terms of the amount of each shipment capacity required at every time interval. In such case, the demand parameter should be restated as $d_{iq}^{pt} \in \mathbb{Z}_+$, $q \in Q$, and the previous equation should be replaced by the following:

$$\sum_k \chi_{ki}^{pt} + b_i^{pt} = \sum_{q \in Q} \mathcal{C}_Q^q d_{iq}^{pt} \qquad \forall i \in D, p \in P, t \in T,$$

$$\sum_k \sum_i \zeta_{kiq}^{pt} = \sum_i d_{iq}^{pt} \qquad \forall i \in D, q \in Q, p \in P, t \in T.$$

Finally, upper and lower bounds on flow and storage capacity must be enforced:

$$l_{ij}^{pt} \leq x_{ij}^{pt} \leq u_{ij}^{pt} \qquad \forall (i,j) \in E_x, p \in P, t \in T, \tag{3.13}$$

$$l_{ij}^{t} \leq \sum_{p \in P} x_{ij}^{pt} \leq u_{ij}^{t} \qquad \forall (i,j) \in E_x, t \in T, \tag{3.14}$$

$$l_{ijq}^{pt} \leq \zeta_{ijq}^{pt} \leq u_{ijq}^{pt} \qquad \forall (i,j) \in E_z^q, q \in Q, p \in P, t \in T, \tag{3.15}$$

$$l_{ij}^{pt} \leq \omega_{ij}^{pt} \leq u_{ij}^{pt} \qquad \forall (i,j) \in E_w, \forall i \in R, \forall p \in (P_I \cup P_F), t \in T, \tag{3.16}$$

$$l_{ijq}^{t} \leq \sum_{p\in P} \zeta_{ijq}^{pt} \leq u_{ijq}^{t} \qquad \forall (i,j) \in E_z^q, q \in Q, p \in P, t \in T, \tag{3.17}$$

$$l_{ij}^{t} \leq \sum_{p\in P} \sum_{q\in Q} \zeta_{ijq}^{pt} \leq u_{ij}^{t} \qquad \forall (i,j) \in E_z, t \in T, \tag{3.18}$$

$$l_{i}^{pt} \leq s_i^{pt} \leq u_i^{pt} \qquad \forall i \in S, p \in P, t \in T, \tag{3.19}$$

$$l_{i}^{t} \leq \sum_{p\in P} s_i^{pt} \leq u_i^{t} \qquad \forall i \in S, t \in T, \tag{3.20}$$

$$\lambda_{ij}^{pt} \leq \gamma_{ij}^{pt} \leq \upsilon_{ij}^{pt} \qquad \forall (i,j) \in E_y, p \in P, t \in T, \tag{3.21}$$

as well as integrality and nonnegativity constraints over the decision variables:

$$x_{ij}^{pt}, s_k^{pt}, b_l^{pt} \geq 0 \qquad \forall (i,j) \in E_x, k \in S, l \in D, p \in P, t \in T, \tag{3.22}$$

$$\chi_{ij}^{pt} \geq 0 \qquad \forall (i,j) \in E_\chi, p \in P, t \in T, \tag{3.23}$$

$$z_{ijk}^{pt}, y_{lm}^{pt}, w_{oq}^{pt} \in \mathbb{Z}_+ \qquad \forall (i,j) \in E_z, k \in Q, (l,m) \in E_y, (o,q) \in E_w, p \in P, t \in T, \tag{3.24}$$

$$a_i^t \in \{0,1\} \qquad \forall i \in V, t \in T, \tag{3.25}$$

$$\gamma_{ij}^{pt} \geq 0 \qquad \forall (i,j) \in E_y, p \in P, t \in T, \tag{3.26}$$

$$\zeta_{ijk}^{pt} \geq 0 \qquad \forall (i,j) \in E_z, p \in P, t \in T. \tag{3.27}$$

3.3 A simple example

In order to evaluate the main features of the proposed model, we present the simplified mining supply chain presented below. Each facility has known production, storage and

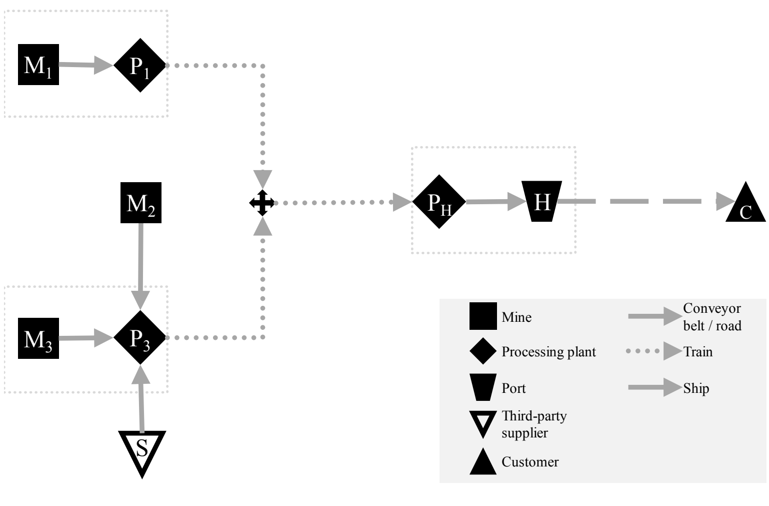

Figure 3.1. A simple example for the S&OP problem.

transfer capacities, as well as predefined product portfolio and operational costs. The supply chain has three mines M_1, M_2 and M_3, one port H and three processing plants — P_1 and P_3 associated to mines M_1 and M_3, respectively, and P_H, a pelletizing plant, associated to the port. Since mine M_2 does not have a processing plant, all its ROM must be processed at P_3. Production and storage capacities and costs are shown in Table 3.1. In this example, we assume that the facilities are not constrained on the shipping capacity; however, standard capacity of trains and ocean carriers shipments is 50 and 100 units, respectively, while trucks may carry any amount of material in each shipment. A single customer C demands products s and p, where product s is formed at the port from a blend of 80% of s_1 and 20% of s_3, and pellets p are formed by agglomerating pellet feed pf at the port's pelletizing plant PH. If demand for s_3 is higher than available capacity, third-party procurement may be employed from supplier S, although at a much higher cost (and again, in this example, with unlimited capacity). For the sake of simplicity, all operational fixed costs are set to \$100, except for the procurement node, which has a fixed cost of \$1000; idle costs are all set to \$10. Also, we assume no losses in any processes $\left(\epsilon_i^{pt} = \eta_{ij}^{pt} = 1\right)$, so that flow balances are always maintained. All facilities have only three production levels: 0%, 50% and 100% of the respective nominal capacities.

Table 3.1. A simple example for the S&OP problem.

Facility	Production	Storage	Unit Cost
Mine M_1	rom : 2000	2000	1.0
Processing plant P_1	s_1 : 1000, pf : 1000	2000	1.0
Mine M_2	rom : 500	1000	5.0
Mine M_3	rom : 1500	2000	5.0
Processing plant P_3	s_3 : 1000, pf : 1000	2000	5.0
Supplier S	s_3 : 500	—	50.0
Railway	—	—	1.0
Processing plant P_H	p : 1000	1000	2.0
Harbor H	s : 4000, p : 1000	7000	2.0

In this example, we assume a scenario of increasing demand, which drives the supply chain to employ different resources at each of the three time periods considered in the planning horizon. At $t = 1$, demand levels are low, so the system directs production towards the lower cost facilities at mines M_1 and M_3 and processing plants P_1 and P_3 until all demand is met. Note that exactly 0.8 parts of s_1 and 0.2 parts of s_3 are needed to deliver each unit of s. At $t = 2$, demand doubles and all mines, including M_2, must be employed to ensure proper delivery. At $t = 3$, demand doubles again and some third-party procurement must be employed at the expense of higher operational costs. At $t = 1$, as seen on Table 3.2, 100 units of s_1 and 400 units of s_3

Table 3.2. Results for the S&OP example at $t = 1$.

$t = 1$; **Demand of** $s : 500$, $p : 500$				
Facility	**Production**	**Shipping**	**Storage**	**Procurement**
Mine M_1	$rom : 1500$	—	—	—
Processing plant P_1	$s_1 : 500$, $pf : 1000$	$s_1 : 100$, $pf : 1000$	$s_1 : 400$	—
Mine M_2	—	—	—	—
Mine M_3	$rom : 500$	—	—	—
Processing plant P_3	$s_3 : 500$	$s_3 : 500$	—	—
Harbor H and P_H	$s : 500$, $p : 1000$	$s : 500$, $p : 500$	$s_3 : 100$, $p : 500$	—

are required to meet the demand of product s. Processing plant P_1 uses 50% of s_1 and 100% of the pf production circuits' capacities, generating 500 units of s_1 and 1000 units of pf, respectively. Of this production, only 100 units of s_1 and all 1000 units of pf are shipped to the harbor, leaving 400 units of s_1 at the stockyard of P_1. Also, since those facilities have the lowest cost, all pf demanded is provided by them. Processing plant P_3's capacity is constrained by Mine M_3 production, thus only employing 50% of its nominal capacity to produce 500 units of s_3 and shipping all of them to the harbor. There, two process take place: the production of 1000 units of p from all the pf received from P_1; and the blending of 500 units s from 100 units of s_1 and 400 units of s_3. Although only half the p produced is needed to meet the demand, the extra production will be needed at the following time periods — also, the lower storage costs of P_1 and P_H compensate the anticipated production. Hence the harbor ships 500 units of s and p and stores the 100 leftover units of s_3 and the 500 units of p.

At $t = 2$, demand doubles and the bottleneck production capacity of M_3 requires Mine M_2 to be employed in producing 500 units of ROM to complement processing plant P_3 supply. In order to meet s demand, 200 units of s_1 and 800 units of s_3 are needed, which are consumed from both received cargo and the port's stockyard. Since 500 units of p were left in the stockyard, only 500 new units from the 1000 units produced are required, thus leaving again a stock of 500 units of p.

Table 3.3. Results for the S&OP example at $t = 2$.

$t = 2$; **Demand of** $s : 1000$, $p : 1000$				
Facility	**Production**	**Shipping**	**Storage**	**Procurement**
Mine M_1	$rom : 1000$	—	—	—
Processing plant P_1	$pf : 1000$	$s_1 : 200$, $pf : 1000$	$s_1 : 200$	—
Mine M_2	$rom : 500$	—	—	—
Mine M_3	$rom : 500$	—	—	—
Processing plant P_3	$s_3 : 1000$	$s_3 : 1000$	—	—
Harbor H and P_H	$s : 1000$, $p : 1000$	$s : 1000$, $p : 1000$	$s_3 : 300$, $p : 500$	—

At $t = 3$, demand increases again and all production and distribution capacity available is not able to meet the delivery goals. Therefore, 500 units of s_3 are procured from the third-party supplier S. Since 400 units of s_1 and 1600 units of s_3 are now required to meet s demand, 1300 units of s_3 are shipped from Processing plant P_3, which are composed of 1000 units of production and 300 units of procurement. Those 1300 units are combined with other 300 units of s_3 left in the harbor stockyard from the previous time period. It is important to point out that, since each processing plant can only operate in discrete production levels, it may not be possible to use up all surplus production, and some stored material may have to be left.

Table 3.4. Results for the S&OP example at $t = 3$.

$t = 3$; **Demand of** $s : 2000$, $p : 1500$				
Facility	**Production**	**Shipping**	**Storage**	**Procurement**
Mine M_1	$rom : 1500$	—	—	—
Processing plant P_1	$s_1 : 500$, $pf : 1000$	$s_1 : 400$, $pf : 1000$	$s_1 : 300$	—
Mine M_2	$rom : 500$	—	—	—
Mine M_3	$rom : 500$	—	—	—
Processing plant P_3	$s_3 : 1000$	$s_3 : 1300$	—	$s_3 : 300$
Harbor H and P_H	$s : 1000$, $p : 1000$	$s : 2000$, $p : 1500$	—	—

Although this is a simple example, some important analyses can be made. For instance, depending on production plans, maintenance schedules may be adjusted to meet a compromise between overall capacity and demand. For that matter, any known variations in nominal capacities may be taken into account when planning production, storage and transfer throughout the supply chain. That scopes both product volumes and quality, since different products may be blended to meet the desired delivery specifications. Also, since production levels are usually set to maximum capacities, production and transfer decisions may have important impacts on stock levels. Bottlenecks within the supply chain could be identified and compensated by the overall capacity. Such analysis could suggest, for example, investment allocation to capacity expansion.

3.4 The Relax&Fix approach

In this Chapter, we are interested in solving S&OP problem instances of practical sizes. However, previous work [Almeida et al., 2009] has shown that the characteristics of the proposed model may pose significant challenges for modern commercial solvers, such as CPLEX. Also, depending on the problem structure, it may be difficult even to find feasible solutions in order to explore the neighborhood solution space. The alternative is thus to develop special-purpose solution algorithms to determine good enough feasible solutions in a reasonable amount of time.

A solution technique that has been successfully applied in large, complex mixed-integer problems is the Relax&Fix strategy [Dillenberger et al., 1994]. Relax&Fix is a general-purpose methodology that solves the original problem in a number of steps, each involving a subproblem of smaller complexity than the original one. That is accomplished first by partitioning the integer variable set into a number of disjunctive subsets. The procedure then tries to solve a sequence of mixed integer subproblems in which the variables of only a single subset remain integer and all others are linearly relaxed. Since the number of integer variables is reduced at each stage, the subproblems can be solved with relative efficiency. As the series advances, each integer variable subset is permanently fixed to the optimal values obtained at the previous stage and the integrality constraint is forced upon the integer variables of the current stage. Solution quality is measured by a gap calculated as:

$$\text{gap}(\%) = 1 - \frac{\iota}{\nu} \tag{3.28}$$

where ι represents the lower bound obtained with the relaxation of the first stage (or with the linear relaxation of the original problem), and ν represents the upper bound obtained with the solution of the last stage. The value ι is clearly a dual bound, since it is obtained by a partial linear relaxation of the original problem. Similarly, ν is clearly a primal bound, since it is a feasible solution obtained by the Relax&Fix solution.

Here we adapt the enhanced Relax&Fix algorithm due to Escudero and Salmeron [2005]. It is important to notice that, since only a portion of the solution space is explored, the algorithm cannot guarantee optimality, nor can it ensure that feasibility is maintained throughout the computations. A backtracking grouping step is usually proposed to deal with such drawbacks, at the cost of increasing the complexity of the infeasible subproblem. In the worst case, however, the algorithm may end up solving the original problem and hence may fail to determine an efficient solution.

The Relax&Fix heuristic has been largely used as a method to obtain good primal bounds (feasible solutions) for hard mixed integer programs either on its own or as part of hybrid algorithms. However, it is clear that the variable set partitioning and the variable fixing criteria have a strong connection with the complexity and with the feasibility of each subproblem. In multi-period or multistage problems, it is a natural choice to partition the variable set according to each period or stage. Escudero and Salmeron [2005] propose a series of partition strategies for a project scheduling problem according to promptness, value, cost, cost-benefit ratio, or even randomly

Algorithm 4: Relax&Fix

Input: Problem $\mathbf{P} : \min \{c(x, y) : x \in X \subseteq \mathbb{Z}^n, y \in Y \subseteq \mathbb{R}^n\}$

Let $V_1, \ldots, V_k$ be a set of disjoint partitions of the integer variable set V.
Let $\mathbf{P}^r$ be a subproblem where only variables V_r are kept integer and all others are relaxed.
$r \leftarrow 1$
Solve $\mathbf{P}^1$
if $\mathbf{P}^1$ is infeasible **then**
  Stop, since $\mathbf{P}$ would also be infeasible.
else
  $\iota(\mathbf{P}) \leftarrow \iota(\mathbf{P}^1)$
end if
while $r \neq k$ **do**
  $r \leftarrow r + 1$
  Solve $\mathbf{P}^r$
  if $\mathbf{P}^r$ is infeasible **then**
    $V_{r-1} \leftarrow V_{r-1} \cup V_r$
    for all $i = r, \ldots, k-1$ **do**
      $V_i \leftarrow V_{i+1}$
    end for
    $k \leftarrow k - 1$
    $r \leftarrow r - 1$
    if $r = 1$ **then**
      Go back to step 4.
    end if
  end if
end while
$\iota(\mathbf{P}) \leftarrow \iota(\mathbf{P}^k)$
Stop, since $\mathbf{P}$ is feasible.

Output: Solution to $\mathbf{P}$.

selecting the variables of each stage. Beraldi et al. [2006] address a stochastic lot-sizing and scheduling problem and develop a Relax&Fix strategy that considers a time partitioning scheme embedded in each node of the problem's scenario tree. Ferreira et al. [2009] evaluate partition strategies based on time, assignment, function, and combinations of those in a soft drink plant integrated lot-sizing and scheduling problem.

3.4.1 Partition strategies

The combined lot-sizing and network flow model we propose in this paper aims at minimizing all operational costs throughout the global mining supply chain. It is important to notice that this model has specific characteristics which impose significant challenges to obtaining feasible solutions — our preliminary experiments have shown that feasibility, and not optimality, is the main issue. This suggests that, when designing Relax&Fix strategies, the partition schemes must be careful not to violate any balance, nor to create any conditions that might violate balances within the network. Also, the assignment restrictions imposed over production facilities must ensure that the system capacity meets the overall demand, especially considering that discrete production levels may generate flows that are significantly below or above the exact demanded quantities or even the transfer and storage capacities.

Table 3.5 lists the Relax&Fix strategies proposed in this work. The PT and

Table 3.5. Relax&Fix strategies.

Strategy	**Partition**	**Stages**
PT	Functional	Production, transportation
TP	Functional	Transportation, production
R	Random	Predefined number of stages
RC	Reduced costs	Predefined number of stages
GF	Geographic	From mine to port
GB	Geographic	From port to mine
TF	Time periods	$t = 1, 2, \ldots, T$
TB	Time periods	$t = T, T-1, \ldots, 1$
TFP	Time periods/products	$t = 1, 2, \ldots, T$; P
TBP	Time periods/products	$t = T, T-1, \ldots, 1$; P
TFPT	Time periods/functional	$t = 1, 2, \ldots, T$; production, transportation
TFTP	Time periods/functional	$t = 1, 2, \ldots, T$; transportation, production
TBPT	Time periods/functional	$t = T, T-1, \ldots, 1$; production, transportation
TBTP	Time periods/functional	$t = T, T-1, \ldots, 1$; transportation, production
ATF	Assignment/time periods	a; $t = 1, 2, \ldots, T$
ATB	Assignment/time periods	a; $t = T, T-1, \ldots, 1$

TP strategies partition the integer variable set according to functional characteristics in two stages: production variables and transportation variables. Depending on the problem size, however, PT and TP may still have to deal with difficult subproblems. Strategy R randomly groups the integer variables into a predefined number of stages, thus making it possible to control the actual number of variables per stage, regardless of time or function. Strategy RC also uses a predefined number of stages but groups the integer variables according to the highest reduced costs given by the linear relaxation of each subproblem. Strategies GF and GB divide the integer variable space according to the main supply chain structures: mines, railway, supply stations and port, solving them in the forward and backward directions, respectively. Such partition, although intuitive, does not always ensure that a reasonable number of variables is present in each stage, especially for longer planning horizons, nor does it ensure that flow conservation is met for any solving sequence. Strategies TF and TB both partition the integer variable space according to the number of time periods in the planning horizon and then traverse it to the forward or backward direction, respectively. The following six strategies all partition the variable set according to time periods, varying direction (forward of backward in time) and further partitioning the subsets according to production/transportation or geographic criteria. Strategies ATF and ATB, on the other hand, try to establish all facilities that should be operational before applying the TF and TB strategies. Combinations of the above strategies in stages and sub-stages are also possible, given that the appropriate consistency conditions are met.

It is important to mention that the set of strategies proposed here is clearly non-exhaustive; although several others can be imagined, we have focused on the ones that seemed to produce the most significant results in both feasibility and efficiency. The following section further evaluates the behavior of the proposed strategies.

3.4.2 Numerical evaluation

In this section we analyze the computational results of the solution strategies proposed before. The CPLEX 12.2 solver was applied both individually and as part of the Relax&Fix framework. In all runs we have limited the solution time to reasonable values in order to allow comparing the different solution strategies against the exact methods employed by the solver. All runs were executed in an Intel Xeon E5520 workstation with four cores (up to eight threads), 2.27GHz and 16GB RAM.

We consider a setting inspired in large-scale mining operations existing in Brazil. Three mining systems A, B and C produce iron ore at different rates, grades and costs, using three different railroads with some level of interoperability. Four ports are used for ore export and seven supply stations — three domestic and four international — are used as advanced commercial stockyards. Some international customers are served by ocean shipments directly from the ports, while others can only be served by shipments from the international supply stations. System A has seven mines and reaches four domestic customers, while System B has nine mines and reaches six domestic customers, and System C has only one mine and reaches three domestic customers from three supply stations. Although each system may operate with relative independence, interoperability may be considered when determining which system will be responsible for satisfying a given international demand. Note that, for clarity reasons, specific elements within each subsystem are not represented in the network — instead, we assume that they are contained within each mine, port or supply station node. Some mines may have one or more processing plants and stockyards, while others may have to transfer their ROM to plants or stockyards in other mines, and some may depend entirely on ROM provided by other mines. Accordingly, each port may have a different number of pelletizing plants, stockyards and piers. The demand centers represent customers, both domestic and international.

Table B.1 in Appendix B describes the characteristics of the main supply chain facilities of each mining system, as considered in the tests. Besides the cost and capac-

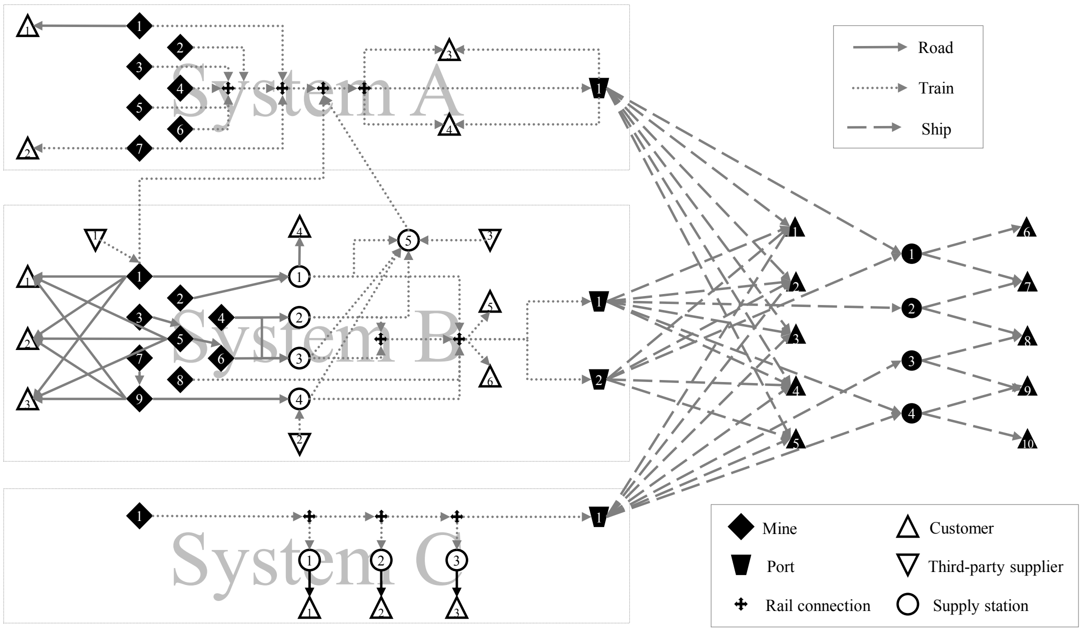

Figure 3.2. Settings for the S&OP problem.

ity parameters, one should pay special attention to the flow loss factors associated to transportation and storage losses, to the efficiency of the processing plants. When those parameters are different than 1, the efficiency of the solution algorithms is severely affected, since a significant part of the flow integrality is lost. Although those parameters may vary for each facility, in average they have the magnitudes presented in Table 3.6.

Table 3.6. Mass and process efficiency parameters.

Parameter	**Averaged Value**
Mass loss due to transportation	0.98
Mass loss due to storage and handling	0.99
Beneficiation process efficiency	0.71
Pelletizing process efficiency (ore)	0.97

The size of the product portfolio is strongly related to market conditions. As it has been stated, when ore product demand and prices are high, the customers tend to prioritize quantity over quality, which drive the mining companies to limit the commercial product portfolio to a more manageable size. This means that the quality ranges that define the ore products tend to be wider, thus leading to a less diversified portfolio. On the other hand, when demand is lower, customers tend to prioritize quality over quantity and hence, the commercial efforts consider a more diversified product portfolio in order to meet specific shipping demands. In such conditions, selecting the appropriate facilities to mine, process and blend the end ore products is of utmost importance. Nevertheless, despite that behavior, in our tests we consider a fixed size portfolio with only five end products. This is done in order to facilitate the comparison between the efficiency of each solution strategy. We provide computational results for several different system conditions, defined by the magnitude of the demand/capacity ratio Γ and the size of the planning horizon T.

We propose a series of problem instances specifically designed to explore the problem structure, the CPLEX solver's performance and the behavior of the Relax&Fix heuristic. The instances are mainly characterized by the behavior of the demand imposed on the supply chain and can be categorized as follows:

- Increasing demand, from 25% to 100% of the system's overall capacity;
- Decreasing demand, from 100% to 25% of the system's overall capacity;
- Seasonally increasing demand, from 25% to 100% of the system's overall capacity, in one or more cycles depending on the planning horizon;
- Seasonally increasing demand, from 100% to 25% of the system's overall capacity, in one or more cycles depending on the planning horizon;
- Uniformly varying demand, increasing variance from 0% up to 125% of the system's overall capacity (with the exceeding demand possibly being met by third-party suppliers) in cycles of two consecutive periods.

Moreover, for each group of problem instances we evaluate different configurations for the fixed operational costs: the standard one shown in B.1 and another one with a general 10-time increase on that parameter. Also, regarding the planning horizon, all tests are executed with $T = \{3, 6, 12, 24, 48\}$ time periods. One should, however, notice that the behavior of demand in real-world scenarios may be rather simpler and may show much smaller variances than the synthetic instances proposed here.

3.4.2.1 General results

As discussed in the previous Chapter, the CPLEX package allows setting up several parameters that control the behavior of its solution algorithms as well as embedded heuristics and pre- and post processors. Here, we have also run trials with different configurations of many combinations of parameters. However, we have observed that, for the model at hand, the emphasis of the mixed-integer Branch&Cut solution procedure, which is controlled by the *MIP emphasis* parameter, is dominant. The default behavior tries to achieve a balance between optimality and feasibility, and may sometimes fail to achieve a feasible solution within a reasonable amount of time. Having observed that, we have configured all our experiments with a higher emphasis on feasibility than on optimality (*MIP emphasis* set to 1). All runs of CPLEX were limited to one hour.

The behavior of each Relax&Fix strategy is evaluated using the same configurations described above. The time limit for each run was also set to one hour, uniformly divided among each iteration. It is important to notice, however, that the proposed strategies display very distinct behaviors. For instance, the random (R) and reduced cost (RC) strategies failed to produce any feasible solutions. This unfortunate, but not unexpected result is most likely due to the infeasible flow balances generated by the order in which the integer variables were solved and fixed. Despite that, the results oriented us in further developing the other strategies.

The GF and GB strategies, on the other hand, organize the integer variables according to their geographic relationships. It is clear that there is a trade-off between the number of stages — and the corresponding amount of integer variables in each stage — and the solution quality. Again, however, determining feasible solutions is the main issue. The solution sequence, that is, the order in which each supply chain facility has its integer variables solved and fixed, has a strong effect on ensuring adequate flow balances on the network. We have tested a number of such configurations and found that not every solution sequence produces feasible results. Nevertheless, we have observed that sequences which start at the upstream facilities (from the ports to the mines) tend to cause less feasibility problems while generating more reasonable solution gaps. In any case, the development of a geographic strategy has proven to be a challenging trial-and-error task.

With the exception of the TF and TB strategies, all other time-related partitions failed to produce feasible results, including the ATF and ATB strategies. Although small enough subproblems can be derived from those partitions, maintaining consistency — especially satisfying the material balances when discrete production levels and transportation capacities are considered — among the larger number of subproblems can be quite challenging. The TF and TB strategies, on the other hand, organize the integer variables according to the time periods in the planning horizon and traverse the variable subsets in the forward and backward directions — from the first period to the last and from the last period to the first —, respectively. Time strategies have shown rather strong results for almost every configuration tested in this Thesis.

3.4.2.2 Analysis

Considering the large number of problem instances — a total of 90 problem instances, detailed in Appendix C — we present below only one of the several cases evaluated, as a means to illustrate the general behavior of the solution approaches. Table 3.7 shows, for each T, the number of variables — emphasizing the number of integer and binary variables —, the number of constraints and the results obtained by using CPLEX and the TF and TB Relax&Fix strategies to solve those five problem instances. From the

Table 3.7. Results of CPLEX and Relax&Fix for increasing planning horizons in a scenario of uniformly decreasing demand.

T	3	6	12	24	48
Variables	4,635	9,270	18,540	37,080	74,160
Integer	1,503	3,006	6,012	12,024	24,048
Binary	804	1,608	3,216	6,432	12,864
Constraints	7,417	14,626	29,044	57,880	115,552
CPLEX LB	1,041,321	1,495,380	2,745,308	5,296,511	10,429,505
CPLEX UB	1,059,840	1,615,756	4,638,513	533,446,174	1,072,093,875
CPLEX gap	2%	7%	41%	99%	99%
CPLEX t(s)	3,600	3,600	3,600	3,600	3,600
TF LB	960,961	1,354,930	2,525,753	4,902,867	9,684,487
TF UB	1,065,732	1,612,831	3,021,139	5,919,244	11,888,284
TF gap	10%	16%	16%	17%	19%
TF t(s)	4	12	30	116	498
TB LB	—	—	2,483,394	—	—
TB UB	—	—	16,362,457	—	—
TB gap	—	—	85%	—	—
TB t(s)	—	—	8	—	—

full listing of results some information may be derived and some important conclusions may be drawn regarding feasibility, optimality and CPU time:

- CPLEX is able to determine feasible solutions in all problem instances, but good results are restricted to shorter planning horizons (at most $T = 12$). Optimality gaps thus range from 1.65% to 99.04%, with much poorer results associated to $T = 24$ and $T = 48$ — in fact, all problem instaces with $T = 48$, show gaps higher than 98%. Also, CPLEX tends to use all CPU time available, a total of 3,600s. Although one could argue that longer solution time limits might allow CPLEX to achieve better solutions, some runs with 28,800s (eight hours) yielded gaps around 20%. On the other hand, it is noticeable that CPLEX's lower bounds are consistently higher than the ones obtained by the TF and TB strategies — what could be expected, since our Relax&Fix algorithm uses the linear relaxation of the oricinal problem as the lower bound.

- The TF strategy is able to determine feasible solutions in all problem instances, with gaps that range from 9.83% (in a $T = 3$ instance) to 35.57% (in a $T = 48$ instance), and demanding as low as 1s to 1,290s of CPU time. However, it is important to notice that the TF strategy takes less than 180s in 76% of the intances, and less than 60s in 62% of the instances. Also, in 92% of the instances the TF strategy achieves optimality gaps below 25%, and in 24% of the instances gaps are below 15%.

- The TB strategy's behavior, on the other hand, is rather disappointing. Only on 29 instances (32% of the total) does it determine feasible solutions — it is true that feasibility is most easily achieved on the decreasing demand instances. However, for those solutions, the gaps are consistently worse, ranging from 14.09% to 89.57% (only 48% of the time achieving figures below 25%). Processing time is otherwise improved, with more than 80% of the instances being executed in less than 60s. That, however, does not compensate the poor upper bounds determined by the strategy.

Regarding the upper bounds obtained by each approach, the data clearly indicates that better results are obtained by the TF strategy when solution time is limited to 3,600s, with special attention to the following:

- In 61 of the 90 problem instances (68% of the the total), TF provides better upper bounds than CPLEX — in fact, that is true for all $T = 24$ and $T = 48$ instances and for 94% of the $T = 12$ instances.

- Strategy TF also beats strategy TB 98% of the time. Moreover, upper bounds up to eight times smaller can be determined by TF.

- For the 29 instances in which strategy TB achieves feasible solutions, on only 24% of them do the upper bounds beat the ones determined by CPLEX.

- The effect of higher fixed costs (FC=10, that is, an increase of ten times on the fixed costs) is clearly noticeable in the problem instances. Applying the TF strategy, the optimality gaps of the FC=10 instances are always higher (1.5 times, in average) than the corresponding ones of the FC=1 instances. The respective CPU times are higher 67% of the time for the FC=10 problem instances.

The above observations may be further illustrated by Figures 3.3 and 3.4, which show the suitability of the proposed solution approaches to the different behaviors of the iron ore demand, as displayed in the problem instances. In those figures, we only

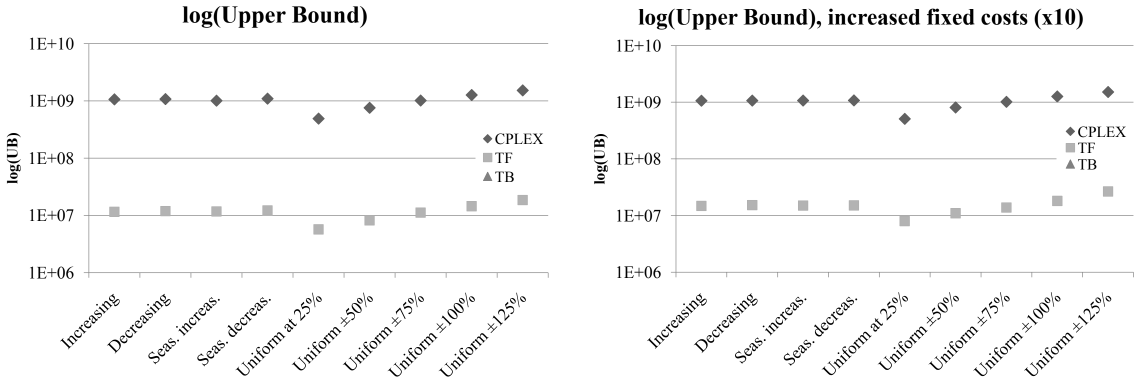

Figure 3.3. Upper bounds obtained by CPLEX and Relax&Fix for the instances with $T = 48$.

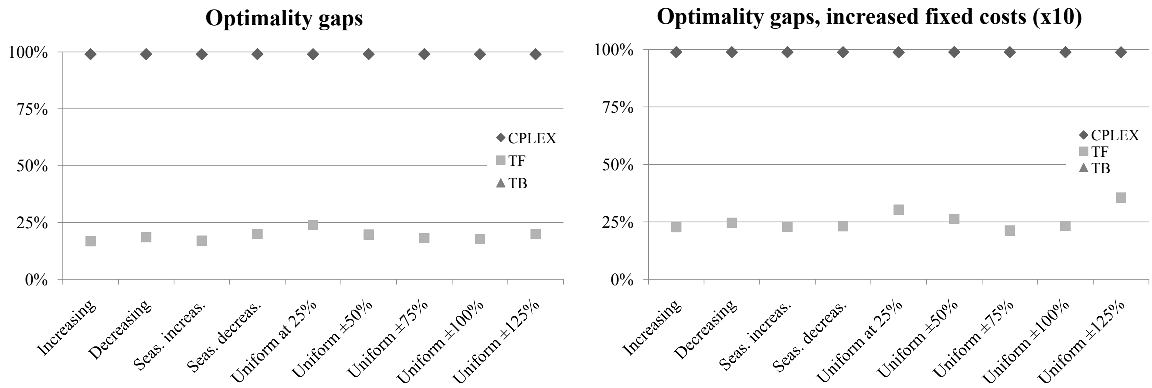

Figure 3.4. Optimality gaps obtained by CPLEX and Relax&Fix for the instances with $T = 48$.

show the results related to $T = 48$, since they are the ones which pose most significant challenges to the solution approaches. There seems to be no particular correlation between the demand magnitude and variation to the complexity of solving each problem instance. One important observation, however, is that the TF strategy seems to organize the integer variables in a more coherent fashion which provides improved feasibility characteristics for each subproblem. This can be explained by the fact that the only connections between stages are established by continuous variables at the stockyard nodes (Equation 3.7), thus facilitating flow balance and hence, feasibility. This condition is not sufficient, of course — as it has been stated before, the Relax&Fix algorithm does not ensure feasibility.

3.5 The soft-fixing local search approach

The obvious difficulties observed in solving larger S&OP problem instances suggest the evaluation of the soft-fixing local search approach presented in Section 2.6. Despite the large number of problem instances, we concentrate our analysis on the one which showed the highest optimality gap for the TF strategy (and possibly the most difficult to solve): uniformly varying demand with ±125%, $T = 48$ and FC=10. As shown in Appendix C, CPLEX was not able to find good feasible solutions for none of the $T = 48$ instances.

Here, we consider $B = \mathcal{B}$, which means that the entire binary variable subset is addressed in the soft-fixing local search algorithm. All runs are again limited to 3,600s. In order to assess the behavior of the procedure, we have run tests with different values of k_0, the initial size of the search neighborhood, but with γ, the relative gap limit for the search subproblems, always set to 0.01. Table 3.8 shows that the soft-fixing local

Table 3.8. Results obtained with the soft-fixing local search approach on the Uniform ±125%, $T = 48$, FC=10, S&OP problem instance. All runs limited to 3,600s with $\gamma = 0.01$.

k_0	**Iter.**	x^*	CPLEX_{UB}	TF_{UB}
25	156	1,396,740,893		
50	56	1,266,848,097		
100	21	**1,245,339,105**	1,516,160,195	26,620,013
200	7	1,267,382,904		
300	3	1,320,805,664		

search can improve the upper bound provided by CPLEX by a factor that ranges from 8% to 21%, depending on the value of k_0. However, the high penalty values associated to unmet demand make the upper bounds obtained by the procedure still quite high and much worse than the ones yielded by Relax&Fix's TF strategy. The evolution of the solution procedure is detailed in Figure 3.5. It can be seen that best results are achieved with $k_0 = 100$, considering the same time limit of 3,600.

Apart from the results shown above, as discussed in the previous section, the soft-fixing local search framework could also be used as a refining procedure to improve the bounds obtained by Relax&Fix. That is, however, left to future work.

3.6 Discussion and future work

This Chapter addressed an integrated sales and operations planning problem at the tactical level of a Global Mining Supply Chain. This is a different application for

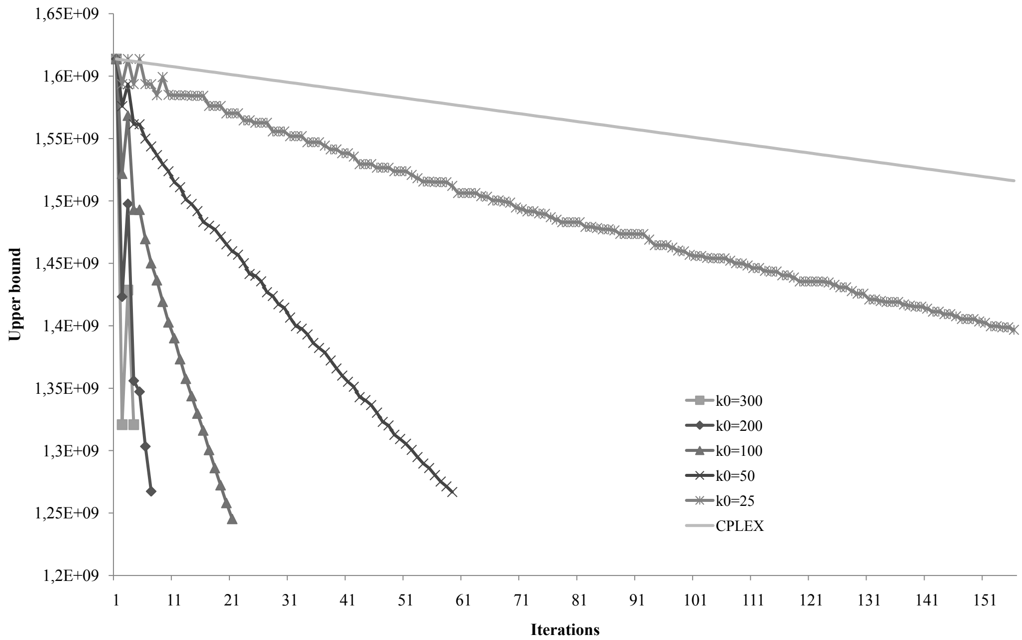

Figure 3.5. Evolution of the soft-fixing local search heuristic for different k_0.

mathematical programming in mining problems, which are usually focused on mine planning and scheduling. We have proposed a mixed-integer programming approach captures relevant physical and logical characteristics observed in real-world, large scale mining operations. In order to deal with the high computational demands, we have developed and evaluated a series of Relax&Fix solution strategies which partition the integer variable space according to reduced costs, geographic relations, time periods and even randomly. Time strategies have shown the best results, especially the TF strategy, in which the subsets are sorted from the first to the last time period. Given the characteristics of the mining supply chain scenario and the behavior of the heuristics, the TF strategy seems to be the fittest to the proposed problem.

We have evaluated the TF and TB strategies for 90 instances with different demand behaviors and increasing planning horizons, always comparing their results against the ones obtained by CPLEX. In general, it may be stated that TF consistently outperforms CPLEX in the larger scenarios, although still not being able to solve the problem to optimality.

Many different Relax&Fix strategies could be envisioned, and many different arrangements between iterations, variable partitioning and computational resources could be developed. Theoretically, the computational resources need not be uniformly divided among the Relax&Fix iterations. Preliminary tests have shown that the first iterations are the ones that consume the most resources. Better gaps on the first iterations should thus improve the overall results. We leave that analysis, as well as the development and evaluation of different strategies, for future work.

It is also important to note that the Relax&Fix solutions, although feasible, can still be far from optimum, as illustrated by the optimality gaps attained. However, the Relax&Fix solutions could still be used as fair primal bounds for other approximation algorithms, such as local search [Bai et al., 2010], or even for the soft-fixing local search procedure depicted here.

Chapter 4

Conclusion

> "All the collected data had come to a final end. Nothing was left to be collected. But all collected data had yet to be completely correlated and put together in all possible relationships.
> A timeless interval was spent in doing that."
>
> *The Last Question, Isaac Asimov*

This Thesis aimed at developing mathematical models and algorithms to address integrated planning problems in a Global Mining Supply Chain. From the first chapter to appendixes, the main contributions were presented and organized around that sole objective. First, we performed a review of mining industry problems and opportunities for the application of mathematical programming approaches. Those problems were then put together under an integrated perspective on strategic, tactical and operational decision levels within the Global Mining Supply Chain concept. Then, we proposed a set of mathematical models and algorithms to address three nontrivial integrated planning problems at each the decision level.

At the strategic level, we have developed a novel multistage stochastic integer programming model to address the capacity planning problem in a Global Mining Supply Chain. The model integrates capacitated facility location and network design decisions with economies of scale on the capacity costs. The formulation is sensitive to both positive and negative variations in demand and allows the establishment of new (or capacity expansion of existing) production facilities and logistics channels, as well as the deactivation of specific facilities and channels in a temporary or permanent basis. In a first study, we have analyzed the characteristics of the problem by means of an empirical evaluation of different CPLEX's settings. Such analysis not only yielded an improved performance of the solver, but also provided pointers to the development of specific algorithms and solution approaches. It is clear that such an approach could also be extended to the other problems discussed in this Thesis. In a second study, we develop a Lagrangian Heuristic as an attempt to improve the dual bounds and derive quality primal bounds for the problem. Results showed significantly better results attained by the heuristic for larger problem instances. We have also proposed a soft-fixing local search approach as an alternative approach to derive improved feasible solutions early in the computation procedure, as demonstrated by numerical experiments. Results also suggested the use of the soft-fixing local search as a refining procedure for the solutions obtained by the Lagrangian Heuristic, which is left to future work.

At the tactical level, we have developed a mixed-integer programming approach to the integrated sales and operations tactical planning problem in a Global Mining Supply Chain. The sales and operations planning problem has been gaining increasing attention in the mining industry, again especially among large-scale organizations. The model has characteristics of a lot sizing problem in a network environment, but presents challenging aspects related to integer flows, discrete production levels and flow conservation violations due to mass losses in processing plants and bulk transportation. A series of Relax&Fix strategies were developed to address realistic sized problem instances and provide good primal bounds for the problem. Additional efforts must be made to further improve those solutions while still maintaining computation time at reasonable figures. The soft-fixing local search approach has also shown improvements on the CPLEX solution, although not nearly as good as the Relax&Fix TF strategy. Again, the use of the soft-fixing local search as a refining procedure for the TF upper bounds is suggested as future work.

At the operational level, we have briefly discussed in the Introduction the idea of a mixed-integer goal programming model to address the integrated short-term programming of iron ore open pits, processing plants, stockyards and shipping operations, which is detailed in Appendix A. The model is inspired in actual conflicts observed between mining, processing plant and stockyard decision makers in a large-scale mining company in Brazil. It is important to emphasize that the actual development, implementation and testing of the model was performed by the graduate student Franklin Assunção Almeida, who was co-advised by the author of this thesis. The main concepts of that work are, nevertheless, an inherent part of the framework established by this Thesis.

One should notice that, as the mathematical models move from operational to strategic decision levels, the more general they become and the easier it gets to extend their application to different industries and scenarios. The integrated sales and operations tactical planning could easily be applied in other process-based industries, such as steel making and chemicals, as well as discrete manufacturing industries, provided the appropriate adaptations were made on the processing restrictions and the material flow characteristics. The stochastic capacity planning could also extend to different industries and even in very different applications, such as telecommunications infrastructure.

It is important to emphasize, however, that we have restricted the integration scope to a single decision level, but spanning through different stages of the Global Mining Supply Chain. Working on the integration of the proposed models among the different decision levels would be a challenging enterprise and would require careful examination of the relationships between decision variables that lie in the interface between each level. Many different approaches could be evaluated, ranging from mere hierarchical [Hax and Meal, 1975] to increasing levels of integration on decision variables [Dauzère-Péres and Lasserre, 1994; Maravelias and Sung, 2009].

One final thought regarding the main contribution of this Thesis. The proposed mathematical models are considerably relevant, both in academic and industry scenarios. They were partly based on a thorough review on the related literature, but mostly inspired by the experience of the author and his advisor in dealing with engineering and technology projects in iron ore mines, railways and ports. We hope the results presented in this text provide a means to further extend the proposed models and algorithms to several possible graduate level studies.

Someday the pain would be gone,
but never the memory.
(The Songs of Distant Earth, Arthur C. Clarke)

Bibliography

Ahmed, S., King, A. J., and Parija, G. (2003). A multi-stage stochastic integer programming approach for capacity expansion under uncertainty. *Journal of Global Optimization*, 26:3–24.

Almeida, F. A. and Pimentel, B. S. (2010). Um modelo matemático para o problema de programação integrada de curto prazo em minas. In *Anais do XLII Simpósio Brasileiro de Pesquisa Operacional*, Bento Golçalves, Brazil. Sociedade Brasileira de Pesquisa Operacional.

Almeida, F. A., Pimentel, B. S., and Mateus, G. R. (2009). Algoritmos para planejamento integrado de produção e transporte de minérios. In *Anais do XLI Simpósio Brasileiro de Pesquisa Operacional*, Porto Seguro, Brazil. Sociedade Brasileira de Pesquisa Operacional.

Azaron, A., Brown, K. N., Tarim, S. A., and Modarres, M. (2008). A multi-objective stochastic programming approach for supply chain design considering risk. *International Journal of Production Economics*, 116:129–138.

Bagajewicz, M. (2008). On the use of net present value in investment capacity planning models. *Industrial & Engineering Chemistry Research*, 47(23):9413–9416.

Bai, R., Kendall, G., and Li, J. (2010). An efficient guided local search approach for service network design problem with asset balancing. In *Proceedings of the International Conference on Logistics Systems and Intelligent Management*, pages 110–115.

Balakrishnan, A. (1984). *Valid inequalities and algorithms for the network design problem.* PhD thesis, Sloan School of Management, MIT.

Beraldi, P., Ghiani, G., Greco, A., and Guerriero, E. (2006). Fix and relax heuristic for a stochastic lot-sizing problem. *Computational Optimization and Applications*, 33:303–318.

Birge, J. R. (1982). The Value of the Stochastic Solution in stochastic linear programs with fixed recourse. *Mathematical Programming*, 24:314–325.

Bradley, S. P., Hax, A. C., and Magnanti, T. L. (1977). *Applied Mathematical Programming*. Addison-Wesley Publishing Company, Reading, Massachusetts.

Danna, E., Rothberg, E., and Pape, C. L. (2005). Exploring relaxation induced neighborhoods to improve MIP solutions. *Mathematical Programming*, 102(1):71–90.

Dauzère-Péres, S. and Lasserre, J.-B. (1994). An integrated approach in production planning and scheduling. *Lecture Notes in Economics and Mathematical Systems*, 411.

de Oliveira, W. L., Sagastizábal, C., Penna, D. D. J., Maceira, M. E. P., and Damázio, J. M. (2010). Optimal scenario tree reduction for stochastic streamflows in power generation planning problems. *Optimization Methods and Software*, 25(6):597–608.

Dias, J., Captivo, M. E., and Clímaco, J. (2006). Capacitated dynamic location problems with opening, closure and reopening of facilities. *IMA Journal of Management Mathematics*, 1(2):317–348.

Dillenberger, C., Escudero, L. F., Wollensak, A., and Zhang, W. (1994). On practical resource allocation for production planning and scheduling with period overlapping setups. *European Journal of Operational Research*, 75:275–286.

Dogan, K. and Goetschalckx, M. (1999). Primal decomposition method for the integrated design of multi-period production distribution systems. *IIE Transactions*, 31:1027–1036.

Dupacová, J. (1999). Portfolio optimization via stochastic programming: Methods of output analysis. *Mathematical Methods of Operations Research*, 50:245–270.

Erlenkotter, D. (1981). A comparative study of approaches to dynamic location problems. *European Journal of Operational Research*, 6:133–143.

Escudero, L. F. and Salmeron, J. (2005). On a fix-and-relax framework for a class of project scheduling problems. *Annals of Operations Research*, 140:163–188.

Ferreira, D., Morabito, R., and Rangel, S. (2009). Solution approaches for the soft drink integrated production lot sizing and scheduling problem. *European Journal of Operational Research*, 196:697–706.

Fischetti, M. and Lodi, A. (2003). Local branching. *Mathematical Programming*, 98(1-3):23–47.

Fisher, M. L. (1981). The lagrangian relaxation method for solving integer programming problems. *Management Science*, 27(1):1–18.

Germain, R., Claycomb, C., and Dröge, C. (2008). Supply chain variability, organizational structure, and performance: The moderating effect of demand unpredictability. *Journal of Operations Management*, 26:557–570.

Guan, Y., Ahmed, S., and Nemhauser, G. L. (2009). Cutting planes for multistage stochastic integer programs. *Operations Research*, 57(2):287–298.

Hax, A. C. and Meal, H. C. (1975). Hierarchical integration of production planning and scheduling. In Geisler, M. A., editor, *Studies in Management Science*, volume 1 of *Logistics*, pages 53–69. Massachusetts Institute of Technology, Cambridge Operations Research Center, Amsterdam.

Heitsch, H. and Römisch, W. (2009). Scenario tree modeling for multistage stochastic programs. *Mathematical Programming: Series A and B*, 118(2):371–406.

Hubbard, D. W. (2010). *How to Measure Anything: Finding the Value of Intangibles in Business*. John Wiley and Sons Inc., Hoboken, New Jersey, 2nd edition.

Ignizio, J. P. (1978). A review of Goal Programming: a tool for multiobjective analysis. *Journal of the Operational Research Society*, 29(11):1109–1119.

Keskin, B. B. and Üster, H. (2007). Meta-heuristic approaches with memory and evolution for a multi-product production/distribution system design problem. *European Journal of Operational Research*, 182(2):663–682.

Kimura, H. (2002). Risk-analysis tools in business strategies. *RAE Electronic FGV*, 1(2):1–14.

Lin, Y.-K. (2009). Two-commodity reliability evaluation of a stochastic-flow network with varying capacity weight in terms of minimal paths. *Computers & Operations Research*, 36:1050–1063.

Lorena, L. A. N. and Senne, E. L. F. (1999). Improving traditional subgradient scheme for lagrangian relaxation: an application to location problems. *International Journal of Mathematical Algorithms*, 1:133–151.

Luce, R. D. and Raiffa, H. (1957). *Games and Decisions.* John Wiley and Sons, Inc., New York.

Luss, H. (1982). Operations research and capacity expansion: A survey. *Operations Research*, 30(5):907–947.

Maravelias, C. T. and Sung, C. (2009). Integration of production planning and scheduling: Overview, challenges and opportunities. *Computers & Chemical Engineering*, 33(12):1919–1930.

Marzano, L. G. B. (2004). *Otimização de Portfólio de Contratos de Energia em Sistemas Hidrotérmicos com Despacho Centralizado.* PhD thesis, Pontifícia Universidade Católica do Rio de Janeiro, Departamento de Engenharia Elétrica.

Melkote, S. and Daskin, M. S. (2001). Capacitated facility location/network design problems. *European Journal of Operational Research*, 129:481–495.

Melo, M. T., Nickel, S., and da Gama, F. S. (2009). Facility location and supply chain management – a review. *European Journal of Operational Research*, 196:401–412.

Parpas, P. and Rustem, B. (2007). Computational assessment of nested benders and augmented lagrangian decomposition for mean-variance multistage stochastic problems. *Informs Journal on Computing*, 19(2):239–247.

Pimentel, B. S., Mateus, G. R., and Almeida, F. A. (2010a). Mathematical models for optimizing the global mining supply chain. In Nag, B., editor, *Intelligent Systems in Operations: Models, Methods and Applications*, pages 133–163. IGI Global, Hershey, Pennsylvania.

Pimentel, B. S., Mateus, G. R., and Almeida, F. A. (2010b). Operations research applied to integrated iron ore production and logistics. In *Operations Research in Mining Seminar*, Santiago, Chile.

Pimentel, B. S., Mateus, G. R., and Almeida, F. A. (2011). Stochastic capacity planning in a global mining supply chain. In *Submitted to IEEE Workshop on Computational Intelligence in Production and Logistics Systems (to appear)*, Paris, France.

Rothberg, E. (2007). An evolutionary algorithm for polishing mixed integer programming solutions. *Informs Journal on Computing*, 19(4):534–541.

Ruszczynski, A. (1997). Decomposition methods in stochastic programming. *Mathematical Programming*, 79:333–353.

Saaty, T. L. (2005). *Theory and Applications of the Analytic Network Process: Decision Making with Benefits, Opportunities, Costs, and Risks*. RWS Publications, 3rd edition.

Santoso, T., Ahmed, S., Goetschalckx, M., and Shapiro, A. (2005). A stochastic programming approach for supply chain network design under uncertainty. *European Journal of Operational Research*, 167:96–115.

Schoemaker, P. J. H. (1995). Scenario planning: A tool for strategic thinking. *Sloan Management Review*, Winter:25–40.

Shapiro, J. E. (2010). Advanced analytics for Sales & Operations Planning. *Analytics Magazine*, May-June:20–26.

Sherali, H. D. and Smith, J. C. (2001). Improving discrete model representations via symmetry considerations. *Management Science*, 47(10):1396–1407.

Shulman, A. (1991). An algorithm for solving dynamic capacitated plant location problems with discrete expansion sizes. *Operations Research*, 39(3):423–436.

Singh, K. J., Philpott, A. B., and Wood, R. K. (2009). Dantzig-wolfe decomposition for solving multistage stochastic capacity-planning problems. *Operations Research*, 57(5):1271–1286.

Tanonkou, G.-A., Benyoucef, L., and Xie, X. (2008). Design of stochastic distribution networks using Lagrangian Relaxation. *IEEE Transactions on Automation Science and Engineering*, 5(4):597–608.

Tsiakis, P. and Papageorgiou, L. G. (2008). Optimal production allocation and distribution supply chain networks. *International Journal of Production Economics*, 111:468–483.

Wagner, M. R., Bhadury, J., and Peng, S. (2009). Risk management in uncapacitated facility location models with random demands. *Computers & Operations Research*, 36:1002–1011.

Wolsey, L. A. (1998). *Integer Programming*. John Wiley & Sons, Inc., New York, first edition.

Appendix A

Integrated short-term mining, processing and shipping operations

> "I am certain there is too much certainty in the world."
>
> *Michael Crichton*

Operational decisions focus on short-term operations programming and sequencing of daily activities of a mining complex. In this Appendix, we propose a mixed-integer goal programming model to address the short-term programming of iron ore open pits, processing plants, stockyards and shipping operations. The model integrates those decision making processes by establishing specific constraints to represent the formal or informal contracts inherent to the interfaces between each department or production unit. We estimate a Value of the Integrated Solution by comparing programs generated by both shared and individualized decisions, thus emphasizing the potential losses due to violations of the interface constraints. The model has specific characteristics which demand specific solution approaches to deal with its inherent complexity.

It is important to emphasize that the work described in this Appendix is the result of the undergraduate monograph by student Franklin Assunção Almeida, who was co-advised by the author of this Thesis. The concepts discussed here are also presented in the following publications:

> Almeida, F. A. and Pimentel, B. S. (2010). Um modelo matemático para o problema de programação integrada de curto prazo em minas. In *Anais do XLII Simpósio Brasileiro de Pesquisa Operacional*, Bento Golçalves, Brazil. Sociedade Brasileira de Pesquisa Operacional.

A.1 Introduction

Iron ore mining operations are known to be extremely cost intensive. Like any commodity, iron ore selling prices are usually constrained by market conditions and hence the pressure over operational performance and productivity can be overwhelming. A typical open pit mining enterprise is usually comprised of an iron ore deposit (the pit, or mine proper), a processing plant, a product stockyard and a train loading system. In large mining companies, the organizational structure often associates each of those stages to specific departments or business units. However, since each department has its own goals (and its own daily operational challenges), operational decisions are commonly made in an isolated, potentially suboptimal fashion [Almeida and Pimentel, 2010].

A.2 Short-Term Mining Operations Planning

At the operational level, the concept of mining supply chain planning and scheduling can be overwhelmingly complex. Short-term planning and scheduling in mines, railroads and ports are complex problems themselves, as it has been shown in previous sections. Detailed plans must be accommodated to specific equipment capacities and time constraints. For instance, the managing of the ship queue depends not only on the availability of berths, but also on the availability of the required ore products, which in turn depends on the delivery schedule of the railroad and, recursively, on the train loading at each mine. Also, stockyard allocation and stockpile formation — the equivalent of warehouse management in manufacturing supply chains — at mines and ports must be performed at a detailed level. End products tonnages must be allocated to specific stockyard areas in order to satisfy loading plans with optimal productivity. Space restrictions and product compatibility must be observed to avoid contamination. Setup costs and time would accrue when a given stockyard area receives a product that is incompatible with the one previously stored. Minimum stock levels must also be determined in order to accommodate uncertain events and, more importantly, to ensure rational utilization of handling equipment. For instance, if many small stockpiles exist, a greater number of different products could be stored, but the reclaimer would have to be repositioned several times in order to fulfill the loading plan. However, the alternative of having large, single-product stockpiles could reduce stockyard flexibility when different ore products are handled. At the ports, stockpiles may be formed according to specific shipments, which may also drive stockpile formation at the mines. At the operational level, blending can be used either to generate composite products or to make last-minute quality corrections, right before loading products into trains or ships. Also, since quality specifications are defined within a certain tolerance range, different ore products may be used to satisfy the same demand. Routing and scheduling decisions, at the mine, at the railway and at the port handling system, must be coordinated with stockyard management in order to satisfy the loading of trucks, trains and ships accordingly. Setup, changeover and variable costs, as well as time and productivity restrictions are the main drivers to operational decisions.

We approach the short-term mining operations planning according to the main stages of the production chain and also according to the interface requirements between those stages.

A.2.1 Self-centered planning

The pit. At the mine proper, a long-term development plan, which establishes the most valuable production schedule over the life time of the iron ore deposit, is assumed to be known. On the short term, however, the mine development plan details which work benches are scheduled for exploitation and how much ore should be removed from each active work bench. The challenge here is to effectively allocate haulage capacity in order to attain maximum productivity while still meeting production requirements. Mass and quality contracts must be then established between the Mine and Processing Plant departments in order to ensure not only minimum quality levels, but also the lowest quality variability possible. Those goals must be achieved both monthly and for each ROM stockpile delivered to the processing plant through the homogenizing stockyard.

The processing plant. On the long term, the processing plant production plan establishes the amount of ore and waste to be generated according to ROM characteristics and equipment availability. Also, the typical split between sinter-feed and pellet-feed coproduction is determined for the planning horizon. On the short term, processing plant operational decisions should aim at achieving the desired production goals and optimizing productivity. However, daily decisions also include (i) performing low-level parameter setups for each production campaign according to ROM characteristics and (ii) scheduling coproduction of specific pellet-feed qualities according to the shipping demand. It is important to notice that, as with any process based industry, the effects of a change in setup at the processing plant may take some time to produce the desired changes in the products delivered. Thus, the number of setups should be kept to a minimum.

The product stockyard and the train loading system. The long-term shipment plan establishes the required products — in both mass and quality — to be delivered by a mine enterprise at each period of the planning horizon. This long-term plan usually considers all connected mine enterprises of a company and is based on a known logistics capacity and availability. At the product stockyard, two important decisions must then be taken in the short term: (i) the allocation of new products to the appropriate stockpiles and (ii) the selection of the correct product qualities and amounts to be reclaimed from the stockpiles and loaded to the trains according to the shipment quality requirements. Alternatively, a train may be loaded with products coming from the stockpiles and directly from the processing plant. The latter option, however appealing to productivity-biased decision makers, must be employed with care, since any quality issues would be left unresolved without the passage through the product stockyard.

A.2.2 Interface contracts and decision integration

Of special interest to this work is the effect of integrating the individual short-term planning decisions to the overall performance of mining enterprise. The following paragraphs detail the main contracts considered in our mathematical model approach.

Mine and Processing Plant. On the short term, although the mine must meet the overall ROM quality requirements, the daily quality variability should be kept to a minimum. For instance, in a short-term planning horizon, a number of work benches would have to be exploited. Since reallocating shovels and trucks to different work benches may have a significant impact on productivity, the smaller the number of reallocations, the better. However, depending on the coproduction scheduling requirements at the processing plant, an analogous scheduling of the active work benches may be performed according to their grade characteristics in order to obtain a tradeoff between mining productivity and ROM specific quality at the homogenizing stockyard.

Processing plant and Product Stockyard. Observing its design restrictions, the processing plant can theoretically generate any product specification given any kind of ROM delivered by the mine. The tradeoff is, of course, the productivity, which is penalized by the amount of waste produced when ROM quality is below average. Thus, given a shipping demand at the product stockyard, the processing plant parameters should be reset to generate products that, combined with the available product stockpiles, would be sufficient to satisfy each train loading demand. However, again, the larger the number of production setups, the worse the processing plant productivity. On the other hand, if the production scheduling at the plant is solely prioritized, the generated products may not exactly match the train loading product quality demand. Again, a tradeoff must be attained.

Train loading system and customer delivery. Since the logistics assets must always be operated at their nominal capacities, a train offered to the loading system must be completely loaded, even if there are no products to provide an exact match of the required shipment qualities. The long-term shipment plan could be adjusted to compensate some part of that quality deviation by reviewing the shipment demand allocated to the company's other mining enterprises. However, in the worst case, each train shipped with an inadequate product quality would incur in contractual penalties charged by the customers. That means that the only way the train loading system can deliver quality products is by having those products available in the product stockyard. A tradeoff between upstream operations' settings and quality product delivery must again be enforced.

A.2.3 A mathematical programming approach

We propose a mixed-integer programming model to address the integrated short-term mining, processing and shipping programming problem [Almeida and Pimentel, 2010]. However, in order to capture the inherent characteristics of this problem, we provide a Goal Programming approach [Ignizio, 1978] to balance the (conflicting) objectives of each business unit within the mining complex. That means that each of the process depicted above must be given a penalty parameter which accounts for the suboptimal effects produced when the interface contracts are taken for granted.

In general terms, the model aims at **minimizing all penalties related to deviation from mass and quality goals in mining, processing and shipping operations, as well as the scheduling cost at the processing plant**, subject to:

- the constraints imposed by the long-term mine production plan, which determines the workbenches to exploit at each time period and how much ROM to remove from them, as well as the operational costs related to the workings of drilling, loading and hauling equipment;

- the mass and quality goals for the homogenizing stockpiles, which also relate to the mass productivity attained at the processing plant;

- the need to establish the best possible production schedule at the processing plant, given that different products may be demanded at the product stockyard and that there may already be quantities of those products previously stored and readily available for shipping;

- the mass and quality contracts with the iron ore customers, which drive the overall balance of produced and shipped ore — with the required quality specifications — available at the product stockyard.

We detail the proposed mathematical formulation for this problem as follows.

Sets

M : Mines within a mining complex;

F^m : Mining workbenches at Mine $m \in M$;

H^m : ROM stockpiles at homogenizing yard of Mine $m \in M$;

P_f^m : Fine ore products generated at Mine's $m \in M$ processing plant;

P_f : Fine ore products generated at the mine complex's processing plants;

P_s^m : Super fine ore products generated at Mine's $m \in M$ processing plants;

P_s : Super fine ore products generated at the mine complex's processing plants;

P^m : All products generated at Mine's $m \in M$ processing plants, $P^m = P_f^m \cup P_s^m$;

P : All products generated at the mine complex's processing plants, $P = P_f \cup P_s$;

K : Quality parameters evaluated in each ore product.

Parameters

T : Time periods in the programming horizon;

O_i : ROM available at workbench $i \in F^m$ of Mine $m \in M$;

L^m : Number of workbenches of mine $m \in M$ that can be simultaneously exploited in each time period;

α_i^- : Penalty for unexploited ROM mass from workbench $i \in F^m$ at mine $m \in M$;

α_i^+ : Penalty for ROM mass exploited from workbench $i \in F^m$ at mine $m \in M$ beyond the amount established by the long-term mining scheduler;

ω_i : ROM mass planned for mining at workbench $i \in F^m$ of mine $m \in M$, as established by the long-term mining scheduler;

π_i : Cost associated to moving drilling and loading equipment to a different workbench $i \in F^m$;

ρ_{ik} : Quality goal for (chemical or physical) parameter $k \in K$ of ROM exploited from workbench $i \in F^m$ of mine $m \in M$;

Q_j^m : Mass goal for each stockpile in homogenizing stockyard H^m of mine $m \in M$;

γ_+^m : Penalty for positive deviation from the mass goal Q_j^m of stockpile j in the homogenizing yard of $m \in M$;

γ_-^m : Penalty for negative deviation from the mass goal Q_j^m of stockpile j in the homogenizing yard of $m \in M$;

δ_{pk}^m : Penalty for positive deviation from the quality goal of parameter $k \in K$ when generating product $p \in P^m$ from each stockpile in the homogenizing stockyard H^m of mine $m \in M$;

η_{pk}^m : Penalty for negative deviation from the quality goal of parameter $k \in K$ when generating product $p \in P^m$ from each stockpile in the homogenizing stockyard H^m of mine $m \in M$;

C^m : Maximum conveying capacity from the homogenizing stockyards and the processing plants of mine $m \in M$;

β^m : Typical mass efficiency of the processing plant of mine $m \in M$;

ϕ_{kp}^m : Mass loss due to every unit of positive deviation from quality parameter $k \in K$ when processing product $p \in P^m$ at the processing plant of mine $m \in M$;

φ_{kp}^m : Mass loss due to every unit of negative deviation from quality parameter $k \in K$ when processing product $p \in P^m$ at the processing plant of mine $m \in M$;

W^m : Maximum production capacity of the processing plant of mine $m \in M$;

m_{pk} : Quality goal for parameter $k \in K$ in product $p \in P$;

D_p^t : Shipping demand for product $p \in P$ at period t;

h_{pb} : Penalty associated to every mass unit of product $b \in P$ used to alternatively meet the demand for product $p \in P$;

n_p^m : Cost associated to the setup of processing plant of mine $m \in M$ when producing $p \in P$;

U_p^1 : Initial stock of product $p \in P$ at the product stockyard;

U_p^T : Minimum final stock of product $p \in P$ at the product stockyard;

U_p : Capacity of product $p \in P$ stockpiles at the product stockyard;

U : Capacity of the product stockyard;

Decision variables

l_i^t : Whether workbench $i \in F^m$ is being exploited at period t, mine $m \in M$;

r_i^t : Whether workbench $i \in F^m$, which was inactive at period $t-1$ is now being exploited at period t, mine $m \in M$;

x_{ij}^t : ROM mass exploited from workbench $i \in F^m$ and used in stockpile $j \in H^m$ at mine $m \in M$, period t;

pl_i^- : Unexploited ROM mass from workbench $i \in F^m$ at mine $m \in M$ throughout the programming horizon;

pl_i^+ : ROM mass exploited beyond from workbench $i \in F^m$ the amount established by the long-term mining scheduler at mine $m \in M$ throughout the programming horizon;

f_j^t : Whether stockpile $j \in H^m$ is formed in the homogenizing stockyard at period t;

c_j^t : Whether stockpile $j \in H^m$ is consumed from the homogenizing stockyard at period t;

dh_{j+}^t : Positive deviation from the mass goal Q_m of stockpile $j \in H^m$ at period t;

dh_{j-}^t : Negative deviation from the mass goal Q_m of stockpile $j \in H^m$ at period t;

dq_{jpk+} : Positive deviation from the quality goal of parameter $k \in K$ for product $p \in P$ generated from stockpile $j \in H^m$, which was formed at period f and consumed at period t;

dq_{jpk-} : Negative deviation from the quality goal of parameter $k \in K$ for product $p \in P$ generated from stockpile $j \in H^m$, which was formed at period f and consumed at period t;

w_{jp}^{fc} : Whether stockpile $j \in H^m$, which was formed at period f in the homogenizing stockyard of mine $minM$, was consumed at period c when generating product $p \in P^m$;

g_p^{mt} : Whether product $p \in P^m$ was generated by the processing plant of mine $m \in M$ at period t;

y_{jp}^{fc} : ROM mass taken from stockpile $j \in H^m$, which as formed at period f and consumed at period c, when generating product $p \in P$;

z_p^{mt} : Mass of product $p \in P$ generated in the processing plant of mine $m \in M$ at period t;

v_p^{mt} : Whether product $p \in P$, which was not being produced in the processing plant of mine $m \in M$ at period $t-1$, is now being generated at period t;

s_{pb}^t : Mass of product $b \in P$, stored in the product stockyard, which was alternatively used to meet the demand of product $p \in P$ at period t;

u_p^t : Stock of product $p \in P$ in the product stockyard at period t.

The objective function seeks minimizing all penalties related to deviation from mass and quality goals in mining, processing and shipping operations, as well as the scheduling cost at the processing plant:

Minimize

$$\sum_{m\in M}\sum_{i\in F^m}\left(\alpha_i^+ pl_i^+ + \alpha_i^- pl_i^-\right) + \sum_{t=1}^{T-1}\sum_{m\in M}\sum_{j\in H^m}\left(\gamma_+^m dh_{j+}^t + \gamma_-^m dh_{j-}^t\right) +$$
$$\sum_{t=1}^{T-1}\sum_{m\in M}\sum_{i\in F^m}\pi_i r_i^t +$$
$$\sum_{t\in 1}^{T-1}\sum_{c=t+1}^{T}\sum_{m\in M}\sum_{j\in H^m}\sum_{p\in P^m}\sum_{k\in K}\left(\delta_{pk}^m dq_{jpk+}^{tc} + \eta_{pk}^m dq_{jpk-}^{tc}\right) +$$
$$\sum_{t=3}^{T}\sum_{m\in M}\sum_{p\in P}\left(n_p v_p^{mt}\right) + \sum_{t=2}^{T}\sum_{p\in P}\sum_{b\in P}\left(h_{pb}s_{pb}^t\right),$$

where several constraints apply, as described in the following paragraphs.

According to a long-term mine production plan, which is assumed to be an input to this model, all ROM available in each workbench should be completely exploited within a given time period:

$$\sum_{t=1}^{T-1}\sum_{j\in H^m} x_{ij}^t + pl_i^- - pl_i^+ = O_i \quad \forall i \in F^m, m \in M. \tag{A.1}$$

However, depending on the short-term decisions, exploitation efforts may be required to switch to a different workbench:

$$r_i^t - l_i^t + l_i^{t-1} \geq 0 \quad \forall i \in F^m, m \in M, t \in 1, 2...T-1. \tag{A.2}$$

Also, the workbenches may not all be simultaneously exploited:

$$\sum_{i \in F^m} l_i^t \leq L^m \quad \forall m \in M, t \in 1, 2...T-1, \tag{A.3}$$

but exploitation must be restricted to active workbenches, that is, workbenches in which there were allocated capacitated drilling and loading equipment:

$$\sum_{j \in H^m} x_{ij}^t \leq O_i l_i^t \quad \forall i \in F^m, m \in M, t \in 1, 2...T-1, \tag{A.4}$$

$$0 \leq x_{ij}^t \leq \omega_i, \quad \forall i \in F_m, j \in H_m, m \in M, t \in T, \tag{A.5}$$

assuming that at least one workbench should be used to form a ROM stockpile at the homogenizing stockyard:

$$\sum_{j \in H^m} x_{ij}^t > l_i^t - 1 \quad \forall i \in F^m, m \in M, t \in 1, 2...T-1. \tag{A.6}$$

We then assume that a ROM stockpile cannot be formed and consumed at the same time period:

$$c_j^t + f_j^t \leq 1, \quad \forall j \in H^m, m \in M, t = 2, 3...T-1, \tag{A.7}$$

and establish a mass goal for each stockpile to be formed:

$$\sum_{i \in F_m} x_{ij}^t - dh_{j+}^t + dh_{j-}^t = Q_m f_j^t, \quad \forall j \in H^m, m \in M, t \in 1, 2...T-1, \tag{A.8}$$

according to the capacity of the stockyard and the productivity of the available stockpiling equipment. However, a given stockpile can only be reclaimed after it has been formed:

$$\sum_{d=2}^{t} c_j^d \leq \sum_{d=1}^{t-1} f_j^d, \quad \forall j \in H^m, \forall m \in M, \forall t \in 2, 3, ...T \tag{A.9}$$

and it can only be formed after the previous stockpile, which occupied the same position in the stockyard, has been reclaimed:

$$\sum_{d=1}^{t} f_j^d \leq \sum_{d=2}^{t} c_j^d + 1, \quad \forall j \in H_m, m \in M, t \in 2, ...T-1, \tag{A.10}$$

The processing plant yields higher productivity if it is fed with ROM which possesses chemical and physical characteristics that are similar to those of the final ore products. Hence, we must establish goals for each quality parameter considered:

$$\sum_{i \in F^m} \rho_{ik} x_{ij}^t - \sum_{c=t+1}^{T} \sum_{p \in P^m} dq_{jpk+}^{tc} + \sum_{c=t+1}^{T} \sum_{p \in P^m} dq_{jpk-}^{tc} = \sum_{c=t+1}^{T} \sum_{p \in P^m} m_{pk} y_{jp}^{tc}$$
$$\forall k \in K, j \in H^m, m \in M, t \in 1, 2...T-1, \tag{A.11}$$

assuming that the deviation variables can only be positive if the corresponding stockpiles have been fed (and previously formed) into the processing plant:

$$dq_{jpk+}^{tc} \leq C^m w_{jp}^{tc}, \quad \forall k \in K, p \in P^m, j \in H^m, t \in 1, 2...T-1, \forall c = t+1, ...T, \tag{A.12}$$
$$dq_{jpk-}^{tc} \leq C^m w_{jp}^{tc}, \quad \forall k \in K, p \in P^m, j \in H^m, t \in 1, 2...T-1. \forall c = t+1, ...T. \tag{A.13}$$

The destination of the ROM stockpiles must be set according to reclamation, formation and to the product generated at the processing plant. This provides the traceability required when assessing the performance of the integrated production plans:

$$\sum_{c=t+1}^{T} w_{jp}^{tc} \leq f_j^t, \quad \forall p \in P^m, j \in H^m, m \in M, t \in 1, 2...T-1, \tag{A.14}$$
$$\sum_{f=1}^{t-1} w_{jp}^{ft} \leq c_j^t, \quad \forall p \in P^m, j \in H^m, m \in M, t \in 2, 3...T, \tag{A.15}$$
$$\sum_{f=1}^{t-1} w_{jp}^{ft} \leq g_p^t, \quad \forall p \in P^m, j \in H^m, m \in M, t \in 2, 3...T. \tag{A.16}$$

Also, there can only be an output of ROM from a stockpile if it has been formed and is to be reclaimed at a given time period:

$$y_{jp}^{fc} \leq C^m w_{jp}^{fc}, \quad \forall j \in H^m, m \in M, f \in 1, 2...T-1, c = f+1, ...T, \tag{A.17}$$

according to the balance of mass of ROM stockpiles:

$$\sum_{i \in F} x_{ij}^{t} - \sum_{c=t+1}^{T} \sum_{p \in P_f^m} y_{jp}^{tc} = 0 \quad \forall j \in H^m, m \in M, t \in 1, 2...T-1, \tag{A.18}$$

$$\sum_{i \in F} x_{ij}^{t} - \sum_{c=t+1}^{T} \sum_{p \in P_s^m} y_{jp}^{tc} = 0 \quad \forall j \in H^m, m \in M, t \in 1, 2...T-1. \tag{A.19}$$

and the balance of mass which occurs in the processing plant:

$$\beta^m \sum_{f=1}^{t-1} \sum_{j \in H^m} y_{jp}^{ft} - z_p^t - \sum_{f=1}^{t-1} \sum_{j \in H^m} \sum_{k \in K} \phi_{pk}^m dq_{jpk+}^{ft} - \sum_{f=1}^{t-1} \sum_{j \in H^m} \sum_{k \in K} \varphi_{pk}^m dq_{jpk-}^{ft} = 0, \quad \forall p \in P_f^m, m \in M, t \in 1, 2...T-1, \tag{A.20}$$

$$(1-\beta^m) \sum_{f=1}^{t-1} \sum_{j \in H^m} y_{jp}^{ft} - z_p^{mt} - \sum_{f=1}^{t-1} \sum_{j \in H^m} \sum_{k \in K} \phi_{pk}^m dq_{jpk+}^{ft} - \sum_{f=1}^{t-1} \sum_{j \in H^m} \sum_{k \in K} \varphi_{pk}^m dq_{jpk-}^{ft} = 0, \quad \forall p \in P_f^m, m \in M, t \in 1, 2...T-1. \tag{A.21}$$

which accounts for the inherent process mass efficiency. It is important to notice that only one product for each family (fines and superfines) may be generated at each time period:

$$\sum_{p \in P_f^m} g_p^{mt} = 1, \quad \forall m \in M, t \in 2, 3...T, \tag{A.22}$$

$$\sum_{p \in P_s^m} g_p^{mt} = 1, \quad \forall m \in M, t \in 2, 3...T, \tag{A.23}$$

and that all production must be limited by the overall capacity of the processing plant:

$$\sum_{p \in P^m} z_p^{mt} \le W^m \quad \forall m \in M, t \in 1, 2...T-1. \tag{A.24}$$

Regarding production scheduling, one must account for every processing plant setup required:

$$v_p^{mt} - g_p^{mt} + g_p^{mt-1} \ge 0, \quad \forall p \in P, t \in 3, 4...T-1 \tag{A.25}$$

And at the product stockyard, one must observe the overall balance of produced and shipped ore:

$$\sum_{m \in M} z_p^{mt} - \sum_{b \in P} s_{bp}^t = u_p^t - u_p^{t-1} \quad \forall p \in P_f^m, t \in T, \tag{A.26}$$

as well as the storage capacity for each product:

$$0 \leq u_p^t \leq U_p, \quad \forall p \in P, t \in T, \tag{A.27}$$

$$0 \leq \sum_{p \in P} u_p^t \leq U, \quad \forall p \in P, t \in T u_p^t \geq U_p^T \quad \forall p \in P, t = T, \tag{A.28}$$

and the shipping demand, which must always be met in terms of mass. That means that every train must leave the mine complex completely loaded, even if the exact required products are not available:

$$\sum_{b \in P} s_{pb}^t = D_p^t, \quad \forall p \in P, t \in T. \tag{A.29}$$

Finally, integrality and nonnegativity constraints apply as follows:

$$pl_i \geq 0, \quad \forall i \in F_m, m \in M, \tag{A.30}$$

$$c_j^t, f_j^t \in \{0,1\}, \quad \forall j \in H^m, m \in M, t \in T, \tag{A.31}$$

$$dh_{j+}^t, dh_{j-}^t \geq 0, \quad \forall j \in H^m, m \in M, t \in T, \tag{A.32}$$

$$y_{jp}^{fc} \geq 0, \quad \forall j \in H, f \in 0,1,...,T, c \in f+1,...,T, \tag{A.33}$$

$$g_p^{mt}, v_p^{mt} \in \{0,1\}, \quad \forall p \in P, t \in T, \tag{A.34}$$

$$w_{pj}^{fc} \in \{0,1\}, \quad \forall p \in P, f \in 0,1,...,T, c \in f+1,...,T, j \in H, \tag{A.35}$$

$$dqf_{jpk+}^{tc}, dqf_{jpk-}^{tc} \geq 0, \quad \forall k \in K, t, c \in T, \tag{A.36}$$

$$s_{pb}^t \geq 0, \quad \forall p \in P, b \in P, t \in T. \tag{A.37}$$

This is a non preemptive goal programming model, since all goals are assumed to be of roughly comparable importance and, being as such, must be addressed simultaneously. It is important to notice that the behavior of such a model can be rather sensitive to the weights and penalties associated to the objectives. In order to coherently prioritize each goal, both weights and penalties should reflect the same species (e.g. cost, profit, or productivity). That may not be a trivial task since, in practice, not all information needed may be available, nor may it be easily converted to different bases. We assume, however, that this is the case for the model developed here. Assessing and managing individual operational costs has been a major concern in modern management and several tools, both conceptual and information technology based, have been proposed to address this issue [Hubbard, 2010]. The Analytic Hierarchy Process (AHP) [Saaty, 2005] could also be used as an alternative approach to determine a coherent prioritization of multiple decision criteria. AHP allows the construction of a robust structure of tangible and intangible criteria and, moreover, makes it possible to convert those criteria to numerical values that can be processed and compared over the entire range of the problem.

The remainder of this Appendix is dedicated to discussing the value of the integrated solution delivered by the model.

A.3 The Value of the Integrated Solution

Large mining operations usually present a rather partitioned organizational structure. In addition, modern management philosophies are based on the establishment of high level goals for the whole company which are broken down into several detailed goals for each business unit, department and, ultimately, employee. That means that all individual goals must be carefully designed and integrated in such a way that some (simple) aggregation function would converge them towards the company's main objective. Although simple in principle, such philosophy can be overwhelmingly complex to implement and manage in practice, especially at the lower decision levels where professional relationships are constantly challenged by daily problems, difficult negotiations and compromises.

In an attempt to provide a quantitative analysis to the integration philosophy, we propose the concept of Value of the Integrated Solution (VIS), which can be defined as the potential benefit from solving the integrated program over solving a set of independent, smaller programs which are probably suboptimal when decisions are actually executed. The approach can be generally described as follows:

1. Solve each subprogram individually, disregarding the terms corresponding to the interface contracts, but still considering all physical constraints;
2. Determine the objective function of the original problem — the integrated approach — using the variables computed for each subproblem with all interface contracts;
3. Solve the integrated problem and determine its objective function
4. The Value of the Integrated Solution is computed from the difference between items (2) and (3) above.

In order to illustrate the VIS concept, we present a simple discussion. Suppose a mining complex composed of one open pit, three workbenches, one homogenizing stockyard that can accommodate four ROM stockpiles, one processing plant and one product stockyard. The processing plant produces four different products (two of type fine and two of type superfine). We assume that only one quality parameter (%Fe) is controlled during a five-period programming horizon. Figure A.1 illustrates the scenario. A detailed specification of the parameters is given in Table A.1. We then

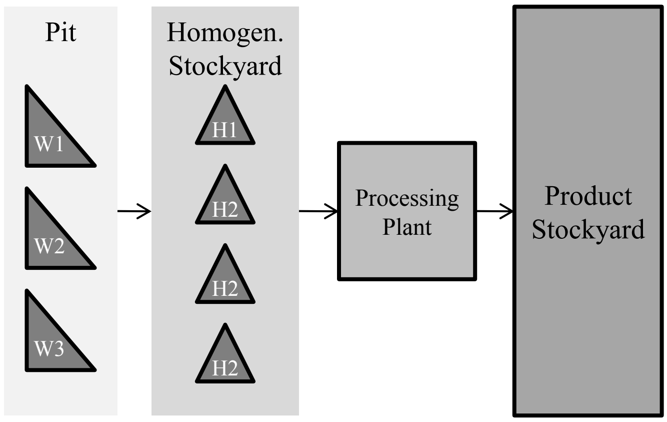

Figure A.1. A simple example for the integrated short-term programming problem.

analyze three different programming situations:

- **Integrated:** All operations work in an integrated fashion, as established by the proposed model.
- **Shipping priority:** Shipping operations will get, whenever possible, the exact product they need, at the cost of increased cost in mining and processing operations.
- **Mining and processing priority:** Mine, homogenizing yard and processing plant operate at optimum levels — some better than the those of the integrated solution —, while shipping quality violations increase.

CPLEX takes only a few seconds to solve this example model (for an analysis of computational results of larger problem instances, please refer to [Almeida and Pimentel, 2010]). Results of executing the integrated short-term programming model on the above scenario derive the VIS results shown in Figure A.2 below. Each of the columns refer to the respective planning situation depicted above. The height of each column represents the objective function as a summation of specific terms: costs associated to

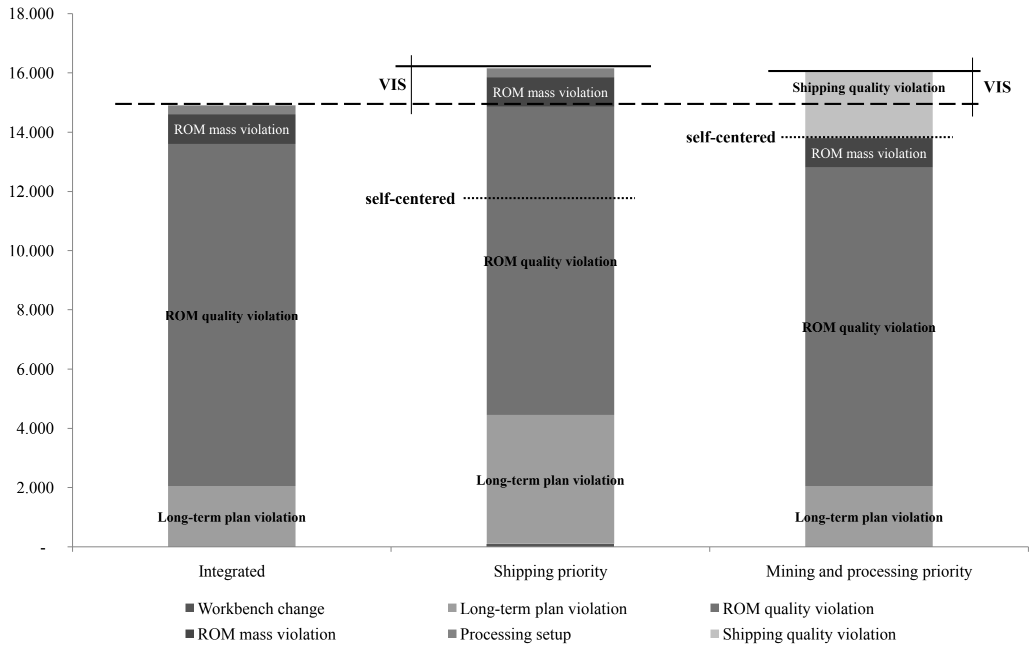

Figure A.2. The Value of the Integrated Solution.

changing the operational workbench — that is, moving the drilling and loading equipment elsewhere in the mine —, penalties associated to violating the long-term mining plan — that is, exploiting more or less than the planned amount at each workbench —, penalties associated to violating the ROM quality and mass contracts between mine and processing plant —, costs associated to setting up the processing plant for every change of product, and penalties associated to violating product quality contracts when shipping it to customers, as illustrated by the chart legend.

When priority is given to shipping operations, mining and processing must operate in a way as to minimize shipping quality deviations, which means that mining equipment should be positioned at the workbenches which would yield ROM quality that approximates the closest the specification of processed products, even if that means violating the long-term mining plan. On the other hand, when priority is given to mining, homogenizing and processing operations, there would be minimum violations on the long-term mining plan, as well as the smallest number possible of plant setups, so that the product stockyard could receive different products than those demanded for shipping.

One can see that all self-centered approaches provide worse results than the integrated solution. All self-centered programs objective functions are calculated using step 2 above and, for each, it is possible to calculate the VIS as the difference regarding the integrated solution objective function. Furthermore, we show for each self-centered program, as a dotted line cutting each column, the value of the objective function as perceived by the decision makers when working individually at each situation considered — calculated using step 1 above.

Table A.1: Dataset for the Integrated Short-Term Mining, Processing and Shipping Programming Problem.

Parameter	**Value**
M	$\{1\}$
F^1	$\{l1, l2, l3\}$
H^1	$\{h1, h2, h3, h4\}$
P_f^1	$\{SF1, SF2\}$
P_s^1	$\{PF1, PF2\}$
K	{%Fe}
T	5
O_1; O_2; O_3	1,150
L^1	2
α_{l1}^-; α_{l2}^-; α_{l3}^-	1.0
α_{l1}^+; α_{l2}^+; α_{l3}^+	1.0
ω_{l1}; ω_{l2}; ω_{l3}	700
π_{l1}; π_{l2}; π_{l3}	100
$\rho_{l1,Fe}$; $\rho_{l2,Fe}$; $\rho_{l3,Fe}$	0.50; 0.40; 0.45
Q_{h1}^1; Q_{h2}^1; Q_{h3}^1; Q_{h4}^1	600
γ_+^1	1.0
γ_-^1	1.0
$\delta_{SF1,Fe}^1$; $\delta_{SF2,Fe}^1$; $\delta_{PF1,Fe}^1$; $\delta_{PF2,Fe}^1$	10
$\eta_{SF1,Fe}^1$; $\delta_{SF2,Fe}^1$; $\delta_{PF1,Fe}^1$; $\delta_{PF2,Fe}^1$	10
C^1	700
β^1	0.8
$\phi_{Fe,SF1}^1$; $\phi_{Fe,SF2}^1$; $\phi_{Fe,PF1}^1$; $\phi_{Fe,PF2}^1$	0.001
$\varphi_{Fe,SF1}^1$; $\varphi_{Fe,SF2}^1$; $\varphi_{Fe,PF1}^1$; $\varphi_{Fe,PF2}^1$	0.001
W^1	1000
$m_{SF1,Fe}$; $m_{SF2,Fe}$; $m_{PF1,Fe}$; $m_{PF2,Fe}$	65,0%; 70,0%; 62,5%; 60.0%
D_{SF1}^2; D_{SF1}^3; D_{SF1}^4; D_{SF1}^5; D_{SF1}^5; D_{SF1}^6; D_{SF1}^7	0; 0; 640; 0; 0; 640
D_{SF2}^2; D_{SF2}^3; D_{SF2}^4; D_{SF2}^5; D_{SF2}^5; D_{SF2}^6; D_{SF2}^7	160; 160; 160; 160; 160; 0
D_{PF1}^2; D_{PF1}^3; D_{PF1}^4; D_{PF1}^5; D_{PF1}^5; D_{PF1}^6; D_{PF1}^7	640; 0; 0; 640; 0; 0
D_{PF2}^2; D_{PF2}^3; D_{PF2}^4; D_{PF2}^5; D_{PF2}^5; D_{PF2}^6; D_{PF2}^7	0; 640; 0; 0; 640; 0
$h_{SF1,PF1}$; $h_{SF1,PF2}$; $h_{SF1,SF2}$	1.0
$h_{SF2,PF1}$; $h_{SF2,PF2}$; $h_{SF2,SF1}$	1.0

Table A.1: Dataset for the Integrated Short-Term Mining, Processing and Shipping Programming Problem.

Parameter	Value
$h_{PF1,SF1}$; $h_{PF1,SF2}$; $h_{PF1,PF2}$	1.0
$h_{PF2,PF1}$; $h_{PF2,SF1}$; $h_{PF2,SF2}$	1.0
n^1_{SF1}; n^1_{SF2}; n^1_{PF1}; n^1_{PF2}	100
U^1_{SF1}; U^1_{SF2}; U^1_{PF1}; U^1_{PF2}	0.0
U^5_{SF1}; U^5_{SF2}; U^5_{PF1}; U^5_{PF2}	0.0
U_{SF1}; U_{SF2}; U_{PF1}; U_{PF2}	1,500
U	1,500

A.4 Discussion and future work

Short-term mining operations are rather complex and an integrated decision framework can provide significant gains, as shown by the Value of the Integrated Solution. It is clear that the goal programming approach is rather sensitive to the magnitudes of the penalty parameters. Hence, providing a sensitivity analysis should be a primary concern. One important discussion, however, is whenever the VIS is always positive for any given positive set of penalty parameters. A formal mathematical verification should be performed and a set of properties of the VIS could be derived for different problem characteristics, as in [Birge, 1982].

The VIS should also be weighed against the computational effort required to determine the integrated solution. Preliminary results have indicated that the integrated problem may be significantly more difficult to solve than the individual subproblems. Also, the solution performance should be evaluated for longer programming horizons (up to 21, for instance, so as to represent a full week of work with three shifts per day). The limit of multipurpose solvers as CPLEX could then be determined and a thorough analysis of the problem characteristics and the solution logs could point out opportunities for developing specific heuristics and decomposition approaches.

The problem could also be addressed as if only one of the objects listed in Section A.2.3 were considered in the objective function, leaving all others to be modeled as constraints. That would require specific lower and upper bounds to be defined for each new constraint. An analysis of the advantages of such an approach to the goal-programming one proposed here are yet to be studied.

Another interesting approach would be to reformulate the model as a multi-objective problem. Pareto-optimal analyses could be performed and different solution frameworks could be developed.

Appendix B

Dataset for the Integrated Sales and Operations Tactical Planning Problem

"It was idle to speculate, to build pyramids of surmise on a foundation of ignorance."

The City and the Stars, Arthur C. Clarke

The data presented below is synthesized from the general characteristics of an actual, massive Brazilian mining company. Since not all elements are described for each facility, the figures represent the most restrictive capacities of each element type[1]. Although system costs and capacities are somewhat equivalent, there is clearly a trend of lower operational unit costs in system C. Fixed costs are relatively higher for offshore operations. Note that only the fixed costs for active nodes are shown; the corresponding fixed costs for idle nodes are, in average, a 10 percent fraction of the active ones. The penalty on unmet demand is set to relatively high values ($b = 2,000$) in order to ensure that only solutions which which fully meet the overall customers' demand are attained (of course, assuming production and distribution capacities are high enough). In practice, however, b values are around the same order of magnitude of ore prices, so that failing to meet all demand may be, at times, a better solution.

[1]Input files and source codes are available from the authors.

Table B.1: Dataset for the integrated S&OP tactical planning problem scenario.

System	Facility	Processing	Storage	Transfer	Unit Cost	Fixed Cost
A	Mine 1	2100	2500	2200	1.17	71.13
	Mine 2	1800	300	1800	1.57	59.98
	Mine 3	400	500	600	1.21	19.66
	Mine 4	2800	1500	2700	1.39	87.87
	Mine 5	1400	200	1300	1.91	42.19
	Mine 6	600	300	700	1.37	22.99
	Mine 7	1200	2000	1400	1.54	45,47
	Rail A	-	-	12000	3.00	393.87
	Port A	2500	7000	10000	1.12	326.01
B	Mine 1	1300	700	1000	1.68	34.96
	Mine 2	100	200	100	1.27	3.50
	Mine 3	-	1500	800	1.85	27.86
	Mine 4	1500	2500	1900	1.25	66.54
	Mine 5	400	1000	400	2.14	14.06
	Mine 6	1200	4000	1300	1.43	44.87
	Mine 7	300	1500	300	1.89	10.40
	Mine 8	700	1200	900	2.07	31.65
	Mine 9	1200	5000	1000	1.99	34.66
	Station 1	-	500	500	0.81	7.12
	Station 2	-	200	1300	1.02	46.13
	Station 3	-	250	3000	1.13	105.23
	Station 4	-	300	1000	0.94	35.43
	Station 5	-	400	1200	2.01	42.59
	Rail B	-	-	7000	2.70	244.32
	Port B1	-	3500	3500	1.21	158.43
	Port B2	-	2000	1500	1.31	88.87
C	Mine 1	9000	8500	8000	0.88	216.34
	Station 1	1000	1000	1000	0.91	11.96
	Station 2	500	500	500	0.94	9.98
	Station 3	250	250	250	1.01	7.39
	Rail C	-	-	8000	2.60	381.22
	Port C	400	4000	7500	1.03	373.82
Offshore	Station 1	-	1500	1500	2.02	115.67
	Station 2	-	1200	1200	2.29	149.21
	Station 3	-	1000	1000	2.25	98.71
	Station 4	-	2000	2000	2.93	198.74

Appendix C

Complete listings of numerical results for the Integrated Tactical Sales and Operations Problem

"Against stupidity the gods themselves contend in vain."

Friedrich von Schiller

Here we present the complete listings of the results obtained with the execution of CPLEX and the TF and TB Relax&Fix strategies over the 90 problem instances covered in Chapter 3. Different behaviors of the iron ore demand are employed to evaluate the response of the heuristic. For each category of demand, we present a graph illustrating the corresponding behavior, followed by a set of tables, one for each value of $T = \{3, 6, 12, 24, 48\}$.

C.1 Tests with increasing demand

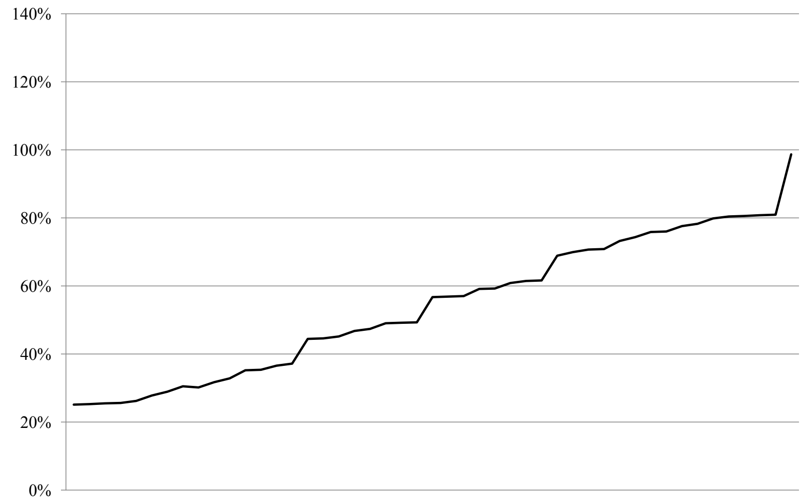

Figure C.1. Increasing demand behavior considered in the tests.

Table C.1. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is increased from 25% to 100% of the system's overall capacity. Planning horizon: 3 time periods.

T=3; Increasing demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	681,453	806,197	15.47%	4
	TB	682,460	794,402	14.09%	2
	CPLEX	755,172	770,275	1.96%	3,600
FC=10	TF	815,404	1,009,245	19.21%	3
	TB	820,175	1,042,508	21.33%	2
	CPLEX	939,750	969,707	3.09%	3,600

Table C.2. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is increased from 25% to 100% of the system's overall capacity. Planning horizon: 6 time periods.

T=6; Increasing demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	1,301,536	1,526,310	14.73%	5
	TB	1,305,236	1,551,747	15.89%	2
	CPLEX	1,421,767	1,506,649	5.63%	3,600
FC=10	TF	1,550,789	1,903,421	18.53%	4
	TB	—	—	—	—
	CPLEX	1,764,967	1,842,925	4.23%	3,600

Table C.3. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is increased from 25% to 100% of the system's overall capacity. Planning horizon: 12 time periods.

T=12; Increasing demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	2,481,508	2,831,818	12.37%	14
	TB	—	—	—	—
	CPLEX	2,666,549	3,046,498	12.47%	3,600
FC=10	TF	2,955,628	3,643,351	18.88%	21
	TB	—	—	—	—
	CPLEX	3,315,800	3,594,487	7.75%	3,600

Table C.4. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is increased from 25% to 100% of the system's overall capacity. Planning horizon: 24 time periods.

T=24; Increasing demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	4,857,210	5,651,259	14.05%	72
	TB	—	—	—	—
	CPLEX	5,224,749	8,323,775	37.23%	3,600
FC=10	TF	5,779,754	7,150,758	19.17%	84
	TB	—	—	—	—
	CPLEX	6,495,617	528,869,188	98.77%	3,600

Table C.5. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is increased from 25% to 100% of the system's overall capacity. Planning horizon: 48 time periods.

T=48; Increasing demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	9,621,579	11,564,394	16.80%	341
	TB	—	—	—	—
	CPLEX	10,345,086	1,059,954,339	99.02%	3,600
FC=10	TF	11,441,503	14,799,148	22.69%	402
	TB	—	—	—	—
	CPLEX	12,852,218	1,062,247,945	98.79%	3,600

C.2 Tests with decreasing demand

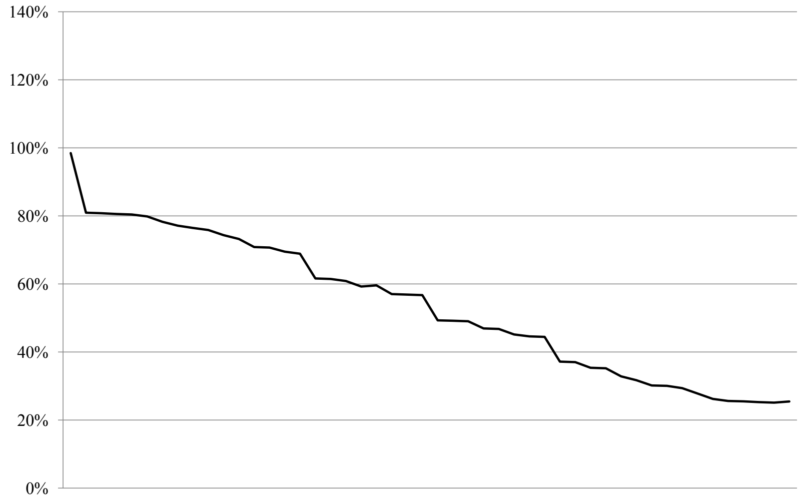

Figure C.2. Decreasing demand behavior considered in the tests.

Table C.6. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is decreased from 100% to 25% of the system's overall capacity. Planning horizon: 3 time periods.

T=3; Decreasing demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	960,961	1,065,732	9.83%	4
	TB	—	—	—	—
	CPLEX	1,041,321	1,059,840	1.75%	3,600
FC=10	TF	1,109,573	1,276,930	13.11%	9
	TB	835,591	8,011,742	89.57%	1,200
	CPLEX	1,228,166	1,253,528	2.02%	3,600

Table C.7. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is decreased from 100% to 25% of the system's overall capacity. Planning horizon: 6 time periods.

T=6; Decreasing demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	1,354,930	1,612,831	15.99%	12
	TB	—	—	—	—
	CPLEX	1,495,380	1,615,756	7.45%	3,600
FC=10	TF	1,617,667	2,016,902	19.79%	10
	TB	1,554,001	10,345,623	84.98%	603
	CPLEX	1,836,565	2,004,646	8.38%	3,600

Table C.8. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is decreased from 100% to 25% of the system's overall capacity. Planning horizon: 12 time periods.

T=12; Decreasing demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	2,525,753	3,021,139	16.40%	30
	TB	2,483,394	16,362,457	84.82%	8
	CPLEX	2,745,308	4,638,513	40.81%	3,600
FC=10	TF	3,014,061	3,782,070	20.31%	39
	TB	2,948,905	17,122,756	82.78%	307
	CPLEX	3,393,866	4,369,306	22.32%	3,600

Table C.9. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is decreased from 100% to 25% of the system's overall capacity. Planning horizon: 24 time periods.

T=24; Decreasing demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	4,902,867	5,919,244	17.17%	116
	TB	—	—	—	—
	CPLEX	5,296,511	533,446,174	99.01%	3,600
FC=10	TF	5,839,586	7,545,787	22.61%	169
	TB	5,772,941	22,429,044	74.26%	40
	CPLEX	6,560,513	520,219,840	98.74%	3,600

Table C.10. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is decreased from 100% to 25% of the system's overall capacity. Planning horizon: 48 time periods.

T=48; Decreasing demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	9,684,487	11,888,284	18.54%	498
	TB	—	—	—	—
	CPLEX	10,429,505	1,072,093,875	99.03%	3,600
FC=10	TF	11,517,279	15,269,757	24.57%	628
	TB	—	—	—	—
	CPLEX	12,923,296	1,067,989,301	98.79%	3,600

C.3 Tests with seasonally increasing demand

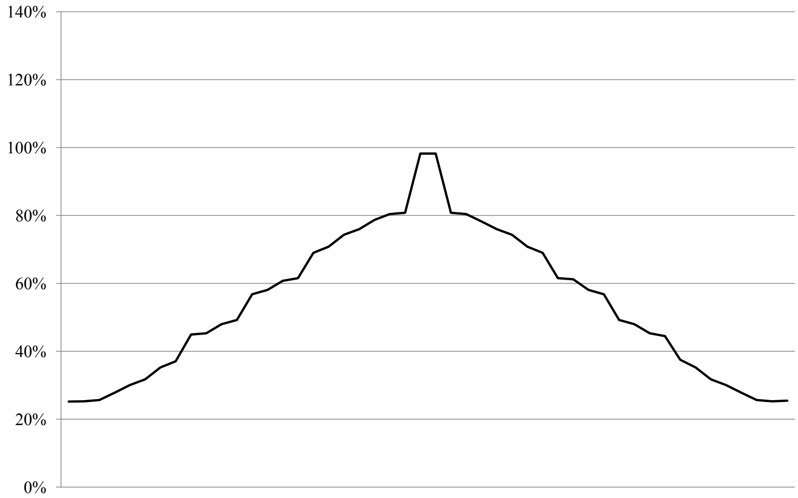

Figure C.3. Seasonally increasing demand behavior considered in the tests.

Table C.11. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is seasonally increased from 25% to 100% of the system's overall capacity. Planning horizon: 3 time periods.

T=3; Seasonally increasing demand		**LB**	**UB**	**Gap**	**t(s)**
	TF	588,871	695,720	15.36%	8
FC=1	TB	584,105	754,264	22.56%	4
	CPLEX	662,447	715,560	7.42%	3,600
	TF	713,685	879,480	18.85%	3
FC=10	TB	700,199	922,383	24.09%	4
	CPLEX	826,741	849,774	2.71%	3,600

Table C.12. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is seasonally increased from 25% to 100% of the system's overall capacity. Planning horizon: 6 time periods.

T=6; Seasonally increasing demand		**LB**	**UB**	**Gap**	**t(s)**
	TF	1,416,377	1,650,811	14.20%	3
FC=1	TB	—	—	—	—
	CPLEX	1,555,714	1,696,397	8.29%	3,600
	TF	1,677,708	2,097,453	20.01%	11
FC=10	TB	—	—	—	—
	CPLEX	1,919,986	2,044,108	6.07%	3,600

Table C.13. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is seasonally increased from 25% to 100% of the system's overall capacity. Planning horizon: 12 time periods.

T=12; Seasonally increasing demand		**LB**	**UB**	**Gap**	**t(s)**
	TF	2,595,421	3,027,976	14.29%	13
FC=1	TB	2,589,338	3,062,313	15.45%	10
	CPLEX	2,787,014	3,185,969	12.52%	3,600
	TF	3,079,842	3,774,151	18.40%	14
FC=10	TB	—	—	—	—
	CPLEX	3,444,954	3,876,076	11.12%	3,600

Table C.14. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is seasonally increased from 25% to 100% of the system's overall capacity. Planning horizon: 24 time periods.

T=24; Seasonally increasing demand		**LB**	**UB**	**Gap**	**t(s)**
	TF	4,958,385	5,919,356	16.23%	93
FC=1	TB	—	—	—	—
	CPLEX	5,327,805	525,942,747	98.99%	3,600
	TF	5,891,783	7,488,953	21.33%	106
FC=10	TB	—	—	—	—
	CPLEX	6,600,344	8,393,980	21.37%	3,600

Table C.15. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is seasonally increased from 25% to 100% of the system's overall capacity. Planning horizon: 48 time periods.

T=48; Seasonally increasing demand		**LB**	**UB**	**Gap**	**t(s)**
	TF	9,732,512	11,730,992	17.04%	357
FC=1	TB	—	—	—	—
	CPLEX	10,459,175	1,005,076,572	98.96%	3,600
	TF	11,561,540	14,971,922	22.78%	507
FC=10	TB	—	—	—	—
	CPLEX	12,963,100	1,070,461,766	98.79%	3,600

C.4 Tests with seasonally decreasing demand

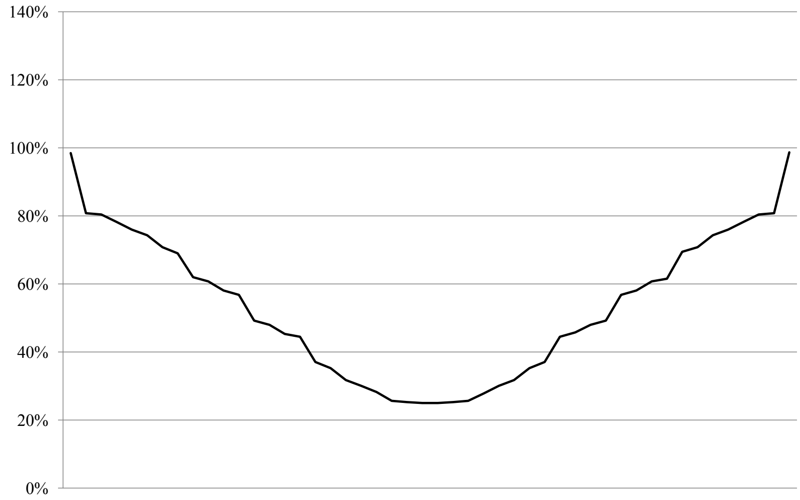

Figure C.4. Seasonally decreasing demand behavior considered in the tests.

Table C.16. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is seasonally decreased from 100% to 25% of the system's overall capacity. Planning horizon: 3 time periods.

T=3; Seasonally decreasing demand		LB	UB	Gap	t(s)
	TF	1,143,461	1,293,832	11.62%	5
FC=1	TB	906,867	3,519,862	74.24%	1
	CPLEX	1,239,664	1,260,402	1.65%	3,600
	TF	1,315,207	1,537,001	14.43%	4
FC=10	TB	1,071,747	9,115,771	88.24%	2
	CPLEX	1,442,393	1,484,169	2.81%	3,600

Table C.17. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is seasonally decreased from 100% to 25% of the system's overall capacity. Planning horizon: 6 time periods.

T=6; Seasonally decreasing demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	1,917,780	2,215,698	13.45%	11
	TB	1,882,483	16,186,266	88.37%	3
	CPLEX	2,090,185	2,309,060	9.48%	3,600
FC=10	TF	2,240,547	2,691,897	16.77%	10
	TB	2,205,358	17,256,897	87.22%	5
	CPLEX	2,530,558	2,888,830	12.40%	3,600

Table C.18. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is seasonally decreased from 100% to 25% of the system's overall capacity. Planning horizon: 12 time periods.

T=12; Seasonally decreasing demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	3,514,118	4,058,417	13.41%	44
	TB	3,485,842	17,812,758	80.43%	8
	CPLEX	3,763,228	4,855,850	22.50%	3,600
FC=10	TF	4,108,078	5,002,573	17.88%	69
	TB	4,075,332	19,392,657	78.99%	100
	CPLEX	4,577,866	243,553,122	98.12%	3,600

Table C.19. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is seasonally decreased from 100% to 25% of the system's overall capacity. Planning horizon: 24 time periods.

T=24; Seasonally decreasing demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	4,996,797	5,948,737	16.00%	134
	TB	4,965,913	19,060,250	73.95%	33
	CPLEX	5,399,442	517,424,422	98.96%	3,600
FC=10	TF	5,942,881	7,529,078	21.07%	141
	TB	—	—	—	—
	CPLEX	6,692,269	487,037,414	98.63%	3,600

Table C.20. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is seasonally decreased from 100% to 25% of the system's overall capacity. Planning horizon: 48 time periods.

T=48; Seasonally decreasing demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	9,767,736	12,198,223	19.92%	603
	TB	—	—	—	—
	CPLEX	10,516,221	1,090,508,353	99.04%	3,600
FC=10	TF	11,613,040	15,097,979	23.08%	627
	TB	—	—	—	—
	CPLEX	13,038,040	1,076,261,087	98.79%	3,600

C.5 Tests with uniformly varying demand

C.5.1 Uniform demand, at 25% of the system's capacity.

Figure C.5. Uniform demand behavior considered in the tests.

Table C.21. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly set to 25% of the system's overall capacity. Planning horizon: 3 time periods.

T=3; Uniform demand		LB	UB	Gap	t(s)
FC=1	TF	289,535	381,429	24.09%	7
	TB	—	—	—	—
	CPLEX	345,062	357,300	3.43%	3,600
FC=10	TF	380,832	511,795	25.59%	55
	TB	—	—	—	—
	CPLEX	457,410	476,158	3.94%	3,600

Table C.22. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly set to 25% of the system's overall capacity. Planning horizon: 6 time periods.

T=6; Uniform demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	558,704	703,010	20.53%	7
	TB	282,294	367,341	23.15%	1
	CPLEX	634,630	689,510	7.96%	3,600
FC=10	TF	722,617	960,251	24.75%	7
	TB	—	—	—	—
	CPLEX	833,456	945,988	11.90%	3,600

Table C.23. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly set to 25% of the system's overall capacity. Planning horizon: 12 time periods.

T=12; Uniform demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	1,097,090	1,377,072	20.33%	14
	TB	—	—	—	—
	CPLEX	1,233,975	1,395,287	11.56%	3,600
FC=10	TF	1,410,983	1,917,663	26.42%	13
	TB	—	—	—	—
	CPLEX	1,637,669	35,209,226	95.35%	3,600

Table C.24. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly set to 25% of the system's overall capacity. Planning horizon: 24 time periods.

T=24; Uniform demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	2,177,839	2,840,227	23.32%	43
	TB	2,170,224	2,988,802	27.39%	30
	CPLEX	2,448,631	3,076,616	20.41%	3,600
FC=10	TF	2,788,673	3,822,438	27.04%	44
	TB	—	—	—	—
	CPLEX	3,255,023	4,105,483	20.72%	3,600

Table C.25. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly set to 25% of the system's overall capacity. Planning horizon: 48 time periods.

T=48; Uniform demand		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	4,340,425	5,703,510	23.90%	207
	TB	—	—	—	—
	CPLEX	4,884,085	488,716,918	99.00%	3,600
FC=10	TF	5,547,226	7,957,822	30.29%	232
	TB	—	—	—	—
	CPLEX	6,475,950	507,794,415	98.72%	3,600

C.5.2 Uniform variance, up to 50% of the system's capacity.

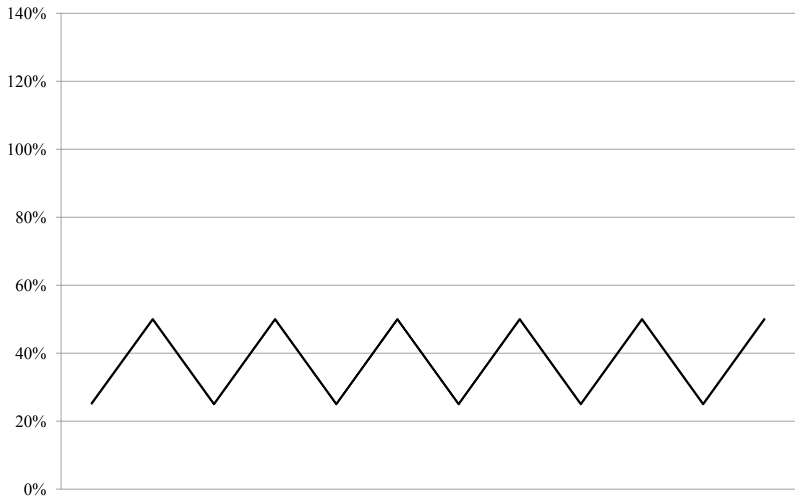

Figure C.6. Uniformly varying (±50%) demand behavior considered in the tests.

Table C.26. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±25% of the system's overall capacity. Planning horizon: 3 time periods.

T=3; Uniformly varying demand (±25%)		LB	UB	Gap	t(s)
	TF	381,412	446,784	14.63%	1
FC=1	TB	—	—	—	—
	CPLEX	425,456	440,358	3.38%	3,600
	TF	485,556	593,678	18.21%	1,255
FC=10	TB	475,531	602,651	21.09%	35
	CPLEX	557,564	574,699	2.98%	3,600

Table C.27. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±25% of the system's overall capacity. Planning horizon: 6 time periods.

T=6; Uniformly varying demand (±25%)		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	838,413	1,011,167	17.08%	5
	TB	836,551	1,037,017	19.33%	2
	CPLEX	931,147	1,008,477	7.67%	3,600
FC=10	TF	1,038,485	1,312,744	20.89%	4
	TB	—	—	—	—
	CPLEX	1,200,756	1,296,617	7.39%	3,600

Table C.28. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±25% of the system's overall capacity. Planning horizon: 12 time periods.

T=12; Uniformly varying demand (±25%)		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	1,658,991	1,997,316	16.94%	14
	TB	—	—	—	—
	CPLEX	1,817,034	2,017,170	9.92%	3,600
FC=10	TF	2,044,582	2,595,591	21.23%	15
	TB	—	—	—	—
	CPLEX	2,325,858	2,635,224	11.74%	3,600

Table C.29. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±25% of the system's overall capacity. Planning horizon: 24 time periods.

T=24; Uniformly varying demand (±25%)		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	3,304,138	4,087,119	19.16%	56
	TB	—	—	—	—
	CPLEX	3,619,253	4,793,566	24.50%	3,600
FC=10	TF	4,060,315	5,355,543	24.18%	56
	TB	—	—	—	—
	CPLEX	4,628,055	336,465,372	98.62%	3,600

Table C.30. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±25% of the system's overall capacity. Planning horizon: 48 time periods.

T=48; Uniformly varying demand (±25%)		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	6,601,484	8,220,554	19.70%	279
	TB	—	—	—	—
	CPLEX	7,213,714	754,541,150	99.04%	3,600
FC=10	TF	8,098,901	10,989,820	26.31%	316
	TB	—	—	—	—
	CPLEX	9,218,604	806,929,370	98.86%	3,600

C.5.3 Uniform variance, up to 75% of the system's capacity.

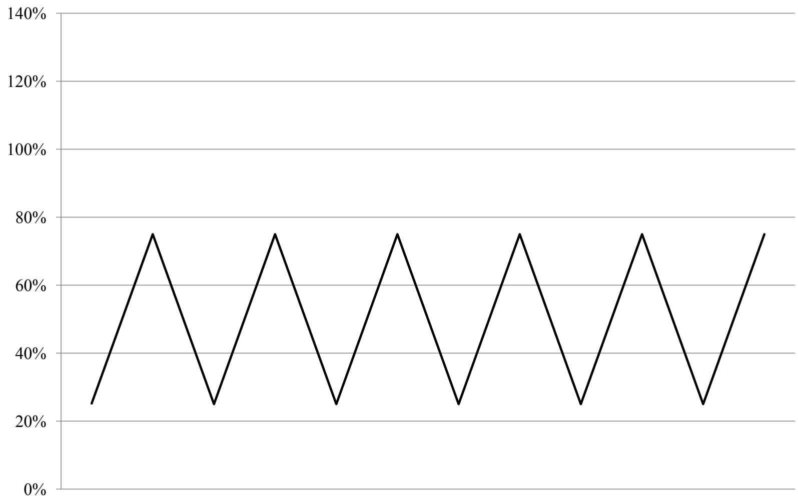

Figure C.7. Uniformly varying (±75%) demand behavior considered in the tests.

Table C.31. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±50% of the system's overall capacity. Planning horizon: 3 time periods.

T=3; Uniformly varying demand (±50%)		LB	UB	Gap	t(s)
FC=1	TF	487,806	582,245	16.22%	15
	TB	487,512	611,769	20.31%	2
	CPLEX	549,320	568,277	3.34%	3,600
FC=10	TF	600,748	751,091	20.02%	10
	TB	—	—	—	—
	CPLEX	700,363	720,501	2.79%	3,600

Table C.32. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±50% of the system's overall capacity. Planning horizon: 6 time periods.

T=6; Uniformly varying demand (±50%)		**LB**	**UB**	**Gap**	**t(s)**
	TF	1,157,478	1,337,662	13.47%	4
FC=1	TB	1,157,010	1,442,183	19.77%	2
	CPLEX	1,274,124	1,345,059	5.27%	3,600
	TF	1,391,404	1,721,964	19.20%	7
FC=10	TB	1,389,724	2,634,475	47.25%	557
	CPLEX	1,587,827	1,685,880	5.82%	3,600

Table C.33. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±50% of the system's overall capacity. Planning horizon: 12 time periods.

T=12; Uniformly varying demand (±50%)		**LB**	**UB**	**Gap**	**t(s)**
	TF	2,297,916	2,677,371	14.17%	11
FC=1	TB	2,296,462	2,827,066	18.77%	7
	CPLEX	2,476,447	2,753,297	10.06%	3,600
	TF	2,748,830	3,396,535	19.07%	13
FC=10	TB	—	—	—	—
	CPLEX	3,083,993	3,478,467	11.34%	3,600

Table C.34. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±50% of the system's overall capacity. Planning horizon: 24 time periods.

T=24; Uniformly varying demand (±50%)		**LB**	**UB**	**Gap**	**t(s)**
	TF	4,583,327	5,388,986	14.95%	55
FC=1	TB	4,583,563	5,614,154	18.36%	28
	CPLEX	4,931,861	499,064,302	99.01%	3,600
	TF	5,470,270	6,992,113	21.77%	71
FC=10	TB	—	—	—	—
	CPLEX	6,130,936	7,612,979	19.47%	3,600

Table C.35. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±50% of the system's overall capacity. Planning horizon: 48 time periods.

T=48; Uniformly varying demand (±50%)		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	9,163,902	11,196,372	18.15%	334
	TB	—	—	—	—
	CPLEX	9,843,832	1,011,340,445	99.03%	2,996
FC=10	TF	10,921,390	13,865,352	21.23%	384
	TB	—	—	—	—
	CPLEX	12,220,176	1,014,839,192	98.80%	1,794

C.5.4 Uniform variance, up to 100% of the system's capacity.

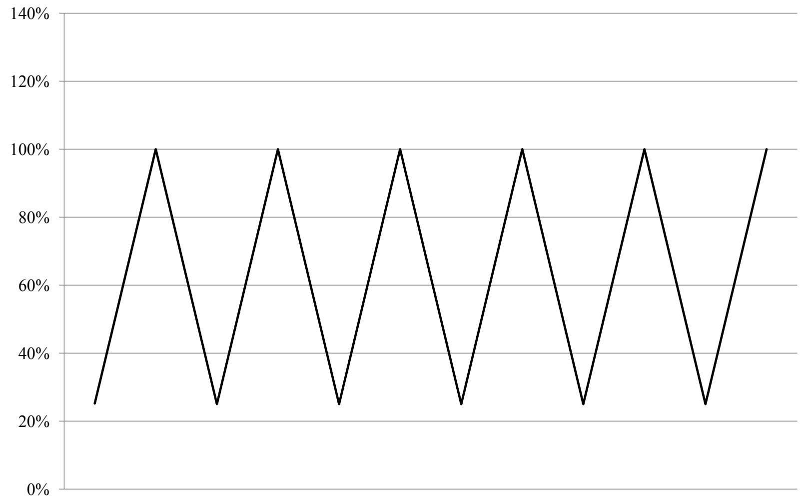

Figure C.8. Uniformly varying (±100%) demand behavior considered in the tests.

Table C.36. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±75% of the system's overall capacity. Planning horizon: 3 time periods.

T=3; Uniformly varying demand (±75%)		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	606,262	710,290	14.65%	5
	TB	601,633	774,790	22.35%	1
	CPLEX	682,351	696,573	2.04%	3,050
FC=10	TF	732,288	897,349	18.39%	18
	TB	—	—	—	—
	CPLEX	846,320	870,316	2.76%	3,122

Table C.37. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±75% of the system's overall capacity. Planning horizon: 6 time periods.

T=6; Uniformly varying demand (±75%)		LB	UB	Gap	t(s)
FC=1	TF	1,502,300	1,727,912	13.06%	4
	TB	—	—	—	—
	CPLEX	1,648,036	1,743,331	5.47%	3,063
FC=10	TF	1,770,530	2,147,961	17.57%	3
	TB	—	—	—	—
	CPLEX	2,004,740	2,130,123	5.89%	3,029

Table C.38. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±75% of the system's overall capacity. Planning horizon: 12 time periods.

T=12; Uniformly varying demand (±75%)		LB	UB	Gap	t(s)
FC=1	TF	2,982,966	3,438,821	13.26%	11
	TB	—	—	—	—
	CPLEX	3,195,750	3,492,561	8.50%	3,019
FC=10	TF	3,500,998	4,275,036	18.11%	15
	TB	—	—	—	—
	CPLEX	3,904,802	4,502,406	13.27%	2,154

Table C.39. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±75% of the system's overall capacity. Planning horizon: 24 time periods.

T=24; Uniformly varying demand (±75%)		LB	UB	Gap	t(s)
FC=1	TF	5,948,687	6,923,416	14.08%	65
	TB	—	—	—	—
	CPLEX	6,361,095	496,731,218	98.72%	2,774
FC=10	TF	6,967,883	8,580,955	18.80%	90
	TB	—	—	—	—
	CPLEX	7,763,934	601,762,416	98.71%	2,595

Table C.40. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±75% of the system's overall capacity. Planning horizon: 48 time periods.

T=48; Uniformly varying demand (±75%)		LB	UB	Gap	t(s)
FC=1	TF	11,891,343	14,461,121	17.77%	375
	TB	—	—	—	—
	CPLEX	12,701,260	1,263,650,655	98.99%	2,865
FC=10	TF	13,910,218	18,111,215	23.20%	473
	TB	—	—	—	—
	CPLEX	15,469,805	1,264,506,937	98.78%	3,008

C.5.5 Uniform variance, up to 125% of the system's capacity.

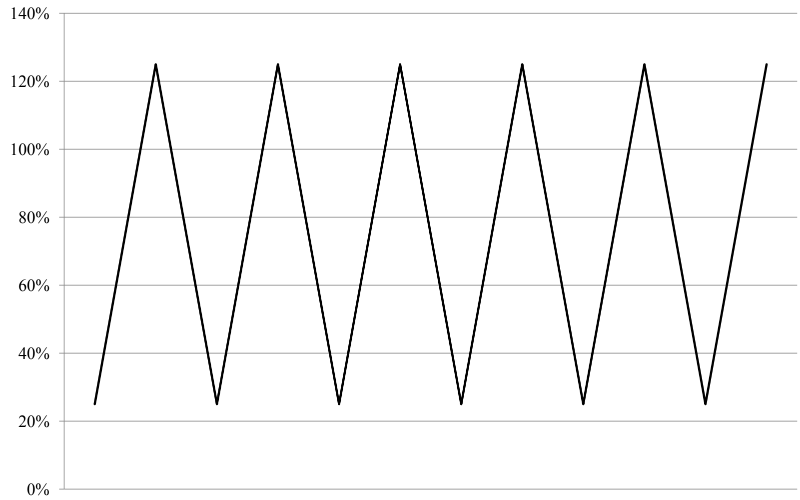

Figure C.9. Uniformly varying (±125%) demand behavior considered in the tests.

Table C.41. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±125% of the system's overall capacity. Planning horizon: 3 time periods.

T=3; Uniformly varying demand (±125%)		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	733,272	861,842	14.92%	8
	TB	—	—	—	—
	CPLEX	827,557	843,754	1.92%	3,096
FC=10	TF	873,650	1,063,588	17.86%	6
	TB	—	—	—	—
	CPLEX	1,015,472	1,038,867	2.25%	3,000

Table C.42. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±125% of the system's overall capacity. Planning horizon: 6 time periods.

T=6; Uniformly varying demand (±125%)		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	1,873,825	2,159,250	13.22%	8
	TB	1,868,875	6,539,925	71.42%	12
	CPLEX	2,055,671	2,118,011	2.94%	3,088
FC=10	TF	2,179,870	2,613,154	16.58%	6
	TB	—	—	—	—
	CPLEX	2,469,435	2,621,922	5.82%	2,087

Table C.43. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±125% of the system's overall capacity. Planning horizon: 12 time periods.

T=12; Uniformly varying demand (±125%)		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	3,722,147	4,270,270	12.84%	22
	TB	—	—	—	—
	CPLEX	3,971,827	268,887,475	98.52%	3,034
FC=10	TF	4,314,623	5,700,046	24.31%	640
	TB	—	—	—	—
	CPLEX	4,796,851	6,163,499	22.17%	2,804

Table C.44. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±125% of the system's overall capacity. Planning horizon: 24 time periods.

T=24; Uniformly varying demand (±125%)		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	7,424,814	8,783,520	15.47%	1,290
	TB	—	—	—	—
	CPLEX	7,903,988	749,032,766	98.94%	2,870
FC=10	TF	8,588,988	11,655,988	26.31%	292
	TB	—	—	—	—
	CPLEX	9,528,358	743,436,906	98.72%	2,509

Table C.45. Results comparing CPLEX and Relax&Fix strategies TF and TB when demand is uniformly varied around ±125% of the system's overall capacity. Planning horizon: 48 time periods.

T=48; Uniformly varying demand (±125%)		**LB**	**UB**	**Gap**	**t(s)**
FC=1	TF	14,843,064	18,536,984	19.93%	551
	TB	—	—	—	—
	CPLEX	15,769,441	1,520,416,607	98.96%	3,081
FC=10	TF	17,150,450	26,620,013	35.57%	752
	TB	—	—	—	—
	CPLEX	18,987,621	1,516,160,195	98.75%	3,054