ELSEVIER

Enterprise-wide modeling & optimization—An overview of emerging research challenges and opportunities

V.A. Varma, G.V. Reklaitis *, G.E. Blau, J.F. Pekny

School of Chemical Engineering, Purdue University, West Lafayette, IN, United States

Abstract

The process systems engineering (PSE) as well as the operations research and management science (ORMS) literature has hitherto focused on disparate processes and functions within the enterprise. These themes have included upstream R&D pipeline management, planning and scheduling in batch and continuous manufacturing systems and more recently supply chain optimization under uncertainty. In reality, the modern process enterprise functions as a cohesive entity involving several degrees of cross-functional co-ordination across enterprise planning and process functions. The complex organizational structures underlying horizontally and vertically integrated process enterprises challenge our understanding of cross-functional co-ordination and its business impact. This article looks at the impact of enterprise-wide cross-functional coordination on enterprise performance, sustainability and growth prospects. Cross-functional coordination is defined as the integration of strategic and tactical decision-making processes involving the control of financial and inventory flows (both internal and external) as well as resource deployments. Initially, we demonstrate the existence of cross-functional decision-making dependencies using an enterprise network model. Subsequently, we discuss interactions between enterprise planning decisions involving project financing, debt-equity balancing, R&D portfolio selection, risk hedging with real derivative instruments, supply chain asset creation and marketing contracts which influence decision-making at the activity/process level. Several case studies are included to re-enforce the point that planning and process decisions need to be integrated.
© 2007 Elsevier Ltd. All rights reserved.

Keywords: Supply chain management; Enterprise modeling; Risk management; Product pipeline management

1. Introduction

Globalization trends have significantly increased the scale and complexity of the modern enterprise. The enterprise has been transformed into a global network consisting of multiple business units and functions. Operational functions include R&D pipelines, production networks (both batch and continuous) and supply chain networks. These functions are supported by financial planning and marketing strategy functions. The enterprise is exposed to internal and external uncertainties. Examples of internal uncertainties include success prospects of R&D projects due to technological risks; production upsets such as batch failures and plant shutdowns. External uncertainties include pricing related uncertainties for raw materials and products (unless the firm is operating in a monosopny or a monopoly), exchange rate fluctuations, market size and demand uncertainties due to competition and macro-economic factors. Process enterprises respond to the evolving business and technology environments by strategic maneuvers involving R&D, manufacturing, supply chain and marketing functions. Strategic maneuvers involving R&D include capital budgeting, R&D project selection and project commercialization. Strategic maneuvers involving the manufacturing function include capital project planning, risk management and financing. Strategic maneuvers involving the supply chain function include distribution asset creation as well as supply chain risk management. Strategic maneuvers involving the marketing function include contracts design and management. Strategic decisions answer the question—"What should the enterprise do to attain strategic goals?". In addition, tactical decisions must answer the question—"How should the enterprise execute strategic decisions?". Tactical R&D decisions include project scheduling and R&D resource management. In recent times, strategic and tactical R&D management has been termed as Innovation Process

* Corresponding author. Tel.: +1 765 494 4075; fax: +1 765 494 0805.
  E-mail address: reklaiti@ecn.purdue.edu (G.V. Reklaitis).

Management (IPM). Tactical manufacturing and supply chain decisions include batch plant scheduling in response to forecasted demands, production asset management and selection of energy feedstock in response to market prices. Traditional PSE and ORMS literature streams tend to focus on sub-sets of these decisions whereas an enterprise functions as a cohesive entity with several degrees of cross-functional co-ordination. Often lack of such cross-functional co-ordination leads to loss of short and long term value. Further, organizational complexity tends to challenge much of our understanding about cross-functional coordination and its business impact. Hence, from the enterprise-wide performance viewpoint, it is sub-optimal to optimize strategic and tactical decisions in a disparate fashion as has been done hitherto in the literature. At the same time, integrated enterprise-wide decision-making is significantly more challenging in comparison to function-specific decision-making. In fact rigorous literature on enterprise-wide modeling is very sparse.

The goal of this paper is to demonstrate that integration of enterprise decision-making leads to substantial value creation. By doing so, we hope to motivate a strong case for development of models that will efficiently integrate decision-making related to R&D, manufacturing, supply chain and marketing functions and help in enhancing our understanding of co-ordination across these functions. The paper does not cover research opportunities in the area of enterprise-wide work practices/systems which is closely related to implementing the enterprise-wide optimization strategies proposed in this paper. These practices/systems include robust and reliable data acquisition systems, development of human skills required to drive the proposed enterprise-wide coordination strategies and real-time modeling and solution systems that can help managers test/develop their insights with formal decision models without the need to understand formulation and algorithmic details. We feel that these issues are too extensive and critical to be covered in a single paper and justify and indeed require separate treatment. The paper is organized as follows: We present an overview of strategic and tactical decision models developed in the Operations Research-Management Science (ORMS) as well as process systems engineering (PSE) literature. The models reviewed are those that in our view are expected to play significant roles as components of enterprise planning architectures. Section 2 presents a critical review of the relevant ORMS literature. Section 3 presents a critical review of the relevant PSE literature. Section 4 presents an enterprise network model that conceptualizes the need for integration of decision-making. Section 5 presents a discussion supported by examples on the integration of capital budgeting and R&D Project Prioritization. Section 6 presents a discussion supported by examples of integrating resource allocation, manufacturing and scheduling decisions under uncertain R&D environments. Section 7 presents several examples on the integration of supply chain components such as risk hedging contracts, integration of production and inventory network decisions, integration of production and capacity planning decisions, integration of production and marketing strategy decisions. Section 8 presents an outlook on computational strategies for enterprise wide modeling and optimization. Finally,

Section 9 summarizes our perspective on the upcoming research area of enterprise-wide modeling and optimization.

2. The ORMS strategic and tactical literature

A comprehensive enterprise-wide modeling framework requires a unification of methodologies developed in corporate finance as well as the operations research and Management Science (ORMS) literature. Hence, we present separate surveys of the strategic and operational literature.

2.1. Strategic enterprise models

Strategic enterprise models (also called capital budgeting models) were devised to build enterprise portfolios that ensure long-term value creation. The earliest capital budgeting models were based on pure economic analysis (Chapman & Ward, 1996). The Discounted Cash Flow (DCF) method remains the most commonly used technique (Krishnan & Ulrich, 2001). However, it is based on expected values of uncertain parameters and is unable to generate quantitative details about the risk associated with a given project (Poh, Ang, & Bai, 2001). While simple from every aspect, the DCF method has been criticized on several counts. The DCF method fails to account for uncertainties in the costs as well as commercial returns. The method simply uses the expected values of the probability distributions modeling these variables. Further, most real investment projects include several decision-making flexibilities embedded within their execution structure. For instance, decision-makers have the flexibility to discontinue funding a project if a competitor captures an unacceptable share of its market segment. Such flexibilities are commonly termed as "Real Options" referring to options on real investments (Pindyk & Dixit, 1994). Failure to incorporate the value generated by such options leads to undervaluation of the investment project. The DCF method also has been criticized for its rigid focus on single criterion decision-making versus more realistic multiple criteria decision-making (Linton, Walsh, & Morabito, 2002). Thus, a method is required that can incorporate project uncertainties, multiple reward and risk criteria as well as embedded options.

Efforts to solve this problem have been collectively termed as decision theoretic methods (Morgan & Henrion, 1990; Nutt, 1998). Decision theory formalizes the key concepts of risk and return by defining the decision-maker's utility function (Markowitz, 1991). Using this formalism, decision theory provides comprehensive portfolio management methods such as decision trees which allow management to undertake complex resource allocation decisions between competing product candidates with full consideration to the possibilities of product failures (Sharpe & Keelin, 1994). The decision tree method also has addressed portfolio management issues such as how many projects to pursue and how many projects to terminate (Ding & Eliashberg, 2002). From a conceptual point of view the Decision tree method is a manifestation of Stochastic Dynamic Programming (Bertsekas, 2000). Roberts (1999) presents an interesting account of market competition related risks to profitability in the pharmaceutical industry. One of the most comprehensive multi-

project analyses using decision trees is presented by Triantis and Childs (2001). They model the present value dynamics of R&D projects using a stochastic partial differential equation which is equivalent to addressing market risks. Subsequently, they discretize the underlying continuous stochastic dynamics using a trinomial model. The trinomial decision tree is used to analyze a two-product case to answer questions such as when to abandon a project given limited resources, how to stagger the development of the two products depending upon where one is in the trinomial tree. In this fashion, optimal dynamic investment policies are formulated. The work by Loch and Bode-Greuel (2001) in the nascent biotech industry is another notable example where Real Options have been applied. While Real Options are penetrating corporate decision-making gradually, several objections have been raised against the decision tree method of analyzing these options. A major criticism points to the occurrence of unmanageably large decision trees, due to significantly rapid increase in the number of project selection and sequencing decisions with the size of the portfolio. Another criticism that is primarily raised by the corporate finance community is the inability of the decision tree method to adapt the discounting rate of return to the non-uniform risk profile of the project (Copeland & Antikarov, 2000). Instead corporate finance researchers recommend the use of financial no-arbitrage theory for purposes of investment project valuation.

The financial no arbitrage principle simply states that a riskless (perfectly hedged) portfolio can fetch a return no different than the risk-free or the 'bank' rate of return. This simple principle has been used to establish fair valuations of financial options. The most famous valuation formula based on the no-arbitrage principle is the Black-Scholes formula (Black & Scholes, 1973) used to value European type call and put options on stocks and commodity assets. Corporate finance researchers have established that several real options embedded in R&D projects show behaviors similar to European and American options (Amram & Kulatilaka, 1999). For example the option to abandon a project in response to market downturns is identical to an American put. The option to defer the development of a product until market studies are undertaken is an example of a European call since the firm will choose to develop the product (at an exercise price = total forward capital requirement) only if the exercise price is lower than the cumulative returns. Such an option is also termed as a deferral option. The option to add capacity after observing the arrival pattern of products into R&D pipelines is similar to an American call and is termed as a growth option. The right to shut down or slow down an operation paying certain fixed costs and the right to restart or speed it later for a different fixed cost is called a 'switching' option. Corporate finance researchers (Pindyk & Dixit, 1994) argue that the similarity of real options and financial options implies that the no-arbitrage based valuation methods developed for financial options may be re-engineered for valuation of R&D projects with embedded real options.

To demonstrate the no-arbitrage valuation of a drug project with an abandonment option, assume that a pharmaceutical enterprise has a commercialized drug product that has only 2 years of patent life remaining. Before the patent runs out the firm is weighing its options to market the drug as a different formulation for the same target (for example as an injection as against a capsule). The firm is concerned about the market risks (technical risks are insignificant). Further the firm has an option to out-license the drug to a generics company for a licensing fee of K ('Salvage' value) in 2 years from now. Let X = Current estimated market value of the re-formulated drug. Assume that after 2 years this value can either move to Xu or Xd where $u = \exp(\sigma\sqrt{dt})$ where $dt = 2$ years and σ = market volatility and $d = 1/u$ (Copeland & Antikarov, 2000). Then the value of the drug in the 'up'-state (i.e. when the market moves up) = V_u ($t = 2$) = max (Xu, K) and its value in the 'down'-state = V_d ($t = 2$) = max (Xd, K). Then according to the financial no-arbitrage theorem the value of the project with the embedded abandonment option (with an existing risk free rate r_{free}) is:

$$V(0) = \left(\frac{u - 1 - r_{free}}{u - d} \right) \frac{\max(Xd, K)}{1 + r_{free}}$$
$$+ \left(\frac{1 + r_{free} - d}{u - d} \right) \frac{\max(Xu, K)}{1 + r_{free}} \tag{1}$$

This result contrasts with the DCF and the decision tree techniques. In the first case each option would have been evaluated independently, while in the second one the factors in parenthesis in Eq. (1) would have disappeared and a subjective discounting factor would have been used. A major drawback of the Real Options methodologies is the combinatorial explosion of options when projects are combined into a portfolio. Hence, while both decision trees and the no-arbitrage methods are feasibly applied for single project valuation their applicability to even moderately sized portfolios runs into problems related to computational complexity.

By contrast, Monte Carlo/Discrete Event simulation methods (Law & Kelton, 2000) can accommodate uncertainties, embedded real options and alternative performance criteria in a computationally feasible manner (Blau & Bunch, 2002). Simulation-based methods have been employed by Adler, Mandelbaum, Nguyen, and Schwerer (1995) to analyze a relatively complex engineering design process, by Blau et al. (2000) in simulating an industrial-scale pharmaceutical new product development (NPD) pipeline, by Repenning (2001) in modeling the control and dynamics of a 2-stage NPD system, by Subramanian, Pekny, and Reklaitis (2003) in studying the effect of activity re-scheduling on portfolio performance and by Blau, Pekny, Varma, and Bunch (2004) in optimizing the reward-risk trade-offs for pharmaceutical R&D portfolios.

2.2. Tactical enterprise models

Several planning and process level enterprise decision problems can be reduced to the so-called Resource Constrained Project Scheduling Problem (RCPSP). Examples include sales force deployment, natural resource development projects, R&D and production scheduling. In the following paragraphs we present an overview of different models and associated algorithms in the deterministic and stochastic RCPSP literature.

The resource constrained project scheduling (RCPS) literature can be classified based on various formulations of the scheduling problem (Brucker, 1995). In terms of the objective, project makespan and Net Present Value (NPV) remain the most widely studied. Another classification criterion is whether the scheduling problem is a single mode or multi-mode (Elmaghraby, 1997) problem. In a single mode problem each project activity has a fixed duration and resource requirement while in the multi-mode version at least one activity has processing durations dependent upon the resources allocated to it. More recently, other classifications such as deterministic or stochastic resource constrained project scheduling, single or multiple project scheduling have emerged. Doersch and Patterson (1977) were among the earliest researchers to show that the general RCPSP is NP-Hard and apply 0–1 integer programming to the single mode RCPS problem. Due to the simplified nature of the single project, single mode RCPS problem, extensive efforts through the 1980s (Patterson, 1984) and 1990s were directed towards this class of problems using exact approaches. In early 1990s, with the availability of significantly greater computing power, implementation of exact branch-and-bound algorithms for reasonably large size instances became possible. Demeulemeester and Herroelen (1992, 1997) proposed branch and bound (B&B) algorithms for the makespan minimization, single mode RCPS problem.

Icmeli and Erenguc (1996a) propose a B&B algorithm for the NPV maximization (MAX-NPV) single mode RCPS problem. The reported performances of several B&B algorithms for the single and multi-modal RCPS problem have been primarily based on a standard set of small to medium test instances proposed by Davis and Patterson (1975) and Sprecher and Kolisch (1996). On the other hand, realistic scheduling problems involve hundreds of activities with special precedence and resource-activity structures. The inability to solve such large realistic problems using B&B has led to a broad literature on project scheduling heuristics. Generally, these heuristics are based on serial and parallel schedule generation schemes (Kolisch, 1996a, 1996b; Naphade, Wu, & Storer, 1997; Salewski, Schirmer, & Drexl, 1997). The heuristics generate either non-delay (parallel schedule generation scheme) or active (serial schedule generation scheme) schedules and typically use priority rules to resolve resource conflicts. For example, Baroum and Patterson (1996) devise weighted cash flow based heuristic procedures for the max-npv single mode, RCPS problem. They claim to obtain improved solutions on the standard Patterson set of test problems (7–50 activities) using their simple cash flow based heuristics. The primary drawback of these procedures is the large and unknown variability in quality of solutions across problem instances.

Heuristics for the multi-mode version of the RCPSP have to contend with another set of combinatorial decisions: the modes of execution of each activity. Kolisch and Drexl (1997) propose a local search based heuristic for the multi-mode RCPS problem, while Hartmann (1997) extends the local search approach to a genetic algorithm (Goldberg, 1989). Icmeli and Erenguc (1996b) present detailed analysis of list scheduling algorithms for generating heuristic solutions to the multi-mode RCPS problem.

They propose a Net Marginal Gain (NMG) heuristic to assign modes to activities during the list scheduling process. They use several priority rules for resolving resource conflicts and use Lagrangian relaxation to establish quality of the solutions generated by their heuristic scheme. Mohring, Schulz, Stork, and Uetz (2003) present a Lagrangian decomposition algorithm for solving a single mode RCPS problem with a general objective function. Their algorithm relaxes all resource constraints and solves the relaxed sub-problem by using a minimum cut network algorithm which runs in polynomial time. Weglarz (1999) gives a comprehensive survey of algorithmic developments in the field of project scheduling.

While there is a rich deterministic project scheduling literature, the literature on stochastic versions of this problem is almost void. Herroelen and Leus (2005) provide an exhaustive survey of the limited volume of works. Stork and Mohring (2000) describe several Branch-and-Bound algorithms that optimize Expected Makespans over a restricted set of scheduling policies.

3. The process systems engineering (PSE) literature review: planning under uncertainty

The process systems engineering (PSE) community has evolved a rich set of tools and methodologies in the area of operational planning and scheduling under uncertainty. The key purpose of this section is to provide a synopsis of PSE modeling and algorithmic efforts in the area of planning and scheduling under uncertainty.

The PSE world has been instrumental in developing stochastic optimization models for planning, scheduling and supply chain problems primarily in response to the inherently uncertain environment into which the process industry is embedded. The earliest works of the 1990s focus on scheduling multi-product batch plants under process and demand uncertainties (Straub & Grossmann, 1992). Carlson and Felder (1992) are among the first in the PSE community to use queuing network analysis for modeling batch production networks effectively integrating a methodology traditionally applied to discrete part manufacturing systems into the PSE area. Ierapetritou and Pistikopoulos (1995) present a two-stage stochastic programming (2-SSP) model that optimizes here-and-now design decisions with full consideration of second stage operating recourse decisions. Ahmed, Tawarmalani, and Sahinidis (2000) describe the application of a novel interval-based branch-and-bound strategy to efficiently solve large scale capacity expansion MILPs under demand uncertainty. Applequist, Pekny, and Reklaitis (2000) present a volume-polytope integration method based on the Lasserre algorithm to evaluate the expectation term of stochastic programs. Ierapetritou, Acevedo, and Pistikopoulos (1996) solve the same problem using a decomposition scheme based on the underlying block angular structure of the problem and hence make their approach more scalable though they do not extend their analysis to more than two stages. Acevado and Pistikopoulos (1997) identify that a significant component of the computational complexity in procedures that numerically evaluate the second stage function lies in the sub-problem solutions.

They propose a multi-parametric approach in which the solution of second stage linear programs is established as function of the uncertain variables by maintaining a list of several bases. Clay and Grossmann (1997) propose a 2-SSP for production planning involving uncertainties in costs and demands. They present theoretical properties of the model as well as compare different decomposition strategies such as Benders decomposition and successive aggregation. Bok, Heeman, and Park (1998) present a 2-SSP capacity expansion model and apply that model to a large scale East Asian supply chain case study.

Bose and Pekny (2000) extend the concept of Model Predictive Control (MPC) to the management of supply chains. They use forecasting and scheduling models within a simulation environment in order to optimize the expected customer satisfaction level (CSL) for centralized, de-centralized and distributed supply chains and study the sensitivity of the CSL to strategic system parameters such as lead times. Gupta, Maranas, and McDonald (2000) present a 2-SSP model that minimizes medium term supply chain costs under demand uncertainty. The key feature of their model is the identification of optimal supply policies under low, intermediate and high demand regimes that enables analytical evaluation of the second stage expectation resulting in a MINLP. Recent research efforts in supply chain (SC) management have focused on the strategic aspects of SC—the SC design problem and de-centralized SC problems. Jung, Blau, Pekny, Reklaitis, and Eversdyk (2004) present a simulation-optimization framework that jointly optimizes the strategic problems of warehouse location and safety stock placement and the tactical problem of re-scheduling production in response to demand uncertainty and events related to the production environment. This framework effectively extends the MPC framework of Bose and Pekny (2000) discussed earlier.

The PSE community has also increasingly recognized the importance of including more detailed consideration of financial components in supply chain and plant operational planning and scheduling. For instance, Oh and Karimi (2004) have incorporated international tax management components such as transfer prices, tax credits, duties and duty drawbacks in supply chain decisions. Badell, Romero, Huertas, and Puigjaner (2004) and Guillen, Badell, Espuna, and Puigjaner (2006) have addressed the timing of short term cash flows, short term financing and short term working capital management in planning and scheduling of batch plants. Finally, Barbaro and Bagajewicz (2004) have considered options in managing risk in production planning under uncertainty.

R&D pipeline management is another research domain that actively engages the PSE community. Honkomp, Reklaitis, and Pekny (1997) proposed an MILP model for selection and sequencing of R&D projects. Their formulation incorporates R&D uncertainties in an expected sense as against a scenario-decomposition approach. Jain and Grossmann (1999) propose a continuous time formulation of the same problem. Maravelias and Grossmann (2001) present a multi-period stochastic programming model that incorporates product selection, manufacturing and testing decisions in a single monolithic model that is solved using Lagrangian decomposition.

The R&D pipeline management goes beyond the issue of scheduling. At a strategic level the problem manifests itself as a portfolio selection problem and as a capacity-planning problem. The strategic portfolio selection problem addresses the issue of optimal selection and prioritization of R&D projects into a portfolio given future technical and market risks. The objective is to maximize an expected net present value (ENPV) at a constrained level of financial risk. Rogers, Gupta, and Maranas (2002) represent a simplified version of the problem as a quadranominal multistage decision tree in which uncertainty in product market value is represented by a Brownian geometric model and technical failure in product development by a binomial. Rogers, Maranas, and Ding (2005) extend this model to include in-licensing and timing of investment decisions. Blau et al. (2004) propose a framework that combines a discrete event simulator and a genetic algorithm to perform strategic portfolio selection. Their framework is able to capture an approximately efficient parabolic reward-risk frontier for a pharmaceutical R&D portfolio. Subramanian et al. (2003) propose a dual loop architecture (SIM-OPT) that enables the automated learning of reactive scheduling policies by studying the responses of a process scheduler to R&D events. Capacity planning aspects have been addressed by Levis and Papageorgiou (2004) who consider the problem in which the development stages and uncertainties are aggregated into lumped product success probabilities while product demand uncertainties are treated explicitly.

In summary the past two decades have seen the development of several key stochastic optimization models and algorithms by PSE researchers. Most of these models were developed to solve operational planning and scheduling problems in domains such as batch and continuous processing, supply chain process management and more recently R&D pipeline management. However, current global trends of process and business integration have resulted in the emergence of opportunities to solve these models within coordinated architectures (Pekny, 2002). The range of uncertainties that need to be handled has expanded tremendously due to such global winds of change. These uncertainties include fluctuating raw material and utility prices, currency exchange rates and the more familiar demand and pricing fluctuations.

In the following section, we propose a general enterprise-wide modeling architecture focused on coordinating enterprise-wide decisions.

4. A network model of the process enterprise

Fig. 1 presents a high level structural map of a typical process enterprise. We emphasize that this is a generic enterprise model without reference to any specific firm. The enterprise may be a group of one or more Strategic Business Units (SBUs). An SBU is typically built out of two major divisions – A Product R&D division and a Commercial Offerings Management division. In addition, de-centralized SBUs may have their own finance and support functions. In the present case, the finance and certain support functions are shown to service all of the SBUs. The rectangular nodes represent individual functions within an SBU. There are typically three types of functions—Planning, Process
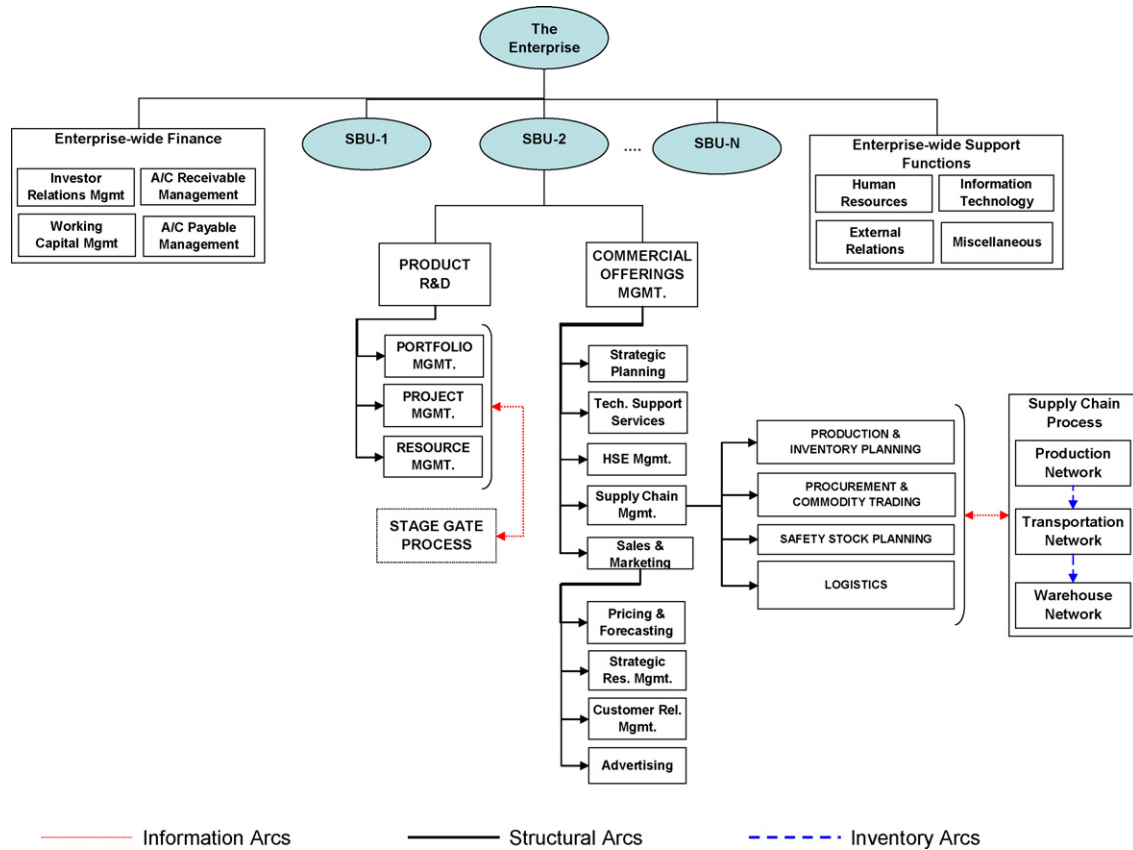
Fig. 1. A structural map of the process enterprise.

and Support. Planning functions are primary decision-making functions that supervise the activities of process functions. The input to a planning function is information about the state of the process function that is being supervised. Planning functions are embedded within all SBU divisions to supervise the processes undertaken by those divisions. In the following paragraphs, we briefly discuss each planning function and its influence upon the financial and inventory flows. The focus will then shift to the often ineffectual coordination across these planning functions and the potential value of coordinated decision-making.

4.1. Product R&D

The Product R&D division manages all major growth platforms that constitute the SBU's R&D portfolio. The primary planning functions within this division include portfolio management, project management and resource management. These functions supervise the functioning of the stage gate process (O'Connor, 1994)—the R&D work process in which products are developed incrementally in a way that allows close monitoring of technical and commercial milestones. Some stage gate processes (especially in the pharmaceutical industry) encompass activities such as process R&D, product design and engineering. Typically these could be standalone activities within some or all stages of development and run in series or parallel with other activities within each stage. For example, in the pharmaceutical industry process R&D and Engineering activities can start

as early as Phase II clinical trials or earlier in order to position the business to meet large scale volume demands once the drug goes into Phase III trials and beyond. The primary goal of these functions is to ensure a large and steady throughput of high value new commercial offerings (NCOs) streaming out of the stage gate process. Market segment forecasting is yet another key planning component of the stage gate process. Kahn (2002) provides an excellent overview of recent forecasting practices for new product development.

The strategic decisions resting within this function include project selection, capital budgeting, project prioritization and resource allocation. Ideally, all these decisions need to be coordinated. However, as we will see in the following sections, this is often not the case, leading to sub-optimal portfolios and ineffective utilization of capital and R&D resources.

4.2. Commercial offerings management (COM)

The COM is a group of functions that ensures highest value extraction from existing product offerings. The functions include strategic product management, technical and engineering support services, Health, Safety and Environment (HSE) management, supply chain management and Sales & Marketing. The COM's actions directly affect a host of outcomes including revenue growth, profit growth, market shares, margins and other metrics. The strategic product management function oversees decisions such as market segmentation, new market and busi-

ness development, production, technical services, supply chain and marketing budgets etc. This function assumes a key role in ensuring an effective allocation of capital and other resources across other COM functions. Technical and Engineering (T&E) support services ensure the smooth functioning of all production and material delivery systems that are critical towards ensuring high product throughput rates and product quality at the least cost inputs. HSE has assumed an ever-increasing role in ensuring high standards of emissions controls, worker safety and health. It is increasingly clear that share-holders and other key capital market players directly associate a firm's intrinsic value with its ability to stay ahead of environmental legislation. Hence, firms have a significant financial incentive to deploy systems that minimize their emission footprints. A case in point is the European Carbon emissions trading regime—a market mechanism that has forced firms to consider CO_2 sequestration technologies in order to reduce long-term emission credit costs. HSE is thus increasingly representative of a process enterprise's corporate ethics. The supply chain management function is composed of a series of sub-functions that include production and inventory planning, procurement and commodity trading (and contractual management), safety stock planning and transportation logistics. As shown in the figure, decisions originating from these functions directly impact the supply chain process operations. Some of these decisions may be strategic (e.g. plant capacity expansions) while the others may be operational in nature (e.g. production scheduling). The PSE community has developed a set of tools to coordinate these decisions (as will be discussed below), but the industrial deployment of these tools has not kept pace with their development. Finally, the Sales and Marketing (S&M) function is responsible for direct customer interface and pricing. S&M also makes decisions on sales force deployment based upon the existing market shares, product offerings, marketing costs and budgets. This decision problem can be looked upon as a type of knapsack problem.

There are other support functions like finance and IT that service the capital, information and other resource requirements

of the main planning and process functions. Closer integration of finance with SBU planning functions is expected to ensure smoother servicing of working capital requirements at significantly reduced costs of capital.

The above discussion should suggest that decisions are fragmented across the enterprise assuming a control strategy that resembles a "cascade control". In certain situations, cascade control is appropriate. For example, the strategic planning function operates on much longer time scales of up to several years. All other planning and tactical decisions need to align with strategic plans involving market positioning and growth strategy. However, in many situations a top-down or "cascade" control strategy can result in lags in response times to market events that can result in less than optimal value extraction from current offerings as well as loss of valuable new product development opportunities. One of the reasons for the lags is the "loop-back" effect in decision-making. For example, in the area of new product development, management may opt for a portfolio of projects to position the firm into preferred market segments only to find later that current resources are heavily overbooked and will not be able to meet commercialization targets. Had strategic portfolio selection and resource planning been undertaken jointly, the portfolio would have been selected to align with resources. While a monolithic decision model that encompasses all major planning and process functions is unrealistic, closely related planning functions can indeed be coordinated. This is the key message of the rest of this paper.

5. Integration of capital budgeting and R&D project prioritization

Capital budgeting is an integral decision originating from the portfolio management function. It involves the selection of R&D and other projects that require long-term investments. Traditionally, reward-risk charting techniques (Cooper, Edgett, & Kleinschmidt, 1998) and discounted cash flow analysis have been deployed. However, Blau et al. (2004) show that this
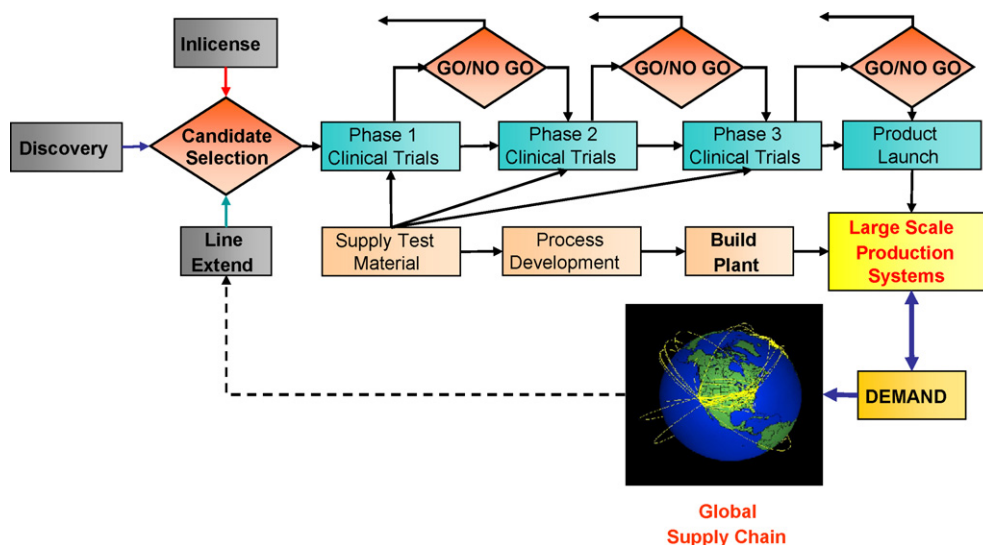

Fig. 2. An (early and advanced stage) R&D activity network of planning and process functions for a pharmaceutical enterprise.
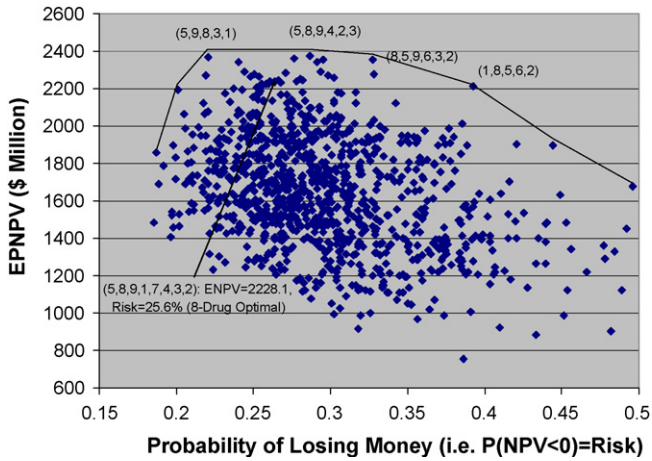
Fig. 3. An approximately efficient Reward-Risk frontier for a candidate set of nine pharmaceutical drugs (adapted from Blau et al. (2004)).

strategy leads to the selection of sub-optimal pharmaceutical portfolios evolving through a stage gate process shown in Fig. 2. Significantly improved financial performance is generated when capital budgeting concepts are integrated with R&D process models. This is because traditional capital budgeting models fail to incorporate physical resource constraints. Hence, either too small a portfolio may be selected which exposes the firm to higher technological, market and financial risks or too large a portfolio may be selected which will stretch the firm's resources in a way that will delay the launch of successful projects. Additionally, project inter-dependencies (Verma & Sinha, 2002) induce non-linear response of reward and risk measures to different portfolio compositions. Fig. 3 shows an approximately efficient economic frontier for a candidate set of nine drugs generated using combined genetic algorithm-discrete event simulation based algorithm architecture (Blau et al., 2004). Notice that the portfolio with the least risk and highest reward is a

five-drug portfolio as against larger portfolios which were recommended by pure financial charting techniques. This is because the five-drug portfolio includes drugs that synergize development features such as success probabilities, costs, durations as well as revenues. Yet the portfolio is small enough not to over-stretch physical resources along the R&D pipeline.

The network model of the enterprise as presented above represents hierarchical control architecture in which higher level planning functions act as supervisory controls on process level functions. Process level functions in turn regulate their operations based on internal decision models. However, the failure of planning models to incorporate the process level regulatory models generates possibilities of un-coordinated decisions. For example, R&D planners must contend with the problem of allocating capital to upstream R&D as well as downstream demands for plant capacity expansions. Thus, there is a need to model financial planning decisions (e.g. allocation of capital), R&D resource allocation and scheduling as well as downstream capacity expansion decisions within an integrated model. It is quite obvious that such a model would present formidable algorithm engineering challenges in view of its size and complexity. Pekny (2002) discusses algorithm architectures that could be used for solving such large-scale, complex cross-functional decision models.

6. Integration of dynamic R&D resource allocation and scheduling decisions

R&D projects can be accelerated by varying financial and physical resource allocation to process development activities. For example in the pharmaceutical industry, consolidated resource allocation to a set of drugs may enable the firm to acquire their clinical sample requirements earlier. Fig. 4 shows a Gantt chart for two resource allocation policies applied to a portfolio of four pharmaceutical drug development projects in
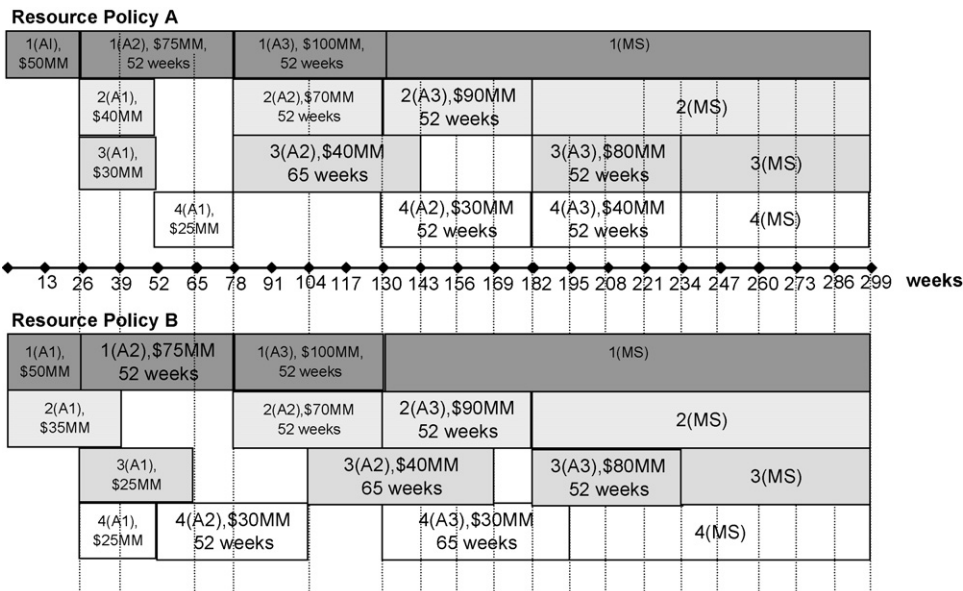

Fig. 4. Effect of adapting resource allocations on a four-drug portfolio.
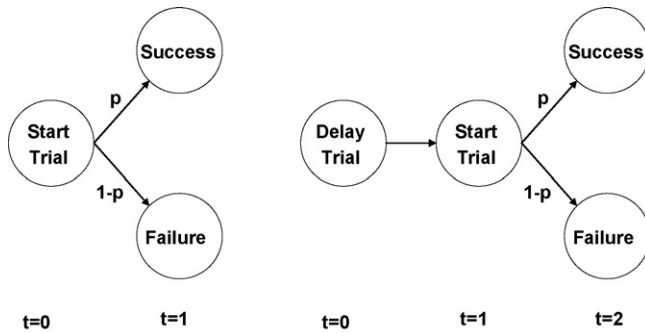
Fig. 5. Left exhibit shows evolution of uncertainty when clinical trial is started here-and-now. Right exhibit shows evolution of uncertainty when clinical trial is delayed until 1 period.

Phase III. The projects are designated as 1, 2, 3 and 4. Each drug must complete three activities AI, AII and AIII (e.g. dosage development, process development, plant construction) before it can be launched (activity MS). In resource allocation policy A, all projects are allocated standard resource levels while in resource policy B certain activities of some projects are allocated lower resources which leads to longer durations. However, the net effect is the re-arrangement of start times in a way that reduces the time to market of drug D4 to 195 weeks without delaying any other drug. A key challenge in the integration of resource allocation and R&D decisions is the issue of estimating process duration sensitivities to changes in the resource allocation levels. Though batch simulations and state-task network models (Kondili, Pantelides, & Sargent, 1993) can be used for that purpose, in doing so, we restrict the possible configurations of the batch network. A fully integrated model of R&D portfolio management and manufacturing processes would be able to select the optimal batch network configurations in response to the evolving state of the portfolio. Similarly, models for process development are required for estimation of sensitivity of its duration to the design strategy. Each design strategy would require a given level of resources.

It can be shown that such project management problems belong to the class of non-Markovian stochastic scheduling and assignment problems (Bertsekas, 2000). Scheduling an activity requires us to ensure that the activity was not started at any previous node of the relevant path from the root of the underlying tree of scenarios. The underlying scenario tree itself adapts structurally to decisions made at each node. A simple example is given in Fig. 5. The left exhibit shows a tree in which a clinical trial activity is scheduled at the current time. The right exhibit shows a tree in which the same activity is scheduled at time $t+1$ so that resolution of uncertainty has changed in response to the decision on starting the activity at time t. The point of this example is that the evolution of uncertain outcomes may be contingent upon decisions made in the past. Contrast this with a commodity's price evolving stochastically over four quarters in which case the price is not a function of production-inventory decisions. In short, business decision-making can either purely react to the evolution of uncertain trends or can actually affect those trends. Project and portfolio management problems are examples of the latter since project and portfolio decisions influence

the outcomes of those investments. For example, higher investments in certain projects may actually increase their probability of success. We see significant opportunities for PSE methods to be applied to such problems. Asset management problems (Goel & Grossmann, 2004) belong to this category since investments into assets like plants, oil fields, etc., can influence their productivity in uncertain ways. Rigorous approaches to such problems include neuro-dynamic programming (Bertsekas, 2000) that approximate value-to-go functions. Several complications arise: the "curse of dimensionality" which limits the number of value-to-go functions that can be stored and learned, the structural adaptation of the scenario tree, and so forth. Goel and Grossmann (2004) propose a Lagrangian decomposition based Branch and Bound approach based on scenario dis-aggregation for a multi-platform development problem in the oil exploration industry. Non-anticipativity constraints that link scenarios are dualized into the ENPV objective so that each scenario is associated with an MILP. The solution of this set of MILPs generates an upper bound at each node. Branching is performed on the violated non-anticipativity constraints. Nodal solutions infeasible with respect to some or all non-anticipativity constraints are transformed into feasible solutions. Global optimality has been demonstrated on industrial scale case examples. This seems as yet to be the only work in the area of non-Markovian decision-making in the PSE community.

Varma, Blau, Reklaitis, and Pekny (2005a) and Varma, Blau, Reklaitis, and Pekny (2005b) present an algorithm architecture for the stochastic scheduling and allocation problem which is non-Markovian.. The architecture searches for efficient R&D project scheduling and resource allocation policies defined over a policy space denoted by Π. The decision problem can be formally stated as follows:

$$\text{Maximize}_{\pi* \in \Pi} \text{ENPV}(\pi)$$

$$\text{Subject to } \text{Risk} \leq \beta_{\text{Risk}} \qquad \text{ATM}_i \leq \beta_{\text{ATM}}(i) \, \forall i \in \{1, \ldots, N\}$$
$$(2)$$

The architecture is flexible enough to handle any non-linear risk measures like Variance of NPV, semi-variance (downside risk), value-at-risk (the spread of returns between the lower 5% quartile of the NPV distribution and the mean value), probability of losing money, etc. The last constraint set limits to the Average time to Market (ATM) for each project. The rationale for limiting project ATMs emerges from loss of first-mover advantages (Cooper et al., 1999) as well as narrower sales windows for products like drugs that bear time-bound intellectual property right protections. The architecture shown in Fig. 6 (Varma et al., 2005a, 2005b), has three components: a process simulator which for an R&D pipeline is a discrete event simulator; a process optimizer which is typically a scheduling MILP with adaptive resource allocation decisions and a preference function learner. The process optimizer is developed and described in Varma, Uzsoy, Blau, and Pekny (in press). The architecture is used to learn efficient stochastic resource allocation policies, i.e., policies that supervise levels of resources to be allocated to pharmaceutical R&D projects as function of the portfolio state. The algorithm is summarized as follows:
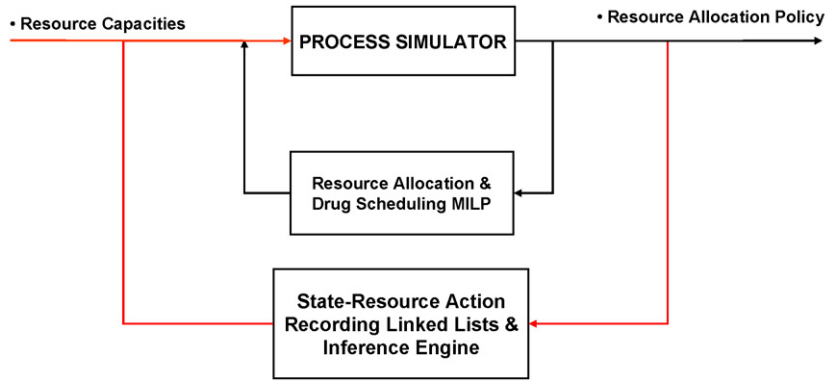
Fig. 6. The algorithm architecture for the stochastic scheduling and resource allocation problem (Varma et al., 2005a, 2005b).

- Step 1: simulation-based learning step.
  - Step 1.1: initialize N linked lists (LL) of drug "States" where state of a Drug $i = s = (DS_i, NLEV_i, NHEV_i)$ where DS_i = development stage of drug i; $NLEV_i$ = number of drugs with expected value lower than Drug i; $NHEV_i$ = number of drugs with expected value larger than Drug i. Declare a function $pf_i(s,m)$ = probability of selection by the MILP of resource allocation level 'm' in a state s for drug $i \in I$.
  - Step 1.2: at the end of each event "measure" the vector of drug states $(s_1, \ldots, s_n)$ and the vector of resource allocations selected by the MILP $(m_1, \ldots, m_n)$. For each drug still in the portfolio, check if the state occurred already exists in its $LL(i)$. If so, update its probability of selection. If not then add a "state-node" to $LL(i)$. At the end of all timelines for each linked list $LL(i)$, for each state s, calculate $pf^*(i,s) = \max_{m \in \{0,1,2\}} pf_i(s,m)$ and $m^*(i,s) = \text{argmax}_{m \in \{0,1,2\}} pf_i(s,m)$.
- Step 2: sort all $LL(i)$, $i = 1, \ldots, n$, in decreasing order of state occurrence probabilities.
- Step 3: simulate the policy stored in $\{LL(i): i = 1, \ldots, N\}$.

The architecture was run on the same nine-drug portfolio as the one used by Blau et al. (2004). It was assumed that the durations of process development and manufacturing activities will contract by 7.5% upon increase in the appropriate resource allocations by 15%. The architecture was used to test four R&D policies: Policy 0, genetic algorithm based static prioritization and non-adaptive resource allocation. The prioritization used was (5, 8, 9, 1, 7, 4, 3, 2, 6)—the economically efficient prioritization sequence determined by Blau et al. (2004); Policy 1, reactive scheduling and non-adaptive resource allocation; Policy 2, reactive scheduling and adaptive resource allocation; Policy 3, GA-based static prioritization and adaptive resource allocation. Adaptive resource allocation was performed based on the resource allocation policy learned using the architecture shown in Fig. 6.

The left exhibit of Fig. 7 depicts the project-wise ENPV comparisons for the four policies while the right exhibit shows the project-wise ATM comparisons for the four policies. Clearly, Policies 2 and 3 uniformly dominate Policies 0 and 1 across all projects. This analysis demonstrates the value of adaptive resource allocation. The main reason for this result is that in Policies 0 and 1 (non-adaptive resource allocation) fractional capacity remains un-utilized so that some projects are forced to wait until "standard" resource levels become available. In Policies 2 and 3 (adaptive resource allocation) fractional capacity is used to sustain the development of less prioritized projects albeit at a slower pace. This analysis clearly establishes that there is significant value embedded in these resource "switching" options.
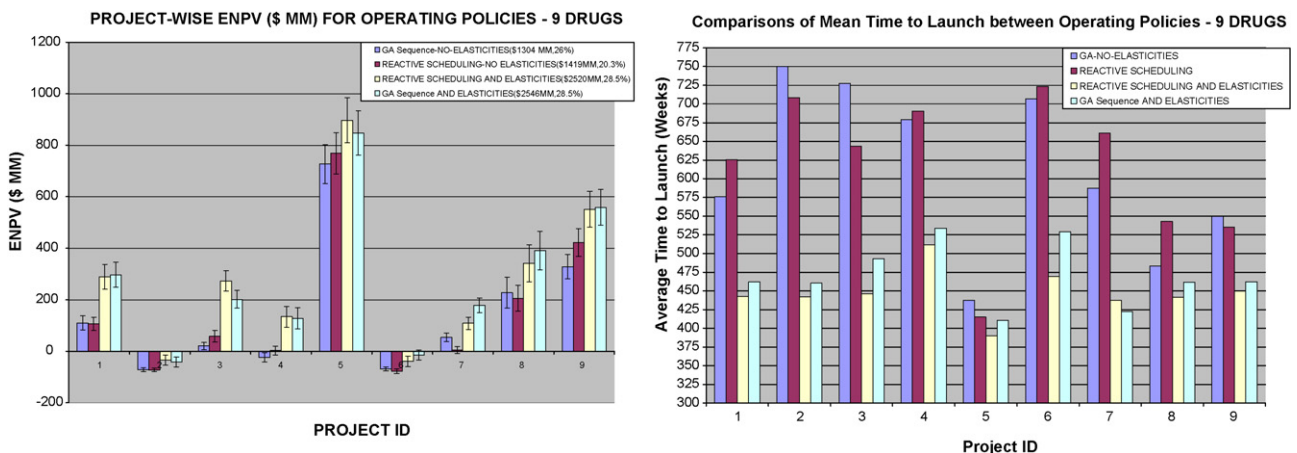

Fig. 7. Effect of dynamically adapting resource allocations on average times to market (ATM).

Further, the ENPV improvements are highlighted in the project ATMs shown in the right exhibit of Fig. 7. Clearly Policies 2 and 3 generate additional value by enabling significantly reduced average times to market. Interestingly, comparison between Policies 0 and 1 and between Policies 2 and 3 reveal that reactive scheduling tends to speed some projects and delay others. The GA-based sequence performs relatively well because of the genetic algorithm's ability to incorporate both reward and risk explicitly into its fitness function. On the other hand, the reactive scheduling MILP only implicitly models risk since non-linear risk terms cannot be incorporated. Nevertheless, this analysis shows the non-linear effects that can arise when adaptive resource allocation and scheduling concerns are modeled in an integrated fashion. Clearly, such models need to be further integrated with capacity planning and financial planning models.

Another domain that needs to be integrated with project management models like the ones discussed above is project financing. Consolidated financing of projects exacts several long-term costs that can over-ride any gains realized in the speedier execution of R&D projects. Large capital expenditures need to be financed by either the firm's liquidity reserves or using debt-equity channels (Smith, 1977). Debt-based capital is serviced by interest payout related costs while equity-based capital is serviced by dividend payouts. These capital service costs are long-term liabilities that can adversely impact individual project valuations (Modigliani & Miller, 1958). A case in point is the pharmaceutical industry where the capital demands keep fluctuating due to the progression of drug products through the pipeline. Further internal political constraints may prevent the firm from fully utilizing its working capital. This raises the specter of over-capitalizing the pipeline which can significantly add to the long term capital service costs. This in turn could lead to severe profitability downgrades due to excessive costs of capital. At the same time, under-capitalizing the pipeline will lead to a retarded R&D pipeline. Hence, the key question that emerges from this discussion is: how to optimally schedule the capitalization of the pipeline as a function of the evolving portfolio? Alternatively, under conditions of pipeline over-capitalization and physical resource limitations, physical capacity expansion needs to be considered. Such expansion would be financed by the excess capital and would enable the firm to speed some of its on-going drug development projects. However, the downside to this strategy is that the firm will begin to incur recurring costs of servicing the added resources by way of manpower as well as equipment maintenance costs. Hence, a model is required to resolve the trade-off between revenue growth generated by increased speed to market and higher physical resource service overheads. Such a model might also serve to determine the optimal capacity expansion plan.

Equity capital leads to floating service liabilities due to fluctuating dividend yields, unless the firm engages in a capital-intensive share buy-back to terminate these floating liabilities, in contrast to the more deterministic nature of debt instruments. Debt can be justified for financing high-risk drug projects on the basis of enabling accurate estimation of total debt servicing liabilities which is not possible with equity capital. Equity cap-
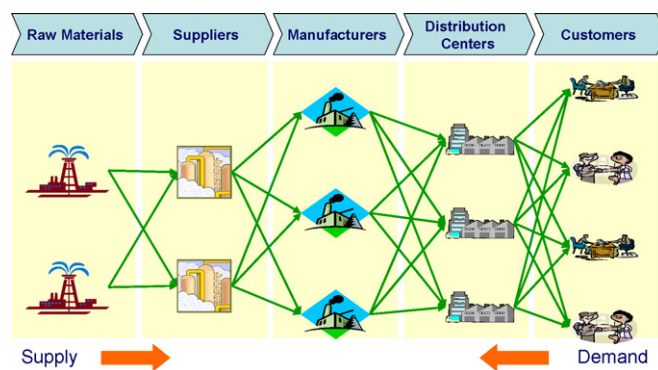

Fig. 8. An integrated supply chain network.

ital for such high-risk projects can be justified on the basis of providing the firm with the flexibility to lower dividend payouts in case the drug fails. Thus, another question would be: how to control the debt-equity ratio as a function of the R&D portfolio's evolution? A comprehensive review of debt and equity financing is available in Boyd and Smith (1998). Further what percentage of sales-generated internal accruals and other receivables should be channeled into the R&D pipeline's capitalization as a function of the portfolio's evolution? Similar questions can be raised in the context of the energy sector that demands large capital investments into natural resource development projects.

7. Integration of supply chain design and operations

Reklaitis and McDonald (2004) discuss the significance of integrated supply chain management (ISCM). Fig. 8 shows an integrated multi-stage supply chain network. Integrated ISCM captures strategic and tactical decisions such as raw material procurement contracts (e.g. hedging via future contracts), routing to plant sites, capacity planning and lead time management (production scheduling), routing of finished products, warehouse positioning, network inventory management and marketing strategies (e.g. revenue sharing contracts). A research goal is to integrate these decisions into one monolithic algorithm architecture. Clearly, much work needs to be done before such a goal can be achieved. In this section, we summarize frameworks that target subsets of these decisions. We start with hedging strategies for raw material cost reductions, and then discuss integration of production-inventory systems and finally integration of production-inventory and marketing strategies.

7.1. Integration of operational planning with currency and commodity risk-hedging

Large commodity chemical producers are constantly exposed to the risk of fluctuating raw material and feedstock prices. This is especially true for the petrochemical industry whose production economics is strongly dependent upon the volatile crude prices. Just as currency contracts (DeRosa, 1998) help companies limit their liquidity procurement costs; similar contracts

help companies limit their commodity procurement costs. These derivative instruments include forwards, futures and options on commodity futures. A commodity forward or future is a contract that obligates the firm to buy a certain quantity of raw material or feedstock that is traded at a pre-determined price known as the "strike" price (Andersen, 1987). A commodity call option is a contract that provides the firm with a right but no obligation to buy a certain quantity at a pre-determined price known as the "exercise price". Similarly, a commodity put option provides the firm with a right but no obligation to sell certain quantity at a pre-determined exercise price. The firm can enter into "long" positions (e.g. long commodity calls) by buying derivatives at a paid premium (debit) or can enter into "short" positions (e.g. short commodity calls) by selling at a gained premium (credit). Since the commodities are traded (e.g. at the Chicago/New York Mercantile Exchanges) their prices are normally assumed to follow multiplicative Brownian motion (Schwartz, 1997) just like stock prices and exchange rates. Using these derivative instruments, firms can hedge against the possibility of high future raw material and feedstock costs. How do these decisions impact process operations? For instance, in the absence of commodity contracts the firm may need to switch to coal based energy when natural gas prices spike. This might lead to enhanced production costs by way of additional costs incurred in pollution treatment operations as exhausts from coal-fired boilers need to be closely monitored and treated. Further, coal based energy generation capacity might need to be expanded at considerable capital investment. In contrast, a company can forego these costs by simply entering into a natural gas forwards contract (at a premium) which will allow the firm to operate its production facilities powered by natural gas even when its price spikes. Commodity forwards, futures and options (long positions) have limited downside by way of contract premiums but involve several decisions on parameters such as maturity dates and strike prices. The maturity date must synchronize with the production schedules. For instance, typical maturity dates must coincide with periods of increased production volumes. This means that commodity derivative pricing models need to be incorporated into production planning and scheduling models failing which will lead to either (1) increased inventory cost by way of premature procurement of raw material or feedstock resulting from short maturity times or (2) a complete ineffectiveness of the contract due to excessively long maturity times. Similarly, strike prices in case of commodity options can have significant impact on profitability. Excessively low strike prices will lead to significantly increased contract premiums especially when commodity prices remain relatively less volatile over the period of maturity. Excessively high strike prices tend to increase the likelihood of the commodity option not being exercised which again leads to ineffective utilization of the working capital component set aside for risk hedging purposes.

7.2. An example of integration of commodity and real options for improved support of process operations

Process plants face fluctuations in raw material and utility costs as well as demands. Often significant cost savings may be realized by operating the same plant in different modes in response to fluctuating utility costs and demands. In this section, we will demonstrate the cost reduction benefits of integrating commodity and real options using a production example.

Energy intensive process operations such as distillation are powered by steam boilers that run on various types of feedstock. The feedstock could be a price-stable feedstock such as coal or a price-volatile feedstock such as natural gas. At first sight coal is more attractive because of its lower cost and price stability. However, there is significant cost-incentive to employ natural gas since exhaust from coal-fired boilers needs to be subjected to extensive treatments before discharge. Moreover, it is possible to enter into natural gas contracts ahead of time in order to hedge against the possibility of price fluctuations. We demonstrate the potential benefits of deploying such a strategy towards operating cost reduction by developing a simple multi-stage stochastic programming model for a hypothetical firm.

A commodity chemicals firm operates N production facilities within a complex. Each facility is equipped with coal fired and/or natural gas fired power generation systems. Natural gas and coal feedstock are shipped to central holding facilities from where they are distributed on demand to the power plants for the N facilities. Further natural gas can be purchased from the open market or by using commodity options. Since coal price (pC) is expected to remain stable, no coal hedging contracts are investigated.

As shown in Fig. 9, the price of natural gas (pNG_k) follows a geometric Brownian motion represented by a binomial lattice where k denotes the index of a scenario node. We assume a finite planning horizon of T periods so that the depth of the scenario tree equals T. The key decision to be made is: how much of the energy requirement to source from natural gas and how much from coal at each node of the binomial lattice? In other words when to switch across feedstock sources and by how much? Further, the firm identifies the option to expand the power generation capacity sourced from either feedstock. However, the capacity expansion budget is constrained by the current debt and equity positions held by the firm. In the real options literature such capacity expansion options are known as "growth options" (Copeland & Antikarov, 2000). Further, the firm can buy options to source natural gas at a lower than market price which can avoid the cost of switching to coal. We call this option a "sourcing option"—a call option on natural gas futures. The firm acquires long positions on such options at the start of the planning horizon. The price of each such sourcing option is determined by Black's (Black, 1976) pricing formula and hence is a function of natural gas price volatility (σ), current price, strike price, and time to maturity. Strike prices and times to maturity are decision variables. The following list displays the other key parameters and decision variables that will be used within the model. The model is defined on the nodes of the binomial lattice:

- Index sets
  - *I*: set of plant sites;
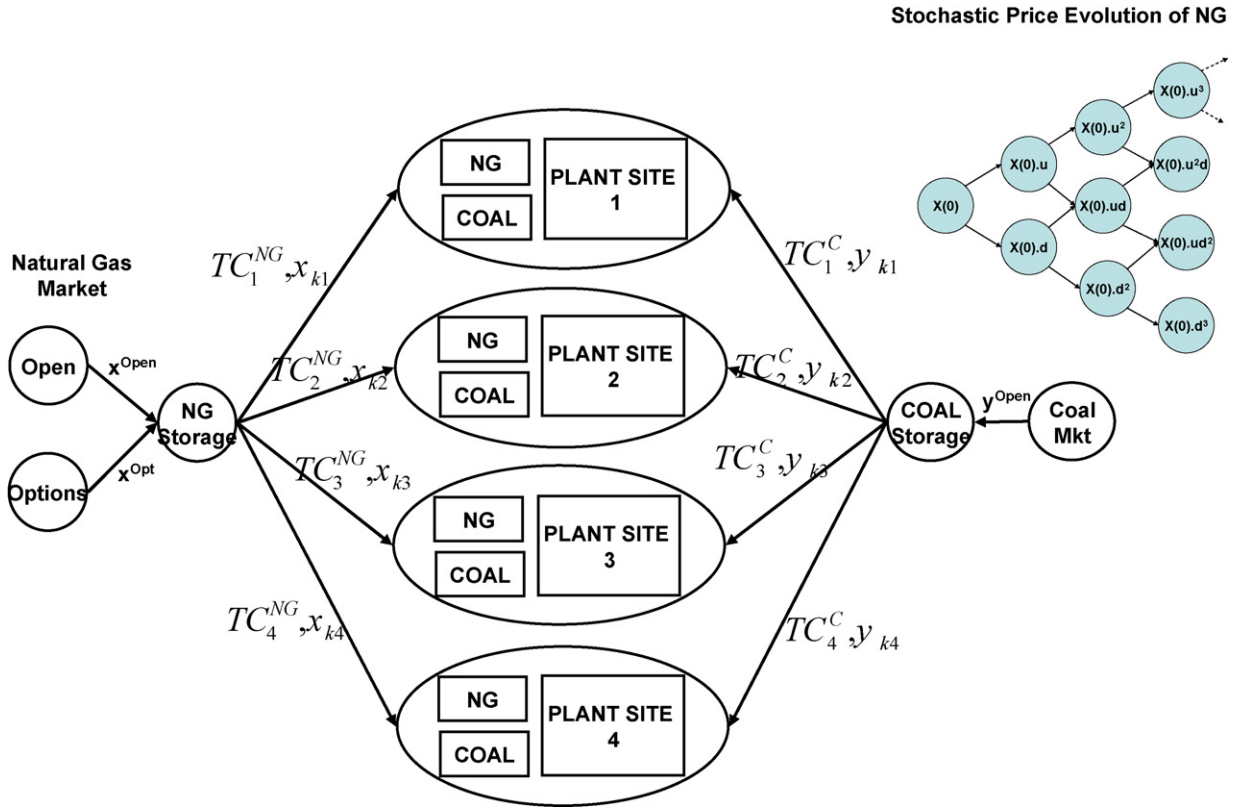 - *K*: set of scenario nodes on the binomial lattice.

Fig. 9. The energy feedstock supply network and the binomial lattice model representing the fluctuations in natural gas prices.

- Parameters
  - E_i = power demand (in kW) at plant site i on a quarterly basis;
  - B = capacity expansion budget;
  - pNG_k = price per mmBTU of natural gas at node k;
  - pC = price per mmBTU of coal;
  - TC_i^{NG} = NG transportation cost per mmBTU from storage to site i;
  - TC_i^C = coal transportation cost per mmBTU from storage to site i;
  - OC_i^{NG} = operating cost overheads per mmBTU of natural gas at site i;
  - OC_i^C = operation cost overheads per mmBTU of coal at site i;
  - HC^{NG} = storage cost per mmBTU of natural gas;
  - HC^C = storage cost per mmBTU of coal;
  - Cx_i = capacity expansion cost per kW of NG capacity added;
  - Cy_i = capacity expansion cost per kW of coal capacity added;
  - $Prob_k$ = probability of node k.
- Decision variables
  - x_{ki} = kW sourced from natural gas generator at site i in scenario node k;
  - x_k^{Opt} = mmBTU of NG sourced by exercising sourcing options at node k;
  - x_k^{Open} = mmBTU of NG sourced from open market at node k;

- y_{ki} = kW sourced from coal generator at site i in node k;
- y_k^{Open} = mmBTU of coal purchased at start of $t(k)$ at node k;
- I_k^{NG} = mmBTU of natural gas left unutilized at the end of period $t(k)$;
- I_k^{Coal} = mmBTU of coal left utilized at the end of period $t(k)$;
- EX_{ki} = capacity (in kW) added to NG-plant at node k;
- EY_{ki} = capacity (in kW) added to COAL-plant at node k;
- $CapX_{ki}$ = NG capacity (in kW) at node k and site i;
- $CapY_{ki}$ = coal capacity (in kW) at node k and site i;
- N_t = number of sourcing options with time to maturity of t periods;
- $pStrike_{NG,t}$ = strike price per mmBTU of natural gas for sourcing option expiring at the start of time period t;
- V_t^{Opt} = total price of NG sourcing options maturing at start of time period t.

$$
\begin{aligned}
\text{Min} \sum_{t=1}^{T} V_t^{Opt} + \sum_{k=1}^{K} \text{Prob}_k \Big\{ & (pNG_k x_k^{Open} + pStrike_{NG,t} x_k^{Opt}) \\
& + pCy_k^{Open} + HC^{NG} I_k^{NG} + HC^C I_k^C + \sum_{i=1}^{N} \{(OC_i^{NG} \\
& + TC_i^{NG}) x_{ki} + (OC_i^C + TC_i^C) y_{ki} + Cx_i EX_{ki} + Cy_i EY_{ki} \} \Big\}
\end{aligned}
\tag{3}
$$

subject to

$$V_t^{\text{Opt}} = N_t(\text{pNG}_0\varphi(d_1) - \text{pStrike}_{\text{NG},t}e^{-r_{\text{free}}t}\varphi(d_2))$$

$$d_1 = \frac{\ln(\text{pNG}_0/\text{pStrike}_{\text{NG},t}) + (\sigma^2/2)t}{\sigma\sqrt{t}}$$

$$d_2 = \frac{\ln(\text{pNG}_0/\text{pStrike}_{\text{NG},t}) - (\sigma^2/2)t}{\sigma\sqrt{t}} \qquad (4)$$

$$I_{k'}^{\text{NG}} + x_k^{\text{Opt}} + x_k^{\text{Open}} = 7.37\sum_{i \in I}x_{ki} + I_k^{\text{NG}}, \; \forall k \in K$$

$$I_{k'}^{\text{C}} + y_k^{\text{Open}} = 7.37\sum_{i \in I}y_{ki} + I_k^{\text{C}}, \; \forall k \in K$$

$$x_{ki} + y_{ki} = E_i, \; \forall i \in I, \; k \in K \qquad x_k^{\text{Opt}} \leq N_{t(k)}, \; \forall k \in K \qquad (5)$$

$$\text{CapX}_{ki} = \text{CapX}_{k'i} + \text{EX}_{k'i}, \; \forall i \in I, \; k \in K$$

$$\text{CapY}_{ki} = \text{CapY}_{k'i} + \text{EY}_{k'i}, \; \forall i \in I, \; k \in K \qquad (6)$$

$$\sum_{i=1}^{N}\text{Cx}_i\text{EX}_{ki} + \sum_{i=1}^{N}\text{Cy}_i\text{EY}_{ki} \leq B, \; \forall k \in K \qquad (7)$$

$$0 \leq x_{ki} \leq \text{CapX}_{ki}, \; 0 \leq y_{ki} \leq \text{CapY}_{ki}, \; \forall i \in I, \; k \in K$$

$$\text{EX}_{ki} \geq 0, \; \text{EY}_{ki} \geq 0, \; \forall i \in I, \; k \in K \qquad (8)$$

The objective function (Eq. (3)) is the sum of the total 'here-and-now' cost incurred to acquire natural gas sourcing option contracts and the second stage expected costs. The second stage expected cost is the sum of the natural gas and coal transportation and procurement costs (by way of exercising sourcing options as well as from the open market), the operating overheads and the capacity expansion costs. Eq. (4) represents Black's pricing (Black, 1976) constraint. These constraints could be eliminated if the strike prices were fixed. Eq. (5) represents the energy balance constraints: The sum of total feedstock inventory from the previous period's node (k'), total inventory sourced from the open market and options must equal the total feedstock utilization and any unutilized inventory, the total natural gas and coal utilization must equal the total energy demand and the natural gas mmBTUs sourced from options must be less than the total number of options for any period since each natural gas sourcing option supplies a single mmBTU. Eq. (6) represents the capacity balance constraints in order to incorporate any intermediate capacity expansions. Eq. (7) represents the budget constraints on capacity expansions. Eq. (8) represent that the total energy supply from natural gas and coal must not exceed its physical generation capacity. The above formulation is the so-called explicit form of a multi-stage stochastic programming formulation. The formulation is an MINLP due to the highly non-linear Black-Scholes pricing constraint and demonstrates the challenges in solving formulations that involve pricing and other financial considerations.

We display the non-linear effects of this problem by solving a $T = 4$-period and $N = 4$-site case instance shown in Fig. 9.

Each period stands for a quarter (=13 weeks). We obtained the price of natural gas futures trading on New York Mercantile Exchange (NYMEX[TM]) closing February 28, 2005 at a spot rate of $6.20 per mmBTU. For purposes of demonstration we fix the strike price at 5.15 per mmBTU for all the four quarters (in-the-money). The short-term volatility rate of natural price movements is estimated at $0.1 per quarter from NYMEX historical pricing data (www.nymex.com). The firm negotiates for delivery of feedstock commodities at the start of each quarter. It is assumed that the firm uses North Appalachia (NAP) coal trading at $62 per ton with a calorific value of 11 mmBTU/ton (www.eia.doe.gov). Hence, the price of coal which is assumed to remain stable is $5.63 per mmBTU. Further, in each period each of the plant sites demand 10,000 kW of energy (1 kW = 56.9 BTU/min). Assume that the initial natural gas based power plant capacity is 8,000 kW and the initial coal based power plant capacity is 2,000 kW. For simplicity an average power generation expansion cost of $12 per kW is assumed from historical data. The total power plant capacity expansion budget is assumed at $600,000. For all scenarios the transportation costs of natural gas (from the Sabine Pipeline Co.'s Henry Hub in Louisiana) are assumed to be constant at $0.45, $0.5, $0.1 and $0.4 per mmBTU to the four plant sites, respectively. For all scenarios the transportation costs of coal (from various locations within the U.S.) are assumed to be constant at $0.25, $0.35, $0.15, $0.30 per mmBTU to the four plant sites, respectively. With this data the above multi-stage stochastic program was solved for a long portfolio of 200,000 natural gas European calls for periods $t = 1$, $t = 2$ and $t = 3$. The formulation was solved using ILOG Cplex[TM] Optimization Suite. Fig. 10 shows the firm's optimal capacity expansion and power switching policies for this option portfolio. The total expected energy cost which includes the option premiums paid upfront as well as the transportation, feedstock procurement costs = $22.40 Million. A significant observation of the solution is that even when the natural gas price moves up in a large number of scenario nodes, natural gas is still used. This is because of the availability of 200,000 natural gas options that help the firm hedge against the natural gas spikes. The premiums for the option maturing at time $t = 1$ (start of Quarter 1) is $260,615, for option maturing at time $t = 2$ is $305,332, the option maturing at time $t = 3$ is $345,085. As the uncertainty increases the option premium increases. Had these options not been available the total energy cost would have been $22.90 Million. Further, the volatility of the total energy cost for the plan with the options is $0.055 million while that for the plan without options is $0.1 million.

Fig. 11 shows the non-linear energy cost vs. strike price curves for different volumes of options. We assume that all options are bought at the same strike price. The trend for each volume of options is explained as follows: At low strike prices (in-the-money) the likelihood of exercising the option is high across all scenario nodes and the energy cost savings justify the higher option premiums. However, at very high or out-of-money strike prices the likelihood of exercising the option is low and renders the option contract ineffective against hedging pricing risk. In the above case, all options with strike prices higher than the current market price are totally ineffective. The fig-
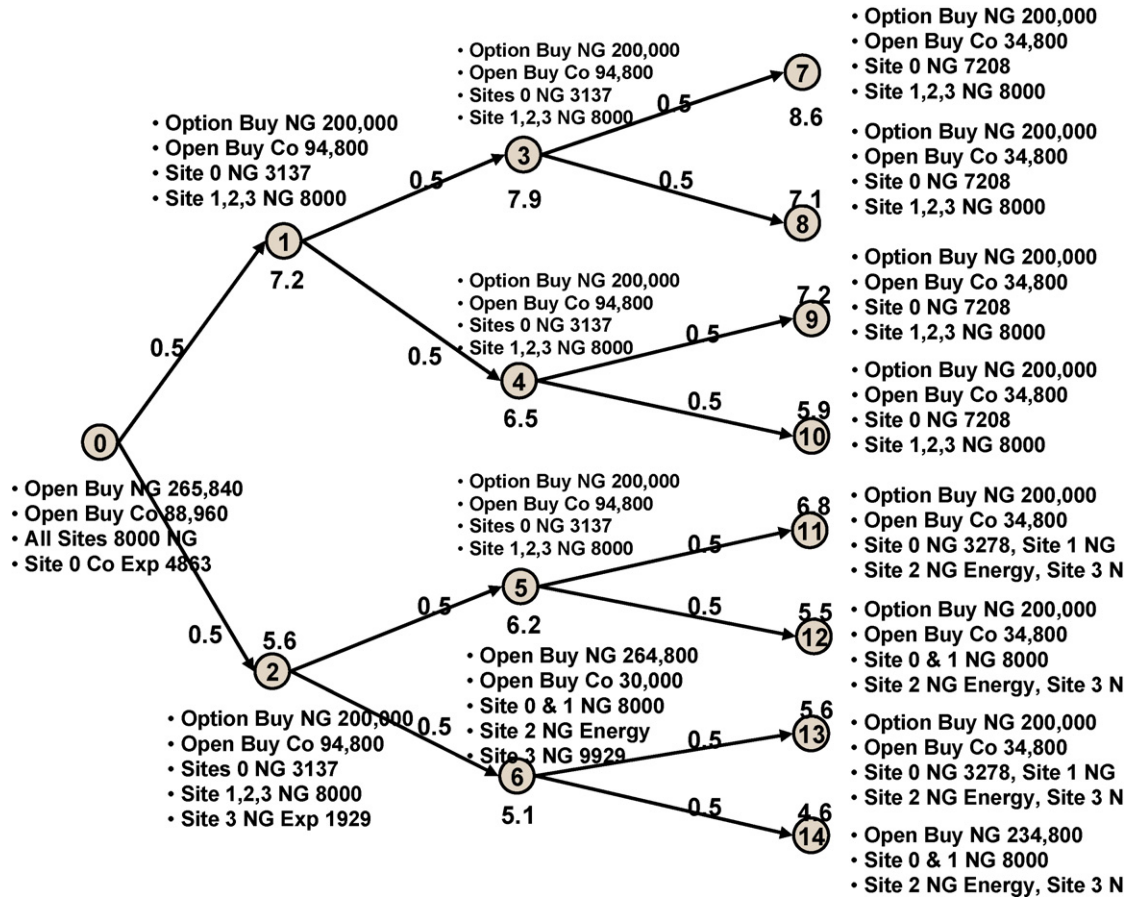
Fig. 10. The scenario tree showing the energy policy of the enterprise—adaptive capacity expansion, energy option trading and feedstock switching.

ure was generated by fixing strike prices and volumes and then applying the energy optimization formulation. The least annualized expected energy cost occurs by buying 300,000 options with strike price set at 30% of current market price i.e. $2.48

per mmBTU. This means that the options have to be deep-in-the-money which leads to high premiums. Also, buying options at-the-money is expected to result in the highest energy cost. It is possible in the above example that an option contract with
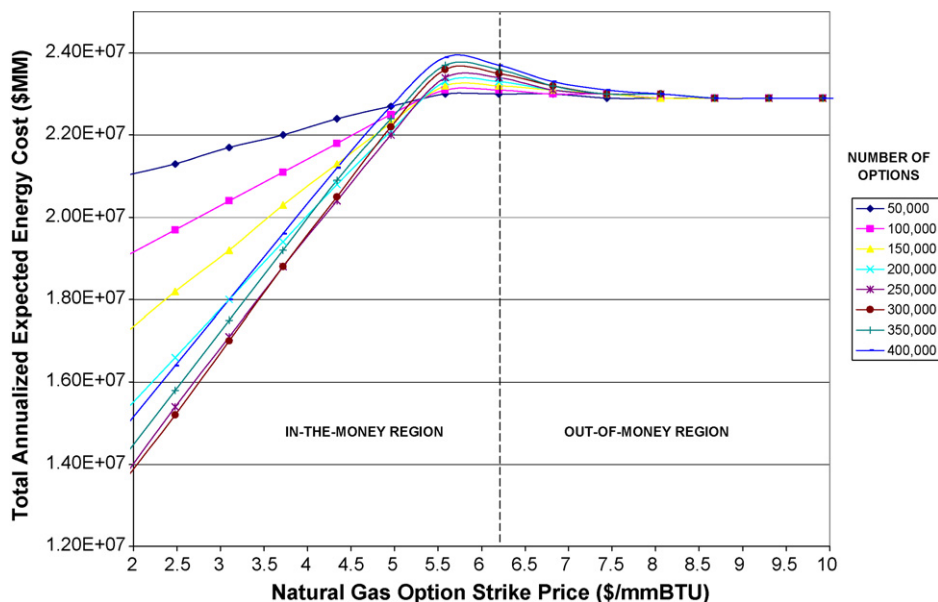

Fig. 11. Total annualized expected energy cost vs. strike price for different volumes of natural gas options (cmp = current market price of natural gas = $6.2 per mmBTU on the NYMEX[TM] 28 February, 2005).
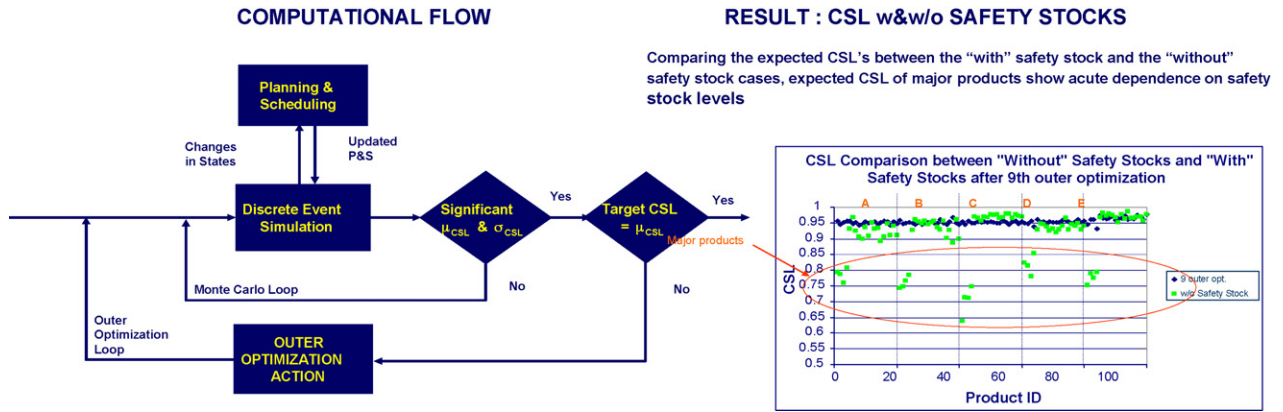
Fig. 12. Left exhibit: the dual loop architecture for optimizing multi-stage network safety stocks and production policies. Right exhibit: the gains in multi-product customer satisfaction levels (CSL) emerging from integration of safety stock production management (Jung et al., 2004).

much higher strike prices but lower annualized energy cost than the current best may exist. Clearly, better search methods are required to explore the highly non-linear space of strike prices instead or enumeration. Pricing formulas tend to be non-convex which implies that local MINLP methods are inadequate for identification of global optima. Future research needs to focus on solving this class of energy policy optimization problem with arbitrary commodity pricing formulas.

7.3. Integration of production and network inventory decisions

Safety stocks provide intermediate and product inventory buffers to hedge against the risk of demand uncertainties. Jung et al. (2004) propose an algorithm architecture that integrates supply chain design by way of warehouse positioning decisions and network safety stock planning with production-inventory planning and scheduling. Their framework solves an aggregate planning model for warehouse positioning. Output from the aggregate model serves as input to a dual loop algorithm architecture shown in the left exhibit of Fig. 12. The outer or the "strategic" loop performs a non-linear stochastic gradient search in the space of multi-product network safety stocks. The inner loop simulates a multi-stage production-inventory network for given network safety stocks while ensuring that production schedules adapt to production and demand realizations. Such regulatory measures are enabled using a multi-period deterministic planning and scheduling MILP model. The architecture thus integrates production-scheduling (operational) decisions while optimizing safety stock (strategic) decisions. The problem can be stated formally as:

$$\text{Minimize}_{\pi* \in \Pi, \theta \in \Theta} \sum_{i=1}^{N} \mu_i(\text{CSL}_i(\theta_i, \pi) - \text{CSL}_i^{\text{Target}}) \quad (9)$$

$$\text{Subject to } \text{CSL}_i(\theta_i, \pi) \geq \text{CSL}_i^{\text{Target}} \, \forall i \in \{1, \ldots, N\}$$

The optimization is performed over the space of network safety stock levels denoted by Θ (handled in the outer loop) and the space of all production planning policies Π (handled in the inner loop). The objective is to minimize the weighted devia-

tions of customer satisfaction levels for N products from their target levels while ensuring that the target levels are met. The right exhibit of Fig. 12 demonstrates significant improvement achieved for the CSLs of all products involved in a large-scale industrial case study using this architecture over the case without any network safety stocks.

7.4. Integration of production and dynamic capacity planning decisions

The next example of enterprise wide planning emerges from the area of multi-product capacity expansion under competitive uncertainties. Excessive capacity addition can lead to unutilized capacity especially when competitor products capture some of the market share. Fig. 13 shows the scenario in which a competitor product arrives into a key market segment which forces the firm to adapt its production system to the reduced demand. Thus, the question is: What must be the firm's capacity expansion or growth option strategy to hedge against the risk of future unutilized capacity? Wan, Reklaitis, and Pekny (2006) formulate this problem as a dynamic stochastic optimization problem with the objective of identifying the optimal capacity expansion strategy. The problem is formulated as a stochastic dynamic program and solved using a neuro-dynamic programming approach in which value-to-go functions are learned using a hybrid Least Squares Support Vector Machine (LSSVM) based architecture (Wan, Reklaitis, & Pekny, 2005). In the first phase, the architecture learns the value-to-go functions at each node of the scenario tree. In the second stage a roll-back of the tree is performed to determine the optimal policy. The left exhibit of Fig. 14
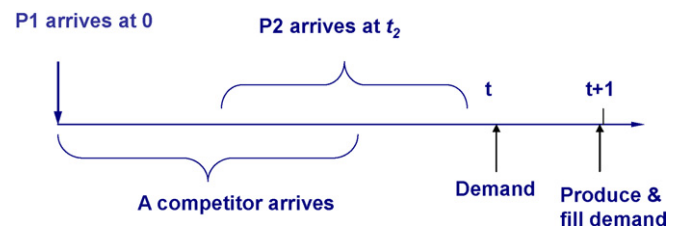

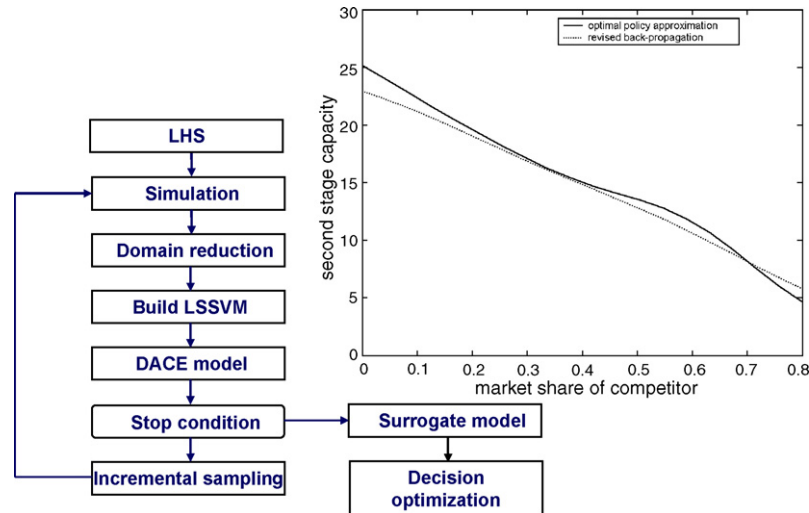Fig. 13. Capacity planning under competitive uncertainties.

Fig. 14. Left exhibit: the LSSVM based capacity management architecture. Right exhibit: the optimal control law of capacity as a function of market share (Wan et al., 2006).

shows the architecture while the right exhibit shows the optimal variation of the second stage capacity for a case study as a function of the competitor's market share as determined by the architecture.

7.5. Integration of marketing strategy and production inventory process

Production-Inventory P&C and marketing strategy functions are linked much more directly. Marketing strategy forecasts short-term demands based on competition and pricing models. The forecasts trigger production-inventory P&C models which output decisions such as multi-product manufacturing volumes, safety stocks, plant and warehouse inventories and even capacity plans. These decisions flow into the production process functions where they trigger short to medium term scheduling and plant operational models. Since plant operations are subject to adverse events such as batch failures or weeping/flooding of distillation columns that reduce product purity, unusual exothermicity generated hot spots in fixed bed reactors etc. production inventories may deviate from target levels. Such deviations and other process information are communicated to the inventory data warehousing function. It is a challenge to the market strategy function to incorporate feedback information about inventory deviations into its strategy models. Strategy models must not only incorporate external market forecast and competition related information but also forecasts of deviations from production targets. For instance, the market strategy model does not need to learn about the probability of a certain distillation column flooding in an operating quarter (13 weeks), however, it is sufficient to learn the probability distribution of the deviation of the product's actual volume from the target volume. The production volumes specified by the marketing strategy models will tend to be higher than those based on pure market forecasts which will compensate for any production mal-functions. Similarly, real time assessments of inventory depletion rates at various warehouses in the supply chain network can allow the

Production-Inventory P&C to better co-ordinate transportation resources across the network.

Supplier–customer contracts play key roles in defining the enterprise market strategy. The de-centralized SC literature has focused on certain types of contracts such as revenue-sharing, QF contracts to reduce overall supply chain costs to all agents. Some of these contracts include mutually coordinated parameters such as supplier-to-retailer pricing, quantity and salvage volumes of unsold inventories. Despite its importance, the literature integrating production-inventory planning and marketing strategy is sparse. An example is the work of Subramanian (2004) who optimize parameters of a supplier–retailer QF-contract for a grocery supply chain with full consideration of detailed production-inventory models for both the entities. The left exhibit in Fig. 15 shows their co-ordination network between the entities. As shown in the right exhibit in Fig. 15 they identify the "win-win" regimes that result in the largest increase in profitability of both entities. This contribution demonstrates the emergence of computational tools like multi-agent simulations to solve problems involving entities with complex behaviors as determined by their attempts to adapt their respective production-inventory systems. These kinds tools are better suited to deal with industrial problems than more theoretical approaches such as the classical game theory which fails to capture such complex behaviors.

8. Outlook on computational strategies

A single monolithic model that can used to jointly and efficiently optimize each and every enterprise decision is unlikely to exist for the foreseeable future. Even if such a model were to exist, decomposition strategies would have to be identified and implemented within appropriate algorithm architectures. Architectures (Pekny, 2002) based on appropriate algorithm components promise significant potential for efficiently and in some cases optimally solving enterprise wide planning problems. In general, optimal algorithm architectures include
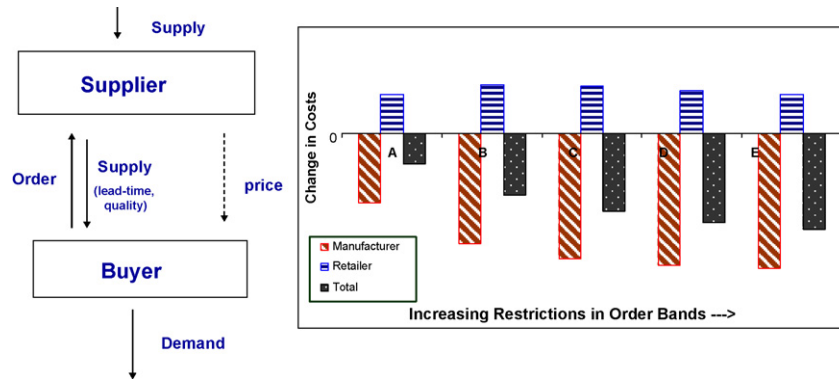
Fig. 15. De-centralized supply chains: left exhibit shows a supplier–retailer co-ordination network. The right exhibit shows the effect of a QF-Contract parameter on operating costs (Subramanian et al., 2004).

synergistic combinations of the best search features of individual algorithm components. In our view, formalisms for identifying efficient algorithm architectures for enterprise-wide planning are in a nascent state of development. Research directed at developing such formalisms would need to address the following questions:

1. Which algorithm components should be selected for the given decision problem based on the specific types of decisions and uncertainties? For example, an algorithm architecture might address multi-stage R&D capacity expansion decisions. Varma et al. (2005b) show that this capacity planning problem can be cast as a non-linear integer program wherein the non-linearities arise from the dependencies of queuing statistics on capacities. The objective is to minimize total cost of capacity expansion while ensuring that products are not unduly delayed. A branch-and-bound algorithm is subsequently proposed which efficiently solves realistic problems to global optimality. This strategy is valid only so long as no resources are shared across multiple projects and that resource allocations to projects are fixed. If these assumptions are relaxed, the search space also includes multi-stage resource allocation decisions in addition to capacity expansion decisions. Since the Branch and Bound algorithm is effective for capacity planning, one could synthesize an algorithm architecture consisting of two modules: a resource allocation module driven by a simulated annealing based search in the space of discrete resource allocations and a capacity expansion module driven by the branch and bound algorithm. Each point explored in the space of resource allocations would be "submitted" to the branch and bound driven capacity module. The challenge for the algorithm designer is: how to utilize information from the capacity module's branch and bound to improve the search efficiency of the simulated annealing based resource allocation module. This example shows the criticality of selecting appropriate algorithm components and mining the information to reinforce the performance of each component.
2. How should the algorithm components be combined into a computationally efficient engine? A related question is: how should the information flow across algorithm components be

controlled in order to achieve a given computational complexity or a given solution quality? The potential pitfall is an "assembly-line" approach to algorithm architecture design in which components are simply bound together. Instead, algorithm components need to be seamlessly integrated and as far as possible event driven, i.e., appropriate components should invoke their functionality in response to demands from other components. Event driven architectures usually limit the computational resource demands.
3. How to build the "algorithmic occam's razor", i.e., parsimonious architectures that efficiently optimize as many enterprise-wide decisions as possible using the least computational resources?

However, to date no comprehensive architectures even exist that integrate project selection, resource allocation and capacity management in the R&D planning area. Architectures are also required for integration of production planning, scheduling, inventory planning and capacity management. These are extremely inter-twined decision problems with highly non-linear and combinatorial decision interactions under production, demand, pricing and competitive uncertainties. Clearly, progress in addressing these problems is badly needed if enterprise-wide optimization is to be achieved, and the "algorithm architectures" approach seems efficient, reliable, robust and adaptive to practitioner concerns.

9. Conclusions and outlook on enterprise-wide modeling

This paper demonstrates the significant research potential in building optimization models and developing algorithms for solving enterprise-wide problems. We presented a conceptual model that views the enterprise as a network of planning and process functions. Each function is associated with a decision model which regulates its operations. The current model architecture suggests that operational planning functions are supervised by strategic planning functions. In turn process functions are supervised by operational planning functions. We point out major drawbacks of this architecture and provide several examples that demonstrate integration between financial and operational decision models. These examples include integration of capi-

tal budgeting (portfolio selection) and R&D pipeline models, integration of scheduling and resource allocation models in the context of controlling project development durations, and integration of commodity derivatives and production-inventory models. However, these are only modest attempts at integrating a small subset of enterprise-wide decision models. Much work remains to be done to target methodologies for coordinating decision models or even building holistic cross-functional models. The algorithmic infrastructure developed by the process systems engineering (PSE), strategic finance and operations research communities should form the basis for research in large-scale enterprise-wide modeling and optimization.

References

Acevado, J., & Pistikopoulos, E. N. (1997). A hybrid paramtric/stochastic programming approach for mixed integer linear programs under uncertainty. *Industrial and Engineering Chemistry Research*, *36*, 2262–2270.

Adler, P. S., Mandelbaum, A., Nguyen, V., & Schwerer, E. (1995). From project to process management: An empirically based framework for analyzing product development time. *Management Science*, *41*(3), 458–484.

Ahmed, S., Tawarmalani, M., & Sahinidis, N. V. (2000). A finite branch and bound algorithm for two stage stochastic integer programs. Working paper, School of Industrial and Systems Engineering, Georgia Institute of Technology.

Amram, M., & Kulatilaka, N. (1999). *Real options: Managing strategic investment in an uncertain world*. Boston, Massachusetts: Harvard Business School Press.

Andersen, T. J. (1987). *Currency and interest rate hedging—a user's guide to options, futures, swaps and forward contracts*. New York Institute of Finance Corp..

Applequist, G. E., Pekny, J. F., & Reklaitis, G. V. (2000). Risk and uncertainty in managing chemical manufacturing supply chains. *Computers and Chemical Engineering*, *24*, 2211–2222.

Badell, M., Romero, J., Huertas, R., & Puigjaner, L. (2004). Planning, scheduling and budgeting value-added chains. *Computers & Chemical Engineering*, *28*, 45–61.

Barbaro, A. F., & Bagajewicz, M. (2004). Use of inventory and option contracts to hedge risk in planning under uncertainty. *AIChE Journal*, *50*(5), 990–998.

Baroum, S., & Patterson, J. (1996). The development of cash flow weight procedures for maximizing the net present value of a project. *Journal of Operations Management*, *14*(3), 209–228.

Bertsekas, D. (2000). *Dynamic programming and optimal control* (2nd ed.). Belmont, MA: Athena Scientific.

Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics*, *3*, 167–179.

Black, F., & Scholes, M. S. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, *81*(3), 637–654.

Blau, G. E., & Bunch, P. (2002). *Process modeling in new product development: PDMA new product development toolbook*. New York: John Wiley & Sons.

Blau, G. E., Mehta, B., Bose, S., Pekny, J. F., Sinclair, G., Kuenker, et al. (2000). Risk management in the development of new products in highly regulated industries. *Computers and Chemical Engineering*, *24*(7), 659–664.

Blau, G. E., Pekny, J. F., Varma, V. A., & Bunch, P. (2004). Selecting a portfolio of dependent new product candidates in the pharmaceutical industry. *Journal of Product Innovation Management*, *21*(4), 227–245.

Bok, J., Heeman, L., & Park, S. (1998). Robust investment model for long-range capacity expansion of chemical processing networks under uncertain demand forecast scenarios. *Computers and Chemical Engineering*, *22*(7), 1037–1049.

Bose, S., & Pekny, J. F. (2000). A model predictive control framework for planning and scheduling problems: A case study of consumer goods supply chain. *Computers and Chemical Engineering*, *24*(2–7), 329–335.

Boyd, J., & Smith, B. (1998). *The evolution of debt and equity markets in economic development*, *12*, 519–560.

Brucker, P. (1995). *Scheduling algorithms*. Berlin: Springer.

Carlson, E. C., & Felder, R. M. (1992). Simulation and queuing network modeling of single-product production campaigns. *Computers and Chemical Engineering*, *16*(7), 707–718.

Chapman, C., & Ward, S. (1996). *Project risk management: Processes, techniques and insights*. Canada: Wiley.

Clay, R. L., & Grossmann, I. E. (1997). A disaggregation algorithm for the optimization of stochastic planning models. *Computers and Chemical Engineering*, *21*(7), 751–774.

Cooper, R. G., Edgett, S. J., & Kleinschmidt, E. J. (1998). Best practices for managing R&D portfolios. *Research Technology Management*, 20–33.

Cooper, R. G., Edgett, S. J., & Kleinschmidt, E. J. (1999). New product portfolio management: Practices and performance. *Journal of Product Innovation Management*, *16*(4), 333–351.

Copeland, T. C., & Antikarov, A. (2000). *Real options: A practitioner's guide*. Texere Publishing Limited.

Davis, E. W., & Patterson, J. H. (1975). A comparison of heuristic and optimum solutions in resource constrained project scheduling. *Management Science*, *21*(8), 944–955.

Demeulemeester, E., & Herroelen, W. (1992). A branch and bound procedure for the multiple resource-constrained project scheduling problem. *Management Science*, *38*(12), 1803–1818.

Demeulemeester, E., & Herroelen, W. (1997). A branch and bound procedure for the generalized resource-constrained project scheduling problem. *Operations Research*, *45*(2), 201–212.

DeRosa, D. F. (1998). *Currency derivatives: Pricing theory, exotic options and hedging applications* (2nd ed.). NY: John Wiley & Sons.

Doersch, R. H., & Patterson, J. H. (1977). Scheduling a project to maximize its net present value: A zero-one programming approach. *Management Science*, *23*(8), 882–889.

Ding, M., & Eliashberg, J. (2002). Structuring the new product development pipeline. *Management Science*, *48*(3), 343–363.

Elmaghraby, S. E. (1997). *Activity networks: Project planning and control by network models*. New York: Wiley Science.

Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. New York: Wiley.

Goel, V., & Grossmann, I. E. (2004). Gas field development planning under uncertainty. AIChE Annual Meeting, Paper 407b, Austin, TX.

Guillen, G., Badell, M., Espuna, A., & Puigjaner, L. (2006). Simultaneous optimization of process operations and financial decisions to enhance the integrated planning/scheduling of chemical supply chains. *Computers & Chemical Engineering*, *30*, 421–436.

Gupta, A., Maranas, C. D., & McDonald, C. M. (2000). Mid-term supply chain under demand uncertainty: Customer demand satisfaction and inventory management. *Computers and Chemical Engineering*, *24*(12), 2613–2621.

Hartmann, S. (1997). Project scheduling with multiple modes: A genetic algorithm. *Annals of Operations Research*, *102*, 111–135.

Herroelen, W., & Leus, R. (2005). Project scheduling under uncertainty, survey and research potentials. *European Journal of Operations Research*, *165*(2), 289–306.

Honkomp, S. J., Reklaitis, G. V., & Pekny, J. F. (1997). Robust planning and scheduling of process development projects under stochastic conditions. AIChE Annual Meeting. Los Angeles, CA.

Ierapetritou, M., Acevedo, J., & Pistikopoulos, E. N. (1996). An optimization approach for process engineering problems under uncertainty. *Computers and Chemical Engineering*, *20*(6/7), 703–709.

Ierapetritou, M., & Pistikopoulos, E. N. (1995). Novel approach for optimal process design under uncertainty. *Computers and Chemical Engineering*, *19*(10), 1089–1110.

Icmeli, O., & Erenguc, S. (1996). A branch and bound procedure for resource constrained project scheduling problem with discounted cash flows. *Management Science*, *42*(10), 1395–1408.

Icmeli, O., & Erenguc, S. S. (1996). The resource-constrained time/cost trade-off project scheduling problem with discounted cash flows. *Journal of Operations Management*, *14*(3.), 255–275.

ILOG. http://www.ilog.com/. ILOG CPLEX Solver Engine.

Jain, V., & Grossmann, I. E. (1999). Resource constrained scheduling of tests in new product development. *Industrial and Engineering Chemistry Research*, *38*, 3013–3026.

Jung, J. Y., Blau, G. E., Pekny, J. F., Reklaitis, G. V., & Eversdyk, D. (2004). A simulation based optimization approach to supply chain management under demand uncertainty. *Computers and Chemical Engineering*, *28*(10), 2087–2106.

Kahn, K. B. (2002). An exploratory investigation of new product forecasting practices. *Journal of Product Innovation Management*, *19*(2), 133–143.

Kolisch, R. (1996a). Efficient priority rules for resource-constrained project scheduling problem. *Journal of Operations Management*, *14*(3), 179–192.

Kolisch, R. (1996b). Serial and parallel resource-constrained project scheduling methods revisited: Theory and computation. *European Journal of Operations Research*, *90*(2), 320–333.

Kolisch, R., & Drexl, A. (1997). Local search for non-preemptive multi-mode resource constrained project scheduling. *IIE Transactions*, *29*(11), 987–999.

Kondili, E., Pantelides, C. C., & Sargent, R. W. H. (1993). A general algorithm for short-term scheduling of batch operations—I MILP formulation. *Computers and Chemical Engineering*, *17*(2), 211–227.

Krishnan, V., & Ulrich, K. T. (2001). Product development decisions: A review of the literature. *Management Science*, *47*(1), 1–21.

Law, A. M., & Kelton, W. D. (2000). *Simulation modeling and analysis* (3rd ed.). McGraw Hill.

Levis, A. A., & Papageorgiou, L. G. (2004). A hierarchical solution approach to multi-site capacity planning under uncertainty in the pharmaceutical industry. *Computers & Chemical Engineering*, *28*, 707–725.

Linton, J. D., Walsh, S. T., & Morabito, J. (2002). Analysis, ranking and selection of R&D projects in a portfolio. *R&D Management*, *32*(2), 139–148.

Loch, C. H., & Bode-Greuel, K. (2001). Evaluating growth options as sources of value for pharmaceutical research projects. *R&D Management*, *31*(2), 231–243.

Maravelias, C. T., & Grossmann, I. E. (2001). Simultaneous planning for new product development and batch manufacturing facilities. *Industrial & Engineering Chemistry Research*, *40*, 6147–6164.

Markowitz, H. M. (1991). *Portfolio selection* (2nd ed.). Cambridge, MA: Blackwell Publishers.

Modigliani, M., & Miller, M. (1958). The cost of capital, corporation finance and the theory of investment. *American Economic Review*, *48*, 261–297.

Mohring, R. H., Schulz, A. S., Stork, F., & Uetz, M. (2003). Solving project scheduling problems by minimum cut computations. *Management Science*, *49*(3), 330–350.

Morgan, G., & Henrion, M. (1990). *Uncertainty: A guide to dealing with uncertainty in quantitative risk and policy analysis*. Cambridge, U.K: Cambridge University Press.

Naphade, K. S., Wu, S. D., & Storer, R. H. (1997). Problem space search algorithms for resource-constrained project scheduling. *Annals of Operations Research*, *70*, 307–326.

Nutt, P. C. (1998). How decision makers evaluate alternatives and the influence of complexity. *Management Science*, *44*(8), 1148–1166.

O'Connor, P. (1994). Implementing a stage-gate process: A multi-company perspective. *Journal of Product Innovation Management*, *11*(3), 183–200.

Oh, H. C., & Karimi, I. A. (2004). *Regulatory factors and capacity expansion planning in global chemical supply chains*. I&EC Research.

Patterson, J. H. (1984). A comparison of exact approaches for solving the multiple constrained resource, project scheduling problem. *Management Science*, *30*(7), 854–867.

Pekny, J. F. (2002). Algorithm architectures to support large scale process systems engineering applications involving combinatorics, uncertainty and risk management. *Computers and Chemical Engineering*, *26*, 239–267.

Pindyk, R. S., & Dixit, A. K. (1994). *Investment under uncertainty*. Princeton University Press.

Poh, K. L., Ang, B. W., & Bai, F. (2001). A comparative evaluation of R&D and project evaluation methods. *R&D Management*, *31*(1), 63–75.

Reklaitis, G. V., & McDonald, C. (2004). Design and operation of supply chains in support of business and financial strategies. Proceedings of FOCAPD 2004, Princeton, NJ.

Repenning, N. P. (2001). Understanding fire fighting in new product development. *Journal of Product Innovation Management*, *18*(5), 285–300.

Roberts, W. R. (1999). Product innovation, product-market competition and persistent profitability in the U.S. pharmaceutical industry. *Strategic Management Journal*, *20*(7), 655–670.

Rogers, M. J., Gupta, A., & Maranas, C. D. (2002). Real options based analysis of optimal pharmaceutical research and development portfolios. *Industrial & Engineering Chemistry Research*, *41*, 6607–6620.

Rogers, M. J., Maranas, C. D., & Ding, M. (2005). Valuation and design of pharmaceutical licensing deals. *AIChE Journal*, *51*, 198–209.

Salewski, F., Schirmer, A., & Drexl, A. (1997). Project scheduling under resource and mode identity constraints: model, complexity, methods, and application. *European Journal of Operations Research*, *102*(1), 88–110.

Schwartz, E. S. (1997). The stochastic nature of commodity prices: Implications for valuation and heding. *Journal of Finance*, *52*, 923–937.

Sharpe, T., & Keelin, T. (1994). How Smithkline-Beecham makes better resource allocation decisions. *Harvard Business Review*, 45–57.

Smith, C. W., Jr. (1977). Alternate methods for raising capital. *Journal of Financial Economics*, 273–307.

Sprecher, A., & Kolisch, R. (1996). PSPLIB—a project scheduling problem library. *European Journal of Operational Research*, *96*, 205–216.

Stork, F., & Mohring, R. (2000). Linear pre-selective policies for stochastic project scheduling. *Mathematical Methods of Operations Research*, *102*, 501–515.

Straub, D. A., & Grossmann, I. E. (1992). Evaluation and optimization of stochastic Flexibility in multi-product batch plants. *Computers and Chemical Engineering*, *16*(2), 69–87.

Subramanian, V. (2004). Analysis of de-centralized supply chain dynamics. PhD Thesis, School of Chemical Engineering, Purdue University, West Lafayette, Indiana.

Subramanian, D. S., Pekny, J. F., & Reklaitis, G. V. (2003). A simulation-optimization framework for stochastic optimization of R&D pipelines. *AIChE Journal*, *49*(1), 96–112.

Triantis, P. D., & Childs, A. J. (2001). Dynamic R&D investment policies. *Management Science*, *45*(10), 1359–1377.

Varma, V. A., Blau, G. E., Reklaitis, G. V., & Pekny, J. F. (2005a). A framework for learning resource allocation policies for pharmaceutical R&D portfolio management. Technical Report. Computer Integrated Process Operations Consortium (CIPAC), School of Chemical Engineering, Purdue University.

Varma, V. A., Blau, G. E., Reklaitis, G. V., & Pekny, J.F. (2005b). A branch-and-bound framework for strategic R&D capacity planning problems in the pharmaceutical industry. Technical Report. Computer Integrated Process Operations Consortium (CIPAC), School of Chemical Engineering, Purdue University.

Verma, D., & Sinha, K. K. (2002). Toward a theory of project interdependencies in high-tech R&D environments. *Journal of Operations Management*, *20*(5), 451–468.

Varma, V. A., Uzsoy, R. M., Blau, G. E., & Pekny, J. F. Lagrangian heuristics for large scale pharmaceutical operations scheduling. *Journal of Heuristics*, in press.

Wan, X., Reklaitis, G. V., & Pekny, J. F. (2005). Simulation-based optimization with surrogate models-application to supply chains. *Computers & Chemical Engineering*, *29*, 1317–1328.

Wan, X., Reklaitis, G. V., & Pekny, J. F. (2006). Simulation-based optimization for risk management in multi-stage capacity expansion. In *Proceedings of 9th international symposium on process systems engineering*

Weglarz, J. (1999). *Project scheduling: Recent models, algorithms and applications*. Dordrecht, The Netherlands: Kluwer Academic Publishers.