A MULTI-HEIRARCHICAL DECISION SUPPORT SYSTEM FOR STRATEGIC AND TACTICAL MANAGEMENT OF R&D PORTFOLIOS

G.E. Blau, J.F. Pekny, V.A. Varma
Purdue University, West Lafayette, IN-47908

Abstract

Two streams of literature exist in the area of pharmaceutical R&D portfolio management: The Strategic literature that focuses on higher level decisions such as project selection, prioritization, in-licensing, capital budgeting and the Tactical literature that focuses on operational decisions such as dynamic resource allocation between competing project activities and detailed activity scheduling. This "split" is unfortunate since strategic decisions tend to exert strong influences on tactical decisions. Our work attempts to bridge the strategic-tactical divide by modeling major decision and policy classes into a multi-hierarchical framework. This framework is composed of three layers of inter-connected optimization components. Each component is driven by suitable algorithms based on its addressed class of decisions. For example, the project selection and prioritization component is driven by the combination of a genetic algorithm and discrete event simulation while the project release time optimizing component is driven by a modified simulated annealing algorithm. The hierarchical architecture is computationally efficient in that it is shown to harness problem structure for pruning the search space of multiple decision classes and allows scalability, inherent parallelism and extensibility to accommodate portfolio planning problems in a wide range of sizes and structures. Appropriate computational case studies are described to demonstrate the utility of this framework to integrated portfolio management.

Keywords

Project Release Times, Multi-modal Project Scheduling, Simulated Annealing, Genetic Algorithm, Risk, Net Present Value, Lead Time.

Introduction

Commercial development of pharmaceutical new drug products involves managing a portfolio of multiple inter-dependent projects under conditions of manpower, equipment and capital constraints in an environment riddled with uncertainties in activity processing durations, resource requirements and product success in clinical trials and in the market. The management of such a portfolio involves decision-making at both strategic and tactical levels. Strategic decisions include selection of drug candidates to be developed from a host of available candidates, prioritization or sequencing of the selected candidates and portfolio investment planning (Pol et al, 2001, Blau et al, 2003). Tactical decisions include dynamic allocation of constrained resources between contending activities of multiple projects and detailed scheduling (Adler et al, 1995). While management scientists have traditionally focused on strategic aspects using economic and decision theories, operational researchers have concentrated their efforts on tactical or operational aspects with objectives such as reducing time to market. This “split” in efforts is unfortunate and may result in un-coordinated and counter-productive efforts towards achieving such common goals as maximizing economic value, reducing economic risk and product development times. Additionally, the prevalence of project dependencies complicates any attempt towards efficient portfolio decision-making. This work is an attempt at bridging this strategic-tactical divide. The rest of the paper is structured as follows: A Bio-pharmaceutical Integrated Strategic and Tactical Planning (BISTAP) framework is proposed and key modeling and computational complexity issues of its components are discussed along with a few computational studies. The paper concludes with an overview of the ongoing work on this framework.

The BISTAP Computational Framework

The BISTAP framework is motivated by a persistent demand from drug development businesses for software systems that can jointly optimize the large number of portfolio management decisions, given the limited computational resources. For example, current practices in most pharmaceutical firms call for approval of most drugs identified as promising by individual research teams. Such practices discourage preemption of product development and encourage under-resourcing of development for certain low priority drugs in the event of an excessively large portfolio. The inherent intuition behind such strategies is to allow as many drugs to develop in view of potential product failures. Several untested but promising alternative operational strategies exist. For instance, the release of projects could be staggered across several quarters instead of releasing a project as soon as it is approved. Projects could be preempted if the large portfolio size starts to overwhelm resources. Nevertheless, while project preemptions and delayed project releases tend to favor product lead times they also tend to increase the portfolio risk by impeding the pipeline's ability to overcome project failures (Figure 1). Besides, even under conditions of staggered or delayed project releases and project preemptions, certain activities can be under or over-resourced leading to alterations of expected durations. Additionally, set up times in re-starting projects can be as long as six months (for typical experimental teams) which calls into question the value of project preemptions.

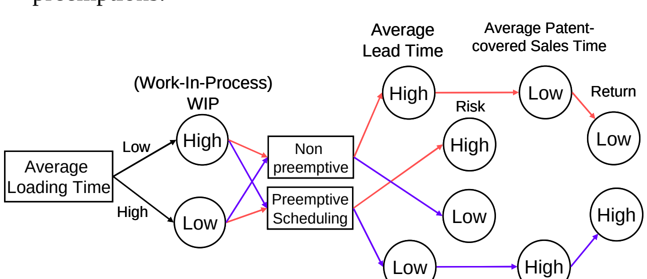

Figure 1. An Influence Diagram showing the Interaction of Project Release/Loading and Scheduling Policies

The BISTAP system provides a problem decomposition based solution towards jointly optimizing these decisions. Figure 2 shows the "core" of the BISTAP system. It is primarily composed of three blocks: a strategic planning optimizer geared towards portfolio risk management in the context of decisions such as project selection, release time and higher-level project prioritization. The second block is called as tactical optimizer, geared towards optimizing measures of product development lead times given the output from the strategic optimizer. The tactical optimizer determines efficient dynamic resource allocation and detailed activity scheduling policies for the given portfolio of prioritized drugs. Finally, there is a capacity planning block called the capacity optimizer which evaluates the queuing and other dynamic characteristics for the given portfolio and associated policies. Based on a combination of simulation and mathematical programming the capacity optimizer recommends a dynamic capacity expansion/reduction policy. It also recommends if re-iteration through the three blocks needs to be undertaken.

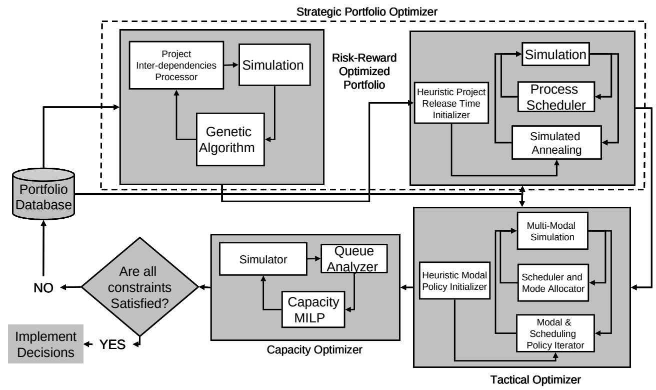

Figure 2: The BISTAP Framework

The Strategic Optimizer

The strategic optimizer generates an efficient portfolio of projects, a higher-level prioritization of projects and an efficient set of project release times (if the user selects such an option). If release time optimization is not sought then a genetic algorithm procedure is invoked (Blau et al, 2003). If release time optimization is required then a reward-risk heuristic is used to recommend a portfolio and project priorities. These are then fed into a Simulated Annealing (SA) based project release time optimizer. The SA, which is a variant of the Markov Chain Monte Carlo (MCMC) procedure, has been implemented as a two-stage algorithm. The first stage of this algorithm applies smaller sample sizes with the goal of spanning a large search space while a search consolidated around the best first stage point fine tunes the best project release time point.

While our current implementation optimizes over all pro-active project release policies (Homem-de-Mello et al, 1999), future implementations will incorporate reactive project release policies in line with the popular CONWIP, KANBAN discrete manufacturing order release policies. The simulation replications included in all algorithms of the strategic optimizer are driven by higher level planning integer programs instead of detailed scheduling programs.

Table 1: Sample Output of Project Release Time-Preemptive Scheduling Bi-level Optimization

Project Index	Initial Solution (weeks)	Stage 1(weeks)	Stage 2 (weeks)
1	5	35	20
2	10	43	39
3	15	53	57
4	20	65	79
5	25	80	101
6	30	101	111
7	35	128	133
8	40	164	224
9	45	210	263
10	50	267	301
Mean (z*) $ MM	5081.64	5718.02	5860.5
Std Dev $ MM	1900.692	2115.431	2220.968
Mean Quantile Risk = P(NPV<E(NPV))	0.4837	0.4233	0.448
Probabilistically Best Scheduling Policy	1	2	2
Probability of Optimality	0.52	0.64	0.67
Policy 0	0.29	0.2301	0.1795
Policy 1	0.52	0.1781	0.1876
Policy 2	0.23	0.6438	0.6694

This forms the basis of the decomposition-oriented paradigm of the BISTAP system in that the strategic optimizer's computational resources are focused towards spanning larger strategic decision search spaces while the tactical optimizer computational resources are focused towards solving harder policy-governing scheduling problems. The genetic algorithm that optimizes the portfolio level decisions has been optionally endowed with the capability to simultaneously optimize pro-active resource-duration trade-off decisions. Such an approach helps the tactical optimizer in converging to effective dynamic resource allocation policies guided by results of the pro-active resource-duration trade-off optimization.

The Tactical Optimizer and Portfolio Search Pruning

The tactical optimizer outputs efficient resource allocation and scheduling policies based on the input portfolio, project prioritizations and release schedules. Scheduling policies include policies governing response actions to events such as project failures, launches and achievement of certain portfolio milestones. Typical response actions include project preemptions and/or resource re-allocation. In our implementation a multi-period, multi-modal resource-constrained scheduling integer program is formulated whenever a significant event occurs in the simulated pipeline. The solution of this integer program provides the current resource allocations to all activities (and hence activity durations) alongside with the sequence in which activities need to be executed. Additionally, if the best resource allocation policy is unable to satisfy the development lead time constraints on a particular project, then a "feasibility cut" is sent back to the strategic optimizer to prevent it from selecting that project in further iterations. This drastically reduces the search space of the portfolio optimization GA.

The strategic and tactical optimizers can be interpreted as the outer and inner optimizers of a bi-level decision-making problem at constant resource capacities. If capacities are perceived to be flexible and controllable then the user can invoke the capacity optimizer.

The Capacity Optimizer

The capacity optimizer invokes procedures based on online analyses of the strategic and tactical optimizer results and outputs a capacity re-distribution policy that is employed in the next strategic-tactical iteration. Further if the probability of resource conflicts due to a certain resource mode is exceptionally high, then a "message" is sent to the tactical optimizer to prevent it from selecting that resource mode for that project. Current implementations have employed queuing related heuristics for generating a modified capacity policy.

Computational Performance on Typical Instances

The computational performance of individual components was established using test instances involving three significantly different structures of the underlying development network. The component optimizers of the framework were tested upon case studies described in Blau et al (2003), Subramanian et al (2003) and Varma et al (2003). The major results related to the strategic optimizer without resource-duration trade-offs and dynamic project release times have been described in Blau et al (2003).

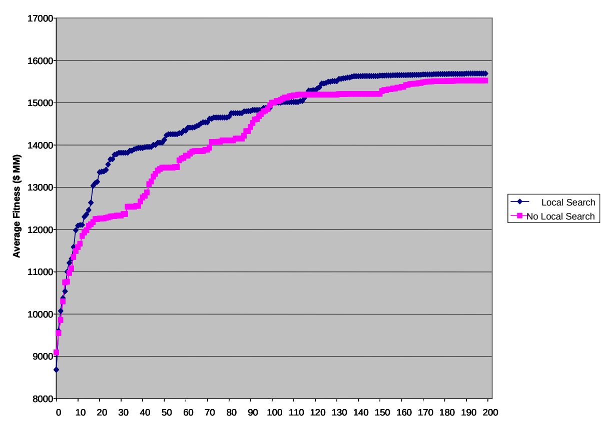

Figure 3. Sample Output from the Project Selection, Sequencing and Pro-active Resource Allocation Optimizer

Figure 3 shows the progressions of two types of genetic algorithms employed within the strategic optimizer. In this case, the framework was adapted to simultaneously optimize project prioritization and pro-active resource-duration trade-offs. The algorithms were parallelized on an RS6000/IBM-SP machine and achieved all convergence criteria in about 13 CPU hours. It can be inferred that the genetic algorithm with intermediate local search components frequently generates superior solutions. For the test instance described in Subramanian et al (2003), an ENPV improvement from $ 9012 MM to $ 15606 MM was realized. This improvement resulted from imparting flexibility of resource allocation (resource-duration trade-offs) to the portfolio and also the ability of the GA to identify the most efficient combination of project priorities and resource allocations. Table 1 shows a sample output from the two-stage Simulated Annealing (SA) algorithm run on a large molecule pipeline test instance. In this run, the framework was adapted to jointly optimize project release times (pro-actively) and scheduling policies. A project sequence and a resource-duration combination were assumed from partial runs of the portfolio prioritization genetic algorithm. The algorithm evaluates the best set of release times along with a probabilistically optimal preemptive scheduling policy.

The tactical optimizer performs detailed scheduling for each of the activities by formulating deterministic multi-modal resource constrained project scheduling (MRCPSP) integer programs. Before incorporating such integer programs into the framework, an algorithmic strategy was required to vastly speed up the MILP solution times by resorting to reliable and efficient heuristics.

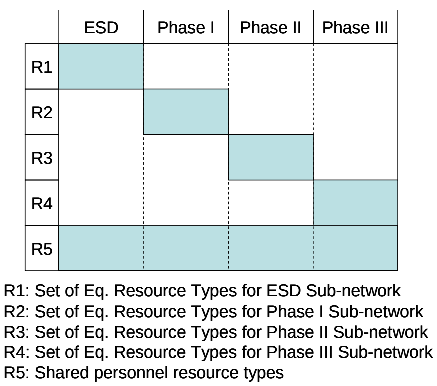

Figure 4: Decomposition Scheme for Multi-modal Scheduling

These heuristics were not only required to speed up the solution but also to supply information about the solution quality. Consequently, a decomposition structure was identified wherein the entire set of projects can be partitioned into groups based on the resource types required by each group. For instance, a group of projects may require low volume equipment as compared to other groups. The only shared resource types are the personnel types. The constraints associated with these resource types were dualized and the resulting Lagrangian problem was solved using a sub-gradient convex optimization method resulting in drastically reduced computational times and average duality gaps varying between 8-25 % for problem sizes as large as 600 activities. Figure 4 illustrates the typical resource-activity matrix responsible for the decomposition based algorithm.

Conclusion and Current Challenges

The BISTAP architecture is endowed with the capabilities to jointly optimize portfolio level strategic and project-level tactical decisions either by optimizing these decisions sequentially and invoking several iterations of this sequence or by adapting the framework to solve a multi-level optimization problem. In the former case (called the “fully decomposed mode of computation”) the main question that arises is how to allocate the computational budget across different optimizers in order to improve the solution quality in the fewest iterations i.e. in order to speed up the convergence in some sense. The latter case (e.g. the bi-level release time and preemptive scheduling problem described earlier, also called the “partially decomposed mode of computation”) involves the solution of nested or inner optimization sub-problems that address lower level or tactical decisions. Consequently, the main question that arises is how much computational budget to allocate to these nested optimizations in order to execute an efficient trade-off between solution qualities of various decisions. These issues, in addition to the algorithm engineering issues of the individual optimizers constitute the set of challenges towards improving the computational efficiency of BISTAP.

References

Adler, P.S., Mandelbaum, A., Nguyen, V., Schwerer, E. (1995). From Project to Process Management. An Empirically-BasedFramework for Analyzing Product Development Time. Management Science, 41(3), 458-484.

Blau G.E., Pekny, J.F., Rajan, K., Bunch P. (2003). A Genetic Algorithm-based Pharmaceutical Portfolio Selection and Scheduling Framework. FOCAPO 2003, Miami.

Homem-de-Mello, T., Shapiro, A., Spearman, .M.L. (1999). Finding Optimal Material Release Times using Simulation Based Optimization. Management Science, 45, 86-102.

Lambrecht, M.R., Ivens, P.L., Vandaele, N.J. (1996). Clips: A Capacity and Lead Time Integrated Procedure for Scheduling. Technical Report NR-9610, Katholieke Universiteit Leuven, Belgium.

Pol, K.L, Ang, B.W., Bai, F. (2001) A Comparative Evaluation of R&D Project Evaluation Methods. R&D Management, 31(1), 63-75.

Subramanian, D., Pekny, J.F., Reklaitis, G.V. (2003). A Simulation-Optimization Framework for Stochastic Optimization of R&D Pipeline Management. AIChE J., 49(1), 96-112.

Varma, V.A., Blau, G.E., Pekny, J.F., Alford, J.S (2003). On the Optimization of Dynamic Project Release and Scheduling Policies in a Bio-pharmaceutical New

Product Development Pipeline. AIChE Annual Meeting, 2003, San Francisco.
Weglarz, J. (1999). Project Scheduling: Recent Models, Algorithms and Applications. Kluwer Academic Publishers.