



Calhoun: The NPS Institutional Archive
DSpace Repository

Faculty and Researchers

Faculty and Researchers' Publications

1986

Optimization of Capital Portfolios

Bradley, G.

Bradley, G., "Optimization of Capital Portfolios," Proceedings of the National Communications Forum 86, 1986, pp. 11-17.

<http://hdl.handle.net/10945/38178>

defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun



Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

OPTIMIZATION OF CAPITAL PORTFOLIOS

Gordon H. Bradley
Professor of Computer Science

Naval Postgraduate School
Monterey, California 93943

ABSTRACT

Capital portfolio optimization for GTE is modeled as a mixed integer linear programming problem. The critical economic tradeoffs between maximizing the long-term market value of the firm's equity and satisfying shorter-term financial constraints, resource limitations and service objectives are incorporated into the model and into the software system that constructs optimal portfolios. The use of the model and the optimization software in the capital program management process is described.

INTRODUCTION

The goal of capital portfolio optimization is to maximize the long-term market value of the firm's equity while satisfying near-term financial constraints, resource limitations and service objectives. Facing an increasingly complex and competitive marketplace, GTE has developed a computer-based system to improve the quality of its capital decisions by using state-of-the-art decision technology to automate the quantifiable aspects of selecting optimal portfolios. This is part of the GTE Capital Program Management System (CPMS). The system allows the effective use of significantly more information at the highest levels of decision making. The additional information comes from requiring more information about capital investment opportunities and from requiring that several alternatives for implementing each be developed. Better decisions are achieved from the increased information by a process that includes using an optimization model to provide a conceptual framework for the problem and an extensive and innovative software system that coordinates the data gathering and constructs optimal portfolios. The automation of the quantifiable aspects of the problem allows management to focus on the critical non-quantifiable factors.

Improving the quality of capital decisions by making better use of information mandates that:

1. Design of portfolios must include an automatic means to guarantee that any selected portfolio meets corporate financial constraints. Manual evaluation of GTE's multi-billion dollar investment program is not feasible.
2. The system must automatically decide when it is justified to modify short-term restrictions in order to improve the long-term value of equity. This is a critical tradeoff between the short-term

restrictions and long-term objectives. Because of the scale and complexity of GTE portfolio decision problems, the evaluation of these tradeoffs must be incorporated into the construction of optimal portfolios.

3. Using all the available quantifiable data, the system must be capable of automatically constructing the best possible (that is, the optimal) capital portfolio. GTE's capital program is so large that the dollar difference between the optimal portfolio and inferior portfolios is huge. The analysis of important non-quantifiable factors depends on the ability to produce optimal portfolios for "What if...?" questions. To decide whether strategy A is better than strategy B, it is necessary to compare the values of the optimal portfolio for each.

Item 1 is handled by incorporating a funds statement for the entire company into the model; this allows financial constraints for the portfolio to be imposed explicitly. Extending the model to allow "penalties" for constraint violations expresses the tradeoffs from 2 by assigning a dollar value to violating the near-term requirements. The mixed integer programming optimization system provides the capability required by 3 and it has demonstrated its ability to construct optimal portfolios in a cost effective and timely manner for GTE capital portfolio problems.

The project database, the model and the software to construct optimal portfolios are built to allow high-level management access to significantly more information about capital projects and their interactions than was previously possible. Studies to explore options, answer questions and investigate tradeoffs can be performed in a timely manner. The system can be used throughout the year as new opportunities and new problems require reconsideration of capital decisions. In a dynamic and demanding business environment the project database represents a valuable corporate asset; the system allows this information to be used effectively in the capital planning process.

A MODEL FOR PORTFOLIO DESIGN

Within GTE, each telephone operating company formulates its own capital portfolio problem. In capital portfolio optimization the candidate projects do not all fit the usual definition of "capital"; they include projects such as multi-year marketing plans. In a GTE telephone operating company the number of candidate projects ranges from many hundreds to several thousands. The candidate projects may begin in any of the first five years.

In the mixed integer linear programming model each candidate project is represented as a variable $x(p)$ where p ranges from 1 to the

maximum number of projects. The value of each variable is the level of participation of the candidate project in the portfolio. Each variable must take on a value between 0 and 1:

$$0 \leq x(p) \leq 1 \text{ for all candidate projects } p.$$

$x(p) = 0$ means the candidate project is not in the portfolio, $x(p) > 0$ means p is in the portfolio and $x(p) = 1$ means p is fully funded.

Most candidate projects are indivisible in that they must either be fully funded or not adopted at all:

$$x(p) = 0 \text{ or } 1 \text{ for all } p \text{ indivisible.}$$

The remaining projects are divisible, they may take on any value between 0 and 1; $x(p)$ is the fraction of the project that will be funded. The term "mixed integer" denotes the presence in the model of both indivisible (=integer) and divisible variables.

Most capital investment opportunities are represented in the model by several candidate projects that are implementation alternatives. The alternatives may represent different timings (for example, accelerated, normal, stretched out), different funding levels (scaled back, normal, expanded), different technologies, different start years, etc. The indivisible alternatives are mutually exclusive, that is, at most one may be in any portfolio. For alternatives p_1, p_2, \dots, p_k this is enforced with a "logical" constraint:

$$x(p_1) + x(p_2) + \dots + x(p_k) \leq 1.$$

There is one such constraint for each group of alternatives. For indivisible candidate projects this logical constraint means that either none or exactly one of the group can be in a portfolio. For divisible projects it is possible that several candidate projects in the group can have a positive fractional value; this represents a "blending" of the candidate projects.

Mutual exclusion among alternatives is one example of a "logical" restriction on the candidate projects that are allowed in a portfolio. Other possible logical restrictions include:

1. Choose at most one candidate project from a set of candidate projects (a generalization of mutual exclusion for alternatives).
2. Choose exactly one candidate project from a set of candidate projects.
3. A specified candidate project cannot be included in the portfolio unless another specified candidate project is also included.
4. Two candidate projects are either both in the portfolio or are both out of the portfolio.

The above description of logical constraints applies to indivisible projects; there is a slightly different interpretation for divisible projects.

The net present value (NPV) of a project is its discounted long-term cash flow. NPV is well accepted as a figure of merit for the long range market value of the firm's equity. Let $NPV(p)$ denote the NPV of project p . The objective is to select the portfolio with the maximum NPV:

$$\text{maximize } \sum_p NPV(p) * x(p)$$

where \sum_p means "sum over all projects p ".

In addition to the logical constraints discussed above, there are short-term (5-year) financial constraints, resource limitations and service objectives that limit the allowable portfolios. These constraints apply to each individual year of the five year short-term planning horizon. Let t be an index for the years 1, 2, ..., 5.

One of the financial constraints is an upper limit $UCAP(t)$ on the total company capital expenditure for each year t . The input data includes the company non-project capital expenditure $OCAP(t)$. Part of the specification of each project is the capital expenditure (actually net plant requirement) of each project in each year $CAP(p,t)$. The financial constraint for each t is:

$$OCAP(t) + \sum_p CAP(p,t) * x(p) \leq UCAP(t).$$

Another financial constraint is a lower limit $LIGF(t)$ on the net funds from internal sources as a percent of the capital requirement. Let X be the vector of all the variables $x(p)$. Let $NFI(X,t)$ denote the net funds from internal sources for the portfolio X in year t and $REQ(X,t)$ denote the capital requirement. The constraint for each t is:

$$NFI(X,t) / REQ(X,t) \geq LIGF(t).$$

$NFI(X,t)$ is calculated by computing the funds statement with the portfolio X . There is a complete funds statement for the company embedded in the portfolio optimization model. The funds statement is based on five years of financial data including information on sources to raise cash (short-term debt, long-term debt, preferred and common stock), various interest rates, corporate tax rates, etc.

There is financial data specified for each project, thus the impact of each project on the funds statement can be calculated. Although it is not practical to develop an explicit constraint for $NFI(X,t)$ or $REQ(X,t)$ in terms of the financial data items, it is possible via implicit calculation of the funds statement to obtain the numbers necessary to explicitly generate the above constraint as a linear function of the project variables $x(p)$.

There is also a lower limit $LNIC(t)$ on the net income to common. $NIC(X,t)$ is the net income to

common of portfolio X in year t. The financial constraint for each t is:

$$NIC(X,t) \geq LNIC(t).$$

Again this constraint is explicitly constructed as a linear function of the project variables by means of an implicit computation of the funds statement and the relevant portion of the income statement.

The above three constraints are the financial restrictions on the portfolio. The embedded funds statement gives the model the potential to constrain the portfolios on any funds statement quantity of interest.

There are critical resources identified by each GTE telephone operating company that in the short term limit its freedom to select capital projects. These presently include upper bounds on labor hours in critical job skills, upper bounds on lines installed to mandate work force leveling over the five years, etc. Each telephone operating company may choose the resources that are most critical for it to control.

For each resource RES(t), a lower limit LRES(t) and/or an upper limit URES(t) may be specified. ORES(t) is the consumption of the resource by company non-project activities. For each project p and year t the consumption of resource r is specified as RES(r,p,t). For each resource r the constraint for each t is:

$$LRES(r,t) \leq ORES(r,t) + \sum_p RES(r,p,t) \\ *x(p) \leq URES(r,t).$$

Service criteria can be identified and objectives established for each criteria. For example, objectives could be set for toll call completion probability, trouble reports per 100 lines, etc. The portfolio can then be constrained to meet all the services objectives for the five years. Analogous to resource constraints, the constraints SER(s) for each t are:

$$LSER(s,t) \leq OSER(s,t) + \sum_p SER(s,p,t) \\ *x(p) \leq USER(s,t).$$

It is difficult to identify service criteria that are universally accepted and it is sometimes difficult to isolate the service contribution of each individual project. At the present time the service part of the system is not being used by the GTE operating companies.

The model features of portfolio optimization include the capability to override the optimization and lock any project into the portfolio ($x(p) = 1$) or lock any project out of the portfolio ($x(p) = 0$). Projects that are divisible can be locked into the portfolio at any value between 0 and 1 or limited to a subinterval between 0 and 1.

Output of the optimization specifies a portfolio by giving the level, $x(p)$, and NPV contribution of each project, $NPV(p)*x(p)$. The output includes for the optimal portfolio the consumption of capital and resources as well as the levels of the service criteria. Also included is the funds statement for the complete company for all five years based on the optimal portfolio.

The description of the model thus far has not included the mechanisms for including the trade-offs between NPV and the near-term financial, resource and service requirements. After motivating the "penalties" approach, the remainder of the model will be presented.

"JUSTIFIABLE" VIOLATIONS OF THE CONSTRAINTS

The constraints of the financial, resource and service requirements over the first five years are the near-term restrictions on the long-term goal of maximizing NPV. The model can be extended to allow these constraints to be violated in a systematic way to achieve portfolios with greater NPV. By contrast, the logical constraints cannot be relaxed; every portfolio must satisfy them. There are, however, several reasons for allowing "slight" violations of the financial, resource and service constraints:

1. The numbers that are specified for the limits are really objectives rather than fixed numbers. This is especially true for the out years (3, 4 or 5) where the values for the limits are not as critical. The constraints can be viewed as "soft" in that the limits can be varied slightly without changing the intent of the restriction.
2. The limitations are objectives that are set relative to the marginal opportunities for investments. That is, the limitations can be stretched to include projects with significant NPV or shrunk if the optimal portfolio contains some low value projects. This is the fundamental question of balancing long-term gain against short-term restrictions. This kind of analysis has become more critical as deregulation has opened up new opportunities and brought new risks.
3. The nature of project design with indivisible projects is such that slight modifications of a few constraints often makes possible a different portfolio with significantly higher NPV.

The last point can be demonstrated with a greatly simplified example with only three indivisible projects and only a one year capital constraint:

| Project | NPV | Capital Requirement |
|---------|--------|---------------------|
| | 1000\$ | 1000\$ |
| P1 | 20 | 30 |
| P2 | 30 | 50 |
| P3 | 50 | 55 |

(the capital limit is 84 thousand dollars.)

This can be written as an integer linear programming problem:

$$\begin{aligned} &\text{maximize} && 20*P1 + 33*P2 + 50*P3 \\ &\text{subject to} && 30*P1 + 50*P2 + 55*P3 \\ &&& \leq 84 \end{aligned}$$

P1, P2 and P3 must be 0 or 1

Of the eight possible portfolios (including the one with no projects), five satisfy the capital constraint; the optimal portfolio contains P1 and P2, has NPV 53 and a capital requirement of 80.

If the capital limit could be violated by one thousand dollars, then a significantly better optimal portfolio is P1 and P3 with NPV 70 and capital requirement 85. It is typical of optimization with indivisible projects that slight modifications of the limits yield portfolios with better NPV.

THE MODEL FOR PORTFOLIO DESIGN WITH PENALTIES

Recognizing that slight violations of some constraints can lead to better capital decisions still leaves a dilemma : 1) With tens or hundreds of constraints, which should be violated ? and 2) what is a "slight" violation ?

The economic principle of marginal prices yields a solution to both difficulties: the financial, resource and service constraints can be violated for a price. This is incorporated into the model by assessing a penalty for violations of the constraints. This penalty is the unit price (in NPV dollars) of allowing the constraint to be violated. For example:

1. The penalty for the capital constraint is the cost to obtain one extra dollar of capital (the "marginal" cost of capital).
2. For a constraint on the number of labor hours consumed in a critical skill, the penalty is the cost to obtain one more labor hour; this could be either the overtime rate or the hourly rate of an outside contractor.
3. The penalty for violating the limit on the number of lines to install could be the cost to hire and train additional staff and purchase additional equipment.

The units of the penalties are dollars of NPV per unit constraint violation. This expresses the economic tradeoff between long-term goals and short-term restrictions.

The example of the previous section can be resolved by leaving the capital limit at 84 and assigning a penalty of one dollar NPV per extra dollar of capital. The objective function becomes:

$$\text{maximize } 20 \cdot P1 + 33 \cdot P2 + 50 \cdot P3 - 1 \cdot (\text{violation})$$

The three portfolios that were previously ignored because they violated the capital constraint are now evaluated. For example, the portfolio with P1 and P3 has objective value $20 + 50 - 1 \cdot (1)$ to give 69. The other two portfolios evaluate to 52 and 62. The five portfolios that do not violate the capital constraint have the same

value that they had before. For this example with penalties, the optimal portfolio contains P1 and P3 with objective value of 69 and capital requirement 85.

The portfolio optimization model that has been implemented at GTE includes penalties for the financial, resource and service constraints and the software constructs the portfolio that maximizes the sum of the NPV minus the sum of the penalties times the violations. In effect the optimization evaluates every portfolio that satisfies the logical constraints and applies a penalty for all portfolios that violate any of the other constraints.

The incorporation of penalties into the model is conceptually similar to what is called "goal programming" in the management science literature. Although conceptually similar, the scale of the GTE capital planning problems requires that the optimal portfolios be constructed in a different and much more efficient manner.

The penalties are chosen by high level executives to quantify their understanding of the economic tradeoffs. In a series of executive level seminars on Portfolio Management being conducted at GTE, the participants have been excited about having a quantified measure of tradeoffs that have always been an implicit (and critical) part of the capital planning decision process.

The results of calculating optimal portfolios using penalties are exactly what would be expected in a stable and well-run company: the optimal portfolios involve either no violations or just minor violations of a few constraints (often in years 4 or 5). This shows that the executives have made good choices for the limits on the constraints. "At the margin" at most a few projects have sufficient NPV to justify violating constraints to get them into the portfolio. Although only few in number the extra projects yield a better portfolio at the price of only a few "slight" violations of some constraints.

USE OF THE OPTIMIZATION IN PORTFOLIO DESIGN

The process of designing a capital portfolio begins when the proponents of the various projects build the database of financial, resource and service information. As the projects are developed the logical restrictions for the projects are developed. The NPV is calculated for each project. The data on the financial impact, resource consumption and service contribution that is not related to the portfolio is gathered. The data necessary to generate the funds statement is collected and validated. This information constitutes the project and financial databases that then remains stable throughout the portfolio design process.

At the beginning of portfolio construction, the executives select limits for the financial, resource and service constraints. They also select the value of the penalties for these constraints.

An optimal portfolio can then be constructed. The work of considering the non-quantifiable factors then begins.

The portfolio construction process includes side studies by the planning staff to explore options, answer questions, investigate tradeoffs and in general to gain an understanding of the projects and their impact on long and short-term objectives. The guidance from the executives might be, for example, to tighten or loosen certain limitations, to lock certain projects in or out of the portfolio, to change one or more of the financial assumptions, etc. Each side study will involve identifying a set of slightly different problems, constructing an optimal portfolio for each and then studying the results. The staff will then present summary results and an explanation for any changes in the optimal portfolio. The main goal of the studies is to gain insight into the company's capital opportunities, the explanations that the several runs produce will be often as important as the numbers.

Throughout the year new opportunities and new problems will prompt a reconsideration of the capital plans. As the year proceeds, the projects that have begun or have been committed to are locked into the portfolio. The new projects then compete against all other projects for a place in the new capital plan. In this process any new project does not have to compete as a budget augmentation, rather each new project can compete on an equal footing with all projects that have not already been locked into the portfolio.

THE OPTIMIZER AND OPTIMAL PORTFOLIOS

The optimizer provides GTE the capacity to construct optimal portfolios for capital planning problems of unprecedented size and detail. A remarkable combination of good capital projects, conscientious data development, sound economic modeling and high technology optimization software works to use all the available quantifiable data to construct the best possible capital portfolio.

The previous paragraph contains a bold claim that the system solves to optimality a problem that is known to be extremely difficult to solve. This section outlines an argument that specifies precisely in what sense the optimizer constructs the optimal (that is, the best possible) portfolio. This section explains in what sense the problem is difficult, discusses the complicating issues of numerical precision and data resolution, gives some details about the state-of-the-art system that performs the optimization, develops some reasons why the optimizer is so successful on this very hard problem, and finally explains in what sense the portfolio that is constructed should be regarded as optimal.

Integer programming is a member of a class of problems that are extremely difficult to solve. Computer scientist theorists believe (but have not yet proven) that there never can be an algorithm to effectively compute optimal solutions for large

problems of this type. For a problem with only 100 indivisible projects, there are 2 to the power 100 (more than 1 followed by 31 zeros) different portfolios. The only known exact algorithms can in the worse case take roughly this many steps. An example of a worse case would be a problem where each different portfolio would have a unique total NPV; thus any algorithm would somehow have to consider each one before declaring it had found the optimal. This somewhat overstates the difficulty because it assumes a rare, pathologically complex problem, but average behavior on randomly generated problems is also very bad.

Almost all commercially available integer programming systems have been designed to solve any possible problem. The usual result of this strategy is a system that requires considerable computer time and displays high variability in solution quality and computer costs from problem to problem.

The size of the GTE capital portfolio problems is significantly beyond the state-of-the-art. The system is presently configured for 2000 projects and 1250 constraints (roughly 250 for financial, resource and service constraints and 1000 for logical constraints). For each project there are usually 25 to 100 different data values (NPV, financial quantities, resource consumption and service contribution). For 2000 projects this yields 50,000 to 200,000 pieces of data. The optimizer uses all this information. Currently, the GTE telephone operating companies are phasing in the use of the optimizer; no company is far enough along to have approached these size limits. The larger companies will eventually go beyond these size limits; we expect to extend the system (and create the necessary technology) to keep pace.

The optimizer is the INSIGHT X-system. The X-system is an experimental testbed that includes all the standard pieces of contemporary optimization methods plus many features that are at the cutting edge of advanced research in the field. In addition to its advanced capabilities, it has features that make it a stable, reliable production system. It is presently installed at seven GTE telephone operating companies.

Despite the gloomy prospects for solving integer programs, the X-system constructs an optimal portfolio for the largest problems in at most a few hundred CPU seconds. How does the optimizer achieve this seemingly impossible feat? The first answer is to customize the software to solve only capital portfolio problems that are structured as the GTE problems. There has also been considerable work done in adjusting the numerous control parameters to tune the software to the characteristics of GTE's problems. There has been extensive experimental work over several years. This, however, is only part of the answer.

Jerry Brown and Glenn Graves, the developers of the X-system, have had considerable experience over the last 20 years solving real problems in many areas for corporate customers and government agencies. Their series of

observations on this topic can be summarized in this hypothesis: careful modeling of real systems that have been shaped by genuine economic forces yield problems that are dramatically easier to solve than random or poorly modeled problems. The GTE capital program exhibits the following relevant characteristics:

1. The existing capital plant base has been shaped over many years by economic forces.
2. The capital projects are not random; they are the result of hundreds of employees carefully developing projects that will benefit the company and fit in with existing and other new projects.
3. CPMS has developed procedures to develop the data in a consistent way and to protect the integrity of the project database.
4. The model of portfolio optimization includes a sophisticated financial model and contains marginal prices that quantify critical economic tradeoffs.
5. The system was developed and tested and tuned with real GTE project data.

All this supports the hypothesis and so predicts success. It should be noted that the above elements are not obvious. There have been four years of development that have included several major changes to the model and the data.

The numerical precision of calculations on a digital computer is another important factor. Calculations are not carried out in the ideal world of exact arithmetic; rather, digital computers implement a finite precision model of arithmetic. The X-system works in double precision arithmetic with 16 decimal digit representations. The X-system uses sophisticated numerical methods to maintain precision and accuracy. There are intricate methods to perform basis factorizations and to maintain a structure of the numerical representation that minimizes numerical difficulties. One of the practical consequences of finite precision arithmetic is that after the extensive calculations to construct portfolios, many portfolios have the same NPV value. Although a problem with 1000 indivisible portfolios could have 2 to the power 1000 portfolios with unique NPV values, on a contemporary digital computer not all can be represented uniquely. This necessitates revising the definition of optimality. Using exact arithmetic there may be a unique portfolio with the maximum NPV value, but on the computer, many portfolios could have the same (maximum) value and thus any one could be regarded as the optimal solution. Thus we can only talk about a portfolio being optimal "within the numerical precision of the computer".

The optimization in the X-system constructs a sequence of better and better portfolios using a branch and bound strategy. Simultaneously, it produces a sequence of decreasing upper bounds on the value of the optimal portfolio. If the upper

bound at some point in the branch and bound tree becomes equal to the value of the current best portfolio, the system stops and the current solution is guaranteed to be optimal ("within the numerical precision of the computer"). This preemptive termination occurs with remarkable frequency. For calculations that don't terminate this way, the system terminates when the percentage difference between the value of the current best solution and the upper bound becomes so small that it is less than the resolution of the project data. In spite of the careful development of costs for all the candidate projects, the data is at best accurate for only the first few most significant digits; the remaining digits are meaningless. The resolution determines a termination criterion for the system. For example, if the difference between the current best solution and the upper bound is \$1000 and no project NPV is accurate to the nearest \$1000, it is time to terminate the system and declare the current best portfolio to be optimal "within the resolution of the data".

The discussion of this section can be summarized:

1. GTE is an operating firm with an existing capital plant that reflects years of careful planning;
2. GTE has many conscientious employees that are developing capital projects that will improve the system;
3. CPMS has incorporated business practices to assure that project data is produced in a consistent way and is correctly gathered and stored;
4. A sophisticated model of capital portfolio design based on sound economic principles has been developed and refined;
5. A state-of-the-art optimization system of advanced design and concept has been customized and then finely tuned on actual GTE portfolio problems;
6. Several years of experience has produced procedures to use the system effectively; and finally
7. A managerial and technical overview is continuously monitoring the system and is prepared to make adjustments when any of the preceding items cease to be true.

We conclude that the CPMS optimizer constructs portfolios for GTE capital portfolio problems that are optimal within the numerical precision of digital computers and the resolution of the project data.

CONCLUSION

The quantifiable aspects of constructing optimal capital portfolios have been effectively automated. The timely and cost effective construction of optimal portfolios by the

optimization software allows "What if..?" questions to be answered and side studies to be executed. This allows management to focus on the critical non-quantifiable factors in capital planning. A particularly important aspect of the process is the ability to relate the long-term financial health of the company to near-term requirements and objectives through the use of economic penalties.

The CPMS optimizer provides GTE with unprecedented capability to make the best use of the available quantitative information to construct optimal portfolios.