Interfaces
Vol. 34, No. 6, November–December 2004, pp. 415–425
ISSN 0092-2102 | EISSN 1526-551X | 04 | 3406 | 0415
DOI 10.1287/inte.1040.0107

Optimizing Military Capital Planning

Gerald G. Brown, Robert F. Dell
Operations Research Department, Naval Postgraduate School, Monterey, California 93943
{gbrown@nps.edu, dell@nps.edu}

Alexandra M. Newman
Division of Economics and Business, Colorado School of Mines, Golden, Colorado 80401, newman@mines.edu

Planning United States military procurement commits a significant portion of our nation's wealth and determines our ability to defend ourselves, our allies, and our principles over the long term. Our military pioneered and has long used mathematical optimization to unravel the distinguishing complexities of military capital planning. The succession of mathematical optimization models we present exhibits increasingly detailed features; such embellishments are always needed for real-world, long-term procurement decision models. Two case studies illustrate practical modeling tricks that are useful in helping decision makers decide how to spend about a trillion dollars.

Key words: finance: capital budgets; military: defense systems.
History: This paper was refereed.

Procurement of materiel has been a concern of the American military since the Revolution. Most early procurement requests amounted to not much more than a field officer's handwritten letter listing "what we must have to accomplish this task." With the exception of goods uniquely critical to the military, such as saltpeter, most supplies were simple, of the sort civilians would buy, and requests were for modest quantities. Debate centered on how to pay for what the military needed, rather than whether its need was real. (To gain an appreciation of the predominant role military procurement played in Revolutionary times, read the writings, annotated diaries, or biographies of those who debated and decided these issues, for example, McCullough's *John Adams* 2001.)

Planning US military procurement remained fairly short term and simple and motivated by apparent need but driven by immediate affordability until after World War II. In 1948, the Hoover Commission required that the military set forth its defense goals and the means by which it would achieve them. In the early 1960s, Secretary of Defense Robert McNamara tried to mitigate the myopia of the single-year budget plan and accurately represent defense systems too complex to be procured and fielded in a single year. He introduced a five-year budget requiring analytical justification. This foundation underlies our military planning today: given a defense requirement that probably is not encumbered by budget concerns, each branch of the military forms a strategy, categorizes it into "mission areas," and translates them into requirements for personnel and materiel. Chambers (1999) provides a detailed history of military funding and its consequences.

Despite early simplicity in justifying military expenditures, US military capital planning has always involved large amounts of resources from many parts of the country, extensive research effort and technology development, huge amounts of money, and the attention of political leaders. In 1794, the US Congress approved the construction of the USS Constitution (Figure 1) and her five sister frigates, costing $800,000 1794 dollars, or about $2.9 billion 2003 dollars (Field 1999), using the newest technology and resources from all the colonies, on the condition that the ships be built exactly as proposed in six different American constituencies.

Modern procurement planning may concern programs that require many years to develop and complete. A program may compete or interact synergistically with others. The program's requirements and costs may change during development. For example, in 1999, the US Navy suffered a $100 million cost overrun in fielding the joint strike fighter (Figure 2) (Ricks 1999).

Because capital planning is important, complex, and expensive, it invites careful analysis. Since the introduction of mathematical programming after World War II, the military and the private sector have used it to solve capital-planning problems, and the resulting decisions have committed trillions of dollars.

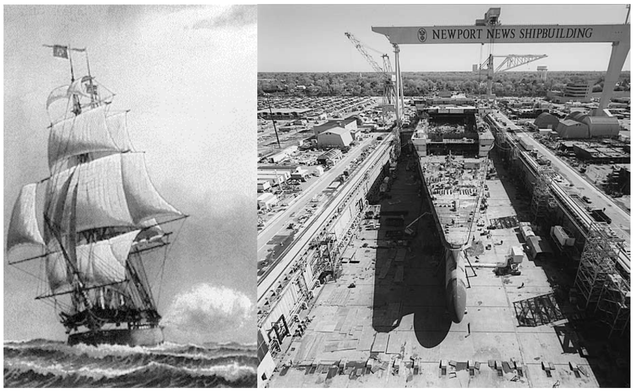

Figure 1: The USS Constitution incorporated innovative naval architecture and the latest armament technology; the highest levels of American government planned and approved its construction in 1794, and it required a huge mobilization of colonial resources. Newport News Shipbuilding, the sole shipyard in the United States capable of building nuclear-powered aircraft carriers, is building the USS Ronald Reagan (CVN76). The cost for the Reagan and her aircraft is expected to be about 10 billion 2004 dollars.

Source, left-hand figure: US Navy, 1997, "USS Constitution, the History," http://www.chinfo.navy.mil/navpalib/allhands/ah0697/jun-pg30.html, June.

Source, right-hand figure: US Navy, 2004, http://www.reagan.navy.mil, January.

Military Capital Planning and Civilian Capital Budgeting

In his recollections, Dantzig (1963, Chapter 2, p. 12) clarifies why the military rushed to develop linear programming just after World War II: "A nation's military establishment, in wartime or in peace, is a complex of economic and military activities requiring almost unbelievably careful coordination in the implementation of plans produced in its many departments."

Figure 2: The Lockheed Martin X35 joint strike fighter has a planned unit cost of about $40 million, with production ramping up to produce 500 planes over fiscal years 2005–2010, deliveries starting in 2008, initial operational capability in 2011, and a total planned production campaign of 3,000 aircraft. The X35 will replace the US Air Force's A-10 and F16, the US Marine's AV-8B and F/A-18, and the US Navy's F/A-18. The three services' long-term capital plans must reflect the influence of this transition across all these aircraft and their weapon systems.

Source: Lockheed Martin Aeronautics Company, 2004, http://www.lmaeronautics.com, January.

In some of the earliest papers in the journals of the new discipline of operations research, authors address military capital planning, for example, Bailey (1953) and Stanley et al. (1954). More recently, Taylor et al. (1983) analyze military aircraft procurement, Brown et al. (1991) develop a large-scale linear integer model for modernizing the US Army's helicopter fleet over a multidecade planning horizon, Brown et al. (1994) describe a nonlinear optimization model the US Air Force used for more than 20 years to recommend purchases of conventional gravity bombs, and Brown et al. (2003) describe a linear integer model the US Air Force has used to select and schedule investments in space-based assets over a 25-year horizon. Salmeron et al. (2002) describe a decision-support system the US Navy can use to decide on capital outlays of about a trillion dollars for navy force structure over a 25-year horizon.

Reports of optimizing civilian capital budgets appear as early as the 1950s (Gunther 1955, Lorie and Savage 1955). In their surveys, Weingartner (1966), Bernhard (1969), and Weingartner (1977) discuss linear and nonlinear models and the use of discount rates. Papers in the civilian literature rarely concern actual applications, with some exceptions. Rychel (1977) presents a multiple-time-period linear integer program to maximize net worth for Cities Service Company. Bradley (1986) describes a model for maximizing short- and long-term net present value for General Telephone and Electronics Corporation subject to financial, resource, and service constraints.

In addition, a body of academic literature focuses more on solution techniques than on solutions of specific capital-budgeting problems: Everett (1963) and Mamer and Shogan (1987) demonstrate the use of Lagrangian relaxation, Kimms (2001) describes the use of Benders decomposition, and Meier et al. (2001) discuss ways to estimate portfolio value uncertainty from samples of real options, with both a heuristic solution and a heuristic bound on solution quality.

Brown et al. (2004) highlight differences between military capital planning and civilian capital budgeting, offering about 75 military and civilian citations. For instance, a military purchase creates no tax event; the government has no concept of depreciation, book value, amortization, or residual return on owners' equity; and there are no external encumbrances on the financial leverage or tax exposure of a military acquisition. Of course, the government operates under a dizzying myriad of self-inflicted political and regulatory restrictions that the private sector does not suffer. While the differences are numerous, the underlying issues are similar, and the private sector can learn a great deal from the military's decades of experience with optimization-based capital planning.

Elements of Optimization Models for Capital Planning

Models for optimizing military capital planning prescribe which weapon systems to procure, when to procure them, and how many to procure.

Portfolio Selection, aka the Knapsack Model

One of the simplest optimization models for military capital planning is a binary knapsack. Given a fixed budget and a set of binary acquisition options, where each option has a value and a cost associated with the procurement of one or more weapon systems, we seek the set of options that has the maximum total value at a portfolio cost no greater than our budget. This model is a linear integer program with a linear objective and a single linear inequality constraint with nonnegative coefficients.

For this simple model, we make standard linear-programming assumptions: additive objective values and additive costs, constant returns to scale, separable options, and deterministic data. In terms of real-world acquisitions, this means that the total portfolio value is just the sum of its component values, and no synergism exists among selected options. The model includes no returns to scale, no volume discounts, and no mutually exclusive or inclusive options, and we have absolutely perfect knowledge a priori of the exact consequences of any action we might choose.

Acquisition Options

In the real world of capital planning, important embellishments go beyond the textbook binary knapsack problem. In particular, we frequently must decide whether or not to buy any units of a weapon system and then decide how many to buy. We may also have several options available for procuring a weapon system. For this generalization, we can use a bounded integer knapsack model (Bertsimas and Tsitsiklis 1997, Chapter 6), with one measure of effectiveness (MOE). Keeney (1992) provides guidance on developing MOEs and on scoring a system's contribution towards an MOE, and Parnell et al. (1998) describe an application. Loerch (1999) presents a linear integer program for the instance in which contribution (or cost) decreases nonlinearly as the integer quantity of the system procured increases. (Such phenomena arise with quantity discounts, learning curves, and diminishing returns.) We can capture this nonlinear contribution (or cost) with a piecewise linear function by using binary selection variables, each of which assumes a value of one if the model chooses the acquisition option in the associated quantity range or a value of zero otherwise.

An embellished knapsack model follows:

1. Indices and Index Sets

a = acquisition option.

w = weapon system.

$w(a)$ = set of weapons system(s) procured under acquisition option a.

2. Parameters [units]

l_{aw} (u_{aw}) = lower (upper) limit on quantity of weapon system $w \in w(a)$ available for purchase under acquisition option a [w-units].

$\mathit{fixedcontr}_a$ = fixed contribution of acquisition option a towards the MOE [value units].

$\mathit{varcontr}_{aw}$ = variable contribution per unit of weapon system $w \in w(a)$ purchased under acquisition option a towards the MOE [value units/w-unit].

$\mathit{fixedcost}_a$ = fixed cost incurred by selecting acquisition option a [$].

$\mathit{varcost}_{aw}$ = variable cost per unit of weapon system $w \in w(a)$ purchased under acquisition option a [$/$w$-unit].

budget = available budget [$].

3. Decision Variables

$\mathit{SELECT}_a = 1$ if any units are purchased under acquisition option a, $= 0$ otherwise [binary].

$\mathit{QUANTITY}_{aw}$ = number of units of weapon system $w \in w(a)$ purchased under acquisition option a [w-units].

4. The Corresponding Linear Integer Program

$$\max \sum_a \Big(\mathit{fixedcontr}_a \mathit{SELECT}_a + \sum_{w \in w(a)} \mathit{varcontr}_{aw} \mathit{QUANTITY}_{aw} \Big)$$

$$\text{s.t.} \sum_a \Big(\mathit{fixedcost}_a \mathit{SELECT}_a + \sum_{w \in w(a)} \mathit{varcost}_{aw} \mathit{QUANTITY}_{aw} \Big) \le \mathit{budget},$$

$$l_{aw} \mathit{SELECT}_a \le \mathit{QUANTITY}_{aw} \le u_{aw} \mathit{SELECT}_a \quad \forall a, w \in w(a),$$

$$\mathit{SELECT}_a \in \{0, 1\} \quad \forall a,$$

$$\mathit{QUANTITY}_{aw} \in \{0, 1, 2, \ldots, u_{aw}\} \quad \forall a, w \in w(a).$$

Restating the effect of the above model, either $\mathit{SELECT}_a = 0$ and $\mathit{QUANTITY}_{aw} = 0$ for all $w \in w(a)$, or $\mathit{SELECT}_a = 1$ and $l_{aw} \le \mathit{QUANTITY}_{aw} \le u_{aw}$ for all $w \in w(a)$. If purchase quantities are sufficiently high, we can reasonably relax the integrality requirement on $\mathit{QUANTITY}_{aw}$. For example, Salmeron et al. (2002) use integer annual quantities of navy ships (for about three ships per year) but continuous annual navy aircraft quantities (for about 100 aircraft per year). For

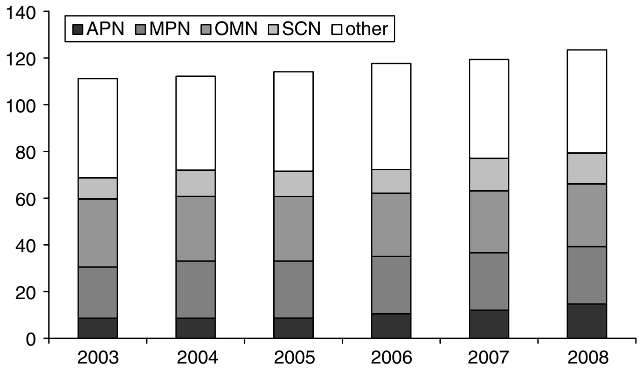

Figure 3: The US Department of the Navy total obligation authority (TOA, a grand total spending limit) was extracted from the Defense Department's future years defense program (FYDP: the multiyear spending plan) as of the January 2003 submission of the president's budget for fiscal year 2004. Shown are fiscal year accounts in constant 2003 billions of dollars and five earmarked colors of money (APN is Aircraft Procurement, Navy, MPN is Military Personnel, Navy, OMN is Operation and Maintenance, Navy, SCN is Shipbuilding and Conversion, Navy, and other represents an aggregation of 20 additional categories). In capital-planning parlance, fiscal year 2003 is the current year, 2004 and 2005 are budget years, 2006 and beyond are out years. In this case, the 2004 Department of Navy spending is as submitted for congressional approval following reviews by the Office of the Secretary of Defense and Office of Management and Budget.

simplicity, we will consider only continuous quantities hereafter.

Colors of Money

Money spent on military assets is often restricted to a specific funding category, or "color of money" (Figure 3). For example, navy money for aircraft is categorized as "Air Procurement, Navy," and ship money is called "Shipbuilding and Conversion, Navy." Each category of money is associated with its own restrictions, such as the time by which and the way in which the money must be spent; additional restrictions may include the rate at which and the assets on which the money can or must be spent. We account for these categories by adding an index for funding category c and modify our budget constraint slightly:

$$\sum_a \left(fixedcost_{ac} SELECT_a + \sum_{w \in w(a)} varcost_{acw} QUANTITY_{aw} \right) \le budget_c \quad \forall c.$$

Interactions Among Decisions

Some acquisition options may require or preclude others. For example, the US Army may have 10 acquisition options for a new tank (weapon system) and may have to select at most one. Newman et al. (2000) give examples, such as a satellite that, if funded, requires a launch vehicle. The options governing the satellite and the launch vehicle are otherwise completely independent.

We define a coercion set as a group of acquisition options that share some restriction associated with selecting each of them.

Common coercion sets include

"select at most, exactly, or at least k of these acquisition options";

"select this acquisition option to be able to select any option in that set"; and

"if you select any acquisition option in this set, then you must also select at least one in that set."

We need these coercions, for example, to keep a shipyard open, maintain redundant sources, exercise a contract option, or limit the number of simultaneous selections.

Synergy

The effects of weapon system contributions are varied and often interact. Given weapon system w procured under acquisition option a and weapon system w' procured under acquisition option a', we can model pairwise interactions that do not depend on the quantity procured. We use an additional binary variable $BOTH_{aa'}$ that has value one when both a and a' are purchased, along with the following linear constraints:

$$BOTH_{aa'} \le SELECT_a,$$
$$BOTH_{aa'} \le SELECT_{a'}, \quad \text{and}$$
$$BOTH_{aa'} \ge SELECT_a + SELECT_{a'} - 1.$$

Such interactions may be synergistic. For instance, a precision weapon and a target designator may each exhibit marginal improvements on their own but together offer dramatically improved effectiveness.

Multiple-Year Planning Horizon

Most capital planning for major weapon systems extends at least over the budget planning horizon of six or eight years or so, if not over the likely lifetime of the systems, but no further than we are willing to risk forecasting the future. When considering a planning horizon as long as 20 or 30 years, we usually keep track of the year in which a weapon system starts service and perhaps the year in which it stops service. The former can correspond to the acquisition decision year, the payment year, or the first service year, or *cohort* year. For ease of exposition, we will initially assume that these years coincide, although reality is more complicated; we later consider the fact that time lags usually separate these events. We also track the weapon systems in inventory (adding new purchases and deducting retirements of old weapon systems), and we can account for operating costs that vary with the service life or age of each system. An acquisition option a is endowed with specific start and stop years as well as minimum and maximum yearly purchase quantities for its associated weapon system(s). For multiple-year planning, converting costs to some base present-value year is an inestimable convenience: For military planning, various organizations publish discount rates, for example, the US Office of Management and Budget (2004).

This gives rise to a generic multiple-year model:

1. Indices and Index Sets

$a =$ acquisition option.

$c =$ color of money.

$w =$ weapon system.

$y =$ year, alias $\underline{y}$.

$w(a) =$ set of weapon system(s) procured under acquisition option a.

2. Parameters [units]

l_{awy} (u_{awy}) = lower (upper) limit on number of weapon system $w \in w(a)$ purchased in year y under acquisition option a [w-units].

$fixedcontr_{awy}$ = fixed contribution of weapon system $w \in w(a)$ towards the MOE in year y under acquisition option a [value units].

$varcontr_{awy}$ = variable contribution per unit purchased of weapon system $w \in w(a)$ towards the MOE in year y under acquisition option a [value units/w-unit].

$fixedcost_{acy}$ = fixed cost in color of money c in year y incurred by selecting acquisition option a [\$].

$varcost_{acwy}$ = variable cost in color of money c in year y per unit purchased of weapon system $w \in w(a)$ under acquisition option a [\$/w-unit].

$budget_{cy}$ = available color of money c budget in year y [\$].

3. Decision Variables

$SELECT_a = 1$ if any units are purchased under acquisition option a, $= 0$ otherwise [binary].

$QUANTITY_{awy}$ = number of units purchased of weapon system $w \in w(a)$ under acquisition option a that begin operation at the end of year y [w-units].

$SERVICE_{wy\underline{y}}$ = number of units of weapon system w in service during year y that first served at the end of year $\underline{y}$ [w-units].

$RETIRE_{wy\underline{y}}$ = number of units of weapon system w taken out of service at the end of year y that first served at the end of year $\underline{y}$ [w-units].

4. The Corresponding Linear Integer Program

$$\max \sum_{a,y,w\in w(a)} \left(fixedcontr_{awy} SELECT_a + varcontr_{awy} QUANTITY_{awy}\right)$$

$$\text{s.t.} \quad \sum_a \left(fixedcost_{acy} SELECT_a + \sum_{w\in w(a)} varcost_{acwy} QUANTITY_{awy} \right) \le budget_{cy} \quad \forall c, y,$$

$$l_{awy} SELECT_a \le QUANTITY_{awy} \le u_{awy} SELECT_a \quad \forall a, w\in w(a), y,$$

$$\sum_{a \mid w\in w(a)} QUANTITY_{awy} = SERVICE_{w,y+1,y} \quad \forall w, y,$$

$$SERVICE_{w,y,\underline{y}} = SERVICE_{w,y+1,\underline{y}} + RETIRE_{w,y,\underline{y}} \quad \forall w, y, \underline{y} < y,$$

$$SELECT_a \in \{0,1\} \quad \forall a,$$

$$QUANTITY_{awy} \ge 0 \quad \forall a, w\in w(a), y,$$

$$SERVICE_{wy\underline{y}} \ge 0 \quad \forall w, y, \underline{y},$$

$$RETIRE_{wy\underline{y}} \ge 0 \quad \forall w, y, \underline{y}.$$

A Single Procurement Can Accrue Multiple-Year Fixed and Variable Costs

When we select an acquisition option a, it may inflict fixed and variable costs over many years. The generalized $fixedcost_{acy}$ and $varcost_{acwy}$ parameters allow us to schedule both the fixed and variable costs for each acquisition option annually and make them payable over many years before and after the weapon system begins service. These cost parameters allow a fixed lag between the time a system is paid for and the time it begins service.

Aged Inventory

Costs, such as operating and maintenance costs, may vary by planning year and also by the age of the system. To this end, we can define $varcost_{acwy\underline{y}}$ as the variable cost in color of money c during year y for a weapon system w under acquisition option a given that the system is $y - \underline{y}$ years old. $SERVICE_{wy\underline{y}}$ represents the units in service during year y that are $y - \underline{y}$ years old (Figure 4).

Similarly, we may need to force overhaul and retirement decisions by constraining maximum service life or some burdened function of service life and usage rate.

Aged inventory is always an issue, so it is curious that textbooks rarely mention it.

Long-Term, Cumulative Budget Goals

We can relax the budget constraint in the above formulation by accumulating both expenditures and the

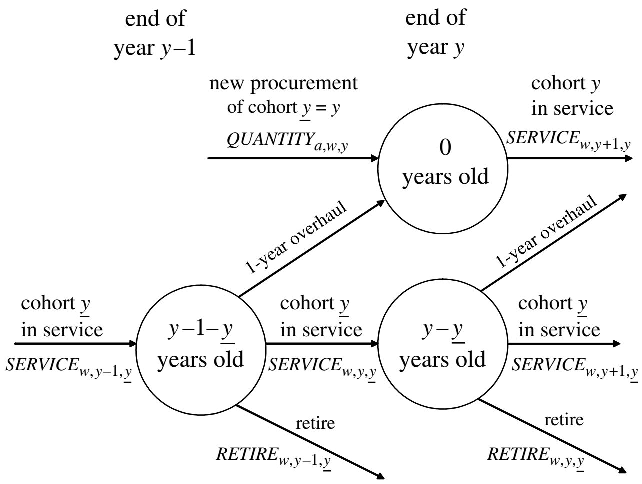

Figure 4: To track the age of an asset, we must account for both its introduction (cohort) year $\underline{y}$ and its service year; each cohort is a unique commodity, and we need conservation-of-flow constraints that retain this distinction. Some service-life-extension actions can overhaul an old asset to create a new one. Inventory aging is an essential complication when, for example, maintenance cost varies by age or if the asset has a maximum planned service life; this can dramatically increase the size of long-term capital-planning models.

budget allowance up to any current period to produce the following set of *cumulative constraints*:

$$\sum_{a}\sum_{y'=1}^{y}\left(fixedcost_{acy'}SELECT_{a} + \sum_{w\in w(a)} varcost_{acwy'}QUANTITY_{awy'}\right) \le \sum_{y'=1}^{y} budget_{cy'} \quad \forall c,\, y.$$

In this way, we can retain unused funds from one year to pay for an acquisition in a subsequent year with greater benefit. In reality, we may not be able to apply past funds to future years: many budgets are granted on a use-or-lose basis. However, by using a model to forecast when we need funds, we may be able to request a priori a distribution of funds to match the optimal requirement.

Because the individual costs of acquisition options are high relative to budgeted funding categories, it can be very difficult to select a portfolio of acquisitions for which the annual outlays fit exactly in each funding category in each planning year of the horizon. In other words, a superficially simple annual budget constraint over a long planning horizon is ridiculous in the real world.

To address this recurring problem, some planning models employ a budget band over the planning horizon made up of yearly lower and upper bounds on each funding category in the short term and with larger bands (greater separation between the upper and lower bounds) farther into the future to reflect planning uncertainties. The use of these bands provides some reasonable degree of freedom as to when funds are spent, even if the total amount spent does not change.

Even with budget bands, we can encounter an optimal solution that is silly, such as leaving a large amount of money unused in some category and year, even though with just a few dollars more we could select an attractive alternative. To avoid such foolishness, we create an *elastic constraint* on the budget. Rather than overlook some solution that is almost, but not quite, feasible, we allow ourselves to violate a budget band, albeit at a high elastic penalty cost per unit of violation. Planners regard an insightful solution with small, cosmetic elastic violations much more positively than a strictly feasible solution that nobody likes.

With elastic yearly budget bands, we can still encounter an optimal solution with myopic elastic violations over the planning years: violations that, after some analysis, we can shift to sooner or later to reduce their number or severity. To capture this in the optimization model, we employ *cumulative elastic constraints*, replacing each (for example, upper) limit each year by the sum of all such limits from the first year up to and including that year. In this case, an elastic budget violation in any year keeps reappearing and inflicting further penalties in later years, unless and until we offset (repay) it through some compensating later event.

State Transitions

Military capital planning exercises monopsony over large industrial sectors, key materiel and components.

To support the US national military strategy, the national technology and industrial base includes plants that are government owned and government operated (GOGO), government owned and contractor operated (GOCO), and contractor owned and contractor operated (COCO). These plants have industrial capability that is critical to the US military, use resources unique to the military, and have capacity far in excess of peacetime needs.

For instance, the US Army manages the country's conventional ammunition production base that meets the peacetime demands of the armed services and maintains capacity to replenish wartime consumption (Bayram 2002). Plants carry out production on individual production lines that make up their production-line complexes, and they occupy pieces of real estate, which may be large when the operations are hazardous. Each individual line, plant, and piece of real estate can be managed as GOGO, GOCO, or

COCO, independent of the management mode of any other component. In addition, components may be intentionally sold (disposed of), idled (mothballed), operated minimally (kept warm), or operated in full shifts.

The US Army's capital planning for ammunition decides not just how much of what to make when and store where, but where to locate this production infrastructure and how to manage and operate it. Over the long term, it can change the locations of plant equipment, the management modes of lines, plants, and real estate, and the operating state of any component.

Some of the greatest discretionary planning costs can accrue from dictating such changes. Accordingly, we need to model state transition costs for changing locations, management modes, and operating states to any feasible subsequent combination of these factors. This is a dramatic generalization of the classic, textbook binary close-open variable multiplied by a fixed cost to open.

One simple way to capture a transition is with a binary variable $OPERATE_{s,s',y}$ that has value one when the component is operating in state s (for example, GOGO mothballed) in year $y-1$ and makes a transition to state s' (for example, GOGO open) in year y. We add the transition costs as functions of these decision variables to the objective function and track the yearly operating state using such constraints as

$$\sum_{s,s'} OPERATE_{s,s',y} = 1 \quad \forall y \quad \text{and}$$

$$\sum_{s'} OPERATE_{s',s,y-1} = \sum_{s'} OPERATE_{s,s',y} \quad \forall s, y.$$

Time Dependencies Among Decisions

Operational considerations give rise to coercion sets that ensure continuity of mission availability between time periods, in addition to pairwise interactions that can be applied to one or more time periods. For example (Figure 5), we can use a coercion subset to denote a weapon system w (or collection of weapon systems), with its corresponding start- and stop-service years, $\underline{y}$ and $\bar{y}$, whose operation is dependent upon another weapon system w', with its own corresponding start- and stop-years, $\underline{y}'$ and $\bar{y}'$.

Weapon system w procured under acquisition option a may be required to operate concurrently with weapon system w' procured under acquisition option a', provided $\underline{y} \le \underline{y}' \wedge \bar{y} \ge \bar{y}'$. We term this *concurrent* operation and impose the constraint

$$SELECT_{a'} \le SELECT_a.$$

We may be required to operate a weapon system w procured under one of the acquisition options $a \in \Omega$ prior to operating weapon system w' procured under acquisition option a', provided $\underline{y} \le \underline{y}'$. We term this *prerequisite operation* and impose the constraint

$$SELECT_{a'} \le \sum_{a \in \Omega} SELECT_a.$$

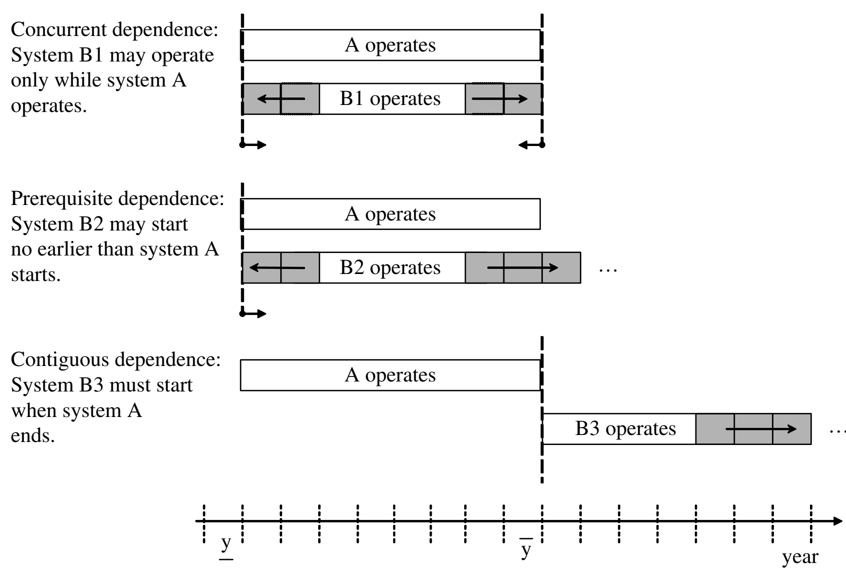

Figure 5: When we begin operating a new system can depend on the timing of the other systems' operations. System A starts operating at the end of year $\underline{y}$ and stops at the end of $\bar{y}$. The concurrent dependence of candidate system B1 on A restricts it to operating only when A does. The prerequisite dependence of B2 on A prevents it from starting before A does. The contiguous dependence of B3 on A requires it to start operating right after A ceases operation.

We may be required to operate system w procured under exactly one of the acquisition options $a \in \Omega$ immediately after weapon system w' procured under acquisition option a', provided $\underline{y}' = \overline{y}$. We term this *contiguous operation* and impose the constraint

$$SELECT_{a'} = \sum_{a \in \Omega} SELECT_a.$$

Persistence

After we refine a plan, and perhaps promulgate it to senior leaders, we may have to revise it to accommodate changes. Optimization has a well-earned reputation for amplifying small changes to input parameters into breathtaking changes to output plans. Often disruptive changes turn out to be unnecessary because they change total costs very little. Brown et al. (1997) explain how to incorporate persistence in optimization-based decision-support models.

For example, suppose that we have a binary legacy capital plan $\mathbf{select}_a^*$ from which we derive the revision $SELECT_a$. The linear expression

$$\sum_{a \mid \mathbf{select}_a^*=0} SELECT_a + \sum_{a \mid \mathbf{select}_a^*=1} (1 - SELECT_a)$$

sums the number of changes (the Hamming distance) between the legacy plan and the revision, which can be constrained or elastically penalized. Restricting the differences between the two solutions, that is, providing an upper bound on this expression, will not reduce costs but frequently reveals alternate solutions that cost little more and exhibit much less turbulence.

End Effects

Long-term capital-planning models have finite planning horizons. They require beginning and ending states as input parameters. When our purpose in long-term planning is to advise how to evolve, it seems paradoxical that we must specify the end state as an input, rather than learn it as an output. Worse, the end state is a long time from now and not easy to reckon. Errors and omissions in determining end states can lead to *end effects*: outrageous behavior at the end of the planning horizon.

We should use common sense in addressing end effects. One way to mitigate end effects is to extend the planning horizon beyond those years actually reported. (Many analysts do this.) The theoretical and practical problem is to determine just how long to extend the time horizon.

When Objectives Are Constraints, and Vice Versa: Dealing with Multiple, Conflicting Measures of Effectiveness

We can use a weighted average to try to coerce hierarchy among component objectives, assuming that we can establish a well-ordered hierarchy. However, doing so is problematic, even for just two objective components. Textbook descriptions of the big-M multiplier method for achieving a feasible and optimal solution illustrate the problem. Just how big does big-M, the objective weight per unit of constraint infeasibility, have to be to guarantee a hierarchical distinction between feasibility and optimality? Our military world record is a model sent to us with 14 hierarchical objectives: even if we give each objective a weight just one order of magnitude higher than the next lower objective, the resulting weighted-average objective would exceed the mantissa length of our floating-point computer arithmetic, so (even without a course in numerical analysis) you can see that we have inflicted a worrisome, if not overwhelming, rounding-error noise on our objective.

We can achieve purely hierarchical solutions without weighted averages (Steuer 1986, Chapter 9). First, we should optimize only with the highest-order objective and then state an aspiration constraint requiring at least the resulting optimal value of this objective function. We then add the aspiration constraint to the existing set of constraints and reoptimize, this time with the second-highest order objective. We repeat this process with each successive lower-level objective.

Pursuing strict hierarchies among conflicting objectives can obscure good trade-offs. We can then relax our aspiration constraint for each objective to an elastic aspiration constraint that expresses some goal for achievement and allows its violation with a linear elastic penalty.

Sometimes we are given only a list of MOEs and are left to determine what the aspiration levels should be. This leads to a series of optimization problems in which each requirement takes its turn in the objective, while its cohorts play the role of elastic aspiration constraints. That is, we empirically discover MOE levels that admit efficient solutions. A common heuristic is to cycle through the MOEs in some priority order, finding the extremal (maximizing or minimizing) value of each, setting some fraction of this value as its aspiration, and continuing to the next MOE.

Bruggeman (2003) uses a *procurement tier*, a set of munition-specific inventory targets that together achieve some level of effectiveness with respect to the (budget-infeasible) mission requirements. He prescribes procurements that achieve all munition levels in a tier before purchasing in another one for any of 40-odd categories of navy munitions over an eight-year planning horizon, keeping track of the influence each procurement has on its civilian supplier, volume price breaks, and future maintenance costs. These procurement plans use available yearly budgets to ramp up the military effectiveness of the inventories achieved, given that the total requirement far exceeds the budgets.

Representing an Uncertain Future with War Plan Scenarios

Military capital acquisitions maintain and strengthen the capability of our forces to successfully complete their missions. Although we do not know which future missions we will actually be ordered to complete, we do have a set of future scenarios (war plans) depicting a variety of representative conflicts worldwide. These scenarios are constantly gamed and maintained by our armed services, and any major capital acquisition will be evaluated by some means to assess its contribution to the urgent necessity to, for instance, prevail in one conflict while forestalling another, then prevail in the second conflict.

For example, Borden (2001) justifies his conclusion that 11 new T-AKE logistics ships are required by demonstrating with an optimization model that schedules such ships supporting urgent deployment of multiple naval strike groups worldwide over a planning horizon of 120 days. He demonstrates exactly how his new fleet would best serve the strike groups in each of about a dozen scenarios.

Although we admire stochastic optimization and teach its virtues, we have no military capital-planning client willing to commit to the required stochastic representation of future military exigencies. It's hard to sell even simple decision theory: military planners worry about what is possible, rather than what is likely.

Lessons from Computational Experience

Military capital-planning problems typically concern numerous assets, for example, weapon systems, munitions, platforms, and vehicles; many years in the planning horizon; and many acquisition options, which are limited only by the procurement planners' imaginations and the competing contractors' responsiveness. As a result, optimization-based decision-support models of military capital-planning problems are large and complex, typically resulting in many thousands of discrete and continuous variables and thousands of constraints. Most of these models require concave cost minimization.

Not all such models can be solved quickly. In our view, a model is tractable only if we can rely on it to produce a useful answer while we still remember the question. But important problems deserve serious analysis. We military analysts have lots of computing power. We think nothing of solving hundreds or thousands of planning scenarios. When the objective is billions of taxpayer dollars and the result is a durable decision to invest in the future defense of our country, analysis is everything.

We have an additional couple of tricks that can make large capital-planning models easier to solve. The US Air Force uses a capital-planning model (Newman et al. 2000) that features about 10,000 variables and about 17,000 constraints; between two and 10 percent of the decision variables must be binary. The model produces answers within two percent of optimality in about three minutes on a Silicon Graphics ONYX2 workstation with four gigabytes of RAM using the CPLEX solver, Version 6.5 (ILOG 2002).

A capital-planning model designed for the US Navy to use in planning procurement of ships and aircraft for the next 25 years (Salmeron et al. 2002) has about 167,000 variables (about 6,000 binary) and about 114,000 constraints. It produces heuristic solutions in a second or two and solutions within 10 percent of optimality in about seven minutes on a 1 GHz Pentium III computer with one gigabyte of RAM using the CPLEX Solver, Version 6.5.

Time Discount Rate and Model Mischief Discount Rate

In any real-world, long-term planning model, analysts make allowances for bad events, unavoidable despite optimization. Given a choice, we prefer to delay the effects of bad news as long as possible into the future by discounting the penalties for such events at a higher rate than their companion costs in the models. We call this the fog-of-far-future planning factor or the model-mischief discount rate. Discounting such penalties attenuates the influence of far-future constraints that can cause trouble. Rarely is the part of the solution corresponding to an out year in a long-term model used as operational guidance in the short term. Therefore, we would rather have a solvable model that provides specific, and good, guidance in the short term, than an unsolvable model, incapable of distinguishing between admissible solutions, that tries to provide a good solution in the far future.

We use present value for all costs and with this reference point use discount rates (Newman et al. 2000) to model capital planning for the US Air Force Space Command. Base-case, nondiscounted model instances require about an hour and a half to solve to an optimality gap (a difference between the value of the best solution found and a bound on the best solution potentially obtainable) of 10 percent. However, applying a 2.5 percent annual discount factor reduces solution times for these same instances to between six minutes and an hour to solve to the same 10 percent optimality gap. Analysis of the discounted model solutions reveals no degradation in solution quality.

Salmeron et al. (2002) express all US Navy procurement, operation, and maintenance costs in constant-year dollars but add an extra inflation factor to realistically represent operation and maintenance costs for older aircraft. The improvement they obtain in solution effort is so dramatic that they no longer attempt nondiscounted base cases. They also allow violations of cumulative budget bands and use a model mischief discount rate to move any such violations as far as possible into the future. Doing so proves worthwhile, because their model also exhibits end effects. For example, near the end of the long planning horizon, ships are forced to retire without alternate replacements on the drawing boards. It is better to deliver workable advice with excellent near-term fidelity than to let this far-term blemish shatter the entire planning exercise.

Relaxation and Aggregation

We seldom decide at the outset to relax problem features or to aggregate detail: these sacrifices of model fidelity are attractive only after daunting computational experience proves them essential.

A decision to select an acquisition option in a particular year may have to be binary in the near term but not in the far term. We must obtain integer acquisition quantities if the quantities are small, especially in the near term, but we frequently relax integer-decision requirements in the intermediate and far terms, which can dramatically improve model responsiveness.

For example, we might replace binary alternatives to select some acquisition option in exactly one out of a set of future years with a relaxation that permits us to select the alternative fractionally during that epoch but fully by its end. We can thus spread planned investments over years in the future, rather than being forced to make them in some particular year (Newman et al. 2000). Salmeron et al. (2002) use continuous quantities for aircraft and all retirements, and they account for the vast majority of the nearly 100,000 decision variables. Some ship quantities may be continuous in the far future, permitting planned procurements to span planning years.

These relaxations are not always easy to express, for example, when interactions exist among far-term decisions. We can use aggregation to limit model sizes and hasten planning cycles while reducing the workload of preparing parameters for far-future alternatives. In the near term, we might want a diversity of alternatives, while in the far term we might just have one alternative to select or reject. In the near term, we might need yearly time fidelity, while in the far term we can aggregate yearly constraints and variables to model an entire decade, reflecting the reality that timing selections in the out years is less precise. As the planning horizon rolls forward, we eventually disaggregate all model features and amplify them to then-present value.

We routinely relax our tolerance of the optimality gap to, say, 10 percent. That may appear too coarse, but many of our optimization models merely compare alternatives. In this case, and as long as the gap (or interval of uncertainty) of the winner is strictly better than that of each loser, we obtain more benefit from a faster response time and a larger optimality gap, than from a smaller gap and no additional information as to which solution we prefer. Even if we are not comparing alternatives, we are contrasting a modus operandi with an optimization solution, and if the objective function of the modus operandi lies outside of the optimality gap interval, then we can conclude that the solution gained through optimization is preferable to the modus operandi. In either case, once we have found a winner, or a portfolio of winners, we can try to reduce the gap of each of these as much as possible, even if this means a lot of computing. Fortunately, there are very few excursions that survive our competitions long enough to require this extra effort.

Conclusion

The military has employed mathematical optimization of capital planning for over 50 years and has made many contributions to this field. While military capital planning differs from private-sector capital budgeting—principally in the sheer amounts of money involved and the long planning horizons—anyone making long-term capital investments likely faces many of the same issues the military does. Many of these recurrent planning issues can be expressed in mathematical models. We recommend the references in our bibliography for more detail.

Acknowledgments

We acknowledge Kevin Wood of the Naval Postgraduate School and J. K. Newman of the University of Illinois for their helpful comments and suggestions and thank Janel Brooks for retrieving some background material.

References

Bailey, R. 1953. Application of operations-research techniques to airborne weapon system planning. *J. Oper. Res. Soc. Amer.* **1**(4) 187–199.

Bayram, V. 2002. Optimizing the capacity and operation of US Army ammunition production facilities. MS thesis, Naval Postgraduate School, Monterey, CA.

Bernhard, R. 1969. Mathematical programming models for capital budgeting—A survey, generalization, and critique. *J. Financial Quantitative Anal.* **4**(2) 111–158.

Bertsimas, D., J. Tsitsiklis. 1997. *Introduction to Linear Optimization*. Athena Scientific, Belmont, MA.

Borden, K. 2001. Optimizing the number and employment of combat logistics force shuttle ships, with a case study of the new T-AKE ship. MS thesis, Naval Postgraduate School, Monterey, CA.

Bradley, G. 1986. Optimization of capital portfolios. *Seminar OPS-01, National Communications Forum* **40**(1), Chicago, IL, 11–17.

Brown, G., D. Coulter, A. Washburn. 1994. Sortie optimization and munitions planning. *Military Oper. Res.* **1**(1) 13–18.

Brown, G., R. Dell, A. Newman. 2004. Optimization-based military capital planning. NPS technical report NPS-OR-04-002, Naval Postgraduate School, Monterey, CA.

Brown, G., R. Dell, R. Wood. 1997. Optimization and persistence. *Interfaces* **27**(5) 15–37.

Brown, G., R. Clemence, W. Teufert, R. Wood. 1991. An optimization model for modernizing the army's helicopter fleet. *Interfaces* **12**(4) 39–52.

Brown, G., R. Dell, H. Holtz, A. Newman. 2003. How US Air Force Space Command optimizes long-term investment in space systems. *Interfaces* **33**(4) 1–14.

Bruggeman, J. 2003. A multi-year ammunition procurement model for Department of the Navy non-nuclear ordnance. MS thesis, Naval Postgraduate School, Monterey, CA.

Chambers, J., ed. 1999. *American Military History*. Oxford University Press, Oxford, U.K.

Dantzig, G. 1963. *Linear Programming and Extensions*. Princeton University Press, Princeton, NJ.

Everett, H. 1963. Generalized Lagrange multiplier method for solving problems of optimal allocation of resources. *Oper. Res.* **11**(3) 399–417.

Field, R. 1999. Planning capital investments in navy forces. MS thesis, Naval Postgraduate School, Monterey, CA.

Gunther, P. 1955. Use of linear programming in capital budgeting. *J. Oper. Res. Soc. Amer.* **3**(2) 219–224.

ILOG. 2002. CPLEX 6.5. Retrieved July 2002 http://www.cplex.com.

Keeney, R. 1992. *Value-Focused Thinking*. Harvard University Press, Cambridge, MA.

Kimms, A. 2001. *Mathematical Programming and Financial Objectives for Scheduling Projects*. Kluwer Academic Publishers, Norwell, MA.

Loerch, A. 1999. Incorporating learning curve costs in acquisition strategy optimization. *Naval Res. Logist.* **46**(3) 255–271.

Lorie, J., L. Savage. 1955. Three problems in rationing capital. *J. Bus.* **28**(4) 229–239.

Mamer, J., A. Shogan. 1987. A constrained capital budgeting problem with applications to repair kit selection. *Management Sci.* **33**(6) 800–806.

McCullough, D. 2001. *John Adams*. Simon and Schuster, New York.

Meier, H., N. Christofides, G. Salkin. 2001. Capital budgeting under uncertainty—An integrated approach using contingent claims analysis and integer programming. *Oper. Res.* **49**(2) 196–206.

Newman, A., G. Brown, R. Dell, A. Giddings, R. Rosenthal. 2000. An integer linear program to plan procurement and deployment of space and missile assets. NPS technical report NPS-OR-00-005, Naval Postgraduate School, Monterey, CA.

Parnell, G., H. Conley, J. Jackson, L. Lehmkuhl, J. Andrew. 1998. Foundations 2025: A value model for evaluating future air and space forces. *Management Sci.* **44**(10) 1336–1350.

Ricks, T. 1999. Navy to stay with joint fighter despite cost overrun at Lockheed Martin. *The Wall Street J.*, Section B, p. 5A, col. 3, February 26.

Rychel, D. 1977. Capital budgeting with mixed integer linear programming: An application. *Financial Management* **6**(4) 11–19.

Salmeron, J., G. Brown, R. Dell, A. Rowe. 2002. An optimization-based decision-support system to plan procurement and retirement of naval platforms. NPS technical report NPS-OR-02-006, Naval Postgraduate School, Monterey, CA.

Stanley, E., D. Honig, L. Gainen. 1954. Linear programming in bid evaluation. *Naval Res. Logist. Quart.* **1**(1) 48–54.

Steuer, R. 1986. *Multiple Criteria Optimization: Theory, Computation, and Application*. John Wiley and Sons, New York.

Taylor, B., A. Keown, A. Greenwood. 1983. An integer goal programming model for determining military aircraft expenditures. *J. Oper. Res. Soc.* **34**(5) 379–390.

United States Department of the Navy. 2003. Future years defense program, January 2003 submission of the president's budget for fiscal year 2004.

United States Office of Management and Budget. 2004. Retrieved January 2004 http://whitehouse.gov/OMB.

Weingartner, H. 1966. Capital budgeting of interrelated projects: Survey and synthesis. *Management Sci.* **12**(7) 485–516.

Weingartner, H. 1977. Capital rationing: n authors in search of a plot. *J. Finance* **32**(5) 1403–1431.