DECISION SCIENCES INSTITUTE

A Portfolio Optimization Model for Investment Planning in the Department of National Defence and Canadian Armed Forces

(Full Paper Submission)

Mark Rempel
Centre for Operational Research and Analysis
Defence Research and Development Canada
mark.rempel@drdc-rddc.gc.ca

Chad Young
Centre for Operational Research and Analysis
Defence Research and Development Canada
chad.young@drdc-rddc.gc.ca

ABSTRACT

Investment planning is a critical activity for military organizations. Canada's Department of National Defence recently began an initiative to more explicitly incorporate project portfolio management principles into their investment planning process. To facilitate the first stages of this transition a decision support system was developed. It combines optimization, visualization and manual revision features to aid in the production of preferred project portfolios. This paper presents the portfolio construction process, the underlying integer programming model, and several of the visualizations employed to enhance decision maker interaction. Design rationale and insights gleaned from its initial use are also introduced.

KEYWORDS: Mathematical Programming, Capital Budgeting, Public Sector, Decision Support Systems

INTRODUCTION

Like most modern defence organizations, Canada's Department of National Defence (DND) is a large, complex and resource intensive enterprise that is required to operate across a diverse set of physical, political, social and financial environments. Managing any large enterprise like DND requires an analytical framework for investment planning that allows executive management to maximize organizational effectiveness within its available resources and risk tolerance.

Many western nations have gravitated toward a Capability-Based Planning (CBP) approach to inform their investment planning process (Government of Australia, 2013; Government of Canada, 2008; Davie, 2013; Davis, 2002; Taylor, 2014). Different nations have implemented CBP principles in different ways. At DND, the introduction of CBP over the last decade has provided an important source of planning information. Among other benefits, CBP has been particularly valuable to senior executive boards for purposes of screening new project proposals from the perspective of the larger enterprise. Additionally, information and analysis concerning cost, resource and policy alignment derived from the CBP process has been useful for purposes of strategic framing. However, in its current implementation at DND, it remains a challenge to create strong links between CBP and the creation of a long-term investment plan.

Recent CBP implementations at DND have followed a waterfall process with three phases: (1) a capability *assessment* phase; (2) a capability *management* phase; and (3) a capability *integration* phase (Taylor et al., 2008; Rempel, 2010). Initially, the last of these phases incorporated a sophisticated optimization model designed to help decision makers choose among the different ways that individual capabilities might be delivered in accordance with enterprise-wide constraints. It was intended that this phase would directly facilitate investment planning. However, since 2010 these sophisticated methods in the last phase have been largely abandoned. Instead, simpler but more labour intensive methods that were more familiar to the staffs responsible for investment planning were utilized. Essentially, the investment planning process reverted to a ranked-list approach. In this approach, candidate investment elements, hereafter referred to as projects, are prioritized and a financial cut-line is used to identify those projects that are suitable for investment (note: Canada's defence spending in 2013 was 1.0% of its gross domestic product, i.e., 18.3 Billion US (World Bank, 2014)). Projects are then scheduled by way of negotiation between project stakeholders, internal programming authorities, financial authorities, and central government agencies. Rescheduling occurs as conditions change. Relatively simple analysis tools such as spreadsheets have been used to facilitate these activities.

On the surface, the ranked-list approach appears to be very practical. However, this approach can become too focused on short-term affordability and often exhibits too much of a project-by-project focus rather than a broader strategic focus. In addition, it is inefficient and resource intensive, in terms of both time and personnel, because it is typically executed as a series of sequential steps performed by separate internal organizations (that evaluate project importance, establish financial requirements, determine organizational capacity requirements and schedule projects) without a consistent view of the whole portfolio. Furthermore, it is often performed for different project categories in isolation from each other (e.g., equipment, infrastructure, informatics). This can effectively lead to the production of multiple "stove-piped" investment plans that must be merged together – after which, there may be reluctance to make alterations even when conditions change.

Mindful of these and other drawbacks, DND recently began new efforts to transition their investment planning process to an approach that incorporates Project Portfolio Management (PPM) principles (Rempel and Young, 2014; Young and Rempel, 2014). This paper introduces a decision support system called Visual Investment Plan Optimization & Revision (VIPOR) that has been designed to support this transition. It leverages advanced visualizations, optimization and other features to interactively adjust a portfolio to facilitate the development of an investment plan. Improving the linkage between CBP and investment planning was an important factor in VIPOR's design. Other factors that led to the development of VIPOR included: (1) a need to enable internal review initiatives that assess and rationalize investments and budgets (Young and Rempel, 2014); (2) a requirement to address issues related to the maintenance of a sustainable defence program amidst a contractionary fiscal climate (Weltman, 2015); and (3) a requirement to support the continuing efforts within the Government of Canada to enhance the management of resources and the achievement of results.

DND investment planning and VIPOR

VIPOR incorporates two main analytical components. The first is an optimization model that takes as its input the information synthesized from CBP and other planning sources. The optimization model facilitates the construction of project portfolios that serve as a basis for creating a Departmental investment plan that contains: (1) a viable 20 year schedule of major projects and planned expenditures; and (2) realistic plans for the introduction, maintenance and/or divestment of military and corporate capabilities. The optimization model aims to help produce project portfolios that rationalize and balance all long-term investments in a way that is traceable, efficient and is in alignment with projections concerning the strategic demand and the supply of available resources. The second analytical component is a set of interactive information visualizations that: (1) facilitate characterization of project portfolios produced by the optimization model; (2) support analytical reasoning; and (3) allow a decision maker to manually adjust a portfolio and assess the impacts of those adjustments.

Outline

In the next section, we introduce background concepts around project portfolio selection and provide a brief overview of VIPOR's high-level design requirements. We then briefly discuss related work in the open literature concerning the use of optimization models to support capital investment planning in military organizations. This is followed by a detailed presentation of the optimization model. Some of the visualizations and interaction methods designed to allow a user to characterize and adjust the portfolio in an iterative fashion are described next. Finally, we provide observations, insights and areas for potential future improvement that have evolved from our efforts to develop VIPOR and facilitate its use with decision makers. The three main contributions of this paper are as follows.

- The development of a decision support system to facilitate investment planning in an organization that is starting to transition its investment planning process from a ranked-list approach to more holistic portfolio approach.
- An enterprise-wide project portfolio construction model that is designed for a military organization and accounts for financial constraints, organizational capacity constraints, and project dependencies.
- A discussion of our observations, experience and insights, many of which confirm lessons identified in some of the related literature pertaining to the application of an optimization model during transitions to an enterprise-wide project portfolio paradigm.

PROJECT PORTFOLIO CONSTRUCTION AND VIPOR'S DESIGN REQUIREMENTS

Project portfolio management (PPM) concepts and best practices have been maturing in their sophistication over the last two decades. The set of projects considered in an application of PPM may be classified into two groups: (1) projects that are being proposed; and (2) projects that are currently at some stage in their development or execution cycle. From the amalgam of these two sets there are several methods that may be used to establish and manage the execution of a subset of projects called the project portfolio (Project Management Institute, 2013a; Microsoft Corporation, 2008). Regardless of the method, the aim is to create and oversee a project portfolio with the "right set of projects" that: (1) best enables the achievement of strategic objectives; (2) creates an efficient match between the projected supply of resources and aggregate project demand; and (3) produces a high degree of comfort among executive leadership with respect to risk.

Although the overall objective – to facilitate the transition to a portfolio-based investment planning process – was clearly defined at the outset of VIPOR's development, the scope and design requirements were unclear. Therefore, in addition to reviewing academic literature concerning PPM, visual analytics, and defence investment planning, the VIPOR design team held discussions with a cross-functional investment planning working group and several executive decision makers. From these discussions, it was determined that, although there are many steps in the recognized project-portfolio management process (Project Management Institute, 2013a), VIPOR would need to support a smaller set of activities. This set is itemized below.

- **Portfolio Construction**: Facilitate the construction of a portfolio of projects that is aligned with strategic intents and respects projected financial resource constraints, organizational capacity constraints, and project dependencies over future time horizons.
- **Portfolio Characterization**: Enable the characterization of a portfolio by its aggregate attributes, such as its compositional balance and utilization of resources. In addition, enable the examination of differences between one or more portfolios.
- **Portfolio Adjustment**: Facilitate manual adjustment (either directly through the addition and/or removal of projects, by constraining available resources, or modifying the criteria weights used to compute a project's value) of a portfolio's composition in order to test the impact of changing decision maker preferences.

The process diagram in Figure 1 is an anchor point around which to understand these three requirements. It combines concepts derived from Archer and Ghasemzadeh (1999), who present an integrated framework for project-portfolio selection, and Keim et al. (2008), who describe a sense-making loop for visual analytics. The diagram shows that data pertaining to each project is first fed into a value model that determines the merit of every project based on several evaluation criteria. This is where the linkage between CBP's output and the optimization model is created. More specifically, CBP assessments are encompassed within the value model and are used to determine each project's merit.

The optimization model facilitates the construction of a feasible project portfolio. The optimization model receives input from the value model, along with information about constraints and the importance weighting assigned to each criterion in the value model.

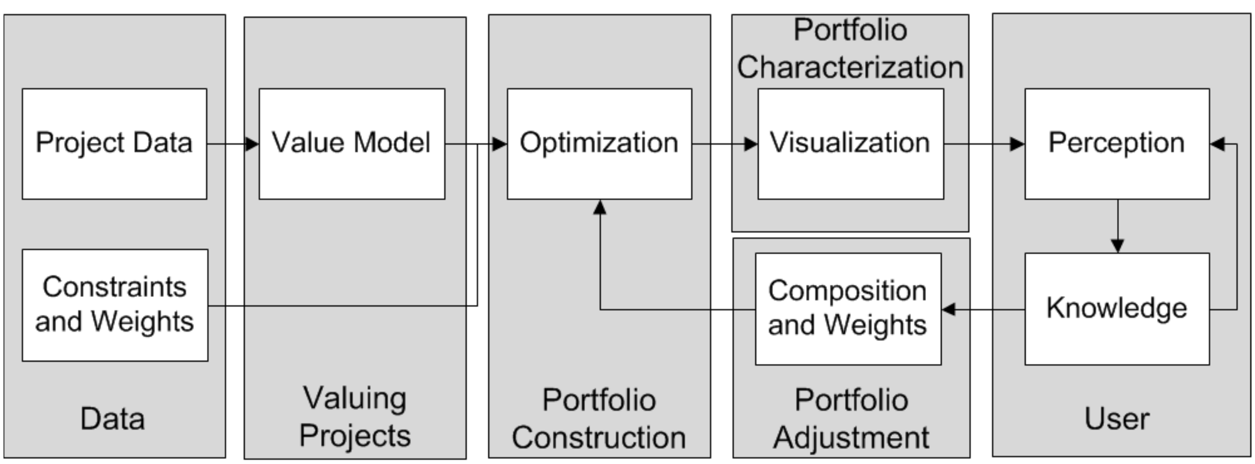

Figure 1: Within this larger process, VIPOR encompasses Portfolio Construction, Portfolio Characterization, Portfolio Adjustment and the flow of data into and between these elements.

Once constructed, a decision maker can characterize a project portfolio by making use of several interactive visualizations. Based on their impression of the portfolio's compositional balance (in addition to aspects not incorporated in the value model and implicit knowledge or fresh insights) the decision maker can then manually adjust the portfolio. These adjustments are typically made by adding and removing projects or by altering the relative importance assigned to criteria in the value model.

Providing decision makers with the ability to interactively and repeatedly adjust a portfolio is a fundamental feature of VIPOR. The adjustments made by the decision makers are incorporated into the optimization model, which is then re-run in order to produce a new portfolio that reflects the decision maker's new preferences and conforms to the pre-existing constraints. The iterative process of *construction*, *characterization*, and *adjustment* can be continued until the decision maker(s) is comfortable with the resulting portfolio. The focus of this paper is on *portfolio construction* aspect of this process. We place less emphasis in this paper on the *portfolio characterization* and *portfolio adjustment* aspects. Details about *project data*, *valuing projects* and the cognitive aspects pertaining to user *perception* and *knowledge* are not discussed.

LITERATURE REVIEW

Previous instantiations of CBP within DND have employed an optimization model to support investment planning activities (Taylor et al., 2008; Rempel, 2010). Like the majority of optimization models employed for this purpose, the model's objective was to determine the set of projects that collectively provide the maximum value to DND subject to resource constraints. Unlike many defence investment planning models, it was based on a multi-objective nonlinear optimization, rather than a single linear objective optimization. The result of this complex implementation was that a large set of candidate portfolios were presented to decision makers. This, in addition to the lack of an associated decision support system, led to portfolios that were treated with suspicion by stakeholders, particularly those who did not agree with the resulting recommendations. Since stakeholder comfort with the investment plan became paramount, optimization was removed from the CBP process and the act of creating an investment plan became "more 'art' than science" (Taylor et al., 2008).

Although optimization techniques have not been widely used in DND to support portfolio construction for investment planning, several examples from other military organizations exist in the *open* literature. Applications identified by Burk and Parnell (2011), plus additional examples, are listed in Table 1. An analysis of these examples reveals three main observations. First, although the examples share common features, each application is context specific. For example, although nearly all the models contain financial constraints, their implementation differs according to the context of the problem at hand (e.g., National Reconnaissance (Parnell, 2002); Army fleet sizing (Walmsley and Hern, 2004); Air Force research and development (Greiner et al., 2004)). Second, apart from Taylor et al. (2008) and Hoehl (2008), the majority of these examples are applications within a single military environment rather than at the enterprise-level. Third, none of these examples describes an associated decision support system for the purpose of optimizing, characterizing and adjusting a portfolio.

PORTFOLIO CONSTRUCTION

The problem of constructing a portfolio of projects from a set of candidate projects is a combinatorial optimization problem that is often referred to as a 0-1 knapsack problem. The basic model consists of: (1) a set of indivisible items, each of which has a value; (2) a single knapsack in which to place items; (3) an objective that represents a need to fill the knapsack with only those items that maximize/minimize its aggregate value; and (4) a single constraint that in some way restricts the combination of items that can be placed in the knapsack. When the basic model is extended to include several constraints and multiple knapsacks the problem becomes known as a 0-1 multiple-knapsack problem with multiple constraints. The portfolio construction problem within VIPOR is this type of problem.

The objective of the portfolio construction problem implemented within VIPOR is to select from the set of potential projects a smaller subset that provides the maximum value to DND. Each project is placed into one of two knapsacks, each knapsack representing one of the two financial sources used to fund projects within DND's hybrid financial system (Solomon and Stone, 2013). The first knapsack corresponds to DND's *accrual funding* envelope. The second knapsack corresponds to DND's *cash-based* funding envelope. With accrual-based funding, project costs are borne over the period during which project outputs are used. With cash-based funding expenses are borne "up-front" during the period in a project when assets are actually purchased. In VIPOR, individual projects can be funded from only one of these two financial envelopes. The project portfolio deemed to have maximal value for DND is composed of the collection of projects in both knapsacks.

Table 1 A selection of portfolio optimization models from the open literature with application within a defence organization. A star in the "Explicitly Iterative Approach" column indicates that a decision maker can choose projects to include in the portfolio prior to running the optimization.

		Portfolio Construction					Portfolio Characterization and Adjustment	
Author and year	**Domain**	**Objective Function**	**Financial Constraints**	**Organizational Capacity Constraints**	**Dependency Constraints**	**Problem Specific Constraints**	**Explicitly Iterative Approach**	**Decision Support System**
Brown et al. (2003)	Space	Minimize elastic penalty	Yearly budget	No	Yes	Yes	Yes	No
Ewing et al. (2006)	Army	Maximize portfolio value	No	No	No	Yes	No*	No
Taylor et al. (2008)	All	Maximize portfolio value	Yes, lump sum	No	Yes	Yes	No	No
Hoehl (2008)	All	Minimize cost	Yes, lump sum	No	Yes	Yes	No*	No
Hamill et al. (2002)	Cyber	Maximize risk mitigated	Yes, lump sum	No	No	No	No	No
Donahue (1992)	Army	Minimize deviation from goals	Yearly budget	No	Yes	Yes	No*	No
Greiner, et al. (2003)	Air Force	Maximize portfolio value	Yes, lump sum	No	No	Yes	No	No
Parnell et al. (2002)	National Recce Office	Maximize portfolio value	Yes, lump sum	No	Yes	No	Yes	No
Walmsley & Hearn (2004)	Army (ACSV fleet)	Three different options	Yes	Yes (supply constraints)	N/A	Yes	No	No

Given that investment planning at DND is conducted based on a 20 year time horizon, and that acquisitions and upgrades of military materiel, infrastructure and systems can have relatively long life spans, each knapsack represents a funding envelope that can extend 40 years past the current year.

Mindful of the funding sources, the long time horizon, and VIPOR's overall objective, the portfolio construction model includes four types of constraints. These are: (1) financial constraints that limit the amount of accrual and cash-based funding that is available in each year; (2) annual organizational capacity constraints that represent the upper limits of the organizational resource base that can be made available to implement projects; (3) dependency constraints that ensure known dependency relationships between individual projects are respected; and (4) user imposed constraints that capture new or altered decision maker preferences identified during portfolio characterization.

The remainder of this section presents more detailed descriptions of the decision variables, objective function and constraints in the portfolio construction model. At the conclusion of this section is a brief description of the methods used to solve for a solution to the modelled optimization problem.

Decision variables

The optimization model includes two decisions variables for each project. The first represents the binary decision to fund a project by accrual funding, and the second represents the binary decision to fund a project by cash-based funding. These are represented by the set of binary variables S, with elements $s_{p,f}$, where p and f are indices that denote the project and funding source respectively. $s_{p,f}$ takes on a value of 1 if project p is selected to be funded by funding source f, and 0 otherwise. In accordance with the way projects are financed at DND, the model will select only one funding source for each project included in the portfolio.

Objective function

The objective of the optimization model is to select the subset of candidate projects that collectively provide the maximum value to DND. In order to ensure that the portfolio construction process was comprehensive and understandable to executive-level stakeholders, the following decisions were made.

- **Project selection**: Projects cannot be subdivided; either a project is selected for inclusion in the portfolio or it is not. Additionally, duplicate copies of a project cannot exist in the portfolio.
- **Project value**: A value model approximates the overall value of each project using three criteria: (1) alignment with National Policy, (2) alignment with institutional/capability needs, and (3) the relative importance placed on the project by its sponsor with DND. First, a linear measureable function is used to convert the project's score against each criterion into a value. Then, an additive measurable value function (Dyer and Sarin, 1979; Kirkwood, 1997) is used to compute the overall value of each project. Decision makers can set the importance weighting for each criterion used in the additive measurable value function.

- **Portfolio value**: The total value of a project portfolio is the sum of the values of the accrual knapsack and the cash-based knapsack. The set of projects that maximize the overall value within constraints is called the optimal portfolio.

It is important to note that there are a variety of approaches that may be used to compute a project's value and the value of a portfolio. Additive, multiplicative, or other non-additive functions may be used (Dyer and Sarin, 1979). We selected an additive model to compute both project value and portfolio value. For project value, we based our decision on the fact that it is the most commonly used model in defence applications. For portfolio value, an additive model was selected because it provides an objective, transparent, and logical structure (Brown et al., 2004; Burk and Parnell, 2011) and was not seen to be overly complicated by decision makers.

The objective function, to maximize the portfolio's value, is stated as follows.

$$\max \sum_{f \in F} \sum_{p \in P} \sum_{c \in C} w_c \cdot v_{p,c} \cdot s_{p,f} \quad (1)$$

where:

- p, P are the index and set of projects;
- c, C are the index and set of criteria (i.e., National Policy, Institutional/Capability, Sponsor priority);
- w_c is the weight of criteria c; all weights must be greater than or equal to 0 and the sum of the weights must be equal to 1;
- f, F are the index and set of funding sources (i.e., accrual, cash-based); and
- $v_{p,c}$ is the value for criterion c of project p as computed in the value model.

Financial constraints

The optimization model contains a series of financial constraints which aim to ensure that the funding required to deliver projects within the portfolio does not exceed the available supply of financial resources. They also ensure that projects are funded by an appropriate funding source (i.e., accrual or cash-based). To achieve these aims the following aspects are incorporated into the model.

- **Multi-year planning horizons**: The planned annual financial demand of a project over its lifetime is used to determine the viability of including it within constraints that limit the available fiscal supply every year (Archer and Ghasemzadeh, 1999; Brown et al., 2004). In addition, to account for "end effects" the planning horizon extends beyond DND's 20-year capital investment planning horizon (Brown et al., 2004).
- **Colours of money**: The source of funding for a project is restricted based on project attributes and the business rules associated with each funding source (Brown et al., 2004). Both accrual and cash-based expenditure profiles are identified for each project. Also associated with each project is an indicator that is used to prescribe the allowable source(s) from which it can draw funds (i.e., cash, accrual or both).
- **Under- and over-programming**: An investment plan that is under-programmed is one that does not consume all available resources; conversely, an over-programmed investment plan sets out to consume more than the available resources. The former

ensures that capacity exists in the future if new investment opportunities arise or if contingency funds are required. The latter reduces the risk of allocated funds going unspent when projects are delayed. The optimization model allows for both of these planning strategies by allowing the selected projects to utilize up to a specified percentage of the available financial resources in each fiscal year.

The financial constraints are stated as follows.

$$\sum_{p \in P} r_{p,f,y} \cdot s_{p,f} \leq w_{y,f} \cdot b_{y,f}, \quad \forall f \in F, y \in Y \tag{2}$$

$$\sum_{f \in F} s_{p,f} \leq 1, \quad \forall p \in P \tag{3}$$

$$\gamma_{p,f} - s_{p,f} \geq 0, \quad \forall p \in P, f \in F \tag{4}$$

where:

- y, Y is the index and set of years;
- $\gamma_{p,f}$ 1 or 0; 1 if project p can be funded by funding source f, and 0 otherwise;
- $r_{p,f,y}$ is the requested funding by project p in funding source f in year y;
- $b_{y,f}$ is the funding limit in funding source f in year y; and
- $w_{y,f}$ is the maximum percentage of available funding in year y in funding source f that should be filled.

Equation (2) ensures that the summation of funds required to deliver projects in the portfolio does not exceed the available cash and accrual budget in each year. The parameter $w_{y,f}$ controls the available envelope of funding in year y from source f. This parameter allows for under (< 1) or over (> 1) programming. Equation (3) ensures that each project is funded by at most a single source. Equation (4) ensures that every project is funded by an appropriate funding source.

Organizational capacity constraints

The optimization model contains a series of constraints to ensure that the organizational capacity required to implement projects in the portfolio does not exceed the available capacity in each fiscal year. This constraint, which is not explicit within defence related investment planning models, ensures that if a project is in the portfolio the capacity required to deliver the project will exist. Details about the inclusion of this constraint in the optimization model are as follows.

- **Project delivery**: The model accounts for the limited delivery capacity of the three main DND organizations that execute major projects, i.e., the materiel, infrastructure, and informatics delivery organizations. A project is delivered by only one of these organizations – the preferred organization is determined according to the nature (i.e., attributes) of the project. In lieu of having information about the supply of human and consumable resources available for project delivery, or the expected demand of each project against this available supply, a fiscal proxy was employed to constrain the delivery capacity for each of the three organizations. These proxies were based on evaluations of each organization's ability to prosecute major projects over the last decade. Effectively these proxies provide an upper limit on the amount of money that an organization has the ability to spend in any given fiscal year.

The equation used to express the delivery capacity constraint is as follows,

$$\sum_{p \in P} \sum_{f \in F} \alpha_{p,o,y} \cdot s_{p,f} \leq k_{o,y}, \ \ \forall o \in O, y \in Y \tag{5}$$

where:

- o, O are the index and set of organizations with capacity limitations regarding implementation of projects;
- $\alpha_{p,o,y}$ is the capacity consumed by project p in organization o in year y; and
- $k_{o,y}$ is the capacity limit in organization o in year y.

Project dependency constraints

Also included within the optimization model are constraints that are intended to capture, in an rudimentary fashion, the dependencies between projects. There are many aspects of project dependency, from the most basic assertion that certain projects depend on others as enablers for delivery, to more advanced concepts pertaining to synergy. The requirements to collect more project data and communicate more advanced theory both increase as dependency models become more intricate. Therefore, so as not to over-complicate the model, only the simplest types of dependencies were included.

- **Project Dependencies**: The model is designed to explicitly enforce one- and two-way dependencies between projects (Archer and Ghasemzadeh, 1999; Brown et al., 2004; Burk et al., 2011). A one-way project dependency indicates that the inclusion of a particular project within a portfolio will necessitate the inclusion of another project, but not vice-versa. A two-way dependency indicates that, for two dependent projects, both must either be included or excluded from the portfolio (i.e., it is not acceptable that one project is included while the other is excluded and vice versa).

Project dependency constraints stated are as follows.

$$\sum_{f \in F} \left(s_{i,f} - s_{j,f} \right) \geq 0, \ \ (i,j) \in I \tag{6}$$

where:

- $(i,j), I$ are a single dependency and the set of project dependencies, where $i, j \in P$.

Equation (6) ensures that project dependencies are respected. I contains many project pairs (i,j). The interpretation of a project pair is that project i can be selected without project j but project j can only be selected if project i is selected. If both (i,j) and (j,i) are in the set I, then i

and j must both be selected or they must both be rejected; it would not be allowable to select one of these projects without the other.

User imposed constraints

Within VIPOR, four types of manual adjustments can be made to alter the composition of a portfolio that has been computed in a previous iteration of the portfolio construction process. A decision maker can: (1) force particular projects into the portfolio and set their funding source; (2) force particular projects into the portfolio but leave the selection of the funding source to be determined by the optimization model; (3) specify those projects that must not be included in the portfolio; and (4) set the funding source for individual projects but allow the optimization model to determine whether those projects are included in the optimal portfolio. Once a decision maker implements an adjustment, an existing constraint in the optimization model is automatically altered, or a new constraint is automatically added, so that these preferences are reflected during computation of the portfolio on the next iteration. Details are enumerated in Table 2.

Table 2 Modifications to the optimization model based on user specified adjustments.

	Adjustment Type	Description	Constraint Added/Modified	
(1)	Specify project p must be in the portfolio and its funding source must be μ (if multiple funding sources are available).	Add a new constraint that states $s_{p,f}$ must be 1.	$s_{p,\mu} = 1$	(7)
(2)	Specify that project p must be in the portfolio.	Modify Equation (1) such that the right-hand side equals 1 for project p.	$\sum_{f \in F} s_{p,f} = 1$	(8)
(3)	Specify project p must not be in the portfolio.	Modify Equation (1) such that the right-hand side equals 0 for project p.	$\sum_{f \in F} s_{p,f} = 0$	(9)
(4)	Select project p's funding source to be μ (if multiple funding sources are available).	Modify Equation (3) such that $\gamma_{p,f} = 0$ for project p and all $f \neq \mu$.	$\gamma_{p,f} - s_{p,f} \geq 0, \ \forall f \in F, f \neq \mu$	(10)

Solution methods

The model described in this section is an integer programming model and is implemented using the GNU Mathematical Programming Language (GMPL) (www.gnu.org/software/glpk/), an open source version of the AMPL modelling language (Fourer et al., 1990). Within VIPOR, the model is solved using two algorithms, each documented by Gass (1975): (1) a branch & bound algorithm is used to compute an integer solution; and (2) a simplex algorithm is used to compute a solution to a Linear Programming (LP) relaxation of the integer problem. The first solution produces a portfolio that becomes the basis for *portfolio characterization* and *adjustment* as per the VIPOR process diagram in Figure 1. By allowing for the fractional inclusion of projects, the second solution provides an upper bound on the highest obtainable aggregate value of a portfolio.

To ensure that a decision maker can interact with the model in a timely manner it is necessary to limit the amount of time allowed to compute an integer solution. However, imposing a time limit can reduce the quality of the integer solution. With an unlimited amount of time, the branch & bound algorithm will find an optimal solution, where the aggregate value of the optimal solution is close to the aggregate value of the LP relaxed solution. If a time limit is set, the algorithm returns the best integer solution found but this may not be an optimal solution. The LP relaxed solution is used to approximate the optimal solution, and therefore gauge the quality of the best integer solution.

PORTFOLIO CHARACTERIZATION AND PORTFOLIO ADJUSTMENT

To support the portfolio characterization and adjustment by a decision-maker, several intermediate tasks were identified for purposes of VIPOR development. Intermediate tasks are specific tasks that a decision maker may want to perform in order to characterize or adjust a portfolio within the iterative portfolio construction process. For each intermediate task, various visualization tasks were then identified using an abstract visualization typology (Brehmer and Munzner, 2013).

Visualization tasks were implemented via seven unique interactive visualizations: (1) a portfolio view, (2) several resource utilization charts, (3) a sensitivity plot, (4) a portfolio content comparator, (5) a portfolio value comparator, (6) a parallel coordinates plot, and (7) an adjustable scatter plot. Each visualization task may be performed using one or more these visualizations, as shown in Table 3. For example, the portfolio view facilitates characterization of the portfolio's composition. It also enables characterization of individual projects and manual adjustment of the portfolio. During the development of the visualizations, several visual encodings and user interaction methods were chosen to ensure that VIPOR would be as intuitive as possible for executive decision makers.

Table 3 VIPOR decision support system – decision domain tasks, visualization tasks, and visualizations.

Decision Domain Tasks			Visualizations						
High-Level Requirement	Intermediate Requirement	Abstract Visualization Task	Portfolio View	Sensitivity Plot	Resource Charts	Portfolio Content Comparator	Portfolio Value Comparator	Parallel Coords	Scatter Plot
Portfolio Characterization	Characterize portfolio composition balance	Identify aggregate characteristics	✓					✓	✓
		Compare specific characteristics of interest	✓					✓	✓
	Characterize resource utilization	Compare trends/profiles			✓				
		Identify disparities			✓				
	Characterize projects	Compare projects within a wider context	✓	✓					✓
		Summarize projects		✓					
	Compare portfolios	Compare portfolios across iterations				✓	✓		
Portfolio Adjustment	Adjust portfolio	Adjust portfolio composition	✓						
		Adjust parameters		✓					

In the remainder of this section, we describe four of the seven visualizations most closely connected with the optimization model: (1) the portfolio view; (2) the portfolio content comparator, (3) the portfolio value comparator, and (4) the sensitivity plot. Collectively, they allow a decision maker to: characterize both portfolios and individual projects; compare portfolios; adjust a portfolio; and prompt re-optimization. It is also important to note that the visualizations present herein show fictitious data.

Portfolio view

The portfolio view depicted in Figure 2 is a primary visualization within VIPOR. It enables a user to: (1) identify aggregate characteristics of a portfolio; (2) compare specific characteristics of interest; (3) compare projects within a wider context; and (4) produce a new portfolio through adjustment of its composition and/or project-specific parameters.

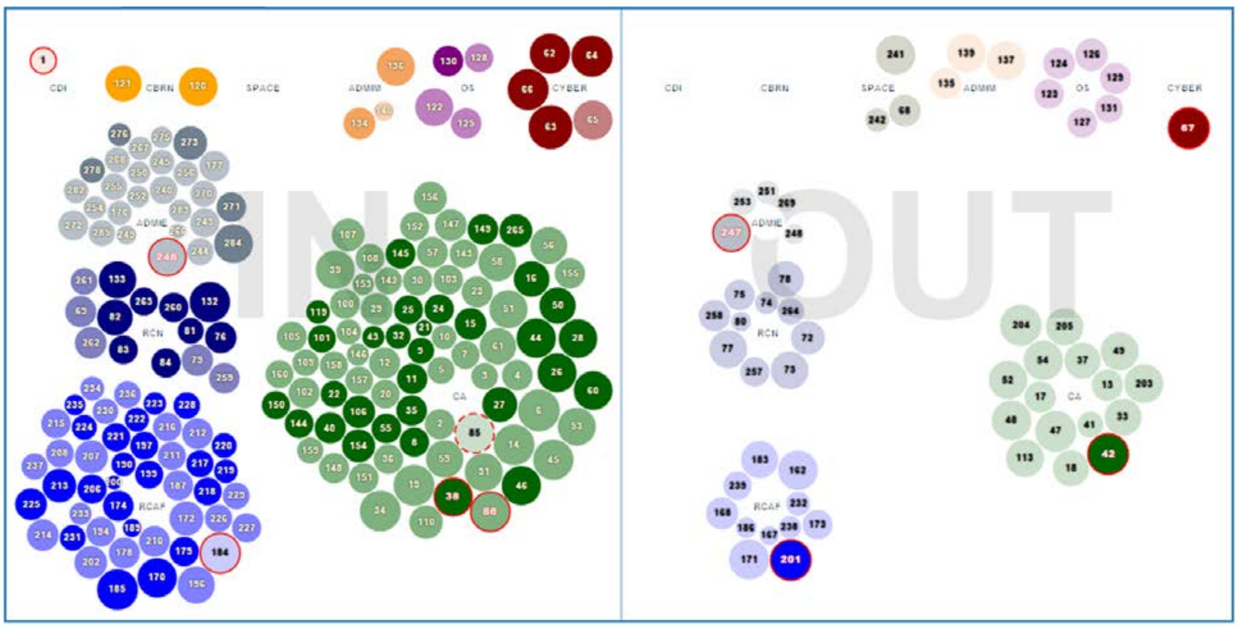

Figure 2 The portfolio view clearly identifies projects that are included in the portfolio (left) as well as those that are excluded (right). Project affiliation with a sponsoring organization is identifiable by the colour coding and clustering; the luminance encodes the funding source. Projects that have been forced into or out of the portfolio have a red border (solid or dashed).

Each project is encoded as a bubble (i.e., a circle). The position of a project encodes whether it is included (left) or excluded (right) from the portfolio. Projects are further clustered by their sponsoring organization. The location of each sponsor cluster can be adjusted via drag-and-drop of the cluster label. In this way, clusters can be compared to one another and/or arranged in accordance to decision maker preferences. In the example shown in Figure 2 the colour of each bubble encodes the project's sponsoring organization. For the most part these colours coincide with those which are most frequently associated with the sponsoring organizations (i.e., Army is green, Air Force is blue, etc.). The area of the bubble encodes a selectable project attribute, such as total cost or value. The luminance of each bubble encodes the funding source for each project. For projects that are included in the portfolio, a brighter circle indicates that the project is to be funded from a cash-based source; a darker circle indicates that an accrual-based funding source is to be used. All projects on the right of the plot are not included in the portfolio, and are encoded with a very low luminance and a black font to indicate that no funding source is assigned to them. The exception to this is where particular projects have been forced out of the portfolio by the user, these projects retain reference to their predetermined funding source (if they have one).

The composition of the portfolio is adjusted via drag-and-drop. In this way individual projects can be forced into or out of the portfolio. Projects are automatically attracted to the sponsor cluster to which it belongs. A project that has been relocated by the user is demarked with a red border. A solid red border indicates: (1) the project has been forced into the portfolio and the decision maker has determined the funding type (user constraint type one), or (2) the project has been forced out of the portfolio (user constraint type three). A dashed red border indicates: (1) the project has been moved into the portfolio, but the funding type has not been selected (user constraint type two); or (2) a project's funding source has been determined by the decision maker but an explicit decision to include or exclude the project has not (user constraint type four). After a user adjusts the composition of a portfolio, it can be reoptimized, characterized and adjusted in an iterative fashion.

As noted above, the portfolio view can be used to perform several tasks. For example, to enable identification of aggregate portfolio characteristics, it is relatively easy to see from Figure 2 that when bubble size encodes project cost and when projects are clustered by project sponsor: (1) the CDI organization (upper left) is an outlier due to the fact that it has only one project, and (2) for Army sponsored projects (green bubbles in the cluster labelled CA), there appears to be a positive correlation between project cost and exclusion from the portfolio. To enable comparison of portfolio characteristics, the visualization illustrates that the Army has the largest number of projects in the portfolio, the Air Force (blue bubbles in cluster labelled RCAF) has the second largest, etc. To enable portfolio adjustment, the visualization shows six projects that have been added to the portfolio by the user (i.e., bubbles on the left with a red border) and four projects that have been forced out of the portfolio (i.e., bubbles on the right with a red border).

Portfolio content and value comparators

Two visualizations in VIPOR enable a user to examine the differences between portfolios generated during the iterative development process. The first of these is the portfolio content comparator shown in Figure 3 (a). It presents multiple portfolios in sequence and indicates which projects are included or excluded from each portfolio. The user can explore the project inclusion patterns within a portfolio, and then compare these patterns across multiple portfolios. In the portfolio content comparator, each project is encoded as a square. Each column of squares encodes the results of one iteration in the portfolio development process. Column labels correspond to the iteration number (i.e., iteration 1, 2, 3, etc.). Green squares denote included projects and black squares denote excluded projects. Squares with a red border identify those projects that have been manually included/excluded from the portfolio by the user. The decision maker can order the projects by one of several attributes, including project value and project cost.

In the example shown in Figure 3 (a) the rows are sorted first by project inclusion/exclusion and then by identification number. This example illustrates that, after the decision maker manually added Project 1 to the portfolio during the second iteration, several projects were subsequently pushed out of the portfolio (e.g., 10, 48, 79) and others were brought into the portfolio (e.g., 131, 171, 239) by the optimization model. With this visualization, it is possible for a decision maker to begin assessing the opportunity cost of including particular projects in the portfolio.

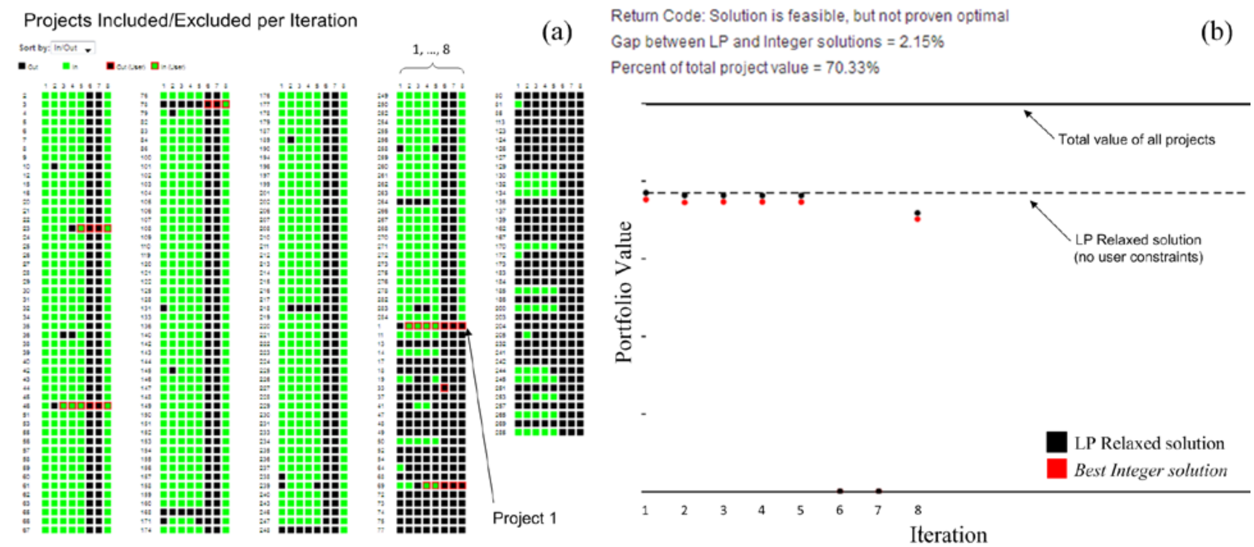

Figure 3 Two VIPOR visualizations depicting a sequence of eight portfolios. Each portfolio corresponds to an iteration in the development process. The addition of user-added constraints and precluded a feasible solution in iterations six and seven. (a) Portfolio *content* comparator displaying project inclusion (green squares) and exclusion (black squares). Squares with a red outline indicate user added constraints. (b) Portfolio *value* comparator displaying the portfolio's value as a function of iteration.

The second visualization that enables portfolio comparison is shown in Figure 3 (b). This portfolio value comparator presents the aggregate value of each portfolio. Both the best integer solution and LP relaxed solution for the iterations are shown. The horizontal axis encodes the iteration number, and the vertical axis encodes the aggregate value of the portfolio. The solid black line across the top represents the total aggregate value of all projects regardless of their inclusion in a particular portfolio. The dashed line shows the LP relaxed solution for the case where no user constraints have been specified. As particular projects are included and excluded, and the portfolio is re-optimized around these preferences, the maximum value that can be achieved will change. Black dots show the LP relaxed solution and red dots show the best integer solution. By combining these elements onto the same plot, a user can explore trends in the aggregate value of the portfolio, and then compare the differences across several iterations.

The portfolio value comparator also demarks cases where a decision maker has attempted to include a collection of projects that, in aggregate, do not comply with constraints. In cases such as these, the aggregate value of the portfolio is set to zero. For example, Figure 3 (b) illustrates that in the 6th and 7th iterations a feasible solution that would conform to decision maker preferences could not be found by the optimization algorithm. On the other hand, the portfolios generated during iterations 1, 2, 3, 4, 5 and 8 produce portfolios where the aggregate value of the best integer solution is within approximately 2% of the solution obtained via LP relaxation. On the 8th iteration however, the cumulative result of forcing several projects into and out of the portfolio has dropped the maximum value achievable to a level that is about 8% lower than what is achievable when the optimal portfolio is computed without accounting for user preferences.

Sensitivity plot

Prior to creating an initial portfolio, a set of 67 optimal portfolios is computed, each corresponding to a different point in the space of possible criteria weights. No optimization time limit or user preferences are imposed. This set of portfolios helps determine the sensitivity of each project's inclusion/exclusion to the criteria weights. Figure 4 is a visualization of the sensitivity analysis for one project. Each dot in the triangle represents one of the 67 portfolios. A green dot indicates that the project is included in the optimal portfolio and a black dot represents that the project is not included. At each location in the triangle, the shading of the background shows the project's overall value as a function of the criteria weights. Iso-value contours of project value are also shown.

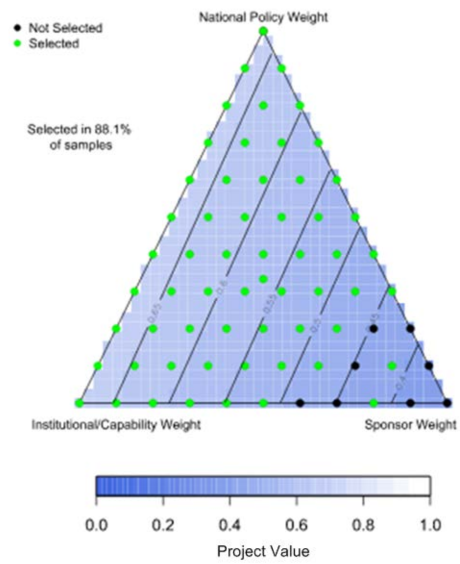

Figure 4 Sensitivity analysis of a project's inclusion/exclusion in a portfolio based on criteria weights.

In this example a decision maker could ascertain that this project is selected in 59 of the 67 cases as weights vary. Points where this project is excluded from the portfolio coincide with cases where the importance of national policy criteria and capability criteria are much lower than the importance assigned to the sponsor criteria. The plot also indicates that when the sponsor criteria weight is high, the project's total value is very low. Thus, a decision maker can gauge the impact of varying weights on the inclusion of this project in the portfolio.

DISCUSSION

A small group of analysts produced VIPOR using an agile development approach. The development was executed in consultation with a cross-functional investment planning working group and several executive decision makers. At the first opportunity, VIPOR was employed within a series of decision maker engagements in order to inform investment planning remits. During the first engagements initial acceptance and familiarity with VIPOR was achieved. These initial engagements were followed by efforts to ascertain appropriate weights which would accurately reflect the importance of each evaluation criteria. In addition, assessments were conducted to ascertain the impacts (e.g., opportunity costs) of including and excluding individual projects that were of particular interest. Finally, several different portfolios were developed to explore the effect of alternate strategic themes.

At the time of writing, efforts to support investment decisions using VIPOR were still underway. No formal assessment of VIPOR's usage and utility has been conducted. However, new efforts to enhance VIPOR are set to begin based on the positive reception it has received at DND. In the remainder of this section, we describe several observations, insights and areas for potential future improvement. These have evolved from our experience developing VIPOR as a decision support system to aid an organization that is just beginning to use PPM concepts for investment planning at the enterprise level.

Executive level support

Gaining and sustaining executive level support of PPM can be a challenge for organizations that are new to this paradigm (pmsolutions; 2013). When faced with this challenge, the VIPOR development team achieved a large degree of success in the lead up to development by using DND's existing data, rather than fictional data, to demonstrate the differences between the incumbent approach and the portfolio approach. Using the existing project data and available budget data, the team computed the cumulative value of projects selected using the ranked-list approach and then compared this against the same value computed based on the portfolio approach. Differences in the value achieved using the two approaches were compared using a graph similar to that shown in Figure 5. For this initial comparison, we used a simple knapsack model where the objective was to maximize overall value while conforming to just a single budget constraint. Without being distracted by a fictional scenario, executives were able to clearly see from this simple analysis that the portfolio approach generates a greater value than does the ranked-list approach for the same budget. Together, with arguments synthesized from the literature about the benefits of PPM and anecdotal evidence about the relative organizational inefficiency of current approaches, this simple analysis was extremely useful for purposes of gaining and maintaining executive level support for VIPOR's development.

Aiding, not replacing human decision making

The intent of VIPOR is to support effective decision making. Like other operational research and business analysis tools, VIPOR was introduced to executive decision makers as a mechanism to provide a better starting-point for developing a project portfolio, and not a "black box" mechanism to determine the final answer. Our experience has been that stressing this aspect is often a key element to gain and maintain stakeholder interest when offering a new tool that amalgamates evidence and enhances insight for purposes of executive decision making. However, VIPOR goes beyond just providing a starting-point. During the initial engagements with decision makers using VIPOR, it was found that features which allow projects to be forced into the portfolio, as well as features which facilitate the iterative development of different portfolios, were much more important than originally envisioned. Decision makers quickly recognized the benefit of merging the power of an optimization model with their own intrinsic knowledge in order to account for realities that are not directly reflected in the value model (e.g., new government investment decisions, new strategic or policy preferences). Additionally, by observing how and why different portfolio options were built, the design team collected ideas about how both the value model and the optimization model could be enhanced.

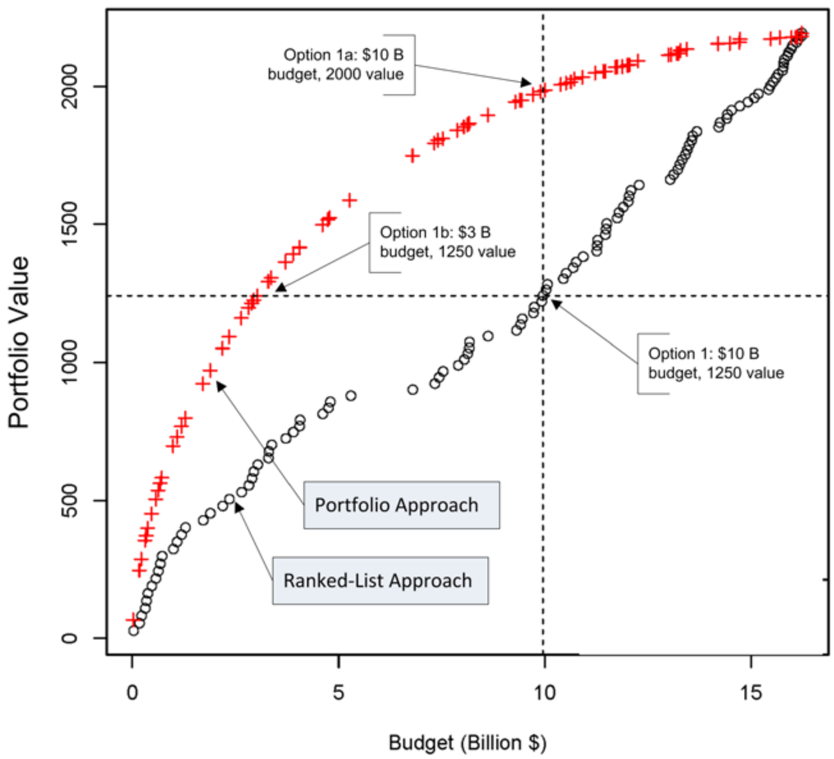

Figure 5 An illustration of the difference between portfolios produced by the ranked-list approach (black circles) and the portfolio approach (red crosses). Each marker is a different portfolio.

The power of visualization

Constructing a preferred portfolio is a complicated endeavor – especially in a large organization where there are many potential projects, several overarching and sometimes conflicting objectives to consider, and a multitude of constraints and conditions to balance. However even though the problem is complicated, care must be taken to ensure that the methods used to solve it are not presented in a way that appears too complex and difficult for decision makers to understand or use. It has long been recognized (e.g., Jackson, 1983) that optimization methods are the most fundamental techniques for determining a project portfolio; but as indicated by Archer and Ghasemzadeh (1999), the techniques and measures for project selection will *not* be used if they *cannot* be understood readily by managerial decision makers.

During the development of VIPOR the design team decided that the complicated nature of constructing an effective project portfolio should not lead to the abandonment of techniques that are most apt to help solve the problem. This led to a focus on intuitive visualizations and an interactive interface. The first aim of incorporating visualization was to ensure that decision makers were presented with information most relevant for evaluating and adjusting the balance in a portfolio – without being overloaded with unnecessary details. Meeting this aim led to the development of the tasks identified in Table 2. The second aim of incorporating visualization was to allow a decision maker to "drill-down" into the data for purposes of verifying a project, constraint and other data fed into the optimization model. The third aim of visualization was to help decision makers develop confidence in the application and an understanding of the portfolio construction paradigm. For purposes of brevity, the visualizations and interactive features to meet this second aim are not presented herein.

To meet the first aim, i.e., *balance*, observations of decision maker interactions with the portfolio view, portfolio content comparator, portfolio value comparator and sensitivity plot (shown in Figures 2 through 4 respectively) were particularly beneficial. For example, the portfolio view was not only useful for both characterizing and adjusting portfolios, it provided a dynamic image that: intrigued many executives and helped anchor the VIPOR application in their memory. It also was particularly valuable for explaining the portfolio optimization process and facilitating acceptance of the optimization concept. The portfolio view has a simple design, but among the visualizations in VIPOR, it also has the greatest degree of aesthetic novelty. Interestingly however, we believe this is not the most likely reason that it seemed to attract the most attention. We propose that its special appeal is because, in combination with resource utilization charts showing the temporal demand placed by the preferred portfolio against the supplies of fiscal and project delivery capacity, this plot provides an aggregate overview of the entire portfolio construction problem. It encapsulates the central idea that portfolio construction is not only about achieving the best *balance* between supply and demand, but it is also about achieving *balance* along a series of secondary, but nonetheless important, dimensions (e.g., project sponsorship, project size, project horizon). Currently, it is largely up to the decision maker to build an appropriate balance along these secondary dimensions by using the interaction and visualization mechanisms to force projects into the portfolio and synthesize the results.

In addition to balance, the other main idea behind project portfolio construction, i.e., *alignment*, is best illustrated by the portfolio content comparator. By ordering projects from the highest to the lowest value, the portfolio content comparator provides visual cues as to how the maximization of portfolio value in accordance with fiscal and other constraints is not necessarily achieved by selecting all of the highest valued projects. While intuitive to some, the explicit illustration of this concept led to more effective planning discussions. In these sessions, emphasis was placed on finding the right combination of projects to meet projected demands across a multi-year time horizon, rather than on determining the specific ranking of particular projects within a prioritized list.

The portfolio content comparator was also instrumental for purposes of gaining and maintaining the interest of senior executives. During the final stages of the development process, VIPOR was used to help characterize the opportunity cost associated with incorporating a new capital project that had only just been identified. To conduct the analysis data for the new project was created (e.g., costs, timelines, resource requirements). The project was then manually forced into the portfolio and an analysis of those projects that were subsequently pushed out of the portfolio was conducted. This analysis led to a use-case for VIPOR, which had been given little attention during the initial development, but later resulted in greater levels of acceptance and understanding of the model. In addition to improving a decision maker’s ability to create portfolios through an iterative process of optimization and interaction, VIPOR could also be used to build an entire portfolio from scratch through the manual inclusion/exclusion of projects. During this manual process, the various visualizations could be used to ensure financial and resource related constraints would not be violated.

To meet the third aim, i.e., *understanding*, it is of interest to note that integration of visualization and interaction features alongside data analysis and advanced modelling algorithms is the hallmark of a relatively new scientific domain called visual analytics (Thomas and Cook, 2005). Indeed, the integration of these two aspects has resulted in acceptance of the portfolio construction approach that VIPOR was built to support. However, not only has the incorporation of advanced visualizations and interaction mechanisms been important for finding solutions to the portfolio construction problem, it has also proved to be useful for helping users develop a better understanding of the problem, and in some cases, it has prompted changes to both how the problem is viewed and which fundamental questions are posed. As an example of the latter, an unintended and yet beneficial consequence of VIPOR was that it led to an executive level request to incorporate features for automatically determining the impacts of financial re-profiling. The request resulted from a visual characterization of the available accrual envelope and the potential benefits of mitigating so-called “bottlenecks” that were seen to put undue restriction on project delivery in certain fiscal years while at the same time leaving financial resources unused in other fiscal years. The request was addressed by altering the optimization model so that, in addition to determining which projects to include in the portfolio, it would also automatically adjust the available financial envelopes according to a set of underlying rules for carrying forward unused funds. These rules were translated into mathematical constraints. In order to demonstrate the impacts of adjusting the profiles of the available financial envelopes, alternative portfolios were created with the auto-profiling feature turned on and off. Expanding portfolio optimization models in this way is not uncommon (Brown et al., 2004), but the realization that this request was initiated as a result of using VIPOR’s interactive visual interface was a signal to the design team of its initial effectiveness.

Sand, small rocks and big boulders

During initial usage of VIPOR, it became apparent that the portfolios it automatically produced usually did not include many of the high cost projects, even though these high cost projects were deemed important. After working with the portfolio view, one senior executive coined what became known as the “sand, small rocks and big boulders” metaphor. On its own, VIPOR creates portfolios with a lot of “sand” and “small rocks”, and only after interacting with the model can “big boulders” be forced into the portfolio.

The apparent bias toward the inclusion of lower cost items is not unique to optimization models of this type. Among others, Greiner et al. (2003) and Parnell et al. (2002) have made similar observations. When VIPOR was first introduced, this effect appeared to leave decision makers with a sense that the model was not computing a correct result. This was not the case. The main reason for this effect is that a linear measurable value function used to convert scores into value. Alternative measurable value functions can also be implemented (e.g., exponential value functions (Kirkwood, 1997)). The effect can be countered in other ways. For example, constraints can be added to the model to limit the total number of projects in the final portfolio. Adaptations like these are straightforward, and they could have been included in VIPOR. However, in consultation with executives, it was decided that the sand, small rocks and big boulders effect actually benefitted initial engagements with senior decision makers. It forced special consideration of both the highest cost projects and the opportunity costs associated with including them in the portfolio.

The deliberate choice *not* to adjust for this effect is consistent with an underlying concept from visual analytics that tools like VIPOR can help a decision maker steer towards a final solution in ways that may not be otherwise achieved by automatic methods alone (Keim et al., 2009).

Further consideration of the sand, small rocks and big boulders effect is planned for future versions of VIPOR.

Limitations and selected areas for future development

As the development of VIPOR began, a decision was made to include only the most fundamental aspects of the portfolio construction problem in the optimization model. The idea was to motivate the inclusion of more sophisticated features through the initial use of VIPOR by decision makers. In part, this decision was taken based on the idea that at DND previous attempts to implement a multi-objective optimization model for balancing investment decisions in connection with CBP had been abandoned (Taylor et al., 2008).

In addition to its design, the scope of VIPOR's initial usage was aligned to help ease the transition to the portfolio approach. In keeping with the idea that the approach was to be applied at the enterprise level, all capital investment projects, regardless of the project type or sponsor, were included in the analysis. However, VIPOR was only used to aid in the selection of projects not yet allocated capital funding with a total cost greater than a specified threshold. Projects already spending money towards delivery were not part of the assessment. In addition, only costs associated with acquisition were incorporated; full lifecycle costs were not considered. These steps to slowly introduce the portfolio approach are in line with observations in a recent PPM survey (pmsolutions; 2013) which indicates that PPM represents the type of cultural change that takes time and executive support to institutionalize – it can take more than 5 years to achieve higher levels of maturity with the approach.

In keeping with the idea that maturity is achieved through steady improvement, the addition of more sophistication is envisioned for the next versions of VIPOR. Potential enhancements include:

- Allowing the optimization model to schedule projects within appropriately specified time bands;
- Using an nonlinear measurable value functions to better model the relative difference between the highest valued projects and the lowest valued projects in criteria;
- Enabling the model to select from alternate versions of the same project, where each version has a different cost and value;
- Incorporating synergistic interactions where the cumulative value of two or more related projects becomes greater than the sum of the individual projects;
- Incorporating persistence constraints to explicitly reduce fluctuations between a starting portfolio and a newly optimized portfolio;
- Adding constraints to explicitly enforce preferences related to a portfolio's overall balance across one or more dimensions (e.g., strategic theme, project sponsor dimension);
- Incorporating uncertainty bands around point estimates for both project related data as well as financial and resource constraints;
- Explicit incorporation of risk concepts at both the project and enterprise levels; and
- Expanded usage to include projects already underway and their full lifecycle costs.

Of course, all of these potential improvements depend on other aspects affecting the transition to the portfolio approach, including cultural, process, and the availability of necessary data.

CONCLUSION

Ghasemzadeh and Archer (2000) outline several factors that can improve the potential for successful implementation of a decision support system for constructing preferred project portfolios. Although it is has been 15 years since their publication, and the domain of PPM has matured much since then, we have found through the development of VIPOR that their recommendations remain especially useful. Both the development and use of VIPOR has benefitted tremendously from a committed senior executive, careful consideration of system and information requirements, and an evolutionary development approach by a committed team. At the time of writing VIPOR is currently fulfilling its aims to establish a more explicit link with current CBP practice and enhance the efficacy of investment planning processes.

Among its many features, two aspects of VIPOR were of particular importance to its success.

- The first aspect was the adoption of an iterative portfolio construction process facilitated by visualization. Not only did this feature capture the attention of decision makers, it enabled VIPOR to take advantage of an automatic optimization algorithm and combine this with human background knowledge and intuition.
- The second aspect that contributed to the success of VIPOR was the explicit consideration of the organization's state of familiarity with PPM concepts, and its cultural legacy and current decision making processes. The optimization model developed for the first implementation of VIPOR included only those features that were deemed most likely to help DND begin its transition to the PPM paradigm at the enterprise level. Not only did this simplify the design, reduce the data collection burden, and shorten computation times, it also helped executives to rapidly gain an appreciation of the model and develop trust in its output. As the level of comfort with PPM increases within the organization, maintaining an appreciation of the concepts is expected to facilitate future efforts in at least four areas: (1) enhancement of PPM decision making processes, (2) improvement of PPM information collection processes, (3) advancing the sophistication of the optimization model, and (4) enriching visualization aspects.

REFERENCES

Archer, N. & Ghasemzadeh, F. (1999). An integrated framework for project portfolio selection. *International Journal of Project Management*, 17(4), 207-216.

Auditor General of Canada (2012). *Spring Report of the Auditor General of Canada to the House of Commons*.

Brehmer, M. & Munzner, T. (2013). A multi-level typology of abstract visualization tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12), 2376-2385.

Brown, G., Dell, R., Holtz, H., & Newman, A. (2003). How US Air Force Space Command Optimizes Long-Term Investment in Space Systems. *Interfaces*, 33(4), 1-14.

Brown, G., Dell, F., & Newman, A. (2004). Optimizing Military Capital Planning. *Interfaces*, 34(6), 415-425.

Burk, R. & Parnell, G. (2011). *Portfolio Decision Analysis: Lessons from Military Applications*, in *Portfolio Decision Analysis: Improved Methods for Resource Allocation*. New York:Springer, 333-357.

Davie, M. (2013). *Comparative Defence Planning: Lessons from New Zealand*. Master's thesis.

Davis, P. K. (2002). Analytic Architecture for Capabilities-Based Planning, Mission-System Analysis, and Transformation. *RAND*.

Donahue, S. (1992). *An Optimization Model for Army Planning and Programming*. Master's Thesis. Naval Postgraduate School.

Dyer, J. & Sarin, R. (1979). Measurable multiattribute value functions. *Operations Research*, 27(4), 810-822.

Ewing, P., Tarantino, W., & Parnell, G. (2006). Use of Decision Analysis in the Army Base Realignment and Closure (BRAC) 2005 Military Value Analysis. *Decision Analysis*, 3(1), 33-49.

Fourer, R., Gay, D., & Kernighan, B. (1990). A modelling language for mathematical programming, *Management Science*, 36(5), 519-554.

Gass, S. (1975). *Linear Programming*, McGraw-Hill.

Ghasemzadeh, F. & Archer, N. (2000). Portfolio selection through decision support. *Decision Support Systems*, 29(1), 73-88.

Government of Canada (2008). *Canada First Defence Strategy*. Retrieved from www.forces.gc.ca/en/about/canada-first-defence-strategy-summary.page, April 15, 2015.

Government of Australia (2013). *Defence White Paper 2013*.

Greiner, M., Fowler, J., Shunk, D., Carlyle, W. & McNutt, R. (2003). A hybrid approach using the analytic hierarchy process and integer programming to screen weapon systems projects. *IEEE Transactions on Engineering Management*, 50(2), 192-203.

Jackson, B. (1983). Decision methods for selecting a portfolio of R&D projects. *Research Management*, September-October, 21-26.

Hamill, J., Deckro, R., Kloeber Jr., J., & Kelso, T. (2002). Risk Management and the Value of Information in a Defence Computer System. *Military Operations Research*, 7(2), 61-81.

Hoehl, M-L. (2008). Force Development: The Linear Programming Approach. *NATO RTO SAS-072 Symposium*.

Keim, D., Andrienko, G., Fekete, J., Gög, C., Kohlhammer, J. & Melançon, G. (2008). Visual Analytics: Definition, Process, and Challenges. *Information Visualization, Lecture Notes in Computer Science*, 4950, 154-175.

Keim, D., Mansmann, F., & Thomas, J. (2009). Visual Analytics: How much visualization and how much analytics? *Sigkdd Explorations*, 11(2), 5-8.

Kirkwood, C. (1997). *Strategic Decision Making*, Duxbury Press.

Microsoft Corporation. (2008). *Product Portfolio Management: Doing the Right Things Right*. Retrieved from www.epmconnect.com, April 15, 2015.

Parnell, G., Bennett, G., Engelbrecht, J., and Szafranski, R. (2002). Improving Resource Allocation within the National Reconnaissance Office. *Interfaces*, 32(3), 77-90.

pmsolutions. (2013). *The State of Project Portfolio Management: A pmsolutions research report*. Retrieved from www.pmsolutions.com, April 15, 2015.

Project Management Institute (2013a). The Standard for Portfolio Management.

Project Management Institute (2013b). The View from Above: The Power of Portfolio Management.

Rempel, M. (2010). Advancing the Capability Based Planning Process for Strategic Decision Making within the Canadian Forces. *Proceedings of the 41st Decisions Sciences Annual Meeting*, 1441-1452.

Rempel, M. & Young, C. (2014). The Portfolio Creation Model developed for the Capital Investment Program Plan Review (CIPPR). *Defence Research and Development Canada,* DRDC-RDDC-2014-L256.

Salmeron, J., Dell, R., Brown, G., & Rowe, A. (2002). Capital Investment Planning Aid (CIPA) - An Optimization-Based Decision-Support Tool to Plan Procurement and Retirement of Naval Platforms. *Naval Postgraduate School*, NPS-OR-02-006.

Solomon, B. & Stone, C. (2013). Accrual budgeting and defence funding: Theory and simulations. *Defence and Peace Economics*, 24(3), 211-227.

Taylor, B., Petryk, R., Blakeney D., Christopher, G., & Beard, A. (2008). Producing an Integrated Capability Roadmap for the Canadian Forces. *NATO RTO SAS-072 Symposium*.

Taylor, B. (2014). Analysis Support to Strategic Planning v2. *The Technical Cooperation Panel Technical Report,* TR-1-JSA-2015.

Thomas, J, & Cook, K. (2005). *Illuminating the Path: Research and Development Agenda for Visual Analytics*. IEEE-Press.

Walmsley, N. & Hearn, P. (2004). Balance of investment in armoured combat support vehicles: An application of mixed integer programming. *Journal of the Operational Research Society*, 55, 403-412.

Weltman, P. (2015). Fiscal Sustainability of Canada's National Defence Program. *Office of the Parliamentary Budget Officer.*

World Bank. (2014). *Military Expenditures*. Retrieved from www.data.worldbank.org/indicator, April 15, 2015.

Young, C. & Rempel, M. (2014). The Portfolio Approach developed to underpin the Capital Investment Program Plan Review (CIPPR). *Defence Research and Development Canada,* DRDC-RDDC-2014-L255.