RC24102 (W0611-029) November 7, 2006
Computer Science

IBM Research Report

Resource Analysis in Project Portfolio Management

Vijay Iyengar, Joseph Kramer, Daniel Milch, Ignacio Terrizzano, John Vergo
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598

Research Division
Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

Resource Analyses in Project Portfolio Management

Vijay Iyengar, Joseph Kramer, Daniel Milch, Ignacio Terrizzano, and John Vergo

***Abstract*—Project portfolio management is an ongoing task in enterprises and systematic methodologies are being developed for this important task to replace ad hoc approaches of the past. Management of project portfolios differs significantly from the same task applied to financial portfolios. Areas of research like real options seek to exploit the similarities in these two domains and address the differences. One of the unique dimensions in project portfolio management is the role of resources (human). Projects are typically proposed with significant uncertainty and lack of specificity in their resource needs. The evaluation of project portfolio options needs to factor in this uncertainty and also handle the heterogeneity in terms of resource types. Evaluation of the resource impact also needs to consider the tradeoff between maximally utilizing the available resources and having cohesive and compact teams to perform the linked (or related) tasks in a project. In this paper we develop models and methods for analyses of the resource dimension in project portfolio management that address its unique characteristics. We then illustrate our methods with experimental results on synthetic examples with varying characteristics.**

I. INTRODUCTION

THERE has been extensive work done in the area of analyzing and evaluating project portfolio options especially in the context of information technology (IT) investments [1, 2]. McFarlan suggested using a value and risk-based approach to manage IT investments as a portfolio [3]. While parallels between financial portfolios and IT portfolios are made frequently, there are important differences that need to be taken into account [4, 5]. The work in the area of Real Options is one approach to address the impact of these differences [4, 5, and 6]. The granularity of IT investments (specifically projects) can be a significant factor in their management. IT projects are also not as tradable or liquid as their counterparts in the financial world [4] and their temporal characteristics add another dimension of complexity to the problem. In addition, projects can be inter-related. Project dependencies are one form of relationships that act as constraints in the portfolio management task [7, 8]. In summary, there are multiple dimensions to be considered during project portfolio management and in this paper we will focus on the role of resources (human) in the analyses of portfolio options. Next, we will discuss the important characteristics of the resources dimension in project portfolio management.

The area of resource allocation for projects has been studied extensively and commercial systems exist to solve large scale problems in this space [9]. Typically, these systems take as input the resource needs of projects and the total resource availability as a constraint to schedule the projects and allocate the resources to them. This process will determine any shortfall in resources due to the needs exceeding the availability. During portfolio management many sets of related ongoing and proposed projects are considered and evaluated. This evaluation involves comparing the tradeoffs in costs, resource needs, benefits and risks of those sets of projects. Typically, the estimates of resource needs for proposed projects have significant amount of uncertainty in them (magnitude and temporal aspects). They also lack specificity in terms of project personnel needed since the project task structure (also referred to the work breakdown structure) is not developed fully at the early stages of a project. Treating the available resources as a hard constraint is not appropriate given the significant levels of uncertainty and lack of specificity. Also, there are differences in the characteristics and the handling across resource types. For example, availability of some resource types could be augmented with outsourcing and external contracting arrangements while for other types (e.g., subject matter experts) this may not be a viable solution. These characteristics need to be factored in when considering the projected shortfall in a particular resource type by any set of projects being considered whose needs exceed the availability for that type. The evaluation will not only involve determination of load and shortfall by resource type but may also require consideration of the temporal breakdown (e.g., time periods with significant shortfall). In this paper, we will develop and use models for representing the resource data that are commensurate with the information (specificity, uncertainty) that is available. We will adapt the work done in the area of resource allocation to develop analyses methods that are suited to the characteristics of the resource data available during project portfolio management. We will illustrate our model and methods by applying them to synthetic examples with varying characteristics. In the next section, we discuss the portfolio management framework in some detail to set the context for resource analyses.

II. PROJECT PORTFOLIO MANAGEMENT FRAMEWORK

A typical project portfolio management framework has an evolving list of projects that are in scope for the management task. Portfolio management is an ongoing process and we will consider a snapshot of this process, say, at a particular review

---

Manuscript created November 10, 2006.

The authors are with the IBM Research Division (contact author: V. Iyengar, phone: 914-784-7339, IBM T.J. Watson Research Center, P.O. Box 704, Yorktown Heights, NY 10598, email: vsi@us.ibm.com).

of the portfolio. At this snapshot, we can group the projects into two classes: committed projects, and projects under consideration. During a particular portfolio review, only the projects under consideration are evaluated and decided upon. Clearly, the needs of the committed projects reduce the availability of enterprise resources for the projects under consideration. Committed projects can be further categorized into two types: those which have resources deterministically assigned, and those that do not. The mechanics of factoring in the needs of committed projects will differ based on this state of their resource assignment.

The framework provides project attributes that are relevant to the decision making process. Project schedule information is provided allowing for representation of uncertainty in the estimates of dates and durations. Functional dependencies between projects need to be modeled in the framework. Any portfolio option that is considered in the decision making process needs to be consistent with the project dependencies [8]. For example, if project A depends on project B, then considering a portfolio option that specifies doing project A and not doing project B would be misleading. Other project attributes include costs and business benefits along with risks. One aspect of cost is the resource need, though it is not always adequate to simply translate it to financial costs (e.g., subject matter experts in short supply). The rest of the paper will describe our approach for analyzing the resource impact of project portfolio options. We start, in the next section, with our model for specifying resource information.

III. RESOURCE MODEL

The model for resources can be described in three parts. First, we will develop taxonomy for specifying the various types of resources. Next, we will utilize it to specify the resource availability. And lastly, we will define a model for resource needs.

A. Resource Taxonomy

The resource taxonomy used for representing resource types in the resource availability and needs data is based on a simple tree model illustrated with a partial example in Figure 1. Clearly, this simple model does not have the flexibility to represent multiple skills of a person in the resource pool and factor them in the resource load and shortfall analysis [10]. Our choice of the simpler model is based on matching it to the typical level of specificity in the estimates of resource needs early in a project's lifecycle. Nodes in the resource taxonomy tree will be referred to as "roles" (short for resource roles) in this paper. Leaf nodes in the taxonomy tree will be called "leaf roles".

B. Resource Availability

Resource availability can be viewed as a two dimensional specification. The first dimension represents individuals in the resource pool. The second dimension represents time broken into adjacent intervals at some level of granularity covering the entire future period of time relevant for the portfolio analysis (e.g., the next 2 years). We will represent this two dimensional specification using the notation A(p, t), which denotes the availability of person p during the time interval t. We also assume specification of the maximum availability possible, Amax, for a person in a time interval. For example, if the time intervals represent work weeks and the maximum availability per week is 40 hours, then A (p, t) can be a quantity from 0 to 40 hours. In addition, the availability model specifies the leaf role for each person, R(p). As we mentioned earlier, we only allow a single leaf role for a person in our model.

The resources needed by committed projects have to be factored out of the available resource pool. For committed projects that have been assigned resources, we compute the effective availability A after subtracting out for each person the time assigned to the committed projects. However, the resources associated with committed projects that have not yet been assigned actual resources need to handled differently as discussed in Section IV.

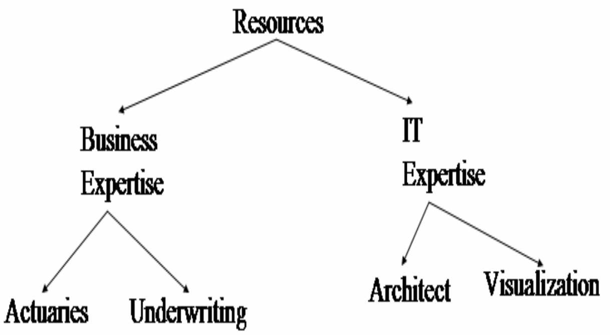

Figure 1. Example of Resource Taxonomy (partial)

C. Resource Needs

The model for specifying the resource needs of each project is discussed next. We will first describe the format of the specification as a set of deterministic quantities and then add the uncertainty model.

For each project q, the resource needs are specified as a list NEEDS(q) of linked needs, L(i). For each linked need L(i), the needed role R(L(i)) (from the resource taxonomy) is specified. Note, the needed role R(L(i)) can be a non-leaf role in the taxonomy implying it can be satisfied by any person from the resource pool with a leaf role in the sub-tree rooted at R(L(i)). Each linked need L(i) quantifies the resource need as a sequence of block needs, denoted by [B(i,j)], where j ranges from 1 to the number of blocked needs in L(i). A block need B(i,j) specifies the need for M(i,j) amount of resources in the time interval from Ts(i,j) to Te(i,j), using the three parameters M, Ts, and Te. The time intervals in a linked need are disjoint and the blocks are sequenced in sorted time order. The temporal specification (Ts and Te) and the magnitude specification (M) use the same granularities that were used in the resource availability model. For example, a linked need for an 'Actuaries' role can specify needing 100 hours during the first 2 weeks in February 2007 and 200 hours during the next 2 weeks (see Table 1). The example in Table 1 also specifies a second linked need of 500 hours for a different role (Underwriting) during the period February 1, 2007 to March 31, 2007.

Table 1. Example of Linked Needs

Linked Need (Role)	Block Needs	
L(1) (Actuaries)	B(1,1): 100 hrs From 2/1/07 to 2/14/07	B(1,2): 200 hrs From 2/15/07 to 2/28/07
L(2) (Underwriting)	B(2,1): 500 hrs From 2/1/07 to 3/31/07	

This model of resource needs is flexible enough to represent the project structures typical at the early stages of a project and to represent more refined project structure determined after project planning. Our model for resource needs has less specificity than models used earlier [11] that specify needs in terms of each person on the project and their fractional utilization (e.g., need 2 full time persons for the first two weeks in February 2007). Again, our coarser model was motivated by what project managers are able to specify (for resource needs) before the detailed task structure for the project is developed. However, each linked need in our formulation is meant to capture a sequence of needs for which the project manager would like continuity in terms of the people assigned to them. Indeed, for a given sequence of block needs B(i, j) corresponding to a linked need L(i), we can readily determine the minimum number of persons, Pmin(i), that would be needed to satisfy the entire need L(i). This minimum number is calculated assuming no other projects are assigned to these persons. Any assignment for the linked need L(i) that involves more than Pmin persons from the resource pool represents some reduction in continuity or "fracturing" beyond the ideal case that is possible with infinite resources. We use this notion of continuity as a more relaxed form of the resource needs specification when compared to earlier models that explicitly specify needs in terms of persons and their fractional utilizations. This will be discussed in detail in the later sections that describe our analyses.

The uncertainty in the resource needs will be modeled by treating the parameters of above model as random variables. We will be using Monte Carlo simulation to evaluate the impact of these uncertainties [12]. Monte Carlo methods are widely used to understand the impact of schedule uncertainties in project networks. Simulation based analysis has the flexibility to use complex probabilistic models for these random variables (e.g., one could choose to specify joint distributions for related parameters). Uncertainties in the magnitudes of resource estimates (M) and in the durations (Te-Ts) are particularly useful to model and simulate.

The approach used to analyze the resource impact due to a project portfolio option O is as follows. For each scenario in the Monte Carlo simulation, we sample parameters in the resource needs model from their specified uncertainty distributions. This will be done for projects in the option O and for projects that have been committed but not assigned. For each scenario, we assign the available resources to satisfy the sampled needs for committed but un-assigned projects and for projects in the option O. The resource needs not met by the assignment of available resources constitute the resource shortfall for this scenario. By analyzing a number of scenarios we can empirically determine the distribution of resource shortfall for the option O.

IV. ANALYSIS WITHOUT CONTINUITY CONSTRAINTS

In this section we will develop simplified analysis utilizing a resource assignment that does not factor in the requirements on continuity in resource needs. This simplified assignment will minimize a linear aggregate shortfall function without any consideration for the fracturing (i.e., lack of continuity) occurring in the assignment. One motivation for discussing this simplified analysis is that the resulting shortfall can be viewed as a lower bound for the analysis in the next section that factors in continuity constraints.

The simplified assignment problem is formulated as follows. Consider resource taxonomy X, with a set of roles Y and a set of leaf roles Z. We will denote the set of projects in the option O as the set Q1 and the set of unassigned committed projects as the set Q2. For each project q, in the set $\{Q1 \cup Q2\}$, the linked resource needs, NEEDS(q), is given. The resource availability A(p,t) is also given for each person, p, in the resource pool and for all the relevant time intervals, t. We will represent the assignment of resources to the projects by the function G(p, t, q, i) which specifies the amount (non-negative) of availability of person p in time interval t that is assigned to satisfy a linked need L(i) of project q. Note that not all combinations of the 4-tuple (p, t, q, i) are meaningful to consider in function G. We only need to consider those combinations that represent a possible legal contribution by a person p at time interval t to the needs L(i) of project q. For example, if the resource role of the person p is not consistent with the needed role R(L(i)) for linked need L(i) of project q then any 4-tuple containing p, q and i can be removed from consideration. Even when the roles are consistent, a 4-tuple (p, t, q, i) can be removed from consideration if the linked need L(i) of project q does not specify any blocked need that encompasses the time interval t. Formally, we can define a notion of contributors to associate a blocked need B(i,j) in the sequence L(i) for a project q to a person p in the resource pool at time interval t. A person p at time interval t can be a contributor to B(i,j) iff the resource role of p is consistent with the needed role for B and if the time interval t is contained within the start (Ts(i,j))and end times (Te(i,j)) for block B(i,j). For this analysis, we will consider the set V of 4-tuples (p, t, q, i) formed by combining for each project $q \in \{Q1 \cup Q2\}$ and for each of its linked needs L(i) all the combinations of persons p and time t that are contributors to one of the blocked needs B(i,j) of L(i).

Since we are ignoring the continuity criterion, the simplified assignment can be formulated as a linear programming problem [13] as described below. An objective function that could be minimized is the aggregate shortfall as measured by the needs of the projects in $\{Q1 \cup Q2\}$ that are not satisfied by the assignment G to the resource pool. For now, we will treat all the resource roles and projects as equally important. Extending this by adding linear weighting functions that indicate the relative importance of projects and resource roles is straightforward. Equivalently, we are maximizing the needs of the specified projects satisfied by the available resource pool. This can be expressed as maximizing the sum of G(p, t, q, i) for all the 4-tuples (p, t, q, i) in V. The following linear constraints are specified in the linear programming formulation.

The first set of constraints specify that the sum of all the assigned contributions for a person p in a time interval t cannot exceed the availability A(p,t). We will denote these as availability constraints (shown below).

$$\sum_{q,i} G(p,t,q,i) \le A(p,t) \text{ for all } (p,t,q,i) \in V$$

The second set of constraints is derived by considering each blocked need in the set of projects $\{Q1 \cup Q2\}$. The sum total of the assigned contributions (from all the contributors) for a particular block should not exceed its need. Let B(i,j) be a blocked need for the linked need L(i) in project q. Let the set C(i,j) of 2-tuples (p, t) denote the contributors to B(i,j) as defined earlier. The block constraint for B(i,j) can be expressed as

$$\sum_{p,t} G(p,t,q,i) \le M(i,j) \text{ for all } (p,t) \in C(i,j).$$

The resultant objective function after solving the linear program is the maximum aggregate resource need that can be satisfied by the available resource pool. From this we can determine the aggregate shortfall for the Monte Carlo scenario being analyzed. Later in the paper, we will use the shortfall achieved by this LP formulation as a lower bound to gain insights into the results achieved.

V. ANALYSIS WITH CONTINUITY CONSTRAINTS

The motivation for using linked needs for projects is to specify where continuity in personnel is relevant and allow some control over the size of the team from the resource pool that is assigned to satisfy each linked need. For each linked need L(i), we had earlier defined Pmin(i) as the minimum number of people needed to satisfy all the blocked needs in L(i). We can certainly achieve this idealized minimum when we can arbitrarily add persons to the resource pool with maximum availability in all the time periods relevant to L(i). However, in real situations, the resource pool availability may severely constrain satisfying a linked need L(i) with teams whose size is even close to Pmin(i). We need a metric to quantify how far an assignment is from ideal that is meaningful to project managers.

Consider a linked need L(i) whose total need W(i) is computed by summing over all its blocked needs. Consider an assignment that satisfies L(i) using N persons from the resource pool with contributions U(k) for k ranging from 1 to N (sorted in increasing order of their contributions). We define a metric F to measure the fracturing in the team assigned to L(i) as follows:

$$F = \begin{cases} 0, \text{ if } N = P\min(i) \\ \dfrac{1}{W(i)} \displaystyle\sum_{k=1}^{N-P\min(i)} U(k), \text{if } N > P\min(i). \end{cases}$$

The formulation of the metric F above, based on the sorted contributions, measures the fraction of the total linked need that is satisfied by the smallest contributors beyond the minimum number. For example, when the linked needs are completely satisfied by the minimum number of persons the metric F has value 0. A value of 0.1 for F indicates that 90% of the linked need is satisfied by the minimum number of top contributors. Note that in the above example when F is 0.1, there is no information on how many additional people were needed to satisfy the remaining 10%. We have chosen to define the metric F in terms of contribution and not in terms of the number of persons assigned for the following reasons. With typical loads for the resource pool we could get long tails in the contributions of the persons assigned to satisfy the project needs. A person-based metric will be accentuated by such long tails and this may overstate the fracturing when we are analyzing resource need estimates with significant uncertainty.

We can now formulate the assignment problem factoring in the metric F. One way would be to extend the formulation in the earlier section by adding a constraint on the achieved F-metric that specifies an upper bound, Fub, for the metric achieved by the assignment (for all linked needs). In order to compute the F metric we need to completely satisfy the project needs by extending the resource pool by adding additional persons with the needed roles. The resource pool will be composed of two parts: the original resource pool used to determine shortfall that is being minimized and the overflow resource pool used to complete the assignment of the excess causing the shortfall and to compute the F metric that is constrained by the upper bound. The formulation for this assignment problem would include the availability and block constraints developed in the previous section. In addition, we can add auxiliary binary integer variables, a(i,p), which indicate whether or not person p is making a non-zero contribution to the linked need L(i). The constraint on the F-metric can be expressed by a non-linear expression for each linked need as shown below.

$$\sum_{p} \text{Contrib}(i,p)\times a(i,p) \geq W(i)\times(1-\text{Fub}), \text{ for each } L(i)$$

$$\text{where } \sum_{p} a(i,p) = P\min(i)$$

The non-linear constraint above uses a short form notation, Contrib(i,p), to represent the sum of all the contributions to the linked need L(i) by the person p. The constraint specifies that the total contributions from the top Pmin(i) contributors should at least equal the fraction (1-Fub) of the total need W(i) for linked need L(i).

The above formulation belongs to the class of mixed integer nonlinear programming (MINLP) [14]. The computational complexity of optimizing using this formulation for the problem sizes we expect to handle motivates us to consider a heuristic solver instead for the constrained assignment problem.

Our heuristic solver is broadly based on the classical "best fit decreasing" heuristic that has been developed for the bin packing problem [15]. A high level sketch of the heuristic is given in Figure 2 below. The major steps in the heuristic are identified by step numbers (e.g., S4) and comments are italicized. The algorithm has two phases and in the first phase the assignment is done by roles (step S1). Roles are considered from bottom to top consistent with the partial order specified in the resource taxonomy. Intuitively, this ordering seeks to satisfy the more specific needs first and then use the choices in availability to satisfy the less specific needs later. In step S2, the linked needs for the role being considered are collected from all projects in the set $\{Q1 \cup Q2\}$ defined earlier. The iterative sequence from S3 to S10 incrementally satisfies the linked needs for the role being considered. In each iteration, a residual target linked need N(i) is selected (step S4). Empirically, we have determined that selecting a target list with the earliest end date for a residual blocked need is a good heuristic. In step S5, selection of a person q from the current resource pool is attempted using a best fit heuristic with look-ahead. The look-ahead feature is used to determine for any candidate person q whether the upper bound on the metric F can be met for the linked need L(i) by some possible future assignments after choosing q. The conservative bound in this look-ahead step is computed by satisfying the remaining needs of L(i) with additional people added to the resource pool. The selection in step S5 also has a preferential order for choosing candidates that are checked by the look-ahead feature. The first choice for the selection process is from the set of people already assigned to the linked need L(i), then it considers availability in the original resource pool and then lastly considers any additional people added to the pool. The intuition for this ordering is straightforward in attempting to minimize the shortfall while satisfying the bound on the F-metric. If all these attempts fail then a new person for the corresponding role is added to the augmented resource pool. In step S9, the selected person q is assigned to L(i) so as to maximally satisfy its remaining needs and the linked needs are updated (step S10). Note that the assignment to q satisfies the maximum possible residual needs across multiple blocks in L(i).

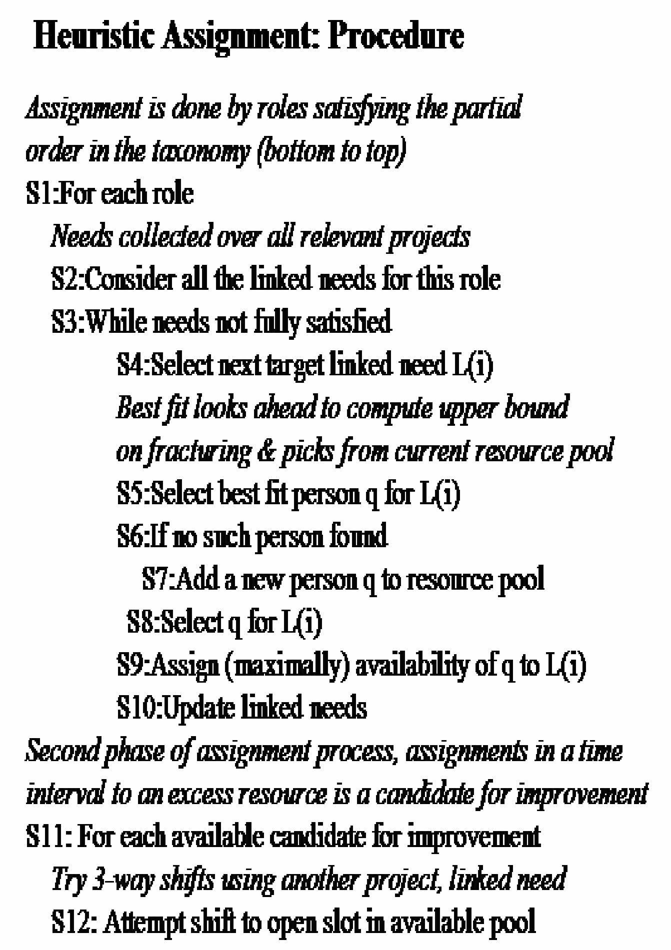

```
Heuristic Assignment: Procedure

Assignment is done by roles satisfying the partial
order in the taxonomy (bottom to top)
S1:For each role
  Needs collected over all relevant projects
  S2:Consider all the linked needs for this role
  S3:While needs not fully satisfied
        S4:Select next target linked need L(i)
        Best fit looks ahead to compute upper bound
        on fracturing & picks from current resource pool
        S5:Select best fit person q for L(i)
        S6:If no such person found
          S7:Add a new person q to resource pool
         S8:Select q for L(i)
        S9:Assign (maximally) availability of q to L(i)
        S10:Update linked needs
Second phase of assignment process, assignments in a time
interval to an excess resource is a candidate for improvement
S11: For each available candidate for improvement
  Try 3-way shifts using another project, linked need
  S12: Attempt shift to open slot in available pool
```

Figure 2. Sketch of heuristic assignment procedure

The first phase of the procedure ends with an assignment that satisfies all the needs with a possibly augmented resource pool (beyond the original resource pool) without violating the constraint on F. The conservative look ahead used in this phase could result in fracturing levels well below what was specified as an upper bound. The second pass improves the shortfall by relaxing the F metric to the allowed levels. It improves shortfall using a greedy approach that considers each time slot contribution in the augmented part of the resource pool. The improvement is attempted using a 3-way shift in resources as sketched next. Let the candidate for improvement correspond to an assignment in time interval T. All other linked needs (for the same role) that are being satisfied by an assignment to the original resource pool in the same time interval T are scanned to attempt the resource shift. The scan checks if any one of them can relinquish its assignment in time interval T by shifting it to another open time interval in the original resource pool. If so, then a three way shift is performed to reduce the shortfall provided that the two linked needs affected do not violate the bound on allowed fracturing.

The assignment after the completion of the second phase represents an achievable value of the shortfall for each scenario. Aggregating the results from multiple Monte Carlo experiments we determine the distributions for resource load and shortfall for any project portfolio option. The shortfall computed by our heuristic should be viewed as an achievable upper bound for the specified constraint on fracturing. The shortfall computed using the LP formulation of the previous section provides a lower bound though it may not be achievable for the specified constraint on fracturing. We will illustrate the use of our analysis with experimental results in the next section.

VI. EXPERIMENTAL RESULTS

The primary goal of these experiments is to illustrate the use of our resource analyses methods during project portfolio management. We will use synthetic examples for projects and resources and empirically explore the impact of various factors like the resource load levels and bounds on the fracturing metric F.

Our synthetic examples are generated by first creating resource taxonomy with some specified parameters. This taxonomy is used to randomly create an initial resource pool with varying availability for the roles in the resource taxonomy. These resources are assigned to a set of projects and linked needs created from this assignment. The actual resource pool used in the example is derived from the initial pool by deleting availability in a randomized fashion to achieve varying amounts of excess project need over the availability. Each example generated in this fashion specifies the nominal values for all the parameters discussed in Section III. We also select the uncertainty distributions to complete the specification of the example. The analysis is then done by using Monte Carlo simulations to sample from these distributions and analyzing each sampled scenario.

In the first set of experiments we will consider 5 examples and analyze the largest portfolio option in each case. The examples are comparable in their total resource needs but the resource availability is increased systematically across the examples (in the order E1, E2, E3, E4, and E5). The resource taxonomy used has two levels with 10 leaf roles. The largest resource pool (in example E5) has 100 persons distributed (non-uniformly) across these ten leaf roles. The set of projects to be analyzed, $\{Q1 \cup Q2\}$, has 20 projects. Linked needs are created for each project in terms of the leaf roles in a randomized fashion. Roughly, 200 linked needs have to be analyzed for each option. In these examples, we model the uncertainty only in the magnitude of resource estimates using a standard triangular distribution with lower and upper limits of -10% and +25% of the nominal value (i.e., the mode). In the Monte Carlo simulations we analyze ten scenarios for each example. For each scenario we determine the load and shortfall while conforming to the upper bound of 0.10 for the F metric using the heuristic assignment procedure in Section V. We also compute the lower bound on shortfall using the LP formulation in Section IV that ignores fracturing (i.e., the constraint on the F metric). Results for this set of experiments are shown in Figure 3 below. The decrease in shortfalls from E1 to E5 reflects the systematic increase in the resource pool we had created in this progression of examples. The gap between the shortfall computed by our heuristic and the LP formulation can be due to two factors. The first factor is the inherent fracturing that is going to occur due to the nature of the effective availability in the resource pool that is not reflected in the LP lower bound. For example, if people in the resource pool have bits of availability scattered over time it may be difficult to achieve low levels of fracturing for new projects. Intuitively, one can expect the first factor to be more significant when the resource pool has adequate, albeit fractured, availability. This explains in part the increase in the gaps as we go from example E1 to E5. The second factor is the sub-optimal shortfall computed by our heuristic. We will see the impact of the second factor in a later experiment.

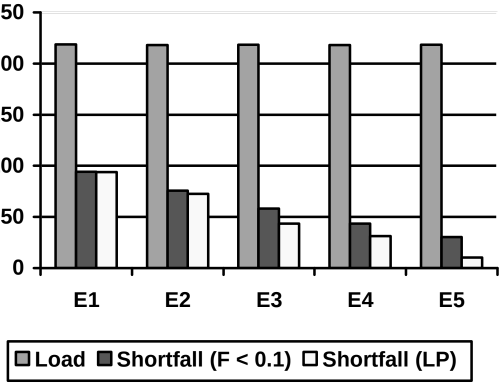

Figure 3. Load and shortfall (in 1000s hours) for 5 examples

We will use example E4 to illustrate how more detailed information can be extracted from this analysis. In order to determine the complete impact of the resource load and shortfall (including the financial impact) we would start with a breakdown of load and shortfall by resource role. For example E4, this breakdown by role is shown in Figure 4. The breakdown by role can be used to determine the financial impact by taking into account the costs for each role, for both existing resources and for excess that have to be acquired (possibly by outsourcing).

As discussed earlier, for some resource roles (e.g., specific subject matter expert), augmenting the resource pool to cover the shortfall may not be a viable option. When the project needs for such resources are specified with significant uncertainty the shortfall is typically analyzed manually to determine the impact and to consider possible mitigating actions in the organization. The plot of load and shortfall for the role R9 in Figure 5 also shows the variability due to the uncertainty in needs specification (error bars indicate the standard deviation around the expected value). The plot can be used to identify the time interval from week 10 to week 30 as having significant shortfall for this role.

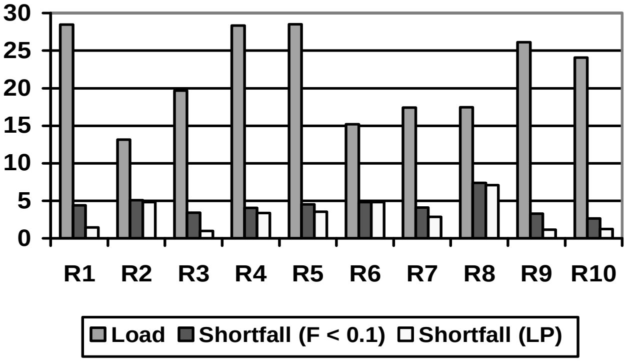

Figure 4. Breakdown of load and shortfall by roles for example E4

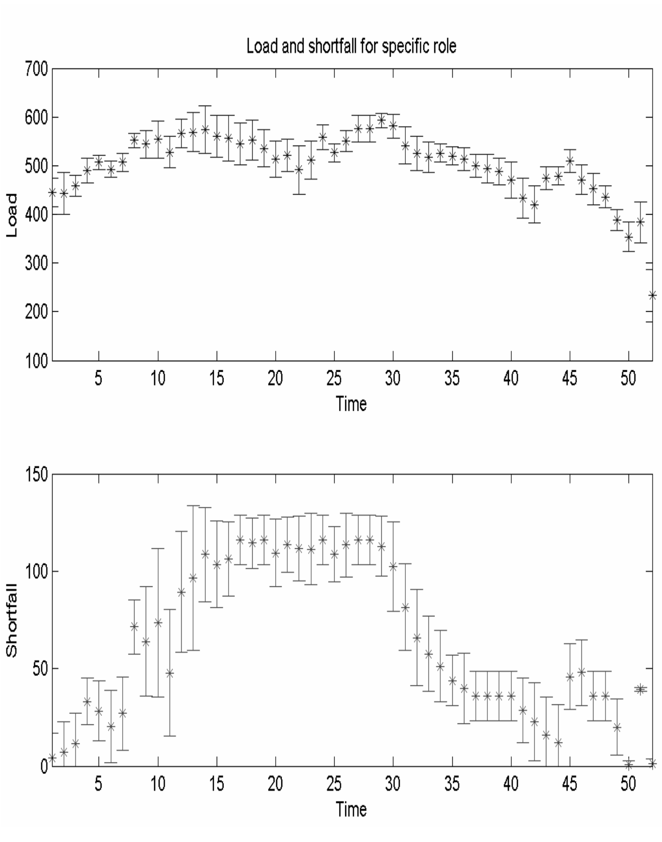

Figure 5. Temporal display of load and shortfall for resource role R9 in example E4

The next experiment illustrates the impact of the chosen level of allowed fracturing (i.e., constraint on F metric) on the shortfall. In Figure 6, we display for example E4 the shortfalls computed (using the assignment heuristic) for various bounds on F (0, 0.01, 0.05, 0.1, 0.25, 0.5, 1). The lower bound on shortfall computed by the LP formulation is also shown in Figure 6. For this example we see a significant impact as the bound on F is set below 0.25 with sharp increases in the shortfall. On the other hand as the bound of F is relaxed, the shortfall approaches the limit value computed by the LP formulation. For this example, our heuristic computes exactly the same shortfall as the LP formulation when the bound is removed (which is not guaranteed in general). Analysis of this tradeoff is important to determine pragmatic levels of fracturing in the portfolio and these can vary across resource roles (based on analyses as in Figure 6 but restricted to each role). We would expect the portfolio manager to determine the fracturing bounds for each role based on the *depth of knowledge* characteristics of the role and the analysis as shown in Figure 6 for that role.

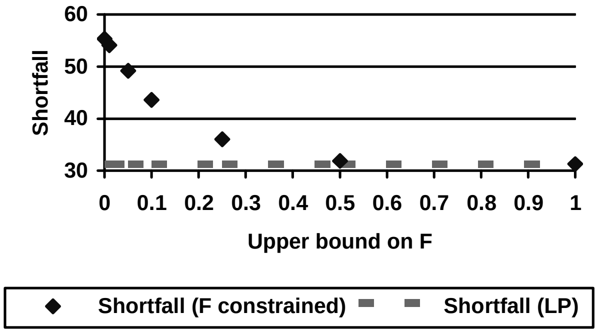

Figure 6. Impact of the F metric constraint on shortfall (example E4)

VII. CONCLUSION

We have presented models and analyses methods for evaluating the resource dimension in project portfolio management. Our approach addresses some of the unique characteristics of this dimension. In particular, our model and methods are tailored to the lack of specificity and the significant uncertainty in the estimates of resource needs typically available in the early stages of projects (e.g., proposed projects). We also factor in practical aspects of project management like controlling the level of fracturing in a team assigned to implement a set of linked tasks. The results of our analyses can be used to feed into established project portfolio management approaches and can also be used to gain insights into the state of resources in an enterprise at various levels of detail.

REFERENCES

[1] P. P. Tallon, K.L. Kraemer, and V. Gurbaxani, “Executives’ Percenptions of the Business Value of Technolgy,” *Journal of Management Information Systems*, vol. 15, no. 1, pp. 145-173, 2000.

[2] P. Weill and M. Broadbent, *Leveraging the New Infrastructure: How Market Leaders Capitalize on IT,* Cambridge, MA: Harvard Business School Press, 1998.

[3] F. W. McFarlan, “Portfolio Appaoach to Information Systems,” *Harvard Business Review*, vol. 59, no. 5, September 1981, pp. 142-150.

[4] P. P. Tallon, R.J. Kauffman, H.C. Lukas, A.B. Whinston, and K. Zhu, “Using Real Options Analysis for Evaluating Uncertain Investments in Information Technology: Insights from the ICIS 2001 Debate,” *Communications of the AIS*, vol. 9, 2002, pp. 136-167.

[5] A. Dixit and R. Pindyck, *Investment under Uncertainty*, Princeton University Press, Princeton, NJ, 1994.

[6] Amran and Kulatilaka, *Real Options, Managing Strategic Investment in an Uncertain World*, Harvard Business School Press, Boston, MA, 1999.

[7] I. Bardhan, S. Bagchi, and R. Sougstad, “A Real Options Approach for Prioritization of a Portfolio of Information Technology Projects: A Case Study of a Utility Company,” Proc. Of the 37th Hawaii International Conference on System Sciences, 2004.

[8] V. Iyengar, D. Flaxer, A. Nigam and J. Vergo, “Evaluation of IT Portfolio Options by Linking to Business Services,” Proc. Of the 2nd International Workshop on Data Engineering Issues in E-Commerce and Services, pp. 66-80, 2006.

[9] C. Mellentien and N. Trautmann, “Resource allocation with project management software,” OR Spectrum, vol. 23, no. 3, pp. 383-394, Aug. 2001.

[10] J.E. Gomar, C.T. Haas, and D.P. Morton, “Assignment and Allocation Optimization of Partially Multiskilled Workforce”, Journal of Construction Engineering and Management, vol. 128, no. 2, pp. 103-109, 2002.

[11] R. Cao, W. Ding, and C. Tian, “Using Resource and Portfolio Management Solution to Align IT Investment with Business,” Proceedings of ICEBE 2005, pp. 368-372.

[12] R.M. Wideman, Project and Program Risk Management, Project Management Institute, Upper Darby, PA, 1992.

[13] D. Bertsimas and J.N. Tsitsiklis, Introduction to Linear Optimization, Athena Scientific, 1997.

[14] C.A. Floudas, Nonlinear and Mixed-Integer Optimization: Fundamentals and Applications, Oxford University Press, 1995.

[15] M.R. Garey and D.S. Johnson, Computers and Intractability, Freeman and Company, 1979.