Methods and Benefits of Integrating the Virtual Design Team with Probabilistic Risk Analysis for Design Project and Program Planning

By

John Chachere

CIFE Working Paper #094
NOVEMBER 2004

STANFORD UNIVERSITY

METHODS AND BENEFITS OF INTEGRATING THE VIRTUAL DESIGN TEAM WITH PROBABILISTIC RISK ANALYSIS FOR DESIGN PROJECT AND PROGRAM PLANNING

by John Chachere, Stanford University

ABSTRACT

This study focuses on complex but routine design, in accordance with VDT's principal experience base. I propose methods by which the Virtual Design Team (VDT) project simulator and Probabilistic Risk Analysis (PRA) technologies may be deeply synthesized to support design project and program planning calculations such as optimal strategy and policy, value of information, and marginal value of resources. VDT and PRA are theoretically and methodologically compatible, and pragmatically complementary in several important ways. PRA has been extended to include human and organizational elements while addressing interdependencies among multiple projects' decisions, but it does not offer substantial guidance on central, challenging problems including project schedule and cost estimation. In practice, the complexity of project behavior routinely eludes expert opinion, the typical PRA data source. VDT offers to compensate for this deficiency by breaking the complexity into manageable and theoretically grounded elements. Conversely, VDT predicts many subtle and important measures of risk, but does not offer sufficient comprehensiveness or precision to address several important programmatic risks. The VDT system also does not model strategic mid-project decisions, and does not directly recommend action, but PRA can provide these functions. In this paper, I describe three methods of integrating PRA and VDT to take advantage of these various complementarities in a practical setting. I explain the key contributions that each of the integration models offers over standalone VDT and PRA, and then detail a combined formulation that employs each of the three integrations to address programs that include multiple, dependent projects. Currently, the proposed tools are particularly well justified for design projects that culminate in a risky endeavor such as a space mission. The paper concludes with an algorithmic complexity assessment, model justification plan, and sample reports.

INTRODUCTION

The fraction of major system failures that can be traced to human and organizational shortcomings is estimated to range from fifty to ninety percent [Cite]. Consequently, organizations that are serious about risk assessment and mitigation must consider their complex projects from a perspective that integrates more diverse sources of error than tradition illustrates. Risk mitigation alternatives, however, can be costly, and can impact schedule and other constraints. In project planning, effective risk management requires understanding the Product, Organization, and Process of design, the Environment in which the design occurs, and- just as importantly- the interactions between these components. In most cases, the same set of elements comes into play once the designed product is placed into operations.

This paper proposes a configuration of analytical tools and use methodologies that will help planners to design and assess a set of projects under different configurations of product, organization, process and environment. The system will enable users to plan
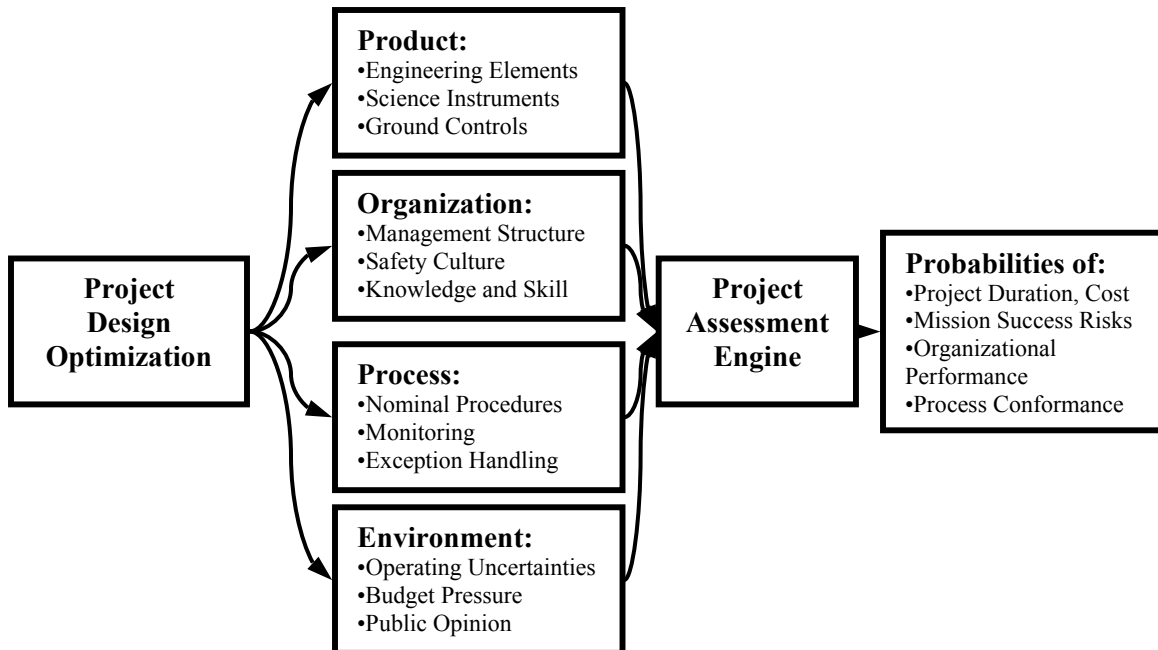
Figure 1 Top-level schematic of the model scope and purpose. Based on a guided assessment of domain expert and decision-maker knowledge, the project design optimization calculates a best configuration of controllable Product, Organization, Process, and Environment (POPE) factors. It also assesses this plan to predict the possible program outcomes. The VDT-PRA formulation also supports the analysis of multiple projects that are dependent on one another in complex ways.

future processes more effectively, early in the design phase, and to explicitly consider the associated risks under a variety of operational scenarios. On the first order, the modeling environment will allow planners to assess uncertainties about how robust their operational Processes are under different scenarios; the actual performance and properties of the Product they are building; the Organization that conducts the mission during its operational phase; and the Environment in which the product and project operate. Of equal importance, the system will allow users to assess the possible interactions between these fundamental project elements. We call this integrated Product – Organization – Process – Environment system a POPE model. Technically, is both complementary to and compatible with the SAM (System-Action-Management) extensions to Probabilistic Risk Assessment.

As diagrammed in Figure 1, we propose a POPE model and risk analysis method that supports assessment of the operational-phase risks in NASA missions as they are planned. The practical value of this research is that program analysts will be better able to plan their suite of projects' product, organization support and processes by considering the risk elements in that design early, and by frequently revisiting the model during program execution. The VDT-PRA approach to POPE modeling offers a much more realistic assessment of programmatic risks than purely engineering-driven models, and also offers a much broader range of testable mitigation strategies. In addition, our formulation will support the early stage assessment of risks associated with the complex detailed design and operations phases, through a parallel methodology.

POINTS OF DEPARTURE

PROJECT AND PROGRAM PLANNING

Project planning attempts to synchronize the range of diverse programmatic decisions that underlie a substantial, compartmentalized venture of limited scope and duration. These decisions include high-level product design and configuration; organizational participants and structure; the processes of design, manufacturing, or operations; and diverse, often less malleable environmental elements as public relations, national culture, facilities, and operating conditions.

Program planning consists of selecting and configuring multiple related projects. One program planning example would be selecting and arranging for the development and release of three software products out of numerous prospects. At another level of abstraction, program planning might address the interactions between conceptual design, detailed design, development, assembly, testing, and operations alternatives in a prospective space mission.

In order to reduce the risk of project or program failure, it is appropriate to apply the basic principle of multiple independent sources. There are many methods of assessing project and program plans, and this paper shows that formal modeling offers both advantages and disadvantages. Therefore, formal models should offer one of several diverse voices in a design conversation.

PRA IN PROJECT AND PROGRAM PLANNING

The Decision Analysis (DA) and Probabilistic Risk Analysis (PRA) communities offer an excellent overarching framework for complex, multi-attribute design decisions. PRA has the longest history in assessing risks of failure in complex products that contain many interdependent components, such as nuclear power plants. The System-Actions-Management formulation of PRA [Murphy and Paté-Cornell 1996] extends this product model to include interactions with human actions, and with management decisions that influence the likelihood of those actions. This perspective is essential in the accurate assessment of operational risks, but SAM provides limited structure for program planning.

Dillon and Paté-Cornell formulate the problem of programmatic risk management in terms of distributing funds among technical product improvements, process budget reserves (for contingency management), product functionality enhancements, and process monitoring procedures. The model calculates distribution of funds among these categories that optimizes technical (product robustness), managerial (development process costs) and strategic (product value to customer) outcome measures. The work offers an important measuring stick against which to compare the proposed formulation.

Dillon, Paté-Cornell and Guikema [2003] offers a similar formulation to address the larger problem of programmatic risk analysis in programs that contain multiple,

dependent projects. The approach also produces a measure of marginal value for the principal budget constraint.

DA, PRA, SAM, and the two described programmatic risk analysis formulations leave a deep understanding of design project performance and risk to domain experts under a guided conversation. Unfortunately, project managers are notoriously error-prone at assessing the extent of human and organizational risks, and at determining the implications of possible mitigation strategies. This is evidenced by the pervasiveness of schedule and cost overruns in a wide range of industries. As a prominent example, NASA's "Faster, better, cheaper" orientation is generally considered to have led to corner cutting as a method of alleviating programmatic risks, and consequently to have contributed to increased in mission risks and consequent failures.

VDT IN PROJECT AND PROGRAM PLANNING

Project managers frequently underestimate the emergent workloads of subordinates whose work is highly interdependent, in part because coordination efforts are not explicit in traditional planning and schedule tracking systems. We use the term "hidden work" to describe coordination and exception handling efforts that produce a substantial fraction of the total labor and schedule pressures in complex projects. Overloaded workers sometimes fail to respond to communications, thereby compounding the information supply problem and compromising others' performance. Complexity and interdependence thus results not simply in additional direct and communication requirements, but also triggers new exceptions and errors. Knowledge of this phenomenon forms the basis of many experienced analysts' skepticism toward ICE performances claims.
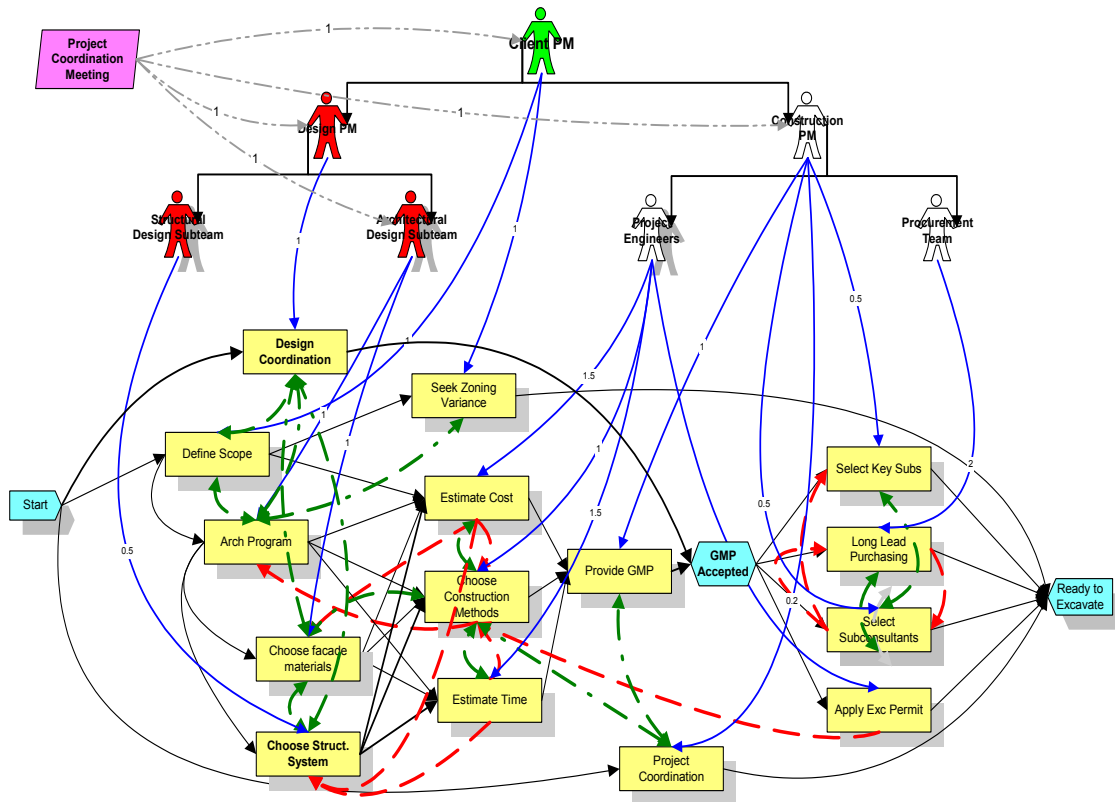
Figure 2 Typical Virtual Design Team (VDT) Project Model. VDT captures design project participants' individual characteristics, organizational exception handling structure, scheduled meetings, task characteristics and precedence, information exchange requirements, and rework dependencies. VDT compares the process's information processing load with the organization's information processing capacity. It produces detailed estimates of a project's emergent cost, schedule, quality, and other measures. We propose to combine VDT's strengths with those of PRA.

Because complex projects are notoriously difficult to predict, we attempt to enhance PRA-based program planning methods with a decision support tool that specifically targets these development risks. The Virtual Design Team simulation system (VDT) was created in part to address project managers' difficulty accurately predicting project behavior. In many industries, it is routine for schedule and cost budgets to be overrun substantially, in spite of careful planning by experienced managers. VDT reduces planners' difficult task of projection down to an estimation of project parameters such as tasks' complexities and workers' skills, in accordance with established organizational

theory. Although the system is not yet calibrated to predict all projects accurately, it has shown some remarkable successes, including in aerospace.

The Computational Organizational Modeling (COM) community quantitatively operationalizes established organizational theories, some of which are relevant to the study of risk. The Virtual Design Team is currently the most feature-rich model of hidden work. For a detailed description of VDT mechanics, see [Jin and Levitt 1996]. This study focuses on routine design projects, in accordance with VDT's principal experience base. VDT's task uncertainty and external rework links match the "complexity and interdependence" that Normal Accident Theory emphasizes, while the emergent backlog and exception handling behaviors map directly to central measures of safety culture.

Because VDT measures likely propensities, rather than certain outcomes, we interpret the system's performance measures to characterize the degrees of *risk* associated with different project aspects. For example, when a project shows high cost risk, large cost overruns are likely unless management responds with proportionally effective interventions. Although VDT the model is not capable of exactly predicting project performance, in the hands of experienced domain experts and modelers the tool has made strikingly accurate predictions of organizational performance [Cite Lockheed study].

As a simulator, the system offers numerous performance metrics, but has model of an integrative objective function. Instead, users who wish to optimize the schedule conceptualize alternatives and test them iteratively, or integrate the model with tools such as genetic algorithms [Cite Bijan].

VDT AND PRA IN PROJECT AND PROGRAM PLANNING
In his dissertation, Pugnetti [1997] applies PRA to project risk analysis to calculate optimal task sequencing under varying utility functions, and produces results that match those of VDT [Chachere 2003]. Pugnetti also couples VDT and PRA, exploiting some of their synergies, but at the same time recognizing that a deeper integration is desirable.

INTEGRATION PARADIGMS

This section describes at a high level three ways of integrating PRA and VDT. Although each combined model has distinctive characteristics, they are not mutually exclusive.

The next section details the methods and capabilities of an algorithm that applies all three integration methods simultaneously.
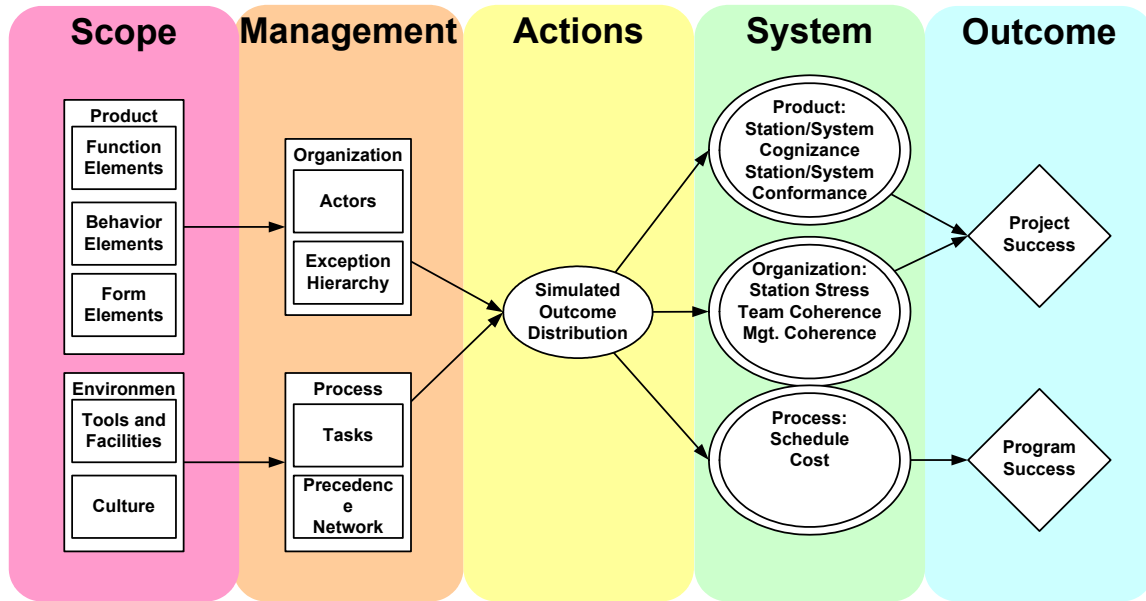

Figure 1: We conceptualize VDT from the perspective of Decision Analysis and Probabilistic Risk Analysis. The figure shows how VDT predicts distinct organization, process, and product risks, which we interpret using PRA in figure 2

PARALLEL PRA AND VDT COORDINATES DESIGN EFFORTS

Definition
This baseline alternative consists of operating VDT to assess the organizational and process design, and using PRA independently to analyze decisions principally relating to product and environment. The two models are then applied independently. It is intuitively clear that an appropriate application and interpretation of these models will improve project and program planning.

Strengths
Partitioning the decisions into those that VDT will assess and those left to PRA decouples and therefore may accelerate the assessment process. Some of this additional time may be used to refine the PRA analysis and include additional alternatives. The time may also be used to compare the VDT and PRA results informally and focus attention on complementary choices. For example, if the best organization and process that VDT analysis has produced contains a substantial risk in a particular subsystem engineering task, the PRA team may choose to compensate by selecting a simpler product design in

that area, in spite of a cost increase that previous risk analyses did not offset. This informal practice is based on the conviction that the greatest risks manifest where organizational and product shortcomings intersect.

VDT uses a fixed set of input parameters to calculate a fixed set of output performance measures. While this is accelerates data collection and improves theoretical verification, it limits the tool from some applications where a different or larger feature set is required. Perhaps the most prominent example of data that VDT does not address directly is the product that designed in an engineering project. VDT does not have an explicit product model, instead providing risks of product degradation based on indirect measurements such as rework and actor skill. When planners wish to evaluate the merits of a product configuration simultaneously with the implications it has for the design process, VDT must be enhanced by a parallel but related effort. This formulation provides this functionality by evaluating VDT and other factors concurrently, and relating them through coordination with the PRA team.

In some applications, VDT offers a second contribution to PRA, introversion in risk assessment. To illustrate, a risk station may determine that the design has low risks in a particular product element. However, this information is based on the information that the engineer domain experts provide. If VDT predicts that these representatives are organizationally or procedurally suspect, then we have reason to doubt the PRA in this domain. We also have reason to believe that later analysis stages will have trouble dealing with their bad assessment data. In these cases, to get the most accurate account of risk, we may turn to outside experts and/or second order uncertainties.

Weaknesses

Without at least a procedural integration, such as periodic discussions to share results, this analysis is not likely to lead to a jointly optimal design. In fact, if the two domains' decisions interact, performing one model's analyses without cognizance of the other model's choices will typically produce inaccurate results. This shortcoming is present in some risk analysis processes that do not include the VDT and PRA models. In addition, several mitigation strategies that employ both of the tools offer benefits beyond the sum of their parts. These strategies include procedural integration, and a synthesis with other integration methods that we describe in the next section.

VDT WITHIN PRA OPTIMIZES PROJECT DESIGNS
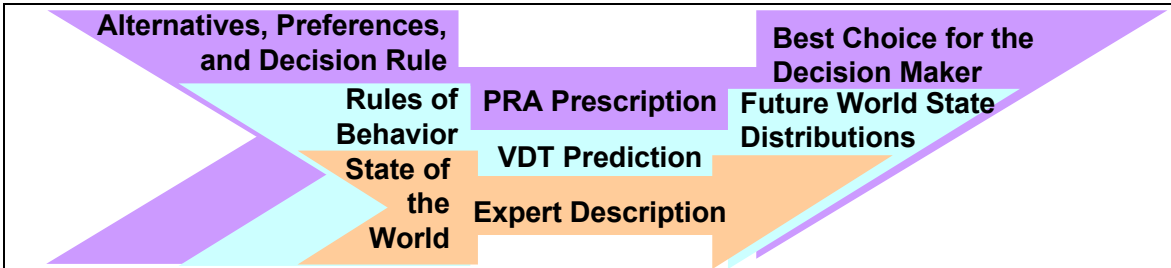
Figure 2 Methods and Purposes of Formal Analysis. Methodologically, PRA and VDT share a reliance on experts to *describe* the project and its tradeoffs. VDT and PRA are complementary in project planning because at the highest level, PRA is a *prescriptive* method that relies on the *predictions* that VDT can provide. PRA therefore rests naturally as an outer shell that uses VDT to predict outcomes as necessary.
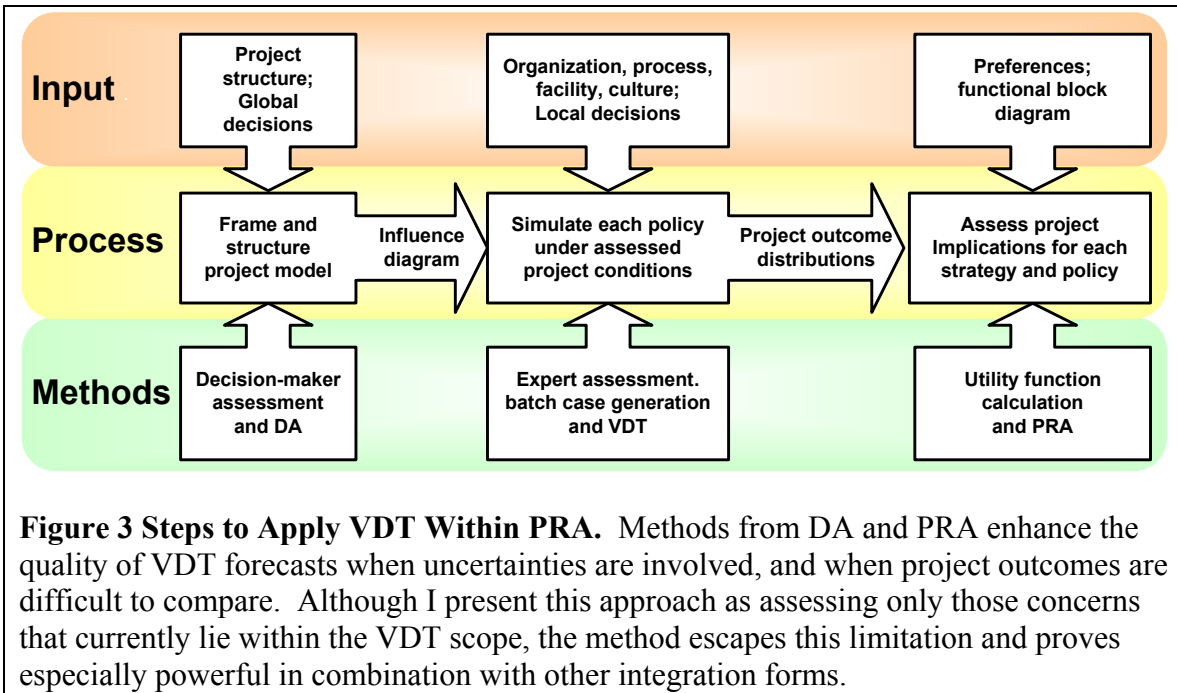
Figure 3 Steps to Apply VDT Within PRA. Methods from DA and PRA enhance the quality of VDT forecasts when uncertainties are involved, and when project outcomes are difficult to compare. Although I present this approach as assessing only those concerns that currently lie within the VDT scope, the method escapes this limitation and proves especially powerful in combination with other integration forms.

Definition

Figure 2 illustrates that the strengths of descriptive, predictive, and prescriptive claims all rely on the methods and qualities of their more fundamental predecessors. A second integration takes advantage of this relationship by using VDT's project outcome predictions within the prescriptive PRA framework. This integration model (in its pure form) addresses only the fixed set of variables in VDT's formulation, but it offers improved precision in most analyses.

Standalone VDT does not allow planners to design their product, organization, and process optimally when the project conditions are uncertain, because it only accepts point estimates as input. PRA and Decision Analysis (DA) solve this problem by evaluating all possibilities (or a statistically representative sample) and weighing them according the probabilities of their occurrence. We can continue this tradition by running VDT on each possibility, and integrating the results.

Figure 3 illustrates the input data, algorithmic process, and off-the-shelf analytic methods that lead from initial project assessment to a recommendation. The first step is to establish with decision makers a decision diagram that captures the project's major decisions and uncertainties. The next step is to work with the appropriate domain experts to model (as constants, uncertainties, or local decisions) VDT's organizational hierarchy, task network, and cultural variables. VDT then predicts the results of each possible configuration of uncertainties and decisions, and using decision-maker preferences and a functional block diagram, these results are assessed to produce a set of optimal decisions.

Strengths

In this approach, PRA takes advantage of VDT's improved ability to predict programmatic behavior in the same way that it leverages experts' projections of project outcomes. Using VDT as an information source inside a PRA methodology enables risk analysts to estimate several eminent but subtle human and organizational risk factors. In combination with other integration methods (as we show later), this facilitates the investigation of relationships to engineering risk.

As input to VDT, experts provide point estimates of actor skill, task complexity, and a range of other variables. The model treats these quantities as accurate and precise characterizations of initial project conditions. The input values influence stochastic intermediate simulation calculations, so that they do not uniquely determine performance. However, due to the microcontingency formulation, and in accordance with the probabilistic mean value theorem, for many cases the variance of simulated outcomes is modest. This small range of outcomes can mislead users into believing that the project outcomes are predictable with considerable certainty. In reality, the modeling process has eliminated some of the project forecasting uncertainty by requiring point estimates initially.

To illustrate, suppose we assess a critical actor's experience level to be high. Simulating this case in off-the-shelf VDT may show a miniscule risk of schedule overrun. However, in the current proposed VDT-PRA process we explicitly discuss the fact that experts are only 90% certain that the desired actor will be available to the project. If simulating the alternative case, with a low-experience actor, shows a catastrophic schedule overrun, this VDT-PRA process will reveal the true 10% schedule risk that VDT alone would not catch.

In order to select among the project design alternatives that are under consideration, this method requires an ability to compare the distributions of possible outcomes that will result from each possible choice. I explain this need in more detail in the next section on a combined algorithm. However it is notable here that we can model each task as aiming to serve a function in a PRA functional block diagram, with the probability of success estimated as a function of the amount of ignored exceptions and other organizational risk factors. Integrating these functional blocks into a fault tree, we can calculate a single measure of risk that represents both human design errors and (given the additional data) fundamental engineering risks. The need for clear objectives thus stimulates a productive, structured conversation on priorities, and the result improves the clarity with which decision-makers may compare outcomes.

Weaknesses

This integration method requires additional investments beyond VDT in the assessment of uncertainties and preferences, and in technologies to help assess the outcomes of multiple scenarios. Overcoming these technical and procedural challenges appears to require substantial but straightforward investments during algorithmic design, technology implementation, and case study.

In isolation, the approach also focuses attention exclusively on the VDT project outcome measures. Few of the decision-makers I target face risks that lay exclusively within the realm that VDT addresses. The effort involved in this study may mislead decision-makers to feel that they have addressed the full range of project risks, when in fact they have only produced a "gold plated" analysis of a subset of the risks. They may have, for example, ignored product design-related, financial or political failure risks. In

isolation, this integration method therefore falls short of meeting our project and program planning goals.

PRA WITHIN VDT SIMULATES PROGRAM ADJUSTMENT
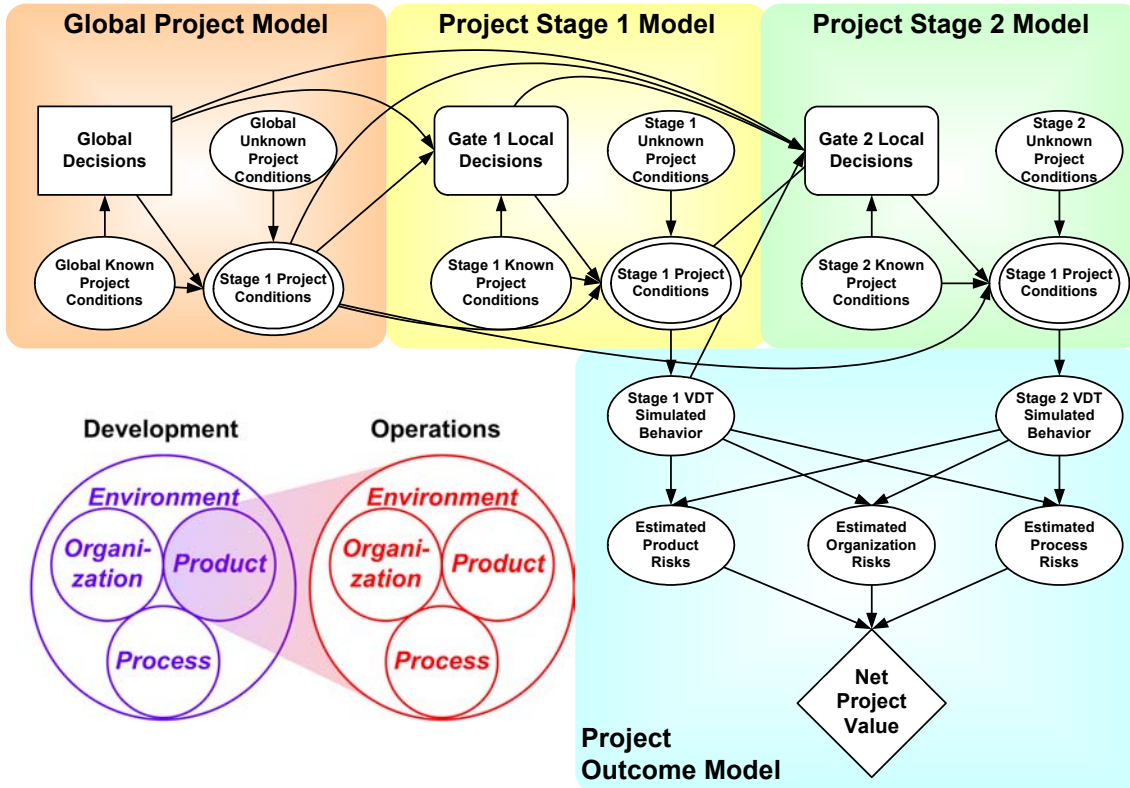
Figure 4 Influence Diagram for a Two-Stage Project. Using PRA, VDT can model mid-project course corrections, instead of assuming an unresponsive management team. The emergent program plan may reflect an optimization of the project portfolio, or it may reflect boundedly rational organizational behavior.

Definition

In many cases, planners wish to assess projects or programs that include important decisions that cannot be made in advance. Some of these mid-project decisions influence features that VDT models. For example, a project's quarterly re-evaluation of staffing levels might increase or reduce full-time-equivalent levels for a particular position.

To represent these conditions, PRA/DA and VDT must hand off execution control. During VDT execution, when a decision point is reached, the DA/PRA system evaluates the decision according to the decision-maker's conditions of rationality (see Weaknesses). If the decision-maker is able to effectively predict future project

performance, the DA system might employ VDT to evaluate the prospects (possibly continuing this method to deeper levels).

An example application of this modeling system addresses projects that are organized in the "stage-gate" fashion that is common among complex engineering disciplines. In a stage-gate project, the total work volume is divided into a sequence of stages with distinct scope and purposes. Each of these periods involves concurrent development among tasks that are designed to be independent of the tasks that occur in other stages. Once all of the work allocated to one stage is complete, the project arrives at a "gate". In the one or more large review meetings that typically constitute a gate, decision-makers review the project's progress and project its prospects. The program will not proceed to the work in the subsequent stage until the gate is passed, that is, until the prior stage's work is reviewed and approved as complete and adequate. This step is called a gate, as opposed to a hurdle, because workers typically must wait for approval to proceed, and a project can be aborted if the review shows its prospects to no longer merit the company's investment. A gate decisions may also accelerate the project by increasing parallelism, or changing to a simpler product design configuration.

Strengths

VDT does not attempt to model actors' midstream monitoring of, and intervention in, project performance. Instead, the system predicts the behavior that would develop if the organization and process were to proceed as it was initially directed to. As the duration of the project or program increases, this assumption becomes decreasingly tenable.

If we consider each simulation experiment to represent one stage, we know that for large numbers of stages no substantial degradation of product, organization, or process is acceptable in each stage. Therefore, project *stage* designers must strive for sustainable results- results that are all within a limited margin. Suppose, for example, that for a sequence of projects, schedule overruns are a chronic possibility. As program planners, we might propose that when the time comes, decision-makers may consider employing an alternative mitigation design that produces accelerated schedule. This would lend project planners additional confidence that, should these schedule delays manifest, there are steps in later stages that may confidently be applied to compensate. The VDT-PRA

integration described here captures the benefits that this conditional plan ("Policy") offers.

Multi-project program optimization is the subject of the APRAM Model [Dillon and Pate'-Cornell 2003]. APRAM shares many similarities to the VDT-PRA synthesis proposed here, including the consideration of multiple interdependent projects. The APRAM algorithm is structurally like this proposal, but here VDT supplies an improved model of programmatic risks (when compared to APRAM's decision tree). In addition, we recognize that program participants make decisions according to a limited amount of information and information processing capacity [Simon, SAM Model]. We therefore consider that mid project course corrections result from boundedly rational decisions (see Weaknesses).

Weaknesses

When predicting actual decision-makers' future choices among alternative project plans, we should consider the many important human and organizational cognitive limits. Organizational research indicates that humans do not make decisions in the stereotypical, "rational" manner that we (as decision analysts do) [March 1994]. They do not have unlimited knowledge of possible alternatives, information, and information processing capability [Simon], and often demonstrate preferences that are very different from those we impute upon them. Even some of the most commonly applied management interventions are misunderstood, and can be counterproductive. For example, VDT may calculate a vastly extended schedule, whereas real managers would likely observe the developing problems and attempt to intervene. Hiring additional workers under these circumstances frequently extends project schedule, rather than reducing it [Brooks 1995]. Great care is necessary to determine whether a DA/PRA formulation of mid-project decision-making captures an accurate prediction of "boundedly rational" (or neo-institutionalistic "identity-based" [Powell and DeMaggio]) choice. In chapter 5 we discuss one justification case study, "VDT v. PRA v. PRA-VDT", that explores the importance of bounds to rationality.

Allowing multiple decisions that interact in defining a VDT model produces an exponential number of simulation demands, which can be challenging to efficiently execute and visualize. This effectively prevents insertion into our VDT model of

decision points at frequent project junctures, even though in many applications adaptation can occur at any time.

In many circumstances (formally detailed under the efficiency heading), we can accommodate this computing power burden by streamlining the model. This involves segmenting the VDT case into different projects, and simulating each alternative in turn. The system can then store the alternative projects' results separately for flexible integrations into subsequent evaluations.

Another method of coping with this issue is to allow VDT to trigger decision points under formally defined circumstances. This may involve calculating (and possibly optimizing) warning levels or review points.

AN INTEGRATED PRA-VDT ALGORITHM

Figure 6 shows how we can implement all three integration methods within the same, flexible algorithm. Although I list the steps in sequence, in practice it will occasionally be appropriate to revisit and refine previous work as learning occurs. This section details each step in this algorithm.

1. DEFINE PROJECT DESIGN PARAMETER DISTINCTIONS

The VDT-PRA system evaluates a broad range of project-related quantities to predict the behavior of complex projects and to recommend how they should be executed. VDT provides one part of our language for communicating the building blocks that determine project performance. The language relates easily between established organizational theory and practitioners' intuition, and helps those familiar with VDT to rapidly understand and communicate new project designs. This language provides modeling distinctions that pass the DA/PRA "Clarity test", and that are difficult and time-consuming to develop from scratch.

The VDT parameters include cultural measures such as centralization, formalization, and matrix strength; Organizational factors such as actors with varying levels of skill and experience; and Process definitions including tasks with varying levels of procedural uncertainty, complexity, and required skills. The relationships among these elements are also clearly defined in the VDT literature, including authority hierarchies that interrelate actors; Primary and secondary task assignments indicating which actors address which

tasks; And links interrelating tasks that have rework and information exchange dependencies.

VDT-PRA offers an opportunity to add parameters that do not originate in VDT, to capture those features that decision-makers feel are important for the specific application. For example, expert assessment may indicate that considerations that are outside the purview of VDT are essential to the decision being made, such as product design configuration, operating environment, or public opinion. If these parameters are orthogonal to VDT behavior, or can be adapted to suit this definition, the core VDT model may suffice with only peripheral embellishment of the decision diagram. However, in cases where these considerations are more complex and significant, the PRA framework and expert assessment can be called upon to greater extent to include decisions and complex behaviors that are dependent upon VDT behavior.

2. CAPTURE DOMAIN EXPERT AND DECISION-MAKER KNOWLEDGE

VDT and PRA are modeling methods that rely upon domain experts for data. When constructing a project model, VDT experts first help domain contacts such as project managers to understand each of the key distinctions of the model. After selecting an appropriate level of abstraction for actors and tasks, modelers conduct a conversation in which domain contacts offer point estimates of actors' skill, exception flow paths, precedence networks, and the other distinctions that VDT makes. The team reviews these parameters and refines the model progressively according to the perceived accuracy of subsequent VDT simulation behavior.

After defining the model structure (see the previous step), PRA-VDT presents two important procedural differences from VDT. The first applies where experts are uncomfortable approximating parameters using point estimates, or where model sensitivity analysis shows heavy dependency of model outcomes. For these variables, the analysts consult with domain experts to assess a distribution of possible values for the important variables. In some cases, analysts help the domain experts to develop joint probability distributions for multiple, dependent variables. For example, a particular task might be estimated to be complex half the time, and simple otherwise.

The second key difference is that the decision-maker identifies some quantities as choices available to be influenced directly, as recommended by the current planning process. For example, project designers might wish to analyze the decision of whether to spend money on improving engineers' skills. This transforms a VDT actor skill characteristic into a decision variable, while the variable also influences cost. As another decision example, the designers may recognize that between the design and development project phases, the team will have the ability to make an informed decision of whether or not to abort the project.

Uncertainty in Organizational Micro-Behavior

As a Monte Carlo simulation, VDT converts qualitative and quantitative project design metrics into irreducible, mathematical distributions of outcome. Specifically, the system describes outcome measure distributions using average results and variances that characterize the model's degree of certainty. For example, VDT does not predict a single quantity of information exchange requests, but instead states a simulated average and variability.

This variability results from the simulated project's sensitivity to conditions that are too minute to predict in advance. For example, whether an actor attends to one memo first, and another second, or vice versa, can determine the course of ever-larger events (as described by the "butterfly effect" principle of chaos theory [Cite]). In general, project designers cannot exactly predict workers' moment-by-moment attention in advance. VDT therefore captures the behavior by choosing at random from distributions that have been found to reflect long run behavior. Although in some cases, the minutiae dampen or cancel out and become less significant, there is no known general method of determining when this is the case. Consequently, the simulation can predict outcomes with a degree of certainty that is limited to a project-specific degree.

Decisions as Project Design Parameters

From the perspective of organizational theory, we interpret this routine type of decision to fall within a "Zone of indifference" [Cite] in which decisions occur somewhat automatically. VDT's random attention rule selection method captures this behavior in a simple stochastic behavioral model. With strategically important and thoughtful

decisions, we substitute a more "Boundedly rational" [March 1994] focus of attention. Even from a neo-institutionalist viewpoint [Powell and DeMaggio, Scott], this behavioral formulation is especially appropriate because rationality is a principal identity of engineers.

Uncertainty in Project Design Parameters

In some cases, project designers may be uncertain about important aspects of project resources. Planners may feel that actors are likely of high skill, but they may recognize a ten percent chance in each case of actors being of medium skill. When there are ten such actors, for example, on average one actor actually has significantly less ability than was is predicted at first order. If the tasks are parallel, this implies that the total project duration will be determined by medium skill, not by high skill. This example illustrates the importance of including uncertainty in simulation parameters.

Procedurally, this requires expert assessment of likelihoods of each input parameter. In addition, it invites a more thorough evaluation of outcome risks. For example, if a project shows a ten percent chance of unacceptable behavior, we will require the ability to trace this result back to possible input parameter uncertainties, such as a ten percent chance of sub-nominal skill in a key role.

Technically, this warrants the design and implementation of an iteration shell that runs VDT on all possible combinations of input parameters. This evaluation leads to a combinatorial explosion of input parameters and corresponding simulation runs. Where computational complexity exceeds the available resources, practitioners can employ Monte Carlo simulation to approximate. Several well-studied techniques attempt to manage this, including the thorough evaluation of all possibilities and post-hoc weighting according to likelihood, and the probabilistic selection of input parameters with the simple averaging of results.

Types of Project Design Parameters

In a stage-gate project, domain experts place each project design parameter into one of six categories:

Global Constants

In a VDT-PRA application, most of the quantities that characterize project behavior are treated as fixed. For example, project planners may choose only to consider a centralized decision-making policy, or they may choose to include weekly, one-hour staff meetings in every stage and in every scenario. Those quantities that are known with certainty, that do not vary among alternatives, and that are invariant among stages are termed global constants.

Global Uncertainties

Some quantities that characterize project behavior are not under project planners' control, but neither are they known with certainty. In these cases, planners must assess the probabilities that a given variable might take on the various possible values. For example, a project might employ the same propulsion team at every stage, but it might not be know whether they are of high, medium or low experience. In this example, the project planners must assess the probabilities of each case, for example assigning a 1/3 probability to each case because they are equally likely. These quantities are called *global uncertainties*.

Global Decisions

Project planners use VDT-PRA to predict the performance of projects, but also to select among different alternative plans. The first decisions that VDT-PRA assists with are global decisions, which involve variations that are consistent among all project phases. For example, planners may choose to employ a formalized structure, but if so, this culture will persist throughout the project. VDT-PRA will evaluate the results of each available alternative for each global decision, and determine the best possible configuration of choices (given the planners' preferences).

Local Constants

It is common for project parameters to be known in advance, and to vary among project phases. For example, the volume of direct work for each task might increase predictably with each phase. Project planners must assign values to each of these *local constants* for every phase in which they appear.

Local Uncertainties

Sometimes planners will expect quantities to vary among phases. For example, the amount of time that a manager allocates to external projects might be expected to lie at 80 percent or 90 percent time, with equal probability, in each phase. VDT-PRA users use domain experts to estimate these uncertainties so that the system can calculate the decision-making implications of each possibility.

Local Decisions

The last class of quantities includes those decisions that review teams make at each gate. Before each phase, for example, teams often make go/no go decisions, sometimes quitting a project that is considered unlikely to be of net benefit. Other examples include the ability to accelerate the project by increasing parallelism, or to select a less complex design process. In the VDT-PRA analysis, these decisions can be made according to the principles of decision analysis, with complete knowledge of all previous design phase outcomes and accurate understanding of the uncertainties and decisions involved in subsequent phases. Alternatively, a substitute model may be designed to approximate the actual behavior of decision-making project participants.

Indirect Observation of Uncertainties by Project Participant Decision-Makers

The distinctions in the stage-gate example are structured in order to prevent a condition that is difficult to resolve without additional assumptions. This occurs when a VDT project variable influences behavior but is only observed indirectly before a subsequent decision. Intuitively, a decision maker is able to observe the indirect impact of an unknown quantity, and will therefore use an informed estimate in making decisions. That is, project decision-makers use the limited observed data to update their estimates of the unknown variable. The challenge lies in defining the precise knowledge available to the decision maker.

Practitioners must either craft an answer to this question that program designers and their appointed domain experts are happy with, or formulate the problem in a way that prevents the condition. One method of preventing mid-project parameter estimation is to guarantee that variables that have influenced project behavior should be revealed at the

next decision point. The decision diagram in Figure 4, for example, describes a two-stage project that does not manifest the issue. If the variables cannot be so arranged, the rationality of the decision maker must be estimated, for example with Bayesian updating.

We can enhance the project decision diagram to include an observed uncertainty that is probabilistically dependent upon the uncertainty of concern. Using a pure influence diagram, a rational decision maker would employ an algorithm such as Bayes-Ball [Cite] to generate an improved estimate of the quantity that the unknown variable has taken. This solution therefore requires an additional assumption, implicit in the indirect observation node's definition, that the emergent consequences of the principal uncertainty's true value (that may be VDT-generated) do not influence the decision-maker's estimation of the unknown quantity.

An alternative is to push the calculation through by generating a continuous approximation to the VDT output parameters. Because VDT is a Monte Carlo method, probability densities for the observed quantities are not available. The VDT simulator therefore does not directly support the calculation required by Bayesian updating. This method adds an assumption about the perception of project behavior, since we have not actually employed continuous estimates in solving the project design problem. Even though this assumption regards decision-maker rationality, and not direct project behavior, the addition of any assumption complicates the model's justification.

3. PREDICT THE INFORMATION PROCESSING BEHAVIOR

In this step, we calculate the set of decision and uncertainty configurations in which each project may occur. If sufficient computing power is available, we run VDT with a large number of different random number seeds, so that we are guaranteed to produce a desired confidence interval on our final utility function (see step 5). If computing power is a constraint, selecting the number of execution times may involve approximation methods. For example, we may execute each VDT configuration a number of times that is estimated to be proportional to the joint probability of its uncertainties, given the configuration's choices.

Figure 2: VDT obeys the mathematical rules of probabilistic risk analysis (PRA), as well as the rule of thumb that complexity should be broken down into fundamental, intuitive pieces. PRA's traditional method, expert assessment of uncertainties, is operationalized with organization theory and simulation.
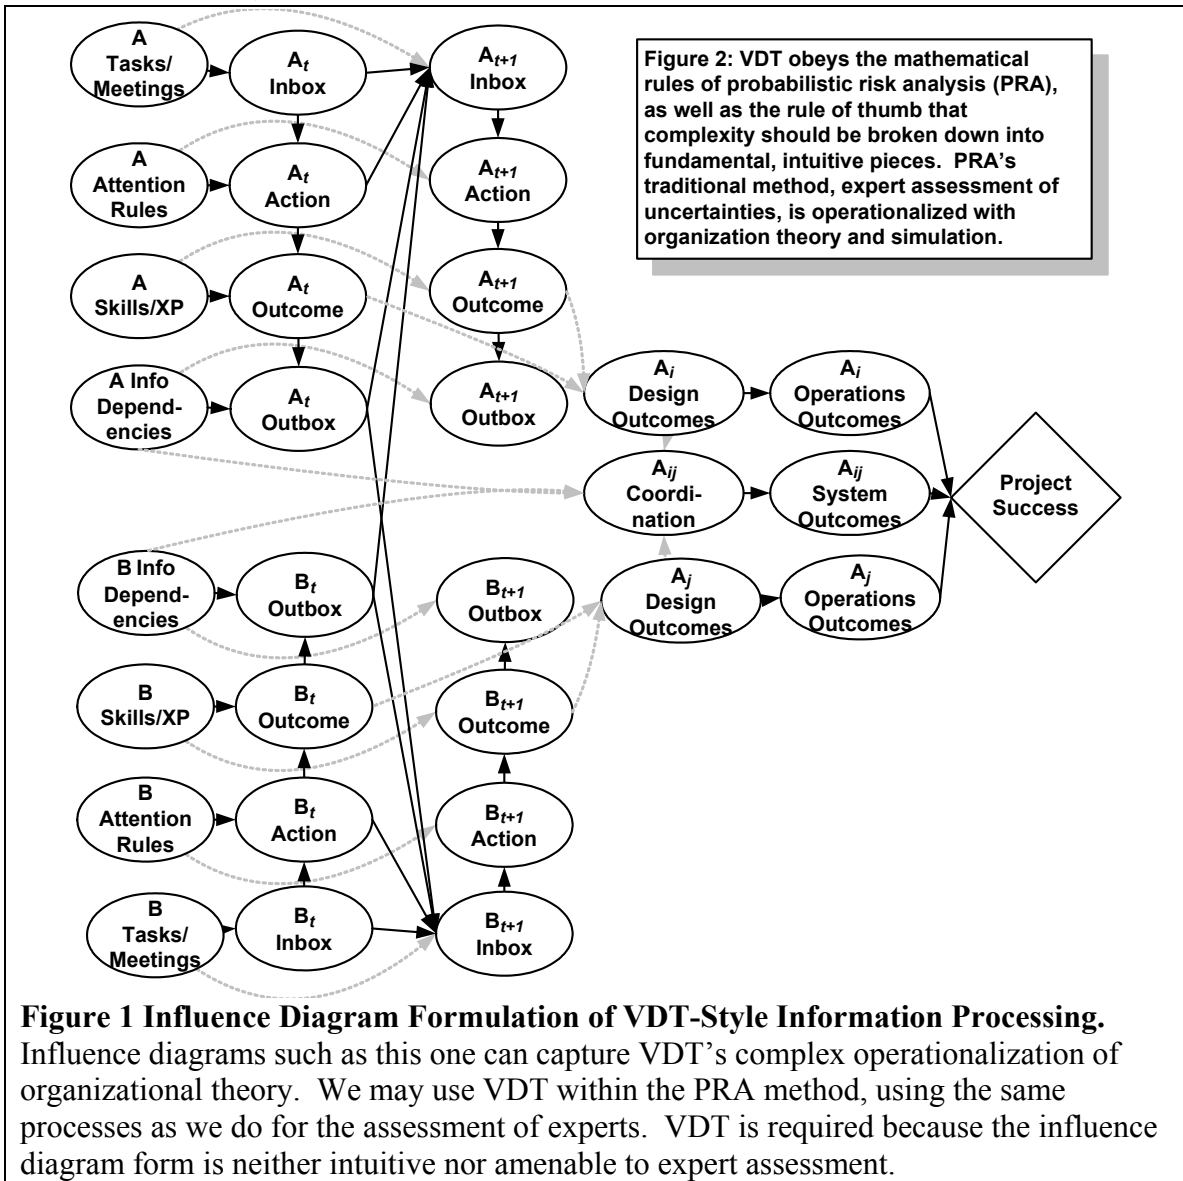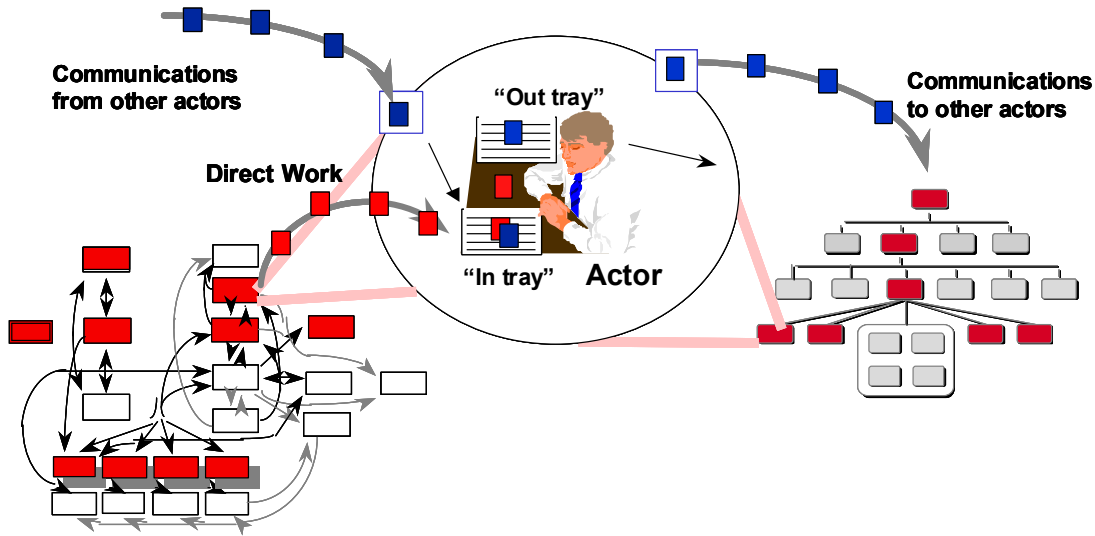
Figure 1 Influence Diagram Formulation of VDT-Style Information Processing. Influence diagrams such as this one can capture VDT's complex operationalization of organizational theory. We may use VDT within the PRA method, using the same processes as we do for the assessment of experts. VDT is required because the influence diagram form is neither intuitive nor amenable to expert assessment.

VDT provides a range of product, organization, and process performance measures, including emergent work volumes, a project schedule, and coordination rates. We describe the VDT outcomes as measures of risk to the mission design product, the design organization, and the conceptual design process. The interpretations in this section extend previous work in the VDT research community by lending the increased interpretive precision that useful and accurate PRA integration requires.

Product Risk

In a design project, I use the phrase *product risk* to describe the likelihood that design choices are functionally inadequate. Accurately estimating product risk is important because it may lead to improper decisions over whether to proceed with a later project stage, or to choices that are needlessly costly, risky, or extended in schedule.

Because VDT does not include an explicit product model, this algorithm relies on PRA to predict the cost, quality, or other measures of a designed product. This algorithm does use emergent VDT project behavior however to predict the likely *accuracy and completeness* of the team's own evaluation of these factors. Our analysis highlights the indirect impact of organizational risk factors on product quality because they are estimated to contribute to 50-75% of major modern catastrophes [Paté-Cornell 1990, Murphy and Paté-Cornell 1996]. Every design organization requires appropriate stations as well as an effective collaborative process to correctly estimate the product's risk, costs, quality, or other features. Benjamin and Pate-Cornell highlight the need for probabilistic risk analysis in a design project setting [2004], and the observed team's new Risk Station testifies to its increasing practical role [Meshkat and Oberto 2004]. Our analysis of product risk is distinct from, and complimentary to these efforts.

In the VDT simulation, overloaded or unqualified actors tend to ignore exceptions and information exchange requests, which contributes to product risk metrics. VDT calculates four types of measures that relate risk to the product design:

- *Functional Exception Risk* measures the rate of rework (or design iteration) that is ordered for individual tasks. High functional risk at a particular station indicates that the station's design is likely to be independently faulty.

- ***Project Exception Risk*** measures the rate of rework or design iteration that is ordered in response to interdependencies among functionally related tasks. When a simulation shows high project risk, this indicates a propensity for failures in the interfaces between two product components.

- ***Information Exchange Risk*** is the fraction of information exchange requests that stations take the time to complete. High communications risk indicates that interrelated tasks are not always sharing information appropriately, which tends to reduce integrated design quality.

- ***Meeting Risk*** measures the rate of attendance at design group meetings. Inadequate attendance at these meetings tends to increase the risk that a lack of global cognizance will lead to the complex but increasingly prevalent "System of System" failures.

We can predict overall design quality using VDT by assessing these metrics at an aggregate project level, or we may drill down to characterize the product in detail. For example, elevated project risk at one design station indicates that other subsystems have not redesigned according to its particular needs, while a high communications risk at another station suggests that its product does not include relevant design details.

Organization Risk

By *organization risk*, we refer to the likelihood and consequences of events that degrade the operating effectiveness of the design team (Team-X) itself. VDT measures several important pressures on the organization that can, especially over time, reduce its operating effectiveness. People who are under time pressure or stress are more likely to make poor decisions [Janis 1982.1], and errors of oversight. In addition, they are more likely to burn out and leave a position, and in complex positions is another risk factor. Projects that produce high organizational risk are likely to burden the participant teams in a way that is not sustainable in the long term.

- ***Backlog*** tracks the amount of work that waits for the attention of each actor, and this can cause the sense of time pressure that researchers have shown to cause errors [Janis 1982.1].

- ***Lateral Strain*** measures actors' frustration with their peers. VDT measures that indicate this include the fraction of information exchange requests that are not granted, and the delay in response.

- ***Vertical Strain*** measures actors' frustration with management. VDT measures the fraction of exceptions that are ignored, and the amount of time that a participant spends waiting for management decisions.

Process Risk

VDT offers two measures of process risk that anticipate the perceived efficiency of the design project:

- ***Cost*** estimates the financial burden of the design project based on the total emergent work volume among all designers and supervisors.

- ***Schedule*** can be viewed as the total project schedule, or time between execution of the first and last work items. Alternatively, VDT calculates detailed schedules including average start and finish times for each task in the process.

4. PREDICT THE CONTEXTUAL BEHAVIOR

For descriptions of over a hundred human and organizational risk factors, and related literature reviews, see Ciavarelli [2003] or Cooke, Gorman and Pedersen [2002]. This algorithm relies on the SAM formulation of PRA [] to augment VDT in important areas it does not address. Examples of factors that VDT does not evaluate include conformity, which decreases the likelihood that individuals will contradict peers' public, erroneous statements [Festinger 1954]. "Groupthink", reduces the likelihood of thorough, critical evaluation of alternatives in a group setting [Janis 1982.2]. Finally, the "Risky shift" phenomenon leads groups to select choices that are more risky than those which any participant would individually choose [Bem et al 1965]. Each of these factors relies on analysis outside VDT, and should act to modify estimates of the quality of the selected design.

   As a second important part of the integration of VDT and DA predictions, VDT-PRA implementations may include explicit choices of design configuration and assessments of

inherent associated risks. This model is complementary to the VDT formulation of design process quality. For example, one product configuration may have an inherent 10 percent failure rate due to uncertainty in the selected materials' quality. At the same time, VDT may conclude that there is a 10 percent chance that the design itself includes errors. This might lead to an integrated assessment that the product is likely to function effectively 81 percent (.9 * .9).

F igure 2 Example functional block diagram and fault tree. We can derive a functional block diagram from process definitions starting with one essential function per design task. For example, a project with three tasks includes three sources of possible unidisciplinary risk. We then transform VDT's measures of information and rework dependencies into instances of interface risk. Finally, VDT measures of project meeting attendance provide information about the subtlest, multiple-system risks. Projects that employ risk stations may develop contingent plans, which also appear in the diagrams.

5. APPLY UTILITY FUNCTIONS

VDT stops short of selecting the optimal choice among project plans because it includes varied and apparently orthogonal outcome measures. PRA and DA methods solve this problem, wherever possible, by first providing the most reasoned possible comparison of outcome measures. For example, VDT-PRA might condense measures of two systems'

reliability into a single measure of total system reliability that we aim to maximize. Similarly, for stage-gate programs we view total schedule as the sum of all stage durations. Once we have defined these aggregations, we assess decision maker's opinion of the relevant tradeoffs, in a process that is similar to uncertainty assessment. For example, a decision maker may view our top-level product, organization, process, and environmental predictions as equally important in selecting the best alternative.

Because the commercial form of VDT presents results in terms of mean and variance, it does not provide enough information to support decision makers who are not neutral to risk. In addition, presenting outcome measures in these aggregated terms excludes information about the correlations between outcome measures. As a consequence, the standard VDT results presentation format is (in general) inappropriate for evaluation by decision-makers whose utilities are nonlinear functions of different outcome measures. To illustrate, VDT may report that half of the time cost was overrun, and half of the time quality was unacceptably low. However, the standard VDT output would not report whether these were the same half of the cases, and we could not predict whether one half, one quarter, or none of the cases would meet the goals of a decision-maker who absolutely required both low cost and high quality.

The solution is to request VDT simulations with one random number seed at a time. Each simulation run is individually evaluated against a (possibly nonlinear) objective function that incorporates PRA risk measures from outside VDT. The individual simulations' resulting utilities are then averaged to the level of decision and uncertainty configuration.

6. *PREDICT MID-PROGRAM DECISION-MAKING POLICY*

The previous section explains that some studies call for mid-program decision-making by project participants. In this step, we use our predictions of project conditions at the decision points and (where appropriate) projections of uncertain prospects to identify the choices that participants will make under each of the anticipated project conditions. The result is a definition of the conditions under which each possible choice will be made.

This VDT-PRA framework algorithm requires that local decisions may be modeled as uncertainties, in alignment (for example) with neo-institutionalist views, or as boundedly rational decision-making. In cases where we approximate the latter using the DA

framework, it is important to provide utility functions that estimate the behavior that project participants will demonstrate, rather than to assume alignment with program management. In particular, local decisions may result from ignorance of emergent quantities, or an interest in outcomes that serve the existing incentive structure (itself a possible decision variable).

7. PRESCRIBE PROGRAM STRATEGY AND POLICY

At this point, we have enough information to identify the optimal strategy- the configuration of decisions that maximizes the decision-maker's expected utility. In addition, we define the policies that will guide future decisions once the project is under way. These calculations of strategy and policy become straightforward implementations of decision analysis, once we recognize that each individual VDT simulation represents one possible resolution of an uncertainty node.

8. ANALYZE SENSITIVITY

In addition to understanding the baseline performance of a project, it is important to know whether the project design has *structural stability*, or whether it is sensitive to small deviations that can be difficult to anticipate. For example, one group might require that a specific station be staffed by an engineer of extraordinary skill, and might stumble when happenstance requires a more average member to substitute. An organization should serve a routine global function only if it is effective both in optimal conditions and under foreseeable organizational and other variations. While some of these uncertainties may be explicit in the VDT-PRA formulation, others that have escaped attention may in fact hold dramatic sway over the project outcome. In particular, VDT includes a number of internal variables that influence predictions. Some of these quantities may not have been previously recognized as critical and calibrated appropriately for the program under study (or modeled as explicit uncertainties).

9. CALCULATE RESOURCES' MARGINAL VALUES

VDT does not explicitly calculate the marginal value of input parameters such as individual actors' skills. However, we can run the VDT-PRA algorithm on both cases and determine the "value of control" over one or more project variables. This is

important because it can inspire and guide decision-making in tradeoffs that extend beyond the formal analysis scope.

10. CALCULATE VALUE OF INFORMATION

Using the standard DA approach, VDT-PRA can measure the value of providing decision makers with information about project uncertainties. In this, modelers should keep in mind however the previous comments under "Indirect Observation of Uncertainties by Project Participant Decision-Makers".

ANALYSIS OF THE ALGORITHM

CONCISE ALGORITHM DEFINITION

This section begins with a condensed definition of the proposed VDT-PRA algorithm, without dwelling on the algorithm's merits.

1. Document the study intent in intuitive terms that can be compared both with outside "laypersons'" documentation and with model structure. Include the project intent, scope of each phase, and action items at each mid-project decision point. Define task network and organizational structure in intuitive terms, if they are constant.

2. Categorize each project parameter as a global constant, global variable, global decision, local constant, local variable, or local decision.

3. Record the values of all global constants, assess the value distributions for global uncertainties, and enumerate the alternatives for each global decision.

4. With the program manager or decision maker, develop a utility function that quantitatively relates the merits of all foreseeable outcomes. This step is complicated by the fact that many product, organization, and process variables must all be compared and combined into a single measure. Unless such a measure can be obtained, it will not (in general) be possible for the system to select the optimal local and global alternatives.

5. For each project phase, record the values of all local constants, and assess the value distributions for local uncertainties, with domain experts. For each local

decision, define the set of alternative values and assess any extraneous implications (such as added cost) associated with the choice.

6. With experts on the projects' gate procedures, develop utility functions that will be used to make local decisions. These may be the same as the global utility function, or they may differ due to self-interest, myopia, bounded rationality, or other factors.

7. For each global strategy (combination of global alternatives)

    a. For each global possibility (combination of global uncertainties)

        i. For each stage 1 strategy (combination of local alternatives)

            1. For each stage 1 possibility (combination of local uncertainties)

                a. For each trial (to get outcome distributions)

                    i. Simulate stage 1

                    ii. Record outcomes of stage 1

                    iii. Proceed to calculate outcomes for all later stages

                b. Calculate the utility of this stage 1 possibility

            2. Calculate the expected utility of this stage 1 strategy (multiply the probability of each stage 1 possibility by its utility, and sum)

        ii. Select the stage 1 strategy that maximizes expected utility for this global strategy and possibility. This is the optimal stage 1 strategy to follow, given the global conditions. It produces the effective utility of this global possibility.

    b. Calculate the expected utility of this global strategy (multiply the probability of each global possibility by its utility, and sum)

8. Select the global strategy with the highest expected utility. This is the optimal global strategy to follow.

9. Review results, including sensitivity to both explicit input and internal modeling parameters

10. Iterate to improve accuracy and refine plans.

11. Calculate value-of-information and marginal resource values to support external initiatives

12. Document Results

13. Implement Recommendations

OPTIMIZING THE ALGORITHM

This algorithm addresses the general circumstances of multiple decisions under uncertainty and progressive revelation. Unfortunately, it can require a great deal of computing power. Specifically, for x ternary global variables (uncertainties plus decisions), and n stages with y ternary local variables each, we would have to run $O(3^x 3^{ny})$ VDT cases, and calculate the objective function $O(r 3^x 3^{ny})$ times (for r runs per case), to calculate the optimal policy. This may be impractical for any large number of variables.

Reducing the Number of Simulations

Fortunately, a few common assumptions that match the stage-gate project structure vastly reduce the required number of simulations. Operations research techniques often rely on the concept of "pinch points" for efficiency. These are moments at which as a result of prior calculations we may disregard large amounts of previously essential knowledge, and characterize a system's state succinctly as one of a relatively small number of possibilities.

Stage-gate projects offer us natural pinch points between each stage because we may typically disregard the details of previous stages and focus only on their outcomes. Often in engineering projects, a great deal of work goes into design documents, and these documents are the principal elements that are reviewed at a gate and passed forward to guide the next stage. From an information processing perspective, the documents create organizational pinch points.

To establish a manageable execution time, our algorithm can allow that decisions at each gate involve a review of prior stages' outcome measures- but not the details of previous stages' decisions or input conditions. More importantly, behavior at each stage is independent of behavior at prior stages, once uncertainties are resolved and decisions are made. In these cases, we may simplify our calculations by simulating behavior in

each stage, for each set of input variables, *independently* of all other stages. This assumption reduces the number of VDT simulation cases from $O(3^x 3^{ny})$ to $O(n3^x 3^y)$, which is linear in the number of stages, though still exponential in the number of variables.

Reducing the Calculations of Objective Function

Unfortunately, implementing the gates as pinch points does not reduce the number of times that we must calculate the objective function ($O(3^x 3^{ny})$). In cases where the stages contribute to the objective function in a linear fashion however, as in the case where they are independent of one another, we can calculate the utilities (and optimal strategies) for each stage independently ($O(n3^x 3^y)$).

BIBLIOGRAPHY

Asch, S. E. (1987, original work published 1952). *Social Psychology*. New York: Oxford University Press.

Bem, D., M. Wallach, and N. Kogan (1965). "Group Decision Under Risk of Aversive Consequences" *Journal of Personality and Social Psychology*, 1(5), 453-460.

Benjamin, J. and M. E. Pate'-Cornell (2004). "Risk Chair for Concurrent Design Engineering: Satellite Swarm Illustration" *Journal of Spacecraft and Rockets* Vol. 41 No. 1 January-February 2004.

Brooks, F. (1995). *The Mythical Man-Month* (20th Anniversary Edition), Addison-Wesley.

Carley, K. (1996). "Validating Computational Models" Working paper prepared at Carnegie Mellon University.

Ciavarelli, A. (2003). "Organizational Risk Assessment" Unpublished manuscript prepared at the Naval Postgraduate School.

Cooke, N., J. Gorman, and H. Pedersen (2002). "Toward a Model of Organizational Risk: Critical Factors at the Team Level" Unpublished manuscript prepared at New Mexico State University and Arizona State University.

Galbraith, J. (1977). *Organization Design*. Reading, MA: Addison-Wesley.

Howard, R. (1992). "Heathens, Heretics and Cults: The Religious Spectrum of Decision Aiding", *Interfaces*, 22, pp 15-27.

Howard, R. and J. Matheson (eds.) (1983). *Readings on the Principles and Applications of Decision Analysis*, Decision Analysis, Strategic Decisions Group, Menlo Park, CA.

Janis, I. (1982.1). *Stress Attitudes, and Decisions: Selected Papers* New York, Praeger Publishers.

Janis, I. (1982.2). *Groupthink: Psychological Studies of Policy Decisions and Fiascoes*, Houghton Mifflin Company, 01 June.

Jin, Y., R. Levitt, T. Christiansen, and J. Kunz. (1995). "The Virtual Design Team: Modeling Organizational Behavior of Concurrent Design Teams," *International Journal of Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, Vol.9, No.2, (April) 145-158.

Jin, Y. and R. Levitt (1996). "The Virtual Design Team: A Computational Model of Project Organizations" *Computational and Mathematical Organization Theory*, 2(3): 171- 196.

Kunz, J., T. Christiansen, G. Cohen, Y. Jin, and R. Levitt (1998). "The Virtual Design Team: A Computational Simulation Model of Project Organizations," *Communications of the Association for Computing Machinery*, (November) pp.84-92.

Lave, C. and J. March (1975). *An Introduction to Models in the Social Sciences*. New York: Harpers and Row.

Levitt, R., J. Thomsen, T. Christiansen, J. Kunz, Y. Jin, and C. Nass (1999). "Simulating Project Work Processes and Organizations: Toward a Micro-Contingency Theory of Organizational Design," *Management Science 45 (11),* November, pp1479-1495.

Luce, R. D. and H. Raiffa (1957). *Games and Decisions: Introduction and Critical Survey*. New York: Wiley. (Reprinted New York: Dover, 1989).

Luce, R. and H. Raiffa (1990). "Utility Theory" In Moser, P.K. (Ed.) *Rationality in Action: Contemporary Approaches* (pp. 19-40) Cambridge University Press: New York, NY.

Marais, K., N. Dulac, and N. Leveson (2004). "Beyond Normal Accidents and High Reliability

Organizations: The Need for an Alternative Approach to Safety in Complex Systems" Working paper prepared at the Massachusetts Institute of Technology.

March, J. and Simon, H. (1958). *Organizations*. New York, John Wiley & Sons, Inc.

March, J. (1994). *A Primer on Decision Making: How Decisions Happen.* New York: Free Press.

Maule, A. and A. Edland (1997). "The Effects of Time Pressure on Human Judgment and Decision Making" in R. Ranyard, W. Crozier, and I. Svenson (Eds.) *Decision Making: Cognitive Models and Explanations* (pp. 189-204) New York: Routledge.

Meshkat, L., and R. Oberto (2004). "Towards a Systems Approach to Risk Considerations for Concurrent Design" Working paper prepared at the Jet Propulsion Laboratory, California Institute of Technology.

Milgrom, P, and J. Roberts (1992). *Economics, Organization & Management* Englewood Cliffs, N.J.: Prentice Hall.

Moder, J. and C. Phillips (1983). *Project Management with CPM, PERT and Precedence Programming* 2nd Edition.

Murphy, D. and M. E. Paté-Cornell (1996). "The SAM Framework: A Systems Analysis Approach to Modeling the Effects of Management on Human Behavior in Risk Analysis", *Risk Analysis*, Vol. 16, No. 4, pp.501-515.

Orasanu, J. (2003). "Organizational Risk Model" Working paper prepared at NASA Ames Research Center.

Paté-Cornell, M.E. (1990). "Organizational Aspects of Engineering System Safety: The Case of Offshore Platforms," *Science*, Vol. 250, November 1990, pp. 1210-1217.

Perrow, C. (1984). *Normal Accidents: Living with High-Risk Technologies* New York: Basic Books.

Perrow, C. (1994). "The Limits of Safety: The Enhancement of a Theory of Accidents" *Journal of Contingencies and Crisis Management*, 2, (4), 212-220.

Powell, W. and P. DiMaggio (Eds.) (1991). *The New Institutionalism in Organizational Analysis* Chicago: University of Chicago Press.

Roberts, K. (1990). "Managing High-Reliability Organizations" *California Management Review*, 32, (4), 101-113.

Scott, W. R. (1998). *Organizations: Rational, Natural, and Open Systems* 4th Edition New Jersey: Prentice-Hall.

Sherif, M., O. J. Harvey, B. J. White, W. Hood, and C. Sherif (1961). *Intergroup conflict and cooperation: the robbers cave experiment* Norman, OK University Book Exchange.

Simon, H. (1977). *The New Science of Management Decision* 3rd revised edition (1st edition 1960) Prentice-Hall, Englewood Cliffs, NJ.

Thompson, J. (1967). "Organizations in Action: Social Science Bases in Administrative Theory", McGraw-Hill, New York.

Tversky, A. and D. Kahneman (1974). "Judgment Under Uncertainty: Heuristics & Biases" in D. Kahneman, P. Slovic, &A. Tversky (Eds.) *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.

APPENDIX B:
PRA-VDT MODEL JUSTIFICATION THROUGH CASE STUDIES

JUSTIFICATION PROCESS

We offer three orthogonal and complementary research elements: *observations* of a radically accelerated project at JPL, formal yet intuitive *theories* that have face validity and offer a straightforward comparison with established social science theories, and *simulation results* that show the combined implications of foundational micro-theories on a project scale.

Our claims are based on simultaneously *validating* theories by comparing them with observations, *verifying* theories' consistent operationalization in a simulation model, and *calibrating* the results' implications against our initial and new observations. Our work is therefore explicitly grounded by consistencies among reality, intuition, and formalism.

VVC of VDT

In this section, we summarize the effectiveness of previous efforts to show that the VDT simulation accurately *predicts* project behaviors. There is a considerable body of published documentation of VDT VVC efforts, including an important but limited body of predictive application. VDT has had some striking success, notably the accurate prediction of schedule delay, and critical organizational faults in the cabling contractor for a Lockheed satellite. Nevertheless valid concerns remain, principally in the area of calibration.

Because VDT does not model mid-course corrections, for long projects it predicts risks of failure rather than failures themselves. Therefore, as is the case with PRA, VDT validation faces the challenge that relating probabilities and the statistics of small samples is not straightforward.

VVC of PRA

In this section, we outline common PRA validation methodologies and indicate where concerns lie most prominently among academics and practitioners. From a theoretical perspective, the "prior" probabilities that PRA elicits from experts are often the most

contentious. Practitioners are also often hesitant to employ PRA because it requires the allocation of scarce resources. For example, while a project manager may believe PRA to be useful, he or she may feel that it is unlikely to be worth the amount of distraction that it will cause some of the project's most busy experts (opportunity costs). In spite of these concerns, PRA/SAM applications of HORM have shown some notable successes, prominently including the prediction of organizational faults supporting the shuttle thermal protection system.

VDT CHANGES

The formulation involves running a case for each permutation of the expert assessed values. This is best executed after a reasonably calibrated baseline is established. It may also include preprocessing to produce a batch file, and overnight execution.

The results of VDT may need to be assessed at the level of each simulation output because otherwise it is not possible to use nonlinear combinations of outcome value measures. For example, risk aversion requires that assessment of result dispersion occur. It is not reasonable simply to assume a normal distribution on outcome measures because we are almost certainly dealing with chaotic phenomena.

Changing the VDT formulation, calibration variables or code implementation makes comparison with other theoretical results from VDT less viable, reducing the value of our results to the academic community. At the same time,

Ideally, we would like to minimize the changes to the behavior matrices, for similar reasons. However, certain of the VDT calibration measures in the matrix file may not be accurate for the aerospace application. Therefore, we will strive to document changes clearly and make the matrices available. This should include a clear description of the intuitive impact- the real-world interpretation. For example, the ICE matrix may have a five-minute communication link, versus a one-hour duration for a traditional project design.

JUSTIFICATION CASES

Within the limited scope of this research, I propose to select cases that justify this approach to a level that is within the reach of well-qualified and adequately resourced practitioners.

Table 2 lists 11 case studies that might contribute to the justification of an integrated PRA-VDT model. I associate with each project a level of effort and earliest start date, as well as measures of expected contribution to verification, validation, and calibration efforts. I also note the potential contribution to the subjects of Project Management (PM), Integrated Concurrent Engineering theory (ICE), and risk analysis (Risk).

Subject	Method	Effort	Start	Verification - Validation Calibration			PM	ICE	Risk
VDT-PRA Proposal	Toy Problem		A04	*	*				
Software v. Hardware	Intellective	**	A04	***	**	*	*		*
PRA v. VDT v. PRA-VDT	Intellective	**	A04	***	**	*	**		**
ICE v. Tradition	Intellective	**	A05	***	***	*	*	***	*
3 Management Projects	Gedanken	**	A04	**	***	**	**		
3 Risk Projects (Observe)	Gedanken	**	Sp05	**	***	***	**		**
Risk Station @ Team-X	Retrospective	*	A04	**	**	***	*	**	**
Team-X Project	Retrospective	*	A04	*	**	*	*	*	
Risk Station at JIMO	Natural History	***	A04	***	**	***	**	*	**
3 CIFE ICE Charettes	Prospective	**	W05	**	***	**	*	**	*
Risk Project (Intervene)	Prospective	**	Sp05	**	***	***	**		**

The Method column indicates the study procedure and potential contributions, as defined by Levitt and Burton []. We build confidence in models by progressing through a series of more refined stages:

1. Toy Problems provide intuitive demonstrations of complex models

2. Intellective studies explore theoretical phenomena on cognitively tractable problems

3. Gedanken experiments compare expert and model predictions for a real-world project

4. Retrospective studies calibrate model predictions against historic documentation

5. Natural History predicts outcomes, then compares them against emergent reality

6. Prospective studies predict, recommend, and intervene in order to benefit a project

SOFTWARE V. HARDWARE

One prominent decision in space mission design regards how much of an information system's functionality is instantiated in a spacecraft's hardware, and how much is coded in software. From a risk management perspective, software has the advantage that even

after a mission is underway, it is possible to alter the programming of software and therefore the behavior of the mission. This option is not available for hardware programming.

To formulate this problem using PRA/VDT, we model the two versions of the POPE structure. In the software version, software engineers are employed to a greater degree in development and operations. During operations, one or more decision points may arise at which, based on the observation of relevant mission uncertainties, the software may be altered in order to adapt. In the hardware case, the appropriate hardware engineering team substitutes for software during design and development, and the decision points are all accelerated to the point at which hardware is printed.

RISK STATION @ TEAM-X, ICE V. TRADITION

This example examines the interactions between two considerations in engineering project design: the adoption of Integrated Concurrent Engineering (ICE) methods, and the incorporation of a risk analysis station. The following table offers some intuitive predictions that the formal model may verify.

Predictions:	Risk Station	No Risk Station
ICE Process	Fastest Project Risks: Low Functional Risks: Low	Fast Project Risks: Low Functional Risks: High
Traditional Process	Slowest Project Risks: High Functional Risks: Low	Slow Project Risks: High Functional Risks: High

VDT V. PRA V. VDT-PRA

Although the benefits are considerable, calculating the advantages of decision support tools is surprisingly uncommon today. The reason for this is sometimes technical- many analysts fail to recognize the opportunity to measure benefits, and those who do understand that the measurement can require considerable technical investment (including customization of the model for a specific implementation).

By simulating the example project in an integrated model, we are able to calculate the benefits of each modeling system for this application. We may apply this methodology to test the commonly held belief that for all but very routine and low-margin projects, early-stage VDT's benefits outweigh its costs. We may also determine whether for all

but routine and reliable products, early-stage PRA's benefits outweigh its costs. We are also able to test the conjecture that VDT-PRA is literally of greater benefit than the sum of its components.

For our goal of calculating the benefits of VDT, of PRA, of PRA-HORM, and of VDT-PRA, we model the choice among these as a global decision variable. Where any one of these is selected, additional actors are added to the project design in each phase. These actors generally burden project costs, schedule, stress and other outcome measures directly. However, they also offer benefits depending on the tools employed.

Adopting these techniques typically requires numerous essential and sometimes difficult process interventions. Previous work in this area (such as the APRAM model) defines the interventions as costs, whereas their impact may be through reduction of available, critical resources (such as decision-makers or domain experts). VDT captures this more complete view of the impact of modeling itself.

Modeling a PRA Chair

Team-X currently uses the Failure Mode Effects Analysis method of ordinal, qualitative measurements of engineering risk likelihoods and consequences. Qualitative methods, however, are not sufficiently precise to serve beyond a transitional stage. Probabilistic Risk Assessment (PRA) offers a far more precise, theoretically grounded, and diverse toolset for these applications. We accept that JPL may find it appropriate to implement this change "under the hood" in a second transitional stage if the current human or system interfaces require a gradual transition to PRA's increased complexity (VDT also uses qualitative surface interfaces with an accessible quantitative underlayment).

Integrating a PRA user role into the project plan can measure the costs of this decision support tool. As a benefit, decisions during the project may be made with an improved understanding of the risks that are being assessed. For example, if the PRA chair assesses and mitigates thermal system risks, outcomes may be improved and measured more accurately. Comparing the results with versus without the chair can produce a measure of benefits.

A PRA chair serves two functions. The first function of a risk chair is to initiate and support risk mitigation. In the model, this activates the utility function's product risk

tree's "Risk Analysis" and "Contingency Design" nodes. The nodes represent the initiative to design and implement contingencies, based on an accurate risk analysis.

A risk chair also consolidates and assesses the as-planned risk involved in each element of a project. This is the probability of failure assuming that the designs are consistent and properly executed. In the model, the station failure probabilities are global uncertainties that provide and upper bound on the product risk component of the objective function. Having a PRA chair provides this information to the decision-makers in the second and subsequent gates. This improves their ability to make an informed decision on whether to abort or continue the project.

Modeling a PRA-HORM Chair

With a VDT-PRA chair, decision-makers are aware of the internal, engineering risks as well as the human and organizational risks associated with a project. Thus, in addition to the engineering mitigations that a PRA chair offers, projects using PRA-HORM can take advantage of local and global alternatives that mitigate project design flaws. For PRA-HORM, these flaws are limited to the management of product risk-influencing elements, which we model as hidden work.

Modeling a VDT Chair

Practitioners who are considering the use of these tools may wish to determine in advance the costs and benefits that VDT offers for their own program. ntegrating a VDT user role into the project plan can measure the costs of this decision support tool. PRA/DA tells us that VDT only offers value where it has the potential to impact decisions. When evaluating mid-project decisions, we can measure the implications of having VDT available by comparing the expected utility with, versus without an accurate estimate of human and organizational project performance measures. VDT benefits the program in the ways that the VDT/PRA system shows it to, by improving choices through improving planners' predictive power. These costs and benefits are complex, and specific to a project and set of available decisions, but measurable.

The Team-X programmatics, cost estimation, and other design procedures may generate a range of pertinent data, but overlook its value for human and organizational risk assessment. A high fidelity model such as VDT may enable Team-X to extend its

holistic and robust project designs to this domain. From an academic perspective, VDT has shown itself amenable, but has not yet been proven capable of supporting an environment as dynamic and information-rich as Team-X. VDT's cost and schedule measures also have not been compared and calibrated against tools as sophisticated as those that Team-X may employ.

With a VDT Chair, decision makers in the simulation are alerted to the organizational and process implications of each alternative design. The team can detect likely cost overruns, schedule overruns, and excessive organizational strain. In the presence of alternatives that can help improve the situation, the VDT chair will ensure that the benefits of these alternatives are properly calculated. The VDT chair does not provide information about the product risks (pending reconsideration).

Modeling a VDT-PRA-HORM Chair

Integrating a VDT/PRA user role into the project plan can measure the costs of this decision support tool. As a benefit, decision-makers can predict project results accurately and decide appropriately, given any specific bounds on their information (such as ignorance of uncertain values).

With a VDT-PRA chair, decision-makers are able to pinpoint the precise value of the utility function for each prospect. This is the ideal condition of "rational" decision-making.

APPENDIX B:
SAMPLE REPORTS

These tables illustrate the utility metric calculation on a common project structure in which stages of activity are punctuated by local decision-making gates. The project involves an initial decision of whether to accelerate the project, followed by the revelation of a task complexity measure, followed by the choice of whether to involve a risk management actor.

Preferences

Project Duration	Labor Cost	Process	Average Backlog	Decision Delay	Rework Volume	InfoExch Delay	Organization	Product Risk	Product	Total
16.5%	16.5%	33.0%	8.3%	8.3%	8.3%	8.3%	33.0%	34.0%	34.0%	100.0%

Risk Baseline

Gate1		Gate2	Stage1							
Tools	Complexity	Decision	Seed	Direct	Rework	Quick Fix	Ignore	Default	Full	Empty
PRA	High	Go	1	0.99	0.99	0.96	0.93	0.9	0.95	0.85
PRA	High	No-go	2	0.99	0.99	0.96	0.93	0.9	0.95	0.85
PRA	Medium	Go	1	0.98	0.98	0.92	0.86	0.8	0.9	0.7
PRA	Medium	No-go	2	0.98	0.98	0.92	0.86	0.8	0.9	0.7
None	High	Go	1	0.99	0.99	0.96	0.93	0.9	0.95	0.85

None	High	No-go	2	0.99	0.99	0.96	0.93	0.9	0.95	0.85
None	Medium	Go	1	0.98	0.98	0.92	0.86	0.8	0.9	0.7
None	Medium	No-go	2	0.98	0.98	0.92	0.86	0.8	0.9	0.7

Task (Function) - Interface (Project) Risk

Gate1		Gate2	Stage 1		Direct	Rework	Quick Fix	Ignore	Default	Risk
Tools	Complexity	Decision	Seed	Risk						
PRA	High	Go	1	Mission	20	4	3	3	0	0.981
PRA	High	Go	1	H/W	20	2	2	4	2	0.974
PRA	High	Go	1	Electrical	20	0	1	5	4	0.934
PRA	High	Go	1	M-H	0	4	3	3	0	0.926
PRA	High	Go	1	H-E	0	2	2	4	2	0.942
PRA	High	Go	1	M-E	0	0	1	5	4	0.921
PRA	High	No-go	2	2	20	4	3	3	0	0.962

System (Meeting) Risk

Gate1		Gate2	Stage1	Baseline		Average		Fraction	System
Tools	Complexity	Decision	Seed	Attend	Absent	Attend	Absent	Attend	Risk
PRA	High	Go	1	0.95	0.85	9	1	0.9	0.940
PRA	High	No-go	2	0.95	0.85	8	2	0.8	0.930
PRA	Medium	Go	1	0.9	0.7	7	3	0.7	0.840
PRA	Medium	No-go	2	0.9	0.7	6	4	0.6	0.820
None	High	Go	1	0.95	0.85	5	5	0.5	0.900
None	High	No-go	2	0.95	0.85	4	6	0.4	0.890
None	Medium	Go	1	0.9	0.7	3	7	0.3	0.760

None	Medium	No-go	2	0.9	0.7	2	8	0.2	0.740

Product DUtility

Gate1		Gate2	Stage1	Task Risk			Interface Risk			System Risk	Product Risk
Tools	Complexity	Decision	Seed	**Mission**	**H/W**	**Electrical**	**M-H**	**H-E**	**M-E**	**All**	
PRA	High	Go	1	1.0	1.0	0.9	0.9	0.9	0.9	0.94	0.674
PRA	High	No-go	2							0.93	
PRA	Medium	Go	1							0.84	
PRA	Medium	No-go	2							0.82	
None	High	Go	1							0.9	
None	High	No-go	2							0.89	
None	Medium	Go	1							0.76	
None	Medium	No-go	2							0.74	

Process DUtility

Gate1		Gate2	Stage 1	Project Duration	Labor Cost	Process DUtility
Tools	Complexity	Decision	Seed			
PRA	High	1	Go	2	4	3.0
PRA	High	2	No-go	4	8	6.0
PRA	Medium	1	Go	6	12	9.0
PRA	Medium	2	No-go	8	16	12.0
None	High	1	Go	10	20	15.0
None	High	2	No-go	12	24	18.0
None	Medium	1	Go	14	28	21.0
None	Medium	2	No-go	16	32	24.0

Organization DUtility

| Gate1 | | Gate2 | Stage1 | Average | Decision | Rework | InfoExch | Organization |
Tools	Complexity	Decision	Seed	Backlog	Delay	Volume	Delay	DUtility
PRA	High	Go	1	1	2	3	4	2.5
PRA	High	No-go	2	2	3	4	5	3.5
PRA	Medium	Go	1	3	4	5	6	4.5
PRA	Medium	No-go	2	4	5	6	7	5.5
None	High	Go	1	5	6	7	8	6.5
None	High	No-go	2	6	7	8	9	7.5
None	Medium	Go	1	7	8	9	10	8.5
None	Medium	No-go	2	8	9	10	11	9.5

Outcome

| Gate1 | | Gate2 | Stage1 | | Process | Product | Organization | |
Tools	Complexity	Decision	Seed	%	DUtility	DUtility	DUtility	Total DUtility
PRA	High	Go	1	25%	3.0	0.7	2.5	2.1
PRA	High	No-go	2	25%	6.0	0.0	3.5	3.2
PRA	Medium	Go	1	25%	9.0	0.0	4.5	4.5
PRA	Medium	No-go	2	25%	12.0	0.0	5.5	5.8
None	High	Go	1	25%	15.0	0.0	6.5	7.2
None	High	No-go	2	25%	18.0	0.0	7.5	8.5
None	Medium	Go	1	25%	21.0	0.0	8.5	9.8
None	Medium	No-go	2	25%	24.0	0.0	9.5	11.2

Prospect

Gate1		Gate2		Expected Product DUtility	Expected Org. DUtility	Expected Process DUtility	Expected Total DUtility
Tools	Complexity	Decision	%				
PRA	High	Go	50%	0.7	2.5	3.0	6.2
PRA	High	No-go	50%	0.0	3.5	6.0	9.5
PRA	Medium	Go	50%	0.0	4.5	9.0	13.5
PRA	Medium	No-go	50%	0.0	5.5	12.0	17.5
PRA	High	**Go**		0.7	2.5	3.0	6.2
PRA	Medium	**Go**		0.0	4.5	9.0	13.5
PRA				0.3	3.5	6.0	9.8
None	High	Go	50%	0.0	6.5	15.0	21.5
None	High	No-go	50%	0.0	7.5	18.0	25.5
None	Medium	Go	50%	0.0	8.5	21.0	29.5
None	Medium	No-go	50%	0.0	9.5	24.0	33.5
None	High	**Go**		0.0	6.5	15.0	21.5
None	Medium	**Go**		0.0	8.5	21.0	29.5
None				0.0	7.5	18.0	25.5
PRA				0.3	3.5	6.0	9.8

Decision

Gate1		Gate2	Expected Product DUtility	Expected Org. DUtility	Expected Process DUtility	Expected Total DUtility
Tools	Complexity	Decision				
PRA	High	Go	6.0	9.8	3.5	6.4
PRA	High	No-go	15.0	21.5	6.5	14.3

PRA	Medium	Go	18.0	25.5	7.5	17.0
PRA	Medium	No-go	21.0	29.5	8.5	19.7
PRA	High	-				
PRA	Medium	-				
PRA	-	-				
None	High	Go	24.0	33.5	9.5	22.3
None	High	No-go	15.0	21.5	6.5	14.3
None	Medium	Go	18.0	25.5	7.5	17.0
None	Medium	No-go	6.0	9.8	3.5	6.4
None	High	-				
None	Medium	-				
None	-	-				

The above table defines the optimal policy. The optimal global decisions are known with certainty based on a calculation of all possibilities using Decision Analysis. Decisions that impact the first stage only are made after global uncertainties become known, and therefore may be contingent upon them. When the global utility function integrates each stage's simulation outcomes in a linear fashion, the optimal strategy at each stage is determined only by global uncertainties. When the stages' outcomes interact, for example by combining process schedule durations to meet a target deadline, decisions in later stages may depend on earlier stages' outcomes. For this reason, local "gate" adaptations are contingent and not so readily explained. For this case, we explain the probabilities of each combination of strategies.