RATIONALE FOR CHANGING GOALS WHILE REPLANNING

Paul D. Collopy, DFM Consulting, Urbana, Illinois

Abstract

In a complex and uncertain world, plans are frequently changed in the course of execution. The change may be an alteration of the course of action, or the goal of the plan may be modified. Previously, goals have been treated as immutable, or changes to goals have been regarded as irrational or beyond the scope of rational behavior. This paper shows that goal changing can be a rational alternative in a decision to modify a plan during execution. The likelihood that replanning may occur has a major impact on the expected cost and benefit of plans, and thus on the evaluation of alternatives when making planning decisions. In particular, the decision to replan is modeled in terms of costs and benefits, leading to a quantitative discussion of when goal change is a preferred alternative. The dynamic evolution of goals and courses of action during plan execution is analyzed.

Introduction

The formal treatment of planning tends to look upon plans as static structures which are built and then executed. In fact, in day-to-day life plans are dynamic structures, constantly changing and reforming actions and goals. While most planning research focuses on planning from scratch, most planned actions in the world result from modifications to plans, plans that may have been changed again and again, to the extent that the original plan is impossible to identify.

Planners who assume that their plans will be executed as design can make serious errors not only in estimating the resources required to execute the plan, but also in estimating the results that will be achieved.

The purpose of this paper is to launch a formal discussion of replanning, the changing of plans, as a rational, analyzable process that is central to human action. The hope is to drive toward a normative theory of replanning and lay the groundwork for practical tools that can advance the state of the art in areas such as technical risk management and cost estimation.

Formalization of Planning and Replanning

Following are a group of definitions that formalize elements of planning for the purpose of discussion.

A **plan** is a course of action intended to achieve a goal. A specific plan can be thought of as an ordered pair as follows:

$$\text{Plan} \equiv (\text{Course of Action, Goal}) \qquad (1)$$

A **course of action** is a partially ordered set of actions, per Sacerdoti (1977). Modeling the world as a series of states (Genesereth and Nilsson, 1987, p. 263), an **action** causes the world to change from the current state to a new state, depending on the current state. To allow for uncertainty, this traditional definition of action is expanded in two ways:

1. An action causes the world to change to a probabilistic distribution of new states, rather than a certain new state.
2. The world may change to a new state without an action (due to outside causes).

To **execute** a course of action (or to execute a plan) is to cause the member actions to occur in a proper order. A **goal** is "a set of desired properties" that may be possessed by a state (Genesereth and Nilsson, 1987, p. 263). A goal is **achieved** when the world reaches a state possessing all the properties of the goal. A plan **failure** occurs when the course of action is completely executed and the goal is not achieved.

For the purposes of this paper, the **planner** is a single actor who creates the plan and executes the course of action. However, the essential points made here can be generalized to situations where a plan is created and executed by multiple actors. The planner **replans** when she changes the plan (the course of action, goal, or both).

Prior Work on Replanning

Planning is a multidisciplinary field, addressed by scholars in public policy, economics, military science, and others. However, two particular schools of thought have addressed the issue of replanning, with radically different focuses and conclusions. One is the field of artificial intelligence, which studies planning as a formal, rational process. The other school includes sociologists and psychologists, who study planning as a human phenomenon. Each will be reviewed in turn.

Artificial Intelligence Views of Replanning

Researchers in artificial intelligence have extensively investigated planning as a complex human activity, which, in some instances, can be formalized and duplicated on machines. In doing so, they have created a normative theory of planning.

The central problem has been to construct a plan, given a goal and a set of allowed actions. Replanning has typically been a side issue, but researches have addressed it. For example, Fikes et al. (1990) developed a planning system that checked flags occasionally to sense whether the plan was proceeding properly. A negative setting of a flag during execution was interpreted as failure of the plan, in which case the planning system constructed a new course of action. Dean and Wellman (1991, p. 14) discuss replanning from the same point of view: "There are ... potential disadvantages in generating sequences of actions. The most obvious disadvantage is that ... the sequence of actions will fail to have the desired effect. Unless the [planner] is really convinced of the accuracy of its model, it will want to check that the plan is proceeding according to its expectations. This checking is referred to as *execution monitoring*." (Authors' emphasis.) Shapiro (1991) discusses a robotic system with an extreme form of execution monitoring: at each time step during plan execution, the robot checks whether the goal has been achieved; if it has not, the current situation is compared to the goal situation, and an action is taken to correct at least part of the discrepancy.

Under one approach to planning, plan modification is central. This is case-based planning, described in Hammond (1989). The planner creates new plans by searching among previously used plans (cases) for the one most similar in goal and context. This best match case, when it is modified to meet the current requirements, becomes the new plan. Case-based planning in Hammond's view does not include replanning during execution, but the approach would seem to accommodate this readily.

Artificial intelligence researchers are uniform in treating goals as immutable and external to the planning problem. Plan failure is usually ignored (that is, plans are required to be failure-proof), or failure is treated as an exceptional condition. Remedies for plan failure have been limited to modifying the course of action, although a considerable literature has developed on techniques for such modifications.

Psychological and Sociological Perspectives

Logically constructing courses of action to achieve goals is the paradigm of rational behavior. In Robert Nozick's philosophical treatise, *The Nature of Rationality*, rational behavior is defined in terms of goal achievement:

> "The beliefs and actions of a rational person are ones generated (and maintained) by a process that reliably achieves certain goals, where reasons play some appropriate role in guiding this process." (Nozick, 1993, p. 177)

However, observers of human and organizational behavior describe a different approach to replanning, goal modification, which does not fit neatly within this notion of rationality. Cyert and March (1959) describe goals as "aspiration levels", and observe that, when an organization fails to meet a goal such as an expense budget, the goal is modified to make it more achievable. Even less rational, at first glance, is the psychological theory of cognitive dissonance, developed by Leon Festinger and reviewed by Aronson:

> "The more a person is committed to a course of action, the more resistant he will be to information that threatens that course." (Aronson, 1973, p. 135)

If a course of action is planned to achieve goal A, but well into execution the planner realizes that the course is leading to state x which does not possess the properties of goal A, the planner may adopt a completely different goal that includes the properties of x. Karl Weick actually asserts that goals normally arise from courses of action:

> "Actions *precede* goal definition ... The common assertion that goal consensus must occur prior to action obscures the fact that consensus is impossible unless there is something tangible around which it can occur. And this 'something tangible' may well turn out to be actions *already completed*. Thus it is entirely possible that goal statements are retrospective rather than prospective." (Weick, 1979, p. 18, author's emphasis)

Thus, people and organizations are observed to create goals based on courses of action, an approach to plan creation completely opposite to that studied in the artificial intelligence community. And, whereas the construction of courses of action to achieve stated goals is the paradigm of rational behavior, retrospective goal creation to fit actions appears to be patently irrational.

Replanning as a Decision

There is a context within which goal change can be understood as a rational process: the decision of a planner to replan. Decision theory provides a normative model for rational decision, and, in the replanning decision, alternatives can include the modification or outright replacement of goals.

A decision is "an irrevocable allocation of resources, it the sense that it would take additional resources, perhaps prohibitive in amount, to change the allocation." (Howard, 1983, p. 23.) The notion of almost-irrevocable commitment is just right for replanning. A planner commits to a course of action as he embarks on it. Allocation of resources to the course of action increases as individual actions are executed. And yet, if the planner perceives that the plan is at risk of failure (that is, the course of action may complete without achieving the goal), the allocation can be revoked and a new allocation made. This new allocation is the replanning decision.

Howard provides a framework for formalizing decisions, the decision basis: "The core of this process is the elicitation or synthesis of the decision basis. ... The basis has three parts: the choices of alternatives the decision-maker faces, the information that is relevant, and the preferences of the decision-maker." (Howard, 1988, p. 681) The replanning decision can be formalized by identifying the alternatives, information, and preferences of the planner.

Alternatives

Using the concept of plan defined above, the planner has two alternatives: he can continue with the current plan, or change the plan. Although there could be many alternative plans (pairs consisting of a goal and a course of action), the discussion below will consider the case where a single alternative is considered, in which either the course of action, the goal, or both are different from the original plan. Generalization to multiple alternatives is straightforward. The construction of the alternative course of action is presumed to be by methods within the state of the art, and is not addressed here.

Information

The relevant information available to the planner includes the following (for the original plan x or any conceived alternative x):

$P_{f\,x}$ = probability of failure of the plan

C_x = estimated cost of the course of action, as planned

s_x = **sunk cost**, the estimated cost of the portion of the course of action already committed to irrevocably

B_x = estimated benefit of the goal, as planned

$\tau_{x,y}$ = **transaction cost**, which is the estimated cost of computation, communication, and persuasion required to change from plan x to plan y (persuasion represents agency costs as described by Jensen and Meckling, 1976)

For purposes of discussion in this paper, C and B are assumed to be quantified in commensurable units (for example, dollars). If they are not, multi-attribute techniques (Keeney and Raiffa, 1976) can be employed to evaluate alternatives. Moreover, C and B are presumed to be net present values, accounting for all time preferences.

Preferences

The planner wishes to maximize benefits and minimize costs. In fact, when it is assumed that costs and benefits are commensurable, the planner wishes to maximize the expected benefits-minus-costs. However, note that whenever $P_f > 0$, the expected costs and benefits (at the end of plan execution) are not equal to the planned costs and benefits (where planning ignores the plan failures and replans). The term **plan value** expresses the expected outcome of a plan in terms of the planner's preferences:

$$V = E[B' - C'] \qquad (2)$$

where the prime indicates the inclusion of changes to the course of action and goal which may occur through replanning.

Evaluating the Replanning Decision

Evaluating the replanning decision comes down to quantifying the expected net value of the modified plan versus the prospect of continuing with the status quo. If the plans are certain to achieve their goals without further replanning, the evaluation, shown in Figure 1, centers on a comparison of the costs of the course of action and the benefits of goals. Note also the term t_1, which addresses sunk costs incurred in the execution of the original plan, and so are included in both prospects; and the transaction cost τ_1, for the direct cost of *changing* from plan 1 to plan 2.

To change from original plan 1 to plan 2 incurs the following loss of value:

$$V_1 - V_2 = R_{1,2} = B_1 - B_2 - (C_1 - s_1 - C_2) + \tau_{1,2} \quad (3)$$

The variable $R_{x,y}$ is the **cost of replanning** from plan x to plan y under certainty.

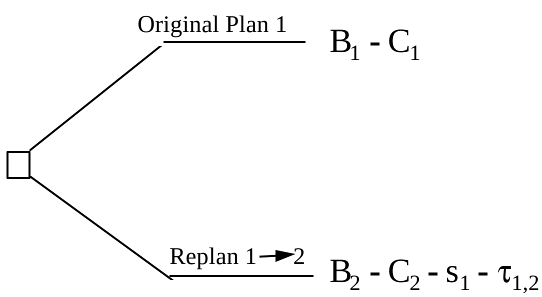

Replan when $R = (B_1 - B_2) - (C_1 - s_1 - C_2) - \tau_{1,2} < 0$

Figure 1: Decision Tree with No Uncertainty

Replanning under Uncertainty

Typically, however, when one replan is being considered, more replans may be possible downstream. Figure 2 shows a decision tree for a base plan with alternative replans, allowing for possible additional replans.

$P_{r\,x}$ in Figure 2 (and in the following discussion) is the probability that plan x will be replanned.

Rather than separately assess the many probabilities and costs in Figure 2, planners may assess alternatives they are directly considering ($P_{r\,1}$, $P_{r\,2}$, $R_{1,2}$, $P_{r\,3}$, $R_{1,3}$, and so on), plus parameters of generic future replans, for example $P_{r\,0}$ and R_0. The reason is that the planner may not have any information that distinguishes one of these future replans from another. Figure 3 illustrates

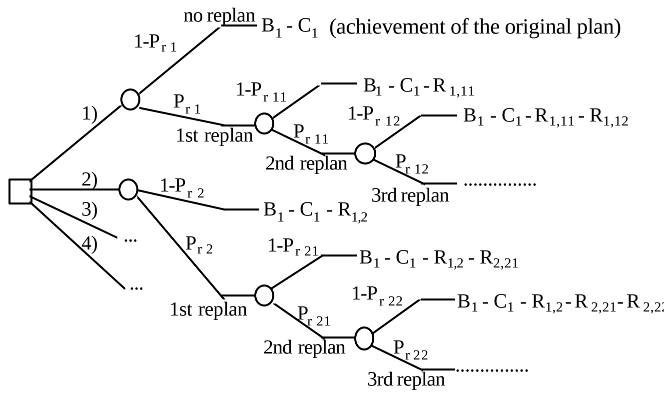

Figure 2: Decision Tree for Replanning under Uncertainty

the replanning decision tree using generic parameters for future replans.

In Figure 3,

$$V_1 = B_1 - C_1 - P_{r1}(R_0 + P_{r0}P_0 + P_{r0}^2 R_0 \cdots) \quad (4)$$

or

$$V_1 = B_1 - C_1 - P_{r1}\frac{R_0}{1 - P_{r0}} \quad (5)$$

The importance of R_0 in this formulation will now be addressed.

R^+ and R^- Planning Domains

Planning situations can be divided into two different domains depending on the expected value of R_0.

The first is the R^+ domain, where $E[R_0] > 0$. In the R^+ domain, value is inversely related to P_r (see equation 5). P_r is referred to as "risk" or "technical risk" with all its negative connotations, and an economic objective is to minimize or eliminate P_r. Most business and engineering organizations operate in R^+ domains, at least in terms of their day-to-day work.

The second is the R^- domain, where $E[R_0] < 0$. Here

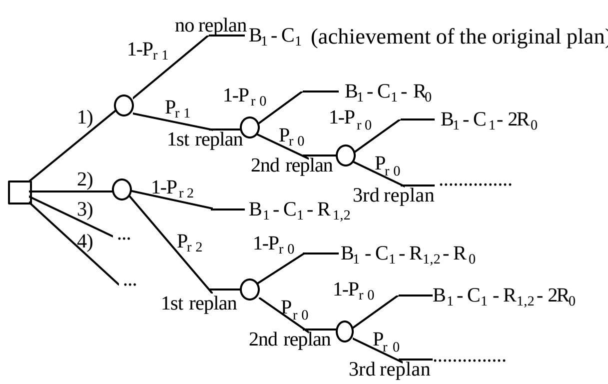

Figure 3: Decision Tree with Generic Assessments

replanning is good and high P_r is desirable. A pure research team or an innovative music group might fit in this category.

The rest of this paper focuses on R^+ domains, in part because this work was inspired by experiences in R^+ domains, and also because the mathematics fits more easily into closed form expressions. Notice that in R^- the optimal P_{r0} for equation (5) is 1, which is an unbounded solution.

<u>Comparison of Expected Outcomes</u>

Returning to Figure 3,

$$V_2 = B_1 - C_1 - R_{1,2} - P_{r2}\frac{R_0}{1-P_{r0}} \qquad (6)$$

The difference in the expected outcomes of alternatives 1 and 2 is then

$$V_1 - V_2 = R_{1,2} - (P_{r1} - P_{r2})\frac{R_0}{1-P_{r0}} \qquad (7)$$

Therefore, the analysis suggests the planner replan when

$$(P_{r1} - P_{r2})\frac{R_0}{1-P_{r0}} > R_{1,2} \qquad (8)$$

or, re-expanding $R_{1,2}$:

$$(P_{r1} - P_{r2})\frac{R_0}{1-P_{r0}} > B_1 - B_2 - (C_1 - s_1 - C_2) + \tau_{1,2} \qquad (9)$$

In other words, the planner replans when the resulting decrease in the probability of replanning (risk reduction) outweighs the cost of replanning (which is positive in R^+).

While there are many ways of assessing the quantities expressed here as R_0 and Pr_0, they will all lead to an inequality of the same form as (9), in which the cost of changing plans (benefit change, course of action change, and transaction cost) is traded against a change in risk (change in P_r weighted with the cost of future replans). In R^+, typically, a replan chooses an inferior plan in order to obtain a reduction in risk. (See "Serendipity and the creation of novel goals" below for situations where this is not the case.)

In Formula (9), the tradeoff between changing a goal versus changing a course of action is clear. s_1, the sunk cost, is constant. Thus, the essence of comparing two options a and b, where a modifies the goal and b modifies the course of action, is comparing the two costs $B_1 - B_a$ and $C_b - C_1$. If

$$B_1 - B_a < C_b - C_1 \qquad (10)$$

then goal change is a better option than the traditional reconstruction of the course of action. Note, however, that the transaction costs $\tau_{1,a}$ and $\tau_{1,b}$ may be quite different and should be considered as follows:

$$(B_1 - B_a) + \tau_{1,a} < (C_b - C_1) + \tau_{1,b} \qquad (11)$$

Also, generally a new plan can include a change of goal and a change of course of action. The point is to illustrate the trade of goal change versus action change to emphasize that this treatment of replanning allows for the rational selection of goals as a normal part of the planning process.

<u>Probability of Failure versus Probability of Replanning</u>

In the previous analysis, the choice of how to replan is highly dependent on P_r, the probability of replanning. P_r can be difficult for planners to assess, in part because it is influenced by their own replanning behavior, and therefore potentially by the assessment itself. However, in many cases the probability of failure, P_f, is a good or very good surrogate for P_r.

Recall that P_f is the probability that, if the course of action is completed as currently planned, the current goal will not be achieved. At any point during plan execution, P_f can be reassessed, which leads to the notion of a stochastic process, $P_f(t)$, the assessment of P_f at any point in time.

Consider the case when a planner allows his course of action to complete when $P_f(t)$ is one (certain failure). Then the goal of his plan is not achieved. He has reached some state x, and the properties of x are then his defacto goal, g_d. This is in fact a replan, in which the course of action has not changed, the transaction cost is zero, but the original goal g_1 has been changed to g_d.

If all plan failures are replans, then the probability of replanning can never be less than the probability of failure.

$$P_r(t) \geq P_f(t), \text{ for all } t \qquad (12)$$

This inequality suggests that all approaches to replanning can be divided into two types:

1. Stubborn strategies, where the planner never replans unless $P_f(t) = 1$. This makes the certain prospect of plan failure a precondition to replanning, so that $P_r(t)$ can never be greater than $P_f(t)$. Combined with inequality (12), this means for stubborn strategies

$$P_r(t) = P_f(t), \text{ for all } t \qquad (13)$$

2. Non-stubborn strategies, where

$$P_r(t) > P_f(t)\text{, for some } t \qquad (14)$$

In order for (14) to hold, there must be some possibility that the planner will replan when failure is not yet certain, that is, when $P_f(t) < 1$. Otherwise, (13) would hold.

To understand this rigorously, observe that $P_f(t)$ and $P_r(t)$ are both martingales over the closed interval [0,1]. This means that, at any particular time t_0, the estimate of any future value

$$E[P_f(t_1)] = t_0\text{, for all } t_1 > t_0 \qquad (15)$$

and similarly for $P_r(t)$.

An immediate result is that the states $P_f(t) = 1$ and $P_r(t) = 1$ are absorbing states, since they are at endpoints of the domain. (For a concise discussion of this property of martingales, see Hoel et al., 1972, pp. 27 - 29.) What this means is that once a planner has assessed that her plan is certain to fail, she will always assess that it is certain to fail from that point forward, and likewise an assessment that a plan is certain to be replanned is locked in. This is analogous to a weather report on Monday that declares a 100% chance of rain on Wednesday. It would then be unreasonable for the same forecaster to say on Tuesday that the chance of rain on Wednesday is only 80%.

Thus,

$$P_f(t_0) = P(P_f(t_1) = 1)\text{, for any } t_1 \geq t_0 \qquad (16)$$

Under the stubborn strategy, $P_f(t) = 1$ is a prerequisite for replanning at or after time t. Furthermore a replan is an inevitable consequence of $P_f(t) = 1$, since the plan will certainly be replanned or fail, and failure is a form of replanning. Therefore,

$$P_r(t_0) = P(P_f(t_1) = 1) = P_f(t_0) \forall t_0, t_1 > t_0 \qquad (17)$$

Thus, stubborn strategies minimize P_r, because $P_r = P_f$, and P_f is the lower bound for P_r. It is observed in Collopy (1998, Chapter 4) that for many formulations of plan value (including (6) and (7) above in R^+), value decreases monotonically with P_r. Thus, value is maximized by a strategy that minimizes P_r, which is the stubborn strategy. Collopy (1998, Appendix B) describes an R^+ case where the stubbornness strategy is not optimal, but even in this case the optimal strategy yields a P_r that is very close to P_f.

The result is that P_f will often be a good or very good surrogate for P_r in evaluating decisions with formulae such as (9) above. Using P_f is an advantage because it is generally easier for the planner to assess than P_r.

Triggers for Replanning

A plan is a commitment to action. Once committed, the planner is relieved of the burden of deciding what to do next -- simply execute the plan until it is complete. Bratman (1987) cites three reasons for planning, that is committing to a set of actions with one decision, as opposed to making separate decisions to execute each action. The reasons are

1. A reduction in decision-making effort (p. 28)
2. A reduction in the uncertainty of the planner's immediate future (pp. 17 - 18)
3. A reduction in the uncertainty of the planner's behavior as perceived by others who know the plan (pp. 17 - 18)

All three reasons depend on the planner remaining committed to the plan. What would induce a planner to revoke such a commitment, and undermine the value of planning in the first place? A planner should not replan because of a spontaneous reassessment of the consequences of the plan versus all alternative goals and courses of action, particularly in R^+. Reason 1 above says that the planner used a plan to avoid just such reassessments. According to Bratman, "We settle in advance on ... plans and tend to reconsider them only when faced with a problem." (1987, p. 28)

This suggests that some less general information from the environment triggers a replan. The stubborn strategy discussed above presents a good candidate for this trigger: the probability of failure of the plan. If the planner is consciously using the stubborn strategy, a more refined trigger can be used: information that indicates that the plan will certainly fail (the truth of the equation $P_f = 1$).

Monitoring probability of failure, or whether the plan is going to fail, may require far less effort than continually contemplating whether a better plan is possible, thus allowing for some replanning while still realizing most of the benefit of reason 1 above.

Trends in the Evolution of Plans under Uncertainty

Figure 4 shows a trend frequently observed in product development products. Over the course of time, the estimated cost to complete the project grows. The project manager hopes that she will at least get additional benefit for the additional money that is being spent. But she is disappointed -- concomitant with the growth in cost is an erosion in benefit.

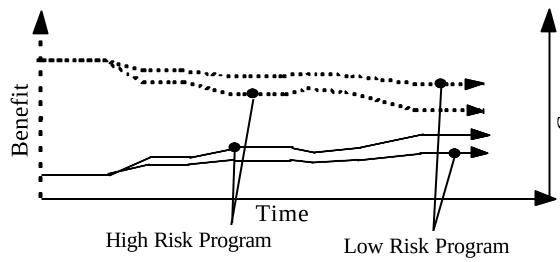

Figure 4: Benefit Erosion and Cost Growth

The phenomenon illustrated in Figure 4 is caused by an underestimation of risk (P_r), or, more often, by an inability to estimate rigorously in the face of uncertainty. Thus, given an original course of action of cost C_1 and an original goal of benefit B_1, the project estimator may have estimated that at the end of the project B_1 would be realized at cost C_1. Compare this to Formula (6).

In fact the current state of the art in cost estimating is to estimate C_1 and add a margin to it, with no formal procedure for determining the margin or relating it to risk. The risk management handbook of the US Department of Defense (Department of Defense, 1985) focuses entirely on eliminating risk, rather than planning in the face of risk. Furthermore, no method in the literature addresses anticipating benefit erosion as a consequence of risk.

Serendipity and the creation of novel goals

In R^+ domains, it is expected that a replan will result in a more expensive course of action, a less beneficial goal, or both. But that is not to say that this is always the case. Sometimes, a replanning decision can find an alternative plan that is a real improvement over the original plan. In R domains, a real improvement is expected. How can this happen? If the new plan is so good, why was it not chosen as the original plan?

There are three possibilities:

1. The new plan has a less expensive course of action or a more beneficial goal, but also has a higher P_r, which made it inferior to the first plan. If the P_r for the first plan rises enough during execution, this plan becomes attractive.

2. The new plan was not practical at the time the planner committed to the previous plan. Some event external to the plan may have happened (a technological discovery or a judicial ruling, for example) that has opened up the possibility of executing the new plan.

3. The earlier planning process failed to discover the new plan, but the current planning process, perhaps prompted by the actions already executed, found it. This has some similarity to the simulated annealing method of optimization, where a jump to a random spot sometimes provides a starting point for a superior search.

Reasons 2 and 3 are occasions of serendipity, where the planner is pleasantly surprised to discover new opportunities. These may include new goals, which could be closely related to the current goal, or could be quite novel. When a novel new goal arises under reason 3, that is, because planning benefited from the course of action parially executed, it is a true instance of the phenomenon cited above by Weick, of a goal arising from retrospective contemplation of actions executed. However, it is not as irrational as Weick suggests.

Conclusions

The study of replanning is in its infancy, but already it offers many promises.

Quantitative Treatment of Risk versus Cost

Using the analysis presented here, backed with empirical assessments, program managers can manage risk at a new level of sophistication, probably with significantly improved results. Risk reduction plans can be quanitatively measured in a cost-benefit framework. Risk reduction will be scaled back or eliminated when the cost of risk reduction exceeds the cost induced by the risk itself.

Quantitative treatment of risk-driven goal erosion

For the first time, this paper presents a quantitative treatment of the erosion of the benefits of goals in the face of risk. This can allow program managers to build margin into specifications quantitatively derived from the level of program risk, or to trade risk reduction plans against loss of system performance.

Theoretical foundation for goal evolution

For the first time, a rational basis is presented for the creation and evolution of plan goals, a process previously thought to be beyond the scope of rational treatment.

Normative theory for replanning

This paper presents the beginnings of a normative theory to guide decision makers through a clear, rational accomplishment of replanning decisions.

References

Aronson, Elliot. "The Rationalizing Animal". Pages 134-144 in *Readings in Managerial Psychology*, Fourth Edition. Edited by Harold J. Leavitt, Louis R. Pondy, and David M. Boje. The University of Chicago Press, Chicago, 1989. Article originally published in 1973.

Bratman, Michael. *Intentions, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA, 1987.

Collopy, Paul D. *Replanning in Organizations*. Dissertation at Stanford University, Stanford, CA, 1998.

Cyert, Richard M. and March, James G. "A Behavioral Theory of Organizational Objectives". Pages 76-90 in Mason Haire, ed., *Modern Organization Theory; A Symposium*. Wiley, NY, 1959.

Dean, Thomas L. and Wellman, Michael P. *Planning and Control*. Morgan Kaufmann Publishers, San Mateo, CA, 1991.

Department of Defense. *Transition from Development to Production ... Solving the Risk Equation*. DoD 4245.7M, Defense Technical Information Center, Cameron Station, Alexandria, VA, 1985.

Fikes, Richard E.; Hart, Peter E.; and Nilsson, Nils J. "Learning and Executing Generalized Robot Plans", pages 189-206 in *Readings in Planning*. Edited by James Allen, James Hendler, and Austin Tate. Morgan Kaufmann, San Mateo, 1990.

Genesereth, Michael R., and Nilsson, Nils J. *Logical Foundations of Artificial Intelligence*. Morgan Kaufmann, Los Altos, CA, 1987.

Hammond, Kristian J. *Case-Based Planning -- Viewing Planning as a Memory Task*. Academic Press, Boston, 1989.

Hoel, Paul G.; Port, Sidney C.; and Stone, Charles J. *Introduction to Stochastic Processes*. Waveland Press, Prospect Heights, IL, 1972.

Howard, Ronald A. "An Introduction to Decision Analysis." Pages 17-56 in *Readings on the Principals* and Applications of Decision Analysis, Vol. 1. Howard, R. A., and Matheson, J. E., editors. Strategic Decision Group, Menlo Park, CA, 1983.

Howard, Ronald A. "Decision Analysis: Practice and Promise". *Management Science*, Vol. 34 Issue 6, pp. 679-695, 1988.

Jensen, Michael and Meckling, William. "Theory of the Firm: Managerial Behavior, Agency Costs, and Ownership Structure". Pages 305-360 in *The Journal of financial Economics*, Vol. 3, 1976.

Nozick, Robert. *The Nature of Rationalism*. Princeton University Press, Princeton, NJ, 1993.

Sacerdoti, Earl D. *A Structure for Plans and Behavior*. Elsevier North-Holland, New York, NY, 1977.

Shapiro, Daniel. *Architectures for Semi-Autonomous Planning: Final Technical Report*. ADS TM 3249-08, Advanced Decision Systems, Mountain View, CA, 1991.

Weick, Karl E. *The Social Psychology of Organizing*. 2nd Edition, Addison-Wesley, Reading, MA, 1979.