A Quantum Approach to Risk Management

Author: Leif Hoglund

Given at the May, 2006 Risk Special Interest Group Meeting (Project Management Institute)

Abstract

There are plenty of tools to help program managers manage all of the data involved in developing and maintaining project plans. There are tools for maintaining task lists, resources, calendars and budgets, and tools that use Gantt charts, network diagrams, status reports and other tracking strategies. But there are no easily implemented tools for evaluating the plan itself, for determining whether plans are likely to be accomplished in the time allocated.
A new method is offered to quantitatively define the quality of project plans in terms of probable success, and to predict the impact of task slippage on planned milestones. The method is based on a model, where tasks and their duration are constructed as though made up of packet waves. The probability amplitudes of these waves are expected to be coherent at milestone dates, so that each task contributes to meeting a given milestone in a way that is directly related to its coherence.

The method makes use of quantum mechanics, and as is typical of quantum mechanics phenomena, the method is able to make predictions about outcome based on rather general input information about each individual task, and to the extent the model reflects actual project task behavior, the method becomes more accurate as the number of tasks increases.

Classical Approach and its Limitations

If the probability distribution is defined for each task, then the milestone probability may be also defined. For example, if the milestone includes two individual tasks with durations y_1 and y_2 and dispersions σ_1 and σ_2, respectively, then the dispersion of the milestone Σ will be defined by a relation:

$$\Sigma^2(y_1 + y_2) = \sigma_1^2 + \sigma_2^2 \pm 2\rho_{12}\sigma_1\sigma_2 \tag{1}$$

where ρ_{12} is the mutual correlation coefficient between tasks y_1 and y_2. For independent tasks, $\rho_{12}=0$. Similar relations may be applied to the milestones including more than two tasks.

If the probability distribution is known for each task, then the probability of performing any milestone and therefore the whole project may be calculated. Multi-path correlations between the tasks complicate the picture, but computer simulations can be used to manage this. In fact, several software packages exist that use task probabilities in an attempt to calculate overall project probabilities. The practical problem with this approach is that it is normally not possible to define task probabilities with the required accuracy. If a simple case is considered where a milestone is comprised of 100 independent tasks of equal duration *D*, and each duration is defined with 10% accuracy, then the milestone probability distribution has dispersion equal to the task duration:

$$\sqrt{\Sigma^2(y_1 + y_2 +..+ y_{100})} = 0.1D\sqrt{100} = D \tag{2}$$

To improve the resolution, each task must be further subdivided into sub-tasks, each sub-task also having similar (10%) duration tolerance. If the number of sub-tasks is 100, the milestone dispersion is decreased by $\sqrt{100} = 10$ times, according conventional rules for random processes. To improve milestone dispersion 10 times more, each sub-task should be further sub-divided into 100 parts etc. Thus, milestone dispersion improves very slowly, and the effort to define and detail tasks becomes higher than the task itself.

For large projects consisting of hundreds and thousands of tasks and dozens of milestones, it can be seen that even moderate dispersion of each task will result in very broad probability distributions for the milestone, and the conventional approach to determining milestone probabilities by detailing tasks is well known to become inadequate.

Milestone probability calculation – new approach

In quantum mechanics, the probability of a certain outcome can often be accurately calculated with very generic knowledge of some input events. In the case of a typical schedule, task durations are often not defined with very high accuracy. For example, a task duration of 1.01 week would be very unusual, and it would be more usual for a 2 week task to be defined with a variance no better than a day or more. This kind of imprecise knowledge of the task is handled quite naturally by quantum mechanics. In fact, a full characterization of event (task) behavior is considered impossible in principle, and events are characterized by a “wave function” ψ, or “probability amplitude”, so that conventional probability P is calculated according the formula $P=|\psi|^2$.

While the application of quantum mechanics may be mathematically pure and compelling, it is not always obvious how to interpret quantum logic. This is particularly the case if there is an attempt to always translate inputs and outputs in classical terms. For example, rather than think in terms of precise measurements, it is useful to consider a system as having certain “properties” which do not need to be exactly measurable to be useful. In fact, in the case of task scheduling, we find many non-physical “forces” acting on the system that cannot be easily accounted for by a classical approach.

Quantum mechanics actually refers to particles that occupy energy states within attracting potential wells or move across repulsive barriers, where each state is characterized by its wave function. While tasks are not quantum particles, they have several characteristics that suggest that quantum modeling could be useful. The most crucial characteristic is the observation that human tasks are not performed at a constant rate over time, but instead are performed in a way that can be modeled using linear combinations of harmonics. Task length and duration are affected by the “forces” of funding profile, by whether there are resources available, the diligence or training of the works, as well as by the milestone dates. These factors affect the outcome, but are often difficult to measure precisely and are therefore often more easily accounted for using a quantum approach.

One of the fundamental quantum mechanical features is interference of amplitudes. Namely, if the system wave function is $\psi = \psi_1 + \psi_2 + \psi_3 + ...$ then the probability P is

$$P = |\psi|^2 = |\psi_1|^2 + |\psi_2|^2 + |\psi_3|^2 + 2\psi_1\psi_2 + 2\psi_1\psi_3 + ... \qquad (3)$$

and mutual coherence between the individual wave functions ψ_i directly defines the output (3). In our model, on-schedule tasks are assumed to be coherent (superimposed in-phase) with the planned milestones. Task slippages will result in phase shift, loss of coherence and therefore a reduction of probability (3). The power of the model is in its ability to define milestone parameters with an accuracy exceeding the variance (dispersion) of the individual tasks. The intent is to allow easy comparison of different plans or task schedules.

The quantum modeling described here provides answers to the question: “If particular tasks slip from the original plan, what impact will such slips have on meeting the originally planned milestone dates?” But if the question is asked “What is the probability that a particular task will slip?” the model does not answer it.

Software Implementation of Quantum Approach

The approach described above has been refined for direct practical application to any project plan.

Obviously, any effort to better understand the tasks involved in a project will always yield rewards, but there is a great deal of information about project planning that can be gleaned from only looking at what might be called the structural aspects of a project. For example, too many parallel paths in a given plan will lead to high management load; too many shared resources will lower the probability of completing certain tasks on time; tasks with many dependencies are riskier than those that do not have so many dependencies; tasks with finish dates very close to milestone dates are riskier than those finished earlier, and so on.

In conventional analysis of plans, the critical path is often used as one structural measure of a program. While critical path is a useful concept, the Quantum approach encompasses many more elements of plan structure, and is therefore able to offer much earlier warning and more detailed diagnostics to managers that plans are going awry.

The analysis output for a given plan includes 3 specific metrics, all of which contribute to the primary goal of determining the impact on milestone date. The metrics are provided for each milestone as follows:

a. Visibility
b. Probability (as a function of Task Tolerance)
c. Task Importance (rank-ordered in terms of ability to affect milestone)

When these metrics are within expected limits, we say the plan is well considered from the point of view of meeting the planned milestone dates.

a) Visibility

A simple example is shown in figure 1. Assuming the top and bottom figures use exactly the same resource over the same planned time schedule, it is useful to understand that from a management perspective the bottom figure is a better plan. This is because as a manager monitors the completion of each task, the bottom plan offers better "visibility". This is illustrated if we consider what happens if any task slips. In the top case, a manager would know the first task has slipped sometime after March. In early May, he has a second opportunity to assess progress but after this, there are no further chances to assess progress before the final milestone in late July. We say the manager has no "visibility" into the plan after May and must therefore count on perfect execution of a 3 month long task to meet his milestone. Because of the low visibility of the plan, we say this plan is "risky".

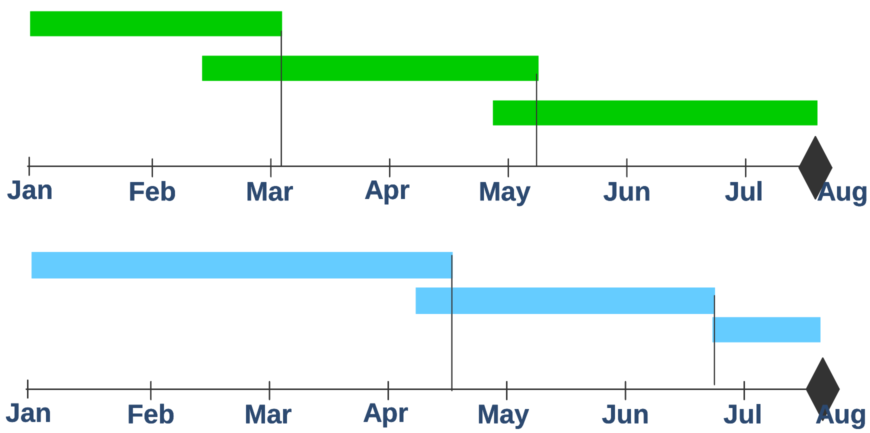

Figure 1: Two Plans, demonstrating visibility

In the bottom figure, because the tasks have been organized differently in terms of their finish dates, a manager can assess the project progress in mid-April and again in late June. Although the actual work performed is the same as in the top figure, the advantage of this scheduling of task reporting is that in terms of the target milestone, progress can be assessed and reported closer to the planned milestone. The manager is said to have better visibility in the second plan. To further improve the plan, the month long task during July may be further subdivided so that the project can be assessed even closer to the milestone date.

The Quantum© output is in the form of a graph (see figure 2), which measures cumulative visibility for each point in time. When the project begins there are no tasks completed, but at project completion all of the tasks should be complete. For a well constructed plan, an ideal graph of visibility would be a straight line extending from 0 to 1, from project beginning to project end.

A well managed project should have no periods where visibility is low, especially if those periods occur immediately before the project is scheduled to be complete.

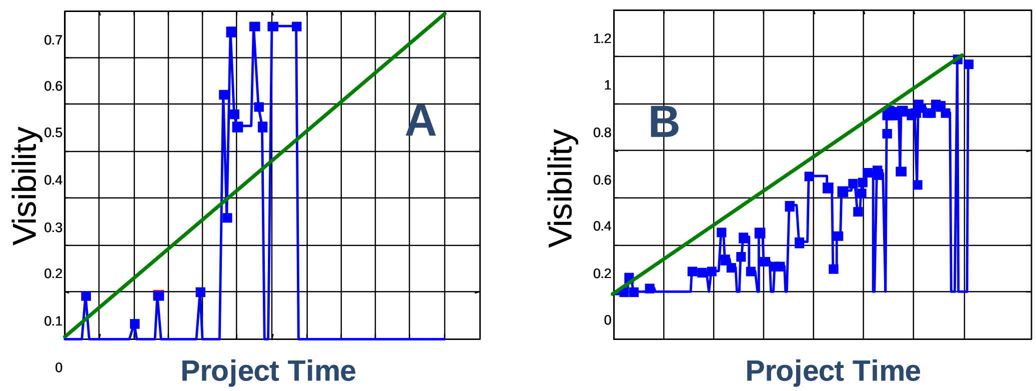

Figure 2: Examples of poor (chart A) and good visibility(chart B)

b) Probability

In the process of deriving a plot of visibility for a given milestone, it is necessary to determine the task tolerance. Task tolerance is the average amount the tasks in a chain may slip before they will affect the milestone date. It is not a measure of the individual amount a specified task may slip but is rather intended to be a relative measure of plan robustness or brittleness. Very low numbers for task tolerance always suggest plans that will need to be revised frequently, whereas higher numbers for task tolerance suggest plans that will tolerate more change without impacting the basic plan.

Figure 3 shows the expected probability of meeting the planned milestone date, plotted against expected task tolerance. For a typical plan having task durations measured in 1-2 weeks, a well considered plan can expect task tolerances of 1-2 days. For an expected task tolerance of 2 days, figure 3 shows that the expected probability is better than 80%. It may be a reasonable expectation that plans are approved only if they have a plan that meet such metrics.

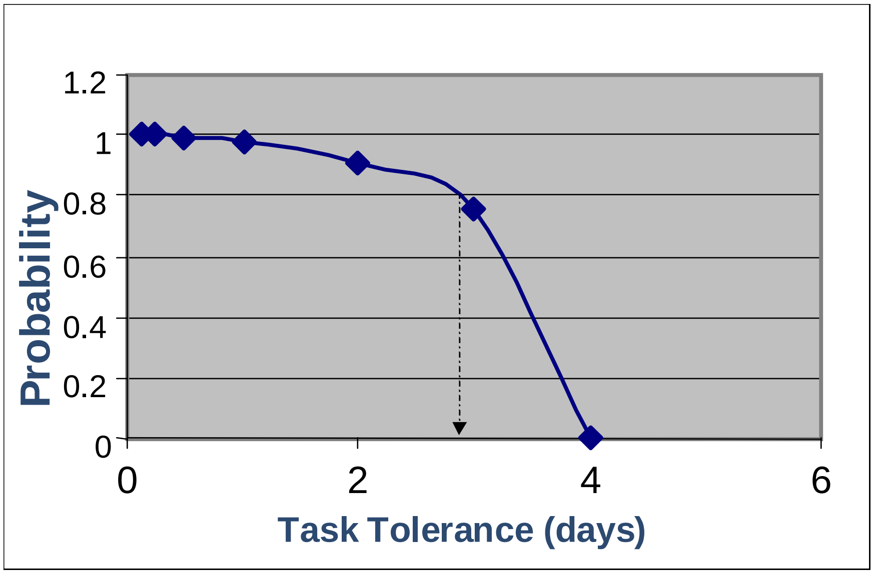

Figure 3: Robust plan with high probability of meeting milestone date

Figure 4 shows plan with very different probability. In fact, the probability falls to zero for even the smallest task tolerance, suggesting that the plan will need frequent re-plan. Even as the milestone date is moved out, the probability curves only improve slightly. Assuming acceptable task tolerance may be 2 days or more, this plan will clearly require a change in expected milestone date.

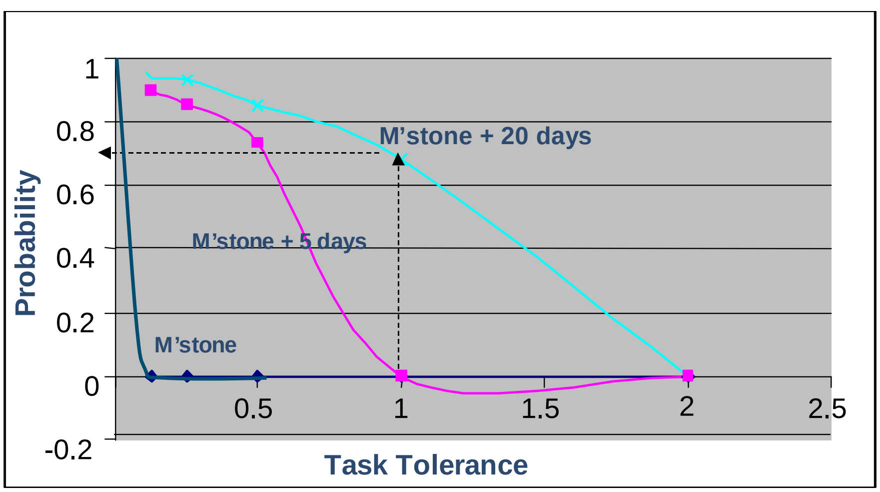

Figure 4: Brittle plan showing low probability of meeting original milestone date.

c) Task Importance

Although a structural approach to assessing projects does not depend on knowledge of the tasks, not all tasks are equal. For example, in figure 5 the top plan shows a task completed well before the milestone date, but the bottom plan shows the same task completed on the milestone date. It is obvious that in terms of meeting the milestone date, any manager would consider the bottom plan riskier, even though the tasks are the same. It is clear that the task at the bottom will always command a higher degree of management attention than the same task executed earlier.

In the lexicon of the Quantum© software, we say the "anticipated risk" of the task that is done later is higher than for the earlier case.

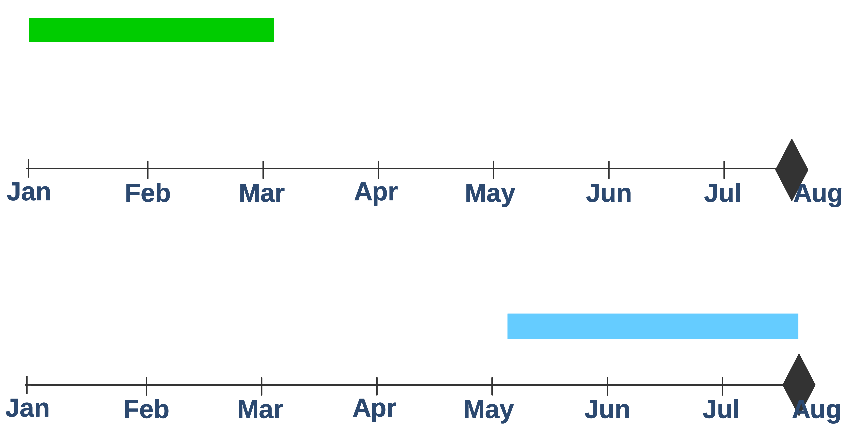

Figure 5: Task position has a large effect on project "risk"

It is interesting to note that while a critical path analysis would also suggest that the later task was more critical than the earlier task, if the later task did not finish exactly on the milestone date, it would rank the two tasks equal. The analysis offered using the new approach is substantially more comprehensive, gradually adjusting the "riskiness" of a task as it is finished closer to the milestone.

In any program, there are key tasks, or task groups, that a manager, executive, engineer or anyone with specific knowledge consider crucial to the program success. These tasks might have difficult schedule, scope or resource constraints, but they are crucial over and above their structural position in a plan. A well-conceived plan should correctly reflect the importance of these tasks by giving them appropriate ranking. Some of the tasks will be important only because they are close to milestones, but if we remove this effect then all of the tasks are important in some way, even those tasks done early in the project. Ideally, a ranking of all tasks by importance should reflect the actual importance to the project. For example, consider the plan in figure 6. If nothing is known about the tasks themselves, which tasks are most "important"?

ID		Task Name	Duration	Start	Finish	Predecessors
1		Mlst1	0 days	Tue 11/1/05	Tue 11/1/05	3,4,6,7,8,5,9,10,
2		Task1	11 days	Tue 8/16/05	Tue 8/30/05	
3		Task2	3 days	Wed 8/31/05	Fri 9/2/05	2
4		Task3	20 days	Mon 9/5/05	Fri 9/30/05	3
5		Task4	53 days	Tue 8/9/05	Thu 10/20/05	
6		Task5	8 days	Tue 8/16/05	Thu 8/25/05	
7		Task6	5 days	Fri 8/26/05	Thu 9/1/05	6
8		Task7	3 days	Fri 9/2/05	Tue 9/6/05	7
9		Task8	6 days	Wed 9/7/05	Wed 9/14/05	8
10		Task9	6 days	Tue 8/23/05	Tue 8/30/05	19
11		Task10	7 days	Wed 8/31/05	Thu 9/8/05	10
12		Task11	5 days	Fri 9/9/05	Thu 9/15/05	11
13		Task12	11 days	Fri 9/16/05	Fri 9/30/05	12,9
14		Task13	18 days	Thu 10/6/05	Mon 10/31/05	13
15		Task14	5 days	Thu 7/14/05	Wed 7/20/05	
16		Task15	10 days	Thu 7/21/05	Wed 8/3/05	
17		Task16	5 days	Thu 8/4/05	Wed 8/10/05	
18		Task17	4 days	Wed 8/17/05	Mon 8/22/05	
19		Mlst2	0 days	Mon 8/22/05	Mon 8/22/05	15,16,17,18

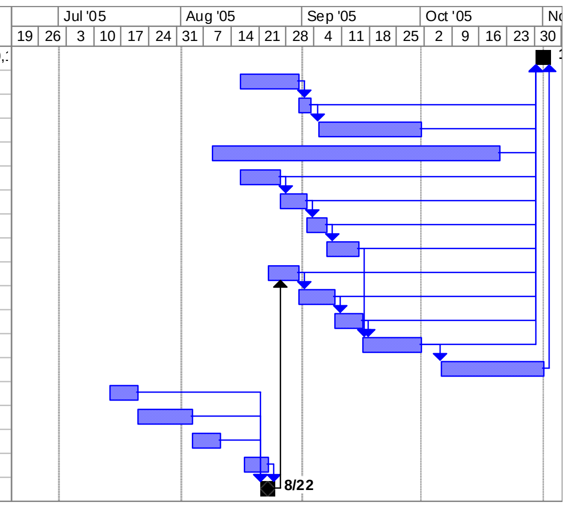

Figure 6: Simple 18 task plan

Figure 7 shows another view of this plan where each task ID is plotted opposite an axis representing measures of task importance. There are two curves, where the black curve is "task visibility", and the red curve is what we call "anticipated risk".

Task visibility is the relative structural importance of the various tasks from the point of view of the manager who must ensure that all of the tasks are completed. That is, those tasks with high dependencies will always appear more important, or "visible" than those tasks without such dependencies.

Anticipated risk is the relative structural importance of the various tasks from the point of view of the manager who is concerned only about meeting milestones. That is, *those tasks which, if slipped, can cause greatest potential slip in milestone date* are ranked highest.

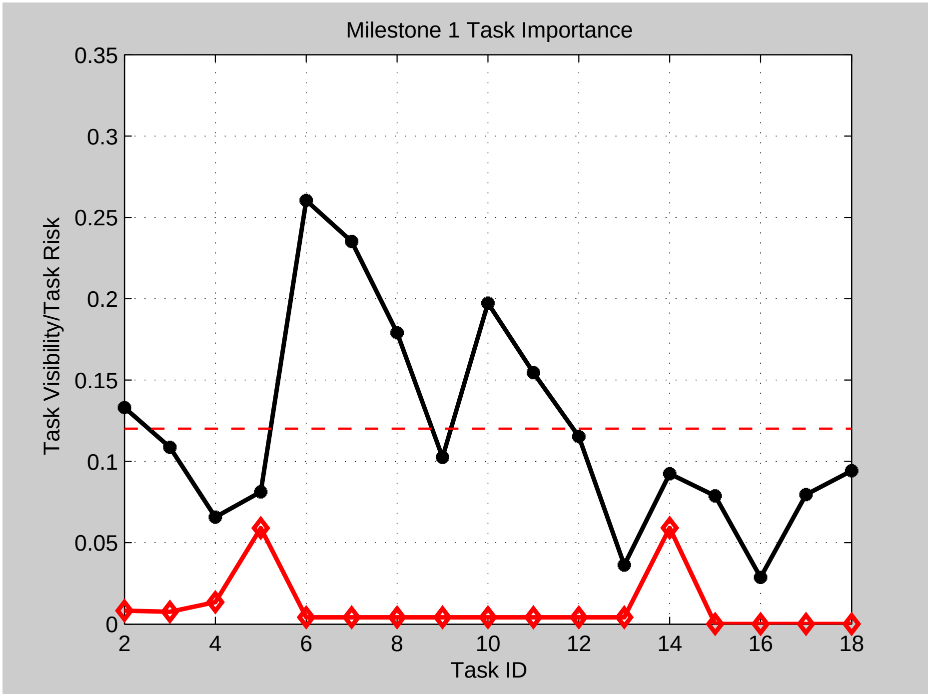

Figure 7: Task Importance Ranking

Using figure 7, tasks #6 and #10 are said to have the highest visibility, in the sense their importance is high in terms of getting the project completed. This makes sense because both these tasks lead off a series of tasks that depend on either #6 or #10 being completed.

The red curve shows the relative contribution of each task's ability to cause the milestone date to be slipped if the task is slipped. In this case, tasks #5 and #14 are important. This makes sense because these tasks are closest to the milestone, so that any slips in these tasks will have the most immediate impact.

If a task shows *both* high visibility and high risk, then it is a candidate for close management.

Comparing Two Plans

If we consider two plans with the same sequence of tasks, where one plan has more aggressive work estimates resulting in an earlier completion date, how would they compare? It is clear that the earlier estimate will be somehow riskier, but with the two plans having very similar structure, is it possible to compare the plans in a way that the regions of increased risk are exposed in a quantitative way?

Two such plans were compared as outlined below:

Plan = 3A
Project Start Date = 04-Feb-2002
Project End Date = 20-Dec-2002
Project Length = 319 days
Total Number of Tasks Excluding Summaries= 35
Number of Milestones = 3

Project = 5A
Project Start Date = 04-Feb-2002
Project End Date = 27-Sep-2002
Project Length = 235 days
Total Number of Tasks Excluding Summaries= 35
Number of Milestones = 3

The plans are identical in task sequence, and have the same start date, but 5A is substantially shorter than 3A, ending almost 2 months before 3A.

Figure 8 compares visibility for the final milestone for the two plans, and in this case the comparison is made for the case where task tolerance for each plan first causes the emergence of what is essentially a critical path, though not in the conventional sense. It is easy to see there are large differences between the two plans, with 5A showing high probability of plan deviation as early as June.

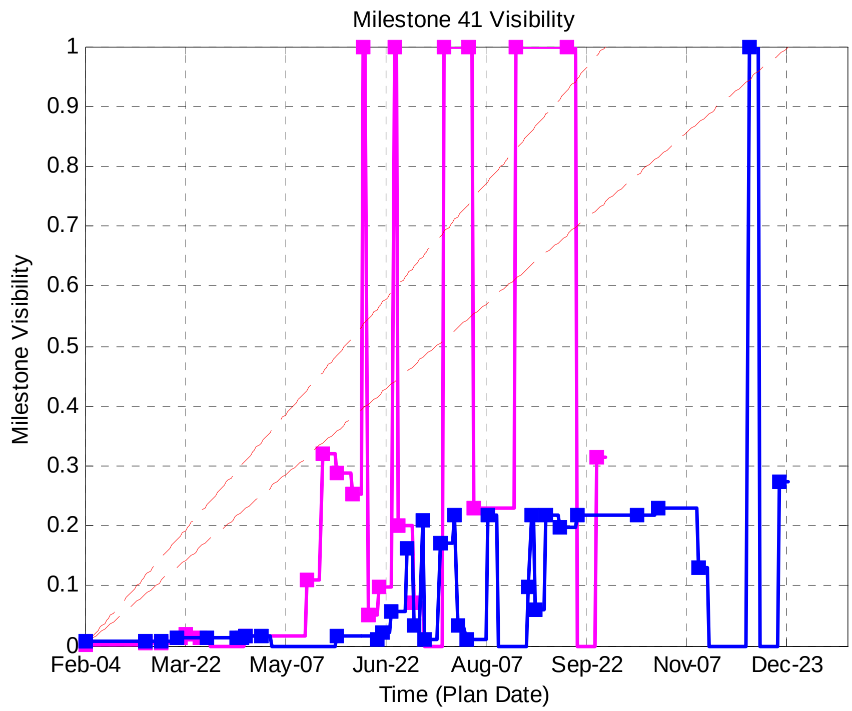

Figure 8: Comparing visibility plots for 5A (mauve) and 3A (blue)

A similar result can be shown if probability plots are superimposed on one another. This is shown in figure 9. In this case, both plans can be accomplished with task tolerances of 1-2 days, but while 3A can tolerate up to 3 day task tolerance, 5A cannot.

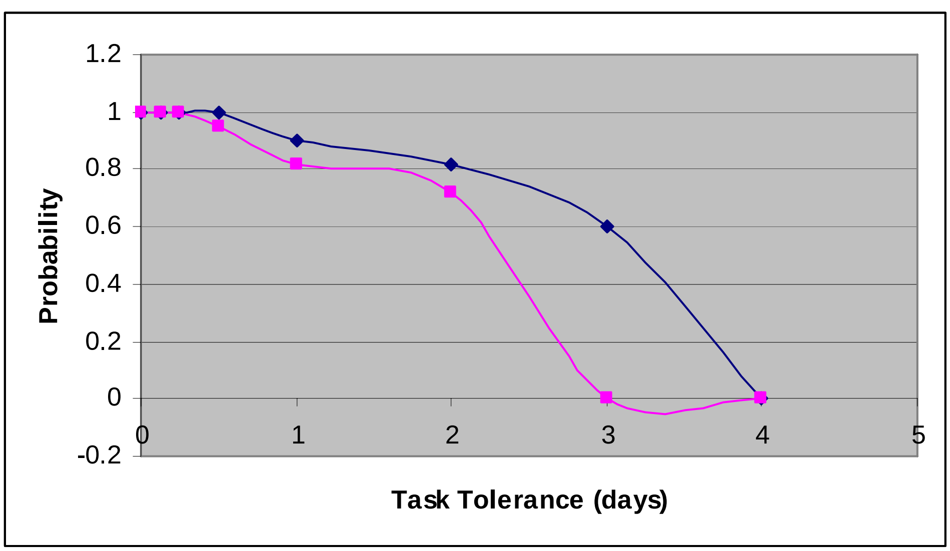

Figure 9: Comparing Probability for 3A (Mauve) and 5A (Blue)

Note there is no notion of whether the work estimates for 3A and 5A are accurate, but 5A shows increased risk none-the-less.

Summary

A new method of evaluating plans in a quantitative way is presented. The method makes use of plan structure, and in this way is entirely consistent with best practices. The method, when automated by software, is shown to quickly “score” plans, to correctly focus management energy on tasks most likely to cause disruption to milestones, and to reasonably compare plans.