PS-2871

A Realistic Look at Float Consumption

Greg Hall, PSP

Abstract

Float is an integral part of CPM scheduling. Without it, project teams would have difficulty identifying the activities requiring attention, or in predicting the significance of a delay or disruption. But mathematically determined float does not cleanly translate to the urgency of completing an operation in the real world.

Nowhere is this lack of definition starker than in the concept of *float consumption*, or the reduction in float of non-critical activities caused by a delay or disruption. In real-world CPM networks, resource limitations are often expressed through changeable Resource-Driven Logic. These 'soft' relationships can be modified to mitigate delays, but this dynamic greatly complicates Float Consumption theory.

This paper will focus on evaluation of more accurate methods to determine Resiliency, or how much stress can be applied to a project schedule before delay can no longer be mitigated. The fallacy of critical path 'purity' on resource-constrained projects will also be explored, as will such projects' resistance to accurate Monte Carlo style predictive analysis. This paper will then push beyond simplistic float concepts and explore the working relationship and communication needed for an owner and contractor to manage potential and realized delay.

Table of Contents

Abstract 1
Table of Contents 2
List of Figures 3
List of Tables and Equations 3
Introduction 4
What is Float? 5
Traditional Float Consumption: Trading Float for Delay 5
A Contractor's Options for Reacting to Delay 7
"Pure Logic Critical Path" vs. "Resource Driven Logic Critical Path" 8
Total Float Basis of Resiliency 8
Additional Measures of Resiliency – Free Float as an Alternative to Total Float 12
Earned Value's Role in Determining Resiliency 15
Measuring Resiliency through RDL to True Logic Ratios 18
Impact of Resource Flexibility on Monte Carlo Analysis 19
Conclusion 21
References 22

List of Figures

Figure 1: Average Total Float reduction from an auto-update. 6
Figure 2: Case study completion date and float consumption.. 10
Figure 3: Percent of activities near critical plot for Case Study project. 11
Figure 4: Relationship Free Float for an activity with two successors.. 12
Figure 5: Average RFF for all non-completed project relationships 13
Figure 6: Average RFF per RDL relationship.. 13
Figure 7: Activity Free Float density plotted against Finish Variance.. 15
Figure 8: Case Study SPI plotted against Finish Variance. 16
Figure 9: Weekly labor hours before and after resource leveling.. 17
Figure 10: Labor Flex compared with Finish Variance.. 18
Figure 11: Ratio of RDL to Total Relationships. 19
Figure 12: Demonstration of RDL's impact on Monte Carlo Simulation 20

List of Tables and Equations

Table 1- Resource Driven Logic on the critical path. 8
Equation 1 - Schedule Density (Moder *et al*)... 13

Introduction

Critical Path Method (CPM) scheduling was developed during the advent of computers. Using mathematics, its creators sought to harness the calculation power of the new machines as an improvement on Gantt Chart scheduling and manager reckoning. [1, p.9] The pervasive nature of CPM scheduling in construction speaks to its usefulness as a tool, but as with many other tools it is not universally the best implement for every task.

Because of the chaos and unpredictability of the real world, it's difficult to fit the rigid, objective mathematics of CPM around every element of a project. However, because CPM is the best tool available, project management teams will make the assumptions necessary to construct their schedules and have confidence in them. Most discussion of the ability of a CPM schedule to accurately model a real-world project assumes:

- An activity's actual duration will be close to its original value, unless some clearly defined and easily modeled delay or disruption event affects it.
- The baseline logic determined by the contractor and shared with the owner represents the best possible sequence to complete the work, based on the knowledge the parties have at time of bid.
- This baseline logic will only be changed as part of a 'revised baseline' process.
- Any progression of work not following the baseline logic is Out of Sequence; the implication is that this progression has compromised the schedule's integrity.

The concept of float in all of its various forms is the mathematical product of CPM calculations. Simplistically, float represents how much an activity's timing can change before effects are observed outside of that activity. If the criteria described above hold true throughout the life of a project, then all of the various types of float could be objectively applied as a project management tool in a much more predictable and orderly way. The concept of *Float Consumption*, where a delay to a noncritical activity would always first consume that activity's float before it affected the rest of the schedule, could always be cleanly applied.

Unfortunately, real-world complications, such as the possibility of changing logic without a revised baseline, force project management teams to apply subjective judgment to the pristine mathematics of CPM. Logic changes made as part of an update process can add subjectivity to how that project's float consumption is perceived. Subjectivity breeds dissent, which, coupled with threats to completion time and budget, often grows into a dispute between owner and contractor.

This paper seeks to apply some order to that subjectivity by temporarily setting Float Consumption aside and correlating other CPM schedule data with a project's ability to absorb delays and disruptions. Through such correlations, analysts may be able to determine whether a delay or disruption can or cannot be mitigated through logic changes.

What is Float?

At the activity level, Float Consumption is commonly equated to Delay Duration minus Total Float equals net delay. While the classic Early Date minus Late Date version of Total Float is important to Float Consumption, other float definitions may come into play.

Most contracts and textbooks define float in this manner:

> Free Float is the length of time the start of an activity can be delayed without delaying the start of a successor activity. Total Float is the length of time along a given network path that the actual start and finish of activity(s) can be delayed without delaying the project completion date. Project Float is the length of time between the [contractor]'s Early Contract Completion and the Contract Completion Date. [2, Provision 3.11.1, p.300]

A contractor incurs time-dependent (usually overhead) costs for every day they are engaged on a project. Their final project cost therefore depends on the attainment of its completion date, and not on the contract-defined deadline for completion. For that reason, the metrics discussed in this paper evaluate Total Float against the scheduled attainment of completion.

Float is a useful metric for more than project delay. Free Float changes can indicate coordination issues among the various supervisors in charge of project elements. Total Float enables a management team to concentrate on the current ‘top threats’ to on-time completion, and to identify ‘opportunities’ to improve it.

Traditional Float Consumption: Trading Float for Delay

Float is often called a commodity or a resource. Many construction contracts further describe float as a “shared commodity to be reasonably used by either party” [2, Provision 3.11.2, p.301]. The concept fits nicely with the belief that as a project progresses, float should reduce, to the point where it diminishes to zero at the end of the project.

In this way of looking at things, float *should* be consumed in the way that a piece of equipment consumes fuel to do its work throughout the day.

Figure 1 below demonstrates this phenomenon. A ten month long baseline schedule was progressed monthly using the Update Progress tool of Primavera P6© to actualize Early Start and Early Finish dates. In addition, remaining duration on each update’s in-progress activities was adjusted such that the baseline dates were followed without deviation. After each of these automated updates, the average value for Total Float for all non-completed activities was calculated and graphed. Since completed activities do not have Total Float values, only in-progress and not started activities are used.

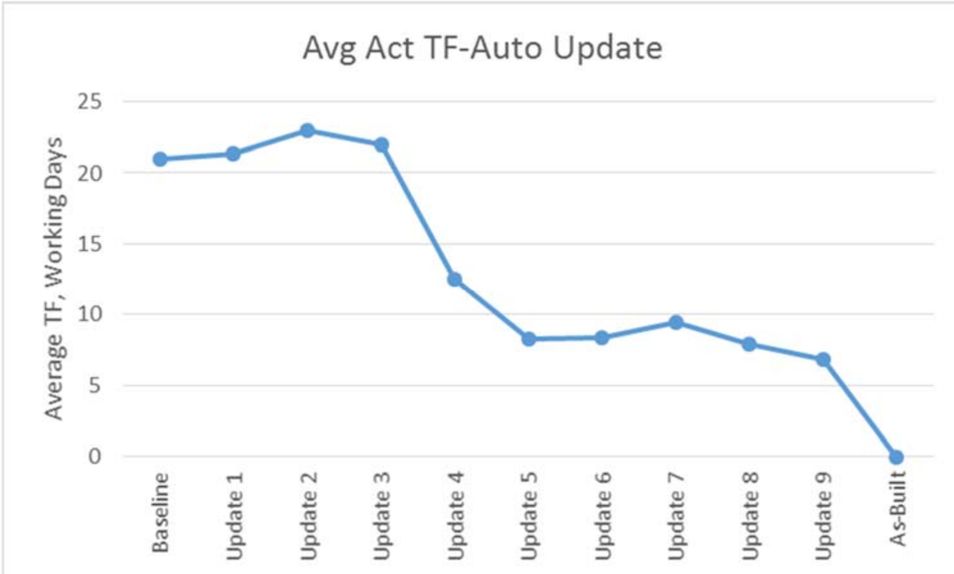

Figure 1: Average Total Float reduction calculated from an auto-update of a ten-month project's baseline schedule.

A comparison of this planned float 'burn' to the average activity Total Float from progressive as-built updates can be an early warning sign of schedule compression. [3, p. 8.] However, logic changes from update to update will reduce the accuracy of this graph, as will the addition of detail which is common to rolling-wave planning. [3, p. 12]

The visualization of float as a consumable resource leads to the concept of *float ownership*, a topic of debate which still isn't conclusively settled. Float ownership theory perhaps takes the 'float as a resource' philosophy one step too far. Using this philosophy, a party who is responsible for a delaying event first has the option to 'consume' the float that they 'own' before they have to expend assets to make the other party whole through acceleration or compensation. [4, p.34]

Float Consumption theory is one of the many ways in which CPM practitioners attempt to approximate, through mathematics, what is really going on: a party who does not affect the critical path does not have to accelerate, or compensate the other party for acceleration, until their delay affects the end date.

One of the things that should be obvious from Figure 1 is that a project with *no delays* will still exhibit a trend of average Total Float that generally reduces with progress. Available float dwindles down to nothing as a function of progress alone; after all, the last activity in a properly constructed CPM network always has zero Total Float. The question becomes: with the knowledge that logic *can* be changed, how far would the actual project's Float Consumption pattern have to deviate from Figure 1 before the end date cannot be recovered? In other words, what is the resource-constrained schedule's "Resiliency" for recovering completion time at little to no cost when subjected to impact?

A Contractor's Options for Reacting to Delay

Contractors are in the business of bringing resources and materials to a worksite to complete the work at the least cost. This is why they create resource-constrained schedules. Algorithm-based resource leveling does not see widespread use. Instead, a contractor usually adds relationships into its CPM schedule that represent a prioritization of activities a certain resource is tasked to perform. [5, p.7] Without such Resource Driven Logic (RDL), the contractor would find it extremely difficult to realistically account for the finite number of people, equipment, and tools brought to the site.

If something prevents a resource from performing the activity to which it is assigned, or a material is not available for that resource to install, a delay has occurred. [6, p.6] Because CPM calculations derive from activity starts and finishes, this delay manifests itself in activity data. Of course, not all activity delays result in project delays; the float of an affected activity often exceeds the delay duration. This is little comfort to the superintendent responsible for the activity, who must find different work for the crew tasked to perform it.

Early in a project, a contractor can often task a delayed resource to another operation. By CPM definition, that contractor has decided to work out of sequence. Out of sequence work is normally referred to in negative terms, but in resource usage terms, it gives a contractor the flexibility needed to eliminate the impact of the delay. More importantly, there is no physical impediment to the contractor performing this work in the revised order.

While moving a resource from a 'Plan A' operation to a 'Plan B' might eliminate delay impact, there may still be a disruption impact in the form of a productivity loss. [4, p.1] However, this paper will not focus on that element of change management.

In cases of a delay where there is not another activity available for affected labor resources, a contractor generally reacts in one of two ways. The first option is tasking the crew with semi- or non-production work, such as prefabricating forms or cleaning and replenishing tool trailers until the problem is solved. If it looks like the delay is lengthy enough, the contractor may send the crew to another project, give the crew some time off, or (as a last resort) lay them off entirely. All of these options have a cost element that can be priced into a change order, but the cost is generally small when compared to the cost of pushing a project's end date.

If the activity is delayed less than its current Total Float value, then the project completion date is maintained. However, the removal of the operation as a piece of productive work available to a resource means that other options must remain available. The risk that a separate delay to unimpacted work could not be solved so cheaply has increased, and the schedule is less "Resilient".

"Pure Logic Critical Path" vs. "Resource Driven Logic Critical Path"
Every scheduler began their career with a basic understanding that a CPM schedule's critical path as the path that "determines the length of the project…the longest path into the last event, since it establishes the last [event time] for the last event."[1, p.103]

In keeping with the theme of 'stipulation for the sake of approximation', most introductory CPM instruction proceeds with an assumption that all relationships are True Logic: that they are not supposed to be changed. In reality, RDL representing a preferential prioritization of resources by the contractor is peppered throughout most project schedules. It does not take much imagination to conclude that RDL might find its way into the critical path of such schedules.

The degree of invasiveness varies highly, but RDL is not difficult to identify and evaluate. A relationship linking two activities which use the same resource but are in different physical locations is generally resource-driven.

Factors which limit the amount of RDL on the critical path are natural progression, such as a tunnel created by a single TBM; phasing, such as the restriction on how many highway lanes or interchanges a contractor can take out of service at any one time; and site access.

Table 1 below summarizes a study of 32 schedules from 17 randomly selected projects. This study was hardly comprehensive, but it does show that regardless of discipline, RDL on the critical path should be no surprise. Of the 32 schedules reviewed, only five had no critical path RDL.

	Schedules Observed	RDL	Total Relationships	Percentage
Dams/Hydropower	4	56	344	16%
Highway	11	55	663	8%
Industrial/OGC	5	104	360	29%
LRT	7	89	343	26%
Water/Wastewater	5	183	496	37%

Table 1-Study of Resource Driven Logic (RDL) on the critical path of randomly selected construction projects.

Because critical paths regularly contain RDL, contractors can change the soft logic to respond to delay. For this reason, and the fact that any experienced scheduler can perform critical path evaluations, the studies performed for this paper concentrated on a holistic evaluation of all activities and logic.

Total Float Basis of Resiliency
Most project teams use Total Float to measure project criticality. Total Float is a factor in checks number 6, 7, 12, and 13 of the "14-point assessment" advocated by the Defense

Contract Management Agency (DCMA) as a measure of schedule quality. [7, p. 7 and 9] Float consumption graphs such as the one shown in Figure 1 can demonstrate if a project is consuming float in a neat and orderly way. Unfortunately, such graphs are only fully reliable if logic remains unchanged from project start to completion. Any of a number of likely scenarios would undermine the usefulness of an average Total Float graph in determining Resiliency:

- The project is delayed on its critical path and a time extension is granted, increasing the float on noncritical (and unimpacted) paths.
- A noncritical, high-budget operation proceeds well ahead of its estimated production rate, skewing earned value metrics and providing a false sense of achievement not echoed by execution of critical activities.
- The project is delayed on its critical path but there is a dispute over entitlement to a time extension, resulting in negative float which the two parties do not address over several months. [8, p. 23]

Most importantly, average Total Float metrics do not answer the most important question affecting project prosecution: whether or not a contractor can restore an on-time completion by rearranging RDL.

A case study was conducted to evaluate alternative metrics for identifying this 'tipping point' on a project. The project studied was a Light Rail Transit (LRT) project in an urban setting. Its schedule had over 5100 relationships at time of baseline, only four percent of which were RDL. Near the end of the project, its final cost is projected to exceed its original price by about 25 percent and its completion date is projected to be over a year late.

These factors led to a significant dispute between owner and contractor, but that is not the focus of the study. Rather, this job was selected because it faced a number of significant delays which began early and continued to worsen. In reaction to these delays, the contractor rearranged RDL and performed other mitigations, and was successful in two instances (the fourth and eighth month) to restore the original completion date before later delays made on-time completion, in the later opinion of the project management team, impossible. By the ninth update period, the completion date deteriorated once more, never to be restored.

Discussions regarding a low-cost recovery of an end date which ultimately would never have been attainable wasted a significant amount of time and energy for both the owner and contractor's managers. Thus, they would have derived a significant savings had they been able to identify the point at which continuous delays consumed all of the project schedule's Resiliency. The case study concentrates on the first twelve update periods of the project, with the objective of capturing when indicators would have informed the owner and contractor that they could no longer use RDL rearrangement to restore an on-time completion.

Figure 2 below shows that merely evaluating average activity Total Float would not indicate to an analyst that anything was terribly amiss. The green line represents average Total Float in the update schedules and the purple line was derived through automatic updates of the baseline schedule in the style of Figure 1. For six months, the actual schedule appeared to have more float in it than expected from the baseline. Actual average Total Float only became less than planned during Month 12, three months after the project, in the retrospective opinion of the project management team, was beyond recovery. Thus, average Total Float did not appear to be predictive of this project's Resiliency.

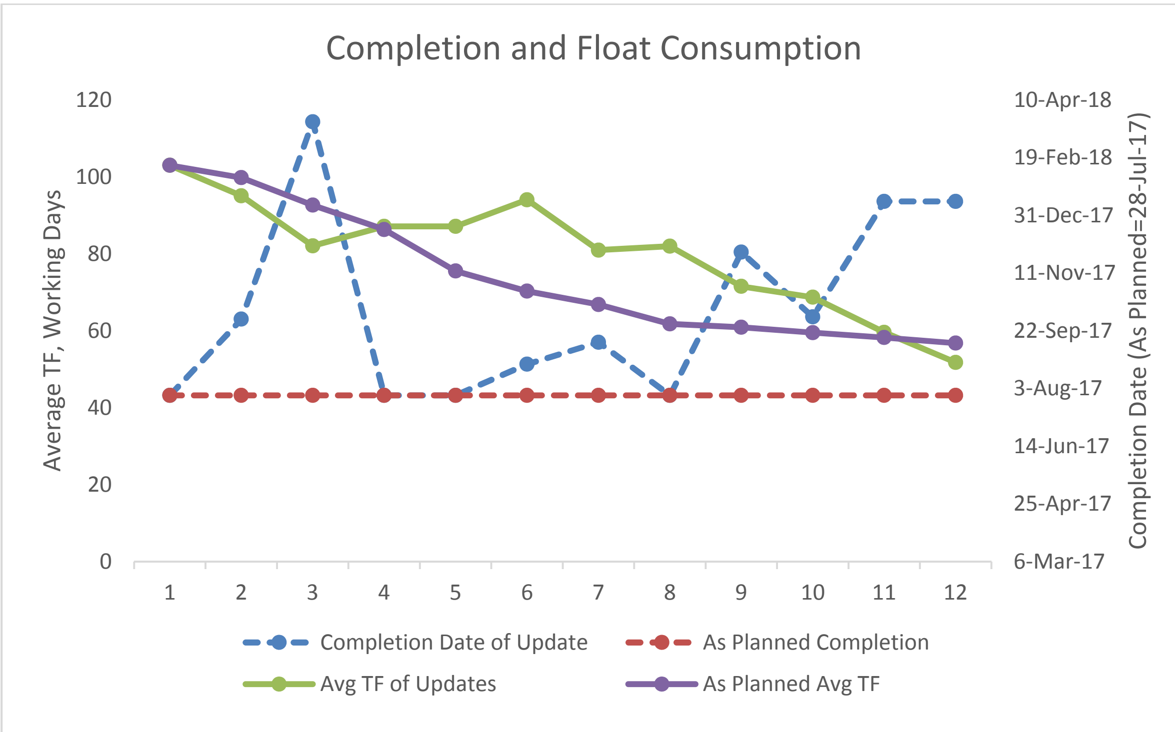

Figure 2: Case Study Completion Date and Float Consumption (reduction of average Total Float of 'in-progress' and 'not started' activities)

The large improvement in completion date from Update 3 to Update 4 may lead one to assume that float was 'traded off' to get the project back on time; in fact, the opposite was true and average activity Total Float improved. Changes to RDL could explain this. Between the two updates, 8 of the 201 RDL relationships were completed or deleted, four were added, and one was changed. In Update 3, the average Total Float of RDL successor activities was 47.6 days. In Update 4, it was 50.6 days. These relatively few changes to RDL were still significant enough to restore on-time completion and to call into question any 'apples to apples' comparison of Total Float between Updates 3 and 4.

Another important observation is that an after the fact forensic analysis (following AACE® International Recommended Practice 29R03, Method Implementation Protocol 3.3) identified six separate logic paths that at one time or another controlled project completion.[10, p.51] Thus, one cannot simply dismiss anomalous data as the result of one heavily delayed path while the rest of the project hummed along without issues. During all of the studied periods, roughly half of all activities in the schedule were at least near-critical.

The discussion of near-critical activities introduces another predictive tool that is reliant on Total Float: Near-Critical Activity Percent. This statistic has been used to establish developing trends with the purpose of warning a project team of impending issues. [3, p.12]. The case study project's values are plotted in Figure 3 below. Both a 'fixed' line (based on a near-critical threshold of 5% of duration remaining to the contractual completion date) and a 'true' line (based on a near-critical threshold of 5% of duration remaining to the calculated completion date) are plotted against each update's completion date. Figure 3 does not seem to predict irrecoverable delay, and instead only reported delay already identified:

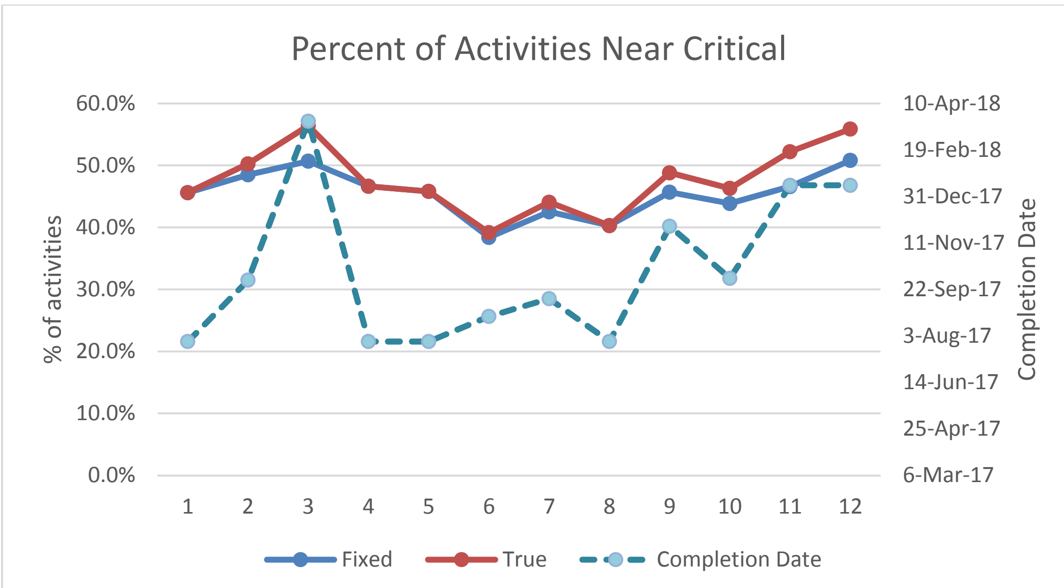

Figure 3: Percent of Activities Near Critical Plot for Case Study Project.

The factors of the case study that worked against the viability of this chart as a predictive tool are:

- The project faced constant delays (many of which affected the critical path) from the beginning.
- The available work fronts on the project meant that there was a significant percentage of RDL that could be changed from month to month.
- From the start, the percentage of activities that were near critical was on the high end of a recommended threshold. Jackson *et al* recommend a range of 20-40%. [3, p. 13]

However, the trend of near-critical activities could indicate that the project had fully consumed its Resiliency and RDL could no longer be rearranged to restore an on-time completion. Periods 10 through 12 showed a marked upward trend. Most notably, the near-critical increase from period 11 to 12, despite an identical completion date, was something not observed in prior periods.

Additional Measures of Resiliency – Free Float as an Alternative to Total Float

Total Float is certainly an indicator of delay to overall completion. However, when looking for a pattern of delay to individual activities, as well as their density and overlap, Free Float may paint a more accurate portrait. [4, p. 35]

Since by definition Free Float is the measure of how much an activity can be delayed before it affects a successor, it is important to differentiate between Relationship Free Float (RFF) and Activity Free Float (AFF). These two measurements can be different.

In Figure 4 below, Activity A has two successors. Activity B is driven by a relationship with a different activity, thus the relationship between A and B has an RFF value of four days. Similarly, the relationship from Activity A to Activity C, which is also driven by a relationship with a different activity, is one day. Each relationship can be described as having its own value for Free Float, but the smaller of the two will be Activity A's value for Activity Free Float.

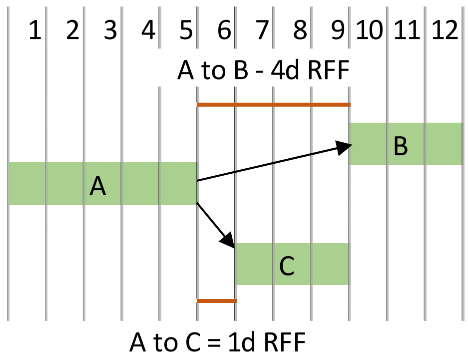

Figure 4: Relationship Free Float for an activity with two non-driving successor relationships.

Evaluation of RFF is especially useful if logic changes frequently. RFF, with its basis in relationships and not activities, might be more effective than Total Float in measuring how efficiently RDL is being used.

Figure 5 below tracks RFF in the case study schedule from update to update. It appears that delays to the end date actually increased average RFF, at least in relationship to the planned gradual 'burn off' of RFF as the project progressed.

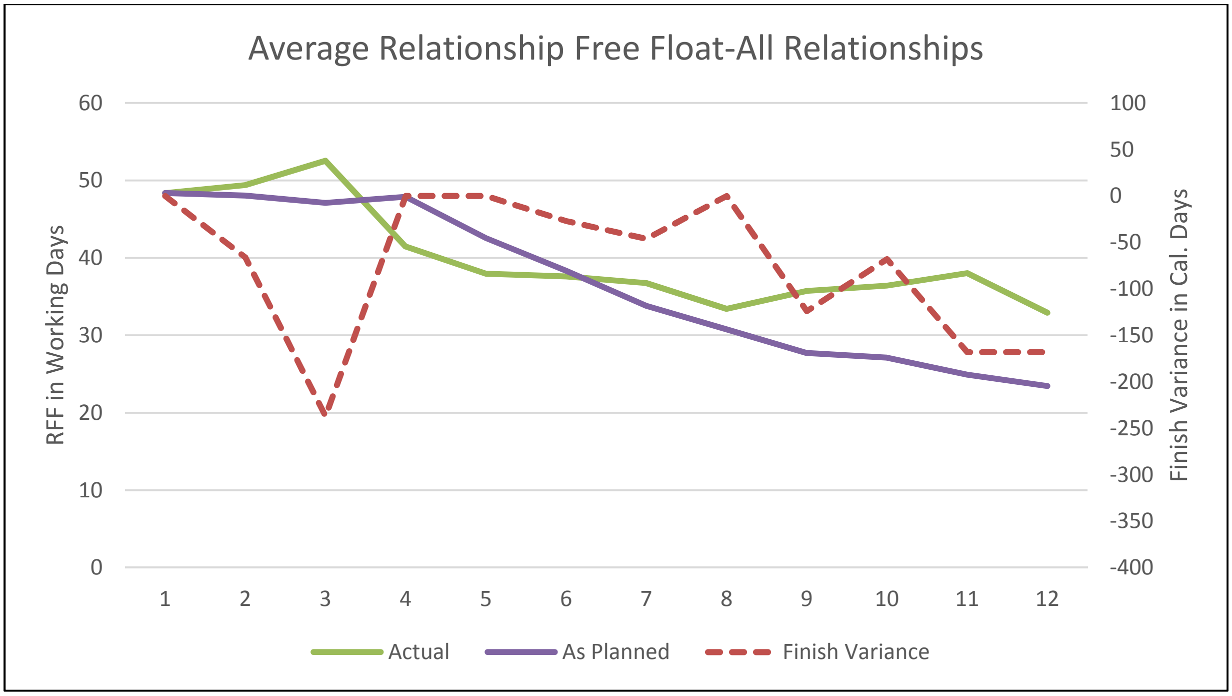

Figure 5: Average RFF for all non-completed project relationships over the first twelve update periods.

A review of RFF from RDL relationships only may be more enlightening. The RFF values were all larger than what had been originally planned.

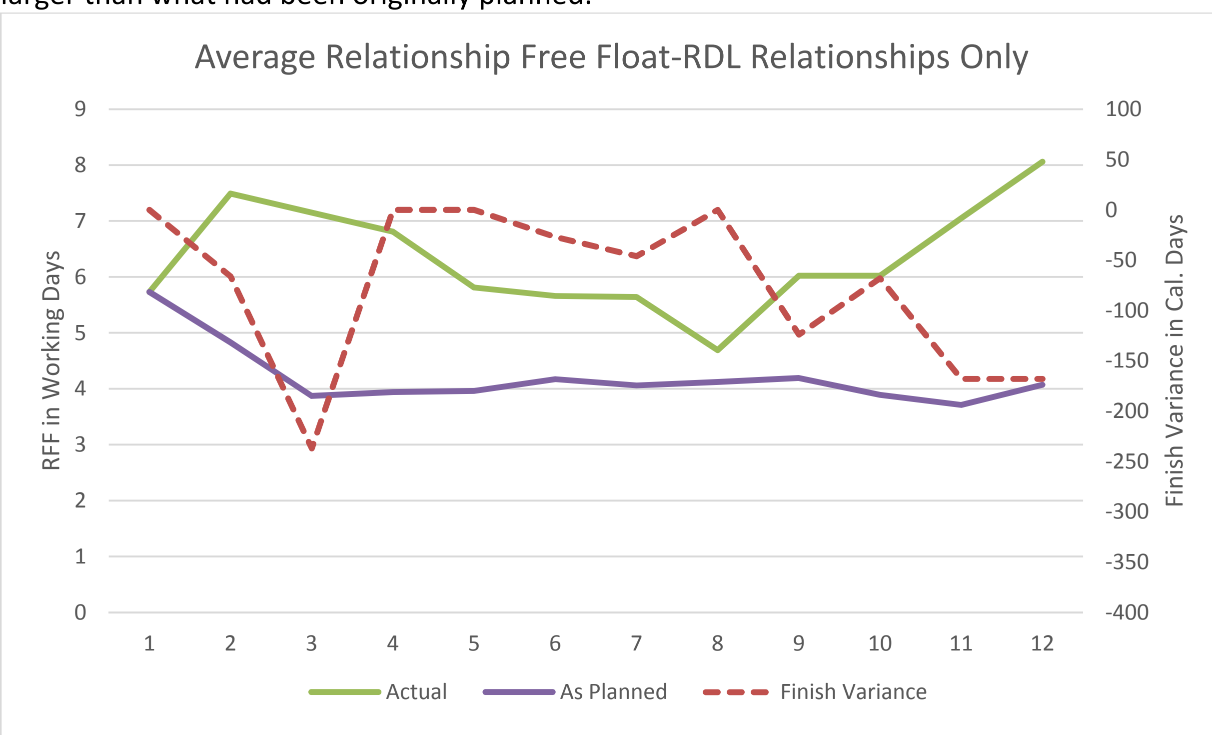

Figure 6: Average RFF per RDL relationship. .

Figure 6 shows the baseline schedule's consistent RFF average to be around 4.0 working days per relationship, indicating the plan for resource deployment was well thought out, but not perfect. Because a relationship with an RFF value of zero is a driving one, an RDL relationship with an RFF value of greater than zero did not need to be added to the schedule.

Therefore, a higher average RFF value would indicate either that RDL had been indiscriminately added, or that existing RDL relationships had been rendered moot as true logic relationships became the driving ones. This case study did not explore if redundant RDL would impact the schedule's accuracy as a predictive tool. It should not affect Total Float calculations, which are based on driving relationships. The author's prior experience is that contractor teams usually remove redundant RDL during a major revision to resource deployment, but otherwise only remove obsolete RDL that later appears on a critical path.

As Figure 6 shows, up-front delays immediately reduced the number of driving RDL relationships. It does not appear that the recovery of the completion date achieved in Period 4 required an extensive re-thinking of resource deployment, as the reduction in average RDL RFF is not significant. However, the final recovery of the completion date achieved in Period 8 correlates with a significant reduction in average RDL RFF. It would be reasonable to conclude the job team invested time and effort that month reviewing and optimizing its resource flow.

The steady increase in average RDL RFF over Periods 9 through 12 may be evidence that the contractor decided that recovery of the end date was no longer tenable, and that its effort was best spent on requesting a project time extension. It's common on projects where timely time extensions are not granted for a contractor to turn their attention from optimizing the schedule and toward the inevitable end-project claim. [8, p. 25] In fact, the contractor was busy compiling the first of four large claims during Period 10, which they transmitted to the owner during Period 11.

RFF tracking requires sophisticated scheduling software, while AFF is more widely available. AFF has been used in the past as a key component to overall schedule density. An equation cited in *Cost Engineering* from Moder *et al* provides a simple means of comparative schedule density [4, p. 36]:

$$Schedule\ Density = \frac{(Sum\ of\ Activity\ Durations)}{(Sum\ of\ Activity\ Durations)+(Sum\ of\ AFF)} \quad \text{Equation 1}$$

As stated above, this is a comparative calculation; there is no identified 'rule of thumb' for what a 'low risk' or 'high risk' value of Schedule Density yet specified. However, the trend of this value provides some meaningful insight to the progress of the project:

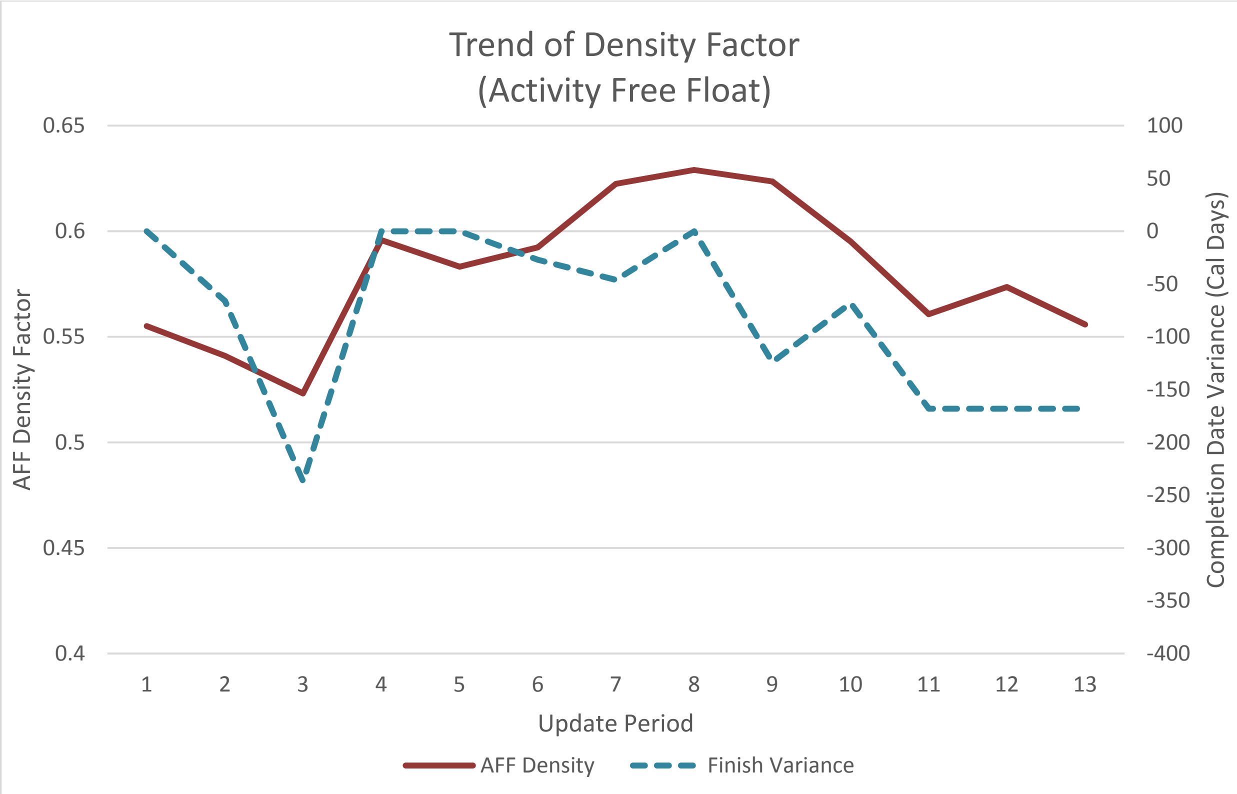

Figure 7: Activity Free Float density plotted against Finish Variance.

During the large Period 3 delay, Figure 7 shows that density reduced, or at least that AFFs increased. This may correspond to the contractor recording the delay, but not reconfiguring their RDL to eliminate it. Period 4's marked increase in density correlates to the recovery of an on-time completion; the contractor made revisions to respond to the delay.

Period 5 showed an improvement in density while completion date remained consistent. While this may not universally correspond to regaining some lost time, it should be noted that Period 5 was a mild winter month and that the contractor performed some operations at a time they had originally planned to be shut down.

The next interesting part of Figure 7 is when the end date deteriorated, but was recovered for the last time in Period 8. This period also saw the highest AFF density. From there, density decreased.

Does this mean that the on-time prediction of Period 8 was untenable? It could certainly be argued that if progress during Period 9 did not follow the timeframes set out in Period 8, the contractor would have a difficult time bringing the end date back.

Earned Value's Role in Determining Resiliency

Earned Schedule and Schedule Performance Index (SPI) have not only been championed as schedule management tools for a long time, they have also emerged from prior comparisons with other schedule metrics as the tool of choice for predicting project outcomes. [10, p.13] However, is SPI meaningful in indicating when a schedule can no longer absorb delay?

There should be a correlation between SPI and float reduction. A large increase in average float could result from performing operations ahead of schedule. It should therefore be accompanied by an SPI (for the current month) with a value greater than 1. An increase in average Total Float is possible with an unsatisfactory current SPI, but only if a) there have been logic changes and/or b) there has been a significant delay to the critical path, but noncritical activities have been unaffected.

Figure 8 shows that SPI for the case study does correlate to the period when recovery was no longer possible. Of note is that SPI remained better than 0.9 until Period 10, and was a moderately successful 0.97 during Period 9, the month after the last on-time update.

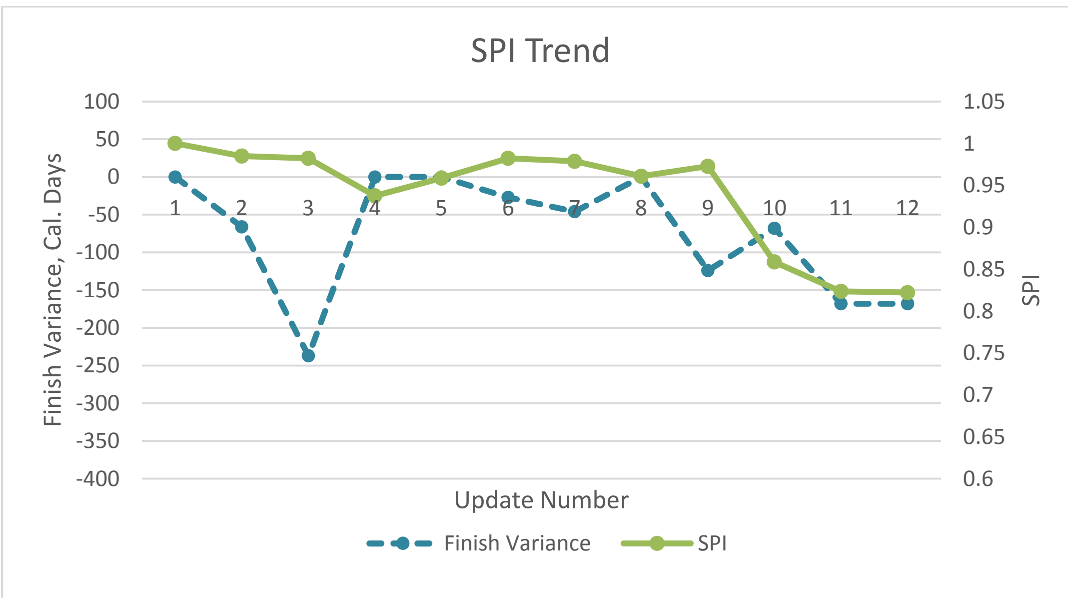

Figure 8: Case Study SPI plotted against finish variance.

There is a possible correlation between the spread of Finish Variance and SPI on a trend chart and the ability for the schedule to absorb delay. For this correlation to be true, there would still have been hope during Period 9 that the project could be kept on time, but subsequent delays eliminated that slim opportunity. During Period 9, the contractor actually thought it could restore an on-time completion, and Period 10's Finish Variance hints at an unsuccessful attempt to once more recover the completion date. The significant drop in SPI shows that on-time completion would have required a significant improvement in monthly earned schedule.

Another time-tested Earned Value metric is the time-distribution curve of cost, or, for ease of apples-to-apples comparison in different labor markets, a time-distribution curve of labor hours. Without RDL or some form of resource leveling, these curves can show a high degree of period-to-period variance and are evidence of an unrealistic plan.

Figure 9 demonstrates one key characteristic of RDL or automatic leveling. Leveling will smooth planned values to something more consistent and achievable, but the scheduled dates of activities are often later than in the un-leveled version. [11, p. 25] Green represents the labor hour requirement before resource leveling, while red represents the labor hour requirement after leveling. The dashed lines represent the date at which half of the budgeted hours are to be earned.

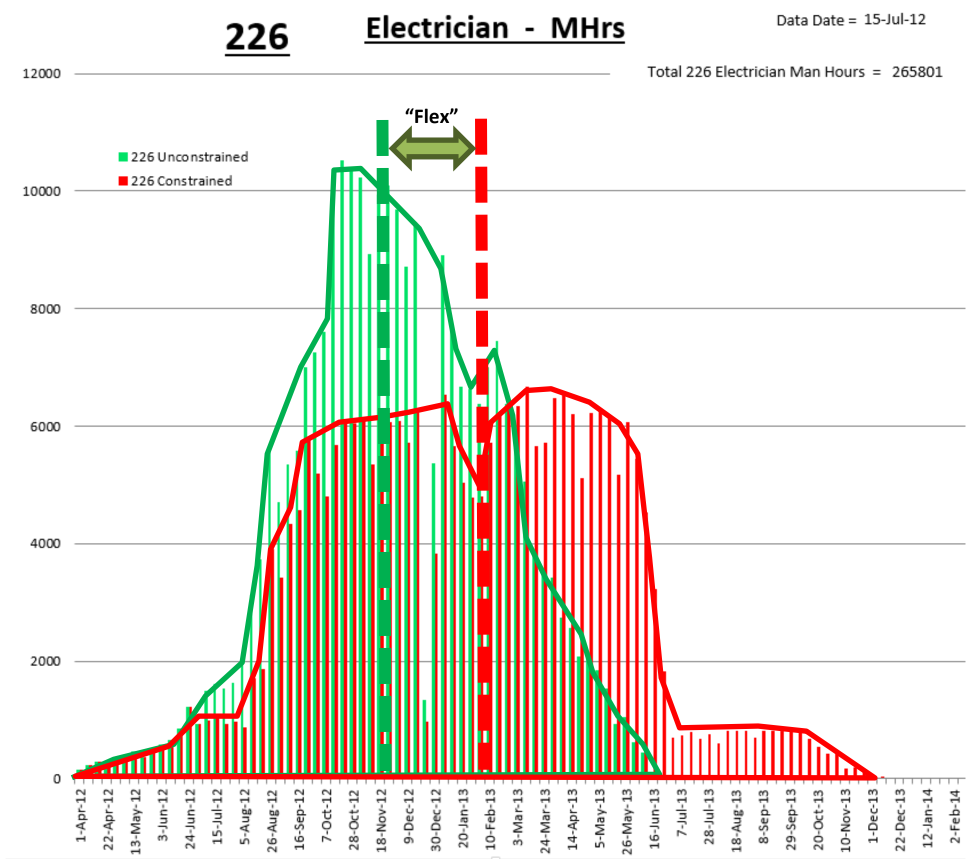

Figure 9: Electrician weekly labor hours from an update of a metal refinery project.

The amount of curve-shifting or "Flex" demonstrates how leveling measures affected the project. It would follow that the greater the amount of RDL in a schedule, the larger the difference between the leveled and unleveled states.

The difference in the date each curve reaches half of its budgeted hours (subsequently referred to as "Labor Flex") might indicate a schedule's ability to absorb change. After all, additional restrictions on the un-leveled schedule would push the distribution of labor hours later – and thus closer to the leveled schedule distribution.

The Labor Flex in the case study was determined by comparing the difference in leveled and unleveled curves in each month's update. To create the unleveled curve it was necessary to remove all RDL in a copy of each schedule. This is no mean feat when evaluating over 4000 relationships, so such a study might not have much of an appeal to project management teams. To determine baseline values, RDL had to be removed in a copy of the baseline which was auto-updated to each month-end date.

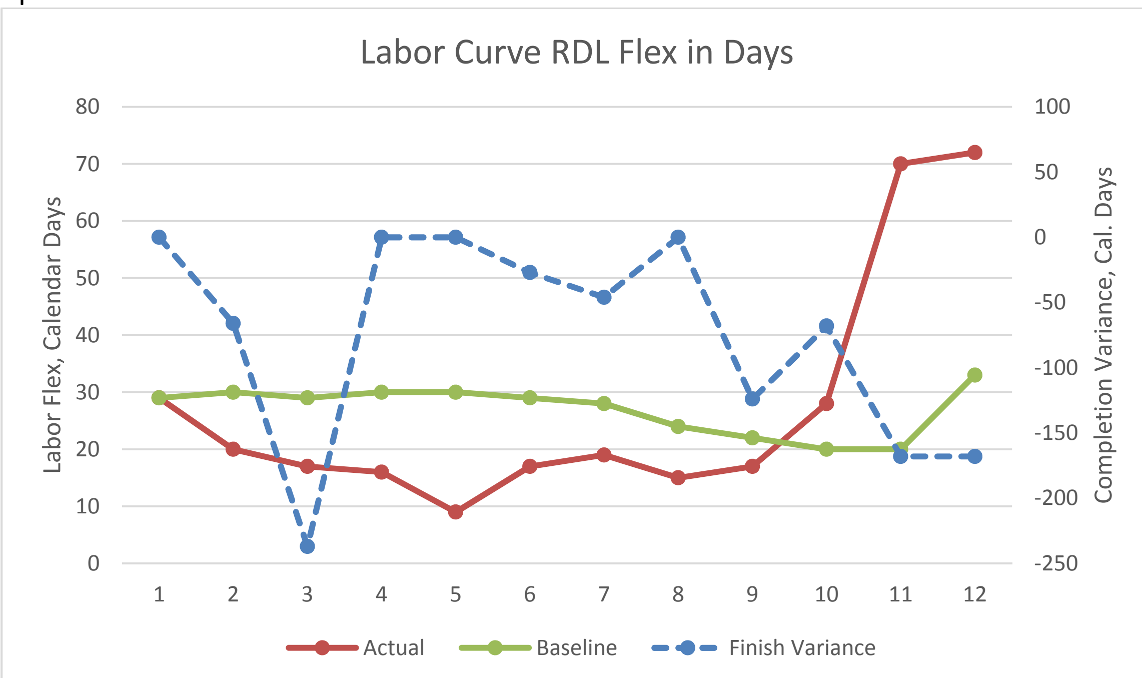

Figure 10: Labor Flex of both the baseline and actual schedules, compared with Finish Variance.

The data in Figure 10 above shows that Labor Flex dropped immediately from its planned levels, hitting a low point in Period 5, before dramatically rising starting in Period 10. Periods 4 and 5 correspond to the periods of maximum density shown in Figure 7, and periods 11 and 12 correspond to the periods where density decreased and stayed relatively low.

Further case studies would be required to determine if this was a function of the schedule being stretched past its breaking point, or of the contractor no longer trying to recover an end date through RDL rearrangement.

As mentioned above, during Period 11 of the case study project, the contractor filed a claim notice and initial cost calculation for the first of four major issues to affect the project during its first year. The owner had already notified the contractor that it saw no merit in the entitlement for these claims; thus the two sides were already at an impasse.

Measuring Resiliency through RDL to True Logic Ratios

The percentage of RDL to Total Relationships might also be predictive of when resource flexibility is exhausted. In theory, RDL relationships could diminish as restrictions and limitations arise. In practice, for that to be true a contractor has to diligently remove all RDL that no longer has the potential to be a driving relationship, and contractors have bigger fish to fry than in cleaning out redundant relationships.

In the case study project, the RDL percentage increased as delays became impossible to work around, as shown in Figure 11. Until Period 10, the actual RDL percentage remained close to what was predicted in the baseline. After that point, the percentage of RDL relationships increased significantly.

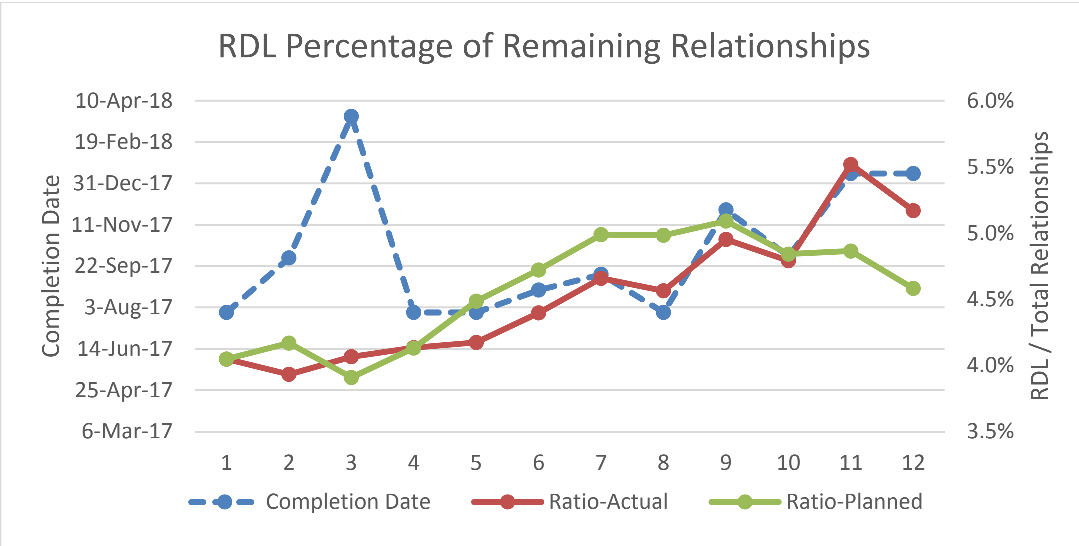

Figure 11: Ratio of RDL to Total Relationships. Only the relationships involving In Progress or Not Started activities are used in calculating the ratio.

Superfluous RDL relationships should increase the average RFF for those relationships, as more non-driving relationships would add more nonzero RFF data points to the calculation. This is confirmed by the similar timing of the actual plot lines shown in both Figure 6 and Figure 11.

Further case study evaluation is necessary before declaring this metric as an indicator of Resiliency. In addition, a nominal ratio for the case study, which is an urban Light Rail Transit project, may not be a nominal ratio for a different discipline, such as a water treatment plant. Therefore, instead of defining risk threshold values in terms of percentage of total relationships, they should probably be expressed as deviation from original.

Impact of Resource Flexibility on Monte Carlo Analysis

Monte Carlo analysis is a powerful tool gaining popularity as both owners and contractors seek to define the risks they're taking on when entering into a project. The last decade has seen a marked improvement in the ability of Monte Carlo Analysis to account for both project-specific risks and systemic risks. [12, p.14]

Monte Carlo Analysis works by varying activity durations in the performance of a series of iterative calculations. With the exception of probabilistic branching, no logic modifications are made during such an analysis. Probabilistic branching is the addition of activities and their accompanying pre-determined logic to a portion of the iterations. [13, p. 10] Figure 12 below depicts the 'missing link' between a Monte Carlo simulation and a fully accurate assessment of risk on a resource constrained project. The simulation necessarily must retain the logic it's given, while during the life of a project RDL may change.

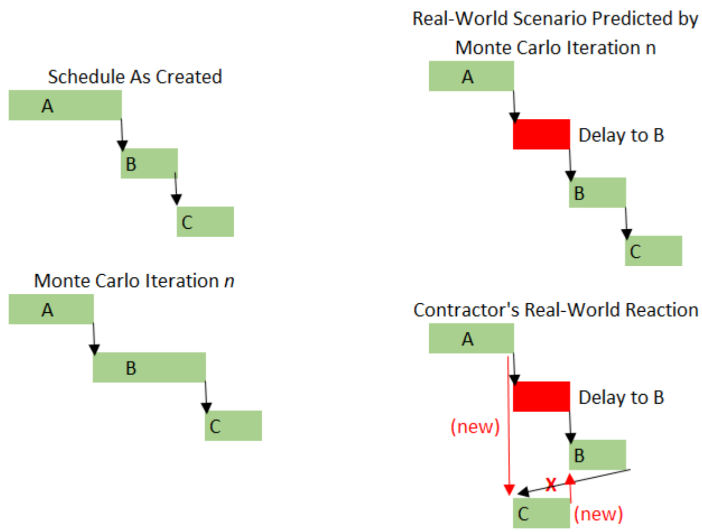

Figure 12: Demonstration of RDL's impact on Monte Carlo Simulation.

The fragnet of activities within Figure 12 all use the same resource and are linked with RDL ties. This fragnet demonstrates that within an iterative calculation of a Monte Carlo simulation, Activity B's duration is toward the pessimistic end of its range of durations. This iteration may actually predict how Activity B proceeds when the project happens—perhaps a two day delay prevents the contractor from moving its crew directly from Activity A to Activity B. However, Activity C would be available to the resource, and contractors prefer to keep crews productively occupied.

Because current Monte Carlo methods do not have a way to simulate the contractor's RDL changes, results of such an analysis may be more pessimistic than they should be. However, other recognized factors such as merge bias tend to also make Monte Carlo calculations more pessimistic and conservative, and experts often take that conservatism into account when making pre-project recommendations.

A Monte Carlo analysis performed on a summary-level schedule (such as a 100 activity Level II schedule that summarizes a 2000 activity Level III schedule) is less exposed to the vagaries of RDL. A typical Level II schedule activity assumes that the scope it captures is performed by a continuously occupied resource, so Level II schedules tend to have a low percentage of RDL.

In addition, conservatism would benefit many projects with a high degree of RDL, such as multi-segment transportation projects, which feature outside risks such as property acquisition and third-party utility relocations. A tendency for a Monte Carlo simulation to overstate completion times and costs would be counterbalanced by the degree of systemic risk these projects carry.

Conclusion

While there are currently no absolute methods of identifying a schedule's Resiliency to delay, this paper's case study has yielded several potential indicators. More research is required to conclude if these indicators are truly predictive of a schedule reaching a saturation point.

Differentiation of RDL relationships against True Logic relationships enables a more comprehensive evaluation of how much Resiliency remains in a schedule. Unfortunately, P6© is a poor tool for RDL evaluation. The software provides no means of designating relationships as RDL or True Logic. Asta Powerproject® does allow a user to identify and track a relationship as resource-driven, and it allows mass-selection of RDL. There is also a way for a user to 'toggle' RDL on and off prior to performing a schedule calculation, a feature that facilitates Labor Flex analysis.

Schedules with a low ratio of RDL to total relationships have less ability to absorb delay through resource rearrangement. More specifically, a schedule with no RDL on its critical path cannot use resource resequencing to recover a critical path delay.

Relationship Free Float on resource-driven relationships may also be indicative of when Resiliency is exhausted. It could be an indicator for when the contractor's current resource usage scheme has broken down.

Even if RDL is not differentiated, tracking the percentage of activities which are critical or near-critical is a useful indicator of how well the schedule absorbs noncritical delay. The month-to-month trend is a more powerful tool than an individual month's statistics.

The owner and contractor on the case study project would also have been better served paying attention to the tried-and-true Earned Value metrics of planned versus earned labor hours and Schedule Performance Index. The SPI for the case study project never exceeded 1, and it dropped significantly at a time where the contractor was concluding that the delays were irrecoverable. For the clearest picture of a project's likelihood of an on-time completion, Resiliency metrics described in this paper should be observed in conjunction with, and not in lieu of, SPI.

The case study seems to indicate that the owner and contractor had an opportunity at Month 8 to recognize what the metrics were indicating: Even though an on-time completion had been restored, Figure 6 showed that the average Relationship Free Float per RDL activity was at its lowest level to date. Figure 7 showed that Activity Free Float density was at its highest to date, at 15% higher than the baseline. And Figure 8 indicated a three-month trend of deteriorating SPI. The combination of those three metrics could have indicated that no more delay could be absorbed.

Unfortunately for both owner and contractor, the schedule put forth at Month 8 could not be followed. Had that been possible, the case study could have yielded interesting data about a project 'living on the edge' but still maintaining an on-time completion.

The above evaluations were performed on one case study project. Before concluding that the metrics are definitive indicators of whether or not a project can absorb a delay and remove its impact through rearrangement of RDL, more research should be performed. The hope is that this paper has sparked interest for follow-on studies and presentations, and that correlations will follow. Each addition to this body of data will enable a project's owner and contractor to more timely shift their energy from solving the unsolvable to determining realistic and equitable acceleration or time extension.

References

1. O'Brien, James & Plotnick, Fredric (2010). *CPM in Construction Management, 7th Ed.* New York: McGraw Hill.

2. United States Army Corps of Engineers. (2012). Specifications for Permanent Canal Closures & Pumps (PCCP). *Section 01320A, Project Schedule*. New Orleans, Louisiana.

3. Jackson, Peter; Vernon, Aaron; & Caddell, Christopher (2012). Schedule Compression Warning Tools. *AACE(R) International Transactions, PS. 1086.* Morgantown, WV.

4. Finke, Michael Ross (2000, June). Schedule Density as a Tool for Pricing Compensable Float Consumption. *Cost Engineering*, pp. 34-37.

5. Hall, Greg (2017). Improving the Update Process: Solutions For a Contentious Industry. *AACE(R) International Transactions, PS.2428.* Morgantown, WV.

6. Fondahl, John W. (1990). The Development of the Construction Engineer: Past Progress and Future Problems. *ASCE Annual Convention.* San Francisco, CA.

7. Patterson, Dr. Dan, (2010). *Taming an Unruly Schedule Using the DCMA 14-Point Assessment.* Austin, TX: Deltek Acumen.

8. Orr, John P. (2011, May). Timely Time Extensions: The Owner's Duty. *Cost Engineering*, pp. 22-27.

9. AACE International Recommended Practice 29R-03, *Forensic Schedule Analysis*, April 25, 2011, pp 51-58.

10. Loop, Robert, (2013). Schedule Metrics Reliability and Predictability. *AACE(R) International Transactions, EVM.1227.* Morgantown, WV.

11. Nosbisch, Michael, & Winter, Ron, (2006, July). Managing Resource Leveling. *Cost Engineering*, pp. 24-34.

12. Hulett, Dr. David & Whitehead, Waylon (2017, July/August). The Monte Carlo method for Modeling and Mitigating Systemic Risk. *Cost Engineering*, pp. 12-25.

13. Hulett, Dr. David & Nosbisch, Michael (2012, November/December). Integrated Cost-Schedule Risk Analysis. *Cost Engineering*, pp. 5-15.