Buffer Management

By Bill McClelland, Partner, AGI

Project Buffer Calculations

At the start of a project, required buffer sizes are determined according to a *Rule of Thumb* using the *Square Root of the Sum of Squares* as a safety net. The underlying principle is as follows: Once work behaviors and incentive systems have been adjusted to eliminate the wasting of safety and once the reporting of early finishes becomes a reality, organizations can count on a canceling effect between tasks that take shorter and tasks that take longer than their estimated ABP times. This means that less safety is required to protect the overall project than would otherwise be required to protect each individual task.

The Central Limit Theorem states that the sum of a large number of independent random variables with identical distributions will have a Normal Distribution. This means that 97.5% of the distribution will lie below 2 Standard Deviations (2σ) above the mean value.

In a perfect world, where all necessary Central Limit Theorem assumptions hold true, the approximate amount of safety required to protect a Critical Chain of length $CC = \sum_i ABP_i$ would be somewhere in the neighborhood of 2σ, where

$$2\sigma = 2\sqrt{\sum_i ((HP_i - ABP_i)/2)^2}$$

which simplifies to

$$2\sigma = \sqrt{\sum_i (HP_i - ABP_i)^2}$$

otherwise known as the *Square Root of the Sum of Squares*.

However, reality shows that the assumptions required by the Central Limit Theorem do not hold true.

1. The underlying distribution of each task is different.
2. Tasks are not truly independent.
3. We cannot count on the instant availability of resources to begin work as soon as a task is ready.

These deviations from the required assumptions must be taken into account in sizing the Project Buffer.

For long chains of tasks, the Rule of Thumb (50% of the safety removed) calculation leads to larger buffers than the S.R.S.S. and has a proven track record for providing adequate protection under these conditions. However, for short chains or, for chains characterized by extreme differences in task variability from one task to another, the Rule of Thumb breaks down. Depending upon the Covariance of the tasks in question, the R.O.T. actually indicates a need for less protection than use of the S.R.S.S. method. The R.O.T. calculation will always be smaller than the S.R.S.S. calculation for chains of 4 or fewer tasks. If the Covariance is large, this crossover point can even occur at 14 or more tasks.

When the S.R.S.S. method calculates a larger buffer requirement than the R.O.T., it is likely to mean that the R.O.T. no longer provides adequate protection. However, the underlying assumptions that led us to the R.O.T. approach haven't changed either. Thus, it is unlikely that the S.R.S.S. calculation will provide adequate protection. Not only are the Central Limit Theorem assumptions not valid but, at this point, the Critical Chain itself is short enough that we need to question whether there will be sufficient "aggregation" or "canceling" effect among tasks that exceed their estimated ABP times and tasks that take less than their estimated ABP times.

When the R.O.T. calculation is lower than the S.R.S.S. calculation, informed judgment should be used to calculate Project Buffer sizes. That is, a careful inspection needs to be made of the task work in question to determine the likelihood that the associated buffer will provide adequate protection. Some projects have only minor variability at the end; others conclude with highly variable tasks; while still others depend upon the simultaneous convergence of several highly variable pathways (where it feels like winning depends upon back to back to back homeruns).

Under these circumstances, consider the S.R.S.S. to be the minimum protection needed. Depending upon the task work under consideration, it might be advisable to increase the Project Buffer size above the S.R.S.S., perhaps up to the total amount of safety removed; $\sum_i (HP_i - ABP_i)$ or even beyond that amount for extreme situations. The insights of the project manager and project team are essential in this determination, as no absolute mathematical answer is possible.

Also note that the above discussion excludes considerations where a Project Buffer needs to be increased to account for anticipated project risk. This entirely valid need is something for consideration over and above any of the Buffer sizing mathematics.

Decreasing Size Buffer Management

As task work is completed along a project's Critical Chain during project execution, there is correspondingly less safety required to protect the committed due date.

The ideal software support would continually adjust the Project Buffer to match the safety required to protect what remains of the Critical Chain. Task color coding would be based on penetration by the Critical Chain into the Green, Yellow, and Red zones of the *currently required Project Buffer*.

Unfortunately, no TOC PM software completely supports this solution, so approximations are required. A typical approximation consists of pre-defining the y-intercepts of the Green, Yellow, and Red zones at two key points: 100% Critical Chain remaining and 0% Critical Chain remaining (see Fig 2, below). This is usually a single setting for the entire portfolio of projects.

The slopes of the Planning and Action Thresholds are calculated based upon an assumption that task variability is uniformly distributed across the entire length of each project, resulting in linear connections between the y-intercepts. If that assumption holds true for all projects in a portfolio and the y-intercepts have been correctly selected, then the software should provide appropriate buffer status signals and appropriate resource assignment signals at the appropriate times. As long as the underlying assumptions hold, this solution is superior to using Fixed Size Buffer Management (where the Project Buffer requirement does not change throughout life of the project) because it avoids the problem of frequent false Planning and Action Threshold crossings toward the latter part of projects.

Consider the following example: A project began with a Critical Chain of 600 days and a Project Buffer of 300 days. 675 days later there are 150 days remaining to be worked on the Critical Chain and 225 days of Project Buffer have been consumed. According to Fixed Size Buffer Management, penetration into the Project Buffer shows the Action Threshold has been crossed. Does this accurately reflect the current status of the project?

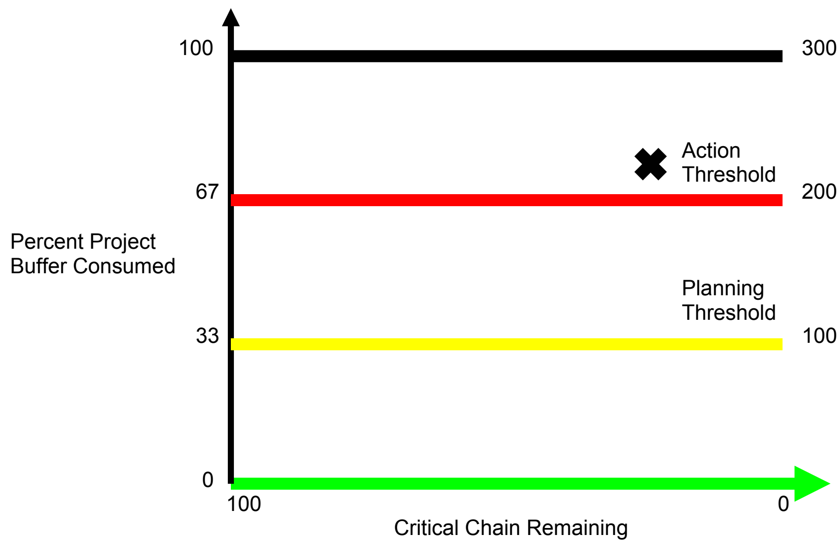

Figure 1 Fixed Size Buffer Management

Figure 1 shows the example's buffer status using Fixed Size Buffer Management. The indication of "red status" is clearly an erroneous signal. The project is currently projecting 150 days remaining along the Critical Chain, which is protected by 75 days of Project Buffer. Assuming that the variability was originally uniformly distributed across the project, 75 days should provide adequate protection for the remainder of the Critical Chain. Reality looks more like Figure 2, below.

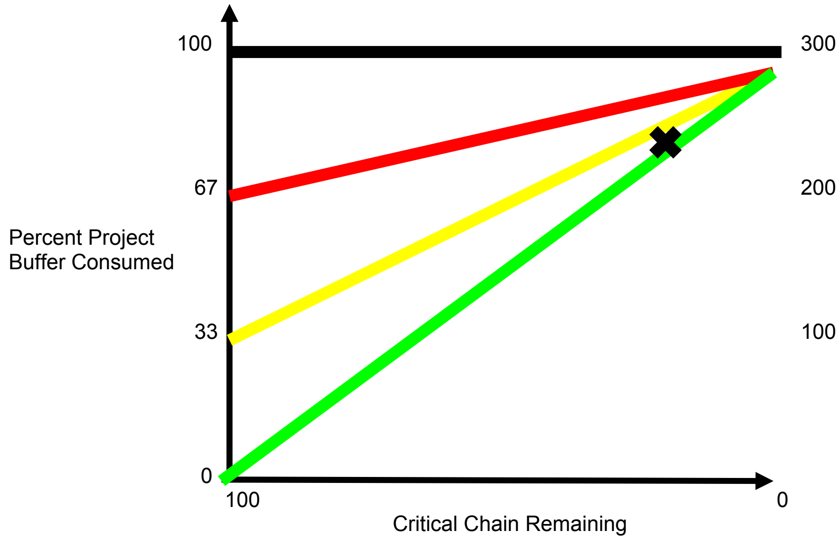

Figure 2 Approximating Decreasing Size Buffer Management

In a system managed according to Fixed Size Buffer Management, the latter part of a project is highly susceptible to increasing numbers of false calls to action. In a system managed according to Decreasing Size Buffer Management, these erroneous signals should be reduced.

The thinking here is built on a premise that in order to be effective, a project management system should strike a balance among three types of intervention risk:

> The risk of *taking action* when none is required
> (α risk, erroneous intervention),
>
> the risk of *not taking action* when action is required
> (β risk, erroneous non-intervention), and
>
> the risk of *taking the wrong action* (taking action in the wrong area…where the action will have no meaningful impact) (λ risk, erroneous intervention action).

When the underlying assumptions are valid (uniformly distributed variability; correct y-intercepts), the use of linear slopes will serve to reduce α risk considerably below that which would be experienced using Fixed Size Buffer Management. β risk will be higher, but tolerably higher, since it approximates 0 for most of a project using Fixed Size Buffer Management.

The key to the actual α and β risks lies with the validity of the underlying assumptions. Years of experience in capturing significant amounts of (ABP, HP) data across many different types of projects in many different types of environments indicates that variability is rarely uniformly distributed and, in fact, can vary significantly from project to project within a given portfolio.

When the variability is skewed within a project and that skew differs from project to project, there cannot be a “system-wide” α and β risk, but rather α and β risks that are localized to each project. However, until software properly supports Decreasing Size Buffer Management, project managers are left with the need to perform manual calculations as a standard part of building buffer recovery plans.

Buffer Recovery When the Critical Chain is the Dominant Chain

Recovery planning should begin when the software signals that Project Buffer consumption has crossed the Planning Threshold (the yellow linear slope). The determination of how much Project Buffer needs to be recovered (and when) should be based upon calculations of Decreasing Size Buffer Management (DSBM). This will provide more accurate signals than the linear slope approximation if the underlying assumptions are not valid for the project in question. If DSBM indicates it is actually too early to begin planning, α risk can be managed lower by putting planning on hold. If DSBM says that the required buffer is even closer to the red zone, β risk can be managed lower by speeding up recovery planning.

During recovery planning, one should also pay attention to the crossover point between the R.O.T. and the S.R.S.S., remembering the recommendation that, to provide full protection, the top of the Green Zone of the Project Buffer should not ever be less than the total amount of safety required ($\sum_i (HP_i - ABP_i)$) to protect all Critical Chain tasks that occur between the project completion date and the crossover point.

Recovery When the Critical Chain is No Longer the Dominant Chain

While delays along the Critical Chain are the most frequent causes of penetration into the Project Buffer, it is possible for delays in non-Critical Chain tasks to cause a chain other than the Critical Chain to become the dominant chain – the chain that is truly consuming the remaining Project Buffer. When this happens, the project schedule is no longer considered to be properly immunized. In fact, it is highly likely that there is insufficient safety and/or that safety is in the wrong place to protect the commitment date.

When such a shift occurs, regardless of whether the Decreasing Size Buffer Management Planning or Action Thresholds have been crossed, the Project Manager must immediately plan and take whatever actions are necessary to make the Critical Chain once again the dominant chain. The penetration of the Critical Chain into the Project Buffer at that point in time may or may not necessitate further planning or action.

If it is impossible to make the Critical Chain the dominant chain, the project should be re-planned immediately (all assumptions revisited, new Critical Chain schedule calculated, new Buffers calculated and placed, etc.) to ensure the proper amount of protection exists in the appropriate places to protect the project commitment date.