

Need for and Benefits of Additional Real-World Project Modeling Capabilities: Part 2

Robert Richards
Stottler Henke Associates, Inc.
1650 S. Amphlett Blvd., Suite 300
San Mateo, CA 94402
Richards@StottlerHenke.com

Abstract— This paper discusses additional useful real-world project modeling capabilities, expanding upon the similarly named paper from the IEEE Aerospace 2023 Conference. Project management tools provide various capabilities for modeling and analyzing projects or a portfolio of projects. These tools provide methods for modeling tasks, their relationships to each other, e.g., finish-to-start constraints, calendars, assignment of resources to tasks, and various other modeling capabilities. These capabilities provide a great foundation for modeling many types of projects, but they often fall short of capturing enough real-world details.

That is, there are many cases where without further real-world project modeling capabilities, the project managers will be overwhelmed by the inaccuracies of the resulting schedule during execution. We explore modeling capabilities that bridge this gap, many of which have been developed through decades of collaboration with NASA which has resulted in capabilities being added to the Aurora intelligent scheduling and project management software product, that NASA has licenses to. One key innovation is the concept of hazard constraints, allowing project managers to model potentially dangerous activities in relation to other tasks, ensuring hazardous operations never coincide with incompatible activities during execution.

The traditional concepts of lead and lag are often imprecise, typically relying solely on time-based measurements that fail to account for real-world delays. For instance, a lag might be defined to allow one task to begin after 20% completion of another task, but this definition is usually based on the original task duration. Consequently, if the task duration changes, the lag doesn't automatically adjust, leading to scheduling inaccuracies. This paper introduces enhanced options for modeling lead and lag relationships that better reflect real-world scenarios. One such improvement is the maximum lag constraint, which initiates a countdown or expiration period when one task starts, setting a time limit for the commencement of a related task.

In addition, sometimes it is important to have certain tasks scheduled as early as possible, while in the same project, other tasks should be scheduled as late as possible, for example, complete as much of the project as possible before delivering a very expensive piece of equipment that needs to be installed as part of the project.

Many projects require risk analysis in the form of Monte Carlo simulation. If a project utilizes one or more real-world project modeling capabilities, the Monte Carlo simulation must of course also support these additional project modeling capabilities to provide accurate results. Fortunately, the Aurora

software that NASA has access to has all the capabilities described in this paper and Part 1 of this paper, including performing full Monte Carlo analysis using any and all of the real-world project modeling capabilities.

By presenting these advanced modeling techniques and their practical applications, this paper aims to bridge the gap between idealized project plans and real-world execution complexities. These capabilities represent a significant advancement in project management, offering tools to help organizations like NASA navigate the challenges of complex, high-stakes projects with greater accuracy and confidence.

TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. SUMMARY OF BENEFICIAL MODELING CAPABILITIES DISCUSSED IN PART 1.....	2
3. SUMMARY OF BENEFICIAL MODELING CAPABILITIES DISCUSSED HEREIN	2
4. ENHANCED TEMPORAL DEPENDENCY CONSTRAINTS.....	3
5. VARIOUS PREFERENCE CONSTRAINTS	4
6. IMPROVED LEAD AND LAG CONSTRAINTS.....	5
7. MONTE CARLO SIMULATION	5
8. BENEFICIAL VISUALIZATIONS.....	6
9. CONCLUSIONS.....	7
REFERENCES	8
BIOGRAPHY	8

1. INTRODUCTION

Scheduling, at its most basic, is the process of assigning tasks to resources over time, with the goal of optimizing the result according to one or more objectives [1]. Scheduling is heavily used in construction, manufacturing, aerospace production, defense, and service industries in order to minimize the time and cost associated with the completion or production of small-to-large, simple to complex projects.

Project management tools provide various capabilities for modeling and analyzing projects or a portfolio of projects, providing methods for modeling tasks, their relationships to each other (e.g., finish-to-start constraints), the resources that

are available, assignment of resources to tasks, calendars for resources and tasks, and various other modeling capabilities. These capabilities provide a great foundation for modeling many types of projects, and in many cases, they provide all the modeling capabilities necessary to execute the project successfully.

However, there are many cases where without further real-world project modeling capabilities, the human project managers will be overwhelmed by the inaccuracies of the resulting schedule during execution since so many real-world details are not taken into consideration.

This paper discusses additional useful real-world project modeling capabilities, expanding upon the similarly named paper from the IEEE Aerospace 2023 Conference [2], drawing upon a diverse range of domains ranging from aerospace and manufacturing to pharmaceutical production and medical resident scheduling, demonstrating how better modeling leads to better outcomes and how generally applicable these modeling options are to most domains. These capabilities have been incorporated into the project management and intelligent scheduling tool, *Aurora*. *Aurora* [3][5] is a software application used by NASA, and NASA now benefits from many of these capabilities.

2. SUMMARY OF BENEFICIAL MODELING CAPABILITIES DISCUSSED IN PART 1

This section provides a brief overview of modeling capabilities that have been found useful across one or more project / production domains and are further discussed in [2].

- Ability to handle physical space constraints, including considering the creation and elimination of the space during the project.
- Ability to model human resources with details beyond just an occupation, such as occupation plus a set of specializations and/or certifications.
- Concurrent Constraints – Specify that two tasks need to happen at the same time in the schedule.
- Exclusivities/non-concurrent constraints – Specify that a task cannot happen at the same time as another task or class of tasks.
- Preferred Resources – Specify a preference order when defining a set of resources that are mostly interchangeable.
- Alternative Resource Combinations – Specify different combinations of requirements that could be used to complete a task, including variants with different durations.
- Ergonomic Constraints – Consider human physical limitations. For example, workers may have

limitations on how long they can work on their knees, both in one sitting and throughout the whole shift.

- Variable Duration Tasks – Specify that a task could use more people and get done more quickly, or fewer people and get done more slowly. Or get done more quickly with a more experienced person.
- Special Manufacturing/Shift Control Properties – This is a set of properties that allows the user to control how tasks interact with shift breaks (can it go between shifts, can it go from one day to another, does it have to complete a certain length of time before a shift ends, etc.)
- Outsource support – Provide graphical and tabular reports that help the user determine what is best to outsource. Productivity has been enhanced via a convenient interface for visualizing what tasks can be outsourced and providing a one-click option to outsource a task that also adjusts the actual model appropriately.
- Capacity Change Constraints – This allows the user to specify a relationship between a task and a resource. Some tasks may make a resource available (e.g., adding a space zone that can subsequently be used for work), while others may make a resource unavailable (e.g., installing panels that block access to a space zone).
- Jig Support – This is specialized support that ensures that a jig is assigned to a series of work and will be retained throughout the work statement.
- Tabular Editor – Most project management tools provide some kind of tabular view. Depending on the functionality, this tabular view can just be a read-only view, or as we have found, a powerful analytical tool and data transfer tool. By providing an Excel-style view of the model and schedule data, the tabular editor allows the user to easily filter, analyze, extract, and enter data. Experience has shown that a “referenced paste” feature that effectively performs a VLookup cross-referencing incoming data with data in the Tabular Editor on key matching criteria provides a robust transfer data option without perfect row alignment.

3. SUMMARY OF BENEFICIAL MODELING CAPABILITIES DISCUSSED HEREIN

A variety of additional project modeling capabilities greatly help project managers better understand their model and schedule.

- Enhanced Temporal Dependency Constraints - Preference Constraints – Beyond the standard

constraints with leads and lags, there are many ways to make even these more informative via justification and notes associated with any constraint, as well as absolute versions of the constraints that will be shown to be necessary for more realistic leads and lags and for other real-world purposes. Section 4 provides details.

- Part 1 [2], discussed a version of Preferred Resources. There are many cases where a preference constraint is also applicable, so preference constraints will be further discussed in this paper.
- Monte Carlo Simulation – This takes advantage of duration distribution information to simulate multiple executions of the schedule to show how things are likely to play out. This provides insight into how brittle the schedule is, how likely it will be to run late, etc.

4. ENHANCED TEMPORAL DEPENDENCY CONSTRAINTS


See Figure 1 regarding many of the enhancements discussed below.

- Normal, Concurrent, Non-Concurrent dropdown: The dropdown shown with ‘Normal’ in Figure 1 also includes the items ‘Concurrent’ and ‘Non-Concurrent,’ allowing a user to thus define concurrent and non-concurrent constraints also.
- <= dropdown: The dropdown also includes ‘=’ & ‘>=’ so all the standard constraint types can be defined, as well as the ‘=’ provides an absolute option. For example, the *absolute finish-to-start* (F=S) constraint means that the successor task must start immediately after the predecessor completes. This is further discussed below.
- Offset (Hours): Offset allows one to introduce lag time into the temporal relationship, with negative numbers corresponding to a lead. For example, Finish <= Start with an offset of one hour means that the start of the second task should occur at least an hour after the finish of the first tasks. Or, mathematically:

$$\text{finish}(\text{first}) + \text{offset} \leq \text{start}(\text{second})$$
- Max Offset (Hours): Represents a situation where the finish of the predecessor starts a countdown/expiration limit for when another task needs to start. This is further discussed below.
- Offset calendar: Most project management tools are inflexible regarding the calendar of the offset, using the calendar of the predecessor for example. This field allows for the offset to have its calendar defined so it is clear how it is determined.

- Justification: This field may be used to explain why this constraint has been built. This may be unnecessary for simple finish-to-start constraints, but as you use more of the options available for constraints it becomes more necessary to provide a justification so others will realize why all the settings are the way they are. This feature provides valuable context for decision-making, facilitating better understanding and informed adjustments during project execution, especially for complex constraints beyond simple finish-to-start relationships.
- Note: The note field is provided so additional notes beyond those that are appropriate for the Justification field may be added.
- Active: The active property simply determines whether or not the constraint is active; un-checking this will cause Aurora to act as if the constraint doesn’t exist.
- Constraint type: The standard is Precedence, and the option is Preference. A Preference constraint means it may be ignored in certain situations. There are many options for Preference constraints which will not be discussed.
- History: This section provides information regarding the edit history of this constraint.
- Mark Constraint as Reviewed: This button allows those with appropriate permission to complete the review process.

Let’s discuss the absolute finish-to-start constraint in a little more detail. It looks like the following in the constraint shown in Figure 1



A task definition includes the resources associated with the task; if the resource needs change during a combined task, then there is more than one task. For example, the painting of a room may be modeled as simply three tasks: a prep task with resources, a painting task with possibly different resources, and a drying task that still requires the room as a resource. The first two tasks will have a normal F<=S link between them. However, the paint task and the drying task will have an F=S link, since the paint will start drying immediately.

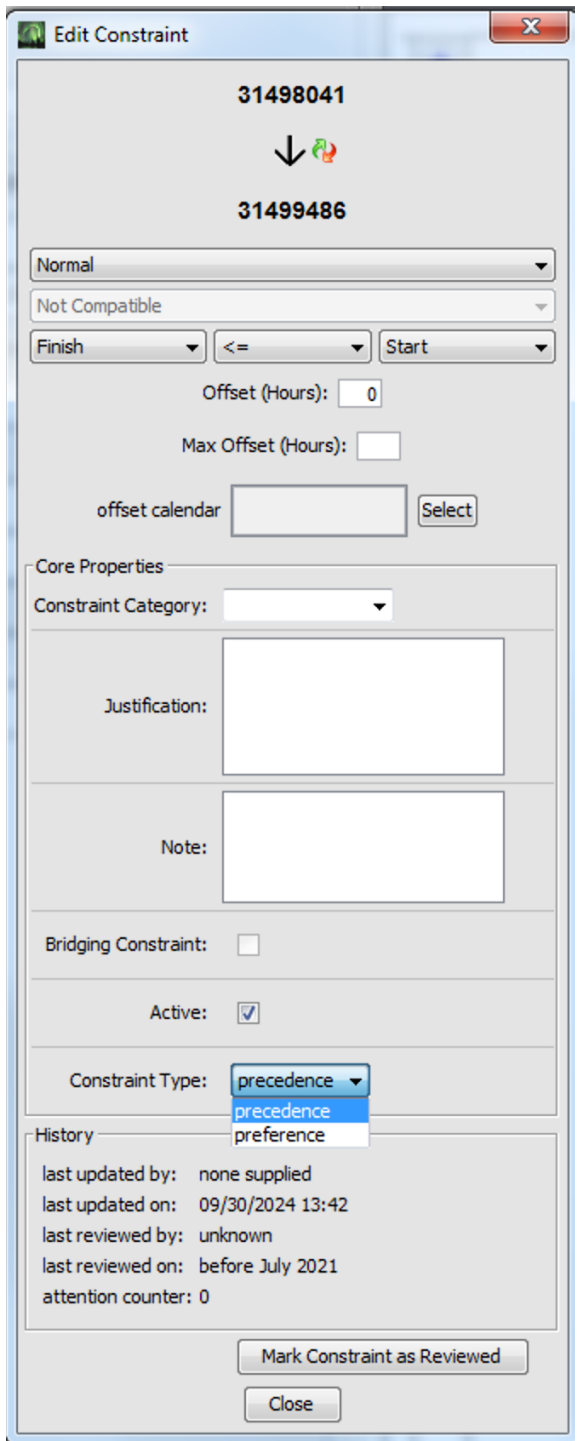


Figure 1. Enhanced Constraint Dialog

Another use for the F=S link is to store more information about the progress of a task at certain intervals of the task. Referring to the Justification and Note fields, these can be used to inform at certain stages of a task, or to further define what a % complete actually requires. Also, as further described below, it can be used to build a lead/lag capability that is more realistic and flexible for certain situations.

As shown in Figure 1, the Max Offset (Hours) option provides the option to define a maximum time that another task has to start after the finish of the predecessor. For example, in Figure 2, the task A3 must start within 25 hours of the completion of the ‘cluster end’ task. This functionality was requested by and initially implemented for Los Alamos National Laboratory to support their project needs.




Figure 2. Maximum offset example with the offset set to 25 hours

Of course, the Aurora scheduler will verify that the max offset is even possible with all the other constraints; for example, if tasks A1 and ‘A critical’ have a total duration of more than 25 hours, the max offset constraint can never be met.

5. VARIOUS PREFERENCE CONSTRAINTS

Possibly, the overall concept of preferences and the lack of their availability in project management tools is the largest contributor to project plan brittleness during execution. An initial plan may be better thought of as a guideline that will obviously need to be changed through time. The entire project plan is the best guideline that can be developed at the time of planning. Just as one may need to improvise during small projects, e.g., I prefer the appropriate socket size with a ratchet to tighten a bolt, if I find it is missing from my toolbox, it may be best to improvise with a crescent wrench or a vise-grip locking pliers, instead of stopping to go to the store to purchase a socket set; projects of all sizes need to adapt to the reality of the current situation to determine how to best move forward.

Of course, one cannot plan for every contingency, so a balance needs to be determined regarding which preferences or options should be modeled and which ones are best left to human decision-making.

For example, there are many cases where an activity may be associated with a specific organization, such as a shop; however, it is not absolutely necessary that the activity be performed by that organization. In this case, a preference constraint could be associated with activities that are preferred to be done with SHOP_A, but if SHOP_A is at capacity and SHOP_B or SHOP_C is not, then the scheduling tool would schedule the activity in SHOP_B or SHOP_C, maximizing utilization and throughput while maintaining activities in their preferred shop as much as possible.

Another example of a preference constraint derives from our implementation of Aurora at General Dynamics Electric Boat (GDEB) per the scheduling various aspects of submarine construction. There is a vast number and variation of components that flow through the machine shop. Oftentimes,

the components needed for an assembly will have a long lead time say, for example, up to six months. These components normally need to go through many different processing steps before they are considered complete; this may include milling, lathing, work in a multi access CNC machine, and welding. The part may be worked on by one machine and then there is a delay before it can be worked on again. GDEB has a preference that the delays between operations should not be “too long.” In addition, the overall time that a component takes to flow through all its steps should also not be “too long.” Working with machine shop personnel, Aurora was adapted to balance the need for overall potential throughput with keeping components, once started, flowing through the process at least as fast as the preferred as determined by actual machine shop personnel responsible for scheduling.

A related situation that occurs in the same GDEB shop is changing the preference for future work on components used in the same assembly if an assembly has been started before all the components have been fabricated. That is, some assemblies may consist of a large enough number of components that work can be started on assembling the components into a complete product even though all the components have not been manufactured. This occurs to keep the assembly workers busy and/or to increase the speed that the final product will be ready. The planning model is built so that all the components are predecessors to the assembly task(s). If the assembly is started before all the components are completed, then all the unfinished components make the assembly task(s) violate the predecessor-successor constraints. To handle the situation, the default configuration is defined as a preference, with assembly started, triggering the other option, in this case what occurs is: 1) the remaining components have their successor changed to the completion of the assembly milestone; and 2) the priority of the remaining components is increased. The priorities are increased, so that, if an assembly is started, it is normally desired to have it completed sooner rather than later.

Providing the capability to model Preference Constraints greatly complicates the scheduling process, which shows why most project management tools do not provide the option.

6. IMPROVED LEAD AND LAG CONSTRAINTS

The traditional concepts of lead and lag are often imprecise, typically relying solely on time-based measurements that fail to account for real-world delays. For instance, a lag might be defined with the intention that one task should begin after 20% completion of another task, but this definition is usually based on the original task duration. For example, if the original duration was 100 units of time, then the lag would be set to $20\% * 100 = 20$ units of time. Consequently, if the task duration changes, the lag doesn't automatically adjust, leading to scheduling inaccuracies. Unfortunately, the consequences of inaccurate lags (or leads) are usually detrimental in both directions. Let us take the case where the

original task is being completed faster than originally expected, so now it should be done in 75 units of time instead of 100. The lag task should now be able to start after 15 units of time instead of 20. With conventional project management software, the software would not realize this, and if the humans do not notice this then the lagged task will start 5 units of time later than necessary, possibly causing an unnecessary delay to the project. Alternatively, if the original task duration grows and 20% of the task is not completed after 20 units of time, the resources that we're expecting to start the lagged task may experience idle time waiting for the task to complete up to 20%.

Significantly more flexible lead and lag capabilities are provided via the already existing absolute-finish-to-start constraint. Using the example above, the 100-unit task would be divided into two tasks: one with a original duration of 20 units of time and the other with 80 units of time, with the tasks connected via an absolute finish-to-start constraint. This way, the original task is still to be worked from finish to end, and there is a distinct attribute in the original task that represents 20% complete. Another benefit of using an absolute finish-to-start constraint is that the constraint itself can include documentation to remind people in the future why it is there and how to understand what 20% complete means. Another benefit is the added flexibility in representing where any duration updates both shorter and longer, should be applied to the overall original task. That is, if the task duration changes, the change can be applied evenly to both subtasks if appropriate, or can be divided differently to each subtask, in either case the point for the original lag moves appropriately. The lag constraint would also need to be re-defined as a normal finish-to-start constraint after the 20% subtask. Therefore, this method eliminates a lag constraint, which is inherently more confusing to keep track because it has been converted to a more standard and understandable F-S constraint. A win-win.

7. MONTE CARLO SIMULATION

Monte Carlo simulation takes advantage of duration distribution information to simulate multiple executions of the schedule to show how things are likely to play out. This provides insight by assessing the impacts of various risks on project cost and timeline, thereby helping identify potential schedule and cost overruns.

Aurora provides both infinite resource and limited resource Monte Carlo analyses. One of Aurora's most powerful capabilities is to solve the non-deterministic polynomial-time hard (NP-hard) problem of resource-loaded scheduling more optimally than other solutions, while the solution time grows linearly with the size of the project [3][4].

The most important consideration is that, if a project uses ANY of the capabilities of Aurora that is not available in other Monte Carlo simulation tools, then the results will be incorrect. This is the main reason Monte Carlo analysis is provided in Aurora since users of Aurora almost always use some unique capability of Aurora and thus would be

prevented from receiving the benefits of accurate Monte Carlo analysis. Boeing is probably the real-world user that pushes this capability the most as they use many of the unique modeling capabilities mentioned herein and in Part 1 [2] and also performs risk analysis leveraging Aurora’s Monte Carlos simulation on these complex models.

One now has the option of building much more realistic models that adapt more correctly to real-world changes and then simulating these changes realistically and rapidly even when considering full resource-loaded schedules. Figure 3 shows the Monte-Carlo Simulation Options dialog for Aurora.

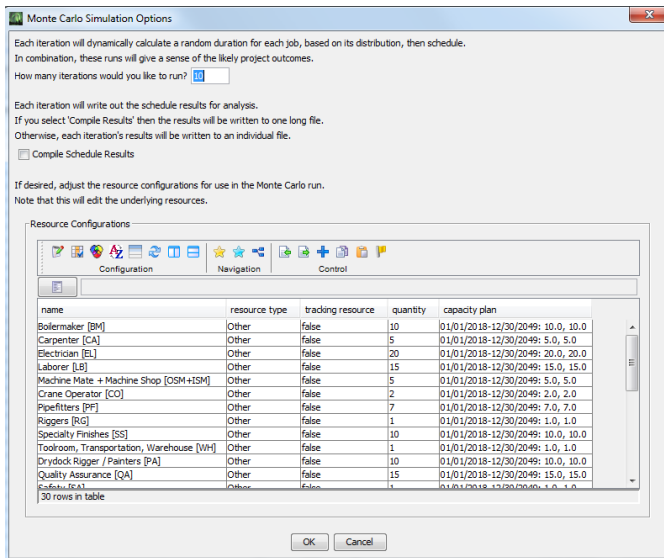


Figure 3. Monte Carlo Simulation Options

8. BENEFICIAL VISUALIZATIONS

As the project modeling capability expands, there is also a need to enhance the visualizations available to better comprehend the project model. This section discusses a subset of visualizations that have proven themselves.

Since we discussed all the constraints that may be required to best model, there needs to be options to distinguish them and to have the option to show or hide. See Figure 4.

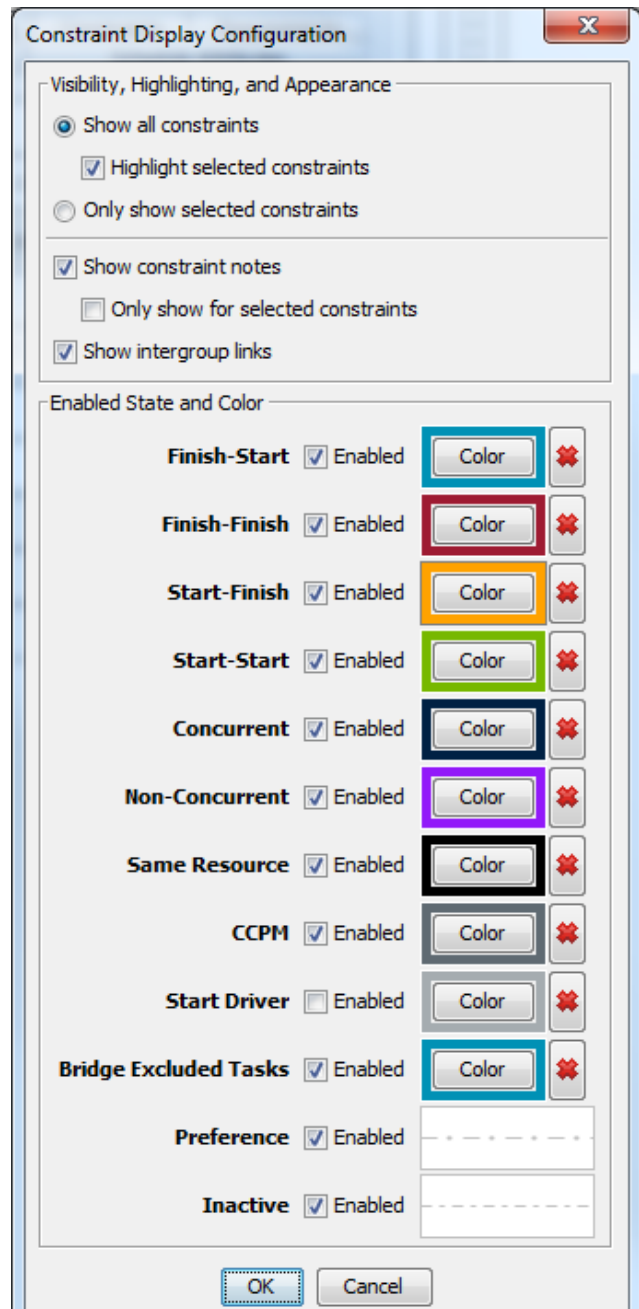


Figure 4. Constraint Display Options

When viewing large projects or multiple projects multiple options for viewing and navigating. Figure 5 shows an example of a large model with the entire model zoomed out to show the entire model in the mini-map overview box in the lower-right of the figure, with a red box (that looks almost like a red line since the model is so large) that represents the portion shown in the main view show behind the overview box. The user is able to drag the red box within this mini-map, and the regular network display will correspond by showing the region featured in the map.




Figure 5. Main Network Diagram with Overview Mini-Map Display Options

At the other extreme is the single-element view as shown in Figure 6. This allows one to view all the relationships relative to a task of interest. The top image is the element highlighted in Figure 5. In Figure 5, it is difficult to see the tasks connected to this task, where the single-element view makes it clear. In the Figure 6, the second image makes it more explicit that tasks related both via temporal constraints and those related to resource constraints are both shown. Navigating the path of interest is trivial also in the single-element view since any task can be clicked and it will move to the center and all of its relationships will be shown.

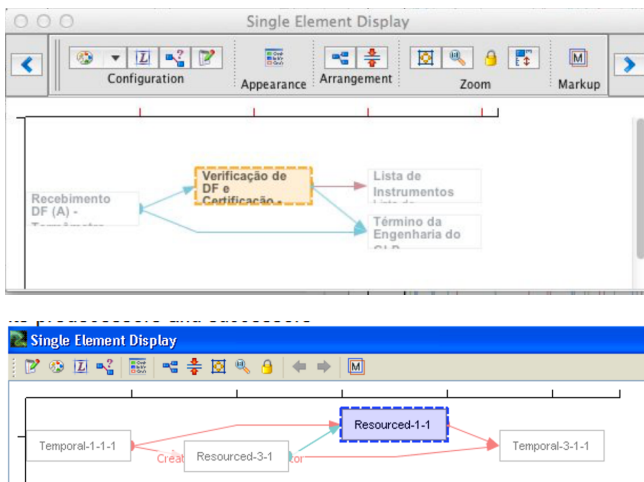


Figure 6. Single Element View

The single-element view is available in Gantt Charts, see Figure 7 and even directly from the Tabular Editor.

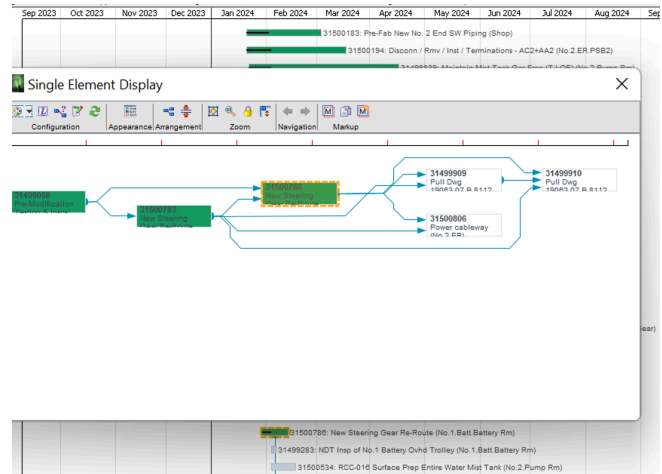


Figure 7. Single-Element View Opened in a Gantt Plot

Similar graphics are provided for the following useful capabilities.

- Upstream/Downstream Task Analysis – These analyses start with a given tasks or tasks and walks up/down the network to find the tasks it is dependent on, or the tasks that are dependent on it. The upstream analysis can help in understanding a key task and what it is dependent on; the downstream analysis can help in understanding a key task and what is impacted by it.
- Point-to-Point Analysis – This finds the path through the network from the first task to the second task (if there is such a path). It can be valuable for analyzing a subset of the network that is connected (e.g., all the work linking Milestone 1 and Milestone 2).

9. CONCLUSIONS

Stottler Henke has had and continues to have the opportunity to work on some of the world’s most complex scheduling challenges with organizations ranging from NASA to Boeing, Siemens, Spirit AeroSystems, the US Air Force, General Dynamics, Los Alamos National Laboratory, Lockheed-Martin, amongst others. Most project management and scheduling tools provide a good foundation of modeling and analysis capabilities, this foundation is unfortunately inadequate in many real-world situations. Also, unfortunate is that it is not easy to discern how inaccurate the model and analysis is when the required capabilities are absent. Thus, it is not clear how inaccurate one’s model and analyses are until experience during the execution phase reveals that the project management results are not accurate.

The human project managers try to deal with these inaccuracies, by constantly having to adjust the schedule during execution. However, since the model does not

represent reality accurately enough to adapt to real-world changes, future predictions regarding how the rest of the project will progress will still be woefully inaccurate.

Again, experience has demonstrated the real-world need for and the diversity of Preference Constraints, and without this modeling capability, project efficiency may be severely degraded. A preference constraint can have many different incarnations, ranging from preferring a more experienced worker in certain situations to preferring one type of machine versus another. The result of having a preference constraint might allow two tasks that both prefer a certain skill level or machine to occur in parallel since the other task may still be allowed to work with the less preferred option. Specifically, the task on the critical path might get the preferred resource, while a task that has few downstream dependences might be allowed to use the less preferred resource since of the duration is longer than expected or some re-work is needed, so this has little chance of affecting the overall project duration in the negative, while doing the non-critical work with the less preferred resource is likely to benefit the project throughput. Without this capability, the resources remain unnecessarily idle, and project completion could be unnecessarily delayed.

Of course, there is always a balance between the amount of detail in a model and what should be left to the human project managers and schedulers. There are so many variables in real-world situations that we need to strike the correct balance between computer modeling and human knowledge. The goal should be to allow the humans to model in the project management/scheduling tool to the level of detail they determine is best, which then frees the humans to perform higher value project management and scheduling.

REFERENCES

- [1] Pinedo, M. L. *Scheduling: Theory, Algorithms, and Systems*. Cham: Springer International Publishing, 2016.
- [2] Richards, R. "Need for and Benefits of Additional Real-World Project Modeling Capabilities". *Proceedings of IEEE Aerospace Conference 2023*. Big Sky, MT, March 4-10, 2023.
- [3] Kalton, A. "Applying an Intelligent Reconfigurable Scheduling System to Large-Scale Production Scheduling," presented at the International Conference on Automated Planning & Scheduling (ICAPS) 2006, Ambleside, The English Lake District, U.K., 2006.
- [4] Kalton, A., R. Richards. (2008) *Advanced Scheduling Technology for Shorter Resource Constrained Project Durations*. AACE International's 52nd Annual Meeting & ICEC's 6th World Congress on Cost Engineering, Project Management and Quantity Surveying. Toronto, Ontario, Canada. June 29 – July 2, 2008.
- [5] Xu, S., R. Stottler, R. Richards, *Intelligent Scheduling at NASA: Application to Ground Operations at Kennedy*

Space Center. *Proceedings of IEEE Aerospace Conference 2017*. Big Sky, MT, March 4-11, 2017.

BIOGRAPHY



Robert Richards received a Ph.D. in Mechanical Engineering from Stanford University. Dr. Richards is managing and has managed multiple projects for both commercial and government clients, including various intelligent scheduling implementations. Dr. Richards is the Principal Scientist and Manager of Stottler Henke's General Dynamics Electric Boat project for scheduling submarine construction. Dr. Richards also leads the project for adapting Aurora to optimize processes at Los Alamos National Laboratory. Dr. Richards has also worked on and continues to work on various projects spanning a wide range of research and application area interests, including training system development; applying artificial intelligence techniques; and decision support tool development for life-critical situations. Dr. Richards has publications in all of these domains