Leveraging PHM in Conjunction with Intelligent Scheduling to Improve Manufacturing Resilience

Robert Richards
Stottler Henke Associates Inc. (SHAI)
1650 S. Amphlett Blvd., Suite 300
San Mateo, CA, 94402
richards@stottlerhenke.com

Jim Ong
Stottler Henke Associates Inc. (SHAI)
1650 S. Amphlett Blvd., Suite 300
San Mateo, CA, 94402
ong@mail.stottlerhenke.com

***Abstract*—The scheduling of a manufacturing facility is a complex endeavor even when the equipment resources are always considered available or at least available 24/7 except for planned maintenance. However, under real-world conditions, the added complexity of unplanned downtime can significantly increase the difficulty of meeting deadlines. More reliable and efficient operations can be achieved by predicting problems and then rescheduling operations to minimize or avoid the problems' adverse effects. Stottler Henke Associates Incorporated (SHAI) has been working with the US Air Force on best practices for leveraging diagnostics, prognostics, and health management in conjunction with intelligent scheduling to improve manufacturing system resilience.**

The goal, then, is for diagnostic systems to identify impending faults quickly and automatically, providing the information needed to the intelligent scheduling system in order to minimize or completely mitigate the issues. Prognostic systems can estimate impending failures or rates of performance degradation; the intelligent scheduling system uses these diagnoses and predictions, along with the manufacturing deadlines and priorities, to develop mitigation strategies to minimize or avoid disruptions. The strategies include scheduling offline operations optimally to minimize the effects of the machines being offline; rescheduling operations due to machines being offline, reconfiguring systems to change their capacity and performance profiles, and/or reduce the usage of critical equipment to lengthen their remaining useful life. In the most uncomplicated cases, the schedule can be adjusted using simple strategies such as reassigning tasks from the faulted equipment to other equipment with similar capabilities. However, in many cases, more global analysis of an adjustment of the schedule is necessary to satisfy the facility's deadlines and other manufacturing goals.

A system for resilience needs to be able to model and simulate the manufacturing system. That is, scenarios need to run to evaluate the manufacturing system's response to various types of problem scenarios and analyze the effectiveness of responses. This evaluation capability can be used to compare resilience strategies that specify optimal policies for employing diagnostic and prognostic capabilities and for responding to current and projected faults via rescheduling and reconfiguration. This will provide insight into the most critical pieces of equipment as related to unplanned downtime. That is, when operating under normal conditions, the most constrained pieces of equipment, may be different from the equipment that has the most significant adverse effect if it was unavailable. Intuitively this may be difficult to understand, take the example that every piece manufactured goes through one of two of the same machines that are running at full speed, but there is a 3rd machine that performs a specialized operation on a subset of very high-value parts. In this case, the two machines are the constraint on overall throughput, but it may not be evident that the 3rd machine, if down, is the constraint on profit. Being able to run scenarios will surface actual cause and effect relationships.

This paper will expand on these ideas and lessons learned from real-world application of these ideas.

978-1-7821-2734-7/20/$31.00 ©2020 IEEE

TABLE OF CONTENTS

1. INTRODUCTION

Resilience is "the ability of a system to withstand potentially high-impact disruptions, and it is characterized by the capability of the system to mitigate or absorb the impact of disruptions, and quickly recover to normal conditions." [1]. By leveraging the results of prognostics and health management (PHM) of equipment in conjunction with intelligent scheduling the resilience of manufacturing processes can be increased.

Prognostics and Health Management (PHM) is a technology to enhance the effective reliability and availability of a product in its life cycle conditions by detection of current and approaching failures. Prognostics is the real-time enhancement of reliability and availability and the prediction of the remaining useful life of the product by assessing the extent of deviation or degradation of a product's monitored parameters from its expected normal operating conditions. Prognostics can yield an advance warning of impending failure in a system, thereby enabling more efficient and effective maintenance and corrective actions [2].

Predictive maintenance [3] is a maintenance management method that leverages PHM to determine maintenance actions by using the actual equipment condition, versus average statistics about equipment failure rates, in order to optimize total plant operation. The common premise of predictive maintenance is that regular monitoring of the actual mechanical condition, operating efficiency, and other indicators of the operating condition of machines and process systems will provide the data required to ensure the maximum interval between repairs and minimize the number and cost of unscheduled outages created by machine failures. That is, since predictive maintenance techniques are designed to help determine the condition of in-service equipment in order to estimate when maintenance should be performed, greater utility and cost savings are possible versus routine or time-based preventive maintenance, as maintenance tasks are only performed when warranted. Predictive maintenance is also known as condition-based maintenance since the actual measured condition of an item is used to determine when to perform maintenance. This contrasts with other maintenance management methods such as:

- Run to failure – Fix the machine only after it breaks down.
- Preventive maintenance – Preventive maintenance tasks are performed, based on elapsed time or hours of equipment operation. This method assumes similar MTBF for machines of each type.

2. SCHEDULING

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 [4], many times the objective is to maximize throughput when using a limited set of resources. Scheduling is heavily used in aircraft maintenance to minimize the time and cost associated with the completion of multiple competing maintenance projects. The Aurora scheduling framework is one example of a general-purpose intelligent scheduler that has been successfully applied to a variety of domains [5], [6], ranging from aircraft production and submarine maintenance. Intelligent scheduling combines graph analysis techniques with heuristic scheduling techniques to quickly produce an effective schedule based on a defined set of tasks and constraints [7]. The constraints include at least the following:

- Temporal: Tasks must be scheduled between the project start and end dates; each task has duration and an optional start date and an optional end date.
- Calendar: Tasks can only be scheduled during working shifts; tasks cannot be scheduled on holidays.
- Ordering: Tasks can optionally be assigned to follow either immediately after/before another task or sometime after/before another task; optionally with a specific offset time in between.
- Resource: Each task can require that resources be available for the task to be scheduled. Examples of resources include people with specific skillsets, e.g., machinist, equipment, e.g., CNC 4-axis mill, and physical space, e.g., the tasks that can only be performed in a specific location.

Human Derived Heuristics: Importance of

Scheduling is an NP-complete problem, that is, the size of the solution space grows exponentially as the model grows linearly and therefore problems of any reasonable size cannot be solved simply mathematically. Most 'solutions', such as resource leveling, use a simple algorithm, and thus result in far suboptimal results. Stottler Henke Associates Inc. (SHAI) has employed a strategy that includes leveraging scheduling heuristics learned from many of the world's best human schedulers in order to solve complex scheduling challenges in reasonable amounts of time. See in Figure 1 that included in the many components involved in intelligent scheduling, is the important Human Directives.

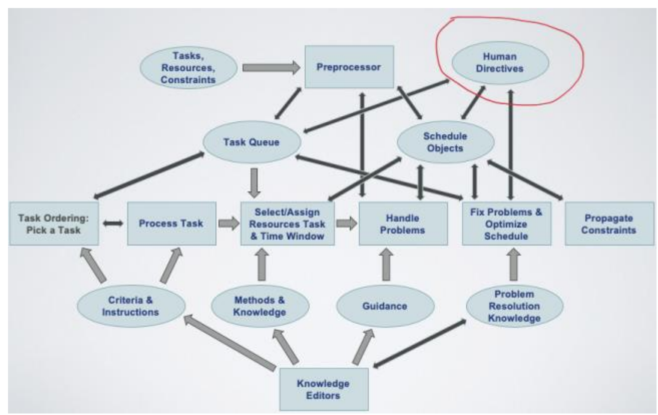

Figure 1. Aurora intelligent scheduler architecture

Consider the following extremely simple example (which is therefore easier to use to illustrate this point) where:

- three activities, called Activity 1, 2, and 3, from three different orders are all competing for time on similar machines in a particular work center.
- The priority is highest (or the due date is soonest) for Activity 1 and lowest for Activity 3.
- Two different machines exist, A which is expensive and precise and B which costs less and has higher throughput.
- Machine A is required for Activity 3, but it can also process activities 1 and 2, though it is not efficient to do so.

Let's look at a solution from a simple scheduler: Activity 1 is chosen first for assignment, since it has the highest priority, and it so happens that at the moment Activity 1 can begin, only Machine A is available, so Machine A is assigned to Activity 1. Activity 2 is assigned to Machine B, which has become available soon after Machine A. Activity 2 is soon completed, owing to Machine B's fast production rate. When Activity 3 is finally examined, its required machine, Machine A, is busy and, worse, busy on an activity that it wasn't essential for. Meanwhile Machine B is idle.

Obviously, this is a suboptimal solution since a different assignment would have prevented Machine B from being idle and prevented expensive Machine A from being assigned to a task that didn't need it. Of course, a more complex scheduler could "look ahead" to see if the cheaper machine might be soon available, but for any such workaround there's a corresponding example that still causes problems. And each of these rules has to be anticipated and created by the scheduling system software developer.

Perhaps a scheduling system could be written that systematically tried every possible solution and selected the best, and therefore optimal, one. In the example above, the number of possible solutions is 2 choices for Activity 1 times 2 choices for Activity 2 times 2 choices for Activity 3 = only 8 possible solutions. However, consider an activity list consisting of only **30** simple resource assignments where (for simplicity's sake) only one resource is required for each activity. Assume on average 4 meaningfully distinct choices (e.g. different machines) for each activity. This means that there are 30 distinct decisions with 4 choices each, so the number of solutions is 4 x 4 x 4 .... x 4 =

$\mathbf{4^{30}}$ = over a million trillion possible solutions,

which are clearly impractical to systematically search. And this calculation was based on an extreme oversimplification. The more realistic, complicated planning problem is much more difficult. This is the essence of NP-Complete problems. The widely recognized and clearly applicable NP-Completeness Theorem states that to guarantee an optimal solution to an NP-Complete problem, it would requires exponential time (e.g. M^N where M is the average number of options per choice and N is the number choices) which is clearly impractical in this case, since N is typically in the thousands. An optimal solution can simply not be guaranteed for this application.

Therefore, to determine near-optimal solutions in reasonable timeframes requires heuristics learned from actual human experts on a large number of situations. We have developed both general heuristics for producing good solutions and the techniques and architecture to incorporate domain specific knowledge and heuristics into the planning system. Our expertise includes substantial experience eliciting the required knowledge and cognitive processes from expert planners, then mimicking those processes in software to create advanced intelligent planning and scheduling systems. To wit, Aurora mimics the *decision-making process of expert schedulers.*

3. INCORPORATING PHM DRIVEN MAINTENANCE WITH OVERALL PLAN SCHEDULING

To be explicit, the PHM required to produce the predictive maintenance knowledge is non-trivial; that is, PHM is itself a very complex endeavor, however, for the purposes of this research we are assuming that the ability to predict future equipment failures is robust and will be assumed as an input.

The situational awareness provided by the predictive maintenance informs the scheduling system of the current and projected state of the various equipment / devices within the plant. Each resource in the scheduling system has a calendar associated with it, and the equipment / devices are just one type of resource. The projected state of each device drives the *calendar* associated with the device. For example, the human resources also have calendars representing such things as known vacation days. Just as the schedule will need to adapt to updates to personnel vacancies, the schedule needs to update based on predictive maintenance information. A benefit of predictive maintenance is that in most cases the prediction will provide a range of time for maintenance. That is, a machine may be shown to need maintenance within the next 100 hours of operation, otherwise the risk of failure is too high. Therefore, the specific time within the next 100 hours when actual maintenance occurs is flexible and the scheduling system can be used to mitigate the disruption to the overall schedule. Furthermore, sometimes when one machine greatly affects the utilization capability of other machines, it can be investigated whether performing upcoming maintenance on the other machines is advantageous. That is, even though other machines may not require maintenance yet, it would still be advantageous to perform routine maintenance early so the clock is reset, and the machines will run farther into the future without requiring maintenance again. Figure 2 shows a sample calendar, and its visual display.

Therefore, the intelligent scheduler is invoked to try to determine alternative schedule(s) that meets the production demands. The alternate schedule might assign different equipment resources, if available, to perform the task, and/or it might reschedule one or more tasks as needed. Although simple changes to the schedule are usually preferred, an intelligent scheduling system can generate feasible (or acceptable) schedules, even when significant changes are necessary.

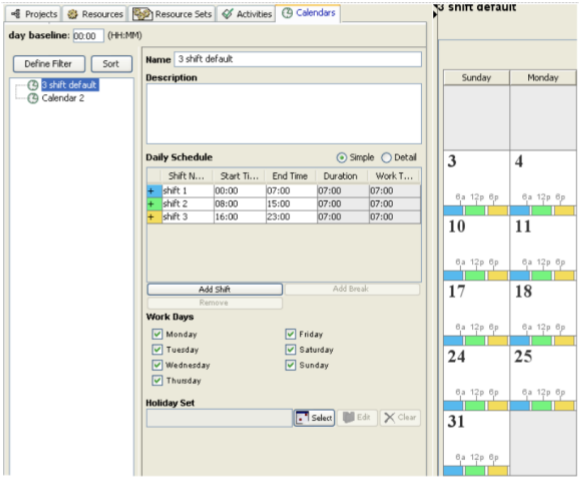

Figure 2. Calendar view

Lessons Learned

SHAI has been working with the US Air Force on best practices for leveraging diagnostics, prognostics, and health management in conjunction with intelligent scheduling to improve manufacturing systems. The lessons learned are mainly related to what aspects of the scheduling system need to be adapted to properly adapt to predictive maintenance knowledge.

Human Derived Heuristics—As described above, intelligent scheduling requires leveraging knowledge learned from human schedulers. When taking into account predictive scheduling, additional heuristics, learned from schedulers in the same domain as where the predictive maintenance occurs is necessary. This is due to the fact that the types of tradeoffs that are preferred cannot be determined mathematically in most situations. Of course, in actual situations, the heuristics in the system can be overridden by the human schedulers at the time, but it greatly improves the utility of the scheduling system if the scheduling system can provide useful suggestions and complete options for the human scheduler to draw from. For example, it may be normally either preferable to work overtime, or to try to outsource certain work, due to machine downtime; so, it is helpful for the automated intelligent scheduler to know which is preferred. Of course, this could get more specific by saying certain machines when they are down, are preferred to be mitigated via overtime, where other machines when they are predicted to be down it is best to be mitigated via outsourcing. Thus, the US Air Force heuristics may be different from the heuristics used in other facilities.

Preference Constraints—We also learned from our experience with the US Air Force that when dealing with an environment that will be taking into account predictive maintenance, it is important to have all of the flexibility that is available in the real world, modeled in the scheduling system. One important aspect of this is what we refer to as *preference constraints*. A preference constraint represents situations where different options are available, but there is a preference for one option over another. For example, there may be a CNC milling machine that is much faster than a manual milling machine, even though the manual milling machine has the capability to do the same operations. So, for our predictive maintenance situation, we would want the CNC milling machine to be the default, and the manual milling machine to act usually as a backup. Therefore, the ability to set a preference between resources is very important. That is, a resource pool could be created for this type of milling operations and include both machines with a preference set for the CNC machine. Now, when the scheduling system sees the situation where the calendar shows that the CNC milling machine is unavailable, it will try to then use the manual milling machine instead. This flexibility may be able to keep the projects on schedule. However, depending on the situation, it may take more changes, such as working extra shifts, or using other slower manual milling machines to keep the projects on time, but again, other capabilities in the scheduling system will help one determine if this is necessary.

Simulation and Monte Carlo Analysis—The significant benefit of predictive maintenance is that there is time to plan for the maintenance and the maintenance itself may be possible to take place over an extended period of time. For example, the maintenance may be necessary in the next eight weeks. This allows one to determine when in the next eight weeks it would be the least disruptive. Let's also assume, for the sake of argument that the maintenance takes a total of 48 hours of calendar time; depending on the ramifications of the downtime, analysis might show that all of it can be performed during the normal day shift, in which case this would take six calendar days to complete. This provides the added flexibility that if six days of downtime cannot be accommodated, the minimum amount of downtime on the calendar could be two calendar days, if the repairs were conducted 24 hours a day. Thus, one can run various simulations to determine what is the cost benefit from different repair schemes. That is, is the extra overtime, or nighttime cost for the repair worth it, due to increased throughput and the potential increased profit. Another benefit of having this known eight week time span is that one can determine where in that time span to conduct the repairs, utilizing this knowledge one might be able to build up buffers, or inventory, working the machine more than normal, so that downstream activities will not be starved.

4. VISUALIZATION

Experience with the US Air Force and others had demonstrated that visualization is very important to help the schedulers better understand the situation, and the options available to them. Since the scheduling system, in many cases will not be able to determine a complete solution on its own, human schedulers will work in conjunction with a scheduling system and its visual output to come up with the best option. For example, in situations where there is a long string of steps to complete a deliverable, one may be able to visualize the work in progress at different stages and visualize the ramifications of delays or stoppages at different stations. Overall, the graphs in the visualization in Figure 3, shows what happens over time when there is a stoppage at a specific step.

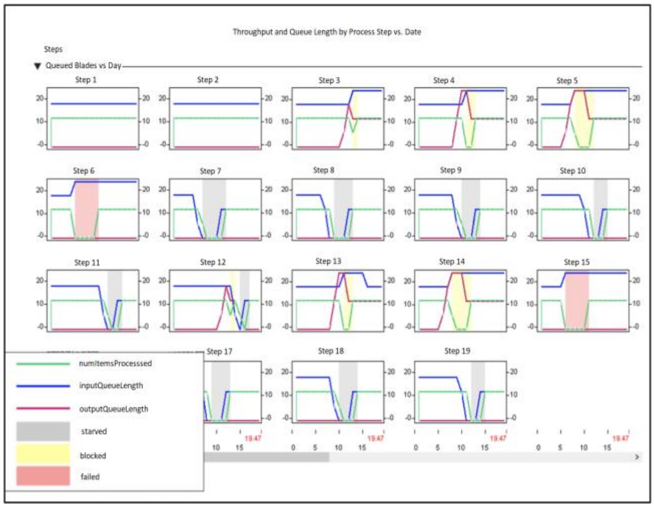

Figure 3. Delays caused at various steps

Visualization shows how steps upstream from the inoperable equipment will eventually be blocked from producing more, in this case because there is nowhere to store their output. Similarly, stations downstream from the blockage will not be able to continue working because they are starved for supplies. If plans can be made ahead of time, knowing that the situation is going to occur, storage areas may be able to be set up, so there is a buffer, possibly before and after the station that will be offline. If, for example, downstream tasks may be faster than the slowest task in the process, which let's assume is upstream from the task that needs to go down from repair, one could build up a buffer to hold inventory before the down machine so that it will build up during the time the machine is down, but not build up so much that it will never be able to clear because when the machine comes back online, the upstream machines will still be producing more products. Also, it's possible the upstream tasks and the steps that will be going down, will try to overproduce so that a buffer after the machine of inventory can be created, so that when the machine is down, this inventory will be drawn down so the downstream tasks can stay partially or completely busy during the machine downtime. Simulation and visualizations will help the factory understand what options are available., and what is the best option to maximize throughput based on detected downtime of a machine.

Understanding how the equipment that is predicted to require maintenance is being utilized overall is important in understanding the ramifications of its unavailability. A histogram helps surface this information, see Figure 4.

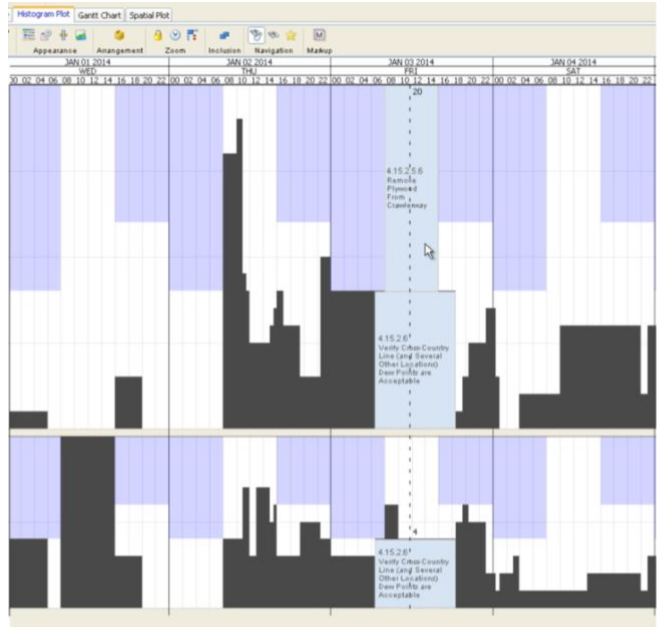

Figure 4. Histogram plot

In addition, the use of filtering and color coding can also help to better explain how the equipment of interest is being used. For example, referring to Figure 5, the items in red may be designate the situation where there is substitute equipment, but the tasks are now scheduled to take longer than when the downed equipment would be used.

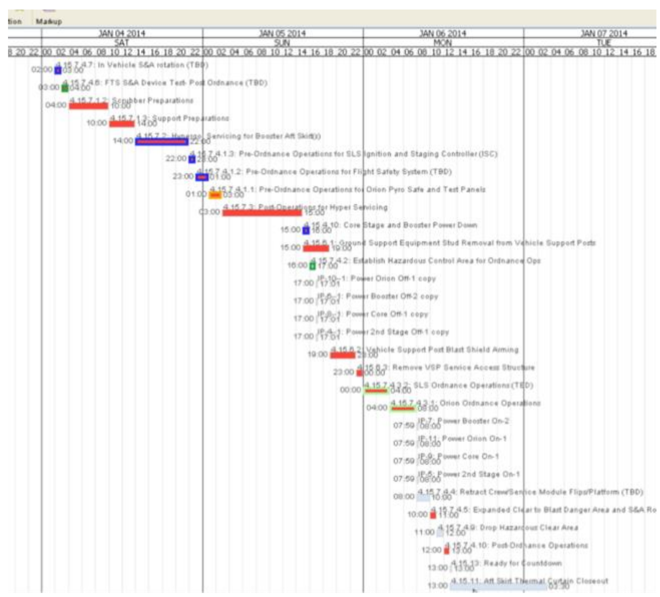

Figure 5. Gantt chart with color coding

Similarly, color coding helps reveal where equipment of interest is currently used in a network diagram. For example, see Figure 6 and Figure 7, looking at the network diagram at different levels of zoom, one could designate tasks that use the equipment of interest in orange to quickly visualize how much it is used overall in the project.

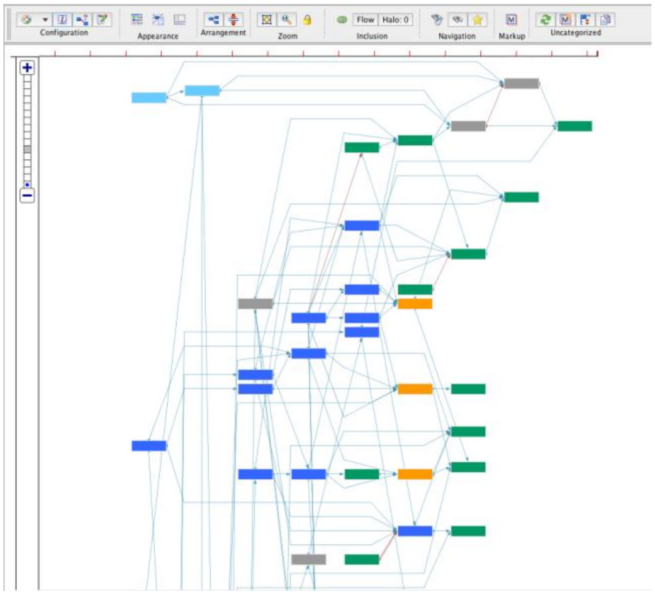

Figure 6. Network diagram w/ color coding

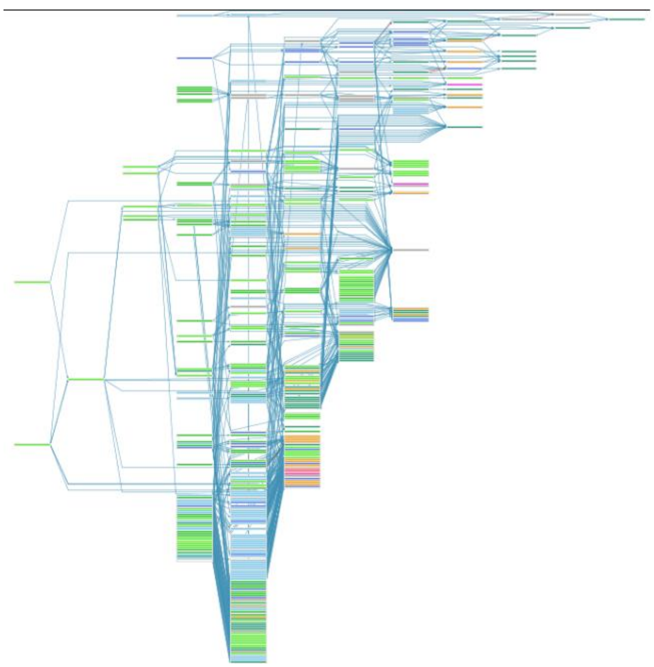

Figure 7. Network diagram w/ color & zoomed out

5. Conclusions

Scheduling a manufacturing, production or maintenance facility, even when assuming all the equipment is available 24/7, is already an incredibly difficult scheduling challenge; incorporating the reality of machinery maintenance only adds to the complexity. Preventative maintenance diagnostic systems help to mitigate some of this complexity by identifying impending faults the intelligent scheduling system has the opportunity to minimize, or completely mitigate the effects, either fully automatically or in conjunction with human schedulers. Prognostic systems can estimate impending failures or rates of performance degradation; the intelligent scheduling system uses these diagnoses and predictions, along with the manufacturing deadlines and priorities, to develop mitigation strategies to minimize or avoid disruptions. The strategies include scheduling offline operations optimally to minimize the effects of the machines being offline; rescheduling operations due to machines being offline, reconfiguring systems to change their capacity and performance profiles, and/or reduce the usage of critical equipment to lengthen their remaining useful life. If the scheduling system cannot mitigate the required maintenance satisfactorily via its own strategies learned from humans and tested via automated Monte Carlo analysis, the results of the scenarios and the visualizations of the results will aid the human scheduler in determining the best plan forward in a mixed initiative manner. By visually surfacing knowledge about the ramifications of different scenarios, the human scheduler will be better informed to decide what future scenarios to test, resulting in maximizing the throughput with the least amount of effort by the human scheduler.

References

[1] Gu, X. et al. (2015), 'Manufacturing System Design for Resilience', *Proceedings CIRP 25th Design Conference Innovative Product Creation* (2015), 135-140. Web site: https://www.researchgate.net/publication/283948933_Manufacturing_System_Design_for_Resilience

[2] IEEE International Conference on Prognostics and Health Management Website: http://phmconf.org/phmtopics.html

[3] Mobley, R. 2002, 'An Introduction to Predictive Maintenance', 2nd ed. Butterworth-Heinemann.

[4] M. L. Pinedo, Scheduling. Cham: Springer International Publishing, 2016.

[5] A. Kalton, "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.

[6] R. Richards, "Critical Chain: Short-Duration Tasks & Intelligent Scheduling in e.g., Medical, Manufacturing & Maintenance," presented at the 2010 Continuous Process Improvement (CPI) Symposium, Cal State University, Channel Islands, 2010.

[7] 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.

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. Dr. Richards is the Principal Scientist and Manager of Stottler Henke's Pfizer project for scheduling pharmaceutical packaging plants. Dr. Richards also lead the project for adapting Aurora to optimize the vehicle testing process by selecting the best vehicle configurations to minimize the vehicle count and overall schedule duration, the result is called Aurora-VT. 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 automation and artificial intelligence techniques; and decision support tool development for life-critical situations. Dr. Richards has publications in all of these domains*

James Ong *received an MBA from Boston University, an MS in computer science (artificial intelligence) from Yale University, an MS in electrical engineering and computer science from UC Berkeley, and a BS in electrical engineering from the Massachusetts Institute of Technology. At Stottler Henke, he leads the development of tools that support automated planning, autonomous systems, and robotics research at NASA. He also leads the development of mission planning, decision/task support, and training applications and tools for the Department of Defense.*