Journal of Innovation
February 2025 | 26th Edition

The Data Centric Architecture of a Factory Digital Twin

Authors:

Steve Harding
Advanced Process Combinatorics, Inc.
harding@ combination.com

Dan Isaacs
Digital Twin Consortium
dan@omg.org

Donald L. Miller
Advanced Process Combinatorics, Inc.
miller @combination.com

Joseph F. Pekny
Advanced Process Combinatorics, Inc.
Purdue University
pekny@purdue.edu

CONTENTS

FIGURES

TABLES

This paper presents a framework for implementing factory digital twins (FDTs), focusing on four applications: (1) debottlenecking, (2) engineering design/retrofit analysis, (3) facility fit, (4) scheduling/planning. These provide substantial value for FDT implementation. The ability to "test-before-invest" fixes mistakes in the virtual realm before committing capital, significantly reducing risks and costs associated with real-world trials.

Our research demonstrates that FDTs serve as effective proxies, allowing complex manufacturing processes to be optimized – both existing and yet-to-be-constructed.

The paper explores how FDTs identify critical data, leverage optimization, and enable rapid scenario testing that would be impractical/impossible in physical systems. We summarize the key role of the Resource Task Network (RTN) which provides a structured means of organizing manufacturing data, describing the manufacturing process, and underpinning the FDT.

Importantly, we address the limitations of FDTs, acknowledging that a single digital twin cannot encompass all aspects of reality. Instead, we use various specialized digital twins to model different aspects of the manufacturing process.

We provide case studies demonstrating how this approach has been successfully implemented, resulting in improved efficiency, reduced downtime, and enhanced innovation.

This article contributes by offering a step-by-step approach to FDT implementation, focusing on applications that are impactful to manufacturing processes. Our findings have significant implications for both practitioners seeking to implement digital twin technologies and researchers exploring the future of digital manufacturing.

Keywords: Factory Digital Twins, Digital Thread, Manufacturing Innovation, Industry 4.0, Optimization, Project Management

1 INTRODUCTION

An FDT accurately replicates the behavior of a manufacturing facility from raw materials to final products. The history of FDTs began in the early 1900s with process charting (e.g. Gantt Charts), but computers in the 1950s enabled FDT optimization and the power of FDTs has grown with the power of computing. Contemporary FDTs became recognizable in the 1970s – for historical perspective see Luyben 1973 [1].

An FDT can generate significant value by rapidly investigating scenarios that would be difficult, costly, or impossible to undertake using the real-world manufacturing facility. An FDT offers the possibility of looking over the multitude of ways a manufacturing facility can be operated to find the best ways that increase capacity, reduce cost, and overcome unexpected events.

An FDT complements digital twins (DT) at the unit operation/manufacturing activity level. A DT intimately depends on the process physics and often involves addressing spatial geometry, partial/ordinary differential equations resulting from the physics, and the underlying technological details. An FDT largely avoids the underlying manufacturing physics by capturing overall behavior using the Resource, Task, Network (see below). A DT can be used to improve process steps identified by the FDT as being critical to overall productivity. In that way FDT and DT form a hierarchy. Similarly, related FDT can be connected to form a supply chain digital twin. This paper focuses on FDT.

An FDT that faithfully reproduces process wide behavior generates value by improving (1) debottlenecking – identifying and reducing process inefficiencies prior to implementation, (2) engineering design/retrofit analysis – reduce or eliminate potential flaws using FDT results before physical implementation, (3) facility fit – use FDTs to determine at which facility a new product should be manufactured, (4) scheduling/planning (Forstrom et al 2023 [17]). Beyond a simulation of the physics of the facility, FDTs reflect operational strategy and facilitate converting business goals into process operation decisions. An FDT designed to optimize profitability or throughput determines the details of process operation to accomplish the objective. FDTs are evolving from advancing computer, modeling, simulation, artificial intelligence, and optimization technology to offer substantial practical value. They have the potential to spread quickly through industry because their value is proportional to the scale of the real-world facility but are inexpensive to develop.

An FDT can be built in an incremental fashion starting with data from a variety of sources. For existing facilities these could include historical data contained in databases, spreadsheets, or from daily operational experience. For new or retrofit/existing projects, a typical source is the requisite engineering designs. An FDT can be iteratively refined by comparing predicted behavior with real world experience and improving the data needed to make them align. This tends to focus FDT data effort on improving the accuracy of data in a goal-oriented fashion. As such, FDT development often results in significant learning about which data is most important and warrants additional effort to improve accuracy.

In addition to process data, and like all digital twins, an FDT uses both historical and real-time data to replicate and optimize a physical thing, in this case the manufacturing facility. At a minimum an FDT requires a final product demand scenario (list of material, amount, due date) to predict behavior, and in real-time applications, requires facility state information to set FDT initial conditions such as inventory levels and in-progress activities. Below, we summarize the general Resource Task Network (RTN[1]) framework that structure the data used to develop FDTs. While the end user of an FDT does not directly interact with an RTN, its role as a framework is central to FDT architecture and is the life blood of the FDT.

Practical FDTs are designed to answer one time and/or regularly occurring questions of interest. An FDT supporting design of a grass roots/new facility will answer questions of interest concerning initial major equipment to be purchased and process capabilities, staffing goals,

[1] RTN, https://en-academic.com/dic.nsf/enwiki/4485962?form=MG0AV3

anticipated incremental expansion, and can serve to ensure robustness of design for a distribution of demand scenarios (see Miller et al 2010 [13]). An FDT supporting facility retrofit will answer many of the same questions but using the existing plant as the starting point. An FDT supporting planning will answer questions pertaining to customer demand satisfaction timing, needed raw material ordering, expected equipment load and labor levels.

An FDT used for detailed scheduling will specify equipment activity sequencing and timing to accomplish the planned production. Using FDTs for facility fit depends on having an FDT for each possible facility at which a new product can be launched. This network of FDTs can be used to determine which of the plants can best undertake the product launch under a range of scenarios. Considering a network of FDTs brings up many tantalizing prospects for driving corporate profitability and detailed cross company supply chain cooperation both in normal times and exceptional times – for example emergency production during a pandemic or a conflict.

The next section reviews the Resource Task Network (RTN) and describes how it must be extended for custom process physics. The following two sections underscore the central importance of timeline management to a FDT and presents a unique innovation for optimizing over the timeline. The following sections review real world FDT characteristics and a step-by-step approach for developing an FDT. Three examples are then presented, and the paper concludes with a discussion of FDT value, a fundamental challenge and a key practical lesson learned. In short, the paper flows from the RTN data to its use in FDT timeline optimization and implementation to examples and concludes with lessons learned.

2 Structuring Data to Develop an FDT: The Resource Task Network (RTN)

The Resource Task Network (RTN) underlies the FDT technology described in the paper. The RTN provides a technology independent means of describing processes and organizing manufacturing data and the presented technology provides the ability to optimize manufacturing operations based upon it. From an intuitive perspective, RTN data can be categorized as follows:

- Master – products, equipment, shift-patterns, and other fundamental data
- Recipe – process activities and bill-of-materials
- Process constraints – data to describe physics that must be obeyed
- Demand – orders, forecasts
- Strategy –safety stock, customer priorities
- State – inventories, activity status
- Output – alerts, decision making information

The life cycle of the data moves from (i) conception – what questions of interest will the FDT answer, (ii) implementation – identify the data sources, (iii) operation – connect to demand and state data and incorporate the FDT into the business process.

The goal of the Resource Task Network (RTN) is to provide a standard description of manufacturing operations in as much detail as necessary to answer questions of interest (Zentner et al 1998 [7], Perez et al 2022 [16]). The RTN provides for data centricity and structure, highlights manufacturing data as a core asset, but is independent of the technology used to implement the FDT. The RTN readily captures bill of material and process flow information but must be extended for specific process physics (see below) and specialized activities, such as quality testing and equipment cleaning that depend on manufacturing technology. Incorporation of and addressing process specific physics is a critical factor to the success and differentiation of FDT technology and starts with the RTN process description.

All value-adding processes involve the transformation of one or more inputs into more valuable products. This is true whether the process involves building automobiles, producing chemicals, fermenting biological products, or managing a complex project with many contributing activities. Modeling any nontrivial process requires a formal framework capable of describing the inputs and the process steps that transform them. The RTN has been a standard tool for describing processes for many years. For the sake of concreteness, we describe the use of the RTN for manufacturing processes below, but it is also applicable to organizing project management and other processes that transform abstract inputs over time.

As the name implies, the RTN is a network of nodes that are interconnected via directed arcs. Each material in a process is represented by a material node. Tasks consume one set of materials (inputs/feeds) and produce another set (products). Figure 2-1 illustrates a simple RTN.

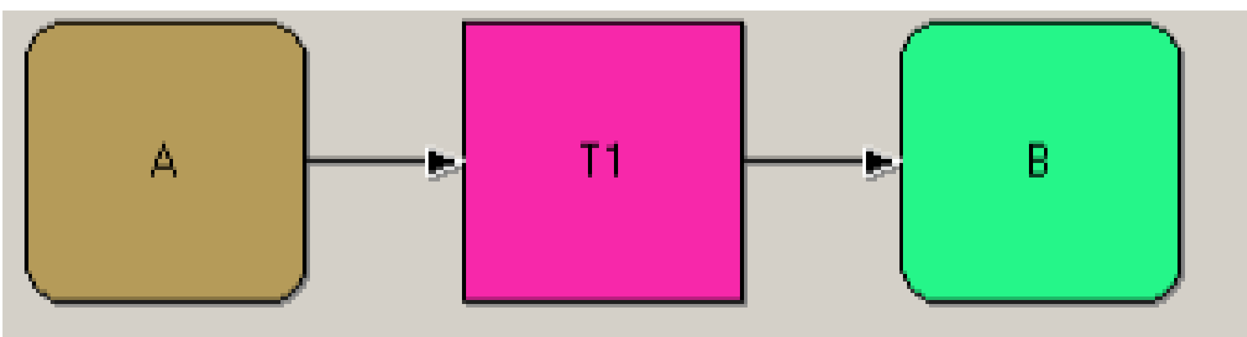

Figure 2-1: A diagram of a simple RTN.

This network RTN contains a single task T1 that converts material A into material B. Figure 2-2 shows a still simple, but more realistic example of an industrial process with eight tasks and five materials. The RTN can easily describe processes of arbitrary depth (the number of processing steps from start to end) as well as process parallelism.

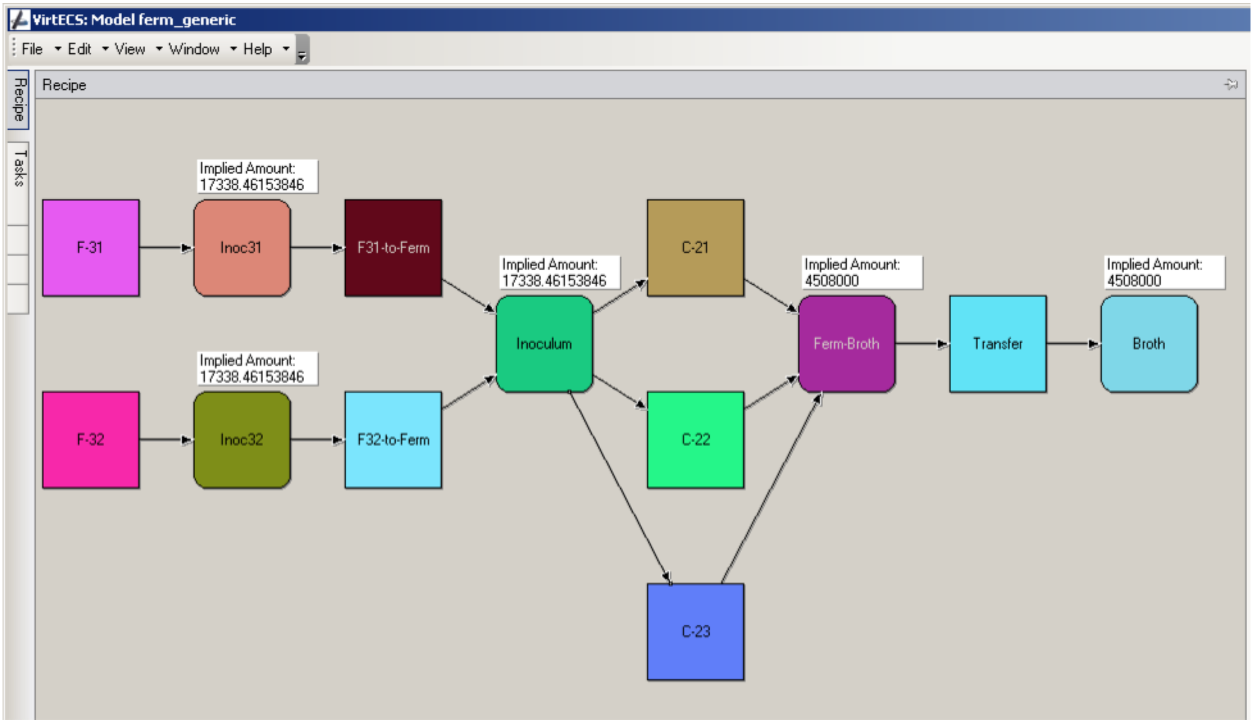

Figure 2-2: An RTN for a multiple step manufacturing process.

The RTN represents a network of nodes alternating between materials and tasks and connects the process ingredients to final products. The RTN is connected to the physical processing equipment and other limited resources such as labor. Each task may be run on one or more pieces of equipment and/or occupy renewable resources, and each material may be stored in one or more pieces of equipment. Thus, we can associate a list of equipment and/or renewable resources with each node in the RTN. Renewable resources in the RTN are those that are occupied during some specified time periods in the execution of a task and then are available for other uses when the task does not need them. Labor is an example of a renewable resource, or different pools of specialized labor.

Specifying a task, a piece of equipment, a start time and a batch size (or extent for rate processes) fully determines an instance of a process task. Each process task instance corresponds to an activity that must be performed and is represented as a box on the Gantt Chart. Materials created by one task at a certain time and consumed by another task later must be stored somewhere during the intervening time.

One of the many features that set advanced FDT technology apart from older technologies is the explicit tracking of material storage. A realistic schedule specifies exactly how much of each material is available and in which vessel it is stored at every point on the timeline or at least guarantees that the implied storage is feasible in practice. The presence of a non-zero amount of material in a storage vessel over a time interval is indicated by storage activity on the Gantt Chart (see Figure 5-1). This feature is of less importance in a plant that makes automobile engines because cylinder heads can be stacked on the factory floor if need be – as such storage is less of a constraint. An FDT designed to handle processing materials that are liquids, slurries, or powders must account for the availability of a storage vessel specifically designed and connected for that purpose. Material oriented FDT must also manage cases where material can only be stored for a limited time, must be stored at least for a specified time, cannot be mixed across batches, or cannot be stored at all. In such processes, managing storage is of paramount importance.

In addition to recipe, equipment, and renewable resource information, specialized manufacturing processes have process physics specific data. For example, in biologics cell culture processes there are specialized cleaning constraints necessary to prevent pathogen contamination. For real world FDT development, the RTN consists of the general information described above plus the specialized data that the FDT must accommodate, which can be highly industry specific. As another example, consider coffee roasting. Once the beans are roasted a certain amount of degassing time is required by the recipe to provide high quality product. A coffee manufacturing FDT must satisfy this degassing time. Honkomp et al 2000 [8] provides a description of a variety of process physics that must be addressed by FDTs.

An important step in developing an FDT process model is to construct the appropriate RTN. This defines the level of detail that the model will be able to express and the data that is needed. A model with insufficient detail will result in unrealistic results, e.g. the model may generate schedules that cannot be executed with the real-world equipment in the process. A typical example would be making a liquid at 8 AM and consuming it at 10 AM when there is no tank available to store it during those two hours.

Too great a level of detail can result in a model that is slow to solve and results that are overly complicated and difficult to understand. As with all modeling activities, much of the art involved in building an FDT model is related to the critical choice of what is to be included and what is to be omitted. In addition, the RTN can evolve over time to increase or decrease the amount of detail according to changes in business requirements. In fact, this flexibility to easily change is one of the advantages of an FDT based on an RTN.

Note that the RTN based FDT describes the manufacturing process at a higher level and does not usually calculate the detailed physics of individual activities/tasks. For example, if a task represents mixing input materials, the FDT does not calculate the liquid flow field in a tank or calculate the degree of mixing as a function of time. Rather, the RTN is specified in terms of how long the mixing usually takes to go from pure inputs to the desired mixed output, perhaps with some modeled process time variability.

For purposes of understanding a mixing step in detail a computational fluid mechanics based DT might be developed. For a chemical reactor, a kinetics model – how fast reactants are converted to products - may be the basis of the activity/task digital twin and supporting the kinetics model might be a quantum chemistry model to gain insight into how reactants are converted to products. In fact, these examples show that real world artifacts – a factory – often need to be represented by a hierarchy of DTs.

3 MANAGING THE TIMELINE IS A KEY CONSIDERATION FOR AN FDT

Awareness of time has always been part of modern business. The adage "time is money" is especially true in the 21st Century. However, in manufacturing, the management of resources through time has never been more complex. Businesses are under pressure to offer increased product differentiation, lower cost, reduced variability, and greater responsiveness by having an accurate awareness of operations.

Complexity in managing resources through time arises for three reasons: causality, combinatorics, and uncertainty (Pekny 2002 [9]). Causality is a simple concept: Tuesday follows Monday, night follows day, etc. However, achieving the goal of providing every customer's services and/or product on time requires that all preparation and precursor steps be completed at the right time. Furthermore, in the real world, a business cannot go back to last week or last month to do something that has become apparent today as necessary for success.

The consideration of causality is often confounded by combinatorics. Businesses have many decisions about how they can operate, and their choices will set the stage for future success or failure. For example, consider making a single decision – should customer one or customer two be satisfied first. The future consists of two timelines, one in which customer one is satisfied followed by customer two and the other in which the order is reversed. Which of these two timelines is better depends on customer expectations and the steps necessary to satisfy the customer.

In real applications, businesses face multiple decisions at a given time and these decisions will interact with subsequent decisions. The number of combinations of different choices for these decisions means that businesses must select one of a huge number of timelines on which they can operate. A key function of an FDT is to optimize among these many timelines and choose the best one whenever significant new information is available. Historical experience can guide choices that appear equivalent.

Twentieth Century physics and business experience show that uncertainty is a fact of life that must be managed by making decisions that work well with a high probability regardless of how uncertainty plays out. In addition to different combinations of decision choices, the uncertainty of process outcomes presents businesses with a huge number of possible timelines on which they could operate, some with very different desirability.

There are several defenses against uncertainty. Carrying inventory, spare capacity, outsourcing contracts, and pricing strategy are all examples of insurance against uncertainty. However, a business can be over-insured. For example, carrying too much inventory incurs excessive cost relative to the probability that it will be needed. Spare capacity and inventory are to some extent interchangeable types of insurance. Inventory is not needed if there is enough time to utilize spare capacity to satisfy an unexpected order. In general, to acquire the right amount of insurance against uncertainty, a business would like to know that their insurance "portfolio" will perform well in a large fraction of the future possible timelines. A key function of an FDT is to provide a data-based approach for determining the right level of “insurance” (Jung et al 2008 [11]).

Different FDT technologies provide a variety of means for addressing the timeline. Spreadsheets are the most popular technology for developing an FDT (Gottfried 2005 [10]). A cell often represents a manufacturing quantity at a particular point in time. Similarly uniform discretization models (UDMs) represent manufacturing quantities at uniform points in time (see below). Discrete event simulator based FDT technology maintains a stack of interesting time points and updates the simulation from earliest to latest time (Law 2024 [20]).

In the next section, we describe a novel technology for optimizing an FDT over the timeline that is built from the RTN.

4 Infinite Dimensional Programming Innovation for Timeline Management

In this section, we summarize the unique treatment of the timeline used by the VirtECS system from APCI (*www.combination.com*) for developing FDTs to enable mathematical programming optimization. VirtECS translates the RTN description to an infinite dimensional mathematical program (Friesz 2010 [12]) that enables branch and bound of the resulting Mixed Integer Linear Program (MILP).

Optimizing process behavior over time involves considering the possible ways that activities could be performed over time to satisfy causality and meet all other physics constraints. Consider the planning of the extent (e.g. batch size, processing rate, or an investment amount) of an activity. If the activity is undertaken (yes or no) at a given time t then the extent must be between allowed minimum and maximum levels. As a mathematical relationship this can be written as:

$$E_{min}\, x_t \le E_t \le E_{max}\, x_t \qquad (1)$$

Where E_{min} is the minimum allowed extent of the activity, E_{max} is the maximum allowed extent of the activity, E_t is the chosen extent of the activity, and x_t is assigned a value of one if the activity occurs at a time t and zero otherwise. If the activity does not occur at time t, then x_t is set to zero, the left and right hand parts of the relationship (1) are zero, and the chosen extent E_t must be zero. Conversely, if the chosen extent at a given time, E_t, is not going to be zero, then the activity must be chosen to occur (x_t must equal one). More complicated process physics can be easily represented mathematically.

By introducing mathematical notation, physics concepts underlying processes can be easily represented using mathematical expressions: material balances, unit allocation constraints, resource limitations, etc. The objective to be optimized can also be specified as a mathematical relationship so that the goal of the problem to be solved becomes choosing the variable values to get the best objective value while satisfying all the constraints.

There are substantial advantages to using equations to describe the behavior of processes. For example, the role of data describing process behavior is well-defined – like the minimum and maximum extent limits in relationship (1). The list of the constraints plus the objective allows a precise statement of the problem to be solved. Uncertainty can also be addressed by allowing parameters like cycle time or the amount produced from an activity to be a random variable with some experimentally observed statistical distribution.

A key complexity management aspect of using equations is the fact that reasoning about the problem description is decoupled from the reasoning about how to get a good or best solution. This promotes rapid evolutionary improvement of the technology. One of the most powerful types of formulation families discovered (Elkamel 1993 [4], Pantelides 1994 [5]) for FDTs, one with very precise and extensible descriptive power is known as the Uniform Discretization Model (UDM).

This type of formulation divides the timeline up into uniform pieces (buckets) and mathematical relationships such as (1) can be written over each bucket. A solution that can be shown on a Gantt Chart (see below) and in other plots is an assignment of variable values that satisfies the constraints over every bucket and results in a good or best possible objective function value. Unfortunately, UDM formulations face two seemingly intractable problems.

First, the number of mathematical relationships that result from realistic sized buckets and an industrial scale time horizon are enormous. The traditional approach to using mathematics has been to generate the equations that are needed and then pass these to a solution algorithm to get an answer. Because practical problems and real data require small bucket sizes for the appropriate realism there can be hundreds of millions or billions of relationships required. Even contemporary and future computers are not sufficiently fast nor do they have enough memory to represent most real world problems. In addition, the number of yes-no variables implies that the potential number of solutions is incredibly large. Thus, the second major difficulty with process management problems is that studies have shown the number of solutions to often be 10^{2500} to 10^{25000} or more due to the combinatorial nature of the yes-no decision variables. Any algorithm that explicitly attempts to look through this vast number of solutions will have a prohibitive execution time on any computer now or in the future.

The work of Miller and Pekny (Miller et al 1991, Miller et al 1995 [6]) demonstrated that highly engineered, custom algorithms could solve even very large combinatorial (yes-no) decision problems. The essence of this algorithm engineering is to develop data structures and mathematical theory highly specific to the class of problem and specific instances of interest. The goal of the mathematical theory is to develop properties that allow implicit search of the solution space, identify regions of that solution space where good or best solutions lie, and focus computational power on exploring these regions.

The other key aspect of the work of Miller and Pekny is to never explicitly generate all the mathematical relationships needed to describe the problem because most of the relationships are trivially satisfied and contain no useful information. Rather only a very small number of the mathematical relationships are generated to describe the solution and prove that it is good or best. VirtECS implements such an approach for FDTs. In particular, VirtECS shrinks the bucket size to be nearly infinitesimal in duration so that its technology actually implicitly generates millions of mathematical variables and relationships per hour of the timeline that is modeled. This highly specialized approach integrates the generation and solution of the mathematical entities. To our knowledge VirtECS is the only FDT system that has implemented such an approach and the tools of infinite dimensional programming available to generate good or even optimal time-based solutions for FDT.

5 Real World FDT

In this section, we summarize real-world examples of FDTs. Using the formalism of the RTN, Figure 5-1 provides examples of several FDT for world scale manufacturing plants in different industries where VirtECS has been used. This data is new, has never been published, and shows that practical FDT have a wide range of characteristics. The FDTs of Figure 5-1 are used on an ongoing basis for planning and scheduling and engineering analysis/retrofit and on an occasional basis for facility fit to support sales processes for contract manufacturers. In addition to the basic RTN statistics,

Figure 5-1 also contains the span of time over which an analysis or schedule has occurred in the example and the amount of time (rightmost column) required for the VirtECS system to obtain the solution (Gantt Chart) on a 3.0 GHz Intel processor computer with 64 gigabytes of memory. Figure 5-1 shows that real world FDT use can predict many months of activity in only a few seconds to few minutes and are practical using inexpensive computers. The longer activity span applications tend to be for engineering analysis/retrofit or facility fit applications and the shortest activity span are for detailed scheduling applications.

The planning uses of FDT cover an activity span of several weeks to many months. Real world applications model from the low hundreds of manufacturing tasks/steps to over one thousand task/steps over all modeled products. Real world FDT typically model dozens to low hundreds of pieces of equipment. Most real world FDT model renewable resources. Even among specific industry applications, such as drug substance, there are a wide range of FDT model sizes reflecting differences in real world plant sizes and the degree of detail of the model for the application.

Name	Digital Twin Type	Tasks	Equipment	Materials	Routings	Storage	Resources	Activity Span (days)	Solve Time (sec)
A	Alcohol Manufacturing	361	218	315	361	289	0	503	198.1
B	Drug Packaging	657	22	701	2478	623	8	77	7.1
C	Drug Product	1097	90	1273	1533	1420	10	357	56.8
D	Drug Product	981	169	1043	5586	1145	10	40	40.7
E	Drug Product	210	8	235	230	234	11	21	3.5
F	Drug Substance	1077	183	1171	1081	1172	35	45	21.5
G	Drug Substance	576	107	636	2111	623	16	132	18.8
H	Drug Substance	543	133	617	948	2587	50	191	90.4
I	Drug Substance	362	189	401	558	413	72	43	20.0
J	Drug Substance	362	92	468	378	500	6	126	67.0
K	Drug Substance	288	112	312	559	319	21	31	6.7
L	Drug Substance	176	55	18	205	301	10	18	5.7
M	Personal Care	847	19	3020	1545	1613	0	18	48.4
N	Specialty Chemicals	717	77	1504	760	944	0	99	24.3

Table 5-1: Real-world FDT and their RTN statistics.

5.1 REALIZING VALUE WITH A DIGITAL TWIN STEP-BY-STEP

The FDT consists of two parts (1) a process model – the cycle times of steps, changeover times, material balance information, etc. structured as an RTN and (2) plant conditions which allow the process model to predict the future reality of the process from the present. The following is the best stepwise practice for developing an FDT. This bootstrap approach to developing an FDT builds confidence, rapidly provides value, and can help streamline business processes. For facility design/retrofit and facility fit applications Step 3 is not needed.

1. The simplest process model should be developed in a manner that allows the digital twin to be used in practice to answer questions of interest. For example, this model can be for evaluating process improvement options, planning, etc. The output of this process model can be viewed on a Gantt Chart (see Figure 5-1).

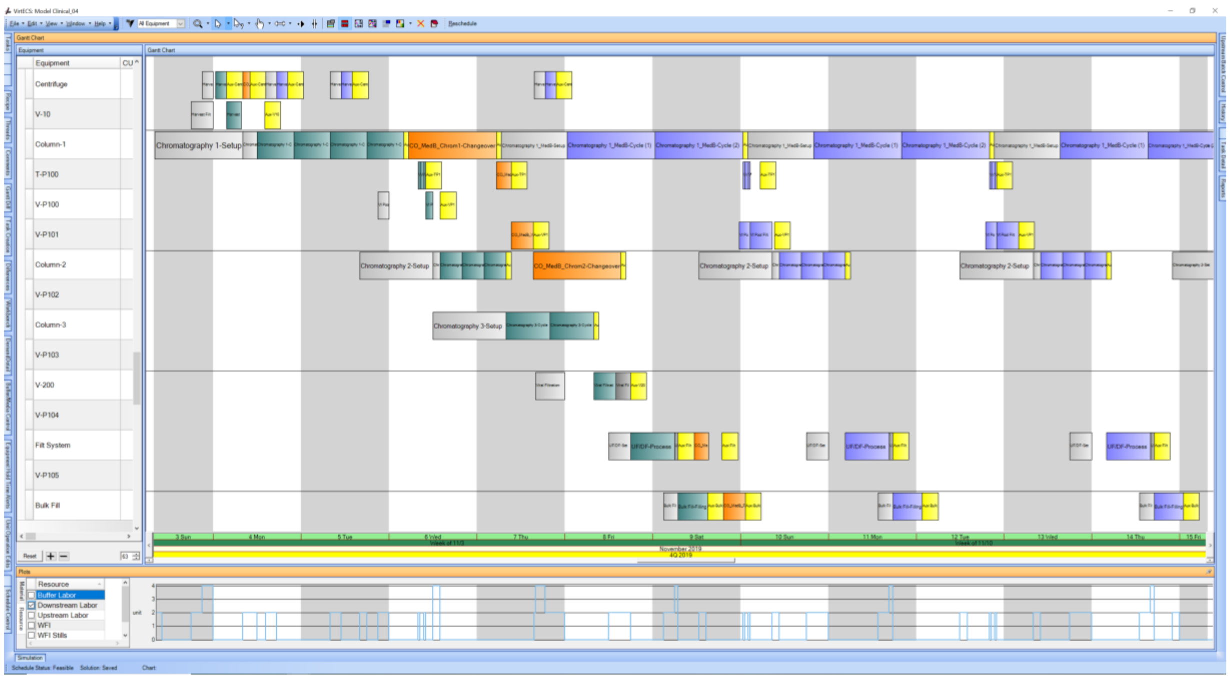

Figure 5-1: An example Gantt chart.

In the example Gantt Chart above, the Y-axis is the set of modeled equipment, and the X-axis is over the time horizon of interest starting with now and out to a point in the future. The Activity Span of Table 5-1 corresponds to the time span of the Gantt chart. The projected inventory levels of raw materials, intermediates, and product are also predicted in the bottom plot of Figure 5-1 using the same X-axis as the above Gantt chart. The Gantt chart text simply provides the names of the activities and materials and is not relevant for purposes of this paper. The key point is that the Gantt Chart is a useful output of the FDT and communicates behavior over the timeline.

2. Based upon routine usage of the starting model in 1, the model can be incrementally expanded to include more detail. Again, the emphasis should be to add only detail that is deemed necessary to address additional questions of interest. Balancing the additional detail is the effort of maintaining the needed data. Perhaps the initial FDT only predicts equipment activity, and the refinement adds labor usage required by each task/activity.

3. After a short period of time over which a useful working model is developed, the predicted activity of the FDT can be published to the manufacturing community – see Figure 5-2. An accurate prediction of the future allows community members to best plan their individual activity to maximize the business impact. The results of the FDT can be published on an intranet website and require minimal training. A blog can develop around a manufacturing activity to collaboratively trouble-shoot or update the community on actual start time and other key process information. A database of predicted versus actual process performance results in continuous improvement.

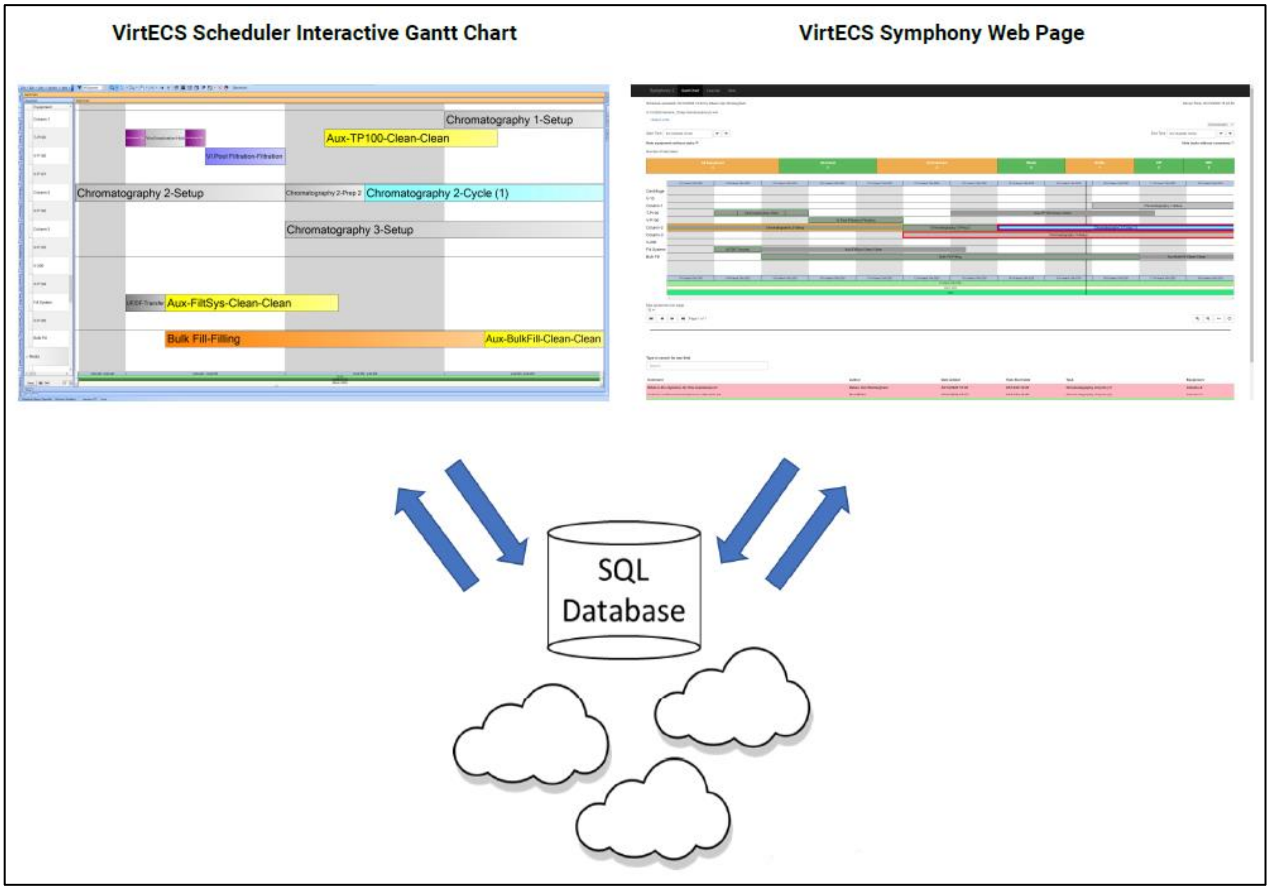

Figure 5-2: The predictions of the FDT can be published in the VirtECS/Symphony web-based environment. (Source: *www.combination.com*)

5.2 Example Applications of an FDT

To illustrate the discussion above, the next three sections provide specific example FDTs. The first two examples consider two case studies involving two-stage processes that make two products. Prod1 and Prod2 are made in a continuous process on equipment R1. The tasks making Prod1 and Prod2 require precursor materials Prec1 and Prec2 respectively, which are made on equipment T1. There are no transitions between Prod1 and Prod2, but switching tasks on T1 requires a 4-hour transition. There is somewhat regular demand for both products. The third example is for a more complex cell culture-based process. The data provided in Table 5-2 and Table 5-3 is the information needed to create the RTN for the first two examples. The third larger example is drawn from biologics drug substance manufacture, which summarizes a case study of an industry partner.

5.2.1 FDT Analysis Example 1

Figure 5-3 shows a two-month schedule satisfying the demand pattern. Because all tasks are continuous and the precursor materials must be available before starting the final tasks, there is an offset between the two production stages.

Task	Material	Rate Kg/hour
Make Prec1	Prec1	250
Make Prod1	Prec1	-250
Make Prod1	Prod1	250
Make Prec2	Prec2	286
Make Prod2	Prec2	-250
Make Prod2	Prod2	250
Make Prec3	Prec3	125
Make Prod3	Prec3	-250
Make Prod3	Prod3	250

Table 5-2: Example 1 process data.

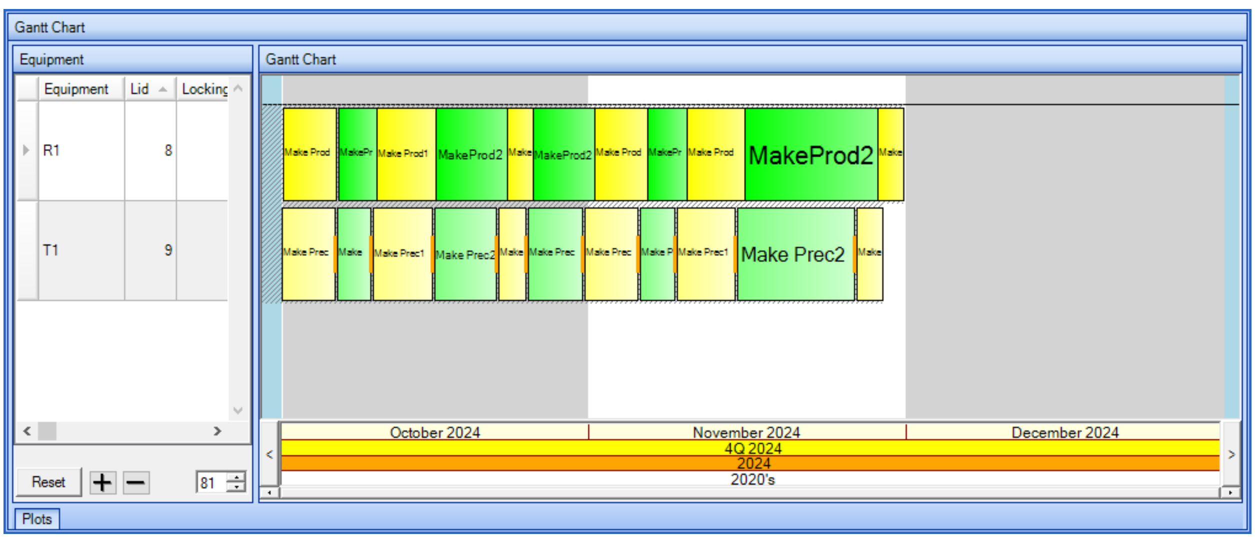

Figure 5-3: Schedule predicted by VirtECS FDT.

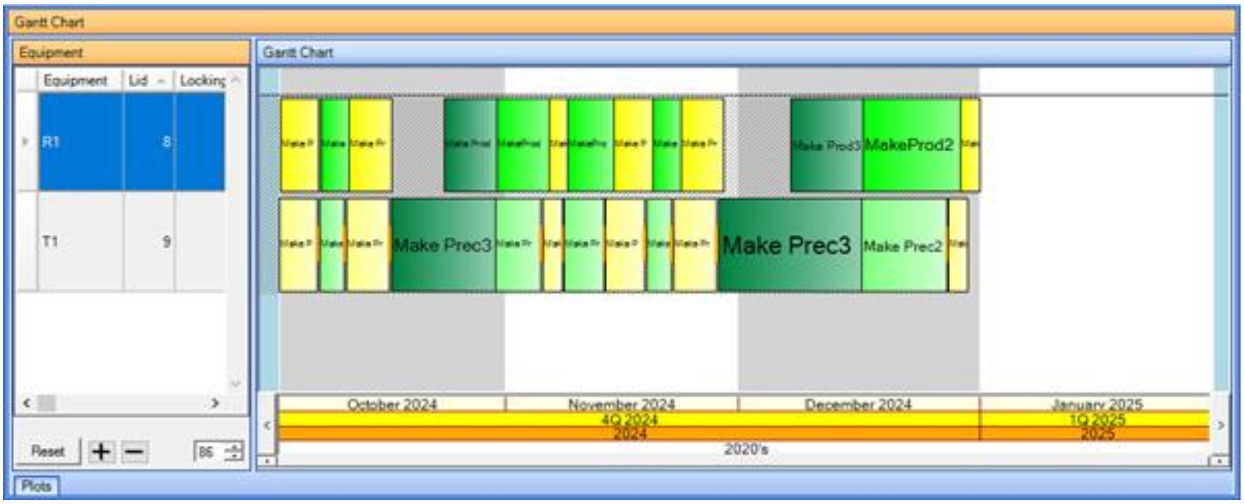

Figure 5-4: Schedule after inserting two demands for product Prod3.

The problem at hand is the introduction of a new third product, Prod3. Prod3 differs in that the relative rate of its precursor material is slower than the final stage. This results in a loss of capacity as the main equipment R1 is idled while awaiting the production of Prec3 in T1. As Figure 5-4 shows, inserting two demands for Prod3 results in two significant periods of idle time on R1, reducing the overall capacity of the plant.

The imbalance between the two stages for the new product can be offset by modifying the plant to add a second parallel equipment (T2) to increase precursor capacity. This illustrates the ability of the FDT to model process modifications and accurately predict the resulting value of those modifications by producing detailed schedules. Figure 5-5 shows the demand scenario from Figure 5-4, but with the addition of a second precursor equipment, T2.

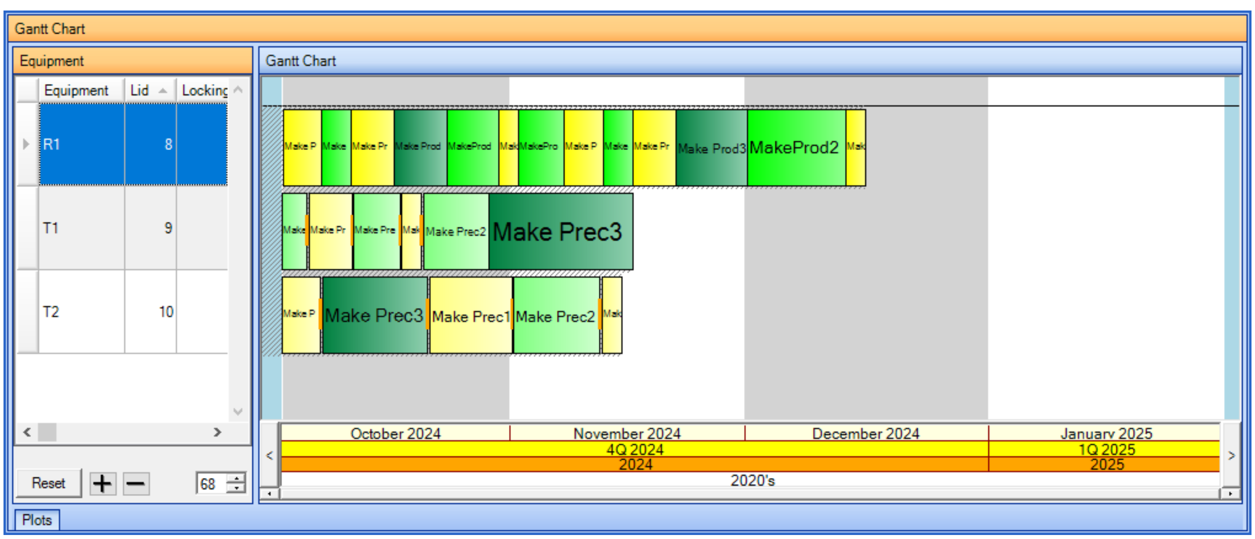

Figure 5-5: The addition of a second precursor equipment T2.

The second precursor equipment provides more than sufficient capacity to feed the main step with the result being that R1 is fully utilized. This reduces the make span of the three-month schedule by about two weeks. This type of automated FDT based analysis becomes critical for highly complex processes where manual/spreadsheet calculation is tedious or impossible.

5.2.2 FDT Analysis Example 2

The next case involves an alternative approach to improving plant capacity with Prod3, not by capital spending, but by changing the way demands for the new product are managed. In this example the new product is well balanced between precursor and final stages but has a significant transition time when run adjacent to the existing products. Adding routine demands for Prod3 into the normal schedule induces transitions that reduce overall plant capacity.

Task	Material	Rate Kg/hour
Make Prec1	Prec1	250
Make Prod1	Prec1	-250
Make Prod1	Prod1	250
Make Prec2	Prec2	286
Make Prod2	Prec2	-250
Make Prod2	Prod2	250
Make Prec3	Prec3	125
Make Prod3	Prec3	-244
Make Prod3	Prod3	244

Table 5-3: Example 2 process data.

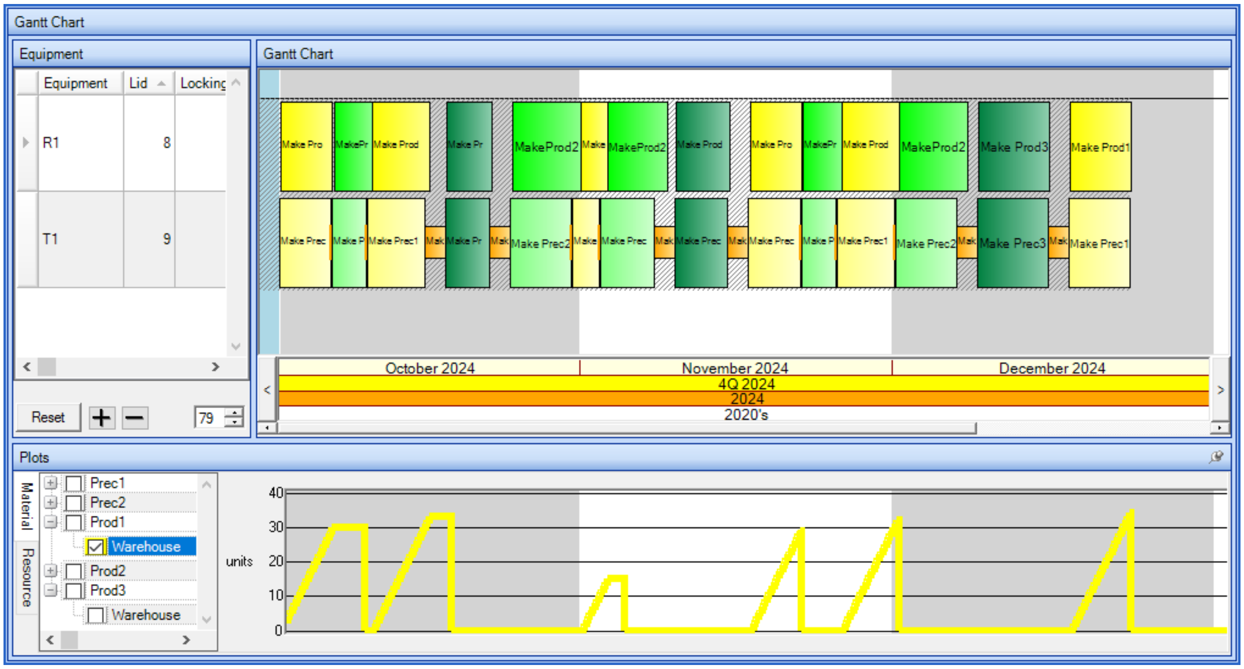

Figure 5-6: Equipment T1 contains transitions (orange bars) that reduce capacity.

The plot at the bottom of Figure 5-6 shows the inventory of Prod1. We can see that production closely matches demand due dates so that final product inventory is kept close to a minimum. But every time a campaign of Prod3 is made, the resulting down time on R1 costs capacity. This effect can be ameliorated by running fewer but long campaigns of Prod3 as shown Figure 5-7.

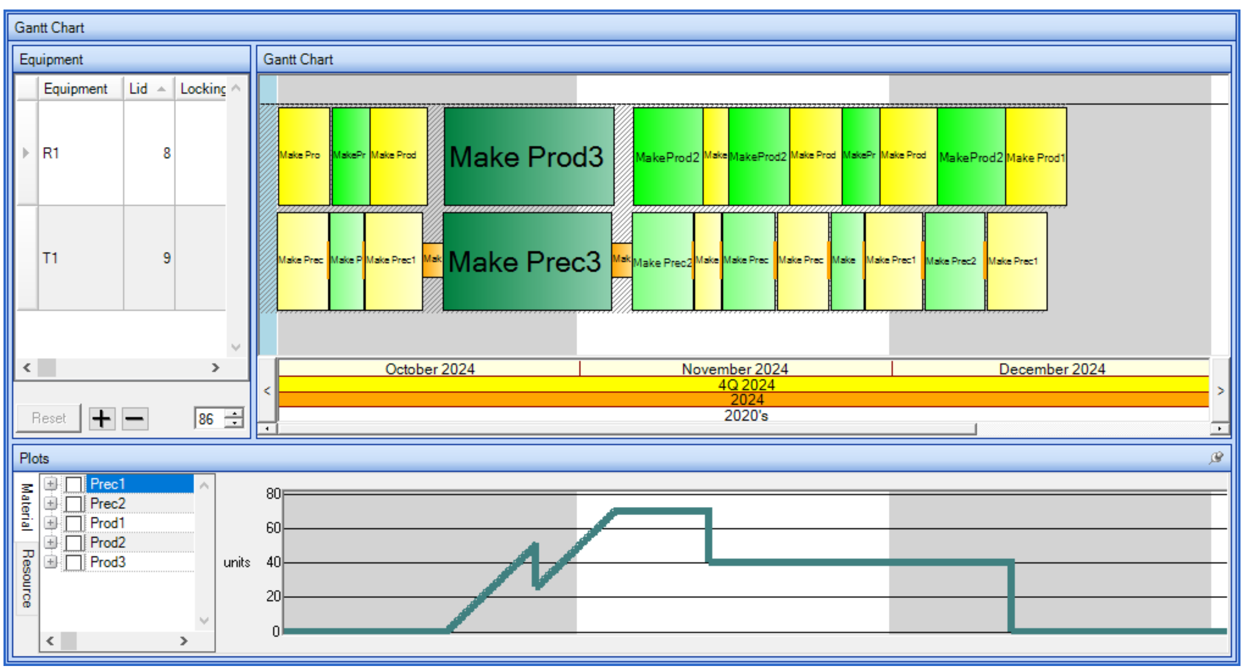

Figure 5-7: Campaigning to reduce transitions.

Here we serve the same three demands for Prod3, but run only a single campaign, early enough for the first demand and large enough to serve all three demands. Of course, this results in carrying more final product inventory, but this solution may be more attractive economically than plant expansion. While reducing inventory has become a watchword for plant efficiency, this example illustrates that tolerating additional inventory can sometimes make good economic sense. An FDT can be used to quickly explore alternatives that result in the best process behavior.

5.2.3 FDT Analysis Example 3

An example biologics process consists of nine processing steps that are divided into a front end and back end. The front end consists of Inoculation followed by three sequential stages of Bioreactors, 750, 2500 and finally 10,000 liters. All stages in the front end have three parallel sets of equipment. The final Bioreactor stage (10,000 liter) is the anticipated process bottleneck. The output from this reactor is stored in centrifuge feed tanks from which it enters the back end of the process. The back end consists of a centrifuge, and three chromatography columns, Anion Wash, Revo, and CatEx. The final step is Filtration.

The chromatography columns have cycle times ranging from 8 to 21 hours. Batch integrity must be preserved so a single batch from the centrifuge is fed to the Anion Wash column, then flows to the Revo Feed Tanks, Revo Column, CatEx Feed Tanks, CatEx Column and on to the Filter.

5.2.4 Specialized Process Constraints

In addition to the normal constraints implied by the physics and chemistry of the system, e.g. material balance, storage level limits, material availability, etc., this process entails three classes of specialized constraints. First is batch integrity. Batches are required to be completely emptied into the next stage and the storage tanks washed out (this takes one hour) before the next batch may be introduced into the storage tank. The FDT must ensure that the requisite hour of time between sequential batches is present so that the schedule the system produces is actually executable in the plant.

The second class of special constraints involves the chromatography columns. Each of the three columns must be periodically repacked after a specified number of cycles. The maximum number of cycles between repacks are Anion Wash 20 cycles, Revo 17 cycles, and CatEx 15 cycles. The repack times vary from 15 to 24 hours. The FDT must track the number of cycles since the last repack on each column and must ensure that the repack tasks are executed at the required frequency.

The third class of specialized constraint results from random delays that can occur on the chromatography columns due to operational variability. A pattern and probability of such delays can be specified as part of the study input. This allows design engineers to experiment with the process before it is built and determine the expected resiliency with respect to expected levels of variability. This is an important tool for design engineers who can use it to ensure that the real-world performance of the plant will live up to nameplate expectations in the face of stochastic disturbances.

5.2.5 Results

Figure 5-8 illustrates the solution to the base case problem, with no random delays on the chromatography columns. Following the process from front end to back end (top to bottom), we see the three parallel flasks that feed the three parallel Bioreactor trains. The initial flask tasks are staggered due to a resource constraint that prevents multiple Inoculation tasks from starting at the same time. As the Bioreactor train starts up, the 10,000-liter bioreactor quickly establishes the rate of the whole train. Material awaiting processing is held in the two smaller bioreactors. Effluent from the 10,000-liter bioreactors is held in the Fuge Feed Tanks until it is processed in the centrifuge.

We can see the three sequential chromatography columns processing with the yellow repack tasks interspersed at the frequency required for each of the columns. The last stage of the process is the relatively fast Filter task at the bottom of the chart.

The 10,000-liter bioreactor is the rate limiting step as we can see by the fact that there are no gaps on this reactor. It is also evident that the Revo column is very nearly rate limiting, only small gaps exist. The pink storage tasks reveal a wealth of information about the nature of this process.

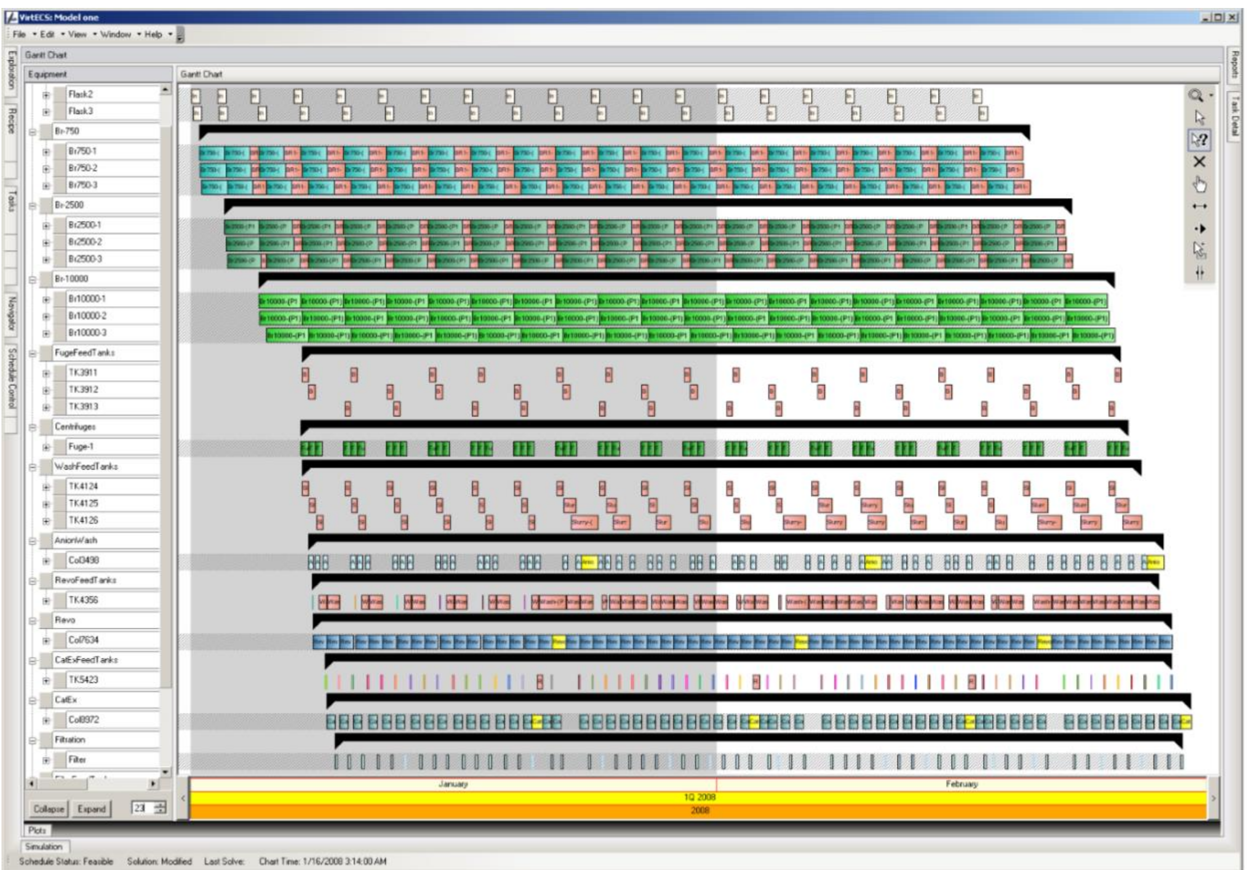

Figure 5-8: The base case solution.

Short duration storage tasks tend to lie upstream of equipment that is not rate limiting. The level of storage activity on the Revo Feed Tanks indicates the tight capacity on the Revo column. Storage patterns on the back end of the process are somewhat irregular because of the uneven pattern of repack tasks (yellow).

Figure 5-9 illustrates a solution when we add a 10% probability of an 8-hour delay to the Revo column (red colored tasks randomly distributed). The Revo column is very nearly rate limiting in the absence of random delays and the addition of these delays puts more stress on the storage system. The RevoFeedTank is starting to become a bottleneck about halfway through the schedule. Note that the VirtECS FDT always maintains a one-hour separation between storage of sequential batches as is prescribed for cleaning the storage tanks.

As the TK4356 becomes saturated, slurry from the centrifuge begins to back up into the three Wash Feed Tanks. Due to this storage capacity, the back end of the process is still able to keep up with the front end, no gaps are induced on the true bottleneck, the 10,000-liter bioreactors. If the available storage had been inadequate, the random delays would have impacted process capacity. This is the sort of complex and subtle effect that can only be foreseen with a highly detailed FDT. If detected in the plant design phase, steps could be taken to determine the most cost-effective process modifications to alleviate the problem.

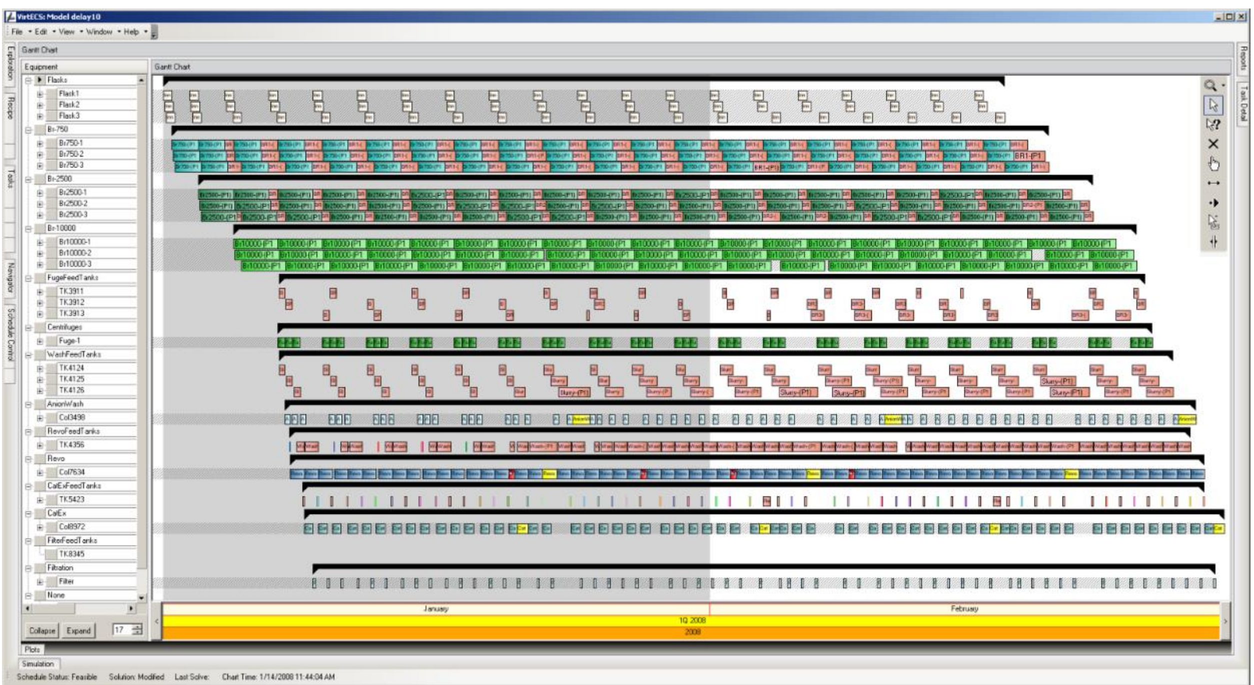

Figure 5-9: Addition of a 10% probability of an 8-hour delay.

6 THE VALUE OF FACTORY DIGITAL TWINS

Determining the value of an FDT is challenging. Essentially two timelines must be sampled for the real process. One timeline in which the FDT is used to manage the real process and an alternate timeline where the FDT is not used. By comparing the performance between the two timelines the value of the FDT is apparent. Real-world processes usually do not have the opportunity to undertake such a controlled study.

Instead, the value can be estimated based upon some real world measures: (1) how busy is the plant, (2) how much revenue does the plant generate, (3) how severe are process pain points such as overtime or how much more product could be sold if it could be produced, (4) how long do manufacturing plans take to develop and how fast can they be changed when the unexpected happens (Stevens et al 2024 [18]). Schregardus 2024 [19] describes a pharmaceutical scheduling application FDT that was justified based on producing just one additional batch a year, which had a value of $1 million.

As an intuitive rule, the busier the plant the more valuable the FDT. Pekny and Miller 1991 [3] provides a controlled study for a simple manufacturing process of the value difference that can occur between different operating strategies for an FDT. Anecdotally, for a busy plant the FDT can increase the throughput by 5-15% even for a world class facility and more when unexpected events become more frequent. The intuition derives from the fact that FDTs can optimize over very many ways the plant can operate, take the best, and can get the details right every time.

A human may be able to match the FDT under ideal conditions, but in real world conditions and with time pressure the best way to proceed is to pair human intuition with the power of an FDT. A useful initial FDT can be developed in several weeks starting from existing data. Over many months an FDT in routine usage can get very sophisticated and catalyze a data-driven modeling culture that pushes the envelope of world scale performance.

In practice the development and use of an FDT catalyzes process improvement in multiple ways. Process data can get corrected or validated in a goal-oriented way. For example, throughput can be predicted by the FDT on a product-by-product basis and the predicted value can be compared against experience. Such simplified model-reality comparisons allow incorrect data to be quickly identified and rectified. Once the model data is validated the FDT can be used to accurately assess changes to the physical process. Process improvement guess work is replaced by data analytics and a high degree of confidence that process changes will yield expected results.

Manufacturing personnel operating under normal circumstances develop strong intuition as to how a process should behave. However, where normal operating intuition does not apply the value of an FDT can be significant in responding to the exceptional. When Hurricane Katrina struck New Orleans in 2005 a VirtECS FDT was being used to manage Folgers coffee production. The plant is above sea level, but the surrounding neighborhoods were flooded which made commuting to the plant difficult and utilities, such as fresh water, were in short supply.

The VirtECS FDT was quickly configured to determine the minimum amount of labor and process water that was needed to begin reduced operations. The National Guard then helicoptered in the needed skeleton crew and delivered the needed process water in trucks. According to plant personnel the FDT predictions helped the Folgers plant to be one of the first New Orleans plants to be back up and running following the hurricane[2].

7 Fundamental Challenge/Practical Lesson Learned

The fundamental challenge of FDTs is that most of the underlying computational problems are NP-Complete (Wigderson 2019 [15], Harel et al 2012 [14]). Known computational theory provides little insight as to how to proceed, but Miller et al 1991 [2] suggests an algorithm engineering approach around similar problems instances works well. Decades of large scale VirtECS FDT experience corroborates this suggestion. Essentially algorithms underlying an FDT can perform poorly in either runtime or solution quality or both.

This applies to high performance algorithms, spreadsheets, or human calculation. However, experience shows that well engineered algorithms designed for specific instance classes can perform well on similar instances. When instances differ substantially from those for which the algorithm was effectively designed, then performance can degrade unacceptably. FDT development that builds on this similar instance approach can be successful and underlying architectures must be designed to rapidly accept new learning from new instances.

A practical lesson learned from real world FDT is that developing the best RTN model is an art. An RTN model that is too detailed performs unnecessarily poorly and is expensive to develop. An RTN model that is not detailed enough is unable to accurately answer questions of interest. Consider a one task RTN model of an entire factory. Raw materials are consumed at time zero of the task and a product is available by summing step process times.

Such a one task model of a factory obviously ignores very many details and is not useful in answering very many questions involving equipment, labor and many other details. However, the other extreme of developing the most detailed model that can be reasonably conceived is often too expensive for most practical needs. Furthermore, the unnecessarily detailed approach is seductive without substantial model building experience. The best practice is to develop the simplest model that will be widely usable and only make that model more complex as practical experience dictates. This is the approach to which all FDT experts converge, though even experts sometimes add too much detail on specific projects until experience suggests how to streamline them.

8 Conclusion

The development and use of an FDT catalyzes process improvement in many ways. The process data needed by an FDT can be structured using the Resource Task Network (RTN) and translated to data structures – for example equations and variables - needed by algorithms that optimize

[2] *https://www.purdue.edu/uns/html3month/2006/060320.APC.Folgers.html*

process behavior. This data gets cleaned and validated in a goal-oriented way by comparing FDT predictions with known process behavior.

Process measurements over time can be used to fine tune the FDT. Process predictions are made which can be used to improve the model, the physical process, or both. Guess work about proposed process changes is replaced by FDT predictions and data analytics. In short, the development and routine use of an FDT can induce a data driven culture and improvement process.

9 REFERENCES

[1] W. L. Luyben, "Process Modeling, Simulation, and Control for Chemical Engineers", McGraw-Hill, 2nd edition, 1973.

[2] D. L. Miller and J. F. Pekny, "Exact Solution of Large Asymmetric Traveling Salesman Problems," *Science*, Vol. 251, pp. 754-761, 1991.

[3] J. F. Pekny and D. L. Miller, "Exact Solution of the No-Wait Flowshop Scheduling Problem With a Comparison to Heuristic Methods", Computer and Chemical Engineering, Vol. 15, No. 11, pp. 741-748, 1991.

[4] A. Elkamel, "Scheduling of Process Operations Using Mathematical Programming Techniques: Towards a Prototype Decision Support System", Ph. D. Thesis, Purdue University, 1993.

[5] C. Pantelides, "Unified Frameworks for Optimal Process Planning and Scheduling", Centre for Process Systems Engineering, Imperial College, 1994.

[6] D. L. Miller and J. F. Pekny, "A Staged Primal-Dual Algorithm for Perfect b-Matching With Edge Capacities", ORSA Journal on Computing, Vol. 7, No. 3, pp. 298-320, 1995.

[7] M. G. Zentner, A. Elkamel, J. F. Pekny, and G. V. Reklaitis, "A Language for Describing Process Scheduling Problems", Computer and Chemical Engineering, Vol. 22, No. 1-2, pp. 125-145, 1998.

[8] S. J. Honkomp, S. Lombardo, O. Rosen, and J. F. Pekny, "The Curse of Reality – Why Process Scheduling Optimization Problems are Difficult in Practice", Computers and Chemical Engineering, Vol. 24, No. 2, pp. 323-328, 2000.

[9] J. F. Pekny, "Algorithm Architectures to Support Large Scale Process Systems Engineering Applications Involving Combinatorics, Uncertainty, and Risk Management", Computers and Chemical Engineering, Vol. 26, No. 2, pp. 239-267, 2002.

[10] B. S. Gottfried, "Spreadsheet Tools for Engineers Using Excel", McGraw-Hill, 3rd edition, 2005.

[11] J. Y. Jung, G. Blau, J. F. Pekny, G. V. Reklaitis, and D. Eversdyk, "Integrated Safety Stock Management for Multi-Stage Supply Chains Under Production Capacity Constraints", Computing and Chemical Engineering, Vol. 32, pp. 2750-2581, 2008.

[12] T. L. Friesz, "Dynamic Optimization and Differential Games", Springer, 2010 – see especially pp. 147-217.

[13] D. Miller, D. Schertz, C. Stevens, and J. Pekny, "Mathematical Programming for the Design and Analysis of a Biologics Facility – Part 1 and 2", BioPharm International, February/March, 2010.

[14] D. Harel and Y. Feldman, "Algorithmics: The Spirit of Computing", Springer, 3rd edition, 2012.

[15] A. Wigderson, "Mathematics and Computation: A Theory Revolutionizing Technology and Science", Princeton University Press, 2019.

[16] H. D. Perez, S. Amaran, S. S. Iyer, J. M. Wassick and I. E. Grossmann, "Applications of the RTN Scheduling Model in the Chemical Industry." Simulation and Optimization in Process Engineering, Elsevier, 1st edition, 2022, pp.365-400, *http://dx.doi.org/10.1016/B978-0-323-85043-8.00006-4.*

[17] J. Forstrom and J. F. Pekny*, "*Achieving Data-Driven Smart Manufacturing & High Impact Manufacturing Social Media*", BioProcess International, Boston, September 2023.*

[18] C. Stevens and J. F. Pekny, "The Economics of Bioprocessing: The Evolution and Value of Digital Twins for Biologics Manufacturing", BioProcess International, San Diego, March 2024.

[19] R. Schregardus, "What is the Return on Investment (ROI) for Advanced Planning and Scheduling (APS) Software", April 17, 2024, *https://www.linkedin.com/pulse/what-return-investment-roi-advanced-planning-aps-rod-schregardus-tnmge/*.

[20] A. M. Law, "Simulation Modeling and Analysis", McGraw-Hill, 6th edition, 2024.

ACKNOWLEDGEMENTS

The views expressed in the *OMG Journal of Innovation* are the author's views and do not necessarily represent the views of their respective employers nor those of the Object Management Group® (OMG®).

- Return to *OMG Journal of Innovation landing page* for more articles and past editions.