SAND2016-3210C

Proceedings of the 2016 Winter Simulation Conference
J. A. Flory, D. D.. Padilla, J. H. Gauthier, A. M. Zwerneman, and S. P. Miller

OPERATIONAL EXCELLENCE THROUGH SCHEDULE OPTIMIZATION AND PRODUCTION SIMULATION OF APPLICATION SPECIFIC INTEGRATED CIRCUITS

John A. Flory
Denise D. Padilla
John H. Gauthier
April M. Zwerneman
Steven P. Miller

Sandia National Laboratories
1515 Eubank SE
Albuquerque, NM 87123, USA

ABSTRACT

Upcoming weapon programs require an aggressive increase in Application Specific Integrated Circuit (ASIC) production at Sandia National Laboratories (SNL). SNL has developed unique modeling and optimization tools that have been instrumental in improving ASIC production productivity and efficiency, identifying optimal operational and tactical execution plans under resource constraints, and providing confidence in successful mission execution. With ten products and unprecedented levels of demand, a single set of shared resources, highly variable processes, and the need for external supplier task synchronization, scheduling is an integral part of successful manufacturing. The scheduler uses an iterative multi-objective genetic algorithm and a multi-dimensional performance evaluator. Schedule feasibility is assessed using a discrete event simulation (DES) that incorporates operational uncertainty, variability, and resource availability. The tools provide rapid scenario assessments and responses to variances in the operational environment, and have been used to inform major equipment investments and workforce planning decisions in multiple SNL facilities.

1 INTRODUCTION

The upcoming weapon programs require an aggressive increase in production demands within the Nuclear Security Enterprise (NSE). In addition, there is a significant portion of development work that causes substantial variability, unbalanced and unstable resource utilization, reactive resource planning, and frequent reshuffling of priorities without quantifying and forecasting impact. A multi-disciplinary team has designed, developed and applied unique modeling, simulation, and optimization tools to support the long-term sustainment and stewardship of the NSE capabilities and resources to ensure successful mission execution. These analytic capabilities have been instrumental in improving Sandia National Laboratories' (SNL) component production productivity and efficiency, identifying optimal operational and tactical execution plans under tight resource constraints, optimizing schedule margin for component availability, and providing confidence in SNL's capability to deliver. This paper discusses the application of this analytic capability to the production of Application Specific Integrated Circuits (ASIC) at SNL. With ten products that have highly variable demand profiles and processes, a single set of shared resources, and the need for external supplier task synchronization, scheduling quickly became an integral part of successful manufacturing. A generalized automated scheduling application, Schedule Management Optimization (SMO), has been developed using an iterative multi-objective genetic algorithm and a multi-dimensional performance evaluator. SMO identifies production, inventory, and resource utilization profiles, as well as external supplier delivery plans required to meet customer demand with level-loaded and stable resource utilization while maintaining safety stock and schedule margin. Schedule feasibility is assessed using a stochastic representation of processes in a discrete event simulation (DES) that incorporates operational uncertainty, variability, and resource reliability. Planned-versus-actuals assessments determine when recovery is necessary and are iterated with SMO and the DES for recovery identification. The suite of tools is easily tailored to many resource-constrained scheduling problems and provides rapid scenario assessments and responses to variances in the operational environment. This analytic capability has been used to inform major equipment investments and workforce planning decisions made in multiple SNL facilities to ensure production readiness and stability.

2 DESCRIPTION OF THE PRODUCTION PROCESS

This section will give a brief description of the ASIC post-fab production process, so that the reader will be able to better understand the constraints under which the problem is formulated.

2.1 Packaging

The first phase of post-fab production is packaging. This starts after wafers are fabricated and includes a series of tasks, tests, and analysis, including a month of processing at an external vendor. Because external vendors are resource constrained and can process a limited number of wafer lots at a time, packaging activities must be carefully scheduled to ensure sufficient packaged parts are available onsite before the testing phase of production begins. At the end of the package and assembly process, one lot of wafers is split into sublots of hundreds or thousands (depending on the product type) of individually packaged die. After package and assembly, the physical construction of a packaged ASIC is complete, and the ASICs are ready to begin the next phase.

2.2 Testing

In the next phase, the ASICs undergo electrical and stress testing to identify die that do not meet specification requirements and to decrease infant mortality. Tasks in this phase include electrical testing, dynamic burn-in, environmental stress screening, visual inspections, and quality review. Parts recirculate through electrical testing more than once, pre and post burn-in. In this phase, a group of several hundred ASICs progress through the production line as a sublot. The sublot quantities are optimized for a given demand profile, rolled throughput yields, availability of resources, and span times for a particular product type. Typically a sublot quantity is much less than the size of a packaging lot due to the constraints of resources that all 10 products are competing for. A sublot must finish testing before it can begin the next phase, qualification. In qualification, samples from the sublot are further tested according to requirements dependent upon the pedigree of the delivery. As with packaging, testing activities must be carefully scheduled to begin only when packaged parts are available, and finish in time for these scheduled qualification tasks.

2.3 Qualification

This phase consists of a series of tasks that qualifies tested ASICs for final delivery. These tests conduct accelerated-life testing on a sample of components from one or more wafer lots thereby consuming the sampled parts. Qualification must complete before parts can be delivered but can start no earlier than the required sample size is available in tested parts inventory. Qualification requirements vary based on the pedigree of the delivery, Process Prove-In (PPI), Qualification Evaluation (QE), or War Reserve (WR). The PPI delivery lots do not require qualification. QE delivery lots, on the other hand, require a relatively large sample size of parts, and WR delivery lots require a smaller sample size. In addition, there are quality reviews performed throughout the production process as well as independent product acceptance tasks to objectively review quality evidence before product is delivered to customers.

3 SCHEDULE MANAGEMENT OPTIMIZATION (SMO)

3.1 Background and Motivation

In this section, we describe a heuristic to optimize ASIC production schedules using a genetic algorithm (GA). The production of ASICs requires synchronizing packaging, production, and qualification activities in order to satisfy scheduled product deliveries. That these activities dynamically utilize multiple resources simultaneously makes optimally scheduling ASIC production a generalization of the classic job shop problem and NP-Hard. The application of traditional heuristics, such as the Shifting-Bottleneck Heuristic (Pinedo 2012), is problematic as production activities can recirculate amongst the same resource multiple times. The optimization is further complicated by feasibility considerations of intermediate production inventories and yield losses. Mixed-integer linear programming (MILP) approaches have been formulated for multi-stage flow shop models of chemical production processes with constraints for availability of intermediate reagents and yield loss (Floudas and Lin 2004, Floudas and Lin 2005). However, these MILP approaches require large-scale problem formulations due to requirements imposed by either the *a priori* discretization of time intervals or the auxiliary variables associated with continuous-time formulations.

To develop a tractable optimization approach that can easily incorporate ASIC production business rules, we propose a GA-based scheduling heuristic. The heuristic utilizes Technology Management Optimization (TMO) software, a GA optimization suite developed at Sandia National Laboratories, to specify a set of decision variables that is used to construct a schedule. These variables are passed from TMO to an external scheduler evaluator, which employs a deterministic algorithm that constructs a feasible schedule. The evaluator assesses the schedule's performance with respect to two objectives and returns the performance values to TMO. Figure 1 provides a high-level depiction of a single iteration of the heuristic. In the next section we describe the algorithm for constructing schedules in detail.

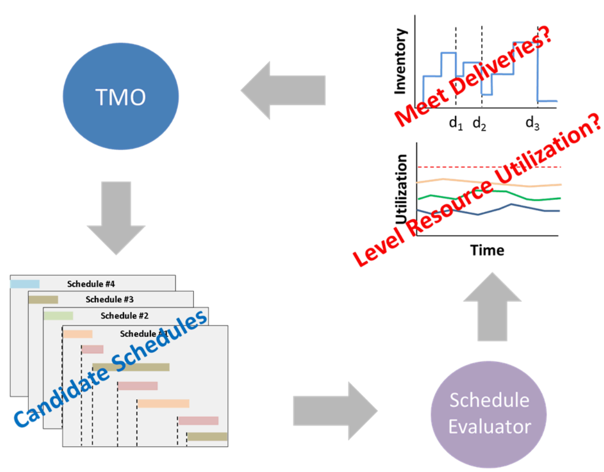

Figure 1: Depiction of a single iteration of the schedule-optimization heuristic.

3.2 Schedule Evaluator Algorithm

Here we describe in detail how schedules are constructed by the schedule evaluator. The schedule evaluator employs a deterministic algorithm to obtain a feasible schedule using decision variables provided by the GA. A key element of this approach is identifying decision variables that are sufficient to characterize an entire production schedule. An important consideration is that packaging, testing, and qualification tasks must be scheduled such that the inventories of each product's packaged and tested parts remain positive. A schedule is considered *inventory infeasible* if a testing task starts when the packaged-parts inventory is less than the test batch quantity, or a qualification task starts when the tested-parts inventory is less than the qualification requirement. Simply defining decision variables to denote the start time of each task is problematic in this approach because it is highly improbable that the GA will ever produce an inventory-feasible solution.

To ensure inventory feasibility, we define decision variables that represent a sequence of test-batch sizes for each product. Let $X_{ij} \in \{0, b_1, b_2, \dots, b_M\}$ denote the quantity of the i^{th} batch of the j^{th} ASIC product tested and b_k denote the k^{th} batch size, where $i = 1,2,\dots N$, $j = 1,2,\dots 10$, and $M, N \geq 0$. We refer to X_{ij} as a *task* and its $M+1$ possible states as *modes*. The purpose of including mode 0 is to temporarily suspend a product's testing to free resources for other products. If $X_{ij} = 0$, then task $X_{i+1,j}$ cannot begin until a one-month period has elapsed. Note that due to yield losses, the quantity of packaged parts entering a testing task is less than the quantity of tested parts produced. However, for illustrative purposes, we will omit consideration of yield losses from testing in the discussion that follows.

Once the sequence of test-batch sizes is known for a given ASIC, the associated packaging tasks to support that test sequence can immediately be determined from the expected number of parts yielded by each packaging task, denoted $Y_p^{(j)}$, $j = 1,2,\dots,10$. To illustrate how packaging tasks are derived from a sequence of testing tasks, consider an example where $Y_p^{(j)} = 300$, $N = 6$, and the following sequence of testing-batch sizes: X_{1j}=100, X_{2j}=300, X_{3j}=200, X_{4j}=300, X_{5j}=100, X_{6j}=200.

Figure 2 illustrates the process by which packaging tasks are determined. Let $Q_{ij}^{(1)}$ denote the quantity of packaged parts in inventory for ASIC type j *just before* the i^{th} testing task. If $Q_{ij}^{(1)} < X_{ij}$, then a packaging task is required before the i^{th} testing task can start. To ensure the required packaging task completes before the associated testing task begins, we impose a strict precedence relationship between the task pair. With this approach, we can guarantee feasibility of packaged inventory.

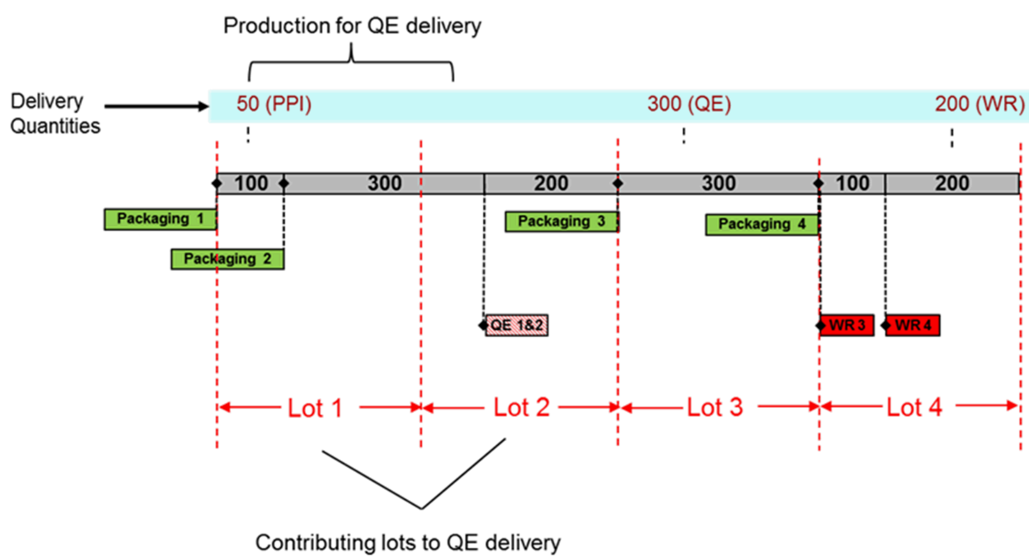

Figure 2: Assigning packaging and qualification tasks to a test-batch sequence.

Maintaining feasibility of tested-parts inventories is more complicated owing to the different pedigrees of product deliveries and their associated qualification requirements. In particular, the qualification task associated with a QE delivery requires a substantial quantity of tested-parts inventory, and the tested-parts obtained from a single packaging task may not be sufficient. WR qualification tasks, on the other hand, require substantially less tested parts. Additionally, qualification tasks should be completed as soon as possible to ensure there are enough qualified inventories to meet deliveries. To balance the requirements of sufficient tested-parts inventories with timely completion of qualification tasks, we employ a scheme to associate qualification tasks with both packaging and testing tasks. In the case of a QE qualification, we identify the last packaging lot that contributes parts to the QE delivery by computing when the cumulative quantity of packaged-parts exceeds the quantity of parts consumed by PPI deliveries, QE deliveries, QE qualification, and expected yield losses during testing. Once the packaging task for the final contributing lot is identified, the first testing task following it is made a precedent of the QE qualification task. WR qualification tasks are assigned precedents in a similar manner. For every packaging lot that contributes strictly to a WR delivery, the first testing task following it is made precedent of the WR qualification task. Figure 2 illustrates the assignment of QE and WR qualification tasks. Note that this scheme of assigning precedence relationships to qualification tasks ameliorates, but does not guarantee, feasibility of tested-parts inventory. If an insufficiently large test-batch size is selected following a packaging task, there may not be enough tested-parts inventory to conduct the subsequent qualification task. To prevent schedules in which such shortages occur, the objective function includes a penalty term for violations of tested-parts inventory and will be further discussed in a subsequent section.

In addition to inventory feasible, the schedule must also be *resource feasible*. That is, the demands on a resource at a given time must not exceed the resource quantity available. The packaging, testing, and qualification tasks dynamically require multiple resources throughout their durations. We represent the dynamic use of resources using right-continuous, piecewise-constant functions termed *profiles*. Figure 3 provides an example task in which an operator performs setup and shutdown on two machines.

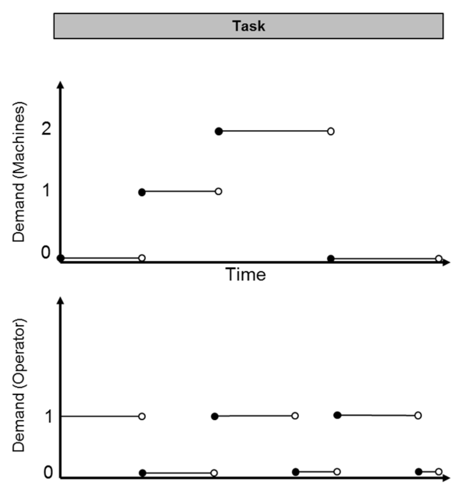

Figure 3: Example profiles of dynamic resource use.

When a task is scheduled, the resource profiles are translated by the task start time. Summing the translated profiles of all scheduled tasks provides a profile of the resource's overall use. Maintaining resource feasibility is straightforward. When a start time is considered for scheduling a task, the translated profile of each task resource is added to the profile of the corresponding resource's overall use given the tasks previously scheduled. If a resource's use exceeds the quantity available, the start time is infeasible. In this case, the next start time considered is the first time after the infeasible start time in which a change occurs in the overall use profile of that resource. Scheduling tasks sequentially in this manner will always maintain resource feasibility, but it does not guarantee precedence relationships are satisfied.

To maintain the precedence relationships between the packaging, production, and qualification tasks, each product's tasks are ordered with respect to their position in the hierarchy of precedence relationships.

This ordering is obtained by creating a directed network in which the tasks are represented bynodes. We denote by $i \rightarrow j$ the precedence requirement of task i on task j, where $i, j \in \{1,2, \dots W\}$ and $W$is the total number of tasks. A directed arc, A_{ij}, connects all node pairs for which $i \rightarrow j$ holds, $i, j \in \{1,2, \dots W\}$ and is given unit length. Applying the Critical Path Method (CPM) algorithm to this network yields the maximum distance to each node, and the corresponding tasks are ordered by increasing maximum distance. Where ties occur, tasks are further ordered by decreasing duration. Figure 4 provides an example where $W = 4$.

Task	Duration (hr.)	Precedents	Maximum Distance	Order
1	1.5		0	2
2	2		0	1
3	3	1, 2	1	3
4	1	2, 3	2	4

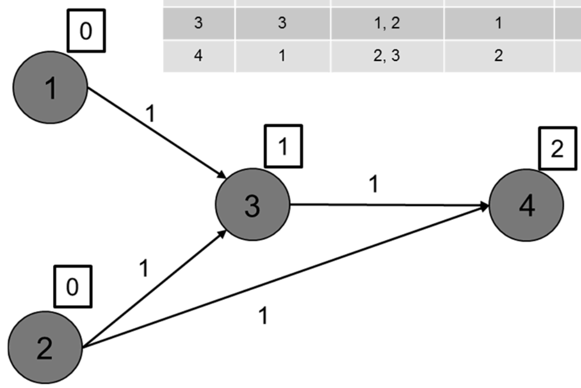

Figure 4: Assigning orders to tasks for scheduling

When tasks are scheduled in order according to the critical path, all the precedent tasks of a given task are already scheduled when the task is considered for inclusion in the schedule. Maintaining precedence relationships merely requires ensuring the task starts no earlier than the latest completion time of all its precedents. There are no precedence relationships between the tasks of different product types, so the ordered task lists of the products are combined into a single list. To maintain determinism of the scheduling algorithm, tasks from different products of equal order are further ordered by their corresponding product indices. Figure 5 summarizes the algorithm for creating an ordered task list for each product, and Figure 6 summarizes the algorithm for scheduling all products' tasks from the combined list. The next section discusses the optimization objectives and fitnessfunctions.

For each ASIC type $j = 1,2,\ldots,10$,

1. Initialize testing tasks of batch size X_{ij}, where $i \to i+1$, $i = 1,2,\ldots N$.
2. Set $k = N+1$. For each i, where $Q_{ij}^{(1)} < X_{ij}$,
   i. Initialize packaging task k, where $k \to i$.

   Next k
3. Initialize QE task l, where $i \to l$ and $k^* \to i$, where k^* is the packaging task associated with the last contributing wafer lot to the QE delivery.
4. Set $l' = 1$. For each packaging task k^{**} associated with wafer lot that strictly contributes to a WR delivery,
   i. Initialize WR task l', where $i \to l'$ and $k^{**} \to i$.

   Next l'
5. Apply CPM to directed network that represents task precedence relationships.
6. Order tasks by increasing maximum distance and decreasing duration.

Next j

Figure 5: Summary of algorithm to create ordered task lists.

1. Let L be the ordered list of all product tasks. Set schedule $S = \emptyset$.
2. For each task $\mathrm{T_i} \in L$,
   i. Compute $\gamma = \max(c_j)$ for all $T_j \in L$ such that $j \to i$, where c_j is the completion time of T_j.
   ii. Set $\theta(S)$ to be the set of all times when profiles change given S.
   iii. While not feasible,
      - Set $t = \min(t' : t' \geq \gamma, t' \in \theta(S))$.
      - Set $\mathrm{s_i} = t$, where $\mathrm{s_i}$ is the start time of T_i.
      - Evaluate resource feasibility of schedule.
      - If not feasible, set $\theta(S) = \theta(S) \setminus \{t\}$.
   iv. Set $S = S \cup \{T_i\}$.

   Next i

Figure 6: Summary of algorithm to schedule tasks.

3.3 Optimization Objectives and Schedule Fitness

The multi-objective optimization seeks schedules that maximize performance with respect to two key characteristics: (1) inventory performance and (2) resource utilization. We employ separate nonnegative performance functions for each characteristic. The functions return a value of *zero* for schedules with the best possible performance and, we seek schedules that minimize each performance function. The result of the multi-objective optimization is a set of Pareto-optimal solutions with respect to both functions. We now discuss each function in turn.

The inventory performance function penalizes shortages in qualified-parts inventory, corresponding to unmet demands, shortages in tested-parts inventory, corresponding to inventory infeasibility, and terminal inventories above required the safety stock, corresponding to overproduction. Because a schedule with a negative tested-parts inventory is physically impossible to execute, the penalty coefficient for tested-part shortages is substantially larger than the coefficients associated with unmet demands and excess production. Likewise, overproduction is preferred to missed demands, so the penalty coefficient for overproduction is less than that for unmet demand. Let $D_j = \{d_{1j}, d_{2j}, \ldots, d_{\beta_j j}\}$, where the d_{ij} is the i^{th} demand quantity of product j, and let $T_j^{(D)} = \{t_{1j}^{(D)}, t_{2j}^{(D)}, \ldots, t_{\beta_j j}^{(D)}\}$, where $t_{ij}^{(D)}$ is the i^{th} demand time of product j, $i = 1,2,\ldots,\beta_j$, $j = 1,2,\ldots,10$, and $\beta_j \geq 0$. To denote inventory levels, let $Q_j^{(2)}(t)$ and $Q_j^{(3)}(t)$ be functions of tested-parts and qualified-parts inventories of product j, respectively, over time, and let $T_j^{(2)}$ and $T_j^{(3)}$ be the set of times at which these functions change values, $j = 1,2,\ldots,10$. The required terminal safety stock of product j is denoted ω_j, $j = 1,2,\ldots,10$, and the duration of the entire schedule is denoted by $\bar{T}$. The inventory performance function, denoted P_I, is as follows:

$$P_I = \gamma_1 \sum_{j=1}^{10} \left(Q_j^{(3)}(\bar{T}) - \omega_j\right)^+ + \gamma_2 \sum_{j=1}^{10} \sum_{i=1}^{\beta_j} \left(d_{ij} - Q_j^{(3)}(t)\right)^+ + \gamma_3 \sum_{j=1}^{10} \sum_{t \in T_j^{(2)}} \left(-Q_j^{(2)}(t)\right)^+,$$

where $\gamma_3 \gg \gamma_2 > \gamma_1 > 0$ and $y^+ \equiv \max\{0, y\}$ for $y \in \boldsymbol{R}$.

The resource utilization function is used to assess the degree to which resource use varies over the schedule duration. Schedules that impose less variation in resource use are favored. To reconcile the differences between the magnitudes of each resource quantity, the performance function utilizes a concept analogous to the notion of coefficient-of-variation for a random variable. Let $U_i(t)$ be a function for the overall use of resource i at time t, $i = 1,2,\ldots.,R$, where R is the total number of resources. The resource-use performance function, denoted P_R, is as follows:

$$P_R = \sum_{i=1}^{R} \frac{(V[U_i(t)])^{1/2}}{E[U_i(t)]}, \text{ where}$$

$$E[U_i(t)] \equiv \bar{T}^{-1} \int_0^{\bar{T}} U_i(t)dt \quad \text{and} \quad V[U_i(t)] \equiv \bar{T}^{-1} \int_0^{\bar{T}} [U_i(t)]^2 dt - E[U_i(t)]^2.$$

4 DISCRETE EVENT SIMULATION & RECOVERY IDENTIFICATION

4.1 Discrete Event Simulation

This section describes the DES model that is used in conjunction with the SMO from Section 3. In general, the ASIC DES model is a detailed stochastic representation of the production processes, resources, and operational business rules such as Military Standards and rework. To complement the scheduler, which is based on a deterministic algorithm, the DES incorporates uncertainty and variability. The stochastic processes and events that play a significant role in tactical execution planning, such as parts failing a test, lower than expected yields, or equipment breaking down, are modeled.

The schedules identified with the SMO are input into the DES and include the packaging, testing, and qualification tasks. Other inputs, such as process flows, times and quantities, required resources, yields, failure modes, and scheduled maintenance, are read in from the same database the scheduler reads from. A single replication of the model runs through the 9 year production program.

Statistics are collected, via multiple replications, on schedule performance, resource utilization and state profiles, and span times by product and functional area. Sensitivity analyses are conducted to quantify the impact of operational uncertainty, variability, and resource reliability on schedule performance. Multiple runs of the stochastic simulation identify schedule risk and resource constraints and are used to inform resource investment and acquisition decisions. Resource utilization profiles identify opportunities for scheduled maintenance activities with the least disrupt to operations. The DES provides a means of dynamically tracking ASICs through production, identifying schedule margin, projecting downstream workload and requirements, and proactive operational planning.

4.2 Recovery Identification

When catastrophic sublot failures and unexpected delays occur, or factors beyond the scope of the model affect production in unforeseen ways, recovery and re-planning may be necessary. Planned-versus-actuals assessments are used to determine when re-planning is necessary and actuals are iterated with the scheduler for recovery identification. The scheduler provides rapid scenario assessments and responses to variances in the operational environment. When new schedules have been identified, they are run through the DES as before to determine the effectiveness of the recovery options. Figure 7 provides a high level depiction of the iterative analysis cycle.

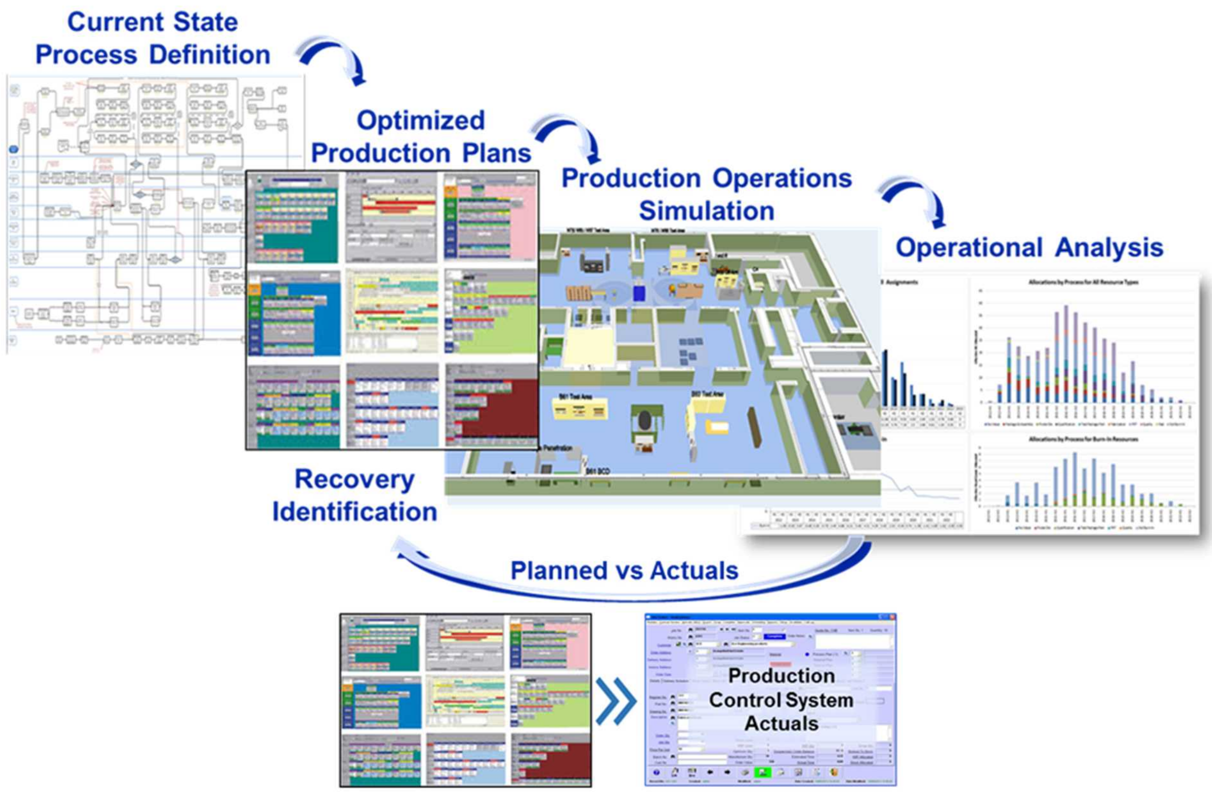

Figure 7: Iterative Analysis Cycle.

5 CONCLUSION

This analytic capability provides model-based decision support and proactive operational and resource planning for Sandia's upcoming production peak and out year ASIC demand profiles. The tools enable Sandia to examine the trade space for the best improvement opportunities to operate optimally and with increased efficiency and productivity, and generate confidence in the ability to deliver. They have been used to inform major, multi-million dollar equipment investments, inform suppliers of expected peaks in production rates, inform workforce planning decisions to include hiring, and have played a vital role in ensuring production readiness and stability.

ACKNOWLEDGMENTS

We would like to acknowledge the following individuals at SNL for their contributions to the development of this capability, for providing application domain knowledge and expertise, for verifying and validating the model and its results, and for providing financial support and resources to make this work possible. Special thanks to David Sandison (Director of Microsystems Science & Technology), Rita Gonzales (Senior Manager of Firing and Embedded Systems), Reno Sanchez (Senior Manager of Microsystem Design/Microsystems & Engineering Sciences Applications Production), Guillermo Loubriel (Manager of Validation & Failure Analysis), Kaila Raby (Manager of Reentry Fusing I), Perry Molley (Manager of Nuclear Weapons (NW) ASIC Product Realization), Nathan Nowlin, David Brumit, Raivo Leeto, and Jamie Levy (NW ASIC Product Realization Team), Stephen Garrett (Microsystems Integration), Dennis Nelson (Packaging and Assembly Operations), Ricky McFarland (Validation & Failure Analysis), Kevin Fox (Validation & Failure Analysis), Frank Loudermilk (Manager of Defect Prevention Engineering), Wenona Ayarbe and Amanda Hawkins (Defect Prevention Engineering), Frederick Sexton (Manager of Advanced Microelectronics and Radiation Effects), Marty Shaneyfelt (Advanced Microelectronics and Radiation Effects), and finally, to Bruce Thompson (Manager of Systems Readiness and Sustainment Technologies), and Craig Lawton (Manager of Systems Analysis and Decision Support).

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

REFERENCES

Floudas C. A., and X. Lin. 2004. "Continuous-Time Versus Discrete-Time Approaches for Scheduling of Chemical Processes: A Review." *Computers & Chemical Engineering* 28: 2109-2129

Floudas C. A., and X. Lin. 2005. "Mixed Integer Linear Programming in Process Scheduling: Modeling, Algorithms, and Applications." *Annals of Operations Research* 139: 131-162.

Pinedo, M. L. 2012. *Scheduling: Theory, Algorithms, and Systems*. 4th ed. New York Dordrecht Heidelberg London: Springer.

AUTHOR BIOGRAPHIES

JOHN A. FLORY received his Ph.D. in Industrial Engineering from the University of Pittsburgh and is a Senior Member of the Technical Staff at Sandia National Laboratories. His research interests include stochastic modeling, optimization, and simulation. His email address is jaflory@sandia.gov.

DENISE D. PADILLA received her M.S. in Electrical Engineering from the University of New Mexico and is a Principal Member of Technical Staff at Sandia National Laboratories. Her research interests include modeling, simulation, optimization, and analysis capabilities for decision support to include production planning, capacity modeling, and scheduling. Her email address is ddpadil@sandia.gov.

JOHN H. GAUTHIER received an M.S. in Computer Science from the University of New Mexico and is a Principle Member of the Technical Staff at Sandia National Laboratories. He is interested in NP problems, including combinatorial optimization and scheduling. His email address is jhgauth@sandia.gov.

APRIL M. ZWERNEMAN received her M.S. in Operations Research from the University of Texas at Austin. She currently works in Systems Analysis & Decision Support at Sandia National Laboratories. Her email address is azwerne@sandia.gov.

STEVEN P. MILLER received his B.S. in Nuclear Engineering from Kansas State University. He is a Principal Software Engineer at Science Applications International Corporation (SAIC), Inc. His interests include probabilistic risk assessments, total system performance assessments, photovoltaic systems analyses, and operations analyses. His email address is spmille@sandia.gov.