

Optimizing Supply-Chain Planning

John W. Bean, Technology and Manufacturing Group, Intel Corporation
Amit Devpura, Technology and Manufacturing Group, Intel Corporation
Michael O'Brien, Technology and Manufacturing Group, Intel Corporation
Shamin Shirodkar, Technology and Manufacturing Group, Intel Corporation

Index words: supply chain, optimization, linear programming, production planning

ABSTRACT

Semiconductor manufacturing is a very capital-intensive endeavor that can return substantial revenues. The production planning process must deliver a build schedule that makes efficient use of Intel's capital resources while satisfying as much demand as possible. This schedule should comprehend the flexibility of production resources, the dynamic nature of supply and demand within Intel's supply chain, as well as the timing of new product releases and production facility improvements.

Previous planning processes relied on spreadsheets for heuristic manual decision making with localized data. With the growing complexity of Intel's products and manufacturing processes, these methods had become inadequate and unsustainable. Upgrading the planning process required better decision algorithms, improved data management, as well as more automated and integrated planning processes.

New tools based on Mathematical Programming were implemented in multiple divisions and stages of Intel's supply chain. The development team worked closely with the users to understand their business and capture their operating logic to create automated decision systems. These tools balance requirements to satisfy demand, achieve inventory targets, and remain within production capacity to reduce costs and satisfy demand across Intel's supply chain. They have been developed to evolve the planning process while maintaining visibility to the logic and data flow to facilitate continuous improvement.

Advances in data management were required to complement decision algorithm improvements. The new tools integrate directly into source data systems while providing planning- and optimization-specific functionality, including mechanisms to track parameter changes and supply dynamic reporting capabilities. These advances allow planners to more easily identify data issues and to better understand the planning

recommendations from the tools. The robust data management infrastructure enables tighter integration of organizations, increased scalability, and more consistent implementation of solutions across business units.

Advances in decision algorithms, data management, and system automation led to improvements in solution quality, data health, and productivity. The new applications allow planners to rapidly perform analyses on multiple business scenarios to produce better solutions and improve collaboration with other organizations. While results reported by the business users over the past four years have proven the stability and value of this decision support technology, there is still work to be done. Plans for extensions and continuous improvement are provided in the last section of this paper.

INTRODUCTION

Semiconductor manufacturing typically proceeds through three major manufacturing stages as shown in Figure 1.




Figure 1: The basic flow in semiconductor manufacturing

First, hundreds of complex devices, each containing millions of transistors, are fabricated on silicon wafers. End-of-line testing sorts functional devices from ones that have manufacturing defects. Second, the wafers are sawn to yield individual devices (called die) that are assembled with substrates that supply physical protection and electrical connectivity. Final testing establishes the detailed performance characteristics of the semi-finished product. Third, after configuration and marking, the final product is packed and shipped to the customer. Basic

planning for these stages includes deciding the timing and quantity of wafer releases into fabrication/sort facilities (F/S), die and substrate releases into assembly/test factories (A/T), and semi-finished goods releases into configure/pack facilities (C/P), as well as assuring the availability of substrates to support the A/T plan. These planning decisions are made more complex by Intel's risk management method of having multiple F/S, A/T, and C/P facilities. Each facility manufactures multiple products, and each product is produced in multiple factories.

Product differentiation must be comprehended in the planning process. As shown in Figure 2, there is a stochastic but measurable distribution of performance among the die that exit F/S. This information along with demand is used to decide the appropriate substrates for specific die. Consider microprocessors for example. The higher speed die should normally be placed in server substrates, the die that consume the least power in laptop substrates, and some of both in desktop substrates. Products exiting the A/T process also exhibit a stochastic but measurable distribution of final speed performance. Faster and slower semi-finished products are segregated. The planning system needs to consider these distributions.




Figure 2: Product differentiation in the production flow

Additional planning decisions must be made for the configuration step as illustrated for desktop products in Figure 2. One feature of semiconductor production is that faster semi-finished products can be configured to run more slowly (6 GHz to 5 GHz and 5 GHz to 4 GHz in Figure 2) depending on demand. Unfortunately, slower products cannot be made to run faster. Furthermore, there are multiple flows that make the same product. Notice in Figure 2 that the slowest of the fast die and the fastest of the low-power die that were put into desktop packages converge to both make the 5 GHz final product. The planning system must comprehend that these two phenomena combine for desktop products to provide

many ways to make the 4 GHz and 5 GHz products. Server and laptop products exhibit similar behavior.

These planning complexities are encountered across Intel's microprocessor, memory, and communications businesses, and over strategic and tactical planning problems. The business problem faced by all is to minimize costs and maximize revenues now and in the future. A number of factors make this a very difficult problem for Intel. On the supply side, there is considerable variability in manufacturing throughput time (TPT), line yields (the proportion of product that survives the manufacturing process), and product performance (of the surviving product, there is a distribution of operating speed, for example). There is also considerable variability on the demand side with customers placing orders that they later alter in quantity, due date, and/or delivery location, or cancel altogether. Manufacturing TPTs are much longer than the time it takes for our customers to change their minds and their orders. Product differentiation, complex and interrelated Bill-Of-Materials (BOMs), and shared capacity among product lines make planning even more complex.

Poor plans can waste valuable manufacturing capacity by either letting it sit idle or by using it to make products that the market decides it doesn't want. A slow planning system can eventually produce a plan that, by the time it is issued to manufacturing, is based on stale data making it unexecutable or sub-optimal. The dual goals of minimizing costs and maximizing demand satisfaction depend on making the right volumes of the right products at the right times. Building efficient plans in a timely fashion is critical to our business success.

This paper begins by describing previous planning approaches and then outlines our development approach. We describe a suite of successful applications across Intel's businesses, quantify the realized benefits, and discuss some of the potential future extensions and improvements.

THE PREVIOUS PLANNING APPROACH

Intel's planning systems and business processes have evolved from the simpler processes that supported limited product segmentation and less complex manufacturing routes in the past. Evolving business conditions create planning requirements that need to consider the high-mix environment where differences in packaging, test parameters, wafer type, product lifecycle, and processing drive different planning decisions. Embedded in the existing systems and processes is business logic referred to as "rules" that has evolved to support operation of Intel's complex supply chain. Many of these rules were held in planners' heads and were never formalized. Extracting these rules and forging them into one coherent

consistent set that all planners would use had never been attempted.

Historically, this complex problem has been attacked with spreadsheets and massive amounts of human capital. Previous planning processes relied on localized data management and heuristic decision making that were not able to consider all relevant factors or respond fast enough to support Intel's planning needs. The spreadsheets supported the heuristic methods that planners developed over the years and handed down through generations of new planners. Long hours were spent, including evenings and weekends, in collecting and repairing data and in executing a cumbersome process.

In the best case, a locally good solution might be found, but just as often time ran out and a partially refined plan was issued. The limitations of an extended analysis time and inadequate accuracy prevented planners from exploring all decision options to develop an optimum build plan. While this might have been adequate with simpler processes and products and a less sophisticated marketplace in the past, given today's conditions, it was clear that this process and its associated tools were inadequate and unsustainable.

OUR DEVELOPMENT APPROACH

Not all development approaches would be successful in transforming planning systems from those described above to solutions that can support both current and future business needs. As with most transformations, the difficulty is as much about moving the business personnel to a new process as it is about accomplishing this transformation while the business continues to move, adapt, improve, and respond in a dynamic market. Experience has shown that a successful approach must have the following characteristics.

- Collaborative: Frequent interactions with "fingers-on-the-keyboard" users to validate design options and verify priorities.
- Incremental: Multiple smaller projects that move forward a step a time, often helping to clarify the overall problem and success metrics.
- Iterative: Frequent development cycles with a validation checkpoint after each development effort and before the next requirements refinement step.

Such an approach starts by seeking to understand the current approach and by developing a solution with limited changes to rules and data feeds. It is only in taking this first step that we can begin to understand the health of the data and the strength of the rules that are used. Getting the data and rules documented provides a starting point for further improvements and brings the advantage of

standardizing how this problem is solved. More benefit will be realized from subsequent efforts when the current process is really challenged, but the automation of the current process is a necessary predecessor for any other improvements. It is not possible to make changes until we understand the details of the current process. This effort follows the steps shown below, which have been refined during many years of decision algorithm pathfinding.

Step 1. Shadow the end-users in the business. Motivated by the understanding that "the devil is in the details," the first step in decision algorithm pathfinding is to shadow the end-users to understand how the current process works, even to the extent that pathfinders would be trained on existing tools as if they were a new planner. Included in this initial analysis is understanding what data are used to make planning decisions and which business rules are used around each set of data. An important stage in pathfinding occurs when pathfinders believe that they understand the user's algorithm. A more important stage is reached when the user believes that the pathfinders understand the algorithm.

Step 2. Develop a prototype and validate it against the current process. The best and perhaps only way to validate that the decision algorithm has been appropriately captured is to transform it into a working prototype and put it into the hands of the users. The expectation is that, when used in close temporal proximity to actual problem solving with the same production data, the user will quickly point out parts of the prototype that are missing or wrong. Rapid iteration is crucial here to hold the users' interest and confidence. Only when the user proclaims that the prototype is producing plans that are as good as (or better than) the current method and producing them as quickly (or quicker) can pathfinding move on to the next step.

Step 3. Make the solution production-worthy. After the prototype is accepted by the planners, it is stabilized to support regular production use. This included code refinement for both the model and the interfaces to the database as well as extensive module and integrated testing. Moving to this step too soon will slow and misdirect prototyping efforts, but never moving to this step will unnecessarily increase business risk and retard continuous improvement.

The desired properties of the tools resulting from this development approach include improvements in both data and decision algorithms delivering benefits in productivity and solution quality. Improved data management including automated data loads would result in fewer errors from manual data input, easier recognition of issues with input data, and fewer planner hours to correct errors in the data. It would also enable rapid evaluation of the resulting plan, support understanding of the sensitivity of

business scenarios to changes in the input data, and streamline dissemination of the results. Decision algorithm automation would guarantee rapid and consistent application of the business rules requiring fewer planner and total hours to execute. It also would provide a foundation for continuous improvement by documenting all the standard (and exception) rules and making them easy to extend and test. Automated decision algorithms would allow planners to explore various business scenarios and drive their recommendations based on that understanding.

Since the purpose of planning is to align supply decisions with forecasted demand, selecting the appropriate algorithmic approach was simple. This type of problem has long been solved in academia and at some other companies with a mathematical technique called mathematical programming, or optimization (Hopp and Spearman 1996, Chopra and Meindl 2001). In a mathematical program, business rules are translated into constraining equations that limit the values of the decisions that the solver is making, and objective functions that quantify the total value of the solution against objectives of the business.

A Linear Program (LP) is a subset of mathematical programming where business rules can be represented as linear equations. An LP solver will make constrained decisions to maximize (or minimize) a linear objective function. Examples of decisions an LP could make are how many wafers to start at a Fab in each week and how much product to allocate to different packaging types and configurations as shown in Figure 2.

Once business rules are represented in linear equations, as either the objective or constraint, an LP can be used to solve very large problems in a short amount of time. For example, a typical LP problem with 150,000 variables can be solved in less than a minute. In addition to rapidly generating an answer, the LP solver generates an “optimum” answer. This means that given the objective function used in the LP, one can be certain that this is the best possible answer that is not influenced by alternate starting points or the order of business rules. The speed and quality of the solution produced by the LP allows planners to explore different input data scenarios (different demand, product priorities, or capacity statements) to understand plan dependency on these factors. An LP also allows evolution of the planning systems through modification of existing rules and the addition of new rules as business needs change.

In this paper, we describe the successful implementation of LP tools in Fab-Sort Manufacturing (FSM), Assembly Test Manufacturing (ATM), and Materials Procurement. This includes work with various product divisions including the Intel Architecture Group (IAG) that

produces microprocessors (CPUs) and support chips (Chipsets), the Flash Products Group (FPG) that manufactures a range of flash memory chips, and the Intel Communications Group (ICG) with their range of networking and communication products.

FAB/SORT MANUFACTURING

The monthly FSM Planning Reset process must deliver a build schedule for the next nine months that makes the best use of Intel’s fungible and expensive capital resources. This plan should comprehend the dynamic nature of supply and demand within the multi-month TPT of Intel’s supply chain, the inherent momentum of Intel’s supply including in-process wafers and die, product and process roadmaps, and manufacturing performance improvement projects. The FSM plan must also consider the product differentiation as shown in Figure 2 to better align planned supply to forecasted demand through modification of Fab wafer start schedules.

The FSM Solver for FPG

The first improvement occurred in the Flash midrange planning process in Q4’01. This project was small in terms of scope and resource involvement to determine the value of this approach and identify any potential roadblocks. The project focused on the automation of the decision process without any significant changes to the way data were entered or the way results were evaluated.

The planner worked with an optimization expert to transform the business rules used in the previous heuristic, manual process into automated decision algorithms. The current business rules that guided the manual process were translated into mathematical programming within the LP solver. Business rules such as “remain within limits of each Fab’s capacity” were translated into constraints (see Equation 1). Strategies such as “minimize missed demand (DemandMiss) and missed inventory targets (InvOver & InvUnder) utilizing relative penalties (MissPen & InvPen)” are translated into objective functions for the LP solver (see Equation 2). Of course these penalties should be translated into true dollar costs and this is on our continuous improvement plan.

$f = \text{fab}; p = \text{product}; r = \text{resource} / \text{process}; t = \text{time}$

$$\sum_{p \in r} \text{FabStarts}_{f,p,t} \leq \text{Capacity}_{f,r,t}$$

Equation 1: Capacity constraint

$$\min \left\{ \begin{array}{l} \sum_{p,t} MissPen_{p,t} \cdot DemandMiss_{p,t} \\ + \sum_{p,t} InvPen_{p,t} \cdot (InvOver_{p,t} + InvUnder_{p,t}) \end{array} \right\}$$

Equation 2: Objective function for Demand and Inventory target misses

The overall planning process using the new optimizing tool requires only 10% of the time required by the old process and achieves 100% of the metrics for a good plan set by manufacturing personnel. The old process scored only 85% on these good plan metrics necessitating a follow-up meeting to negotiate changes. The success of the new tool eliminated the need for this meeting and indicated that it was possible to capture all the decision rules used in solving FPG's FSM planning problem. It demonstrated that these rules could be represented in a set of linear equations and that the solver could be controlled with priorities and penalties that would make sense to planners and support business requirements. The amount of improvement in this project, both in terms of productivity and plan quality, indicated the potential for this approach to improving FSM planning in other areas of Intel's supply network.

There was a lateral implementation of this same planning tool in Q2'02 to the IAG Chipset division. The same data management system, which allows for loading of the input data from existing spreadsheets and inserting the results back into that same spreadsheet, was migrated to Chipsets.

FSM Solver for IAG CPUs

Building from the success in FPG and IAG Chipsets, the decision was made in Q1'02 to develop a similar tool within IAG CPU. The CPU planning problem had the additional complexities of product binning, product configuration, and more complex product mapping as shown in Figure 2.

The wider scope of this project inherited additional complexities but also provided more insights into the data management side of planning solutions. Through this project, we were able to evaluate the current health of planning data and characterize the nature of planning data problems. These problems included the complexity of distributed manual data ownership, integrating data from multiple systems, evaluating the results of an automated decision algorithm, managing data fallout, and maintaining traceability to changes to the data.

Utilizing the CPU solution as a starting point, a similar solution was implemented to support planning for some parts of ICG. The database/solver architecture will allow portions of ICG to modify and add business rules in the

form of additional constraints, objectives, and their supporting data feeds to supplement the existing application. The ability to translate the CPU solution for use in ICG again demonstrates the scope of the solution and the ability of the tool to respond to the changing needs of different divisions. Performance results for IAG and ICG in both planner productivity and plan quality were similar to those achieved for the initial FSM FPG tool.

FSM SOLVERS on SQL with .Net

Taking the next step in IAG CPU planning in the area of data management to enable further and more complex decision automation, it was recognized in Q1'04 that there was a need to transfer the application to a more robust SQL platform. SQL provides a data management platform that allows storage of more data and faster, more complex data manipulations. This data management improvement allows storage and comparison of multiple versions and any manual modifications to the data. The multi-user SQL environment allows parallel examination and manipulation of the data to speed up the reset process and enable more thorough investigation of the solution space through multiple solves.

As part of this transition, the team also took advantage of the benefits of .Net through development of interfaces and architecture in .Net. The .Net architecture allows more rapid prototypes and development going forward. It facilitates rapid changes to different data sources if business decisions require it. The .Net components allow for rapid development of reports and interfaces as well as a more robust interaction with the solver application including faster and better data transfer from and to solver applications. Overall, the time required for data management was decreased to 15% of its previous level.

Building off the success of the SQL/.Net transition into CPU, efforts began to migrate this solution to other product divisions. These migrations included modifications to both decision and data components of the application to comprehend division-specific requirements. Each division may have a different set of business rules based on the value of inventory and accuracy of forecasts. These efforts verify that the FSM wafer start planning process is similar enough between divisions to share the same application platform. This widespread implementation enables knowledge and process sharing between divisions.

Figure 3 provides an overview of the extent and focus of this sequence of FSM projects. Each project is shown in terms of what was gained in the area of both data management and decision automation. These efforts demonstrated the level of interdependence between data management and decision automation. The amount of decision automation is limited at some point by the level

of data management. A complex decision algorithm is of limited value if results can't be quickly and accurately evaluated. A planning solution that requires too much time to collect and prepare data will prevent planners from having time to generate a good plan and increase the risk of data issues impacting the timing and/or accuracy of that solution.

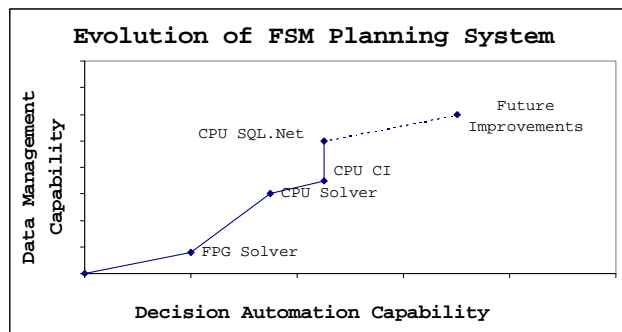


Figure 3: Data management and decision automation evolution of FSM systems

ASSEMBLY/TEST MANUFACTURING

ATM planning has the responsibility to match wafer supply from Fab to demand in the market by routing the right wafers to the right A/T factories and planning the right starts into the A/T factories. The ATM planning process had become very complex not only due to the increase in the number of products, but also the number of A/T factories and Fabs. The planning process for the ATM factories had become so complex for the manual process that it was not only divided by products but also decomposed into a request response process. The product division planners would do a detailed analysis on demand including a high-level analysis on supply resulting in a preliminary plan. They would then submit their preliminary plan and supply request to the ATM factory planners, who would respond based on a detailed analysis of the supply including splitting the builds to various A/T factories. A lot of time and energy was spent in this process including regular iteration of this request and response cycle.

The optimization techniques for planning were extremely well received as a tool that would prevent planning failures by automating the planning process, creating an effective work-life balance, and producing a better plan. The approach taken was to incrementally solve the problem by first automating the response process and then step by step integrating the whole request-response process.

A major challenge for the project was getting clean and consistent data. A lot of the Plan of Record (POR) data was not clean, and the clean-up effort had been

undertaken as a multi-year project. To provide good quality data, it was decided to get some data from the POR systems and some from the manual system used by the planners, which included formalization and maintenance of data that were stored and maintained in planners' heads. A plan to move from the manual data sources to the POR data sources was also formalized to intercept the data clean-up efforts.

The team had extensive user involvement and commitment that helped in finding the requirements faster. The technical team was very flexible with the requirements and created prototypes to test and confirm the business rules. With exceptional teamwork and a dedicated effort the team was able to put the solver tool in production for the response process for one product in four months in Q4'03. After that the response solver tool was proliferated to five products including the Intel® Pentium® 4 processor and products built on Intel® Centrino™ mobile technology, and will continue for new products.

After the successful implementation of the solvers for the response process, the team took the next step of integrating the request-response process. The solution needed to comprehend the different time zones of different planners, so that it would be robust for multiple users, and have traceability and scalability. The platform chosen for developing the solution was .Net with an SQL database. The team incorporated POR data sources for demand, but still used the manually maintained mappings of Figure 2. With dedicated team effort and dedication of the users, the tool was successfully implemented for one product in Q1'05. With the help of automation and optimization the business process has been completely reengineered to integrate the demand and supply planning process. This tool is now the POR tool for one of the Pentium products. It will be proliferated to all products for which we have the response tool, and it will be used to plan any future products including the multi-core products.

ATM planning is an example of how operations research has been used at Intel to not only build better plans but also re-engineer the business process. Integration of the request-response process would not have been possible without the optimization and automation tools. This implementation and integration of the planning processes

® Intel and Pentium are registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

™ Centrino is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

serves as an example for future enhancements and integration of the supply-chain planning for Intel.

These tools have not only provided strategic value by paving the way for future planning processes and systems, but have also provided tactical benefits. The response tools have provided an ROI of tens of millions of dollars per year due to better plans and have improved the work-life balance for the planners. The need to work on the weekends has been completely eliminated, working late has been reduced considerably, and work is more evenly distributed over the week. The integrated solver has reduced the time required for reconciliation between the central and A/T factory planning by 75%, and has improved the quality of the plan with better utilization of capacity and better demand support by 5% to 10%.

MATERIALS PROCUREMENT

Supplying Intel's A/T factories that build microprocessors and chipsets with the appropriate substrates is a complex task from a number of different perspectives. The combinatorial complexity of this planning problem stems from the large number of suppliers, package families and individual substrates, and A/T factories that are spread across many geographies, as shown in Figure 4. Note that each supplier and each factory use many substrates, and each substrate is made and used in many places to hedge uncertainty. Financial complexity is due partly to the very large expenditures made in substrate manufacturing and consumption, and partly to the complex contractual arrangements that surround substrates. These contracts protect all parties against short-term demand shifts between products and against uncertain future markets as new substrates are developed. Moreover, there is evolutionary complexity as the number of suppliers, packages, and factories grows over time, as the lifecycles of products decrease, and as the ramp rates for supplier and factory qualification increase.




Figure 4: Supplier, substrate, buyer complexity

In past years, before this problem became so complex, planning individual substrates by a trial and error methodology was common using multiple custom spreadsheets and extensive manual data management. Business rules, some embedded in contracts, some passed down the management chain, were contained in planners' heads and applied during spreadsheet runs. Understanding of suppliers' capacity was based on experience. Planning was undertaken in parallel for many substrates with the risk of simultaneous calls to the same supplier requesting conflicting actions. This process could include multiple iterations between suppliers and Intel planners and still produce a suboptimal plan for everyone involved.

This materials procurement process has been dramatically improved by the development of a centralized data base and a Web-based optimization tool. The data base provides a single repository for data about all suppliers, package families and individual substrates, and A/T factories. More importantly, it contains capacity models co-developed by Intel and its suppliers that accurately represent flexibility between substrates. The optimization tool is based on an LP formulation of the planning problem and addresses all suppliers, substrates, and factory needs simultaneously. Very significantly, it encodes the wide variety of business rules that must be considered.

This new business process is now executed in less than 10% of the time required by the old process. The resulting plan is capacity feasible without involving extensive iterations with suppliers and has saved Intel tens of millions of dollars on its expenditures since the system was implemented in Q3'03. Given solver run times that are less than five minutes, and a Web-based interface that provides overall system and data transparency, further savings are expected as the system is employed beyond simple planning. For example, the tool can support what-ifs around the capacity needed for substrates in development for future products as well as for future demand scenarios. In addition, it can be used to explore various pricing and contracting combinations before and during business negotiations. This successful approach to substrate planning is being proliferated to other materials planning problems in business groups across Intel.

BENEFITS

These efforts have provided qualitative and quantitative benefits to the organization, from improvements in data management to decision automation. The benefits of these efforts were realized in areas of productivity, data quality, solution quality, and continuous improvement.

Productivity

Copy and paste data collection tasks and subsequent checking were replaced by direct database links. Faster plan generation was provided through automation of decision algorithms with mathematical programming solvers. Faster plan analysis was realized through increased visibility into plan details and simplifies consideration of alternate scenarios. Dynamic reporting allowed for comparison of scenarios (session to session and run to run) to more quickly identify dataset differences and assess the quality of each solution. And of course, documentation and automation of business rules make it more efficient to train new planners.

Data Quality

Copy and paste data collection errors were eliminated by the direct linking to external data sources. The auditing of collected data for predetermined data quality issues was automated and formalizing the traceability of required manual adjustments resulted in more rapid issue resolution. There was improved identification of root causes of data issues with data quality checks and reporting visibility.

Solution Quality

There was a much more uniform understanding of business rules across each business through documentation and subsequent translation into mathematical programming constraints and objectives. Encoding the rules into the LP tools guaranteed a consistent application of business rules from plan to plan and from planner to planner that was never previously achieved. The optimization solver generates the best solution given the supplied data, as opposed to manual heuristics that previously risked stopping at a feasible solution. Furthermore, better data management support and faster solvers speed up the scenario setup-plan generation-solution analysis cycle to allow more complete consideration of alternate solution options.

Continuous Improvement

This transparent system has improved collaboration and coordination between various planning arms. Now different planning groups have better visibility and understanding of each others' capabilities and limitations. A standard tool with comparable rules across divisions has made it easier to recognize differences and similarities for continuous improvement. The business rules can be changed more quickly and more reliably to reflect changes in Intel's strategies, changes in Intel's markets, and improvement ideas from the planners using the tools on a regular basis. This has improved collaboration to improve the planning process. From a broader perspective, these

new transparent planning tools have started to increase collaboration between different functions in the company. For example, Sales and Marketing now knows the business rules employed in planning and can use this information to better meet the needs of customers.

FUTURE IMPROVEMENTS

Both the old and the new planning processes include the assumption that the forecast parameter values supplied to the planning tool (e.g., demand, capacity, yield, TPT) are all suitable to be used in planning. In fact, all of the parameters are measurements made on stochastic processes. The forecasts are constructed by looking at historical data and current improvement projects to estimate what will happen in the future. Some of the stochastic processes on the supply side are well understood. For example, it is possible to reproduce the distribution of TPTs of items passing through a factory by building a discrete event simulation that includes random machine breakdowns, the unavailability of equipment technicians during breaks, and so on. For the demand side, the situation is much different since the underlying stochastic processes are much more difficult to characterize. Our investigations of forecast error (forecasts vs. actuals) indicates that planning systems should consider the limits of our ability to make accurate forecasts of the expected supply that Intel can produce and the expected demand that customers will purchase. Future efforts to improve planning system design and operation given this variability will make Intel's planning systems more valuable in determining how to utilize capacity resources to effectively satisfy market demand.

CONCLUSION

Development of automated data systems and optimization-based tools has revolutionized Intel's supply-chain planning processes. Benefits have been realized in data and solution quality as well as planner productivity across all product divisions and all manufacturing organizations. Planning time has decreased dramatically, supply costs have been reduced, and demand satisfaction has improved. But perhaps the most important contribution of the efforts described here is the facilitation of continuous improvement. With better data, documented business rules, and fast planning and analysis tools, Intel planners have the time and facilities to define world-class performance.

ACKNOWLEDGMENTS

We acknowledge our manager, Karl Kempf, who has always encouraged and supported us to challenge the status quo and push for continuous improvement. We also acknowledge the numerous planners whose willingness to

acknowledge the numerous planners whose willingness to share the “secrets” of the current process and accept improvements allowed us to realize these successes.

REFERENCES

- [1] W. J. Hopp and M. L. Spearman, *Factory Physics: Foundations of Manufacturing Management*, New York: McGraw Hill, 1996, Ch. 16 (Aggregate and Workforce Planning).
- [2] S. Chopra and P. Meindl, *Supply Chain Management: Strategy, Planning, and Operation*, Part 2 (Planning Demand and Supply in a supply-demand network), Prentice-Hall, Upper Saddle River, NJ, 2001.

AUTHORS' BIOGRAPHIES

John W. Bean is a senior supply-chain engineer in the Decision Technologies department in Intel, Chandler, Arizona. He graduated in 1994 with a B.S. degree in Manufacturing Engineering from Brigham Young University, and in 1997 he graduated from the Leaders for Manufacturing (LFM) program at the Massachusetts Institute of Technology (MIT) with M.S. degrees in Material Science and Management. He has worked on WIP management, information systems development, and optimization solutions in FSM Industrial Engineering and Planning. His e-mail is john.w.bean at intel.com.

Amit Devpura is a senior supply-chain engineer in the Decision Technologies department in Chandler, Arizona. He graduated with a Bachelors degree from IIT Delhi, India in 1995. He received an M.S. degree in Mechanical Engineering and a Ph.D. degree in Industrial Engineering from Arizona State University (ASU) in 2000 and 2003, respectively. He started working at Intel in the Q&R organization in 2000 and has been working in Decision Technology since 2003, where his focus has been on implementing decision science tools in ATM planning. His e-mail is amit.devpura at intel.com.

Michael C. O'Brien is a senior supply-chain engineer in the Decision Technologies department in Chandler, Arizona. He graduated in 1991 with a B.S. degree in Systems Engineering and in 1996 with an M.S. degree in Industrial Engineering from the University of Arizona. He has worked on shop-floor scheduling projects, tactical and strategic planning systems, and has operated as an internal consultant providing simulation, optimization, and statistical support in a wide variety of business environments. His e-mail is michael.obrien at intel.com.

Shamin Shirodkar is a senior supply-chain engineer in the Decision Technologies department in Chandler, Arizona. He has a B.S. and M.S. degree in Industrial Engineering and is currently pursuing his Ph.D. in Applied Mathematics at Arizona State University. He has worked

on applying the Theory of Constraints to Ramping Semiconductor Fabrication facilities and on implementing a variety of decision science tools in the supply-chain arena both internal to Intel as well as collaboratively between Intel, its material suppliers, and manufacturing subcontractors. His e-mail is shamin.a.shirodkar at intel.com.

Copyright © Intel Corporation 2005. This publication was downloaded from <http://developer.intel.com/>.

Legal notices at <http://www.intel.com/sites/corporate/tradmarx.htm>