Robust Process Control at Cerestar's Refineries

Kumar Rajaram

Anderson Graduate School of Management, UCLA, Los Angeles, California 90095, USA

R. Jaikumar

Graduate School of Business, Harvard University, Cambridge, Massachusetts 02138, USA

Franz Behlau

Cerestar Benelux BV, Nijverheidsstraat 1, 4551 LA, Sas Van Gent, The Netherlands

Frans van Esch

Cerestar Benelux BV

Corrie Heynen

Cerestar Benelux BV

Ralph Kaiser

Cerestar Deutschland GmbH, Dusseldorfer Straße 191, D- 47809 Krefeld, Germany

Albert Kuttner

Cerestar Deutschland GmbH

Izaak van de Wege

Cerestar Deutschland GmbH

With annual sales of over $2 billion, Cerestar is Europe's leading manufacturer of made-to-order wheat- and corn-based starch products. Cerestar relies on refineries that are highly automated and require large fixed investments. Starting in 1993, we developed Robust Process Control (RPC) to increase average throughput and reduce throughput variation by combining engineering principles with OR/MS techniques. RPC includes a mathematical-programming model to reduce downtimes due to product switchovers, models for process optimization, and dynamic control models for process-flow synchronization. Cerestar implemented the resulting decision support system at eight refineries in six countries. It has increased average daily throughputs by 20 percent and reduced average throughput variation by 50 percent. Concomitantly, the refineries have reduced their consumption of supplies and utilities. In addition to over $11 million in annual benefits, RPC has had major strategic and organizational impact.

Large-scale process-based operations are prevalent in almost every major industry. For instance, most processes in the intermediate-food-processing, pharmaceutical, paper, and petrochemical industries fall into this category. They are constructed to reliably produce large volumes of well-established product, using conventional and tested technology.

Cerestar is Europe's leading manufacturer of made-to-order wheat- and corn-based starch products, such as glucose, sorbitol, dextrose, and gluten, with annual sales exceeding $2 billion. From 1927 to 1987, it was part of the Corn Products Refining Company (CPC), New York. In 1987, it was incorporated as a separate company under the Italian group Ferruzi, and in 1994, due to corporate restructuring, it became a company of the French group Eridania Beghin-Say. During this time, Cerestar has grown to become a major producer of these products, which are used extensively as components in the food-processing industries (for example, by breweries, confectioneries, and bakeries), the consumer-product industries (for example, cosmetics and toothpaste), and such other industries as paper, pharmaceuticals, textiles, and specialty chemicals.

To produce these products, Cerestar operates over 20 different types of industrial-scale processes in 16 plants located in nine countries. These can be broadly classified into physical processes, such as refining, separation, grinding and extracting, and chemical processes, such as hydrogenating and modifying starch products. Since building these processes requires major capital investments, it is crucial that they constantly produce high volumes of output at the correct quality level. To achieve this goal, these processes rely on high degrees of automation, operate continuously, and usually shut down only a few times a year for scheduled maintenance. As product switchovers result in long downtimes, products are produced in long campaigns.

Cerestar has a rich tradition of developing innovative approaches to process management. During the mid-80s, it implemented a framework for process improvement called Social Technical Systems (STSs). STSs have been widely recognized as benchmarks for process improvement, and they have been adapted for similar processes in many industries. Under these systems, automation at these processes was standardized and unified in a single centralized control room. This led directly to drastic reductions in work forces and increases in productivity. Existing operators were retrained to perform new tasks or reassigned to new processes.

Upside variation indicates a potential for increasing throughput.

In particular, adoption of STSs increased daily average throughputs to levels well above the capacities estimated in engineering specifications. However, there was still significant variation in day-to-day outputs at several processes. Downside variation is extremely expensive due to several factors. It requires larger quantities of finished product inventory and typically leads to greater per-unit costs for energy, supplies, and environmental degradation. In addition, such variation causes greater attrition of physical components and lower

customer-service levels, and it impedes the development of process knowledge as valuable operator time is spent on "fire fighting."

Upside variation indicates a potential for increasing throughput. Such increases are crucial in this industry because products are commodities with market-defined prices. Consequently, firms can increase their margins only by reducing the unit cost of production. Recent trends in this industry include consolidation, increased demand for products due to growing markets in east Europe and Asia, and tighter product specifications. In light of these trends, Cerestar needed to quickly increase both the capacity and the capabilities of its processes to maintain its dominance in the industry. Realizing these

Cerestar operates over 20 types of industrial-scale processes in 16 plants in nine countries.

goals through the traditional strategy of building new plants was no longer a viable option, because the costs of automation, information, and control, a significant proportion of total costs, had greatly increased. In addition, demand in these new markets was not yet stable, and large investments could be very risky. Finally, Cerestar might lose a large portion of market share during the time it would take to start up these new processes. These factors compelled Cerestar to explore methods of increasing the reliability and output of its existing process without making significant investments in new process technology.

To achieve these objectives, David Challenor, the executive vice president in charge of manufacturing, put together an internal team in September 1992. Members agreed that Cerestar's production planning and process control were excessively complex and lacked a scientific and systematic approach to problem solving. The process-automation systems provided data in huge volumes. However, these data were not being used efficiently to improve process productivity because there were only a few common standards, insufficient diffusion of ideas across plants, and low retention of knowledge and experience within the organization. The team members thought that these factors contributed to the tremendous variation in outputs across several processes.

When the team reported its findings in May 1993, Cerestar decided to collaborate with Jai Jaikumar, at the Harvard Business School, and Kumar Rajaram, then at the Wharton School. Their research focused on improving the productivity of large-scale industrial processes without significant investment in new process technology. They based their approach on combining engineering principles with OR/MS-based techniques. In June 1993, Cerestar formed a task force comprising the authors to improve the performance of these processes by applying these techniques. We focused on the refining processes because refined products accounted for a large part of total profits. We chose the glucose refinery at Sas Van Gent in The Netherlands as our test site because it was the flagship refinery of the company with the most sophisticated process and automation technology. If we could achieve improvements at

this refinery, our techniques could clearly result in improvements at all the other refineries.

The Operating Environment

Glucose is the generic term given to a wide variety of sugars produced from starch. The glucose-refining process at Cerestar converts starch slurry to glucose. The slurry first goes through an incubation stage, where it is treated with enzymes to convert it to glucose at a certain sugar or dextrose-equivalent (DE) level. The glucose is then purified by filtration, made colorless with decolorizers, made tasteless by polishers, and reduced in water content through evaporation. Individual refineries at Cerestar vary in the type of technology they use to achieve this conversion and in the number of parallel stages at each step.

These glucose products are used as intermediate products in various food-processing industries to produce such products as candy, beer, and soft drinks and also in other industries to produce such products as paper and pharmaceuticals. Customers specify their requirements by DE level. Rather than produce small runs for each customer and incur extensive downtimes due to switchovers, Cerestar developed many years ago a simple but then revolutionary idea of producing a few types of products across a range of DE levels. It then meets customer requirements by blending these glucoses, called *basic grades*. This enabled Cerestar to conduct long production campaigns without compromising its ability to accurately meet customer demand. Cerestar produces basic grades in a predefined sequence, stores them in tanks and draws them from the tanks and mixes them in blenders to meet demand. Since switching from one basic grade to another still causes significant downtimes, minimizing switchovers is crucial. In minimizing switchovers, Cerestar had to consider how much storage to provide, how to allocate storage across the basic grades, and when to conduct which campaign for basic grade production.

Control of the stages of the refinery is automated through centralized process-control systems. These systems are broad ranging and flexible, allowing the user to set and change many parameters at each stage. Broadly speaking, based upon their functionality, we can partition these parameters into two groups. The primary function of the parameters in the first group is to provide a stable operating environment for the process stage. These parameters, such as temperature and pressure, are set within a well-defined and narrow range determined by design and safety considerations. The process-control system uses regulators to automatically maintain these settings within this range. The second group of parameters specify the rate at which each stage should be operated. The process-control system uses control variables to set these rates. Operators determine the setting of these control variables to best match production requirements.

The major focus of engineering paradigms of control at any stage is to determine the partition of parameters into regulators and control variables and to determine the optimal setting for the regulators. To determine these, we developed scientific models based on the physics, chemistry, and mechanics of that stage to

explain its operation while assuming the controlled environment of the laboratory. Under such an assumption, setting the control variables to meet a production rate is straightforward. However, in practice, we may have to alter this scientific model to account for the interaction effects between the stages forming the process and for the operating environment at the plant. In this environment, stages typically lag in reacting to changes in the settings of control variables. Capacity imbalances may exist within a process because the components in the process may not be made to the same specifications, and some may require periodic downtimes because of the nature of the chemical processes involved. In addition, operators must act quickly to meet short-term production requirements.

Since September 1994, the automated system has run the refinery 95 percent of the time.

Because of these circumstances, operators have difficulty understanding the process well enough to choose and set control variables to achieve high levels of output consistently and to base their choices on scientific and systematic approaches to problem solving. Instead, they rely on subjective expertise, may be preoccupied while setting control variables, and pay insufficient attention to routine maintenance and to the manual portions of the operations. These practices may lead to significant variation in outputs. For instance, at the Sas refinery, although average daily throughput was 345 tons, well in excess of the 300 tons per day capacity estimated

from engineering specifications, this output varied between 445 and 90 tons per day.

Our approach was to blend engineering principles with OR/MS techniques in the operating environment of the plant. To do so, we developed a framework called Robust Process Control designed to increase average throughputs and reduce throughput variation without major investments in new process technology. Cerestar first implemented this framework in Sas Van Gent between 1993 and 1994 and subsequently implemented it at seven other refining processes in five countries between 1994 and 1997.

Robust Process Control

Robust Process Control (RPC) consists of four steps. We first simplify the process, then we stabilize the bottleneck, synchronize the other stages with this bottleneck, and standardize procedures to ensure that we maintain gains.

Step 1: Process simplification. To simplify the process, we reduce its operational complexity by investigating existing procedures and performing tests to reduce rework, recycles, and transients at individual stages of the process. Reducing these disruptions makes the process more predictable and increases the duration of the steady state. In these processes, it is difficult to switch from one basic grade to another because regulator and control-variable settings at each stage must be reconfigured in a coordinated way. Consequently, minimizing the number of switchovers diminishes operational complexity. To reduce the number of switchovers, we used optimization-based models (described in the appendix) to determine the
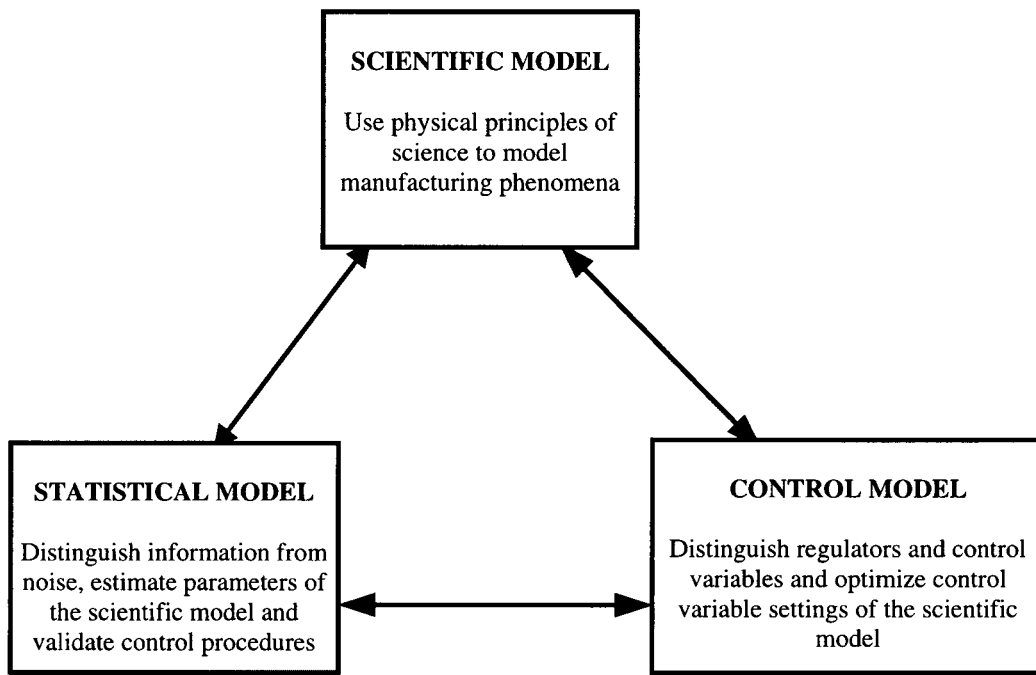
Figure 1: The process simplification step of Robust Process Control consists of a scientific model, a statistical model, and a control model. These models are jointly developed in the operating environment of the plant.

size of the storage tanks, the allocation of basic grades to tanks, and when to switch from one grade to another [Rajaram 1998]. A simplified process is a transparent process. It is now easier to detect the bottleneck, and optimization models based on simple first-order physical relationships are sufficient to stabilize it.

Step 2: Process stabilization. To stabilize the process, we develop an optimization model to stabilize the bottleneck, translate a production target to the required flows at the bottleneck, and calculate the setting of the control variable that best achieves this flow. The architecture of this optimization model consists of a scientific model, a statistical model, and a control model (Figure 1). The scientific model describes the operation of a stage based on a com-

prehensive examination of the operating mechanism at this stage. To develop this model, we focus on simplicity and first-order interactions and ensure that we can accurately measure the variables in the model. These variables need to be part of the existing control system and compatible with the safety tolerances of the process. To estimate the parameters for the scientific model, we develop a statistical model based on data gathered in independent samples at the correct physical points in the process. The control model is used to choose the control variable and its settings to best achieve the flows needed to meet the production target. This model is developed from the validated scientific model [Jaikumar and Rajaram 1996].

Step 3: Process synchronization. To syn-

chronize the stages in the process, we ensure that the requirements at the bottleneck are always met by stabilizing operations at the other stages and synchronizing the flows at each of these stages to those at the bottleneck. We do this by using a procedure that calculates the buffers and the steady state flow required in real time at each stage to compensate for regeneration downtimes (appendix).

Step 4: Process standardization. The first three steps of this framework ensure a simplified, stabilized process with a few control variables set to ensure that flow settings are synchronized to the bottleneck. To standardize the process, we develop a decision support system, which automates the computation required in the previous steps. In effect, during steady state the process completely runs itself like the autopilot in an aircraft. The role of the operator is to maintain a log of activities to record abnormalities, to understand why they occur, and to deal with extreme contingencies. Operators have more time to do maintenance, to detect faults in small nonautomated parts of the process, such as pumps and motors whose failure could be extremely disruptive, and to develop an understanding of the process, key to further technological innovation.

Implementation

We first implemented RPC at Cerestar's glucose refinery at Sas Van Gent. This refinery produces over 150 types of products, blended from six basic grades. The refinery at Sas is expected to operate continuously, except during the four planned maintenance shutdowns each year. Aver-

age throughput was around 345 tons per day, well above the capacity of 300 tons per day estimated from engineering specifications. However, day-to-day variations in this throughput were around 25 percent of this average. Since demand is quite stable, this variation is largely due to the production-planning and process-control techniques used at this refinery.

The refining process at this plant comprises 10 stages controlled sequentially and concurrently by three operators, each responsible for a fixed and prespecified group of stages (Figure 2). The first operator controls the process through to the first buffer tank, the second from the decolorization until the second buffer tank, and the third the remaining portion of this refinery. In the production-planning and process-control architecture prior to implementation of RPC (Figure 3), the production planner transformed customer demand to demand for basic grades and determined the duration of the production campaigns required by specifying production-planning targets. Each operator then translated this plan into flow requirements and chose the control variables and settings for his or her parts of the process. The architecture after implementation of our four-step approach (Figure 4) consists of production-planning, input, prescriptive, and output modules. Each module was developed during the different stages of RPC and integrated to form the automated system, which since September 1994 has actually run the refinery 95 percent of the time. The remainder of the time represents downtimes due to maintenance shutdowns or mechanical breakdowns. During these nontypical periods, operators
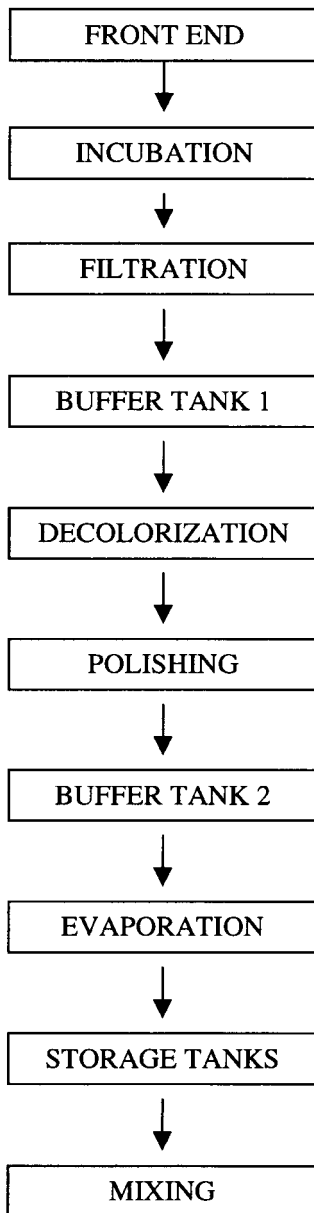
```
┌─────────────────────┐
│     FRONT END       │
└─────────────────────┘
          │
          ▼
┌─────────────────────┐
│     INCUBATION      │
└─────────────────────┘
          │
          ▼
┌─────────────────────┐
│     FILTRATION      │
└─────────────────────┘
          │
          ▼
┌─────────────────────┐
│   BUFFER TANK 1     │
└─────────────────────┘
          │
          ▼
┌─────────────────────┐
│   DECOLORIZATION    │
└─────────────────────┘
          │
          ▼
┌─────────────────────┐
│     POLISHING       │
└─────────────────────┘
          │
          ▼
┌─────────────────────┐
│   BUFFER TANK 2     │
└─────────────────────┘
          │
          ▼
┌─────────────────────┐
│    EVAPORATION      │
└─────────────────────┘
          │
          ▼
┌─────────────────────┐
│   STORAGE TANKS     │
└─────────────────────┘
          │
          ▼
┌─────────────────────┐
│      MIXING         │
└─────────────────────┘
```

Figure 2: The glucose-refining process at Sas Van Gent comprises 10 stages controlled sequentially and concurrently by three operators, each responsible for a fixed and prespecified group of stages. The first operator controls the process through to the first buffer tank, the second from the decolorization until the second buffer tank, and the third the remaining portion of this process.

control the process manually.

The first step we took in developing this system was to simplify the process by reducing rework at the decolorizers and polishers, eliminating recycles at the filters, and reducing switchovers between basic grades. Finally, we used the model described in the appendix to determine the total storage volume required at the end of the process and its allocation across basic grades. The portion of this model that dynamically determines when to switch to a different campaign is part of the production-planning module (Figure 4). This information is then passed on to the input module. By simplifying the process, we ensured that disruptions arising from switches of products, filters, and ion-exchange regeneration at the decolorizer and polishers are predictable and minimized in frequency and duration. With increased stability, we identified the second-stage evaporator as the bottleneck.

To stabilize the process, we increased the throughputs and reduced throughput variation in this evaporator based upon the optimization model described by Jaikumar and Rajaram [1996]. Using this model, we can translate the production target into flow required at the evaporator and determine the setting of the best control variable to achieve these flows. We stabilized the decolorizers, polishers, and filters using similar models and synchronized their flows to the evaporator using the procedure outlined in the appendix. Jaikumar and Rajaram [1997] discuss these models and the estimation of the timing and duration of regeneration downtimes in detail. These actions reduced the number of control variables from 44 to four.
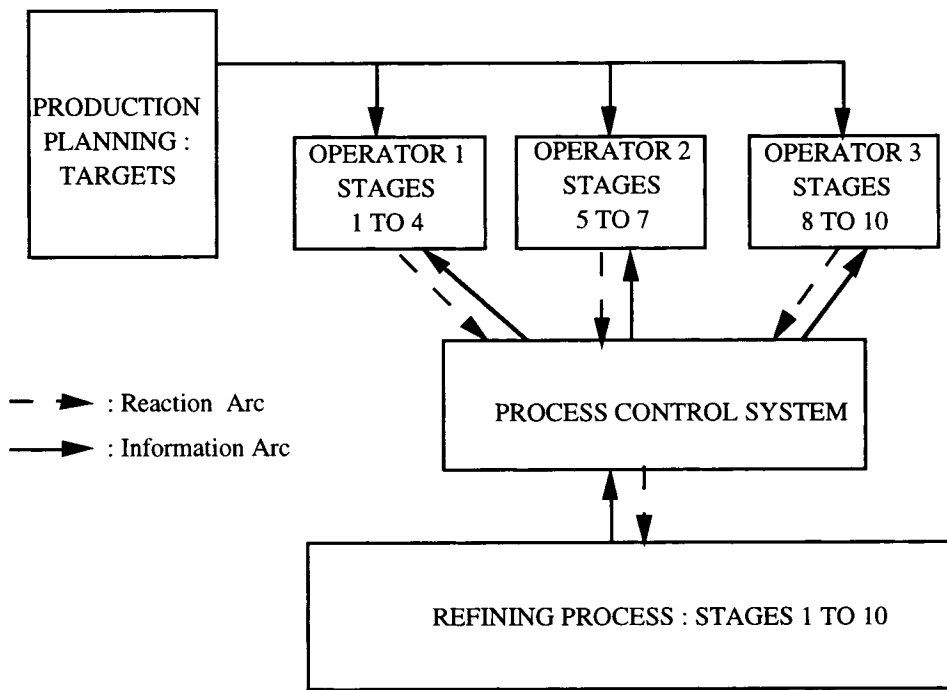
Figure 3: In the production-planning and process-control architecture prior to implementation of Robust Process Control at the Sas Van Gent glucose refinery, the production planner transformed customer demand to demand for basic grades and determined the duration of the production campaigns required by specifying production-planning targets. Each operator then translated these targets into flow requirements and chose the control variables and settings for his or her parts of the process.

The control variables eliminated were changed to regulators and set at their optimum levels.

This reduction in control variables had a profound impact on the operational complexity at this refinery. It now employs fixed standards during all shifts and no longer depends so heavily on the subjective expertise of the operators. The operators can now monitor the remaining control variables more effectively and better understand the cause and effect in the operation of each stage. This knowledge and the reduction of control variables means fewer changes during the operation of this process, which drastically reduces the variation at each stage and at the output.

To standardize and preserve these changes, we incorporated all the calculations required by these models in the prescriptive module of the architecture (Figure 4). The input module collects the data to drive these models from the process-control system based upon a half-hourly sampling frequency. The output module passes the flow requirements and the control settings on to the process-control system. We developed these modules using the C^{++} programming language and linked them to Extend [1992], a software that offered a user-friendly interface and

the capability to perform off-line simulations. This provided us with an opportunity to simulate and test this system before implementation.

Since implementation, downtimes due to unplanned disruptions at this refinery have been reduced by over 70 percent. One can better appreciate this change by observing the flows at the filters, polishers, and evaporators before and after implementation (Figures 5 to 7) during a randomly sampled period of four days. Average daily throughput at the refinery has increased by over 18 percent from a base level of 345 tons per day and day-to-day throughput variation has been reduced by around 60 percent. To understand the impact of these changes, compare the daily

throughput during a two-month period before Cerestar initiated this project with the daily throughput during a two-month period after implementation (Figure 8).

A major challenge we faced during implementation was to understand and redefine the role of the operator in this new system. It was crucial that the operators believed in the system and did not feel threatened by it. Consequently, during each stage of development, we actively sought and incorporated their suggestions. We elected to provide them with the flexibility to overrule the recommendations of the model if they prepared a detailed report explaining the reasons for their actions. While the operators were initially skeptical about the ability of this system,
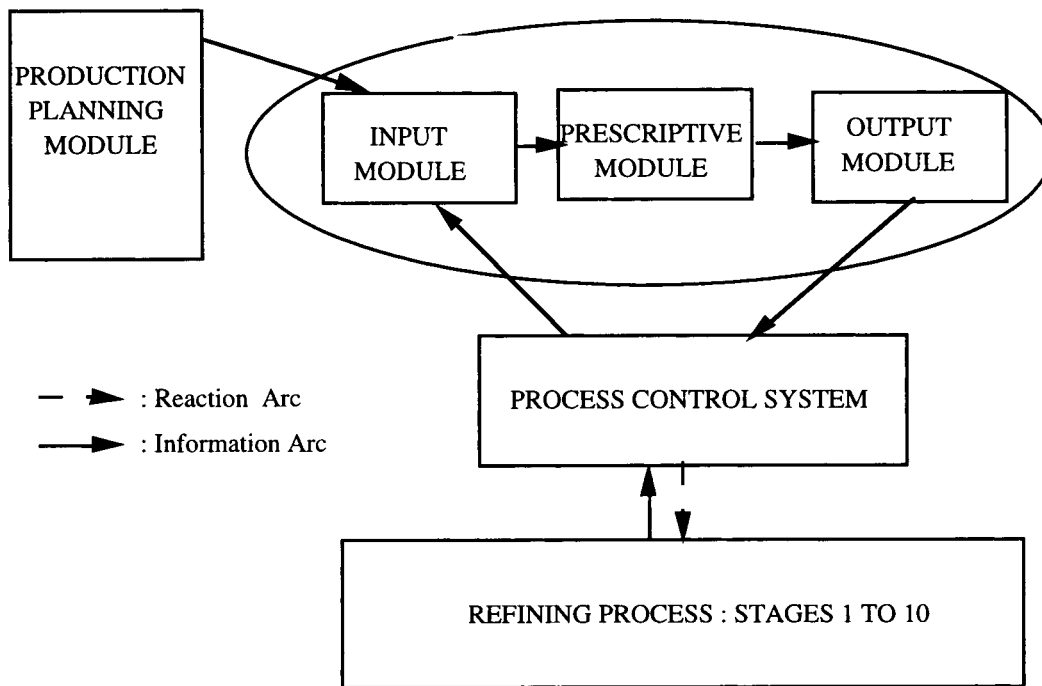

Figure 4: The production-planning and process-control architecture after implementation of Robust Process Control at the Sas Van Gent glucose refinery consists of production-planning, input, prescriptive, and output modules, which are integrated to form an automated system. Since September 1994, this system has actually run the refinery 95 percent of the time.
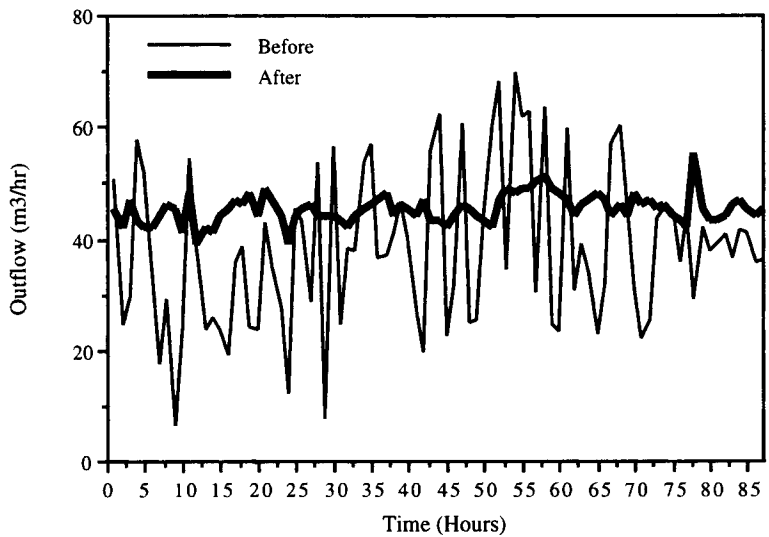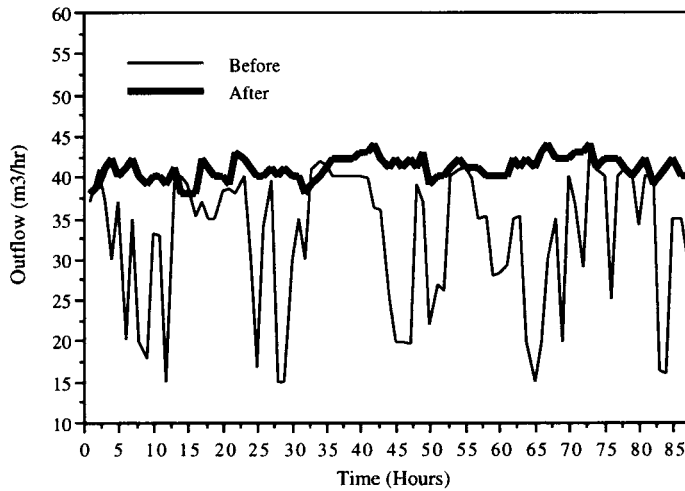
Figure 5: Outflow variation at the filtration stage of the glucose-refining process at Sas Van Gent has been significantly reduced after implementation of Robust Process Control.

they embraced this approach once it was proven in the plant. They were extremely pleased that they no longer needed to spend their time determining settings to the control variables and fire fighting. The current role of the operator is to improve this system, to develop better understanding of the process, to spend more time performing routine maintenance, and to monitor the smaller nonautomated portions of the process. This strategy has had handsome dividends. Disruptions due to mechanical breakdowns have been reduced by more than 50 percent. Better process understanding has led to further technological innovations, which has lowered the consumption of supplies and utilities at several process stages.

In September 1994, the system for RPC became operational at the Sas refinery. Since then we have been implementing this concept in Cerestar's refineries in Spain, France, the United Kingdom, and Italy, and at its three sorbitol refineries in

Germany (Figure 9). By October 1996, Cerestar was using our system to run all these processes. In carrying our approach to the seven refineries, we observed that the operations people at these processes had similar engineering backgrounds to those of Sas and, like them, had a lot of know-how about the process. Thus, we were not surprised to find their attitudes toward the new system comparable to those encountered at Sas. Here again, operators used a lot of unproved theories to defend certain concepts. We came across older and skeptical process pioneers who argued these theories based upon seniority instead of proven data from the plant. To convince these people that our approach was valid, we often resorted to describing examples and showing the specific results from other projects.

In addition, we found it crucial to understand the organizational differences in these countries to understand what we could accomplish and when. In particular,

the cultural differences at these plants played an important role in determining the character of the implementation teams. Broadly speaking, teams ranged from self-standing groups of people completely committed to this approach to groups of people who wanted to look for compromises, who thus required more guidelines. The biggest hurdle in implementation turned out to be in configuring the software to account for differences in automation capacity at each plant. We overcame this problem through the support of the board of directors at the corporate headquarters in Brussels and of the manufacturing directors responsible for each plant who redeployed software and automation engineers from other projects to work on

Figure 6: Outflow variation at the decolorization stage of the glucose-refining process at Sas Van Gent has been significantly reduced after implementation of Robust Process Control.
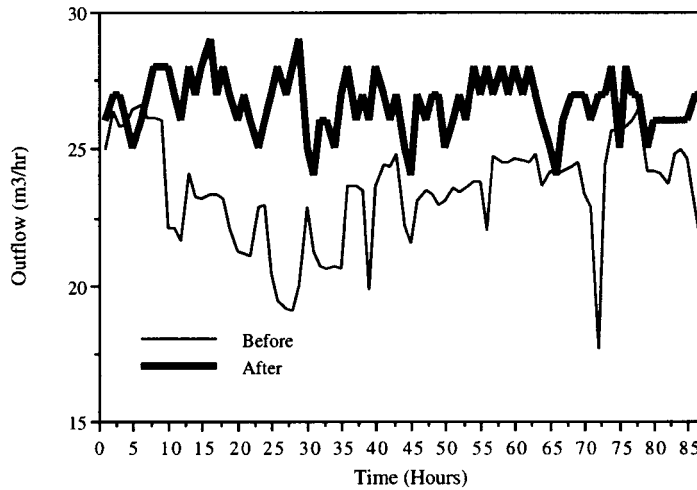
Figure 7: Outflow variation at the evaporation stage of the glucose-refining process at Sas Van Gent has been significantly reduced after implementation of Robust Process Control.
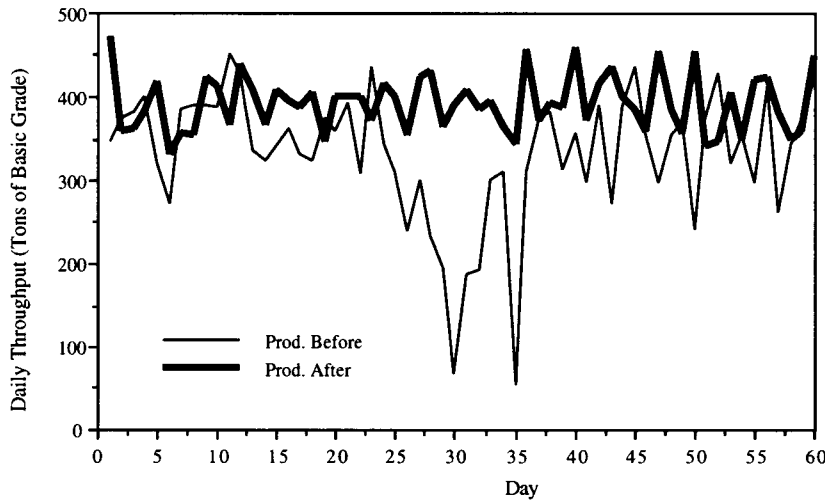
Figure 8: Average daily throughput has increased while throughput variation has decreased after implementation of Robust Process Control at the Sas Van Gent glucose refinery.
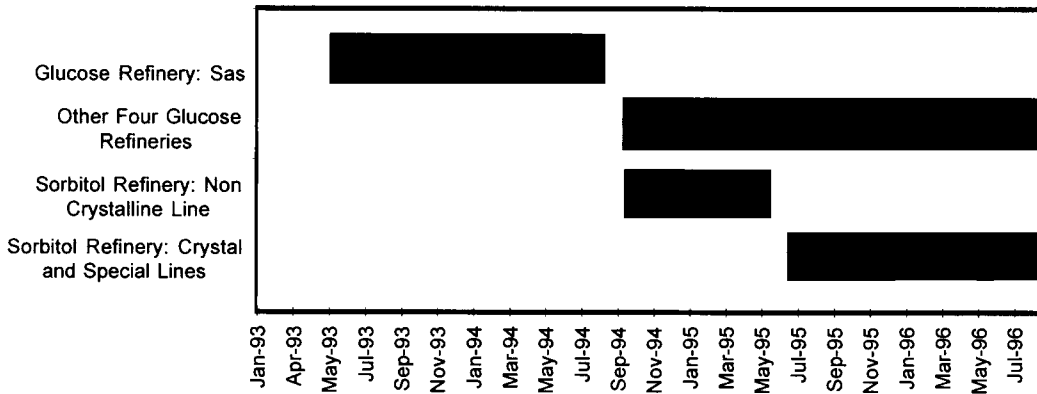

Figure 9: The implementation schedule of Robust Process Control at eight Cerestar refineries from June 1993 to September 1996.

this project. In the final analysis, this level of support was one of the most important factors in overcoming the not-invented-here syndrome; it was the key to achieving the results.

Results

Since October 1996, RPC systems have run eight refining processes at Cerestar at least 90 percent of the time. Basic grade switches have been reduced by an average of around 40 percent, resulting in an average reduction of downtime by over 15 per-

cent. The number of control-variable options for the operators has been reduced by an order of magnitude. Disruptions arising from incorrect specification of control-variable settings (leading to congestion and stops at process stages) or from mechanical breakdowns have decreased by more than 60 percent across all these processes. On average, daily throughput has increased by more than 20 percent and throughput variation has diminished by around 50 percent. Results at

the individual processes are summarized in Tables 1 and 2.

Benefits

To calculate the economic gains from the increase in throughputs, we must first convert percentage gains to actual tonnage. This calculation suggests that there is an increase of 500 tons per day across all eight refineries. This additional tonnage can be sold, as capacity at these plants is usually lower than demand. In effect, RPC has achieved these gains with no new fixed investments. If Cerestar had had to build a new refinery to realize these gains, the cost of constructing a 500-ton-per-day refinery would have been at least $60 million. If we use the standard depreciation rates and the manpower costs employed at Cerestar, the annual costs of producing these 500 tons would be at least $8 million.

The reduction of variation at these refineries has had several additional benefits. It

Site (Refinery Type)	Basic Grade Switches		Control Variables		Percentage Reduction in Disruptions
	Before	After	Before	After	
Holland (glucose)	8	6	44	4	70
UK (glucose)	14	7	65	4	60
Spain (glucose)	12	6	70	5	85
Italy (glucose)	9	7	45	4	90
France (glucose)	11	6	65	3	65
Germany (noncrystalline sorbitol)	14	8	15	2	70
Germany (crystalline sorbitol)	11	7	17	2	60
Germany (special sorbitol)	12	9	13	2	75

Table 1: The number of basic grade switches and control variables has been notably reduced after implementation of Robust Process Control at eight of Cerestar's refineries. This has drastically reduced operational complexity and resulted in dramatic reduction in disruptions.

Site (Refinery Type)	Percentage Increase in Average Daily Throughput	Percentage Reduction in Coefficient of Variation of Daily Throughput
Holland (glucose)	18	60
UK (glucose)	10	50
Spain (glucose)	35	75
Italy (glucose)	50	85
France (glucose)	15	60
Germany (noncrystalline sorbitol)	15	80
Germany (crystalline sorbitol)	19	55
Germany (special sorbitol)	19	65

Table 2: After Robust Process Control was implemented at eight of Cerestar's refineries, average daily throughput has increased significantly, while the coefficient of variation at these eight refineries has been notably reduced.

has reduced the consumption of supplies, such as enzymes, reagents, catalysts, and other chemicals used in operating these refineries. In addition, savings in such utilities as energy and water have been estimated to be around $3.5 million annually. Finished goods inventory has been reduced by over 30 percent. Service levels have been increased, and customers have rewarded Cerestar with larger contracts over extended periods. Cerestar is now applying the same concepts in other processes. Implementation at a wheat-grinding process has already increased yields by over five percent by reducing consumption of raw materials by around 15 percent. This process has also reduced its consumption of water and energy by over 15 percent. These improvements have been valued at over $2.5 million annually. Promising results are expected from projects started at four starch-modification processes and five corn-grinding processes.

The reduction of variation has reduced the consumption of supplies.

The strategic impact of RPC on Cerestar has been substantial. RPC has provided Cerestar with the ability to produce specialty products at the cost of commodities. It has improved the tolerances of basic commodity grades to such an extent that Cerestar can now make specialty products effectively by using existing blending procedures. In addition, because RPC can run existing processes with higher capacity and better precision, Cerestar has been able to buy plants and run their processes

more efficiently, a better alternative than constructing new, more expensive processes. This strategy was strongly affirmed during January 1996. Cerestar purchased American Maize Products, a company that runs a network of starch-processing plants. It expects to run the existing processes at these plants more profitably using this approach. Currently, Cerestar is implementing RPC in several of the processes at this company and has already obtained remarkable results. Finally, because RPC minimizes automation, control, and information technology, Cerestar has undertaken major capacity expansions of current processes in a cost-efficient manner. For example, Cerestar is planning to expand the capacity of the wheat-grinding process at Sas by around 250 percent by investing $150 million; building a new process to achieve comparable capacity gains would cost around $250 million.

The organizational impact of RPC has been tremendous. In his annual presentation to the board in 1996, David Challenor (the international manufacturing director) said, "Robust process control is transforming us into a learning organization. Every process problem is now viewed as an opportunity for learning and process improvement. We do not just gather data; we convert these data to valuable information for process analysis. Decisions are based on this information instead of allowing opinions or anecdotal evidence to dictate future actions. Standardization of operational procedures is now a priority: proven control concepts are implemented and are not subject to personal interpretation. Now, in all our projects, we strive toward simplicity. We would rather be near

optimal and stable during a long period than try to be optimal always and achieve optimality for extremely short periods. This said, I must emphasize that we are never satisfied with the status quo and are always striving for improvement without adding complexity. In this effort, we are being aided by all our operators. This system has transformed our operators from fire fighters to process innovators" [Challenor 1996, p. 5].

The success of this project has motivated manufacturing to look more closely at other processes and other areas, including production planning, product costing, and process design. This has fostered its close cooperation with several other areas in the organization, including marketing, finance, and engineering. As Challenor notes, "In addition to the obvious economic benefits and the impact on our thinking, this work has improved the spirit of team work and communication in our multinational organization. We are extremely optimistic about the future of this project, not only due to its promise in other types of processes, but to its potential to help us focus on the managerial and strategic decisions required to guide us through the next millennium" [Challenor 1996, p. 5].

In summary, RPC has had a major economic, strategic, and organizational impact at this company. Cerestar expects to maintain the gains we described and to increase them continuously several years into the future.

Dedication

This paper is dedicated to the memory of R. Jaikumar, who passed away while mountain climbing in Ecuador on Febru-ary 10, 1998. Words cannot express our sorrow and deep appreciation of Jai's contribution to this work.

APPENDIX

Models for Optimizing Product Switches

To minimize the number of product switches, we calculate the total end-process buffers (tanks) required and determine how to allocate products (basic grades) to these tanks to minimize basic grade switches based on a long-run average of demand. We also develop a procedure to correct for deviations from this av-

erage and determine when to switch and to which grade to switch while running the process in real time. To model the problem of choosing the total number of tanks and the allocation of basic grades to tanks, we define the following variables:

$I = \{1,\ldots m\}$: Index set of basic grades,
$J = \{1,\ldots n\}$: Index set of tanks,
D_i: Long-term average of demand for basic grade i per time period,
V_j: Size of tank j in volume units,
V: Total available volume,
C_b: Cost per unit volume (including space, installation, tank and basic grade holding costs),
S_i: Switchover cost for product I.

The volume selection problem (VSP) is represented as follows:

$$(\text{VSP}) \ Z \ = \ \text{Min} \ S_i Z(V) \ + \ C_b V$$
$$V \geq 0.$$

In this problem, we trade off the cost of volume with the costs due to downtimes because of basic grade switches. The number of switches $Z(V)$ is the solution to the volume allocation problem (VAP) defined as follows:

$$(\text{VAP}) \ Z(V) \ = \ \text{Min} \ \sum_{i=1}^{M} \frac{D_i}{\sum_{j=1}^{m} \alpha_{ij} V_j}$$
$$\sum_{i=1}^{m} \sum_{j=1}^{n} \alpha_{ij} V_j \leq V$$
$$\alpha_{ij} \in \{0,1\}, \ \forall i,j.$$

During each production campaign, we would ideally produce enough of a basic grade to fill up the allocated tanks before initiating a switch. This would ensure that product switches are minimized. In our application, all possible switches between basic grades are feasible, and the switching times between these grades are identical. Violation of this assumption would require us to make significant modifications to include the effects of sequencing in this

model [Rajaram 1998].

It is easy to solve VSP once we solve the subproblem VAP. However, this problem is highly intractable due to the integer variables that are nonlinear in the objective. Consequently, we elected to develop a heuristic to address this problem. This heuristic consists of two phases. In the first phase, we solve the following continuous version of this problem (CVAP). Note that this provides a lower bound on VAP:

$$(\text{CVAP}) \ Z^{ub}(V) \ = \ \text{Min} \ \sum_{i=1}^{m} \frac{D_i}{V_i}$$
$$\sum_{i=1}^{m} V_i \leq V$$
$$V_i \geq 0.$$

This problem is easily solved by setting the tank allocation to each basic grade as

$$V_i \ = \ \frac{V}{1 \ + \ \sum_{\substack{p=1 \\ P \neq i}}^{n} \sqrt{\dfrac{D_p}{D_i}}}.$$

However, the continuous solution may be infeasible due to tank-batch-size constraints enforced by V_j. To derive a feasible solution, we rank order set I in increasing order of demand to form set I'. We construct two sets A and B, which form a partition on I', and either $|A| = |B|$ or $|A| + 1 = |B|$. The basic grades in A have lower demand and are therefore more sensitive to demand variation. Consequently, we provide more safety stock for these grades by rounding up the continuous solution to the next feasible solution. Conversely, we round down the continuous solution for the basic grades in B. This approach performed remarkably well for the parameters defined by these nine processes. In all these cases, the solution provided by this heuristic was within two percent of the upper bound provided by the continuous approximation. As Rajaram

[1996] discussed, this method performs favorably with other randomly generated data sets.

Once we have determined the volume and allocation, we correct for deviations from the long-run average demand and determine when to switch grades and to which grade to switch on a real-time basis using the following method. To develop a precise definition of this procedure, we define the following:

$V_{i,\max}$ (respectively, $V_{i,\min}$): The maximum (respectively, minimum) permissible volume for the ith basic grade, defined based upon the sizes of the allocated tanks,
$V_{i,a}$: The actual total volume of the basic grade at these tanks,
$D_{i,\max}$ (respectively, $D_{i,\min}$): The maximum (respectively minimum) daily demand for the ith basic grade, estimated from the short-term planning horizon,
λ_i: The daily production rate for the ith basic grade,
$t_{i,st}$: Time required to start up the ith basic grade including switchover times (in hours) (this is independent of the current grade under production),
$t_{i,sh}$: Time required to shut down the ith basic grade under production.

For the ith basic grade, we develop the following disjunctive constraints:

$$V_{i,a} + \frac{(\lambda_i - D_{j,\min})t_{i,sh}}{24} \leq V_{i,\max} . \tag{1}$$

$$V_{i,a} + \frac{D_{i,\max}t_{i,st}}{24} \geq V_{i,\min}. \tag{2}$$

Constraint (1) enforces the condition that while this grade is being produced, actual volumes at the tank and the maximum expected buildup during shutdown should always be lower than its maximum permissible volume. Conversely, while another grade is being produced, the actual volume of this grade and the maximum expected volume depletion during its start-up should always be greater than its

minimum permissible volume. To make this procedure operational, we would keep producing the ith basic grade until Constraint (1) is violated for this grade or Constraint (2) is violated for any other grade, whichever occurs earlier. At that instant, we would switch to that grade for which Constraint (2) is violated first. It is important to recognize that we minimize the number of basic grade switches and maximize the duration of a production campaign by initiating a switch only when these boundary conditions are violated.

Procedures for Flow Synchronization

We consider an n-stage sequential process. In this process, we consider the ith stage, which operates for a known and constant duration (steady state) and is then periodically regenerated for a fixed period (transient state). We construct a buffer in front of this stage to ensure sufficient storage to keep the bottleneck operating during the transient state. We also determine the flow on a real-time basis at this stage to keep the bottleneck operational. We term this flow the *synchronized flow.* To determine the size of this buffer and the synchronized flow, we define the following variables:

$I = \{1,. . .n\}$: Index the set of stages,
$F_{i,s}$: Flow at the ith stage during steady state,
$F_{i,t}$: Flow at the ith stage during the transient state,
$t_{i,s}$: Duration of steady state at stage i,
$t_{i,t}$: Duration of the transient state at stage i.

We define $F_{i,e}$, the effective flow at the ith stage, as follows:

$$F_{i,e} = \frac{F_{i,s}t_{i,s} + F_{i,t}t_{i,t}}{t_{i,s} + t_{i,t}}.$$

The bottleneck is the stage with the lowest effective flow. Without loss of generality, let us assume that this is the kth stage. To ensure that the ith stage always meets the flow requirements at the bottleneck, we would set $F_{i,s} = F_{k,e} + (F_{k,e} - F_{i,t})t_{i,t}/$

$t_{i,s}$. This equation suggests that during steady state we would operate this stage at that flow which meets the effective flow at the bottleneck (i.e., $F_{i,e}$) and also builds up sufficient volume to account for the shortfall during transience (i.e., $(F_{k,e} - F_{i,t})t_{i,t}/t_{i,s}$). This buildup should be the size of the in-process buffer after this stage.

At any instant, we monitor the following variables from the process:

$V_{i,a}$: The actual volume at this buffer,
F_k: The required flow at the bottleneck,
t_i: Time remaining before the next transient at the ith stage.

We determine F_i, the synchronized flow at the ith stage, as $F_i = F_k + (F_k t_{i,t} - V_{i,a})/t_i$. This equation implies that synchronized flow is equal to the sum of the required bottleneck flow and the buildup required in excess of available volume during steady state used to compensate for the impending transience.

References

Challenor, David 1996, "Cerestar 2000: Strategy and actions for competitive success in the new millennium," presentation to the board of directors, Cerestar Headquarters, Brussels, Belgium.

Extend 1992, "Performance modeling for decision support," Imagine That, Inc., San Jose, California.

Jaikumar, R. and Rajaram, K. 1996, "Incorporating operator—process interactions in process control: A framework and an application to glucose refining," Working paper #96-070, Harvard Business School.

Jaikumar, R. and Rajaram, K. 1997, "A decision support system for operationally intelligent control," Working paper, Department of Operations and Technology Management, Anderson School, UCLA.

Rajaram, K. 1998, "Lot sizing in large scale industrial processes: Models, applications and analysis," Working paper, Department of Operations and Technology Management, Anderson School, UCLA.