Improving Car Body Production at PSA Peugeot Citroën

Alain Patchong • Thierry Lemoine • Gilles Kern
PSA Peugeot Citroën, Route de Gisy, 78943 Vélizy, France
alain.patchong@mpsa.com • thierry.lemoine@mpsa.com • gilles.kern@mpsa.com

---

In 1998, following a change in top management, the new CEO of PSA Peugeot Citroën decided to adopt a triple-axis strategy of growth, innovation, and profitability and set an ambitious target for each. To meet these objectives, PSA decided to focus on the car-body shops, which were the bottlenecks of its plants. An R&D team conducted a project to support car-body production for PSA Peugeot Citroën. PSA manufactures over 75 percent of its cars on lines designed and continually improved with the team's new analytic operations research tools. These OR tools, which combine simulation and Markov-chain models of series-parallel systems, have improved throughput with minimal capital investment and no compromise in quality—contributing US $130 million to the bottom line in 2001 alone. The impact of this project went beyond the boundaries of PSA as its suppliers acquired the tools without being requested to do so.

(*Industries: automotive. Production/scheduling: applications*.)

The Body Shop, the Paint Shop, and the Final Assembly Shop

PSA Peugeot Citroën, established through a merger of the two prominent French carmakers Peugeot and Citroën in 1976, is the sixth-largest carmaker in the world and the second-largest in Europe (PSA stands for Peugeot Société Anonyme). PSA is on a strong growth trajectory; it is, in fact, the fastest-growing car company in the world. Production for 2004 is projected to be 3.5 million cars, up from 3.1 million in 2001 and 2.3 million in 1998. PSA will launch 25 new car models between 2001 and 2004. The financials are also very robust, with net profits of $2.33 billion in 2001 on sales of $45.49 billion. PSA's sustained profitability and growth are strongly supported by its 16,000-person research and development organization.

A typical PSA car plant contains three shops (Figure 1). The body shop has the most complex manufacturing flow and is located at the beginning of the assembly process. Welders in the body shop assemble the raw materials, which are stamped parts, creating the body in white. After leaving the body shop, the body in white goes to the paint shop. The last step is the final assembly shop, where the painted body is merged with the remaining parts to make a marketable product. These parts range in importance from vital (for example, the engine system) to decorative (for example, the hubcaps). The car that leaves the final assembly shop is shipped to retailers.

PSA Strategy and Business Needs

In 1998, following a change in top management, the new CEO of PSA Peugeot Citroën decided to adopt a triple-axis strategy of growth, innovation, and profitability. His goals for these axes were (1) to produce 3.5 million cars in 2004, (2) to introduce 25 new models between 2001 and 2004, and (3) to increase profits on sales in 2002 to five percent, up from the target of 4.6 percent in 2001.

PSA management decided to pursue these goals through the so-called common-platform policy. A common-platform is a base used in several vehicles. Two vehicles based on the same platform share the same subframe, engines and gearboxes, front and rear

Interfaces, © 2003 INFORMS
Vol. 33, No. 1, January–February 2003, pp. 36–49

0092-2102/03/3301/0036$05.00
1526-551X electronic ISSN

Figure 1: The body shop is at the upstream end of the work flow in a typical car plant.

suspension, and a number of other parts except style and personality features. The goal of the common platform policy is to reach a 60 percent level of common parts among all cars assembled on the same platform, regardless of their brand—Peugeot or Citroën (Figure 2). This policy, which would reduce investment by relying on a flexible process at low cost, would affect the manufacturing system.

PSA encountered two problems in implementing its common-platform policy in car-body shops: (1) Its single-flow architecture, meant that every car model went through the same sequence of stations, limiting PSA's ability to handle diverse models. (2) Its production-line design was based on intuition and tradition, with some ingrained beliefs and practices causing inefficiency. So, the process of redesigning the production line for new car models was lengthy and erratic. As a result, PSA suffered from painfully slow ramp-ups and production shortfalls.

We needed a new architecture for the car-body production line that could handle model diversity and new car launches easily and quickly and an accurate method of sustaining quick innovation without overinvestment.

The Car-Body Assembly

The term *car body* designates the naked steel shell of the car, including the openings (doors, hood, and trunk door). It does not include the engine system or decorative items. The assembly process begins in the off-flow area where workers preassemble raw materials (mainly stamped steel parts) using simple machines and put them in containers. Carriers move these containers to the assembly area and place them along the assembly lines. Robots load and weld the parts stored in containers on assembly lines. At the start, robots weld preassembled parts on specific lines to form separately the front and the rear parts of the understructure, respectively called the front unit and the rear unit. Then, robots weld the two parts together with an upper part (called the *upper front*) to form the car's understructure, which is welded to the body's sides and the roof to produce the car body less the openings. The car body less the openings goes to the final assembly line, where workers perform final finishing operations and add the opening parts, such as the exhaust, hood, and doors. The resulting car body then goes to the paint shop.

The Original Car-Body Shop

Before we implemented the new assembly process, the main disadvantage of the car-body manufacturing system was that it consisted of only one flow (Figure 3). Because all products manufactured in the system had the same route, the shop operated under rigid constraints:

—To handle the diverse car models made at a plant, PSA needed a complex and rigorous control policy. It resequenced models whenever the parts schedule changed. When the control system was down, the shop ceased operations, reducing its efficiency. These problems damaged PSA's growth and profitability, two of the three axes of PSA's strategy.

—PSA had to redesign and reconstruct all the production lines in a shop when it launched a new product because all its car models shared manufacturing systems. This problem hindered innovation, the third axis of PSA's strategy.

To improve the shops' efficiency, we decided to simplify and minimize the extent of PSA's plant control policy. We redesigned the shops' structure to make them more agile in handling new products.

A New Standard for Shop Structure

To implement PSA's triple-axis strategy we needed to transform the manufacturing facilities to produce as many different models as possible using common

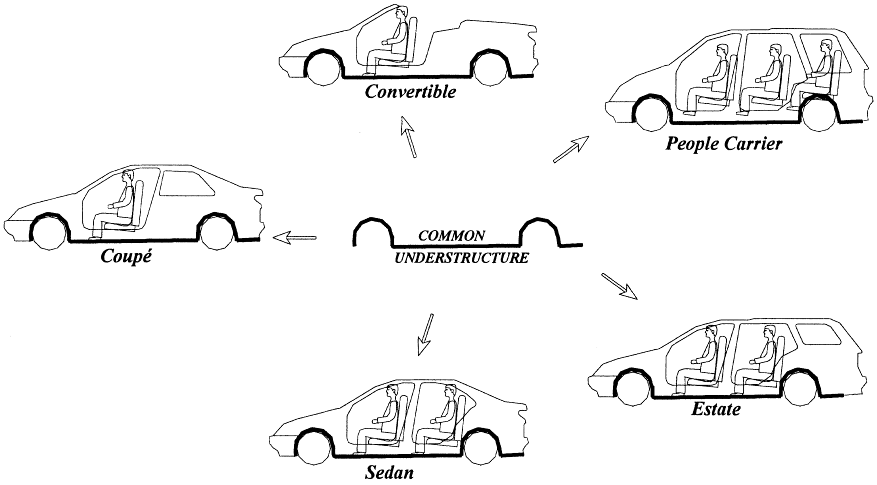

Figure 2: PSA's new strategy emphasizes using a common understructure for multiple car-body models, produced in the shop regardless of brand. The objective is to merge parts wherever possible without altering the identity or personality of the brands. For all the models, namely sedans, coupés, convertibles, estates, and people carriers, each brand brings a distinctive personality.

parts. PSA decided to reorganize all 12 original Peugeot and Citroën production sites by platform, regardless of brand. It dedicated each plant to a single platform—small, medium, or large. Together the plants cover all PSA's offerings.

To launch 25 new models in four years, PSA needed plants that could be reconfigured quickly and easily (Figure 4). We developed a new standard shop architecture to achieve this flexibility by treating the two kinds of welding stations on the line differently—the frame-welding stations and the finish-welding stations. Frame welding is directly related to the geometry of a car and, in a way, defines the car's shape. The framing tools used for this type of welding are difficult to modify for a new model. Finish welding strengthens the car's structure and, unlike frame welding, is easy to move or modify simply by reprogramming robot trajectories. Our rationale was to put as many framing stations as possible on specific lines and to put the finishing stations only on common lines (Figure 4). By doing so, we achieved two key objectives: (1) quick change because we needed to build only specific additional lines, and we were able to do so without disrupting the production, and (2) greater model diversity, because the common lines' robots just needed to load the corresponding programs to work on different types of parts.

After redesigning the production line, we also needed tools the staff could use to accurately and quickly design the manufacturing systems for new cars.

Designing Manufacturing Shops

A well-designed shop produces cars the marketing department requests. The costs include the money to build the shop and the money to run it. Designers of manufacturing systems must determinate (1) the parameters that affect shop efficiency: buffer sizes and locations, assembly lines' efficiency, and cycle time; and (2) a control policy that allows the shop to deliver the right products at the right time without degrading its efficiency.

Solution Approach

To realize its ambitious forecasts, PSA needed tools R and D personnel could use to design shops accurately and quickly. The main tools we knew about for evaluating performance were simulation and analytic methods (Dallery and Gershwin 1992) (Table 1).

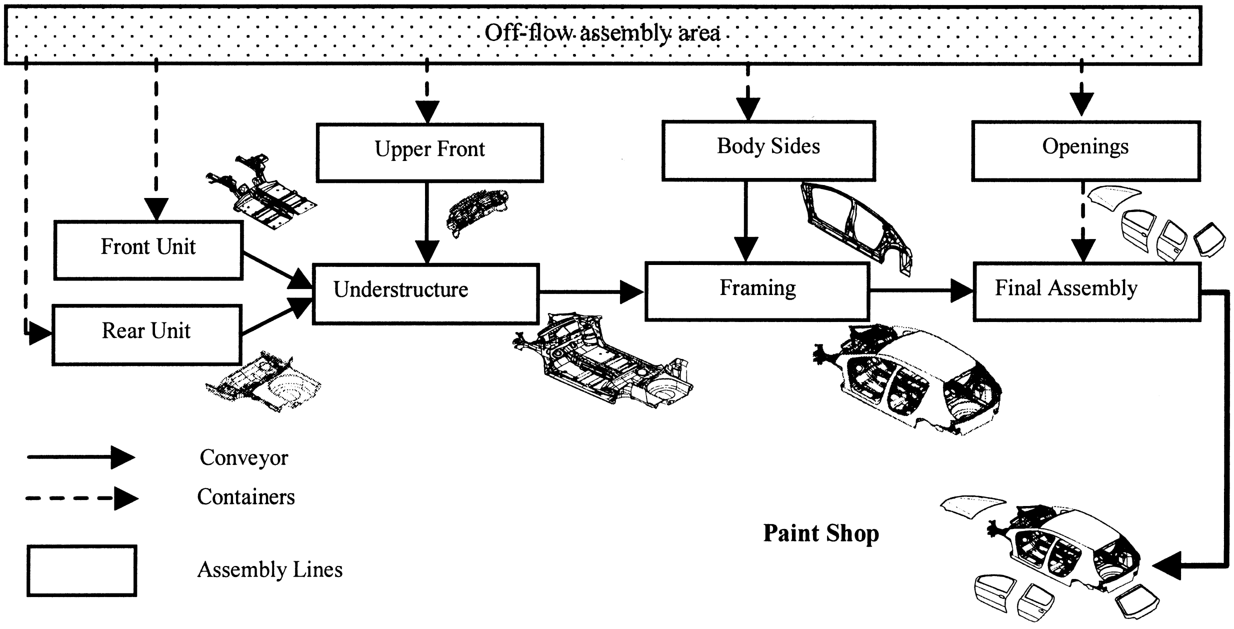

Figure 3: The car body is only the naked steel shell of the car, including openings. It does not include the engine system or decorative items. In the typical workflow in the original PSA body shop, workers preassembled raw materials in the off-flow assembly area using simple machines and routed the resulting part in containers (dashed line arrow) to the assembly area. There, robots assembled the rear unit, the front unit, and the upper front of the car separately on specific lines and sent them by conveyors (solid line arrow) to the understructure line, where other robots assembled them and sent the resulting part to the framing line, also by conveyors. On the framing line, robots assembled the understructure, the body's sides, and the roof to produce the car body less openings, which went to the final assembly line. Workers added openings. Then, the resulting car body went to the paint shop.

Simulation is too time consuming to be completely satisfactory. Analytic methods alone, however, could not be used to model some complex details or observe dynamic effects accurately. Therefore, we opted for an approach encompassing both simulation and analytic methods. We developed an iterative three-step design method that takes advantage of the speed of analytic methods and the accuracy of simulation. To insure some flexibility, we expressly chose not to automate the implementation and the transition between the steps. The design engineers manage those activities.

—In step 1, initialization, we calculate the initial target values for the cycle time and efficiency of all the production lines based on the most recent values for similar existing lines. We use analytic methods for this step.

—In step 2, macrosimulation, using analytic methods or simulation software (for complex lines), we estimate the efficiency of the production lines. We adjust the cycle times and the efficiency of production lines and modify their processes accordingly until we obtain estimates that are as close as possible to targets set in step 1.

—In step 3, validation, we estimate production-line performance using simulation software. We plug the result of step 2 into a detailed simulation model for validation. We can make slight modifications, if necessary, to make the shop's estimated production rate meet the objective forecast.

Simulation

At PSA, we perform discrete-flow simulation using off-the-shelf commercial software called Arena. Depending

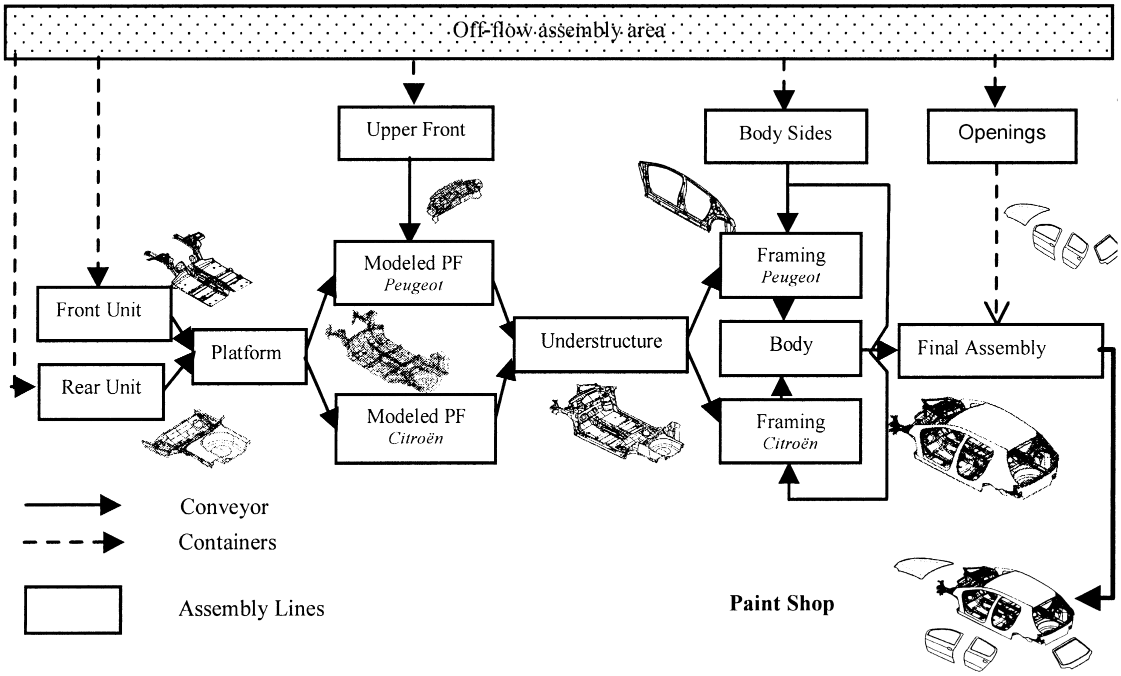

Figure 4: In the new PSA body shop, we put as many framing stations as possible on specific lines (Modeled Platform and Framing), while we put the finishing stations only on common lines (Understructure and Body). By doing so, we achieved two key objectives: (1) quick change, because we needed to build only additional specific lines and we were able to do so without disrupting the production, and (2) greater model diversity, because common lines' robots just needed to load the corresponding programs to work on different types of parts.

on the complexity of the process, the simulation lasts from 400,000 to 800,000 time units. During the first 50,000 time units, we do not collect statistics. We use this warm-up period to reduce the influence of the initial transient phase on the final results.

To speed up the modeling process, we developed specific modules for modeling the most common elements or sets of elements encountered in our shops. We constantly enriched this library of most-common elements and keep it up to date. We also developed generic framework models for the most common architectures of assembly lines. Experienced engineers validated the simulation model. They studied the simulation's animations and numerical results in a variety of situations: purely random failures, no failures, failures on specific machines at specific times, and multiple simultaneous failures, and they determined that the simulated system's response was realistic in every case. After this work, people did not have to redevelop the most common components every time they were needed, nor did they have to start from zero when modeling, whatever the assembly line. We thus reduced the time spent in developing models and minimized the risk of error. We calculate that modeling that used to take three weeks now takes a single day, almost a 95 percent reduction in development time.

Uptimes and downtimes of the stations on the line

	Advantage	Weakness
Simulation	Accurate modeling Dynamic aspects	Time consuming Specialist needed
Analytic methods	Easy to use Instant results	No dynamic aspect Limited in modeling

Table 1: Simulation is time consuming and requires specialists, while analytic methods are quick and easy to use (not to develop) but cannot handle complex modeling.

are the random variables of our simulations. We used exponential distributions after conducting research that showed they were very accurate. We found out that adding another moment to the downtime distributions in our plants (as in, for example, the hyperexponential distribution) would bring more complexity without really improving the accuracy of our results.

The results we provide in our simulation report are hourly production, average level of buffers, and shop efficiency. We also provide curves depicting the evolution of the buffer level and production rate versus time.

The Analytic Methods

In developing our general approach, we were guided by the principle that the end product had to be easy for our technicians to learn and use. We took a systems-analysis approach of isolating each problem, finding a solution, and then choosing carefully between developing a separate tool for the problem or integrating the solution into a larger set of tools. Whenever possible, we sought to adapt existing technology to meet our needs.

Because no commercial software was available for analytic methods, we had to develop our own software. To do so, we needed to collect reliability data on the most common machines (or elements), develop an easy-to-use interface, and adapt or develop analytic equations for the body shop.

Series-Parallel Production Lines

Figure 5 depicts a small body shop's production line with three stages. Because this line contains two stations in parallel (stations 21 and 22), we called it a series-parallel flow line. A series-parallel system is similar to a classical flow line, the only difference being that a given stage (the second in our example) may consist of parallel machines. As we did not find a satisfactory solution for a series-parallel line in the literature, we had to develop a new technique (Patchong and Willaeys 2001).

A line can exhibit the following problematic features that we had to deal with in our project:

—Unlike the classical system generally described in the literature, buffers can fail. We developed a method to handle such failures.

—In PSA's lines, human operators perform such tasks as loading, welding, and drilling. The impact of human operators was not addressed in the literature, so we developed a method to take it into account in our analytic models.

—Functional stops for changing tools require optimal organization and human resources management to reduce their effect on the lines' efficiency. We developed a solution for this issue.

—Various processes have different working times. We used the law of conservation of flow to solve this issue.

Modeling and Analyzing Series-Parallel Production Lines

We developed a method for modeling and analyzing series-parallel flow lines that amounts to replacing each parallel-machine stage by a single equivalent machine to obtain a classically-structured flow line with machines in series. We defined the parameters of the equivalent machine to make its behavior in the line match that of the actual set of parallel machines. We developed our model based on the Markov chain of the states of a machine in a production line (Appendix). After this transformation, we were able to use classical analytic methods for long lines to solve the resulting production line. We used a variant of analytic methods for long lines based on the work of Le Bihan (1998). Patchong and Willaeys (2001) give the details.

Unreliable Buffers

Stanley Gershwin (1987) developed an analytic method for long production lines. He used a decomposition technique to transform a problem with many machines into a two-machine problem whose exact solution can be used as a close approximation to the original line's solution. In our new technique for accommodating buffer unreliability, we replace the unreliable buffer with a new machine coupled with a reliable buffer. The new machine's reliability matches that of the unreliable buffer it represents. With this structure, we can then combine Gershwin's (1994) two-machine solution method with Buzacott's (1967) zero-buffer model to obtain an analytic solution (Figure 6).

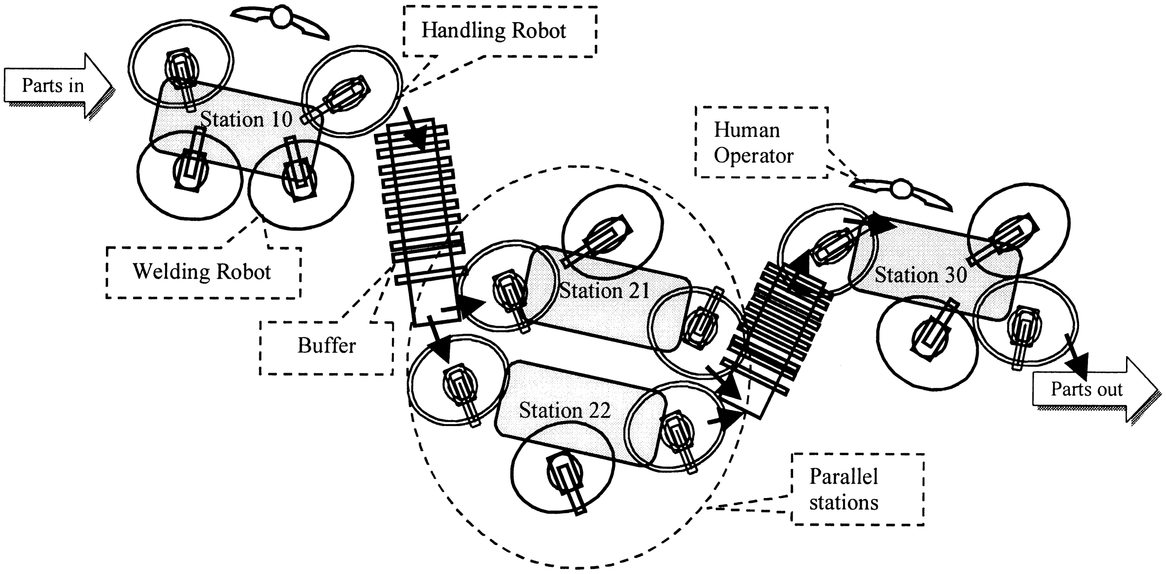

Figure 5: The series-parallel line comprises three stages separated by intermediate buffers. The first and the second stages have one station each (station 10 and station 30), while the second stage has two stations in parallel (stations 21 and 22). The part arrives in the system on the station 10. When its operation is finished a robot moves the part to the downstream buffer where it has two possible ways to go. Depending on their availability, the part can go to station 21 or station 22, which perform the same operation. The part that leaves station 21 or station 22 goes to the downstream buffer. A robot moves this part from the buffer to the station 30. After the completion of the operation on station 30, the part goes out of the system.

Processes with Different Working Time

If shops, which operate in three shifts, are to produce the maximum number of parts, we must design the production lines of all the shops to work as long as possible in a 24-hour day, while taking into account two major constraints: (1) idle time for preventive maintenance (three hours per day for areas with complex processes and two hours for the rest of the shop), and (2) worker breaks imposed by labor legislation (lines that need human operators to produce cease during breaks, but highly automated lines, because they need very few people to work, operate during breaks). As a consequence, a typical flow shop has three types of flow sectors (Figure 7): (1) The upstream flow includes complex processes (need for three-hour maintenance) and human operators (one-hour break (three 20 minutes breaks) for the three shifts) and produces 20 hours a day. The upstream flow receives preassembled parts only. This sector comprises the lines for the front unit, rear unit, front upper, and body sides (Figure 4). (2) The main flow includes complex processes (need for three-hour maintenance) and is highly automated (no need for human operators) and produces 21 hours a day. The main flow receives the parts coming from the upstream flow. This sector comprises the platform, modeled platform, understructure, framing, and body lines (Figure 4). (3) The other flow sector includes simple operations, namely preassembly machines and the final assembly line (need for two-hour maintenance). This sector, which also uses many human operators (one-hour break for the three shifts), produces 21 hours a day and comprises the final assembly lines and off-flow preassembly (Figure 4).

The following examples illustrate the ideas used in design and analysis when contiguous areas, say A1 and A2, have different working times.

—The cycle times of A1 and A2 and their areas should be roughly proportional to their planned working times.

—If A1 respects worker breaks while A2 observes no

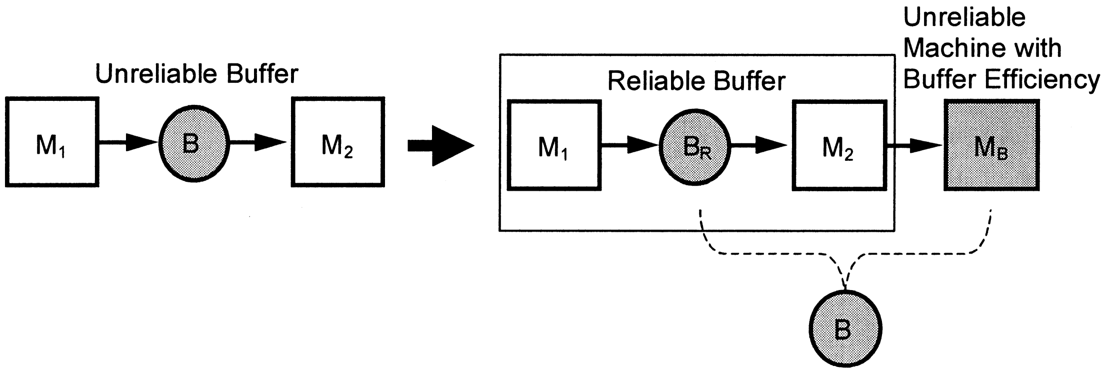

Figure 6: The solution technique for unreliable buffers consists of replacing the unreliable buffer with an unreliable machine coupled with a reliable buffer.

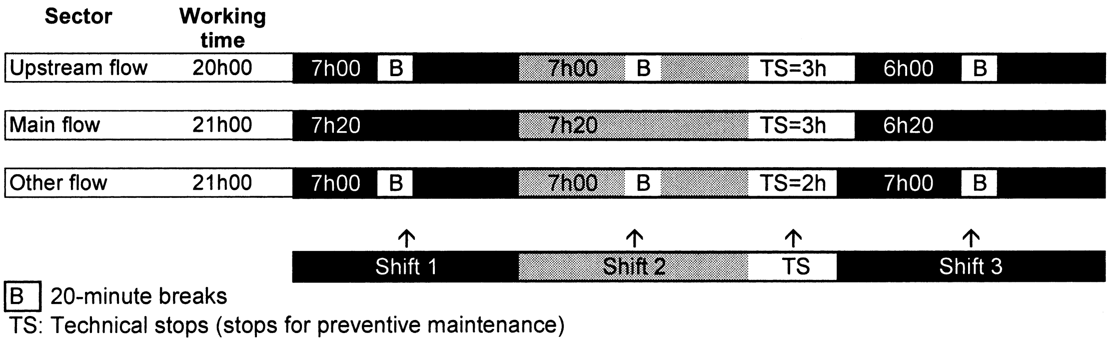

Figure 7: The upstream flow, which includes complex processes (three-hour maintenance needed) and human operators (3 shifts x 20-minute breaks), produces 20 hours a day. The upstream flow receives preassembled parts only. The main flow, which includes complex processes (three-hour maintenance needed) and is highly automated (no need for human operators), produces 21 hours a day. The main flow receives parts coming from the upstream flow. The other flow sector includes preassembly machines and the final assembly line (two-hour maintenance needed). This area, which also uses many human operators (3 shifts x 20-minute breaks), produces 21 hours a day.

breaks, the intermediate buffer should have added to its capacity the amount that would be produced during the break, if the two areas worked synchronously.

Human Factors

Experience has shown that the more complex a machine is, the less reliable it tends to be. It therefore seems preferable to use a human operator instead of an automated and complex machine to perform intricate tasks; hence, PSA uses manufacturing systems comprising both automated and manual stations in its plants. Patchong et al. (1997) proposed a cost-based approach that confirms use of such mixed systems. Although the literature abounds with papers on methods for analyzing manufacturing systems, few concern mixed manufacturing systems. We performed research in this area on site at a manufacturing plant within PSA Peugeot-Citroën. Based on this work, we defined a machinelike reliability (determination of a mean time to repair and a mean time to fail) for a human operator (documented in internal, unpublished reports).

Unlike the processing time of an automated station, which varies very slightly from part to part (Figure 8),

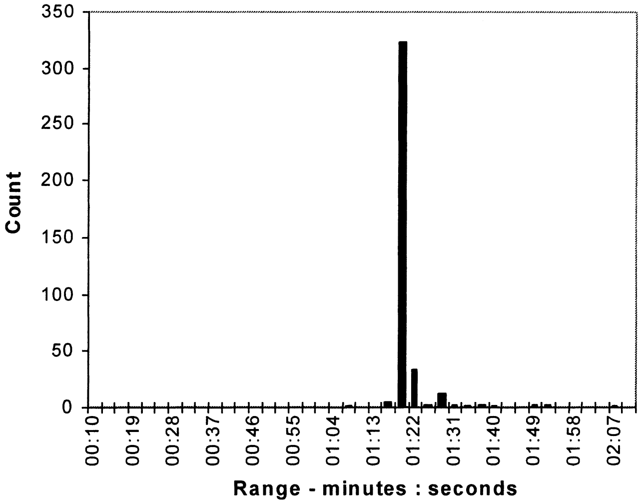

Figure 8: A record of the processing time of an automated station, which is represented on the x-axis here, shows that it varies very slightly from part to part.

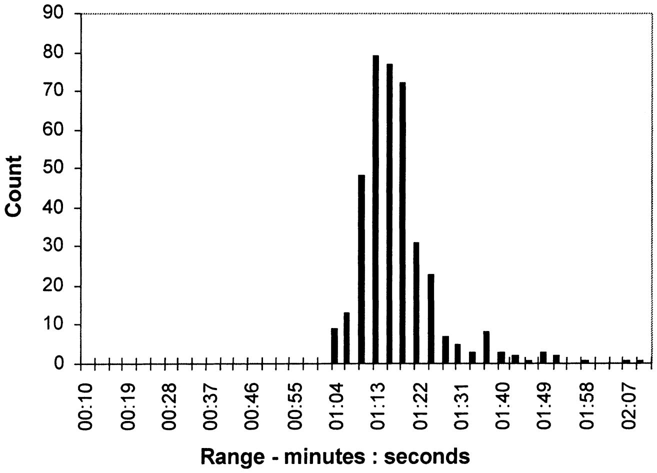

Figure 9: The x–axis represents the record of the processing time of a manual station. It shows that, unlike the processing time of an automated station, which varies very slightly from part to part (Figure 8), the processing time of a manual station is highly dispersed.

the processing time of a manual station is highly dispersed (Figure 9). We defined *overtime* for a manual station as the occurrence during a cycle of processing time greater than the expected processing time (Patchong 1997). We took overtime into account using a variant of Figure 13 in the Appendix. We tried to define for a human operator the same parameters as for machines so that we could use the human operator as a machine in existing models or in models we were developing. We did this by using an analogy, which consists of considering overtime as failure and, symmetrically, resumption as repairing. The operational state becomes a macro state made of two states: the *overtime state* and the *normal working state,* which corresponds to work without overtime (Figure 10). The probability rates of transition between these two states are *d* (*average overtime rate*) and *r* (*average resumption rate* or average rate of return to the normal working state). In the normal working state, machines process without overtime or are idle (waiting for a part to process or evacuating a finished part). Given that we now considered overtimes to be failures, we derived the three main parameters for the "human machine" (average failure rate, average repair rate, and processing rate) and used them in the analytic methods we developed. The curve depicted in Figure 11 proved very valuable. It predicts an operator's efficiency as a function of the extra time he has to perform a task.

We collected data to verify the model results empirically. It showed that the production rate of the model yielded by our method was very close to the actual production rate (for example, a 0.7 percent error for a manufacturing system with no intermediate buffers).

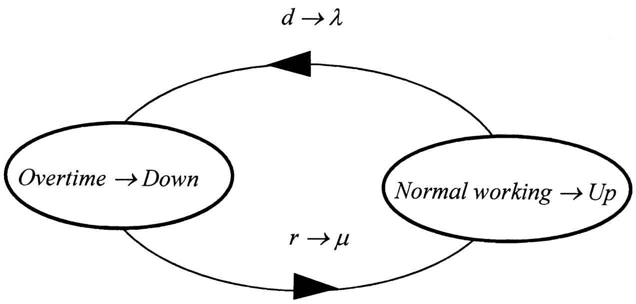

Figure 10: We defined *overtime* for a manual station as the occurrence during a cycle of processing time greater than the expected processing time. Using a variant of Figure 13 in the appendix, we define for a human operator the same parameters as for machines based on an analogy, which consists of considering overtime as failure and, symmetrically, resumption as repairing. The operational state becomes a macro state made of two states: the *overtime state* and the *normal working state*, which corresponds to work without overtime. The probability rates of transition between these two states are *d* and *r*.

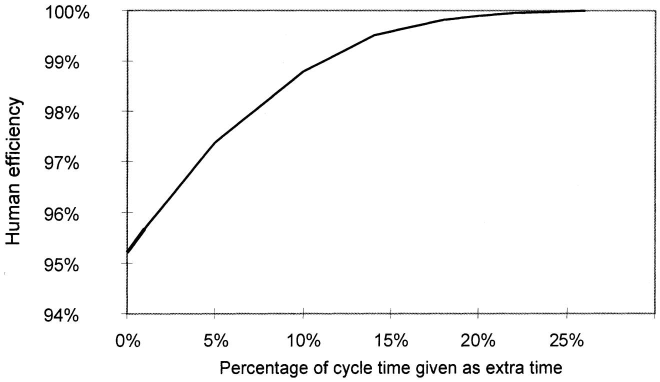

Figure 11: Based on the data collected in our manufacturing plants, we modeled human efficiency as dependant on the extra time the operator had to perform a task. This curve shows that the more extra time an operator had the greater his efficiency. For example, an operator with two percent of extra time would have a 96 percent efficiency, whereas with 10 percent of extra time, he would have had an efficiency of 99 percent.

Other Issues

We also analyzed the design of kanban buffers and operating policies for the manufacturing systems in the upstream flow sector. Our goal was to guarantee independence between the diversities to reduce the disorder in the production schedule. We took functional stops (for example, stops for grinding welding-gun tips, changing welding-gun tips or welding guns, and cleaning tools) into account. (We discuss the work on kanban and functional stops in internal, unpublished documents.)

After developing the design method, we had to collect reliability data. We developed a unified process for collecting accurate reliability data and conducted a campaign in the plants to explain why it was so important to maintain the process. Before collecting data, we first determined the most common elements (machines) in the manufacturing process. We decided to focus on the 20 percent of elements that constituted 80 percent of the manufacturing systems in the body shop. This approach helped us save money while guaranteeing a good level of accuracy for our results. Then, we collected data within PSA's major factories to find the MTTR (mean time to repair) and the MTTF (mean time to failure) of these most common elements.

After implementing the new method, we conducted training sessions in PSA's plants and research and development divisions on use of the method and software and new shop standards. During these sessions, we taught employees how to design manufacturing systems.

Training

We believe that the best operations research projects go beyond tool building and analysis. They should have lasting value beyond the solution of the current problem. To educate and help R and D staff members design better lines, we created new teaching materials. Armed with lessons learned from this project, they are creating efficient lines and responding quickly to car-model changes.

An Example of Lesson

In *The Goal,* Eli Goldratt and Jeff Cox (1986) stressed that eliminating (or protecting, if not) bottlenecks is the key to improving manufacturing systems. PSA did a great job over the years of eliminating bottlenecks, so we often faced the question, "How do we improve a line that has no bottlenecks?" The answer is to work next on improving the performance in the neighborhood of the buffer closest to half full.

For example, to determine which machine to improve first in a line with 11 identical machines in series and no bottleneck, we consider improving each of the 11 machines one at a time. We have three ways to improve a machine's performance (Figure 12): (1) by reducing the mean time to repair (MTTR), (2) by increasing the mean time to failure (MTTF), or (3) by reducing the cycle time. Regardless of which type of improvement we made, our model showed that the line's productivity would be greatest when the improved machine was at the middle of the line where the buffers are half full. So the answer to our question appeared to be, in practice, the application of a rule we set, which improved first the performance of the station whose buffer is closest to half full.

Impact

Using OR methods we created

—User-friendly software called DispO,

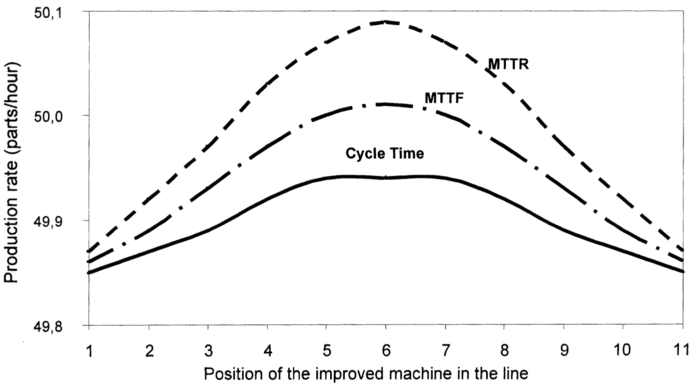

Figure 12: To improve a line with 11 identical machines in series and no bottleneck, on which of the 11 machines must we focus first? We have three ways to improve a machine's performance: (1) reducing the mean time to repair (MTTR), (2) increasing the mean time to failure (MTTF), or (3) reducing the cycle time. This graph depicts system performance as predicted by our analytic models for each type of improvement applied to each of the 11 machines. The position of the machine improved is represented by the x-axis. The graph shows that, regardless of which type of improvement we made, the line's productivity would be greatest when the improved machine was at the middle of the line where the buffers are half full. This is consistent with our teaching with respect to improving first the performance of the station whose buffer is closest to half full.

—Methods people could use to deal with specific issues, and

—Training documents to help staff members design production lines quickly.

Our work has helped PSA and its suppliers, and other companies using the new tools. The methods we developed have gained the attention of some academics.

Impact on PSA

We wanted to create software that production-line designers could use on their own without supervision. Those people, who are at the lowest level on a project team, perform many tasks. We spent a lot of time on the human-machine interface to make the software user-friendly and easy to learn, a necessity to insure its viability. In addition, we assert that our software tool consistently enables users to produce fundamentally better results that they would have otherwise. They gain greater mastery of their assembly line and therefore develop better designs quickly while drawing on fewer human resources.

The new methods and software put to rest some incorrect but deeply ingrained beliefs and practices. For instance, manufacturing engineers used to ask us to build shops in which each machine had a lower cycle time than that of the downstream machine(s). They called this philosophy *pushed flow*. We proved that, in a balanced shop, the machines to improve first are those in the middle of the shop (Figure 12). So there was no reason to add potential to upstream machines.

The following are some lessons people learned from the tools we developed:

—People used to think that the capacity of buffers that are always full must be increased so that there would be enough place to store more material for the good of the production. We proved that one must focus on half-full buffers and then, whenever possible, reduce the capacity of buffers that are full all of the time to increase the capacity of half-full buffers.

—People used to believe that buffer allocation did not really matter. We showed that given equal total buffer space, several smaller buffers are better than a few bigger buffers.

—People used to think that the action that paid back the most was decreasing cycle time. We demonstrated that for equivalent impact, the most profitable actions were, in order: (1) decreasing MTTR, (2) increasing MTTF, and (3) decreasing cycle time (Figure 12).

—Some manufacturing people used to calculate the equivalent cycle time of a set of parallel machines as equal to the mean of their cycle times. We showed that the inverse of the equivalent cycle time of a set of parallel machines is the mean of the inverse of their cycle time.

—It was commonly believed that the resulting efficiency of a set of machines in a series without an intermediate buffer is the product of their efficiency. This is inaccurate, and for the kinds of systems we dealt with, the difference with the accurate formula is over four percent. Buzacott (1967) gives the accurate formula.

The new methods and software demonstrated that sustained improvement was possible for all the manufacturing systems in PSA's body shops. PSA raised its standards accordingly.

Credibility for Operations Research

PSA personnel, initially skeptical about operations research (OR), had the opportunity to compare the results the analytical methods predicted with the actual outcomes. Persuaded by the accuracy of the forecasts, they have come to respect OR. Consequently, other PSA divisions have adopted the tools we developed and have initiated further OR projects. The success of the project helped management to understand the value of OR in our daily mission. PSA management decided to create a new section dedicated to shop design and performance improvement.

Modeling that used to take three weeks now takes a single day.

Profit Increase

Because we could improve production-line designs rapidly, the plants could ramp up quickly and satisfy strong demand in the market without pressuring the workforce to work overtime. These improvements led to a better social climate and fewer human errors, which implied better-quality products. We took careful measurements of productivity before and after the changes. In 2001, the average increase in productivity was four percent and 75 percent of the cars PSA manufactured worldwide were directly affected by this OR work. This percentage will be over 90 percent before long.

Knowing that the car shop is the bottleneck of PSA's plants, we based our calculation of the benefit on the belief that improving the plants' efficiency implies increased volume which implies additional profit. We assume that PSA sells all the cars produced, and that is precisely the case. PSA's lines are running at their full capacity, and its waiting list does not decrease. We conservatively estimate that, in 2001, our work brought PSA US $130 million of additional profit! US $130 million dollars is about 6.5 percent of PSA's total profit from selling cars in 2001. This is money that the company would not have made without OR.

Impact on Other Companies

By using the new tools and the accompanying training documents, PSA's suppliers gained knowledge about manufacturing system engineering. As a result, they designed accurate systems quickly. Furthermore, using the same tools led PSA and its suppliers to use the same standards, reducing the time they spent negotiating production schedules. Almost all of PSA's suppliers now use the new tools, even though PSA does not require them to do so. We expect the impact of our work to go beyond PSA and its suppliers as other companies acquire the software. PSA management is evaluating an outside company's proposal to market the software.

In 2001, our work brought PSA US $130 million of additional profit.

Impact on Academia

As the lead analyst on this project, Alain Patchong was invited to present this work at Massachusetts Institute of Technology, and at some top grandes écoles (French engineering schools) before audiences consisting mainly of tomorrow's decision makers. His presentations, which have become annual events at some schools, illustrate how applied OR can bring great benefits to organizations.

Acknowledgments
We thank Stanley B. Gershwin, senior research scientist at MIT, for his helpful suggestions. We also thank Richard Rosenthal for his valuable advice.

Executive summaries of Edelman award papers are presented here. The complete article was published in the INFORMS journal Interfaces [2003, 33:1, 36-49]. Full text is available by subscription at http://www.extenza-eps.com/extenza/contentviewing/viewJournal.do?journalId=5