Proceedings of the 2008 Winter Simulation Conference
S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, J. W. Fowler eds.

EMBEDDING HUMAN SCHEDULING IN A STEEL PLANT SIMULATION

David Briggs

Business Improvement and Supply Chain
Corus Long Products
Brigg Road
Scunthorpe, North Lincolnshire, DN16 1BP, UK

ABSTRACT

A simulation was commissioned to understand the interactions that constrain the capacity of a steel plant. The aim was for this to become a reusable tool that could evaluate the effect of future changes to market requirements and operational practices. This paper describes how a simulation model incorporating human decision-making was conceived and constructed. The use of simulation as a tool for knowledge capture in scheduling is considered. The resulting tool has been in use for four years and has acted as a driver to reconsider where the real processing bottlenecks are and what part scheduling can play in managing them.

1 INTRODUCTION

In 2003 Corus' Scunthorpe steel plant's goal was to identify the configuration and product mix requirements to achieve 4.6 million tons/annum (mtpa) liquid steel. Concern was raised that this declared strategy required a yielded 4.35 mtpa of continually cast steel compared to a plant record of 3.68 mtpa. The operational research department traditionally had answered questions of this nature by using a combination of:

- Deterministic approaches to individual plant capacities
- Scheduling key components using 'ideal' conditions assuming no inter-plant dependencies
- Simulations of individual plant units or small groups of plant using ideal timings and queue based approaches

The output from these static models, scheduling tools and 'local' simulations had proved sufficient when applied to small self-contained plant. However, the summations of the outputs had historically shown little resemblance to reality. From this it was recognized that the total plant casting capacity is less than the sum of its four casters. To help the steel plant meet their production goals we needed to understand the interactions that constrained the capacity. Our conclusion was that we would have to simulate the whole steel plant. As a result we undertook to provide a strategic tool to enable Scunthorpe steel plant's capability to be understood. Specifically to:

1. Determine what tonnage can be produced
2. Determine what product mix can be produced and the effect of the mix on tonnage
3. Understand where plant constraints exist
4. Be adaptable to changes in plant configuration and process parameters
5. Demonstrate the effect of modifying or removing constraints

1.1 The Scunthorpe Steel Plant

Steel is produced at Scunthorpe from liquid iron through the Basic Oxygen (BOS) process. Iron arrives at the BOS plant and is poured into circa 300 tonne ladles, which are then desulphurised through one of two units. It is then charged into one of three BOS Vessels where the iron is converted to steel. At this stage alloying elements are added to control the finished steel properties. This 300t unit of steel is called a cast. Additional Secondary Steelmaking processes are then carried out depending on the grade (selected according to the end use of the steel being made). Typically around 50 out of over 1300 grades are produced each week. These processes are carried out at one of three Ladle Arc Furnaces (LAF) and two Vacuum Degassers (VDG). The chemical analysis, homogeneity and temperature of the steel must all be closely controlled to ensure that the steel is fit for purpose. After steelmaking has been completed the cast is sent to one of four continuous casting machines (Concast) where it is cast into a precise solid section for dispatch for further processing or end use. Groups of casts, called sequences, of identical or similar grade are processed through the casters without break. The timing of arrival at the casters is critical, if a cast is delayed the sequence may be broken incurring a costly and time-consuming machine reset and the logistical problem of holding or recycling the delayed cast of steel.

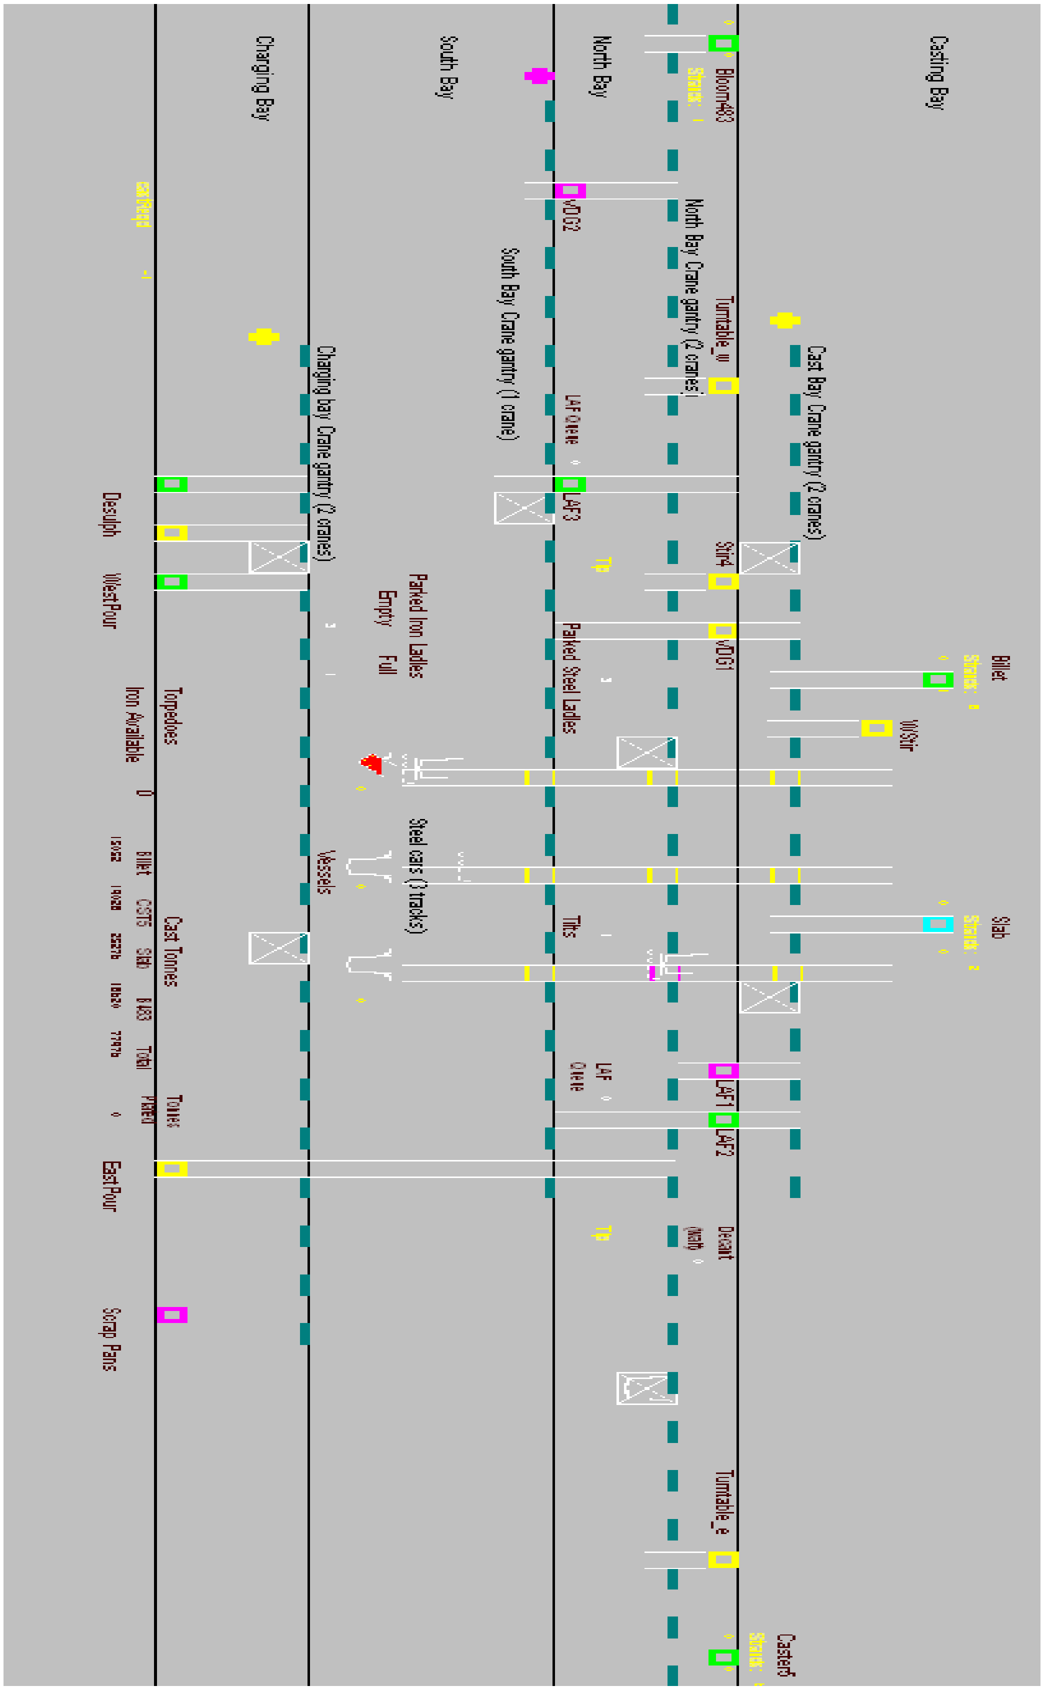

Figure 1: Layout of the Steel plant

The steel plant is arranged in a series of bays. Movement of both empty and full ladles is carried out by cranes (within bay movements) and steel cars (between bay movements). Figure 1 shows a snapshot of the simulation model. The model uses a plan view of the layout of the plant.

1.2 The Essential Role of the Schedule

A shift planner using a bespoke decision support system called the Scunthorpe Shift Scheduling System (SSSS) actively schedules the Scunthorpe steel plant. The SSSS is connected to the plant's process control systems and is continually updated with production events. It then reflects the impact of these events onto the forward schedule enabling the planner to modify the schedule as required. The planner also receives updates from the operators of the different production units about forecast availabilities, processes that are beginning to deviate from the standard and any other information that may require a change to the schedule. It is down to the experience of the planner to decide which information is significant and needs to be input into the SSSS to revise the schedule. Typically the schedule is not robust to change, with relatively minor deviations to timing having a large impact on its feasibility.

Traditionally, simulation models reveal processing bottlenecks by blocking and queuing material at pinch points or by rejecting or scrapping material that cannot be processed. In the steel plant neither of these methods is appropriate. Liquid steel can only be held for a limited time without changing its properties, not least its liquid state. The planner has to decide how to divert delayed steel through other processes to minimize the impact of any processing bottleneck. If rejecting the steel is the best alternative, it too has knock on impact as the same logistical resources (cranes and steel cars) that are used for production have to be diverted from their normal tasks to handle the rejected steel. Consideration was given to solving these problems by 'rewinding' the simulation when a problem occurred and repeating the simulation with an alternative schedule to try to avoid the problem. This is similar to the method used by Ramakrishnan, Lee, and Wysk (2002), though in this case was rejected because the aim was to understand actual operation and the steel plant capacity, where replanning with hindsight was not an option.

This dependency on the schedule was recognized during the scoping of the Steel Plant Simulation Project. Any attempt to mimic the steel plant's operation would have to include the ability to change the schedule on casts already part way through production. It was also recognized that there was no existing definition of a 'right' way to reschedule following a perturbation to a plan. The reasons for this are twofold; firstly the planner's expertise has not been adequately captured and secondly the objectives are fluid, changing with commercial and production priorities. This means quite literally that today's best solution may be unacceptable if the identical situation occurred tomorrow.

The scheduling cycle, in Figure 2, shows the processing cycles that occur in the scheduling and production at the steel plant. The scope of the simulation was to include each of these steps.

2 HYBRIDISATION/HUMAN COMPUTER INTERACTION (HCI)

The operation of the Scunthorpe steel plant is heavily dependent on the schedule. The schedule in turn depends on the planner responding to uncontrolled events in conjunction with both tactical and strategic priorities.

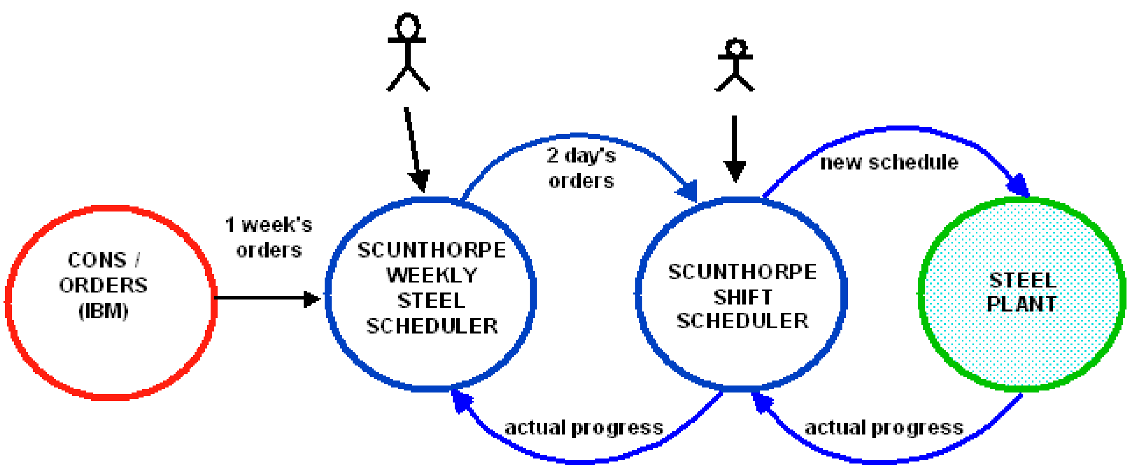

Figure 2: The scheduling cycle

As a project team the question we asked ourselves was could a simulation of the steel plant be validated if it did not include the impact of the schedule and the planner? We concluded that it could not. We further decided that within the scope of the simulation project we could not adequately capture the planner's expertise to construct an automated scheduling system within the simulation. This was a pragmatic decision based on our perception of the complexity of scheduling and the absence of any clear metrics to assess between good and bad scheduling decisions. The project team had been responsible for development and support of SSSS for five years (the system itself had been evolving for nearly 20 years) and so had familiarity with the hard operational constraints. We determined that our simulation should include an enhanced version of SSSS that solved the simpler scheduling conflicts and called on planner support when key scheduling decisions were required. We would also give the planner overall control to 'correct' any of these 'simpler' decisions that he disagreed with.

This created a requirement to have access to a planner for operational runs of the simulation. The expected consequence of this would be to greatly limit the number of simulation experiments that could be carried out. This was because part of the simulation process would be for the planner to understand a simulated plant condition sufficiently to make a 'realistic' scheduling decision.

Having decided that the simulation would interact with an operator it was then necessary to consider the HCI. We quickly decided that offline versions of the existing scheduling systems would provide the simulation operator with tools that have both the same functionality and the same interfaces that the planner has, removing a need for training in a new interface.

3 THE FORM OF THE SOLUTION

The simulation was constructed as a series of communicating modules each replicating part of the process from order entry through scheduling to production, with the core simulation engine representing the steel plant. Table 1 lists the steps that relate back to the cycles shown in Figure 2.

The Simulate event/Update Schedule cycle repeats until either a conflict occurs, requiring manual intervention or the end of day is reached and the Daily Cycle is repeated. The planner can intervene at any time to modify the schedule.

Table 1: Simulation process steps

Step	Process	System
1	Order entry	Access database
2	**Daily cycle**	
2.1	Sequencing	SWSS*
2.2	2 Day schedule	SWSS
2.3	Confirm plan	SWSS
2.4	**Shift cycle**	
2.4.1	Shift (c. 8 hour) schedule	SSSS
2.4.2	Simulate event	Simulation engine
2.4.3	Update schedule	SSSS
2.4.3.1	Detect conflict	SSSS
2.4.3.2	Resolve conflict	Planner

* Scunthorpe Weekly Scheduling System

4 MODEL DESIGN

4.1 Order Entry

An Access database was constructed to represent the IBM mainframe order entry system. Orders were held at the cast level with all the details that define processing requirements and timings. The same database was used as a utility system to store and maintain simulation parameters such as the distributions used for breakdowns, repairs and process times. This resulted in approximately 150 data distributions, each based on observed data. Figure 3 shows an example the process time data. Statistical tests were applied to verify, for example, that crane speeds in different bays could be modeled using a single distribution. An interesting observation is that the database held both the 'standard' processing times used in scheduling and distributions based on extensive collection of actual times. As a side benefit of the project the standard scheduling times were reviewed based on this data.

4.2 Offline Scheduling Systems

Offline simulation versions of the two existing online scheduling systems, SWSS and SSSS were constructed. SWSS is used to build coarse week length schedules on a daily basis to take account of planned outages and the competing requirements of different steel qualities on the operational plant. SSSS is used to continually schedule to a greater detail in a shorter time horizon, typically eight to 12 hours. In the simulation a single operator interacts with both systems whereas different operators, a day planner for SWSS and a shift planner for SSSS use the live systems. The scheuling systems operate to the granularity of the minute.

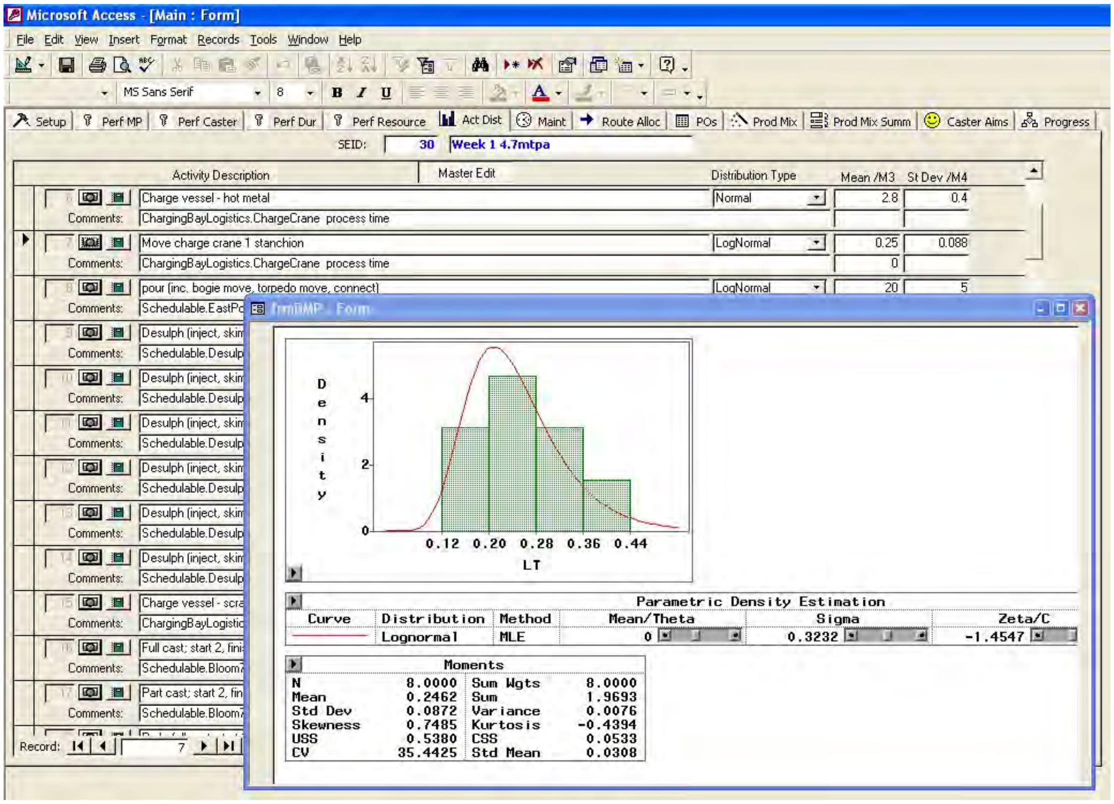

Figure 3: Process time interface

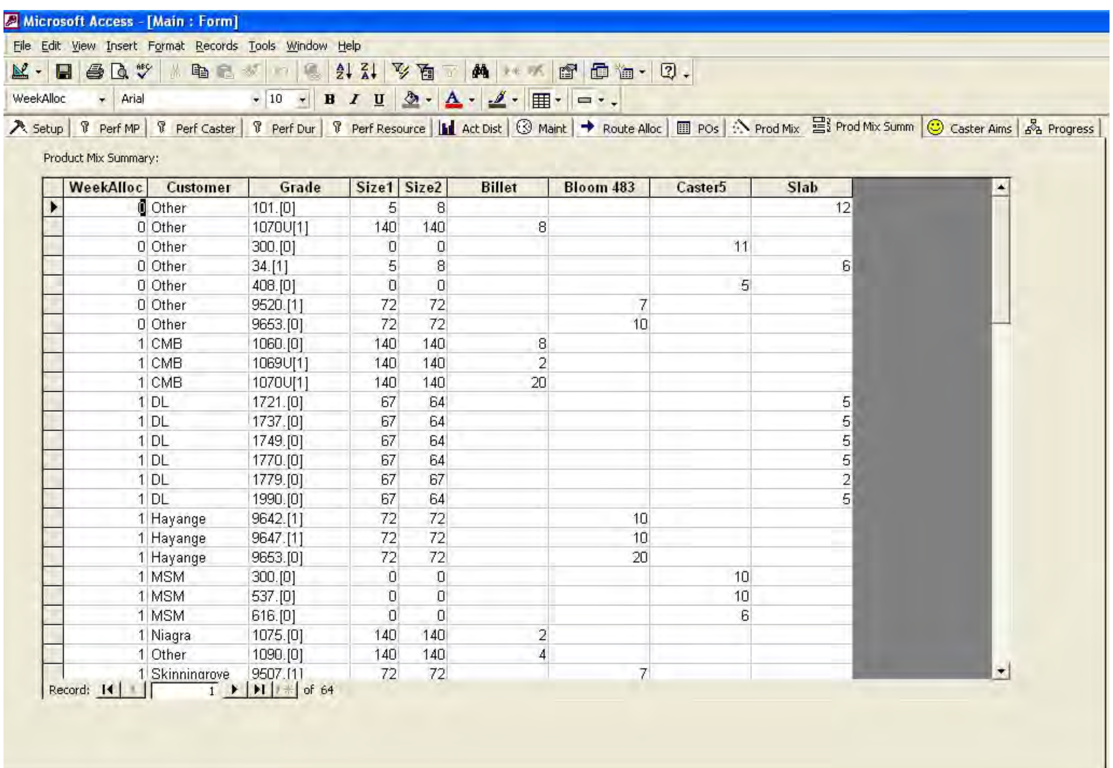
Product Mix Summary:

WeekAlloc	Customer	Grade	Size1	Size2	Billet	Bloom 483	Caster5	Slab
0	Other	101.[0]	5	8				12
0	Other	1070U[1]	140	140	8			
0	Other	300.[0]	0	0			11	
0	Other	34.[1]	5	8				6
0	Other	408.[0]	0	0			5	
0	Other	9520.[1]	72	72		7		
0	Other	9653.[0]	72	72		10		
1	CMB	1060.[0]	140	140	8			
1	CMB	1069U[1]	140	140	2			
1	CMB	1070U[1]	140	140	20			
1	DL	1721.[0]	67	64				5
1	DL	1737.[0]	67	64				5
1	DL	1749.[0]	67	64				5
1	DL	1770.[0]	67	64				5
1	DL	1779.[0]	67	67				2
1	DL	1990.[0]	67	64				5
1	Hayange	9642.[1]	72	72		10		
1	Hayange	9647.[1]	72	72		10		
1	Hayange	9653.[0]	72	72		20		
1	MSM	300.[0]	0	0			10	
1	MSM	537.[0]	0	0			10	
1	MSM	616.[0]	0	0			6	
1	Niagra	1075.[0]	140	140	2			
1	Other	1090.[0]	140	140	4			
1	Skinningrove	9507.[1]	72	72		7		

Record: 1 of 64

Figure 4: A simulated weeks order mix

4.3 Simulation Engine

The simulation engine was constructed using Lanner's Witness simulation package. Witness was chosen due to the teams experience with the tool and because it had both an inbuilt database interface and a facility to callout to bespoke operational code (written in the C programming language). This allowed us to write a cross platform message system to communicate between the subsystems that could interact with the core Witness simulation engine.

The simulation engine required the ability to change the next planned event on a cast. Usually within a discrete event simulation (DES) once an event has been determined it must occur. It was necessary to make a distinction between events that were deterministic, e.g. once started casting a cast will complete after a (randomly) determined time, and those that were conditional, e.g. a cast will be sent to the VDG plant if it becomes free in the next 20 minutes, otherwise it will be sent direct to a caster (as a result of a reschedule in SSSS). A job list mechanism was written in Witness to store the planned jobs that could be updated by the SSSS. When the time expired on a planned job, it would be committed to and any subsequent schedule update would be disregarded. In this way a clear responsibility could be defined between the simulation engine and the SSSS. At every update SSSS would send the next job for every planned and in-progress cast. The simulation engine would update the job list for every cast that did not have a committed next event.

4.4 Connectivity and Processing

The scheduling subsystems both run on DEC Alpha Workstations under the VMS operating system and the simulation engine and database are on a PC running Windows XP. A messaging system to communicate between the subsystems was written. Most communication is between the simulation engine and the SSSS. The simulation engine needs to pass 'progress events' to the SSSS and then wait for a 'reschedule' event. The SSSS processes the progress event and updates the schedule as a consequence of the new information. Where necessary the operator would be called on to resolve an implied conflict. The revised schedule of the next planned event on each simulated cast is then sent back to the simulation. It updates its job list to reflect the target times and in some cases to change the next process. This cycle continues through the life of the simulation experiment.

5 MODEL CONSTRUCTION AND VALIDATION

The modeling phase of the project began in April 2003 and was completed in March 2004 consuming approximately three man-years of effort. The effort was split fairly evenly between:

- The core simulation engine
- Modifying the scheduling systems and writing the order entry software
- Data collection and analysis

The model validation was carried out using a sample week's operation with a known product mix and constraining the simulation operator with the actual operational priorities and scheduled outages. Validation was judged by comparison between scheduled and actual throughput, utilization and availability and by the subjective measure of the operator's judgment of the behavior of the simulation.

During validation runs we were able to simulate approximately one days operation in 1½ hours. However this speed reduced as we introduced higher production rates with consequentially more complex schedules that required greater manual intervention.

6 EXPERIMENTATION AND ANALYSIS

In the year since the simulation was commissioned the steel plant had itself commissioned major enhancements that were shortly due to come on-stream. The simulation was used to understand the consequence of the new configuration to allow better planning of their operation.

6.1 Experiments, Robustness, and Replications

Simulation experiments consisted of fabricated weeks worth of order data and plant availabilities to examine the impact on throughput and the moving bottleneck points. Figure 4 shows a summary of a weeks order mix. Figure 5 shows the simulated progress of a single day of this order mix through the model. The long run times of the simulation, requiring an experience planner to operate the model full time, restricted our opportunity to perform many experiments or replications within experiments.

This has driven the analysis to focus on understanding the behavior of the steel plant rather than on confident quantitative measures of capacity. The only formal replications that have been carried out were during the validation stage. This has led to the question of what does a replication of the simulation mean when there is an embedded operator in the decision making process? Traditionally a replication would be a repeat of the simulation with all deterministic values held fixed whilst the random factors were varied. Where some decision-making depends on human experience and intuition should this be considered a random or a deterministic process? Anecdotally we have discovered that there is an operator effect on the simulation output leading us to anticipate a similar effect in actual plant operation.

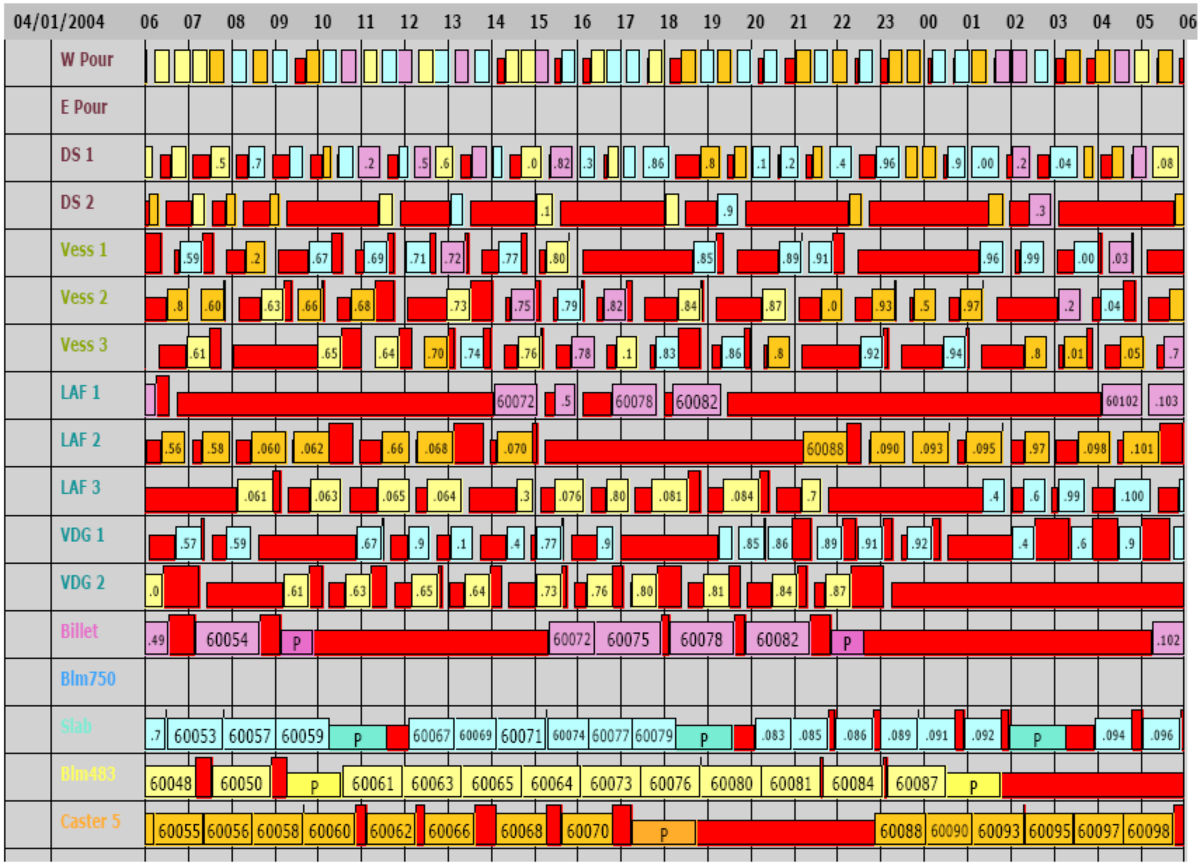

Figure 5: A simulated days production

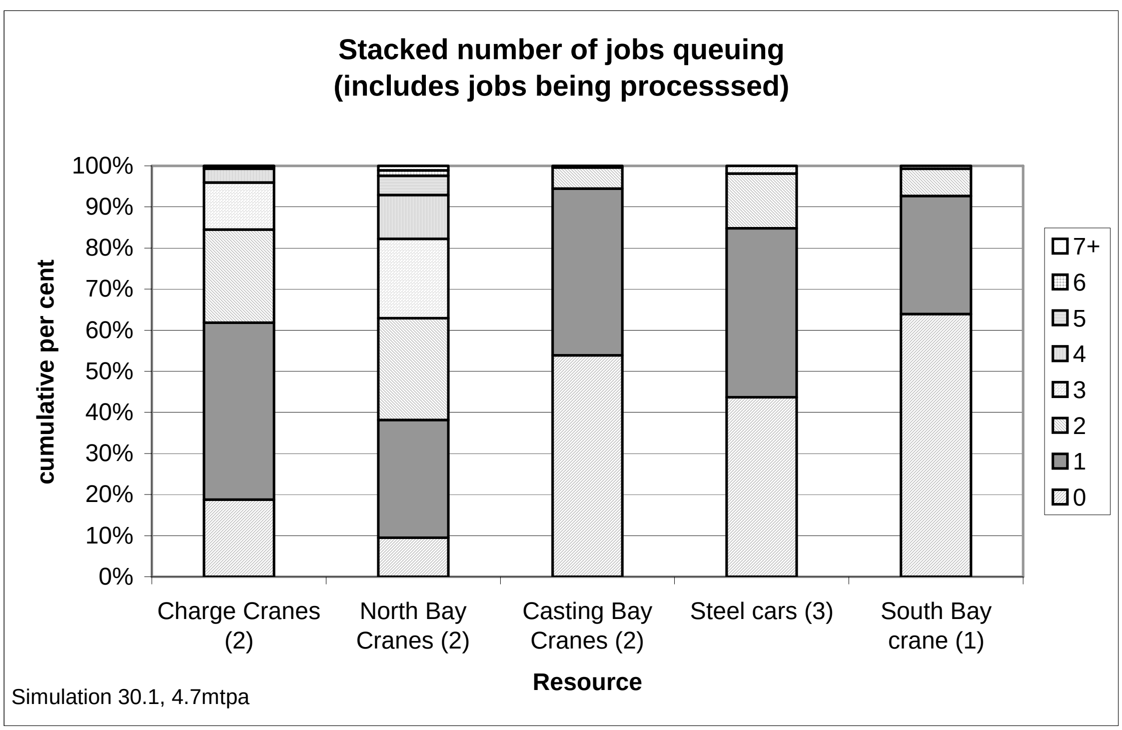

Figure 6: Logistical unit queues

6.2 Logistics

Prior to the simulation model, assessment of the capacity of the steel plant had always focused on the production units (e.g. vessel, secondary steelmaking facilities and casters). The simulation revealed the importance of the logistical resources - the cranes and steel cars - that move the steel ladles between them. Figure 6 shows the percentage of time that each set of logistical resources has a number of queued jobs. It highlights the congestion on the two North Bay Cranes where nearly 35% of the time at least one job was delayed.

The SSSS and the planner have also focused on the production units. Including the logistical resources formally in SSSS and hence giving visibility of where these are the constraint may give an operational improvement.

6.3 Planning and Scheduling

In the simulation the planner takes all the operational decisions, from building the coarse weekly schedule down to the detail of how to resolve each schedule conflict. Having a single consistent controller of the plant has allowed the simulation to achieve higher than actual throughput. Initially this appeared to be a deficiency of the simulation however it does reflect how the principle of having an active, revised schedule should work. If anticipated deviations are fed into the SSSS routinely then it can be used to revise the schedule. Currently each operational unit can make its own, possibly sub-optimal revision to suit their own requirements. Done independently this fails to assess its impact on the other operational units. This has highlighted the perhaps intuitive but often ignored conclusion that a schedule has most value if it is adhered to.

7 BENEFITS AND LEARNINGS

To assess the success of the simulation project we need to reconsider the five stated aims, to:

1. Determine what tonnage can be produced
2. Determine what product mix can be produced and the effect of the mix on tonnage
3. Understand where plant constraints exist
4. Be adaptable to changes in plant configuration and process parameters
5. Demonstrate the effect of modifying or removing constraints

The first aim is the easiest to state but the most complex to satisfy. It can only be answered within the context of the other four aims. The simulation model has become a useful communication tool to explain why 'what is the capacity' is a relative question that requires context. It has helped focus the attention more on 'if we want to achieve x what must be put in place?'

The simulation model has been used to investigate a number of 'what if' scenarios. This has identified the characteristics of order mixes and plant operation that impact on capacity. It has also allowed the relative value of operational changes to be assessed, e.g. the reduction in cycle times or improvement in reliability.

The model has proved itself to be adaptable with constant updates being made as plant capability has changed. This has included minor parameter changes, e.g. in process times or reliability, and major logic changes where different logistical flows have been evaluated including the commissioning of new plant. The model continues to be used as further evaluations are identified to assist in the ever-changing demands on the steel plant.

7.1 Removing The Man – A Better Schedule?

An experienced planner using a bespoke decision support system carries out the scheduling of the steel plant. Can the experience of the simulation model assist in improving the quality of the decision support? Can we remove the reliance on an individual's experience and make the scheduling process faster and more robust?

The simulation could be used as a platform for knowledge capture of the scheduling process. It could be used to identify what information a planner needs in able to make scheduling decisions, to build up a case library of scheduling decisions in defined situations and to assess the quality/robustness of these decisions. We discovered that planners have different approaches to scheduling. This is definitely a factor that needs to be considered when identifying scheduling rules. The simulation could be used to assess any scheduling rules that results from this research.

In considering how to improve the schedule it has been useful to classify why the schedule is constantly changing. Explicitly, why does the steel plant deviate from plan?

When a schedule is constructed it uses the best estimates of process timing and availability. However it is done with full knowledge that these are just estimates and so the schedule as constructed is unlikely to be realized. These deviations from the schedule are due to one or more forms of uncertainty or dynamics, as described by Ali and Seifoddini (2006).

Many of the processes have a deterministic (process dependent) variability in that the processing time depends on attributes of the steel being processed. Accurate information on these attributes would allow a better standard process time to be determined.

There may be deliberate deviations from the schedule. These may reflect real or imagined inadequacies in the schedule. However the deviations are in themselves likely to be sub-optimal in that SSSS is the only tool we have to analyze the consequences of these controllable but unpredictable deviations.

Each timing estimate is susceptible to minor random deviations. This is the most common use of the term variability.

Unpredictable events that are large enough to always cause a significant disruption to the schedule, e.g. major plant failure should be classified separately to random noise because they are by definition events that no schedule can be robust to. Elimination of these events is outside of the scope of scheduling and must be tackled by alternative means, e.g. maintenance regimes.

The construction of a better schedule depends on which factors we are aiming to improve. By providing the schedule with timely information updates the deterministic variability can be reduced. By making every effort to follow the schedule, and where this is not possible then feeding the extent of this conflict into the schedule for revision, the controllable variability can be reduced. The schedule cannot eliminate random noise, however measures of robustness may be possible.

The simulation model has already contributed by guiding our focus on to the real controlling processes of the steel plant. The next challenge for the simulation is to help guide us on how to schedule the steel plant better.

REFERENCES

Ali S. A., and H. Seifoddini. 2006. Simulation Intelligence and Modeling for Manufacturing Uncertainties. In *Proceedings of the 2006 Winter Simulation Conference*, 1920-1928.

Ramakrishnan S., S. Lee, and R. A. Wysk. 2002. Implementation of a Simulation-Based Control Architecture for Supply Chain Interactions. In *Proceedings of the 2002 Winter Simulation Conference*, 1667-1674.

AUTHOR BIOGRAPHY

DAVID BRIGGS is a business analyst in the Corus Long Products division. He has been with Corus and its predecessor companies for 18 years where he has worked on a wide range of Operational Research projects, though concentrating on Simulation and Scheduling applications. He is a member of the UK's Royal Statistical Society and Operational Research Society.