Maximizing profit by balancing throughput, quality, and efficiency in short-term production planning

Short-term production planning in the steel industry is still often carried out line by line, without considering the strong interdependence of melt shop, caster, and rolling mill. The real value, however, emerges when these processes are synchronized as one continuous flow. This article discusses how modern automatic scheduling systems achieve this integration by combining mixed-integer linear programming (MILP), heuristic rules, and machine learning-based adaptation. The model balances multiple objectives, throughput, quality, feasibility, and site-specific KPIs, while reacting dynamically to disruptions such as equipment downtime or changing campaign conditions. Real examples from steel plants show how reactive rescheduling and synchronized optimization can increase direct charging rates, reduce caster idle time,and save energy.

By transforming scheduling from manual decision-making to data-driven optimization, steel producers can move from reactive planning to proactive, profit-oriented control. The result is a measurable increase in throughput and stability, lower operational cost, and a significant step toward a fully digital and connected steel plant.

Author: Michael Peintinger PhD*, Smart Steel Technologies Inc.*

INTRODUCTION

Short-term scheduling in a steel mill is still too often handled process by process, with each line planned in isolation, as if the melt shop, caster, and rolling mill were independent systems. In reality, these operations are tightly interlinked and interdependent. A delay in the melt shop immediately disrupts casting, which in turn affects rolling. The true value lies in synchronizing all production steps into one continuous flow.

In modern mini mills (CSP, QSP, ESP), this synchronization is crucial: a single delayed ladle or missed casting window can cascade through the entire process chain. In integrated mills, additional complexity arises from managing slab yard operations and maintaining high direct-charging rates. Here, congestion in the rolling mill, or coil yard can force the caster to idle, wasting both energy and heat. Effective short-term production planning must therefore go beyond throughput optimization at individual units. It must coordinate all processes to maximize efficiency and profit while respecting operational KPIs and reacting dynamically to real-world disruptions across the entire steelmaking chain.

SCHEDULING PRIORITIES

The question arises: what should a modern automatic scheduling system truly focus on? *Figure 1* shows an anecdotal, yet representative, distribution of responses gathered from scheduling experts during sales discussions and consulting projects. It highlights the practical weighting that experienced planners assign to different objectives such as throughput, KPI adherence, feasibility, and reactive rescheduling.

While human schedulers have decades of experience creating executable programs, their “tribal knowledge” intuition for what constitutes a feasible and efficient schedule, must first be captured, documented and formalized within the model. Building an effective scheduling engine therefore begins with learning from expert practice: understanding the implicit rules, priorities, and constraints that define how a good schedule looks and behaves.

Once this knowledge is encoded, the system

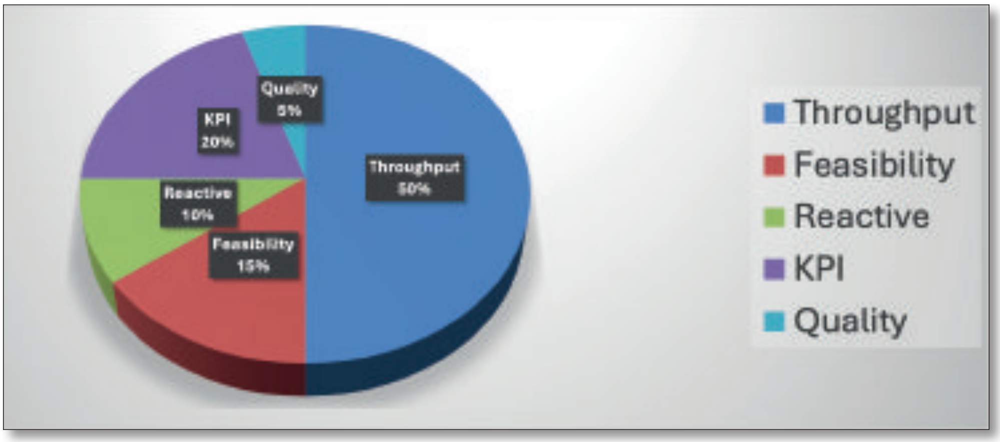

Fig 1 **Responses from experienced schedulers indicating which objectives a modern automatic scheduling system should prioritize**

must be able to balance four primary objectives simultaneously and dynamically adjust them depending on market conditions, resource availability, and process status. For instance, when a degasser is offline or congestion builds up at the galvanizing line, this information is often still communicated verbally over the phone. An intelligent scheduling system must therefore allow the user to guide, or restrict the model interactively, blocking certain process steps, or prioritizing others, through predefined templates.

Templates serve as scenario-based frameworks that reflect typical operating modes of the plant. In some cases, roughly half of the scheduling effort is directed toward maximizing throughput. In others, the focus shifts to protecting scarce grades or dimensions critical for width or thickness transitions (*Figure 2*), depending on plant configuration, product portfolio, or even seasonal mix. Additional templates may target site-specific KPIs such as optimizing entry temperature consistency, minimizing furnace energy use, or maintaining quality metrics within narrow statistical ranges.

Finally, there are special case templates that handle unique operational situations which defy standard logic but are vital for successful campaigns, for example, start-up sequences, end of campaign rolling, or handling leftover slabs and coils. Reactive rescheduling is often considered the most complex task and thus is frequently treated as a secondary, or unfeasible function. However, this capability generates the greatest financial impact. Preventing a single production program from halting, or avoiding the generation of geometrically inconsistent stock can save tens of thousands of dollars in direct and opportunity cost.

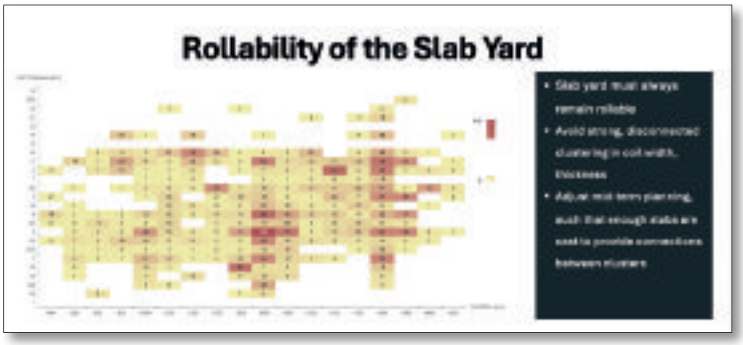

Fig 2 **The slab yard remains rollable as long as there are no gaps between the clusters**

MODEL ARCHITECTURE AND OPTIMIZATION ALGORITHMS

At the core of an advanced automatic scheduling system lies a mathematical model capable of representing the entire steelmaking and rolling process as a single, interconnected optimization problem (*Figure 3*). Each order, heat, or coil is characterized by a set of attributes, grade, thickness, width, target quality, and due date, and each process step imposes specific technological and capacity constraints.

To handle this complexity, the model is typically formulated as a mixed-integer linear programming (MILP) problem. The MILP structure ensures that discrete decisions, such as sequencing, batch grouping, or campaign length, are explicitly modeled, while continuous →

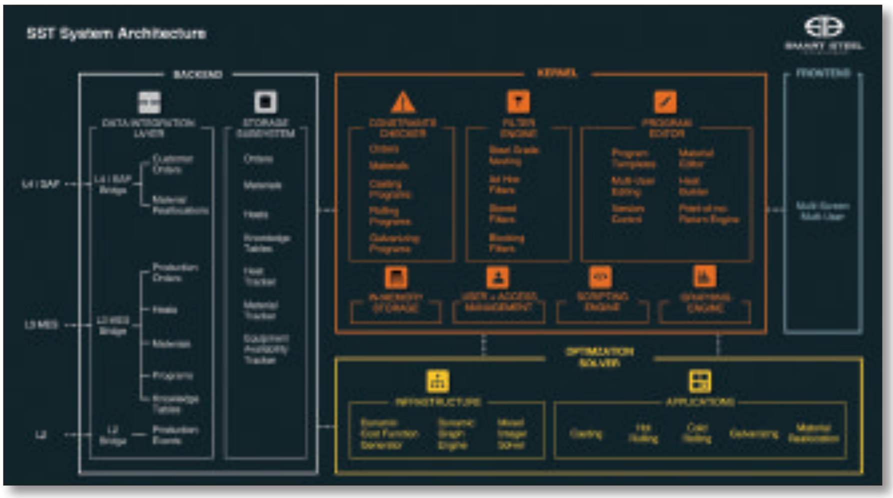

Fig 3 **SST system architecture**

parameters, such as process times or target temperatures, are optimized simultaneously. However, a pure MILP approach alone is often computationally intensive when applied to full-scale plant data. Therefore, modern schedulers employ hybrid architectures, combining optimization algorithms with heuristic and rule-based logic:

- Rule-based preprocessing filters out unfeasible or undesired combinations before optimization (e.g., forbidden grade transitions or width jumps). MILP optimization computes the best possible schedule within defined constraints and objectives, balancing throughput, delay penalties, and quality metrics.
- Post-optimization heuristics refine results to ensure plant-level executability, e.g., alignment with crane logistics, yard layout, or operator shift changes.

Machine learning can enhance this process by continuously tuning parameters, learning realistic process times, or predicting disruptions such as caster slowdowns, or rolling mill bottlenecks. The integration of data driven prediction with mathematically grounded optimization enables the scheduler to react quickly and efficiently to real-world conditions. This is a core requirement for short-term production planning.

OPERATIONAL CONSTRAINTS AND THE ROLE OF SYNCHRONIZATION

The rolling mill is often where bottlenecks are most visible. Furnaces typically run slower than the rolling mill and caster, making furnace output the limiting factor. Techniques such as hot charging, optimized charging blocks, and improved walking beam utilization can relieve this constraint. Many plants operate with uneven furnace setups, for example, several reheating furnaces with only a few walking beams and units frequently in maintenance, makes scheduling even more complex.

Program structure strongly influences mill efficiency. Too many small slabs reduce pacing, while short programs force frequent roll changes. Aggressive due dates can create congestion in the coil yard, and products requiring rework, such as trapezoidal slabs needing taper correction, must be carefully distributed across the schedule to avoid new bottlenecks. The rolling mill cannot be optimized in isolation. Without synchronization with the melt shop and caster, furnaces may sit empty, or slab deliveries may arrive late. When rolling programs are sequenced to accept slabs directly from the caster, energy is saved, heat loss is reduced and throughput increases substantially.

Feasibility remains a cross-cutting requirement. The scheduler's goal is to construct programs with minimal constraint violations while preserving scarce materials for special campaigns. Rare grades or input materials must be reserved even when their temporary use could simplify the plan. Likewise, some thicknesses, or product families cannot be processed for extended periods due to maintenance or tooling limitations. Special cases are common in modern product portfolios. Bake-hardening steels, ultra-low-carbon grades, ultra-thin gauges, dynamo steels, and surface-critical products often require dedicated sequencing and careful material preservation to integrate into the broader campaign logic. Managing these cases successfully is as important as throughput because such products typically carry the highest commercial, or technical value.

Site-specific KPIs account for a significant share of scheduling logic. These metrics may not directly increase capacity but stabilize operations. Entry temperature control is a prime example: smoother and more consistent entry conditions reduce mechanical stress on mill stands and improve product uniformity. Quality-related KPIs, even when secondary, can substantially reduce rework if transitions remain within narrow geometric or metallurgical bands. Balancing these objectives produces smoother flow from the melt shop to the rolling mill and prevents local instabilities from growing into plant-wide disruptions.

Reactive rescheduling completes this picture. No short-term schedule survives unchanged. Heats may be delayed, slabs misplaced, or caster sequences interrupted. Direct-charging programs are particularly vulnerable because they rely on precise timing between furnace, caster, and rolling. Once a program is live, changes must remain actionable and executable on the shop floor. Reactive rescheduling therefore creates disproportionate value: by reallocating orders, reassigning slabs, or modifying rolling programs in real time, schedulers can prevent program collapse. The financial impact is significant. Saving even a single rolling program from failure can avoid losses well above ten thousand dollars. Synchronization across melt shop, caster, and rolling enables the system to absorb disruptions rather than amplify them.

REACTIVE RESCHEDULING AND DYNAMIC OPTIMIZATION

In steelmaking, a ladle delay, degasser outage, or quality deviation at the caster can invalidate hours of pre-calculated production sequences. This is where reactive rescheduling becomes a decisive capability.

In a modern scheduling system, each process step continuously feeds real-time status information back into the optimization model. The scheduler identifies when deviations exceed defined tolerance bands, for example, when casting speed drops below a threshold or a hot-charging sequence risks interruption. When such events occur, the system automatically triggers a local re-optimization, recalculating only the affected portion of the program instead of regenerating the entire plan. Such deviations can also arise from the melt shop itself. In practice, the actual steel weight in a ladle often differs from the planned tonnage. If there is less steel than expected, the missing tonnage must be reassigned, typically by re-sequencing or merging subsequent heats to ensure campaign continuity. If there is excess steel, the system must quickly find compatible orders that can absorb the additional tonnage. This requires instant validation of grade, width, and thickness constraints and rapid transfer of new geometries to the Level 2 (L2) automation system as soon as the operator requests new orders for the surplus steel. By reacting automatically to such deviations, the scheduler preserves productivity and prevents quality or dimensional inconsistencies. The local re-optimization ensures all downstream processes such as caster, rolling mill, and finishing lines operate smoothly without manual intervention.

Reactive scheduling also enables profit-based decision-making in real time. When a disruption occurs, the system can simulate alternative scenarios within seconds, such as whether to delay the heat sequence, switch campaigns, or reassign coils to different finishing lines. Each option is evaluated not only by feasibility but also by its financial impact, allowing planners to select the most profitable recovery strategy.

By logging every rescheduling event and its outcome, the system learns patterns of disruption frequency, recovery cost, and constraint sensitivity. These data continuously refine the model's predictive accuracy and responsiveness.

CASE EXAMPLE AND OUTLOOK

In a typical flat-products mini mill, operating with continuous casting and hot rolling, the introduction of an automated, optimization-based scheduler reduced idle time in the caster. Direct charging rates increased as casting and rolling programs became better aligned. Improved furnace sequencing reduced specific energy consumption per ton, and more stable

→

rolling campaigns lowered roll wear and scrap rates. Together, these effects translated into significant financial gains.

In integrated plants, similar results are achieved by optimizing slab yard operations and increasing hot-charging utilization. Fewer reheating cycles and more consistent casting to rolling synchronization not only improve yield but also contribute to a lower carbon footprint, making the system an enabler for both economic and environmental performance.

The true value, however, lies beyond individual efficiency gains. Modern scheduling systems transform the decision-making process itself. They provide schedulers and operators with the ability to evaluate scenarios in real time, understand the financial impact of each alternative, and maintain stability across the entire production chain. This shift from reactive firefighting to proactive, model-based decision support marks a fundamental step toward the fully connected steel plant.

Looking ahead, the next stage of development will combine scheduling optimization with predictive analytics and self learning models. By integrating process data from Level 2 and quality systems, schedulers will evolve from rule based tools to adaptive systems that continuously refine their logic. This will further improve responsiveness to disruptions, optimize resource allocation, and maximize profitability.

In short, by balancing throughput, quality, and efficiency in real time, digital scheduling systems close the loop between planning and execution. They enable steel producers to make every ton count, every minute productive, and every schedule profitable. **MS**

Dr Michael Peintinger is Managing Director Americas at Smart Steel Technologies Inc., based in Cincinnati, USA

CONTACT: michael.peintinger@smart-steel-technologies.com

I SMART STEEL TECHNOLOGIES

SST Synchronized Scheduling

#STOP

CAST AND ROLL SCHEDULING

PING PONG

17 % Energy savings

23 % Less incidents

31 % Scheduling cost reduction

< 1 min Schedule generation time

SST MID-TERM PLANNER

Comprehensive supply chain planning solution. Raw materials supply, production planning, scenario creation.

SST SHORT-TERM SCHEDULER

Most advanced scheduling software. Automated generation of casting and rolling programs. Expert filtering, editing and scripting. Covers all constraints, KPIs and mill-specific rules.

SST LONG PRODUCT SCHEDULER

Centralized planning system for flexible production cycles. Covers diverse steel grades, profiles, finishing mills and cutting optimization.

SST MATERIAL ALLOCATOR

Automatic reallocation of unallocated slabs, bars and coils to new orders. Ranks according to grade, geometry, processing step. Maximizes revenue.

Smart Steel Technologies GmbH I Berlin, Germany
Smart Steel Technologies Inc. I Cincinnati, USA

www.smart-steel-technologies.com
request@smart-steel-technologies.com