



Discrete Optimization

Optimizing steel coil production schedules under continuous casting and hot rolling

Nelson Torres^a, Gus Greivel^b, Joshua Betz^c, Eduardo Moreno^a, Alexandra Newman^{b,*}, Brian Thomas^c

^a Faculty of Engineering and Science, Universidad Adolfo Ibáñez, 7941169 Peñalolen, Santiago, Chile

^b Colorado School of Mines, Operations Research with Engineering, Golden, CO 80401, USA

^c Colorado School of Mines, Department of Mechanical Engineering, Golden, CO 80401, USA

ARTICLE INFO

Keywords:

Integer programming
Large-scale optimization
Scheduling
Steel production
Continuous casting

ABSTRACT

In continuous steel casting operations, heats of molten steel are alloyed and refined in ladles, continuously cast and cut into slabs, and hot-rolled into coils. We present a mixed-integer program that produces a daily casting schedule and that is solved using state-of-the-art software for a 100% direct-charge steel mill; two casters concurrently produce slabs, which are rolled into coils at a single hot rolling mill. This model minimizes penalties incurred by violating plant best practices while strictly adhering to safety and logical constraints to manage risk associated with manufacturing incidents. An efficient formulation, combined with variable reduction and cutting planes, expedites solutions for small instances containing hundreds of variables and thousands of constraints by factors of at least two or three (and sometimes even 100); instances an order of magnitude larger along both problem dimensions suggest solutions that reduce costs incurred using plant best practices by as much as 40%.

1. Background

Producing almost two billion tons per year, modern steel plants create a wide range of products of which 97% is by steelmaking, continuous casting and rolling operations (WorldSteel Association, 2022). Each plant has its own considerations when scheduling customer orders, and, correspondingly, there is no industry-wide accepted optimization tool at the time of this writing. This paper investigates the scheduling of continuously cast coils produced in minimills from heats of molten steel. Thin slabs are direct-charged into a tunnel furnace and immediately hot rolled, i.e., not drawn from work-in-process inventory. We focus on the Nucor Steel facility in Decatur, Alabama, USA which runs a minimill operation with two continuous casters that directly feed a single hot rolling mill. A mixed-integer program creates schedules for both the heat and coil sequences.

1.1. Steelmaking

In conventional steelmaking, iron is produced from iron ore in a blast furnace, and transformed to steel in a basic-oxygen furnace or similar vessel. Minimill steelmaking operations recycle scrap (see, e.g., Rong & Lahdelma, 2008) purchased on a secondary market, which is melted in an electric furnace. In both conventional (integrated) and

minimill operations, the composition, or “grade” of the molten metal is created by alloying and ladle-refining. Fig. 1(a) shows an example of the resulting batch or “heat” of molten steel in its ladle, which holds about 160 tons, and contains a single, specific grade with a uniform composition. Each heat is continuously-cast and cut into slabs, which are rolled into coils. In addition to the steel grade, each customer order consists of the coil dimensions (thickness or “gauge,” width, and length) which define the “order weight.” Due to possible edge trimming, at additional cost, the final width may be narrower than the cast width; the larger cast width and weight must be used when defining the weight for the heat.

Because the molten metal becomes mixed during the continuous-casting process, the extent of the composition difference between the grades of prior and subsequent heats is important. The coils at the end of one heat and at the beginning of the next contain “intermixed” steel. The quantity of the steel, the severity of the intermix, and the ease with which the intermixed steel can be sold on a secondary market or stocked for later sale affect the quality of a given schedule. Models predict the number of tons that must be scrapped for a consecutive pair of steel grades, casting conditions, and customer tolerances which define the grade transition (Thomas, 1997).

* Corresponding author.

E-mail address: anewman@mines.edu (A. Newman).



(a) Ladle of molten steel tapped from electric arc furnace (Roy, 2018)



(b) Cutting a slab from the strand exiting a continuous-caster (Primetals Technologies, 2019)

Fig. 1. Steelmaking and casting processes (Primetals Technologies, 2019; Roy, 2018).

1.2. Continuous casting

Each ladle containing a heat of steel is transported to the continuous caster, where the molten metal flows through a holding vessel into a water-cooled, bottomless mold that creates the strand. Although the thickness is fixed for a given operation (3.6 inches at Nucor), the mold side-walls are adjustable during casting to produce the desired width. The solidified strand exiting the caster is cut into thin slabs, as shown in Fig. 1(b). At most minimills, including at Nucor, the slabs exiting each caster immediately enter a tunnel furnace, which homogenizes their temperature and delivers them to the hot rolling mill. This process is known as 100% “direct charging,” in which there is no slab inventory. This policy contrasts with most integrated steel plants, which can transport and store slabs, so that scheduling the heats produced in steelmaking and continuous-casting operations can be separated from scheduling the coils produced in the hot-rolling operation.

1.3. Hot rolling

Slabs exiting the tunnel furnaces are delivered to the hot rolling mill, alternating between the two casters at the Nucor minimill. After decreasing its thickness to about 1.18 inches via a roughing mill, each slab passes through a five-stand hot finishing mill, as illustrated in Fig. 2(a). Each roll stand decreases the thickness of the slab, now called “strip,” according to the compressive roll force applied by the pair of work rolls. Exiting the final roll stand, the strip thickness reflects the gauge ordered by the customer. The strip product is cooled on a runout table, and formed into coils to be sold or sent for further downstream processing.

As thickness decreases during each hot rolling stand, the strip speed should be increased with increasing length to maintain constant volume. Errors in the roll force applied to achieve the desired intermediate

gauge cause a mismatch in speed, which may result in strip shape defects or cause a catastrophic accumulation between two roll stands, known as a “cobble” (see Fig. 2(b)), which shuts down the entire operation and presents a major safety hazard. Because roll force is adjusted using a sophisticated artificial-intelligence system, errors (and associated cobbles) are more likely when the percent gauge change between consecutively cast coils is larger, or the strip is thinner.

As slabs are rolled into coils, the frictional forces gradually wear down the rolls. The strip edges tend to cause slight grooves, which generate surface defects in subsequent coils. To avoid this problem, the width of consecutive slabs should not increase. Thus, slabs are generally cast from wide gradually to narrow, and also from thick gradually to thin. Eventually, the worn surfaces of the rolls cause too many surface quality problems on the product, so the rolls are changed, after which the width and gauge are “reset,” i.e., coils can be wide and thick again. Although the last roll stand usually experiences higher forces and wears down first, operators usually take advantage of the down time to simultaneously change all of the rolls. We term this a “roller change,” where roller refers to all of the rolls in the five stands together. Our optimization model enforces operational rules (constraints) designed to lower the risk of cobbles and quality problems.

1.4. Downstream processing

Many coils are treated with downstream processing operations to meet a wide range of possible product requirements. These operations include “pickling” with acid to remove scale and rust; cold rolling to further decrease gauge and increase strength; galvanizing, or coating the coil with zinc to improve corrosion resistance; and trimming, to decrease width and sharpen the edges. The only downstream process considered in our scheduling model is trimming, which provides an



(a) A steel slab passes through several rolling stands in the hot-rolling process (Board & Vellum, 2015)



(b) A cobble resulting from a mismatch in roll force, gauge, and speed between rolling stands (Industrial History, 2021)

Fig. 2. Steel Rolling Process (Board & Vellum, 2015; Industrial History, 2021).

opportunity to schedule the “cast width” of a slab to be wider than the final “order width” of the coil required by the customer. This is often done to accommodate other scheduling constraints, such as rolling from wide to narrow, creating full heats, minimizing intermix between grades, and avoiding cobble risk. However, trimming also incurs additional cost. Sierra-Paradinas et al. (2021) present a mixed-integer program to address cutting to suit customer orders. Another method to accommodate scheduling constraints is to insert a “stock coil” (associated with no immediate customer) into the schedule. Such coils incur an opportunity cost of foregoing an ordered coil, and an inventory storage cost. All other coils are shipped immediately to fulfill a customer order (see Fig. 3).



Fig. 3. Finished coils processed for shipping from Nucor Steel (Nucor, 2022).

1.5. Example process summary

Our scheduling optimization model focuses on operations at the Nucor Steel mill located in Decatur, AL. This minimill features two electric-arc furnaces that melt scrap and alloys into ladles, each containing a heat of 150–170 tons of molten steel with a specific grade. A heat is refined in one of two liquid metallurgical facilities that serve two continuous casters operating in parallel. All slabs produced on each caster are direct-charged into their respective tunnel furnaces and delivered to the hot rolling mill. While both casters are operational, the hot rolling mill typically alternates between the two casters from which it receives coils, automatically adjusting its rolling parameters based on the previous coil processed by the caster from which the current coil arrives. This enables significant alternating differences between successive coils, according to the slab's originating caster. Regardless of the originating caster, successive coils processed through the rolling mill should be relatively similar in width. After decreasing the slab's thickness to its final gauge in the hot rolling mill, and coiling it, the product may be sent for downstream processing such as trimming, or delivered to the customer. Fig. 4 depicts the complete process of Nucor's operation.




Fig. 4. Complete casting process at Nucor Steel Decatur, Alabama, USA.

2. Literature review and contribution

Most steel manufacturing operations lack sufficiently detailed software to make their scheduling decisions; as a result, experienced operators are left to generate manual schedules (Graves, 1981; Maccarthy & Liu, 1993). Steelmaking, and, in fact, various facets in the industry, possess specific rules related to creating good schedules (Lee et al., 1996). As optimization modeling has become more prevalent and solvers have rendered instances more tractable, some steel companies have formulated formal optimization models, thereby increasing productivity, efficiency and throughput.

Tang et al. (2000a) schedule just the coils through the hot rolling mill at the Shanghai Baoshan Iron & Steel Complex using a genetic algorithm. Tang et al. (2000b) determine the time at which to schedule a batch (or ladle's worth of steel) using a linearized model originally posed as a nonlinear program; the detailed, discrete decisions we consider are absent, but timing considerations and temperature levels are emphasized. Mao et al. (2014) add discrete decisions, casting the model in a hybrid flowshop framework, but omit considerations specific to an order such as the geometry and chemistry of a coil. Pan (2016) relaxes the former authors' assumption that the sequence of batches is determined prior to optimization, but do not consider detailed coil scheduling.

A notable example that considers a greater level of detail is Baosteel, a company that uses a mixed-integer program to schedule heats into sequences (Tang, Meng et al., 2014; Tang & Wang, 2008). The authors pose a p -median capacitated model to allocate slabs to batches, and a general, multi-objective mixed-integer program to sequence heats. Because of their intractability, the models are solved using a dynamic-programming heuristic and tabu search, respectively. These particular Baosteel facilities, however, decouple the casting and hot rolling processes by moving cast products to a storage facility from which the hot rolling mill draws to rehear and roll. This type of *inventory-based rolling mill* allows casting and rolling constraints to be considered independently. Specifically, the casting schedule only incorporates a coil's steel grade and casting width when assigning it to a heat (Tang et al., 2011), resulting in model instances with thousands (versus tens or hundreds of thousands) of variables and constraints, rendering the resulting computation tractable with reformulation (in this case, as a set-packing model that lends itself to column generation) and state-of-the-art solvers (Tang, Wang et al., 2014). This advantage in scheduling, however, comes at the cost of losing the energy efficiency associated with a homogenizing furnace (rather than a reheating furnace) and the extra costs associated with transporting, storing and managing slab inventory.

A steel manufacturing facility in Austria has also adopted a mixed-integer program for its scheduling; however, it too can place cast slabs into inventory before sending them to the hot mill. Their approach seeks to minimize the sum of penalties associated with down time incurred at the caster, changes in casting speed throughout a sequence, scrap produced, and missed due dates (Missbauer et al., 2009). Downstream of casting, ArcelorMittal's Dofasco facility has used mixed-integer programming to schedule its hot rolling mill, which is

also decoupled from the caster (Lopez et al., 1998). The modeling approach is based on the traveling salesman problem in which each coil represents a node, and there exists an arc between any two coils that can be rolled consecutively. The cost associated with this arc is a function of the difference in properties of the coil, which consist of gauge, width, and hardness. Whereas in the classical formulation of this problem, each node must be visited exactly once and subtours are precluded, in their application, subtours (containing a single node each) represent uncast coils, and the number of such subtours is constrained.

Tan et al. (2013) examine scheduling over the combination of the casting and rolling processes. They assume that a slab yard can hold an infinite quantity of work-in-process before it is sent to the rolling mill. In this way, the processes are completely decoupled, and the sequence of different geometries does not need to be coordinated with the various steel compositions of the continuous-cast slabs. The objective maximizes the amount of throughput, and the authors use a decomposition strategy that combines math- and constraint programming to solve instances of their model with hundreds of slabs. Mattik et al. (2014) integrate these models by first posing one for the continuous casting process to produce a schedule with the objective of minimizing total slab production time (makespan) while considering the chemical compatibility of the steel grades. A second model for the rolling process minimizes waiting times between continuous casting and rolling. The aggregate model quickly returns near-optimal solutions with the benefit of reducing energy consumption, but lacks the detailed geometric considerations we incorporate for the sake of our rolling process.

Silaen et al. (2017) report on ArcelorMittal-USA's efforts to assign heats to a caster that integrates casting and hot rolling. Their solution methodology separates the model into three sub-problems (i.e., order selection, assignment of orders to casters, and scheduling the orders on the casters), the first two of which are solved using a heuristic while the remaining sub-problem is solved to optimality. In essence, sets of “good” heats are first determined, and then the heats are grouped into schedules. Because decisions (upon which the quality of the solution to the previous sub-problem depends) are heuristically determined, global optimality cannot be proven; however, a near-optimal solution to the monolith may be quickly constituted from the solutions of several sub-problems (Harjunkoski & Grossmann, 2001).

Riccardi et al. (2015) consider environmental incentives for production both in integrated facilities and in minimills, but no detailed production scheduling models are determined. At a systems level, Ray and Kim (1995) analyze costs in the steel industry. Regarding distribution, Rosyidi et al. (2021) examine a supply chain problem in the steel industry consisting of the integration of supplier-side and buyer-side decisions. Their model maximizes profit while adhering to constraints on capacities, supplies, demands, and inventory balance. No detailed operational decisions at the scheduling stage are considered, however. Similarly, Fukuyama et al. (2021) present a high-level survey regarding the utilization of iron and steelmaking facilities in China. A complete review of scheduling models related to integrated steel production can be found in Tang et al. (2001), and Özgür et al. (2021) survey planning and scheduling models for hot rolling mills.

This paper presents the first *coil-based model* of steelmaking, continuous casting and rolling with 100% direct-charge operations. By contrast, *heat-based models*, which dominate the literature at the time of this writing, lack the ability to effectively intertwine the macro-level, grade-based decisions for each heat that are important at the caster, and the various geometric decisions at the coil level that the rollers must consider. In fact, we show in Section 5.2 that trying to optimize these processes separately, while fast, leads to suboptimal, or even infeasible, solutions. Our contributions, therefore, lie in the development of methodologies to expedite solutions to this more complex *coil-based model* in which both chemistry-related decisions at the caster and geometric decisions for the rollers must be made in tandem. Section 3 introduces a model to determine the order in which coils are cast such that they constitute valid heats, consider intermix between

grades, and conform to rules regarding geometries of consecutive and adjacent coils. In Section 4, we introduce several techniques to improve model tractability, including: (i) variable elimination and efficient formulation, (ii) deployment of appropriate big-M values, and (iii) cutting planes. Section 5 reports numerical results regarding solution time and quality, and compares schedules obtained with our model, a heuristic, and two decoupled approaches for Nucor Steel's operations. Section 6 concludes.

3. Coil-based model formulation

To construct a model sufficiently general to schedule both steel-making and casting in conjunction with hot rolling in a 100%-direct-charge-facility, sequencing rules between adjacent coils must hold; this requires a coil-based model, in addition to heat-based constraint satisfaction enforced in previous heat-based models. Each sequence is partitioned into separate rolling campaigns, and each time a roller stand change occurs, we assume that all roll stands are changed to reset the capacity of the finishing mill, and therefore the geometric coil constraints. Relationships for defining the order of production of a sequence of coils are established by segmenting it into notional, spatial “slots” – equal to half of the number of coils, in which coils, given an arbitrary index number, can be placed. A slot contains two coils, each cast at the same time on one of the two casters to allow for the comparison of parameter values of adjacently cast coils (see, e.g., Pinto & Grossmann, 1998 for this framework). In this section, we mathematically develop the scheduling model for the process described in Section 1.

Before progressing to the full model formulation, we present a simple example with ten coils and two casters to illustrate the complexity of this scheduling problem. Although adjacently cast coils may be of different grades as prescribed by the casting process, these same coils must adhere to the geometric rules enforced in the rolling process; this follows from the fact that the mill contains two independently-run casters, but a single set of rollers onto which both casters alternate depositing heats of in-progress steel. Therefore, between roller changes, coil widths and gauges are expected to be monotonically non-increasing and should include no exceptionally large decreases. Consider a “simple” order of ten coils, consisting of three grades (see Table 1).

Rules in the casting process dictate that we cannot switch grades within a heat on either caster. Moreover, switching grades between heats, if even feasible, incurs a cost that depends on the two adjacent grades. Because our example involves three grades, with four coils of one grade and three coils each of the other two, at least one grade change must occur on each of the casters in order to fill the five slots with coils and construct a ten-coil schedule. Based on the interaction between the two casters and single roller, the coils exiting each caster must merge into a feasible rolling sequence (based on geometric constraints), with the ability to reset the geometric constraints (i.e., increase width and/or gauge) after a roller change. Coils in our small example (see Fig. 5) are labeled with their coil number from Table 1; grades are indicated by color (grade 1 = blue; grade 2 =

Table 1
Small ten-coil example.

Coil	Grade	Width (in)	Gauge (in)
1	1	50	0.08
2	1	47	0.06
3	1	49	0.08
4	1	48	0.06
5	2	48	0.06
6	2	49	0.08
7	2	49	0.08
8	3	48	0.06
9	3	49	0.08
10	3	49	0.08




Fig. 5. Visualization of the small 10-coil, 2-caster, 3-grade example as given spatially; note that time is irrelevant for the purposes of the schedules we provide.

Table 2

Sets	Description
$\mathcal{C}$	coils
$\mathcal{G}$	grades
$\mathcal{K} = \{1, 2\}$	casters
$S = \{1, \dots, \mathcal{C} /2\}$	slots on each caster
$\mathcal{C}^{mill}, \mathcal{C}^{HRB}, \mathcal{C}^{cut}$	subset of coils with edge-type <i>mill</i> , <i>HRB</i> , and <i>cut</i> , respectively
$\hat{\mathcal{C}}$	subset of coils for which a roller change could occur

red; grade 3 = violet); order width is indicated by the base width of the slabs; gauge is indicated by the height of the slabs; grade changes in the casters and roller changes in the rolling mill are indicated by vertical dashed lines; and, sequencing from the two casters into the rolling schedule is indicated by the arrows. Our depiction provides a relative ordering between the coils, but omits the element of time, because the mill operators regard the problem spatially; indeed, the operation may be sped up or slowed down depending on last-minute changes required in the properties of the steel (outside the scope of our work). Although Fig. 5 illustrates a feasible schedule (where feasibility is nontrivial based on the need to adhere to several different geometric sequencing rules), it incurs seven grade changes among ten coils. Although not obvious in this small example, a reorganization of coils in larger schedules might reduce the number of grade changes while still satisfying constraints related to geometrical considerations. If not, the mill is willing to incur *trim loss* at a penalty by casting coils wider than their order width to preserve the monotonically-non-increasing coil width rule. (Note that this same rule cannot be applied to gauge and that, when applied to width, the exact implementation depends on three different width flexibility specifications from the customer.) The consideration of these various chemical and geometric rules, combined with the flexibility for the rules to be broken with a roller change and customer width specifications, renders this problem exceedingly complex, even for a modest number, e.g., 40, of coils in a rolling campaign.

The level of detail in our scheduling problem requires the introduction of the following notation: (i) sets (in Table 2) categorize steel grades and edge requirements specified in customer orders, casters and rollers in the mill, and slots available to be scheduled; and, (ii) parameters (in Table 3) prescribe the characteristics of each customer order such as the finished coil grade, order width, gauge, and length (which together define the dependent parameter coil weight according to the geometry and steel density). Finally, Table 4 introduces continuous, non-negative variables (in black font) describing coil width and gauge characteristics, and binary variables (in red font) defining coil sequencing and coil order characteristics.

In summary, we consider a set of $|\mathcal{C}|$ coils, $|\mathcal{G}|$ grades, and $|\mathcal{K}|$ casters, where each caster contains $|S|$ slots. Each coil, as a result of customer order specifications, possesses a desired grade, gauge (thickness), and width, where the latter characteristic is complicated by the fact that the order width may differ from the cast width based on a

Table 3

Parameter	Description	Units
Width-Based		
w_c	order width of coil c	[in]
$\hat{w}_c$	allowed width excess for each coil	[in]
$\hat{w}$	extra allowed width excess for cut-edge coils	[in]
$\underline{w}$	additional minimum width excess for mill-edge coils when using alternative bounds	[in]
$\bar{w}$	additional maximum width excess for mill-edge coils when using alternative bounds	[in]
w_c^U	upper limit on the cast width of coil c	[in]
$\bar{w}$	allowed inter-caster width difference without penalty	[in]
Degradation-Based		
d_c	degradation of the roller associated with coil c	[fraction]
$\bar{d}$	maximum degradation allowed on a roller	$[\leq 1]$
Heat-Based		
$\bar{T}, \underline{T}$	upper and lower weight limits for a heat, respectively	[tons]
t_c	weight of coil c	[tons]
Chemistry-Based		
$\gamma_{gg'}$	penalty cost of a transition from grade g to g'	[\$]
$g_c \in \mathcal{G}$	grade of coil c	[–]
Gauge-Based		
a_c	gauge of coil c	[in]
$\underline{a}_c$	limit on minimum gauge for the subsequently cast coil without a penalty, given by the following formula:	[in]
$\underline{a}_c = \begin{cases} 0.5 \cdot a_c, & a_c \geq 0.4 \\ 0.75 \cdot a_c, & 0.4 > a_c \geq 0.123 \\ 0.9 \cdot a_c, & a_c < 0.123 \end{cases}$		

Table 4

Variable	Description	Units
Continuous Non-negative Variables:		
D_s	cumulative degradation of current roller by slot s	[%]
T_{ks}	cumulative tonnage of current heat by slot s on caster k	[tons]
W_{ks}	cast width of the coil assigned to caster k slot s	[in]
Binary Variables:		
B_c	1 if a mill-edge coil c uses its alternative cast-width bounds, 0 otherwise	
G_{ks}	1 if a grade change occurs between the coils assigned to caster k slot s and $s+1$, 0 otherwise	
H_{ks}	1 if there is a heat change on caster k in slot s , 0 otherwise	
X_{ksc}	1 if coils c is assigned to caster k in slot s , 0 otherwise	
Y_s	1 if there is a roll change at slot s , 0 otherwise	

“trim condition,” which includes three varieties: (a) “cut edge” coils, in which the coil must be trimmed, (b) “HRB” (or hot-rolled-black) coils, which are hot-rolled and shipped with as-cast black edges, and

(c) “mill edge” coils for which trimming can be chosen, as desired, to satisfy other scheduling constraints. At a system level, rollers degrade as a function of coils passing through them, and we track this, and the resulting change-outs. Additionally, each heat must adhere to a minimum and maximum tonnage. The quality of a schedule is characterized by the process dynamics: (a) operational limitations, such as ladle and roller capacity, (b) product quality constraints, such as those involving width variations between consecutively cast coils on the same caster, and adjacently cast coils between the two casters, (c) penalties associated with grade intermix and trim loss, and (d) penalties associated with risk, such as excessive gauge drops that could lead to rolling-mill cobbles.

We now state the full model, which we term (M), in which the objective is implicit. That is, we minimize the weighted sum of violations of each soft constraint, i.e., constraint that can be violated using a continuous-valued *elastic* variable, where these are denoted by a dot over the relational operator. Ultimately, the objective function consists of a weighted sum of these elastic variables, where the relative weights are given in Table 7 of Section 5. Any constraints containing a variable with an index of $s - 1$ hold only for the case in which the slot number is strictly greater than 1.

(M) subject to:

(see Section 3.1: Valid Assignment)

$$\sum_{k \in \mathcal{K}} \sum_{s \in \mathcal{S}} X_{ksc} = 1 \quad \forall c \in \mathcal{C} \quad (1a)$$

$$\sum_{c \in \mathcal{C}} X_{ksc} = 1 \quad \forall k \in \mathcal{K}, s \in \mathcal{S} \quad (1b)$$

(see Section 3.2: Cast Width)

$$W_{ks} - \hat{w} \leq W_{k,s+1} \leq W_{ks} \quad \forall k \in \mathcal{K}, s \in \mathcal{S} : s < |\mathcal{S}| \quad (2a)$$

$$|W_{1s} - W_{2s}| \leq \bar{w} \quad \forall s \in \mathcal{S} \quad (2b)$$

$$\sum_{c \in \mathcal{C}} w_c \cdot X_{ksc} \leq W_{ks} \leq \sum_{c \in \mathcal{C}} w_c^U \cdot X_{ksc} \quad \forall k \in \mathcal{K}, s \in \mathcal{S} \quad (2c)$$

$$w_c \cdot X_{ksc} + \underline{w} \cdot B_c \leq W_{ks} \leq w_c \cdot X_{ksc} + \bar{w} \cdot B_c + \hat{w} + M(1 - X_{ksc}) \quad \forall k \in \mathcal{K}, s \in \mathcal{S}, c \in \mathcal{C}^{mill} \quad (2d)$$

(see Section 3.3: Roller Changes)

$$Y_1 = 1 \quad (3a)$$

$$Y_s \leq \sum_{c \in \mathcal{C}} X_{ksc} \quad \forall k \in \mathcal{K}, s \in \mathcal{S} \quad (3b)$$

$$\sum_{s \in \mathcal{S}} Y_s \leq 1 \quad (3c)$$

$$D_s \geq \max\{D_{s-1} + \sum_{k \in \mathcal{K}} \sum_{c \in \mathcal{C}} d_c \cdot X_{ksc} - M \cdot Y_s, \sum_{k \in \mathcal{K}} \sum_{c \in \mathcal{C}} d_c \cdot X_{ksc}\} \quad \forall s \in \mathcal{S} \quad (3d)$$

(see Section 3.4: Heats)

$$T_{k,1} = \sum_{c \in \mathcal{C}} t_c \cdot X_{k,1,c} \quad \forall k \in \mathcal{K} \quad (4a)$$

$$\sum_{c \in \mathcal{C}} t_c \cdot X_{ksc} + T_{k,s-1} - M \cdot H_{ks} \leq T_{ks} \leq$$

$$\sum_{c \in \mathcal{C}} t_c \cdot X_{ksc} + \min\{T_{k,s-1}, M \cdot (1 - H_{ks})\} \quad \forall k \in \mathcal{K}, s \in \mathcal{S} \quad (4b)$$

$$\underline{T} \cdot H_{ks} \leq T_{k,s-1} \leq \bar{T} + (T_{ks} - \sum_{c \in \mathcal{C}} t_c \cdot X_{ksc}) \quad \forall k \in \mathcal{K}, s \in \mathcal{S} \quad (4c)$$

(see Section 3.5: Grade Changes)

$$\gamma_{g_1, g_2} \cdot \left(\sum_{c \in \mathcal{C}} X_{ksc} + \sum_{\substack{c' \in \mathcal{C} \\ g_{c'} = g_2}} X_{k,s+1,c'} - 1 \right) \doteq 0 \quad \forall k \in \mathcal{K}, s \in \mathcal{S}; g_1, g_2 \in \mathcal{G} : g_1 \neq g_2 \quad (5a)$$

$$G_{ks} \geq \sum_{\substack{c \in \mathcal{C} \\ g_c = g}} X_{ksc} + \sum_{\substack{c' \in \mathcal{C} \\ g_{c'} \neq g}} X_{k,s+1,c'} - 1 \quad \forall k \in \mathcal{K}, s \in \mathcal{S}; g \in \mathcal{G} \quad (5b)$$

$$G_{ks} \leq H_{ks} \quad \forall k \in \mathcal{K}, s \in \mathcal{S} \quad (5c)$$

$$T_{ks} \geq \sum_{c \in \mathcal{C}} t_c \cdot X_{ksc} - M \cdot (1 - G_{ks}) \quad \forall k \in \mathcal{K}, s \in \mathcal{S} \quad (5d)$$

(see Section 3.6: Excessive Gauge Decrease)

$$\sum_{c \in \mathcal{C}} a_c \cdot X_{ksc} - \sum_{c' \in \mathcal{C}} a_{c'} \cdot X_{k,s+1,c'} \doteq 0 \quad \forall k \in \mathcal{K}, s \in \mathcal{S} \quad (6)$$

(see Section 3.6.1: Variable Bounds)

$$0 \leq D_s \leq \bar{d} \quad \forall s \in \mathcal{S} \quad (7a)$$

$$0 \leq T_{ks} \leq \bar{T} \quad \forall k \in \mathcal{K}, s \in \mathcal{S} \quad (7b)$$

$$W_{ks} \geq 0 \quad \forall k \in \mathcal{K}, s \in \mathcal{S} \quad (7c)$$

$$B_c, G_{ks}, H_{ks}, X_{ksc}, Y_s \text{ binary} \quad \forall c \in \mathcal{C}, k \in \mathcal{K}, s \in \mathcal{S} \quad (7d)$$

We attempt as compact a formulation as possible, recognizing that both the number of slots and the number of coils can be large for realistic instances and, therefore, any binary variable containing both of these instances is likely to contribute significantly to the intractability of the model. A natural formulation might use binaries that track adjacent coils and/or that consider roller-slot-coil-caster combinations, but we seek to avoid such formulations. Instead, we do not include a roller index on any variable; rather, we implicitly account for “badness” between any two consecutively cast coils using continuous variables and low-indexed binaries.

3.1. Valid assignment

Constraint (1a) ensures that each coil is cast on exactly one caster and appears in exactly one slot. Constraint (1b) guarantees that each caster-slot pair possesses a coil, which assumes that both casters are operational when the schedule is implemented.

3.2. Cast width

These constraints relate to the characteristics of a coil within consecutive slots (such as a decrease in width with increasing slot number) and to coil characteristics (such as lower and upper bounds on the cast and order width). We use a continuous variable to monitor the width of a coil in a slot on a given caster by simply tracking the width of a caster-slot pair (without the coil index). Each coil $c \in \mathcal{C}$ has an *order width* w_c . All coils, except cut edge, can be cast within the standard range $[w_c, w_c + \hat{w}]$. Coils requiring a cut-edge must be cast with a relaxed upper bound on width of $w_c + \hat{w} + \bar{w}$. By contrast, mill-edge coils can be either within the standard range $[w_c, w_c + \hat{w}]$ or within an alternative range $[w_c + \underline{w}, w_c + \hat{w} + \bar{w}]$, which includes the minimum trim of $\bar{w}$ required by the guide rolls during the trimming operation. (Hot-rolled-black edge coils must be cast strictly within the American Society for Testing and Materials standards, which allow for final width to increase relative to order width by up to $\hat{w}$.)

We impose the following two relationships between casters and slots: (i) constraint (2a) ensures that, for a given caster and slot, the cast width in that slot must be between a maximum allowable width decrease and the width contained in the previous slot; and, (ii) constraint (2b) requires that for the same slot on the two casters (i.e., adjacent slots) the width difference not exceed a given tolerance. We linearize this latter constraint in the implementation, and allow it to be violated (up to a limit) with a penalty. Constraint (2c) relates the width for a slot-caster combination to the product of the binary representing whether or not a specific coil is cast in that slot and the parameter representing the width of a coil. The lower limit is elasticized

at a penalty; the upper limit enforces a hard constraint where the coefficient is given by:

$$w_c^U = \begin{cases} w_c + \hat{w} & c \in \mathcal{C}^{HRB} \\ w_c + \hat{w} + \hat{w} & c \in \mathcal{C}^{cut} \\ w_c + \hat{w} + \bar{w} & c \in \mathcal{C}^{mill} \end{cases}$$

The cast width of a mill-edge coil must be between a lower limit $\underline{w}$ and an upper limit $\bar{w}$ according to constraint (2d), as long as a given coil is assigned to that caster-slot location. The number of constraints considering both inequalities in (2d) is $\mathcal{O}(|\mathcal{C}|^2)$, which produces a large number of constraints if most coils are mill edge.

3.3. Roller changes

Constraint (3a) enforces the initial condition that there is a roller change in the first slot. Constraint (3b) allows a roller change only in slots where coils that belong to the set $\hat{\mathcal{C}}$ are placed. Elastic constraint (3c) counts each roller change after the first and penalizes it accordingly. For a degradation model that depends solely on the number of coils passed through the rollers, this is a constant. For an accumulated degradation model that depends on the sequence of coils, this number can vary (Udofia et al., 2023). Constraint (3d) records degradation for a given slot as a function of the coil in it across both casters; if there is no roller change in that slot (i.e., $Y_s = 0$), roller degradation is accumulated. However, if there is a roller change, the first argument of the maximization function is void and the roller degradation is determined only by the coil in the slot following the change, given as the second argument.

3.4. Heats

Constraint (4a) records the weight (tonnage) of the coil in slot 1 for each caster. With no grade change, for all slots except the first, constraint (4b) equates aggregate tonnage on a given caster with the sum of that accumulated in the previous slots and the contribution from the currently considered slot; otherwise, it imposes an upper bound which is then reset to the tonnage from the currently considered slot. The tonnage of a given grade of steel at heat boundaries is accounted for in constraint (4c). Specifically, the lower limit forces a minimum tonnage of a given grade before a heat change occurs to preclude underweight heats. The upper limit accounts for “spillover” between heats of the same grade; overcast tonnage of a given heat in slot s can be offset if there exists capacity in the previous slot; this part of the constraint is obviated when a grade change occurs.

3.5. Grade changes

Computationally expensive constraints (5a) record a grade change for coils with differing grades cast in adjacent slots. Constraint (5b) counts grade changes on each caster and in each slot as a function of the coils cast. Constraint (5c) forces a heat change with a grade change. Constraint (5d) overrides (4c), which, together with (4b), imply that, post grade change, the heat tonnage in the slot matches the tonnage of its assigned coils, precluding spillover from the previous heat.

3.6. Excessive gauge decrease

Constraint (6) penalizes excessive gauge decrease. There is no penalty on gauge increase between coils in consecutive slots.

3.6.1. Variable declarations

Finally, we require that the domains of our model variables be restricted appropriately.

4. Methodology

We introduce several techniques to improve the tractability of (M), including: (i) variable elimination and efficient formulation, (ii) the deployment of appropriate “big-M” values, and (iii) cutting planes. An initial feasible solution provided prior to the initiation of the branch-and-bound algorithm can result in tractability gains (Klotz & Newman, 2013). Allen et al. (2022) have constructed such a solution using a heuristic approach. However, this solution did not universally aid the branch-and-bound search, despite resulting in a feasible solution for the monolith. Other attempts based on a fix-and-relax heuristic (Pochet & Wolsey, 2006) did not necessarily yield even a feasible solution. Therefore, our numerical experiments do not warm-start the solver.

4.1. Variable elimination

A natural sequencing formulation would explicitly track not only the slot and caster on which each coil is cast, but also the adjacent coil in that location. In this way, it is easy to account for the difference in geometric, as well as composition, processing, and operational risk characteristics that induce penalties. However, the corresponding variable would result in $\mathcal{O}(|\mathcal{C}|^2|\mathcal{S}||\mathcal{K}|)$ binaries, which would populate most constraints. Additionally, it would be more natural to account for the roller on which a coil is cast, in a given slot and on a given caster. This would result in $\mathcal{O}(|\mathcal{C}||\mathcal{R}||\mathcal{S}||\mathcal{K}|)$ variables, where $\mathcal{R}$ defines the set of rollers. For even a 100-coil model with two casters and several rollers, the “more natural” formulation would result in millions of variables and constraints. We therefore: (i) eliminate the roller index altogether and explicitly track degradation; (ii) remove consecutive coil articulation from the indices and implicitly account for “badness” between adjacent coils using continuous variables, and (iii) retain a minimal set of binaries that account for roller, heat and grade changes between slots (where the former is only a function of slot and not of caster), differences in cast- versus order-width, and the location of a coil in a caster and slot. We can further reduce the number of variables using so-called “indexed sets” that a priori preclude: (i) grade changes in certain, e.g., “early” or “late,” slots; (ii) heat changes between certain slots, e.g., when there are no grade changes; (iii) roll changes early or late in the campaign; and, (iv) coil assignment to certain slots based on factory rules.

Specifically, regarding plant rules, experience dictates that the highest quality products are cast after the first, and before the last, few heats of a sequence. The head and the tail of a cast sequence have higher defect rates based on steel cleanliness studies (Zhang & Thomas, 2003). Therefore, coils with lower quality and cleanliness requirements are scheduled early and late in the sequence (Luecking & Jannasch, 2008). At Nucor, categories in which coils can be placed are dictated by the quality requirements that an end-user specifies for a particular coil.

4.2. Big-M determination

Based on physical limitations of the casting and rolling processes, as well as American Society for Testing and Materials standards, we are able to tune big-M values as given in Table 5.

Table 5
Values for M.

Constraint number	Value for M	Rationale
(2d)	w_c	Maximum coil width when $X_{ksc} = 0$ as given by the expression for w_c^U in Section 3.2
(3d)	1	Maximum degradation on a roller before changing
(4b)	$\bar{T}$	Maximum tonnage in a heat
(5d)	$\bar{T}$	Maximum tonnage in a heat

4.3. Cutting planes

Finally, we derive valid inequalities to strengthen the model, necessitating the introduction of additional notation (shown in Table 6).

Table 6

Sets, parameters and variables for cutting planes.

Object	Description	Units
Sets		
$\bar{\mathcal{G}}$	subset of grades for which we compute a minimum grade transition cost	
Ω	set of potential values for arbitrary widths used in clique constraints	
$\mathcal{C}_\omega, \bar{\mathcal{C}}_\omega$	sets of coils for which ω is an upper or lower bound on the cast width limits, respectively	
Parameters		
$\gamma_{\bar{\mathcal{G}}}$	minimum grade transition cost over the subset $\bar{\mathcal{G}}$	[\$]
ω	arbitrary width used in clique constraints	[in]
Variables		
Γ_{ks}	grade change cost for slot s on caster k	[\$]
$\bar{G}_{ksg_1g_2}$	1 if a grade change occurs between the coils assigned to caster k slot s and $s+1$ from grade g_1 to grade g_2 , 0 otherwise	
Z_{kg}	1 if a coil of grade g is assigned to caster k , 0 otherwise	
P_s^ω	1 if any coil in $\bar{\mathcal{C}}_\omega$ is scheduled in slot s , 0 otherwise	

4.3.1. Trivial bounds

$$\sum_{s \in S} Y_s \geq \left\lceil \frac{\sum_{c \in \mathcal{C}} d_c}{\bar{d}} \right\rceil \quad (8)$$

$$\sum_{k \in \mathcal{K}} \sum_{s \in S} G_{ks} \geq |\mathcal{G}| - |\mathcal{K}| \quad (9)$$

$$\sum_{s \in S} H_{ks} \geq \left\lceil \frac{\sum_{i=1}^{|S|} t_{(i)}}{\bar{T}} \right\rceil - 1 \quad \forall k \in \mathcal{K} \quad (10)$$

$$\sum_{s \in S} H_{ks} \leq \left\lceil \frac{\sum_{i=s+1}^{|S|} t_{(i)}}{\bar{T}} \right\rceil - 1 \quad \forall k \in \mathcal{K} \quad (11)$$

Valid inequality (8) represents the minimum number of roller changes necessary based on the total degradation caused by all the coils scheduled in the rolling campaign. Similarly, (9) imposes a lower bound on the number of grade changes as a function of the number of grades in the campaign and casters in the mill. Cut (10) limits the minimum number of heats based on the weight of each coil where $t_{(i)}$ are the coil tonnages, sorted *increasing* by weight; because (10) imposes a bound for each caster, we account for the fact that the heaviest coils could all be on either caster. We place an upper bound on the same sum as in (11). The minimum and maximum number of grade changes in each case is one fewer than the bound on the total number of heats.

4.3.2. Improving the bounds on grade changes

Because the cost to change grades constitutes the majority of the penalty in the objective, we find it advantageous to better align the linear programming relaxation objective with that of the integer program using additional binary variables. To measure the cost explicitly, we let:

$$\Gamma_{ks} \geq \gamma_{g_1g_2} \cdot \left(\sum_{c \in \mathcal{C}} X_{ksc} + \sum_{c' \in \mathcal{C}} X_{k,s+1,c'} - 1 \right) \quad \forall k \in \mathcal{K}, s \in S; g_1, g_2 \in \mathcal{G} : g_1 \neq g_2 \quad (12)$$

and introduce extended binaries $\bar{G}_{ksg_1g_2} \in \{0,1\}$ indicating whether a grade change from grade g_1 to g_2 occurs in slot s on caster k . We then replace constraint (12) with:

$$\bar{G}_{ksg_1g_2} \geq \sum_{c \in \mathcal{C}} X_{ksc} + \sum_{c' \in \mathcal{C}} X_{k,s+1,c'} - 1$$

$$\forall k \in \mathcal{K}; s \in S; g_1, g_2 \in \mathcal{G} : g_1 \neq g_2 \quad (13)$$

and explicitly measure the grade change cost $\Gamma_{ks} \geq 0$ for each caster and slot in which the grades in consecutive slots are different (see constraint (5a)) as:

$$\Gamma_{ks} = \sum_{\substack{g_1, g_2 \in \mathcal{G} \\ g_1 \neq g_2}} \gamma_{g_1g_2} \cdot \bar{G}_{ksg_1g_2} \quad \forall k \in \mathcal{K}, s \in S$$

Hence, we enforce a change for each grade. Unless there are exactly $|\mathcal{S}|$ coils of the same grade g (which can be cast on the same caster), then there must be a grade transition from and/or to g :

$$\sum_{k \in \mathcal{K}} \sum_{s \in S} \sum_{g' \in \mathcal{G}} \bar{G}_{ksgg'} + \sum_{k \in \mathcal{K}} \sum_{c \in \mathcal{C}} X_{k,|S|,c} \geq 1 \quad \forall g \in \mathcal{G} \quad (14)$$

$$\sum_{k \in \mathcal{K}} \sum_{s \in S} \sum_{g' \in \mathcal{G}} \bar{G}_{ksg'g} + \sum_{k \in \mathcal{K}} \sum_{c \in \mathcal{C}} X_{k,1,c} \geq 1 \quad \forall g \in \mathcal{G} \quad (15)$$

That is, either a coil with grade g is assigned to the end (beginning) of a caster, or there is a transition from (to) grade g to (from) another grade at some slot.

Letting Z_{kg} be a variable equaling 1 if a coil of grade g is assigned to caster k , we introduce:

$$X_{ksc} \leq Z_{kg} \quad \forall k \in \mathcal{K}, s \in S, c \in \mathcal{C} : g = g_c \quad (16)$$

$$\sum_{k \in \mathcal{K}} Z_{kg} \geq 1 \quad \forall g \in \mathcal{G} \quad (17)$$

The number of inequalities in (16) can be reduced with the following reformulation:

$$\sum_{\substack{c \in \mathcal{C} \\ g_c = g}} X_{ksc} \leq m \cdot Z_{kg} \quad \forall k \in \mathcal{K}, s \in S, g \in \mathcal{G} : g = g_c$$

where $m = |\mathcal{C}| \ni g_c = g$. Hence, we can disaggregate constraints (14) and (15) for each caster as follows:

$$\sum_{s \in S} \sum_{g' \in \mathcal{G}} \bar{G}_{ksgg'} + \sum_{\substack{c \in \mathcal{C} \\ g_c = g}} X_{k,|S|,c} \geq Z_{kg} \quad \forall k \in \mathcal{K}, \forall g \in \mathcal{G} \quad (18)$$

$$\sum_{s \in S} \sum_{g' \in \mathcal{G}} \bar{G}_{ksg'g} + \sum_{\substack{c \in \mathcal{C} \\ g_c = g}} X_{k,1,c} \geq Z_{kg} \quad \forall k \in \mathcal{K}, \forall g \in \mathcal{G} \quad (19)$$

Rather than forcing grade changes for each grade, because the number of grades in a campaign is rather small (i.e., in practical applications between two and five), we can force grade changes for subsets of grades. To this end, we compute all possible combinations of grades potentially appearing on a caster and the corresponding minimum cost for each such sequence. For example, if there are three grades $\{g_1, g_2, g_3\}$ in a campaign but on a caster coils are only assigned with grades g_1 and g_2 (but not g_3), we can compute the minimum grade-change cost γ_{g_1, g_2, g_3} , and add a constraint:

$$\sum_{s \in S} \Gamma_{ks} \geq \gamma_{g_1, g_2, g_3} \cdot \left(1 - (1 - Z_{kg_1}) - (1 - Z_{kg_2}) - Z_{kg_3} \right) \quad \forall k \in \mathcal{K}$$

In general, for any subset of grades $\bar{\mathcal{G}}$, we can compute its associated minimum cost:

$$\gamma_{\bar{\mathcal{G}}} := \min_{\pi} \sum_{i=1}^{|\bar{\mathcal{G}}|-1} \gamma_{g_{\pi_i} g_{\pi_{i+1}}}$$

where this minimum is computed over all permutations π of the elements of $\bar{\mathcal{G}}$. Then, we can add the constraints:

$$\sum_{s \in S} \Gamma_{ks} \geq \gamma_{\bar{\mathcal{G}}} \cdot \left(1 - \sum_{g \in \bar{\mathcal{G}}} (1 - Z_{kg}) - \sum_{g \notin \bar{\mathcal{G}}} Z_{kg} \right) \quad \forall k \in \mathcal{K}; \bar{\mathcal{G}} \subseteq \mathcal{G} : |\bar{\mathcal{G}}| > 1 \quad (20)$$

to improve the lower bound. Although there are $2^{|\mathcal{G}|-1}$ possible subsets $\bar{\mathcal{G}} \subseteq \mathcal{G}$, because, in practice, $|\mathcal{G}| \leq 5$, the exponential nature of the size does not hinder implementability.

In addition to improving the bound on the grade change cost Γ_{ks} , we can force a grade change for each grade appearing in a subset $\bar{\mathcal{G}}$, strengthening valid inequalities (18) and (19):

$$\sum_{s \in S} \sum_{\substack{g' \in \bar{\mathcal{G}} \\ g' \neq g}} \bar{G}_{ksgg'} + \sum_{\substack{c \in C \\ g_c = g}} X_{k,|S|,c} \geq \left(1 - \sum_{g' \in \bar{\mathcal{G}}} (1 - Z_{kg'}) - \sum_{g'' \notin \bar{\mathcal{G}}} Z_{kg''} \right) \quad \forall k \in \mathcal{K}, \forall g \in \bar{\mathcal{G}} \quad (21)$$

$$\sum_{s \in S} \sum_{\substack{g' \in \bar{\mathcal{G}} \\ g' \neq g}} \bar{G}_{ksgg'} + \sum_{\substack{c \in C \\ g_c = g}} X_{k,1,c} \geq \left(1 - \sum_{g' \in \bar{\mathcal{G}}} (1 - Z_{kg'}) - \sum_{g'' \notin \bar{\mathcal{G}}} Z_{kg''} \right) \quad \forall k \in \mathcal{K}, \forall g \in \bar{\mathcal{G}} \quad (22)$$

That is, if grade g appears in the subset $\bar{\mathcal{G}}$, then there must be a transition from (to) grade g to (from) another grade g' also in $\bar{\mathcal{G}}$ on the same caster, unless a coil of grade g is in the last (first) slot of the caster.

4.3.3. Improving the bounds on trim loss

Given the non-decreasing cast-width at each caster, and because each coil c can have a cast-width between w_c and w_c^U , for each pair of coils (c, c') such that $w_c > w_{c'}^U$, if both are assigned on the same caster, then c must be sequenced before c' . Hence, the following *clique constraint* forbids coil c' from being assigned before c :

$$X_{ksc} + X_{k,s',c'} \leq 1 \quad \forall k \in \mathcal{K}; s, s' \in S : s' < s; c, c' \in C : w_c > w_{c'}^U \quad (23)$$

Valid inequality (23) can be strengthened as follows: Given an arbitrary width ω , (23) is also valid replacing c' by any coil with an upper bound $w_{c'}^U < \omega$ and replacing c by any coil with an order width $w_c \geq \omega$. Let Ω be the set of potential values for ω and let us denote $\underline{C}_\omega = \{c \in C : w_c^U < \omega\}$ and $\bar{C}_\omega = \{c \in C : w_c \geq \omega\}$. Hence, the following constraints are valid:

$$\sum_{c \in \bar{C}_\omega} X_{ksc} + \sum_{c' \in \underline{C}_\omega} X_{k,s',c'} \leq 1 \quad \forall k \in \mathcal{K}; s, s' \in S : s' < s; \omega \in \Omega \quad (24)$$

However, including this requires $\mathcal{O}(|S|^2)$ constraints for each ω . To avoid this, we use an additional variable $P_s^\omega \in [0, 1]$ for each potential width $\omega \in \Omega$ and we add the following constraints:

$$\sum_{c \in \bar{C}_\omega} X_{ksc} \leq P_s^\omega \quad \forall k \in \mathcal{K}, s \in S, \omega \in \Omega \quad (25)$$

$$\sum_{c \in \underline{C}_\omega} X_{ksc} \leq 1 - P_s^\omega \quad \forall k \in \mathcal{K}, s \in S, \omega \in \Omega \quad (26)$$

$$P_{s+1}^\omega \leq P_s^\omega \quad \forall s \in S, \omega \in \Omega \quad (27)$$

So, if any coil $c \in \bar{C}_\omega$ is scheduled in slot s' then $P_s^\omega = 1$ for $s \leq s'$. If this occurs, then no coil $c' \in \underline{C}_\omega$ can be scheduled in slots $s < s'$. Hence, (25) and (26) imply (24), but with a considerable reduction in the number of constraints.

We can compute Ω , the set of *critical pairs* of widths, in the following way: let $\bar{\rho}$ be an ordering of C such that $w_{\bar{\rho}_i} \geq w_{\bar{\rho}_{i+1}}$, and let $\underline{\rho}$ be an ordering of C such that $w_{\underline{\rho}_i}^U \leq w_{\underline{\rho}_{i+1}}^U$. Given a coil c with order width w_c , we find the first coil in the ordering $\underline{\rho}$ such that $w_c > w_{\underline{\rho}_*}^U$. Then, we find the last coil in the ordering of $\bar{\rho}$ such that $w_{\bar{\rho}_*} > w_{\underline{\rho}_*}^U$. We call this pair of coils $(\bar{\rho}_*, \underline{\rho}_*)$ a *critical pair* and we add $\bar{\rho}_*$ to the set Ω . In the worst case, the set Ω can be equal to the set of order widths $\{w_c : c \in C\}$, resulting in constraints similar to the number in (24). However, in practice, the number is much smaller (between 4 to 12 different values in our instances).

Table 7

Values for objective function penalties, scaled relative to the smallest, trim loss.

Penalty name	Value of penalty	Units	Constraint
Roller change	13.3	\$ per change	(3c)
Trim loss	1	\$ per inch	(2c)
Intercaster width difference	6.6	\$ per inch	(2b)
Grade-to-grade change	6.6–50	\$ per grade change	(5a)
Excessive gauge decrease	166.6	\$ per inch	(6)

Table 8

Width specifications, given in inches.

$\bar{w}$	$\bar{w}$	$\bar{w}$	$\bar{w}$	$\bar{w}$
2	0.5	5	1.5	6

Table 9

Description of cut experiments.

Type	Description	Cut numbers
Grade-change Cuts		
TYPE 0	(M) from Section 3	–
TYPE 1	(M) + $\bar{G}_{ksg, \bar{g}_2}$ variables and enforcing grade changes	(14)–(15)
TYPE 2	As TYPE 1, but + Z_{kg} variables	(18)–(19)
TYPE 3	(M) + Z_{kg} variables, and improving costs for subsets of grades	(20)
TYPE 4	As TYPE 2, and improving costs for subsets of grades	(21)–(22)
Trim-loss Cuts		
CRITICAL PAIRS	Cliques and using P_s^ω variables	(25)–(27)

5. Computational experiments

We code *Model (M)* in PYTHON 3.7 and solve its instances using GUROBI 9.0.2. We use the solver parameter settings MIPFocus=2 and SYMMETRY=2, as suggested by GUROBI's tuning procedure on the small instances of (M). Additionally, we turn off solver-generated cuts; without the confounding effects, solve times for those instances that attain optimality can be up to twice as fast. The computers employed for the experiments use CentOS Linux v7.6.1810 on x86 64 architecture, with four eight-core Intel Xeon E5-2670 processors and 128 GB of RAM. The model and sample data sets can be found here: <https://github.com/borelian/TorresEtAl2023>.

Scaled objective function penalties are given in Table 7, and width information pertinent to all runs is given in Table 8; tonnages of each coil range between 20 and 40 and, as mentioned previously, complete heats contain between 150 and 170 tons. We compare the relative efficacy of the cuts derived in Section 4.3 according to the combinations given in Table 9.

5.1. Small instances

We construct small instances containing between 28 and 38 coils using subsets from larger instances (containing about 150 coils) for which there are known, feasible solutions, in particular, complete heats. That is, the larger solutions contain sufficient coils such that complete heats can be constituted and coil characteristics are relatively comparable, e.g., there are no extreme differences in gauge or width. These instances contain hundreds of constraints and thousands of variables (both before and after presolve). We provide details regarding the number of grades, the number of each coil cut type, the difference between the smallest and largest widths and gauges, and the number of critical pairs in Table 10. These small studies compare the impacts of including the various cuts for instances that solve within an hour to the solver's default optimality tolerance to determine how much of an impact the cuts have specifically in tightening the lower bound, which, we believe, even for the larger instances which cannot solve to optimality, affects the provable quality of the solution to a greater extent than the solution itself.




Fig. 6. Upper and lower bounds, normalized with respect to the optimal objective value, as a function of solution time for the small instances (with critical-pair cuts).

Table 10

Description of small instances.

		$ G $	$ C^{mill} $	$ C^{HRB} $	$ C^{cut} $	$\min_e \{w_e\}$	$\max_e \{w_e\}$	$\min_e \{a_e\}$	$\max_e \{a_e\}$	$ \Omega $
Instances	I1	1	8	0	22	37.6	47.5	0.08	0.19	1
	I2	2	2	0	28	49.5	64.4	0.10	0.20	3
	I3	2	7	0	23	38.0	54.6	0.07	0.26	2
	I4	3	25	0	5	60.1	65.4	0.08	0.31	0
	I5	2	5	25	0	43.9	54.8	0.08	0.24	3
	I6	2	15	11	4	48.4	55.4	0.07	0.21	3
	I7	2	19	0	19	50.1	50.3	0.10	0.18	0
	I8	3	1	0	27	51.4	57.2	0.12	0.24	1
	I9	2	7	0	21	57.9	62.7	0.08	0.27	0

Table 11

Solve time for small instances, normalized with respect to the absolute solve time given in the second column.

		without CRITICAL-PAIR CUTS						with CRITICAL-PAIRS CUTS					
Solution		Grade-change cut type:						Grade-change cut type:					
time (sec.)		0	1	2	3	4		0	1	2	3	4	
Instances	I1	77	1.00	0.96	0.99	0.95	0.96	<u>0.52</u>	0.54	0.54	<u>0.52</u>	<u>0.52</u>	
	I2	651	1.00	0.21	0.25	0.32	0.37	0.02	0.06	0.02	<u>0.01</u>	0.02	
	I3	55	1.00	3.03	0.87	2.28	0.90	<u>0.19</u>	1.10	0.82	0.28	0.87	
	I4	1,618	1.00	4.18	<u>0.02</u>	4.45	0.09	1.02	3.95	2.44	4.45	1.41	
	I5	12	1.00	1.98	5.96	3.87	2.11	<u>0.35</u>	0.48	0.49	0.49	0.47	
	I6	30	1.00	<u>0.56</u>	0.64	1.32	0.96	0.79	0.74	0.79	1.11	1.04	
	I7	7,200 ^a	1.00	0.68	1.00	<u>0.23</u>	0.66	1.00	0.95	0.45	0.24	0.49	
	I8	7,200 ^a	1.00	0.52	0.33	5.96	0.21	3.23	0.29	0.36	1.07	<u>0.10</u>	
	I9	35	<u>1.00</u>	2.15	1.07	2.59	1.06	1.17	1.98	1.24	1.25	1.24	
Arithmetic Mean		1.00	1.09	0.55	1.56	0.59		0.50	0.66	0.50	0.49	<u>0.42</u>	
Geometric Mean		1.00	1.59	1.24	2.44	0.81		0.92	1.12	0.80	1.05	<u>0.68</u>	

^a Time limit.

Solution time to optimality is given in Table 11, where the results are normalized relative to times obtained from the corresponding instance of (M) from Section 3, i.e., TYPE 0. Each instance is associated with ten different solution variations: (i) five grade-change cut variations (including the default of no such cuts added) *without* the critical-pair cuts associated with trim loss; and, (ii) five grade-change cut variations (including the default of no such cuts added) *with* the critical-pair cuts associated with trim loss. For each of these

ten variations, we underline the one with the fastest solve time, i.e., the one with the smallest normalized value. Fig. 6 shows the upper and lower bounds in the branch-and-bound tree over time for the six instances invoking different grade-change cuts (TYPE 0 - TYPE 4) using the critical-pair cuts.

With respect to the small instances, we observe that the solve times vary considerably from between just over 10 s to more than 1,000 s. Generally speaking, the inclusion of critical-pair cuts improves

performance with the exception of INSTANCES 4, 7 AND 9 for which no critical pairs exist (see last column of Table 10); there are mixed results for INSTANCE 6 in which there are smaller width differences between coils, relative to the other instances. We also note that the effect of including cut TYPES 1–4 is minimal (with or without critical-pair cuts) relative to their lack of inclusion on INSTANCE 1 because this instance contains only one grade (see second column of Table 10). For the other instances that contain multiple grades, critical-pair cuts without additions, i.e., grade-change cuts of TYPE 0, appear to be the most effective, when assessed in aggregate (either by arithmetic or geometric mean), though the presence of grade-type cuts does effect a smaller gap early on in the solves, leading one to conclude that their presence may benefit larger instances in which the solve is stopped short of optimality. Fig. 6 demonstrates this phenomenon and shows that the majority of the solution time is spent tightening the bounds, rather than determining good solutions. In the absence of critical-pair cuts, grade-change cuts of TYPES 2 and 4 perform best (when considering both types of means), indicating that when $|\Omega| = 0$, these are an effective option.

5.2. Large instances

These instances are indicative of the greatest number of coils that the mill would simultaneously schedule and contain, on average, 5,000 constraints and 16,500 variables; presolve eliminates a trivial number of constraints. This set of computational experiments compares the benefits of using the optimization model relative to a heuristic, which, in and of itself, is already an improvement over the manual scheduling practice that the steel mill had used (Allen et al., 2022). Additionally, we show how optimizing the casting and rolling processes in tandem creates schedules that are superior to those resulting from isolating these phases and solving them in succession. Table 12 provides details on the number of grades, the number of each coil cut type, the minimum and maximum order widths and gauges, and the number of critical pairs. We impose a time limit of 12 h per instance, after which gaps are still large (about 43%, on average) but, in all likelihood– based on the results of the computational experiments on the small instances, owing to a weak bound.

Rather than focusing on the relative improvements from the various cut types, we instead depict the improvements from the optimization model (*Model (M)*) over schedules produced by solving the casting and rolling processes in isolation, and by the heuristic (which we call *Reference*) in Fig. 7, separated by cost type and normalized to *Reference*. We execute the problem separation between the casters (which emphasize chemistry) and the rollers (which emphasize geometry) in two manners: (i) we remove the cast-width (W_{ks} and B_c) and implicit gauge decrease variables; the resulting model minimizes grade and roller changes without considering the geometry required by the rolling process; and (ii) we fix the values for the assignment of a coil in a slot on a caster, X_{ksc} , resulting from the heuristic (*Reference*) and optimize over the resulting variables; this approach invariably fixes the chemistry while optimizing the degree of freedom in the geometry, the width. We refer to these two contrived methods as “variable elimination” and “variable fixing,” respectively.

For the “variable elimination method,” in which constraints (2a)–(2d) and (6) are removed, grade change costs are emphasized and the problem is considerably easier to solve; the four large instances yield a result in a matter of minutes. However, fixing the variables to this problem and attempting to obtain a solution for the monolith (*Model*




Fig. 7. Relative contribution of the different costs for the *Reference*, the “variable fixing” (*fixed*), and *Model (M)* solutions, where the magnitude of each cost is given as the product of the weight from Table 7 and the value of the corresponding elastic variable.

(*M*) yielded infeasibilities in three of the four instances due to the model’s inability to satisfy the cast width constraints. Table 13 displays the grade change and gauge decrease costs (normalized to those given by *Model (M)*); trim loss and intercaster width difference costs cannot be computed in the absence of a fully feasible solution. As expected, grade change costs are lower (all less than 1.0), but at the expense of gauge decrease costs (averaging more than five times those in the reference solutions) and feasibility for the monolith.

The “variable fixing” approach generates the same sequences as when solving for the unfixed variables (i.e., those other than X_{ksc}). The widths conform to a feasible schedule with minimized trim loss conditioned on the values for X_{ksc} . These instances also solve in minutes, and, while they are feasible for the monolith, they are all more expensive than those produced by *Model (M)*.

Regarding the heuristic *Reference*, which also generally produces feasible solutions, reductions in total objective function value (cost) obtained from *Model (M)* relative to *Reference* range between 20% and 40%. Greater savings in general result from reductions in the gauge and trim loss penalties relative to the other penalties; roller change cost is constant across solutions for all instances based on the simplistic wear model we use, which is independent of the sequence of coils.

To understand how *Model (M)* obtains the savings over *Reference*, we analyze INSTANCE C in detail, where the upper panel of Fig. 8 shows the solution to *Reference* and the lower panel depicts that from *Model (M)*. Purple vertical bars indicate a roller change. For each graph, the colored squares on top denote the grade of each coil, where those associated with the north caster are shown above those associated with the south caster. The *Model (M)* solution adds grade changes on the south caster (highlighted by the red ovals), but compensates for this penalty by lowering those associated with gauge and width. On the primary axis, circles and crosses denote order widths on the north and south casters, respectively, while the dashed and solid lines denote the cast widths on the north and south casters, respectively. Before the roller change in particular, the optimized solution with respect to the north caster saves considerably on trim loss associated with mismatched order and cast widths. Finally, the blue and red lines on the lower part of each graph represent the gauge, where the red oval highlights a particularly excessive gauge decrease (red lines in general denote gauge decrease), and the dashed and solid lines represent the north and south casters, respectively. *Model (M)*’s solution reduces the associated cost, for example, by moving large-gauge coils with small order widths to the end of the sequence.

Table 12
Description of large instances.

		$ G $	$ C^{mill} $	$ C^{HRRB} $	$ C^{cut} $	$\min_c \{w_c\}$	$\max_c \{w_c\}$	$\min_c \{a_c\}$	$\max_c \{a_c\}$	$ \Omega $
Instances	A	3	12	0	122	49.5	64.4	0.08	0.27	9
	B	3	31	0	93	38.0	54.6	0.06	0.29	11
	C	3	85	0	85	58.4	65.4	0.08	0.37	1
	D	2	22	96	6	43.9	54.8	0.08	0.24	6

Table 13
Costs, relative to *Model (M)*, obtained from the “variable elimination” method.

Instance	Grade change cost	Gauge decrease cost
A	0.33	3.94
B	0.36	7.73
C	0.86	6.55
D	0.50	2.98




Fig. 8. Details of scheduled solutions for INSTANCE C, where the purple vertical bars indicate a roller change, the squares denote grades differentiated by color with the top string corresponding to the north caster and the bottom string corresponding to the south caster. Circles and crosses denote order widths on the north and south casters, respectively, while the dashed and solid lines denote the same for the cast widths. The blue and red curves at the bottom of the graph denote the gauge on the north (solid) and south (dashed) casters, respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

We see how the dominating decisions of grade often obfuscate the heuristic’s ability to make trade-offs between the more nuanced decisions of width and gauge; on the other hand, although not provably optimal, *Model (M)* is able to take advantage of its clairvoyance to determine better solutions with respect to geometry, often, though not in INSTANCE C, when the grade changes are more favorable as well.

6. Conclusions

We pose here an optimization model that produces sequences of coils which satisfy constraints associated with heat sequences, as well as several operational limitations and quality considerations based on coil geometry and steel composition. Because the mill is 100% direct-charge, the casting and rolling operations are integrated without the ability to hold work-in-process. This results in more inextricably intertwined decisions than accounted for in the existing literature on steel-production optimization models. In order to generate optimal schedules for small instances and good solutions for larger ones, we propose a non-intuitive, but efficient, formulation and suggest cuts which tighten it further. We demonstrate not only the efficacy of the cuts, but also how solutions for larger instances would improve with the use of an optimization model over a heuristic or an approach that decouples casting and rolling decisions. The larger instances benefit in almost all aspects of the schedule quality, but particularly with respect to non-obvious tradeoffs that the optimization model can make involving the coil geometries.

While our cuts close the optimality gap on the smaller instances, provable optimality on the larger instances remains elusive owing to their size and complexity. Future work should examine more ways to tighten these gaps, perhaps through reformulation that would provide more structure to the model with respect to possible grade changes, which we hypothesize contribute most significantly to the weak bounds. Nonetheless, the solutions that *Model (M)* yields are useful in an

operational setting in that they can be obtained in an overnight run for the next day’s worth of coils to be sequenced at a large commercial minimill in North America; the optimization model improves not only the manually generated schedules, but those obtained via heuristics as well.

Acknowledgments

This work was partially funded by the Continuous Casting Consortium of the Colorado School of Mines and of the University of Illinois at Urbana-Champaign and by ANID/CONICYT-Fondecyt Regular 1200809. The authors wish to thank their industry contacts at Nucor Steel, Decatur, Alabama, USA, in particular, Dr. Bryan Petrus, for their help with and insights on all industry-specific aspects of this project. We acknowledge the aid of Mykel Allen, Walter Behalo, Samuel Billingham, Gideon Kukoyi, and Dr. D. Antony Tarvin regarding related modeling efforts. Finally, the comments of an anonymous review team helped to strengthen this paper.

References

Allen, M., Tarvin, D. A., Newman, A. M., & Thomas, B. G. (2022). Nucor steel improves production scheduling. *Working paper*.

Board & Vellum (2015). A visit to the NUCOR steel plant. <https://www.boardandvellum.com/blog/a-visit-to-the-nucor-steel-plant/>, (Last accessed on 2022-11-23).

Fukuyama, H., Liu, H.-h., Song, Y.-y., & Yang, G.-l. (2021). Measuring the capacity utilization of the 48 largest iron and steel enterprises in China. *European Journal of Operational Research*, 288(2), 648–665.

Graves, S. C. (1981). A review of production scheduling. *Operations Research*, 29(4), 646–675.

Harjunkoski, I., & Grossmann, I. E. (2001). A decomposition approach for the scheduling of a steel plant production. *Computers & Chemical Engineering*, 25(11–12), 1647–1660.

Industrial History (2021). Steel mill “cobble”. <http://industrialscenery.blogspot.com/2021/09/steel-mill-cobble.html>, (Last accessed on 2022-11-23).

Klotz, E., & Newman, A. M. (2013). Practical guidelines for solving difficult mixed integer linear programs. *Surveys in Operations Research and Management Science*, 18(1–2), 18–32.

Lee, H. S., Murthy, S. S., Haider, S. W., & Morse, D. V. (1996). Primary production scheduling at steelmaking industries. *IBM Journal of Research and Development*, 40(2), 231–252.

Lopez, L., Carter, M. W., & Gendreau, M. (1998). The hot strip mill production scheduling problem: A tabu search approach. *European Journal of Operational Research*, 106(2–3), 317–335.

Luecking, F., & Jannasch, E. (2008). New material reassignment process increases yield. In *AISTech 2008: Proceedings of the iron and steel technology conference*.

Maccarthy, B. L., & Liu, J. (1993). Addressing the gap in scheduling research: a review of optimization and heuristic methods in production scheduling. *The International Journal of Production Research*, 31(1), 59–79.

Mao, K., Pan, Q.-k., Pang, X., & Chai, T. (2014). A novel Lagrangian relaxation approach for a hybrid flowshop scheduling problem in the steelmaking-continuous casting process. *European Journal of Operational Research*, 236(1), 51–60.

Mattik, I., Amorim, P., & Günther, H.-O. (2014). Hierarchical scheduling of continuous casters and hot strip mills in the steel industry: A block planning application. *International Journal of Production Research*, 52(9), 2576–2591.

Missbauer, H., Hauber, W., & Stadler, W. (2009). A scheduling system for the steelmaking-continuous casting process. a case study from the steel-making industry. *International Journal of Production Research*, 47(15), 4147–4172.

Nucor (2022). Nucor product website. <https://nucor.com/products/steel/sheet>, (Last accessed on 2022-11-23).

Özgür, A., Uygun, Y., & Hütt, M.-T. (2021). A review of planning and scheduling methods for hot rolling mills in steel production. *Computers & Industrial Engineering*, 151, Article 106606.

Pan, Q.-K. (2016). An effective co-evolutionary artificial bee colony algorithm for steelmaking-continuous casting scheduling. *European Journal of Operational Research*, 250(3), 702–714.

Pinto, J. M., & Grossmann, I. E. (1998). Assignment and sequencing models for the scheduling of process systems. *Annals of Operations Research*, 81, 433–466.

Pochet, Y., & Wolsey, L. A. (2006). *Production planning by mixed integer programming*. Springer Science & Business Media.

Primetals Technologies (2019). Primetals Technologies receives first SRD segment slab caster order in China from Angang Iron & Steel. <https://www.primetals.com/press-media/news/primetals-technologies-receives-first-srd-segment-slab-caster-order-in-china-from-angang-iron-steel> (Last accessed on 2022-11-23).

- Ray, S. C., & Kim, H. J. (1995). Cost efficiency in the US steel industry: A nonparametric analysis using data envelopment analysis. *European Journal of Operational Research*, 80(3), 654–671.
- Riccardi, R., Bonenti, F., Allevi, E., Avanzi, C., & Gnudi, A. (2015). The steel industry: A mathematical model under environmental regulations. *European Journal of Operational Research*, 242(3), 1017–1027.
- Rong, A., & Lahdelma, R. (2008). Fuzzy chance constrained linear programming model for optimizing the scrap charge in steel production. *European Journal of Operational Research*, 186(3), 953–964.
- Rosyidi, C. N., Hapsari, S. N., & Jauhari, W. A. (2021). An integrated optimization model of production plan in a large steel manufacturing company. *Journal of Industrial and Production Engineering*, 38(3), 186–196.
- Roy, P. (2018). Electric arc furnace tapping inside ladle. <https://www.youtube.com/watch?v=kla0h0M6XgE>.
- Sierra-Paradinas, M., Soto-Sánchez, Ó., Alonso-Ayuso, A., Martín-Campo, F. J., & Gallego, M. (2021). An exact model for a slitting problem in the steel industry. *European Journal of Operational Research*, 295(1), 336–347.
- Silaen, A., Afsar, A., Zhao, H., Wu, B., Trimble, J., Grenough, D., Wilke, M., Fabina, L., Mengel, J., & Zhou, C. (2017). Development of a caster scheduling model for process optimization. In *AISTech 2017: Proceedings of the iron and steel technology conference* (pp. 3315–3320). AISTech.
- Tan, Y., Liu, S., & Huang, Y. (2013). A hybrid approach for the integrated scheduling of steel plants. *ISIJ International*, 53(5), 848–853.
- Tang, L., Liu, J., Rong, A., & Yang, Z. (2000a). A multiple traveling salesman problem model for hot rolling scheduling in Shanghai Baoshan Iron & Steel Complex. *European Journal of Operational Research*, 124(2), 267–282.
- Tang, L., Liu, J., Rong, A., & Yang, Z. (2000b). A mathematical programming model for scheduling steelmaking-continuous casting production. *European Journal of Operational Research*, 120(2), 423–435.
- Tang, L., Liu, J., Rong, A., & Yang, Z. (2001). A review of planning and scheduling systems and methods for integrated steel production. *European Journal of Operational Research*, 133(1), 1–20.
- Tang, L., Meng, Y., Wang, G., Chen, Z.-L., Liu, J., Hu, G., Chen, L., & Zhang, B. (2014). Operations research transforms Baosteel's operations. *Interfaces*, 44(1), 22–38.
- Tang, L., & Wang, G. (2008). Decision support system for the batching problems of steelmaking and continuous-casting production. *Omega*, 36(6), 976–991.
- Tang, L., Wang, G., & Chen, Z.-L. (2014). Integrated charge batching and casting width selection at Baosteel. *Operations Research*, 62(4), 772–787.
- Tang, L., Wang, G., Liu, J., & Liu, J. (2011). A combination of Lagrangian relaxation and column generation for order batching in steelmaking and continuous-casting production. *Naval Research Logistics*, 58(4), 370–388.
- Thomas, B. (1997). Modeling study of intermixing in tundish and strand during a continuous-casting grade transition. *Iron & Steelmaker*, 24(11), 83–96.
- Udofia, E., Greivel, G., Newman, A. M., & Thomas, B. G. (2023). Statistical roll force model for the prediction of roller wear in a hot-strip steel mill. *Working paper*.
- WorldSteel Association (2022). <https://worldsteel.org/steel-topics/statistics/world-steel-in-figures-2022/>.
- Zhang, L., & Thomas, B. G. (2003). State of the art in evaluation and control of steel cleanliness. *ISIJ International*, 43(3), 271–291.