



Iowa Research Online
The University of Iowa's Institutional Repository

Tippie College of Business Publications

9-11-2013

Sugarcane Harvest Logistics in Brazil

Kamal Lamsal
Emporia State University

Philip C. Jones
University of Iowa

Barrett W. Thomas
University of Iowa

Copyright © 2013, by Kamal Lamsal, Phillip C. Jones, Barrett W. Thomas.

Transportation Science Vol. 51, No. 2, May 2017, pp. 771–789. <https://doi.org/10.1287/trsc.2015.0650>

Comments

The authors would like to thank John Deere for financial support.

Hosted by [Iowa Research Online](https://www.lib-ir@uiowa.edu). For more information please contact: lib-ir@uiowa.edu.

Sugarcane Harvest Logistics in Brazil

KAMAL LAMSAL

Department of Management Sciences, University of Iowa
kamal-lamsal@uiowa.edu

PHILIP C. JONES

Department of Management Sciences, University of Iowa
philip-c-jones@uiowa.edu

BARRETT W. THOMAS

Department of Management Sciences, University of Iowa
barrett-thomas@uiowa.edu

Abstract

Sugar mills in Brazil represent significant capital investments. To maintain appropriate returns on their investment, sugar companies seek to run the mills at capacity over the entire nine months of the sugarcane harvest season. Because the sugar content of cane degrades considerably once it is cut, maintaining inventories of cut cane is undesirable. Instead, mills want to coordinate the arrival of cut cane with production. In this paper, we present a model of the sugarcane harvest logistics problem in Brazil. We introduce a series of valid inequalities for the model, introduce heuristics for finding an initial feasible solution, and for lifting the lower bound. Computational results demonstrate the effectiveness of the inequalities and heuristics. In addition, we explore the value of allowing trucks to serve multiple rather than single locations and demonstrate the value of allowing the harvest speed to vary.

I. INTRODUCTION

Since 1989, the monthly global price of raw sugar has averaged US\$ 0.1191 per pound. Starting in 2008, however, the price of raw sugar has steadily risen, peaking at US\$ 0.3209 per pound in January of 2011. While prices have fallen from those highs, raw sugar was trading at US\$ 0.2039 per pound in October 2012, the last month for which aggregate data is available (Economic Research Service 2012). There are two key sources of this rise in world sugar prices. First, consumption is increasing in countries such as India, China, Indonesia, and Turkey (McConnell et al. 2010). Second, ethanol production diverts raw sugar from consumers. Between 2000 and 2010, world ethanol use increased by 300%, and as a result of increasing oil prices, economic growth, and new government mandates, the growth is expected to continue (Valdes 2011). Importantly, the recent lifting of an US import tariff on foreign-made ethanol led to a ninefold increase in US imports of Brazilian sugar-based ethanol in 2012. Additional growth is expected in 2013 (Wexler December 17, 2012).

In the face of these high sugar prices, both consumers and producers have a keen interest in increasing world sugar supplies. As the world's largest exporter of both raw and refined sugar and the country whose production costs drive world sugar prices (McConnell et al. 2010), Brazil deserves particular focus. Yet surprisingly, the authors are not aware of any research that addresses logistics in the Brazilian sugarcane industry, and the work from other countries does not apply due to significant industry differences.

The focal point of the Brazilian industry is its sugar mills that crush raw cane to extract the juice from which raw sugar is eventually made. Sugar mills represent significant capital investments, and to maintain appropriate returns on their investment, sugar companies seek to run the mills near or at capacity over the entire nine months of the sugarcane harvest season.

Running the mills at capacity requires an adequate supply a sugar cane. To overcome the challenges of coordinating harvest and transport operations with the mill needs, the obvious solution is to decouple the mill operation from the supply operation by carrying a stock of raw cane. However, raw cane presents a complication. Because of evaporation and bacterial growth, the sugar content of cut cane degrades considerably over time (Salassi et al. 2004, Saska et al. 2009, Saxena et al. 2010). In areas of the world where cane is cut whole stalk, the cut stalk can last several days without significant lost of sugar content. However, in Brazil and some other countries, cane is primarily mechanically harvested, resulting in 12 - 18 inch billets of cane stalk. The multiple exposed ends increase the degradation in comparison to whole stalk cane (Salassi et al. 2004).

Because of the degradation, sugar producers want to reduce the cut-to-crush time, the time between when the sugar cane is cut in the field and when it is crushed at the mill. According to our conversations with our industry partners, a conservative estimate of average cut-to-crush time in a typical mill area in Brazil is three hours, even though the average travel time between the fronts and the mill is less than an hour (Personal Communication with Jose Coelho, Sugar Cane Segment Manager, John Deere, April 13, 2010). Improved logistics coordination, particularly coordinating the rate of harvest with the availability of trucks, offers an opportunity to reduce the cut-to-crush time and improve sugar yields while maintaining required service levels at the mill.

Estimating the cost of sugar loss is challenging. Sugar loss is affected by the cane variety, temperature, whether the cane is cut whole stock or billeted, whether it is burnt or clean, and the amount of debris entering the cane stalk during cutting and storage. These attributes differ by country, and only limited research explores the issue. As a rough estimate, for 2012/2013, Brazil is expected to crush 570 metric tons of sugarcane (Barros 2012). According to Saska et al. (2009), sugar loss per hour is linear over 24 hours and depends on the temperature range of the cane. Saska et al. (2009) finds that, for billeted cane (the case in Brazil) in the temperature range of 22-27 °C a reasonable range for temperatures during the harvest season in Brazil's largest sugar producing region, the center-south, is 0.03 tons per 100 tons of sugarcane per hour. At \$0.20 per pound of sugar, the cost of sucrose loss alone is over \$225 million and over \$330 million if prices were again to rise near \$0.30 per pound. However, if consider the possibility that the storage temperatures are much higher than the ambient temperature, notably over 27 °C the sugar loss increases to 0.32 tons per 100 tons of sugarcane per hour. In that scenario, the losses increase to \$2.4 billion and \$3.6 billion at \$0.20 and \$0.30 per pound of sugar, respectively. As noted in Salassi et al. (2004), dextran formation over time can additionally impact sugar value. Regardless, even small reductions in cut-to-crush time are likely to have an important monetary impact.

The sugarcane harvest is composed of three operations that must be coordinated: infield operations, over-the-road transport, and the mill operations. The infield operations usually occur in several pre-specified fronts a day. A front is a cluster of geographically close, but not necessarily contiguous, fields. The infield operations have several components. First, the cane is cut in the field, usually using a machine known as a harvester that processes the cane into uniformly sized billets (12-18 inches). While in operation, the harvester continuously feeds billets into an infield storage unit known as a cart. The cart is pulled by an infield transporter. This infield transporter and cart combination runs along side the harvester during the harvest operations, and when the cart is filled, the transporter and cart combination must be rotated with another infield vehicle and its associated cart to allow for continuous harvest operations. Filled carts are transported to

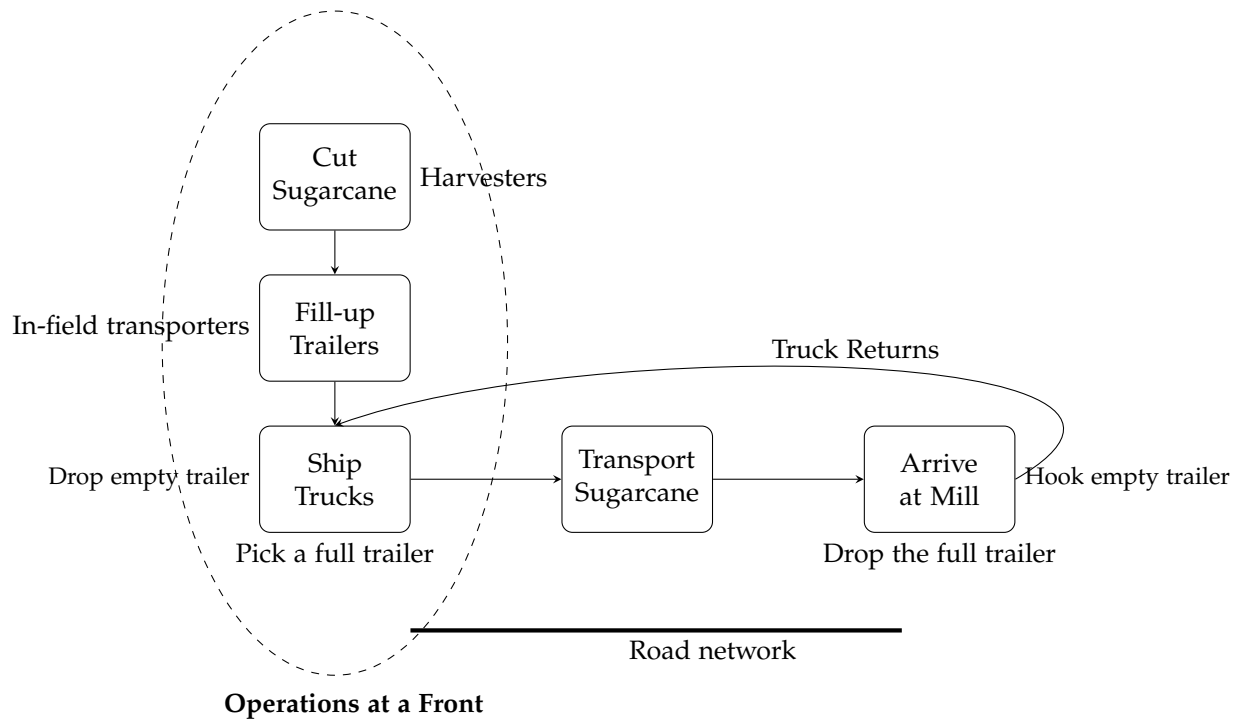


Figure 1: Mill Area Operation

an area known as the trans-loading zone that serves all of the fields in a front.

The second operation of the harvest begins at the trans-loading zone. At the trans loading zone in the fronts, the contents of the filled carts are transferred to over-the-road transport vehicles. These vehicles take the harvested cane from the fronts to the mill. The final operation of the harvest takes place at the mill where the over-the-road vehicles are unloaded. Once an over-the-road vehicle is unloaded, it can return to a front for its next load. Figure 1 illustrates the sugarcane harvest logistics problem. Note that, in practice, there are multiple fronts at varying distances from the mill even though only one representative front is shown in the figure below.

In this paper, we present a mixed linear integer programming model for the deterministic sugarcane harvest logistics problem in Brazil. The decision variables are the speed of harvesters at the fronts and the assignment of trucks to each of the loads produced at the fronts. Our objective is to minimize the cut-to-crush time of the sugarcane subject to the constraint that the mill is never starved of raw material. We present valid inequalities that allow us to strengthen the original formulation of the problem and also a heuristic for finding an initial feasible solution. In addition, we introduce a heuristic for lifting the lower bound of the linear relaxation. The result is particularly important in proving optimality as the linear relaxation of the original model often returns an objective value of zero. With the valid inequalities and heuristics, we demonstrate the ability to solve real-world sized problems in reasonable time. We demonstrate the value of our approach by contrasting it with the potentially managerially desirable approach of allowing trucks to serve a single front throughout the day. We also demonstrate the value of allowing the harvest rates to vary.

The rest of the paper is organized as follows. In Section II, we survey related literature.

Section III presents our math programming formulation. Section IV.2 presents the valid inequalities that strengthen the formulation, develops a heuristic for generating a feasible solution to the math program, and also introduces a heuristic for lifting the lower bound of the linear relaxation. Section V presents our experimental design and includes a description of the datasets on which we perform our computational experiments. The section also presents bounds on the number trucks needed to serve a particular mill area. These bounds are useful in designing our experiments. In Section VI, we present the results of our computational experiments. Section VII provides conclusions and presents future work.

II. LITERATURE REVIEW

Agricultural applications have a long history in the Operations Research literature, beginning with Heady (1954) that introduced the use of linear programming to the agricultural sector. Recent reviews of this broad field can be found in Lowe and Preckel (2004) and Ahumada and Villalobos (2009). The supply of feedstock to sugar mills has garnered significant attention in the academic literature. Giles (2009) gives an overview of logistics issues stressing the need for the coordination between harvesting, transport, and storage for the smooth operation and the profitability of a sugar mill. Yet, to the best of the authors' knowledge, none of the literature addresses the degradation of sucrose during the cut-to-crush delay, a key concern in Brazil.

Most work on the sugarcane industry has been done in the last 15 years. The work can be more or less divided into value chain optimization and harvest scheduling. The only work related to the Brazilian sugarcane industry of which the authors are aware focuses on value chain optimization. Kawamura et al. (2006) and Paiva and Morabito (2009) focus on production planning across multiple sugar mills, and da Silva et al. (2013) considers production planning at the individual mill level. Jena and Poggi (2013) present an optimization model for scheduling fields for harvest. Additional value-chain related literature includes Higgins et al. (1998), Grunow et al. (2008), and Kostin et al. (2011), which cover the Australian, Venezuelan, and Argentinian industries, respectively.

In this paper, we focus on harvest and scheduling. This subset can be divided by country, notably Australia, South Africa, Brazil, Cuba and Thailand. There is very little similarity among the work from different countries because of differences in industry structure and operation. The harvesting practices and resultantly the transportation of the cane from the field is influenced by the terrain, the relative cost of labor and capital, the state of road and rail networks, and the distribution of sugarcane fields relative to the mill. In Brazil, in 2009, 70.0% of the cane supply was harvested from mill controlled fields, with 55.5% being mill owned (Neves et al. 2010), giving mill operators significant control over the supply process for the mill. In this situation, the mill does not face the problem of being oversupplied during the peak season as in Thailand nor is the issue of integrating the decision making process between growers and harvesters of the importance that it is in South Africa. Further, harvesting in Brazil is done 24 hours a day, and as a result, storage is not an issue as it is Australia, Cuba, and the United States.

Most closely related to the work presented here is Salassi et al. (2009b). Salassi et al. (2009b) focus on sugar harvest logistics in Louisiana in the United States. As in the Brazilian case, each farm supplies a pre-determined number of truckloads each day. A key difference is that, once assigned to a farm, a truck serves loads at that farm until the harvest is completed at the farm. Further, harvest rates are fixed, and because harvesting only takes place over 14 hours in a day, the objective is designed to reduce truck congestion at the mill rather than to minimize cut-to-crush delay. One similarity is that Salassi et al. (2009b) groups the many farms harvested in a day into farm groups. These farm groups reduce the problem size as fronts do in the research presented

here. Previously, Salassi et al. (2004) had explored the value of extending the hours of harvesting in the day. The goal was to reduce the degradation of cane resulting from storage.

In South Africa, the mill neither owns nor controls a big share of the farms. Hansen et al. (2002) develop a simulation model and conduct sensitivity analysis to investigate and reduce the delays in the South African sugarcane harvest and delivery systems. Their study shows that an integrated system comprising the harvest, transport, and mill process can lead to significant reductions in delay times and cost. Le Gal et al. (2009) develop a simulation to investigate the impact of increased mechanization of the sugarcane harvest. Lejars et al. (2008) also develop a simulation model to see the effects of centralized decision-making among the various stakeholders (sugar cane growers, harvesters, haulers, and millers) in the South African industry versus decentralized decision-making. McDonald et al. (2008) also develop a similar simulation model to simulate the sugarcane harvesting, transport and mill-yard activities for a mill supply area.

In Thailand, the sugar industry has a large number of small-sized, independent farms. This industry configuration causes uneven supplies throughout the harvesting season. Thus, research addresses supply issues to Thai sugar mills. Supsomboon and Yosnual (2004) present a stochastic model that helps mills optimize their order quantities, given the uncertainties of farmers' delivery lead times and quantities. Prichanont et al. (2005) use discrete-event simulation to demonstrate that the number of trucks should be reduced in order to avoid excess supply. They show that the excessive transportation cost is due to inefficient cane delivery truck utilization and extensive truck waiting time at the mill. Using the simulation model for one of the mills they studied, they show that as much as 600 of the existing 1000 trucks can be eliminated while the mill's needs remain statistically unchanged.

Diaz and Perez (2000) present a simulation model for the transportation of sugarcane in Cuba. For the same country, Lopez Milan et al. (2006) develop a linear programming model to pick up cane from different farms and storage locations to minimize transportation cost. The model accounts for more than one mode of transport. Besides road transport, they consider inter-modal transports which include first, transportation via trucks to warehouses, and then subsequent shipments to the mill by train. In Cuba, the train system is a cheaper alternative to road transport and also serves as a buffer because the mill operates for 24 hours a day while harvesting is done only 14 hours a day. Lopez-Milan and Pla-Aragones (2013) introduce a decision support system that builds on the work in Lopez Milan et al. (2006). Unlike in this paper, neither Lopez Milan et al. (2006) nor Lopez-Milan and Pla-Aragones (2013) directly address the operational level scheduling of individual transport vehicles.

Historically, the Australian sugar industry has been the most exposed to the "world market price" because Australia neither has a large domestic demand like India nor access to protected high-priced European Union or United States domestic markets (Hildebrand 2002). To stay competitive in the market, the Australian sugar industry focused on lowering the transportation cost, and this is the main theme of the literature related to the Australian sugar industry. While some areas in the Australian sugar industry use rail transport (Higgins and Postma 2004, Higgins and Davies 2005, Higgins and Laredo 2006), most relevant to the discussion in this paper is the use of truck transportation. Unlike Brazil, in Australia, harvesting is limited to daylight hours. Higgins (2006) develops an approach for scheduling the individual vehicles in a road-bound transportation system. Two alternative solution methods based on meta-heuristics are proposed for the model to solve problems of practical sizes. Both meta-heuristics were able to find solutions with an average reduction in vehicle queue time of about 90% compared to the manual methods used by the traffic officers in the mills. This helped reduce the required number of vehicles. Unlike the work presented in this paper, Higgins (2006) is concerned with reducing the waiting time for the trucks only at the unloading zone at the mill. We are concerned with decreasing the overall

cut to crush time for the cane, while fulfilling the mill requirements at all times. We also have the option of controlling the harvester rate. Unlike Higgins, we provide an exact solution method.

III. MODEL

In this section, we present a linear, mixed integer programming formulation for coordinating sugarcane logistics harvests. For ease of exposition, we assume that all trucks start at the mill and harvesting has not yet started at the fronts. We also assume that the number of loads available at the fronts exactly matches the needs of the mill for the time horizon in question. We first introduce the notation for the problem and then present the math program.

III.1 Notation

\mathcal{N} : Set of needs at the mill, where the cardinality of \mathcal{N} is N .

\mathcal{T} : Set of trucks available to service loads, where the cardinality of \mathcal{T} is T .

\mathcal{F} : Set of fronts, where the cardinality of \mathcal{F} is F .

n_i : Time at which a load is needed at the mill to ensure the mill can continue operating at the desired rate, $i = 1, \dots, N$.

m : Time between two consecutive mill needs, assumed constant.

d_i : Dispatch time to pick up the i^{th} need, $i = 1, \dots, N$.

h_i : Time at which the harvest of the i^{th} mill need is completed (ready time), $i = 1, \dots, N$.

p_i : Pick-up Time of the i^{th} need, $i = 1, \dots, N$.

a_i : Arrival time of the i^{th} load to the mill, $i = 1, \dots, N$.

β_f : Earliest time that harvesting can begin at front f , $f = 1, \dots, F$.

L_f : Number of loads to be produced from front f , $f = 1, \dots, F$.

t_f : Travel time to the front f , $f = 1, \dots, F$.

s_f : Travel time to mill from front f , $f = 1, \dots, F$.

r_f : Round trip time for front f , $f = 1, \dots, F$.

\underline{H}_f : Minimum time needed to harvest a load at front f , $f = 1, \dots, F$.

b_t : Earliest time time that truck t can be dispatched, $t = 1, \dots, T$.

$$x_{it} = \begin{cases} 1 & \text{if truck } t \text{ brings the } i^{\text{th}} \text{ need to the mill.} \\ 0 & \text{otherwise} \end{cases}$$

$$y_{if} = \begin{cases} 1 & \text{if the } i^{\text{th}} \text{ need is fulfilled by front } f \\ 0 & \text{otherwise} \end{cases}$$

III.2 Formulation

$$\begin{aligned}
 \text{(P) minimize } & \sum_{i \in \mathcal{M}} [(p_i - h_i) + (n_i - a_i)] & (1) \\
 \text{subject to } & d_i \leq p_i \leq a_i \leq n_i & , \forall i \in \mathcal{N} & (2) \\
 & h_i \leq p_i & , \forall i \in \mathcal{N} & (3) \\
 & \sum_{t \in \mathcal{T}} x_{it} = 1 & , \forall i \in \mathcal{N} & (4) \\
 & \sum_{f \in \mathcal{F}} y_{if} = 1 & , \forall i \in \mathcal{N} & (5) \\
 & d_i + \sum_{f \in \mathcal{F}} [t_f * y_{if}] \leq p_i & , \forall i \in \mathcal{N} & (6) \\
 & p_i + \sum_{f \in \mathcal{F}} [s_f * y_{if}] = a_i & , \forall i \in \mathcal{N} & (7) \\
 & a_i - d_{i'} \leq (2 - x_{it} - x_{i't}) \times M & , \forall i, i' \in \mathcal{N}, i' > i, \forall t \in \mathcal{T} & (8) \\
 & h_i + \underline{H}_f - h_{i'} \leq (2 - y_{if} - y_{i'f}) \times M, \forall i, i' \in \mathcal{N}, i' > i, \forall f \in \mathcal{F} & (9) \\
 & \sum_{i \in \mathcal{M}} y_{if} \leq L_f & , \forall f \in \mathcal{F} & (10) \\
 & d_t \geq b_t & , \forall t \in \mathcal{T} & (11) \\
 & h_i - \sum_{f \in \mathcal{F}} \underline{H}_f y_{if} \geq \sum_{f \in \mathcal{F}} \beta_f y_{if} & , \forall i \in \mathcal{N} & (12) \\
 & x_{it} \text{ binary}, y_{if} \text{ binary} & , \forall i \in \mathcal{N}, \forall f \in \mathcal{F}, \forall t \in \mathcal{T}. & (13)
 \end{aligned}$$

Equation (1) is the problem objective. We note that we can ignore travel time as, to maintain feasibility, all of the loads from each front must be transported to the mill and by assumption there are only enough loads available at each front such that the total is required to meet the mill's needs. Constraints (2) enforce the order of events associated with each load. For the load that fulfills the i^{th} need, the dispatch time of the truck that services the i^{th} need must occur before the actual pickup, and the load that fulfills the i^{th} need at the mill should arrive to the mill before the i^{th} need. Similarly, pickup at the front must occur before the arrival at the mill. Constraints (3) stipulate that harvesting for a load must be completed before it is picked up. Constraints (4) require that each load is picked up only once and constraints (5) say each load is harvested only once. Constraints (6) connect dispatch time and pick-up time, and the constraints (7) link pick-up time with arrival time at the mill. Constraints (8) say that, if two loads are picked up by the same truck, the dispatch time of the latter load is at least as large as the previous load's arrival time at the mill. Constraints (9) are the analogy of 8 when two loads are from the same front. i.e. if two loads are harvested at the same front, we can start harvesting the latter only after the first one is harvested. For both constraints (8) and (9), M is a large number. Constraints (10) enforce the harvest cap on each front. Constraints (11) require that no truck is dispatched before it is available. Constraints (12) ensure that harvesting does not begin at a front before that front is available. Constraints (13) ensure binary decision variables.

IV. VALID INEQUALITIES, INITIAL FEASIBLE SOLUTIONS, AND LIFTED LOWER BOUNDS

While the model has obvious similarities to a constrained assignment problem, the addition of the variable harvest rates and thus variable ready times, creates challenges. Preliminary work

demonstrated that it was not possible to solve instances of the above model using commercial solvers. In fact, it was often not possible to find even a feasible solution in a reasonable time. Thus, in this section, we present valid inequalities, a heuristic for finding a feasible solution, and a heuristic for lifting the lower bound of the linear relaxation. As we demonstrate in our computational results, these enhancements allow us to solve realistically-sized problems in reasonable time.

IV.1 Valid Inequalities

In this section, we present results that help strengthen the formulation in the presented previously. We first note that we can strengthen the formulation by replacing Constraints (2) with

$$(d_i + \underline{t}) \leq (p_i + \underline{s}) \leq a_i \leq n_i, \forall i \in \mathcal{N}, \quad (14)$$

where $\underline{t} = \min_{f \in \mathcal{F}} [t_f]$ and $\underline{s} = \min_{f \in \mathcal{F}} [s_f]$.

Next, we state and prove a proposition that demonstrates the loads arriving to the mill can be processed in a first-in-first-out manner.

Proposition 1 (FIFO Arrival Times). *Let a_i and a_j denote the arrival times of loads i and j , respectively, at the mill. Let load i be assigned to the mill need occurring at time n_i and let load j be assigned to the mill need occurring at time n_j . If $a_i < a_j$ and the preceding assignment is feasible, then we may assume without loss of generality that $n_i \leq n_j$.*

Proof. Because a load cannot fill a mill need that occurs prior to its arrival, it follows that $a_i \leq n_i$ and $a_j \leq n_j$ as we have assumed feasibility. Note that the total waiting time for the two loads at the mill is $(n_i - a_i) + (n_j - a_j)$. With $a_i < a_j$, it follows that $a_i < n_j$. If it were the case that $n_j < n_i$, then it would also follow that $a_j < n_i$, and it would be possible to switch the loads so that load i is assigned to mill need n_j and load j is assigned to mill need n_i . The total waiting time for the two loads would be $(n_j - a_i) + (n_i - a_j)$. Because the total waiting time is the same given either assignment, the proposition follows. \square

The result offers a way to break symmetry among the arrival of loads to the mill. Symmetry occurs when a group of variables forms a ‘‘symmetry group,’’ a group of variables can be permuted without changing the value of the solution. As noted in (Margot 2010), breaking symmetry can turn a computationally intractable problem into one that is easily solved. Consequently, we make use of the Proposition 1 and add the following constraints to the formulation:

$$a_i \leq a_{i+1}, \forall i \in \mathcal{N} \setminus N. \quad (15)$$

The next result bounds from below the arrival time values for a solution to **P**. To facilitate presentation, we first present additional notation. Let \underline{H}_f be the minimum time required to harvest a load at front f and let α_{fl} be the arrival time to the mill of the l^{th} load harvested at front f . Then, for a front f , consider the sequence of arrival times such that $\alpha_{f1} = \beta_f + \underline{H}_f + s_f$, $\alpha_{f2} = (\beta_f + \underline{H}_f) + \underline{H}_f + s_f = \alpha_{f1} + \underline{H}_f$, and thus $\alpha_{fl} = \alpha_{f,l-1} + \underline{H}_f$ for $l = 2, \dots, L_f$. Let \mathcal{LB}_f be the ordered set of such arrival times for front f and let $\mathcal{LB} = \cup_{f \in \mathcal{F}} \mathcal{LB}_f$, ordered in ascending order. We let a'_i be the i^{th} element in \mathcal{LB} . The following result follows directly from the construction of \mathcal{LB} and Proposition 1.

Proposition 2. *The value a'_i is a lower bound on a_i for all $i \in \mathcal{N}$ resulting from a solution of **P**.*

As a result of Proposition 2, we can add the following constraints to \mathbf{P} :

$$a'_i \leq a_i, \forall i \in \mathcal{N}. \quad (16)$$

The following corollary follows from the construction of \mathcal{LB} and Proposition 2.

Corollary 1. *If there exists a need $i \in \mathcal{N}$ such that $a'_i > n_i$, then \mathbf{P} is infeasible.*

We next present an upper bound on the arrival time values for a solution to \mathbf{P} . Again, we introduce new notation. Let η_{fl} be the time at which harvesting of the l^{th} load at front f is completed and let $\hat{\alpha}_{fl}$ be the arrival time to the mill of the l^{th} load harvested at front f . For each front $f \in \mathcal{F}$, let $\eta_{fL_f} = n_N - s_f$ and $\hat{\alpha}_{fL_f} = \eta_{fL_f} + s_f$. Then, $\eta_{fl} = \eta_{f,l+1} - \underline{H}_f$ and $\hat{\alpha}_{fl} = \eta_{f,l+1} - \underline{H}_f + s_f = \hat{\alpha}_{f,l+1} - \underline{H}_f$, for $l = 1, \dots, L-1$. Let \mathcal{UB}_f be the ordered set of such arrival times for front f and let $\mathcal{UB} = \cup_{f \in \mathcal{F}} \mathcal{UB}_f$, ordered in ascending order. We let a''_i be the i^{th} element in \mathcal{UB} . The following result also follows directly from the construction of \mathcal{UB} and Proposition 1.

Proposition 3. *The value a''_i is an upper bound on a_i for all $i \in \mathcal{N}$ resulting from a solution of \mathbf{P} .*

As a result of Proposition 3, we can add the following constraints to \mathbf{P} :

$$a''_i > a_i, \forall i \in \mathcal{N}. \quad (17)$$

In addition to the FIFO result and bounds on the earliest and latest arrivals of loads to the mill, we can also fix some of the initial assignments of trucks to loads. Because we assume a homogeneous fleet of vehicles, we begin by noting that, without loss of generality, we can enforce:

$$x_{11} = 1 \text{ and} \quad (18)$$

$$x_{1t} = 0, \forall t \in \mathcal{T}, t \neq 1. \quad (19)$$

As a consequence of Proposition 2, we know that the earliest time at which truck 1 can possibly return back to the depot after picking up load 1 is a'_1 . Based on this information, we introduce Proposition 4 that characterizes an initial set of truck assignments.

First, it is useful to define the following notation. Recall that the round trip time for front f is r_f . Let $\rho = \min_{f \in \mathcal{F}} r_f$. Define the set \mathcal{S} such that $\mathcal{S} = \{i : i \in \mathcal{N}, n_i < a'_1 + \rho\}$.

Proposition 4. *Every i in \mathcal{S} requires a different truck.*

Proof. As ρ is the smallest possible round trip time, having arrived back to the mill at the earliest at time a'_1 , truck 1 cannot arrive back to the mill with its second load any earlier than $a'_1 + \rho$. Further, no truck serving a load i such that $n_i < a'_1 + \rho$, the loads in \mathcal{S} can return from serving its second load before $a'_1 + \rho$. Consequently, if a truck t in \mathcal{T} serves a load i in \mathcal{S} , t cannot also serve a load i' in \mathcal{S} . \square

As a consequence of Proposition 4, we add the follow constraints to \mathbf{P} , incorporating constraints (18) and (19):

$$x_{ii} = 1, \forall i \in \mathcal{S} \text{ and} \quad (20)$$

$$x_{ij} = 0, \forall i \in \mathcal{S}, \forall t \in \mathcal{T}, t \neq i. \quad (21)$$

Even with the strengthened formulation, the lower bound provided by the linear relaxation of the math program above is not tight. We note that the value of the relaxation can be raised by solving a relaxed integer program that excludes the truck assignment constraints (constraints 4 and 8). This integer program is easier to solve than that presented above. The lower bound thus calculated gives the total wait if we could pick up all the loads at their ready times. Any constraint on number of trucks only increases the wait time.

IV.2 A Heuristic for Generating a Feasible Solution

Starting a branch-and-bound procedure with an initial feasible solution often improves computation time. In this section, we develop a heuristic that is capable of generating feasible solution to problem **P**. The heuristic operates by decomposing problem **P** into its harvest and truck assignment components. While the heuristic does not guarantee a feasible solution, our computational experiments demonstrate that the method does so for all of our test cases and that the feasible solutions that are found are effective in improving computation time.

We begin by presenting an algorithm that generates a set of assignments of loads to fronts and corresponding arrival times to the mill. Our approach is motivated by the construction of the set \mathcal{UB} in the previous section. It is straightforward to see that if the values in \mathcal{UB} are feasible, then the arrival times in \mathcal{UB} are optimal.

The algorithm can be found in Algorithm 1. The algorithm returns a set of assignments of fronts to mill needs, denoted \bar{y} , and a set of arrival times for each mill need i , denoted \bar{a} . Throughout, using the variable l_f , the algorithm tracks the number of loads at front f that are remaining to be assigned. For each front f , the algorithm uses the value β_f to track the next time at which a load from front f can reach the mill.

The algorithm begins with the last mill need and assigns a front to meet that need and seeks to assign loads in descending order of the time at which the load is needed at the mill. To help maintain feasibility, before choosing a front to meet the N^{th} and final need of the mill, we test all fronts f to ensure that, if front f is not assigned the N^{th} load, then f still has enough time to harvest all L_f loads required by front f . To meet the m^{th} need, the algorithm chooses the front that can deliver the load as close to n_m as possible. In the case of the need to break a tie, the algorithm chooses the front f that minimizes the ratio $\frac{n_m}{\underline{H}_f \times l_f}$, where \underline{H}_f is the minimum time required to harvest a load at front f . The ratio is a measure of a front's flexibility to meet future loads and choosing the front with the maximum ratio chooses the least flexible front to break the tie. Once the assignment for the m^{th} need is chosen, \bar{a} and \bar{y} are updated. The algorithm then updates the latest time at which the front chosen to fill the m^{th} need can supply a feasible load to the mill. This updated time reflects the fact that load m arriving at time \bar{a}_m and supplied by front f must have been harvested by time $\bar{a}_m - s_f$. So, the latest time at which front f could finish harvesting its next load is $\bar{a}_m - s_f - \underline{H}_f$. Then, the latest time at which a load from front f could arrive at the mill after supplying the m^{th} load is $\bar{a}_m - s_f - \underline{H}_f + s_f = \bar{a}_m - \underline{H}_f$. In the next step, the algorithm updates the next available delivery time of all fronts to reflect the the time at which the $(m - 1)^{\text{st}}$ need is required by the mill.

An alternative to Algorithm 1 is to solve a relaxation of the **P** that removes constraints 4 and 8. The same relaxation is discussed in the previous section for finding an improved root node bound. However, preliminary experiments found that the proposed algorithm more often leads to feasible and better solutions when coupled with the following truck assignment phase.

We next present an algorithm for assigning the loads in \bar{a} to vehicles. From \bar{a} and \bar{y} returned by Algorithm 1, it is straightforward to compute a set of corresponding harvest completion times \bar{h} . The algorithm is presented formally in Algorithm 2. Throughout, the algorithm uses the value

Algorithm 1 Assignment of Fronts to Mill Needs

Output:

A vector of assignments of fronts to needs, \bar{y} , a vector a arrival times for the mill needs, \bar{a}

Initialization:

Set $m = M$, $\beta_f = n_N \forall f \in \mathcal{F}$, $l_f = L_f \forall f \in \mathcal{F}$, and $\bar{y}_{if} = 0 \forall i \in \mathcal{N}, f \in \mathcal{F}$

while $m \neq 0$ **do**

$temp \leftarrow \operatorname{argmax}_{f \in \mathcal{F}, l_f \neq 0} \{\beta_f\}$

$\bar{a}_m \leftarrow \beta_{temp}$

$\bar{y}_{m,temp} \leftarrow 1$

$\beta_{temp} \leftarrow \bar{a}_m - H_{temp}$

$\beta_f \leftarrow \min\{\beta_f, n_{m-1}\} \quad \forall f \in \mathcal{F}$

$l_{temp} \leftarrow l_{temp-1}$

$m \leftarrow m - 1$

end while

δ_t to represent the time at which vehicle t is available to service its next load. For each load $i \in \mathcal{N}$, the algorithm chooses the vehicle that has been at the mill the longest. Ties are broken arbitrarily. Once a vehicle has been selected for assignment, the algorithm updates \bar{x} accordingly, and then for the assigned vehicle t , the algorithm updates δ_t by computing the time at which the vehicle will return to the depot.

Algorithm 2 Assignment of Vehicles to Harvested Loads

Output: Assignment of vehicles to loads, \bar{x} .

Initialization:

$\delta_t = 0 \forall t \in \mathcal{T}$ and $\bar{x}_{it} = 0 \forall i \in \mathcal{N}, t \in \mathcal{T}$

for $i = 1$ to N **do**

$temp \leftarrow \operatorname{argmin}_{t \in \mathcal{T}} \{\delta_t\}$

$\bar{x}_{i,temp} \leftarrow 1$

$\delta_{temp} \leftarrow \max\{\bar{h}_i, \delta_{temp} + \sum_{f \in \mathcal{F}} (\bar{y}_{if} + t_f)\} + \sum_{f \in \mathcal{F}} (\bar{y}_{if} + s_f)$

$i \leftarrow i + 1$

end for

IV.3 A Heuristic to Lift the Lower Bound

As the linear relaxation of the proposed model, which we call \mathbf{P}' , is often zero, proving optimality is challenging. We develop a heuristic that exploits the structure of the solution to raise the lower bound. For ease of exposition, we describe the heuristic with regard to the front-assignment variables, the y variables. Algorithm 3 formally presents the heuristic. The algorithm for the truck-assignment variables, the x variables, is analogous. A lower bound can be found by taking the minimum of the wait times computed from the two.

The method takes as input a solution to \mathbf{P}' , with assignment variables y_{if} , which we will refer to as y' to represent that the binary condition has been relaxed. For each $f = 1, \dots, F$, we let $\mathcal{R}_f = \{i \mid i \in \mathcal{N}, y'_{if} \notin \{0, 1\}\}$, the set of mill needs for which the front f fulfills a partial load in the solution to \mathbf{P}' . We assume that the set \mathcal{R}_f is ordered from smallest to largest, and for algorithmic convenience, we assume that 0 is an element of \mathcal{R}_f for every front f . Thus,

given any two consecutive elements in \mathcal{R}_f , say τ_j and τ_k , there exists a sequence of mill needs $i = \tau_j + 1, \dots, \tau_k - 1$ such that $y'_{if} \in \{0, 1\}$.

Now, consider some sequence $i = \tau_j + 1, \dots, \tau_k - 1$. Let \bar{i} be the largest value in $i = \tau_j + 1, \dots, \tau_k - 1$ such that $y'_{if} = 1$. Using an idea similar to that used in the heuristic for finding a feasible solution, we recognize that the latest possible time at which the harvest of the \bar{i}^{th} mill need could be completed while still maintaining the feasibility of the solution is $n_{\bar{i}} - s_f$. Ignoring any constraints on vehicles, this completion time would also imply that the \bar{i}^{th} mill need is satisfied with no wait.

Now, if it exists, we let $\bar{i} - 1$ be the next largest $i = \tau_j + 1, \dots, \tau_k - 1$ such that $y'_{if} = 1$. Given the previous logic, the load $\bar{i} - 1$ could not have been harvested any later than $n_{\bar{i}} - s_f - \bar{H}_f$. Thus, if front f satisfied mill need \bar{i} by harvesting the load at $n_{\bar{i}} - s_f$, the latest that front f can harvest load $\bar{i} - 1$ is $\min\{n_{\bar{i}} - s_f - \bar{H}_f, n_{\bar{i}-1} - s_f\}$. If the minimum is obtained by $n_{\bar{i}} - s_f - \bar{H}_f$, then the $(\bar{i} - 1)^{\text{th}}$ load incurs a minimum wait of $n_{\bar{i}-1} - s_f - [n_{\bar{i}} - s_f - \bar{H}_f]$ time units. The correctness of the lifted lower bound follows from these feasibility arguments. The algorithm sums these waits over all fronts and all mill needs to compute the lifted lower bound.

V. EXPERIMENTAL DESIGN

In this section, we describe the computational experiments designed to test our approach as well as to gain insight into the sugarcane harvest logistics problem. To aid the description of the experiments, we first describe how we determine the number of trucks that should be used in a dataset and then the datasets that we use.

V.1 Bounds on the Number of Trucks

As the number of available trucks is fixed in the short term, the number of trucks is a parameter in our sugarcane logistics model. We test our approach on a range of the number of trucks. It is obvious that, if there are too few trucks available, there is no feasible solution to our problem. The number of trucks necessary for feasibility depends on how far the fronts are from the mill and the times when the mill needs the loads. As the number of available trucks increases, we can find better solutions in that the average wait time per load decreases. A sufficient increase in the number of trucks guarantees we can pick up all the loads at their ready times so each load arrives at the mill exactly when needed. Any further increase in the number of trucks beyond that point cannot further reduce the wait time.

Thus, we present bounds on the number of trucks required to achieve feasibility and the number of trucks necessary to achieve a wait time of zero. A corollary of the upper bound is that it guarantees a solution for which each truck need serve only one front. This number of trucks provides a managerially attractive solution, but requires more trucks to achieve the same level of average wait time per load than the alternative in which we allow trucks to serve multiple fronts.

We begin by bounding the number of trucks necessary for feasibility. The amount of truck time associated with each load i depends on the front at which the load is harvested. For each truck t , t 's shift ends when the last load it transports arrives at the mill. Let end_t be that time for truck t such that $end_t = a_{(i|i \text{ is the last load picked up by truck } t)}$. The start time for the truck t is denoted by b_t .

Proposition 5. *The number of trucks needed to fulfill all the mill needs at the mill need times, $\{n_1, n_2 \dots n_N\}$, is bounded on the lower side by a positive number k such that k is the smallest integer satisfying*

$$\sum_{t \in 1 \dots k} (n_{N-t+1}) \geq \sum_{f \in \mathcal{F}} L_f r_f.$$

Algorithm 3 Heuristics to raise lower bound

Output: Lower bound to the integer solution
Initialization: $Wait = 0$
Input: A solution to P'
for $f = 1$ to F **do**
 for $k = 1, \dots, |\mathcal{R}_f|$ **do**
 if $k \neq |\mathcal{R}_f|$ **then**
 $i \leftarrow \tau_{k+1} - 1$
 else
 $i \leftarrow N$
 end if
 $Flag \leftarrow TRUE$
 while $i \neq \tau_k$ && $Flag = TRUE$ **do**
 $temp \leftarrow \infty$
 if $y'_{if} = 1$ **then**
 $Flag \leftarrow FALSE$
 $Next \leftarrow n_i - s_f - \bar{H}_f$
 for $j = i - 1, \dots, \tau_k + 1$ **do**
 if $y'_{jf} = 1$ **then**
 if $Next < n_j - s_f$ **then**
 $Wait \leftarrow Wait + n_j - s_f - Next$
 $Next \leftarrow Next - \bar{H}_f$
 else
 $Next \leftarrow n_j - s_f - \bar{H}_f$
 end if
 end if
 end for
 end if
 $i \leftarrow i - 1$
 end while
 end for
end for

Proof. Because of the assumption that the total number of loads harvested across all fronts is exactly the number of loads needed to fulfill the mill needs, $\sum_{f \in \mathcal{F}} L_f r_f$ is the same regardless of what need is served by what front. Further, we note that the minimum possible time needed to serve all loads, the time when the trucks do not wait at either the mill or the fronts, is the same sum $\sum_{f \in \mathcal{F}} L_f r_f$.

Then, for k available trucks, the maximum available truck time is

$$\begin{aligned} \sum_{t \in 1 \dots k} n_{(N-t+1)} &\geq \sum_{t \in 1 \dots k} n_{(N-t+1)} - \sum_{t \in 1 \dots k} b_t \\ &\geq \sum_{t \in 1 \dots k} a_{(N-t+1)} - \sum_{t \in 1 \dots k} b_t = \sum_{t \in 1 \dots k} \text{end}_t - \sum_{t \in 1 \dots k} b_t \\ &= \sum_{t \in 1 \dots k} (\text{end}_t - b_t). \end{aligned}$$

The second inequality follows from the fact that a load must arrive before the mill need time that it satisfies. The first equality holds by definition of end_t . The second equality rearranges terms and represents the time that k trucks needed to serve all of the mill needs.

Thus, the smallest k such that $\sum_{t \in 1 \dots k} n_{(N-t+1)} \geq \sum_{i \in 1 \dots N} \rho_i$ is the smallest number of trucks that could cover the minimum possible time needed to serve all loads. \square

If the mill needs are evenly spaced, we can rewrite $\sum_{t \in 1 \dots k} n_{(N-t+1)}$ as $k \times \left(\frac{n_N + n_{(N-k+1)}}{2} \right)$.

We also note that, by treating each front as a separate mill area, Proposition 5 allows us to compute a lower bound on the number of trucks required to serve each front with dedicated vehicles. We call this separability. We formally present the result in Corollary 2. The proof follows directly from Proposition 5 and is omitted.

Corollary 2. For $f \in \mathcal{F}$, let k_f be the smallest integer that satisfies

$$\sum_{t \in 1 \dots k_f} n_{(N-t+1)} \geq L_f r_f.$$

Then, $k = \sum_{(f \in \mathcal{F})} k_f$ is the minimum number of trucks that could possibly serve each front separately.

We next present an upper bound on the number of vehicles required to serve the loads. The proof of the upper bound requires the realization that, if each load is picked up at its ready time, the solution can be improved only by changing the harvest times.

Proposition 6. The upper bound on the number of trucks needed meet all mill needs among fronts is given by \hat{k} such that $\hat{k} = \sum_{(f \in \mathcal{F})} \hat{k}_f$ where $\hat{k}_f = \min \{L_f, \lceil \frac{r_f}{\bar{H}_f} \rceil\}$.

Proof. The shortest possible time between the ready times of two consecutive loads at front f is \bar{H}_f . To serve at their ready times loads whose ready times differ by \bar{H}_f , it is clear that we need at least $\hat{k} = \lceil \frac{r_f}{\bar{H}_f} \rceil$ trucks. For any front, however, we never need more than $\hat{k}_f = L_f$ trucks, because L_f trucks is enough to serve each load at front f with its own vehicle. If each load is served by its own vehicle, then a vehicle is always capable to serving the load at its ready time. \square

V.2 Datasets

Our computational analysis uses several datasets generated based on our conversations with our industry partners (Personal Communication with Jose Coelho, Sugar Cane Segment Manager, John Deere, April 13, 2010; Personal Communication with Craig Wenzel, Staff Engineer, Worksite Systems and Productivity Group, John Deere, May 28, 2010). All of the data is available from myweb.uiowa.edu/bthoa/iowa/Research.html.

Each instance represents a typical mill area operations in Brazil. A mill area has four to eight fronts. We generate two mill areas each for four through eight fronts.

The front to mill travel time is on average one and half times higher than the mill to front travel time. The increase in return trip time reflects the impact of a load trailer. In each mill area, the closest front has the trip time between 30 minutes and 50 minutes whereas the farthest front has the round trip time between 70 minutes and 120 minutes.

Each front also has a harvest quota, the number of loads to be produced from the front in the given problem instance. The sum of the harvest quotas across all fronts in an instance is equal to the number of the total mill needs. The closest front serves 5% to 10% of the total mill needs and the farthest front serves between 35% and 60%.

The minimum time required to harvest a load at each front is based on the number of loads each front serves. We first compute the time that would elapse if loads from a front were evenly spaced. We then assume that the minimum time is 70% of that time. The 70% reflects the fact that, in practice, a capacity cushion is used as reactive capacity. Given the previously described harvest quotas, the closest front has the highest minimum time required to harvest a load. The minimum harvest time of the closest front ranges from 15 minutes to 50 minutes. The farthest front has the lowest minimum harvest time which ranges between 10 minutes and 30 minutes.

In all our instances, all the fronts and the trucks are allowed to start their operations 100 minutes before the the mill's first need. If the mill's first need is at the 100th minute, all the fronts can potentially start harvesting at 0th minute and trucks are available to leave mill for a pick up at 0th minute. We call this gap between when fronts can start harvesting and the time of mill's first need the "warm-up period." This warm-up period is analogous to having a setup time in manufacturing setting and eases the construction of feasible datasets.

Each geography is solved for four different inter-mill need times, three, four, five, and 10 minutes, for a total of 480 loads. These inter-mill need times reflect the range of values that might be encountered. For each instance and inter-mill need time combination, we also solve for the number of trucks in the range (lower bound to upper bound) as calculated in Section V.1. Altogether, we generate 204 instances.

V.3 Experiments

In this paper, we seek to address three questions. First, what is the computational value of the proposed valid inequalities and initial solution heuristic in finding optimal solutions in reasonable computation times? Two, what is the value of coordinating vehicles across fronts rather than assigning vehicles to specific fronts for the entire horizon? Third, what is the value of allowing variable harvest rates?

We address the first question by running our datasets using the initial model. We then add the valid inequalities and finally combine the valid inequalities with the initial feasible solution. We note that we tried various subsets of the valid inequalities, but were able to achieve provably optimal solutions in reasonable runtime only by using the entire set. The results of these experiments are discussed in Section VI.1.

Our second question is motivated by our review of the literature. We found that, in sugarcane harvesting and transportation operations, one common practice was to assign the trucks to a single front for the entire horizon. The practice arises particularly in countries in which the industry is not vertically integrated. In such countries, the growers are responsible for the transport of cane to the mill, and consequently growers employ a dedicated set of vehicles. In Brazil and as modeled in this paper, however, the high level of vertical integration allows the coordination across the fronts. While coordination can reduce the number of trucks, separability decreases the managerial complexity of the operation. The results of these experiments are presented in Section VI.2.

Our third question is motivated by the knowledge, that in practice, harvesters are run at approximately 70% of their capacity. As noted previously, this extra capacity is used as reactive capacity, but on average this reactive capacity is unused. Thus, on average, the capacity cushion could also be used facilitate coordination, as we have modeled in this paper. However, there is a managerial challenge to such coordination, particularly when decisions are being made without automated decision support.

To assess the value of varying harvest rates to reduce cut-to-crush time, we consider a case in which harvest rates are fixed. We create fixed harvest rates by setting the harvest time at each front to the average harvest time required to meet each front's quota of loads. This choice of harvest rates causes infeasibility in most instances. We overcome this issue by increasing the warm-up period. To determine this warm-up period, we iterated through warm-up times seeking the lowest warm-up time that achieved feasibility. We did not consider the cases where only a subset of fronts' warm-up time is increased. Consequently, the average wait time that we report for the fixed harvest rates is using the minimum warm-up period needed to meet all the mill needs. The results of these experiments are presented in Section VI.3.

VI. COMPUTATIONAL RESULTS

This section presents the results of our computational experiments. We first demonstrate the value of our valid inequalities and the use of an initial feasible solution. We then explore the cost of the managerially attractive separable solution. Finally, we explore the value of variable harvest rates.

The math programs are solved using GUROBI OPTIMIZER 5.1. The experiments were performed on a 3.40 GHz Intel Core i7-3770 CPU running the Ubuntu 12.04 operating system. For all of the reported results, we implement the lifting heuristic described in Section IV.3 and Gurobi's relaxation-induced-neighborhood-search routine, called RINS. The lifting heuristic was implemented in C++ and communication with Gurobi was achieved through Gurobi's C++ Interface. In our initial experiments, these heuristics alone did not improve the performance of branch-and-bound, but proved valuable in proving optimality once the valid inequalities of Section IV.1 and the initial solution of Section IV.2 were implemented. The heuristics are run on all relaxed solutions for which 5% or fewer of the binary variables are fractional. The 5% value represents a compromise between the runtime of the heuristic and their value in reducing overall runtime. Branching was set to give priority to the front assignments or y variables. All runs were terminated when the optimality gap was 1% or less.

VI.1 Algorithmic Performance

Tables 1 through 4 present the results of the experiments testing the value of the valid inequalities and the initial feasible solution. As noted previously, without the valid inequalities, the instances rarely found a feasible solution and never return an optimal solution, even with significant runtime. Thus, the tables report the computation times for runs with just the valid inequalities and then

runs with both the valid inequalities and an initial feasible solution. We label these computation times as "VI" and "Both," respectively. For each instance, the table also reports the average wait time per load. We also report the runtimes for each instance. In almost all cases, we were able to prove optimality with reasonable runtimes. Instances marked with "*" are instances for which we could not prove optimality even with 10,000 seconds of runtime. In those cases, the reported result is the best found feasible solution with the integer gap reported in the brackets after 1,000 seconds of runtime.

Finally, for some instances, such as instance 5a with 24 trucks, we do not report any values. In these cases, the reported number of trucks is the number of trucks that achieves the lower bound on the number of trucks for the instance. However, even after 10,000 seconds of runtime, we were unable to find a feasible solution for such instances. Given that the lower bound presented in Section V.1 does not guarantee feasibility, we believe that it is likely that the lower bound is infeasible in these instances, but we were also unable to prove infeasibility. We mark these instances with "-".

The results in the tables demonstrate that valid inequalities alone are almost always able to achieve optimal solutions. In only ten cases out of 204 did the solver return a solution without being able to prove optimally in 10,000 seconds. Only four instances require more than 1,000 seconds to prove optimality.

In most cases, the addition of the initial feasible solution has a positive impact on runtimes. Using results for the instances with an inter-mill need time of 3, the initial feasible solution improves runtime by an average of almost 17% for the instances proved to optimality. We are also able to prove optimality for three cases for which it was not previously proven.

In terms of problem characteristics, run times increase in the number of fronts but decrease in the number of trucks. This result is not surprising. The problem size grows as the number of fronts grows, increasing runtime. However, the fronts are not the challenge in determining feasibility. Feasibility is driven by the trucks and their assignments. Thus, while the problem size also grows in the number of trucks, the computation time associated with the growth in problem size is overcome by the decrease that comes from the reduced challenge in finding feasible solutions and the resulting increase in fathoming in the branch-and-bound algorithm.

Finally, from a managerial perspective, it is also valuable to consider the trade-off between the number of trucks and the average wait of each load. As an example, Figure 2 presents a graph of the average waiting time per load and the number of trucks. It is clear that most of the reduction in wait time comes from the addition of one truck over the number required for feasibility. This relationship is evident throughout the instances. Given the lower bound is analytically computable and generally feasible, the result offers a potential "rule of thumb" for planning purposes.

Geography	# of trucks	Average wait	VI	Both
4a	29	25.00	* (7.4 %)	* (6%)
	30	22.63	728	600
	31	18.38	716	537
	32	11.225	599	454
	33	8.53	421	414
	34	4.79	297	274
	35	4.2	206	239
	36	1.36	193	181
	37	0.2	184	180
	4b	30	34.2	* (4.5%)
31		30.1	928	680
32		14.7	730	543
33		5.03	695	435
34		3.31	547	356
35		0.5375	539	281
36		0.2625	356	248

	37	0.09	310	217
5a	24	-	-	-
	25	23.05	846	827
	26	19.15	784	741
	27	13.31	516	497
	28	9.375	425	408
	29	5.13	394	342
	30	0.10	341	325
5b	16	17.02	909	855
	17	15.2	823	693
	18	7.6	735	672
	19	2.1	556	422
	20	0.5	556	523
	21	0.3	380	347
6a	17	14.41	768	711
	18	9.78	832	645
	19	4.25	564	518
	20	1.20	510	454
	21	0.425	398	386
	22	0.25	319	366
6b	17	17.39	660	646
	18	9.70	485	431
	19	4.17	470	422
	20	0.41	400	356
	21	0	205	190
7a	16	8.1	*(5.6%)	*(2.4%)
	17	6.98	*(1.1%)	956
	18	1.46	951	895
	19	1.10	813	608
	20	0.12	728	340
	21	0.033	567	361
	22	0	507	226
7b	26	-	-	-
	27	6.27	946	811
	28	4.13	909	840
	29	1.56	871	729
	30	0.67	723	626
	31	0.31	701	597
	32	0.17	664	501
8a	23	21.81	*(7.5%)	*(2.6%)
	24	12.3	992	889
	25	9.22	831	714
	26	6.19	982	648
	27	4.27	765	682
	28	1.29	722	561
	29	0.54	710	513
	30	0.44	557	340
	31	0.18	522	296
8b	31	9.56	1536	1077
	32	6.725	937	918
	33	2.8	639	581
	34	1.70	599	394
	35	0.95	564	336
	36	0.67	465	432

Table 1: Intermill-Needtime = 3, 480 loads

Geography	# of trucks	Average wait	VI	Both
4a	22	27.39	*(5.3%)	*(2.2%)
	23	7.51	924	876
	24	3.26	780	764
	25	2.84	543	510
	26	1.45	388	334
	27	0.416	306	276
	28	0.0625	289	254
4b	23	26.32	*(1.1%)	1030
	24	10.9	952	873
	25	8	928	715

	26	5.07	837	680
	27	0.996	782	490
	28	0.375	650	406
	29	0.083	468	362
	30	0	355	227
5a	18	-	-	-
	19	22.34	* (1.9%)	* (1.7%)
	20	16.28	471	401
	21	8.125	418	380
	22	4.08	391	358
	23	0.13	184	179
5b	12	18.81	879	776
	13	7.56	742	708
	14	0	194	182
6a	13	17.70	878	637
	14	11.28	707	581
	15	1.25	462	497
	16	0	256	201
6b	13	10.90	985	917
	14	1.50	446	372
	15	0.19	258	229
	16	0	218	231
7a	13	5.99	* (3.2%)	* (1.6%)
	14	4.01	1103	905
	15	2.64	845	706
	16	1.27	615	543
	17	0.96	638	566
	18	0.85	517	421
7b	20	22.9	* (1.5%)	812
	21	9.09	885	711
	22	1.85	685	547
	23	0.57	600	433
	24	0.29	610	306
8a	17	25.05	* (2.1 %)	* (1.8 %)
	18	16.96	903	877
	19	4.97	850	762
	20	1.575	776	655
	21	0.85	680	519
	22	0.33	585	458
	23	0.21	617	364
8b	21	30.86	* (4.7 %)	* (3.8 %)
	22	22.02	879	764
	23	11.18	709	676
	24	8.15	760	619
	25	2.47	469	246
	26	1.17	342	221
	27	0.09	305	197

Table 2: Intermill-Needtime = 4, 480 loads

Geography	# of trucks	Average wait	VI	Both
4a	17	28.44	970	879
	18	14.92	575	617
	19	8.89	538	512
	20	4.95	421	418
	21	0	272	259
4b	19	26.02	1811	1750
	20	15.01	945	932
	21	2.93	863	606
	22	0.5875	717	596
	23	0.35	532	356
	24	0	547	301
5a	15	21.09	946	803
	16	15.37	734	613
	17	0	342	303
5b	9	-	-	-
	10	10.89	818	713

6a	11	0.47	602	510
	11	10.59	917	835
	12	1.08	771	625
	13	0.16	337	299
6b	14	0.04	261	240
	10	26.01	963	872
	11	17.15	876	774
	12	8.68	677	531
7a	13	0.22	668	594
	11	4.15	*(1.5%)	*(1.2%)
	12	2.21	630	512
	13	0.575	587	369
7b	14	0.125	557	290
	16	19.62	1708	1286
	17	12.41	962	855
	18	1.83	878	729
	19	1.22	749	561
	20	1.05	701	473
8a	21	0.34	623	414
	13	19.87	1025	901
	14	14.80	816	785
	15	4.66	728	632
	16	2.24	714	563
8b	17	0.64	698	501
	18	0.18	688	317
	17	12.66	1757	1491
	18	9.56	932	859
	19	5.39	888	702
	20	1.67	867	623
	21	1.01	607	434
22	0.2	556	406	

Table 3: Intermill-Needtime = 5, 480 loads

Geography	# of trucks	Average wait	VI	Both
4a	9	51.875	788	714
	10	0.5	335	321
	11	0	206	194
4b	9	54.8375	992	717
	10	16.06	863	600
	11	0.746	620	596
	12	0.416	477	218
5a	8	17.37	833	621
	9	0	247	202
5b	5	4.10	768	605
	6	0	200	196
6a	6	7.51	451	319
	7	0.67	298	284
	8	0.32	290	264
	9	0.16	244	231
6b	5	17.75	953	786
	6	1.47	813	802
	7	0.1	342	343
7a	5	7.03	*(6.4 %)	*(2.1 %)
	6	1.86	812	726
	7	0.325	875	574
	8	0	465	257
7b	10	10.35	706	634
	11	2.6	517	369
	12	0.02	408	321
8a	7	3.22	917	573
	8	0	267	202
8b	9	9.56	856	632
	10	3.7	800	665
	11	0.48	711	533

Table 4: Intermill-Needtime =10, 480 loads




Figure 2: Relationship between Number of available trucks and average wait times (4a- 4 mins)

VI.2 Cost of a Separable Solution

Table 5 presents the results of the experiment to demonstrate the value of coordination. For each geography and intermill need time, the table presents the number of trucks needed to find a feasible separable solution, or a solution for which the trucks serve only one front, and the number of trucks needed to find a feasible solution when coordination across fronts is allowed. These columns are labeled "Separable" and "Coordination," respectively.

As the table shows, on average, when coordination is allowed, almost 27% fewer vehicles are needed to achieve feasibility compared to the number needed to achieve feasibility when requiring separability. This result is consistent regardless of the number of fronts. However, the value of coordination increases as the time between mill needs increases. When the intermill need time is three, the average difference in the number of trucks is 18.1%. At an intermill need time of 10, the difference increases to 40.3%.

Geography	InterMill Need Time	Separable	Coordination
4a	3	32	29
	4	27	22
	5	21	17
	10	14	9
4b	3	34	30
	4	26	23
	5	23	19
	10	13	9
5a	3	32	25
	4	24	19
	5	20	15

	10	11	8
5b	3	20	16
	4	17	12
	5	13	10
6a	10	10	5
	3	22	17
	4	17	13
6b	5	14	11
	10	10	6
	3	21	17
7a	4	17	13
	5	15	11
	10	9	5
7b	3	19	16
	4	18	13
	5	16	11
8a	10	9	5
	3	33	27
	4	25	20
8b	5	22	16
	10	16	10
	3	39	23
8a	4	24	17
	5	19	13
	10	13	7
8b	3	38	31
	4	29	21
	5	23	17
	10	16	9

Table 5: Number of trucks needed for separable

VI.3 Value of Variable Harvest Rates

Table 6 presents the results of our experiments designed to demonstrate the value of variable harvest rates. For each instance, geography and intermill need time, we present the number of trucks required to achieve a feasible solution when harvest rates are constant, called "Const Trucks." We also presented the average wait time per load associated with this number of trucks, called "Const Avg Wait." Finally, We In Table 6 we present the number of trucks required to achieve a zero wait time using the variable harvest rates, called "Var Trucks."

As the table shows, allowing variable harvest rates leads to solutions that have often considerably less wait time per load, almost 24 minutes on average, than the constant harvest rate solutions. Further, the constant harvest rate solutions require 12% more trucks than the variable harvest rate case to pick up all of the loads at their ready times. In only one case does the constant harvest rate case achieve feasibility with fewer trucks than the number required to achieve no wait with variable harvest rates.

Geography	InterMill Need Time	Const Trucks	Const Avg Wait	Var Trucks
4a	3	39	43	37
	4	31	19	28
	5	24	26	21
	10	13	21	11
4b	3	42	20	39
	4	31	44	30
	5	24	16	21
5a	10	14	15	12
	3	34	64	30
	4	25	38	23
	5	20	26	17
	10	11	45	9

5b	3	22	23	21
	4	17	15	14
	5	14	11	11
	10	9	19	6
6a	3	24	9	22
	4	19	9	16
	5	14	5	14
	10	8	18	9
6b	3	23	13	21
	4	18	17	16
	5	15	23	13
	10	9	29	7
7a	3	24	28	22
	4	19	13	18
	5	16	20	14
	10	8	33	8
7b	3	34	10	32
	4	26	17	24
	5	22	25	21
	10	13	35	12
8a	3	34	20	31
	4	25	18	23
	5	20	15	18
	10	12	25	8
8b	3	43	36	36
	4	33	29	27
	5	25	23	22
	10	13	15	11

Table 6: Number of trucks needed to pick up loads at ready times and the corresponding average wait times

VII. CONCLUSION

In this paper, we introduce the sugarcane harvest logistics problem for Brazil. We introduce valid inequalities and a heuristic to generate a feasible initial solution. Our results show that the presented valid inequalities and heuristic for finding feasible solutions provide significant computational advantages. We also introduce a lifting heuristic that was advantageous in proving optimality. We further demonstrate that coordinating the trucks across fronts offers the opportunity to dramatically reduce the number of vehicles needed to serve the mills. Finally, our results show that variable harvest rates reduce cut-to-crush times while also reducing the number of vehicles needed to serve the loads.

There are two important directions for future work. The first direction builds on the work of Salassi et al. (2009a) that explores a sugarcane harvest logistics problem in the United States. While there are key differences between Brazilian and US-based operations, most notably the fact that harvests in Brazil run 24 hours a day, the US operations have many more farms involved in a daily harvest operation. Even with the enhancements introduced in this paper, the size of the US problem requires additional research.

Second, the problem introduced here is deterministic. Not surprisingly, a harvest operation has a number of uncertainties. Importantly, there are isolated weather events and breakdowns that require attention. Yet, this research is an important foundation for the stochastic case. For one, the research presented here can be used to solve for perfect information solutions for the stochastic case. These solutions can be helpful for evaluating heuristic approaches. Second, the methods here can be valuable in a rolling horizon procedure. When one considers a horizon of four to six hours rather than 24 as we have here, the proposed solution method runs in a few seconds and can thus be amenable for use in real time.

Acknowledgements We would like to thank John Deere for financial support. We are also grateful to Craig Wenzel and Brian Gilmore for their support of this research and their help in developing our knowledge of sugarcane harvests and logistics. We also thank Jean-Francois Cordeau for C++ code that provided helpful insight for our own implementation.

REFERENCES

- O. Ahumada and J.R. Villalobos. Application of planning models in the agri-food supply chain: A review. *European Journal of Operational Research*, 196(1):1–20, 2009. ISSN 0377-2217.
- Sergio Barros. GAIN Report: Brazil Semi-Annual Sugar 2012. Technical Report BR12014, USDA Foreign Agricultural Service, October 2012. URL http://gain.fas.usda.gov/RecentGAINPublications/SugarSemi-annual_SaoPauloATO_Brazil_10-1-2012.pdf.
- Aneirson Francisco da Silva, Fernando Augusto Silva Marins, and José Arnaldo Barra Montevechi. Multi-choice mixed integer goal programming optimization for real problems in a sugar and ethanol milling company. *Applied Mathematical Modelling*, 37(9):6146–6162, May 2013. ISSN 0307904X. doi: 10.1016/j.apm.2012.12.022. URL <http://linkinghub.elsevier.com/retrieve/pii/S0307904X13000061>.
- J.A. Diaz and H.G. Perez. Simulation and optimization of sugar cane transportation in harvest season. *Winter Simulation Conference*, 2:1114–1117, 2000.
- Economic Research Service. Sugar and sweeteners yearbook tables. Technical report, United States Department of Agriculture, 2012. URL <http://www.ers.usda.gov/data-products/sugar-and-sweeteners-yearbook-tables.aspx>.
- RC Giles. A simulation study of cane transport system improvements in the Sezela Mill area. Master’s thesis, School of Bioresources Engineering and Environmental Hydrology, University of KwaZulu-Natal, Pietermaritzburg, South Africa, December 2009.
- M. Grunow, H.-O. Gunther, and R. Westinner. Supply optimization for the production of raw sugar. *International Journal of Production Economics*, 110:224–239, 2008.
- AC Hansen, AJ Barnes, and PWL Lyne. Simulation modeling of sugarcane harvest-to-mill delivery systems. *Transactions of the ASABE*, 45(3):531–538, 2002.
- E.O. Heady. Simplified presentation and logical aspects of linear programming technique. *Journal of Farm Economics*, 36(5):1035–1048, 1954.
- A. Higgins. Scheduling of road vehicles in sugarcane transport: A case study at an Australian sugar mill. *European journal of operational research*, 170(3):987–1000, 2006. ISSN 0377-2217.
- A. Higgins and I. Davies. A simulation model for capacity planning in sugarcane transport. *Computers and Electronics in Agriculture*, 47:85–102, 2005.
- a J Higgins and L a Laredo. Improving harvesting and transport planning within a sugar value chain. *Journal of the Operational Research Society*, 57(4):367–376, July 2006. ISSN 0160-5682. doi: 10.1057/palgrave.jors.2602024. URL <http://www.palgrave-journals.com/doi/10.1057/palgrave.jors.2602024>.

- A J Higgins, R C Muchow, A V Rudd, and A W Ford. Optimising harvest date in sugar production : A case study for the Mossman mill region in Australia I . Development of operations research model and solution. *Field Crops Research*, 57:153 – 162, 1998.
- Andrew J. Higgins and Steve Postma. Australian sugar mills optimise siding rosters to increase profitability. *Annals of Operations Research*, 128:235–249, 2004. 10.1023/B:ANOR.0000019107.68291.36.
- C. Hildebrand. Independent assessment of the sugar industry. Report to the Hon. Warren Truss MP, Minister for Agriculture, Fisheries and Forestry, available from http://www.apec.org.au/docs/06_TP_Sugar/04_report.pdf, June 2002.
- Sanjay Dominik Jena and Marcus Poggi. Harvest planning in the brazilian sugar cane industry via mixed integer programming. Technical report, Departamento de Informatica, Pontificia Universidade Catolica do Rio de Janeiro, Rio De Janeiro, 2013. URL ftp://139.82.16.194/pub/docs/techreports/13_02_jena.pdf.
- M S Kawamura, D P Ronconi, and H Yoshizaki. Optimizing transportation and storage of final products in the sugar and ethanol industry : a case study. *International Transactions in Operational Research*, 13:425–439, 2006.
- A.M. Kostin, G. Guillén-Gosálbez, F.D. Mele, M.J. Bagajewicz, and L. Jiménez. A novel rolling horizon strategy for the strategic planning of supply chains. Application to the sugar cane industry of Argentina. *Computers & Chemical Engineering*, 35(11):2540–2563, November 2011. ISSN 00981354. doi: 10.1016/j.compchemeng.2011.04.006. URL <http://linkinghub.elsevier.com/retrieve/pii/S0098135411001402>.
- P.Y. Le Gal, J. Le Masson, C. N. Bezuidenhout, and L. F. Lagrange. Coupled modelling of sugarcane supply planning and logistics as a management tool. *Computers and Electronics in Agriculture*, 68:168–177, 2009.
- C. Lejars, P.Y. Le Gal, and S. Auzoux. A decision support approach for cane supply management within a sugar mill area. *Computers and Electronics in Agriculture*, 60(2):239–249, 2008. ISSN 0168-1699.
- E. Lopez Milan, S. Miquel Fernandez, and L.M. Pla Aragones. Sugar cane transportation in Cuba, a case study. *European Journal of Operational Research*, 174(1):374–386, 2006.
- Esteban Lopez-Milan and Lluís Pla-Aragones. A decision support system to manage the supply chain of sugar cane. *Annals of Operations Research*, to appear, 2013.
- T.J. Lowe and P.V. Preckel. Decision technologies for agribusiness problems: A brief review of selected literature and a call for research. *Manufacturing & Service Operations Management*, 6(3):201–208, 2004.
- Francois Margot. Symmetry in integer linear programming. In Michael Junger, Thomas M. Lieblich, Dennis Naddef, George L. Nemhauser, William R. Pulleybank, Gerhard Reinelt, Giovanni Rinaldi, and Laurence A. Wolsey, editors, *50 Years of Integer Programming*, chapter 17, pages 647–686. Springer-Verlag, Berlin, 2010.
- M. McConnell, E. Dohlman, and S. Haley. World sugar price volatility intensified by market and policy factors. *Amber Waves*, 8(3):28 – 35, 2010.

- B. C. McDonald, E. Dube, and C. N. Bezuidenhout. TranSwarm: A sugarcane transport simulation model based on behavioural logistics. In *Proceedings of the South African Sugar Technologists' Association*, volume 81, pages 434–438, 2008.
- Marco Fava Neves, Vinicius Gustavo Trombin, and Matheus Alberto Consoli. Measurement of sugar cane chain in brazil. *International Food and Agribusiness Management Review*, 13(3):37–54, 2010.
- Rafael P. O. Paiva and Reinaldo Morabito. An optimization model for the aggregate production planning of a Brazilian sugar and ethanol milling company. *Annals of Operations Research*, 169(1):117–130, September 2009. ISSN 0254-5330. URL <http://www.springerlink.com/index/10.1007/s10479-008-0428-9>.
- K. Prichanont, S. Prichanont, and N. Buransri. Improvement guidelines for sugar cane delivery systems. In *35th International Conference on Computers and Industrial Engineering*, pages 1585–1590, 2005.
- M. E. Salassi, F. G. Barker, and M. A. Deliberto. Optimal scheduling of sugarcane harvest and mill delivery. *Sugar Cane International*, 27(3):87 – 94, 2009a.
- M. E. Salassi, F. G. Barker, and M. A. Deliberto. Optimal scheduling of sugarcane harvest and mill delivery. *Sugar Cane International*, 27(3):87–90, 2009b.
- Michael E Salassi, Mercedes Garcia, and Janis B Breaux. Impact of Sugarcane Delivery Schedule on Product Value at Raw Sugar Factories. *Journal of Agribusiness*, 22(1):61–75, 2004.
- M. Saska, S. L. Goudeau, and M. Marquette. Determination of sucrose lose in clean unburnt billet cane. *Journal of the American Society of Sugar Cane Technologies*, 9:53 – 77, 2009.
- P. Saxena, R. P. Srivastava, and M. L. Sharma. Impact of cut to crush delay and bio-chemical changes in sugarcane. *Australian Journal of Crop Science*, 4(9):692–699, 2010.
- S. Supsomboon and J. Yosnual. A stochastic model for optimal order quantity for sugarcane factory with uncertainties. In *Proceedings of the Fifth Asia Pacific Industrial Engineering and Management Systems Conference*, pages 21.2.1–21.2.11, 2004.
- Constanza Valdes. Can brazil meet the world's growing need for ethanol? *Amber Waves*, 9(4):38 – 45, 2011.
- Alexandra Wexler. U.S. Throws Gas on Sugar Market. *The Wall Street Journal*, page C1, December 17, 2012.