Research in Transportation Economics 42 (2013) 3–12

Contents lists available at SciVerse ScienceDirect

Research in Transportation Economics

journal homepage: www.elsevier.com/locate/retrec

Modeling yard crane operators as reinforcement learning agents

Fateme Fotuhi [a], Nathan Huynh [a,*], Jose M. Vidal [b], Yuanchang Xie [c]

[a] Department of Civil and Environmental Engineering, University of South Carolina, 300 Main Street, Columbia, SC 29208, USA
[b] Department of Computer Science and Engineering, University of South Carolina, Swearingen Center, Columbia, SC 29208, USA
[c] Department of Civil and Environmental Engineering, University of Massachusetts Lowell, One University Avenue, Lowell, MA 01854, USA

ARTICLE INFO

Article history:
Available online 3 January 2013

JEL:
R41
R42
L92
O18

Keywords:
Reinforcement learning
Q-learning
Multi-agent systems
Yard crane scheduling
Drayage operations

ABSTRACT

Due to the importance of drayage operations, operators at marine container terminals are increasingly looking to reduce the time a truck spends at the terminal to complete a transaction. This study introduces an agent-based approach to model yard cranes for the analysis of truck turn time. The objective of the model is to solve the yard crane scheduling problem (i.e. determining the sequence of drayage trucks to serve to minimize their waiting time). It is accomplished by modeling the yard crane operators as agents that employ reinforcement learning; specifically, q-learning. The proposed agent-based, q-learning model is developed using Netlogo. Experimental results show that the q-learning model is very effective in assisting the yard crane operator to select the next best move. Thus, the proposed q-learning model could potentially be integrated into existing yard management systems to automate the truck selection process and thereby improve yard operations.

1. Introduction

Truck turn time refers to the time it takes a drayage truck to complete a transaction such as picking up an import container or dropping off an export container at a terminal; drayage refers to the trucking movements linking seaport terminals with importers, exporters, and rail terminals. Specifically, truck turn time is the difference between the truck's exit time and the truck's entry time. It is a measure of a terminal's efficiency in receiving and delivering containers. For this reason, terminals are constantly striving to reduce their terminals' truck turn time. Fig. 1 shows the breakdown of truck turn time by process: entry gate, yard, and exit gate, for a major container terminal in the U.S. These data highlight the need to improve container terminals' yard operations. This paper discusses a methodology for improving the yard operations by enhancing the yard cranes' service strategy. The challenges associated with this particular endeavor are the limited number of yard cranes available to serve a high number drayage trucks, the random nature of drayage truck arrivals, and the variation in work load among yard blocks.

Improving the yard operations requires the combination of (1) decreasing the ratio of trucks to yard cranes, (2) improving the yard crane service strategy, and (3) improving the yard crane rehandling strategy. The ratio of trucks to yard cranes can be reduced by purchasing more yard cranes; however, the high initial investment, plus expensive maintenance and operating costs of these cranes often prohibit terminals from pursuing this option. Improving the yard crane service strategy and rehandling strategy is more practical and less costly for terminal operators. This study aims to improve the service strategy; that is, the yard crane scheduling problem (YCSP). The YCSP addressed in this study deals with the strategy the cranes employ in selecting the next truck to provide service to. Consider the scenario illustrated in Fig. 2, there are several practical ways in which the crane could go about servicing the four trucks; there are technically 24 possible sequences, but some of these sequences are not practical (e.g. serving truck 2, 3, 1, and then 4). One possible sequence is Truck 2, Truck 3, Truck 4, and Truck 1. Another possible sequence is Truck 2, Truck 1, Truck 3, and Truck 4. It is also possible for the crane to serve the trucks in this order: Truck 4, Truck 3, Truck 2, and Truck 1. It is evident that some sequences will be more effective than others in minimizing the total waiting time of the four trucks. Thus, the objective of the YCSP is to find the best sequence of trucks to serve to minimize the total truck waiting time.

The yard cranes are operated by operators who are given some general guidelines, but who have the freedom to make judgment

* Corresponding author. Tel.: +1 803 777 8947; fax: +1 803 777 0670.
E-mail addresses: fateme.fotuhi@gmail.com (F. Fotuhi), huynhn@cec.sc.edu, HUYNHN@mailbox.sc.edu (N. Huynh), jmvidal@gmail.com (J.M. Vidal), yuanchang_xie@uml.edu (Y. Xie).

0739-8859/$ – see front matter
http://dx.doi.org/10.1016/j.retrec.2012.11.001

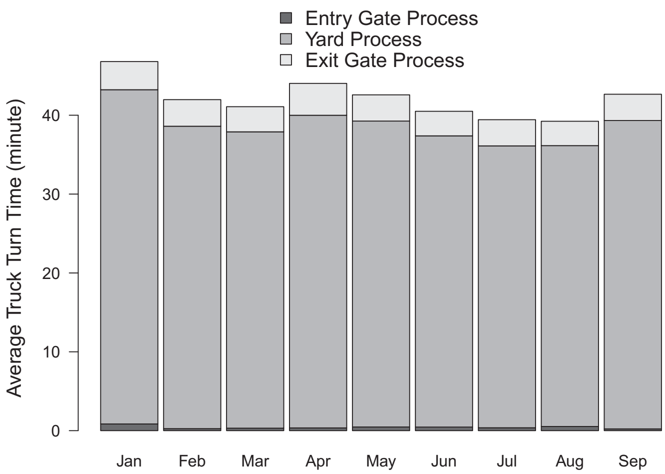

Fig. 1. Truck turn time by process for a major U.S. container terminal in 2009.

calls on how to move about the yard serving drayage trucks. The reason crane operators are given flexibility in deciding their work sequence is because at any given time, the yard cranes support multiple operations: vessel, drayage trucks, warehouse moves, rail moves, and inspection moves. There is not a centralized system that coordinates all of these movements and it is not known in advance when these activities will occur during the day. For these reasons, terminal operators will let crane operators decide what's best based on their knowledge of what is happening at the moment. Crane operators are directed by the foreman when they need to deviate from their usual plans. At the Port of Houston, the crane operators in general follow a distance-based strategy. That is, they favor serving the closest waiting truck. As illustrated in Fig. 3, cranes at Port of Houston will first look to serve trucks waiting in the same yard block. As the crane travels in one direction (toward a truck) it will serve all trucks in the direction of travel. If there are no trucks waiting in the same yard block the crane is on, the crane operator first checks with the container yard management system to see if there are trucks currently at the gate coming to his block. If there are, he stays put. If not, he will look for trucks in other blocks waiting for service. He will only travel to the truck in another yard block if he is the closest crane that is available.

At the Port of Charleston, the yard cranes follow a similar service strategy as that at the Port of Houston. However, there are two distinct differences. The first difference is that at Charleston, the waiting time of a truck is factored into the crane operator's decision; crane operators can see when a truck is gated in via their computers that run the terminal's yard management system. That is, the crane operators at Charleston may skip nearby trucks to go and serve a truck that has excessively long waiting time. The second difference is that at Charleston, the cranes follow a"sweeping" strategy, meaning they prefer to move continuously from one end of the yard block to the other and serve most trucks along the way.

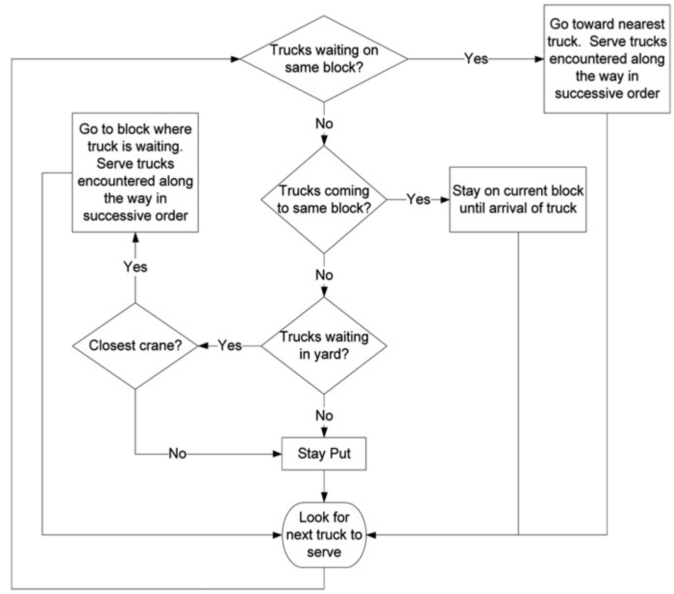

Fig. 3. Port of Houston yard crane service strategy.

Note that not all trucks in the sweeping direction of the cranes will be serviced. The exceptions are depicted in Fig. 4. The two scenarios shown in Fig. 4 illustrate the situations when the operator would abandon the sweeping strategy (to serve the longer waiting trucks first). When the operator does these exceptions, he is following a time-based strategy (i.e. serve the longest waiting truck first).

The objective of this paper is to develop a model of yard cranes that can learn about their environment and adapt to changing conditions such that over time they will converge to an optimal service strategy. It is motivated by the need for a guidance system that instructs each crane operator which drayage truck to serve next. Such a centralized system is expected to yield better overall performance than the system where each crane operator makes autonomous and independent decisions. The developed model draws on theories from multi-agent systems and reinforcement learning. It implements a multi-agent system in which each crane operator is an agent who utilizes reinforcement learning; specifically, a q-learning model. Q-learning is a technique that works by learning an action-value function that gives the expected utility of taking a given action in a given state. It has been proven to be effective for models with finite state and action space. Q-learning models have been applied in various industries such as robotics, elevator control, network routing and finance, but its application has never been explored in marine terminals. To this end, this study seeks to investigate the applicability and usefulness of q-learning models as a method to solve the YCSP. The developed framework can be extended to other applications at marine terminals as well as inland intermodal facilities.

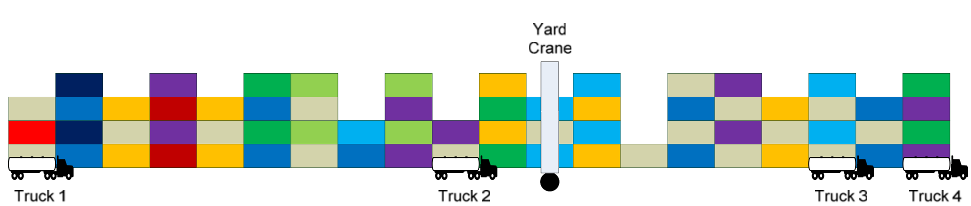

Fig. 2. Illustration of yard crane scheduling problem.