Computers & Operations Research 66 (2016) 20–28

Contents lists available at ScienceDirect

Computers & Operations Research

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

A rescheduling and cost allocation mechanism for delayed arrivals

Edward C. Rosenthal *, Eric M. Eisenstein

Department of Marketing and Supply Chain Management, Fox School of Business, Temple University, Philadelphia, PA 19122 USA

ARTICLE INFO

Available online 12 August 2015

Keywords:
Production scheduling
Vickrey–Clarke–Groves mechanism
Assignment problem
Combinatorial exchange

ABSTRACT

We propose a solution to the problem of rescheduling a sequence of arrivals that are subject to a delay event at a common destination. Such situations include jobs arriving at a single production facility, aircraft whose landings are postponed, and ships that are inbound to a dock or lightering facility. Each arrival faces a nonlinear cost due to the delay, but the delay costs can be mitigated by allowing the arrivals to be reordered. We optimize the reordering process by designing a Vickrey–Clarke–Groves (VCG) mechanism to construct a payoff matrix describing the amounts necessary to move the currently assigned arrival slots either earlier or later. Using this payoff matrix, we compute the optimal reordering of the arrivals by utilizing the well-known solution to the assignment problem, which maximizes the benefit in a computationally efficient fashion. The VCG mechanism is strategyproof, that is, no arrival has an incentive to misreport the value of moving up or down in the sequence. We also show that participating in the centralized process is to no arrival's disadvantage. Because VCG procedures in general are subject to budget deficits, we provide alternative mechanisms to overcome this difficulty. Finally, we carry out computational experiments demonstrating that the VCG mechanism can be implemented for realistically-sized problem sets and that the cost savings are significant.

1. Introduction

The problem of rescheduling arrivals at a processing facility occurs frequently in a number of applications. In manufacturing, there may exist multiple customers that have contracted for items that are produced sequentially at a single plant, where each customer relies on a prespecified delivery date. In the airline industry, inclement weather and mechanical failure cause flight arrivals to be delayed. The maritime industry faces similar problems with delays at loading docks and lightering facilities. In all of these applications, the arrivals typically face a nonlinear cost over the delay period, where the costs for some arrivals mount more rapidly than for others.

Given a significant delay, when some (if not all) of the arrivals are unable to be processed at their prespecified times, their differences in cost structures allow certain arrivals to benefit more than others from a reordering, or permutation, of the original sequence. Although the problem of how to best reorder the arrival sequence has been studied in the literature, we propose a new approach, based on soliciting monetary bids from the arrivals to pay to move up, or to be paid to move down, in the arrival sequence. There are four main contributions of this work: one, the introduction of a Vickrey–Clarke–Groves (VCG) market mechanism utilizing the solicited bids, which is group optimal and individually rational; two, that disjoint sets of arrivals have an incentive to join together, because the joint benefit that results is at least as large as the sum of their separate benefits; three, that our procedure is strategyproof, i.e., that no arrival has any incentive to misrepresent its cost structure; and four; the implementation of a standard, computationally efficient assignment algorithm (as part of the VCG procedure) to reschedule the arrivals so as to minimize the total cost of the delay.

The remainder of this paper is organized as follows. After covering the relevant literature in the rest of this section, in Section 2 we define the arrival rescheduling problem as well as the notion of a feasible solution to this problem. We then investigate the underlying economic structure of the problem and use cooperative game theory to show, contrary to typical auction situations, that there exists a well-defined synergy among the different customers. In Section 3 we define the information structure to be used between the various agents and the controller and then present the VCG procedure. In Section 4 we describe alternative procedures that compensate for budget shortfalls that can occur under VCG. In Section 5 we report on a set of computational runs of the mechanism to demonstrate its feasibility and cost savings. Finally, in Section 6, we discuss some potential pitfalls of our approach along with suggestions on how to best implement it, followed by some concluding remarks.

* Corresponding author. Tel.: +1 215 204 8177.
E-mail address: edward.rosenthal@temple.edu (E.C. Rosenthal).

http://dx.doi.org/10.1016/j.cor.2015.07.016
0305-0548/

1.1. Background and sequencing literature

There is a great variety of literature from different fields that speaks to the general problem of scheduling arrivals so as to minimize their costs. First, we note the cooperative nature of the problem. Consider a set of arrivals, processed according to a prespecified sequence, which are subject to a delay. If *side payments* – monetary transfers – among the arrivals are allowed, then it is possible that they can all benefit if some arrivals that were originally scheduled later in the sequence were allowed, through such payments, to move up in the queue by switching places with some of the earlier ones.

The literature on *sequencing games* treats some of the issues in question. Curiel et al. [1] and Slikker [2] examine scheduling situations with one processor and n sequential users. Assuming linear costs, these works utilize cooperative game theory to allocate the benefits and reorder the users, but do not consider the incentive problems of misreporting the cost functions.

Specific industry examples of our problem have been studied. In manufacturing, there may exist multiple customers that have contracted for items with prespecified delivery dates. Van Mieghem [3] finds that delay cost functions in such situations (holding costs, lost future sales, damaged reputation) are typically nonlinear. In the maritime industry, delays can occur at ports of call, where arriving vessels face demurrage and other costs which can greatly vary. Kao et al. [4], Kao and Lee [5], and Suh and Lee [6] have considered scheduling vessels in order to minimize demurrage costs but they do not treat the incentive problem of truthful reporting.

There is a considerable literature on airport recovery from a delay. Vasquez-Marquez [7] presents a decision support procedure to reschedule delayed flight arrivals. However, the procedure cannot be extended to multiple carriers without facing the problem of truthful cost reporting. Filar et al. [8] examine the literature from a number of standpoints. The categories most relevant to the present paper are aircraft landing sequence and perturbation due to ground delay. Numerous papers [9–16] have studied flight delay and recovery, but none of these papers consider transfer payments among the carriers or mention the problems inherent with private information from different carriers.

1.2. Incentives in queues

A large body of research considers incentive problems for customers in a queue (or arriving to join one) to be processed sequentially at a single facility [17–20]. In general these papers assume linear costs and treat the queue as one in which random arrivals join or bid to join upon observing queue length, unlike the present study.

Dolan [21] was the first to propose the use of a VCG mechanism in an arrival context, but once again, customers arrive according to an underlying stochastic process and are not already scheduled. This stochastic arrival process also appears in [22–26]. In [27,28] carriers negotiate landing slots prior to the event. In our study, however, the queue is actually a predetermined sequence in which arrivals have a fixed position at the outset. Whether the set of jobs is fixed or is allowed to grow, all of the previous studies allow the jobs to be arbitrarily reordered. While we do not assume that the processing facility is contractually obligated to meet the prespecified deadlines, we do assume that the delivery times were meant to be honored. Therefore, unlike the previous literature, we assume that no arrival can be "bumped" from its position in the processing sequence without its consent, and this further implies that if an arrival is additionally delayed, then they must be compensated through a side payment.

1.3. Auctions, airline delay, and combinatorial exchanges

The sealed-bid process that we use has its origins in the well-known result of Vickrey [29], who originated the second-price, sealed-bid auction where the highest bidder wins but only has to pay the second highest price for a good. He showed that truthful bidding is a dominant strategy for such auctions. Clarke [30] and Groves [31] generalized the Vickrey procedure to auctions with multiple items. The idea behind such generalized Vickrey auctions, or VCG auctions, is that each bidder will pay the externality that his bidding imposes on the other bidders, or, put differently, each winning bidder will pay the amount of her bid but will be refunded the increase in the value of the objective function that is due to her participation [32]. The outcomes of VCG mechanisms are individually rational [33], but Rothkopf [32] provides a number of reasons for their lack of practicality.

There has been some recent, relevant work concerning ground delay for airlines, as well as auctions and combinatorial exchanges. For ground delay problems, Vossen and Ball [34] use a centralized procedure to assign arrivals to arrival slots; they point out, however, that their algorithm is not a market procedure with side payments and they in fact encourage further research to this end. Ball et al. [28] consider the use of an auction to allocate arrival slots to particular carriers on a given day to cope with ground delays. But they do not offer a systematic mechanism that will adjust arrival schedules dynamically, i.e., in real time, or incorporate side payments, for a subset of carriers. Castelli et al. [35] develop a mathematical program to assign flights to landing slots for European air traffic management problems in which there is an imbalance between landing slot supply and demand. Their market-based mechanism is applied to a quite similar problem to the one we study, but is not strategyproof and therefore the reports their mechanism solicits may be manipulated. Other literature on ground delay programs and related problems comes under the name of collaborative decision making, or CDM (see, for example [36]), but we are not aware of any such work that develops a market mechanism as in the present paper.

A number of papers study VCG mechanisms in general combinatorial exchanges. Parkes et al. [37] develop a VCG-based approach to treat a broad class of combinatorial exchange problems, but they do not seek strategyproofness, and it is unclear how to implement their approach in a real-time environment such as the one we study. Other papers in this area are [33,38–40]. These papers deal with surplus redistribution under VCG schemes in general. In general, it is hard to predict how these mechanisms would work for the arrival rescheduling problem.

2. Problem definition and economic structure

Given a set of arrivals with a predetermined processing sequence, we will elicit monetary amounts that arrivals are willing to pay and the amounts that they require to receive in order to move up or down, respectively, in the sequence. We restrict our rescheduling mechanism to use private ("sealed-bid") information which is communicated to a central controller that will determine the revised processing sequence. This is because the various arrivals and their agents may not want to publicly divulge their cost structures. Since the various arrivals may have incentive to misreport their cost structures, it is of paramount importance that rescheduling mechanisms are strategyproof. Our mechanism is intended to run in real time during a delay event.

To formalize this process, consider a set $N=\{1, 2,\ldots, n\}$ of tasks, or *arrivals*, that are scheduled for processing at facility F. Each arrival is associated with a particular *agent*, and we note that an agent may represent more than one of the arrivals. Because of the