

New Auction Algorithms for the Assignment Problem and Extensions †

by

Dimitri Bertsekas ‡

Abstract

We consider the classical linear assignment problem, and we introduce new auction algorithms for its optimal and suboptimal solution. The algorithms are founded on duality theory, and are related to ideas of competitive bidding by persons for objects and the attendant market equilibrium, which underlie real-life auction processes. We distinguish between two fundamentally different types of bidding mechanisms: *aggressive* and *cooperative*. Mathematically, aggressive bidding relies on a notion of approximate coordinate descent in dual space, an ϵ -complementary slackness condition to regulate the amount of descent approximation, and the idea of ϵ -scaling to resolve efficiently the price wars that occur naturally as multiple bidders compete for a smaller number of valuable objects. Cooperative bidding avoids price wars through detection and cooperative resolution of any competitive impasse that involves a group of persons.

We discuss the relations between the aggressive and the cooperative bidding approaches, we derive new algorithms and variations that combine ideas from both of them, and we also make connections with other primal-dual methods, including the Hungarian method. Furthermore, our discussion points the way to algorithmic extensions that apply more broadly to network optimization, including shortest path, max-flow, transportation, and minimum cost flow problems with both linear and convex cost functions.

† Many thanks are due to Yuchao Li for extensive helpful comments.

‡ Fulton Professor of Computational Decision Making, School of Computing and Augmented Intelligence, Arizona State University, Tempe, AZ.

1. INTRODUCTION

In this paper, we discuss auction algorithms for solving numerically the classical assignment (aka weighted bipartite matching) problem, where there are n persons, denoted by $i = 1, \dots, n$, and n objects, denoted by $j = 1, \dots, n$, which we have to match on a one-to-one basis. Each person i may be matched to any object j within a given subset $A(i) \subset \{1, \dots, n\}$. By a *complete assignment* we mean a set of person-object pairs $(1, j_1), \dots, (n, j_n)$, such that $j_i \in A(i)$ for all $i = 1 \dots, n$, while the objects j_i are all distinct. There is a known value a_{ij} for matching person i with object $j \in A(i)$, and we want to find a complete assignment that maximizes the total value

$$\sum_{i=1}^n a_{ij_i}.$$

The assignment problem has received a lot of attention since the 1950s. It arises in many practical settings, the most obvious ones being resource allocation problems, such as assigning personnel to jobs, resources to tasks, and related contexts, such as scheduling and data association. The assignment problem also appears often as a subproblem in various methods for solving more complex problems.

Recent applications of the assignment problem include:

- (a) *Optimal transport* (arising in cosmology among others; see e.g., Brenier et al. [BFH03], Frisch and Sobolevskii [FrS06], Lavaux [Lav08], Villani [Vil09], [Vil21], Santambrogio [San15], Galichon [Gal16], Metivier et al. [MBM19], Schmitzer [Sch16], [Sch19], Walsh and Dieci [WaD17], [WaD19], Peyre and Cuturi [PeC19], Levy, Mohayae, and von Hausegger [LMH21], Merigot and Thibert [MeT21]).
- (b) *Graph similarity problems* (arising in computational biology among others; see e.g., Kollias et al. [KSS14], Erciyes [Erc15]).
- (c) *Graph neural networks* (see e.g., Zhou et al. [ZCH20], Aironi, Cornell, and Squartini [ACS22], Nurlanov, Schmidt, and Bernard [NSB23]),
- (d) *Combinatorial auctions* (see e.g., Parkes and Ungar [PaU99], De Vries and Vohra [DeV03]).
- (e) *Computational physics* (see e.g., Kosowsky and Yuille [KoY94], Jacobs, Merkurjev, and Esedoglu [JME18], Bertozzi and Merkurjev [BeM19], [Mer20]).
- (f) *A variety of dynamic task allocation, resource allocation, scheduling, multiagent, and multi-robot problems* (see e.g., Bayati et al. [BPS07], Bayati, Shah, and Sharma [BSS08], Choi, Brunet, and How [CBH09], Liu and Shell [LiS13], Luo, Chakraborty, and Sycara [LCS14], Morgan et al. [MSC16], Tang et al. [TZG18], Duan et al. [DLT19], Huang, Zhang, and Xiao [HZX19], Luzak et al. [LuM20], [LGO20], Otte, Kuhlman, and Sofge [OKS20], Aziz et al. [APP22], Wang et al. [WMW22], Garces et al. [GBG23], Li et al. [LZX23], Wang, Li, and Yao [WLY23], and Li et al. [LZC24]).

The assignment problem is also of great theoretical significance because, despite its simplicity, it embodies a fundamental linear programming structure. In particular, the important single commodity linear cost network flow problem can be reduced to an assignment problem by means of a simple reformulation. Thus, any method for solving the assignment problem can be generalized to solve the linear network flow problem, and in fact this approach is particularly helpful in understanding the extensions of auction algorithms to network flow problems that are more general than assignment. Detailed discussions can be found in the author's network optimization textbooks [Ber91a], [Ber98], and are very relevant to the research directions presented in this paper.

Duality Theory for the Assignment Problem

To develop an intuitive understanding of auction algorithms, it is helpful to introduce an economic equilibrium problem that turns out to be equivalent to the assignment problem. Let us consider the possibility of matching the n objects with the n persons through a market mechanism, viewing each person as an economic agent acting in his/her own best interest. Suppose that object j has a price p_j and that the person who acquires the object must pay the price p_j . Then the (net) profit of object j for person i is $a_{ij} - p_j$ and each person i would logically want to be assigned to a maximal profit object $j_i \in A(i)$, i.e., one satisfying

$$a_{ij_i} - p_{j_i} = \max_{j \in A(i)} \{a_{ij} - p_j\}.$$

A set of prices $p = (p_1, \dots, p_n)$ and a set of assigned pairs $\mathcal{A} = \{(i_1, j_1), \dots, (i_k, j_k)\}$ where each assigned person satisfies the preceding condition, i.e.,

$$a_{i_m j_m} - p_{j_m} = \max_{j \in A(i_m)} \{a_{i_m j} - p_j\}, \quad \text{for all assigned pairs } (i_m, j_m) \in \mathcal{A}, \quad (1.1)$$

are said to satisfy *complementary slackness* (CS for short). When CS holds, for a set of prices p and a complete assignment \mathcal{A} (i.e., one where $k = n$), we have a form of economic equilibrium whereby each person is assigned to an object that offers maximum profit, and has no incentive to switch to a different object.

A fundamental duality theorem states that a complete assignment that satisfies the CS condition (1.1) together with some set of prices, is optimal, i.e., it offers maximum total value. Moreover, the corresponding set of prices solves an associated dual optimization problem, which is to minimize over $p = (p_1, \dots, p_n)$ the dual cost function

$$\sum_{i=1}^n \pi_i + \sum_{j=1}^n p_j, \quad (1.2)$$

where π_i is the maximum profit that is attainable for person i under the set of prices p :

$$\pi_i = \max_{j \in A(i)} \{a_{ij} - p_j\}, \quad i = 1, \dots, n. \quad (1.3)$$

Mathematically, we can view this as a consequence of the celebrated duality theorem of linear programming, whereby the assignment optimization is viewed as the primal problem and the minimization of the cost (1.2)-(1.3) is the dual problem.†

Algorithms for Solving the Assignment Problem

There are several iterative algorithms for the solution of the assignment problem, which are described in detail in several sources, including the linear programming textbook by Bertsimas and Tsitsiklis [BeT97], the network optimization books by Bertsekas [Ber91a], [Ber98], and Burkard, Dell’Amico, and Martello [BDM12], and the extensive surveys by Ahuja, Magnanti, and Orlin [AMO88], [AMO89], and Burkard and Cela [BuC99], among others. In particular, two classical types of methods are:

- (a) *Primal simplex methods*, which start with some feasible assignment (a primal basic solution) and iteratively increase the value of the assignment by using the mechanism of the simplex method, suitably adapted to take advantage of the underlying graph structure.
- (b) *Dual cost improvement methods*, which include Kuhn’s Hungarian method [Kuh55], and the relaxation algorithms by Bertsekas [Ber81], [Ber85], and Bertsekas and Tseng [BeT88]. These methods start from a dual solution (a set of object prices) and iteratively modify the prices along dual descent directions, thus generating a cost improving sequence of dual solutions.

A third and distinct class of iterative methods for the assignment problem is *auction algorithms*, the subject of this paper. These methods resemble real-life auctions and can be loosely interpreted as approximate coordinate descent methods for solving the dual problem. The approximation is controlled by a parameter $\epsilon > 0$, which may be reduced in the course of the algorithm. Auction algorithms differ from primal methods and dual methods in a fundamental way: they may deteriorate both the primal and the dual objectives at any one iteration by an amount that depends on ϵ . Still, with appropriate implementation and control of the size of ϵ , they find an optimal primal and dual solution pair.

Aims and Contributions of the Paper

The present paper focuses on three related types of auction algorithms, *conservative*, *aggressive*, and *cooper-*

† The proof is very simple. For any set of prices (p_1, \dots, p_n) and any complete assignment $(1, j_1), \dots, (n, j_n)$, using the definition (1.3) of the profit π_i , we have

$$\sum_{i=1}^n \pi_i + \sum_{j=1}^n p_j = \sum_{i=1}^n \max_{j \in A(i)} \{a_{ij} - p_j\} + \sum_{j=1}^n p_j \geq \sum_{i=1}^n \{a_{ij_i} - p_{j_i}\} + \sum_{j=1}^n p_j = \sum_{i=1}^n a_{ij_i}.$$

Under the CS condition (1.1), equality holds in the above relation. Thus when CS is satisfied, (p_1, \dots, p_n) attains the minimum of the dual cost on the left side above, while $(1, j_1), \dots, (n, j_n)$ attains the maximum of the right side.

ative, which aim to find a set of prices and a complete assignment that attain the market equilibrium noted earlier, and hence solve the corresponding dual and primal problems. The aggressive auction algorithm was first proposed by the author in the paper [Ber79], and was followed by a proposal of a cooperative auction algorithm in the paper [Ber81]. The conservative auction algorithm, which is a limiting form of the aggressive auction algorithm, was also discussed in these papers, and in fact it was suggested as an effective initialization of some of the cooperative algorithms of the paper [Ber81], despite the fact that in general it does not guarantee convergence to an optimal assignment.†

The distinction between the conservative and aggressive auction algorithms can be described in terms of a critical parameter ϵ that characterizes the “intensity” of competition between the persons for the objects: in conservative auction $\epsilon = 0$, while in aggressive auction $\epsilon > 0$. The cooperative auction algorithm, as given in [Ber81], uses $\epsilon = 0$, so it has a conservative character. The present paper extends substantially the cooperative auction framework by allowing $\epsilon > 0$ and by integrating the three different types of auction into a single method, aiming to combine their best characteristics. In particular, the extension to the case where $\epsilon > 0$ involves qualitatively significant changes in the algorithm’s character, and appears to be substantially faster for many problems. The new ideas of this paper also point the way towards extensions to network optimization problems that are more general than assignment.

We first review in Section 2 some of the known ideas relating to conservative and aggressive auctions, and the principal challenges that they face due to what we will call *competitive impasses* and *price wars*. In

† The term “naive auction” was used instead of “conservative auction” in these and other subsequent works. We will avoid the term “naive” in this paper: it is somewhat misleading because conservative auction embodies interesting ideas, and is useful both conceptually and practically, despite the fact that it does not guarantee convergence to an optimal assignment. The paper [Ber81] also proposed and tested a two-phase algorithm, whereby conservative/naive auction was used in the first phase to initialize a Hungarian algorithm used in the second phase. The code of Jonker and Volgenant [JoV87], often referred to as the “JV code,” is very similar. It uses the conservative auction algorithm to initialize a Hungarian-like sequential shortest path method, but starts conservative auction with the classical choice for initial prices: p_j is set to $\min_i a_{ij}$, rather than $p_j = 0$, the author’s choice in the code of [Ber81] (in fact the authors of [JoV87] developed their code working from a printout of the author’s 1981 code). The JV code has been used widely, as it clearly performs better than codes based on the classical Hungarian method, thanks to its conservative auction initialization. On the other hand, aggressive auction codes seem to outperform the JV code, and other Hungarian-related codes, for many types of problems, although assessments differ on this issue; see e.g., Bertsekas and Eckstein [BeE88], Castañón [Cas93], Zaki [Zak95], Malkoff [Mal97]. Aggressive auction codes also seem to outperform codes that are inspired by preflow-push ideas (whose mechanism can be viewed as mathematically equivalent to the one of the aggressive auction algorithm); see the papers by Bertsekas [Ber93], Naporstek and Leshem [NaL16], Alfaro et al. [APV22], and the textbook [Ber98] (Section 7.3.3).

Section 3, we propose a new cooperative auction algorithm, which aims to provide a mechanism for addressing price wars. We discuss several variations, including the *expanding coalitions variant* of cooperative auction, which provides a conceptual vehicle for bridging the ideas of auction and Hungarian methods. The algorithm is structured so that it can combine harmoniously conservative, aggressive, and cooperative auction ideas. A combination of this type was given in the paper [Ber81] for the special case where $\epsilon = 0$. We provide a similar combination, but one where $\epsilon > 0$. In Section 4, we discuss additional variations of the algorithms of Sections 2 and 3, as well as the role of ϵ -scaling within the broader cooperative auction framework of the paper.

In the present paper we will focus on the algorithmic ideas underlying auction algorithms for the assignment problem, particularly the new cooperative versions. In a future report, we will provide results of computational experimentation and describe how our auction ideas can be extended to other linear network flow problems, such as shortest path, max-flow, transportation, and transshipment problems. We will also extend our algorithms of Sections 3 and 4 to single commodity network flow problems with separable convex cost functions, building on auction algorithmic ideas presented in the papers by Bertsekas, Polymenakos, and Tseng [BPT97], [BPT98], and discussed in more detail in Chapter 9 of the book [Ber98].

2. CONSERVATIVE AND AGGRESSIVE AUCTIONS

Let us first establish some terminology. In what follows, by an *assignment* we mean a set of person-object pairs $(i_1, j_1), \dots, (i_k, j_k)$, such that $j_1 \in A(i_1), \dots, j_k \in A(i_k)$, while i_1, \dots, i_k are distinct persons and j_1, \dots, j_k are distinct objects. If $k = n$ the assignment is called *complete*, and if $k < n$ the assignment is called *partial* (or *incomplete*). The empty assignment, where there are no assigned persons or objects, is also considered to be a partial assignment. Generally, assignments (complete, incomplete, or empty) will be denoted by \mathcal{A} . We assume throughout that there exists at least one complete assignment for our given problem. Also for simplicity in describing algorithms, and without loss of generality, we assume that $A(i)$, the set of objects to which person i can be assigned, contains at least two elements.

A common characteristic of all auction algorithms is that they maintain at all times a partial assignment \mathcal{A} and a set of object prices $p = (p_1, \dots, p_n)$, which satisfy an approximate form of the CS condition (1.1) that involves a parameter $\epsilon \geq 0$. The partial assignment grows progressively to become a complete assignment, at which time the auction algorithm terminates.

The central mechanism of an auction algorithm is a bid by an unassigned person i for his/her “best” object j_i (the one that maximizes the person’s profit):

$$a_{ij_i} - p_{j_i} = \max_{j \in A(i)} \{a_{ij} - p_j\}.$$

In particular, person i bids for j_i by raising its price from p_{j_i} to \bar{p}_{j_i} given by

$$\bar{p}_{j_i} = \epsilon + \text{the price level that makes the profit of } j_i \text{ equal to the second best profit } \max_{j \in A(i), j \neq j_i} \{a_{ij} - p_j\}.$$

Depending on whether $\epsilon = 0$ or $\epsilon > 0$, the auction is called *conservative* or *aggressive*, respectively. Thus in an aggressive auction the object prices are raised by larger increments. In this section we will review these two different types of auction and their properties. For detailed discussions, which include additional topics, such as parallel and asynchronous distributed implementations, we refer to the textbooks [BeT89], [Ber91a], and [Ber98], and the tutorial papers [Ber90] and [Ber92]. No new research is presented in this section.

2.1 Conservative Auction

As in real-life auctions, a person needs to balance two competing considerations when determining a proper bid size: a high bid for his/her preferred object discourages bids of other persons for that object, but also diminishes his/her profit upon acquiring that object. Thus it makes sense for a person i to maximize the bid for a preferred object j_i subject to the constraint that this object continues to offer maximum profit, i.e., to raise the price of j_i to

$$\bar{p}_{j_i} = a_{ij_i} - w_i,$$

where w_i is the “second best” profit,

$$w_i = \max_{j \in A(i), j \neq j_i} \{a_{ij} - p_j\},$$

thereby bringing the profit

$$a_{ij_i} - \bar{p}_{j_i},$$

of the best object j_i to the level of the profit w_i of the second best object; see Fig. 2.1. We view this auction mechanism as conservative because when selecting a bid, person i takes no risk, in the sense that he/she will never end up with a non-maximum profit object.†

Let us describe the conservative auction algorithm more precisely. The algorithm proceeds in iterations and throughout its operation, maintains a set of prices $p = (p_1, \dots, p_n)$ and a partial assignment \mathcal{A} where each assigned person is assigned to a maximal profit object, i.e., the CS condition (1.1) is satisfied. It terminates when following an iteration, the assignment obtained is complete. The algorithm starts with any set of prices and partial assignment that satisfy CS; for example it may start with an arbitrary set of prices and the empty assignment. Given the current set of object prices p and partial assignment \mathcal{A} , a conservative auction iteration generates a new set of prices and a new assignment as follows.

† Price rises below the maximum level $a_{ij_i} - w_i$ also have this property, but larger price rises tend to accelerate the termination of the auction, and are therefore better suited for our algorithmic purposes.

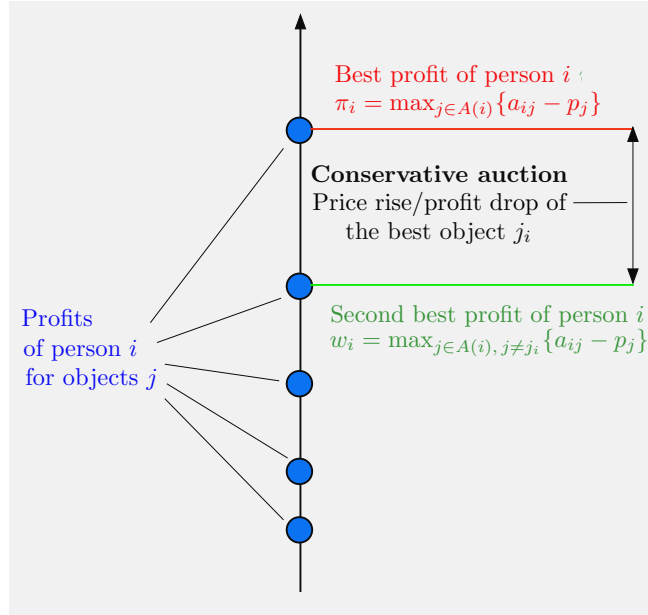


Figure 2.1 Illustration of the price rise of the best object j_i of an unassigned person i in the conservative auction algorithm. The price of j_i is increased by $\pi_i - w_i$, while the profit of j_i is made equal to w_i .

Conservative Auction Iteration

We select an unassigned person i and an object j_i that offers maximum profit for i under the given prices,

$$a_{ij_i} - p_{j_i} = \max_{j \in A(i)} \{a_{ij} - p_j\}. \quad (2.1)$$

We set the price of j_i to


$$\bar{p}_{j_i} = a_{ij_i} - w_i, \quad (2.2)$$

where w_i is the “second best” profit,

$$w_i = \max_{j \in A(i), j \neq j_i} \{a_{ij} - p_j\}. \quad (2.3)$$

Finally, we add to the assignment \mathcal{A} the pair (i, j_i) , and if j_i was assigned to some other person \tilde{i} , we remove from \mathcal{A} the pair (\tilde{i}, j_i) , thus forming a new assignment $\bar{\mathcal{A}}$.

It can be seen that the conservative auction algorithm maintains the CS condition (1.1) throughout its operation, and generates a sequence of partial assignments whose cardinalities are not decreasing, so if it terminates, the complete assignment obtained at termination is optimal, while the corresponding final prices are an optimal solution to the dual problem, by the duality theorem noted earlier.



At Start of Iteration #	Object Prices	Assigned Pairs	Bidding Person	Best Object	Bidding Increment
1	(0, 0, 0)	(1, 1), (2, 2)	3	2	0
2	(0, 0, 0)	(1, 1), (3, 2)	2	2	0
3	(0, 0, 0)	(1, 1), (2, 2)	3	2	0

Figure 2.2 Illustration of how the conservative auction algorithm may never terminate for a 3×3 assignment problem. Here objects 1 and 2 have value $C > 0$ for all persons, and object 3 has value 0 for all persons. The algorithm starts from the initial prices $p = (0, 0, 0)$ and the partial assignment $\{(1, 1), (2, 2)\}$. There is a competitive impasse involving persons 1, 2, and 3, and objects 1 and 2. The algorithm cycles as persons 2 and 3 alternately bid for object 2 (or object 1) without changing its price because they prefer equally object 1 and object 2.

On the other hand, conservative auction offers no guarantee of termination: we may end up with a situation where the object prices stop changing, while the cardinality of the current assignment stops growing, as some persons simply change their assigned objects in some way. In particular, by Eqs. (2.1)-(2.3), the new price \bar{p}_{j_i} of the preferred object j_i cannot decrease, i.e.,

$$\bar{p}_{j_i} \geq p_{j_i},$$

and it will increase strictly (i.e., $\bar{p}_{j_i} > p_{j_i}$) if and only if the profit $a_{ij_i} - p_{j_i}$ of j_i is strictly larger than the second best profit w_i ; cf. Fig. 2.1. Thus neither the object prices nor the cardinality of the current assignment will change if there are multiple objects that offer maximum profit for person i , and all of these objects are assigned.

Typically, the cause of nontermination of conservative auction can be traced to what we will call a *competitive impasse*. We can somewhat loosely describe competitive impasse as a situation where there is a set of persons that compete for a smaller number of (more than one) equally desirable objects, and there is no apparent way to allocate objects to persons without leaving some person(s) dissatisfied in the end; see the example of Fig. 2.2. In practice, however, conservative auction can quickly succeed in assigning a substantial number of objects, and for this reason it can be used for effective initialization of other assignment algorithms, as was noted in the papers [Ber81] and [JoV87].

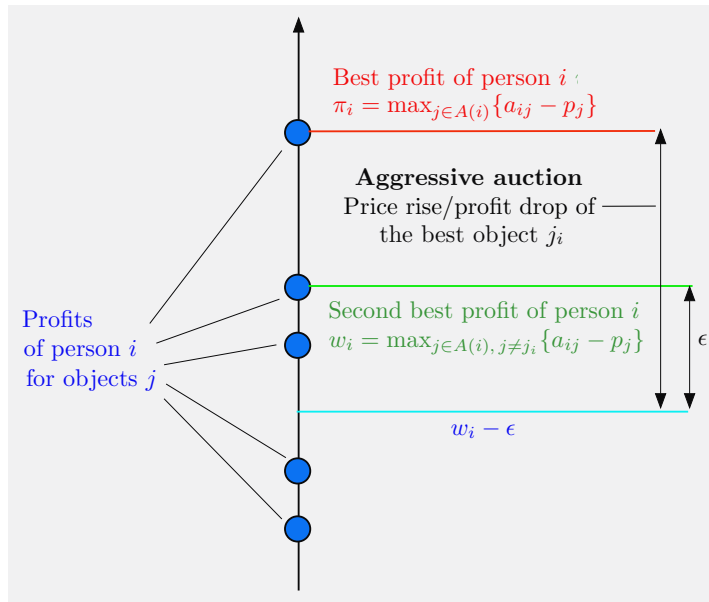


Figure 2.3 Illustration of the price rise of the best object j_i of an unassigned person i in the aggressive auction algorithm. The price of j_i is increased by $\pi_i - w_i + \epsilon$, while the profit of j_i is strictly decreased to $w_i - \epsilon$.

2.2 Aggressive Auction

The aggressive auction algorithm is similar to its conservative counterpart, but guarantees convergence to a complete assignment. In particular, a competitive impasse is resolved by requiring that a bid by an unassigned person i for the best object j_i increases the price of j_i by at least some positive increment ϵ . In particular, person i raises the price of the best object j_i by the amount

$$\pi_i - w_i + \epsilon,$$

where $\pi_i = \max_{j \in A(i)} \{a_{ij} - p_j\}$ is the profit of the best object, given by Eq. (1.3), and w_i is the second best profit, given by Eq. (2.3); see Fig. 2.3. We refer to this type of auction as *aggressive*, because in contrast to the conservative type, it is guaranteed to apply positive price rises (at least ϵ), and it may produce a complete assignment where some of the persons are assigned to a non-maximum profit object. We will also contrast aggressive auction with the cooperative type of auction algorithm (to be discussed shortly), which aims to first detect a competitive impasse and then resolve it through a process of mutual agreement among the competing persons.

In summary, given a set of object prices (p_1, \dots, p_n) and a partial assignment \mathcal{A} , an aggressive auction iteration generates a new set of prices and a new assignment as described below. The algorithm terminates when following an iteration, the assignment obtained is complete.

Aggressive Auction Iteration

We select an unassigned person i and an object j_i that offers maximum profit for i under the given prices,

$$a_{ij_i} - p_{j_i} = \max_{j \in A(i)} \{a_{ij} - p_j\}. \quad (2.4)$$

We set the price of j_i to

$$\bar{p}_{j_i} = a_{ij_i} - w_i + \epsilon \quad (2.5)$$

where w_i is the “second best” profit,

$$w_i = \max_{j \in A(i), j \neq j_i} \{a_{ij} - p_j\}.$$

Finally, we add to the assignment \mathcal{A} the pair (i, j_i) , and if j_i was assigned to some other person \tilde{i} , we remove from \mathcal{A} the pair (\tilde{i}, j_i) , thus forming a new assignment $\bar{\mathcal{A}}$.


It can be shown that the algorithm is guaranteed to terminate (under our assumption that there exists at least one complete assignment; see the original paper [Ber79], or the books [BeT89], [Ber91a], [Ber98] for a proof). Intuitively, the reason is that each bid by a person i is guaranteed to increase the price of his/her best object j_i by at least the positive increment ϵ , thus making j_i “less attractive” for other persons. If the auction did not terminate, the prices of the assigned objects would eventually increase to sufficiently high levels to make some of the unassigned objects attractive enough to receive bids and join the assignment. This is similar to what happens in real-life auctions.

The aggressive auction algorithm is designed to maintain the following relaxed form of the CS condition (1.1), called ϵ -complementary slackness (ϵ -CS for short):

$$a_{ij_i} - p_{j_i} \geq \max_{j \in A(i)} \{a_{ij} - p_j\} - \epsilon, \quad \text{for all assigned pairs } (i, j_i), \quad (2.6)$$

provided the initial set of prices and partial assignment satisfy this condition. One possibility to satisfy the ϵ -CS condition initially is to start with an arbitrary set of prices and the empty assignment. There are also other more sophisticated possibilities for selecting favorable initial conditions.

Thanks to the ϵ -CS condition, it can be shown that the final assignment obtained is optimal within $n\epsilon$, and hence exactly optimal if the values a_{ij} are integers and $\epsilon < 1/n$. To see this, note that the complete assignment and set of prices obtained at termination *satisfy CS for a fictitious/slightly perturbed problem* where all values a_{ij} are the same as before, except for the values a_{ij_i} of the n assigned pairs (i, j_i) , which are modified by an amount of no more than ϵ ; cf. Eq. (2.6). The final complete assignment is optimal for



At Start of Iteration #	Object Prices	Assigned Pairs	Bidding Person	Best Object	Bidding Increment
1	(0, 0, 0)	(1, 1), (2, 2)	3	2	ϵ
2	(0, ϵ , 0)	(1, 1), (3, 2)	2	1	2ϵ
3	(2ϵ , ϵ , 0)	(2, 1), (3, 2)	1	2	2ϵ
4	(2ϵ , 3ϵ , 0)	(1, 2), (2, 1)	3	1	2ϵ
5	(4ϵ , 3ϵ , 0)	(1, 2), (3, 1)	2	2	2ϵ
6

Figure 2.4 Illustration of how the aggressive auction algorithm overcomes the competitive impasse problem for the 3×3 example of Fig. 2.2 by making the bidding increment at least equal to ϵ . The table shows one possible sequence of bids and assignments generated by the auction algorithm, starting with all prices equal to 0 and the partial assignment $\{(1, 1), (2, 2)\}$. At each iteration except the last, the unassigned person bids for either object 1 or 2, increasing its price by ϵ in the first iteration and by 2ϵ in each subsequent iteration. In the last iteration, after the prices 1 and 2 rise to or above C , object 3 receives a bid and the auction terminates. The number of iterations for this to happen is roughly C/ϵ .

this perturbed problem, and therefore also optimal within $n\epsilon$ for the original (unperturbed) problem. Thus thanks to the extra ϵ bidding increment, the aggressive auction algorithm succeeds in terminating with a complete assignment, at the risk of some persons ending up with a non-maximum profit object (by as much as ϵ), and an attendant error of at most $n\epsilon$ from optimality.

Unfortunately, the aggressive auction algorithm runs into another difficulty, which can also be traced to a competitive impasse. This difficulty, called a *price war*, refers to a protracted sequence of small price rises of order ϵ , which results from groups of persons competing for a smaller number of two or more objects that are more or less equally desirable. An example of a price war in the case of a 3×3 assignment problem is given in Fig. 2.4, and it can be seen that it degrades computational efficiency. In particular, the number of iterations in this example is proportional to C/ϵ , and a similar example (given as Exercise 7.4b in the book [Ber98]) shows that the number of iterations needed to resolve a price war can be as high as nC/ϵ .

Generally, the complexity of the algorithm can be shown to be proportional to C/ϵ , where

$$C = \max_{i=1, \dots, n, j \in A(i)} |a_{ij}| \tag{2.7}$$

is the range of object values. Thus, the complexity is pseudopolynomial and is often unacceptable. In actual use of the aggressive auction algorithm, price wars are common, particularly when the range C is large and the assignment problem is sparse, i.e., each person can be assigned to only a small subset of objects).

One way to overcome the detrimental effect of price wars is ϵ -scaling, a natural computational idea that was noted in the original aggressive auction proposal of the paper [Ber79]. Here the algorithm is first run for a fairly large initial value of ϵ , to converge quickly and yield good object price estimates. These estimates are used to initialize an aggressive auction with a reduced value of ϵ . After several successive rounds of ϵ -reduction by some constant factor, this process will bring ϵ to a sufficiently low level to produce an optimal assignment. It can be shown that the (worst-case) computational complexity of aggressive auction with ϵ -scaling is polynomial, $O(nm \log(nC))$, where m is the number of arcs of the bipartite graph representing the assignment problem and C is the range of values, given by Eq. (2.7). This estimate was derived in the author's textbook [BeT89] (Section 5.4) and paper [Ber88], following a progression of related complexity analyses for the max-flow and the minimum cost flow problem involving several works (Karzanov [Kar74], Shiloach and Vishkin [ShV82], Goldberg and Tarjan [GoT86], [GoT90], Bertsekas and Eckstein [BeE87], [BeE88], Ahuja, Magnanti, and Orlin [AMO88], [AMO89], Ahuja and Orlin [AhO89], Cheriyan and Maheshvari [ChM89], Orlin and Ahuja [OrA92]). The recent papers by Naparstek and Leshem [NaL16], and Khosla and Anand [KhA21] provide probabilistic complexity analyses.

For an account of the computational complexity aspects of the aggressive auction algorithm with ϵ -scaling, see the textbooks [BeT89] (Section 5.4) and [Ber98] (Section 7.1.2). The latter textbook also contains detailed discussions (including computational complexity) of extensions of the auction algorithm to related problems, such as asymmetric assignment problems, max-flow, minimum cost flow, with both linear (in Chapter 7) and convex separable cost (in Chapter 9).

Aggressive auction with an efficient ϵ -scaling implementation is widely recognized as one of the most effective assignment algorithms.† Several code implementations are publicly available, including some (written in FORTRAN and dating from the early 90s) that can be found in the author's website. A recent code, written in MATLAB, has been made available by Bernard [Ber23a]. The algorithm typically outperforms its

† The experimental verification of the advantages of the aggressive auction algorithm took a long time to establish, owing in part to the primitive state of computer technology at the time. Indeed, given that the aggressive auction algorithm appeared to be radically different from the established assignment algorithms in 1979, like primal simplex and Hungarian, and lacking a thorough computational comparison, the author harbored deep doubts about its effectiveness. In fact, these doubts prompted the development of an alternative cooperative algorithm (with $\epsilon = 0$), which appeared to be conceptually closer to the Hungarian method, the most popular assignment algorithm at the time; see [Ber81]. The story of the discovery of the aggressive auction algorithm is recounted near the end of a videolecture by the author that can be found at <https://www.youtube.com/watch?v=T-fSmSqzcqE>

competitors by a wide margin, as has been shown convincingly by many computational studies. Its advantage is particularly pronounced when good initial object price estimates are available. As a result, the method is very efficient in situations where many similar assignment problems are solved with small variations in their data. Then the final prices for a given problem solution can be used as starting prices for solution of other similar problems, often with impressive computational savings.†

Another advantage of the aggressive auction algorithm and its extensions to other network flow problems is that it is well-suited for parallel computation, and it is valid even when it is implemented as a distributed asynchronous algorithm. This has been established in the book by Bertsekas and Tsitsiklis [BeT89] (Sections 5.3 and 6.5), as well as in several related computational studies: Bertsekas and Castañon [BeC91], Wein and Zenios [WeZ91], Amini [Ami94], Bertsekas et al. [BCE95], Beraldi, Guerriero, and Muzumano [BGM97], [BGM01], [BeG97], Zavlanos, Spesivtsev, and Pappas [ZSP08], Bus and Tvrdk [BuT09], Sathe, Schenk, and Burkhart [SSB12], Nascimento et al. [NVJ16], Naparstek and Leshem [NaL16], Sena, Silva, and Nascimento [SSN21]. The cooperative auction algorithms to be discussed next, can also use good initial price estimates with advantage, but they are not as well suited for distributed computation.

3. COOPERATIVE PRICE RISES AND COOPERATIVE AUCTION

We will now consider an alternative approach for dealing with competitive impasses and price wars. The key characteristic that differentiates it from the aggressive auction approach is the use of multiple-object price rises that aim to forestall price wars. In effect, a group of persons recognize that they are caught up in a multi-object competitive impasse, and rather than engage in a time consuming price war, they collectively agree to raise the prices of the relevant objects by a large common increment, thus preparing to bid for additional objects without violating ϵ -CS.

We call such multi-person bid mechanisms *cooperative*, and we will show that they can be combined harmoniously with the aggressive and conservative auction mechanisms that involve bids by a single person. This idea dates to the paper [Ber81], which included combinations of cooperative multi-person bids with conservative single-person bids, and experimentally demonstrated the potential advantages of such combinations.

To understand the cooperative auction mechanism, let us consider the 3×3 assignment problem of Figs. 2.2 and 2.4. There, starting with zero prices, persons 1, 2, and 3 compete for valuable objects 1 and 2 (value

† Such situations arise often in practice. An example is data association contexts, where related two-dimensional assignment problems are solved repeatedly; see the author’s monograph [Ber20a] (Section 3.4.2) and paper [Ber20b], and references on multi-target tracking, such as Blackman [Bla86], Bar-Shalom and Fortman [BaF88], Bar-Shalom [Bar90], Castañon [Cas92], Pattipati, Deb, Bar-Shalom, and Washburn [PDB92], Poore [Poo94], Poore and Robertson [PoR97], Popp, Pattipati, and Bar-Shalom [PPB01], and Emami et al. [EPE20].

C), and aim to avoid assignment to the valueless object 3. As we have seen in Fig. 2.2, conservative auction fails for this problem, due to a competitive impasse created by perpetual zero-increment bids by persons 1, 2, and 3, for the two desirable objects 1 and 2. Aggressive auction succeeds in finding the optimal assignment after a protracted price war that lasts for about C/ϵ iterations, as illustrated in Fig. 2.4. Cooperative auction, aims instead to detect the competitive impasse, to identify the set of persons that are involved in it, and to *form a coalition of these persons for the purpose of performing a cooperative price rise* to resolve quickly the impasse within the coalition. In particular, persons 1, 2, and 3 agree to raise the prices of objects 1 and 2 from 0 to $C + \epsilon$, preserving ϵ -CS, while allowing object 3 to be assigned at the next iteration, thus resolving the competitive impasse without a price war.† This example also illustrates that *price wars involve more than one object*. This motivates the use of an aggressive bidding approach when the ϵ -zone of the bidding person contains only one object. We will return to this theme later in this section.

We will now extend the idea just described to the general $n \times n$ assignment problem. To this end we need to address the following issues:

- (a) The algorithm should maintain a partial assignment and a set of prices that satisfy CS (or ϵ -CS). Thus, once the algorithm terminates, the complete assignment obtained at termination is optimal (or optimal within $n\epsilon$, respectively).
- (b) As in the case of conservative and aggressive auctions, the algorithm should aim to enlarge the current partial assignment as long as this is done without violating CS or ϵ -CS.
- (c) The algorithm needs an explicit or implicit mechanism to detect that there is a competitive impasse or price war going on. It also needs a mechanism to identify the coalition of persons that are involved in the price war; this coalition will involve a single unassigned person and $m > 1$ assigned persons, and the corresponding assigned m objects for which the $m + 1$ persons compete.
- (d) Once a coalition of $m + 1$ persons is detected, the prices of the corresponding m assigned objects should be raised simultaneously through a cooperative price rise that does not violate CS or ϵ -CS. An efficient mechanism to calculate the cooperative price rise level should be incorporated into the algorithm.

In what follows in this paper, we will aim to design a broad class of algorithms and variations thereof, which are based on the preceding considerations, and mitigate the occurrences of competitive impasses and price wars. To this end we introduce some definitions, all of which refer to a specific set of prices

† In a real auction the person that is ultimately assigned to the valueless object 3 may need to be compensated by prior agreement with his/her coalition partners. This issue is not addressed in this paper, because our objective is computational efficiency in solving the assignment problem, and not the design of fair real-life auction mechanisms. Some possibilities include consideration of profit sharing between persons, or randomized solutions, whereby persons can acquire fractional amounts of multiple objects, with the fractions adding to 1 for each person.




Figure 3.1 Illustration of the ϵ -zone of a person i . It consists of all the objects whose profit is within ϵ of the maximum profit $\pi_i = \max_{j \in A(i)} \{a_{ij} - p_j\}$.

$p = (p_1, \dots, p_n)$ and partial assignment \mathcal{A} satisfying ϵ -CS for some fixed $\epsilon \geq 0$ (note that $\epsilon = 0$ is a possibility). If $\epsilon > 0$, the algorithm can be combined with ϵ -scaling, i.e., applying the algorithm with larger values of ϵ to obtain good starting prices for applying the algorithm with smaller values of ϵ . However, the algorithm works even with $\epsilon = 0$.

Preliminary Concepts

We first introduce the concept of ϵ -zone of a person, a new idea that plays a central role in this paper.

Definition 3.1 (ϵ -Zone of a Person): Given a set of prices p , the *maximum profit* of a person i , denoted π_i , is defined as

$$\pi_i = \max_{j \in A(i)} \{a_{ij} - p_j\}.$$

For a given $\epsilon \geq 0$, the ϵ -zone of a person i , denoted $\mathcal{Z}(i)$, is the set of objects j whose profit for i is within ϵ of being maximal:

$$\mathcal{Z}(i) = \{j \in A(i) \mid a_{ij} - p_j \geq \pi_i - \epsilon\}.$$

Figure 3.1 illustrates the above definition. Note that the ϵ -zone $\mathcal{Z}(i)$, roughly speaking, consists of the “almost best” objects of person i (those whose profit is within ϵ of being best). It always contains the maximum profit object(s) for person i (and only those if $\epsilon = 0$). Moreover, if a person i is assigned to an object j while ϵ -CS holds, then j belongs to the ϵ -zone $\mathcal{Z}(i)$.

Definition 3.2 (Alternating Path): Let a set of prices p and a partial assignment \mathcal{A} satisfying ϵ -CS be given. An *alternating path* is a person sequence (i, i_1, \dots, i_k) and corresponding object sequence (j_1, \dots, j_k) , $k \geq 1$, such that:

- (a) The person i is unassigned, while the persons i_1, \dots, i_k are assigned to objects j_1, \dots, j_k , respectively.
- (b) The object j_1 belongs to the ϵ -zone of person i , while for $m = 2, \dots, k$, the object j_m belongs to the ϵ -zone of person i_{m-1} .

Definition 3.3 (Augmenting Path): Let a set of prices p and a partial assignment \mathcal{A} satisfying ϵ -CS be given. An *augmenting path* is an alternating path (i, i_1, \dots, i_k) , as per Definition 3.2, together with an unassigned object j that belongs to the ϵ -zone of i_k . Given such a path, a corresponding *augmentation* consists of assigning person i to j_1 , reassigning person i_k to j , and reassigning persons i_1, \dots, i_{k-1} to objects j_2, \dots, j_k , respectively (thereby increasing the cardinality of the assignment by one, while maintaining ϵ -CS). Assigning an unassigned person i to an unassigned object j within his/her ϵ -zone $\mathcal{Z}(i)$ is also viewed as an augmentation.

An augmenting path as defined above, is denoted by (i, i_1, \dots, i_k, j) , while the corresponding alternating path is denoted by (i, i_1, \dots, i_k) [in the case where i is assigned to j , the augmenting path is denoted (i, j)]. Figures 3.2 and 3.3 illustrate alternating and augmenting paths. Key observations here are that:

- (a) An augmenting path starts with an unassigned person i and ends with an unassigned object j , while all other persons and objects in the path are assigned.
- (b) Person i and object j can get assigned through an augmentation, which reassigns objects to persons, while maintaining ϵ -CS. This augmentation makes progress towards obtaining a complete assignment.

The concepts of alternating and augmenting paths are well-known (for the case $\epsilon = 0$) in the theory of assignment, matching, and max-flow algorithms. In particular, an augmentation increases the cardinality of




Figure 3.2 Illustration of an alternating path (i, i_1, i_2) , consisting of the three persons in a 3×3 assignment graph. The first person is unassigned and the subsequent persons are assigned. The objects in the alternating path must belong to the ϵ -zones of the corresponding persons in the path, i.e., $j_1 \in Z(i)$ and $j_2 \in Z(i_1)$ [in addition to $j_1 \in Z(i_1)$ and $j_2 \in Z(i_2)$, which is true by ϵ -CS]. Another alternating path is (i, i_1) .

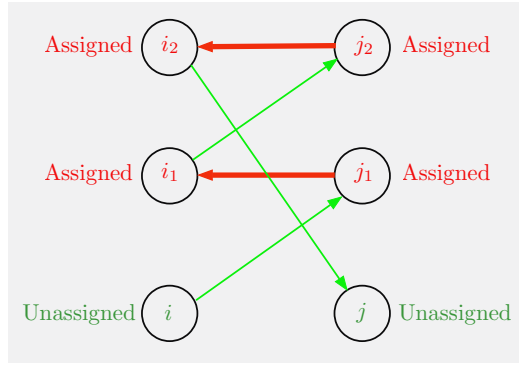


Figure 3.3 Illustration of an augmenting path (i, i_1, i_2, j) in a 3×3 assignment graph. It consists of the alternating path (i, i_1, i_2) (cf. Fig. 3.2), followed by the unassigned object j , which belongs to the ϵ -zone of person i_2 .

the assignment by 1, while changing the maximal profits of the persons of the augmenting path by no more than ϵ . Thus an augmentation makes intuitive sense for small values of ϵ .

We now introduce a notion of coalition of persons, which is central in cooperative auction.

Definition 3.4 (Coalition Partners of an Unassigned Person): Let a set of prices p and a partial assignment \mathcal{A} satisfying ϵ -CS be given, and let i be an unassigned person. A person i' is said to be a *coalition partner* of i if there is an alternating path that starts with i and ends with i' . The set of persons consisting of person i together with all his/her coalition partners is called the *coalition* of i and is denoted by $\mathcal{C}(i)$.

Figure 3.4 provides illustrations of $\mathcal{C}(i)$, the coalition of i . Generally, $\mathcal{C}(i)$ consists of a single unassigned person, namely i , together with $m \geq 0$ assigned coalition partners. It consists of the single person i

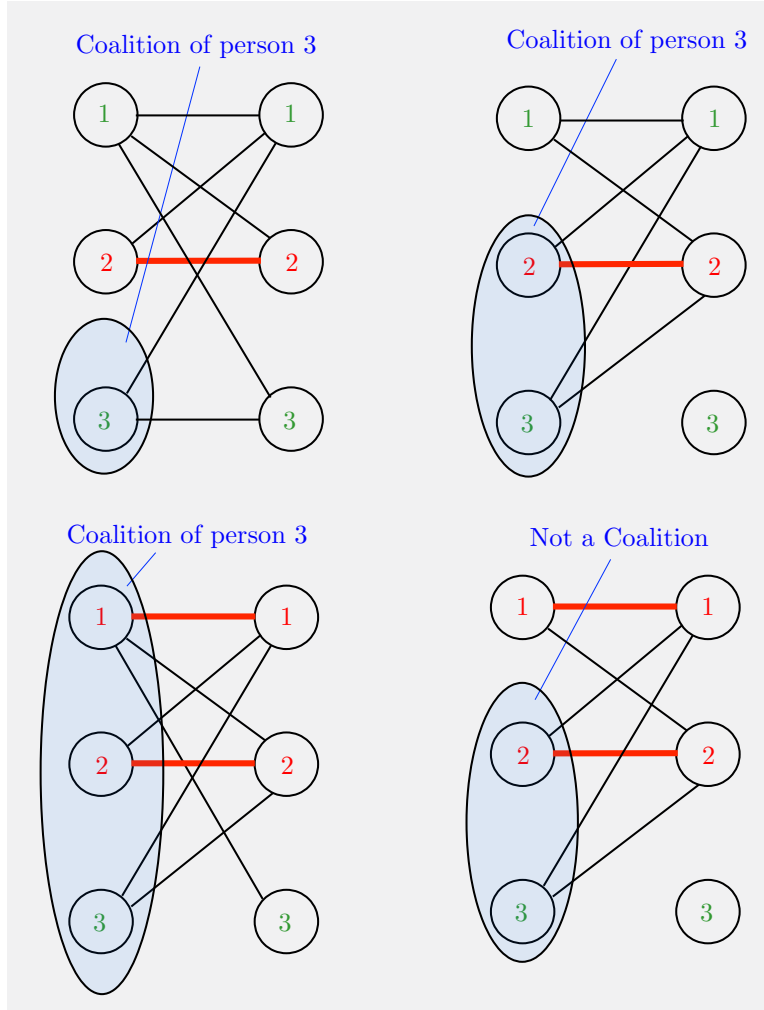


Figure 3.4 Illustrations of different cases of $\mathcal{C}(3)$, the coalition of the unassigned person 3 in a 3×3 assignment problem. In each of the four cases, an arc (i, j) indicates membership of object i in the ϵ -zone of person i (other arcs are not shown). Red arcs correspond to assigned pairs, black arcs to unassigned pairs.

$[\mathcal{C}(i) = \{i\}]$ if and only if the ϵ -zone of i does not include any assigned objects (cf. the top left graph of Fig. 3.4). We note that we can obtain $\mathcal{C}(i)$ by using a form of forward search that progressively generates a tree of alternating paths starting from i , until no more assigned persons can be found; see the implementation details given later in this section.

Cooperative Auction Iteration

An auction iteration involving a cooperative price rise can now be described in words. We are given a set of object prices $p = (p_1, \dots, p_n)$ and a partial assignment \mathcal{A} satisfying ϵ -CS. The iteration starts with an unassigned person i and tries to generate $\mathcal{C}(i)$, the coalition of i . When $\mathcal{C}(i)$ is obtained without intermediate discovery of an augmenting path, the prices of all the objects involved in the coalition will be simultaneously

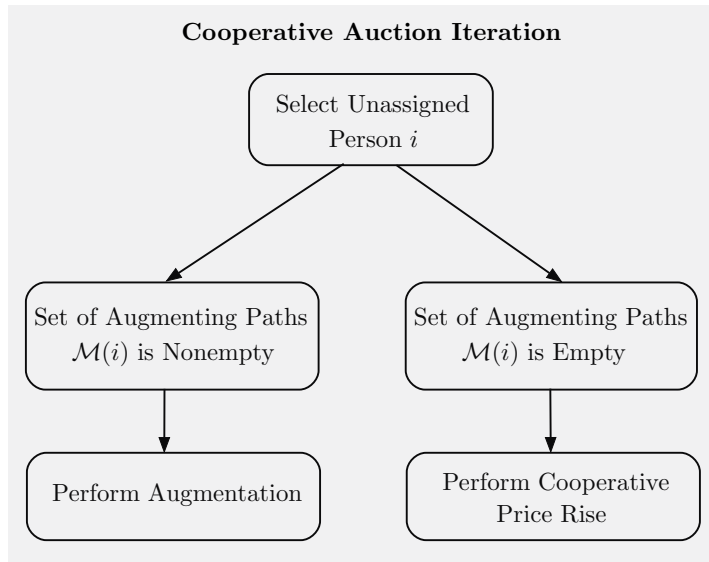


Figure 3.5 Block diagram of a cooperative auction iteration.

raised. Similar to the aggressive auction iteration, the price rise amount exceeds ϵ , and is the maximum possible that preserves ϵ -CS.

We will now state in detail the iteration just described in summary; see the block diagram of Fig. 3.5. We may call this iteration “purely” cooperative, to distinguish it from a method that involves a combination with the conservative and aggressive iterations. We will describe this combined method later in this section.

Cooperative Auction Iteration

Given a set of object prices $p = (p_1, \dots, p_n)$ and a partial assignment \mathcal{A} satisfying ϵ -CS, select an unassigned person i . Let $\mathcal{M}(i)$ be the set of augmenting paths that start with i .

- If $\mathcal{M}(i)$ is nonempty, perform an augmentation along some augmenting path from $\mathcal{M}(i)$, increase the price of the last object in this augmenting path by the maximum amount that will not violate ϵ -CS, and go to the next iteration.
- If $\mathcal{M}(i)$ is empty, let $\mathcal{O}(i)$ denote the set of objects that are assigned to some coalition partner of i . Raise the prices of the objects in $\mathcal{O}(i)$ by the maximum common amount for which the ϵ -zone of every person i' in $\mathcal{C}(i)$ is a subset of the ϵ -zone of the same person i' after the price rise.

The preceding iteration description of the cooperative auction algorithm leaves out the details of the computations of the sets $\mathcal{M}(i)$, $\mathcal{C}(i)$, and $\mathcal{O}(i)$, and the price rise amount. To implement efficiently the iteration, it is necessary to properly organize and streamline these computations. The data structures and

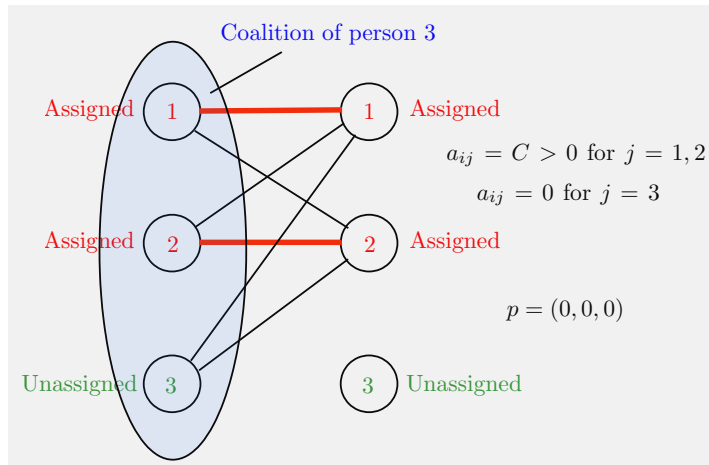


Figure 3.6 Illustration of the cooperative auction iteration for the example of Figs. 3.4, assuming that $C > \epsilon$. Here the coalition partners of the unassigned person 3 are the persons 1 and 2. The cooperative auction iteration consists of a price rise of objects 1 and 2 from 0 to $C + \epsilon$, followed by object 3 coming into the ϵ -zones of all the persons, and allowing an augmentation along $(3, 3)$ that completes the assignment.

procedures for doing so are similar to well-known implementations of auction, Hungarian, and dual descent algorithms, and will be discussed later (cf. Sections 3.3 and 3.4).

Let us illustrate the steps of the cooperative auction iteration with an example.

Example 3.1

Consider the 3×3 assignment example of Figs. 2.2 and 2.4. We will describe a single iteration of the cooperative auction algorithm, starting with set of prices $p = (0, 0, 0)$ and partial assignment $\{(1, 1), (2, 2)\}$.

The iteration starts with person 3, the only one left unassigned. We assume that $C > \epsilon$, so the ϵ -zone $Z(3)$ is the set of objects $\{1, 2\}$. Thus we need to construct the coalition of person 3 with a view towards a cooperative price rise. The alternating paths are $(3, 1)$, $(3, 2)$, $(3, 1, 2)$, and $(3, 2, 1)$, so the coalition partners of person 3 are persons 1 and 2, as illustrated in Fig. 3.6. No augmenting path can be found, i.e., $\mathcal{M}(3)$ is empty, so we increase the prices of objects 1 and 2 by the maximum amount that will not violate ϵ -CS. Thus the prices of objects 1 and 2 are raised to $C + \epsilon$, adding object 3 to the ϵ -zones of persons 1, 2, and 3. At the next iteration, the augmenting path $(3, 3)$ will be discovered, and the algorithm will terminate with an augmentation along $(3, 3)$.

Thus the cooperative auction algorithm terminates very quickly in this example. By contrast, the conservative auction algorithm would not terminate at all because of a competitive impasse (cf. Fig. 2.2), while the aggressive auction algorithm would require about C/ϵ iterations because of a price war (cf. Fig. 2.4).

Variants and Modifications

Note that in the preceding example, the augmenting path $(3, 3)$ is created immediately following the

price rise, so the corresponding augmentation can be done right away. This suggests a modification of the cooperative auction algorithm so that when an augmenting path is discovered following a price rise, the corresponding augmentation is done right away, rather than wait for another iteration. The expanding coalition variant of the algorithm, which will be discussed shortly, embodies this modification.

If on the other hand an augmenting path is not discovered immediately following a cooperative price rise, there is also a possibility to assign person i through a reassignment of the coalition partners of i . We can view this as a somewhat more aggressive form of collective bid of the coalition $\mathcal{C}(i)$, which aims to acquire a new object for the coalition, at the expense of deassigning a person from outside the coalition. It leads to the person reassignment variant of the cooperative auction algorithm, which will be discussed in Section 4.

Similarities with Noncooperative Auction Iterations

Some similarities with the noncooperative auction iterations, which suggest interesting algorithmic variants, are noteworthy. In particular, assume that the ϵ -zone $\mathcal{Z}(i)$ contains a *single unassigned* object (by necessity the maximum profit object). Then the cooperative auction iteration will produce identical results with the conservative iteration (if $\epsilon = 0$) and with the aggressive iteration (if $\epsilon > 0$): it will assign i to that object and raise its price by an amount that exceeds ϵ . If on the other hand the ϵ -zone $\mathcal{Z}(i)$ contains a *single assigned* object j , the results will be different, because the person assigned to j is a coalition partner of i , and this will trigger the mechanism for computing and enlarging the coalition of i .

In what follows (Section 3.2), we will discuss a variant of the cooperative algorithm that behaves identically with the aggressive auction iteration when $\mathcal{Z}(i)$ contains a single object (assigned or unassigned), and is much faster, both in theory and in practice. The motivation for this variant is that the aggressive auction iteration is known to work very fast in the absence of price wars, which involve competition for multiple objects, so *a potential price war is not an issue when $\mathcal{Z}(i)$ consists of a single object.*[†] In Section 4 we will describe still another variant of the cooperative algorithm, which behaves identically with the aggressive auction iteration when there is at most one coalition partner of i [rather than $\mathcal{Z}(i)$ containing a single object]. This is the variant noted earlier, which involves reassignment of the coalition partners of i immediately following a cooperative price rise.

3.1 Cooperative Auction Iteration With Coalition Expansions

When the augmenting path set $\mathcal{M}(i)$ is empty (which will happen when all coalition partners of i are

[†] For an illustration of why price wars involve at least two objects, consider the 3×3 problem of Fig. 2.4. If there were only one valuable object (value C for all persons) and the other two objects were valueless, the type of price war illustrated in the figure would not occur. For an illustration of how a price war can be generated subsequent to aggressive auction iterations, consider the same example with a fourth person added (with identical values as the other three persons) and a fourth object added offering value -1 for all four persons.

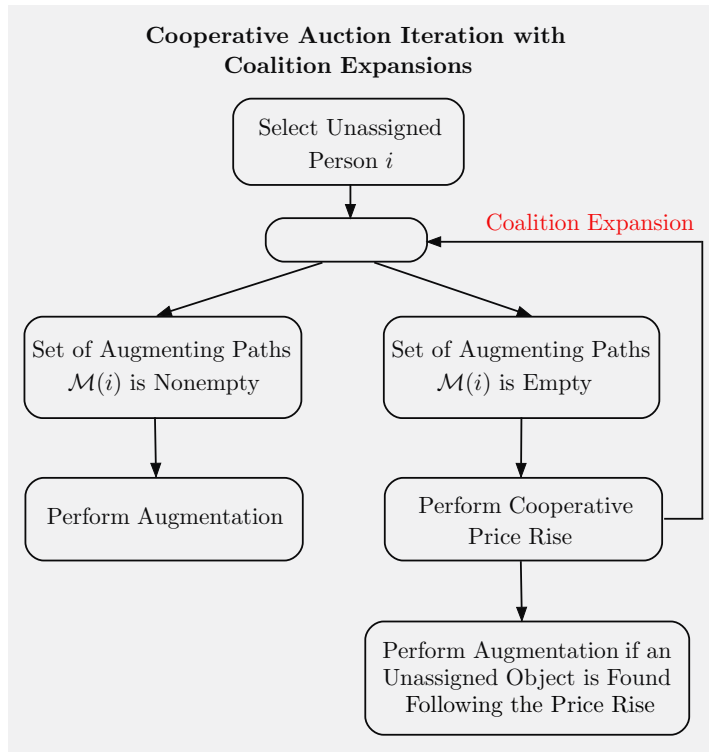


Figure 3.7 Block diagram of a cooperative auction iteration with coalition expansions. The iteration always terminates with an augmentation, possibly following multiple coalition expansions.

assigned), the initial unassigned person i will remain unassigned at the end of the iteration. In this case, since the choice of the unassigned person to start the next iteration is unrestricted, we have the option to start with the same person i . Then the new set of coalition partners of i will include the preceding set of coalition partners, so the *coalition* $\mathcal{C}(i)$ will be simply expanded and need not be rebuilt from scratch [by design the ϵ -zone of every person i' in $\mathcal{C}(i)$ before the price rise is a subset of the ϵ -zone of the same person i' after the price rise].

This observation motivates an interesting variant of the cooperative auction iteration, which involves multiple successive coalition expansions started by the same single person i , up to the point where an augmentation takes place; see Fig. 3.7. We call this the *expanding coalitions variant*, and we note that it will always terminate with an augmentation, resulting in assignment of the starting person i , and an increase of the cardinality of the current assignment by 1. Thus, it will produce a complete assignment in exactly n iterations, starting from an empty assignment, while maintaining ϵ -CS throughout the process (under our assumption that the problem is feasible so a complete assignment exists).

Example 3.2

To illustrate the coalition expansion process, let us consider a 4×4 version of the 3×3 problem of Figs. 2.2 and 2.4. Here, in addition to the three persons and objects of these figures, there are a fourth person 4 and




Figure 3.8 Illustration of the variant of the cooperative auction iteration that involves an expanding coalition. We consider a 4×4 version of the 3×3 problem of Figs. 2.2 and 2.4, as shown above. The initial prices are $p = (0, 0, 0, 0)$ and the initial partial assignment is $\{(1, 1), (2, 2), (4, 3)\}$. We assume that $\epsilon < 1/n = 1/4$ (to guarantee that the final assignment is optimal).

The cooperative auction iteration starts with the unassigned person 3, and constructs the coalition of persons 1, 2, and 3, similar to Fig. 3.4. The prices of objects 1 and 2 rise to $C + \epsilon$, thus bringing the assigned object 3 into the ϵ -zone of the coalition partners 1, 2, and 3. In the expanding coalition variant of the cooperative iteration, the search for coalition partners continues, adding person 4 to the coalition. A new price rise of objects 1, 2, and 3 (by $1 + \epsilon$ units) is then performed. This brings object 4 into the ϵ -zone of person 4 and allows the augmentation $(3, 4, 4)$, and termination with the assignment $\{(1, 1), (2, 2), (3, 3), (4, 4)\}$ and prices $\bar{p} = (C + 1 + 2\epsilon, C + 1 + 2\epsilon, 1 + \epsilon, 0)$. It can be seen that the final assignment and prices satisfy ϵ -CS.

object 4, as shown in Fig. 3.8. Person 4 can be assigned to object 3 with value 0 and to object 4 with value -1. Every feasible assignment must include the pair $(4, 4)$, so the optimal assignments are the ones of the 3×3 problem, augmented with $(4, 4)$, such as for example

$$\{(1, 1), (2, 2), (3, 3), (4, 4)\}. \quad (3.1)$$

Let the initial prices be $p = (0, 0, 0, 0)$ and the initial partial assignment be

$$\{(1, 1), (2, 2), (4, 3)\},$$

as shown in Fig. 3.8. The cooperative auction iteration starts with the unassigned person 3, and constructs the coalition $\mathcal{C}(3) = \{1, 2, 3\}$. In the expanding coalition variant, the search for coalition partners continues after

the price rise of objects 1 and 2 (by the amount $C + \epsilon$), adding person 4 to the coalition, which brings object 4 into the ϵ -zone of person 4, and allows the augmentation $(3, 4, 4)$ and termination with the assignment (3.1) .

Here is a more complicated example, which also demonstrates the potentially significant computational savings for reusing the computation of previous coalitions to save in the computation of subsequent coalitions.

Example 3.3 (Computational Advantage of Expanding Coalitions)

Consider the $n \times n$ assignment example of Fig. 3.9 (the values a_{ij} are shown above the lines connecting persons and objects). All persons are assigned as shown, except for person i , who initiates a cooperative auction iteration with $\epsilon = 0$. The starting object prices are $p = (0, \dots, 0)$, and satisfy CS together with the partial assignment shown.

Let us apply the cooperative auction algorithm with expanding coalitions and $\epsilon < 0.5$. The starting coalition is $\mathcal{C}(i) = \{i, i_1, i_2\}$ and the price rise of the set of objects $\mathcal{O}(i) = \{j_1, j_2\}$ is $r = 0.5$, which brings object j_3 into the 0-zone of person i_2 . A new iteration is started by person i , with coalition $\mathcal{C}(i) = \{i, i_1, i_2, i_3\}$ and price rise of the set of objects $\mathcal{O}(i) = \{j_1, j_2, j_3\}$ equal to $r = 0.5$, which brings object j_4 into the 0-zone of person i_3 . This coalition expansion process continues for $n - 3$ iterations, up to when person i_{n-1} is included in the coalition, the prices of objects j_1, \dots, j_{n-1} rise by 0.5, which brings object j into the 0-zone of person i_{n-1} , with an augmentation ensuing along the augmenting path $(i, i_2, i_3, \dots, i_{n-1}, j)$. The assignment thus obtained is complete and optimal.

This process requires $n - 3$ iterations, and $O(n^2)$ computation (because each of the $n - 3$ coalitions is rebuilt from scratch). If it is carried out with the expanding coalition variant, it requires a single iteration with $n - 3$ coalition expansions, and $O(n)$ computation.

Suppose now that we use $\epsilon \geq 0.5$. Then every object is contained in the ϵ -zone of some person, the starting coalition $\mathcal{C}(i)$ is the entire person set $\{i, i_1, \dots, i_{n-1}\}$, and the algorithm terminates in one iteration, without any coalition expansion. There are two possible augmentations from i to j :

$$(i, i_1, i_2, i_3, \dots, i_{n-1}, j) \quad \text{and} \quad (i, i_2, i_3, \dots, i_{n-1}, j),$$

and two corresponding complete assignments. The first of these is suboptimal while the second is optimal. The solution generated depends on the order in which persons i_1 and i_2 enter $\mathcal{C}(i)$. This illustrates how the number of coalition expansions may be reduced with larger values of ϵ .

3.2 Combinations with Noncooperative Auction Algorithms

We will now explore the possibility of combining the noncooperative auction algorithms (both conservative and aggressive) with the cooperative algorithm. In particular, we are given a set of object prices $p = (p_1, \dots, p_n)$ and a partial assignment \mathcal{A} satisfying ϵ -CS. The iteration starts with an unassigned person i and tries to generate the set of coalition partners of i . In the process it will perform an aggressive (or conservative)




Figure 3.9 Illustration of the multiple coalition expansions algorithm with $\epsilon = 0$; cf. Example 3.3. The values a_{ij} are shown above the lines connecting persons and objects, and the initial prices are all 0. The algorithm with multiple coalition expansions requires a single iteration (with multiple coalition expansions) and $O(n)$ computation to assign person i , while the algorithm without coalition expansions (cf. Section 3) requires n iterations and $O(n^2)$ computation.

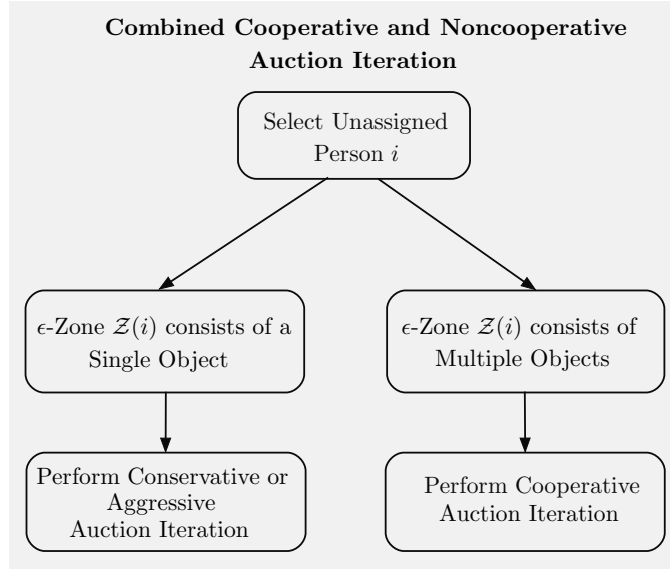


Figure 3.10 Block diagram of a combined cooperative and noncooperative auction iteration.

auction iteration if $\mathcal{Z}(i)$, the ϵ -zone of i , contains a single object and $\epsilon > 0$ (or $\epsilon = 0$, respectively), and a cooperative auction iteration otherwise; see Fig. 3.10. The intuitive idea is that *when $\mathcal{Z}(i)$ consists of a single object, there can be a most one coalition partner of i , so a price war is not possible*. This favors the use of a noncooperative auction iteration.

The iteration just described in summary is stated in detail as follows.

Combined Cooperative and Noncooperative Auction Iteration

Given a set of object prices $p = (p_1, \dots, p_n)$ and a partial assignment \mathcal{A} satisfying ϵ -CS, select an unassigned person i . If the ϵ -zone $\mathcal{Z}(i)$ contains a single object perform a noncooperative auction iteration (conservative if $\epsilon = 0$ or aggressive if $\epsilon > 0$). Otherwise perform a cooperative auction iteration.

Note that the iteration can optionally be used with or without coalition expansions. In the former case a cooperative auction iteration is simply continued starting from the same person i , up to the point where an augmentation takes place. It should be noted that the cooperative auction iteration with coalition expansions and $\epsilon = 0$ bears similarity to the Hungarian method, which is typically inferior both in theory and in practice to efficiently implemented aggressive auction iterations. Moreover, its theoretical complexity is known to be inferior to the one of the aggressive auction algorithm. On the other hand, combinations of conservative auction and the Hungarian method have worked well in practice, as verified by the computations given in the author's paper [Ber81], and by the experience with the JV code [JoV87]. The combined iteration

given above, with or without coalition expansions, is new for $\epsilon > 0$, and has not been adequately tested, but with proper implementation, is expected to work more efficiently than either one of its cooperative and noncooperative components working in isolation.

The evaluation of the performance of the combined aggressive and cooperative auction iteration, with $\epsilon > 0$ and the expanding coalition process, in conjunction with ϵ -scaling, is an issue of great interest, both theoretically and experimentally. A relevant fact here is that two-phase algorithms, which involve aggressive auction ($\epsilon > 0$) in the first phase and a Hungarian-like algorithm ($\epsilon = 0$) in the second phase after most of the objects have been assigned, have been shown to have computational complexity that is superior to either aggressive auction or the Hungarian method in isolation of each other. In particular, Orlin and Ahuja [OrA92] have derived a related

$$O(\sqrt{n} m \log(nC)) \quad (3.2)$$

worst-case complexity result for a two-phase algorithm of this type, with the threshold for switching between the two phases skillfully chosen (see also Chapter 5, Exercise 4.5 of the book [BeT89], with solution included in the internet-posted version of the book). The use of $\epsilon > 0$ together with ϵ -scaling, requires fewer coalition expansions, as can be seen from Example 3.3, and seems to be a natural alternative way to deal with a large number of coalition expansions for many problems. Thus it is reasonable to conjecture that a complexity estimate like the one of Eq. (3.2) can be proved for some version of the combined aggressive and cooperative auction iteration.

3.3 Properties of the Cooperative Auction Algorithm

In this section we will discuss some general properties and implementations of cooperative auction. We first note that if there is no augmenting path starting from i [i.e., $\mathcal{M}(i)$ is empty], the set of objects that are assigned to some coalition partner of i , is the union of the ϵ -zones of the persons in the coalition of i :

$$\mathcal{O}(i) = \cup_{i' \in \mathcal{C}(i)} \mathcal{Z}(i'). \quad (3.3)$$

To see this, note that when $\mathcal{M}(i)$ is empty, all objects in the ϵ -zones of i and his/her coalition partners must be assigned to some coalition partner of i (otherwise an augmentation would be performed).

Let us now provide an explicit formula for the common price rise for the case where $\mathcal{M}(i)$ is empty. For each person $i' \in \mathcal{C}(i)$, consider the scalar

$$r_{i'} = \epsilon + \min_{j \in \mathcal{Z}(i')} \{a_{i'j} - p_j\} - \max_{j \notin \mathcal{O}(i), j \in A(i')} \{a_{i'j} - p_j\}, \quad (3.4)$$

(by convention, the maximum above is $-\infty$ if the set over which the maximum is taken is empty). It can be seen from Fig. 3.11 that $r_{i'}$ is the maximum price rise of the objects in $\mathcal{O}(i)$ that will keep every object in $\mathcal{Z}(i')$, the ϵ -zone of i' , within $\mathcal{Z}(i')$ following the price rise. To keep *all* the objects in $\mathcal{O}(i) = \cup_{i' \in \mathcal{C}(i)} \mathcal{Z}(i')$

within the ϵ -zone of either i or some coalition partner of i , the common price rise should not exceed any one of the amounts $r_{i'}$, $i' \in \mathcal{C}(i)$. Thus the maximum possible common price rise is

$$r = \min_{i' \in \mathcal{C}(i)} r_{i'}. \quad (3.5)$$

Moreover, following the price rise, the union of the ϵ -zones of the persons in $\mathcal{C}(i)$ consists of $\mathcal{O}(i)$ and a nonempty set $\overline{\mathcal{O}}(i)$ of additional objects. This is the set of objects that attain the maximum in the maximization of Eq. (3.4), $\max_{j \notin \mathcal{O}(i), j \in A(i'')} \{a_{i''j} - p_j\}$, while i'' attains the minimum in Eq. (3.5). This shows that the set of objects $\overline{\mathcal{O}}(i)$ obtained at the end of the iteration is nonempty and that r is finite (otherwise the existence of a complete assignment assumption would be violated). Note that if any of the objects within $\overline{\mathcal{O}}(i)$, say object j , is unassigned, we can perform an augmentation that starts at i and ends at j , and increase the price of j by the maximum amount that will not violate ϵ -CS. This can be done efficiently, and it is generally recommended, as it increases the cardinality of the assignment by one, but for simplicity, we have not stated this explicitly. Alternatively, we may suitably modify the cooperative iteration description, so an augmentation is automatically performed, if possible, following a price rise. Figure 3.12 provides an illustration.

We summarize the principal conclusions from the preceding discussion in the following proposition.

Proposition 3.1: Consider the cooperative auction iteration under the assumption that there is no augmenting path starting from i , i.e., $\mathcal{M}(i)$ is empty. Consider also $\mathcal{O}(i)$, the set of objects that are assigned to some coalition partner of i . Then following a price rise:

- (a) $\mathcal{O}(i)$ is equal to the union of the ϵ -zones of all persons in $\mathcal{C}(i)$.
- (b) The prices of the objects in $\mathcal{O}(i)$ are raised by the common increment

$$r = \min_{i' \in \mathcal{C}(i)} r_{i'},$$

where $r_{i'}$ is given by Eq. (3.4), and we have $r > \epsilon$.

- (c) Following the price rise, the union of the ϵ -zones of the persons in $\mathcal{C}(i)$ consists of $\mathcal{O}(i)$ and a nonempty set $\overline{\mathcal{O}}(i)$ of additional objects.
- (d) If any of the objects within $\overline{\mathcal{O}}(i)$, say object j , is unassigned, an augmentation that starts at i and ends at j can be performed. Moreover, the prices and assignment obtained following this augmentation satisfy ϵ -CS.

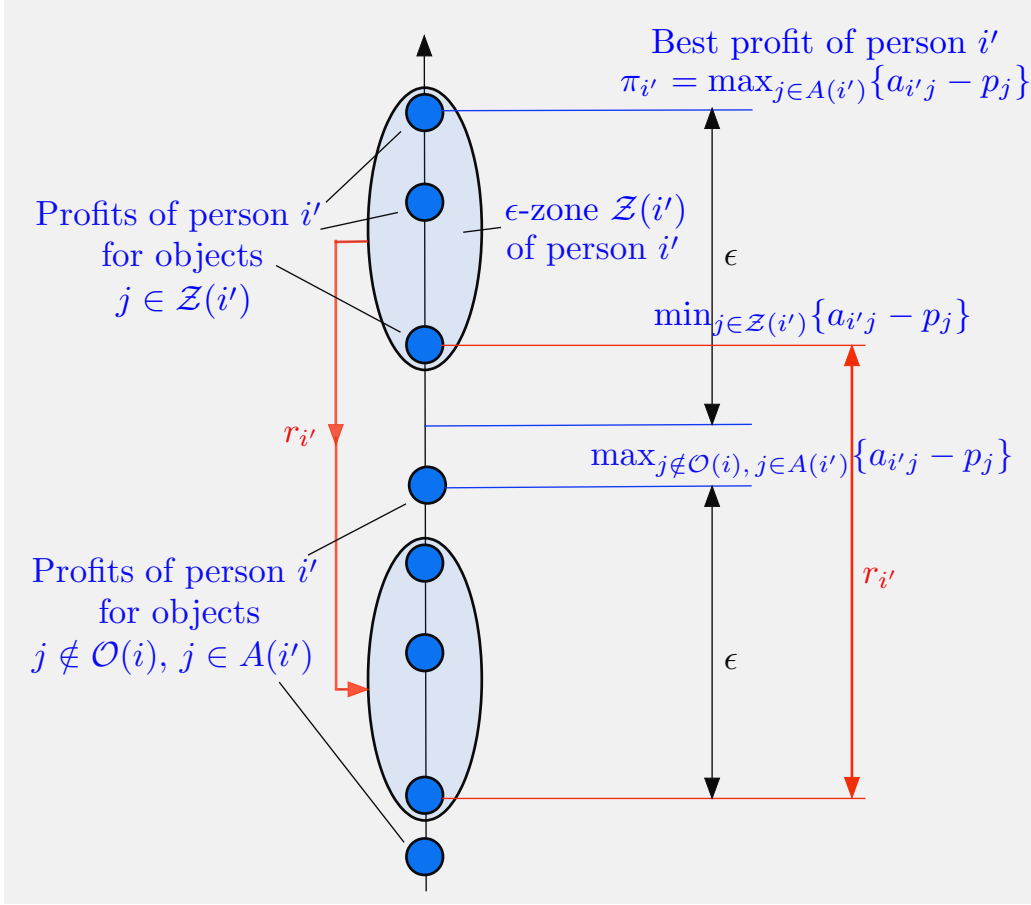


Figure 3.11 Illustration of a price rise (and corresponding profit drop) of the objects in the ϵ -zone $\mathcal{Z}(i')$ of a person $i' \in \mathcal{C}(i)$. The figure shows the maximum amount $r_{i'}$ by which we can raise the prices of objects j in the ϵ -zone $\mathcal{Z}(i')$, while guaranteeing that all $j \in \mathcal{Z}(i')$ will stay within the ϵ -zone of i' following a cooperative price rise. It is given by

$$r_{i'} = \epsilon + \min_{j \in \mathcal{Z}(i')} \{a_{i'j} - p_j\} - \max_{j \notin \mathcal{O}(i), j \in A(i')} \{a_{i'j} - p_j\},$$

cf. Eq. (3.4).

In the above figure the ϵ -zone $\mathcal{Z}(i')$ consists of the three most profitable objects of person i' with profits within the top ellipse. The figure also shows the profits of the objects $j \notin \mathcal{O}(i)$ with $j \in A(i')$ [the profits of any additional objects $j \in \mathcal{O}(i)$ with $j \in A(i')$ but $j \notin \mathcal{Z}(i')$ are not shown]. The figure assumes that the set

$$\{j \mid j \notin \mathcal{O}(i), j \in A(i')\}$$

is nonempty; if it is not, we have $r_{i'} = \infty$.

After we raise the prices of the objects in $\mathcal{O}(i)$ by $r_{i'}$, the profits of objects in $\mathcal{Z}(i')$ move downward by $r_{i'}$, just within ϵ of the fourth object, which now becomes the most profitable. Price rises by amounts smaller than $r_{i'}$ still keep the three most profitable objects within the ϵ -zone $\mathcal{Z}(i')$, but may not be sufficient to bring the fourth object into $\mathcal{Z}(i')$.

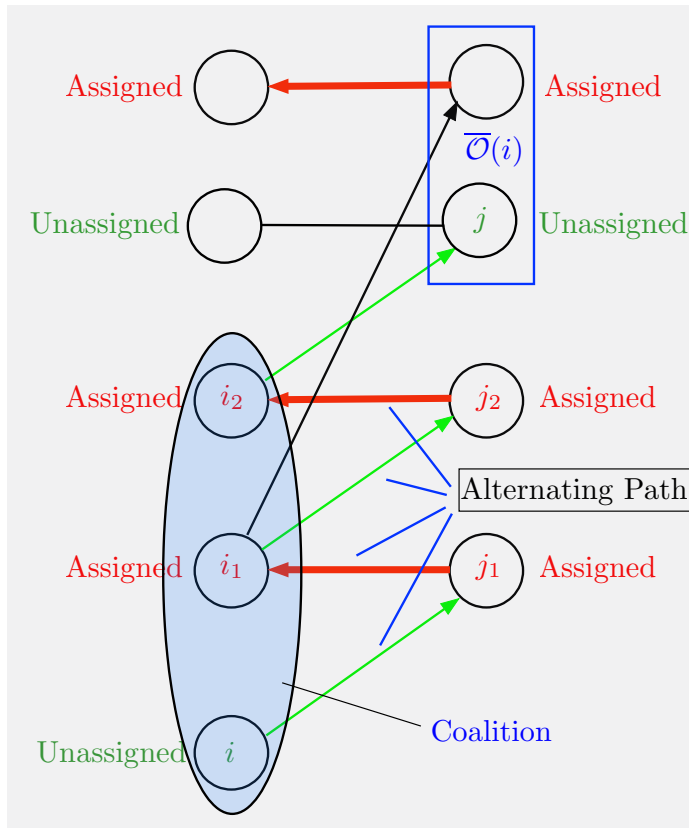


Figure 3.12 Illustration of the cooperative auction iteration, where after raising the prices, there is an object in the set $\bar{O}(i)$ that is unassigned. Then, we may optionally perform an augmentation, shown in green, along a corresponding augmenting path $[(i, i_1, i_2, j)]$ in the figure]. In particular, persons i , i_1 , and i_2 get assigned to j_1 , j_2 , and j , respectively.

The computational complexity of the algorithm is not expected to be better than the one of the aggressive auction algorithm $[O(nm \log(nC))]$, where m is the number of arcs of the bipartite graph representing the assignment problem and C is the range of values, given by Eq. (2.7)]. However, depending on the implementation and the type of problem addressed, it appears that the cooperative auction algorithm, as given in this section, can outperform the aggressive auction algorithm, particularly in situations where price wars are likely.

3.4 Common Price Increment Computation

We will now focus on the most complicated part of a cooperative auction iteration, namely the computation of the common price rise increment r of Eqs. (3.4)-(3.5), and the new set of objects $\bar{O}(i)$ that are subsequently brought into the coalition of i , when there is no augmenting path starting from i . We will first describe one possible implementation that can be interpreted graphically, and we will subsequently provide a more general implementation in pseudocode. We assume that we are given a set of object prices $p = (p_1, \dots, p_n)$

and a partial assignment \mathcal{A} satisfying ϵ -CS, together with an unassigned person i to start the iteration. We also assume that the ϵ -zone $\mathcal{Z}(i)$ contains multiple objects all of which are assigned, and that no augmenting path starting from i exists.

In particular, we will use the layered graph shown in Fig. 3.13 to illustrate the computation of:

- (a) The set of coalition persons $\mathcal{C}(i)$.
- (b) The set of coalition objects $\mathcal{O}(i)$, i.e., the set of objects assigned to the persons in $\mathcal{C}(i)$.
- (c) The common price rise r of the objects in $\mathcal{O}(i)$.

In this computation we break down the sets $\mathcal{O}(i)$ and $\mathcal{C}(i)$ into layers of disjoint subsets $\mathcal{O}_1, \mathcal{C}_1, \dots, \mathcal{O}_k, \mathcal{C}_k$, where for some positive integer $k < n$, and for $m = 1, \dots, k$:

The m th person layer \mathcal{C}_m is the set of persons i' such that every alternating path that starts at i and ends at i' contains at least m persons other than i .

The m th object layer \mathcal{O}_m is the set of objects that are assigned to the persons in \mathcal{C}_m .

The layers \mathcal{O}_m and \mathcal{C}_m are computed successively, and can be visualized in terms of the tree of alternating paths shown in Fig. 3.13. The details of the computation are as follows:

Layer Construction

- (a) We construct \mathcal{O}_1 , which is the set of objects in the ϵ -zone $\mathcal{Z}(i)$ of person i , and then \mathcal{C}_1 , which is the set of persons assigned to the objects in \mathcal{O}_1 .
- (b) Given \mathcal{C}_m , we construct \mathcal{O}_{m+1} as the set of objects $j \notin \mathcal{O}_1 \cup \dots \cup \mathcal{O}_m$ that belong to the ϵ -zone of at least one person in \mathcal{C}_m ; if \mathcal{O}_{m+1} is empty, then we stop (i.e., $m = k$), having computed $\mathcal{O}(i)$ and $\mathcal{C}(i)$ according to

$$\mathcal{O}(i) = \mathcal{O}_1 \cup \dots \cup \mathcal{O}_k, \quad \mathcal{C}(i) = \{i\} \cup \mathcal{C}_1 \cup \dots \cup \mathcal{C}_k.$$

In the process of constructing the layers $\mathcal{O}_1, \mathcal{C}_1, \dots, \mathcal{O}_k, \mathcal{C}_k$, we obtain the set of *border objects*, denoted \mathcal{B} , and consisting of the objects that do not belong to $\mathcal{O}(i)$ but can be matched with a person in the coalition $\mathcal{C}(i)$, i.e.,

$$\mathcal{B} = \{j \notin \mathcal{O}(i) \mid j \in A(i') \text{ for some } i' \in \mathcal{C}(i)\};$$

see Fig. 3.13. The border objects are obtained during the process of constructing the sets $\mathcal{O}(i)$ and $\mathcal{C}(i)$ as described earlier.

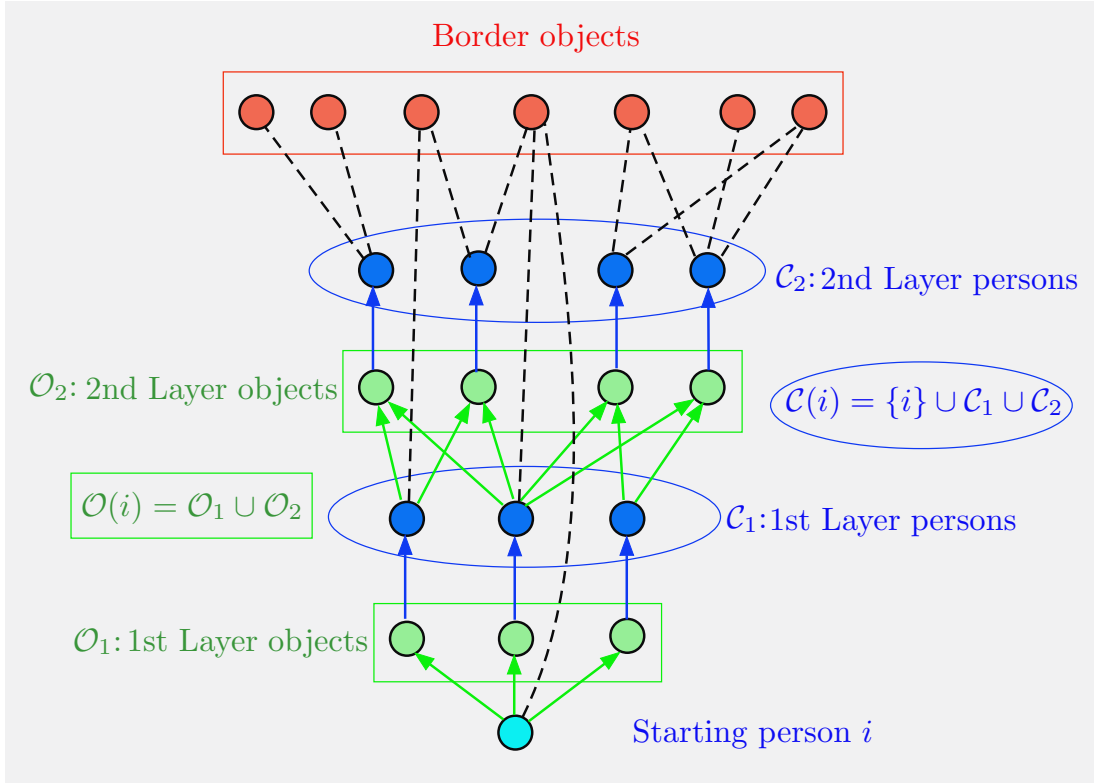


Figure 3.13 Illustration of a search tree to construct the coalition of i ,

$$\mathcal{C}(i) = \{i\} \cup \mathcal{C}_1 \cup \mathcal{C}_2,$$

assuming no augmentation occurs during the cooperative auction iteration (the figure assumes two object layers $\mathcal{O}_1, \mathcal{O}_2$, and two person layers $\mathcal{C}_1, \mathcal{C}_2$). The objects in \mathcal{O}_1 are the ones in the ϵ -zone $\mathcal{Z}(i)$. The objects in \mathcal{O}_2 are the ones that do not belong to \mathcal{O}_1 but belong to the ϵ -zone $\mathcal{Z}(i')$ of some person of \mathcal{C}_1 . The set \mathcal{B} of border objects consists of all objects j that do not belong to

$$\mathcal{O}(i) = \mathcal{O}_1 \cup \mathcal{O}_2,$$

but can be matched with some person $i' \in \mathcal{C}(i)$, i.e.,

$$\mathcal{B} = \{j \notin \mathcal{O}(i) \mid j \in A(i') \text{ for some } i' \in \mathcal{C}(i)\}.$$

Green arrows indicate pairs (i', j') such that $i' \in \mathcal{C}(i)$ and $j' \in \mathcal{Z}(i')$. Broken lines indicate pairs (i', j') such that $i' \in \mathcal{C}(i)$, $j' \in A(i')$ but $j' \notin \mathcal{O}(i)$. The paths from i to the (blue) nodes in $\mathcal{C}_1 \cup \mathcal{C}_2$ are the shortest alternating paths.

Simultaneously with the computation of $\mathcal{O}(i)$, $\mathcal{C}(i)$, and \mathcal{B} as described above, we can also compute the cooperative price rise amount of the iteration using Eq. (3.5):

$$r = \epsilon + \min_{i' \in \mathcal{C}(i)} \left\{ \hat{\pi}_{i'} + \min_{j \in \mathcal{B}, j \in A(i')} \{p_j - a_{i'j}\} \right\}, \quad (3.6)$$

where $\hat{\pi}_{i'}$ is given by

$$\hat{\pi}_{i'} = \min_{j \in \mathcal{Z}(i')} \{a_{i'j} - p_j\}.$$

Combining the preceding equations with Eq. (3.5) and interchanging the order of minimizations in Eq. (3.6), we obtain

$$r = \epsilon + \min_{j \in \mathcal{B}} \min_{i' \in \mathcal{C}(i), j \in A(i')} \{\hat{\pi}_{i'} + p_j - a_{i'j}\} = \epsilon + \min_{j \in \mathcal{B}} d_j, \quad (3.7)$$

where for all $j \in \mathcal{B}$

$$d_j = \begin{cases} \min_{i' \in \mathcal{C}(i)} \{\hat{\pi}_{i'} + p_j - a_{i'j}\} & \text{if } j \in A(i') \text{ for some } i' \in \mathcal{C}(i), \\ \infty & \text{otherwise.} \end{cases}$$

To understand the intuitive meaning of d_j , we first note that $\hat{\pi}_{i'}$ is the profit of person i' , assuming i' is awarded the least profitable of the objects in his/her ϵ -zone. Then we can view d_j as a *profit loss* incurred when person i' is reassigned to j from his/her least profitable object within $\mathcal{Z}(i')$. The common price rise r of Eq. (3.6) can be interpreted as ϵ plus the minimum possible profit loss some person i' is reassigned to some $j \in \mathcal{B}$ from his/her least profitable object in $\mathcal{Z}(i')$. Note also that each reassignment of a person $i' \in \mathcal{C}(i)$ to an object in $\mathcal{O}(i)$, in the course of an augmentation, involves a loss or gain in profit of at most ϵ , since the objects assigned to i' before and after the augmentation both belong to the ϵ -zone $\mathcal{Z}(i')$.

Note that $\hat{\pi}_{i'}$ can be computed while we go over the set of associated objects $A(i')$ of person i' , to determine whether they can be added to $\mathcal{O}(i)$. Thus the computation of r can be organized progressively: first update the quantity d_j , as new persons i' are added to the coalition $\mathcal{C}(i)$, and then at the end of the iteration, after $\mathcal{C}(i)$ and \mathcal{B} are obtained, take the minimum over $j \in \mathcal{B}$ of d_j to obtain r ; cf. Eq. (3.7). Also the set of objects $\overline{\mathcal{O}}(i)$ that enter the ϵ -zone of at least one person in the coalition $\mathcal{C}(i)$ following the price rise, include the ones that attain the minimum of $d_{j'}$ over $j' \in \mathcal{B}$.

A More General Implementation of the Coalition Construction Process

Let us now provide pseudocode for a more general implementation of the cooperative auction iteration that constructs the sets $\mathcal{C}(i)$, \mathcal{B} , the scalar r of Eq. (3.7), and an augmenting path (if one is discovered in the course of the iteration). The code uses two temporary lists of persons C and C' . At the end of the iteration, $C = \mathcal{C}(i)$ and $C' = \emptyset$. The code also uses an array \mathcal{L} of labels, with label $\mathcal{L}(j)$ corresponding to object j . The labels $\mathcal{L}(j)$ are initially set to 0, and are updated to record a predecessor person i' of j in a potential augmenting path, i.e., $j \in \mathcal{Z}(i')$. The labels are used to trace backwards an augmenting path from an unassigned object to the unassigned person i that starts the iteration [if one is discovered, in which case we go to the next iteration, without completing the computation of $\mathcal{C}(i)$ and \mathcal{B}].

Pseudocode to Construct the Sets $\mathcal{C}(i)$ and \mathcal{B} , or Construct an Augmenting Path

Initialization: $C = \emptyset$, $C' = \{i\}$, $\mathcal{B} = \{1, \dots, n\}$, $d_j = \infty$ and $\mathcal{L}(j) = 0$ for all $j \in \{1, \dots, n\}$.

Until $C' = \emptyset$:

Remove a person i' from C' and add it to C . Let $\hat{\pi}_{i'} = \min_{j \in \mathcal{Z}(i')} \{a_{i'j} - p_j\}$. For all $j \in A(i') \cap \mathcal{B}$:

- If $j \in \mathcal{Z}(i')$ and j is assigned to a person i'' , remove j from \mathcal{B} , and add i'' to C' if it is not already in C' . Moreover, if $\mathcal{L}(j) = 0$, set $\mathcal{L}(j) = i'$.
- If $j \in \mathcal{Z}(i')$ and j is unassigned, perform an augmentation that starts at i and ends at j , by tracing labels backwards from j to the unassigned person i , along the augmenting path defined as follows:

$$j \rightarrow i' \rightarrow j_1 = \text{Object Assigned to } i' \rightarrow i_1 = \mathcal{L}(j_1) \\ \rightarrow j_2 = \text{Object Assigned to } i_1 \rightarrow i_2 = \mathcal{L}(j_2) \rightarrow \dots \rightarrow j_k \rightarrow i,$$

where j_k is an object in the ϵ -zone $\mathcal{Z}(i)$, so $i = \mathcal{L}(j_k)$; see Fig. 3.14. Go to the next iteration.

- If $j \notin \mathcal{Z}(i')$, set $d_j \leftarrow \min\{d_j, \hat{\pi}_{i'} + p_j - a_{i'j}\}$.

Set $\mathcal{B} \leftarrow \{j \in \mathcal{B} \mid d_j < \infty\}$, $\mathcal{C}(i) = C$, $r = \epsilon + \min_{j \in \mathcal{B}} d_j$.

It can be verified that different rules for choosing the person i' to be removed from C' will lead to the same sets $\mathcal{C}(i)$ and \mathcal{B} , and the same price rise r . On the other hand, one may or may not obtain the layered structure illustrated in Fig. 3.13, which corresponds to a special rule for choosing i' . This is the rule that removes the persons i' from C' in the same order in which they entered C' . Other rules may also be considered based on a heuristic or more principled rationale in a given problem.

4. ADDITIONAL AUCTION VARIANTS

There are a number of interesting variations of the cooperative auction algorithm, in addition to those we have discussed so far. Most of these variations are aimed at accelerating convergence, mitigating as much as possible the effects of price wars, and enhancing the suitability for parallel computation. Several of these variations have similar theoretical properties. However, their practical performance may be significantly affected by the character of the specific problem that is being solved, such as graph density/sparsity, large/small range of values a_{ij} , and special characteristics of the graph's structure, such as large/small "diameter" (a measure of the average number of hops between two randomly chosen persons).

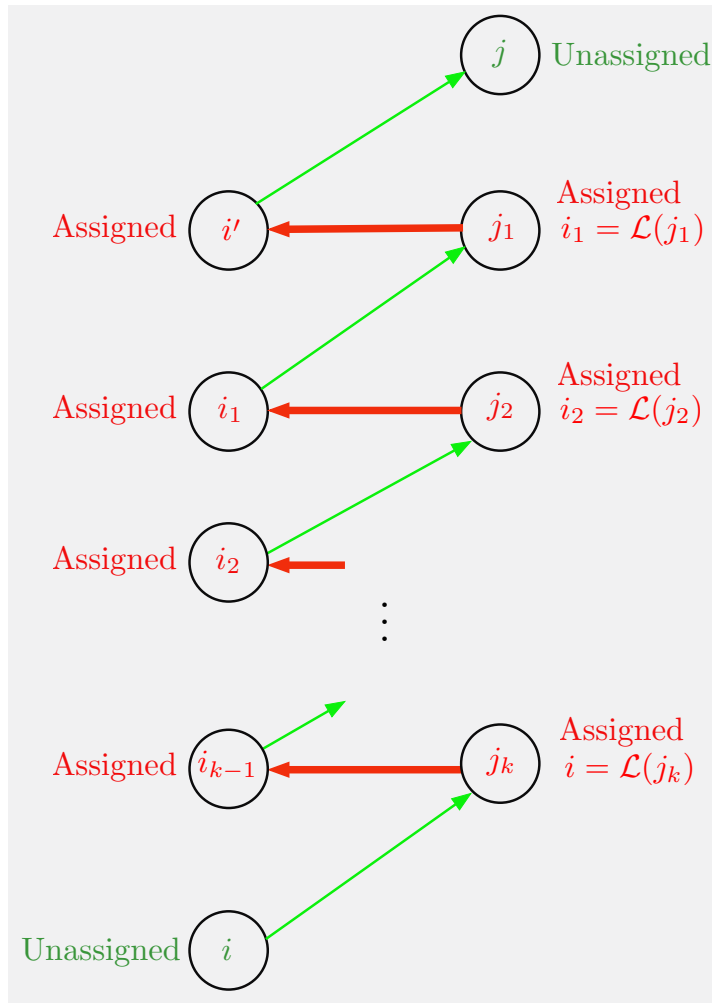


Figure 3.14 Illustration of an augmenting path from the unassigned object j to the unassigned person i ,

$$j \rightarrow i' \rightarrow j_1 \rightarrow i_1 \rightarrow j_2 \rightarrow i_2 \rightarrow \cdots i_{k-1} \rightarrow j_k \rightarrow i,$$

which is discovered during an iteration that starts with the unassigned person i .

In what follows in this section we will review a number of algorithmic ideas that form the basis for variants of conservative, aggressive, and cooperative algorithms, and their combinations. The wide spectrum of possibilities suggests a view of an auction algorithmic landscape where there is no universal best choice that works optimally for all problems. This view is supported by extensive computational results in the paper [Ber95b], which tested comparatively some (but by no means all) of the algorithmic ideas discussed in the present paper within a broader context of network optimization problems.

Cooperative Auction Iteration With Collective Bidding and Person Reassignments

This variant of the cooperative auction iteration aims to bring it closer to the aggressive auction iteration, at the expense of foregoing the option of expanding coalitions. Consider the cooperative iteration for the

case where there is no augmenting path [$\mathcal{M}(i)$ is empty]. Then after the subsequent collective price rise, the union of the ϵ -zones of the persons in $\mathcal{C}(i)$ consists of $\mathcal{O}(i)$ and a nonempty set $\overline{\mathcal{O}}(i)$ of additional objects, as we have discussed in Section 3.3. There are now two possibilities:

- (a) *There is an unassigned object \bar{j} within the set $\overline{\mathcal{O}}(i)$.* Then as we discussed earlier, an augmenting path is created following the price rise, which starts at i and ends at \bar{j} (cf. Example 3.1 and Fig. 3.12). This augmentation can be performed immediately, without waiting for the next iteration to discover it.
- (b) *All objects in the set $\overline{\mathcal{O}}(i)$ are assigned.* In this case, the cooperative auction algorithm first presented in Section 3 simply goes to the next iteration. However, there is also a possibility to assign person i through a reassignment of the coalition persons and a rearrangement of the corresponding assigned pairs. This is illustrated in Fig. 4.1, which should be contrasted with Fig. 3.12.

The corresponding auction algorithm variant is identical to the cooperative auction iteration of Section 3, except for the additional person reassignment process, which is performed following a price rise that does not result in a subsequent augmentation; see Fig. 4.2. We state this variant formally as follows:

Cooperative Auction Iteration With Person Reassignments

Given a set of object prices $p = (p_1, \dots, p_n)$ and a partial assignment \mathcal{A} satisfying ϵ -CS, select an unassigned person i . Let $\mathcal{M}(i)$ be the set of augmenting paths that start with i .

- If $\mathcal{M}(i)$ is nonempty, perform an augmentation along some augmenting path from $\mathcal{M}(i)$, increase the price of the last object in this augmenting path by the maximum amount that will not violate ϵ -CS, and go to the next iteration.
- If $\mathcal{M}(i)$ is empty, let $\mathcal{O}(i)$ denote the set of objects that are assigned to some coalition partner of i . Raise the prices of the objects in $\mathcal{O}(i)$ by the maximum common amount for which the ϵ -zone of every person i' in $\mathcal{C}(i)$ is a subset of the ϵ -zone of the same person i' after the price rise.

Let $\overline{\mathcal{O}}(i)$ denote the set of objects $j \notin \mathcal{O}(i)$, which following the price rise, belong to the ϵ -zone of a person in $\mathcal{C}(i)$. Select an object $\bar{j} \in \overline{\mathcal{O}}(i)$, with preference given to unassigned objects. Let (i, i_1, \dots, i_k) be an alternating path such that \bar{j} is in the ϵ -zone of person i_k following the price rise of the coalition objects $\mathcal{O}(i)$. Let also j_1, \dots, j_k be the objects that are assigned to the persons i_1, \dots, i_k in the current assignment \mathcal{A} . Then change \mathcal{A} by assigning i to j_1 , i_m to j_{m+1} for $m = 1, \dots, k-1$, and i_k to \bar{j} ; any person assigned to \bar{j} under \mathcal{A} becomes unassigned. Finally, raise the price of \bar{j} by the maximum amount that will not violate ϵ -CS, and go to the next iteration.

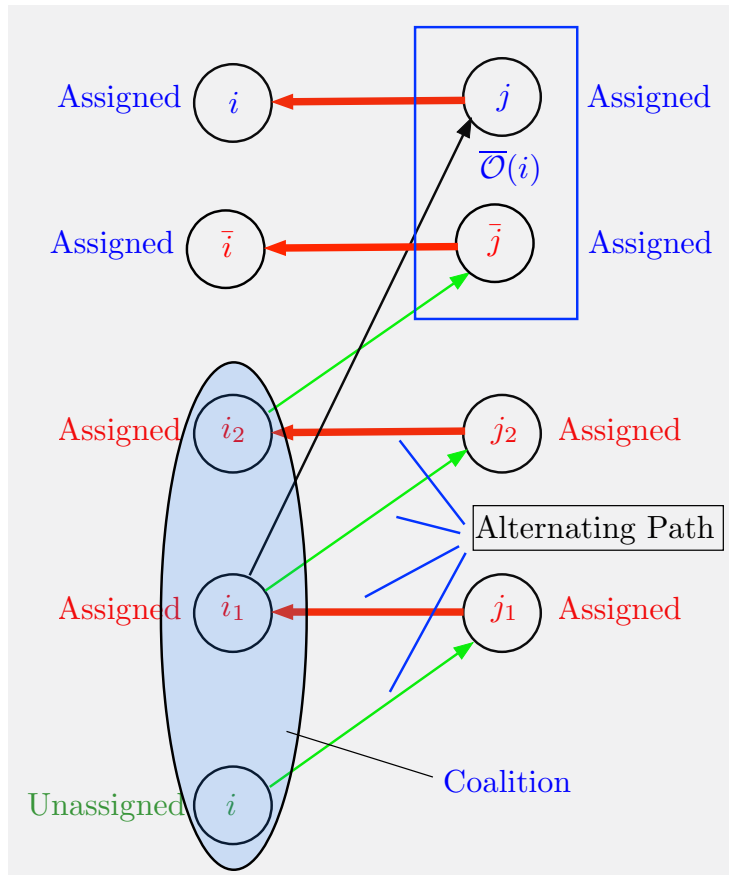


Figure 4.1 Illustration of the cooperative auction iteration with person reassignments when the set of augmenting paths $\mathcal{M}(i)$ is empty, and all the objects in the set $\overline{\mathcal{O}}(i)$ are assigned. Then, we choose an object $\bar{j} \in \overline{\mathcal{O}}(i)$ and perform a reassignment of persons to objects, shown in green, along a corresponding alternating path $[(i, i_1, i_2)$ in the figure]. In particular, persons i , i_1 , and i_2 get assigned to j_1 , j_2 , and \bar{j} , respectively, while the person \bar{i} that is assigned to \bar{j} under \mathcal{A} becomes unassigned. Note that the object \bar{j} is not unique: any object in $\overline{\mathcal{O}}(i)$ (such as j in the figure) and alternating path corresponding to that object [such as (i, i_1) in the figure] can be used.

The person reassignments in the preceding iteration can be viewed as a collective bid, which aims to acquire a new object for the coalition $\mathcal{C}(i)$, at the expense of deassigning a person from outside the coalition. The reassignments bring the iteration closer in spirit to the aggressive auction algorithm. In particular, it can be seen that *if $\mathcal{Z}(i)$ consists of a single object (assigned or unassigned), the preceding iteration behaves identically with the aggressive auction iteration.* On the other hand, we should also note that person reassignments do not allow the use of coalition expansions.

ϵ -Scaling Variations and Auction Initialization

The use of ϵ -scaling may or may not be necessary for the cooperative auction algorithm of Section 3 and its variants. After all, with $\epsilon = 0$ the cooperative algorithm is known to be reliable and to perform well for many problems, particularly those involving a dense assignment graph (this has been established by a

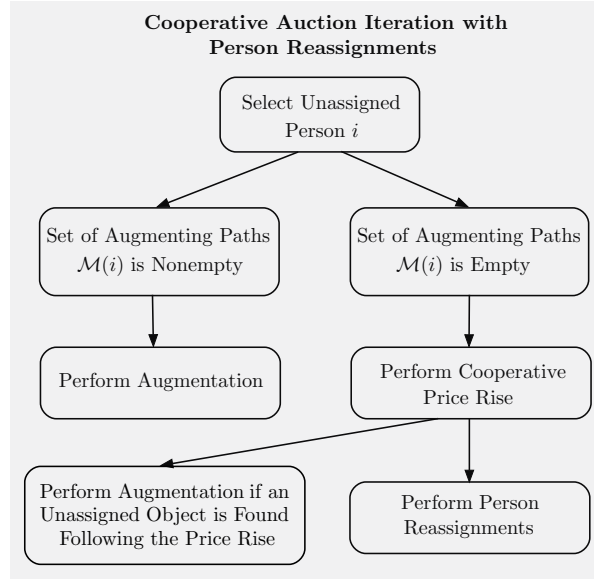


Figure 4.2 Block diagram of a cooperative auction iteration with person reassignments. Here, if an unassigned object is not discovered following a cooperative price rise, person i assigned through a reassignment of the coalition persons and a rearrangement of the corresponding assigned pairs.

number of studies starting with the original paper [Ber81]). On the other hand, ϵ -scaling may be needed to improve the robustness and the performance of both the aggressive and the cooperative algorithms for the case of a sparse assignment graph.

A critical step in ϵ -scaling is when a complete assignment is obtained with some value of ϵ and then, to run the algorithm with a smaller value $\bar{\epsilon} < \epsilon$, one must discard from the assignment those pairs that do not satisfy $\bar{\epsilon}$ -CS. An alternative possibility is to use a variant of the auction algorithm that does not require that the initial price and assignment satisfy $\bar{\epsilon}$ -CS. In this variant, we try to execute the cooperative and noncooperative iterations as if $\bar{\epsilon}$ -CS were satisfied, and when assigned pairs (i, j) not satisfying $\bar{\epsilon}$ -CS are encountered, to discard these pairs from the assignment as needed, while making sure that all newly assigned pairs satisfy $\bar{\epsilon}$ -CS. With this somewhat speculative mode of operation, progress can be made towards satisfying $\bar{\epsilon}$ -CS as the algorithm is running, with potentially significant computational savings. An additional advantage of this type of scheme is that it can be operated as an “anytime algorithm,” i.e., an algorithm that progressively improves on a feasible solution, and returns a complete assignment even if it is interrupted because a computational budget limit or other real-time constraint has been reached.

Variations of ϵ -scaling such as the preceding one can also be helpful when favorable initial conditions prices and assignment pairs are known, which, however, do not satisfy ϵ -CS. For example, good initial conditions may be available by using a trained neural network, which accepts the problem data and provides an approximately optimal (complete or partial) assignment, and corresponding prices, which together may not fully satisfy ϵ -CS for a desired value of ϵ .

In a related context, which is very common in practice, assignment problems are solved repeatedly with small variations in the problem's data (such as small changes in the problem's graph or values). Then there is much to be gained by reusing information in the form of prices and assignment pairs, even if they do not satisfy ϵ -CS. As an example, the author's paper [Ber22] has introduced auction algorithms for path construction and shortest path problems, where the initial conditions need not satisfy ϵ -CS, but are progressively rectified in the course of the algorithm. This is particularly convenient in on-line applications where the problem data changes and maintaining ϵ -CS at all times is difficult (a knowledge graph context of this type is considered by Agarwal, Bertsekas, and Liu [ABL23]). The ideas of the paper [Ber22] (and related ideas from an earlier max-flow paper by the author [Ber95a]) can be extended to the algorithms of the present paper for solving assignment problems as well as other network optimization problems.

Adaptive ϵ -Scaling

One possibility to improve the performance of ϵ -scaling schemes is to introduce adaptivity, whereby the value of ϵ is modified in the course of the algorithm, depending on algorithmic progress. In particular, we may start with a small value of ϵ and suitably increase it if some heuristic criterion suggests that a price war is underway (a simple heuristic of this type is implemented in the author's FORTRAN codes noted earlier).

Another possibility is to *use a person-dependent value of ϵ* , so each person has his/her own value that determines the size of his/her ϵ -zone. In particular, if the parameter value ϵ_i is used by person i , we may increase ϵ_i by some factor (up to some upper bound), each time i submits an aggressive auction single-person bid, thereby expanding the ϵ -zone $\mathcal{Z}(i)$. This enhances the cooperative character of iterations that involve repeat bidders, such as the ones participating in a price war. Intuitively, in this form of adaptive ϵ -scaling, *a person i that submits an aggressive bid repeatedly, only to be outbid later by some other person, seeks coalition partners by increasing ϵ_i in order to get through a price war more quickly.*

Reverse Iterations, Similar Persons and Objects, Third Best Profit Test

We mention some additional variants of the auction algorithm, which have been proposed in the literature, and can be adapted to the cooperative framework of this paper and its extensions to other network optimization problems. An important variation involves the use of reverse iterations (see Bertsekas, Castañón, and Tsaknakis [BCT93], and the books [Ber91a], Section 4.2, and [Ber98], Section 7.2). In the (forward) auction iterations that we have described so far, persons compete for objects by bidding and raising the prices of objects. In reverse auction iterations, roughly speaking, *the objects compete for persons by essentially offering discounts.*

We can describe reverse auction in two equivalent ways: one where unassigned objects lower their prices as much as possible to attract an unassigned person or to lure a person away from its currently held object without violating ϵ -CS, and another where unassigned objects select a best person and raise his/her profit

as much as possible without violating ϵ -CS.

Mathematically, reverse auction is equivalent to forward auction with the roles of persons and objects, and the roles of profits and prices interchanged. On the other hand a typically more effective algorithmic scheme is obtained when forward and reverse auction are combined in an algorithm that switches from forward to reverse auction and back at suitable times. Such a combined algorithm simultaneously maintains a partial assignment, a price vector p , and a profit vector π satisfying the following ϵ -CS condition:

$$\pi_i + p_j \geq a_{ij} - \epsilon, \quad \text{for all } i \text{ and } j \in A(i), \quad (4.1)$$

$$\pi_i + p_j = a_{ij}, \quad \text{for all assigned pairs } (i, j). \quad (4.2)$$

When forward iterations are used, the prices of objects are increased according to the rules that we have described, while the profits of persons are subsequently reduced to maintain the conditions (4.1)-(4.2). When reverse iterations are used, the profits of persons are increased according to rules that can be viewed as “reverse” from the rules we have described, while the prices of objects are subsequently reduced to maintain the conditions (4.1)-(4.2).

A simple way to guarantee the validity of such a combined algorithm is to refrain from switching from one type of the auction to the other until the number of assigned person-object pairs increases by at least one. We refer to the books [Ber91a], Section 4.2, and [Ber98], Section 7.2, for related analysis and discussion. In practice, combined forward/reverse auction algorithms are affected less by price wars and often work substantially faster than the forward versions. Price wars can still occur in combined forward/reverse algorithms, but they arise through more complex and unlikely problem structures than in forward algorithms. For this reason combined forward/reverse auction algorithms depend less on ϵ -scaling for good performance than their forward counterparts. In fact, starting with $\epsilon < 1/n$, thus bypassing ϵ -scaling, is sometimes the best choice.

Some variations that are important from both the algorithmic and the theoretical/conceptual point of view deal with problems where there are many “similar” persons and objects [many persons i with identical object sets $A(i)$ and values a_{ij} , $j \in A(i)$]. Problems of this type are particularly susceptible to price wars; see the books [Ber91a], Section 4.2, [Ber98], Chapter 7. The paper by Bertsekas and Castañón [BeC89], and the more recent papers by Walsh and Dieci [WaD17], [WaD19] propose related auction algorithms in the context of transportation problems, which can be converted into assignment problems with many similar persons and objects. Walsh has also written publicly available auction codes for transportation problems; see <https://github.com/jdwalsh03/auction>. Alternatively, transportation problems may be viewed as special cases of linear single commodity network problems, and they can be addressed by corresponding natural extensions of auction algorithms.

Finally, we mention another variation of the aggressive auction iteration, which is based on the “third best” profit test (see Exercise 1.7, Section 4.1 of the book [Ber91a], or Exercise 7.7 of the book [Ber98]).

The motivation here is that frequently in the auction algorithm the two best objects for a given person do not change between two successive bids of that person. The third best test aims to exploit this fact by checking whether the two best objects from the preceding bid continue to be best. If the test is passed, the computation of the profits $a_{ij} - p_j$ of the remaining objects j is unnecessary.

In particular, suppose that at a given aggressive auction iteration that starts with the unassigned person i , we compute the best and second best profits

$$\max_{j \in A(i)} \{a_{ij} - p_j\}, \quad \max_{j \in A(i), j \neq j_1} \{a_{ij} - p_j\},$$

the corresponding best and second best objects

$$j_1 = \arg \max_{j \in A(i)} \{a_{ij} - p_j\}, \quad j_2 = \arg \max_{j \in A(i), j \neq j_1} \{a_{ij} - p_j\},$$

and the *third best profit*

$$y_i = \max_{j \in A(i), j \neq j_1, j \neq j_2} \{a_{ij} - p_j\}.$$

Suppose that at a subsequent iteration when person i bids based on an updated price vector \bar{p} , we have

$$a_{ij_1} - \bar{p}_{j_1} \geq y_i, \quad a_{ij_2} - \bar{p}_{j_2} \geq y_i.$$

Because person i 's profits can only decrease when passing from the price vector p to the price vector \bar{p} , this guarantees that j_1 and j_2 continue to be the two best objects for i (although j_1 may become worse than j_2 , based on the updated price vector \bar{p}). As a result no further computation is needed to execute the aggressive auction iteration. The third best profit test has proved to be quite effective in practice, requires minimal additional overhead, and has been implemented in the author's FORTRAN codes.

Special Choices of Unassigned Persons

All the algorithms that we have discussed so far, except for the ones involving expanding coalitions, leave open the choice of the unassigned person i that initiates the auction iteration. However, problems with special structure may lend themselves to special/favorable choices of i . For example in assignment problems that have a path construction structure, such as shortest path-type or max-flow-type problems, it may be beneficial to choose unassigned persons in a sequence that corresponds to a candidate solution path or candidate augmenting path; see the author's paper [Ber22] for related auction algorithmic ideas.

In the context of the assignment problem, a special choice of this type corresponds to choosing the person i that starts an auction iteration to be one that has just lost his/her assigned object due to an aggressive bid by another person. We will not go into further details, and instead refer to the papers [Ber91b], [Ber95a], [Ber95b], [Ber22], and the books [Ber91a], [Ber98] for discussion of such possibilities and the intuition behind them.

Heuristic Criteria for Switching to Cooperative Auction

An issue that arises in combinations of conservative/aggressive and cooperative auction is how to control the switch from one type of auction to another. One possibility is to forgo the aggressive iteration and do a cooperative iteration instead, if some heuristic criterion suggests that a price war is underway; for example, a relatively large number of aggressive iterations that do not produce an augmentation. This is similar to what is done in two-phase auction algorithms with $\epsilon = 0$, which start as single-person/conservative auction and switch to a cooperative auction if price wars persist, e.g., the algorithms of [Ber81] and [JoV87].

Dealing with Infeasibility

Let us consider the case of an infeasible problem, where there does not exist a complete assignment. In this case, the auction algorithm cannot possibly terminate. It will keep on increasing the prices of some objects by increments of at least ϵ . Furthermore, some persons will be submitting bids infinitely often, and the corresponding profits will be decreasing toward $-\infty$. Methods to detect infeasibility of a given problem have been developed and have been discussed in several of the author's works; see for example [Ber91a], [Ber92], [Ber98]. These methods can be easily incorporated into the algorithmic framework of this paper.

A simple method to deal with infeasibility is to convert the problem to an equivalent feasible problem by adding a set of artificial person-object pairs to the original set of pairs. The values of these pairs should be very small, so that none of them participates in an optimal assignment unless the problem is infeasible. We refer to Section 3.3 of the tutorial paper [Ber92] for further discussion. An alternative possibility is to first check for feasibility of the problem (before attempting to solve it) by using a low complexity bipartite matching algorithm for infeasibility detection.

Finally, let us note that if the expanding coalitions variant is used, the detection of infeasibility is simple: the problem is infeasible if and only if in the course of some cooperative iteration (with coalition expansion) we encounter an empty set of border nodes; this can only happen if there is no complete assignment cf. Prop. 3.1(c).

5. CONCLUDING REMARKS

We have introduced a new cooperative auction iteration, and variations thereof, for symmetric linear assignment problems, which may use a positive value of ϵ , and can resolve competitive impasses and price wars without requiring the use of ϵ -scaling (although it can be used in conjunction with ϵ -scaling). The iteration is recommended when the ϵ -zone of the starting unassigned person consists of multiple objects, all of which are assigned, an indication of the possibility of a price war; otherwise the classical aggressive form of the auction iteration is typically preferable. The variant of the cooperative auction iteration that involves person reassignments actually coincides with the aggressive auction iteration when the ϵ -zone of the starting person

consists of a single object.

The auction iterations described in this paper admit extensions to other classical network optimization problems such as asymmetric assignment, multiassignment, shortest path, k -shortest path, max-flow, and transportation problems. All of these problems can in turn be viewed as special cases of the general single commodity linear network flow problem, which is commonly referred to as the *minimum cost flow problem* (MCNF for short) in the literature.

We plan to discuss extensions of the cooperative auction algorithm and its variants to other network flow problems in future publications. However, it is worth mentioning here some connections between the assignment algorithms of the present paper and algorithms for the MCNF problem, which point the way to future work:

- (a) Conservative auction, when generalized to the MCNF problem, becomes the single node relaxation method described in Section 6.3 of the book [Ber98].
- (b) Aggressive auction, when generalized to the MCNF problem, becomes the ϵ -relaxation method first proposed by the author in the paper [Ber86], and described and analyzed in detail in the books [BeT89] (Sections 5.3, 5.4), [Ber91a] (Section 4.5), and [Ber98] (Section 7.4). This method is also closely related to preflow-push methods, as noted earlier.
- (c) The cooperative auction algorithm with $\epsilon = 0$ and no coalition expansions, when generalized to the MCNF problem, becomes the relaxation method of Bertsekas [Ber85], and Bertsekas and Tseng [BeT88]; see also the books [Ber91a] (Section 3.3) and [Ber98] (Section 6.3).
- (d) The cooperative auction algorithm with $\epsilon = 0$ and coalition expansions, when generalized to the MCNF problem, becomes the classical primal-dual (sequential shortest path) method; see also the books [Ber91a] (Section 3.2) and [Ber98] (Section 6.2).
- (e) The variant of the cooperative auction algorithm that was first presented in Section 4 (person reassignments along an alternating path), when generalized to the MCNF problem with $\epsilon = 0$, becomes a variant of the relaxation method described by the author in the paper [Ber95b] under the name “early flow augmentations.”
- (f) An auction algorithm for the max-flow problem, given by the author in the paper [Ber95a], combines several of the variations of aggressive and cooperative auction algorithms that we have discussed. Of course, the max-flow problem has special structure (such as zero arc costs and hence no need for ϵ -scaling), which can be exploited when specializing the algorithms of the present paper to its context.
- (g) Cooperative auction with $\epsilon > 0$, and its variants with and without coalition expansions and person

reassignments, are new algorithms, which generalize without much difficulty to the MCNF problem and its special cases noted earlier. Early ideas in this regard can be found in the paper by Bertsekas and Castañón [BeC93c], and the book [Ber98], Section 9.6.

Another form of extension to a MCNF problem that involves a convex (rather than linear) separable cost function, is also possible. It can be based on related problem transformation ideas (see the papers by Bertsekas, Polymenakos, and Tseng [BPT97], [BPT98], and the textbook [Ber98], Chapter 9).

A basic mechanism for extension of auction algorithms to MCNF problems and special cases thereof is to first convert such problems to assignment problems, by using well known transformations, then apply one of the algorithms of the present paper, and then streamline the computations for efficiency. However, as a practical matter one should not try to literally convert one of the assignment algorithms of the present paper to a new problem structure. Instead one should aim to combine and adapt the principal algorithmic ideas presented in this paper, in sensible ways that experimentally can be shown to work well for the given type of problem. These ideas are conservative, aggressive, and cooperative price rises and augmentations, under the umbrella of the mathematically fundamental approximation framework of ϵ -CS, and the intuitive framework of auction-based economic competition.

Let us also mention extensions of the cooperative auction iteration (possibly in combination with aggressive auction iterations) that may involve multiple unassigned persons. These persons may submit bids in parallel or distributed, possibly asynchronous, fashion. Extensions of this type are not considered in the present paper. We refer to the book [BeT89], Sections 5.3 and 6.5, for related discussions of distributed asynchronous aggressive auction algorithms, and also the papers by Bertsekas and Castañón [BeC93a], [BeC93b] for distributed asynchronous implementations of the Hungarian method and primal-dual methods.

We note that beyond their use in addressing the MCNF problem, our algorithmic ideas lend themselves well for incorporation in heuristics for assignment-like problems, which are more difficult than the linear assignment problem that we have considered in this paper. Such problems include multi-dimensional assignment, combinatorial auctions, dynamic task allocation, and multiagent/multi-robot problems. A noteworthy approach is to use an auction algorithm as a base heuristic for a rollout algorithm; see the books [Ber98] (Section 10.5), [Ber20a] (Section 3.4.2), and [Ber23b] (Chapter 2).

Finally, let us mention an interesting connection with reinforcement learning. One of the important favorable characteristics of auction algorithms is that the final prices obtained from solution of a given assignment problem can be used as initial prices for applying the algorithms to other problems, which are structurally similar. This suggests that one may try to “learn” favorable initial prices from data, which encode this knowledge into a neural network that can supply on demand good initial prices for a given problem and (possibly) an associated favorable partial assignment. Work on machine learning and neural network approaches towards assignment problems is at a very early stage at present; see e.g., Lee et al.

[LXY18], Emami et al. [EPE20], Aironi, Cornell, and Squartini [ACS22], and Liu et al. [LWL22]. It is reasonable to expect that auction algorithms and their intuitive economic competition-like mechanism lend themselves well to this line of research.

6. REFERENCES

- [ABL23] Agrawal, G., Bertsekas, D., Liu, H., 2023. “Auction-Based Learning for Question Answering over Knowledge Graphs,” *Information*, Vol. 14, 336, <https://doi.org/10.3390/info14060336>.
- [ACS22] Aironi, C., Cornell, S., and Squartini, S., 2022. “Tackling the Linear Sum Assignment Problem with Graph Neural Networks,” in *International Conference on Applied Intelligence and Informatics*, Springer, pp. 90-101.
- [AMO88] Ahuja, R. K., Magnanti, T. L., and Orlin, J. B., 1988. *Network Flows*, dspace.mit.edu.
- [AMO89] Ahuja, R. K., Magnanti, T. L., and Orlin, J. B., 1989. “Network Flows,” in *Handbooks in Operations Research and Management Science*, Vol. 1, Optimization, Nemhauser, G. L., et al. (eds.), North-Holland, Amsterdam, pp. 211-369.
- [APP22] Aziz, H., Pal, A., Pourmiri, A., Ramezani, F., and Sims, B., 2022. “Task Allocation Using a Team of Robots,” *Current Robotics Reports*, Vol. 3, pp. 227-238.
- [APV22] Alfaro, C. A., Perez, S. L., Valencia, C. E., and Vargas, M. C., 2022. “The Assignment Problem Revisited,” *Optimization Letters*, Vol. 16, pp. 1531-1548.
- [AhO89] Ahuja, R. K., and Orlin, J. B., 1989. “A Fast and Simple Algorithm for the Maximum Flow Problem,” *Operations Research*, Vol. 37, pp. 748-759.
- [Ami94] Amini, M. M., 1994. “Vectorization of an Auction Algorithm for Linear Cost Assignment Problem,” *Comput. Ind. Eng.*, Vol. 26, pp. 141-149.
- [BCE95] Bertsekas, D. P., Castañón, D. A., Eckstein, J., and Zenios, S., 1995. “Parallel Computing in Network Optimization,” *Handbooks in OR and MS*, Ball, M. O., Magnanti, T. L., Monma, C. L., and Nemhauser, G. L. (eds.), Vol. 7, North-Holland, Amsterdam, pp. 331-399.
- [BCT93] Bertsekas, D. P., Castañón, D. A., and Tsaknakis, H., 1993. “Reverse Auction and the Solution of Inequality Constrained Assignment Problems,” *SIAM J. on Optimization*, Vol. 3, pp. 268-299.
- [BDM12] Burkard, R., Dell’Amico, M., and Martello, S., 2012. *Assignment Problems*, Society for Industrial and Applied Mathematics.
- [BFH03] Brenier, Y., Frisch, U., Henon, M., Loeper, G., Matarrese, S., Mohayaee, R., and Sobolevskii, A., 2003. “Reconstruction of the Early Universe as a Convex Optimization Problem,” *Monthly Notices of the Royal Astronomical Society*, Vol. 346, pp. 501-524.
- [BGM97] Beraldi, P., Guerriero, F., and Musmanno, R., 1997. “Efficient Parallel Algorithms for the Minimum

- Cost Flow Problem,” *Journal of Optimization Theory and Applications*, Vol. 95, pp. 501-530.
- [BGM01] Beraldi, P., Guerriero, F., and Musmanno, R., 2001. “Parallel Algorithms for Solving the Convex Minimum Cost Flow Problem,” *Computational Optimization and Applications*, Vol. 18, pp. 175-190.
- [BPS07] Bayati, M., Prabhakar, B., Shah, D., and Sharma, M., 2007. “Iterative Scheduling Algorithms,” in *IEEE INFOCOM 2007-26th IEEE International Conf. on Computer Communications*, pp. 445-453.
- [BPT97] Bertsekas, D. P., Polymenakos, L. C., and Tseng, P., 1997. “An Epsilon-Relaxation Method for Convex Network Optimization Problems,” *SIAM J. on Optimization*, Vol. 7, pp. 853-870.
- [BPT98] Bertsekas, D. P., Polymenakos, L. C., and Tseng, P., 1998. “Epsilon-Relaxation and Auction Methods for Separable Convex Cost Network Flow Problems,” in *Network Optimization*, by P. M. Pardalos, D. W. Hearn, and W. W. Hager (eds.), *Lecture Notes in Economics and Mathematical Systems*, Springer-Verlag, N.Y., pp. 103-126.
- [BSS08] Bayati, M., Shah, D., and Sharma, M., 2008. “Max-Product for Maximum Weight Matching: Convergence, Correctness, and LP Duality,” *IEEE Trans. on Information Theory*, Vol. 54, pp. 1241-1251.
- [BaF88] Bar-Shalom, Y., and Fortman, T. E., 1988. *Tracking and Data Association*, Academic Press, N. Y.
- [Bar90] Bar-Shalom, Y., 1990. *Multitarget-Multisensor Tracking: Advanced Applications*, Artech House, Norwood, MA.
- [BeC89] Bertsekas, D. P., and Castañón, D. A., 1989. “The Auction Algorithm for the Transportation Problem,” *Annals of Operations Research*, Vol. 20, pp. 67-96.
- [BeC91] Bertsekas, D. P., and Castañón, D. A., 1991. “Parallel Synchronous and Asynchronous Implementations of the Auction Algorithm,” *Parallel Computing*, Vol. 17, pp. 707-732.
- [BeC93a] Bertsekas, D. P., and Castañón, D. A., 1993. “Parallel Asynchronous Hungarian Methods for the Assignment Problem,” *ORSA J. on Computing*, Vol. 5, pp. 261-274.
- [BeC93b] Bertsekas, D. P., and Castañón, D. A., 1993. “Parallel Primal-Dual Methods for the Minimum Cost Flow Problem,” *Computational Optimization and Applications*, Vol. 2, pp. 317-336.
- [BeC93c] Bertsekas, D. P., and Castañón, D. A., 1993. “A Generic Auction Algorithm for the Minimum Cost Network Flow Problem,” *Computational Optimization and Applications*, Vol. 2, pp. 229-260.
- [BeE87] Bertsekas, D. P., and Eckstein, J., 1987. “Distributed Asynchronous Relaxation Methods for Linear Network Flow Problems,” *IFAC Proceedings*, Vol. 20, pp. 103-114.
- [BeE88] Bertsekas, D. P., and Eckstein, J., 1988. “Dual Coordinate Step Methods for Linear Network Flow

Problems,” *Math. Programming, Series B*, Vol. 42, pp. 203-243.

[BeG97] Beraldi, P., and Guerriero, F., 1997. “A Parallel Asynchronous Implementation of the ϵ -Relaxation Method for the Linear Minimum Cost Flow Problem,” *Parallel Computing*, Vol. 23, pp. 1021-1044.

[BeM19] Bertozzi, A. L., and Merkurjev, E., 2019. “Graph-Based Optimization Approaches for Machine Learning, Uncertainty Quantification and Networks,” in *Handbook of Numerical Analysis*, Vol. 20, pp 503-531.

[BeT88] Bertsekas, D. P., and Tseng, P., 1988. “Relaxation Methods for Minimum Cost Ordinary and Generalized Network Flow Problems,” *Operations Research*, Vol. 36, pp. 93-114.

[BeT89] Bertsekas, D. P., and Tsitsiklis, J. N., 1989. *Parallel and Distributed Computation: Numerical Methods*, Prentice-Hall, Engl. Cliffs, N. J. (can be downloaded from the author’s website).

[BeT94] Bertsekas, D. P., and Tseng, P., 1994. “RELAX-IV: A Faster Version of the RELAX Code for Solving Minimum Cost Flow Problems,” Report LIDS-P-2276, MIT.

[BeT97] Bertsimas, D., and Tsitsiklis, J. N., 1997. *Introduction to Linear Optimization*, Athena Scientific, Belmont, MA.

[Ber79] Bertsekas, D. P., 1979. “A Distributed Algorithm for the Assignment Problem,” Lab. for Information and Decision Systems Report, MIT, May 1979.

[Ber81] Bertsekas, D. P., 1981. “A New Algorithm for the Assignment Problem,” *Math. Programming*, Vol. 21, pp. 152-171.

[Ber85] Bertsekas, D. P., 1985. “A Unified Framework for Minimum Cost Network Flow Problems,” *Math. Programming*, Vol. 32, pp. 125-145.

[Ber86] Bertsekas, D. P., 1986. “Distributed Asynchronous Relaxation Methods for Linear Network Flow Problems,” 1986 25th IEEE Conference on Decision and Control, pp. 2101-2106.

[Ber88] Bertsekas, D. P., 1988. “The Auction Algorithm: A Distributed Relaxation Method for the Assignment Problem,” *Annals of Operations Research*, Vol. 14, pp. 105-123.

[Ber90] Bertsekas, D. P., 1990. “The Auction Algorithm for Assignment and Other Network Flow Problems: A Tutorial,” *Interfaces*, Vol. 20, pp. 133-149.

[Ber91a] Bertsekas, D. P., 1991. *Linear Network Optimization*, MIT Press, Cambridge, MA.

[Ber91b] Bertsekas, D. P., 1991. “An Auction Algorithm for Shortest Paths,” *SIAM J. on Optimization*, Vol. 1, pp. 425-447.

- [Ber92] Bertsekas, D. P., 1992. “Auction Algorithms for Network Flow Problems: A Tutorial Introduction,” *Computational Optimization and Applications*, Vol. 1, pp. 7-66.
- [Ber93] Bertsekas, D. P., 1993. “Mathematical Equivalence of the Auction Algorithm for Assignment and the Epsilon-Relaxation (Preflow-Push) Method for Min Cost Flow,” in *Large Scale Optimization*, Hager W. W., Hearn D. W., Pardalos P. M. (eds), Springer, Boston, MA.
- [Ber95a] Bertsekas, D. P., 1995. “An Auction Algorithm for the Max-Flow Problem,” *J. of Optimization Theory and Applications*, Vol. 87, pp. 69-101.
- [Ber95b] Bertsekas, D. P., 1995. “An Auction/Sequential Shortest Path Algorithm for the Minimum Cost Network Flow Problem,” Report LIDS-P-2146, MIT.
- [Ber98] Bertsekas, D. P., 1998. *Network Optimization: Continuous and Discrete Models*, Athena Scientific, Belmont, MA (can be downloaded from the author’s website).
- [Ber20a] Bertsekas, D. P., 2020. *Rollout, Policy Iteration, and Distributed Reinforcement Learning*, Athena Scientific, Belmont, MA.
- [Ber20b] Bertsekas, D. P., 2020. “Constrained Multiagent Rollout and Multidimensional Assignment with the Auction Algorithm,” arXiv preprint, arXiv:2002.07407.
- [Ber22] Bertsekas, D. P., 2022. “Auction Algorithms for Path Planning, Network Transport, and Reinforcement Learning,” Arizona State University/SCAI Report; arXiv:22207.09588.
- [Ber23a] Bernard, F., 2023. “Fast Linear Assignment Problem using Auction Algorithm,” MATLAB Central File Exchange,” <https://www.mathworks.com/matlabcentral/fileexchange/48448-fast-linear-assignment-problem-using-auction-algorithm-mex>.
- [Ber23b] Bertsekas, D. P., 2023. *A Course in Reinforcement Learning*, Athena Scientific, Belmont, MA.
- [Bla86] Blackman, S. S., 1986. *Multi-Target Tracking with Radar Applications*, Artech House, Dehdam, MA.
- [BuC99] Burkard, R. E., and Cela, E., 1999. “Linear Assignment Problems and Extensions,” in *Handbook of Combinatorial Optimization*: Springer, Boston, Supplement Vol. A, pp. 75-149.
- [BuT09] Bus, L., and Tvrdk, P., 2009. “Towards Auction Algorithms for Large Dense Assignment Problems,” *Computational Optimization and Applications*, Vol. 43, pp. 411-436.
- [CBH09] Choi, H. L., Brunet, L., and How, J. P., 2009. “Consensus-Based Decentralized Auctions for Robust Task Allocation,” *IEEE Transactions on Robotics*, Vol. 25, pp. 912-926.
- [Cas92] Castañón, D. A., 1992. “New Assignment Algorithms for Data Association,” in *Signal and Data*

Processing of Small Targets, Vol. 1698, pp. 313-323.

[Cas93] Castañón, D. A., 1993. “Reverse Auction Algorithms for Assignment Problems,” in Algorithms for Network Flows and Matching,” Johnson, D. S., and McGeoch, C. C., (eds.), American Math. Soc., Providence, RI, pp. 407-429.

[CeZ97] Censor, Y., and Zenios, S. A., 1997. Parallel Optimization: Theory, Algorithms, and Applications, Oxford University Press.

[ChM89] Cheriyan, J., and Maheshwari, S. N., 1989. “Analysis of Preflow Push Algorithms for Maximum Network Flow,” SIAM J. Computing, Vol. 18, pp. 1057-1086.

[DLT19] Duan, X., Liu, H., Tang, H., Cai, Q., Zhang, F., and Han, X., 2019. “A Novel Hybrid Auction Algorithm for Multi-UAVs Dynamic Task Assignment,” IEEE Access, 8, pp. 86207-86222.

[DeV03] De Vries, S., and Vohra, R. V., 2003. “Combinatorial Auctions: A Survey,” INFORMS Journal on Computing, Vol. 15, pp. 284-309.

[EPE20] Emami, P., Pardalos, P. M., Elefteriadou, L., and Ranka, S., 2020. “Machine Learning Methods for Data Association in Multi-Object Tracking,” ACM Computing Surveys (CSUR), Vol. 53, pp. 1-34.

[Erc15] Erciyes, K., 2015. Distributed and Sequential Algorithms for Bioinformatics, Springer.

[FrS06] Frisch, U., and Sobolevskii, A., 2006. “Application of Optimal Transport Theory to Reconstruction of the Early Universe,” Journal of Mathematical Sciences, Vol. 133, pp. 1539-1542.

[GBG23] Garces, D., Bhattacharya, S., Gil, S., and Bertsekas, D. P., 2023. “Multiagent Reinforcement Learning for Autonomous Routing and Pickup Problem with Adaptation to Variable Demand,” 2023 IEEE International Conference on Robotics and Automation (ICRA), pp. 3524-3531.

[Gal16] Galichon, A., 2016. Optimal Transport Methods in Economics, Princeton University Press.

[GoT86] Goldberg, A. V., and Tarjan, R. E., 1986. “A New Approach to the Maximum Flow Problem,” Proc. 18th ACM STOC, pp. 136-146.

[GoT90] Goldberg, A. V., and Tarjan, R. E., 1990. “Solving Minimum Cost Flow Problems by Successive Approximation,” Math. of Operations Research, Vol. 15, pp. 430-466.

[HZX19] Huang, Y., Zhang, Y., and Xiao, H., 2019. “Multi-Robot System Task Allocation Mechanism for Smart Factory,” in 2019 IEEE 8th Joint International Information Technology and Artificial Intelligence Conference (ITAIC), pp. 587-591.

[JME18] Jacobs, M., Merkurjev, E., and Esedoglu, S., 2018. “Auction Dynamics: A Volume Constrained

- MBO Scheme,” *Journal of Computational Physics*, Vol. 354, pp. 288-310.
- [JoV87] Jonker, R., and Volgenant, T., 1987. “A Shortest Augmenting Path Algorithm for Dense and Sparse Linear Assignment Problems,” *Computing*, Vol. 38.
- [KSS14] Kollias, G., Sathe, M., Schenk, O., and Grama, A., 2014. “Fast Parallel Algorithms for Graph Similarity and Matching,” *Journal of Parallel and Distributed Computing*, Vol. 74, pp. 2400-2410.
- [Kar74] Karzanov, A. V., 1974. “Determining the Maximal Flow in a Network with the Method of Preflows,” *Soviet Math Dokl.*, Vol. 15, pp. 1277-1280.
- [KhA21] Khosla, M., and Anand, A., 2021. “Revisiting the Auction Algorithm for Weighted Bipartite Perfect Matchings,” *arXiv preprint arXiv:2101.07155*.
- [KoY94] Kosowsky, J. J., and Yuille, A. L., 1994. “The Invisible Hand Algorithm: Solving the Assignment Problem with Statistical Physics,” *Neural Networks*, Vol. 7, pp. 477-490.
- [Kuh55] Kuhn, H. W., 1955. “The Hungarian Method for the Assignment Problem,” *Naval Research Logistics Quarterly*, Vol. 2, pp. 83-97.
- [LCS14] Luo, L., Chakraborty, N., and Sycara, K., 2014. “Provably-Good Distributed Algorithm for Constrained Multi-Robot Task Assignment for Grouped Tasks,” *IEEE Transactions on Robotics*, Vol. 31, pp. 19-30.
- [LGO20] Lujak, M., Giordani, S., Omicini, A., and Ossowski, S., 2020. “Decentralizing Coordination in Open Vehicle Fleets for Scalable and Dynamic Task Allocation,” *Complexity*, pp. 1-21.
- [LZC24] Li, H., Zhu, D., Chen, M., Wang, T., and Zhu, H., 2024. “An Improved Reeds?Shepp and Distributed Auction Algorithm for Task Allocation in Multi-AUV System with Both Specific Positional and Directional Requirements,” *J. of Marine Science and Engineering*, Vol. 12, p. 486.
- [LMH21] Levy, B., Mohayaee, R., and von Hausegger, S., 2021. “A Fast Semidiscrete Optimal Transport Algorithm for a Unique Reconstruction of the Early Universe,” *Monthly Notices of the Royal Astronomical Society*, Vol. 506, pp. 1165-1185.
- [LWL22] Liu, H., Wang, T., Lang, C., Feng, S., Jin, Y., and Li, Y., 2022. “GLAN: A Graph-Based Linear Assignment Network,” *arXiv preprint arXiv:2201.02057*.
- [LXY18] Lee, M., Xiong, Y., Yu, G., Li, G. Y., 2018. “Deep Neural Networks for Linear Sum Assignment Problems,” *IEEE Wirel. Commun. Lett.*, Vol. 7, pp. 962-965.
- [LZX23] Li, H., Zhu, H., Xu, D., Lin, X., Jiao, G., Song, Y., and Huang, M., 2023. “Dynamic Task Al-

location Based on Auction in Robotic Mobile Fulfilment System,” *Journal of Industrial and Management Optimization*, Vol. 19.

[Lav08] Lavaux, G., 2008. “Lagrangian Reconstruction of Cosmic Velocity Fields,” *Physica D: Nonlinear Phenomena*, Vol. 237, pp. 2139-2144.

[LiS13] Liu, L., and Shell, D. A., 2013. “Optimal Market-Based Multi-Robot Task Allocation via Strategic Pricing,” *Robotics: Science and Systems*, Vol. 9, pp. 33-40.

[LuM20] Lujak, M., and Matezovic, M., 2020. “On Efficiency in Dynamic Multi-Robot Task Allocation,” in *AIRO Proceedings*, pp. 49-53.

[Mal97] Malkoff, D. B., 1997. “Evaluation of the Jonker-Volgenant-Castañon (JVC) Assignment Algorithm for Track Association,” *Proc. SPIE 3068, Signal Processing, Sensor Fusion, and Target Recognition*; <https://doi.org/10.1117/12.280801>.

[MBM19] Metivier, L., Brossier, R., Merigot, Q., and Oudet, E., 2019. “Graph Space Optimal Transport for FWI: Auction Algorithm, Application to the 2d Valhall Case Study,” in *81st EAGE Conference and Exhibition, European Association of Geoscientists and Engineers*.

[MSC16] Morgan, D., Subramanian, G. P., Chung, S. J., and Hadaegh, F. Y., 2016. “Swarm Assignment and Trajectory Optimization Using Variable-Swarm, Distributed Auction Assignment and Sequential Convex Programming,” *The International Journal of Robotics Research*, Vo. 35, pp. 1261-1285.

[MeT21] Merigot, Q., and Thibert, B., 2021. “Optimal Transport: Discretization and Algorithms,” in *Handbook of Numerical Analysis*, Elsevier, Vol. 22, pp. 133-212.

[Mer20] Merkurjev, E., 2020. “A Fast Graph-Based Data Classification Method with Applications to 3D Sensory Data in the Form of Point Clouds,” *Pattern Recognition Letters*, Vol. 136, pp.154-160.

[NSB23] Nurlanov, Z., Schmidt, F. R., and Bernard, F., 2023. “Universe Points Representation Learning for Partial Multi-Graph Matching,” in *Proc. of the AAAI Conference on Artificial Intelligence*, Vol. 37, pp. 1984-1992.

[NVJ16] Nascimento, A. D. P., Vasconcelos, C. N., Jamel, F. S., and Sena, A. C., 2016. “A Hybrid Parallel Algorithm for the Auction Algorithm in Multicore Systems,” in *2016 International Symposium on Computer Architecture and High Performance Computing Workshops (SBAC-PADW)*, pp. 73-78.

[NaL16] Naparstek, O., and Leshem, A., 2016. “Expected Time Complexity of the Auction Algorithm and the Push Relabel Algorithm for Maximum Bipartite Matching on Random Graphs,” *Random Structures and Algorithms*, Vol. 48, pp. 384-395.

- [OKS20] Otte, M., Kuhlman, M. J., and Sofge, D., 2020. “Auctions for Multi-Robot Task Allocation in Communication Limited Environments,” *Autonomous Robots*, Vol. 44, pp. 547-584.
- [OrA92] Orlin, J. B., and Ahuja, R. K., 1992. “New Scaling Algorithms for the Assignment and Minimum Mean Cycle Problems,” *Math. Programming*, Vol. 54, pp. 41-56.
- [PDB92] Pattipati, K. R., Deb, S., Bar-Shalom, Y., and Washburn, R. B., 1992. “A New Relaxation Algorithm and Passive Sensor Data Association,” *IEEE Trans. Automatic Control*, Vol. 37, pp. 198-213.
- [PPB01] Popp, R. L., Pattipati, K. R., and Bar-Shalom, Y., 2001. “ m -Best SD Assignment Algorithm with Application to Multitarget Tracking,” *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 37, pp. 22-39.
- [PaU99] Parkes, D. C., and Ungar, L. H., 2000. “Iterative Combinatorial Auctions: Theory and Practice,” *Aaai/iaai*, 7481, p. 53.
- [PeC19] Peyre, G., and Cuturi, M., 2019. “Computational Optimal Transport: With Applications to Data Science,” *Foundations and Trends in Machine Learning*,” Vol. 11, pp. 355-607.
- [PoR97] Poore, A. B., and Robertson, A. J. A., 1997. “New Lagrangian Relaxation Based Algorithm for a Class of Multidimensional Assignment Problems,” *Computational Optimization and Applications*, Vol. 8, pp. 129-150.
- [Poo94] Poore, A. B., 1994. “Multidimensional Assignment Formulation of Data Association Problems Arising from Multitarget Tracking and Multisensor Data Fusion,” *Computational Optimization and Applications*, Vol. 3, pp. 27-57.
- [SSB12] Sathe, M., Schenk, O., and Burkhart, H., 2012. “An Auction-Based Weighted Matching Implementation on Massively Parallel Architectures,” *Parallel Computing*, Vol. 38, pp. 595-614.
- [SSN21] Sena, A. C., Silva, M. N., and Nascimento, A. P., 2021. “An Efficient Vectorized Auction Algorithm for Many-Core and Multicore Architectures,” in *Latin American High Performance Computing Conference*, Springer, pp. 76-90.
- [San15] Santambrogio, F., 2015. *Optimal Transport for Applied Mathematicians*, Springer Intern. Publ.
- [Sch16] Schmitzer, B., 2016. “A Sparse Multiscale Algorithm for Dense Optimal Transport,” *J. of Mathematical Imaging and Vision*, Vol. 56, pp. 238-259.
- [Sch19] Schmitzer, B., 2019. “Stabilized Sparse Scaling Algorithms for Entropy Regularized Transport Problems,” *SIAM Journal on Scientific Computing*, Vol. 41, pp. A1443-A1481.

- [ShV82] Shiloach, Y., and Vishkin, U., 1982. “An $O(n^2 \log n)$ Parallel Max-Flow Algorithm,” *J. Algorithms*, Vol. 3, pp. 128-146.
- [TZG18] Tang, J., Zhu, K., Guo, H., Gong, C., Liao, C., and Zhang, S., 2018. “Using Auction-Based Task Allocation Scheme for Simulation Optimization of Search and Rescue in Disaster Relief,” *Simulation Modelling Practice and Theory*, Vol. 82, pp. 132-146.
- [TsB00] Tseng, P., and Bertsekas, D. P., 2000. “An ϵ -Relaxation Method for Separable Convex Cost Generalized Network Flow Problems,” *Mathematical Programming*, Vol. 88, pp. 85-104.
- [Vil09] Villani, C., 2009. *Optimal Transport: Old and New*, Springer, Berlin.
- [Vil21] Villani, C., 2021. *Topics in Optimal Transportation*, American Mathematical Society.
- [WLY23] Wang, Y., Li, H. and Yao, Y., 2023. “An Adaptive Distributed Auction Algorithm and its Application to Multi-AUV Task Assignment,” *Science China Technological Sciences*, pp. 1-10.
- [WMW22] Wang, C., Mei, D., Wang, Y., Yu, X., Sun, W., Wang, D., and Chen, J., 2022. “Task Allocation for Multi-AUV System: A Review,” *Ocean Engineering*, Vol. 266, p. 112911.
- [WaD17] Walsh, J. D., and Dieci, L., 2017. “General Auction Method for Real-Valued Optimal Transport,” arXiv preprint arXiv:1705.06379.
- [WaD19] Walsh, J. D., and Dieci, L., 2019. “A Real-Valued Auction Algorithm for Optimal Transport,” *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 12(6), pp. 514-533.
- [WeZ91] Wein, J., and Zenios, S. A., 1991. “On the Massively Parallel Solution of the Assignment Problem,” *J. of Parallel and Distributed Computing*, Vol. 13, pp. 228-236.
- [ZCH20] Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C., and Sun, M., 2020. “Graph Neural Networks: A Review of Methods and Applications,” *AI Open*, Vol. 1, pp. 57-81.
- [ZSP08] Zavlanos, M. M., Spesivtsev, L., and Pappas, G. J., 2008. “A Distributed Auction Algorithm for the Assignment Problem,” *Proc. 47th IEEE Conference on Decision and Control*, pp. 1212-1217.
- [Zak95] Zaki, H. A., 1995. “A Comparison of Two Algorithms for the Assignment Problem,” *Computational Optimization and Applications*, Vol. 4, pp. 23-45.