

Parallel Auction Algorithm

Haolin Guo (hg2691), Yuanqing Lei (yl5457), Ava Penn (ap4315)

Dec. 16, 2024

1 Introduction

In this project, we implemented the sequential and parallel versions of an auction algorithm in Haskell. The auction algorithm is an optimization technique used for solving linear assignment problems, where the goal is to match agents to tasks in a way that minimizes or maximizes a total cost. The graph implementation of this problem is given a bipartite graph $G = (V, E)$ with bipartition (A, B) and weight function $w : E \rightarrow \mathbb{R}$, find a matching of maximum weight, where the weight of a matching M is given by $w(M) = \sum_{e \in M} w(e)$.

The sequential implementation of this algorithm is inspired by economic principles where agents bid for items (similar to a second-price auction), leading to iterative improvements in the set of prices until the total payoff is maximized. We chose to focus on this algorithm because it becomes computationally infeasible at a large number of bidders and items, and because steps 2 and 3 as indicated below are suitable for running in parallel—that is, they are mostly done independent of other tasks.

We focus on two approaches: the Jacobi implementation and the Gauss-Seidel implementation, and compare their runtime efficiencies. These approaches are adapted from Jin [1], which implements a similar algorithm in C.

2 The Assignment Problem

Consider the following example. There are three bidders (B_1, B_2, B_3) and three items (I_1, I_2, I_3). The payoffs of assigning each item to each bidders are represented in the following payoff matrix:

$$\begin{bmatrix} 2 & 4 & 3 \\ 3 & 2 & 1 \\ 3 & 4 & 10 \end{bmatrix}$$

Here, the entry in row i and column j (e.g., 4 in the top left) represents the payoff of assigning item I_i to bidder J_j . In the context of auctions, payoff is similar to utility, or the value that the bidder assigns to a specific item (how much they are willing to pay).

The goal is to find an assignment where each item is assigned to exactly one unique bidder, such that the total payoff is maximized.

To maximize the total payoff in the example above:

- Assign I_1 to B_2 (payoff 4),
- Assign I_2 to B_1 (payoff 3),
- Assign I_3 to B_3 (payoff 10).

The total payoff of this assignment is:

$$4 + 3 + 10 = 17$$

This optimal solution can be obtained using algorithms like the auction algorithm. This algorithmic approach has applications in many types of allocation/linear assignment problems.

3 The Auction Algorithm

3.1 Brute-force Implementation

Using a brute-force sequential approach involves generating all possible permutations of assignments and calculating the total pay-off for each permutation to identify the maximum. This approach is exponential in the size of the input matrix, and is thus intractable. However, we used this brute-force approach to test the correctness of our sequential implementation on small matrices. What follows is our implementation in Haskell.

```
1 optimalAssignment :: PayoffMatrix -> Assignment
2 optimalAssignment matrix = maximumBy (comparing totalPayoff) assignments
3   where
4     bidders = [0 .. length matrix - 1]
5     items = bidders -- assume square matrix
6     assignments = [Map.fromList (zip items perm) | perm <- permutations bidders]
7     totalPayoff assignment = sum [matrix !! b !! i | (i,b) <- Map.toList assignment
   ]
```

3.2 Algorithm

The auction algorithm is taken from Jin [1]. It is pseudo-polynomial in that it also depends on the largest element of the payoff matrix. The worst-case performance is $O(n^3)$ or $O(C \cdot n^2)$, but on average, it is expected to perform in $O(n^2 \log n)$. The $O(n^3)$ Hungarian algorithm is more difficult to implement in parallel, however, and it has been found in practice that the auction algorithm often outperforms the Hungarian algorithm.

1. Start with a set U of all bidders. U denotes the set of all unassigned bidders. Initialize a set of prices to zero and any structure that stores the current tentative (partial) assignment.

```
1 initialUnassigned = [0 .. numBidders - 1]
2 initialPrices = Map.fromList [(j, 0) | j <- [0 .. numItems - 1]]
```

2. Pick any bidder i from U . Search for the item j that gives the highest net payoff $A_{ij} - p_j$, and also an item k that gives the second highest net payoff.

```
1 -- calculate net payoffs for all items
2 netPayoffs = [(j, netPayoff i j prices) | j <- [0 .. numItems - 1]]
3
4 -- find the best and second-best items
5 (bestItem, maxPayoff) = maximumBy (comparing snd) netPayoffs
6 secondMaxPayoff = if length netPayoffs > 1
7                   then maximum [ p | (j,p) <- netPayoffs, j /= bestItem ]
8                   else maxPayoff - epsilon
```

3. Update the price p_j of item j as:

$$p_j \leftarrow p_j + (A_{ij} - p_j - (A_{ik} - p_k)). \quad (1)$$

This update ensures that the updated prices satisfy $A_{ij} - p_j = A_{ik} - p_k$ (it makes it so that the bidder is indifferent to buying the two items).

```
1 -- update the price of the best item
2 newPrice = (prices Map.! bestItem) + (maxPayoff - secondMaxPayoff + epsilon)
3 updatedPrices = Map.insert bestItem newPrice prices
```

- Assign item j to bidder i . If item j was previously assigned to another bidder s , remove that assignment and add s back to U .

```

1 -- handle previous assignment of the item
2 (newAssignment, remainingUnassigned) =
3   case Map.lookup bestItem assignment of
4     Just prevBidder ->
5       -- since bestItem was assigned to prevBidder, remove that assignment and
6         add prevBidder back into U
7       let updatedAssignment = Map.insert bestItem i assignment -- reassign item
8         to current bidder i
9         updatedUnassigned = prevBidder : unassignedBidders
10      in (updatedAssignment, updatedUnassigned)
11     Nothing ->
12      (Map.insert bestItem i assignment, unassignedBidders)
13 in go remainingUnassigned updatedPrices newAssignment

```

- If U becomes empty, the algorithm terminates; otherwise, return to Step 2.

4 Parallelization: Gauss-Seidel

The Gauss-Seidel version focuses on parallelizing step 2 of the auction algorithm, where each bidder searches for the best and second-best items to bid on. This parallelization divides the items among p threads, allowing each thread to search its partition independently.

This seems like the most obvious and intuitive way to implement parallelization, since the bid on an item i does not affect the bid on an item j . However, there is a loss of efficiency to overhead for the following reasons:

- Synchronization Costs:** After the individual searches, the results must be merged to determine the overall best and second best items.
- Load Imbalance:** If the item partitions are not evenly distributed, or if the complexity varies due to the variation in item values, some threads may finish earlier, leaving others idle.

4.1 Parallelization Choices

We implemented the algorithm in parallel using Haskell's `Control.Parallel.Strategies` library. Specifically, the `parMap` function is used to divide the workload across multiple threads, with each thread independently processing a partition of items.

- `rpar` sparks the evaluations in parallel. `parMap` starts the evaluation of each chunk in the list in parallel.
- `chunkItems` function: Items are split into p chunks, where each chunk is processed by a separate thread. This distributes independent computations of selecting the best bidder per item.

Here is a snippet of code that uses parallelization. Please see the appendix for the entirety of the code:

```

1 -- parallelize the search for best and second-best items
2 partitions = chunkItems 1600 netPayoffs -- change this number iteratively to find
3   the best size chunk
4 -- for chunks: tested 2, 4, 8, 20, 100, 400, 1600, 6400, 10000, 20000
5 partialResults = parMap rpar findBestAndSecond partitions
6 (bestItem, maxPayoff, secondMaxPayoff) = mergeResults partialResults epsilon

```

We tested the code with varying numbers for p , seeing which number of chunks resulted in the best sparks output (i.e. wanting to keep garbage collected and fizzled sparks low). We ended up choosing 1600 chunks because we thought it offered a good balance between the number of threads and the number of work done on each thread, like the painters on the wall analogy. Also observe that the total number of sparks doesn't increase after 1600—because of the problem size. Since we tested this first with matrices of 1000x1000 and didn't observe a great speedup, we didn't try to optimize on larger problems and thus didn't end up needing to adjust this variable for larger problems. Here are the results of testing p chunks on four cores:

Number of Threads	Total Sparks	Converted	Overflowed	GC'd	dud	Fizzled
2	48305	5879	0	0	12468	29958
4	96611	15163	0	0	28712	52736
8	193203	42612	0	0	41620	108971
20	483003	100362	0	0	164934	217707
100	2415022	611727	0	0	1103680	699615
400	9660052	4538608	0	0	4081373	1040071
1600	24150254	24014601	0	0	2796	132857
6400	24150416	23863965	0	0	8463	277988
10000	24150184	24065539	0	0	3140	81505
20000	24150248	23986340	0	0	4805	159103

4.2 Gauss-Seidel Runtime/speedup analysis

We only tested this parallel approach on a 1000x1000 matrix (with randomly generated doubles between 0 and 100) due to realizing the Jacobi version offered more interesting results.

The following analysis shows that even though the GS version is simple, its merging overhead bottlenecks the speed.

Number of Cores	Runtime (s)	Speedup
1	176.316	1.00×
2	256.689	0.69×
3	301.140	0.59×
4	298.948	0.59×
5	267.741	0.66×
6	298.785	0.59×
7	298.913	0.59×
8	196.961	0.90×

Table 1: Runtime and speedup across different numbers of cores.


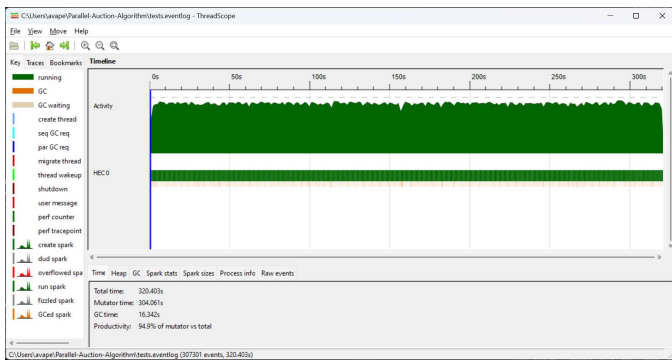
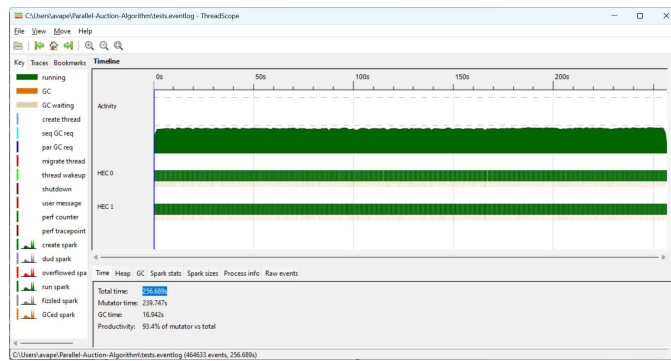


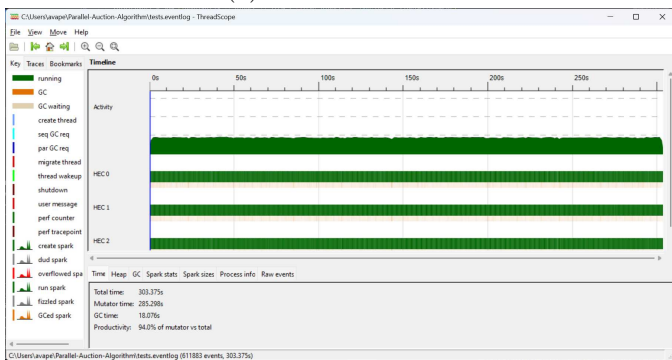
Figure 1: Gauss-Seidel speedup for 1000x1000 matrix



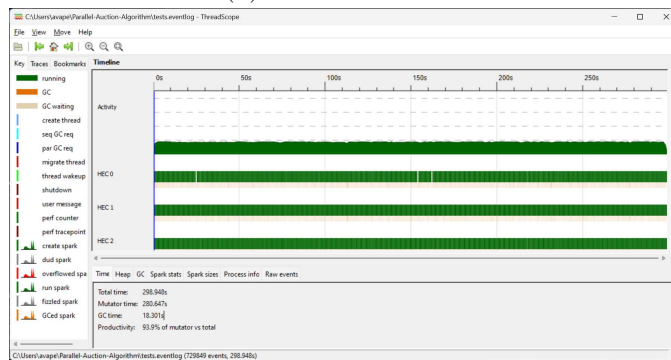
(a) 1 core: 176.316s



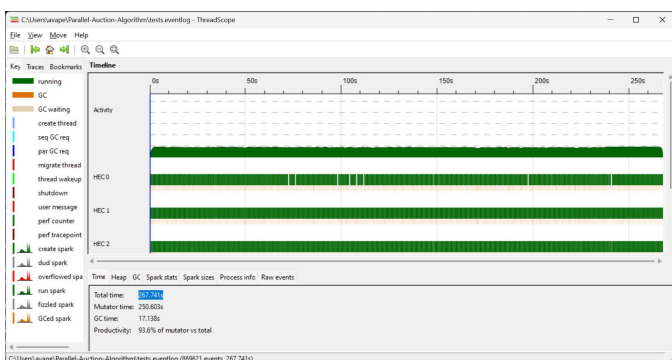
(b) 2 cores: 256.689s



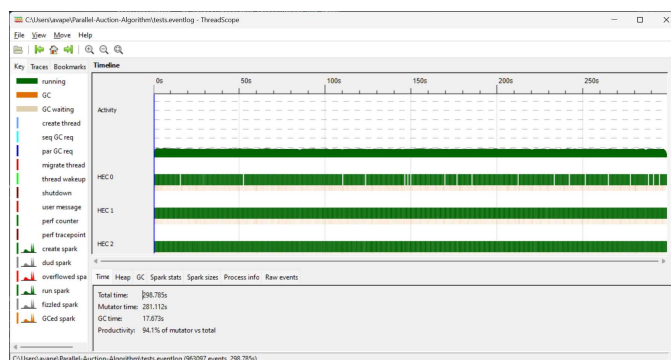
(c) 3 cores: 301.140s



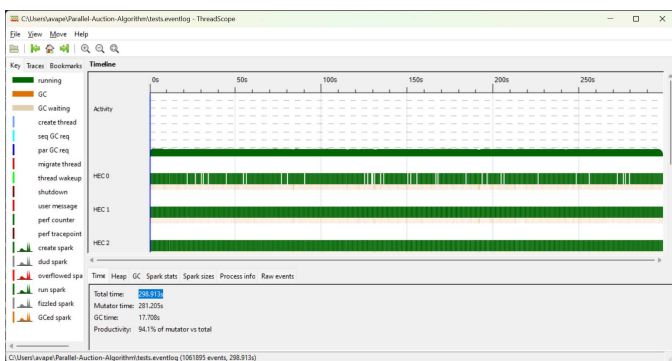
(d) 4 cores: 298.948s



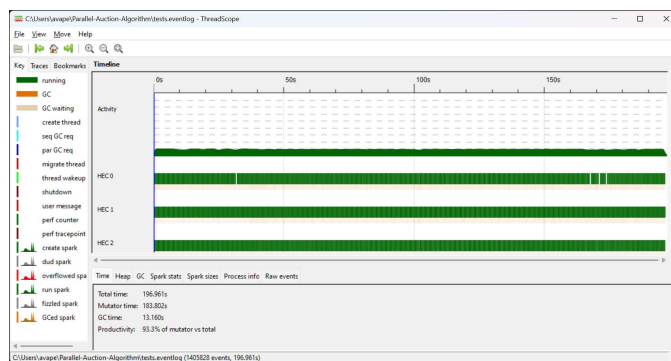
(e) 5 cores: 267.741s



(f) 6 cores: 298.785s



(g) 7 cores: 298.913s



(h) 8 cores: 196.961s

Figure 2: Execution times and event logs for different core counts

5 Parallelization: Jacobi

The Jacobi version parallelizes step 3 of the algorithm. It allows multiple bidders to search for their bids simultaneously. Through parallelization, each core handles a portion of the total bidders awaiting, reducing the runtime. Each thread handles one bidder. It may happen that two or more bidders make bids for the same item in parallel; in this case, we can only make one of them the tentative owner of the item. There is also one synchronization stage at the end of every iteration: we have to make sure several bidders bidding for the same item do not conflict, since the prices used to search for the best item may be outdated. It has been proven though that even with outdated prices during the search, updating the price as long as the new price is higher than the original (but latest) price is still correct.

By focusing on bidders rather than items, the Jacobi version avoids the merging overhead present in the Gauss-Seidel version, offering more interesting results.

5.1 Parallelization Choices

We implemented the algorithm in parallel using Haskell's `Control.Parallel.Strategies` library. Specifically, the `parMap` function is used to divide the workload across multiple threads, with each thread independently processing a partition of items.

Here is a snippet of code that uses parallelization. Please see the appendix for the entirety of the code:

```
1 synchronizedParallelBidding :: [Bidder] -> Prices -> [(Bidder, Item, Double)]
2 synchronizedParallelBidding bidders prices =
3   map (bestBid prices) bidders 'using' parList rdeepseq
```

- `using` applies the parallel evaluation strategy (`parList rdeepseq`) to a list of bidders.
- The `parList` strategy evaluates each element of a list in parallel.
- The `rdeepseq` strategy ensures that each element in the list is fully evaluated to normal form before being returned—it's used because the bid computation must be fully carried out before results can be merged.

Essentially, what this does is it creates a spark for each element in the list returned by `map (bestBid prices) bidders`. This is the same as each spark corresponding to finding the best item and bid price for a single bidder. This level of granularity was chosen because it was just the first implementation we tried and it happened to distribute the workload well. The total number of sparks however changes problem to problem, since the number of bidders in the subset U at any given iteration is variable depending on the payoff matrix. It changes even more drastically when the size of the matrix changes. For some measure of the problem size and how well it parallelizes we include the sparks information for a 1000x1000 matrix and a 3000x3000 matrix:

Size of matrix	Total sparks	Converted	Overflowed	GC'd	dud	Fizzled
1000x1000	1411	1407	0	0	0	4
3000x3000	4324	0	0	0	0	0

5.2 Jacobi Runtime/speedup analysis (1000x1000)

The table below summarizes the runtime of the auction algorithm executed on different numbers of cores, for a test case of 1000x1000.

Number of Cores	Runtime (s/ms)	Speedup
1	3.840 s	1.00×
2	3.709 s	1.04×
3	2.230 s	1.72×
4	2.351 s	1.63×
5	1.356 s	2.83×
6	1.292 s	2.97×
7	1.406 s	2.73×
8	969.87 ms	3.96×

Table 2: Runtime and speedup across different numbers of cores.

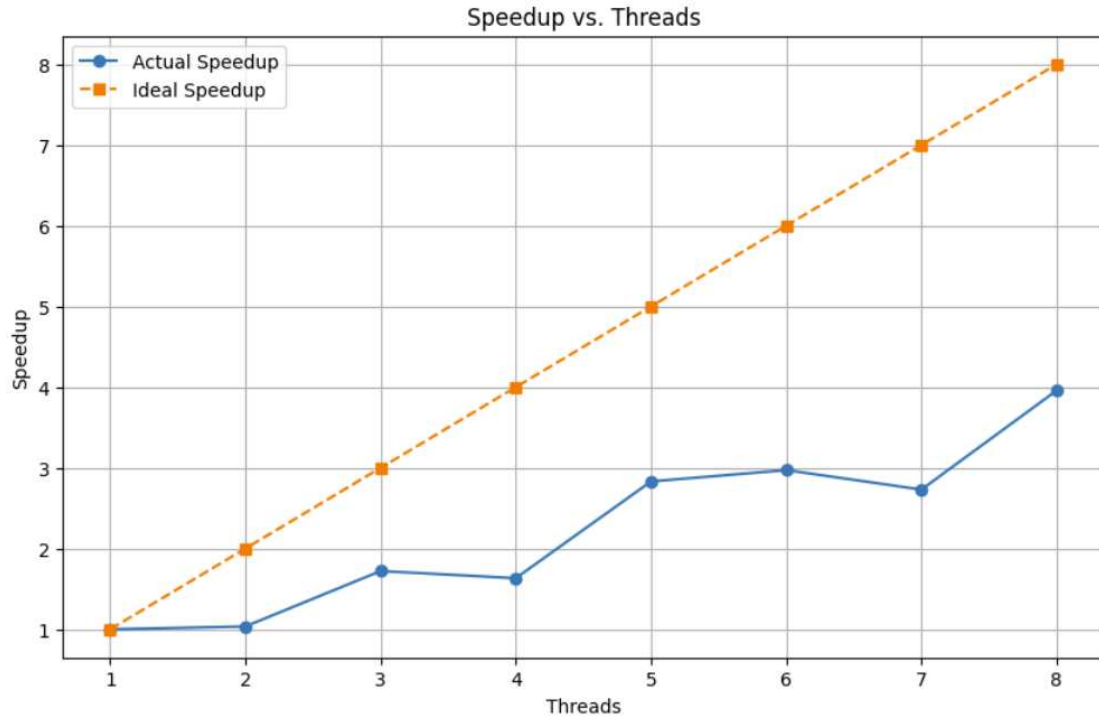
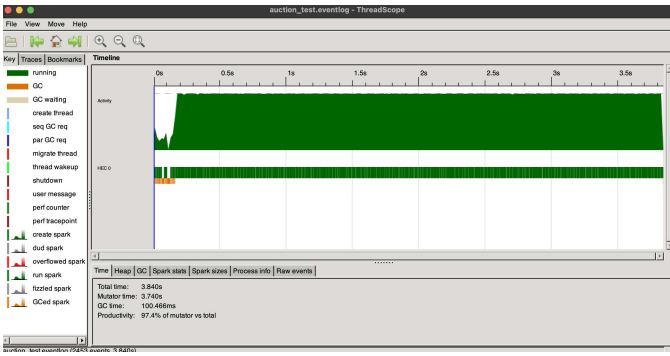
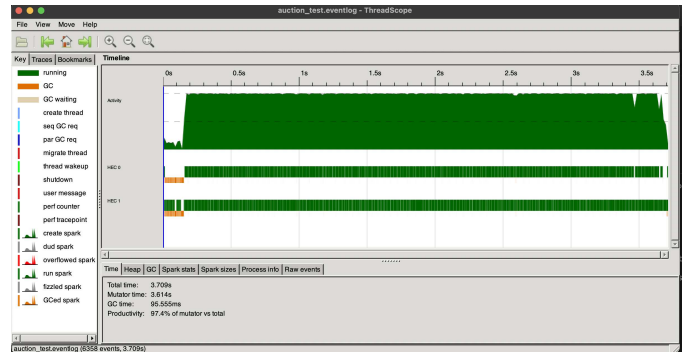


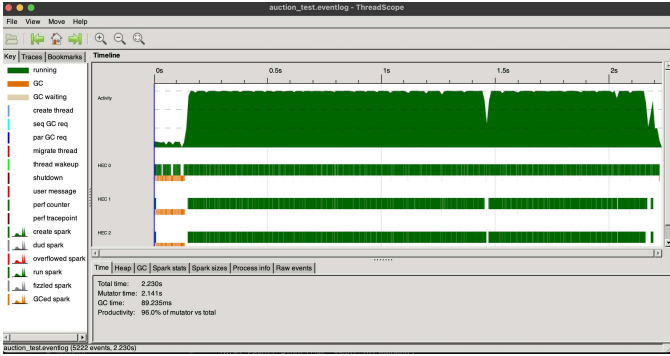
Figure 3: Actual speedup and ideal speedup



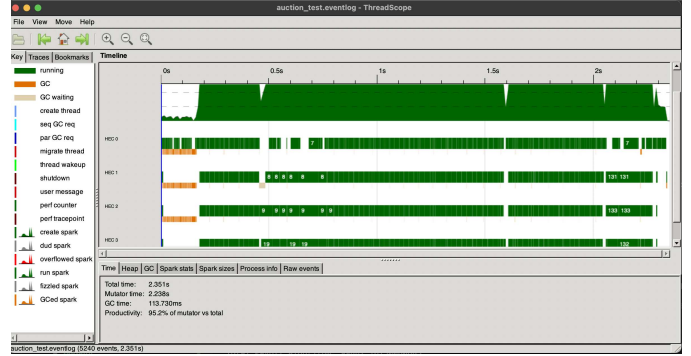
(a) -N1 eventlog



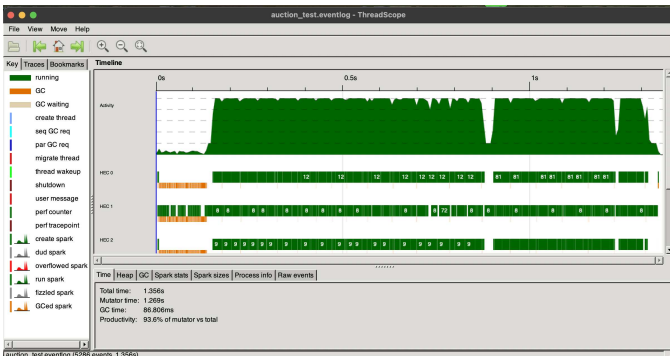
(b) -N2 eventlog



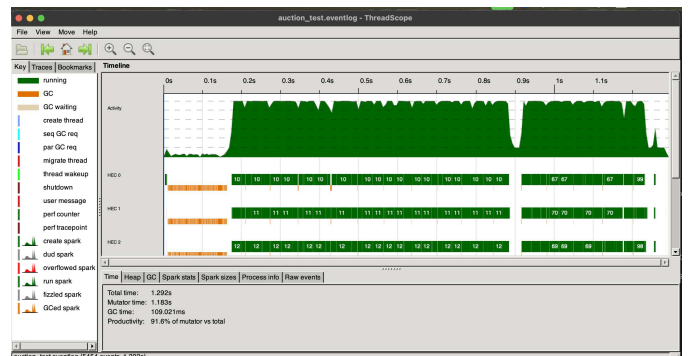
(c) -N3 eventlog



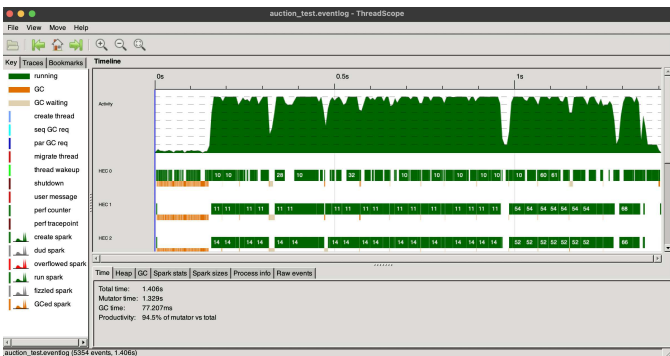
(d) -N4 eventlog



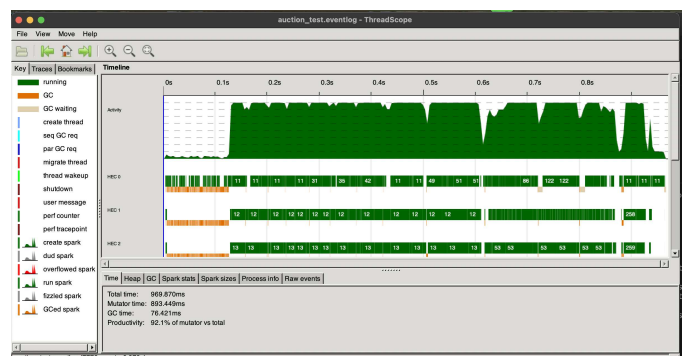
(e) -N5 eventlog



(f) -N6 eventlog



(g) -N7 eventlog



(h) -N8 eventlog

Figure 4: Jacobi Algorithm Eventlog for 1000 x 1000 matrix

As one can observe, the productivity measures for all numbers of cores are above 90%, signaling efficient core usage. However, the speedup is not ideal, since the test dataset is not large enough. We will test a larger use case of 3000 x 3000 to demonstrate the parallel algorithm's speedup ability.

5.3 Jacobi Runtime/speedup analysis (3000x3000)

The table below shows the runtime and speedup of the auction algorithm for a larger test case with a 3000x3000 matrix.

As the matrix size becomes larger, the speedup is more apparent. In this test case of 3000x3000, where runtime takes a measure of minutes, the speedup is closer to perfect.

Number of Cores	Runtime (s/ms)	Speedup
1	151.657 s	1.00×
2	82.560 s	1.83×
3	53.064 s	2.85×
4	40.421 s	3.75×
5	30.587 s	4.95×
6	27.942 s	5.42×
7	24.039 s	6.31×
8	21.226 s	7.14×

Table 3: Runtime and speedup across different numbers of cores for a 3000 x 3000 matrix.

As one observes, the speedup is more ideal as the matrix size gets larger. This is because as the matrix grows larger, the computations become more significant than the spark overhead. The speedup diagram demonstrates that we can achieve near-ideal speedup using the Jacobi algorithm.


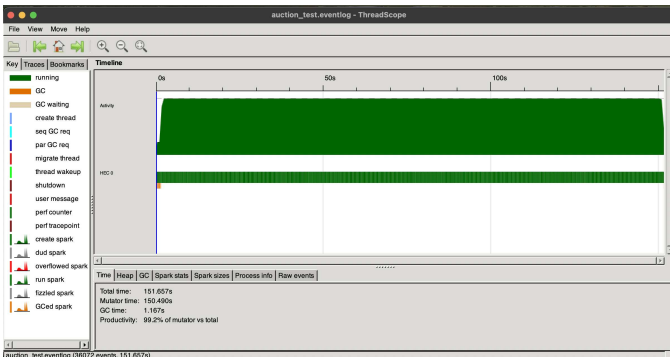
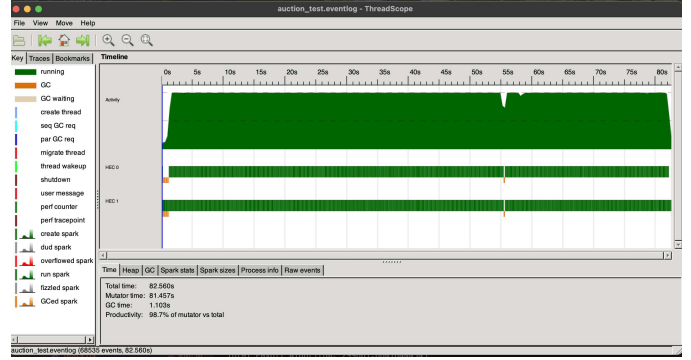


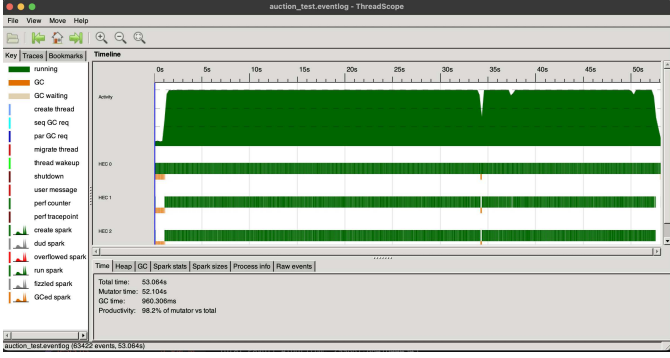
Figure 5: Actual speedup and ideal speedup



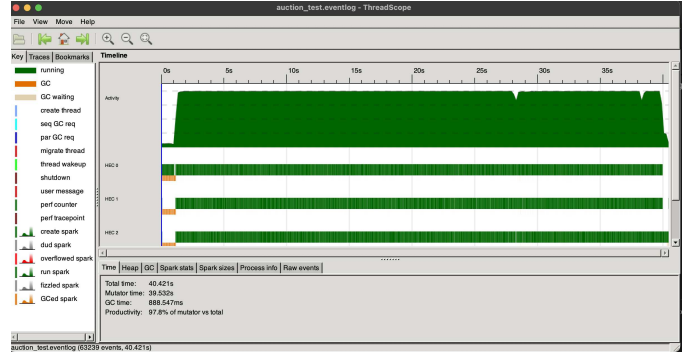
(a) -N1 eventlog



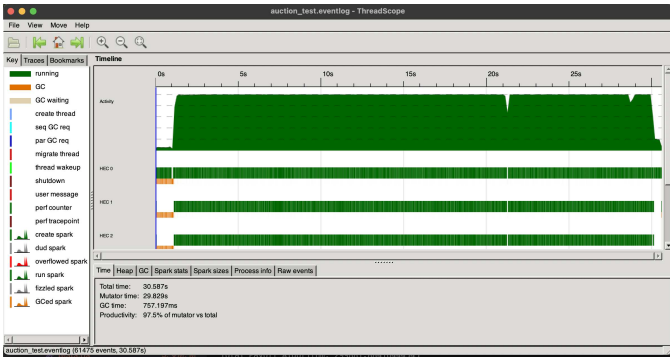
(b) -N2 eventlog



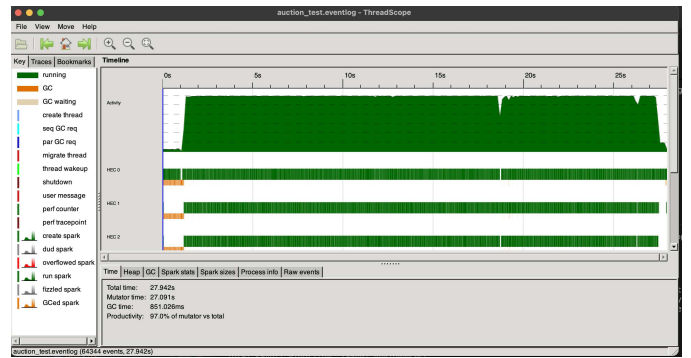
(c) -N3 eventlog



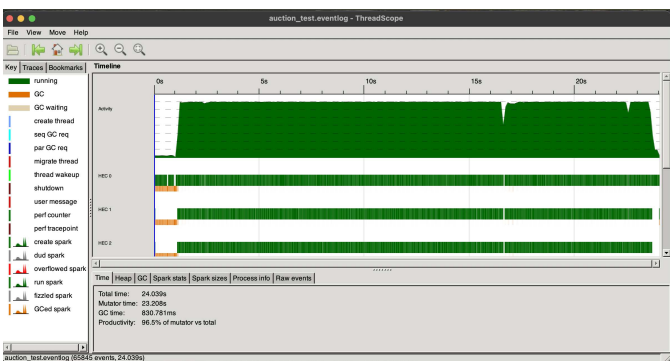
(d) -N4 eventlog



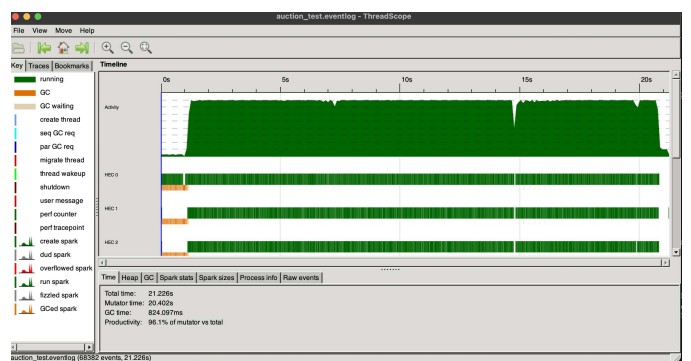
(e) -N5 eventlog



(f) -N6 eventlog



(g) -N7 eventlog



(h) -N8 eventlog

Figure 6: Jacobi Algorithm Eventlog for 3000 x 3000 matrix

6 Conclusion

- The assignment problem becomes more computationally practical in parallel!
- The Gauss-Seidel implementation faces significant synchronization overhead and load imbalance issues, resulting in substantially slower performance compared to the Jacobi version.

- As the data size increases, the Jacobi algorithm demonstrates near-ideal scalability, making it a highly effective approach for parallelization.

Note about testing: Testing was initially done with seven small matrices (with dimensions less than 6x6) to verify correctness. Once the algorithm was verified, random generation was introduced with the ability to adjust the size of the matrix through a command-line argument. Please see the test file and README.md for usage.

7 References

[1] Jin, J. (2016). Parallel Auction Algorithm for Linear Assignment Problem.

8 Appendix

gs_auction.hs

```

1 module GSAuction (gsAuctionAlgorithm) where
2
3 import Control.Parallel.Strategies
4 import qualified Data.Map as Map
5 import Data.List
6 import Data.Maybe
7 import Data.Ord (comparing, Down(..))
8
9 type PayoffMatrix = [[Double]]
10 type Bidder = Int
11 type Item = Int
12 type Prices = Map.Map Item Double
13 type Assignment = Map.Map Bidder Item
14
15 gsAuctionAlgorithm :: Double -> PayoffMatrix -> (Assignment, Double)
16 gsAuctionAlgorithm epsilon inputMatrix = (finalAssignment, totalPayoff)
17   where
18     numItems = length (head inputMatrix)
19     initialUnassigned = [0 .. length inputMatrix - 1]
20     initialPrices = Map.fromList [(j, 0) | j <- [0 .. numItems - 1]]
21
22     -- get the resulting assignment and also the total payoff, to return
23     finalAssignment = go initialUnassigned initialPrices Map.empty
24     totalPayoff = sum [inputMatrix !! bidder !! item | (item, bidder) <- Map.toList
25                       finalAssignment]
26
27     go :: [Bidder] -> Prices -> Assignment -> Assignment
28     go [] _ assignment = assignment
29     go (i : unassignedBidders) prices assignment =
30       let
31         -- calculate net payoffs for all items
32         netPayoffs = [(j, netPayoff i j prices) | j <- [0 .. numItems - 1]]
33
34         -- parallelize the search for best and second-best items
35         partitions = chunkItems 1600 netPayoffs -- change this number iteratively
36                   to find the best size chunk
37         -- for chunks: tested 2, 4, 8, 20, 100, 400, 1600, 6400, 10000, 20000
38         partialResults = parMap rpar findBestAndSecond partitions
39         (bestItem, maxPayoff, secondMaxPayoff) = mergeResults partialResults
40         epsilon
41
42         -- update price according to the auction algorithm description
43         newPrice = (prices Map.! bestItem) + (maxPayoff - secondMaxPayoff + epsilon
44           )
45         updatedPrices = Map.insert bestItem newPrice prices

```

```

43     -- handle previous assignment of the item
44     (newAssignment, remainingUnassigned) =
45         case Map.lookup bestItem assignment of
46             Just prevBidder ->
47                 let updatedAssignment = Map.insert bestItem i assignment
48                     updatedUnassigned = prevBidder : unassignedBidders
49                     in (updatedAssignment, updatedUnassigned)
50             Nothing ->
51                 (Map.insert bestItem i assignment, unassignedBidders)
52     in go remainingUnassigned updatedPrices newAssignment
53
54 -- calculate net payoff for a bidder for a specific item
55 netPayoff :: Bidder -> Item -> Prices -> Double
56 netPayoff i j prices = inputMatrix !! i !! j - (prices Map.! j)
57
58 -- find the best and second-best items in a partition
59 findBestAndSecond :: [(Item, Double)] -> (Item, Double, Maybe Double)
60 findBestAndSecond payoffs =
61     let (bestItem, maxPayoff) = maximumBy (comparing snd) payoffs
62         secondMaxPayoff = if length payoffs > 1
63                             then Just $ maximum $ map snd (filter ((/= bestItem) .
64                                 fst) payoffs)
65                             else Nothing
66     in (bestItem, maxPayoff, secondMaxPayoff)
67
68 -- merge results from all partitions
69 mergeResults :: [(Item, Double, Maybe Double)] -> Double -> (Item, Double,
70     Double)
71 mergeResults results epsilon =
72     let
73         allPayoffsWithItems = concatMap (\(item, p, ms) -> [(item, p), (item,
74             fromMaybe (-1 / 0) ms)]) results
75         sortedPayoffsWithItems = sortBy (comparing (Down . snd))
76             allPayoffsWithItems
77         (bestItem, maxPayoff) = head sortedPayoffsWithItems
78         secondMaxPayoff = if length sortedPayoffsWithItems > 1
79                             then snd (sortedPayoffsWithItems !! 1)
80                             else maxPayoff - epsilon
81     in (bestItem, maxPayoff, secondMaxPayoff)
82
83 -- split items into equal-sized chunks for parallel processing
84 chunkItems :: Int -> [a] -> [[a]]
85 chunkItems n items = let (q, r) = length items `quotRem` n
86                       in goChunks q r items
87
88     where
89         goChunks _ 0 [] = []
90         goChunks q r xs = let (chunk, rest) = splitAt (q + if r > 0 then 1 else 0)
91                             xs
92                       in chunk : goChunks q (max 0 (r - 1)) rest

```

jacobi_auction.hs

```

1  module JacobiAuction (jacobiAuctionAlgorithm) where
2
3  import Control.Parallel.Strategies (parList, rdeepseq, using)
4  import Data.List (maximumBy, foldl')
5  import Data.Ord (comparing)
6  import qualified Data.Map as Map
7
8  type Bidder = Int
9  type Item = Int
10 type Prices = Map.Map Item Double
11 type Assignment = Map.Map Item Bidder -- mapping from item to bidder (to correspond
12                                         to implementation in paper)
13 type PayoffMatrix = [[Double]]

```

```

13
14 jacobiAuctionAlgorithm :: Double -> PayoffMatrix -> (Assignment, Double)
15 jacobiAuctionAlgorithm epsilon inputMatrix = (finalAssignment, totalPayoff)
16   where
17     numItems = length (head inputMatrix)
18     initialUnassigned = [0 .. length inputMatrix - 1]
19     initialPrices = Map.fromList [(j, 0) | j <- [0 .. numItems - 1]]
20
21     -- get the resulting assignment and also the total payoff, to return
22     (finalAssignment, _) = runSynchronizedAuction initialUnassigned initialPrices
23     Map.empty
24     totalPayoff = sum [inputMatrix !! bidder !! item | (item, bidder) <- Map.toList
25                       finalAssignment]
26
27     runSynchronizedAuction :: [Bidder] -> Prices -> Assignment -> (Assignment, [
28       Bidder])
29     runSynchronizedAuction [] _ assignment = (assignment, [])
30     runSynchronizedAuction unassignedBidders prices assignment =
31       let
32         bidResults = synchronizedParallelBidding unassignedBidders prices
33         updatedPrices = foldl' updatePrices prices bidResults
34         (newAssignment, newUnassigned) = resolveConflicts bidResults assignment
35       in
36         if null newUnassigned
37         then (newAssignment, newUnassigned)
38         else runSynchronizedAuction newUnassigned updatedPrices newAssignment
39
40     synchronizedParallelBidding :: [Bidder] -> Prices -> [(Bidder, Item, Double)]
41     synchronizedParallelBidding bidders prices =
42       map (bestBid prices) bidders 'using' parList rdeepseq
43
44     -- find the best item and second-best payoff for a bidder
45     bestBid :: Prices -> Bidder -> (Bidder, Item, Double)
46     bestBid prices i =
47       let
48         netPayoffs = [(j, netPayoff i j prices) | j <- [0 .. numItems - 1]]
49         (bestItem, maxPayoff) = maximumBy (comparing snd) netPayoffs
50         secondMaxPayoff = if length netPayoffs > 1
51                           then maximum $ map snd (filter ((/= bestItem) . fst)
52                             netPayoffs)
53                           else maxPayoff - epsilon
54         bidPrice = (prices Map.! bestItem) + (maxPayoff - secondMaxPayoff + epsilon)
55       in (i, bestItem, bidPrice)
56
57     -- resolve conflicts: only one bidder can win an item
58     -- paper states that this will still result in the optimal assignment, even if
59     -- prices are outdated
60     resolveConflicts :: [(Bidder, Item, Double)] -> Assignment -> (Assignment, [
61       Bidder])
62     resolveConflicts bids assignment =
63       let
64         groupedBids = Map.fromListWith (++) [(item, [(bidder, bidPrice)]) | (bidder
65           , item, bidPrice) <- bids]
66         resolvedAssignments =
67           Map.mapWithKey (\_ bidders -> fst $ maximumBy (comparing snd) bidders)
68             groupedBids
69         newAssignment =
70           foldl' (\acc (item, bidder) -> Map.insert item bidder acc) assignment (
71             Map.toList resolvedAssignments)
72         unassignedBidders =
73           [bidder | (_, bidders) <- Map.toList groupedBids, (bidder, _) <- bidders,
74             bidder 'notElem' Map.elms newAssignment]
75       in (newAssignment, unassignedBidders)

```

```

66
67 -- update prices for items based on the winning bids
68 updatePrices :: Prices -> (Bidder, Item, Double) -> Prices
69 updatePrices prices (_, item, bidPrice) =
70     let currentPrice = Map.findWithDefault 0 item prices
71     in Map.insert item (max currentPrice bidPrice) prices
72
73 -- calculate net payoff for a bidder for a specific item
74 netPayoff :: Bidder -> Item -> Prices -> Double
75 netPayoff i j prices = inputMatrix !! i !! j - (prices Map.! j)

```

sequential_auction.hs

```

1 module SequentialAuction (auctionAlgorithm, optimalAssignment) where
2
3 import Data.List (maximumBy, permutations)
4 import Data.Ord (comparing)
5 import qualified Data.Map as Map
6
7 type Bidder = Int
8 type Item = Int
9 type Prices = Map.Map Item Double
10
11 -- item is the key, bidder is the value, for consistency with the algorithm from
12 -- the paper
13 type Assignment = Map.Map Item Bidder
14 type PayoffMatrix = [[Double]]
15
16 auctionAlgorithm :: Double -> PayoffMatrix -> (Assignment, Double)
17 auctionAlgorithm epsilon inputMatrix = (finalAssignment, totalPayoff)
18     where
19         numItems = length (head inputMatrix)
20         numBidders = length inputMatrix
21
22         initialUnassigned = [0 .. numBidders - 1]
23         initialPrices = Map.fromList [(j, 0) | j <- [0 .. numItems - 1]]
24
25         finalAssignment = go initialUnassigned initialPrices Map.empty
26         totalPayoff = sum [inputMatrix !! bidder !! item | (item, bidder) <- Map.toList
27             finalAssignment]
28
29 go :: [Bidder] -> Prices -> Assignment -> Assignment
30 go [] _ assignment = assignment
31 go (i : unassignedBidders) prices assignment =
32     let
33         -- calculate net payoffs for all items
34         netPayoffs = [(j, netPayoff i j prices) | j <- [0 .. numItems - 1]]
35
36         -- find the best and second-best items
37         (bestItem, maxPayoff) = maximumBy (comparing snd) netPayoffs
38         secondMaxPayoff = if length netPayoffs > 1
39             then maximum [ p | (j,p) <- netPayoffs, j /= bestItem ]
40             else maxPayoff - epsilon
41
42         -- update the price of the best item
43         newPrice = (prices Map.! bestItem) + (maxPayoff - secondMaxPayoff + epsilon)
44         updatedPrices = Map.insert bestItem newPrice prices
45
46         -- handle previous assignment of the item
47         (newAssignment, remainingUnassigned) =
48             case Map.lookup bestItem assignment of
49             Just prevBidder ->

```

```

        and add prevBidder back into U
50     let updatedAssignment = Map.insert bestItem i assignment -- reassign
        item to current bidder i
51         updatedUnassigned = prevBidder : unassignedBidders
52     in (updatedAssignment, updatedUnassigned)
53     Nothing ->
54     (Map.insert bestItem i assignment, unassignedBidders)
55     in go remainingUnassigned updatedPrices newAssignment
56
57 netPayoff :: Bidder -> Item -> Prices -> Double
58 netPayoff i j prices = inputMatrix !! i !! j - (prices Map.! j)
59
60 -- find the optimal assignment by brute force (adjusted to return item->bidder)
61 optimalAssignment :: PayoffMatrix -> Assignment
62 optimalAssignment matrix = maximumBy (comparing totalPayoff) assignments
63 where
64     bidders = [0 .. length matrix - 1]
65     items = bidders -- assume square matrix
66     assignments = [Map.fromList (zip items perm) | perm <- permutations bidders]
67     totalPayoff assignment = sum [matrix !! b !! i | (i,b) <- Map.toList assignment
    ]

```

tests.hs

```

1  module Main (main) where
2
3  import SequentialAuction (auctionAlgorithm)
4  import JacobiAuction (jacobiAuctionAlgorithm)
5  import GSAuction (gsAuctionAlgorithm)
6  import qualified Data.Map as Map
7  import Control.Monad (unless)
8  import System.Random (mkStdGen, randomRs, StdGen, split)
9  import System.Environment (getArgs, getProgName)
10 import System.Exit (die)
11 import System.IO.Error (catchIOError)
12
13 type Bidder = Int
14 type Item = Int
15 type PayoffMatrix = [[Double]]
16 type Assignment = Map.Map Item Bidder
17
18 roundToTenths :: Double -> Double
19 roundToTenths x = fromIntegral (round (x * 10)) / 10
20
21 printMatrix :: PayoffMatrix -> IO ()
22 printMatrix m = do
23     putStrLn "Price matrix:"
24     mapM_ (putStrLn . formatRow . map roundToTenths) m
25     where
26         formatRow :: [Double] -> String
27         formatRow row = "[" ++ unwords (map show row) ++ "]"
28
29
30 printAuctionResults :: PayoffMatrix -> Assignment -> Double -> IO ()
31 printAuctionResults matrix assignment totalPayoff = do
32     putStrLn "\nAssignments and payoffs:"
33     let payoffBreakdown = [(item, bidder, matrix !! bidder !! item) | (item, bidder
34         ) <- Map.toList assignment] -- !! is same as matrix[bidder][item]
35     mapM_ (\(i, b, p) -> putStrLn $ "Item " ++ show i ++ " -> Bidder " ++ show b ++
36         " : " ++ show (roundToTenths p)) payoffBreakdown
37     putStrLn $ "\nTotal payoff: " ++ show (roundToTenths totalPayoff)
38
39 runAlgorithm :: (Double -> PayoffMatrix -> (Assignment, Double)) -> PayoffMatrix ->
40     IO (Assignment, Double)
41 runAlgorithm algorithm matrix = do

```

```

39     let (assignment, totalPayoff) = algorithm 0.01 matrix -- always assume 0.01 is
        sufficient for epsilon
40     return (assignment, totalPayoff)
41
42 main :: IO ()
43 main = runProgram 'catchIOError' \_ ->
44     die "ERROR, try making sure the command-line arguments are formatted correctly"
45
46 runProgram :: IO ()
47 runProgram = do
48     args <- getArgs
49     case args of
50     [sizeStr, algStr] -> do
51         let maybeSize = reads sizeStr :: [(Int, String)] -- read the input, and
            cast as Int, String tuple
52         case maybeSize of
53         [(n, "")] -> do
54             algFunc <- case algStr of
55                 "seq"    -> return auctionAlgorithm
56                 "gs"     -> return gsAuctionAlgorithm
57                 "jacobi" -> return jacobiAuctionAlgorithm
58                 -        -> die "ERROR, please enter 'seq', 'gs'
                    ', or 'jacobi'"
59             let seed = 100 -- causes generated matrix to stay the same if
                file isn't reloaded
                gen = mkStdGen seed
                (matrix, _) = generateMatrix gen n n
                (assignment, totalPayoff) <- runAlgorithm algFunc matrix
60
61             -- show the results when the matrix is resonably small
62             -- assume that this case is used for testing correctness by
63             hand
64             if n < 6
65             then do
66                 printMatrix matrix
67                 printAuctionResults matrix assignment totalPayoff
68
69             -- assume that this case is used for testing runtime on
70             large matrices
71             else do
72                 putStrLn $ "Total payoff: " ++ show (roundToTenths
73                     totalPayoff)
74         _ -> do
75             pn <- getProgName
76             die $ "ERROR, Invalid command line arguments. Usage: " ++ pn
77     _ -> do
78         pn <- getProgName
79         die $ "ERROR, Usage: " ++ pn
80
81
82 generateMatrix :: StdGen -> Int -> Int -> (PayoffMatrix, StdGen)
83 generateMatrix gen rows cols = (matrix, finalGen)
84     where randomNumbers = take (rows * cols) $ randomRs (0.0, 100.0) gen
85         (finalGen, _) = split gen
86         matrix = chunksOf cols randomNumbers
87
88 -- splits the randomNumbers list into chunks of size cols
89 chunksOf :: Int -> [a] -> [[a]]
90 chunksOf _ [] = []
91 chunksOf n xs = let (ys, zs) = splitAt n xs in ys : chunksOf n zs

```