

Lower bounds for selection of maximum,  
minimum, second smallest element and the  
median  
DM19 notes fall 2006

Jørgen Bang-Jensen  
Imada, SDU

30. august 2006

## 1 Introduction

The purpose of this note is to prove lower bounds for the number of comparisons of numbers which every algorithm must make in order to determine both the maximum and the minimum element, respectively the second smallest element, respectively the median. Such lower bounds are already discussed in the first part of the notes “Noter til DM19” (the copy from the book by Baase). Here the lower bounds are established by showing how an adversary may generate bad input for any algorithm by following certain rules (which among other things involves assignment of values to the elements and possibly changing these values while the data set is still under construction).

One weakness of this way of formulating the adversary’s strategy is that it may be confusing that the values of the elements is changing during the construction. Hence, even though it is intuitively clear that the method works, one is left with a feeling of not really having seen a real proof for the lower bound. In these notes we do obtain a precise proof of the lower bounds by formulating the adversary’s strategy as that of forming an acyclic partial orientation of a complete graph. The orientation is constructed concurrently with the execution of the algorithm for which the lower bound is been established (that is, one arc is oriented for each comparison that is made by

the algorithm) and no orientation will ever be changed after it is assigned to an edge. Only at the very end of the process the actual values of the elements are decided, based on the final acyclic partial orientation when the algorithm terminates. As we shall show below, using the orientation at that time, the adversary can easily assign values to the elements in such a way that all comparisons made by the algorithm receive a valid answer. Thus we can obtain a lower bound for the number of comparisons which must be made by the algorithm, provided that we can argue that this many arcs must always be oriented for a unique answer to exist.

## 2 Some terminology and notation

We often denote a digraph  $D = (V, A)$  where  $V$  is the set of vertices and  $A$  the set of arcs. We may also use the notation  $V(D)$  for the set of vertices and  $A(D)$  for the set of arcs of  $D$ . An arc from a vertex  $u$  to a vertex  $v$  is denoted  $u \rightarrow v$ . We say that the arc *starts* in  $u$  and *ends* in  $v$ . For every vertex  $x \in V(D)$  we denote by  $d_D^+(x)$  ( $d_D^-(x)$ ) the number of arcs which start (end) in  $x$ . We also use  $d_D(x)$  for the total number of arcs in  $D$  which are incident with the vertex  $x$ . Thus  $d_D(x) = d_D^+(x) + d_D^-(x)$ . The *complete graph*  $K_n$  is the unique graph (up to isomorphism) on  $n$  vertices in which every pair of distinct vertices is joined by an edge.

An *out-arborescence* is an orientation of a tree such that every vertex except one (called the *root*) has precisely one arc ending in it. No arc ends in the root and it is easy to show that if  $r$  is the root in an out-arborescence  $T$ , then  $r$  can reach every other vertex  $x \in V(T)$  by a directed path (Exercise 1).

## 3 Sorting and acyclic partial orientations of $K_n$

In the notes on tournaments we have seen that given  $n$  distinct real numbers  $x_1, x_2, \dots, x_n$  we can define a tournament  $TT_n = (V, A)$  where  $V = \{x_1, x_2, \dots, x_n\}$  and  $A = \{x_i \rightarrow x_j : x_i < x_j\}$ . Thus we translate the property that  $x_i < x_j$  to an arc from  $x_i$  to  $x_j$  in  $TT_n$ . The tournament  $TT_n$  is acyclic and has the property that if  $x_i \rightarrow x_j$  and  $x_j \rightarrow x_k$  then we also have that  $x_i \rightarrow x_k$ . In other words, the relation  $R = \{(x_i, x_j) : x_i \rightarrow x_j \in A\}$  is transitive.

It is easy to show that up to isomorphism there is precisely one transitive tournament on  $n$  vertices (Exercise 2).

We know from DM02 that every acyclic digraph  $D$  has a topological ordering (below we shall use the more appropriate name *acyclic ordering*) that is, an ordering  $x_{i_1}, x_{i_2}, \dots, x_{i_n}$  such that  $D$  has no arc  $x_{i_q} \rightarrow x_{i_p}$  in  $D$  with the property that  $q > p$ . Clearly the opposite is also true: if  $D$  has an acyclic ordering, then  $D$  must be acyclic (all arcs go forward with respect to the ordering). The proof of the following small observation is left to the reader as Exercise 3.

**Lemma 3.1** *Let  $M = (V, A \cup E)$  be obtained from the complete graph  $K_n$  by orienting a subset  $A$  of the edges of  $K_n$  (the set  $E$  denotes those edges which still have no orientation assigned to them). Let  $D = (V, A)$  be the digraph consisting of all the vertices of  $V$  and the arcs  $A$ . Then one can orient the edges in  $E$  such that the resulting tournament  $T = (V, A \cup A')$  is acyclic (that is, it is isomorphic to  $TT_n$ ) if and only if  $D$  is acyclic (here  $A'$  denotes the set of arcs we obtain by orienting the edges of  $E$ ).*

From the discussion above it follows that given any acyclic digraph  $D = (V, A)$  with  $n$  vertices  $V = \{v_1, v_2, \dots, v_n\}$  we can assign a set  $S$  of  $n$  distinct real numbers (one for each vertex)  $x_1, x_2, \dots, x_n$ , with the property that  $x_i > x_j$  if and only if there is no directed path from  $v_i$  to  $v_j$  in  $D$ . To accomplish this, simply start from an acyclic ordering  $\mathcal{L}$  of  $V(D)$  and let the numbers in  $S$  be arbitrary distinct numbers which satisfy that  $x_i > x_j$  if and only if  $v_i$  is after  $v_j$  in  $\mathcal{L}$ . In particular, we may always choose  $S$  to be the set  $\{1, 2, \dots, n\}$  (by associating the number  $i$  with the  $i$ 'th vertex in  $\mathcal{L}$ ).

Thus we have shown that a partial orientation  $M = (V, A \cup E)$  of the complete graph  $K_n$  corresponds to a permutation of the numbers  $1, 2, \dots, n$  if and only if the oriented part  $D = (V, A)$  is acyclic. It is also clear that every permutation of  $\{1, 2, \dots, n\}$  can be realized in this way. We will make use of these facts below by constructing partial orientations of  $K_n$  which will always be acyclic and hence will correspond to some suitably chosen permutation of  $\{1, 2, \dots, n\}$ .

Our arguments below concerning partial orientations  $D$  of  $K_n$  can be translated to concrete inputs which will force the actual algorithm to perform the postulated number of comparisons, simply by replacing the vertices of  $D$  by the numbers  $1, 2, \dots, n$  in such a way that the number  $x_k$  (the  $k$ 'th number of the input) is assigned the value  $i$  precisely if the vertex  $v_k \in V$  is assigned the number (place)  $i$  in the current acyclic ordering  $\mathcal{L}$  of  $D$ .

We denote by  $Q(x, y)$  the query which asks for a comparison of  $x$  and  $y$  by the algorithm. We shall always “answer” such a query by assigning an orientation to the edge  $xy$  of  $K_n$ . Since no optimal algorithm would compare the same numbers twice, we may assume that  $x$  and  $y$  have not been compared earlier by the algorithm, nor will they be so at a later stage of the algorithm.

We will always use  $D$  to denote the actual oriented graph with  $n$  vertices and those edges which have so far been oriented (that is in the beginning when no edge has been oriented we have  $D = (V, \emptyset)$ ) and we will let  $\mathcal{L}(D)$  denote the current acyclic ordering of  $D$ . Furthermore, we denote by  $\mathcal{L}_D(x)$  the position of the vertex  $x$  in  $\mathcal{L}(D)$  (that is,  $\mathcal{L}_D(x) = i$  if and only if  $x$  is the  $i$ 'th vertex in  $\mathcal{L}(D)$ ). The adversary will update  $\mathcal{L}(D)$  as new arcs are oriented so that at any time  $\mathcal{L}(D)$  is an acyclic ordering of the current digraph  $D$ .

## 4 Lower bound for finding both the maximum and the minimum element among $n$ numbers

The adversary will build a partial orientation of  $K_n$  which is always acyclic by always answering the query  $Q(x, y)$  according to the following strategy.

1. If  $d_D^+(x) = 0 = d_D^+(y)$ , then add the arc  $x \rightarrow y$  to  $D$  if  $d_D^-(y) > 0$  and add the arc  $y \rightarrow x$  to  $D$  if  $d_D^-(y) = 0$ .
2. If  $d_D^+(x) = 0 < d_D^+(y)$ , then add the arc  $y \rightarrow x$  to  $D$ .
3. If  $d_D^+(x) > 0$  and either  $d_D^+(y) = 0$ , or  $d_D^-(x) = 0$ , then add the arc  $x \rightarrow y$  to  $D$ .
4. If  $d_D^+(x)d_D^-(x) > 0$  and  $d_D^+(y) > 0$ , but  $d_D^-(y) = 0$ , then add the arc  $y \rightarrow x$  to  $D$ .
5. If  $d_D^+(x)d_D^-(x), d_D^+(y)d_D^-(y) > 0$ , then add the arc  $x \rightarrow y$  to  $D$  if  $\mathcal{L}_D(x) < \mathcal{L}_D(y)$  and otherwise add the arc  $y \rightarrow x$  to  $D$ .

**Lemma 4.1** *If  $D$  is acyclic and the edge  $xy$  is oriented as above then the resulting digraph is also acyclic.*

**Proof:** Clearly it is enough to show the the new arc that is added cannot be part of a cycle. In 1. the arc  $x \rightarrow y$  cannot be part of a cycle as  $d_D^+(y) = 0$ . The same argument shows that  $y \rightarrow x$  cannot be part of a cycle in 2. In 3. we either have  $d_D^+(y) = 0$ , or  $d_D^-(x) = 0$ , implying that the arc  $x \rightarrow y$  cannot be part of a cycle. In 4. the arc  $y \rightarrow x$  cannot be part of a cycle as  $d_D^-(y) = 0$ . Finally, in 5. we add the arc such that it goes forward with respect to the acyclic ordering of  $D$  and hence the new arc cannot be part of a cycle.  $\diamond$

Every algorithm for finding both the maximum and the minimum number among  $n$  distinct numbers must obtain at least  $2n - 2$  units of information. Namely, it has to exclude  $n - 1$  numbers from being the maximum and  $n - 1$  numbers from being the minimum among the  $n$  numbers. The only step in which the adversary is forced to deliver 2 units of new information is when rule 1. is applied and we have  $d_D^-(x) = d_D^-(y) = 0$ . When rules 2.-4. are applied at most one unit of new information is given and when rule 5. is applied no new information is given.

**Theorem 4.2** *Every algorithm which finds both the maximum and the minimum element among  $n$  distinct numbers must perform at least  $\lceil 3n/2 \rceil - 2$  comparisons.*

**Proof:** Let  $\mathcal{A}$  be an arbitrary algorithm for the problem. By orienting  $K_n$  concurrently with the execution of  $\mathcal{A}$  according to the rules above (which tell the adversary which orientation to choose when  $\mathcal{A}$  asks the query  $Q(x, y)$ ) the adversary will deliver 2 pieces of new information (that is one of  $x$  and  $y$  is excluded from being the maximum and the other from being the minimum) at most  $\lfloor n/2 \rfloor$  times. Thus in order for  $\mathcal{A}$  to collect the required information  $\mathcal{A}$  must make at least  $2n - 2 - \lfloor n/2 \rfloor = \lceil 3n/2 \rceil - 2$  comparisons.  $\diamond$

## 5 Lower bounds for finding the second smallest element among $n$ distinct numbers

Obviously the problem is precisely as difficult as that of finding the second largest element among  $n$  distinct numbers. In the first part of these notes (Baase) it is shown that one can solve that problem using  $n + \lceil \log n \rceil - 2$  comparisons by using the so-called tournament method (not to be confused

with the mathematical concept of a tournament). It can easily be seen that when we use this method for finding the second smallest element, every element which has only been compared with the final minimum element could potentially still be the second smallest. Thus the adversary's strategy will be to ensure that the minimum element is compared with  $\lceil \log n \rceil$  different numbers.

The adversary constructs an acyclic partial orientation of  $K_n$ , starting from  $D = (V, \emptyset)$  by answering the queries of the form  $Q(x, y)$  as follows. The adversary maintains a collection of disjoint out-arborescences  $T_1, T_2, \dots, T_k$ . Every vertex  $x$  for which we currently have  $d_D^-(x) = 0$  is the root of one of these arborescences (we also denote this arborescence by  $T(x)$ ). At the beginning (when no edges have been oriented) every vertex is the root of its own private out-arborescence which consists only of the vertex itself and has no arcs. In order to distinguish between arcs which are part of an out-arborescence and arcs which are not we partition the arc set of  $D$  into two sets: black arcs corresponding to arcs in the arborescences and red arcs corresponding to the rest of the arcs of  $D$  (see Figure 1).

1. If  $d_D^-(x) = d_D^-(y) = 0$  (that is  $x$  and  $y$  are both the roots of out-arborescences), then add a black arc  $x \rightarrow y$  to  $D$  if  $|V(T(x))| > |V(T(y))|$  and otherwise add a black arc  $y \rightarrow x$  to  $D$ .
2. If  $d_D^-(x) = 0 < d_D^-(y)$ , then add a red arc  $x \rightarrow y$  to  $D$ .
3. If  $d_D^-(y) = 0 < d_D^-(x)$ , then add a red arc  $y \rightarrow x$  to  $D$ .
4. If  $d_D^-(x), d_D^-(y) > 0$ , then add a red arc  $x \rightarrow y$  to  $D$  if  $\mathcal{L}_D(x) < \mathcal{L}_D(y)$  and otherwise add a red arc  $y \rightarrow x$  to  $D$ .

**Lemma 5.1** *If  $D$  is acyclic and the edge  $xy$  is oriented according to the rules above, then the resulting digraph (obtained by adding the new arc to  $D$ ) is still acyclic and the black arcs form a forest of out-arborescences.*

**Proof:** The last claim follows from the fact that a new black arc will always start in the root of a black out-arborescence to the root of another one (implying that the later ceases being the root of an out-arborescence and that the resulting out-arborescence contains all the vertices of the two original out-arborescences). None of the rules 1.-3. can result in a cycle after




Figure 1: A time picture of the oriented graph  $D$  which is constructed by the adversary. Dotted arcs correspond to red arcs and full arcs to black arcs. There are currently 5 disjoint out-arborescences with respectively 8,4,2,1 and 1 vertices.

adding the new arc as the arcs will always start in a vertex  $z$  with  $d_D^-(z) = 0$ . In rule 4, we orient in such a way that  $\mathcal{L}(D)$  is still an acyclic ordering of the resulting digraph after adding the new arc. Hence no cycle can be created by rule 4.  $\diamond$

**Theorem 5.2** *By applying the above strategy the adversary will force any comparison based algorithm  $\mathcal{B}$  to use at least  $n + \lceil \log n \rceil - 2$  comparisons in order to determine the second smallest element.*

**Proof:** Let  $\mathcal{B}$  be an arbitrary comparison based algorithm for the problem. As long as a vertex  $x$  has  $d_D^-(x) = 0$ ,  $x$  is a possible candidate for being the minimum element and as soon as  $x$  has  $d_D^-(x) > 0$ , then  $x$  can no longer be the minimum (according to the way we assign values to the element of  $D$  upon termination of the algorithm). Hence when  $\mathcal{B}$  terminates there will be precisely one vertex  $x$  with  $d_D^-(x) = 0$ . It follows from the rules above that a vertex  $z$  can only change from having  $d_D^-(z) = 0$  to having  $d_D^-(z) > 0$  when a black arc is added into the vertex hence it follows (by induction) that the unique vertex  $x$  which has  $d_D^-(x) = 0$  when  $\mathcal{B}$  terminates is precisely the root of the resulting spanning out-arborescence which contains all  $n$  vertices of  $D$ . All vertices in this out-arborescence which are children of  $x$  are candidates for being the second smallest element after making all the comparisons corresponding to the black arcs. This follows from the fact that the only vertex that can reach these (besides themselves) by a directed path using only black arcs is  $x$ . Hence such a vertex can only be excluded from ending up at the second place in the final acyclic ordering by having a red arc into it. Every time a black arc is added, the size of the resulting black out-arborescence is

increased by at most two (see rule 1.). Thus  $x$  has been the root of at least  $\lceil \log n \rceil$  different out-arborescences during the execution of  $\mathcal{B}$  and hence  $x$  has at least  $\lceil \log n \rceil$  children in the final out-arborescence.

It follows from rules 1.-4. that, at any time during the execution of  $\mathcal{B}$ , if  $z$  has a red arc into it, then  $z$  also has a black arc into it. In order to ensure that only one of the children of  $x$  can be number two in the final acyclic ordering at least  $t - 1$  red arcs must be generated where  $t$  denotes the number of children of  $x$  in the final out-arborescence (which contains precisely  $n - 1$  black arcs at termination of  $\mathcal{B}$ ). Thus  $\mathcal{B}$  must make at least  $(n - 1) + (\lceil \log n \rceil - 1) = n + \lceil \log n \rceil - 2$  comparisons in order to determine the second largest element.  $\diamond$

## 6 A lower bound for the number of comparisons need to find the median of $n$ numbers

For simplicity we assume below that the numbers are all distinct and that  $n$  is odd, implying that the median is unique.

Every algorithm for determining the median  $m$  among  $n$  distinct numbers must determine, for each of the other  $n - 1$  numbers, whether they are smaller or larger than the median. This must be accomplished either by a direct comparison with the median or indirectly through a chain of comparisons.

Again the adversary will use partial orientations of  $K_n$  to force any algorithm to perform many comparisons before the median is found. Observe that if we already know that  $x < m < y$ , then the information  $x < y$  does not help us find the median. We call such a comparison (of  $x$  and  $y$ ) a *useless* comparison. The aim of the adversary is to convey as much useless information as possible while at the same time preserving consistency. This is achieved by maintaining the acyclic orientation  $D$  such that as many queries as possible can be answered by useless information while still preserving consistency. Below we describe the adversary's strategy.

At any time we maintain three sets  $S, L, U$  which form a partitioning of  $V$ . These are defined by

$$S = \{x \in V : d_D(x) > 0 \text{ and } \mathcal{L}_D(x) < (n + 1)/2\},$$

$$L = \{x \in V : d_D(x) > 0 \text{ and } \mathcal{L}_D(x) > (n + 1)/2\},$$

$$U = V - S - L.$$



Note that at the beginning we have  $U = V$  as no edges have been oriented yet.

The adversary answers the query  $Q(x, y)$  according to the rules below with the additional restriction that we must always have  $|S|, |L| \leq (n-1)/2$ . If we reach a point where  $|S| = (n-1)/2$  and  $|L| < (n-1)/2$ , then all remaining vertices  $z \in U$  except one are placed in  $L$  and are moved to the end of  $\mathcal{L}(D)$ . Similarly if we reach a point where  $|L| = (n-1)/2$  and  $|S| < (n-1)/2$ , then all remaining vertices  $z \in U$  except one are placed in  $S$  and are moved to the beginning of  $\mathcal{L}(D)$ . When only one vertex  $z$  is left in  $U$  the adversary moves this element to position  $(n+1)/2$  in the current acyclic ordering. This is the element which will finally become the median (denoted  $m$  below). As of this point  $m$  will remain on position number  $(n+1)/2$  in the acyclic ordering until the algorithm terminates. Note also that the element  $m$  is considered undefined until  $U$  becomes empty.

In order to be able to distinguish comparisons which are known to be useless from (potentially) useful comparisons we denote the first ones by red arcs and the others by black arcs.

1. If  $x, y \in U$ , then  $S := S + x$ ,  $L := L + y$ , add a red arc  $x \rightarrow y$ , move  $x$  to position 1 and  $y$  to position  $n$  in the new acyclic ordering.
2. If  $x \in L$  and  $y \in U$ , then  $S := S + y$ , add a red arc  $y \rightarrow x$  move  $y$  to position 1 in the new acyclic ordering.
3. If  $x \in S$  and  $y \in U$ , then  $L := L + y$ , add a red arc  $x \rightarrow y$  move  $y$  to position  $n$  in the new acyclic ordering.
4. If  $y \in L$  and  $x \in U$ , then  $S := S + x$ , add a red arc  $x \rightarrow y$  move  $x$  to position 1 in the new acyclic ordering.
5. If  $y \in S$  and  $x \in U$ , then  $L := L + x$ , add a red arc  $y \rightarrow x$  move  $x$  to position  $n$  in the new acyclic ordering.
6. If  $x \in S, y \in L$ , then add a red arc  $x \rightarrow y$ .
7. If  $x \in L, y \in S$ , then add a red arc  $y \rightarrow x$ .
8. If  $\mathcal{L}_D(x) < \mathcal{L}_D(y)$  and  $x, y \in S + m$ , or  $x, y \in L + m$ , then add a black arc  $x \rightarrow y$ .




Figure 2: A possible structure of  $D$  at the time when the median  $m$  has been determined (that is it is the only possible candidate left). The dotted arcs are red arcs and the others are black arcs. To ease visibility only a minimal set of black arcs has been included

9. If  $\mathcal{L}_D(y) < \mathcal{L}_D(x)$  and  $x, y \in S + m$ , or  $x, y \in L + m$ , then add a black arc  $y \rightarrow x$ .

Note that in the last two rules we only allow  $x$  or  $y$  to equal  $m$  if we have  $U = \emptyset$  at this point.

**Lemma 6.1** *If  $D$  is acyclic and a new arc between  $x$  and  $y$  is oriented according to the rules above and added to  $D$  then the resulting digraph is acyclic and every red arc  $u \rightarrow v$  satisfies  $\mathcal{L}_D(u) < (n + 1)/2 < \mathcal{L}_D(v)$ .*

**Proof:** Exercise 6. ◇

Figure 2 shows an example of the current digraph  $D$  at the time when the median  $m$  has been found.

**Theorem 6.2** *By applying the above strategy the adversary will force every comparison based algorithm  $\mathcal{C}$  to perform at least  $3n/2 - 3/2$  comparisons before the median is determined among  $n$  distinct numbers.*

**Proof:** Note that if  $D$  is an acyclic digraph with an acyclic ordering  $\mathcal{L}(D)$ , where  $\mathcal{L}_D(x) < \mathcal{L}_D(y)$  and there is no directed path from  $x$  to  $y$  in

$D$ , then  $\mathcal{L}(D)$  can be changed into a new acyclic ordering in which  $y$  appears before  $x$  (Exercise 7). This implies that when  $\mathcal{C}$  terminates there must exist paths from  $x$  to  $m$  and from  $m$  to  $y$  in  $D$  for arbitrary  $x, y \in V(D)$  such that  $\mathcal{L}_D(x) < (n+1)/2 < \mathcal{L}_D(y)$ . No red arc can be on such a path as every red arc  $u \rightarrow v$  satisfies  $\mathcal{L}_D(u) < (n+1)/2 < \mathcal{L}_D(v)$  and all arcs of  $D$  go forward with respect to  $\mathcal{L}(D)$ . This implies that every vertex  $x$  with  $\mathcal{L}_D(x) < (n+1)/2$  must have at least one black arc out of it and every vertex  $y$  with  $\mathcal{L}_D(y) > (n+1)/2$  has at least one black arc into it. Thus there are at least  $n-1$  black arcs since all arcs  $u \rightarrow v$  with  $\mathcal{L}_D(u) < (n+1)/2 < \mathcal{L}_D(v)$  are red. At the same time at least  $(n-1)/2$  red arcs can be added by the adversary since the rules 1.-5. may be applied at least  $(n-1)/2$  times (until either  $|S| = (n-1)/2$ , or  $|L| = (n-1)/2$ ). This shows that in total  $\mathcal{C}$  must perform at least  $n-1 + (n-1)/2 = 3n/2 - 3/2$  comparisons.  $\diamond$

## 7 Exercises

1. Let  $T$  be an out-arborescence with root  $r$ . Prove that  $T$  contains a directed path from  $r$  to  $x$  for every vertex  $x$  of  $T$ . Hint: apply induction.
2. An *isomorphism* between two digraphs  $D = (V, A)$  and  $H = (V', A')$  is a mapping  $f : V \rightarrow V'$  which is 1-1 and onto and which preserves arcs, that is, if  $u \rightarrow v \in A$  then  $f(u) \rightarrow f(v) \in A'$ . Prove that any two transitive tournaments on the same number of vertices are isomorphic. Then derive a linear algorithm for finding such an isomorphism given two transitive tournaments on  $n$  vertices.
3. Prove Lemma 3.1
4. Give an ordered set of comparisons which an algorithm  $\mathcal{B}$  (for finding the second smallest element among  $n$ ) could perform and which would result in the adversary constructing the digraph in Figure 1.
5. Argue that if  $\mathcal{B}$  is the tournament method, described on page 129 in the first part of these notes (Baase) then the adversary will generate precisely  $\lceil \log n \rceil - 1$  red arcs.
6. Prove Lemma 6.1.

7. Suppose that  $D$  is an acyclic digraph with two vertices  $x, y$  such that  $D$  contains no path from  $x$  to  $y$ . Let  $\mathcal{L}(D)$  be an acyclic ordering of  $D$  such that  $\mathcal{L}_D(x) < \mathcal{L}_D(y)$ . Prove that there exists another acyclic ordering of  $D$  in which  $y$  is before  $x$ . Hint: it is not always enough just to interchange the positions of  $x$  and  $y$ .
8. Find an acyclic ordering of the digraph Figure 2. Then describe a sequence of comparisons that will result in the adversary constructing precisely this digraph  $D$ .