

*An adversary means opposition and competition,  
but not having an adversary means grief and loneliness.*

— Zhuangzi (Chuang-tsu) c. 300 BC

*It is possible that the operator could be hit by an asteroid and your \$20 could fall off his cardboard box and land on the ground, and while you were picking it up, \$5 could blow into your hand. You therefore could win \$5 by a simple twist of fate.*

— Penn Jillette, explaining how to win at Three-Card Monte (1999)

## 20 Adversary Arguments

### 20.1 Three-Card Monte

Until Times Square was sanitized into TimesSquareLand™ by Mayor Rudy Guiliani, you could often find dealers stealing tourists' money using a game called 'Three Card Monte' or 'Spot the Lady'. The dealer has three cards, say the Queen of Hearts and the two and three of clubs. The dealer shuffles the cards face down on a table (usually slowly enough that you can follow the Queen), and then asks the tourist to bet on which card is the Queen. In principle, the tourist's odds of winning are at least one in three.

In practice, however, the tourist *never*<sup>1</sup> wins, because the dealer cheats. There are actually *four* cards; before he even starts shuffling the cards, the dealer palms the queen or sticks it up his sleeve. No matter what card the tourist bets on, the dealer turns over a black card. If the tourist gives up, the dealer slides the queen under one of the cards and turns it over, showing the tourist 'where the queen was all along'. If the dealer is really good, the tourist won't see the dealer changing the cards and will think maybe the queen *was* there all along and he just wasn't smart enough to figure that out. As long as the dealer doesn't reveal all the black cards at once, the tourist has no way to prove that the dealer cheated!<sup>2</sup>

### 20.2 $n$ -Card Monte

Now let's consider a similar game, but with an algorithm acting as the tourist and with bits instead of cards. Suppose we have an array of  $n$  bits and we want to determine if any of them is a 1. Obviously we can figure this out by just looking at every bit, but can we do better? Is there maybe some complicated tricky algorithm to answer the question "Any ones?" without looking at every bit? Well, of course not, but how do we prove it?

The simplest proof technique is called an *adversary* argument. The idea is that an all-powerful malicious adversary (the dealer) *pretends* to choose an input for the algorithm (the tourist). When the algorithm wants looks at a bit (a card), the adversary sets that bit to whatever value will make the algorithm do the most work. If the algorithm does not look at enough bits before terminating, then there will be several different inputs, each consistent with the bits already seen, the should result in different outputs. Whatever the algorithm outputs, the adversary can 'reveal' an input that is has all the examined bits but contradicts the algorithm's output, and then claim that that was the input that he was using all along. Since the only information the algorithm has is the set of bits it examined, the algorithm cannot distinguish between a malicious adversary and an honest user who actually chooses an input in advance and answers all queries truthfully.

For the  $n$ -card monte problem, the adversary originally pretends that the input array is all zeros—whenever the algorithm looks at a bit, it sees a 0. Now suppose the algorithms stops before looking at all

<sup>1</sup>What, never? No, **NEVER**. Anyone you see winning at Three Card Monte is a shill.

<sup>2</sup>Even if the dealer isn't very good, he cheats anyway. The shills will protect him from any angry tourists who realize they've been ripped off, and shake down any tourist who refuses to pay. You *cannot* win this game.



Similarly, if the algorithm doesn't look at every bit to the left of  $\ell$ , the adversary could replace some unexamined bit with a zero. Finally, if there are any unexamined bits between  $\ell$  and  $r$ , there must be at least two such bits (since  $r - \ell$  is always even) and the adversary can put a 01 in the gap.

In general, we say that a bit pattern is *evasive* if we have to look at every bit to decide if a string of  $n$  bits contains the pattern. So the pattern 1 is evasive for all  $n$ , and the pattern 01 is evasive if and only if  $n$  is even. It turns out that the *only* patterns that are evasive for *all* values of  $n$  are the one-bit patterns 0 and 1.

## 20.4 Evasive Graph Properties

Another class of problems for which adversary arguments give good lower bounds is graph problems where the graph is represented by an adjacency matrix, rather than an adjacency list. Recall that the adjacency matrix of an undirected  $n$ -vertex graph  $G = (V, E)$  is an  $n \times n$  matrix  $A$ , where  $A[i, j] = [(i, j) \in E]$ . We are interested in deciding whether an undirected graph has or does not have a certain *property*. For example, is the input graph connected? Acyclic? Planar? Complete? A tree? We call a graph property *evasive* if we have to look at all  $\binom{n}{2}$  entries in the adjacency matrix to decide whether a graph has that property.

An obvious example of an evasive graph property is *emptiness*: Does the graph have any edges at all? We can show that emptiness is evasive using the following simple adversary strategy. The adversary maintains *two* graphs  $E$  and  $G$ .  $E$  is just the empty graph with  $n$  vertices. Initially  $G$  is the complete graph on  $n$  vertices. Whenever the algorithm asks about an edge, the adversary removes that edge from  $G$  (unless it's already gone) and answers 'no'. If the algorithm terminates without examining every edge, then  $G$  is not empty. Since both  $G$  and  $E$  are consistent with all the adversary's answers, the algorithm must give the wrong answer for one of the two graphs.

## 20.5 Connectedness Is Evasive

Now let me give a more complicated example, *connectedness*. Once again, the adversary maintains two graphs,  $Y$  and  $M$  ('yes' and 'maybe').  $Y$  contains all the edges that the algorithm knows are definitely in the input graph.  $M$  contains all the edges that the algorithm thinks *might* be in the input graph, or in other words, all the edges of  $Y$  plus all the unexamined edges. Initially,  $Y$  is empty and  $M$  is complete.

Here's the strategy that adversary follows when the adversary asks whether the input graph contains the edge  $e$ . I'll assume that whenever an algorithm examines an edge, it's in  $M$  but not in  $Y$ ; in other words, algorithms never ask about the same edge more than once.

```

HIDECONNECTEDNESS( $e$ ):
  if  $M \setminus \{e\}$  is connected
    remove  $(i, j)$  from  $M$ 
    return 0
  else
    add  $e$  to  $Y$ 
    return 1

```

Notice that the graphs  $Y$  and  $M$  are both consistent with the adversary's answers at all times. The adversary strategy maintains a few other simple invariants.

- **$Y$  is a subgraph of  $M$ .** This is obvious.
- **$M$  is connected.** This is also obvious.

- **If  $M$  has a cycle, none of its edges are in  $Y$ .** If  $M$  has a cycle, then deleting any edge in that cycle leaves  $M$  connected.
- **$Y$  is acyclic.** This follows directly from the previous invariant.
- **If  $Y \neq M$ , then  $Y$  is disconnected.** The only connected acyclic graph is a tree. Suppose  $Y$  is a tree and some edge  $e$  is in  $M$  but not in  $Y$ . Then there is a cycle in  $M$  that contains  $e$ , all of whose other edges are in  $Y$ . This violated our third invariant.

We can also think about the adversary strategy in terms of minimum spanning trees. Recall the anti-Kruskal algorithm for computing the *maximum* spanning tree of a graph: Consider the edges one at a time in increasing order of length. If removing an edge would disconnect the graph, declare it part of the spanning tree (by adding it to  $Y$ ); otherwise, throw it away (by removing it from  $M$ ). If the algorithm examines all  $\binom{n}{2}$  possible edges, then  $Y$  and  $M$  are both equal to the maximum spanning tree of the complete  $n$ -vertex graph, where the weight of an edge is the time when the algorithm asked about it.

Now, if an algorithm terminates before examining all  $\binom{n}{2}$  edges, then there is at least one edge in  $M$  that is not in  $Y$ . Since the algorithm cannot distinguish between  $M$  and  $Y$ , even though  $M$  is connected and  $Y$  is not, the algorithm cannot possibly give the correct output for both graphs. Thus, in order to be correct, any algorithm must examine every edge—*Connectedness is evasive!*

## 20.6 An Evasive Conjecture

A graph property is *nontrivial* if there is at least one graph with the property and at least one graph without the property. (The only trivial properties are ‘Yes’ and ‘No’.) A graph property is *monotone* if it is closed under taking subgraphs — if  $G$  has the property, then any subgraph of  $G$  has the property. For example, emptiness, planarity, acyclicity, and *non*-connectedness are monotone. The properties of being a tree and of having a vertex of degree 3 are not monotone.

**Conjecture 1 (Aanderra, Karp, and Rosenberg).** *Every nontrivial monotone property of  $n$ -vertex graphs is evasive.*

The Aanderra-Karp-Rosenberg conjecture has been proven when  $n = p^e$  for some prime  $p$  and positive integer exponent  $e$ —the proof uses some interesting results from algebraic topology<sup>3</sup>—but it is still open for other values of  $n$ .

There are non-trivial non-evasive graph properties, but all known examples are non-monotone. One such property—‘scorpionhood’—is described in an exercise at the end of this lecture note.

## 20.7 Finding the Minimum and Maximum

Last time, we saw an adversary argument that finding the largest element of an unsorted set of  $n$  numbers requires at least  $n - 1$  comparisons. Let’s consider the complexity of finding the largest *and* smallest elements. More formally:

Given a sequence  $X = \langle x_1, x_2, \dots, x_n \rangle$  of  $n$  distinct numbers, find indices  $i$  and  $j$  such that  $x_i = \min X$  and  $x_j = \max X$ .

<sup>3</sup>Let  $\Delta$  be a contractible simplicial complex whose automorphism group  $\text{Aut}(\Delta)$  is vertex-transitive, and let  $\Gamma$  be a vertex-transitive subgroup of  $\text{Aut}(\Delta)$ . If there are normal subgroups  $\Gamma_1 \triangleleft \Gamma_2 \triangleleft \Gamma$  such that  $|\Gamma_1| = p^\alpha$  for some prime  $p$  and integer  $\alpha$ ,  $|\Gamma/\Gamma_2| = q^\beta$  for some prime  $q$  and integer  $\beta$ , and  $\Gamma_2/\Gamma_1$  is cyclic, then  $\Delta$  is a simplex.

No, this will not be on the final exam.

How many comparisons do we need to solve this problem? An upper bound of  $2n - 3$  is easy: find the minimum in  $n - 1$  comparisons, and then find the maximum of everything else in  $n - 2$  comparisons. Similarly, a lower bound of  $n - 1$  is easy, since any algorithm that finds the min and the max certainly finds the max.

We can improve both the upper and the lower bound to  $\lceil 3n/2 \rceil - 2$ . The upper bound is established by the following algorithm. Compare all  $\lfloor n/2 \rfloor$  consecutive pairs of elements  $x_{2i-1}$  and  $x_{2i}$ , and put the smaller element into a set  $S$  and the larger element into a set  $L$ . If  $n$  is odd, put  $x_n$  into both  $L$  and  $S$ . Then find the smallest element of  $S$  and the largest element of  $L$ . The total number of comparisons is at most

$$\underbrace{\lfloor \frac{n}{2} \rfloor}_{\text{build } S \text{ and } L} + \underbrace{\lfloor \frac{n}{2} \rfloor - 1}_{\text{compute min } S} + \underbrace{\lfloor \frac{n}{2} \rfloor - 1}_{\text{compute max } L} = \lceil \frac{3n}{2} \rceil - 2.$$

For the lower bound, we use an adversary argument. The adversary marks each element  $+$  if it *might* be the maximum element, and  $-$  if it *might* be the minimum element. Initially, the adversary puts both marks  $+$  and  $-$  on every element. If the algorithm compares two double-marked elements, then the adversary declares one smaller, removes the  $+$  mark from the smaller element, and removes the  $-$  mark from the larger one. In every other case, the adversary can answer so that at most one mark needs to be removed. For example, if the algorithm compares a double marked element against one labeled  $-$ , the adversary says the one labeled  $-$  is smaller and removes the  $-$  mark from the other. If the algorithm compares to  $+$ 's, the adversary must unmark one of the two.

Initially, there are  $2n$  marks. At the end, in order to be correct, exactly one item must be marked  $+$  and exactly one other must be marked  $-$ , since the adversary can make any  $+$  the maximum and any  $-$  the minimum. Thus, the algorithm must force the adversary to remove  $2n - 2$  marks. At most  $\lfloor n/2 \rfloor$  comparisons remove two marks; every other comparison removes at most one mark. Thus, the adversary strategy forces any algorithm to perform at least  $2n - 2 - \lfloor n/2 \rfloor = \lceil 3n/2 \rceil - 2$  comparisons.

## 20.8 Finding the Median

Finally, let's consider the *median* problem: Given an unsorted array  $X$  of  $n$  numbers, find its  $n/2$ th largest entry. (I'll assume that  $n$  is even to eliminate pesky floors and ceilings.) More formally:

Given a sequence  $\langle x_1, x_2, \dots, x_n \rangle$  of  $n$  distinct numbers, find the index  $m$  such that  $x_m$  is the  $n/2$ th largest element in the sequence.

To prove a lower bound for this problem, we can use a combination of information theory and two adversary arguments. We use one adversary argument to prove the following simple lemma:

**Lemma 1.** *Any comparison tree that correctly finds the median element also identifies the elements smaller than the median and larger than the median.*

**Proof:** Suppose we reach a leaf of a decision tree that chooses the median element  $x_m$ , and there is still some element  $x_i$  that isn't known to be larger or smaller than  $x_m$ . In other words, we cannot decide based on the comparisons that we've already performed whether  $x_i < x_m$  or  $x_i > x_m$ . Then in particular no element is known to lie between  $x_i$  and  $x_m$ . This means that there must be an input that is consistent with the comparisons we've performed, in which  $x_i$  and  $x_m$  are adjacent in sorted order. But then we can swap  $x_i$  and  $x_m$ , without changing the result of any comparison, and obtain a different consistent input in which  $x_i$  is the median, not  $x_m$ . Our decision tree gives the wrong answer for this 'swapped' input.  $\square$

This lemma lets us rephrase the median-finding problem yet again.

Given a sequence  $X = \langle x_1, x_2, \dots, x_n \rangle$  of  $n$  distinct numbers, find the indices of its  $n/2 - 1$  largest elements  $L$  and its  $n/2$ th largest element  $x_m$ .

Now suppose a ‘little birdie’ tells us the set  $L$  of elements larger than the median. This information fixes the outcomes of certain comparisons—any item in  $L$  is bigger than any element not in  $L$ —so we can ‘prune’ those comparisons from the comparison tree. The pruned tree finds the largest element of  $X \setminus L$  (the median of  $X$ ), and thus must have depth at least  $n/2 - 1$ . In fact, the adversary argument in the last lecture implies that *every* leaf in the pruned tree must have depth at least  $n/2 - 1$ , so the pruned tree has at least  $2^{n/2-1}$  leaves.

There are  $\binom{n}{n/2-1} \approx 2^n / \sqrt{n/2}$  possible choices for the set  $L$ . Every leaf in the original comparison tree is also a leaf in exactly one of the  $\binom{n}{n/2-1}$  pruned trees, so the original comparison tree must have at least  $\binom{n}{n/2-1} 2^{n/2-1} \approx 2^{3n/2} / \sqrt{n/2}$  leaves. Thus, any comparison tree that finds the median must have depth at least

$$\left\lceil \frac{n}{2} - 1 + \lg \binom{n}{n/2-1} \right\rceil = \frac{3n}{2} - O(\log n).$$

A more complicated adversary argument (also involving pruning the comparison tree with little birdies) improves this lower bound to  $2n - o(n)$ .

A similar argument implies that at least  $n - k + \lceil \lg \binom{n}{k-1} \rceil = \Omega((n - k) + k \log(n/k))$  comparisons are required to find the  $k$ th largest element in an  $n$ -element set. This bound is tight up to constant factors for all  $k \leq n/2$ ; there is an algorithm that uses at most  $O(n + k \log(n/k))$  comparisons. Moreover, this lower bound is *exactly* tight when  $k = 1$  or  $k = 2$ . In fact, these are the *only* values of  $k \leq n/2$  for which the exact complexity of the selection problem is known. Even the case  $k = 3$  is still open!

## Exercises

1. (a) Let  $X$  be a set containing an odd number of  $n$ -bit strings. Prove that any algorithm that decides whether a given  $n$ -bit string is an element of  $X$  *must* examine every bit of the input string in the worst case.
  - (b) Give a one-line proof that the bit pattern 01 is evasive for all even  $n$ .
  - (c) Prove that the bit pattern 11 is evasive if and only if  $n \bmod 3 = 1$ .
  - \* (d) Prove that the bit pattern 111 is evasive if and only if  $n \bmod 4 = 0$  or 3.
2. Suppose we are given the adjacency matrix of a *directed* graph  $G$  with  $n$  vertices. Describe an algorithm that determines whether  $G$  has a *sink* by probing only  $O(n)$  bits in the input matrix. A sink is a vertex that has an incoming edge from every other vertex, but no outgoing edges.
- \*3. A *scorpion* is an undirected graph with three special vertices: the *sting*, the *tail*, and the *body*. The sting is connected only to the tail; the tail is connected only to the sting and the body; and the body is connected to every vertex except the sting. The rest of the vertices (the head, eyes, legs, antennae, teeth, gills, flippers, wheels, etc.) can be connected arbitrarily. Describe an algorithm that determines whether a given  $n$ -vertex graph is a scorpion by probing only  $O(n)$  entries in the adjacency matrix.

4. Prove using an adversary argument that acyclicity is an evasive graph property. [*Hint: Kruskal.*]
5. Prove that finding the second largest element in an  $n$ -element array requires *exactly*  $n - 2 + \lceil \lg n \rceil$  comparisons in the worst case. Prove the upper bound by describing and analyzing an algorithm; prove the lower bound using an adversary argument.
6. Let  $T$  be a perfect ternary tree where every leaf has depth  $\ell$ . Suppose each of the  $3^\ell$  leaves of  $T$  is labeled with a bit, either 0 or 1, and each internal node is labeled with a bit that agrees with the *majority* of its children.
  - (a) Prove that any deterministic algorithm that determines the label of the root must examine all  $3^\ell$  leaf bits in the worst case.
  - (b) Describe and analyze a *randomized* algorithm that determines the root label, such that the expected number of leaves examined is  $o(3^\ell)$ . (You may want to review the notes on randomized algorithms.)
- \*7. UIUC has just finished constructing the new Reingold Building, the tallest dormitory on campus. In order to determine how much insurance to buy, the university administration needs to determine the highest safe floor in the building. A floor is considered *safe* if ~~a drunk student~~ **an egg** can fall from a window on that floor and land without breaking; if the egg breaks, the floor is considered *unsafe*. Any floor that is higher than an unsafe floor is also considered unsafe. The only way to determine whether a floor is safe is to drop an egg from a window on that floor.

You would like to find the lowest unsafe floor  $L$  by performing as few tests as possible; unfortunately, you have only a very limited supply of eggs.

  - (a) Prove that if you have only one egg, you can find the lowest unsafe floor with  $L$  tests. [*Hint: Yes, this is trivial.*]
  - (b) Prove that if you have only one egg, you must perform at least  $L$  tests in the worst case. In other words, prove that your algorithm from part (a) is optimal. [*Hint: Use an adversary argument.*]
  - (c) Describe an algorithm to find the lowest unsafe floor using *two* eggs and only  $O(\sqrt{L})$  tests. [*Hint: Ideally, each egg should be dropped the same number of times. How many floors can you test with  $n$  drops?*]
  - (d) Prove that if you start with two eggs, you must perform at least  $\Omega(\sqrt{L})$  tests in the worst case. In other words, prove that your algorithm from part (c) is optimal.
  - \**(e)* Describe an algorithm to find the lowest unsafe floor using  $k$  eggs, using as few tests as possible, and prove your algorithm is optimal for all values of  $k$ .