Lecture Notes

This week we return to the problem of finding the cardinality of a finite set. (No more infinite sets!) We saw one way to determine the number of elements in a finite set S last week. Namely, if there is a bijection between S and some set T of known cardinality, then we can conclude that S has the same number of elements as T. In this lecture, we will see many more ways to determine the cardinality of a finite set.

1 The Cardinality of a Union of Sets

Suppose that we know the sizes of some sets $A_1, A_2, \ldots A_n$. How many elements are in the union? If the sets are disjoint, then the size of the union is given by the simple *Sum Rule*. If the sets are not necessarily disjoint, then the size is given by the more general, but more complicated *Inclusion-Exclusion Principle*.

1.1 The Sum Rule

Theorem 1.1 (Sum Rule) If $A_1, A_2, \ldots A_n$ are disjoint sets, then:

$$|A_1 \cup A_2 \cup \ldots \cup A_n| = |A_1| + |A_2| + \ldots + |A_n|$$

The Sum Rule says that the number of elements in a union of disjoint sets is equal to the sum of the sizes of all the sets. The Sum Rule can be proven by induction on the number of sets, but the proof is omitted here.

As an example of the Sum Rule, suppose that MIT graduates 60 majors in math, 200 majors in EECS, and 40 majors in physics. How many students graduate from MIT in these three departments? Let A_1 be the set of math majors, A_2 be the set of EECS majors, and A_3 be the set of physics majors. The set of graduating students in these three departments is $A_1 \cup A_2 \cup A_3$. Assume for now that these sets are disjoint; that is, there are no double or triple majors. Then we can apply the Sum Rule to determine the total number of graduating students.

$$\begin{aligned} |A_1 \cup A_2 \cup A_3| &= |A_1| + |A_2| + |A_3| \\ &= 60 + 200 + 40 \\ &= 300 \end{aligned}$$

1.2 Inclusion-Exclusion Principle (special cases)

The Sum Rule gives the cardinality of a union of disjoint sets. The Inclusion-Exclusion Principle gives the cardinality of a union of sets that may intersect. The Inclusion-Exclusion Principle for n sets is messy to write down, so we'll start with the simple special cases $n = 2$ and $n = 3$.

Theorem 1.2 (Inclusion-Exclusion Principle for 2 sets) Let A and B be sets, not necessarily disjoint.

$$|A \cup B| = |A| + |B| - |A \cap B|$$

Proof. Items in the intersection of A and B are counted twice in $|A| + |B|$. Therefore, by subtracting $|A \cap B|$, every element is counted once overall. ∎

Theorem 1.3 (Inclusion-Exclusion Principle for 3 sets) Let A, B, and C be sets, not necessarily disjoint.

$$\begin{aligned} |A \cup B \cup C| &= |A| + |B| + |C| \\ &\quad - |A \cap B| - |A \cap C| - |B \cap C| \\ &\quad + |A \cap B \cap C| \end{aligned}$$

Though this formula contains many terms, the general pattern is easy to remember: *add* the sizes of individual sets (first line), *subtract* intersections of pairs of sets (second line), and *add* the intersection of all three sets (third line).

Proof. Items contained in just one of the sets A, B, or C are counted once on the first line. Since these items are not contained in any intersection of sets, they are not subtracted away on the second line or counted again on the third line. In total, these items are counted just once.

Items contained in two of the sets A, B, and C are counted twice on the first line, subtracted away once on the second line, and not counted again on the third line. Again, in total, these items are counted just once.

Items contained in all three sets A, B, and C are counted three times on the first line, subtracted away three times on the second line, and added back once on the third line. Since $3 - 3 + 1 = 1$, these items are also counted just once overall. ∎

The name "Inclusion-Exclusion" comes from the way items are counted, subtracted away, counted again, etc.

Earlier we applied the Sum Rule to count MIT graduates in math, EECS, and physics, assuming that there were no double or triple majors. Using Inclusion-Exclusion, we can count the number of graduates even if some people have multiple majors. (After all, this is MIT, where some students even have 4 majors...) Suppose the numbers are as follows:

60	math majors	$(\lvert A\rvert = 60)$
200	EECS majors	$(\lvert B\rvert = 200)$
400	physics majors	$(\lvert C\rvert = 40)$
10	math/EECS double majors	$(\lvert A\rvert \cap \lvert B\rvert = 10)$
4	math/physics double majors	$(\lvert A\rvert \cap \lvert C\rvert = 4)$
15	EECS/physics double majors	$(\lvert B\rvert \cap \lvert C\rvert = 15)$
2	triple majors	$(\lvert A\rvert \cap \lvert B\rvert \cap \lvert C\rvert = 2)$

We can compute the total number of graduates, $|A \cup B \cup C|$, by plugging numbers into the Inclusion-Exclusion formula:

$$\begin{aligned}
|A \cup B \cup C| &= |A| + |B| + |C| - |A \cap B| - |A \cap C| - |B \cap C| + |A \cap B \cap C| \\
&= 60 + 200 + 40 - 10 - 4 - 15 + 2 \\
&= 273
\end{aligned}$$

1.3 Inclusion-Exclusion Principle (general case)

Here is the nasty Inclusion-Exclusion formula for n sets. The formula is given both in words and in symbols. The word version is the one to remember, but look over the symbolic version to make sure that you really know what the theorem says. The proof is by induction on n, but omitted here.

Theorem 1.4 (Inclusion-Exclusion for n sets) Let $A_1, A_2, \ldots, A_n$ be sets, not necessarily disjoint. The cardinality of the union is computed as follows:

add the sizes of all individual sets
subtract the sizes of all two-way intersection
add the sizes of all three-way intersections
subtract the sizes of all four-way intersections
add the sizes of all five-way intersections
etc.

Restated in symbols, the cardinality of the union is:

$$
\begin{aligned}
|A_1 \cup A_2 \cup \ldots \cup A_n| &= \sum^{1 \le i \le n} |A_i| \\
&\quad - \sum^{1 \le i < j \le n} |A_i \cap A_j| \\
&\quad + \sum^{1 \le i < j < k \le n} |A_i \cap A_j \cap A_k| \\
&\quad \ldots \\
&\quad + (-1)^{n+1} |A_1 \cap A_2 \cap \ldots \cap A_n| \\
&= \sum_{k=1}^{n} (-1)^{k+1} \left(\sum^{S \subseteq [1,n], |S|=k} |\bigcap_{i \in S} A_i| \right)
\end{aligned}
$$

1.4 Counting Primes

How many of the numbers $1, 2, \ldots, 100$ are prime? One way to answer this question is to test each number up to 100 for primality and keep a count. This requires considerable effort. (Is 57 prime? How about 67?)

Another approach is to use the Inclusion-Exclusion Principle. This requires one trick: to determine the number of primes, we will first count the number of *non-primes*. We can then find the number of primes by subtraction. We will use this trick of "counting the complement" several times in coming weeks.

Reducing the Problem to the Cardinality of a Union

The set of non-primes in the range $1, \ldots, 100$ consists of the set C of composite numbers in this range $(4, 6, 8, 9, \ldots, 99, 100)$ and the number 1, which is neither prime nor composite. The main job is to determine the size of the set C of composite numbers. For this purpose, define:

$$A_p = \{x \mid 1 \le x \le 100, p \mid x, x \ne p\}$$

In words, A_p is the set of numbers in the range $1, \ldots, 100$ that are divisible by p, but not equal to p. For example, $A_2 = \{4, 6, 8, \ldots, 100\}$.

Claim 1.5 $C = A_2 \cup A_3 \cup A_5 \cup A_7$

The claim explains the point of these funny A_p sets: we can write the set C of composite numbers as a union of them. We can compute the cardinality of the union using Inclusion-Exclusion, and this will tell us the number of composite numbers in the range $1, \ldots, 100$.

Proof. We prove the two sets equal by showing that each contains the other.

First, we show that $A_2 \cup A_3 \cup A_5 \cup A_7 \subseteq C$. Let n be an element of $A_2 \cup A_3 \cup A_5 \cup A_7$. Then $n \in A_p$ for $p = 2, 3, 5$ or 7. This implies that n is in the range $1, \ldots, 100$, n is divisible by p, and n is not equal to p. This implies that n is a composite in the range $1, \ldots, 100$, and so $n \in C$.

Second, we show that $C \subseteq A_2 \cup A_3 \cup A_5 \cup A_7$. Let n be an element of C. Then n is a composite number in the range $1, \ldots, 100$. This means that n has at least two prime factors p and q. One of these must be 2, 3, 5, or 7. (Otherwise, both p and q are at least 11, and so $n \geq pq \geq 11 \cdot 11 = 121$, a contradiction.) This implies that n is an element of A_2, A_3, A_5, or A_7, and so $n \in A_2 \cup A_3 \cup A_5 \cup A_7$. ■

Corollary 1.6 $|C| = |A_2 \cup A_3 \cup A_5 \cup A_7|$

Computing the Cardinality of the Union

We have reduced the problem of counting primes to a problem about the cardinality of a union of sets. Specifically, we must evaluate $|A_2 \cup A_3 \cup A_5 \cup A_7|$. This will give us the number of composites in the range $1, \ldots, 100$, and from this we can figure out the number of primes. As a stepping stone, we can compute the cardinality of each set A_p:

$$|A_p| = \left\lfloor \frac{100}{p} \right\rfloor - 1$$

The first term, $\lfloor \frac{100}{p} \rfloor$, is the number of values in the range $1, \ldots, 100$ that are divisible by p. The second term, -1, arises because we defined A_p to exclude p itself. This formula gives:

$$\begin{aligned}
|A_2| &= \lfloor \tfrac{100}{2} \rfloor - 1 = & 49 \\
|A_3| &= \lfloor \tfrac{100}{3} \rfloor - 1 = & 32 \\
|A_5| &= \lfloor \tfrac{100}{5} \rfloor - 1 = & 19 \\
|A_7| &= \lfloor \tfrac{100}{7} \rfloor - 1 = & 13
\end{aligned}$$

Here is an erroneous way to compute the number of composites in the range 1 to 100.

$$\begin{aligned}
|C| &= |A_2 \cup A_3 \cup A_5 \cup A_7| \\
&= |A_2| + |A_3| + |A_5| + |A_7| \\
&= 49 + 32 + 19 + 13 \\
&= 113
\end{aligned}$$

In the first step, we applied the Sum Rule and the rest was substitution and simplification. The result is obviously wrong because there are only 100 numbers in the range 1 to 100! The problem is that the Sum Rule is inapplicable; the Sum Rule requires that the sets A_2, A_3, A_5, and A_7 be disjoint. However, 6 is in both A_2 and A_3, for example. Since the sets intersect, we must use the Inclusion-Exclusion Principle instead:

$$\begin{aligned}
|C| &= |A_2 \cup A_3 \cup A_5 \cup A_7| \\
&= |A_2| + |A_3| + |A_5| + |A_7| \\
&\quad - |A_2 \cap A_3| - |A_2 \cap A_5| - |A_2 \cap A_7| - |A_3 \cap A_5| - |A_3 \cap A_7| - |A_5 \cap A_7| \\
&\quad + |A_2 \cap A_3 \cap A_5| + |A_2 \cap A_3 \cap A_7| + |A_2 \cap A_5 \cap A_7| + |A_3 \cap A_5 \cap A_7| \\
&\quad - |A_2 \cap A_3 \cap A_5 \cap A_7|
\end{aligned}$$

There are a lot of terms here! Fortunately, all of them are easy to evaluate. For example, $|A_3 \cap A_7|$ is the number of multiples of $3 \cdot 7 = 21$ in the range 1 to 100, which is $\lfloor \frac{100}{21} \rfloor = 4$. (Note that there is no reason to subtract 1 as we did when evaluating $|A_p|$ above.) Substituting values for all of the terms above gives:

$$\begin{aligned}
|C| &= 49 + 32 + 19 + 13 \\
&\quad - 16 - 10 - 7 - 6 - 4 - 2 \\
&\quad + 3 + 2 + 1 + 0 \\
&\quad - 0 \\
&= 74
\end{aligned}$$

This calculation shows that there are 74 composite numbers in the range 1 to 100. Since the number 1 is neither composite nor prime, there are $100 - 74 - 1 = 25$ primes in this range.

In retrospect, checking each number from 1 to 100 for primality and keeping a count of primes might have been easier! However, the Inclusion-Exclusion approach used here is asymptotically faster as the range of numbers grows large. The naive strategy requires N runs of a primality test. The Inclusion-Exclusion approach seems to require summing an immense number of terms, but fewer than N of these are non-zero and the rest can be ignored.

2 Products of Sets

Recall the definition of the product of two sets:

$$A \times B = \{(a, b) \mid a \in A, b \in B\}$$

For example, if $A = \{x, y, z\}$ and $B = \{1, 2\}$, then $A \times B = \{(x,1),(y,1),(z,1),(x,2),(y,2),(z,2)\}$. In this case, A contains three items, B contains two items, and the product contains $3 \cdot 2 = 6$ items. More generally, the size of a product of sets is given by the Product Rule:

Theorem 2.1 (Product Rule) If $A_1, A_2, \ldots A_n$ are sets, then:

$$|A_1 \times A_2 \times \ldots \times A_n| = |A_1| \cdot |A_2| \cdot \ldots \cdot |A_n|$$

(The proof is by induction on n, but not given here.)

Here is another way to look at the product rule: the number of ways to pick one item from A_1, 1 item from A_2, ..., and 1 item from A_n is $\prod_{i=1}^{n} |A_i|$.

Four-Course Meals

As an example of the Product Rule, suppose that a menu lists 15 appetizers, 6 pastas, 10 main courses, and 4 desserts. How many four course-meals are possible? In this case, A_1 is the set of appetizers, A_2 is the sets of pastas, A_3 is the set of main courses, and A_4 is the set of desserts. The number of four-course meals is the number of ways of picking one item from each set. By the Product Rule this is:

$$\begin{aligned} |A_1 \times A_2 \times A_3 \times A_4| &= |A_1| \cdot |A_2| \cdot |A_3| \cdot |A_4| \\ &= 15 \cdot 6 \cdot 10 \cdot 4 \\ &= 3600 \end{aligned}$$

Binary Strings

How many n-bit binary strings are there? If we let $B = \{0, 1\}$, then the set of n-bit binary strings is:

$$\underbrace{B \times B \times \ldots \times B}_{n \text{ terms}}$$

By the Product Rule, the number of binary strings is $|B|^n = 2^n$.

Telephone Numbers

How many telephone numbers are there? There are actually two formats for telephone numbers, an old one and a new one. The formats can be defined in terms of three sets of digits:

$$\begin{aligned} A &= \{0, 1, \ldots, 9\} \\ B &= \{2, 3, \ldots, 9\} \\ C &= \{0, 1\} \end{aligned}$$

The old format is (BCA) $BBA - AAAA$ and the new format is (BAA) $BAA - AAAA$. This gives:

$$\begin{aligned}
\text{old numbers} &= |B| \cdot |C| \cdot |A| \;\cdot\; |B| \cdot |B| \cdot |A| \;\cdot\; |A| \cdot |A| \cdot |A| \cdot |A| \\
&= 2 \cdot 8^3 \cdot 10^6 \\
&= 1.024 \text{ billion} \\
\text{new numbers} &= |B| \cdot |A| \cdot |A| \;\cdot\; |B| \cdot |A| \cdot |A| \;\cdot\; |A| \cdot |A| \cdot |A| \cdot |A| \\
&= 8^2 \cdot 10^8 \\
&= 6.4 \text{ billion}
\end{aligned}$$

The reason for the new format is that there are more numbers. In particular, the number of area codes is increased from $8 \cdot 2 \cdot 10 = 160$ to $8 \cdot 10 \cdot 10 = 800$.

Passwords

Suppose that a password consists of 8 characters where each character is either a number or a lowercase letter. A password is *legal* if there is at least one number and at least one letter. How many legal passwords are there?

We can solve this problem by "counting the complement". We used this method earlier to find the number of primes in the range 1 to 100 by counting the composites in this range. In this case, we can find the number of legal passwords by counting the illegal passwords. An illegal password has either all numbers or all letters. By the Product Rule, the number of passwords with all numbers is 10^8, and the number of passwords with all letters is 26^8. Also by the Product Rule, the total number of passwords (both legal and illegal) is 36^8. Therefore, the total number of legal passwords is $36^8 - 26^8 - 10^8$.

3 Tree Diagrams

The Sum and Product Rules are both useful in counting the number of ways that something can be done. In more complicated problems, we need to combine these rules. In such cases, *tree diagrams* are helpful.

3.1 A Problem Getting Dressed

Suppose we have 3 blue shirts, 2 red shirts, and 1 green shirt. We also have 2 gray pants and 3 brown pants. How many outfits are possible? (Two pieces of clothing with the same color are still considered distinct; assume that they have slightly different shades.)

This problem is easier than it sounds. Let S be the set of shirts and P be the set of pants. We form an outfit by picking one shirt and one pair of pants. By the Product Rule, there are $6 \cdot 5 = 30$ outfits. All the color information is irrelevant.

However, suppose that gray pants "go with" only blue and green shirts and brown pants "go with" only blue and red shirts. How many *matching* outfits are there?

One approach to this problem is to count the complement set. That is, we count the number of matching outfits by counting the number of mismatching outfits. By the Product Rule, there are $2 \cdot 2 = 4$ mismatching gray-red outfits and $3 \cdot 1 = 3$ mismatching brown-green outfits. Therefore, there are $30 - 4 - 3 = 23$ matching outfits.

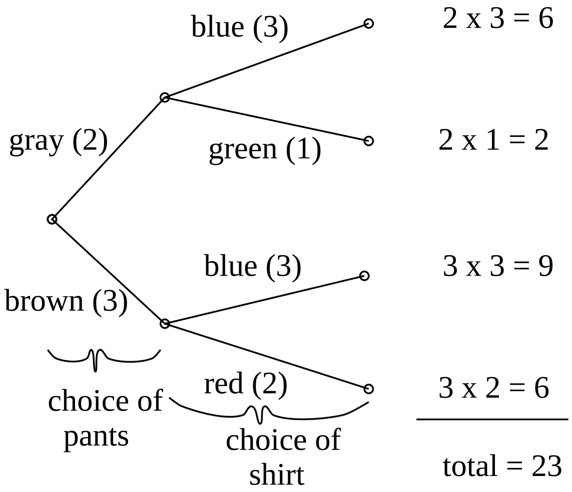

Figure 1: *This is a tree diagram for the matching outfits problem. Edge weights are shown in parentheses.*

A second approach is to use a tree diagram like the one shown in Figure 1. Tree diagrams are useful in counting the number of possible outcomes arising from a sequence of decisions. In the simplest case, each internal node of a tree diagram represents a decision, each child corresponds to an available choice, and each leaf is associated with an outcome. The choice selected at one internal node determines what decisions must be made subsequently. For example, if we choose gray pants, then we must subsequently decide on a blue or green shirt; however, if we choose brown pants, then we must subsequently decide on a blue or red shirt. The tree diagram in Figure 1 is a little more complicated because edges have weights; these correspond to distinct choices that lead to the same set of future decisions. For example, we can choose either of the two pairs of gray pants without affecting our subsequent choice of a shirt.

The number of matching outfits can be computed from the tree diagram as follows. A root-to-leaf path corresponds to a sequence of choices such as "gray pants, blue shirt". The number of distinct outcomes of this type can be computed by the Product Rule; we just multiply the weights of the edges on the path. For example, the number of ways to choose a pair of gray pants and a blue shirt is $2 \cdot 3 = 6$. Then we can use the Sum Rule to compute the number of outcomes of all types; we just add up all the products. This gives $6 + 2 + 9 + 6 = 23$ matching outfits. This is the same number we found by the method of counting the complement. (Good thing!)

3.2 Playoff Outcomes

In how many ways can a 5-game playoff series be decided? For example, one possibility is that the home team wins a game, then loses one, and then wins two more. In this case, there would be no need for a fifth game.

We can solve this problem with the tree diagram shown in Figure 2. In this case, there are no edge weights. (Equivalently, one can say that the edges all have weight 1.) The number of outcomes is equal to the number of leaves, which is 20.

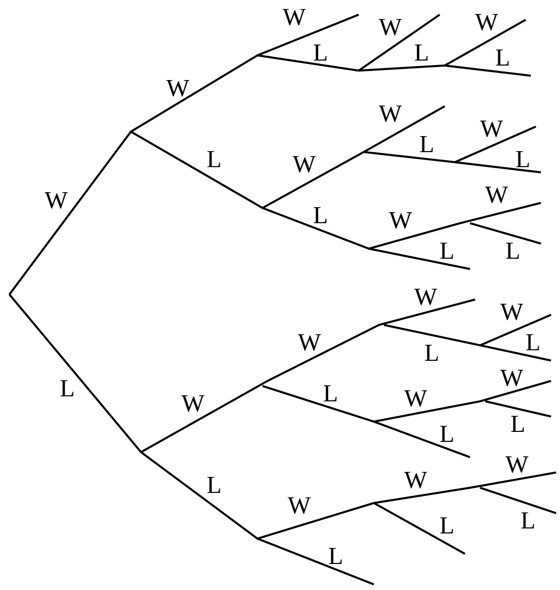

Figure 2: *This is a tree diagram for the number of ways to decide a 5-game playoff. W and L indicates wins and losses for the home team.*

In a larger problem, like the 7-game World Series, the tree diagram would be unmanageable. As we will soon see, however, sometimes we can observe a pattern and count the number of leaves without drawing out the complete tree.

4 Permutations

In how many ways can n items be ordered in a line? Each such ordering is called a *permutation*. So the question can be restated as, "How many permutations are there of n items?"

It is unmanageable to solve this problem by drawing a complete tree diagram. However, the partial diagram shown in Figure 3, suggests that there are n choices for the first item, $n-1$ choices for the second item given the first, $n-2$ choices for the third item given the first two, etc. Therefore, the total number of leaves in the tree is $n(n-1)(n-2)\ldots 2\cdot 1 = n!$. This gives an important fact; never forget it!

Fact 1 The number of permutations of n items is $n!$.

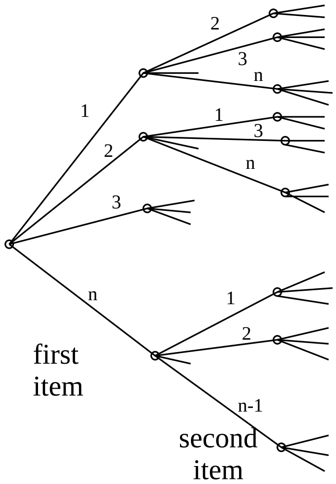

Figure 3: *This is a partial tree diagram for counting the number of ways n items can be ordered in a line. The diagram suggests that there are n choices for the first item, $n-1$ choices for the second, $n-2$ for the third, and so forth.*

For example, the number of ways to order a deck of cards is 52!. This is clearly a very big number; we can estimate how big with Stirling's Formula:

$$\begin{aligned} 52! &\geq \frac{1}{\sqrt{2\pi 52}} \left(\frac{52}{e}\right)^{52} e^{\frac{1}{12 \cdot 52 + 1}} \\ &= 2.46 \ldots \cdot 10^{65} \end{aligned}$$

This is more than the number of atoms in the universe! This is more than the number of induction proofs in 6.042! Such rapid growth in the number of possibilities in combinatorial problems gives rise to the term "combinatorial explosion". There is no way a computer could ever try all these arrangements!

4.1 A Lower Bound for Sorting

The fact that there are $n!$ ways to order n items can be combined with the Pigeonhole Principle to prove a nice lower bound on the number of comparisons needed to sort a list of n items. In particular, we can prove that the Merge Sort algorithm that we covered three weeks ago is very close to optimal. (Recall that Merge Sort required at most $n \log_2 n - n + 1$ comparisons.)

Theorem 4.1 Any *comparison-based* sorting algorithm must make at least the following number of comparisons to sort n items in the worst case:

$$n \log_2 n - n \log_2 e$$

In a comparison-based sorting algorithm, we can compare two items and move items around in memory. All other operations on items, such as looking at certain bits or performing arithmetic on items, are ruled out. For example, there is a sorting algorithm called "bucket sort" that is not comparison-based and for which this theorem does not apply.

The idea behind the proof is the same as in the ball-and-balance problems that we saw two weeks ago. There we compared the number of possible locations of the odd-weight ball with the number of leaves in a decision tree defined by our weighing scheme. Here we compare the number of initial permutations of items with the number of leaves in a computation tree defined by the sorting algorithm.

Proof. Let $x_1, \ldots, x_n$ be the items to be sorted. Since we are proving a lower bound, we can assume that the items are all distinct.

Let $\mathcal{A}$ be any sorting algorithm. Algorithm $\mathcal{A}$ must permute the items $x_1, \ldots, x_n$ to put them in the correct order. Since a different permutation is required for each different initial ordering of the items, the number of different permutations that might be required is $n!$.

The computations performed by algorithm $\mathcal{A}$ can be modeled with a tree diagram as shown in Figure 4. Each internal node corresponds to a comparison of two items. The two subtrees beneath each internal node define the computations performed based on the result of the comparison. A run of the algorithm corresponds to a root-to-leaf path. Associated with each leaf is the permutation of $x_1, \ldots, x_n$ that places the items in sorted order.

We can assume that every node in the tree diagram is reached for some input. If a node were unreachable (say, because both $x_1 < x_2$ and $x_1 > x_2$ must hold to reach it), then we could remove that node from the tree.

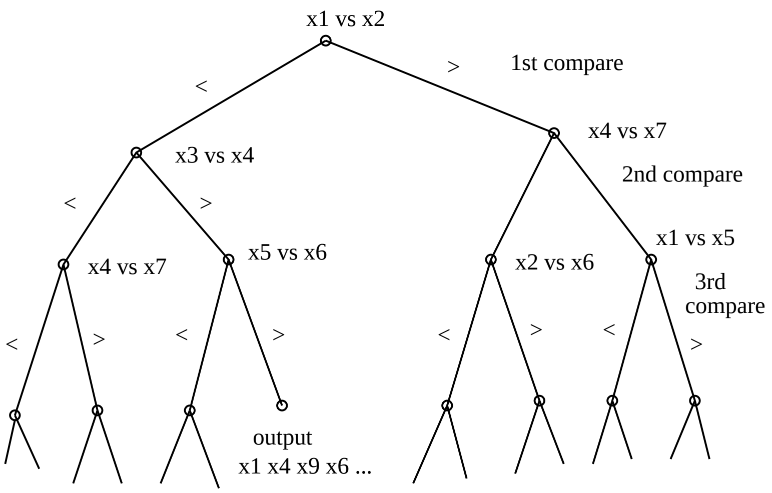

Figure 4: *This is how the computation tree of a sorting algorithm might look.*

Every one of the $n!$ permutations of $x_1, \ldots, x_n$ must appear at some leaf. Otherwise, if items are ordered according to some missing permutation, then the algorithm can not possibly output the correct answer.

Let L be the number of leaves in the computation tree. We will obtain both an upper bound and a lower bound on L.

Permutations are assigned in some way to leaves. Regard the permutations as pigeons and the leaves as holes. If $L < n!$, then by the Pigeonhole Principle, two permutations are assigned to the same leaf. This is a contradiction, since the algorithm always permutes the items one particular way. Therefore $L \geq n!$.

On the other hand, the number of leaves is limited by the depth of the computation tree. Let D be the depth of the tree; that is, D is the length of the longest root-to-leaf path. Early in the course, we proved that the number of leaves in a binary tree with depth D is at most 2^D. Therefore, $L \leq 2^D$.

Since D is the length of the longest root-to-leaf path in the computation tree, D is the number of comparisons used by algorithm $\mathcal{A}$ in the worst case. We can get a lower bound on D by putting together the upper and lower bounds on the number of leaves L.

$$\begin{aligned} 2^D \geq L &\geq n! \\ D &\geq \log_2 n! \end{aligned}$$

Now we can use Stirling's Formula to make sense of $\log_2 n!$:

$$\begin{aligned} D &\geq \log_2 \left(\sqrt{2\pi n} \left(\frac{n}{e}\right)^n e^{\frac{1}{12n+1}} \right) \\ &= n \log_2 n - n \log_2 e + \frac{1}{2} \log_2 (2\pi n) + \frac{\log_2 e}{12n+1} \\ &\geq n \log_2 n - n \log_2 e \end{aligned}$$

■

Merge Sort requires at most $n \log_2 n - n + 1$ comparisons. From the theorem, at least $n \log_2 n - 1.45n$ comparisons are necessary in the worst case. The difference is only about $n/2$, a low-order term. Merge Sort is very nearly optimal!

4.2 r-Permutations

In a horserace, the first-place horse is said to "win", the second-place horse is said to "place", and the third-place horse is said to "show". A bet in which one guesses exactly which horses will win, place, and show is called a Trifecta. How many possible Trifecta bets are there in a race with 10 horses? This problem is easy enough that we do not need to draw out the tree diagram; there are 10 choices to win, 9 choices to place given the winner, and 8 choices to show given the first two finishers. This gives $10 \cdot 9 \cdot 8 = 720$ possible Trifecta bets.

This is a special case of a standard problem: counting r-permutations of a set.

Definition An *r-permutation* of a set is an ordering of r distinct items from the set.

For example, the 2-permutations of the set $\{A, B, C\}$ are:

$$\begin{array}{ccc} (A,B) & (A,C) & (B,C) \\ (B,A) & (C,A) & (C,B) \end{array}$$

The number of r-permutations of an n-element set comes up so often in combinatorial problems that there is a special notation.

Definition $P(n, r)$ denotes the number of r-permutations of an n-element set. In other words,

$$P(n,r) = n(n-1)\dots(n-r+1) = \frac{n!}{(n-r)!}$$

For example, the number of Trifecta bets in a 10 horse race is $P(10, 3) = \frac{10!}{7!} = 720$.

How many strings of 5 letters are there with no repetitions? For example, ZYGML, is such a string. In general, these strings are exactly the 5-permutations of the set of letters. Therefore, there are $P(26, 5) = \frac{26!}{21!}$, which is 7,893,600.