Some induction solutions

Introduction

I'm going to present solutions to some of the induction problems at the end of Jeff Erickson's notes. I will do some of the PECS structure, but for exam purposes you probably want to concentrate on the E. I will no doubt have some minor asides.

Problem 1

Prove that given an unlimited supply of 6-cent coins, 10-cent coins, and 15-cent coins, one can make any amount of change larger than 29 cents.

Preparation

The proposition $P(n)$ that we're trying to verify is 'We can make change for n cents', and we have to show that it's true for $n \geqslant 30$. It's obvious to try a few values to begin with

$$\begin{aligned} 30 &= 2 \times 15 \\ 31 &= 1 \times 15 + 1 \times 10 + 1 \times 6 \\ 32 &= 2 \times 2 + 2 \times 6 \end{aligned}$$

That's odd – there were other possibilities for the first case but there seems to be no obvious pattern between the three that we could lift for an induction. Time to think a bit more.

Well, as soon as we get six consecutive amounts we can make change for, then we can lift up by 6 by using one more 6-cent coin. That seems to be the core idea, so let's just finish off the next three cases:

$$\begin{aligned} 33 &= 1 \times 15 + 3 \times 6 \\ 34 &= 1 \times 10 + 4 \times 6 \\ 35 &= 1 \times 15 + 2 \times 10 \end{aligned}$$

So that gets the six in a row, and now we just add enough extra 6 cent coins to do the whole thing.

Execution

We will prove by induction that for $n \geqslant 30$ it is possible to make change for n cents using 6-cent, 10-cent and 15-cent coins.

The computations above show that this is possible if $30 \leqslant n \leqslant 35$. So, suppose that $n \geqslant 36$ and that the result is true for all k with $30 \leqslant k < n$.

In particular, since $30 \leqslant n - 6$ it is possible to make change for $n - 6$ cents. Do so, and then add one 6-cent coin. This shows it is possible to make change for n cents.

By induction, for $n \geqslant 30$ it is possible to make change for n cents using 6-cent, 10-cent and 15-cent coins.

Satisfaction

I don't actually find this a very satisfactory example since I think there's a better way to understand the result. I wouldn't ask you to prove by induction that you can make change for n dollars using 2-dollar coins and at most one 1-dollar coin would I? You understand naturally the principle. So, I'd tend to lift that understanding to this problem as follows:

- First, compute as above that we can do it for 30 through 35 cents.
- Now suppose we're given any $n > 30$. Write $n = 6k + r$ where $0 \leqslant r < 6$. Since $n > 30$, $k \geqslant 5$.
- So, $n = 6(k - 5) + (30 + r)$ and we can make change for n using $(k - 5)$ 6-cent coins and whatever coins we need for $30 + r$.

Aside: Is there still an induction in that form of the argument? Not directly, but there is an appeal to the *division algorithm* that we can do division with a unique remainder lying between 0 and the divisor. To prove that's possible *does* actually require induction, even though it's arithmetic you learn in elementary school.

Problem 2

Prove that

$$\sum_{i=0}^{n} r^i = \frac{1 - r^{n+1}}{1 - r}$$

for every non-negative integer n and every $r \neq 1$.

Preparation

A review of notation is in order - the Σ stands for "sum" and should (or at least can) be thought of as a for loop – add up the quantities to its right as the index variable (i) runs from 0 to n inclusive.

The base case ($n = 0$) is clear - there's only one term in the sum and it's r^0 which is 1 while the right hand side is $(1 - r)/(1 - r)$ and that's also 1. But even the next case is a little mysterious as the left hand side is $1 + r$ and the right hand side is $(1 - r^2)/(1 - r)$. But, it is indeed the case that $1 - r^2 = (1 - r) \times (1 + r)$ so that works too.

I think we'll just have to dive in and trust the algebra to work out.

Execution

Define $S(n) = \sum_{i=0}^{n} r^i$. We will prove by induction that, if $r \neq 1$ then for all non-negative integers n, $S(n) = (1 - r^{n+1})/(1 - r)$.

This is true when $n = 0$ since $S(0) = 1$ while $(1 - r)/(1 - r) = 1$ also.

Suppose that $n > 0$ and that the result holds at $n - 1$, i.e., that

$$S(n-1) = \frac{1 - r^n}{1 - r}.$$

Note that $S(n) = S(n - 1) + r^n$. Using this, and the formula for $S(n - 1)$ we compute:

$$\begin{aligned}
S(n) &= S(n-1) + r^n \\
&= \frac{1 - r^n}{1 - r} + r^n \\
&= \frac{1 - r^n + r^n \times (1 - r)}{1 - r} \\
&= \frac{1 - r^n + r^n - r^{n+1}}{1 - r} \\
&= \frac{1 - r^{n+1}}{1 - r}.
\end{aligned}$$

This is what we wished to prove.

By induction, we conclude that $S(n) = (1 - r^{n+1})/(1 - r)$ for all non-negative integers n provided that $r \neq 1$ (which is needed to avoid division by zero).

Satisfaction

This is a bit better than the previous one, but I still don't like it much. The proof I was taught in secondary school goes like this:

$$S(n) = 1 + r + r^2 + \cdots + r^n$$
$$r \times S(n) = r + r^2 + \cdots + r^n + r^{n+1}.$$

Take the difference of these two and note that all but the first term in the top line cancel with all but the last term in the bottom line. That is:

$$(1 - r) \times S(n) = 1 - r^{n+1}.$$

Now, if $r \neq 1$ we can divide both sides by $1 - r$ and so we're finished.

> **Aside:** So this one is actually quite interesting – is there a hidden induction in that argument? There is, but a very subtle one. The basic rules of arithmetic that allow us to rearrange sums are the commutative and associative laws, i.e., that $a + b = b + a$ and $a + (b + c) = (a + b) + c$ for any a, b, and c. How can we conclude that in an arbitrary arithmetic expression such as the one we get from $S(n) - r \times S(n)$ which has $2n + 2$ terms we can rearrange them and cancel out all the matching terms? By an inductive application of the basic rules for sums of two and three terms! That is, we can prove by induction that if A and B are two arithmetic expressions that differ only in the order of the terms somehow, then their values are equal (the general commutative law) and that this value is independent of the order in which we add up pairs of terms (the general associative law).

Problem 4a and 4b

Recall the standard recursive definition of the Fibonacci numbers:

$$F_0 = F_1 = 1, \textit{ and, } F_n = F_{n-1} + F_{n-2} \textit{ for } n \geqslant 2.$$

Prove that $\sum_{i=0}^{n} F_i = F_{n+2} - 1$ *for all* $n \geqslant 0$ *and that*

$$F_n^2 - F_{n+1}F_{n-1} = (-1)^n$$

for all $n \geqslant 1$.

Note that I've changed the indexing on the Fibonacci numbers to be consistent with our usage in COSC201, which accounts for the change in the exponent on the right hand side of the second expression.

Preparation

This really looks like it's going to be an algebra slog. The first part doesn't look too bad since when we increase n by one we can replace the sum on the left hand side by the inductive expression for all but its final term, and the final term, and hopefully that will work.

The second one looks a bit messy since the inductive hypothesis will have things like F_{n-1}^2 in it, while we'll want to be working with F_n^2. It's a matter of scribbling out a few things and hoping for the best.

Execution

We will prove both results by induction.

The first is true at $n = 0$ since the left hand side is $F_0 = 1$ while the right hand side if $F_2 - 1 = 2 - 1 = 1$. So the result holds for $n = 0$.

Suppose now that $n > 0$ and the result holds at $n - 1$. That is, $\sum_{i=0}^{n-1} F_i = F_{n-1+2} - 1 = F_{n+1} - 1$.

Then,

$$\begin{aligned}
\sum_{i=0}^{n} F_i &= \left(\sum_{i=0}^{n-1} F_i\right) + F_n \\
&= F_{n+1} - 1 + F_n \text{ (using the inductive hypothesis)} \\
&= F_n + F_{n+1} - 1 \\
&= F_{n+2} - 1,
\end{aligned}$$

which is just what we wanted.

So, by induction, $\sum_{i=0}^{n} F_i = F_{n+2} - 1$ for all $n \geqslant 0$.

For the second result we confirm that for $n = 1$

$$F_1^2 - F_2 F_0 = 1^2 - 2 \times 1 = -1 = (-1)^1$$

so the result holds at $n = 1$.

Now suppose that $n > 1$ and that the result holds at $n-1$. Consider:

$$\begin{aligned}
F_n^2 - F_{n+1}F_{n-1} &= F_n \times (F_{n-1} + F_{n-2}) - (F_n + F_{n-1}) \times F_{n-1} \\
&= F_nF_{n-1} + F_nF_{n-2} - F_nF_{n-1} - F_{n-1}^2 \\
&= F_nF_{n-2} - F_{n-1}^2 \\
&= -\left(F_{n-1}^2 - F_nF_{n-2}\right) \\
&= -(-1)^{n-1} \quad \text{(using the inductive hypothesis)} \\
&= (-1)^n,
\end{aligned}$$

which is just what we wanted.

So, by induction, $F_n^2 - F_{n+1}F_{n-1} = (-1)^n$ for all $n \geqslant 1$.

Satisfication

The Fibonacci numbers (and related sequences) are a rich source of identities of this type. Ones which just involve sums are usually pretty straight forward. Ones involving products can be a bit more involved. If you know a bit of matrix algebra, you should be able to show that the entries of the matrix

$$\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}^n$$

are all closely related Fibonacci numbers, and many identities follow from this.

Question 5, paraphrased slightly

Prove that every integer (positive, negative, or zero) can be expressed as a sum of distinct terms of the form $\pm 3^i$ where i can be any non-negative integer. By convention, the sum of an empty set of terms is 0.

Preparation

This repays doing some examples, but the first thing to notice is that since we can change the sign of every term, we can just consider non-negative integers. Of course 0 is easy from the note at the end of the question.

$$
\begin{aligned}
1 &= 3^0 \\
2 &= 3^1 - 3^0 \\
3 &= 3^1 \\
4 &= 3^1 + 3^0 \\
5 &= 3^2 - 3^1 - 3^0 \\
6 &= 3^2 - 3^1 \\
7 &= 3^2 - 3^1 + 3^0 \\
8 &= 3^2 - 3^0 \\
9 &= 3^2.
\end{aligned}
$$

The interesting thing that this suggests to me is the point at which negative terms are needed - the first case is 2, the second case is 5. Each of these is half (rounded up) of the next power of 3. I think this suggests an inductive/recursive strategy to represent a positive integer n. Namely, if n is greater than or equal to 3^k but less than $3^{k+1}/2$, write it as 3^k plus the representation of $n - 3^k$, while if it is greater than $3^{k+1}/2$ but less than 3^{k+1} then write it as 3^{k+1} minus the representation of $3^{k+1} - n$.

The one thing to keep an eye on is to make sure that we get *distinct* terms when we carry out this strategy.

Execution

We will prove by induction that every non-negative integer can be expressed as a sum of distinct terms of the form $\pm 3^i$ where i can be any non-negative integer. In fact, we will prove more, that if $n < 3^{k+1}/2$ for some k then there is such an expression where in each term $i <\leqslant k$.

> **Aside:** We need the "more" part in order to make the induction go smoothly and get sums with distinct terms – this is a bit unsatisfactory and will be dealt with in the "Satisfaction" section!

From this the result will follow because 0 is expressible as an empty sum, and we can express a negative integer by changing the signs of all the terms in the representation of its absolute value.

From the above, we see that the result is certainly true for $n = 0$ (in fact for $0 \leqslant n \leq 9$). Suppose now that $n > 1$ and the result holds for all $0 \leqslant k < n$.

Choose k so that $3^k \leqslant n < 3^{k+1}$. If $n < 3^{k+1}/2$ then:

$$
n - 3^k < 3^{k+1}/2 - 3^k = \frac{3 \times 3^k - 2 \times 3^k}{2} = \frac{3^k}{2}.
$$

In this case, by the inductive hypothesis, $n - 3^k$ can be represented as a sum of terms of the required type where the maximum exponent occurring is less than k. Taking this sum, and adding 3^k gives the required representation of n.

In the other case, $n > 3^{k+1}/2$ (it can't be equal because $3^{k+1}/2$ is not an integer!). Then:

$$3^{k+1} - n < 3^{k+1} - \frac{3^{k+1}}{2} = \frac{2 \times 3^{k+1} - 3^{k+1}}{2} = \frac{3^{k+1}}{2}.$$

So, by the inductive hypotheses $3^{k+1} - n$ can be represented as a sum of terms of the required type where the maximum exponent occurring is less than $k+1$. But then n can be expressed as 3^{k+1} plus the negative of that representation.

In either case we can find a representation of n of the required form (and with the requisite conditions on the powers of 3 that occur), and so by induction the result is true.

Satisfaction

Not very. That's messy and awkward. Perhaps it can point us to a better way. And, I think it does (to two better ways in fact). Namely, the argument shows that we're really thinking about "What integers can we represent using distinct terms of the form $\pm 3^i$ where $0 \leqslant i \leqslant k$?" (where k is now the parameter). And the argument shows that it is every integer in the range from $-3^{k+1}/2$ to $3^{k+1}/2$ (I know those aren't integers - that's fine, it's the integers in the range). That gives a new statement to prove by induction!

We will prove by induction that, for every non-negative integer k, any integer a with $-3^{k+1}/2 < a < 3^{k+1}/2$ can be written as a sum of distinct terms of the form $\pm 3^i$ where $0 \leqslant i \leqslant k$.

The result is true when $k = 0$ because the only values of a that we need to to represent are -1, 0, and 1, and we can do this as -3^0, the empty sum, and 3^0.

So, suppose that $k > 0$ and that the result is true for all smaller values of k. We need to show that every a with $-3^{k+1}/2 < a < 3^{k+1}/2$ can be written as a sum of distinct terms of the form $\pm 3^i$ where $0 \leqslant i \leqslant k$. We can assume that a is positive (if a is negative, change its sign, find a representation for that, and then change the sign of all the terms). If $a < 3^k/2$ then by the inductive hypothesis, we can find such a sum using only $i < k$ which is good enough. So, the remaining cases are where

$$\frac{3^k}{2} < a < \frac{3^{k+1}}{2}.$$

Our plan is to use 3^k and terms with smaller exponents to represent a. So we compute:

$$\begin{aligned}\tfrac{3^k}{2} - 3^k < a - 3^k < \tfrac{3^{k+1}}{2} - 3^k\\ -\tfrac{3^k}{2} < a - 3^k < \tfrac{3\times 3^k - 2\times 3^k}{2}\\ -\tfrac{3^k}{2} < a - 3^k < \tfrac{3^k}{2}\end{aligned}$$

So, we can represent $a - 3^k$ using terms of the form $\pm 3^i$ with $0 \leqslant i < k$, and hence can represent a as 3^k plus this representation.

Therefore, by induction, the result holds.

Finally (and informally), I know how to write (positive) numbers in base 3 (using digits 0, 1 and 2). I'm trying to write them (effectively) using digits 0, 1 and -1. But

$$2 \times 3^i = 3^{i+1} - 3^i$$

So I could convert a 2 digit somewhere to a -1 digit, adding one to the next digit along. What if that makes that digit 3? Well then I can convert it to a 0 adding 1 to the next digit along. Eventually (working from least significant to most significant digit) I run out and I have a representation of the kind I want. Here's what I mean:

$$\begin{aligned}464_{10} &= 122012_3\\ &= 12202(-1)_3\\ &= 1221(-1)(-1)_3\\ &= 13(-1)1(-1)(-1)_3\\ &= 20(-1)1(-1)(-1)_3\\ &= 1(-1)0(-1)1(-1)(-1)_3\\ &= 729 - 243 - 27 + 9 - 3 - 1\\ &= 464\end{aligned}$$

A number of the remaining problems in the section are also about weird bases to represent numbers in (6, 7, 19).

Problem 13, paraphrased

The d-dimensional hypercube, Q_d, is a graph with 2^d vertices. The vertices are all the strings of 0's and 1's of length d and two vertices are adjacent if they differ in exactly one position. Prove that for every $d \geqslant 2$, Q_d has a Hamilton cycle (a walk, starting and finishing at the same vertex and visiting every other vertex exactly once).

Preparation

The graph Q_2 is a square and there's a pretty obvious Hamilton cycle (HC) for instance

$$00 \to 01 \to 11 \to 10 \to 00$$

The graph Q_3 looks like the vertices of a cube. We have two cycles around opposite faces based on the one from Q_2 (I'm doing this of course because I already have induction in mind)

$$000 \to 001 \to 011 \to 010 \to 000$$
$$100 \to 101 \to 111 \to 110 \to 100$$

We'd like to sew these together into an HC for Q_3. That might take a bit of thinking - but the basic idea is to go almost all the way round one face, jump up to the other, follow the corresponding path backwards, and then drop back down. That is:

$$000 \to 001 \to 011 \to 010 \to 110 \to 111 \to 101 \to 100 \to 000.$$

And that's the inductive step in general too - we're ready to go.

Execution

We will prove by induction that, for $d \geqslant 2$, the d-dimensional hypercube, Q_d has a Hamilton cycle.

As noted above, the result is clearly true for $d = 2$ since Q_2 is just a cycle of length four.

Suppose then that the $d > 2$ and the result holds for $d - 1$. Let $v_0, v_1, v_2, \dots, v_k$ be an HC for Q_{d-1} (so $k = 2^{d-1}$ and $v_k = v_0$ but all other vertices are distinct).

The vertices of Q_d are of the form $0v$ and $1v$ where v is a vertex of Q_{d-1}. Vertices with the same first bit are adjacent if the corresponding vertices in Q_{d-1} are adjacent, and vertices with differing first bits are adjacent only if the remaining parts are the same. The following is an HC of Q_d:

$$0v_0 \to 0v_1 \to 0v_2 \to \cdots \to 0v_{k-1} \to 1v_{k-1} \to 1v_{k-2} \to \cdots \to 1v_2 \to 1v_1 \to 1v_0 \to 0v_0.$$

Therefore, Q_d has an HC in this case, and by induction this is true for all $d \geqslant 2$.

Satisfaction

It's nice to notice that we've actually proved a bit of a theorem. For any graph G we can define the graph $2 \cdot G$ to be a graph containing two disjoint copies of G with each vertex in one copy joined to the corresponding vertex in the other copy (we have, e.g.,

$Q_d = 2 \cdot Q_{d-1}$). Then what our argument really shows is that whenever G has an HC, so does $2 \cdot G$.

There are still open problems about HC's in Q_d - see the wikipedia page. For instance, it's unknown how many different HC there are in Q_d in general. The exact values for $d \leqslant 6$ are known (the last one determined only in 2010), but not beyond that. They are related to Gray codes, and to solutions of the Towers of Hanoi puzzle.

Problem 14

A *tournament* is a directed graph with exactly one directed edge between each pair of vertices. That is, for any vertices v and w, a tournament contains either an edge $v \to w$ or an edge $w \to v$, but not both. The reason for the name is that we can think of the edges as representing the outcomes of the matches in a single-round-robin tournament where the vertices represent players (and ties are not possible).

A Hamiltonian path in a directed graph G is a directed path that visits every vertex of G exactly once.

Prove that every tournament has a Hamiltonian path. In fact, prove that every tournament either has a unique Hamiltonian path or contains a directed triangle - three vertices a, b and c with edges $a \to b \to c \to a$.

Preparation

Let's start with the first part. Small examples don't help much here - the result is completely trivial in the case of one or two vertices, and you can check all the possibilities for three vertices, but I don't think they help much in thinking how the induction will go.

So let me think of this a little more abstractly. I have a tournament, and a particular player a (I'm going to use the player rather than vertex terminology). If there's a Hamiltonian path (HP) then that path goes through a at some point. The predecessor of a (if any) is a player who beat a, and the follower (if any) is someone whom a beat. But every player either beat a or was beaten by a and if we look at just the players who beat a then, among themselves, they played a smaller tournament (and of course the same applies to the players a beat). So we can just take HP in the sub-tournament of players who beat a, followed by a, followed by HP in the sub-tournament of players whom a beat. There's the induction!

OK, what about the second part? One case where we'd know for sure where an HP starts is the case where there's a champion who didn't lose any matches. After some scribbling of examples, I think perhaps that if there's no directed triangle there has to be such a champion. Let me see why. Consider any player a whose overall record (number of wins) is at least as good as everyone else's. Suppose that a lost a match, say to c. Then c cannot have beaten all the players whom a beat, or c would have a better record. So there's some player b who a beat, but who also beat c, i.e., a directed triangle.

So if there's no directed triangle there's a grand champion and they have to be first in the HP. But now the sub-tournament on all the remaining players also has a grand champion, who must be second in the HP and so on.

Execution

We will prove by induction that, in any tournament, there is an HP.

A tournament on 0, 1 or 2 vertices has an HP trivially (completely empty, just the vertex, or just the edge of the tournament). So suppose that we have a tournament on n vertices for some $n > 2$ and the result holds for all tournaments on k vertices for all $k < n$.

Choose a vertex a and let L be the set of vertices x where $x \to a$ is an edge (i.e., the set of players whom a lost to), and let B be the set of vertices x where $a \to x$ is an edge (i.e., the set of players whom a beat). Considering just edges between vertices of L we have a tournament with fewer than n vertices and hence an HP, PL. Likewise, considering just edges between vertices of B we have a tournament with fewer than n vertices and hence an HP, PB. But then PL followed by a followed by PB is an HP of the original tournament.

By induction, every tournament has an HP.

For the second part, see the preparation section above!

Satisfaction

I'm satisfied that I finished writing these solutions up!