CS 70 Spring 2005 — Discrete Mathematics for CS, Clancy/Wagner — Final Solns

Problem 1. (8 points)

Let $a(n)$ denote the number of ternary strings of length n with no pair of adjacent 2's. Example: "2011" is a ternary string with no pair of adjacent 2's.

(a) What are $a(1)$ and $a(2)$?

Solution: $a(1) = 3$, $a(2) = 8$.

Grading: 1 pt for correct answer.

Solution #1: $a(n) = 2a(n-1) + 2a(n-2)$ for $n \geq 3$.
Justification: If the first digit is 0 or 1 (2 choices), then the remaining $n-1$ digits can be any ternary string with no 22s ($a(n-1)$ choices). If the first digit is 2 (1 choice), then the second digit can be 0 or 1 (2 choices), and the remaining $n-2$ digits can be any ternary string with no 22s ($a(n-2)$ choices).

Solution #2: $a(n) = 3a(n-1) - 2a(n-3)$ for $n \geq 4$.
Justification: Take any ternary string of length $n-1$ with no 22s, and append 0, 1, or 2 ($3a(n-1)$ choices). We need to exclude strings of the form $\cdots 022$ and $\cdots 122$ ($2a(n-3)$ choices); once we've eliminated these, what remains must be exactly the set of ternary strings of length n with no 22s.

Common mistakes: Overcounting. For example, some said $a(n) = 3^n - 3^{n-2}(n-1)$, reasoning that there are $3^{n-2}(n-1)$ ternary strings with a 22, since there are $n-1$ places where the 22 could start and 3^{n-2} options for the other $n-2$ digits in the string. The problem with this reasoning is that it overcounts strings with multiple 22s; for instance, the string 22022 is counted twice, and the string 22222 is counted four times.

Grading: 3 pts for a correct recurrence; 1 pt for something with a partial explanation.

(b) Prove that $a(n) \geq 2^{n+1}$ for all $n \geq 2$.

Solution #1: Proof by strong induction.
Base cases: $a(2) = 8 \geq 2^{2+1}$. $a(3) = 22 \geq 2^{3+1}$.
Induction step: Assume $a(k) \geq 2^{k+1}$ for $k = 2, 3, \ldots, n$. Then for $n \geq 3$,

$$\begin{aligned} a(n+1) &= 2a(n) + 2a(n-1) && \text{(by part (b))} \\ &\geq 2 \cdot 2^{n+1} + 2 \cdot 2^{n-1+1} && \text{(by the inductive hypothesis)} \\ &\geq 2 \cdot 2^{n+1} = 2^{(n+1)+1}. \end{aligned}$$

Solution #2: For each ternary string of length n with no 22s, we can get two unique ternary strings of length $n+1$ with no 22s by appending either a 0 or a 1. This shows that $a(n+1) \geq 2a(n)$ for all $n \geq 1$. Now use induction to show the desired result.

Solution #3: There are 2^n ternary strings (of length n) with no 2. There are $n \times 2^{n-1}$ more that have exactly one 2. Thus there are at least $2^n + n \times 2^{n-1}$ ternary strings of length n with no 22. Moreover, $2^n + n \times 2^{n-1} \geq 2^n + 2 \times 2^{n-1} \geq 2^{n+1}$ for $n \geq 2$.

Grading: 4 pts for an impeccable proof. For induction proofs: 1 pt for all needed base cases and a correctly stated induction predicate; 3 pts for a correct and fully justified inductive step. -1 pt for backwards reasoning, for steps with no justification, for small algebraic errors, for confusing what you can assume vs. what you must prove, or for confusion between integers and booleans (e.g., $a(n)$ vs. $P(n)$).

Problem 2. (4 points)

Solve the following system of equations modulo 7 for x and y. Show your work.

$$y \equiv 5x - 3 \pmod 7$$
$$y \equiv 3x + 2 \pmod 7$$

Solution 1: First solve for y. Multiply the first equation by 3 and the second equation by 5, giving

$$3y \equiv 15x - 9 \pmod 7$$
$$5y \equiv 15x + 10 \pmod 7$$

Subtract the first equation from the second, giving $2y \equiv 19 \pmod 7$, or equivalently $2y \equiv 5 \pmod 7$. Since $2^{-1} \equiv 4 \pmod 7$, we see that $y \equiv 2^{-1} \cdot 5 \equiv 4 \cdot 5 \equiv 6 \pmod 7$.

Substituting into the first equation gives $6 \equiv 5x - 3 \pmod 7$, then $9 \equiv 5x \pmod 7$, and finally $2 \equiv 5x \pmod 7$. Since $5^{-1} \equiv 3 \pmod 7$, we have $2 \cdot 3 \equiv x \pmod 7 \equiv 6$.

Solution 2: First solve for x. Subtract the second equation from the first. We get $0 \equiv 2x - 5 \pmod 7$, or $2x \equiv 5 \pmod 7$. Since $2^{-1} \equiv 4 \pmod 7$, we see that $x \equiv 2^{-1} \cdot 5 \equiv 4 \cdot 5 \equiv 6 \pmod 7$.

Plugging in to the second equation yields $y \equiv 3 \cdot 6 + 2 \equiv 20 \equiv 6 \pmod 7$.

Final answer: $x \equiv 6 \pmod 7$, $y \equiv 6 \pmod 7$.

Common mistakes: Writing $x = 6$ rather than $x \equiv 6 \pmod 7$: in fact, any of $x = \ldots, -8, -1, 6, 13, 20, \ldots$ works. Same for y. Also, some ended up with fractions, when they should have used modular inverses: e.g., claims that the solution to $2x \equiv 5 \pmod 7$ is $x \equiv 2.5 \pmod 7$.

Grading: 4 pts for a fully correct answer; -1 pt if final answer doesn't include all solutions or isn't expressed modulo 7. 1 pt for partial progress towards an answer.

Problem 3. (4 points)

Suppose there are two males and two females to be paired. Prove that in every stable pairing, at least two individuals get their first-choice mate.

Solution #1: I will prove the contrapositive. Consider any stable pairing where less than two individuals get their first-choice mate. I will demonstrate that this pairing is not stable. Without loss of generality, assume Alex is paired with Carol in this pairing, and Bill with Dianne. There are two cases:

- Exactly one person gets his/her first-choice mate.
  (Without loss of generality, we may assume Alex is the one who gets his first-choice mate.)
- No one gets his/her first-choice mate.

In either case, both Bill and Carol got their second-choice mate. This means that their preference lists must be Bill: Carol $>$ Dianne; and Carol: Bill $>$ Alex. But this means that Bill $\leftrightarrow$ Carol form a rogue couple, so this pairing cannot be stable.

Solution #2: Suppose not. At least three individuals are "unhappy" (fail to get their first-choice mate), so by the pigeonhole principle, at least one whole gender is unhappy. Also, there is at least one unhappy person of the other gender, call him/her Sam. Then Sam's non-spouse prefers Sam, and Sam prefers his/her non-spouse, so they form a rogue couple.

Solution #3: Consider any stable pairing, with couples (m_1, f_1) and (m_2, f_2) paired up. Then since (m_1, f_2) is not a rogue couple, we know that at least one of $\{m_1, f_2\}$ received his/her first-choice mate. Similarly, since (m_2, f_1) is not a rogue couple, at least one of $\{m_2, f_1\}$ received their first-choice mate. In total, we have at least $1+1=2$ individuals who received their first-choice mate.

Common mistakes: Some used a case analysis, but ruled out only the case where exactly one person gets their first-choice mate, and forgot to do the case of zero first-choice mates. Many proved the claim only for TMA pairings. But, some stable pairings cannot be output by the TMA (namely, pairings that aren't male-optimal).

Grading: For proof by cases: +2 pts for each of the two cases. 2 pts if you only proved the theorem for the pairing output by the TMA.

Problem 4. (10 points)

Shuffle a standard 52-card deck, and deal out a hand of 13 cards. You get 4 points for each Ace in the hand, 3 points for each King, 2 for each Queen, 1 for each Jack, and nothing for the other cards. Let the random variable X denote the total number of points in your hand.

(a) Calculate $\Pr[X=0]$. You may leave your answer as an unevaluated expression. Show your work.

Solution #1: $\Pr[X=0] = \frac{36}{52} \times \frac{35}{51} \times \cdots \times \frac{24}{40}$.
Justification: $39/52$ is the probability that the first card is worth 0 points, since there are $52-16=36$ 0-point cards. $35/51$ is the probability that the second card is worth 0 points given that the first is, etc.

Solution #2: There are $\binom{36}{13}$ ways to choose a subset of 13 cards among all 36 0-point cards (ignoring the order in which cards are chosen). This is out of a sample space of $\binom{52}{13}$ ways to choose a 13-card hand (again ignoring order). Since the probability distribution is uniform, we can divide to get $\Pr[X=0] = \binom{36}{13}/\binom{52}{13}$.

Common mistakes: $(36/52)^{13}$ (computed for selection with replacement). Some tried to compute Pr[no Aces and no Kings and ...] or $\Pr[X \neq 0]$ using the sum rule (but the sets aren't disjoint).

Grading: 3 pts for correct answer. 1 pt for partial progress, e.g., computing Pr[no Aces] correctly. Nothing deducted for minor arithmetic errors (e.g., miscounted # of face cards).

(b) Calculate $\mathbf{E}[X]$. This time we want a number; again, show your work.

Solution #1: $X = X_1 + \cdots + X_{13}$, where $X_i =$ the number of points you receive for the ith card in your hand.

Also $\Pr[X_i = 0] = 36/52$ and $\Pr[X_i = 1] = \Pr[X_i = 2] = \Pr[X_i = 3] = \Pr[X_i = 4] = 4/52$, since there are 36 0-point cards and four 1-point, 2-point, 3-point, and 4-point cards. This means $\mathbf{E}[X_i] = 0 \times 36/52 + 1 \times 4/52 + 2 \times 4/52 + 3 \times 4/52 + 4 \times 4/52 = 40/52$.

Finally, by linearity of expectation, $\mathbf{E}[X] = \mathbf{E}[X_1] + \cdots + \mathbf{E}[X_{13}] = 13 \times 40/52 = 10$.

Solution #2: Deal out all 52 cards to four players, each player receiving 13 cards. Let the random variables W, X, Y, Z denote the number of points each player receives. Note that there are 40 points total in the deck, so we know $W + X + Y + Z = 40$. By linearity of expectation, $\mathbf{E}[W] + \mathbf{E}[X] + \mathbf{E}[Y] + \mathbf{E}[Z] = 40$. But $\mathbf{E}[W] = \mathbf{E}[X] = \mathbf{E}[Y] = \mathbf{E}[Z]$ (by symmetry), so $4 \times \mathbf{E}[X] = 40$, or $E[X] = 40/4 = 10$. Intuitively: "there are 40 points in the deck, and each player receives one-fourth of them, on average."

Grading: 4 pts for correct answer. 3 pts for correct answer without explaining where answer came from (no mention of linearity of expectation, X_i r.v.'s, ...). 1 pt for setting up problem correctly (e.g., defined X_i's and noted $\mathbf{E}[X] = \sum_i \mathbf{E}[X_i]$, but couldn't compute $\mathbf{E}[X_i]$).

(c) Suppose we tell you that your hand has 6 spades, 2 hearts, 3 diamonds, and 2 clubs. Now consider the expected value of X, conditioned on this fact. Does this number increase, decrease, or stay the same, compared to your answer in part (b)? Briefly justify your answer.

Solution: Stays the same.

As in solution #1 to part (b), write $X = X_1 + \cdots + X_{13}$, where cards 1–6 are the spades, 7–8 are the hearts, 9–11 are the diamonds, and 12-13 are the clubs.

Now $\Pr[X_1 = 0] = 9/13$ and $\Pr[X_i = 1] = \Pr[X_i = 2] = \Pr[X_i = 3] = \Pr[X_i = 4] = 1/13$, since there are 13 0-point spade cards and four 1-point, 2-point, 3-point, and 4-point spade cards. Note that $9/13 = 36/52$ and $1/13 = 4/52$. We see that this extra knowledge does not affect the distribution of X_1 at all, so $\mathbf{E}[X_1]$ remains the same. The same goes for all of the $\mathbf{E}[X_i]$.

Therefore, $\mathbf{E}[X]$ is unaffected by this extra knowledge.

Common mistakes: Many thought it would decrease: at most four of the six spades can be face cards, so at least two of the spades cannot possibly be face cards. The flaw in this reasoning becomes evident if I tell you that you have a hand with 13 spades: then you are guaranteed to get 10 points, so the expectation stays the same, even though nine of your cards aren't face cards. Intuitively, each of the spade cards "has a chance" to be a face card; the fifth spade cannot be a face card only if the first four spades were all face cards.

Grading: 3 pts for "stays the same", with correct intuition: e.g., "card rank/value/point-count is independent of suit", or "probability of getting a face card is independent of suit", at a minimum. 2 pts for correct reasoning, but forgot to state final answer ("stays the same"). 1 pt for "stays the same", with a handwave or weak justification. 0 pts for "stays the same", with a totally wrong justification. 0 pts for "decreases".

Problem 5. (9 points)

We have a room full of m people. Assume that each person's birth date is uniformly and independently distributed among the set of 365 possibilities. Let the random variable X denote the number of different birth dates among all the people in the room.

Example: Alice's birth date is April 5, Bob's is June 23, Carol's is April 5. There are two different birth dates, so $X = 2$ in this case.

(a) Calculate $\mathbf{E}[X]$. Your answer should be a simple function of m.

Solution #1: Let $Y_i = 1$ if no one has a birthday on day i, and $Y_i = 0$ otherwise. Then $X = 365 - Y_1 - Y_2 - \cdots - Y_{365}$.

Also $\Pr[Y_i = 1] = (1 - \frac{1}{365})^m$, since the probability that the jth person's birthday is on day i is $1 - \frac{1}{365}$, and these probabilities are independent. Since Y_i is an indicator variable, this means $\mathbf{E}[Y_i] = (1 - \frac{1}{365})^m$.

Finally, by linearity of expectation, $\mathbf{E}[X] = 365 - \mathbf{E}[Y_1] - \cdots - \mathbf{E}[Y_{365}] = 365 - 365 \times (1 - \frac{1}{365})^m$.

Solution #2: "Balls and bins." Think of each date as a "bin", and each person as a "ball", so there are m balls and 365 bins. Let Y denote the number of empty bins after tossing m balls (i.e., the number of birthdays that don't appear in a room of m people). By the calculation at the end of Notes 20, $\mathbf{E}[Y] = 365 \times (1 - \frac{1}{365})^m$. Finally, $X = 365 - Y$, so $\mathbf{E}[X] = 365 - 365 \times (1 - \frac{1}{365})^m$.

Common mistakes: Many tried to use the definition of expectation, $\mathbf{E}[X] = \sum_i i \times \Pr[X = i]$. The problem is that $\Pr[X = i]$ is very hard to compute exactly (see part (b)).

Grading: 5 pts for fully correct answer. 4 pts if general approach is correct, but calculation or final answer has some errors. 1 pt for evidence of either applying the definition of expected value, or use of indicator variables.

(Note: both #5(a) and #5(b) were graded out of a max of 5 pts, but total score for #5 was capped at 9 pts.)

(b) Determine how many ways there are to assign a birth date to each of the m people in the room, so that we end up with exactly 3 different birthdays among them all. Show your work.

Solution: Choose a set of three days that will be allowed as birthdays. There are $\binom{365}{3}$ ways to do this.

Now think of labelling each person with the number 1, 2, or 3, according to which birthday he or she has. With any such labelling, we can compute the set of labels that appears; let n_S denote the number of ways to label people so that the set of labels is exactly S. Let's compute $n_{\{1,2,3\}}$. Do it in stages:

- $n_{\{1\}} = 1$: there is only one way to get a labelling where 1 is the only number that appears, namely, each person receives the label 1. Similarly, $n_{\{2\}} = 1$ and $n_{\{3\}} = 1$.

- $n_{\{1,2\}} = 2^m - n_{\{1\}} - n_{\{2\}} = 2^m - 2$: there are 2^m ways to assign each person the label 1 or 2, but this overcounts by including the case where everyone receives the same label (which we do not want to include). So, we subtract $n_{\{1\}} = 1$ for the labelling where everyone gets the label 1, and subtract $n_{\{2\}} = 1$ more for the case where everyone gets the label 2.

  Similarly, $n_{\{1,3\}} = 2^m - 2$ and $n_{\{2,3\}} = 2^m - 2$.
- $n_{\{1,2,3\}} = 3^m - n_{\{1,2\}} - n_{\{1,3\}} - n_{\{2,3\}} - n_{\{1\}} - n_{\{2\}} - n_{\{3\}} = 3^m - 3\times(2^m - 2) - 3 = 3^m - 3\times 2^m + 3$: there are 3^m ways to assign each person a label of 1, 2, or 3, but this overcounts because it includes labellings where not all 3 birthdays appear. Therefore, we subtract to account for all the cases where we end up with less than 3 different birthdays among the m people.

Final answer: $\binom{365}{3} \times n_{\{1,2,3\}} = \binom{365}{3} \times (3^m - 3\times 2^m + 3)$.

Common mistakes: Computing number of ways to end up with *at most* three different birthdays: e.g., $\binom{365}{3} \times 3^m$. But we wanted *exactly* 3, not *at most* 3.

Many computed the answer assuming people are indistinguishable, treating (Alice born on April 5, Bob on June 23) the same as (Alice on June 23, Bob on April 5). This makes the problem easier, but gives the wrong answer.

Grading: 5 pts for correct answer, and correctly treated people as distinct. 3 pts for getting $\binom{365}{3}$ (or $365\times 364\times 363$) in the answer.

For those who (incorrectly) assumed people are indistinguishable: 3 pts for the correct answer of $\binom{365}{3} \times \binom{m-1}{2}$, with explanation (e.g., stars and bars). -1 pt for forgetting the factor of $\binom{365}{3}$. -1 pt if stars and bars reasoning had errors.

(Note: both #5(a) and #5(b) were graded out of a max of 5 pts, but total score for #5 was capped at 9 pts.)

Problem 6. (7 points)

Eric is feeling generous today. Earlier, he secretly chose a number X uniformly at random from the set $\{1,2,3,\ldots,100\}$ and put \$$X$ into one box and \$$2X$ into another (identical) box.

You randomly choose a box by flipping a fair coin. Eric is going to show you what is in the box you chose, and then you will have the option of either taking the box you chose or trading boxes.

(a) Suppose you see \$8 in the box you chose. Given this, what is the probability that the other box contains \$16?

Solution: Let Y denote the value you see. We want to compute $\Pr[X = 8 | Y = 8]$, since if $X = 8$ and we see \$8 then the other box will contain \$16.

Compute: $\Pr[Y = 8 | X = 8] = 1/2$; $\Pr[X = 8] = 1/100$; so $\Pr[Y = 8 \wedge X = 8] = 1/200$. Similarly, $\Pr[Y = 8 \wedge X = 4] = 1/200$. By the sum rule, $\Pr[Y = 8] = \Pr[Y = 8 \wedge X = 8] + \Pr[Y = 8 \wedge X = 4] = 1/200 + 1/200 = 1/100$.

Finally, by the definition of conditional probability, $\Pr[X = 8 | Y = 8] = \Pr[X = 8 \wedge Y = 8]/\Pr[Y = 8] = (1/200)/(1/100) = 1/2$. Final answer: $1/2$.

Common mistakes: Most failed to carefully compute $\Pr[Y = 8]$ or $\Pr[X = 8 | Y = 8]$, and just hand-waved that since there are two possibilities for what is in the other box, both possibilities must have probability $1/2$. That reasoning is broken (it fails if the original distribution on X wasn't uniform), but we were exceptionally generous in the grading and gave this faulty answer full credit even though it is not a very good answer.

Grading: 2 pts for correct solution, which applied the definition of conditional probability. 2 pts for correct answer, which said the other box can contain \$4 or \$16 and claimed both are equally likely (possibly without further justification). 1 pt for correct answer, with a justification that would have (wrongly) also led to the same answer if you saw \$7 instead of \$8. -1 pt for small errors.

(b) Suppose you see \$8 in the box you chose. What is the expected value of your winnings, if you choose to stick with the box you chose?

Solution: \$8.

Grading: 1 pt for correct answer. No partial credit.

(c) Suppose you see \$8 in the box you chose. What is the expected value of your winnings, if you choose to swap boxes? Should you stick or switch?

Solution: Given that we see \$8, the other box contains \$16 with probability $1/2$ (see part (a)), and will contain \$4 with probability $1/2$ (by a similar argument). Thus, the expected value is $\$16 \times 1/2 + \$4 \times 1/2 = \$10$.

Yes, you should switch. Switching increases your expected value, so in the long run it is a better bet. (Put another way: The gain from when the other box contains \$16 exceeds the loss from when the other box contains \$4.)

Grading: 2 pts for correct answer. -1 pt for errors in calculation. Otherwise, no partial credit.

(d) Suppose instead that you see \$7 in your chosen box. Should you stick or switch? Briefly justify your answer.

Solution: Definitely switch. In this case, the only possibility is that $X = 7$, so the other box must contain \$14. ($X = 3.5$ is impossible, since X has to be an integer.)

Common mistakes: Some of you, misreading the problem, thought that there were two possibilities for what is in the other box, \$14 and $\$3\frac{1}{2}$.

Grading: 2 pts for correct answer ("switch") and justification. 1 pt for correct answer ("switch"), but thought other box could contain \$14 or $\$3\frac{1}{2}$.

Problem 7. (6 points)

Let F_k denote the kth Fibonacci number, defined by $F_0 = F_1 = 1$ and $F_k = F_{k-1} + F_{k-2}$. Let $\textsc{ExtEuclid}(x,y)$ denote the result of running the extended Euclidean algorithm on inputs x and y.

What does $\textsc{ExtEuclid}(F_{k+1}, F_k)$ output? Show how you got your result.

Solution: First, recall the definition of $\textsc{ExtEuclid}(x,y)$:

```
ExtEuclid (x,y):
    if y = 0 then return(x, 1, 0)
    else
        (d, a, b) := ExtEuclid (y, x mod y);
        return((d, b, a - (x/y) * b))
        // integer division is used for x/y
```

Note that for $x = F_{k+1}$ and $y = F_k$, $\lfloor x/y \rfloor = \lfloor F_{k+1}/F_k \rfloor = \lfloor (F_k + F_{k-1})/F_k \rfloor = 1$ and $x \mod y = F_{k-1}$.

We start by generating small cases.

EXTEUCLID$(1,0)$ returns $(1,1,0)$.

EXTEUCLID(F_1,F_0) computes $(1,1,0)$, then returns $(1,0,1)$.

EXTEUCLID(F_2,F_1) computes $(1,0,1)$, then returns $(1,1,-1)$.

EXTEUCLID(F_3,F_2) computes $(1,1,-1)$, then returns $(1,-1,2)$.

EXTEUCLID(F_4,F_3) computes $(1,-1,2)$, then returns $(1,2,-3)$.

EXTEUCLID(F_5,F_4) computes $(1,2,-3)$, then returns $(1,-3,5)$.

EXTEUCLID(F_6,F_5) computes $(1,-3,5)$, then returns $(1,5,-8)$.

EXTEUCLID(F_7,F_6) computes $(1,5,-8)$, then returns $(1,-8,13)$.

A pattern emerges: EXTEUCLID$(F_{k+1},F_k) = (1, F_{k-1}(-1)^{k-1} + F_k(-1)^k)$.
Check with $k = 5, F_5 = 8, F_6 = 13, F_4 = 5 : F_6F_4 - F_5F_5 = 65 - 64 = 1$.
Check with $k = 6, F_6 = 13, F_7 = 21, F_5 = 8 : -F_7F_5 + F_6F_6 = -168 + 169 = 1$.

Another formula that works is EXTEUCLID$(F_{k+1},F_k) = (1, F_{k-2}(-1)^k, F_{k-1}(-1)^{k+1})$.
Check with $k = 5, F_5 = 8, F_6 = 13, F_4 = 5, F_3 = 3 : -F_6F_3 + F_5F_4 = -39 + 40 = 1$.
Check with $k = 6, F_6 = 13, F_7 = 21, F_5 = 8, F_4 = 5 : F_7F_4 - F_6F_5 = 105 - 104 = 1$.

Common mistakes: By far, the most common error was to compute only the gcd, forgetting that EXTEUCLID returns a triple. Several of you computed a and b, but *not* the gcd.

Grading: 6 pts for a correct answer. 2 pts for only computing the gcd; 5 pts for computing everything but the gcd. -1 pt for small errors (e.g., off-by-one, getting the sign wrong). The exam definition of the Fibonacci sequence differed slightly from that given in homework and in Rosen (by specifying that $F_0 = 1$ rather than 0); if your solution showed clearly that you were using the alternative definition, you were not penalized for an off-by-one answer.

Problem 8. (12 points)

Mike challenges David to think of a polynomial in one variable— $P(x) = a_nx^n + a_{n-1}x^{n-1} + \cdots + a_1x + a_0$— whose coefficients are all nonnegative integers. Mike claims that he will ask for just two evaluations of the polynomial, and then he will tell David its coefficients. Here's how.

1. Mike asks David to evaluate $P(1)$. He calls that value r.

2. Mike asks David to evaluate $P(r+1)$. After receiving the answer and giving it some thought, he reads off the coefficients of P.

Here's an example. Mike asks for $P(1)$, and David says 9. Mike then asks for $P(10)$. David says 342. Mike correctly identifies P as $3x^2+4x+2$.

(a) Demonstrate that knowing $r = P(1)$ and $y = P(r)$ does *not* provide you with enough information to uniquely determine P.

Solution: This strategy cannot distinguish $Q(x) = 1$ from $Q'(x) = x^k$ for any $k \geq 1$: since $Q(1) = Q'(1) = 1$, we see that $r = 1$ and $y = 1$ no matter whether $P = Q$ or $P = Q'$.

Common mistakes: Claiming that you need at least $n+1$ points to determine a polynomial of degree n (this answer received 0 pts)—see part (b) for why this answer is wrong.

Grading: 2 pts for a correct answer. 1 pt for solutions that mentioned $P(1) = 1$, but used an inappropriate example (e.g. $Q(x) = x^2 - 2x + 2$: negative coefficients weren't allowed in the candidate polynomials) as a polynomial that couldn't be distinguished from $P(1) = 1$.

(b) Explain, in terms detailed enough for another CS 70 student to be able to implement your algorithm immediately, how to determine P efficiently. (Announced at the exam: you are given *only* $r = P(1)$ and $y = P(r+1)$.)

Solution: The idea is that we will express y in base $r+1$. Simply reading off the digits of y reveals the coefficients of P. In pseudocode:

FINDP(r, y):
1. Set $n := 0$.
2. Set $a_n := y \bmod (r+1)$. Set $y := \lfloor y/(r+1) \rfloor$.
3. If $y \neq 0$, set $n := n+1$ and go to step 2.
4. Output $(a_n, \ldots, a_0)$ as the coefficients of P, i.e., $P(x) = a_n x^n + a_{n-1}x^{n-1} + \cdots + a_1 x + a_0$.

Discussion: Note that this does *not* contradict the interpolation theorem we showed in class ("you need $n+1$ points to uniquely determine a polynomial of degree n"). The difference is that the theorem was talking about arithmetic modulo a prime p, whereas this problem is about arithmetic in $\mathbf{N}$. In $\mathbf{N}$, it is sometimes possible for just two points to determine a polynomial of high degree, if those two points are chosen just right.

Grading: 6 pts for a correct and complete answer.

Positive points for solutions earning 3 or less: $+3$ pts for computing n; $+2$ pts for computing a_0; $+1$ pts for saying mod, or computing $P(r+1) - P(1)$.

Some common deductions: -1 pt for small areas of vagueness such as not saying when to stop the loop; -1 pt for modding/dividing by r instead of $r+1$; -2 pts for increasing inefficiency unnecessarily, e.g. by computing x^{r+1} with a loop for each a_i instead of incrementally; -3 pts for grossly increasing inefficiency by checking each way to split $P(1)$ into r pieces (actually everyone who tried this had other errors as well).

(c) Give a big-Oh estimate that's as accurate as possible for the number of arithmetic operations on integers required to execute the algorithm you described in part (b). Justify your answer.

Solution: The algorithm executes $\mathbf{O}(\log_{r+1} y) = \mathbf{O}(\frac{\lg y}{\lg r+1})$ iterations, and in each iteration it performs $\mathbf{O}(1)$ integer operations, for a total of $\mathbf{O}(\frac{\lg y}{\lg r+1})$ arithmetic operations on integers. This is also $\mathbf{O}(n)$.

Grading: 4 pts for a correct answer. If answer to part (b) was incorrect, efficiency estimate had to match the answer to (b).