Massachusetts Institute of Technology
6.042J/18.062J, Fall '05: Mathematics for Computer Science — October 26
Prof. Albert R. Meyer and Prof. Ronitt Rubinfeld — revised November 5, 2005, 1389 minutes

Solutions to Problem Set 6

Problem 1. You've seen how the RSA encryption scheme works, but why is it hard to break? In this problem, you will see that finding secret keys is as hard as finding the prime factorizations of integers. Since there is a general consensus in the crypto community (enough to persuade many large financial institutions, for example) that factoring numbers with a few hundred digits requires astronomical computing resources, we can therefore be sure it will take the same kind of overwhelming effort to find RSA secret keys of a few hundred digits. This means we can be confident the private RSA keys are not somehow revealed by the public keys [1]

For this problem, assume that $n = p \cdot q$ where p, q are both *odd* primes and that e is the public key and d the secret key of the RSA protocol as described in Week 6 Notes. Let $x ::= e \cdot d - 1$.

(a) Show that $\phi(n)$ divides x.

Solution. $ed \equiv 1 \pmod{\phi}(n)$ by definition of d, so $\phi(n)$ divides x by definition of $\equiv$ mod $\phi(n)$. ■

(b) Conclude that 4 divides x.

Solution. Since p, q are odd, both $p - 1$ and $q - 1$ are even. Thus 4 divides $(p-1)(q-1) = \phi(n)$, so by part (a), 4 also divides x. ■

(c) Show that if $\gcd(r, n) = 1$, then $r^x \equiv 1 \pmod{n}$.

Solution. By Euler's Theorem, $r^{\phi(n)} \equiv 1 \pmod{n}$. By part (a), $x = k\phi(n)$ for some integer, k, so

$$r^x = r^{k\phi(n)} = \left(r^{\phi(n)}\right)^k \equiv 1^k \equiv 1 \pmod{n}.$$

■

[1] This is a very weak kind of "security" property, because it doesn't even rule out the possibility of deciphering RSA encoded messages by some method that did not require knowing the secret key. Nevertheless, oveer twenty years experience supports the security of RSA in practice.

A *square root* of m modulo n is a nonnegative integer $s < n$ such that $s^2 \equiv m \pmod{n}$. Here is a nice fact to know: when n is a product of two odd primes, then every number m such that $\gcd(m, n) = 1$ has 4 square roots modulo n.

In particular, the number 1 has four square roots modulo n. The two trivial ones are 1 and $n - 1$ (which is $\equiv -1 \pmod{n}$). The other two are called the *nontrivial* square roots of 1.

(d) Since you know x, then for any integer, r, you can also compute the remainder, y, of $r^{x/2}$ divided by n. So $y^2 \equiv r^x \pmod{n}$. Now if r is relatively prime to n, then y will be a square root of 1 modulo n by part (c).

Show that if y turns out to be a *nontrivial* root of 1 modulo n, then you can factor n. *Hint:* From the fact that $y^2 - 1 = (y + 1)(y - 1)$, show that $y + 1$ must be divisible by exactly one of q and p.

Solution. Since y is a square root of 1 modulo n, we know that n divides $y^2 - 1 = (y + 1)(y - 1)$. So p must divide either $y + 1$ or $y - 1$, and likewise q must divide either $y + 1$ or $y - 1$.

But if y is nontrivial, then $y + 1$ and $y - 1$ are positive and smaller than n, so if $y + 1$ is divisible by p it can't also be divisible by q, and likewise, if it is divisible by q it can't also be divisible by p. So $y + 1$ must be divisible by exactly one of p and q. So $\gcd(y + 1, n)$ must equal p or q. ■

(e) It turns out that at least half the positive integers $r < n$ that are relatively prime to n will yield y's in part (d) that are nontrivial roots of 1. Conclude that if, in addition to n and the public key, e, you also knew the secret key d, then you can be sure of being able to factor n.

Solution. Keep choosing r's at random. Most r's wil be relatively prime to n and at least half of these will yield nontrivial y's in part (d), so you can be sure to turn up the needed nontrivial y in not very many tries. ■

Problem 2. The Massachusetts Turnpike Authority is concerned about the integrity of the new Zakim bridge. Their consulting architect has warned that the bridge may collapse if more than 1000 cars are on it at the same time. The Authority has also been warned by their traffic consultants that the rate of accidents from cars speeding across bridges has been increasing.

Both to lighten traffic and to discourage speeding, the Authority has decided to make the bridge *one-way* and to put tolls at *both* ends of the bridge (don't laugh, this is Massachusetts). So cars will pay tolls both on entering and exiting the bridge, but the tolls will be different. In particular, a car will pay \$3 to enter onto the bridge and will pay \$2 to exit. To be sure that there are never too many cars on the bridge, the Authority will let a car onto the bridge only if the difference between the amount of money currently at the entry toll booth minus the amount at the exit toll booth is strictly less than a certain threshold amount of $\$T_0$.

The consultants have decided to model this scenario with a state machine whose states are triples of natural numbers, (A, B, C), where

- A is an amount of money at the entry booth,
- B is an amount of money at the exit booth, and
- C is a number of cars on the bridge.

Any state with $C > 1000$ is called a *collapsed* state, which the Authority dearly hopes to avoid. There will be no transition out of a collapsed state.

Since the toll booth collectors may need to start off with some amount of money in order to make change, and there may also be some number of "official" cars already on the bridge when it is opened to the public, the consultants must be ready to analyze the system started at *any* state. So let A_0 be the initial number of dollars at the entrance toll booth, B_0 the initial number of dollars at the exit toll booth, and C_0 the number of official cars on the bridge when it is opened. The Authority will be careful to ensure that C_0 is not large enough to cause a collapse. You should assume that even official cars pay tolls on exiting or entering the bridge after the bridge is opened.

(a) Give a mathematical model of the Authority's system for letting cars on and off the bridge by specifying a transition relation between states of the form (A, B, C) above.

Solution. State (A, B, C) goes to state

(i) $(A+3, B, C+1)$, provided that $A - B < T_0$ and $C < 1000$. This transition models the case where a car enters the bridge.

(ii) $(A, B+2, C-1)$, provided that $0 < C \leq 1000$. This transition models the case where a car leaves the bridge.

■

(b) Characterize each of the following derived variables

$$A, B, A+B, A-B, 3C-A, 2A-3B, B+3C, 2A-3B-6C, 2A-2B-3C$$

as one of the following

constant	C
strictly increasing	SI
strictly decreasing	SD
weakly increasing but not constant	WI
weakly decreasing but not constant	WD
none of the above	N

and briefly explain your reasoning.

Solution. In every transition, at least one of A and B increases. So their sum is strictly increasing. $2A - 3B$ can fluctuate, going up on (i) and down on (ii).

The difference $3C - A$ doesn't change under transitions of type (i), but decreases under transitions of type (ii); so is weakly decreasing. Likewise, $B + 3C$ doesn't change under transitions of type (ii), but increases under transitions of type (i); so is weakly increasing.

On the other hand, $6C$ and $2A - 3B$ simultaneously increase by 6 under transition (i) or simultaneously decrease by 6 under transition (ii), which makes their difference constant.

Finally, under (i), $2A$ increases by 6, B is unchanged, and $3C$ increases by 3, so $2A-2B-3C$ increases by $6 - 3 = 3$. However, under (ii), A is unchanged, $3C$ decreases by 3 and $2B$ increases by 4, so $2A - 2B - 3C$ decreases by $-(-4) - 3 = 1$.

The completed table follows.

A	WI
B	WI
$A + B$	SI
$A - B$	N
$3C - A$	WD
$2A - 3B$	N
$B + 3C$	N
$2A - 3B - 6C$	C
$2A - 2B - 3C$	N

■

The Authority has asked their engineering consultants to determine T and to verify that this policy will keep the number of cars from exceeding 1000.

The consultants reason that if A_0 is the initial number of dollars at the entrance toll booth, B_0 is the initial number of dollars at the exit toll booth, and C_0 is the number of official cars on the bridge when it is opened, then an additional $1000 - C_0$ cars can be allowed on the bridge, so as long as $A - B$ has not increased by $3(1000 - C_0)$ there shouldn't more than 1000 cars on the bridge. So they recommend defining

$$T_0 ::= 3(1000 - C_0) + (A_0 - B_0).$$

(c) Use the results of part (b) to define a simple predicate, P, on states of the transition system which is satisfied by the start state, that is $P(A_0, B_0, C_0)$ holds, is not satisfied by any collapsed state, and is an *invariant* of the system. Verify that the P you define has these properties.

Solution. Let $D_0 ::= 2A_0 - 3B_0 - 6C_0$.

Invariant: $[2A - 3B - 6C = D_0] \wedge [C \leq 1000]$.

The invariant obviously holds at the state (A_0, B_0, C_0) because we know that $C_0 \leq 1000$. It does not hold in any collapsed state. To verify the Invariant, assume (A, B, C) satisfies the Invariant and has a transition to (A', B', C'). We check that (A', B', C') satisfies the Invariant by considering the two kinds of transitions.

Transition (i) (a car enters the bridge): so

$$6C' = 6(C+1) = 6C + 6 = (2A - 3B - D_0) + 6 = 2(A+3) - 3B - D_0 = 2A' - 3B' - D_0,$$

which implies that

$$2A' - 3B' - 6C' = D_0, \tag{1}$$

as required.

Also, the transition is possible only if $A - B < T_0$. But this implies

$$\begin{aligned}
6C' &= 2A' - 3B' - D_0 && \text{(by (1))}\\
&= 2(A' - B') - B' - D_0\\
&= 2((A+3) - B) - B - D_0 && \text{(since } A' = A+3,\ B' = B\text{)}\\
&= 2(A - B) - B - D_0 + 6\\
&\leq 2(A - B) - B_0 - D_0 + 6 && \text{(since } B \text{ is WI)}\\
&\leq 2(T_0 - 1) - B_0 - D_0 + 6 && \text{(since } A - B \leq T_0 - 1\text{)}\\
&= 2[3(1000 - C_0) + (A_0 - B_0)] - B_0 - D_0 + 4\\
&= 6000 - 6C_0 + 2A_0 - 3B_0 - D_0 + 4\\
&= 6004,
\end{aligned}$$

and so $C' \leq \lfloor 6004/6 \rfloor = 1000$, as required.

Transition (ii) (a car leaves the bridge): so

$$6C' = 6(C-1) = 6C - 6 = 2A - 3B - 6 = 2A - 3(B+3) = 2A' - 3B'.$$

In addition, $C' < C \leq 1000$ so $C' \leq 1000$. ■

(d) A clever MIT intern working for the Turnpike Authority agrees that the Turnpike's bridge management policy will be *safe*: the bridge will not collapse. But she warns her boss that the policy will lead to *deadlock*— a situation where traffic can't move on the bridge even though the bridge has not collapsed.

Explain more precisely in terms of system transitions what the intern means, and briefly, but clearly, justify her claim.

Solution. The intern means that any sequence of transitions will arrive at a state in which no transition is possible, even though there are no cars on the bridge. This happens because every time a car enters and then exits the bridge the value of $A - B$ increases by 1. So after 3000 cars have crossed the bridge, no further car can enter the bridge because

$$A - B \geq 3000 + A_0 - B_0 \geq 3(1000 - C_0) + (A_0 - B_0) = T_0.$$

After that, cars can only exit the bridge. So after at most 3000+1000 transitions, the system deadlocks with the bridge empty but no cars allowed onto the bridge. ■

Problem 3. Vertices u, v in a digraph are said to be *unconnected* when there is no path either from u to v or from v to u. The following procedure can be applied to any digraph, G:

Pick two vertices u, v such that either

1. there is an edge (u, v) of G and there is also a path from u to v which does *not* include this edge; in this case, delete the edge (u, v), or
2. u and v are unconnected; in this case, add the edge (u, v).

Repeat these operations until it is no longer possible to find vertices u, v to which an operation applies.

This procedure can be modelled as a state machine. The start state is G, and the states are all possible digraphs with the same vertices as G. The final states are the digraphs on which no operation is possible.

(a) For any state, G, let e be its number of edges, and p its number of pairs of unconnected vertices. Define a decreasing natural number valued derived variable that is a function of e and p. Conclude that the procedure terminates started on any finite digraph, G.

Solution. Let G be a state and G' be the resulting state after a transition. We show this derived variable is strictly decreasing by showing $2p + e$ of G' is less than $2p + e$ of G.

Suppose G' is the result of the first transition. Clearly, e decreases by 1. Thus, all we need to show is that p never increases (we do not care if it decreases or remains the same). We show this by showing if two arbitrary vertices u and v are connected in G, then they are connected in G'. Consider the connected vertices u and v. Call the deleted edge $x - y$ going from some vertex x to some vertex y. The only way for u and v to not be connected in G' is if *all* paths between the vertices contain the edge $x - y$. However, this cannot be because for each path containing $x - y$, there is also a path containing the alternate path from x to y that we know must exist (by the condition which we deleted $x - y$).

Now suppose G' is the result of the second transition. Clearly, adding an edge cannot cause two vertices that were previously connected to become unconnected. By the condition stated in the transition, we know we have connected two vertices that were previously unconnected. Thus, because no pairs of connected vertices may become unconnected, and because we *know* of a pair of vertices that were unconnected are becoming connected, we can conclude p decreases by *at least* one, so $2p$ decreases by at least 2. Yet, e only increases by exactly 1. Thus, in this case, $2p + e$ is decreasing as well. ■

(b) Prove that the set of final states reachable from DAG start states are the *line graphs*.

Solution. First, we observe every line graph is a final state. This follows because there are no pairs of unconnected vertices and there is some path between every pair of vertices.

Second, we show that if a digraph is not a line graph, there must be some transition possible. If some pair of vertices are not connected, then a transition is possible, so we may assume all are connected. Suppose there is a vertex, u, with out-degree at least two. So u has edges going to $v \neq w$. But v and w are connected, so without loss of generality, we may assume a path goes from w to v. But then a transition removing edge (u, v) is possible. A similar argument shows that a transition is possible if there is a vertex with in-degree at least two. So every vertex has out-degree and in-degree at most 1 and every pair of vertices is connected. This implies the graph is a line graph, as the reader may verify. ■

(c) Prove that the property of being a DAG is an invariant of this procedure.

Solution. To show the property of being a DAG is invariant, we show if G is a DAG, and G' is the result of applying one of the two operations 1 and 2 described above to G, then G' is a DAG.

So suppose G' is the result of applying operation 1. Since removing an edge cannot create a cycle, G' remains a DAG.

Next, suppose G' is the result of applying operation 2. We prove G' is acyclic by contradiction. Assume some cycle exists. The edge (u, v) must be an edge in the cycle, otherwise the cycle would have existed in G as well. Let u be the start vertex of the cycle (we can pick the start vertex arbitrarily from any vertices in the cycle). Let the first edge be (u, v).

Note there must be a path from v to u to complete the cycle, thus there must have been a path from v to u in the graph G. This contradicts the condition on which we added an edge from u to v. ■

(d) Prove that if G is a DAG, the procedure terminates with a line graph whose path relation is a topological sort of the partial order defined by G. *Hint:* Strengthen the DAG invariant in the previous part.

Solution. The invariant is that the current DAG is a refinement of the starting DAG, where relation R_2 is a *refinement* of relation R_1 iff they have the same domain and codomain, and xR_1y implies xR_2y for all x, y. ■

Problem 4. (a) Give an example of a stable match between 3 boys and 3 girls where no person gets their first choice.

Solution. Call the boys $1, 2, 3$ and the girls a, b, c. Consider the following preference list:

choice	1st	2nd	3rd	choice	1st	2nd	3rd
1	a	b	c	a	2	3	1
2	b	c	a	b	3	1	2
3	c	a	b	c	1	2	3

The matching $(1, b), (2, c), (3, a)$ is stable even though no person gets their first choice.

To see the intuition behind this solution, notice first that the first choice of any boy has that boy as her last choice and vice versa. Second, notice that everyone ends up with their second choice.

Since we show a pairing where everyone has their second choice, this is stable because the only way to have a rogue pair is for a boy or girl to want their first choice, but their first choice always likes them least so will never want to leave their current partner. Therefore, we end up with a stable pairing where no one gets their first choice. ■

(b) Describe a simple procedure to determine whether or not a stable marriage problem has a unique solution, that is, only one possible stable marriage assignment.

Solution. See if the Mating algorithm with Boys as suitors yields the same solution as the algorithm with Girls as suitors. These two marriage assignments are boy-optimal and girl-optimal, respectively, so they agree iff there is a unique solution. ■

Problem 5. A Harvard BS graduates and starts with an annual salary of $140,000, with a $25,000 raise guaranteed every year. An MIT SB graduate starts with $100,000, with a guaranteed 15% raise every year. Assume the bankrate is a fixed 3% per year. That is, the bank will pay $1.03 a year from now if you deposit $1.00 today.

(a) Suppose both graduates retire after the same number of years. Use the fact that $x = o((1+\epsilon)^x)$ to explain why the MIT SB must come out ahead if they work for enough years. (You should *not* make use of closed forms for various sums in your explanation.)

Solution. The Harvard graduate only gets a fixed annual income raise, so his income grows linearly, *i.e.*, $\Theta(n)$ where n is the number of years worked. The MIT graduate's raise is based on a certain percentage of his annual income. Therefore, his income grows exponentially, *i.e.*, $\Theta(1.15^n)$. Since $n = o(1.15^n)$, we conclude that the MIT grad's salary will eventually grow and stay larger than the Harvard grad's by more than any fixed constant factor. But anyway, once we realize his salary stays larger, it's clear he will eventually catch up and then exceed the total salary paid the Harvard grad. The main question is whether this will occur in their lifetimes!

We have not considered the diminished present value of salary paid years from now. As long the bankrate, b, remains lower than the 1.15 factor by which the MIT grad's salary increases annually, then the present value of the year n salary still grows exponentially, though more slowly than $\Theta(1.15^n)$. Namely its growth is $\Theta((1+\epsilon)^n)$ where $1+\epsilon ::= 1.15/b$. So the argument of the preceding paragraph still holds. We can safely assume the MIT grad is savvy enough not to accept a job whose annual percent increase was lower than the bankrate. ■

(b) Suppose both graduates retire after n years. For which values of n is the MIT graduate's salary package better that the Harvard grad's?

Solution. One dollar after year i is worth r^i in today's currency, where

$$r ::= \frac{1}{1.03} = 0.970\,873\ \ldots.$$

So

$$\begin{aligned}
\text{Hvd}_n &= \sum_{i=0}^{n}(140000 + 25000i)r^i \\
&= 140000 \cdot \sum_{i=0}^{n} r^i + 25000 \cdot \sum_{i=0}^{n} ir^i, \\
\text{MIT}_n &= 100000 \cdot \sum_{i=0}^{n} 1.15^i r^i \\
&= 100000 \cdot \sum_{i=0}^{n} (1.15r)^i
\end{aligned}$$

But

$$\sum_{i=0}^{n} ir^i = \frac{r-(n+1)r^{n+1}+nr^{n+2}}{(1-r)^2},$$

so

$$\begin{aligned}
\text{Hvd}_n &= 140000\frac{(1-r^{n+1})}{1-r}+25000\frac{(r-(n+1)r^{n+1}+nr^{n+2})}{(1-r)^2} \\
&= \frac{1000\cdot(140\cdot(1-r-r^{n+1}+r^{n+2})+25\cdot(r-(n+1)r^{n+1}+nr^{n+2}))}{(1-r)^2} \\
&= \frac{1000\cdot(29.2r-(165+25n)\cdot r^{n+1}+(140+25n)\cdot r^{n+2})}{(1-r)^2} \\
\text{MIT}_n &= \frac{100000(1-(1.15r)^{n+1})}{1-1.15r} \\
&= \frac{103000}{0.12}\cdot((1.15r)^{n+1}-1)
\end{aligned}$$

and for $n = 15$,

$$\begin{aligned}
\text{Hvd}_{15} &= \frac{1000\cdot(29.2r-(165+25\cdot 15)\cdot r^{16}+(140+25\cdot 15)\cdot r^{17})}{(1-r)^2} = 4,034,764 \\
\text{MIT}_{15} &= \frac{103000}{0.12}\cdot((1.15r)^{16}-1) = 4,146,917.
\end{aligned}$$

Thus the the overall income of the MIT graduate is less the first 14 years and more after the 15th year when compared to the Harvard graduate! (Longterm, a MIT degree is worth more – but we knew *that* already.) ■

Problem 6. Books Books and more Books! If the 6.042 staff is to stand a chance at the Book Extension Stacking Challenge, we have to consider all the angles!

Recall the basic book stacking challenge from the course notes where you have an unlimited supply of books to stack that are all the same weight.

(a) What if instead of all books weighing the same, you have a book that weighs 1 pound, a book that weighs $\frac{1}{2}$ pounds, a book that weighs $\frac{1}{4}$ pounds, where each successive book weighs half as much as the previous book. Say you had n such books, and also that the you have a duplicate of the lightest book. How far out can you stack the books? Note that all books are still the same size, just different weights. *Hint:* Where should the heaviest books be?

Solution. The lightest books will be on top. Notice that every time you add a heavier book to the bottom, it's weight will equal the weight of all the books already in the stack. Thus, you can model this problem as stacking two books of equal weight.

Every time you add a heavier book to the bottom, you can extend the stack out by the same constant amount. The top of the stack will have its center of gravity at the edge of the new bottom book, and the whole stack will have its center of gravity at the edge of the table. Thus, the bottom book will stick out 1/4 of its length past the edge of the table. Every book you add will add 1/4 of a book length to the extension of the stack! ■

(b) What if you had to stack such that the lightest books were on the bottom of the stack and the heaviest books were on top of the stack. How far out can you stack an infinite number of books where each book is twice as heavy as the book below it (we're looking for either infinitely far or finitely far)? Justify your answer.

Solution. Books can only be stacked finitely far. Start with the original book of weight 1. Note that this book weighs as much as the rest of the books combined. For contradiction, assume that its center of gravity could be more than a half a book length past the edge of the table. We know that the center of gravity of the whole system must be before the edge of the table, so we know that at least as much weight must be strictly half a book length behind the edge of the table (since the total weight available is only equal to the weight of the first book). However, that would require that all the other books are fully on the table without overhang, and the top book can't possibly be on top of another book. ■

(c) What if the books were Harmonically weighted: $1, \frac{1}{2}, \frac{1}{3}$, *etc.*, *etc.*, and the heaviest book had to be on top. Would it be possible for the top of the stack to be arbitrarily far past the edge of the table?

Solution. Yes.

We know that if all books weighed the same that it would be possible to stack infinitely far past the end of the table. We also know that we can stack books on top of eachother until they way an arbitrary amount because the Harmonic Series diverges. Therefore, we can stack enough books together such that each stack weighs at least twice as much as the previous stack. Treating each stack like a book, part a of this problem illustrates how you can stack arbitrarily far out from the edge of the table. ■

Problem 7. Use the integral method to find upper and lower bounds for the following summation that differ by at most 0.05.

$$\sum_{i=1}^{\infty} \frac{1}{i^3}$$

Hint: Try adding the first few terms explicitly and then use integrals to bound the sum of the remaining terms.

Solution. We can bound the summation above as follows:

$$\begin{aligned}
\sum_{i=1}^{\infty} \frac{1}{i^3} \quad &\leq \quad \frac{1}{1} + \frac{1}{8} + \frac{1}{27} + \int_4^{\infty} \frac{1}{(x-1)^3}\,dx \\
&= \frac{1}{1} + \frac{1}{8} + \frac{1}{27} + \left(-\frac{1}{2\cdot(x-1)^2}\right)_4^{\infty} \\
&= \frac{1}{1} + \frac{1}{8} + \frac{1}{27} + \frac{1}{18} \\
&= 1.2176\ldots
\end{aligned}$$

We can bound the summation below similarly:

$$\begin{aligned}
\sum_{i=1}^{\infty} \frac{1}{i^3} \quad &\geq \quad \frac{1}{1} + \frac{1}{8} + \frac{1}{27} + \int_4^{\infty} \frac{1}{x^3}\,dx \\
&= \quad \frac{1}{1} + \frac{1}{8} + \frac{1}{27} + \left(-\frac{1}{2x^2}\right)_4^{\infty} \\
&= \quad \frac{1}{1} + \frac{1}{8} + \frac{1}{27} + \frac{1}{32} \\
&= \quad 1.1933\ldots
\end{aligned}$$

■

Problem 8. **(a)** Given that $f(x) = O(g(x))$, prove that $f(x)^2 = O(g(x)^2)$

Solution. Since $f(x) = O(g(x))$,

$$\exists x_0, c\ \forall x \geq x_0, |f(x)| \leq c\,|g(x)|\,.$$

So, $\forall x \geq x_0\,[f(x)^2 \leq c^2 g(x)^2]$. Therefore, there exist $x_0' ::= x_0$ and $c' ::= c^2$ such that $\forall x \geq x_0'\,[f(x)^2 \leq c' g(x)^2]$. So $f(x)^2 = O(g(x)^2)$. ■

(b) Let $f(x) ::= 2x$ and $g(x) ::= x$, so $f(x) = O(g(x))$. Prove that $2^{g(x)} = o(2^{f(x)})$, so $2^{f(x)} \neq O(2^{g(x)})$.

Solution. Now $2^{f(x)} = 2^{2x} = 4^x$, and $2^{g(x)} = 2^x$. So $2^{g(x)}/2^{f(x)} = 2^x/4^x = 2^{-x}$, so the ratio goes to zero as x goes to infinity, proving that $2^{g(x)} = o(2^{f(x)})$. But $h_1 = o(h_2)$ implies $h_2 \neq O(h_1)$ for any functions h_1, h_2, so $2^{f(x)} \neq O(2^{g(x)})$. ■