Solutions to Problem Set 7

Problem 29: Coin-flipping

(a) Assume y is a square both mod p and mod q, i.e. $b_p^2 \equiv y \pmod{p}$ and $b_q^2 \equiv y \pmod{q}$. Since $\gcd(p, q) = 1$, according to Chinese Remainder Theorem, there exists exactly one solution b so that $b \equiv b_p \pmod{p}$ and $b \equiv b_q \pmod{q}$. So $b^2 \equiv y \pmod{p}$ and $b^2 \equiv y \pmod{q}$. Thus $p|(b^2 - y)$ and $q|(b^2 - y)$, which means $pq|(b^2 - y)$. So we conclude that $b^2 \equiv y \pmod{n}$, which contradicts the fact that y doesn't have a square root mod n. The same proof applies to $-y$. So neither y nor $-y$ can be a square both mod p and mod q.

(b) Since $q \equiv 3 \pmod{4}$, then $q = 4m + 3$ and $(q - 1)/2 = 2m + 1$ is odd, where m is an integer. So $(-1)^{\frac{q-1}{2}} \equiv -1 \pmod{q}$. And because y is not a square mod q, according to Euler Criterion, $y^{\frac{q-1}{2}} \not\equiv -1 \pmod{q}$. But $y^{q-1} \equiv 1 \pmod{q}$ and 1's two square roots are ± 1, so $y^{\frac{q-1}{2}} \equiv -1 \pmod{q}$. Thus we have

$$(-y)^{\frac{q-1}{2}} \equiv (-1)^{\frac{q-1}{2}} y^{\frac{q-1}{2}} \equiv 1 \pmod{q}.$$

Again according to Euler Criterion, we know that $-y$ is a square mod q.

(c) The result immediately follows from (a) and (b). To be specific, from (a), neither y and $-y$ can be a square both mod p and mod q. We discuss it in two cases:

Case 1: If y is not a square mod p, then according to b, $-y$ must be a square mod p. And from (a) again, we know $-y$ can't be a square mod q at the same time, so y must be a square mod q according to (b).

Case 2: If y is not a square mod q, then according to b, $-y$ must be a square mod q. And from (a) again, we know $-y$ can't be a square mod p at the same time, so y must be a square mod p according to (b).

From the above description, we know that y can't be a quadratic non-residue of both p and q. So the above description is sufficent and the proposition is proven.

(d)

$$\begin{aligned} b^2 \equiv y \pmod{p} &\Rightarrow p|(b^2 - y) \\ b^2 \equiv -y \pmod{q} &\Rightarrow q|(b^2 + y), \end{aligned}$$

so we have $pq|(b^2 - y)(b^2 + y)$, i.e. $n|(b^4 - y^2)$. Since $b^2 \not\equiv \pm y \pmod{n}$, otherwise, it contradicts that neither y nor $-y$ has a square root mod n, according to The Quadratic Sieve, $\gcd(b^2 - y, n)$ gives a nontrivial factor of n, i.e. p or q. Hence Bob successfully factors n and claims victory.

Problem 30: Indistinguishability

Since U_n is the uniform distribution on binary strings of length n, U_n can be regarded as coming from n flips of a fair coin. And the only way we know so far to distinguish two probability distributions is to predict their results and find the difference on their prediction accuracy. So the following solutions are both seeking the proper prediction methods under this knowledge.

(a) $J(z)$ tries to distinguishes U_1 from X_1 by predicting their outputs, i.e. 1("heads") or 0("tails"). Since U_1 comes from a flip of a fair coin, predicting either 1 or 0 provides the same success/failure probability of $1/2$. And for X_1 using a biased coin, predicting 1 provides the maximum success probability of $2/3$ while predicting 0 provides the maximum failure probability of $1 - 1/3 = 2/3$. So $J(z)$ can output 1 either to indicate success prediction or failure prediction, and both generate the same distinguishibility:

$$|\Pr[J(U_1) = 1] - \Pr[J(X_1) = 1]| = 2/3 - 1/2 = 1/6.$$

And this is the largest value for ϵ using prediction method for $J(z)$.

(b) When we want to distinguish U_n from X_n, we can focus on some specific patterns of strings, or focus on general statistics. Since specific patterns appear with low probability especially when $n \to \infty$, we focus on the latter. Consider the number of 1 which appear in the two strings. Assume $J(z)$ will output 1 if the number of 1 is no less than m, in which $0 \leq m \leq n$, then we have

$$\begin{aligned} \Pr[J(U_n) = 1] &= \sum_{i=m}^{n} \mathrm{C}_n^i \left(\frac{1}{2}\right)^n \\ \Pr[J(X_n) = 1] &= \sum_{i=m}^{n} \mathrm{C}_n^i \left(\frac{2}{3}\right)^i \left(\frac{1}{3}\right)^{n-i}. \end{aligned}$$

So we obtain

$$|\Pr[J(U_n) = 1] - \Pr[J(X_n) = 1]| = \left| \sum_{i=m}^{n} \mathrm{C}_n^i \left(\frac{2^i}{3^n} - \frac{1}{2^n}\right) \right|.$$

We get rid of negative items in the series so as to maximize the above value, i.e. $\frac{2^m}{3^n} - \frac{1}{2^n} > 0$, which induces $m > \log_2(3n/2) \approx 0.6n$. Figure 1 shows that ϵ almost strictly increases when n increases. And the minimum value is $1/6$ when $n = 1$ as in (a). Actually, when $n \to \infty$, $\Pr[J(U_n) = 1] \to 0$ while $\Pr[J(X_n) = 1] \to 1$. This can be proven using Central Limit Theorem, whose details will be skipped here.

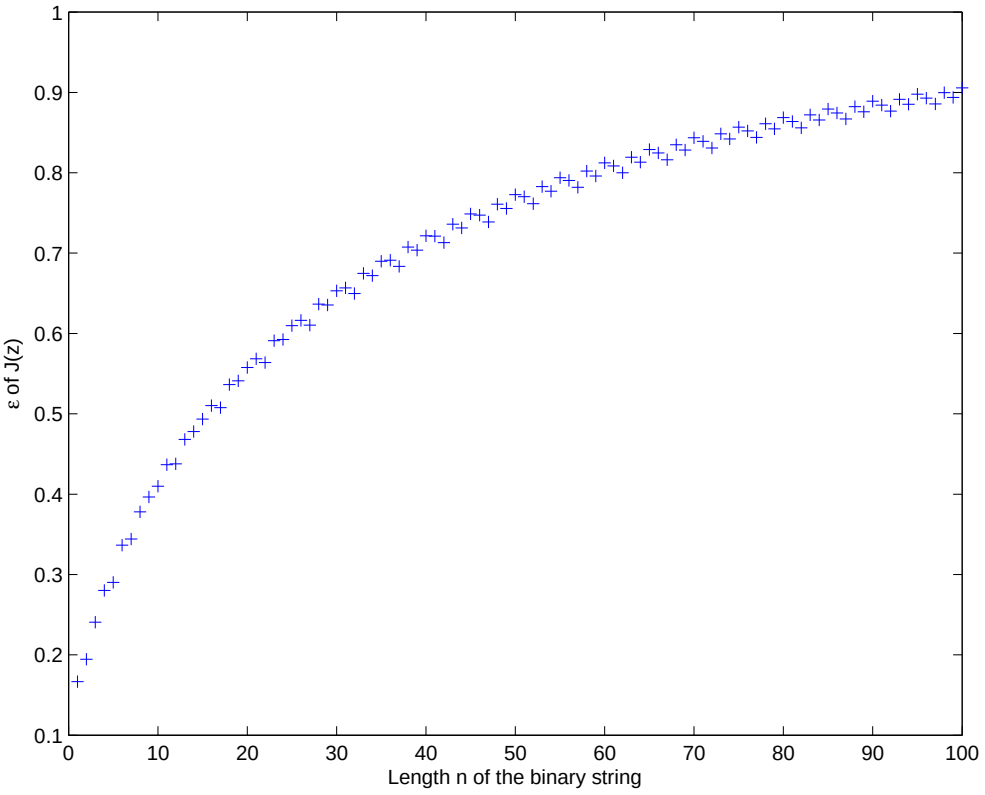

Figure 1: Degree ϵ by which $J(z)$ distinguishes U_n from X_n as a function of n.