HW1 Solutions

With thanks to Adrian Mettler and Li Zhuang for some solutions

1 Long Answer: Computational indistinguishability

Let $\mathcal{A}_R$ denote the set of algorithms running with resources R. Then

$$d_R(D, D') = \max_{A \in \mathcal{A}_R} \text{Adv}\, A = \max_{A \in \mathcal{A}_R} \left| \Pr[X \stackrel{\$}{\leftarrow} D; A(X) = 1] - \Pr[X' \stackrel{\$}{\leftarrow} D'; A(X') = 1] \right|.$$

(a) *Data Processing Inequality.* Let L be any (possibly randomized) algorithm running with resources R', and let $L(D)$ denote the distribution obtained by picking x from D and then outputting $L(x)$. Prove $d_R(L(D), L(D')) \leq D_{R'+R}(D, D')$.

$$\begin{aligned} d_R(L(D), L(D')) &= \max_{A \in \mathcal{A}_R} \left| \Pr[X \stackrel{\$}{\leftarrow} L(D); A(X) = 1] - \Pr[X' \stackrel{\$}{\leftarrow} L(D'); A(X') = 1] \right| \\ &= \max_{A \in \mathcal{A}_R} \left| \Pr[X \stackrel{\$}{\leftarrow} D; A(L(X)) = 1] - \Pr[X' \stackrel{\$}{\leftarrow} D'; A(L(X')) = 1] \right| \\ &\leq \max_{B \in \mathcal{A}_{R+R'}} \left| \Pr[X \stackrel{\$}{\leftarrow} D; B(X) = 1] - \Pr[X' \stackrel{\$}{\leftarrow} D'; B(X') = 1] \right| \\ &= d_{R+R'}(D, D') \end{aligned}$$

using the fact that $B = A \circ L$ is in $\mathcal{A}_{R+R'}$ if A is in $\mathcal{A}_R$.

(b) *Triangle Inequality.* Prove $d_R(D, D'') \leq d_R(D, D') + d_R(D', D'')$ for any distributions D, D', D''.

Note first that $|\alpha + \beta| \leq |\alpha| + |\beta|$ always holds (the triangle inequality for real numbers). Now

$$\begin{aligned} d_R(D, D'') &= \max_{A \in \mathcal{A}_R} \left| \Pr[X \stackrel{\$}{\leftarrow} D; A(X) = 1] - \Pr[X'' \stackrel{\$}{\leftarrow} D''; A(X'') = 1] \right| \\ &= \max_{A \in \mathcal{A}_R} \Big| \Pr[X \stackrel{\$}{\leftarrow} D; A(X) = 1] - \Pr[X' \stackrel{\$}{\leftarrow} D'; A(X') = 1] \\ &\qquad + \Pr[X' \stackrel{\$}{\leftarrow} D'; A(X') = 1] - \Pr[X'' \stackrel{\$}{\leftarrow} D''; A(X'') = 1] \Big| \\ &\leq \max_{A \in \mathcal{A}_R} \left| \Pr[X \stackrel{\$}{\leftarrow} D; A(X) = 1] - \Pr[X' \stackrel{\$}{\leftarrow} D'; A(X') = 1] \right| \\ &\qquad + \left| \Pr[X' \stackrel{\$}{\leftarrow} D'; A(X') = 1] - \Pr[X'' \stackrel{\$}{\leftarrow} D''; A(X'') = 1] \right| \\ &\leq \max_{A \in \mathcal{A}_R} \left| \Pr[X \stackrel{\$}{\leftarrow} D; A(X) = 1] - \Pr[X' \stackrel{\$}{\leftarrow} D'; A(X') = 1] \right| \\ &\quad + \max_{A \in \mathcal{A}_R} \left| \Pr[X' \stackrel{\$}{\leftarrow} D'; A(X') = 1] - \Pr[X'' \stackrel{\$}{\leftarrow} D''; A(X'') = 1] \right| \\ &= d_R(D, D') + d_R(D', D'') \end{aligned}$$

(c) *Multiple Samples.* Prove $d_R(D \times D, D' \times D') \leq 2 d_R(D, D')$.

Let L be a randomized algorithm that, on input x, picks a random $y \stackrel{\$}{\leftarrow} D$ and outputs (x, y). We can assume L takes $O(1)$ resources. Note that $L(D) = D \times D$, and $L(D') = D' \times D$. By the data processing inequality (part (a)),

$$d_{R+O(1)}(D \times D, D' \times D) = d_{R+O(1)}(L(D), L(D')) \leq d_R(D, D').$$

Similarly,

$$d_{R+O(1)}(D' \times D, D' \times D') \leq d_R(D, D').$$

Applying the triangle inequality (part (b)), we get $d_R(D \times D, D' \times D') \leq d_{R+O(1)}(D \times D, D' \times D) + d_{R+O(1)}(D' \times D, D' \times D') \leq 2d_R(D, D')$.

(d) *Variation Distance.* The variation distance is $V(D_1, D_2) \equiv \frac{1}{2}\sum_x |\Pr_{D_1}[x] - \Pr_{D_2}[x]|$. Prove $d_\infty(D, D') = V(D, D')$.

(Recall: $\Pr_{D_1}[x]$ is short-hand for $\Pr[X \stackrel{\$}{\leftarrow} D_1; X = x]$.)

Let $f(x) = \Pr_{D_1}[x] - \Pr_{D_2}[x]$. Of course, $\sum_x f(x) = \sum_x \Pr_{D_1}[x] - \sum_x \Pr_{D_2}[x] = 1 - 1 = 0$. Let S^+ denote the set $\{x : f(x) \geq 0\}$, and S^- the complementary set $\{x : f(x) < 0\}$. Then $\sum_{x \in S^+} f(x) + \sum_{x \in S^-} f(x) = \sum_x f(x) = 0$, so $\sum_{x \in S^+} f(x) = \sum_{x \in S^-} |f(x)|$. Now we're ready to calculate. We have

$$\begin{aligned}
d_\infty(D_1, D_2) &= \max_{A \in \mathcal{A}_R} \left| \Pr[X \stackrel{\$}{\leftarrow} D_1; A(X) = 1] - \Pr[X \stackrel{\$}{\leftarrow} D_2; A(X) = 1] \right| \\
&= \max_{A \in \mathcal{A}_R} \left| \sum_x \Pr[A(x) = 1] \Pr_{D_1}[x] - \sum_x \Pr[A(x) = 1] \Pr_{D_2}[x] \right| \\
&= \max_{A \in \mathcal{A}_R} \left| \sum_x \Pr[A(x) = 1] \cdot f(x) \right| \\
&= \max_{A \in \mathcal{A}_R} \left| \sum_{x \in S^+} \Pr[A(x) = 1] \cdot f(x) - \sum_{x \in S^-} \Pr[A(x) = 1] \cdot |f(x)| \right|
\end{aligned}$$

Notice that all the terms above are positive. Thus, the maximum is attained either by $A^+(x) =$ if $x \in S^+$ then 1 else 0 or by $A^-(x) =$ if $x \in S^-$ then 1 else 0. In fact, it doesn't matter which you choose, since $\mathrm{Adv}\, A^+ = \sum_{x \in S^+} f(x) = \sum_{x \in S^-} |f(x)| = \mathrm{Adv}\, A^-$, so we might as well choose A^+. Then

$$\begin{aligned}
d_\infty(D_1, D_2) &= \mathrm{Adv}\, A^+ \\
&= \sum_{x \in S^+} f(x) \\
&= \frac{1}{2} \sum_{x \in S^+} f(x) + \frac{1}{2} \sum_{x \in S^-} |f(x)| \\
&= \frac{1}{2} \sum_x |f(x)| \\
&= V(D_1, D_2).
\end{aligned}$$

(e) *Information Theoretic Security.* Prove $d_R(D, D') \le d_\infty(D, D')$.

This one is easy. We use the fact that $\max_{x \in S} f(x) \le \max_{x \in T} f(x)$ whenever $S \subseteq T$. Thus:

$$d_R(D_1, D_2) = \max_{A \in \mathcal{A}_R} \operatorname{Adv} A \le \max_{A \in \mathcal{A}_\infty} \operatorname{Adv} A = d_\infty(D_1, D_2).$$

(f) *Conditioning.* Let $D|E$ represent the distribution of D conditioned on the event E occuring. Prove $d_R(D, D') \le d_R(D|E, D'|E) + \Pr[\neg E]$.

Consider any adversary A using resources R. Then

$$\begin{aligned}
\operatorname{Adv} A &= \Big|\Pr[X \xleftarrow{\$} D; A(X) = 1] - \Pr[X' \xleftarrow{\$} D; A(X') = 1]\Big| \\
&= \Big|\Pr[X \xleftarrow{\$} D; A(X) = 1|E] \Pr[E] + \Pr[X \xleftarrow{\$} D; A(X) = 1|\neg E] \Pr[\neg E] \\
&\quad - \Pr[X' \xleftarrow{\$} D; A(X') = 1|E] \Pr[E] - \Pr[X' \xleftarrow{\$} D; A(X') = 1|\neg E] \Pr[\neg E]\Big| \\
&= \Big|(\Pr[X \xleftarrow{\$} D; A(X) = 1|E] - \Pr[X' \xleftarrow{\$} D; A(X') = 1|E]) \cdot \Pr[E] \\
&\quad + (\Pr[X \xleftarrow{\$} D; A(X) = 1|\neg E] - \Pr[X' \xleftarrow{\$} D; A(X') = 1|\neg E]) \cdot \Pr[\neg E]\Big| \\
&\le \Big|(\Pr[X \xleftarrow{\$} D; A(X) = 1|E] - \Pr[X' \xleftarrow{\$} D; A(X') = 1|E]) \cdot \Pr[E]\Big| \\
&\quad + \Big|(\Pr[X \xleftarrow{\$} D; A(X) = 1|\neg E] - \Pr[X' \xleftarrow{\$} D; A(X') = 1|\neg E]) \cdot \Pr[\neg E]\Big| \\
&= \underbrace{\Big|\Pr[X \xleftarrow{\$} D; A(X) = 1|E] - \Pr[X' \xleftarrow{\$} D; A(X') = 1|E]\Big|}_{\le d_R(D|E, D'|E)} \cdot \underbrace{\Pr[E]}_{\le 1} \\
&\quad + \underbrace{\Big|\Pr[X \xleftarrow{\$} D; A(X) = 1|\neg E] - \Pr[X' \xleftarrow{\$} D; A(X') = 1|\neg E]\Big|}_{\le 1} \cdot \Pr[\neg E] \\
&\le d_R(D|E, D'|E) + \Pr[\neg E].
\end{aligned}$$

Therefore

$$d_R(D, D') = \max_{A \in \mathcal{A}_R} \operatorname{Adv} A \le d_R(D|E, D'|E) + \Pr[\neg E].$$

Remark. This corresponds to "conditioning on the absence of the bad event." Commonly, we'll let E be the event that the Bad flag does not get set. In the special case where $D|E \equiv D'|E$ (i.e., if the Bad flag wasn't set, then the distributions are identical), we learn that $d(D, D') \le \Pr[\neg E]$, i.e., $d(D, D') \le \Pr[\text{Bad flag gets set}]$.

(g) *An XOR Lemma.* Prove $d_R(D \oplus D', U) \le 2 d_\infty(D, U) d_\infty(D', U)$.

We'll use part (d) and a judicious application of the triangle inequality:

$$\begin{aligned}
&2d_\infty(D,U)d_\infty(D',U)\\
&= 2\cdot\left(\frac{1}{2}V(D,U)\right)\cdot\left(\frac{1}{2}V(D',U)\right)\\
&= \frac{1}{2}\left(\sum_{x_1}\left|\Pr_D[x_1]-\frac{1}{2^k}\right|\right)\left(\sum_{x_2}\left|\Pr_{D'}[x_2]-\frac{1}{2^k}\right|\right)\\
&= \frac{1}{2}\sum_{(x_1,x_2)}\left|\Pr_D[x_1]-\frac{1}{2^k}\right|\left|\Pr_{D'}[x_2]-\frac{1}{2^k}\right|\\
&= \frac{1}{2}\sum_{(x_1,x_2)}\left|\Pr_D[x_1]\Pr_{D'}[x_2]-\frac{1}{2^k}\Pr_D[x_1]-\frac{1}{2^k}\Pr_{D'}[x_2]+\frac{1}{2^{2k}}\right|\\
&= \frac{1}{2}\sum_x\sum_{\substack{(x_1,x_2)\\ x_1\oplus x_2=x}}\left|\Pr_D[x_1]\Pr_{D'}[x_2]-\frac{1}{2^k}\Pr_D[x_1]-\frac{1}{2^k}\Pr_{D'}[x_2]+\frac{1}{2^{2k}}\right|\\
&\geq \frac{1}{2}\sum_x\left|\sum_{\substack{(x_1,x_2)\\ x_1\oplus x_2=x}}\Pr_D[x_1]\Pr_{D'}[x_2]-\frac{1}{2^k}\Pr_D[x_1]-\frac{1}{2^k}\Pr_{D'}[x_2]+\frac{1}{2^{2k}}\right|\\
&= \frac{1}{2}\sum_x\left|\sum_{\substack{(x_1,x_2)\\ x_1\oplus x_2=x}}\Pr_D[x_1]\Pr_{D'}[x_2]-\frac{1}{2^k}\underbrace{\sum_{\substack{(x_1,x_2)\\ x_1\oplus x_2=x}}\Pr_D[x_1]}_{=1}-\frac{1}{2^k}\underbrace{\sum_{\substack{(x_1,x_2)\\ x_1\oplus x_2=x}}\Pr_{D'}[x_2]}_{=1}+\underbrace{\sum_{\substack{(x_1,x_2)\\ x_1\oplus x_2=x}}\frac{1}{2^{2k}}}_{=1/2^k}\right|\\
&= \frac{1}{2}\sum_x\left|\sum_{\substack{(x_1,x_2)\\ x_1\oplus x_2=x}}\Pr_D[x_1]\Pr_{D'}[x_2]-\frac{1}{2^k}\right|\\
&= V(D\oplus D',U)=d_\infty(D\oplus D',U)\geq d_R(D\oplus D',U).
\end{aligned}$$

2 Short Answer: Indistinguishability and guessing

Answer. $c=1$.

Justification. We have $\Pr[b=0]=\Pr[b=1]=\frac{1}{2}$. For any algorithm A,

$$\begin{aligned}
&\Pr[b\xleftarrow{\$}\{0,1\};x\xleftarrow{\$}D_b;A(x)=b]\\
&= \Pr[b\xleftarrow{\$}\{0,1\};x\xleftarrow{\$}D_b;A(x)=b|b=0]\Pr[b=0]\\
&\quad +\Pr[b\xleftarrow{\$}\{0,1\};x\xleftarrow{\$}D_b;A(x)=b|b=1]\Pr[b=1]\\
&= \Pr[x_0\xleftarrow{\$}D_0;A(x_0)=0]\Pr[b=0]+\Pr[x_1\xleftarrow{\$}D_1;A(x_1)=1]\Pr[b=1]\\
&= \frac{1}{2}(1-\Pr[x_0\xleftarrow{\$}D_0;A(x_0)=1])+\frac{1}{2}\Pr[x_1\xleftarrow{\$}D_1;A(x_1)=1]\\
&= \frac{1}{2}-\frac{1}{2}(\Pr[x_0\xleftarrow{\$}D_0;A(x_0)=1]-\Pr[x_1\xleftarrow{\$}D_1;A(x_1)=1])\\
&= \frac{1}{2}\pm\frac{1}{2}\operatorname{Adv}A
\end{aligned}$$

Hence

$$2|\Pr[b \xleftarrow{\$} \{0,1\}; x \xleftarrow{\$} D_b; A(x) = b] - \frac{1}{2}| = \mathrm{Adv}\, A \leq d_R(D_0, D_1).$$

In other words, the definition of (t, ϵ)-indistinguishable and (t, ϵ)-inseparable are equivalent.

Remark. (t, ϵ)-indistinguishability captures the intuitively appealing idea that an adversary who is given either a value sampled from D_0 or from D_1 (where we've chosen randomly which type of value to give the adversary), cannot guess which type of value it was given. The equivalence of indistinguishability and indistinguishability means that we can freely use the definition of $\mathrm{Adv}\, A$ given in class when performing formal calculations, while we can stick to our intuition of guessing which distribution the value came with, without fear of paradox.

3 Short Answer: Correlation

This is false for arbitrary joint distributions X, Y, Z. For example, suppose they have been generated by the process $X \xleftarrow{\$} \{0,1\}^n, Y \xleftarrow{\$} \{0,1\}^n, Z = Y$. Then $X \sim Y$ but $(X, Z) \not\sim (Y, Z)$, as (X, Z) is a uniform $2n$-bit value, while (Y, Z) always contains the same n-bit value repeated twice.

Remark. This exposes a fallacy that I often see when teaching cryptography. Here's one version of the fallacy that might seem intuitively appealing (but is wrong nonetheless):

> **Claim (bogus).** If $X \sim Y$, then $(X, Z) \sim (Y, Z)$.
>
> **"Proof" (bogus).** The only difference between (X, Z) and (Y, Z) is in the first half of the pair. Consequently, any attacker can only distinguish between these two by distinguishing between X and Y. But we assumed that X and Y are indistinguishable, hence the claim follows.

The above reasoning is intuitively appealing, but fatally flawed. Maybe the distinguisher can somehow exploit the correlation between the first and second half of the pair; the above reasoning fails to take that possibility into account.

Typically, the reason people fall for this is that they omit the reduction. If you try to write down an explicit reduction, you quickly discover the flaw in the reasoning. For reference, here is the corrected version of the above.

> **Lemma (corrected).** Let X, Y, Z be random variables, where Z is independent of both X and Y, and where there is an efficient algorithm to sample from the distribution of Z in time t'. If X and Y are (t, ϵ)-indistinguishable, then (X, Z) and (Y, Z) are $(t - t' - O(1), \epsilon)$-indistinguishable.
>
> **Proof.** Consider any algorithm B that runs in time $t - t' - O(1)$ and tries to distinguish between (X, Z) and (Y, Z). Define the algorithm A by
>
> $A(s)$:
> 1. Choose t randomly according to the distribution on Z.
> 2. Run $B(s, t)$, and return whatever B does.
>
> Notice that $A(X)$ has the same distribution as $B(X, Z)$, and $A(Y)$ has the same distribution as $B(Y, Z)$. Also, A runs in time at most t. Then we find
>
> $$\mathrm{Adv}\, A = \mathrm{Adv}\, B \leq \epsilon.$$