

Exercise (Standard simulator for the Fiat-Shamir protocol). Let $\mathcal{S}(\beta)$ be the canonical simulator for the Fiat-Shamir sigma protocol, which takes in the challenge β and outputs α, γ such that for uniformly chosen $\beta \leftarrow \{0, 1\}$ the distribution of (α, β, γ) coincides with the message distribution between honest prover and verifier. Show that the following simulator construction

$$\mathcal{V}_\circ(\phi) \begin{cases} \beta_\circ \leftarrow \{0, 1\} \\ (\alpha, \gamma) \leftarrow \mathcal{S}(\beta_\circ) \\ \beta \leftarrow \mathcal{V}_*(\phi, \alpha) \\ \text{if } \beta \neq \beta_\circ \text{ then return } \perp \\ \text{else return } \mathcal{V}_*(\gamma) \end{cases}$$

creates an output distribution ψ_\circ that is statistically $\frac{1}{2}$ -distant from the output distribution of malicious verifier \mathcal{V}_* that interacts with the honest prover.

Solution. There are many ways how to prove this seemingly evident fact. In this treatment we consider two most revealing approaches.

PROOF BASED ON EVENT TREES. The simplest way to visualise what happens in the simulation is to draw the event trees for the real protocol execution and for the simulation and compare them. For clarity, let us first consider adversaries \mathcal{V}_* that do not abort in the real protocol run, i.e., do not output \perp . Also, let k be the number of all possible commitment messages α . In the Fiat-Shamir protocol, $k = \frac{\phi_e(N)}{4}$ where $\phi_e(\cdot)$ in this context is Euler's totient function. The latter follows from the fact that α is a quadratic residue.




Figure 1: Event trees for real protocol execution and for the protocol simulation.

Figure 1 shows that there are many ways to obtain ψ as an output. During the protocol execution, the prover \mathcal{P} chooses uniformly a challenge α by computing $\alpha \leftarrow r^2$ for $r \leftarrow \mathbb{Z}_N^*$. Next, \mathcal{V}_* chooses β based on its previous knowledge ϕ and the commitment α . After that the prover \mathcal{P} computes the unique reply $\gamma \leftarrow rs^\beta$. The verifier \mathcal{V}_* then processes the obtained information and produces an output. With some probability for each path in the event tree the outcome is ψ . During the simulation β_\circ is chosen uniformly. After that the sub-simulator \mathcal{S} creates matching α, γ values. Note that the number of possible commitments α that simulator outputs is still k . Moreover, the distribution of α values is still uniform for fixed β_\circ , since the simulator must sample protocol transcripts with the same probability that they occur in the protocol where

both parties are honest. After that the simulator feeds ϕ and α value to \mathcal{V}_* to obtain β . If $\beta = \beta_o$, the matching response is fed into \mathcal{V}_* to obtain the output. If $\beta \neq \beta_o$, the simulator halts abruptly.

At first glance, you may get a deceptive feeling that all paths are equally provable in the real protocol run. This is not true because $\mathcal{V}_*(\phi, \alpha)$ may bias the output β depending on the value of α . For the similar reason, the claim that in path during the simulation we reach \perp is $\frac{1}{2}$ is incorrect. However, if we consider path pairs that go through α_i in the simulation, the potential bias cancels out. Thus, we can straightforwardly prove

$$\Pr \left[\begin{array}{l} \beta_o \leftarrow_u \{0, 1\}, \alpha, \gamma \leftarrow \mathcal{S}(\beta_o), \\ \beta \leftarrow \mathcal{V}_*(\phi, \alpha) : \beta \neq \beta_o \wedge \alpha = \alpha_o \end{array} \right] = \frac{1}{2} \cdot \frac{1}{k} \cdot \Pr[\mathcal{V}_*(\phi, \alpha_o) = 0] + \frac{1}{2} \cdot \frac{1}{k} \cdot \Pr[\mathcal{V}_*(\phi, \alpha_o) = 1] = \frac{1}{2k}$$

for any α_o . Consequently,

$$\Pr \left[\begin{array}{l} \beta_o \leftarrow_u \{0, 1\}, \alpha, \gamma \leftarrow \mathcal{S}(\beta_o), \\ \beta \leftarrow \mathcal{V}_*(\phi, \alpha) : \beta \neq \beta_o \end{array} \right] = \sum_{\alpha_o} \Pr \left[\begin{array}{l} \beta_o \leftarrow_u \{0, 1\}, \alpha, \gamma \leftarrow \mathcal{S}(\beta_o), \\ \beta \leftarrow \mathcal{V}_*(\phi, \alpha) : \beta \neq \beta_o \wedge \alpha = \alpha_o \end{array} \right] = \frac{1}{2}.$$

As the second step note that there is a direct one-to-one correspondence of event paths in the real protocol execution and event paths without abortion in the simulation. Namely, each path is determined by α and β value. Let us now establish how the probabilities of these paths are related:

$$\Pr \left[\begin{array}{l} \beta_o \leftarrow_u \{0, 1\}, \alpha, \gamma \leftarrow \mathcal{S}(\beta_o) : \beta_o = \beta_0 \\ \mathcal{V}_*(\phi, \alpha) = \beta_0 \wedge \alpha = \alpha_0 \wedge \gamma = \gamma(\alpha_0, \beta_0) \end{array} \right] = \frac{1}{2} \cdot \Pr \left[\begin{array}{l} \alpha, \gamma \leftarrow \mathcal{S}(\beta_0) : \mathcal{V}_*(\phi, \alpha_0) = \beta_0 \\ \alpha = \alpha_0 \wedge \gamma = \gamma(\alpha_0, \beta_0) \end{array} \right]$$

where $\gamma(\alpha_0, \beta_0)$ denotes the correct response for α_0 and β_0 . As the simulation always yields a correct response γ and the first message is uniformly distributed even for fixed β_0 , we get

$$\Pr \left[\begin{array}{l} \beta_o \leftarrow_u \{0, 1\}, \alpha, \gamma \leftarrow \mathcal{S}(\beta_o) : \beta_o = \beta_0 \\ \mathcal{V}_*(\phi, \alpha) = \beta_0 \wedge \alpha = \alpha_0 \wedge \gamma = \gamma(\alpha_0, \beta_0) \end{array} \right] = \frac{1}{2k} \cdot \Pr[\mathcal{V}_*(\phi, \alpha_0) = \beta_0].$$

The probability that we choose the corresponding path in the real protocol execution is

$$\Pr \left[\begin{array}{l} \alpha \leftarrow \mathcal{P}, \beta \leftarrow \mathcal{V}_*(\phi, \alpha) \\ \alpha = \alpha_0 \wedge \beta = \beta_0 \end{array} \right] = \frac{1}{k} \cdot \Pr[\mathcal{V}_*(\phi, \alpha_0) = \beta_0].$$

Thus, we have proven that the probability of simulated paths is exactly twice lower than in the real protocol.

These two facts make the computation of the statistical distance straightforward. Recall that the statistical distance is the maximal advantage of any distinguisher. Also, recall that the bound is achieved by Neyman-Pearson distinguisher, which for each input x outputs the class for which the probability of getting x is larger. In our case,

$$\mathcal{A}_{\text{np}}(\psi_0) = \begin{cases} 0 & \text{if } \psi_0 \neq \perp, \\ 1 & \text{if } \psi_0 = \perp, \end{cases}$$

since probability of obtaining $\psi \neq \perp$ must be twice as low in the simulation as in the real protocol. As the corresponding advantage

$$|\Pr[\mathcal{A}_{\text{np}}(\psi_0) = 1 | \mathcal{H}_0] - \Pr[\mathcal{A}_{\text{np}}(\psi_0) = 1 | \mathcal{H}_1]| = \frac{1}{2},$$

we have proven the claim.

PROOF BASED ON GAME REWRITING. Another alternative for proving the desired claim is a systematic rewriting of the simulator constructor until we reach the code from which the claim is more or less evident. First, recall that the zero-knowledge property of sigma protocols is defined through the equivalence of

following games:

$$\mathcal{Q}_0^{\mathcal{B}} \qquad \mathcal{Q}_1^{\mathcal{B}}$$

$$\left[\begin{array}{l} \alpha \leftarrow \mathcal{P} \\ \beta \leftarrow \{0, 1\} \\ \gamma \leftarrow \mathcal{P}(\beta) \\ \mathbf{return} \mathcal{B}(\alpha, \beta, \gamma) \end{array} \right. \qquad \left[\begin{array}{l} \beta \leftarrow_{\mathcal{U}} \{0, 1\} \\ \alpha, \gamma \leftarrow \mathcal{S}(\beta) \\ \mathbf{return} \mathcal{B}(\alpha, \beta, \gamma) . \end{array} \right.$$

That is, for any distinguishing algorithm \mathcal{B} the distinguishing advantage is zero:

$$\text{Adv}_{\mathcal{Q}_0, \mathcal{Q}_1}^{\text{ind}}(\mathcal{B}) = |\Pr[\mathcal{Q}_0^{\mathcal{B}} = 1] - \Pr[\mathcal{Q}_1^{\mathcal{B}} = 1]| = 0 .$$

In other terms, the distributions of (α, β, γ) generated in both games coincide. As a consequence, we can rewrite the simulator constructor in the following way:

$$\mathcal{V}_o(\phi) \qquad \mathcal{V}_o(\phi)$$

$$\left[\begin{array}{l} \beta_o \leftarrow_{\mathcal{U}} \{0, 1\} \\ (\alpha, \gamma) \leftarrow \mathcal{S}(\beta_o) \\ \beta \leftarrow \mathcal{V}_*(\phi, \alpha) \\ \mathbf{if} \beta \neq \beta_o \mathbf{then return} \perp \\ \mathbf{else return} \mathcal{V}_*(\gamma) \end{array} \right. \rightsquigarrow \left[\begin{array}{l} \alpha \leftarrow \mathcal{P} \\ \beta_o \leftarrow \{0, 1\} \\ \gamma \leftarrow \mathcal{P}(\beta_o) \\ \beta \leftarrow \mathcal{V}_*(\phi, \alpha) \\ \mathbf{if} \beta \neq \beta_o \mathbf{then return} \perp \\ \mathbf{else return} \mathcal{V}_*(\gamma) . \end{array} \right.$$

Since \mathcal{V}_* response β depends only on α , we can move the statement upwards together with the if block:

$$\mathcal{V}_o(\phi) \qquad \mathcal{V}_o(\phi)$$

$$\left[\begin{array}{l} \alpha \leftarrow \mathcal{P} \\ \beta_o \leftarrow \{0, 1\} \\ \gamma \leftarrow \mathcal{P}(\beta_o) \\ \beta \leftarrow \mathcal{V}_*(\phi, \alpha) \\ \mathbf{if} \beta \neq \beta_o \mathbf{then return} \perp \\ \mathbf{else return} \mathcal{V}_*(\gamma) \end{array} \right. \rightsquigarrow \left[\begin{array}{l} \alpha \leftarrow \mathcal{P} \\ \beta \leftarrow \mathcal{V}_*(\phi, \alpha) \\ \beta_o \leftarrow \{0, 1\} \\ \mathbf{if} \beta \neq \beta_o \mathbf{then return} \perp \\ \gamma \leftarrow \mathcal{P}(\beta) \\ \mathbf{else return} \mathcal{V}_*(\gamma) . \end{array} \right.$$

From this it is clear that the simulator outputs \perp with the probability $\frac{1}{2}$, since β_o is generated independently after the β is fixed. The reasoning is fundamentally the same we obtained by analysing the event trees. The code rewriting steps that do not alter the semantics of the program literally reshape the event trees without altering the output distribution and thus implicitly mimic the analytic reasoning without nasty details.

As the refactored simulator code outputs \perp with probability $\frac{1}{2}$ we can simplify the code even further:

$$\mathcal{V}_o(\phi) \qquad \mathcal{V}_o(\phi) \qquad \mathcal{V}_o(\phi)$$

$$\left[\begin{array}{l} \alpha \leftarrow \mathcal{P} \\ \beta \leftarrow \mathcal{V}_*(\phi, \alpha) \\ \beta_o \leftarrow \{0, 1\} \\ \mathbf{if} \beta \neq \beta_o \mathbf{then return} \perp \\ \gamma \leftarrow \mathcal{P}(\beta) \\ \mathbf{else return} \mathcal{V}_*(\gamma) . \end{array} \right. \rightsquigarrow \left[\begin{array}{l} \alpha \leftarrow \mathcal{P} \\ \beta \leftarrow \mathcal{V}_*(\phi, \alpha) \\ b \leftarrow_{\mathcal{U}} \{0, 1\} \\ \mathbf{if} b = 0 \mathbf{then return} \perp \\ \gamma \leftarrow \mathcal{P}(\beta) \\ \mathbf{else return} \mathcal{V}_*(\gamma) . \end{array} \right. \rightsquigarrow \left[\begin{array}{l} b \leftarrow_{\mathcal{U}} \{0, 1\} \\ \mathbf{if} b = 0 \mathbf{then return} \perp \\ \alpha \leftarrow \mathcal{P} \\ \beta \leftarrow \mathcal{V}_*(\phi, \alpha) \\ \gamma \leftarrow \mathcal{P}(\beta) \\ \mathbf{else return} \mathcal{V}_*(\gamma) . \end{array} \right.$$

From this representation it is also evident that

$$\Pr[\mathcal{V}_o(\phi) = \psi] = \frac{1}{2} \cdot \Pr[\alpha \leftarrow \mathcal{P}, \beta \leftarrow \mathcal{V}_*(\phi, \alpha), \gamma \leftarrow \mathcal{P}(\alpha) : \mathcal{V}_*(\gamma) = \psi] .$$

After that the reasoning about the statistical distance is identical to the approach used in the first proof.