
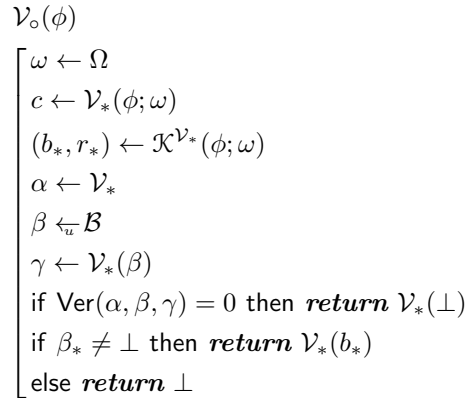


**Exercise (Standard simulator for the QNR-ZK protocol).** Let  $n$  be a composite number with a factorisation  $n = pq$  known to the prover  $\mathcal{P}$ . Let  $v \in \mathbb{Z}_n^*$  be a number for which the prover wants to prove that it is quadratic non-residue. Show that the the following zero-knowledge protocol



where the verifier  $\mathcal{V}$  uses sigma protocol  $\text{POK}[\exists r, \beta : c = r^2 v^b]$  to prove the knowledge of  $b$  and  $r$  is simulatable. More precisely, let  $\mathcal{K}^{\mathcal{V}^*}$  standard knowledge extractor for the protocol  $\text{POK}[\exists r, \beta : c = r^2 v^b]$ . Then show that the following simulator construction



can create an output distribution  $\psi_o$  that is statistically  $\varepsilon_o$ -distant from the output distribution of malicious verifier  $\mathcal{V}_*$  that interacts with the honest prover. Also, estimate how the running-time of the simulator depends on the desired statistical distance  $\varepsilon_o$ .

**Solution.** For the analysis, let us contrast the simulation with the real-world execution. For that we observe when the execution of a simulator diverges form the real-world interaction between  $\mathcal{V}(\phi)$  and  $\mathcal{P}$  provided that the randomness of the verifier  $\omega \leftarrow \Omega$  is same. Again, it is instructive to see the event trees first, see Figure 1. On the left you see the execution corresponding to the simulation. On the right you see the real protocol execution augmented with the run of the knowledge extractor. Here, the protocol run does not use knowledge-extraction for anything, it just provides the same split of events.

First, observe that the simulation and real-world execution are identical till the last verification step of  $\text{POK}[\exists r, \beta : c = r^2 v^b]$ , since all messages are created exactly the same way in both runs. If the verification fails, then the verifier gets the same input in both executions. Hence, we must consider the remaining case. Clearly, executions diverge when the knowledge-extraction fails. In the remaining case,  $\mathcal{V}$  receives  $\bar{b}$  in the real execution and  $b_*$  in the simulation. Hence, we must show that  $\bar{b} = b_*$ .

The latter reveals a small but important detail in the knowledge proof. Namely, it is straightforward to see that any invertible element  $c$  one has only two possible decompositions:  $c = r^2 v^b$  and  $c = (-r)^2 v^b$  if the Jacobi symbols of  $v$  and  $c$  match. Thus,  $b_*$  must be equal to  $\bar{b}$ . However, if the Jacobi symbols of  $v$  and  $c$  are not both ones, the existence of the decomposition is not guaranteed and we cannot conclude that  $b_* = \bar{b}$ .

Fortunately, in the first step of the proof of knowledge the verifying party checks that  $v$  and  $c$  are in the suitable form and the same is done in the knowledge extractor. Hence, the claim  $b_* = \bar{b}$  holds.




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

These observations combined with the the following lemma allow us to estimate the statistical distance by computing the probability of the event that verification succeeds while knowledge error fails.

**Lemma.** *If two strategies start to diverge only after the event  $\text{div}$  and are identical before that the the statistical distance between the output distributions is at most  $\Pr[\text{div}]$ .*

*Proof.* Let  $\mathcal{S}_1$  and  $\mathcal{S}_2$  be the strategies in question and  $\psi_1$  and  $\psi_2$  corresponding output distributions. Then by the definition

$$\begin{aligned}
\text{sd}(\psi_1, \psi_2) &= \frac{1}{2} \cdot \sum_{\phi} |\Pr[\phi_1 \leftarrow \mathcal{S}_1 : \phi_1 = \phi] - \Pr[\phi_2 \leftarrow \mathcal{S}_2 : \phi_2 = \phi]| \\
&= \frac{1}{2} \cdot \sum_{\phi} |\Pr[\phi_1 \leftarrow \mathcal{S}_1 : \phi_1 = \phi \wedge \text{div}] - \Pr[\phi_2 \leftarrow \mathcal{S}_2 : \phi_2 = \phi \wedge \text{div}]| \\
&\quad + \frac{1}{2} \cdot \sum_{\phi} |\Pr[\phi_1 \leftarrow \mathcal{S}_1 : \phi_1 = \phi \wedge \neg \text{div}] - \Pr[\phi_2 \leftarrow \mathcal{S}_2 : \phi_2 = \phi \wedge \neg \text{div}]| \\
&= \frac{1}{2} \cdot \sum_{\phi} |\Pr[\phi_1 \leftarrow \mathcal{S}_1 : \phi_1 = \phi \wedge \text{div}] - \Pr[\phi_2 \leftarrow \mathcal{S}_2 : \phi_2 = \phi \wedge \text{div}]| \\
&\leq \frac{1}{2} \cdot \sum_{\phi} \Pr[\phi_1 \leftarrow \mathcal{S}_1 : \phi_1 = \phi \wedge \text{div}] + \frac{1}{2} \cdot \sum_{\phi} \Pr[\phi_2 \leftarrow \mathcal{S}_2 : \phi_2 = \phi \wedge \text{div}] \\
&\leq \frac{1}{2} \cdot \Pr[\phi_1 \leftarrow \mathcal{S}_1 : \text{div}] + \frac{1}{2} \cdot \Pr[\phi_2 \leftarrow \mathcal{S}_2 : \text{div}] \leq \Pr[\text{div}]
\end{aligned}$$

Note that by the construction the event  $\text{div}$  has the same probability for both strategies since up to this procedure both systems are identical and thus  $\Pr[\text{div}]$  is well defined.  $\square$

For brevity, let  $\text{SimFail}$  denote the event that verification succeeds while knowledge error fails. As the probability of this event may depend on the randomness  $\omega$ , we must configure the knowledge extractor so that the corresponding probability is negligible on average over all possible  $\omega \in \Omega$ . First all all the failure

probability can be high for fixed  $\omega$  only if verification succeeds and knowledge error fails with non-negligible probability. Intuitively, achieving such a bad case is a careful balancing act, since higher success probability on verification also increases the success of  $\mathcal{K}^{\mathcal{V}}(\phi; \omega)$ . Figure 2 illustrates the corresponding behaviour.




Figure 2: Knowledge-extraction failure as function of verification success. The  $x$ -axis corresponds to the verification success and  $y$ -axis corresponds to the failure probability. The upper graph corresponds to the knowledge extraction failure and the lower graph corresponds to the simulation failure. The plot is drawn for the knowledge extractor that does 10 probes. In reasonable simulation constructions  $\ell \gg 10$  and the drop after the  $\frac{1}{2}$  is extremely sharp and the optimum does not change for fixed running-time simulators.

Figure 2 clearly shows that the worst configuration is such that  $\Pr[\text{Ver}(\alpha, \beta, \gamma) = 1] = \frac{1}{2}$  for which the simulation failure `SimFail` occurs with probability  $\frac{1}{2}$ . The latter holds for fixed randomness  $\omega \in \Omega$ . Thus, the worst adversarial configuration for the simulation is such that it will succeed the proof of knowledge with probability  $\frac{1}{2}$ . As a result,

$$\Pr[\text{SimFail}] = \sum_{\omega \in \Omega} \Pr[\omega] \cdot \Pr[\text{SimFail}|\omega] = \frac{1}{2}$$

and the lemma provides a poor bound on the statistical distance:

$$\text{sd}(\psi_1, \psi_2) \leq \Pr[\text{SimFail}] = \frac{1}{2} .$$

The bound is sharp. If the malicious verifier chooses a random challenge or a challenges obtained prior to the beginning of the protocol. Then the verifier can cheat in the proof of knowledge by first guessing  $\beta$  and then simulating the triple using the canonical simulator  $\mathcal{S}$ . As a result, the verifier will learn the quadratic residuosity of  $c$  with probability  $\frac{1}{2}$ . Under the standard assumptions detecting quadratic residuosity is intractable problem and thus no efficient simulator that achieves  $\text{sd}(\psi_1, \psi_2) \ll \frac{1}{2}$  can exist.

**NAIVE PROTOCOL UPGRADE.** The large discrepancy in the simulation is doused by the high knowledge error. The simplest solution is to repeat the proof-of-knowledge protocol  $k$  times. As a result, the knowledge-error  $\kappa$  decreases from  $\frac{1}{2}$  to  $2^{-k}$ . However, the lower knowledge error alone does not guarantee a small simulation failure, the the knowledge-extraction must be successful enough when  $\varepsilon > \kappa$ . To understand the logic behind the simulation of multi-stage proof-knowledge, it is instructive to look at Figure 3.

By the construction the events on each stage are independent form the previous stages. The stage fails if the verification succeeds but knowledge extraction fails. This stage failure depends on the adversaries behaviour. However, as the stage failure is identical to the simulation failure of the knowledge proof, it is easy to guarantee that the stage failure probability is below  $\frac{1}{2}$ . Thus, the statistical distance between verifiers outputs in the simulation and real protocol execution is well bounded:

$$\text{sd}(\psi_1, \psi_2) \leq \Pr[\text{SimFail}] \cdot \dots \cdot \Pr[\text{SimFail}] \leq 2^{-k} .$$




Figure 3: Combined simulation and knowledge extraction in case of a multi-stage proof of knowledge. Black dots represent verification stages dotting the simulation. Grey dots represent knowledge-extraction phases following successful verification stages. Dots with crosses denote states where perfect simulation is achievable either because a verification stage failed or because a knowledge-extraction yielded the desired witness. In these cases, perfect simulation of the protocol becomes trivial.

This bound leads to an interesting observation about knowledge-extraction strategy. The knowledge extraction does not have to be very successful to get a simulation profile that is bounded above by  $\frac{1}{2}$ . By simple analysis we can show that the failure probability of a standard knowledge extractor is

$$\Pr [\mathcal{K}^{\mathcal{V}}(\phi; \omega) = \perp] = \begin{cases} 1, & \text{if } \varepsilon \leq \frac{1}{2}, \\ (2 - 2\varepsilon)^\ell, & \text{if } \varepsilon > \frac{1}{2}, \end{cases}$$

where the number of rewinds is  $2\ell$ . As the success probability in the verification stage is  $\varepsilon$ , the simulation failure can be expressed as follows

$$\Pr [\text{SimFail}] = \begin{cases} \varepsilon, & \text{if } \varepsilon \leq \frac{1}{2}, \\ \varepsilon(2 - 2\varepsilon)^\ell, & \text{if } \varepsilon > \frac{1}{2}. \end{cases}$$

The corresponding graph illustrated in Figure 4 shows that the optimal number of rewinds is 2.




Figure 4: Knowledge-extraction failure as function of verification success. The  $x$ -axis corresponds to the verification success and  $y$ -axis corresponds to the failure probability. The upper collection of graphs corresponds to the knowledge extraction failure for  $\ell \in \{1, \dots, 10\}$  and the lower graph corresponds to the simulation failure for  $\ell \in \{1, \dots, 10\}$ . The optimal configuration where  $\ell = 1$  is drawn in bold. By increasing the value of  $\ell$ , the simulation failure drops faster when  $\varepsilon > \kappa$  but this does not reduce the worst-case bound on the simulation failure.

**ROUND-EFFICIENT SOLUTION.** More efficient way to reduce the knowledge error is to run several protocols in parallel. The latter does not change the soundness of the protocol, since the latter depends on witness

indistinguishability that is preserved under parallel composition. It is also straightforward to see that the knowledge error indeed drops  $2^{-k}$ . There are  $2^k$  potential challenges and if the verifier is able to answer at least two of them correctly, then at least two sub-protocol runs are successful with different challenges and we can extract the witness. In brief, the knowledge extractor has to find two successful verification runs in order to extract the secret. As the simulation is already successful, it is enough to use random challenges in order to reveal the second successful verification run with the same randomness.

If we probe the verifier with random challenges, the probability that we get a successful run that is different from the one revealed in simulation is  $(\varepsilon - \kappa)$  where  $\varepsilon$  is the unknown success probability for verification. Thus, the probability of continuous stream of failures after  $\ell$  tries is

$$\Pr[\mathcal{K}^{\mathcal{V}}(\phi, \kappa) = \perp] = \begin{cases} 1, & \text{if } \varepsilon \leq \kappa, \\ (1 + \kappa - \varepsilon)^\ell & \text{if } \varepsilon > \kappa. \end{cases}$$

Hence, the probability of simulation failure if

$$\Pr[\text{SimFail}] = \begin{cases} \varepsilon, & \text{if } \varepsilon \leq \kappa, \\ \varepsilon(1 + \kappa - \varepsilon)^\ell & \text{if } \varepsilon > \kappa. \end{cases}$$

Figure 5 shows that this setting offers different tradeoffs in terms of rewindings.




Figure 5: Simulation failure as function of verification success when  $\kappa = \frac{1}{8}$ . The  $x$ -axis corresponds to the verification success and  $y$ -axis corresponds to the simulation failure probability. The collection of graphs corresponds to  $\ell \in \{1, \dots, 10\}$ . The choice of the optimal value of  $\ell$  is not a clear cut. The maximal value for simulation failure is above  $\frac{1}{8}$  when  $\ell \in \{1, \dots, 9\}$  but there is no rational reasons why the number of rewindings should be large enough to guarantee that the simulation failure is equal to  $\kappa$ .

To find the maximal value on the simulation failure profile note that logarithm is a monotone function and thus we can minimise  $\mathcal{L} = \log \Pr[\text{SimFail}]$  instead. It is straightforward to see that

$$\frac{\partial \mathcal{L}}{\partial \varepsilon} = \frac{1 + \kappa - \varepsilon(\ell + 1)}{\varepsilon(1 + \kappa - \varepsilon)} \quad \text{for } \varepsilon \geq \kappa$$

and thus the function is bound to have a single maximum in the location

$$\varepsilon_* = \frac{1 + \kappa}{1 + \ell}$$

after which the function starts to decrease. Consequently, if the number of rewindings

$$\ell \geq \frac{1}{\kappa},$$

the simulation failure is quarantined to decrease after  $\varepsilon \geq \kappa$  and simulation failure coincides with knowledge error. For other choices of  $\ell$  the maximal simulation failure is

$$p_* = \max_{\varepsilon \in [0,1]} \Pr [\text{SimFail}] = \frac{(1 + \kappa)^{\ell+1}}{2 + \ell} \cdot \left(1 + \frac{1}{\ell + 1}\right)^{\ell+1} \approx \frac{(1 + \kappa)^{\ell+1}}{2 + \ell} \cdot e$$

To illustrate the asymptotic behaviour in the process  $\kappa \rightarrow 0$ , we can try different ways to increase  $\ell$ . By choosing  $\ell = \Theta(\frac{1}{\kappa})$  we get

$$p_* \approx \frac{(1 + \kappa)^{\frac{1}{c\kappa} + 1}}{2 + \frac{1}{c\kappa}} \cdot e \approx e^{1 + \frac{1}{c}} \cdot \frac{(1 + \kappa)c\kappa}{2c\kappa + 1} \approx \alpha\kappa \ .$$

In other words, if  $\ell = \Theta(\frac{1}{\kappa})$  then the simulation failure is of order  $\Theta(\kappa)$ . The latter is not very good for getting tight simulation bounds. In order to get statistical difference  $2^{-80}$  we have to do around  $2^{80}$  rewindings, which makes the simulator  $2^{80}$  times slower than the original attack code.

The simulation failure does not have to be of same order as the knowledge error  $\kappa$ . As the knowledge error decreases exponentially in terms of the protocol efficiency, we could in principle allow different tradeoffs between  $p_*, \kappa, \ell$ . Ideally, we would like that  $p_*$  would decrease exponentially wrt  $\ell$ . As the second term in the approximation

$$\log p_* \approx 1 - \log(2 + \ell) + (\ell + 1) \log(1 + \kappa) \approx 1 - \log(2 + \ell) + (\ell + 1)\kappa$$

cannot be of order  $\ell$ , the exponential decay of simulation failure is impossible. In fact the simulation failure can decrease as a function of  $\ell$  only if

$$\kappa \leq \frac{\log(2 + \ell)}{\ell + 1} \ .$$

Note that even if the third term remains constant and does not grow, the simulation failure decreases at most linearly in  $\log(\ell)$  and thus for any choice of  $\kappa$  the simulation failure is still  $\Omega(\frac{1}{\ell})$ . In that sense, inversely proportional slowdown cannot be escaped with conventional methods.