Composable Deniability

Jörn Müller-Quade and Dominique Unruh

Institute for Algorithms and Cognitive Systems/E.I.S.S.
Fakultät für Informatik, Universität Karlsruhe, Am Fasanengarten 5,
76131 Karlsruhe, Germany

Simulation based definitions of security, like the UC framework or Reactive Simulatability, seem to cover all possible aspects of security. However, one important security property is not covered by these models: The ability to deny secret data even when confronted with evidence the adversary has gathered during a protocol run.

We present a framework to define composable deniability. To do this we introduce two additional machines to the UC framework: In the ideal model we add a deceiver, which reflects the ability to deny in the ideal model (which is easy, because of the absence of cryptographic mechanisms). For a corrupted party P_i trying to deny secret data the deceiver can modify the inputs and outputs of this party P_i. In the real model we introduce a deceiver-simulator, which is able to modify all protocol messages of the party P_i. This deceiver-simulator must be able to fake the messages according to the lies of the deceiver in the ideal model, so that the environment cannot distinguish. A deceiver simulator which is successful even in presence of all the cryptographic mechanisms used in the real protocol yields a denial strategy.

Definition 1. *A protocol ρ is said to implement an ideal functionality $\mathcal{F}$ with* composable deniability *if for every real adversary $\mathcal{A}$ there exists a simulator $\mathcal{A}_\mathcal{S}$ such that for every deceiver $\mathcal{D}$ there exists a deceiver-simulator $\mathcal{D}_\mathcal{S}$ such that for all environment $\mathcal{Z}$ we have indistinguishability of the output of $\mathcal{Z}$ when interacting with the real model or the ideal model.*

The notion of composable deniability is the first definition of deniability which exhibits a composition theorem. I.e., it is possible to replace ideal functionalities used in a larger protocol by real protocols without losing security or the property of deniability.

The model allows to formulate different flavours of deniability: In the strictest version the denial strategy must be independent of the set of parties controlled by the adversary and each party must be able to deny "on-line". This is in contrast to a weaker definition where denial must be possible after termination of the protocol. In its strongest form the model even covers the situation where a party deliberately deviates from the protocol to thwart its own deniability (as for example in a vote buying attack). Furthermore different communication channels can be considered: It is useless to deny the existence of a message the adversary could see on the channel.

As an example we give a protocol for deniable authentication in this very strict framework.