The Subliminal Channels in the U.S. Digital Signature Algorithm (DSA)

Gustavus J. Simmons
Sandia Park, NM 87047, USA

Abstract

Since the DSA is derivative from El Gamal's digital signature scheme—which Simmons showed in 1985 permitted a subliminal channel—it should come as no surprise that the DSA also permits a similar channel. The subliminal channel in the El Gamal scheme, however, had several shortcomings. In order for the subliminal receiver to be able to recover the subliminal message, it was necessary for him to know the transmitter's secret key. This meant that the subliminal receiver had the capability to utter undetectable forgeries of the transmitter's signature. Also, only a subset of the desired message set could be communicated subliminally ($\varphi(p-1)$ messages instead of p-1) and some of those that could be transmitted were computationally infeasible for the subliminal receiver to recover.

The subliminal channels in the DSA avoid all of these difficulties! In fairness, it should be mentioned that not all are avoided at the same time. The channel in the DSA analogous to the one Simmons demonstrated in the El Gamal scheme can communicate messages conveying the full $\log_2 |\mathcal{K}|$ bits, where $\mathcal{K}$ is the set of session keys; all of which are easily recovered by the subliminal receiver. However, this broadband channel still requires the subliminal receiver to know the transmitter's secret key. There are two other narrowband ($<< \log_2 |\mathcal{K}|$) subliminal channels in the DSA that do not give the subliminal receiver any better chance of forging the transmitter's signature than an outsider has. The price one pays for this integrity for the transmitter's signature is a reduced bandwidth for the subliminal channel and a difficult but feasible (dependent on the bandwidth actually used) amount of computation needed to use the channel. Two quite different such channels have been devised: one places the computational load almost entirely on the transmitter, the other almost entirely on the subliminal receiver. Since the total computation is essentially the same in either case, the choice of a particular channel would be based on which end is best equipped to do the necessary computation.

To make clear what a remarkable coincidence it is that the apparently inherent shortcomings of subliminal channels using the El Gamal scheme can all be overcome in the DSA, we will analyze each of the channels implemented in both schemes. The inescapable conclusion, though, is that the DSA provides the most hospitable setting for subliminal communications discovered to date.

Introduction

In 1983 Simmons introduced the notion of a subliminal channel existing in an encrypted communication channel by pointing out that if for each plaintext there existed two or more corresponding cipher texts, the identity of the cipher used to communicate a plaintext could convey information additional to that revealed by the decryption of the cipher [5]. In particular, in a public-key based authentication scheme in which the decryption key must be public information in order for public receivers to be able to decrypt cipher texts and verify the authenticity of the encrypted plaintext, this raised the possibility of there also being subliminal receivers who could recover information hidden from