

Beefing up public-key encryption

February 18 2013



Most financial transactions on the Internet are safeguarded by a cryptographic technique called public-key encryption. Where traditional encryption relies on a single secret key, shared by both sender and recipient, public-key encryption uses two keys that are mathematically related. One, the public key, is published on the Internet, and any sender can use it to encrypt a message; the second, the private key, is known only to the recipient and is required for decryption.

Standard public-key [encryption](#) is secure as long as an attacker knows nothing other than the public key. But financial institutions and other

large organizations seek security against more sophisticated attacks, called chosen-ciphertext attacks (CCAs), in which the attacker also has examples of successful decryption.

Unfortunately, public-key [encryption schemes](#) that are resilient against CCAs are hard to devise. Their complexity means that software implementations are prone to small errors that can introduce both vulnerabilities and inaccuracies during decryption.

At the International Conference on the Theory and Applications of Cryptographic Techniques this spring, a pair of postdocs at MIT's Computer Science and Artificial Intelligence Laboratory describe a new technique for taking one of these vulnerable, error-prone CCA schemes and turning it into a secure CCA scheme. The result could be of practical use, in the development of more-secure encryption protocols, but it could also provide theoretical insight into the very nature of [cryptographic security](#).

Playing the odds

In cryptography circles, a message to be encoded is called a plaintext; the encrypted version of it is called a ciphertext. An encryption scheme is considered secure if even someone who knows two plaintexts in advance would find it virtually impossible to deduce which of two ciphertexts encodes which.

In the type of weak-CCA schemes that the MIT researchers—Huijia Lin and Stefano Tessaro—consider, the probability of distinguishing the ciphertexts is non-negligible. It may not be very big, but it's big enough to be a cause for concern. Similarly, there's also a non-negligible probability of errors during decryption.

Lin and Tessaro's result hinges on the observation that while, in the

average case, the probability of distinguishing weakly encrypted ciphertexts may be unacceptably high, in some particular cases, it's negligible. Moreover, it's possible to compute the probability that the encrypted version of a randomly generated plaintext will be secure.

As it happens, combining a weakly encrypted ciphertext with a strongly encrypted one produces a strongly encrypted hybrid. In essence, Lin and Tessaro's scheme entails encrypting enough random plaintexts that, probabilistically, at least a few of them will be secure. Then they're all combined.

Lin and Tessaro's technique doesn't just secure transmissions against attackers who have some examples of successful decryption; it secures them against adversaries who have a black box that, without disclosing the [secret key](#), can decrypt any ciphertext they feed it—except the one that's under attack.

"In real life, maybe it seems more plausible that people would just get a couple of examples of ciphertexts and messages, but as cryptographers, we always want to prevent the worst possible scenario," Lin says. "Even if we want to handle the case where you just have examples of ciphertexts and the messages, it's hard to exhaust all possible scenarios. By considering the strongest attack, we automatically become immune to all possible scenarios, which are hard to enumerate."

"The question they ask"—how to improve the security of vulnerable encryption schemes—"is a very natural and intriguing one," says Abhi Shalat, an assistant professor of computer science at the University of Virginia, who studies encryption. "And the techniques that they use to show that demonstrate a lot of creativity and elegance."

In information theory, Shalat says, it's well established that "if you and I are trying to communicate, and we have a slightly better chance of less

noise than the adversary has, you can sort of amplify that advantage so that we can actually talk privately." But "that technique doesn't necessarily work when the adversary is adaptive—when the adversary can query a [black-box decryption] oracle and be actively malicious, as is the case with encryption on the Internet," Shalat says. "One of the really nice ideas in the paper is the extension of this information-amplification idea to this adaptive setting."

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Beefing up public-key encryption (2013, February 18) retrieved 20 March 2024 from <https://phys.org/news/2013-02-beefing-public-key-encryption.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--