DIMACS
Center for Discrete Mathematics & Theoretical Computer Science
Founded as a National Science Foundation Science and Technology Center

Is Your Database Secure?: New Attacks and Leaks Discovered

[July, 2017] Rutgers computer science graduate student Betül Durak recently helped to uncover an attack on the FF3 encryption method [1] recommended by NIST (the National Institute of Standards and Technology) for securing information in databases. FF3 is a "format-preserving" encryption (FPE) method, which means that the output of the encryption (called the ciphertext) has the same format as the input. FPE makes it possible to encrypt data in legacy business applications built on structured data models to add security in a relatively seamless way. For example, when a legacy retail system expects a 16-digit credit card number, FPE provides a ciphertext that looks like a credit card number, but with the added security of encryption. The "tweakability" feature of FF3 is intended to enhance its security against dictionary attacks and also allow for cases where there may be reasons to distinguish different instances of the same plaintext (such as two customers both named "John Doe"). NIST recommended FF3 for use in commercial transactions because it was seen as providing the strong security guarantees while also abiding important practical constraints.

CS graduate student Betül Durak

Durak's recent work with her collaborator Serge Vaudenay (EPFL) [2] challenged this belief by presenting a new attack on the FF3 scheme that is practical when the domain size is small. They shared their findings with NIST, leading to an April 12, 2017 news release by NIST stating, "NIST has concluded that FF3 is no longer suitable as a general-purpose FPE method." The computational complexity of the attack depends on the size of the domain that the implementation is designed to encrypt. When that domain consists of the middle six digits of a credit card number, NIST concluded that the attack may be practical to execute.

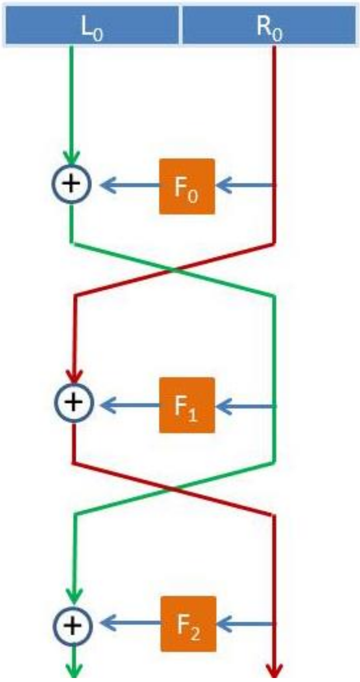

Three-round Feistel Network with group operation +

FF3 is based on a Feistel network (FN) construction, which is a widely used approach in block ciphers. Encryption based on a Feistel network consists of multiple rounds of processing of the plaintext after it is broken into left and right parts. Each round consists of a transformation step followed by a permutation step. More rounds in a FN lead to more security, but at the cost of slower encryption and decryption. The FF3 construction is an 8-round FN that uses a "tweak" parameter XORed with a round counter as an input to the block cipher (the F's in the figure). The XOR operation guarantees that round functions are pairwise different, a property referred to as "domain separation". The attack exploits "bad" domain separation in FF3. Namely, a specific design choice of FF3 allows permuting the round functions by changing the tweak, enabling a so-called slide attack using only two tweaks. The attack gives a round-function-recovery attack on FF3 when when the message domain is small and the adversary has power to choose both the messages to encrypt and the tweaks. Durak and Vaudenay's paper, which will appear in the 37th Annual International Cryptology Conference (Crypto 2017) in August, describes the attack and also provides a simple patch to thwart it. The NIST release discusses the attack's practical viability.

This work on FF3 is part of Durak's growing body of research on encryption methods that provide security guarantees while also satisfying additional practical requirements. Preserving the format of the input is one such practical constraint. Other times these constraints require the encryption to allow certain things about the plaintexts to be computable from the ciphertexts without a secret key. For instance, order-revealing encryption (ORE) enables determining the order of two plaintexts from their ciphertexts. Allowing this information – but nothing more – to be revealed is difficult to guarantee, and sometimes more is revealed than intended. In recent work with her advisor David Cash (Rutgers) and Thomas DuBuisson (Galois) [3], Durak examined such leakage in ORE when applied on data that might be encountered in practice. They noted that columns of data in a table are often correlated because rows in a table usually correspond to individual records. This observation opens up a new avenue for an attack in which an attacker attempts to extract information from multiple columns simultaneously rather than from the columns individually. Durak and her collaborators showed that when multiple columns of correlated data are encrypted with ORE, attacks can use the encrypted columns together to reveal an alarming level of information even with provably secure ORE constructions that prevent information leakage in attacks on individual columns.

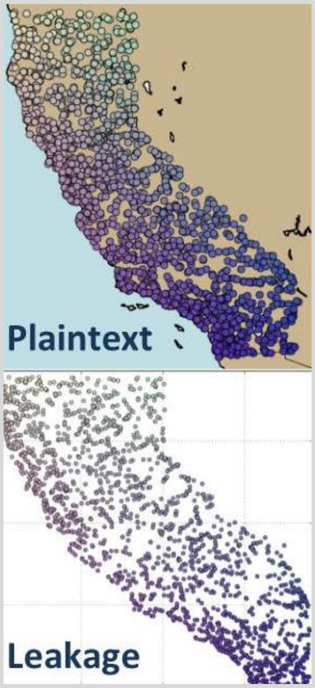

The plaintext contains locations of California road intersections

They studied the effect of correlation using simple multicolumn versions of attacks on ideal ORE and analyzed them using visualizations and measurements of accuracy on geographic datasets where latitude and longitude were correlated. Results showed that the rough shape of the geographic distributions is present in the ciphertext leakage. An example is shown in the figure at left in which the plaintext consisted of locations of road intersection in California. The images suggest a concerning loss of secrecy. Furthermore, the authors conjecture that security in practice may be much worse than revealed by their experiments conducted on relatively small datasets because leakage grows as the dataset grows. This research appeared at the 2016 Conference on Computer and Communications Security (CCS 2016).

Durak's work on ORE and secure databases more broadly is part of the DIMACS-led component of the Jana project led by Galois. Jana is funded through DARPA's Brandeis Program which seeks to harness the value of big data while protecting the privacy of individuals. The experiments with ORE led Durak and her collaborators to conclude that its security would be inadequate for safeguarding information such as personal location histories contained in cell phone data.

References:

[1] M. Dworkin, "Recommendation for Block Cipher Modes of Operation: Methods for Format-Preserving Encryption," NIST Special Publication 800-38G, 2016.

[2] F. B. Durak and S. Vaudenay, "Breaking the FF3 Format-Preserving Encryption Standard Over Small Domains," Cryptology ePrint Archive, Report 2017/521.

[3] F. B. Durak, T. DuBuisson, and D. Cash, "What Else is Revealed by Order-Revealing Encryption?" in Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security.