Lecture 8: Timing attacks

- Homework 4 has been posted
- There is no lecture on Tuesday 2/18
- In-class test on Thursday 2/20, covering the following material:
 - All lectures up to (and including) this one
 - All Friday discussion sections
 - All concepts within the first 4 homework assignments
 - All required textbook reading
- Nicolas' discussion section tomorrow will be test preparation

Crypto = Scientific field at intersection of many disciplines

Algorithms

Known for cipher design. Primarily found in European academia.

Engineering

Known for software dev and side channel attacks. Primarily found in industry.

$A \Rightarrow B$

This class

Complexity theory

Known for reductions. Primarily found in American academia.

Mathematics

Known for cryptanalysis. Primarily found in government.



Cryptography



Cryptanalysis

Physics of implementation

Math of algorithm



Breaking data confidentiality



message **P**



fakefiles

•

(something)



b key **K**



Side channel attacks on crypto implementations

- Crypto security definitions ensure that the output is "harmless"



• But, crypto implementations can reveal more than its desired outputs! These side channels of information weren't captured in our definitions

• Focus for this week: side channels on AES \Rightarrow and thus any AES mode



Divide and conquer

- Break 1 byte of the message or key at a time
- For each byte: guess all 256 values and check which works
- (Think: how you see crypto broken in any Hollywood movie)





Last time: Power analysis of AES in hardware



Mallory: oscilloscope to measure power



Alice: FPGA that runs AES

Today: Timing attacks on AES in software



Question: what might affect the runtime of AES?

Answer: the S-box! Let's look at simplified first and last rounds of AES



Simplified picture of first round of AES



- The first round of AES begins with xor of key followed by S-box lookups
- Let A and B denote the 16-byte state before/after the first round S-box
- Claim: if Eve knows input X along with A (or B), she can recover the secret key



Simplified picture of last round of AES



- Similar attack on last round: if Eve knows Y and A/B, can recover key
- There is a bijection between the first and last round keys of AES
- (Note: picture omits the final round ShiftRows, which permutes bytes)

AES code has table lookups

static con	st u8 Te	4[256]	= {				
0×63U,	0×7cU,	0×77U,	0×7bU,	0×f2U,	0×6bU,	0×6fU,	0×c5U,
0×30U,	0×01U,	0×67U,	0×2bU,	0×feU,	0×d7U,	0×abU,	0×76U,
0×caU,	0×82U,	0×c9U,	0×7dU,	0×faU,	0×59U,	0×47U,	0×f0U,
0×adU,	0×d4U,	0×a2U,	0×afU,	0×9cU,	0×a4U,	0×72U,	0×c0U,
0×b7U,	0×fdU,	0×93U,	0×26U,	0×36U,	0×3fU,	0×f7U,	0×ccU,
0×34U,	0×a5U,	0×e5U,	0×f1U,	0×71U,	0×d8U,	0×31U,	0×15U,
0×04U,	0×c7U,	0×23U,	0×c3U,	0×18U,	0×96U,	0×05U,	0×9aU,
0×07U,	0×12U,	0×80U,	0×e2U,	0×ebU,	0×27U,	0×b2U,	0×75U,
0×09U,	0×83U,	0×2cU,	0×1aU,	0×1bU,	0×6eU,	0×5aU,	0×a0U,
0×52U,	0×3bU,	0×d6U,	0×b3U,	0×29U,	0×e3U,	0×2fU,	0×84U,
0×53U,	0×d1U,	0×00U,	0×edU,	0×20U,	0×fcU,	0×b1U,	0×5bU,
0×6aU,	0×cbU,	0×beU,	0×39U,	0×4aU,	0×4cU,	0×58U,	0×cfU,
0×d0U,	0×efU,	0×aaU,	0×fbU,	0×43U,	0×4dU,	0×33U,	0×85U,
0×45U,	0×f9U,	0×02U,	0×7fU,	0×50U,	0×3cU,	0×9fU,	0×a8U,
0×51U,	0×a3U,	0×40U,	0×8fU,	0×92U,	0×9dU,	0×38U,	0×f5U,
0×bcU,	0×b6U,	0×daU,	0×21U,	0×10U,	0×ffU,	0×f3U,	0×d2U,
0×cdU,	0×0cU,	0×13U,	0×ecU,	0×5fU,	0×97U,	0×44U,	0×17U,
0×c4U,	0×a7U,	0×7eU,	0×3dU,	0×64U,	0×5dU,	0×19U,	0×73U,
0×60U,	0×81U,	0×4fU,	0×dcU,	0×22U,	0×2aU,	0×90U,	0×88U,
0×46U,	0×eeU,	0×b8U,	0×14U,	0×deU,	0×5eU,	0×0bU,	0×dbU,
0×e0U,	0×32U,	0×3aU,	0×0aU,	0×49U,	0×06U,	0×24U,	0×5cU,
0×c2U,	0×d3U,	0×acU,	0×62U,	0×91U,	0×95U,	0×e4U,	0×79U,
0×e7U,	0×c8U,	0×37U,	0×6dU,	0×8dU,	0×d5U,	0×4eU,	0×a9U,
0×6cU,	0×56U,	0×f4U,	0×eaU,	0×65U,	0×7aU,	0×aeU,	0×08U,
0×baU,	0×78U,	0×25U,	0×2eU,	0×1cU,	0×a6U,	0×b4U,	0×c6U,
0×e8U,	0×ddU,	0×74U,	0×1fU,	0×4bU,	0×bdU,	0×8bU,	0×8aU,
0×70U,	0×3eU,	0×b5U,	0×66U,	0×48U,	0×03U,	0×f6U,	0×0eU,
0×61U,	0×35U,	0×57U,	0×b9U,	0×86U,	0×c1U,	0×1dU,	0×9eU,
0×e1U,	0×f8U,	0×98U,	0×11U,	0×69U,	0×d9U,	0×8eU,	0×94U,
0×9bU,	0×1eU,	0×87U,	0×e9U,	0×ceU,	0×55U,	0×28U,	0×dfU,
0×8cU,	0×a1U,	0×89U,	0×0dU,	0×bfU,	0×e6U,	0×42U,	0×68U,
0×41U,	0×99U,	0×2dU,	0×0fU,	0×b0U,	0×54U,	0×bbU,	0×16U

Source: github.com/openssl/openssl/blob/master/crypto/aes/aes_core.c

Computer caching



Computers *cache* recently-accessed data, assuming that if you wanted it before, then you may want it again

- Response of array lookup depends upon whether the value is already in cache
- This, in turn, depends on whether you've already looked up this value in the past



Cache attack: Mallory co-resident with Alice

- For now, suppose Mallory has a presence on Alice's machine
 - Co-located VMs on the cloud
 - Unprivileged account on a multi-tenant Unix machine with full-disk encryption
 - Unprivileged application that can write files to a phone with disk encryption
- Cache is shared between all tenants on a machine
- Ergo, Mallory can influence the state of Alice's cache! [Osvik, Shamir, Tromer 2006]

How the cache works

- Fixed mapping between locations in memory & cache
- When victim Alice runs AES, selected portions of the S-box move into cache
- If Mallory controls a large region of memory (~size of cache), she can fill in the cache with her own contents



Cache

Main memory

Prime + Probe attack

Algorithm:

- Mallory fills the cache with a large array A 1.
- 2. Wait for Alice to execute an AES encipher/decipher operation
- Mallory re-reads array A, records time to retrieve each byte 3.

Strength: Find key byte with ~800 samples over 65ms

Countermeasure: check for scans of large arrays?





Upshot: AES evicted some lines of Mallory's cache based upon secret key

Evict + Time attack

Algorithm:

- 1. Mallory creates a large array A but does not read it yet
- 2. Trigger an AES encipher/decipher with known input x/output y
- 3. Mallory reads a few bytes of array A
- 4. Trigger another AES encipher/decipher with the same x/y
- Upshot: 2nd AES is slower iff Mallory evicted the right cacheline

Strength: Find key byte with ~50k s a really large array



Strength: Find key byte with ~50k samples over ~30s, without ever reading

		r -			
		L			
_		÷			
-		r			
		L			
_		÷		 	
_					
	_				
		1.1			

Timing attack: Mallory observes Alice over network

- Suppose Alice kicks Mallory off of her machine
 - Mallory cannot tamper with Alice's cache
 - Mallory doesn't get to observe Alice's cache directly
- Still, timing information may be viewable remotely!
 - Mallory can observe response times to Alice's TLS packets over the internet
 - Mallory can use this info to find Alice's key (albeit with many more samples)

Timing of first round table lookups



- In the first round, AES code makes 16 S-box table lookups
- If $a_1 == a_2$, then S[a_2] table lookup will be fast since the value is in cache
- In general, 1st round running time \propto # of distinct intermediate values

How can Mallory exploit speed differences?

- the first two S-box lookups
- Then, Mallory knows that

 $a_1 = a_2$

• Let's say Mallory tells Alice to encipher an input with $x_1 = 01$, $x_2 = 02$

Suppose for now that Mallory magically learns that a cache hit occurs in

 $key_1 \oplus x_1 = key_2 \oplus x_2$

 $key_1 \oplus key_2 = x_1 \oplus x_2 = 03$

How can Mallory find Alice's key?

- Even knowing key₁ \oplus key₂ = $x_1 \oplus x_2 = 03$ doesn't tell you key₁ or key₂
- What if all 16 input bytes caused collisions?
- Then Mallory can also compute $key_1 \oplus key_3$, $key_1 \oplus key_4$, ..., $key_1 \oplus key_1 \oplus key_{16}$
- I claim that Mallory has effectively learned 120 of the 128 bits of key!
 - There are 256 choices for key₁
 - Each choice gives a unique remaining option for key₂, key₃, ..., key₁₆
- Brute-force the rest if you have an (x, y) pair

Making the attack more realistic

Simplifying assumptions so far

- 1 input $\mathbf{x} \rightarrow$ many cache collisions
- Can tell which bytes of **a** collide
- Timing measurement corresponds precisely to first round runtime, which is exactly proportional to # of *a* collisions

How to remove these assumptions

- View time for many colliding x (stronger signal)
- Vary **x** samples only in certain locations (more precise *signal*)
- Collect even more samples to overcome *noise*

Tactic 1: Collect more samples

- Strategy
 - Don't assume the existence of a single "magical" **x** with many collisions
 - Instead, simply try many possible **x**
- If **x** is chosen randomly, then the probability that: • a given pair of bytes (e.g., bytes 1 and 2) collide = 1/256 • byte 1 collides with some other byte $\approx 1/16$

- there exists a collision = 1 (256 choose 16)/(256¹⁶) ≈ 1 10⁻¹⁴
- Just as before, each collision yields a constraint on the key
- Sample enough **x** until we observe 15 independent constraints

Tactic 2: Strategically vary x

- Mallory needs to (1) observe a collision and (2) know where it occurs
- We can determine which bytes collide by fixing part of x
- Example: take average timing over several inputs with
 - $x_1 = 0, x_2 = 0$, and the other 14 bytes randomly chosen
 - $x_1 = 0, x_2 = 1$, and the other 14 bytes randomly chosen
 - •
 - $x_1 = 0, x_2 = 255$, and the other 14 bytes randomly chosen
- For whichever bucket is consistently faster, $x_1 \oplus x_2 = key_1 \oplus key_2$

Tactic 3: Repeat to overcome noise

- We know that Time(AES) is smaller when $a_1 = a_2$ than when they differ
- Mallory's measurement of AES runtime depends on many other factors
 - Other bytes in the same cache line
 - Other bytes of the 1st round (or last round if Mallory knows the output y)
 - Other rounds
 - Network latency (if you're conducting this attack remotely)
- With enough samples, we can average over this noise!
- Bin running times by $x_1 \oplus x_2$, see which is smallest

Countermeasures to (cache) timing attacks

- 1. Don't have table lookups
 - Hardware implementations of AES are not vulnerable
 - There exist other ciphers that are designed to avoid the need for table lookups (e.g., we will see later in the course that SHA-3 doesn't have any)
- 2. Look up the entire table
 - Pre-load the entire S-box into the cache before beginning AES
 - Then the timing doesn't depend on the particular values that you look up
 - Precarious because you might get interrupted in the middle of execution

Side channels \Rightarrow difficult to implement crypto securely

Foot-Shooting Prevention Agreement

I, _____, promise that once Your Name I see how simple AES really is, I will <u>not</u> implement it in production code even though it would be really fun. This agreement shall be in effect until the undersigned creates a meaningful interpretive dance that compares and contrasts cache-based, timing, and other side channel attacks and their countermeasures.



Source: moserware.com/2009/09/stick-figure-guide-to-advanced.html



What you should do

Validate code for timing independence

📮 agl / ctgrind		O
♦ Code ① Issues 0 ⑦ F	Pull requests 1 🛛 🔟 Projects 0 🔸 Puls	se 📊 Graphs
Checking that functions are cons	stant time with Valgrind	
3 commits	🛿 1 branch	O releases
Branch: master New pull request		
Adam Langley C++ support and	constify pointers	
Makefile	Initial import	
	A couple of typos	
Ctgrind.c	C++ support and constify pointers	
Ctgrind.h	C++ support and constify pointers	
🖹 test.c	Initial import	
valgrind.patch	Initial import	
Checking that functions	s are constant time with Valgrind.	

Source: github.com/agl/ctgrind

Use good crypto coding conventions

This page lists coding rules with for each a description of the problem addressed (with a concrete example of failure), and then one or more solutions (with example code snippets).

Contents [hide]
1 Compare secret strings in constant time
1.1 Problem
1.2 Solution
2 Avoid branchings controlled by secret data
2.1 Problem
2.2 Solution
3 Avoid table look-ups indexed by secret data
3.1 Problem
3.2 Solution
4 Avoid secret-dependent loop bounds
4.1 Problem
4.2 Solution
5 Prevent compiler interference with security-critical operations
5.1 Problem
5.2 Solution
6 Prevent confusion between secure and insecure APIs
6.1 Problem
6.2 Bad Solutions
6.3 Solution

Source: cryptocoding.net/index.php/Coding_rules

