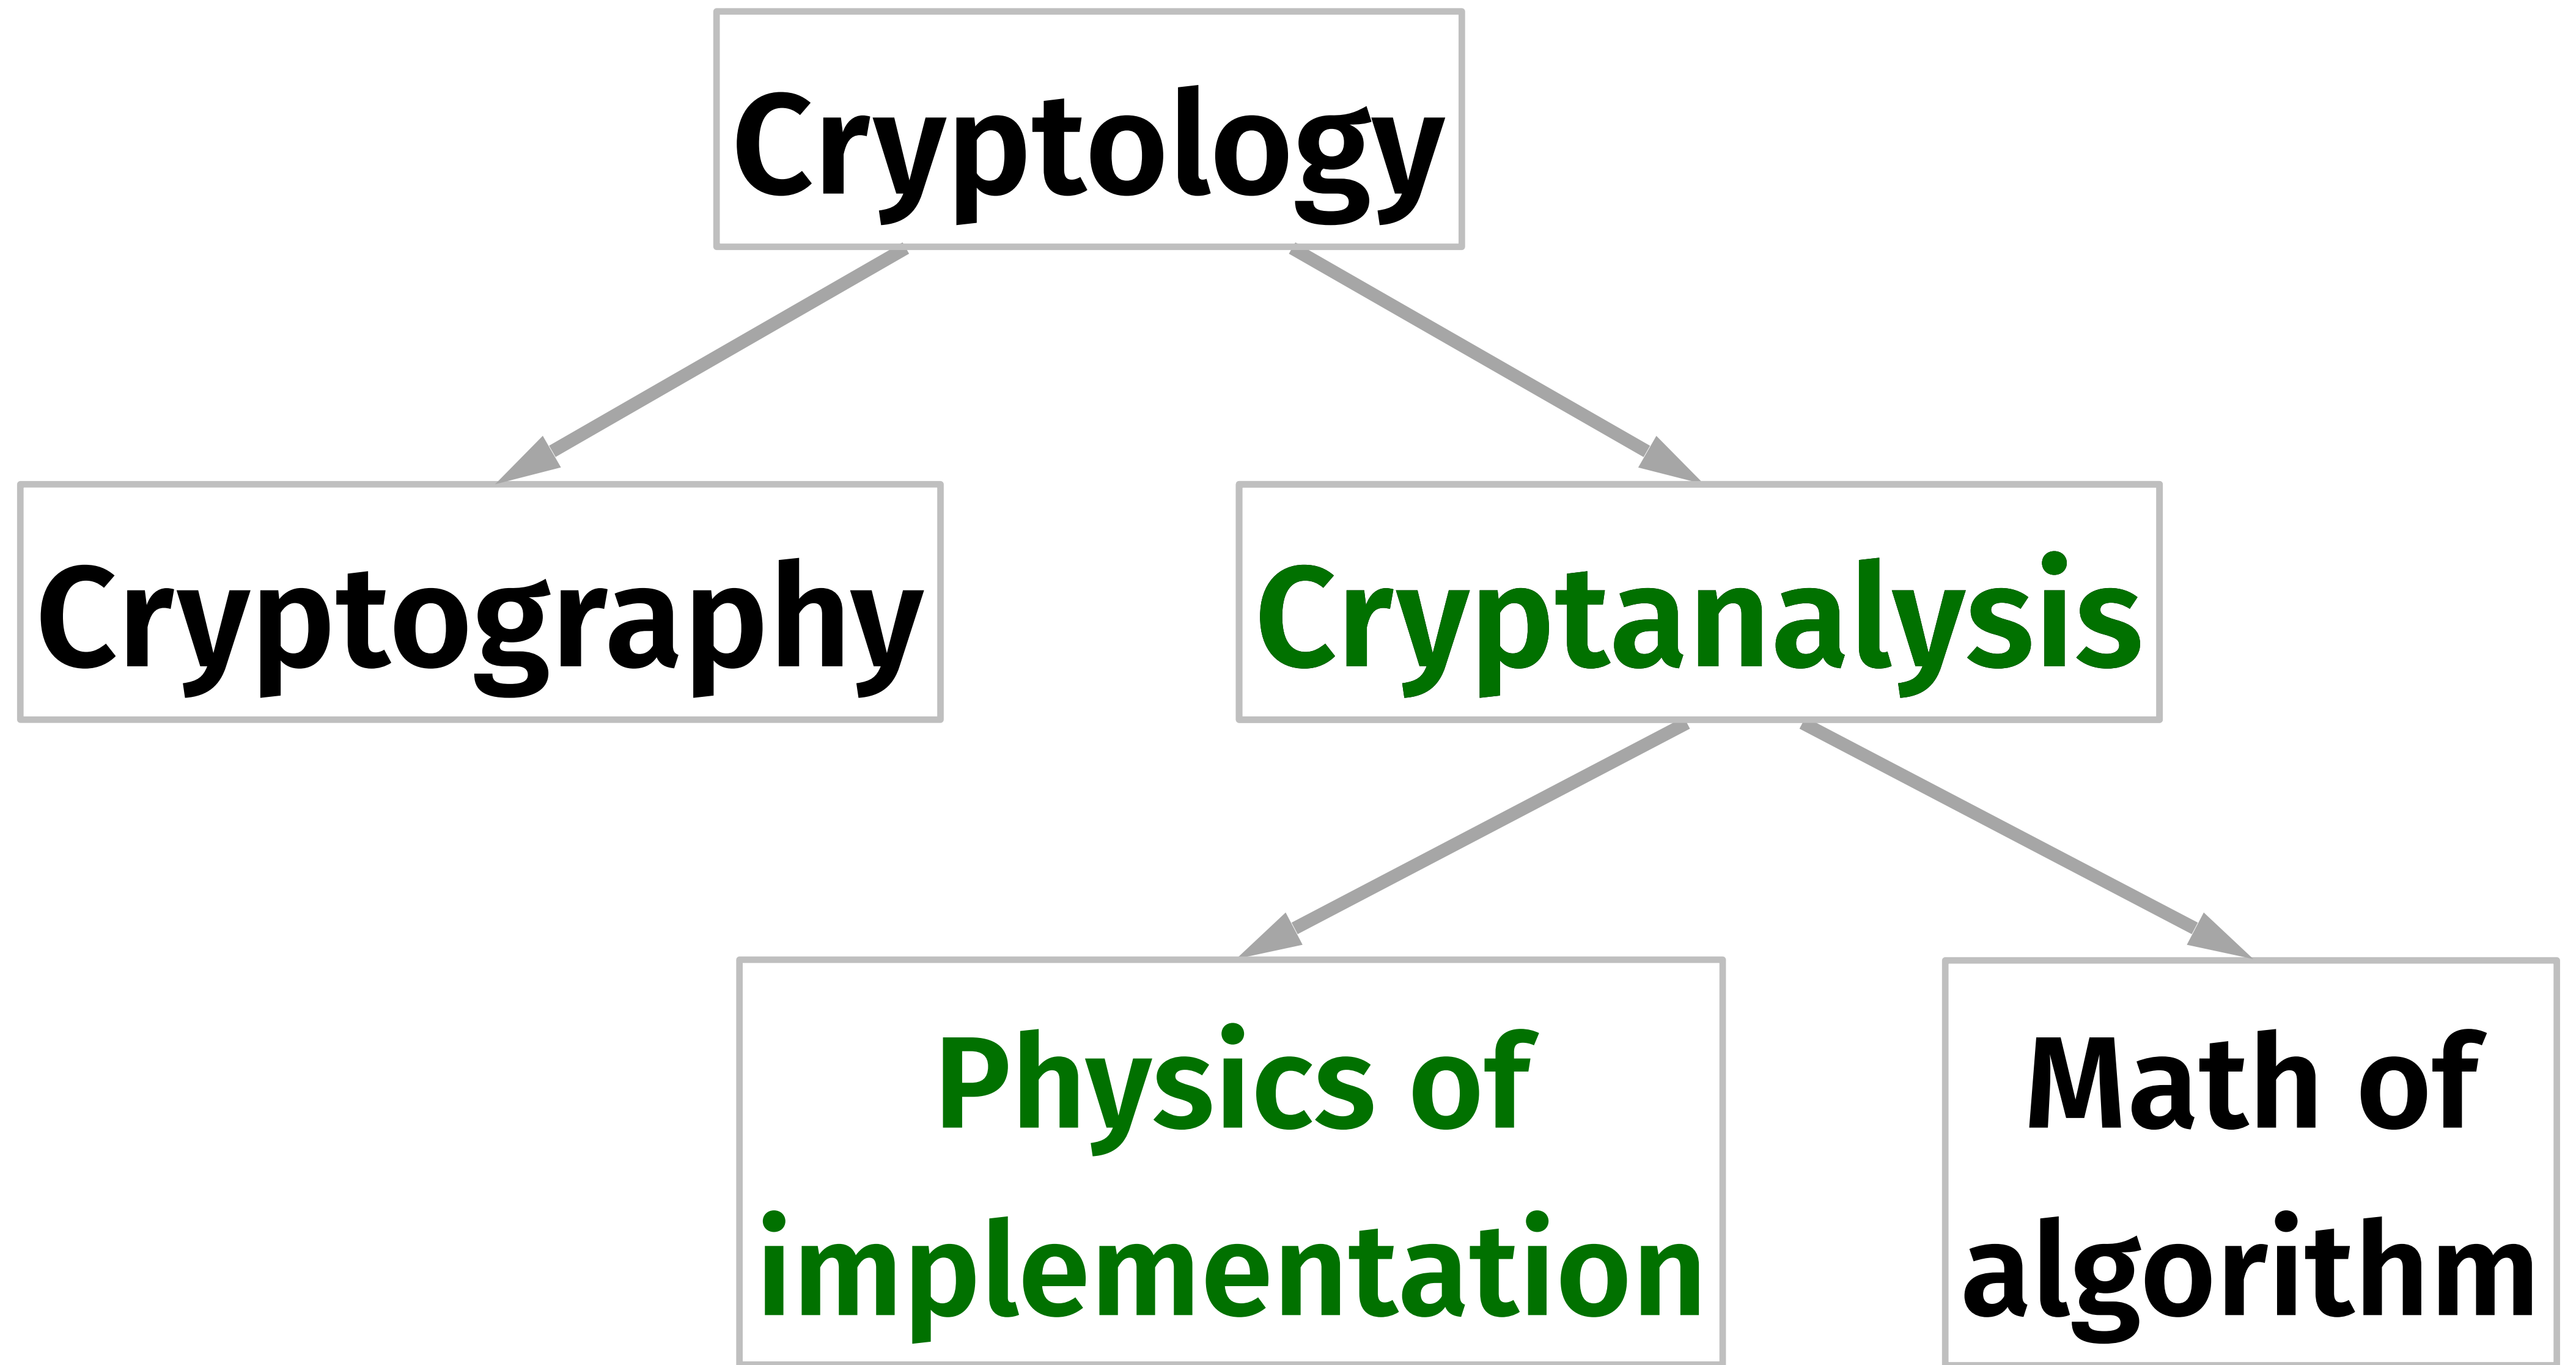


Lecture 12: Authenticated Encryption

- Lab 5 due Friday at 11pm
- No discussion sessions tomorrow
- Nicolas will hold office hours tomorrow at the usual time
- Have a good spring break!



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
not implement it in production code
even though it would be really fun.

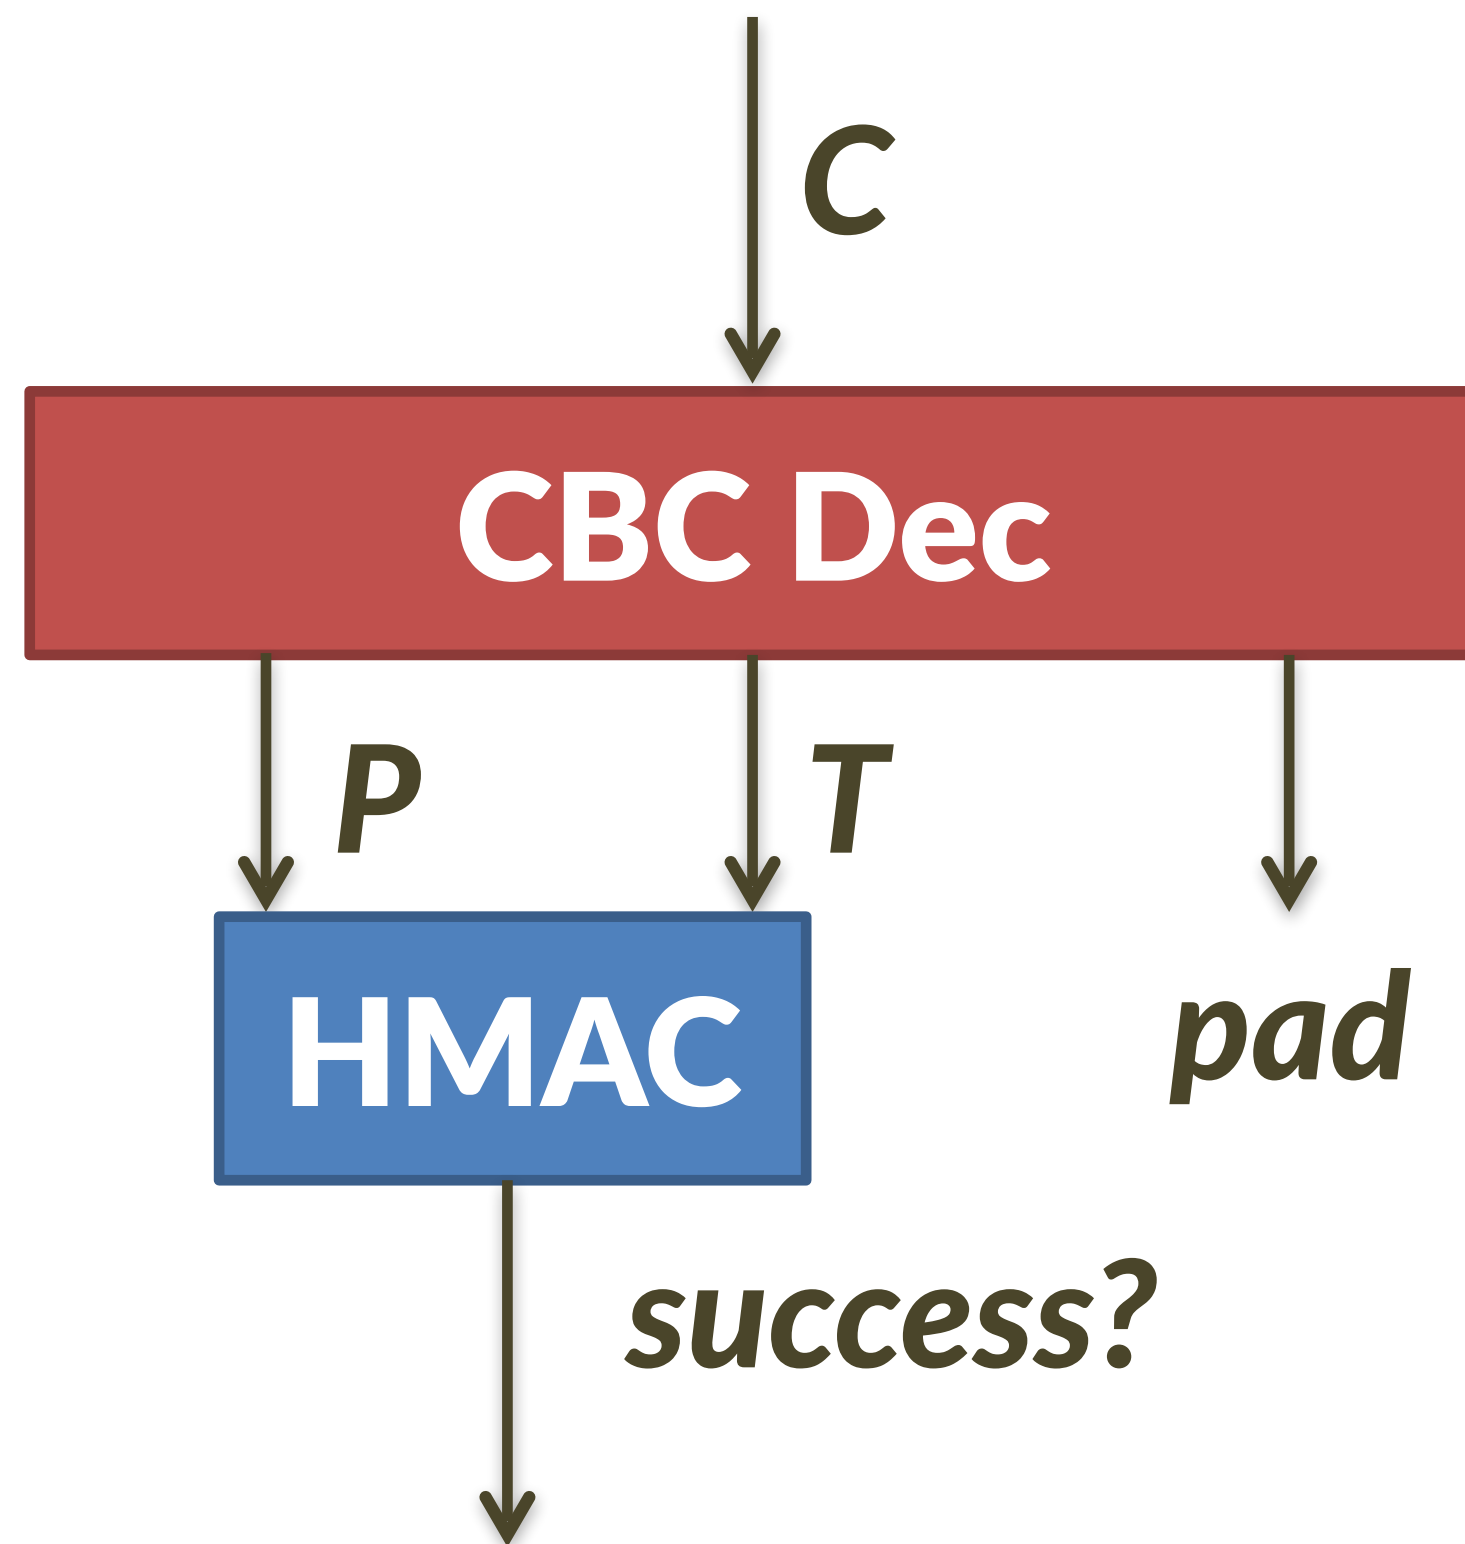
This agreement shall be in effect
until the undersigned creates a
meaningful interpretive dance that
compares and contrasts cache-based,
timing, and other side channel attacks
and their countermeasures.

X _____
Signature Date

Source:

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

Last time: Padding oracle attack



Outcomes

1. Invalid padding
2. Valid padding, wrong HMAC
3. Valid padding, right HMAC

What to do in cases #1 and #2?

- Typical answer: return error message
- We can use error messages to find P !

How can we fix this?

- Remember the three cases

1. Invalid padding
2. Valid padding, wrong HMAC
3. Valid padding, right HMAC

Required effort

- ⇒ Read the padding bytes
- ⇒ Read padding bytes, compute the HMAC
- ⇒ Read padding bytes, compute the HMAC

- Bob's solution: return the *same* error message in cases #1 and #2

- Mallory's countermeasure: can still distinguish the two cases by observing the *time* that the MAC-then-Encrypt system takes to execute!

- Bob's new solution: ensure crypto software's runtime is *independent* of input (i.e., perform the HMAC test whether the padding is correct or not)

- This won't work; Mallory can exploit timing variations within HMAC itself 😞

Software is hard!

- Timing independence is hard

OpenSSL Fact @OpenSSLFact

Jul 24, 2013

/*The aim of right-shifting md_size is so that the compiler doesn't figure out that it can remove div_spoiler...which I hope is beyond it.*/

- So is software in general

OpenSSL Fact @OpenSSLFact

Sep 3, 2012

/* [we should] obviate the ugly and illegal kludge in CRYPTO_mem_leaks_cb. Otherwise the code police will come and get us.*/

OpenSSL Fact @OpenSSLFact

Jan 22, 2013

/* EEK! Experimental code starts */

OpenSSL Fact @OpenSSLFact

Sep 5, 2012

/* BIG UGLY WARNING! This is so damn ugly I wanna puke ... ARGH! ARGH! ARGH! Let's get rid of this macro package. Please? */


- So are compilers in general



Mudge @dotMudge · Jan 25

Modern compilers make a lot of optimizations and perform advanced heuristics to determine what to emit. The resulting binaries have many (attack-able) components you cannot learn from the source alone.

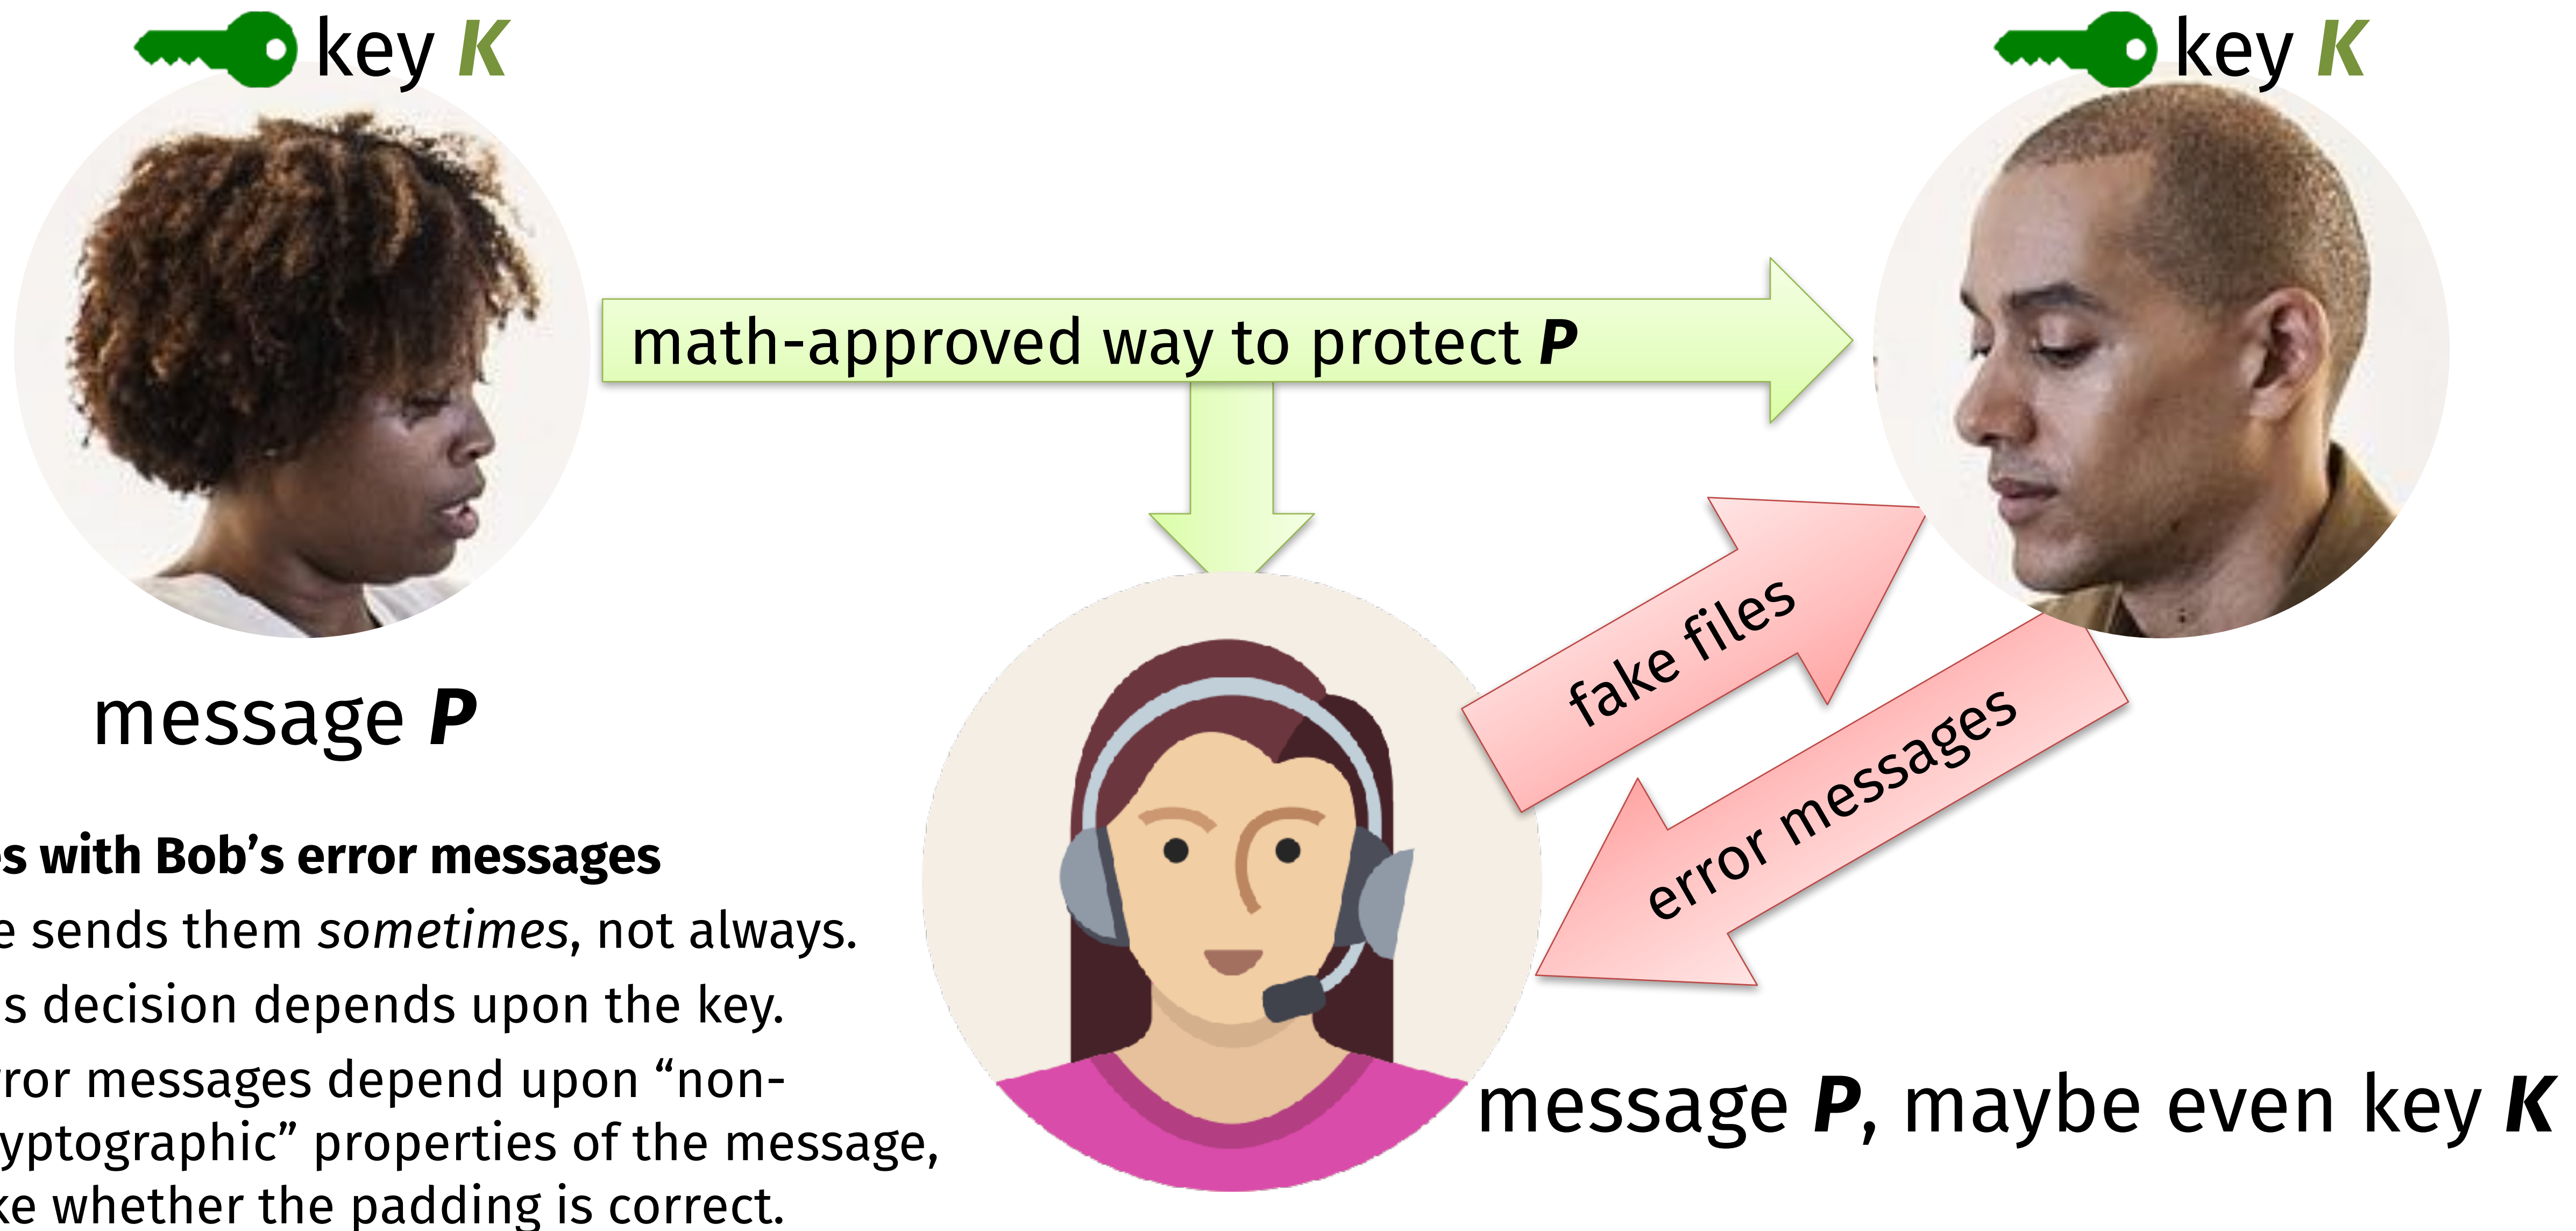
Source is the intent, the binary is reality.



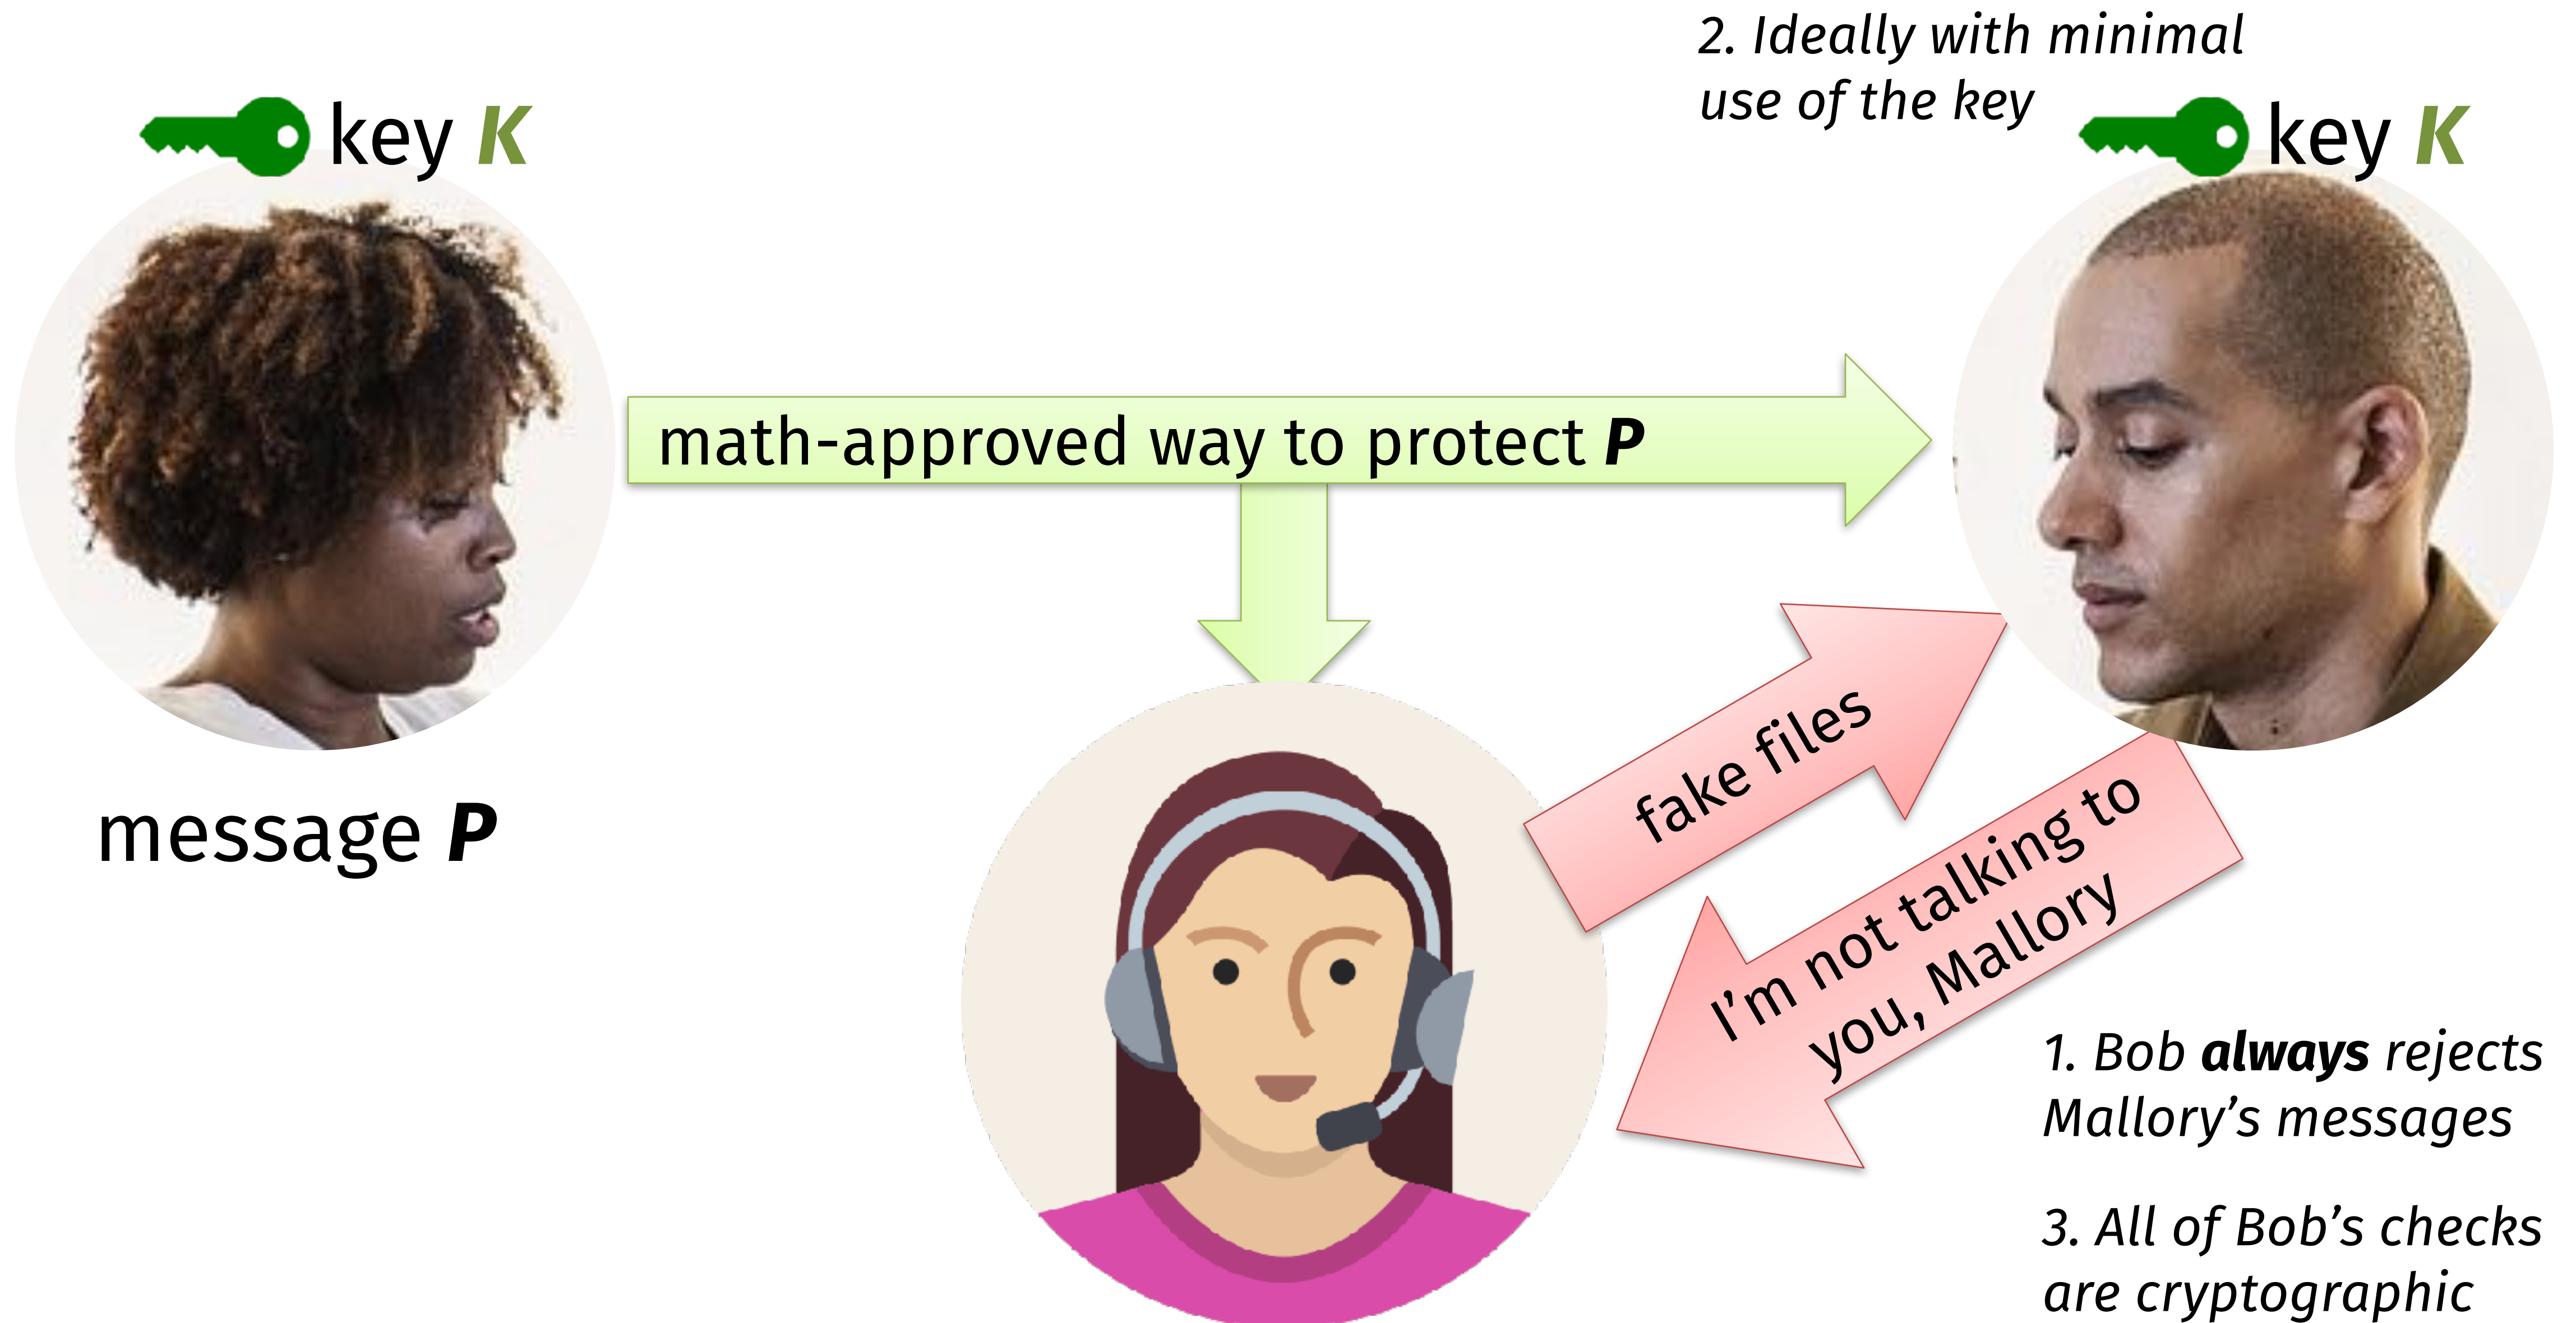
Steven Bellovin @SteveBellovin · Jan 25

My favorite is how hard it is to zero out a cryptographic key that you're done with--the optimizer says "this variable is never used again", so it deletes the zeroize operation.

Part 2: Breaking crypto via side channels



Our desired countermeasure



“Confidentiality xor authenticity is not possible. If you don't have both, often you don't have either.”

– *Prof. Matthew Green, Johns Hopkins*


“If you have to perform any cryptographic operation before verifying the MAC on a message you’ve received, it will somehow inevitably lead to doom!”

– *Moxie Marlinspike*

Encryption xor Authentication

Privacy

IND\$-CPA against
nonce-respecting Eve




Authenticity

Even after viewing many (A, T) pairs,
Mallory cannot forge a new one



First, let's strengthen privacy

IND\$-CPA



IND\$-CCA

Same thing, but now Mallory has access to encryption *and* decryption oracles



What is the connection to padding oracles?


Formalizing IND\$-CCA

Comprises 3 algorithms:

- $\text{KeyGen}(\lambda)$ outputs a key $K \leftarrow \{0,1\}^\lambda$
- $\text{Encrypt}_K(\text{message } P, \text{nonce } N) \rightarrow C$
- $\text{Decrypt}_K(\text{ciphertext } C, \text{nonce } N) \rightarrow P$

Satisfies 3 constraints

- **Performance:** all 3 algorithms are efficiently computable
- **Correctness:** $\text{Dec}_K^{-1}(\text{Enc}_K(P, N)) = P$ for all $K \in \{0,1\}^\lambda$, $N \in \{0,1\}^\mu$, and $P \in \{0,1\}^*$




- **(q, t, ε)-IND\$-CCA:** for every *nonce-respecting* adversary **A** who makes $\leq q$ queries and runs in time $\leq t$,

$$A^{\text{Enc}_K, \text{Dec}_K} \approx_{q,t,\epsilon} A^{\$, \$^{-1}}$$


where $\$$ responds randomly and so does $\$^{-1}$ subject to consistency with $\$$

Combining Enc and MAC *generically*


Enc and MAC



MAC then Enc



Enc then MAC




Intuitive concerns with MAC then Enc


- The private data P is authenticated, but C is not!
- Recipient must perform decryption before knowing whether the message is authentic

Combining Enc and MAC *generically*


Enc and MAC



MAC then Enc



Enc then MAC



Confidentiality	None	CPA	CCA!
Integrity	Plaintext integrity: Cannot make CT that decrypts to message that sender never encrypted		Ciphertext integrity: Cannot make new valid CTs, only know sender-made ones

Formalizing ciphertext integrity

- Goal: Mallory cannot make a valid CT that wasn't previously made by sender
- Imagine that Mallory is trying to perform a padding oracle attack
- If she spams Bob with malformed CTs, now he simply rejects them all!



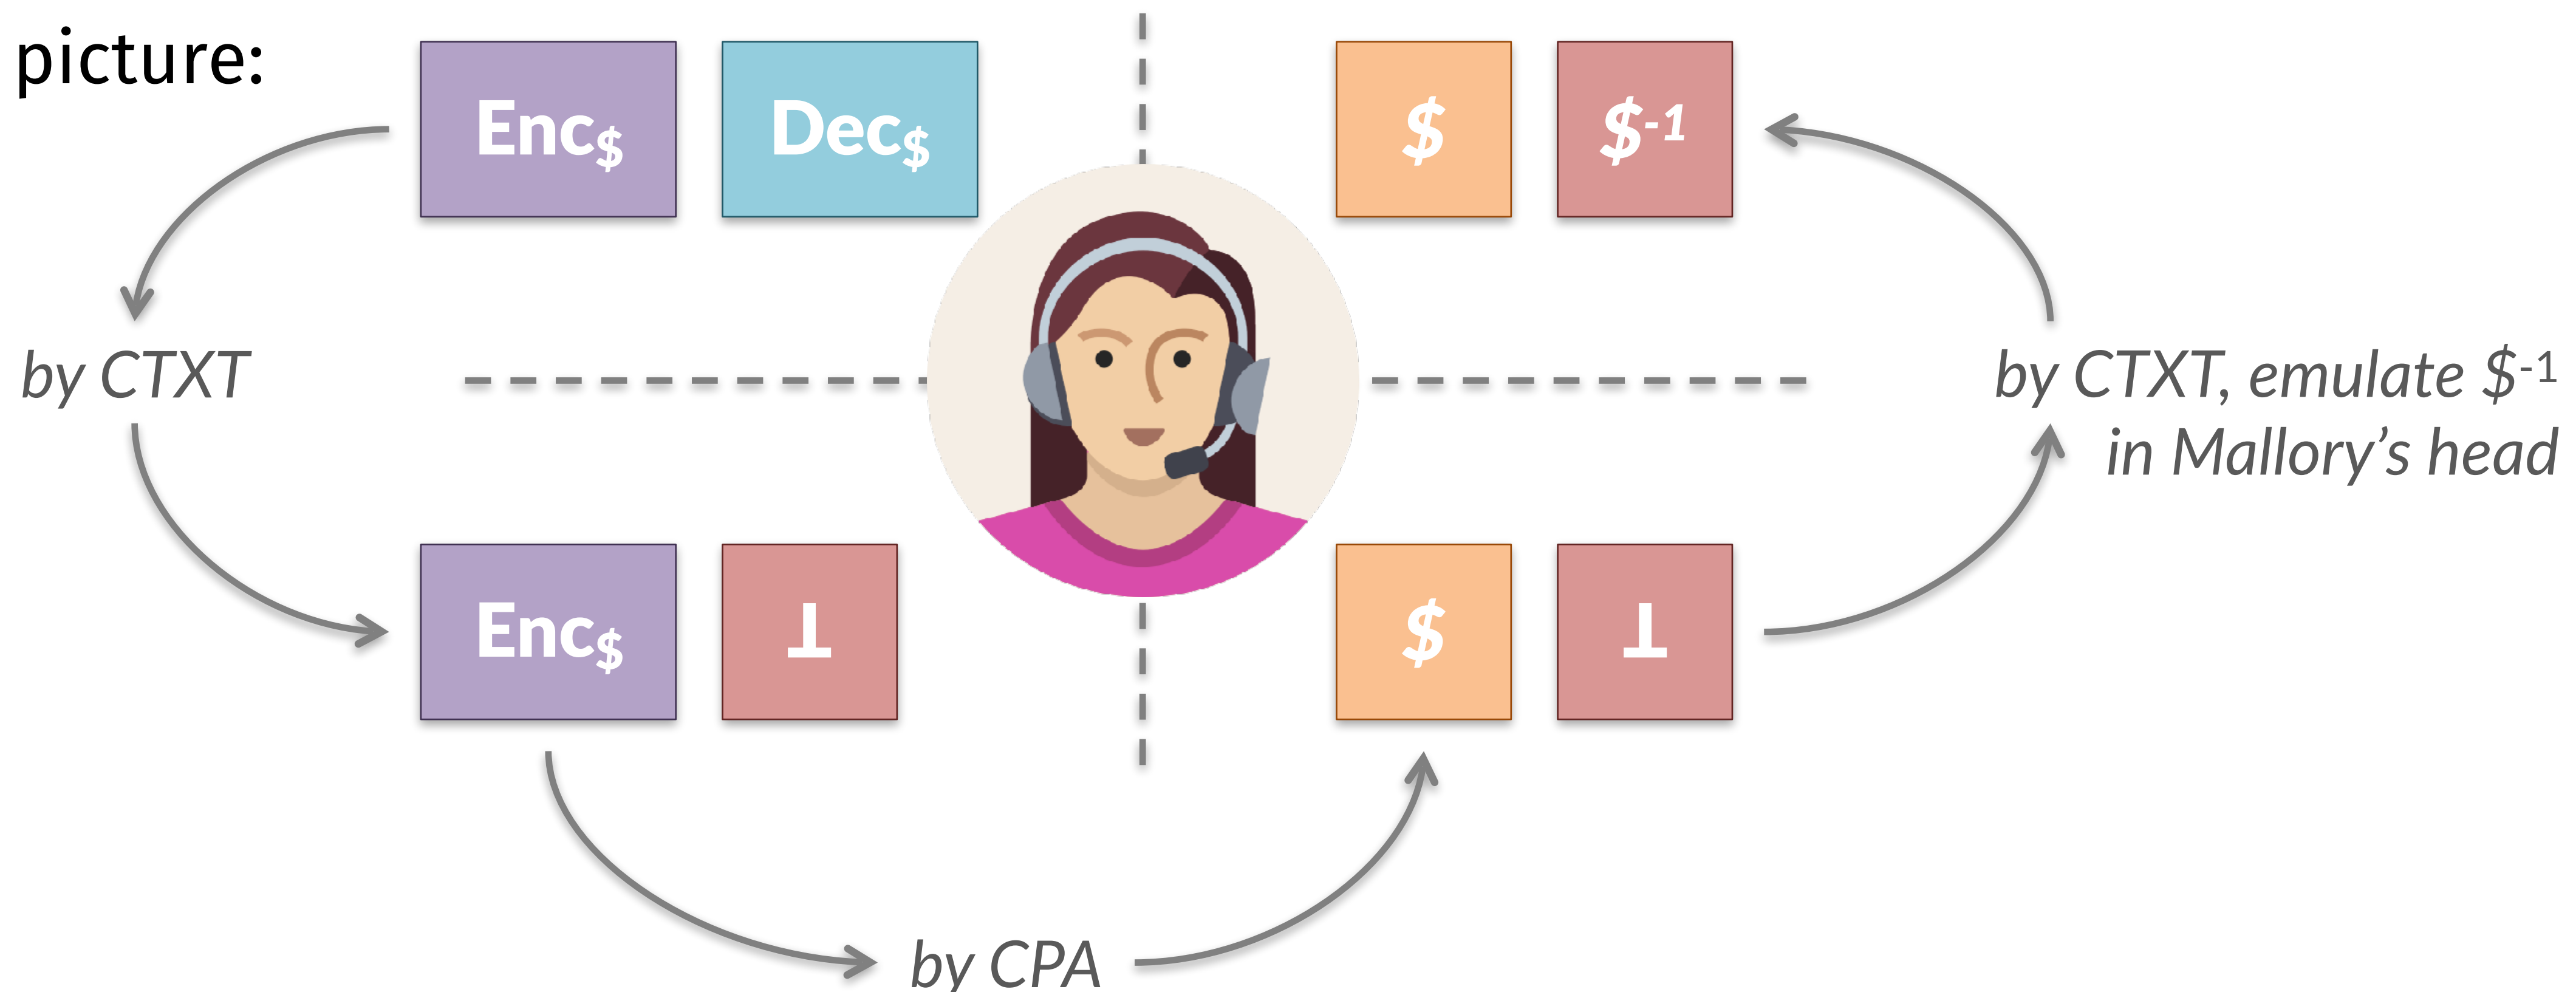
Operation: This box returns a single “integrity failure” error message no matter what Mallory submits!

Restriction: Mallory cannot attempt to decrypt ciphertexts that are the result of prior encryptions.

Relating integrity and confidentiality

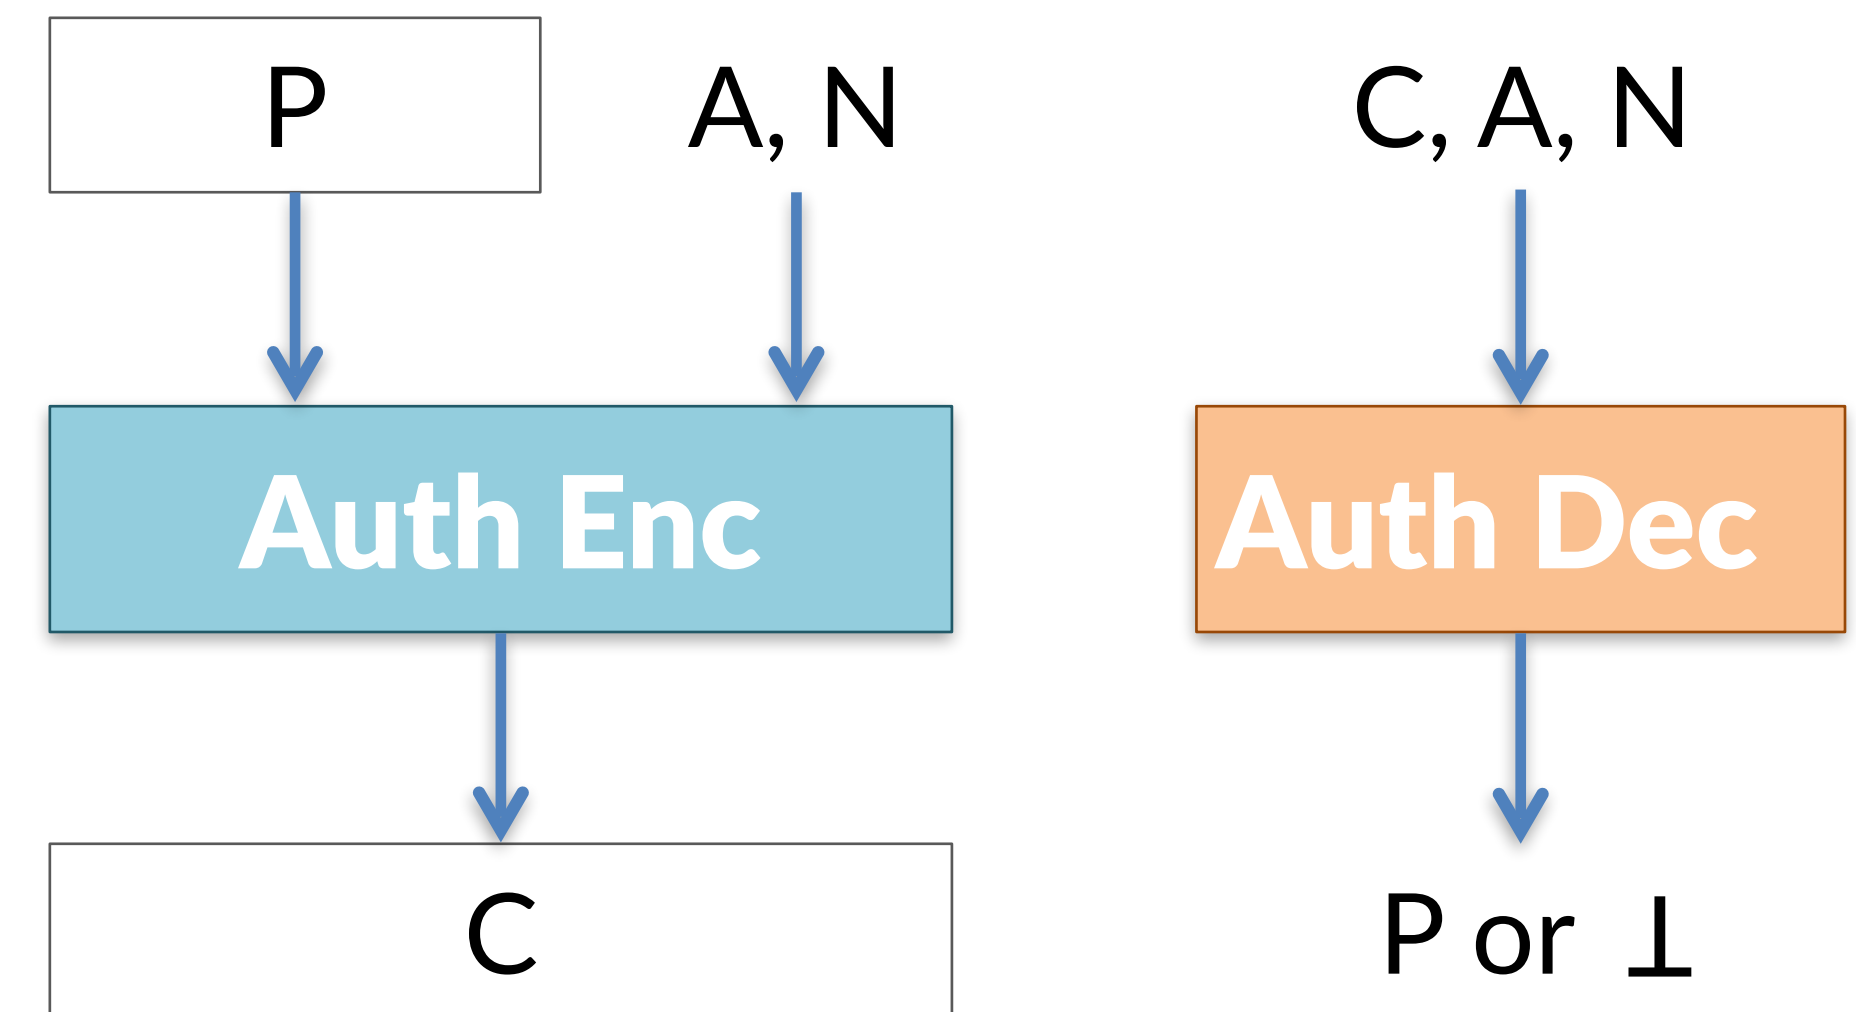
- **Thm.** Suppose that an encryption scheme provides (q, t, ϵ_1) -CPA privacy and (q, t, ϵ_2) -ciphertext integrity. Then, it also provides $(q, t, \epsilon_1 + 2\epsilon_2)$ -CCA privacy.
- Intuition: If Mallory can't forge new messages, then Dec oracle useless to her

- Proof by picture:



Def. Authenticated Encryption with Associated Data (AEAD)

- **KeyGen**: randomly chooses key, as always
- **Enc**(authenticated data A, private + auth data P, nonce N) \rightarrow ciphertext C of length $|C| > |P|$
- **Dec**(C, A, N) \rightarrow P or \perp



Why combine authentication and encryption?

- Better security: resist some of these physical side channel attacks
- Simplicity: developers have fewer decisions (i.e., opportunities for mistakes)
- Performance: save in time + space costs, also often only need 1 key

AEAD as a picture

