CIS 4930:
Secure IoT

Lecture 4

Prof. Kaushal Kafle

Derived from slides by Adwait Nadkarni, William Enck, Micah Sherr and Patrick McDaniel

Class Notes

- **Project proposals + team due today!**
  - Required to visit me personally with your group asap (within the next week).
- Project has 2 sections:
  - **Section 1**: Focus on understanding an IoT Platform (09/05-10/10)
  - **Section 2**: Security analysis of real-world IoT apps (10/17-12/12)
- Proposed deadline changes:
  - Section 1: 09/05 – 10/22
  - Section 2: ~~10/29~~ 10/24 – 12/12

Class Notes

- Homework 2 assigned today after class. Due on 09/19
- Bug bounty:
  - Include a short paper summary when you pitch your bug
  - The paper proposes X. They perform Y. The results/findings are Z.

- **Meet the TA - Mohamed Mossad**
- *Email*: mohamed15@usf.edu
- *Office hours*: Mondays 1-2pm
- *Office space*: Link

Block ciphers: Generic Block Encryption

- Converts one input plaintext block of fixed size b bits to an output ciphertext block also of *b* bits
- Benefits of large *b*? of short *b*?
  - Think cryptanalysis!
- Block and key size are *separate parameters*
- E.g., AES, DES

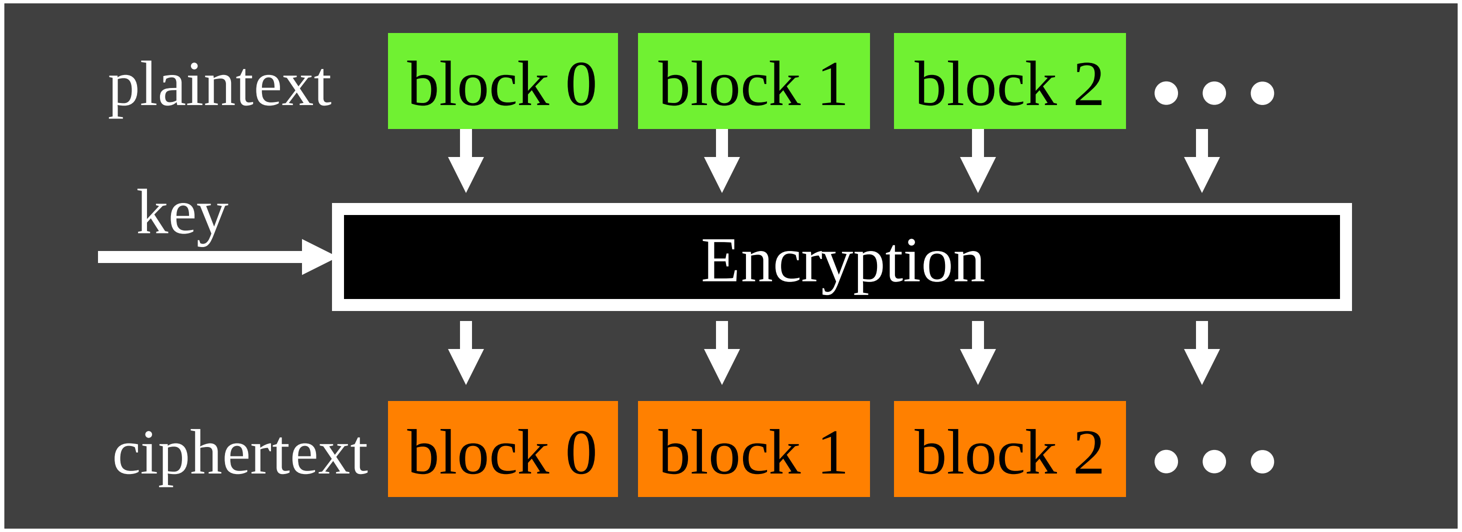

Two Principles for Cipher Design

- *Confusion*: Make the relationship between the <plaintext, key> input and the <ciphertext> output as complex (non-linear) as possible
  - Mainly accomplished by *substitution*
- *Diffusion*: Spread the influence of each input bit across many output bits
  - Mainly accomplished by *permutation*
- Idea: use *multiple*, *alternating* permutations and subsitutions
  - *S→P→S→P→S→… or P→S→P→S→P→…*
  - Does it have to alternate?, e.g., S→S→S→P→P→P→S→S→…

-> Stream ciphers do not have diffusion! Why?

Two Principles for Cipher Design

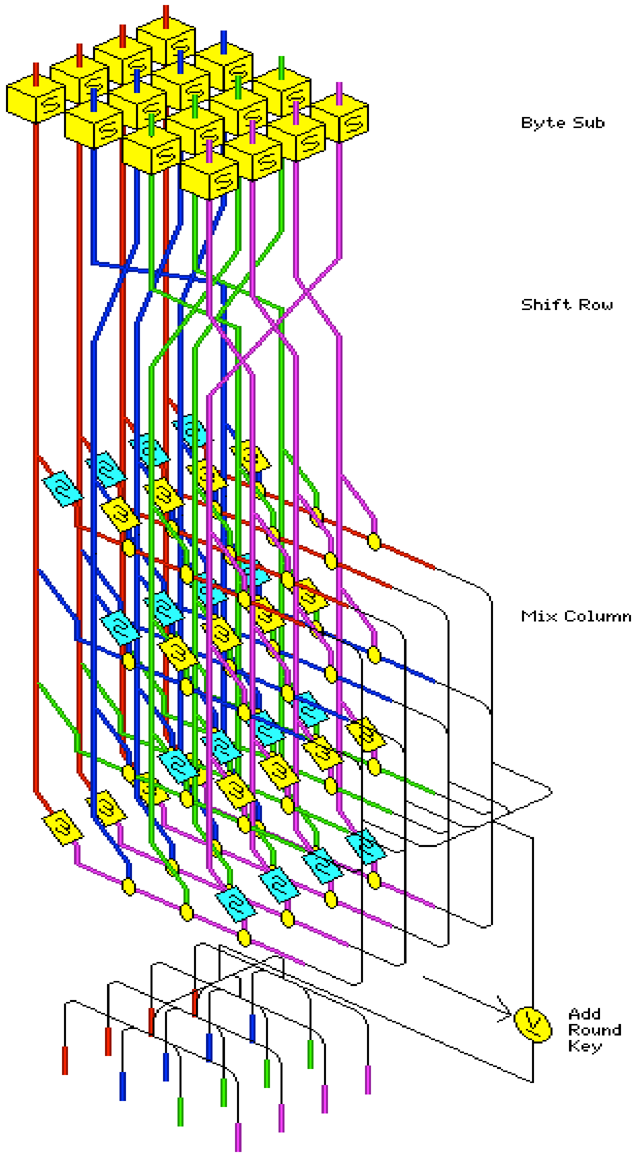

Two Principles for Cipher Design

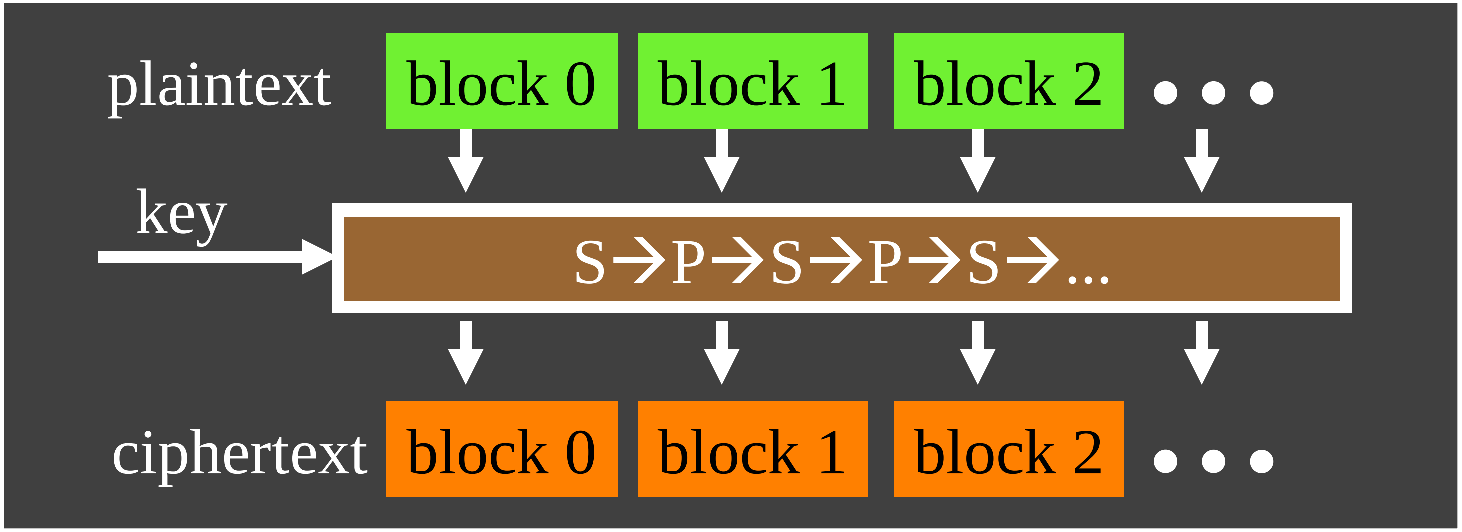

- Can I *predictably* change the plaintext, by changing the ciphertext?
  - No. The relationship is too complex.

Modes of Operation

- Most ciphers work on blocks of fixed (small) size
- How to encrypt long messages?
- Modes of operation
  - ECB (Electronic Code Book)
  - CBC (Cipher Block Chaining)
  - CTR (Counter)
  - (there are many more; we will look at 3 for a bare minimum understanding)

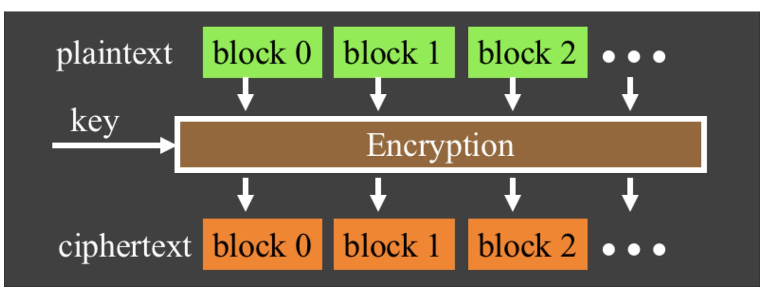

Issues for Block Chaining Modes

- *Information leakage*: Does it reveal info about the plaintext blocks?
- *Ciphertext manipulation*: Can an attacker modify ciphertext block(s) in a way that will produce a predictable/desired change in the decrypted plaintext block(s)?
  - Note: assume the structure of the plaintext is known, e.g., first block is employee #1 salary, second block is employee #2 salary, etc.
- *Parallel/Sequential*: Can blocks of plaintext (ciphertext) be encrypted (decrypted) in parallel?
- *Error Propagation*: If there is an error in a plaintext (ciphertext) block, will there be an encryption (decryption) error in more than one ciphertext (plaintext) block?

Electronic Code Book (ECB)

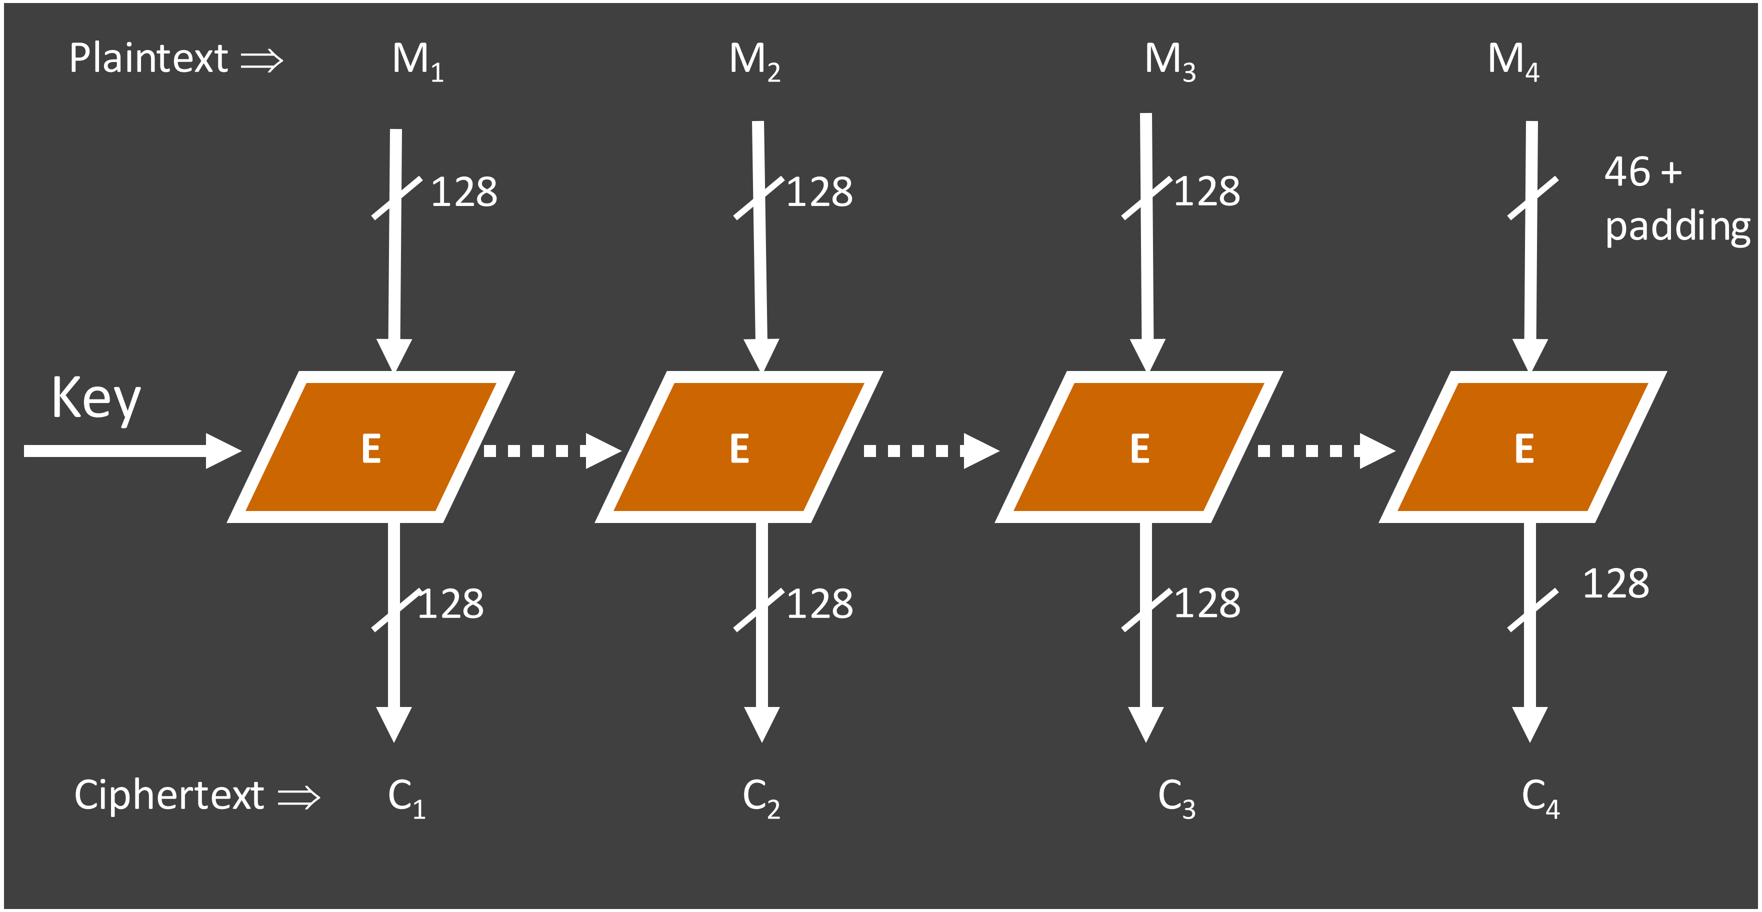

- The easiest mode of operation; each block is independently encrypted

ECB Decryption

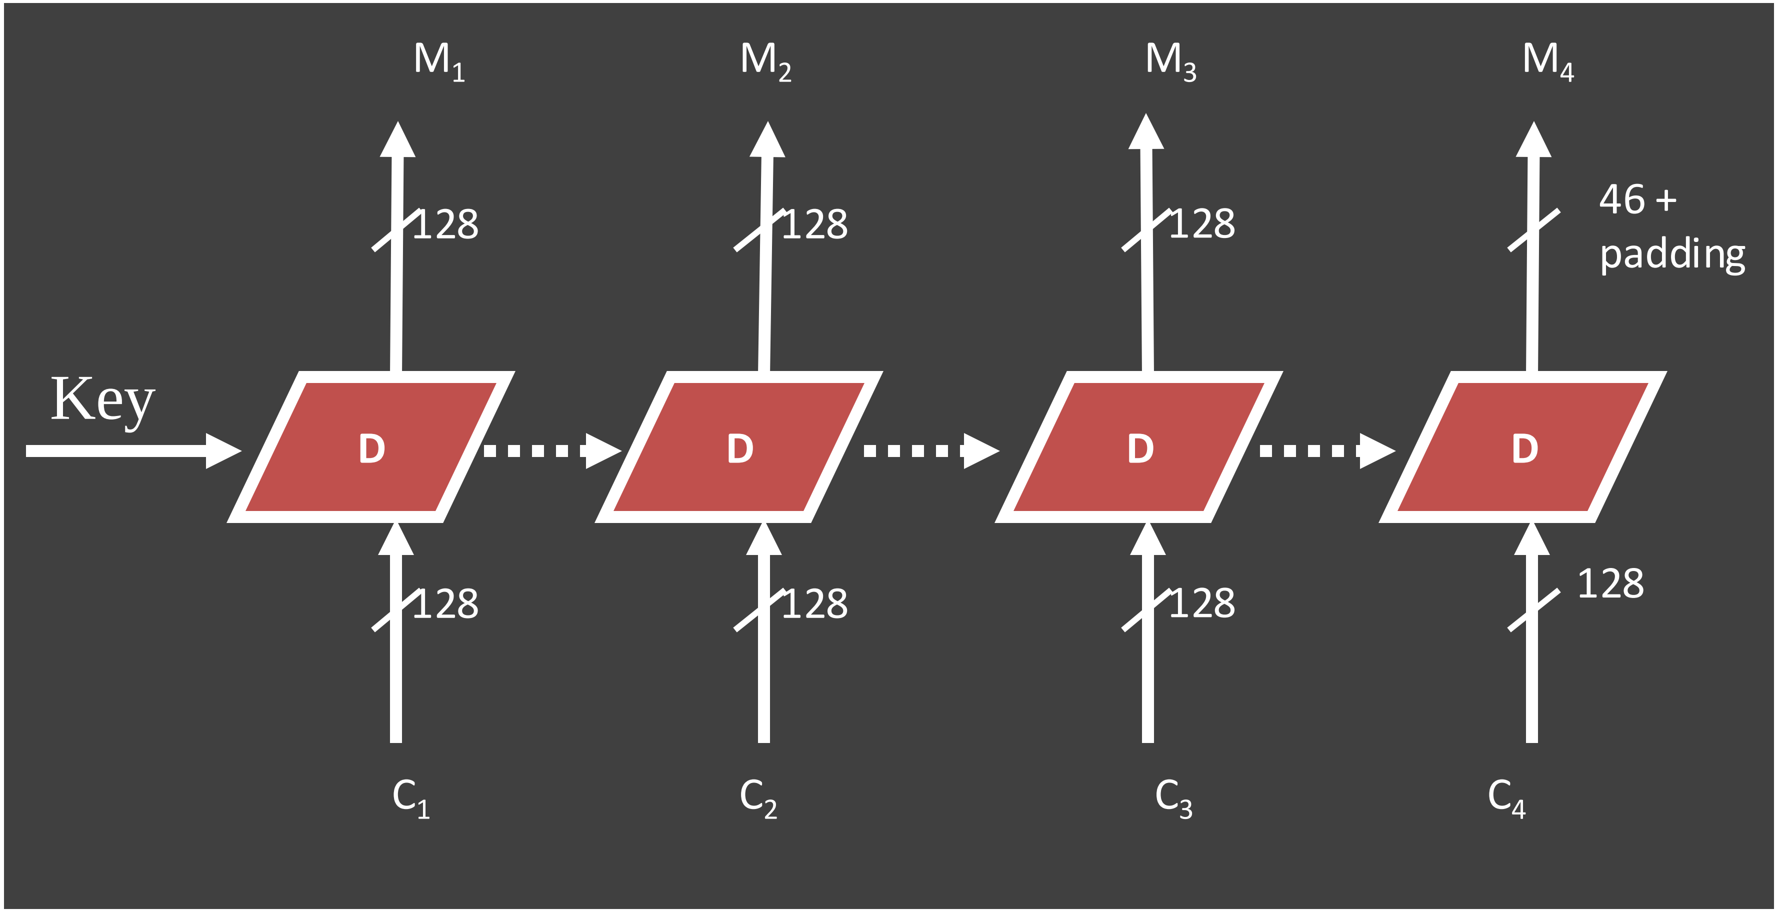

- Each block is independently decrypted

ECB Issues

- *Information leaks*: two ciphertext blocks that are the same
- *Manipulation*: switch ciphertext with predictable results on plaintext (e.g., shuffle).
- *Parallel*: yes
- *Error Propagate*: no

Plaintext

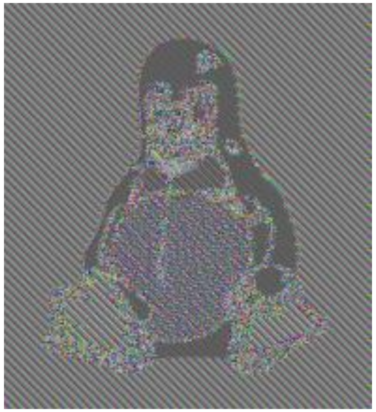

ECB

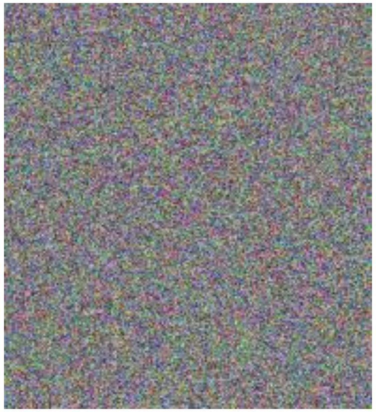

Other modes

4:13 Ode to ECB

by Ben Nagy

Oh little one, you're growing up
You'll soon be writing C
You'll treat your ints as pointers
You'll nest the ternary
You'll cut and paste from github
And try cryptography
But even in your darkest hour
Do not use ECB

CBC's BEASTly when padding's abused
And CTR's fine til a nonce is reused
Some say it's a CRIME to compress then encrypt
Or store keys in the browser (or use javascript)
Diffie Hellman will collapse if hackers choose your g
And RSA is full of traps when e is set to 3
Whiten! Blind! In constant time! Don't write an RNG!
But failing all, and listen well: Do not use ECB

They'll say "It's like a one-time-pad!
The data's short, it's not so bad
the keys are long—they're iron clad
I have a PhD!"
And then you're front page Hacker News
Your passwords cracked—Adobe Blues.
Don't leave your penguin showing through,
Do not use ECB

Sometimes it can seem like there's ECB everywhere. ECB on TV, ECB in music, it's endless. But that doesn't make it safe. Or right. So tune out and avoid ECB, no matter what your friends, the TV, or your favourite cryptographer tells you.

Canadian Joke Council

True Bugs Wait ♡

@natashenka
#truebugswait

Cipher Block Chaining (CBC)

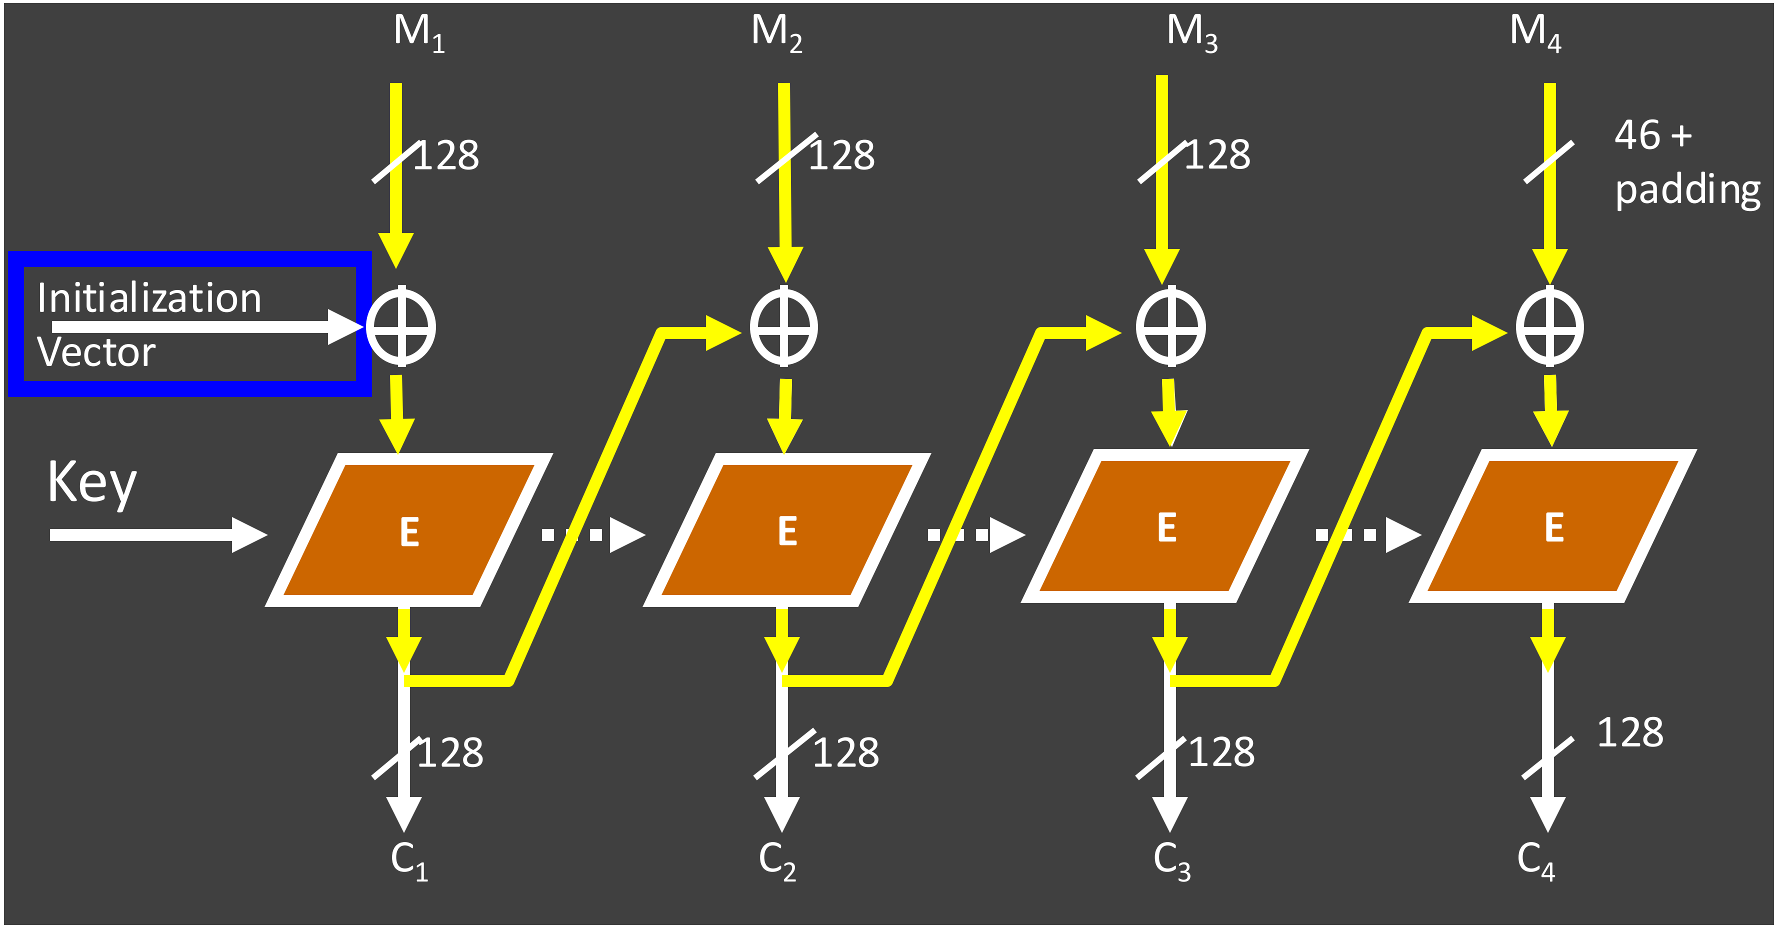

- Chaining dependency: each ciphertext block depends on all preceding plaintext blocks

Initialization Vectors

- Initialization Vector (IV)
  - Used along with the key; not secret
  - For a given plaintext, changing either the key, or the IV, will produce a different ciphertext
  - Why is that useful?
- IV generation and sharing
  - Random; may transmit with the ciphertext
  - Incremental; predictable by receivers

CBC Decryption

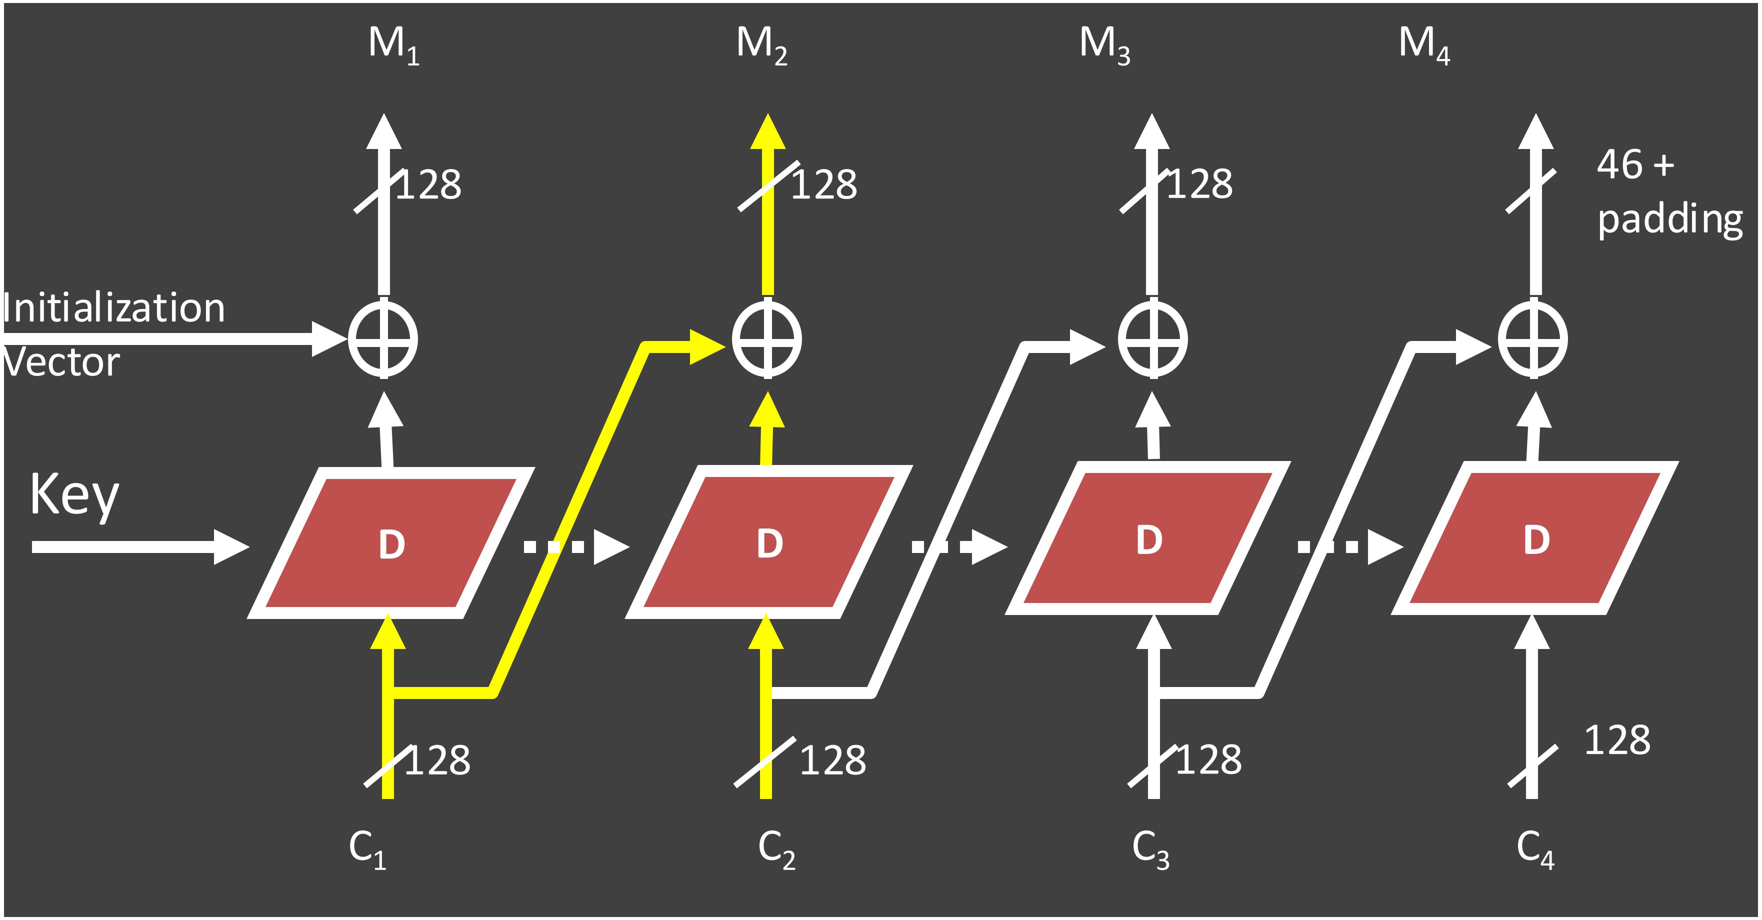

- How many ciphertext blocks does each plaintext block depend on?

CBC Properties

- Does information leak?
  - Identical plaintext blocks will produce different ciphertext blocks
- Can ciphertext be manipulated profitably?
  - Yes
- Parallel processing possible?
  - no (encryption), yes (decryption)
- Do ciphertext errors propagate?
  - yes (encryption), a little (decryption)

Counter Mode (CTR)

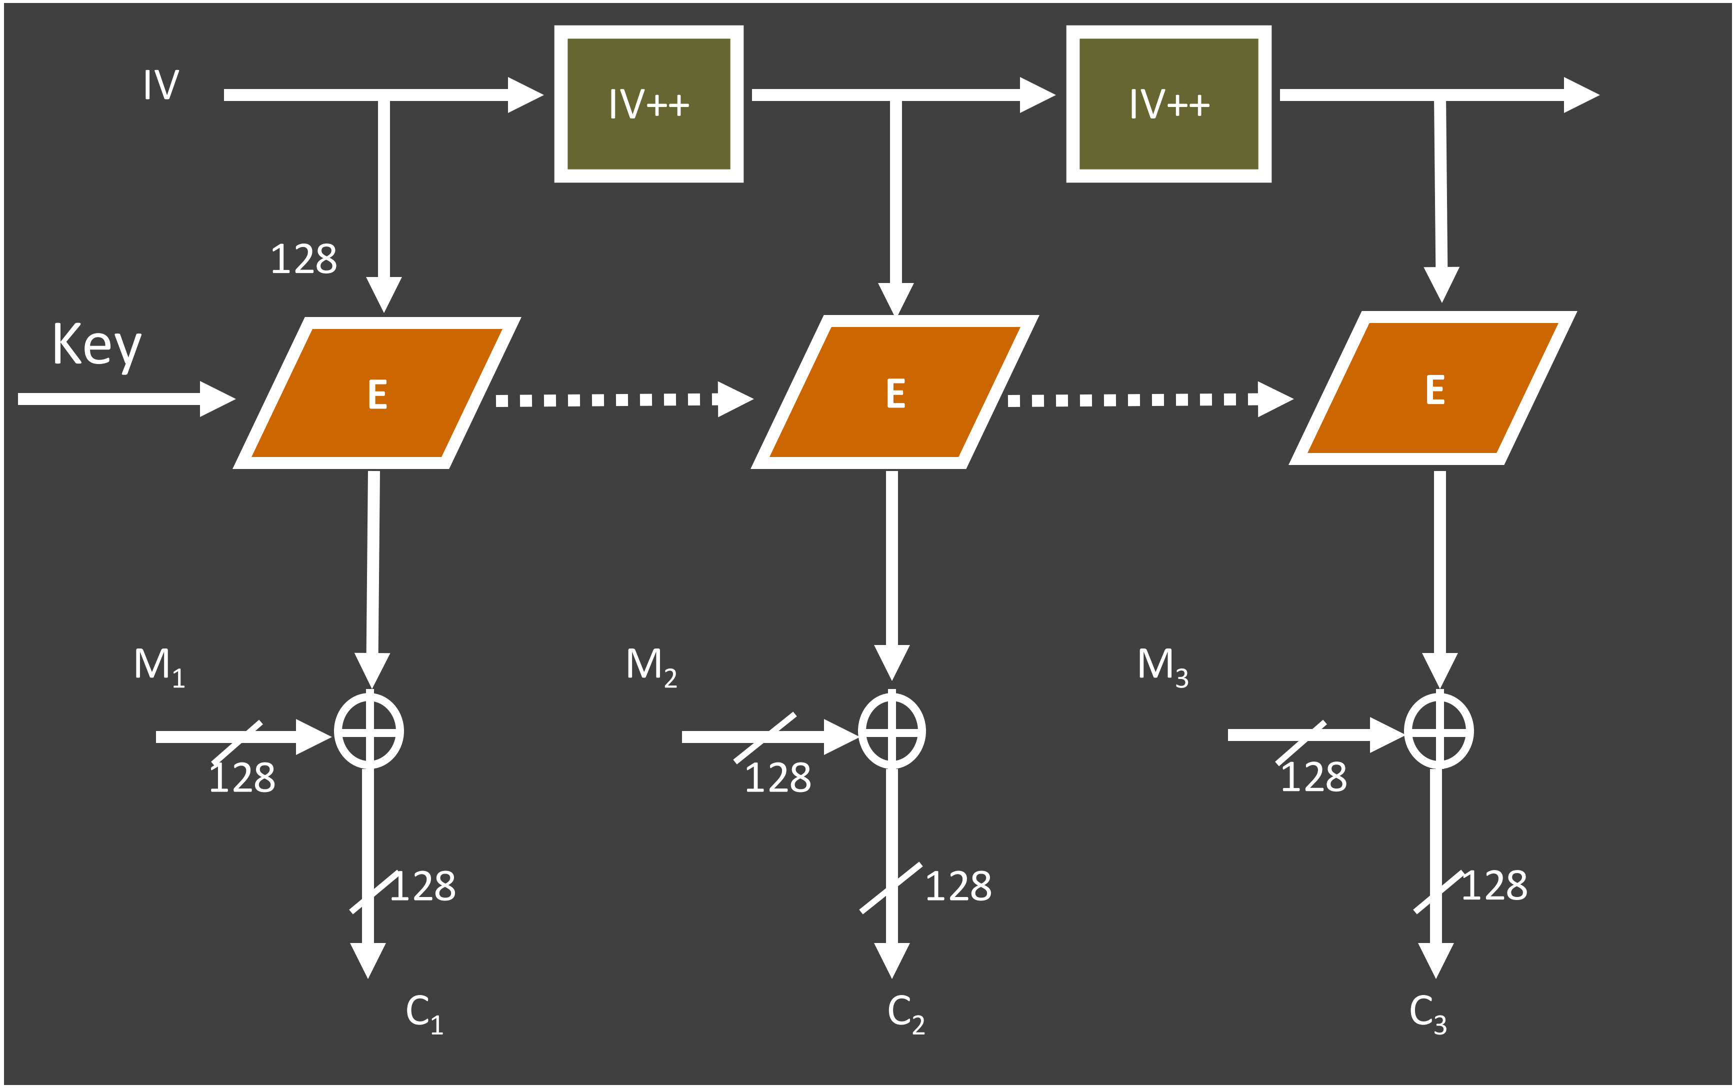

CTR Mode Properties

- Does information leak?
  - Identical plaintext block produce different ciphertext blocks
- Can ciphertext be manipulated profitably
  - Yes!
- Parallel processing possible
  - Yes (both generating pad and XORing)
- Do ciphertext errors propagate?
  - No.
- Allow decryption the ciphertext at any location
  - Ideal for random access to ciphertext

What encryption does and does not

- Does:
  - confidentiality
- Doesn’t do:
  - data integrity
  - source authentication
- **Need:** ensure that data is not altered and is from an authenticated source

Principals

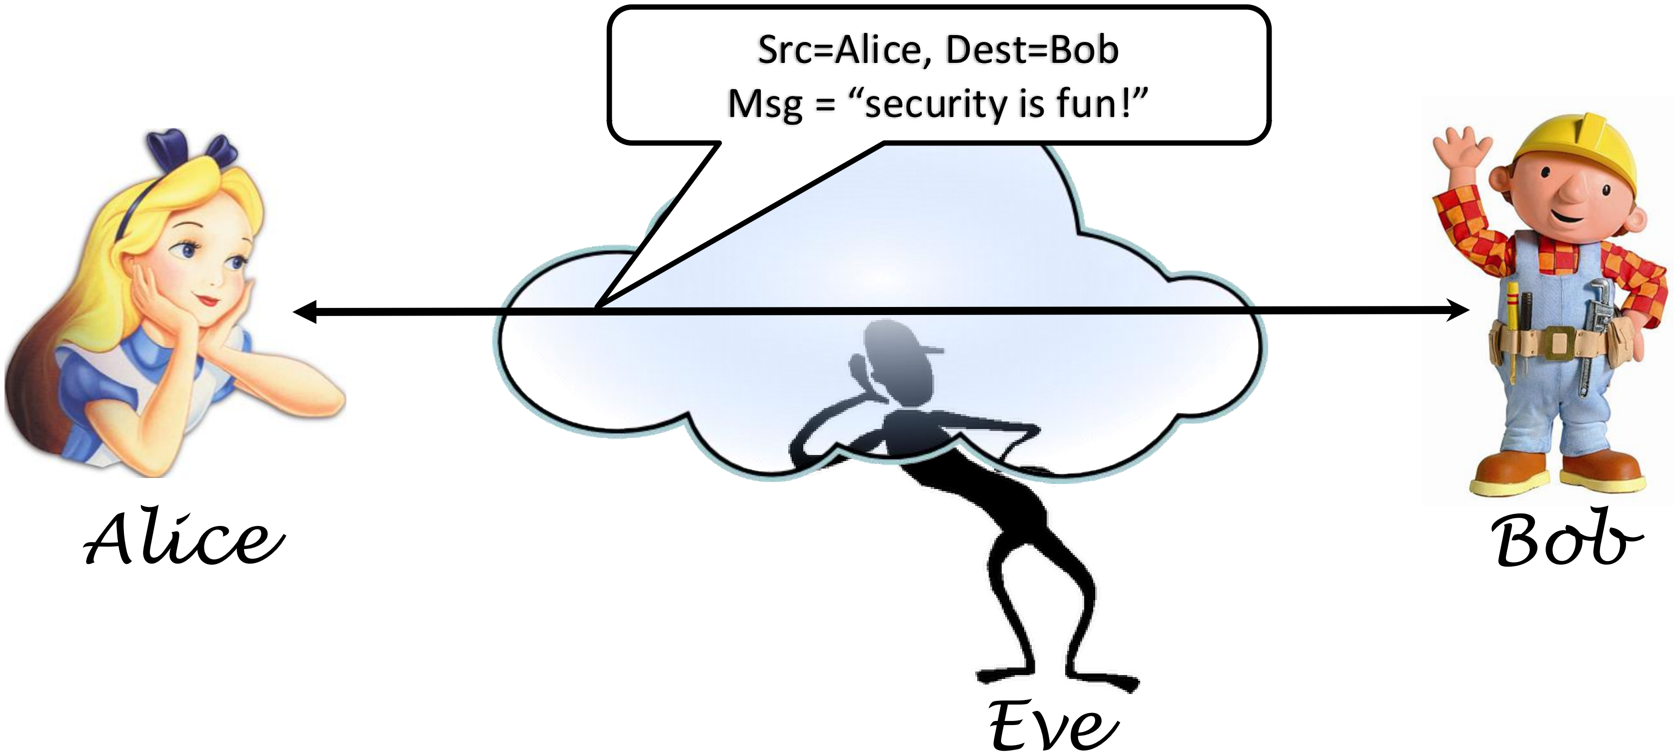

Man-in-the-Middle (MitM) attack

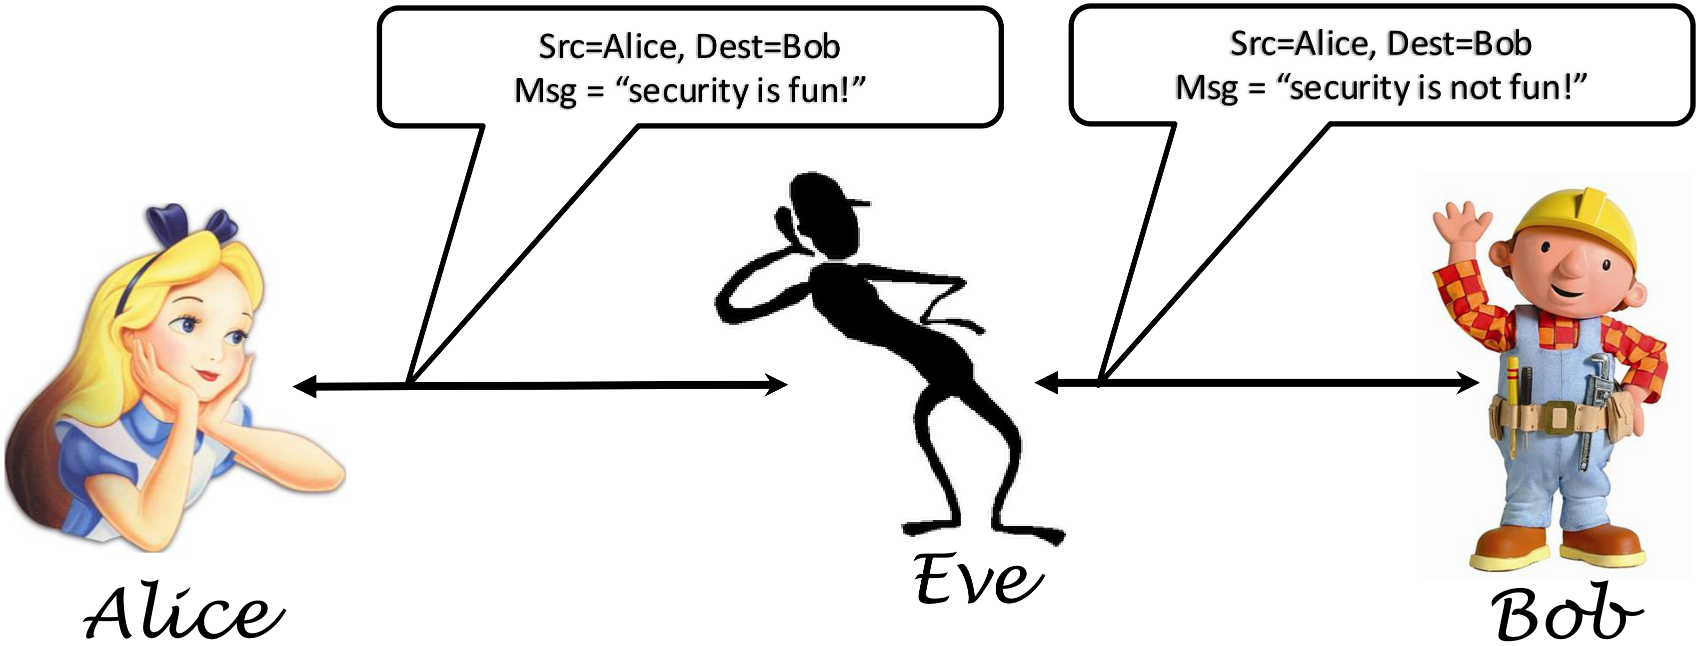

- For confidentiality, just encrypt.
- How do we ensure integrity?

Message Authentication Codes (MACs)

- MACs provide message **integrity** and **authenticity**
- $MAC_K(M)$ – use symmetric encryption to produce **short sequence of bits** that depends on both the message (M) and the key (K)
- MACs should be resistant to **existential forgery**: Eve should not be able to produce a valid MAC for a message M' without knowing K
- To provide confidentiality, authenticity, and integrity of a message, Alice sends
  - **$E_K(M, MAC_K(M))$** where $E_K(X)$ is the encryption of X using key K
- Proves that M was encrypted (confidentiality and integrity) by someone who knew K (authenticity)

Why are we sending M?

Message Authenticity

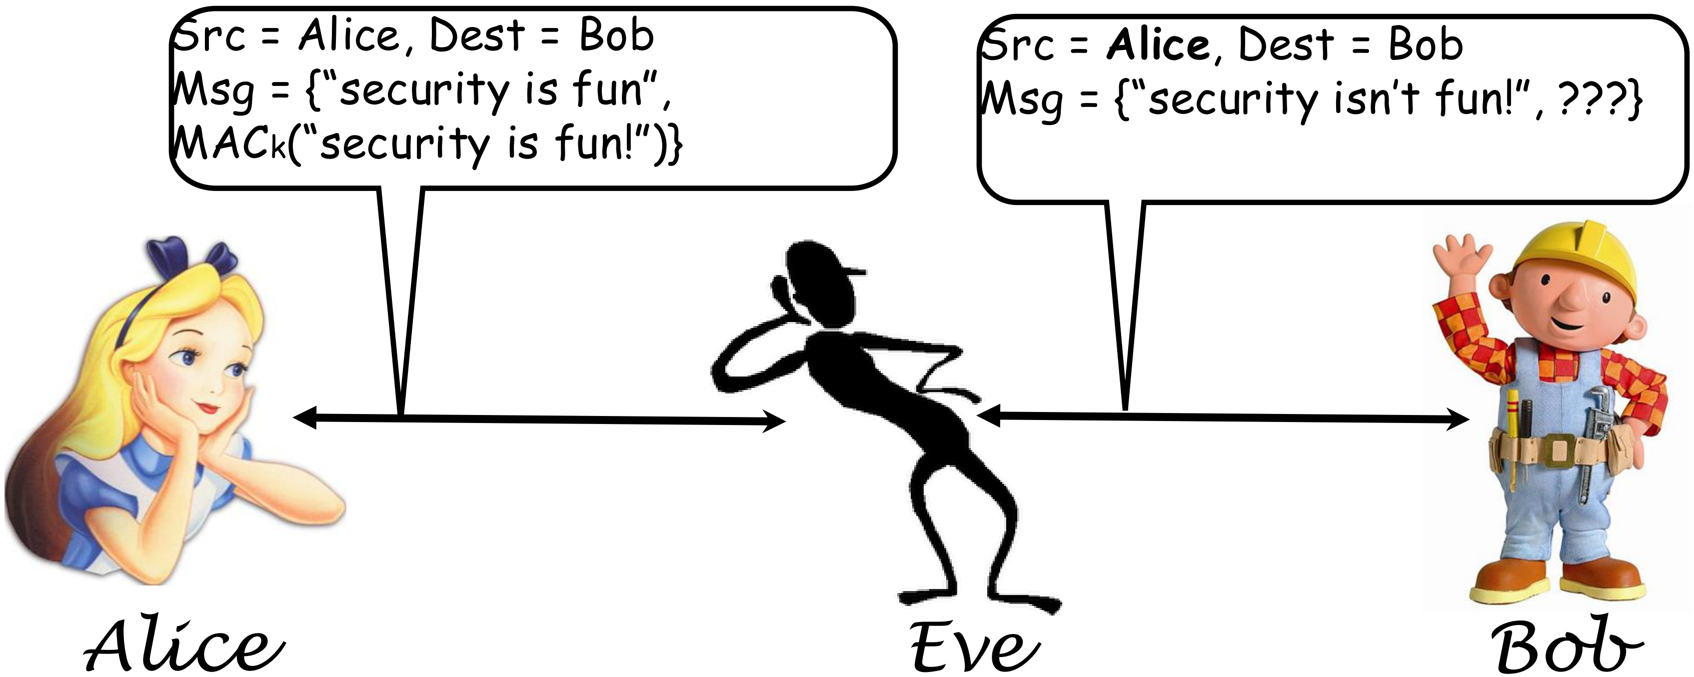

Without knowledge of *k*, Eve can't compute a valid MAC for her forged message!

Encryption and Message Authenticity

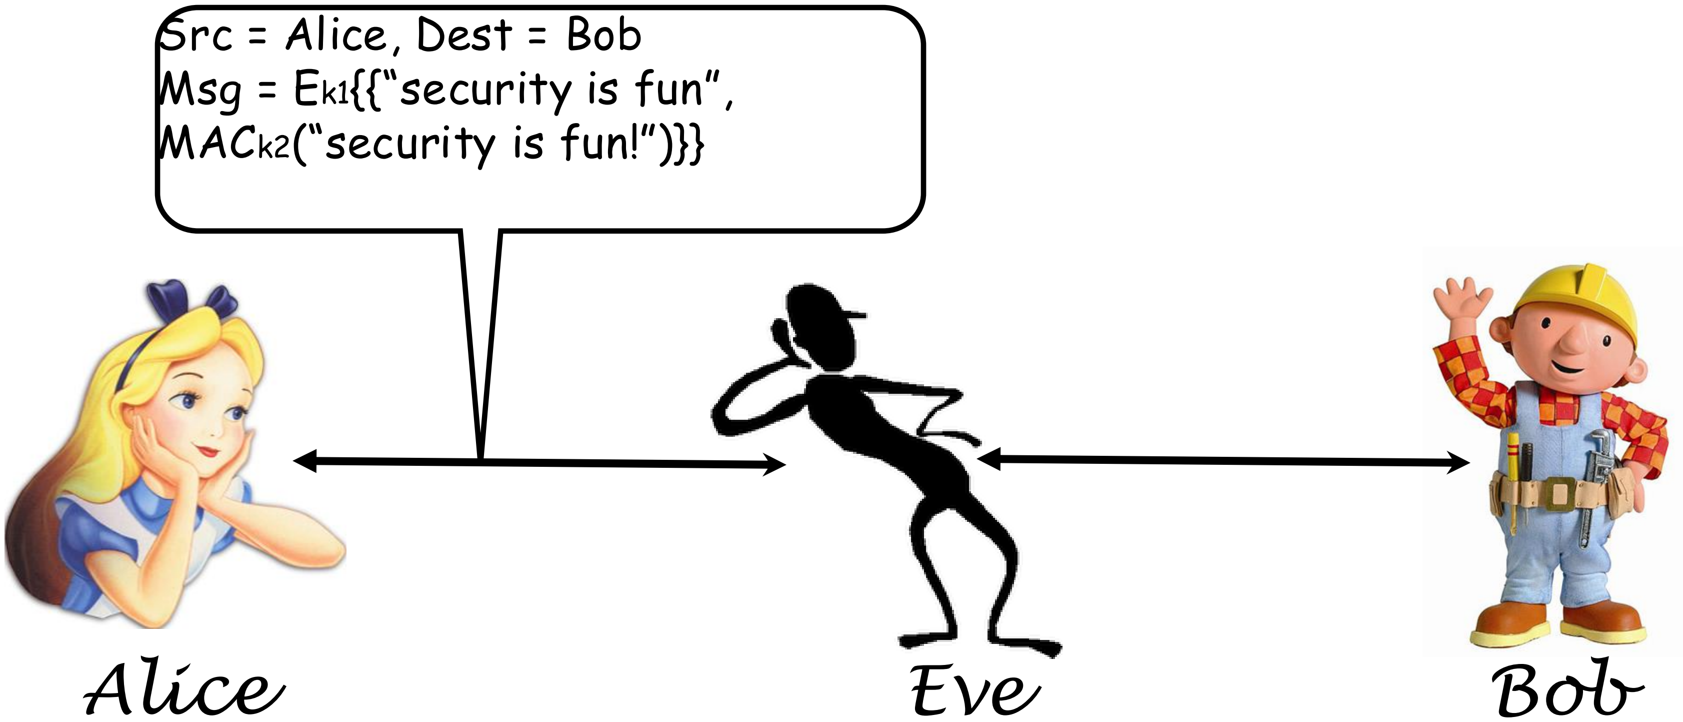

Without knowing *k1*,
Eve can't read Alice's message.

Without knowing *k2*, Eve can't compute a valid MAC for her forged message!

Cryptographic Hash Functions

- **Hash function** `h`: deterministic one-way function that takes as input an arbitrary message `M` (sometimes called a *preimage*) and returns as output `h(M)`, a small fixed length *hash* (sometimes called a *digest*)
- Hash functions should have the following two properties:
  - *compression*: reduces arbitrary length string to fixed length hash
  - *ease of computation*: given message `M`, `h(M)` is easy to compute

Hash functions are usually fairly inexpensive

(i.e., compared with public key cryptography)

```
~ openssl speed sha256
Doing sha256 ops for 3s on 16 size blocks: 33515432 sha256 ops in 3.00s
Doing sha256 ops for 3s on 64 size blocks: 29299431 sha256 ops in 2.99s
Doing sha256 ops for 3s on 256 size blocks: 19059503 sha256 ops in 3.00s
Doing sha256 ops for 3s on 1024 size blocks: 7433662 sha256 ops in 3.00s
Doing sha256 ops for 3s on 8192 size blocks: 1104810 sha256 ops in 3.00s
Doing sha256 ops for 3s on 16384 size blocks: 559814 sha256 ops in 2.99s
version: 3.3.1
built on: Tue Jun  4 12:53:04 2024 UTC
options: bn(64,64)
compiler: clang -fPIC -arch arm64 -O3 -Wall -DL_ENDIAN -DOPENSSL_PIC -D_REENTRANT
CPUINFO: OPENSSL_armcap=0x987d
The 'numbers' are in 1000s of bytes per second processed.
type             16 bytes     64 bytes    256 bytes   1024 bytes   8192 bytes  163
sha256          178748.97k   627145.01k  1626410.92k  2537356.63k  3016867.84k  30

~
```

Why might hashes be useful?

- **Message authentication codes (MACs):**
  - e.g.: $\text{MAC}_K(M) = h(K|M)$
    (but don't do this, use HMAC instead)
- **Modification detection codes:**
  - detect modification of data
  - any change in data will cause change in hash

Prof. Pedantic proposes the following hash function, arguing that it offers both compression and ease of computation.

- h(M) = 0 if the number of 0s in M is divisible by 3
- h(M) = 1 otherwise

Why is this a lousy crypto hash function?

Cryptographic Hash Functions

- Properties of good cryptographic hash functions:
  - **preimage resistance:** given digest y, computationally infeasible to find preimage x' such that h(x')=y
    (also called “one-way property”)
  - **2nd-preimage resistance:** given preimage x, computationally infeasible to find preimage x' such that h(x)=h(x')
    (also called “weak collision resistance”)
  - **collision resistance:** computationally infeasible to find preimages i,j such that h(i)=h(j)
    (also called “strong collision resistance”)

Birthday Attack

- **Birthday Paradox:** chances that 2+ people share birthday in group of 23 is > 50%.
- General formulation
  - function f() whose output is uniformly distributed over H possible outputs
  - Number of experiments Q(H) until we find a collision is approximately:

$$Q(H) \approx \sqrt{\frac{\pi}{2} H}$$

  - E.g.,

$$Q(365) \approx \sqrt{\frac{\pi}{2} 365} = 23.94$$

- Why is this relevant to hash sizes?

See: https://betterexplained.com/articles/understanding-the-birthday-paradox/

Practical Implications

- Choosing two messages that have the same hash h(x) = h(x') is more practical than you might think.
- Example attack: secretary is asked to write a "bad" letter, but wants to replace with a "good" letter.
  - Boss signs the letter after reading
- Find collision between 2^37 'good' vs 2^37 'bad' letters

Dear Anthony,

{This letter is / I am writing} to introduce {you to / to you} {Mr. / --} Alfred {P. / --} Barton, the {new / newly appointed} {chief / senior} jewellery buyer for {our / the} Northern {European / Europe} {area / division}. He {will take / has taken} over {the / --} responsibility for {all / the whole of} our interests in {watches and jewellery / jewellery and watches} in the {area / region}. Please {afford / give} him {every / all the} help he {may need / needs} to {seek out / find} the most {modern / up to date} lines for the {top / high} end of the market. He is {empowered / authorized} to receive on our behalf {samples / specimens} of the {latest / newest} {watch and jewellery / jewellery and watch} products, {up / subject} to a {limit / maximum} of ten thousand dollars. He will {carry / hold} a signed copy of this {letter / document} as proof of identity. An order with his signature, which is {appended / attached} {authorizes / allows} you to charge the cost to this company at the {above / head office} address. We {fully / --} expect that our {level / volume} of orders will increase in the {following / next} year and {trust / hope} that the new appointment will {be / prove} {advantageous / an advantage} to both our companies.

Figure 11.7 A Letter in 2^{37} Variations

(from Stallings, Crypto and Net Security)

Some common cryptographic hash functions

- MD5 (128-bit digest) [don’t use this]
- SHA-1 (160-bit digest) [stop using this*]
- SHA-256 (256-bit digest)
- SHA-512 (512-bit digest)
- SHA-3 [recent competition winner]

General Structure of Hash

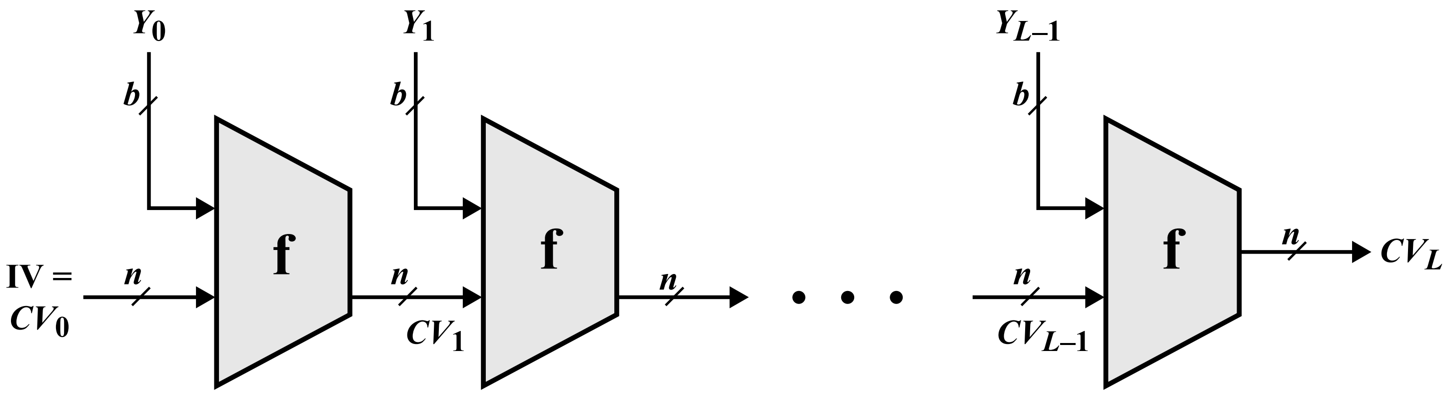

IV	=	Initial value	L	=	number of input blocks
CV_i	=	chaining variable	n	=	length of hash code
Y_i	=	*i*th input block	b	=	length of input block
f	=	compression algorithm			

(from Stallings, Crypto and Net Security)

Message Extension Attack

- Why is $MAC_k(M) = H(k|M)$ bad?
- How can Eve append M' to M?
  - Goal: compute H(k|M|M') without knowing k
- Solution: Use H(k|M) as IV for next f iteration in H()

A Better MAC

- Objectives
  - Use available hash functions without modification
  - Easily replace embedded hash function as more secure ones are found
  - Preserve original performance of hash function
  - Easy to use