



STORAGE DEVELOPER CONFERENCE

SNIA ■ SANTA CLARA, 2016

Fun with Linearity:

How encryption and erasure codes are intimately related


Jason Resch

IBM

Presentation Overview

- ❑ Linear Functions
- ❑ Combining Linear Functions
- ❑ Linearity of Erasure Codes
- ❑ Erasure Codes and Encryption
- ❑ Erasure Codes and Integrity Checks
- ❑ Exploiting Linearity for New Applications

Linear Functions



Linear Functions Defined

- A function ***F*** is “*linear*” if it satisfies the following:
 - $\mathbf{F(x) + F(y) = F(x + y)}$
- In short, a function is linear if an operation applied to the inputs yields the same result as performing that operation on the outputs.

Examples of Linear Functions

- ❑ Consider the function **MulFive(X)**
 - ❑ It multiplies any input by 5
 - ❑ Is this a linear function?
- ❑ Let's consider two inputs, $X=3$ and $Y=7$:
 - ❑ $\text{MulFive}(3) = 15$
 - ❑ $\text{MulFive}(7) = 35$
 - ❑ $\text{MulFive}(3 + 7) = \text{MulFive}(3) + \text{MulFive}(7)$ ✓

CRC as a Linear Function

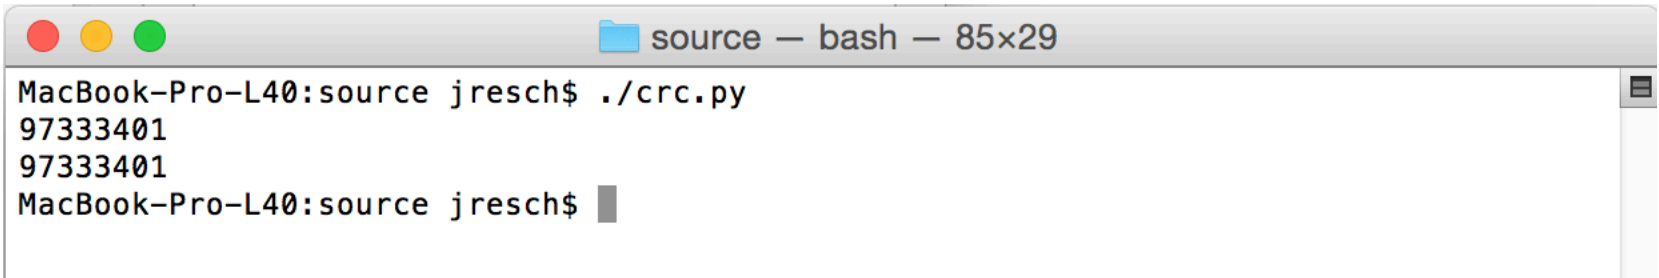
- ❑ Many functions have this property, among them:
 - ❑ The Cyclic Redundancy Check (CRC)
 - ❑ CRC is commonly used for integrity checks
- ❑ CRC is a “polynomial division” in a finite field
 - ❑ Addition in this field is “ \oplus ” rather than “+”

Consequences of CRC's Linearity

- ❑ Given the definition of linearity, it follows that:
 - ❑ $\text{CRC}(x) \oplus \text{CRC}(y) = \text{CRC}(x \oplus y)$
- ❑ In actual implementations of CRC, there is a caveat, we must also add $\text{CRC}(\text{zeros})$
 - ❑ Where 'zeros' is a binary string consisting of all zeros and is equal in length to 'x' and 'y'
 - ❑ $\text{CRC}(x) \oplus \text{CRC}(y) \oplus \text{CRC}(\text{zeros}) = \text{CRC}(x \oplus y \oplus \text{zeros})$

Example of zlib CRC-32's linearity

```
1  |#!/usr/bin/python
2
3  |from zlib import crc32
4
5  |▼ def xor(x, y):
6  |   |return ''.join(chr(ord(a) ^ ord(b)) for a,b in zip(x,y))
7
8  |▼ def calc_combined_crc(x, y):
9  |   |zeros = chr(0) * len(x)
10 |   |return crc32(x) ^ crc32(y) ^ crc32(zeros)
11
12 |print calc_combined_crc('foo', 'bar')
13
14 |foobar = xor('foo', 'bar')
15 |print crc32(foobar)
```



A terminal window titled "source — bash — 85x29" on a MacBook-Pro-L40. The window shows the execution of a script named "crc.py". The prompt is "MacBook-Pro-L40:source jresch\$". The command entered is "./crc.py". The output of the script is "97333401" on two separate lines. The prompt then returns to "MacBook-Pro-L40:source jresch\$".

```
MacBook-Pro-L40:source jresch$ ./crc.py
97333401
97333401
MacBook-Pro-L40:source jresch$
```

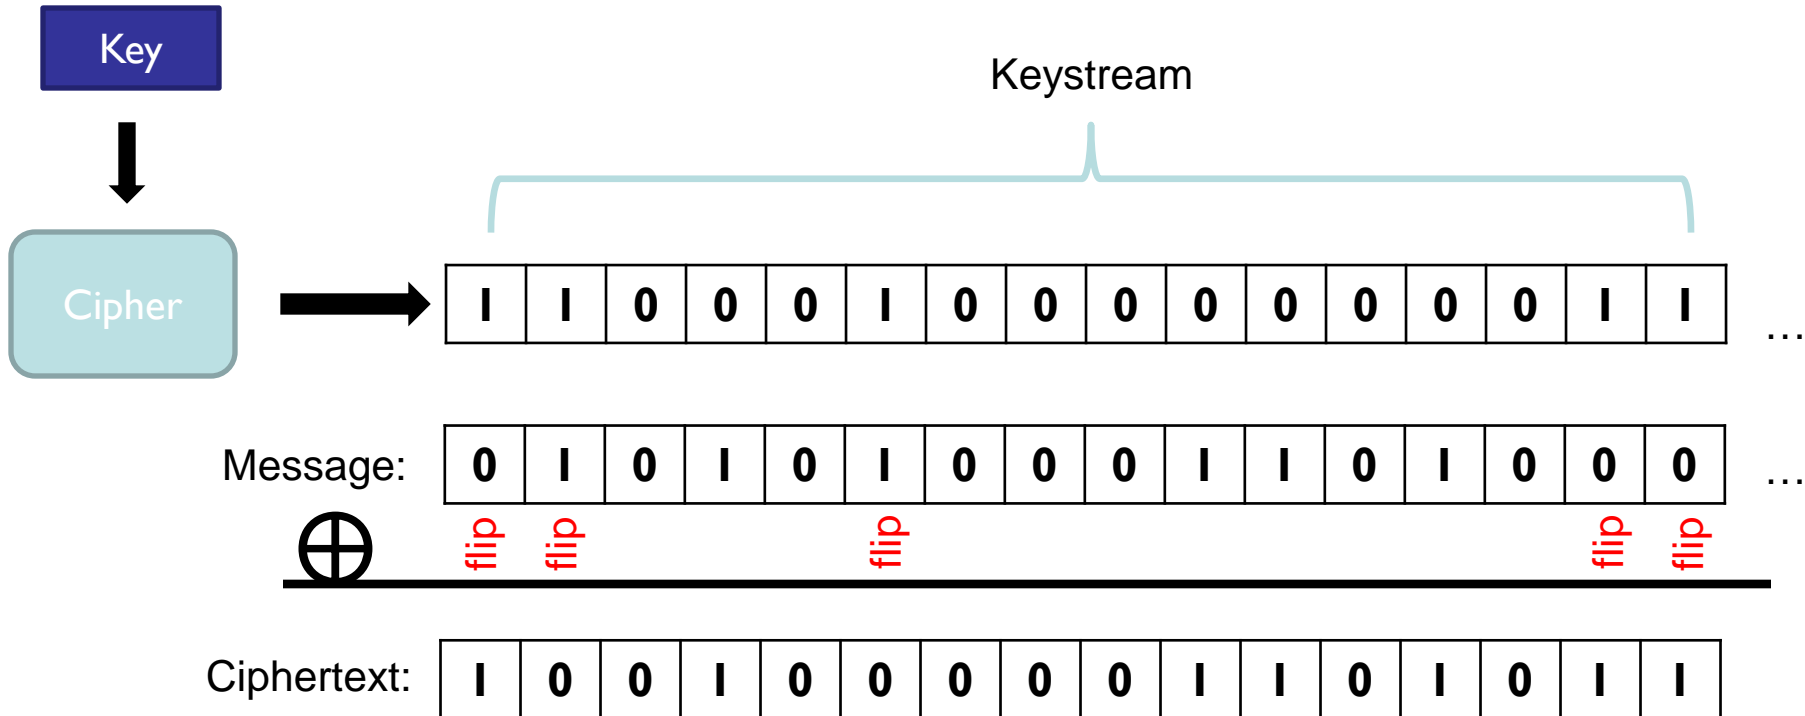

Combining Linear Functions



Encryption Functions

- ❑ Encryption functions, in their most basic form, are simply random number generators
 - ❑ Key acts as a “seed” to generate an arbitrarily long set of random output—the “keystream”
 - ❑ Keystream is XORed with Plaintext to encrypt


Encryption Visualized



Encryption and CRC

- ❑ The XOR operation used in encryption is the same as addition in the finite field of CRC...
 - ❑ $\text{Ciphertext} = \text{Plaintext} \oplus \text{Keystream}$
- ❑ Since the Ciphertext is equal to the Plaintext added to the Keystream it follows that:
 - ❑ $\text{CRC}(\text{Ciphertext}) = \text{CRC}(\text{Plaintext}) \oplus \text{CRC}(\text{Keystream})$

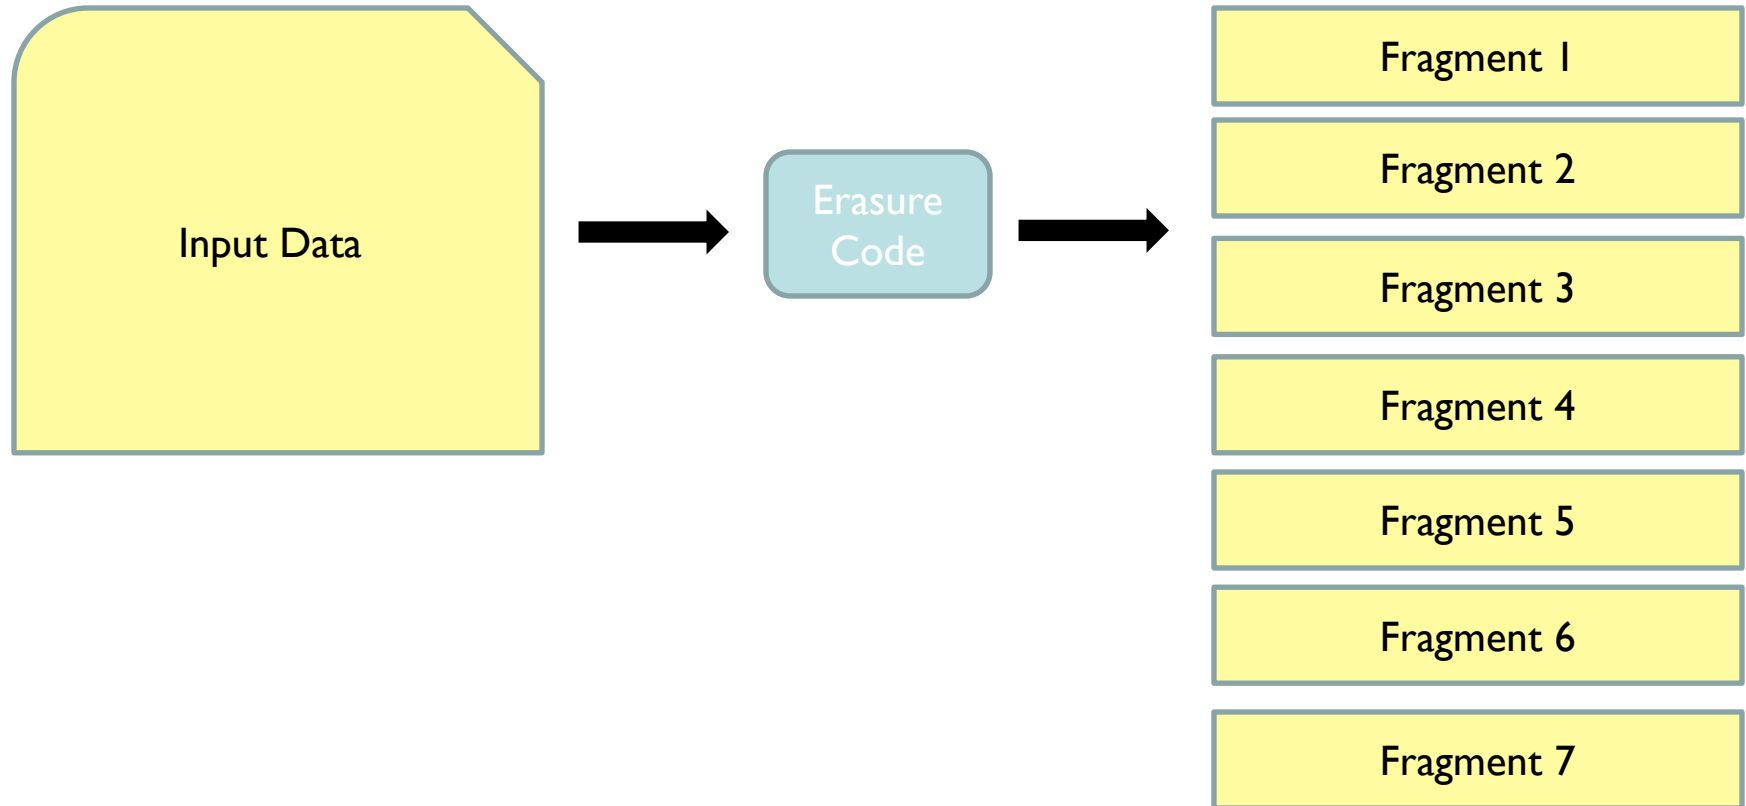
Linearity of Erasure Codes



Erasure Codes

- ❑ Erasure Codes encode T inputs into N outputs
- ❑ Can recover original input from any T of outputs
- ❑ Examples:
 - ❑ RAID 5, RAID 6, Reed-Solomon, Rabin's Information Dispersal, Shamir Secret Sharing

Visualizing Erasure Codes



Applications of Erasure Codes

- ❑ Erasure Codes are often used to achieve durability and availability in storage systems
 - ❑ Files split into T fragments, redundancy expands these into N fragments
 - ❑ Each fragment stored to different node/drive
 - ❑ Reliable: Data can survive $(N - T)$ faults
 - ❑ Efficient: Overhead equal to (N / T)

Erasure Codes are Linear

- ❑ Like CRC and Encryption, Erasure Codes operate within a finite field
- ❑ Encoding and decoding are implemented via addition and multiplication in a field
- ❑ Erasure Codes are linear:
 - ❑ $EC(x) \oplus EC(y) = EC(x \oplus y)$

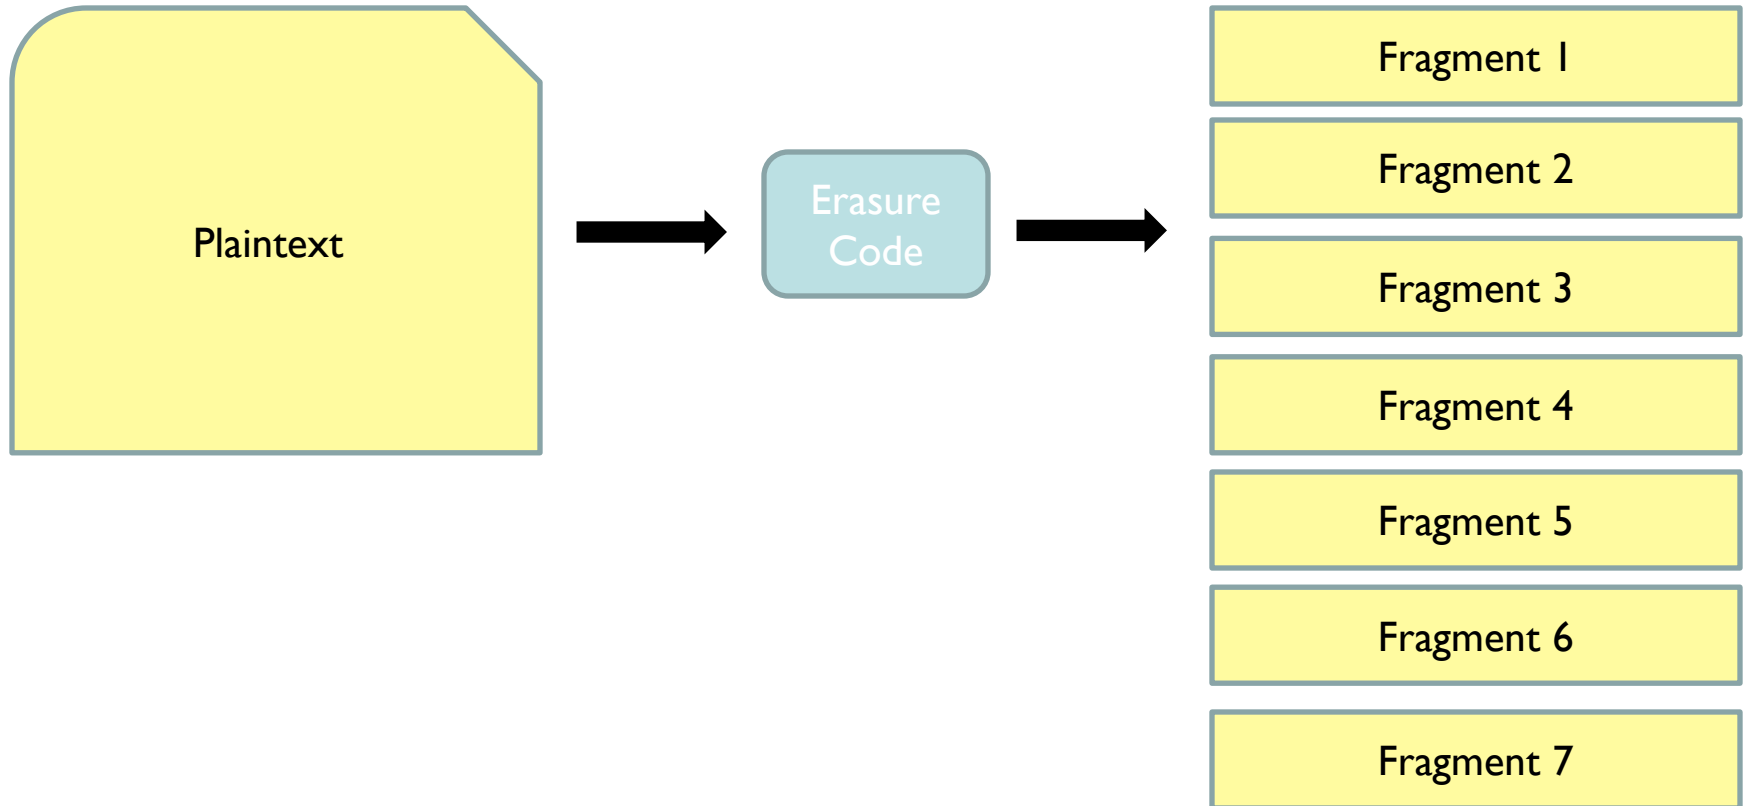
Erasure Codes and Encryption



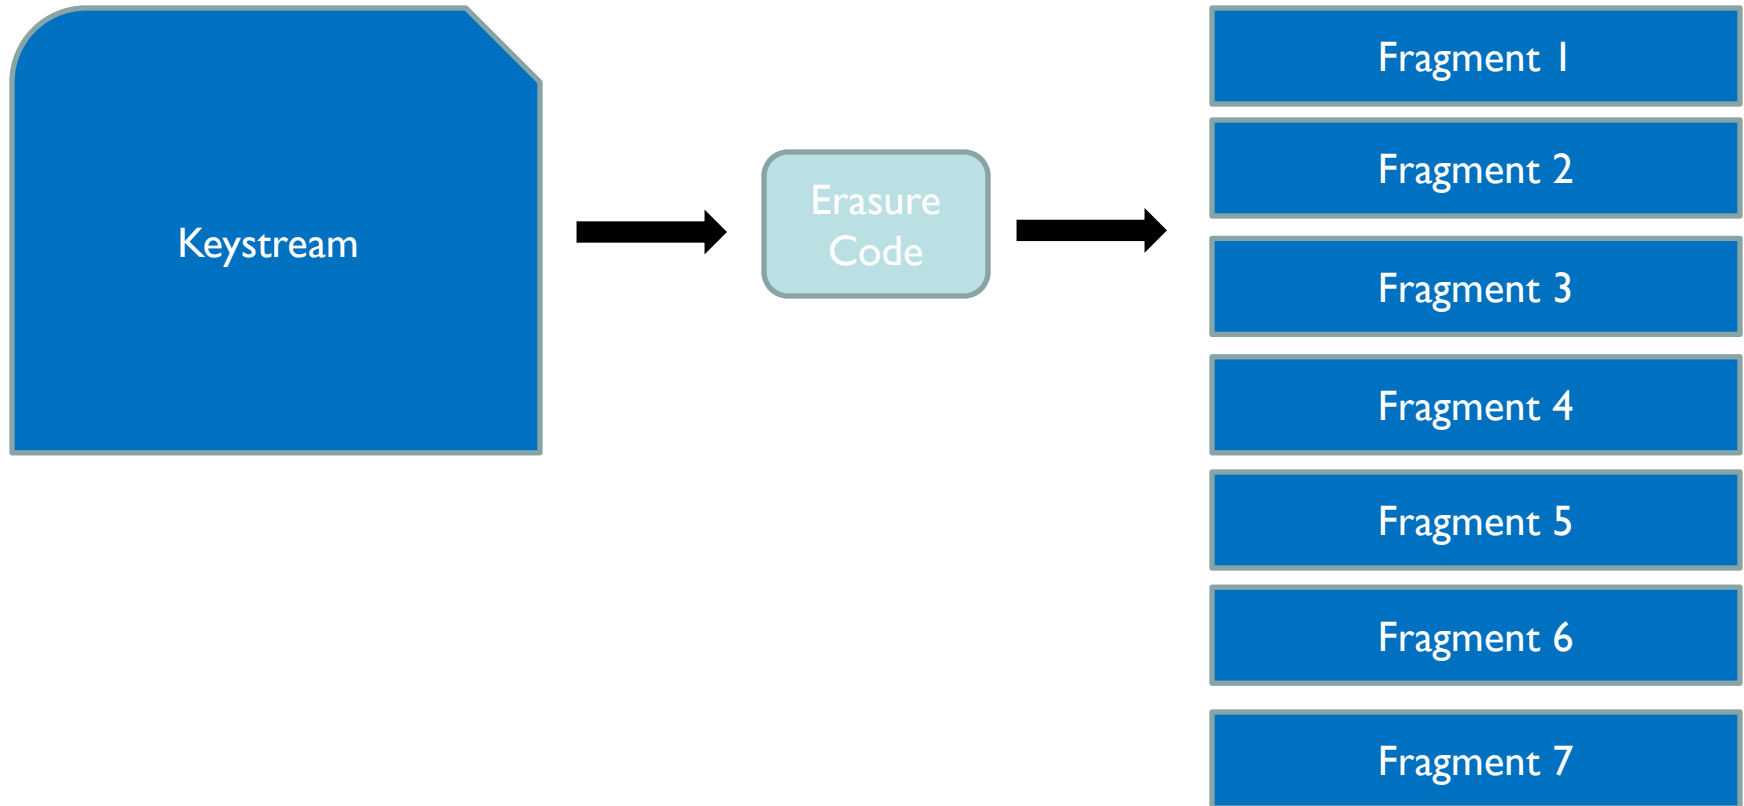
Erasure Codes and Encryption

- Since encryption is addition, and because Erasure Codes are linear functions:
 - $EC(\text{Ciphertext}) = EC(\text{Plaintext}) \oplus EC(\text{Keystream})$
- If we erasure code encrypted data, we will get output fragments that will be identical to if we erasure code the plaintext, and add those fragments to the fragments resulting from erasure coding the keystream.


Erasure Coding Plaintext



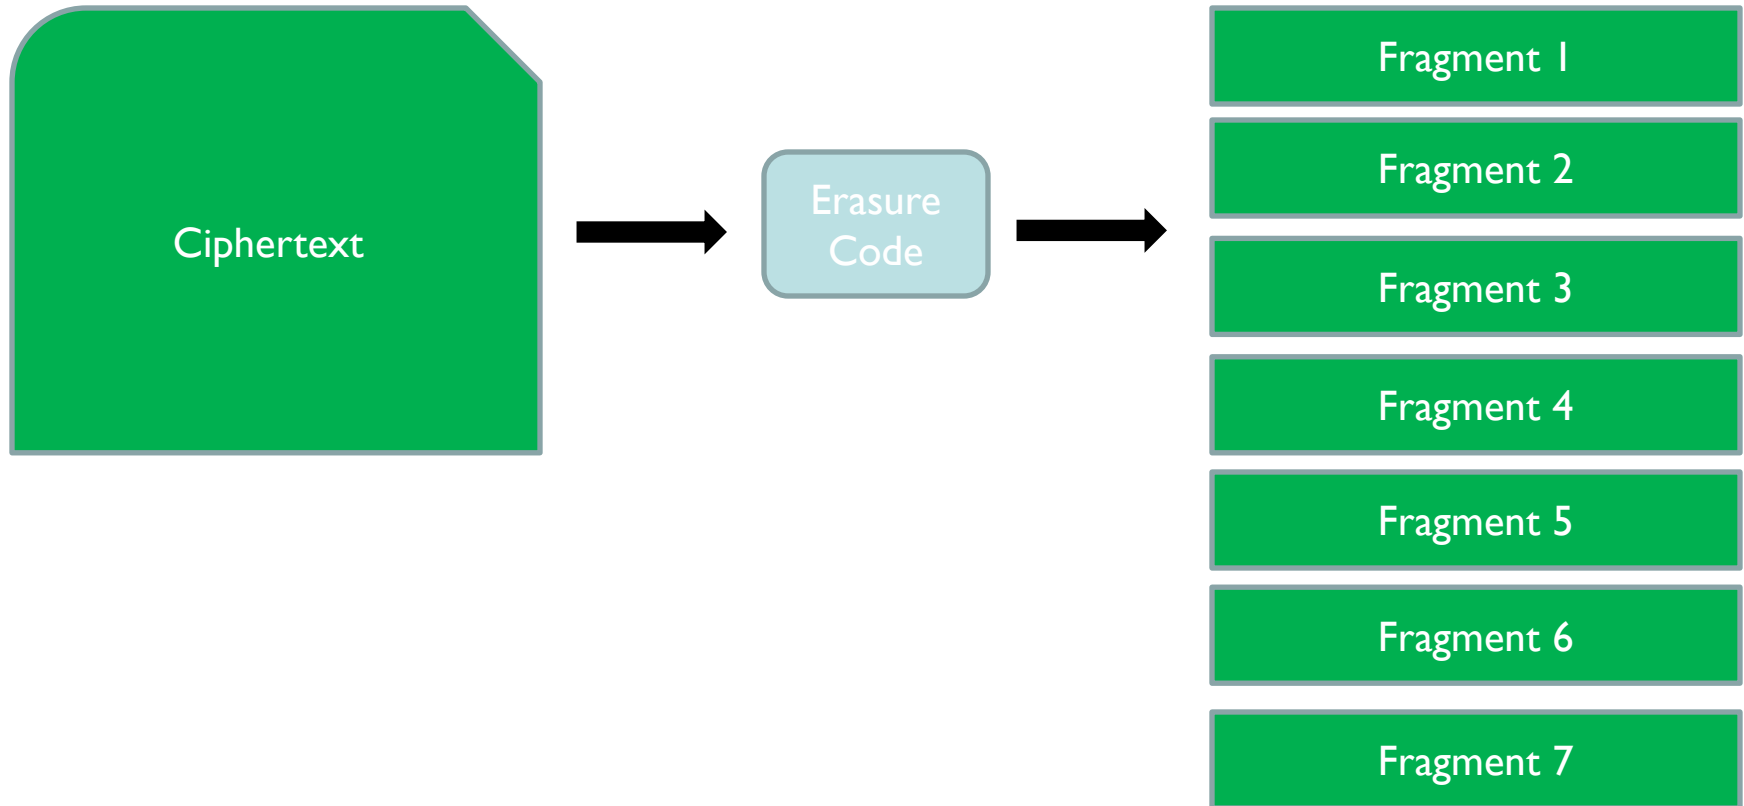
Erasure Coding Keystream



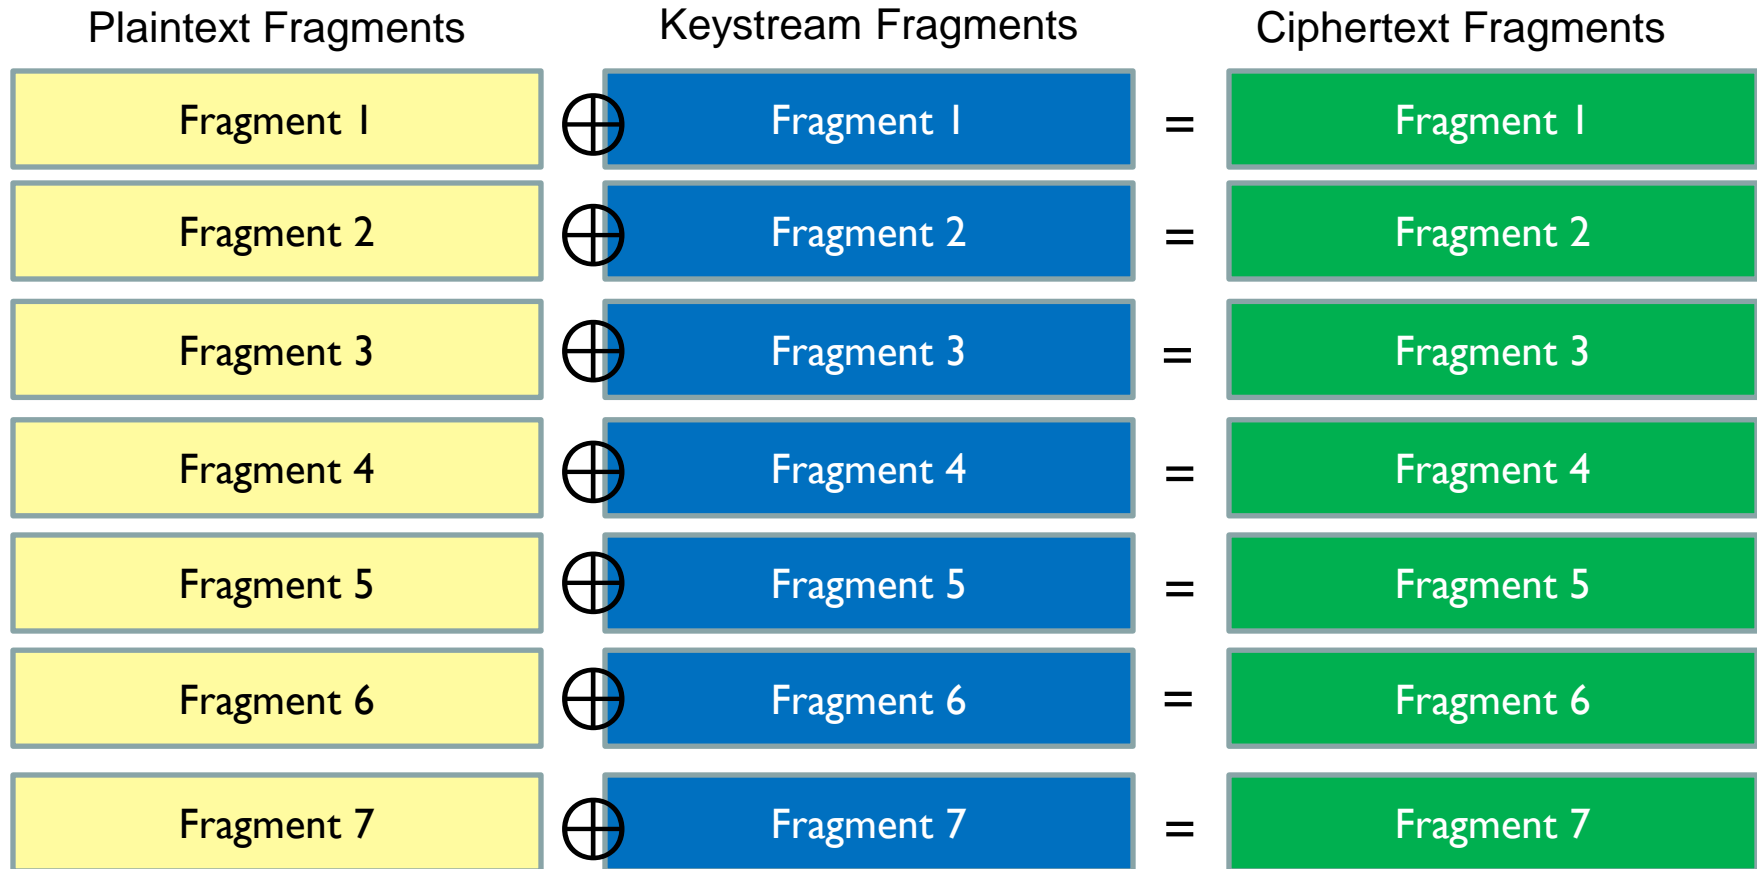
Relating the Erasure Code Inputs



Erasure Coding Ciphertext



Relating the Erasure Code Outputs



Using this property in practice

- ❑ In a distributed erasure coded system, we can:
 - ❑ Encrypt previously unencrypted fragments
 - ❑ Decrypt previously encrypted fragments
 - ❑ Rekey encrypted fragments
- ❑ But we can do so without any network transfer!
- ❑ We only need to send keys to storage nodes

Illustrating Rekeying of Fragments

What's currently stored

$$\boxed{\text{EC}(\text{Plaintext} \oplus \text{Keystream}_1)} = \boxed{\text{EC}(\text{Plaintext})} \oplus \boxed{\text{EC}(\text{Keystream}_1)}$$

What we want stored

$$\boxed{\text{EC}(\text{Plaintext} \oplus \text{Keystream}_2)} = \boxed{\text{EC}(\text{Plaintext})} \oplus \boxed{\text{EC}(\text{Keystream}_2)}$$

What we want stored

$$\boxed{\text{EC}(\text{Plaintext} \oplus \text{Keystream}_2)} = \boxed{\text{EC}(\text{Plaintext} \oplus \text{Keystream}_1)} \oplus \boxed{\text{EC}(\text{Keystream}_1 \oplus \text{Keystream}_2)}$$

Erasure Codes and Integrity Checks



Erasure Codes and CRC

- ❑ We know Erasure Codes are linear
- ❑ We know Cyclic Redundancy Checks are linear
 - ❑ How can we combine them for practical uses?
- ❑ We will need to explore the internal workings of Erasure Codes to see how this is possible..

Erasure Codes: Encoding

$N \times T$ "Encoding Matrix"

T inputs

\times

$=$

N outputs

Erasure Codes: Encoding (Color Coded)

$N \times T$ "Encoding Matrix"


T inputs

\times

$=$

N outputs

Erasure Codes: Delete Irrelevant Rows



Erasure Codes: Getting the Inputs Back

$T \times T$ "Truncated Matrix"

T inputs


\times

$=$


T outputs



Erasure Codes: Matrix Cancellation




Erasure Codes: Recovered Inputs



Combining CRCs with Erasure Codes

- ❑ With this background on how Erasure Codes encode and decode information, lets see what else we can do..
 - ❑ We observed decoding as happening all at once, given all T of the outputs
 - ❑ But it doesn't have to be like this, the decode operation can be done in T distinct steps

Decomposing the Output Vector



Decoding with one output at a time

$(\text{Truncated Matrix})^{-1}$ output partially decoded result

The diagram shows a 5x5 matrix of gray squares, a multiplication symbol, a 5x1 column of white squares with the top square highlighted in purple, an equals sign, and a 5x1 column of light purple squares.

Decoding with one output at a time

$(\text{Truncated Matrix})^{-1}$ output partially decoded result

The diagram shows a 5x5 grid of gray squares representing the $(\text{Truncated Matrix})^{-1}$. To its right is a vertical column of five squares representing the 'output' vector; the second square from the top is highlighted in yellow. To the right of the output vector is an equals sign. To the right of the equals sign is a vertical column of five yellow squares representing the 'partially decoded result' vector.

Decoding with one output at a time

$(\text{Truncated Matrix})^{-1}$ output partially decoded result

The diagram shows a 5x5 grid of gray squares representing the $(\text{Truncated Matrix})^{-1}$. To its right is a multiplication symbol \times , followed by a 5x1 column of white squares representing the 'output' vector, where the third square from the top is highlighted in red. To the right of the output vector is an equals sign $=$, followed by a 5x1 column of light red squares representing the 'partially decoded result' vector.

Decoding with one output at a time

$(\text{Truncated Matrix})^{-1}$ output partially decoded result


The diagram shows a 5x5 grid of gray squares representing the inverse of a truncated matrix. This is multiplied (indicated by a 'x' symbol) by a 5x1 column of white squares representing the current output. The 4th square in this column is highlighted in blue. The result is a 5x1 column of blue squares, representing the partially decoded result. An equals sign '=' is placed between the multiplication and the result.

Decoding with one output at a time

$(\text{Truncated Matrix})^{-1}$ output partially decoded result

The diagram shows a 5x5 grid of gray squares representing the inverse of a truncated matrix. This is multiplied (indicated by a 'x' symbol) by a 5x1 column of white squares representing the output, where the bottom square is highlighted in orange. The result is a 5x1 column of orange squares representing the partially decoded result, indicated by an equals sign.

Putting them all together





9/7/2016

Observation

- Just as the ciphertext is the sum of a plaintext and a keystream, the original input is a sum of the ***T*** partially decoded results
 - This means the sum of the CRCs of the partially decoded results is the data's CRC!

Computing CRCs of Data from fragments



9/7/2016

Exploiting Linearity for New Applications



Secure Rebuilding

- ❑ Linearity also has applications for rebuilding:
 - ❑ It enables rebuilding of lost fragments, without having to decode the data first
 - ❑ Saves significant bandwidth in some topologies
 - ❑ It can be further extended to allow rebuilding without exposing any other fragment
 - ❑ Sounds paradoxical and impossible, but its not!
 - ❑ Very useful in cases of secret sharing

Conclusions

- ❑ Linear functions leave much room for innovation
- ❑ So far it has provided the ability to:
 - ❑ Encrypt, decrypt and rekey erasure coded data without any network transfer
 - ❑ Verify integrity of erasure coded data without having to download or retrieve it
 - ❑ Rebuild fragments with minimal network transfer and without having to decode data

Questions



Secure and Efficient Rebuilding:

<http://dimacs.rutgers.edu/Workshops/SecureNetworking/Slides/resch.pptx>