

TLS Crypto Seminar : 28 February 2019

Partially Specified Channels

The TLS 1.3 Record Layer without elision

Christopher Patton,
Tom Shrimpton
CCS 2018

Presented by:
Vivek Arte

UC San Diego

Overview

- ❖ Motivation
- ❖ Partially specified channels
- ❖ The TLS 1.3 Record Layer

Motivation

- ❖ Protocols are often only **partially specified**.
- ❖ Standard:
 - ❖ collection of implementations with a shared core set of behaviors
- ❖ Challenge for provable security
 - ❖ **What is relevant to security?!**

Multiplexing in TLS 1.3

- ❖ The TLS 1.3 Record Layer handles streams for three distinct sub-protocols
 - ❖ handshake
 - ❖ alert
 - ❖ application-data
- ❖ Each sub-protocol has side-effects on the sender and receiver state, and thus could affect security

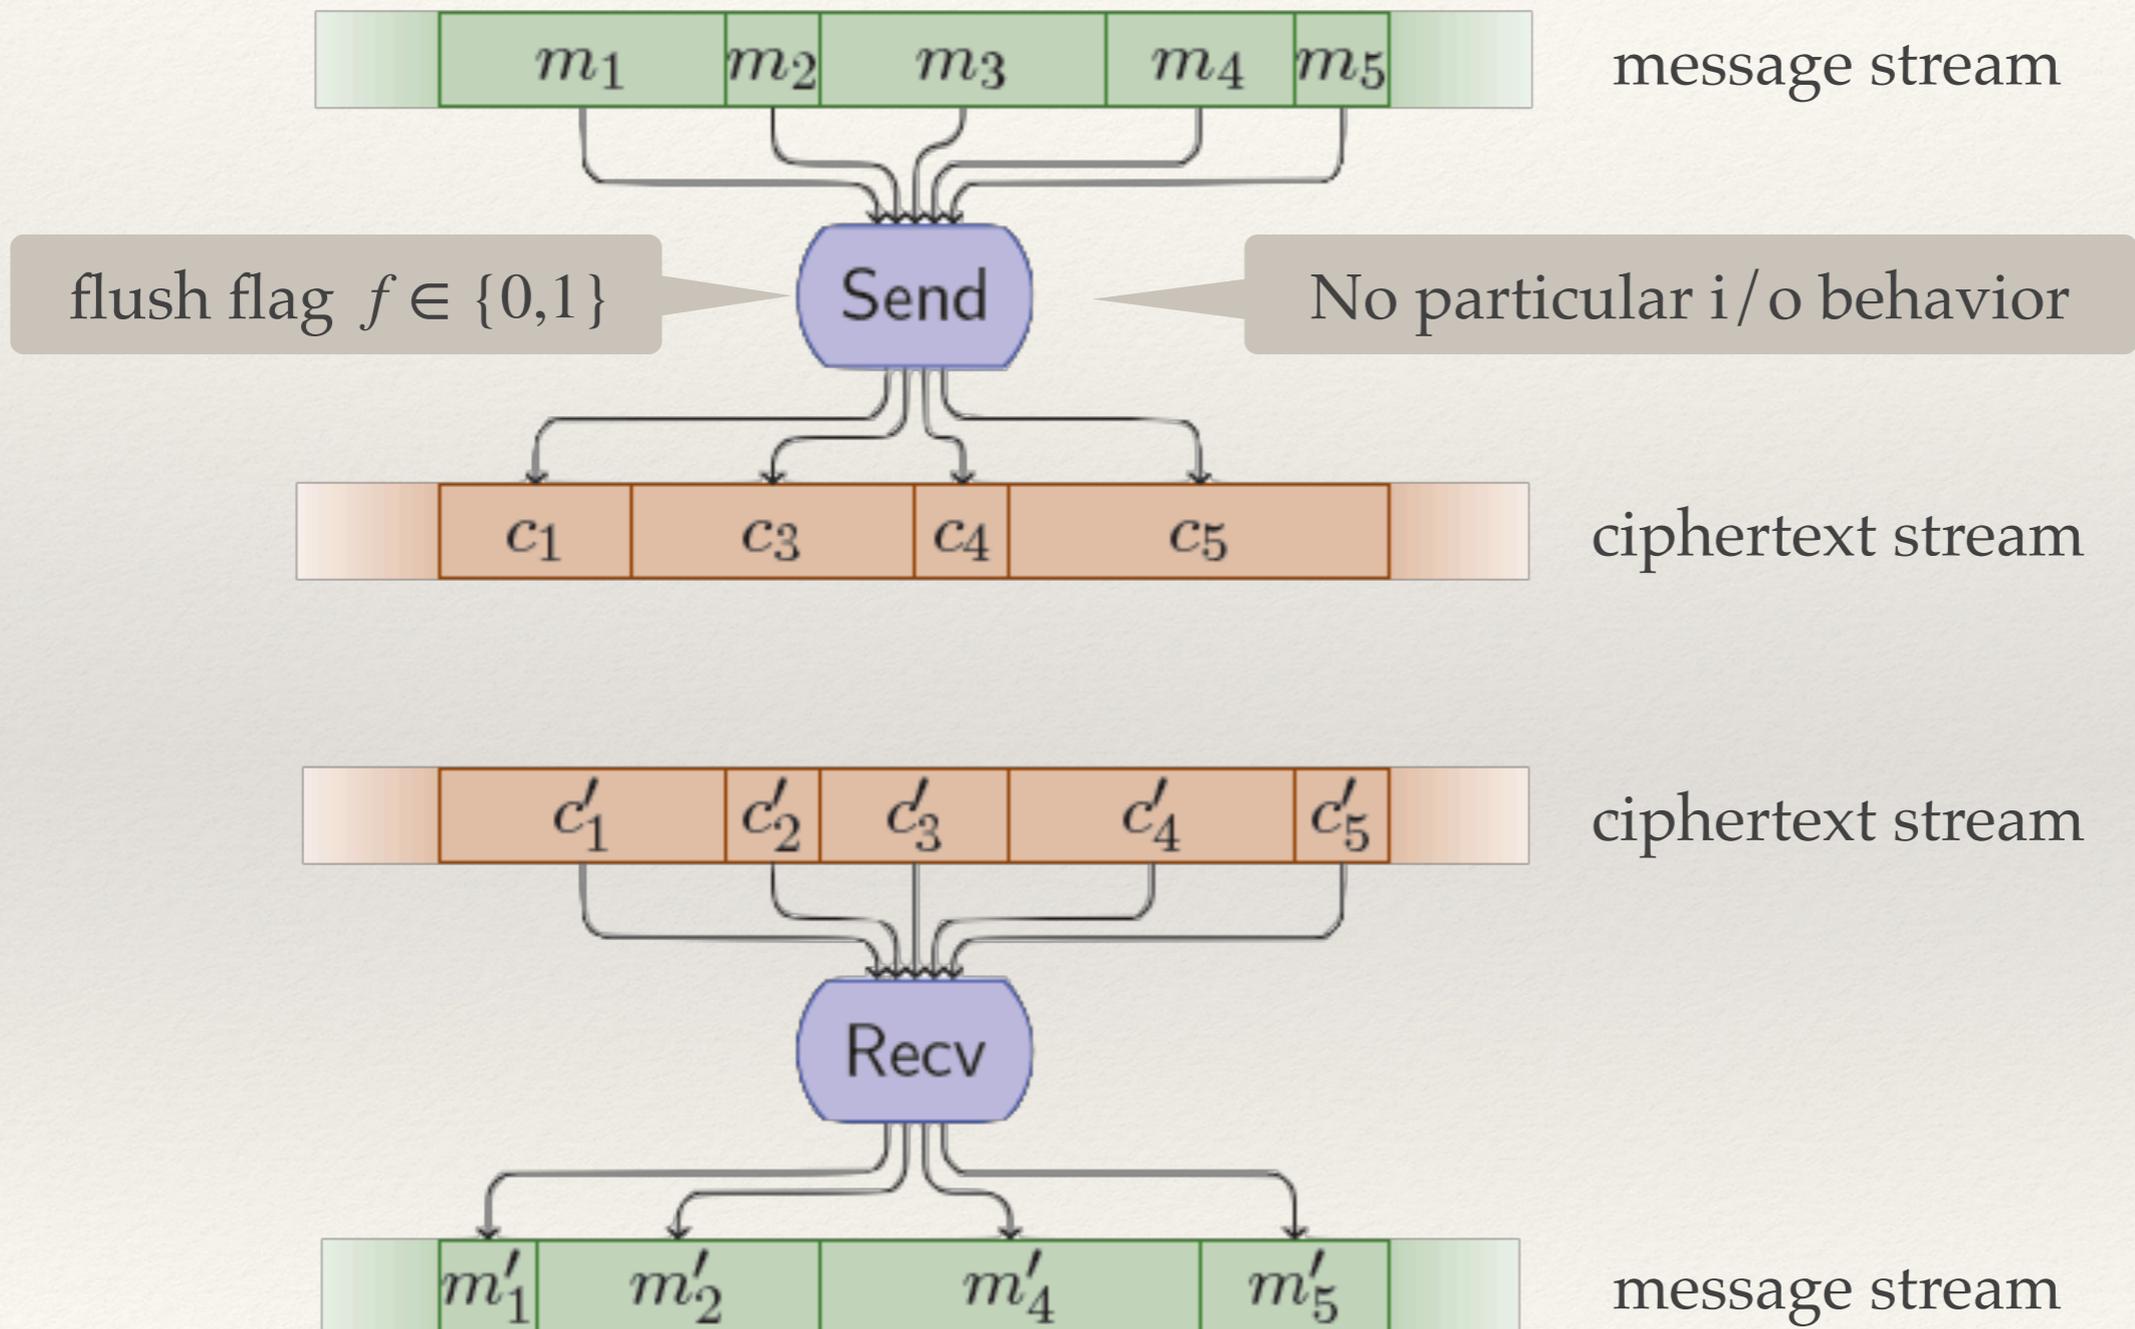
Stateful Authenticated Encryption [BKN02]

- ❖ Defined security notions of **confidentiality** and **integrity** for stateful symmetric encryption
- ❖ Accounts for **replay** and **out-of-order delivery** attacks
- ❖ Ciphertexts are **atomic!**

Stream-based Channels [FGMP15]

- ❖ TLS - provides a **streaming interface** for applications
 - ❖ Never promised to treat messages atomically!
- ❖ Fragmentation at the sender and receiver ends could differ
- ❖ [FGMP15] gives specifications and security notions for stream-based channels

Stream-based Channels [FGMP15]



Correctness [FGMP15]

$$(st_{S,0}, st_{R,0}) \leftarrow_{\$} \text{Init}(1^\lambda)$$

$$(st'_S, c) \leftarrow_{\$} \text{Send}(st_S, m, f)$$

$$(st'_R, m) \leftarrow_{\$} \text{Recv}(st_R, c)$$

❖ Correctness

- ❖ No matter how ciphertexts are fragmented at the sender side, and re-fragmented at the receiver side, the returned message stream is a prefix of the initial message stream

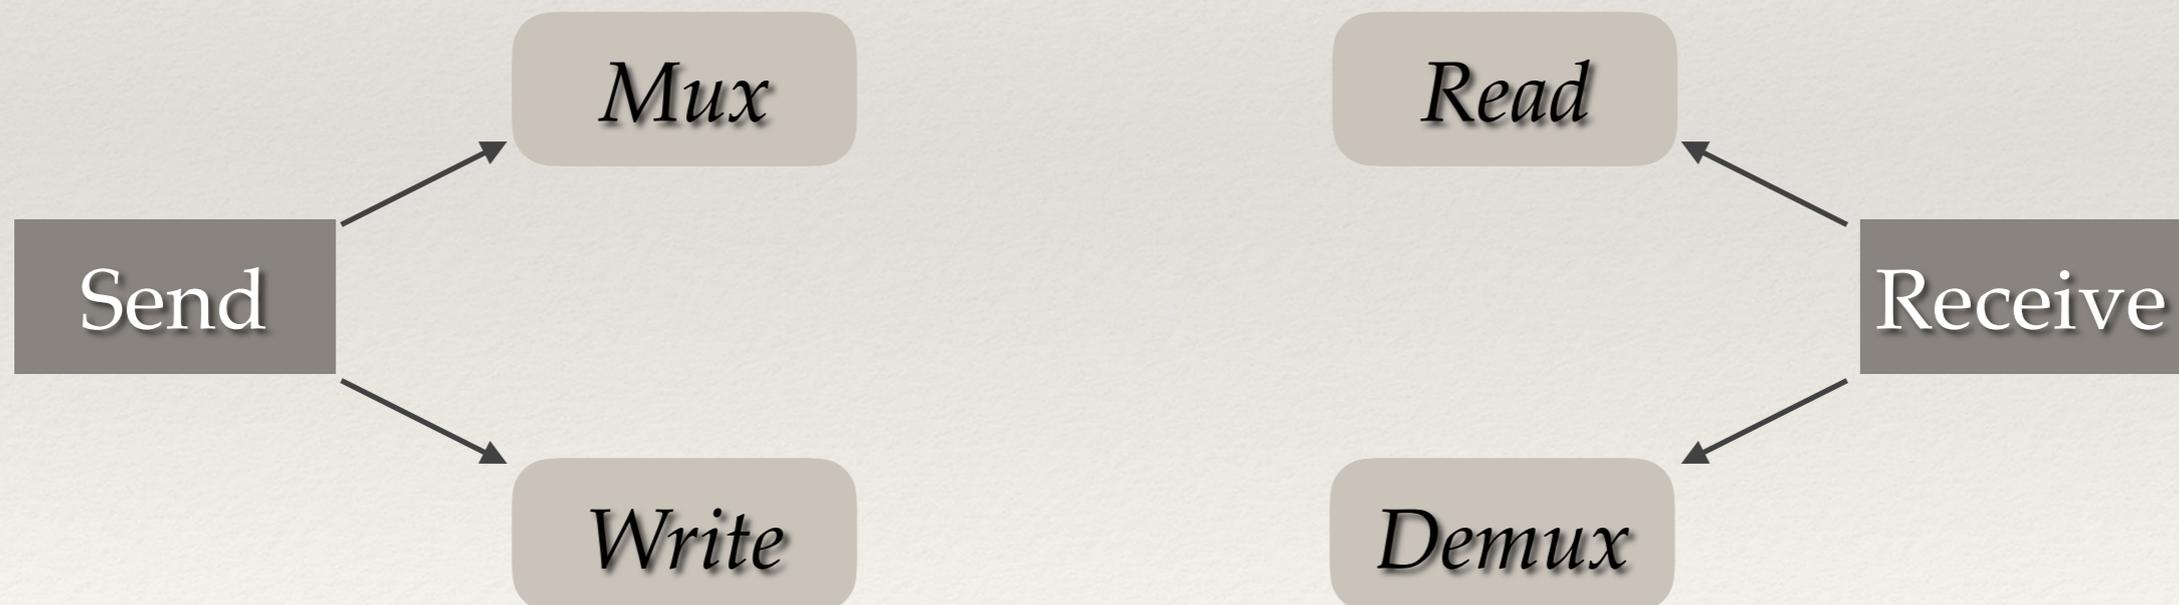
$$\|c[1, \dots, i]\| \preceq \|c'\| \preceq \|c\| \implies \|m[1, \dots, i]\| \preceq \|m'\| \preceq \|m\|.$$

$$i \in \{0\} \cup \{j : f_j = 1\}$$

everything up to last flush

Multiplexing [PS18]

- ❖ Streams are of the form $(M_1, sc_1), (M_2, sc_2), \dots$



Partially Specified Channels

things that are mandated and explicitly described

[RS09]

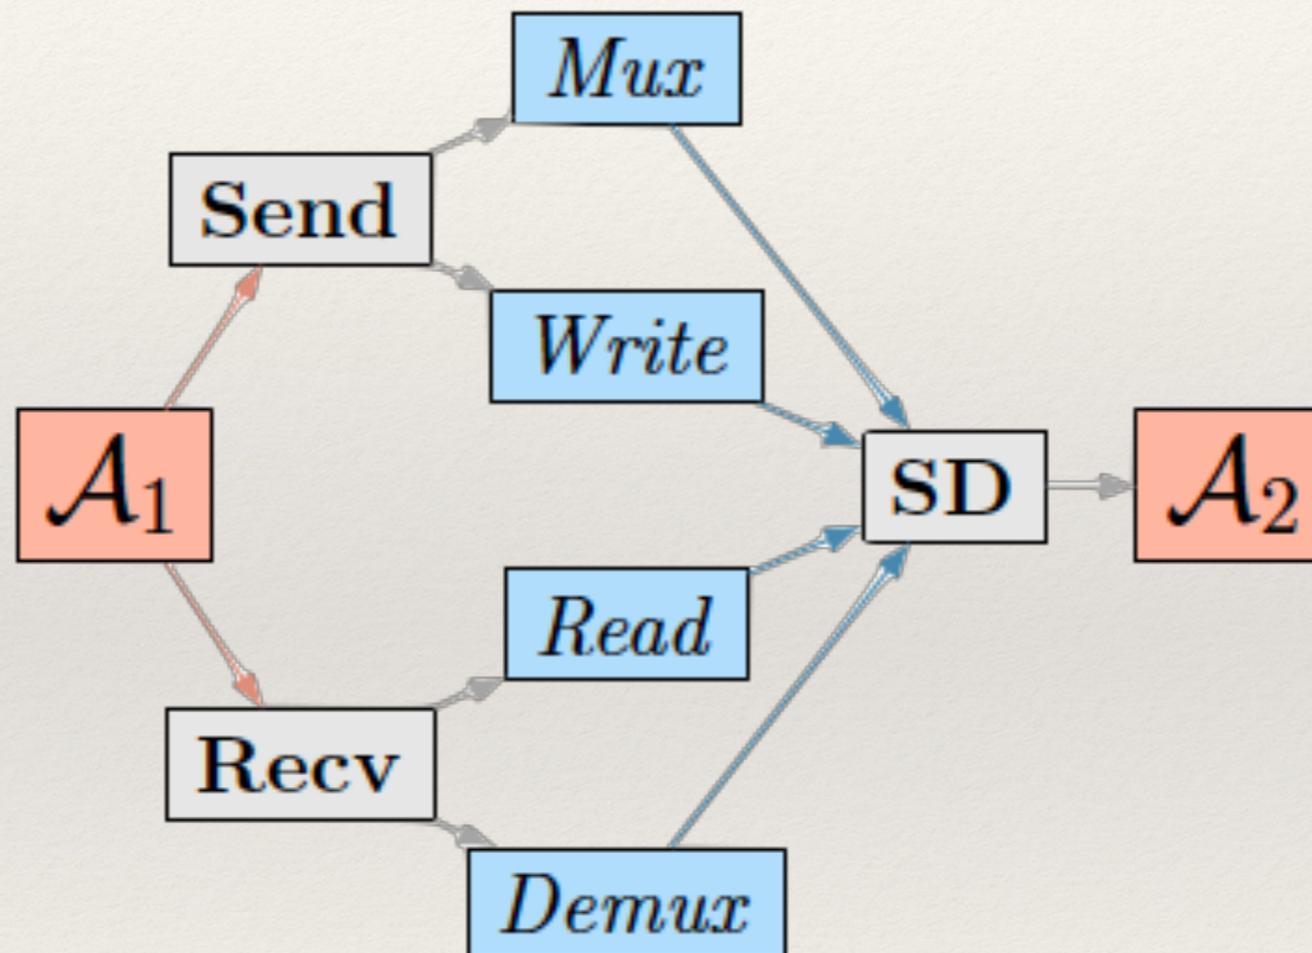
- ❖ Standard = partial specification + additional details

everything else

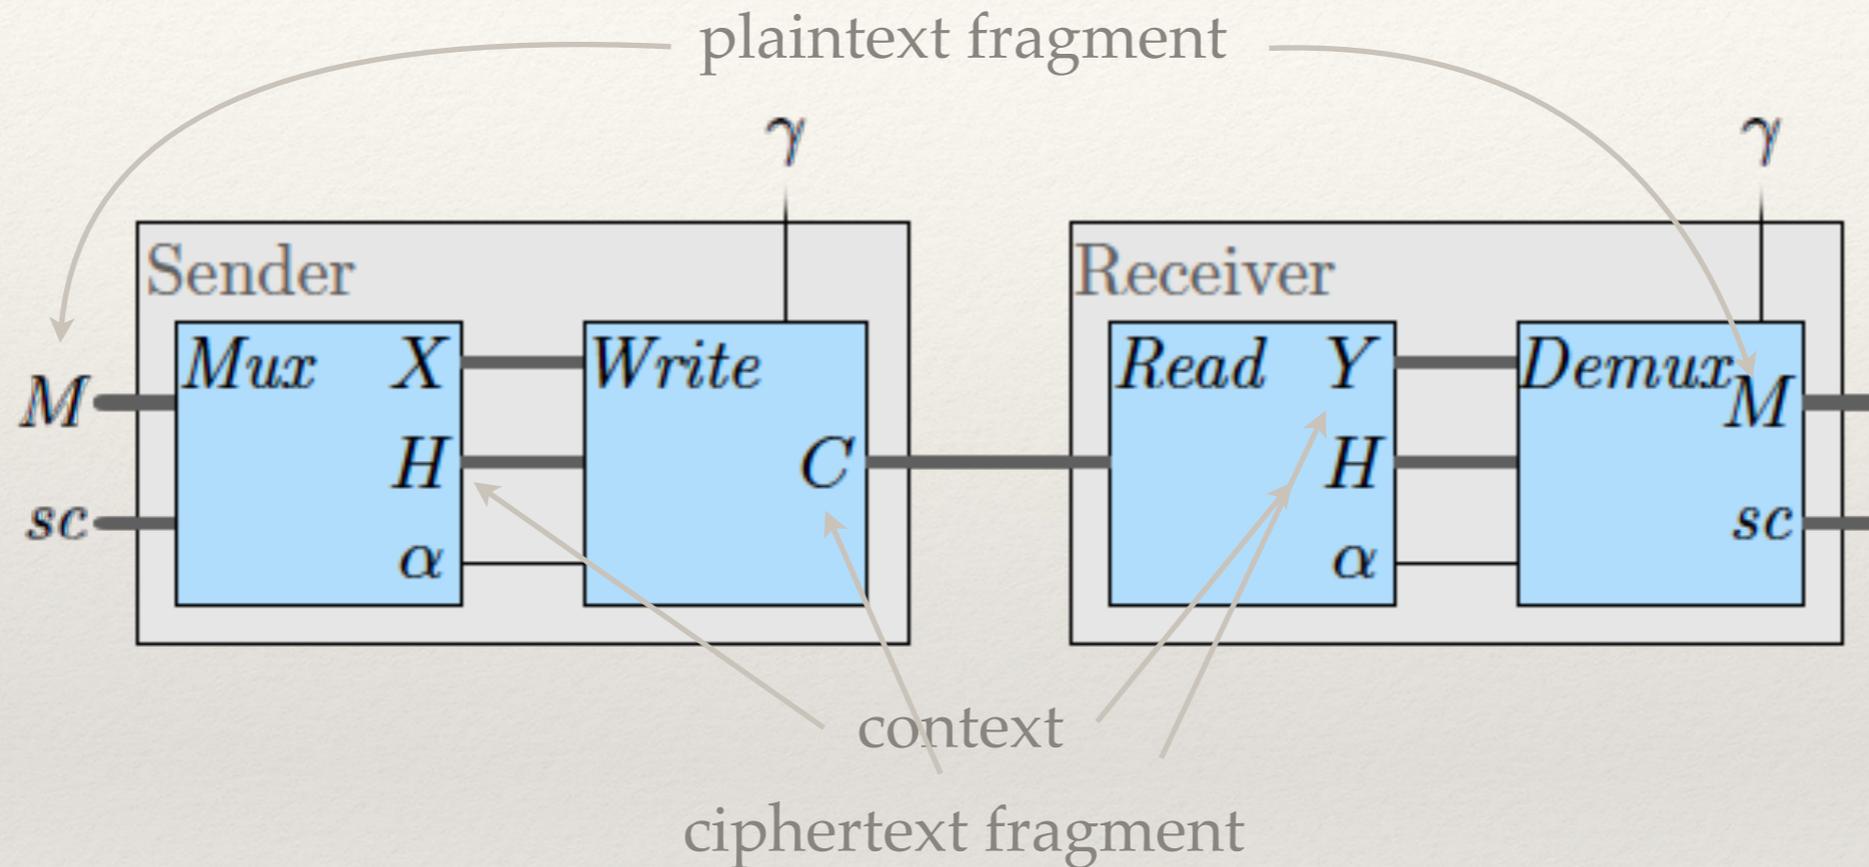
- ❖ *Mux, Write, Read, Demux* : fully specified
- ❖ The rest of the details are formalized as an oracle given to each algorithm

In security games, queries made to the oracle are serviced by the adversary

Execution Model



Syntax



$Init() \rightarrow (Mu, Wr, Re, De)$

$Mux^{\circ}(M, sc, Mu) \rightarrow (X, H, \alpha)$

$Read^{\circ}(C, Re) \rightarrow (Y, H, \alpha)$

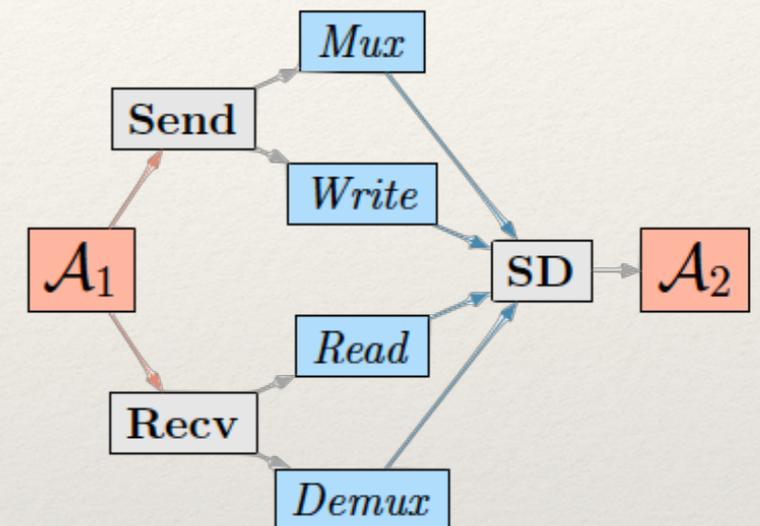
$Write^{\circ}(X, H, \alpha, Wr) \rightarrow (C, \gamma)$

$Demux^{\circ}(Y, H, \alpha, De) \rightarrow (M, sc, \gamma)$

Privacy Notions

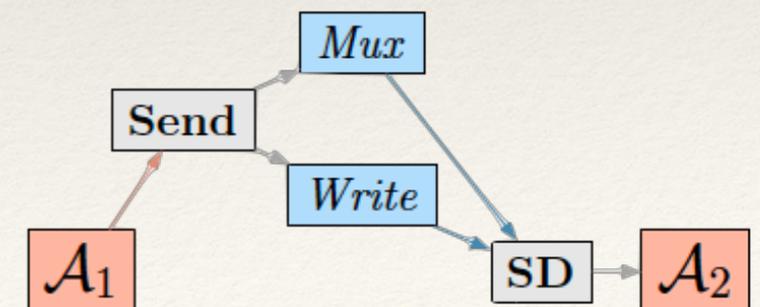
❖ PRIV-SR

- ❖ **Send** : allows adversary to provide the sender with arbitrary message fragments and stream contexts
- ❖ **Recv** : allows adversary to deliver arbitrary ciphertext fragments to the receiver



❖ PRIV-S

- ❖ PRIV-SR without access to the Recv oracle

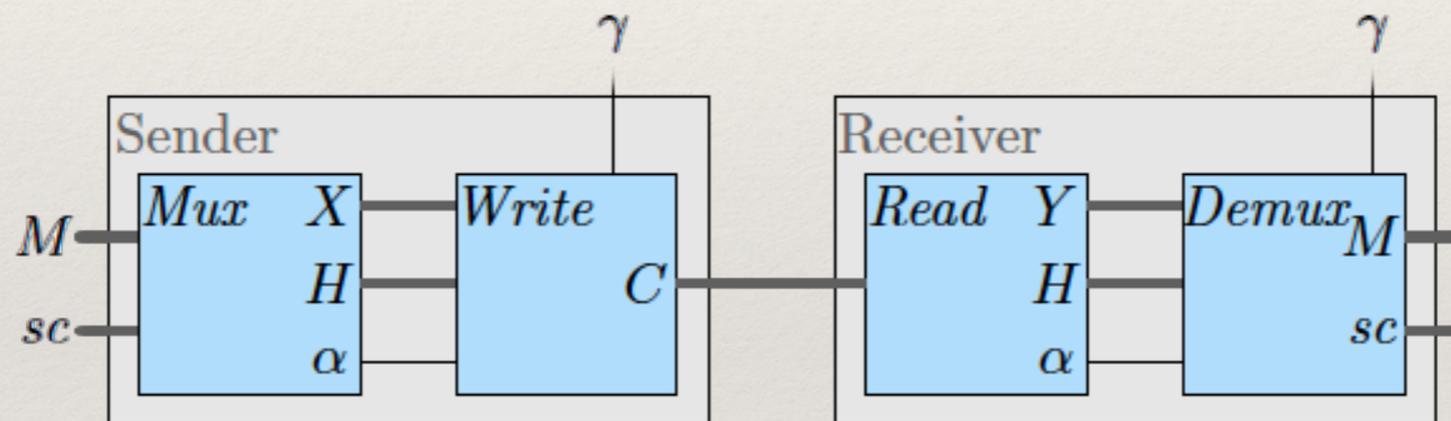


Privacy Notions

- ❖ Privacy is in terms of **left-or-right indistinguishability** of ciphertexts
 - ❖ PRIV-SR : must suppress the output of **Recv** in situations that will give trivial distinguishing attacks
 - ❖ These “situations” are when the channel is **in-sync**

Channel Synchronization

- ❖ *Read* : models the receiver side buffering and defragmentation



- ❖ Channel is **in-sync** as long as the ciphertext fragments Y output by Read remain a prefix of the ciphertext stream C transmitted by the sender

PRIV-SR Security Notion

Exp_{CH,l,b}^{priv-sr}(\mathcal{A})

```
1 declare str  $S$ ,  $Env$ , bool  $sync$ 
2 ( $Mu$ ,  $Wr$ ,  $Re$ ,  $De$ )  $\leftarrow$   $Init()$ 
3  $sync \leftarrow 1$ 
4  $b' \leftarrow \mathcal{A}_1^{Send,Recv}(\mathbf{var} \mathit{Env})$ 
5 return  $b'$ 
```

Send(M_0, sc_0, M_1, sc_1)

```
6  $L_0 \leftarrow leak(\ell, M_0, sc_0)$ 
7  $L_1 \leftarrow leak(\ell, M_1, sc_1)$ 
8 if  $L_0 \neq L_1$  then return  $(\perp, \perp)$ 
9 ( $X, H, \alpha$ )  $\leftarrow Mux^{SD}(M_b, sc_b, \mathbf{var} \mathit{Mu})$ 
10 ( $C, \gamma$ )  $\leftarrow Write^{SD}(X, H, \alpha, \mathbf{var} \mathit{Wr})$ 
11  $S \leftarrow S \parallel C$ 
12 return ( $C, \gamma$ )
```

Recv(C)

```
13 ( $Y, H, \alpha$ )  $\leftarrow Read^{SD}(C, \mathbf{var} \mathit{Re})$ 
14 ( $M, sc, \gamma$ )  $\leftarrow Demux^{SD}(Y, H, \alpha, \mathbf{var} \mathit{De})$ 
15 if  $sync$  and  $Y \preceq S$  then
16    $S \leftarrow S \% Y$ ;  $M, sc \leftarrow \perp$ 
17 else  $sync \leftarrow 0$ 
18 return ( $M, sc, \gamma$ )
```

SD(I)

```
19  $O \leftarrow \mathcal{A}_2(I, \mathbf{var} \mathit{Env})$ ; return  $O$ 
```

$leak(\ell, M, sc)$

```
20 switch ( $\ell$ )
```

```
21   case  $lensc$ : return  $\langle |M|, sc \rangle$ 
```

```
22   case  $len$ : return  $\langle |M|, |sc| \rangle$ 
```

```
23   case  $none$ : return  $\varepsilon$ 
```

$$\mathbf{Adv}_{CH,l}^{\text{priv-sr}}(\mathcal{A}) = 2\Pr_b[\mathbf{Exp}_{CH,l,b}^{\text{priv-sr}}(\mathcal{A}) = b] - 1$$

Integrity Notions

❖ INT-CS

- ❖ Requires that the channel (i.e. the **ciphertext stream**) should remain in-sync
- ❖ The adversary wins if it can make the out-of-sync **Recv** oracle output a valid message fragment and context

❖ INT-PS

- ❖ Requires that the **plaintext streams** carried by the channel should remain in-sync
- ❖ The adversary wins if at any point in the game, the output plaintext stream is not a prefix of the input plaintext stream

INT-CS and INT-PS Notions

Exp_{CH}^{int-cs}(\mathcal{A})

```

1 declare str Env, S, bool sync, win
2 (Mu, Wr, Re, De) ← Init()
3 sync ← 1;  $\mathcal{A}_1^{\text{Send,Recv}}$ (var Env)
4 return win

```

Send(M, sc)

```

5 (X, H,  $\alpha$ ) ← MuxSD(M, sc, var Mu)
6 (C,  $\gamma$ ) ← WriteSD(X, H,  $\alpha$ , var Wr)
7 S ← S || C
8 return (C,  $\gamma$ )

```

Recv(C)

```

9 (Y, H,  $\alpha$ ) ← ReadSD(C, var Re)
10 (M, sc,  $\gamma$ ) ← DemuxSD(Y, H,  $\alpha$ , var De)
11 if sync and  $Y \preceq S$  then S ← S % Y
12 else sync ← 0
13 win ← win ∨ (M ≠ ⊥ ∧ sc ≠ ⊥)
14 return (M, sc,  $\gamma$ )

```

SD(I)

```

15 O ←  $\mathcal{A}_2(I, \text{var Env})$ ; return O

```

Exp_{CH}^{int-ps}(\mathcal{A})

```

16 declare str Env, S[], str R[], bool win
17 (Mu, Wr, Re, De) ← Init()
18  $\mathcal{A}_1^{\text{Send,Recv}}$ (var Env)
19 return win

```

Send(M, sc)

```

20 (X, H,  $\alpha$ ) ← MuxSD(M, sc, var Mu)
21 (C,  $\gamma$ ) ← WriteSD(X, H,  $\alpha$ , var Wr)
22 Ssc ← Ssc || M
23 return (C,  $\gamma$ )

```

Recv(C)

```

24 (Y, H,  $\alpha$ ) ← ReadSD(C, var Re)
25 (M, sc,  $\gamma$ ) ← DemuxSD(Y, H,  $\alpha$ , var De)
26 if M ≠ ⊥ and sc ≠ ⊥ then
27   Rsc ← Rsc || M
28 if Rsc  $\not\preceq$  Ssc then win ← 1
29 return (M, sc,  $\gamma$ )

```

SD(I)

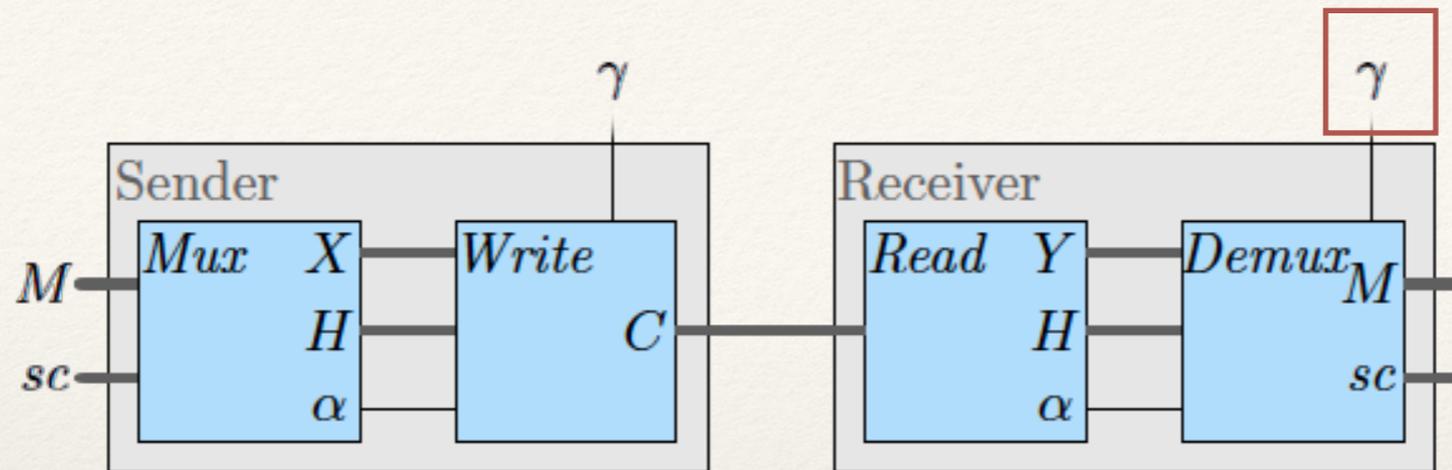
```

30 O ←  $\mathcal{A}_2(I, \text{var Env})$ ; return O

```

$$\text{Adv}_{\text{CH}}^{\text{int-cs}}(\mathcal{A}) = \Pr[\mathbf{Exp}_{\text{CH}}^{\text{int-cs}}(\mathcal{A}) = 1]$$

$$\text{Adv}_{\text{CH}}^{\text{int-ps}}(\mathcal{A}) = \Pr[\mathbf{Exp}_{\text{CH}}^{\text{int-ps}}(\mathcal{A}) = 1]$$



❖ PRIV-S \wedge INT-CS \implies ? PRIV-SR



Receiver-status Simulatability

❖ SIM-STAT

- ❖ This notion captures what the adversary learns from the receiver's state by observing the status messages output
- ❖ Simulation-based game : for every efficient adversary, efficient simulator such that real status messages are indistinguishable from fake ones

Exp _{\mathcal{CH}, S, b} ^{sim-stat}(\mathcal{A})

```
1 declare str Env, S
2 (Mu, Wr, Re, De) ← Init()
3 b' ←  $\mathcal{A}_1^{\text{Send,Recv}}$ (var Env)
4 return b'
```

Send(M, sc)

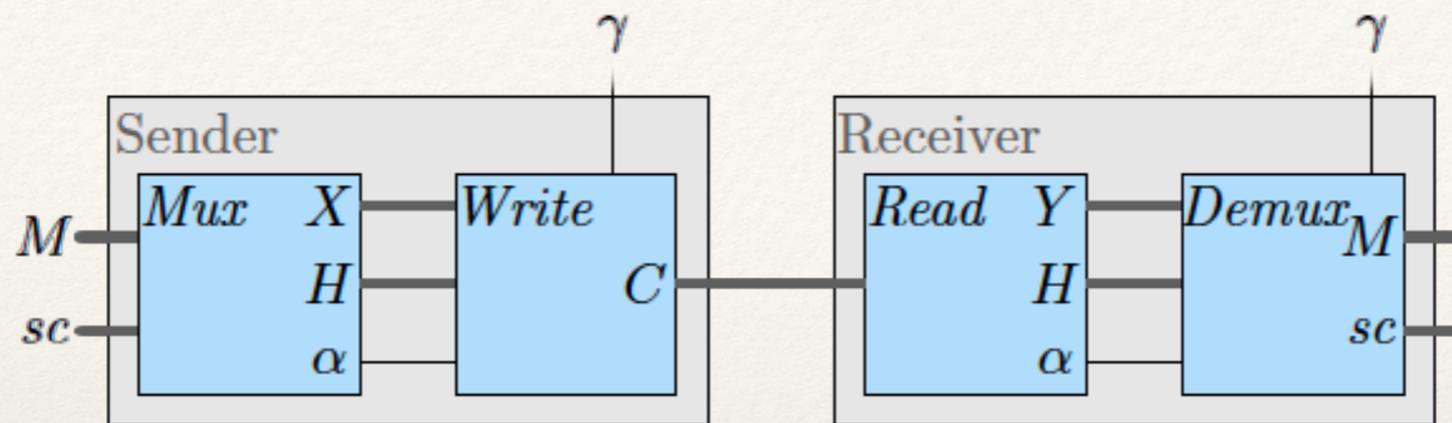
```
5 (X, H, α) ←  $\text{Mux}^{\text{SD}}$ (M, sc, var Mu)
6 (C, γ) ←  $\text{Write}^{\text{SD}}$ (X, H, α, var Wr)
7 S ← S || C
8 return (C, γ)
```

Recv(C)

```
9 if b = 1 then
10 (Y, H, α) ←  $\text{Read}^{\text{SD}}$ (C, var Re)
11 (*, *, γ) ←  $\text{Demux}^{\text{SD}}$ (Y, H, α, var De)
12 else γ ←  $\mathcal{S}^{\text{SD}}$ (C, S)
13 return γ
```

SD(I)

```
14 O ←  $\mathcal{A}_2(I, \text{var Env})$ ; return O
```



❖ $\text{PRIV-S} \wedge \text{INT-CS} \wedge \text{SIM-STAT} \implies \text{PRIV-SR}$

The TLS 1.3 Record Layer

- ❖ Three client-server protocols executing concurrently
 - ❖ **handshake** : (re-)initialization of the channel
 - ❖ **record** : exchange application data
 - ❖ **alert** : close the channel

- ❖ Each flow is authenticated and encrypted as soon as client and server exchange key material

TLS 1.3 Records

- ❖ Plaintext records encode:
 - ❖ content type
 - ❖ stream fragment
 - ❖ length of fragment ($< 2^{14}$ bytes)
 - ❖ *legacy_record_version* (for backward compatibility)
- ❖ Streams of data are transformed into a sequence of records
- ❖ Record boundaries are subject to certain rules

Record Boundary Rules

- ❖ Handshake : no interleaving
- ❖ Handshake : no spanning a key change
- ❖ Handshake and Alert : no zero length messages
- ❖ One alert per record

The Core Components

- ❖ Which fully specified components can be altered without affecting security?
- ❖ Which unspecified or partially specified components are critical to security?

Observations

- ❖ Record boundaries may leak the content type!
 - ❖ Hiding the content and the type unachievable in general due to the record boundary rules
- ❖ Associated data is unauthenticated

Record Header Authentication

- ❖ Header : *opaque_type, legacy_record_version, length*
- ❖ What if the header is different than specified?
 - ❖ *length* changed : invalid with high probability
 - ❖ If the others are changed, it should be alright since it doesn't affect decryption - it is left optional in the spec
 - ❖ But this is an INT-CS attack!
 - ❖ We **must** authenticate the header
- ❖ To formalize that the value should not affect security, we allow the specification details to choose the bits

Is the model too strong?

- ❖ One point of view is that this does not constitute a “real attack” on privacy or integrity, since inputs to decryption were not affected
- ❖ This is correct only if down-stream handling of the plaintext is independent of these values

The Core Components

- ❖ **Stream Multiplexer**

- ❖ Transforms data streams into records
- ❖ Captures the non-cryptographic functionality
- ❖ Consider it to be partially specified

- ❖ **AEAD scheme**

- ❖ **Nonce generator**

- ❖ Both these are core cryptographic functionalities
- ❖ Required to be fully specified

Partially-specified Multiplexers

❖ mPRIV-S

- ❖ Captures the adversary's ability to discern information about the inputs to *Mux* given its outputs
- ❖ Like the PRIV-S game earlier, except:
 - ❖ No Write oracle.
 - ❖ Rather than (X, γ) , it returns γ and the length of X

❖ SIM-mSTAT

- ❖ Captures simulatability of the status message output by *Demux*.

AEAD Scheme

- ❖ Encryption and Decryption are both deterministic
- ❖ Standard security notions are as follows:
 - ❖ PRIV
 - ❖ Indistinguishability under Chosen Plaintext Attack
 - ❖ INT
 - ❖ Integrity of Ciphertexts

Nonce Generator

- ❖ It consists of a pair of algorithms:
 - ❖ $Init() \rightarrow ng$ (randomized, initializes the state)
 - ❖ $Next(ng) \rightarrow N$ (computes the next nonce and updates state)
- ❖ $Coll$: outputs 1 if there is a nonce-reuse

Partially Specified Record Layer

- ❖ $(\text{PRIV}_{\text{AEAD}}) \wedge (\text{mPRIV-S}_{\text{Mux}}) \implies (\text{PRIV-S}_{\text{CH}})$
- ❖ $\text{INT}_{\text{AEAD}} \implies \text{INT-CS}_{\text{CH}}$
- ❖ The SIM-STAT security of the channel reduces to the SIM-mSTAT security of the multiplexer and the integrity of the AEAD scheme
- ❖ Can combine all these with the earlier result regarding PRIV-SR security

Conclusion

- ❖ Partial specification of protocols is simple and flexible
- ❖ Allows us to think formally about what the protocol must get right, and what it may get wrong
- ❖ Helps point out which matters are security-critical

Thank You!