

Computer Networks

X_400487

Lecture 3

Chapter 3: The Data Link Layer—Part 1



Lecturer: Jesse Donkervliet






Recap of the Physical Layer

Responsible for transferring *bits* over a *wire-like* medium.

Maximum data rate determined by *bandwidth* and *signal-to-noise ratio*.

$$R = B \times \log_2 \left(1 + \frac{S}{N} \right)$$

Physical layer responsible for translating bits to and from waves; sharing channel with multiple users



0100111101010101100110101100110101010110110110110000110101010011001
01011010110010111010100101010011110101010110011010110011010101011101
00111101010101001010101011011011011000011010101001100101011101010010
0100101010011110101010100110101011001110101101101101100001101010100
10101011011011011000011010101001100110011010101101010101011011001101
01010101011011001101010101001010101110101001010100111101010101100110
10100110011001101010110101010101101100110101000101111010101011001101
01101100110101010100101110101010100110101010110110011010101010010101

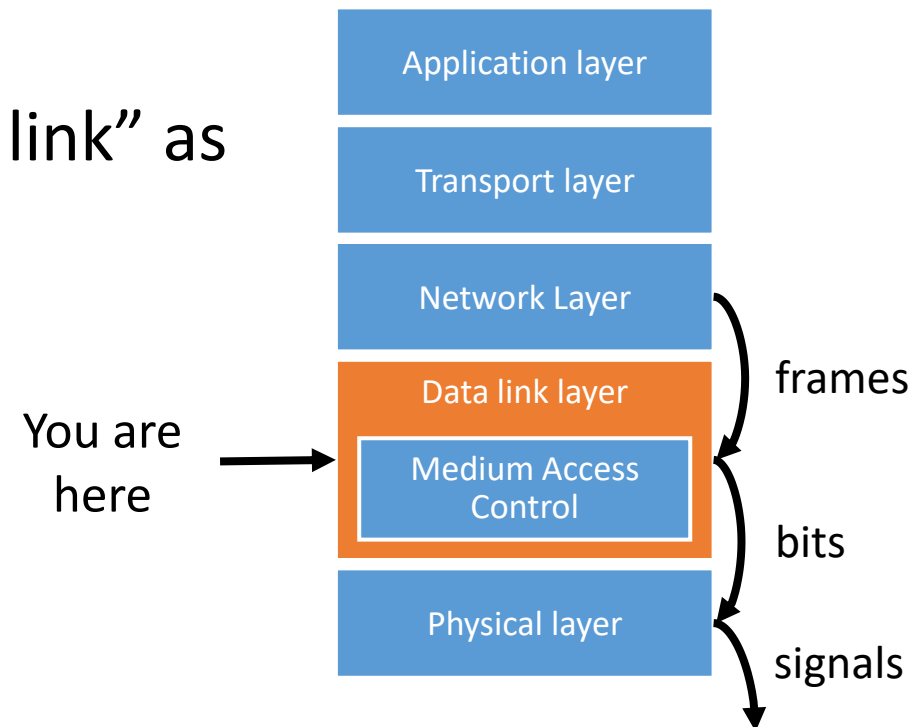
Where are the frames?

10101001111010101011001101011001101010101101101101100001101010100110
100110010011010101101010101101100110101010100101101101010101101101
00110101100110101010110110110110000101110101001010100111101010101100
1001101011101010010101001111010101011001101011001101010101101101101
01101011001101010101101101100001101010100110011001101010110101010
10101011011001101010101001010011001101010110101010101101100110101010
10101011010101010110110011010101010010101101101101100001101010100110
10010101110101001010100111101010101100110101100110101010110110110110
011001101010101001010111101010010101001111010101011001101011001101010
01111010101011001010111010100101010011110101010110011010110011010101

The Data Link layer

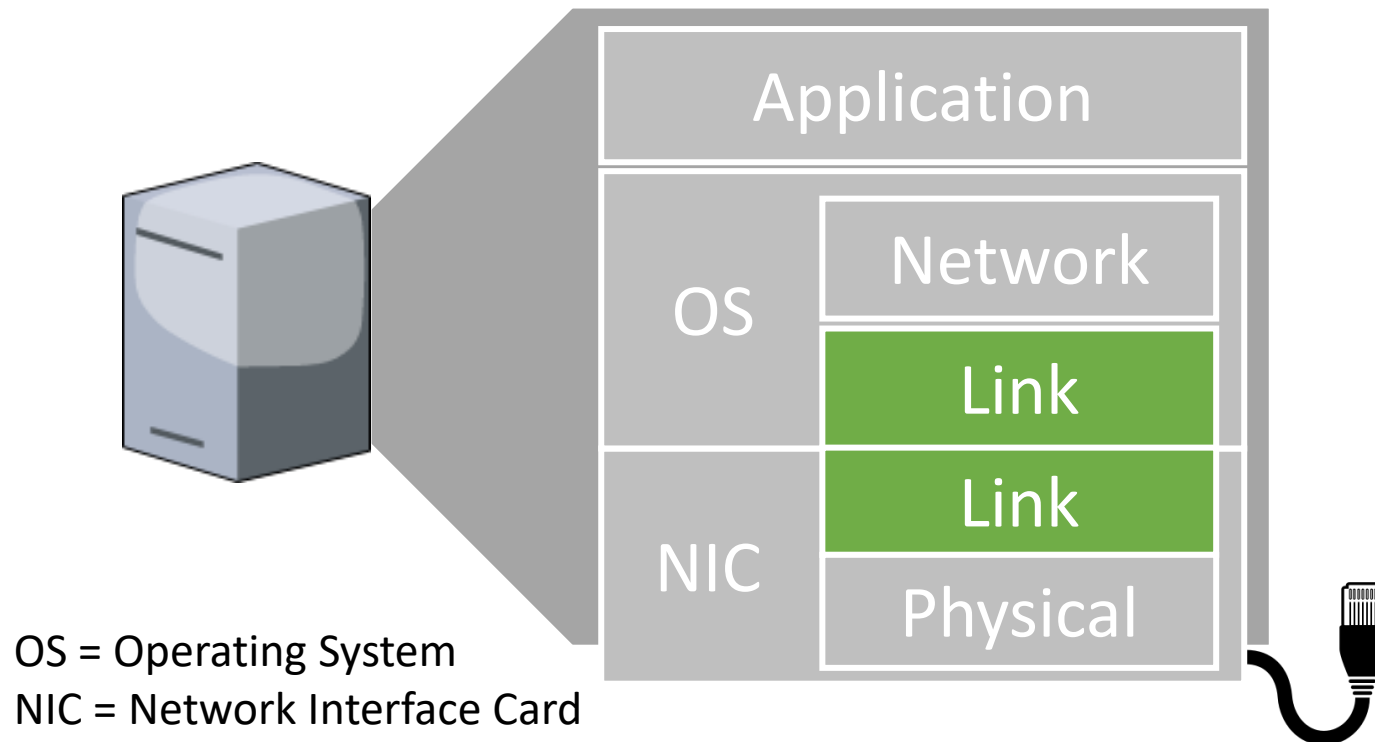
Responsible for transferring *frames* over a single link

1. Groups bits into frames
2. Offers “sending frames over a link” as a *service* to the network layer
3. Handles **Q: Why needed?** transmission errors
4. Regulates data flow



Link layer environment

Commonly implemented as NICs and OS drivers; network layer (IP) is often OS software.



Data Link Layer — Roadmap

Part 1


- **Framing**
- **Flow Control**
- **Guaranteed Delivery**
- **Sliding Window Protocols**

Part 2

- Error detection
- Error correction

Framing

From Bit Stream to Discrete Units of Information




Framing Methods

1. Byte count.
2. Flag bytes with byte stuffing.
3. Flag bits with bit stuffing.
4. Use special symbols in physical layer.

'Cheating' because physical layer does not know about frames (according to *our* model)

Use for '1' and '0'

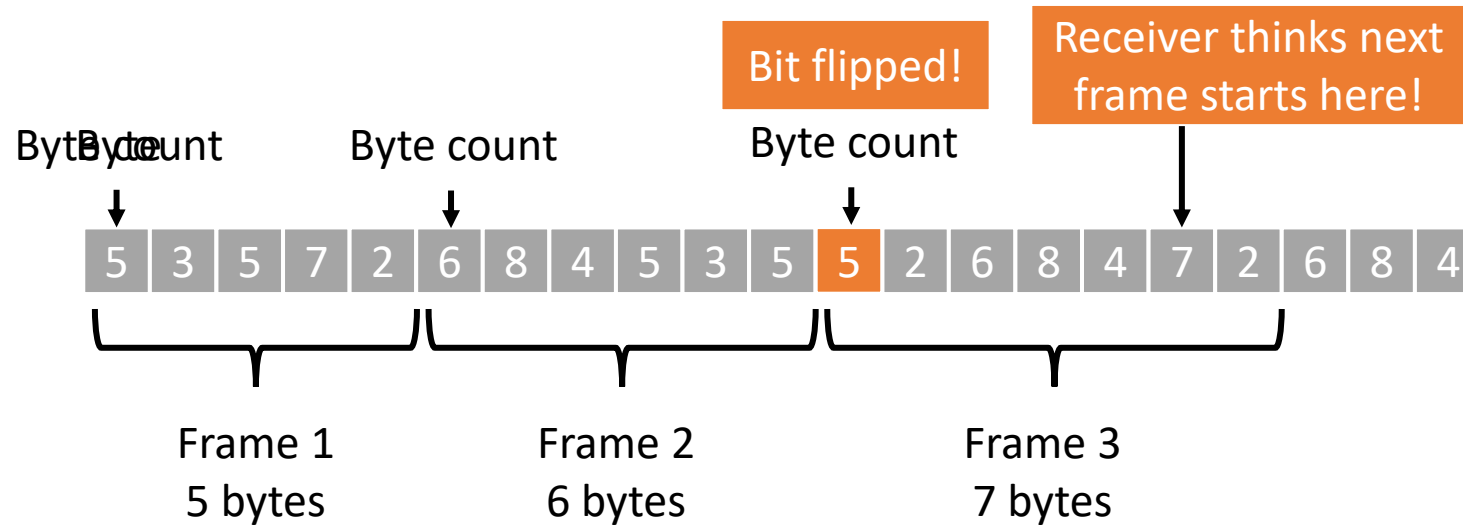


Example of method 4 using Phase Shift Keying

Framing

Byte count

Q: Advantage?
Disadvantage?




Framing

Byte stuffing

Q: Disadvantage?

Use a 'flag' byte to indicate start and end of frame.

Let's say our flag byte is 00000111_2 (7_{10}).



Framing


Byte stuffing

Character	Escape Character
%	%25

What if the data contains a flag byte?

Use an 'escape' byte to ignore certain flag bytes.

Let's say our escape byte is 00001001_2 (9_{10}).



Q: Algorithm on the receiving side?


Q: Are we done?

Framing

Byte stuffing

Escape bytes can also occur in data!

Let's use letters for generality.
Flag byte = F, Escape byte = E.



Escape both 'escape' and 'flag' bytes.

Q: What is the overhead of this approach?

Framing

Bit stuffing

Byte stuffing can be space inefficient.

Byte stuffing

Flag byte →

Escape byte →

Bit stuffing

Bit pattern

Insert single bit

Example:

Bit pattern:

01111110

Insert bit if pattern in data:

011111010

Add one extra bit



Bit stuffing example

Receiver

Bit pattern: 01111110

```
1100010100011110111100011100110110000001101110001111101000101101111
0100011010100010010100111110100001110010000011001001010000001011111
0101110000111110100110001110011111001100110000011101111001111001010
0001110111100001110111000110010110011001100011111001001110011010111
0101001110000111110110101000000100011111001101011100110011000010100
0101010000011010010111011010100101001001111100010010100100000010010
1101111101001100100111000011101111101111000010100110100011111100111
1110000000100100001001011001110101101100000011000110101100000000011
0101101001100001100000111101100100101001100110110110011001101101100
0101000111101011101000111110110010101000001010011011100110101101000
0110101111100011100010000010110101010100101111100101100010010010000
0110101110111011100000010001100100000001111101110010010111011010010
0010111110010101001011100011001010011001000100011010100100101010011
0111001111000111101000101001010011111000111110001001111101100010000
1010001101010111001011100010010011110101101001111100001111101101001
1011001011111010001010110000011111000111
```


Q: Are we done?

Bit pattern: 01111110

Bit stuffing example

Receiver

F1 {

```
1100010100011110111100011100110110000001101110001111101000101101111
0100011010100010010100111110100001110010000011001001010000001011111
0101110000111110100110001110011111001100110000011101111001111001010
0001110111100001110111000110010110011001100011111001001110011010111
0101001110000111110110101000000100011111001101011100110011000010100
0101010000011010010111011010100101001001111100010010100100000010010
1101111101001100100111000011101111101111000010100110100
      0000001001000010010110011101011011000000110001101011000000000011
0101101001100001100000111101100100101001100110110110011001101101100
0101000111101011101000111110110010101000001010011011100110101101000
0110101111100011100010000010110101010100101111100101100010010010000
0110101110111011100000010001100100000001111101110010010111011010010
0010111110010101001011100011001010011001000100011010100100101010011
0111001111000111101000101001010011111000111110001001111101100010000
1010001101010111001011100010010011110101101001111100001111101101001
1011001011111010001010110000011111000111
```

F2 }

Bit stuffing example

Receiver

Bit pattern: 01111110

F1 {

```
11000101000111101111000111001101100000011011100011111 1000101101111
010001101010001001010011111 100001110010000011001001010000001011111
 10111000011111 100110001110011111 01100110000011101111001111001010
0001110111100001110111000110010110011001100011111 01001110011010111
010100111000011111 110101000000100011111 01101011100110011000010100
01010100000110100101110110101001010010011111 0010010100100000010010
11011111 10011001001110000111011111 1111000010100110100
  000000100100001001011001110101101100000011000110101100000000011
0101101001100001100000111101100100101001100110110110011001101101100
010100011110101110100011111 110010101000001010011011100110101101000
01101011111 00111000100000101101010101001011111 0101100010010010000
01101011101110111000000100011001000000011111 1110010010111011010010
001011111 010101001011100011001010011001000100011010100100101010011
0111001111000111101000101001010011111 0011111 0010011111 1100010000
10100011010101110010111000100100111101011010011111 00011111 1101001
1011001011111 10001010110000011111 00111
```


F2 }

Bit pattern in data

Flow Control

Could you speak more slowly, please?

A Resource Management problem

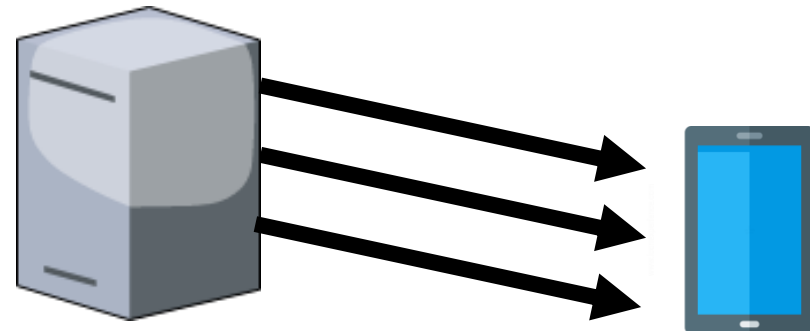


Utopian simplex protocol

The ideal case

...


```
while True:  
    packet = from_network_layer()  
    frame.payload = packet  
    to_physical_layer(frame)
```




Stop-and-wait for error-free channel

...

```
while True:  
    packet = from_network_layer()  
    frame.payload = packet  
    to_physical_layer(frame)  
    event = wait_for_event()
```




Example of Utopian Simplex Protocol



Guaranteed Delivery

Acknowledgments, Sequence Numbers, and Retransmissions



How Can We Know If a Frame Gets Lost?



Ask a different question

Q: How can we know if a frame *arrives*?

Send a message back:
“I got your message!”


Q: When do we want
to retransmit data?

Q: What if the acknowledgment gets lost?

We assume our original
message did not arrive

It depends on ...
the application!


Acknowledgments let the sender know
it does *not* need to retransmit data.



Many protocols either use or don't use acknowledgments. Different approach: Support acknowledgments, but let the application decide if it needs to use acknowledgments or not.



To acknowledge, or not to acknowledge




To acknowledge, or not to acknowledge

It depends on ...
the physical medium!

It depends on ...
the application!

This problem will return later
(in the transport layer)




Automatic Repeat ReQuest (ARQ) Guaranteed Delivery over Unreliable Channel


Same as stop-and-wait, except:

1. Keep track of frames using sequence numbers.
2. Wait until previous frame has been accepted.


ARQ Example



ARQ Example

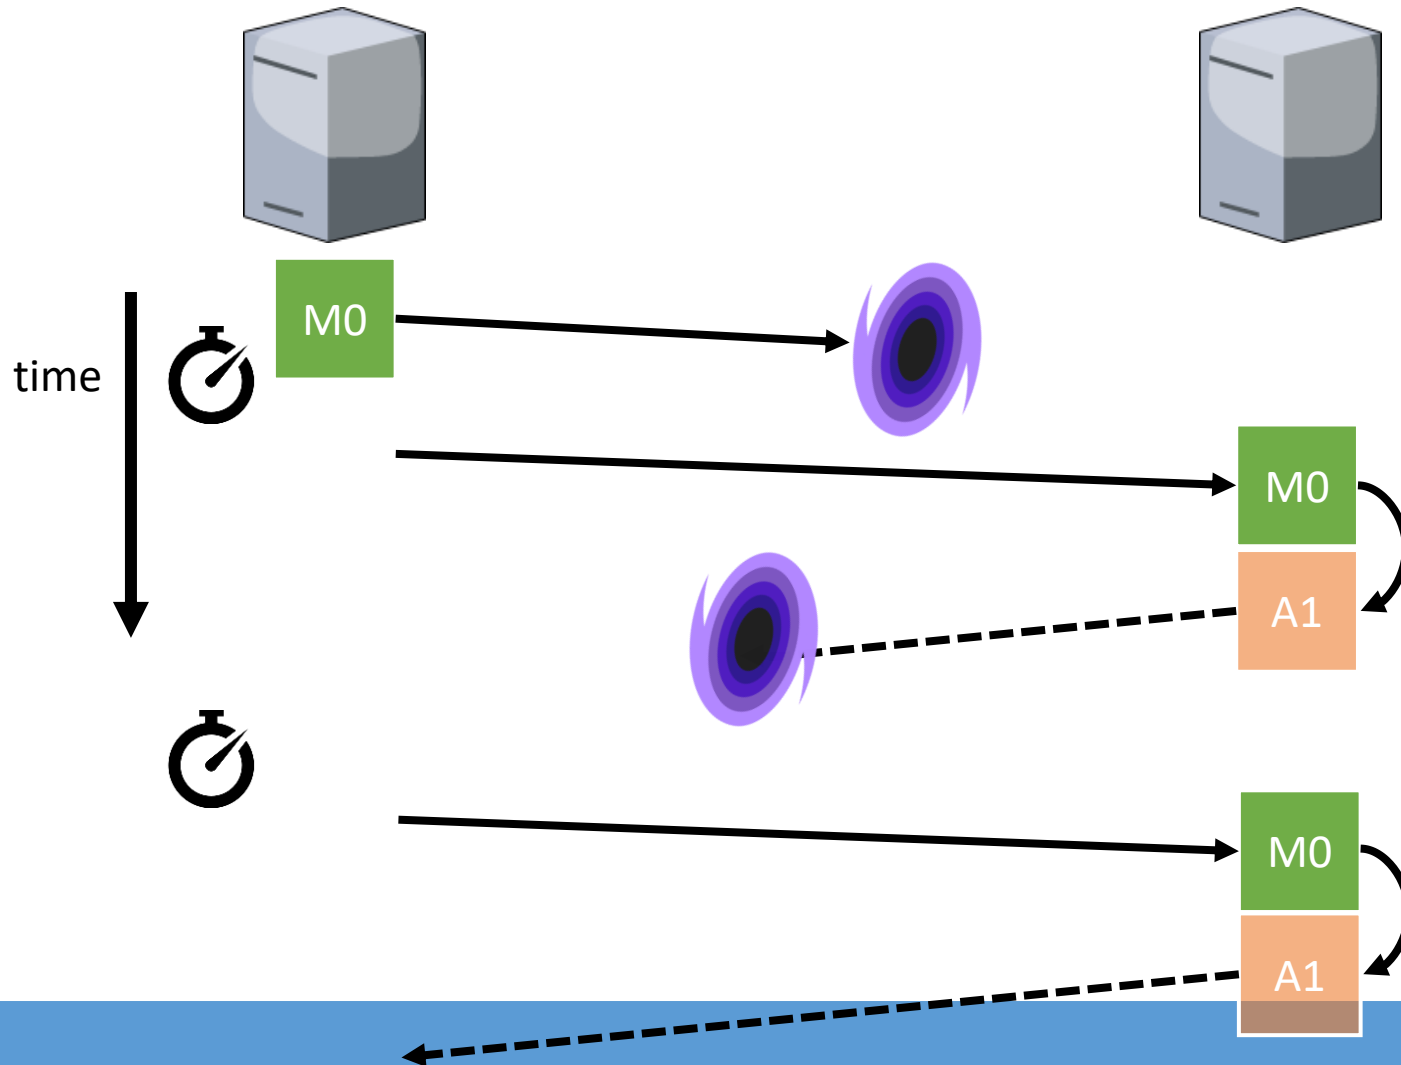


ARQ Example



Sequence numbers needed to differentiate between retransmission and next frame

ARQ Example



Sequence numbers needed to differentiate between retransmission and next frame

Automatic Repeat ReQuest (ARQ) Guaranteed Delivery over Unreliable Channel

ARQ adds error control

Receiver acks frames that are correctly delivered.

Sender sets timer and resends frame if no ack.

Q: How long should we wait?

Q: What can go wrong?

Frames and acks must be identifiable (e.g., with sequence number)

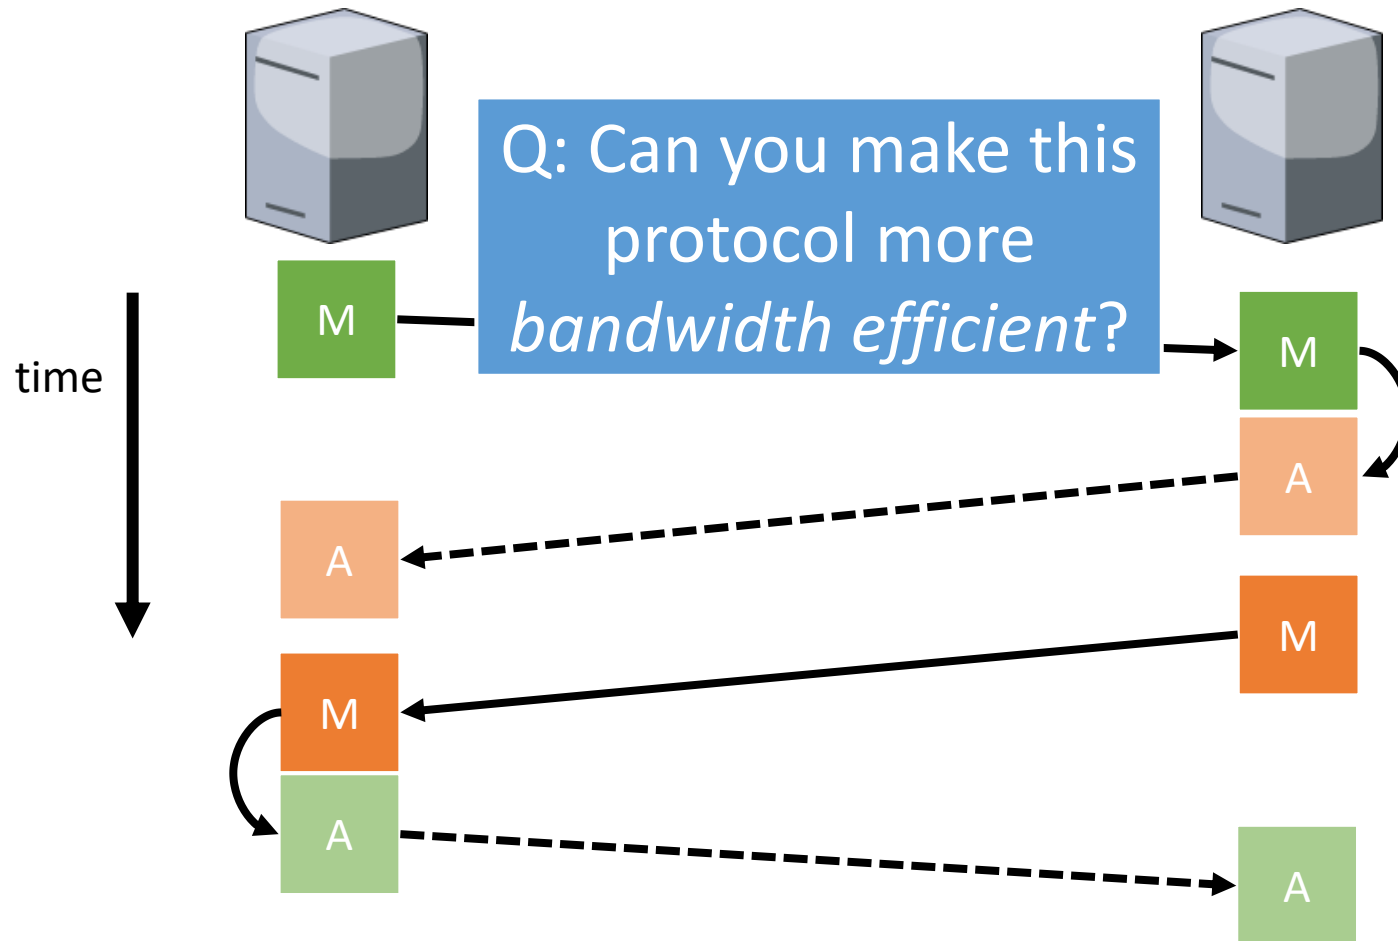
Else receiver cannot tell retransmission (due to lost ack or early timer) from new frame.

For stop-and-wait, 2 numbers (1 bit) are sufficient.

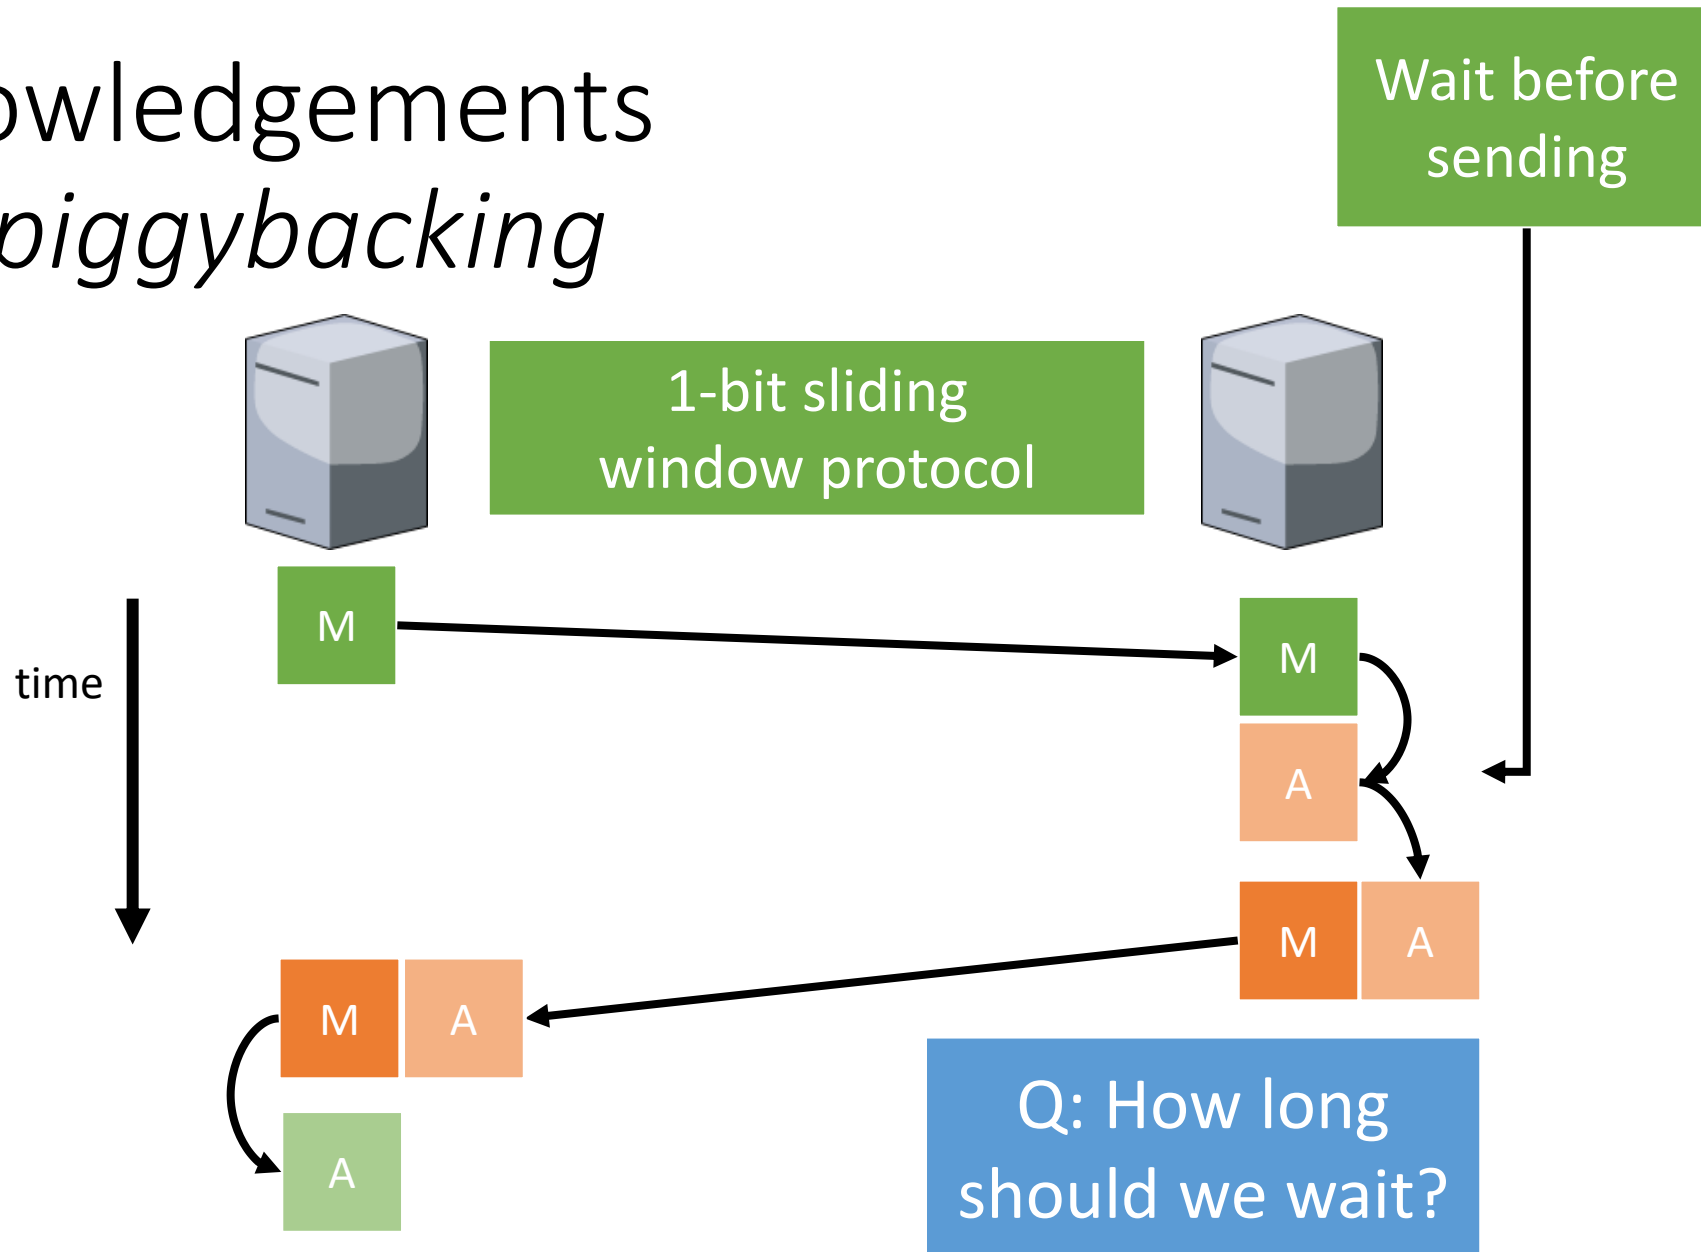
Q: Why sufficient?

Acknowledgements

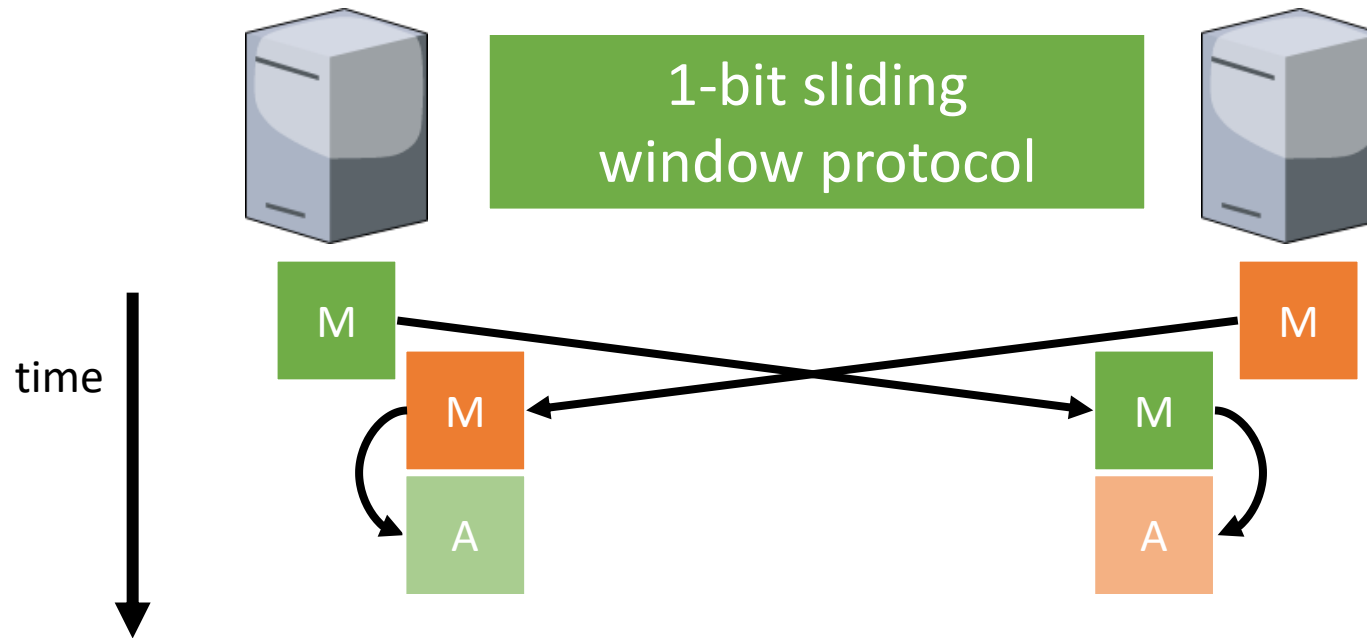
Bidirectional communication



Acknowledgements With *piggybacking*



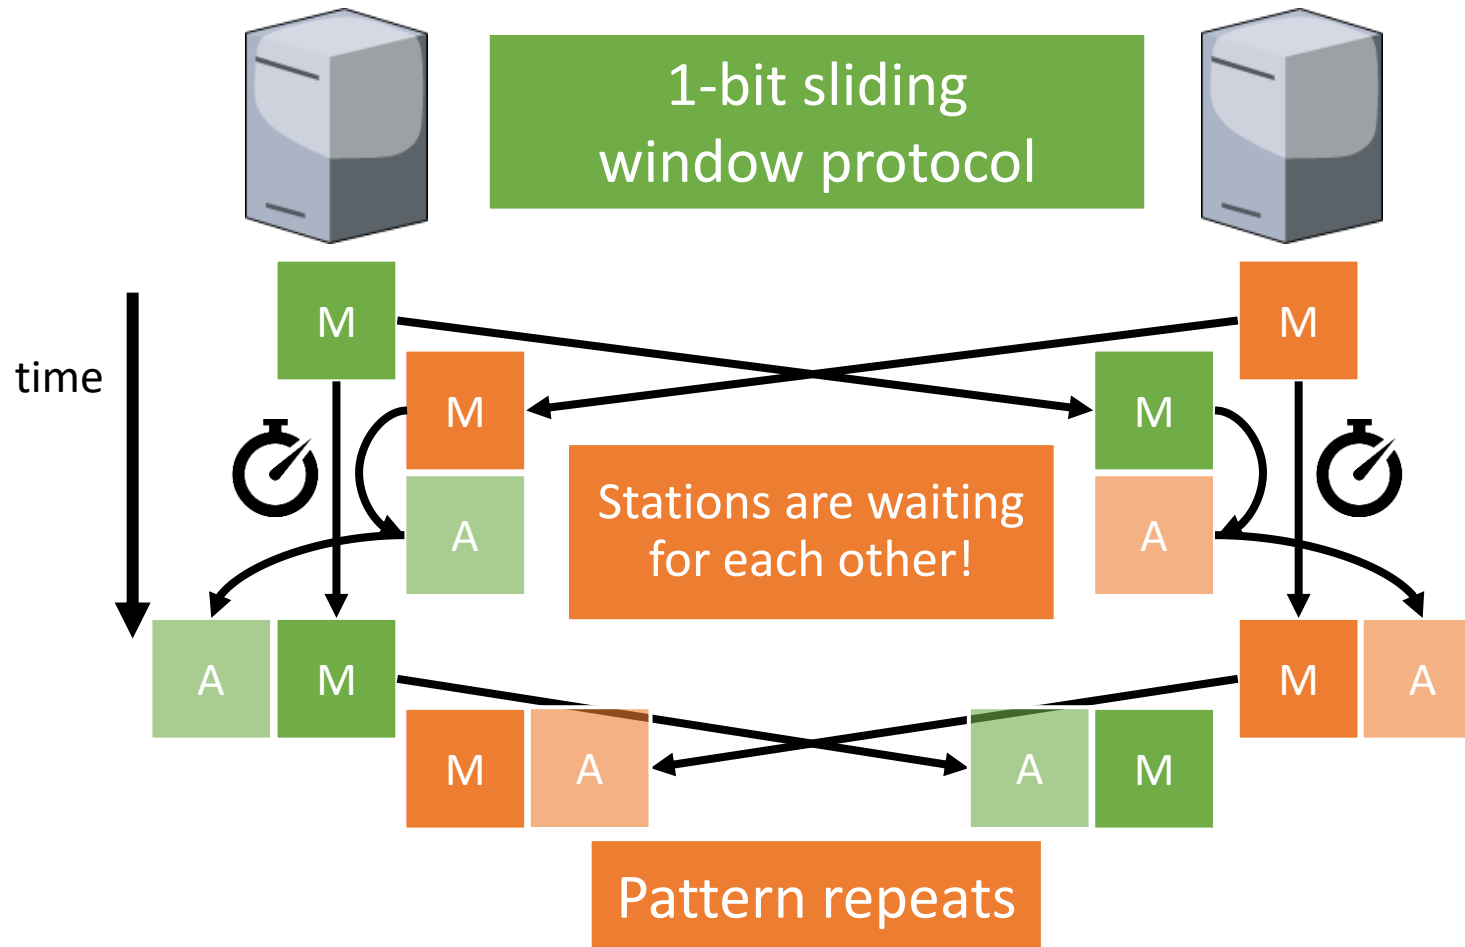
Stop-and-Wait



Q: What happens now?

Stop-and-Wait Special Case

Even without errors, half the bandwidth is wasted on retransmissions




Sliding Window Protocols


Improving Performance using *Pipelining*



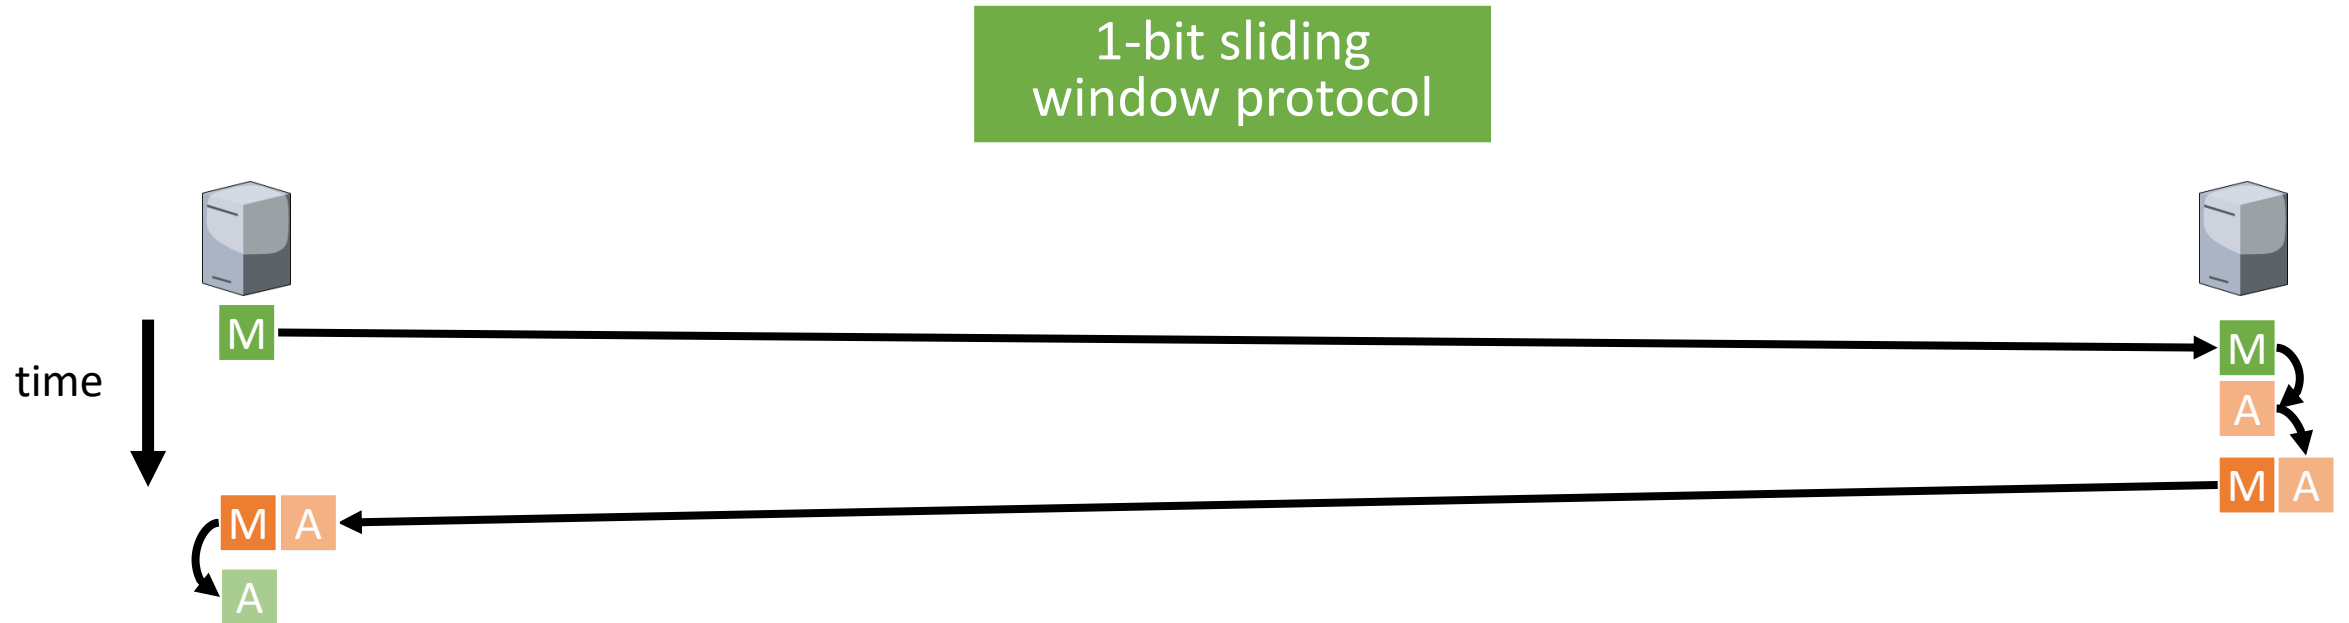
Stop-and-Wait: A 1-Bit Sliding Window Protocol



Stop-and-Wait: A 1-Bit Sliding Window Protocol



Stop-and-Wait: A 1-Bit Sliding Window Protocol



Bandwidth inefficient for high-latency channels

Q: Which properties cause performance to decrease?

Sliding window protocols

When using stop-and-wait, *data rate decreases* when:

- Latency increases
- Frame size decreases

Solution

Send next frame while waiting for acknowledgment of current frame

Sender window specifies how many frames a sender is allowed to send before waiting for an acknowledgement.

Receiver window specifies the range of frames that the receiver is allowed to accept.

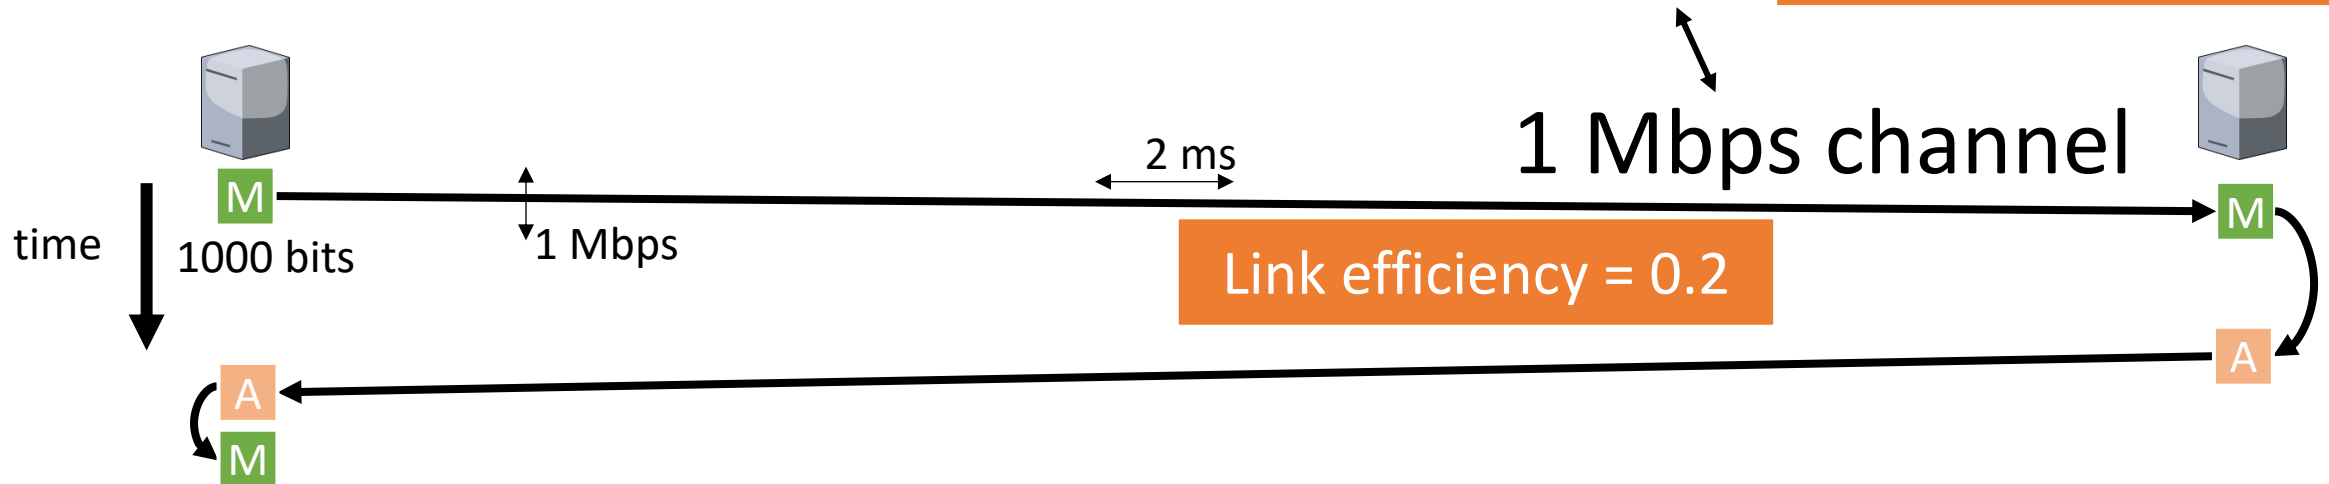
Stop-and-Wait / ARQ: A 1-Bit Sliding Window Protocol

It takes $\frac{1,000}{1,000,000} = 0.001$ seconds to send frame

It takes 2 ms for the frame to arrive at the receiver, takes 2 ms for the (0-bit) acknowledgment to come back at the sender

1 frame per $0.001+0.002+0.002$ seconds = 200 kbps

Small window reduces performance



Stop-and-Wait: A 1-Bit Sliding Window Protocol


- Frame size (in bits/bytes): f
- Window size (in frames): w
- Bandwidth (max. data rate of physical channel): B_p
- Bandwidth (frames per second): B_f
- Propagation delay (in seconds): D

It takes $\frac{f}{B_p}$ seconds to send frame, $\frac{B_p}{f} = B_f$

It takes D s for the frame to arrive at the receiver, takes D s for the (0-bit) acknowledgment to come back at the sender

1 frame per $\frac{f}{B_p} + 2 \times D$ seconds

Link utilization $\leq \frac{w}{1 + 2B_f D}$




Go-Back-N




	IF	ID	EX	MEM	WB				
i		IF	ID	EX	MEM	WB			
$i+1$			IF	ID	EX	MEM	WB		
$i+2$				IF	ID	EX	MEM	WB	
$i+3$					IF	ID	EX	MEM	WB

Q: Do you recognize this type of optimization?

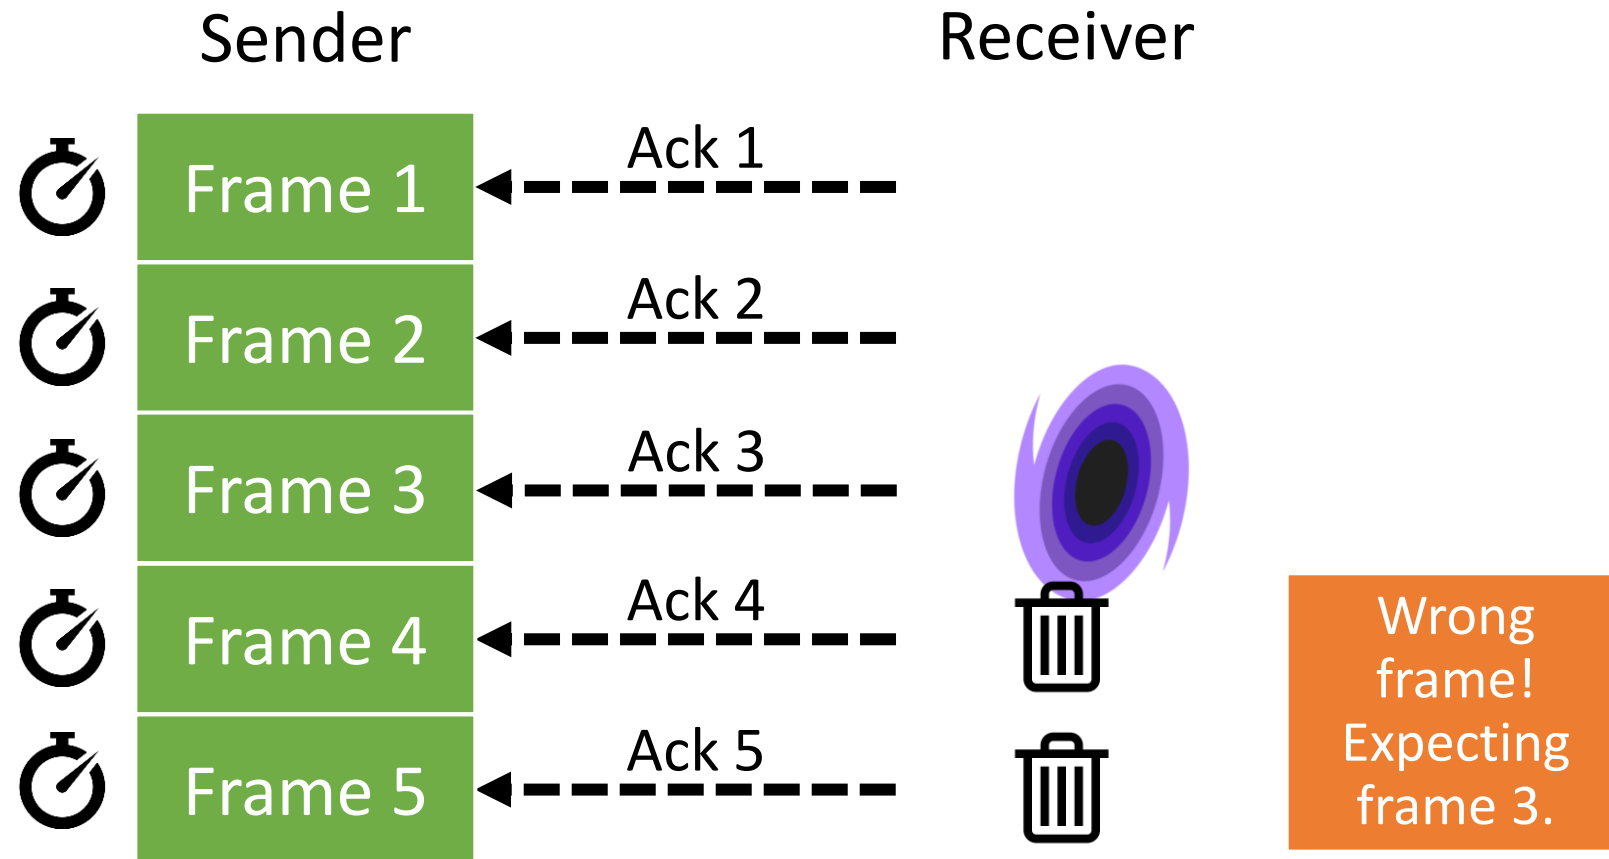


time
↓



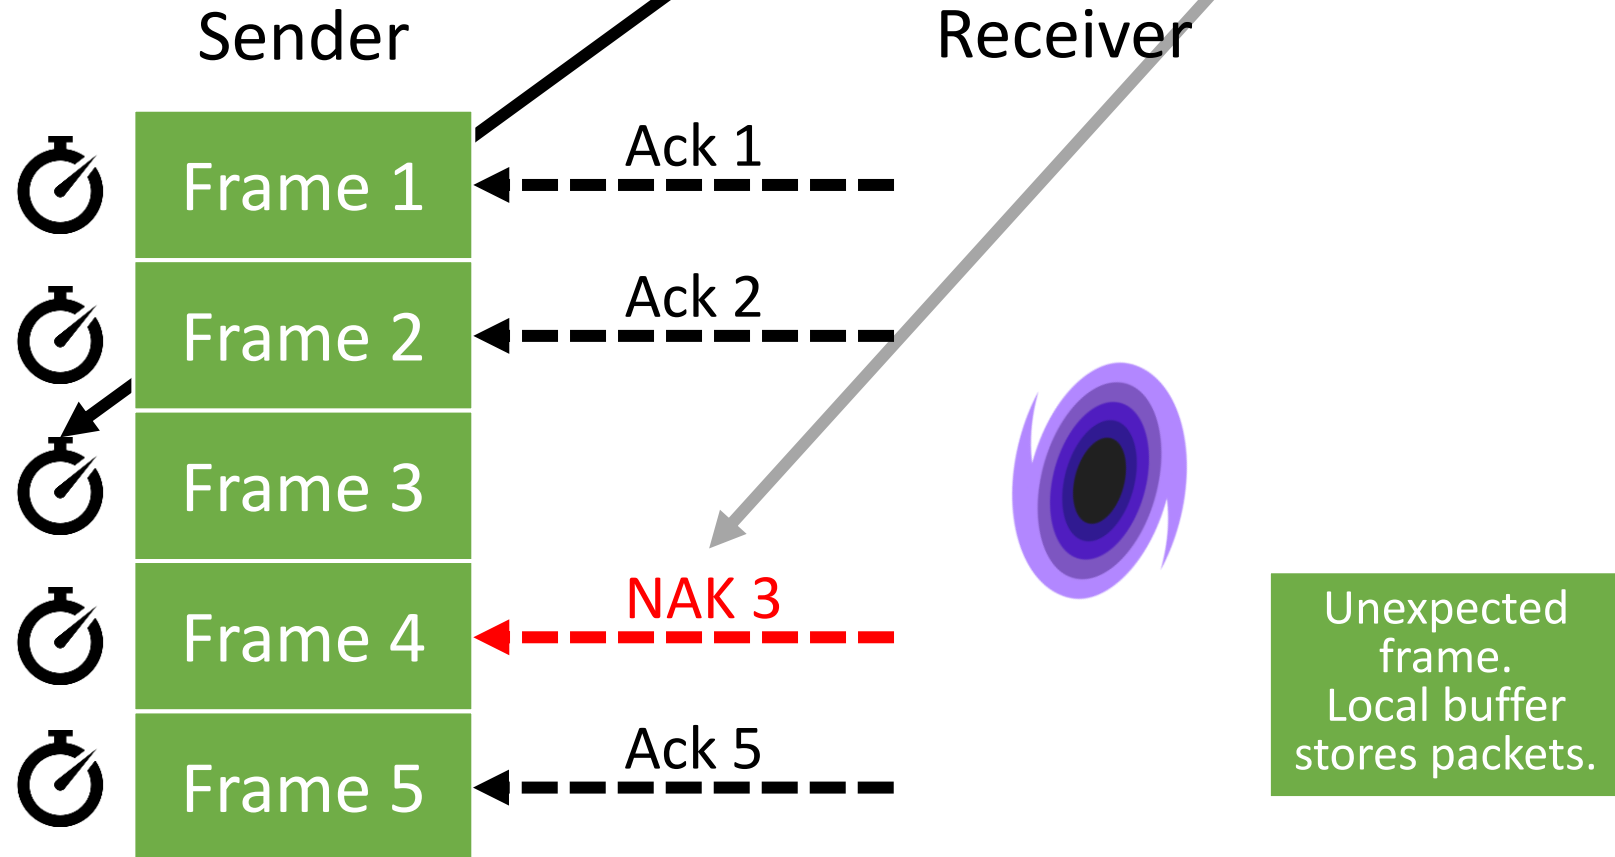
Q: What is the size of the receiver window?

Go-Back-N



Selective repeat

Selective repeat can use timers, negative acknowledgements, or both.



Delivering frames in the order in which they were sent is easy after adding sequence numbers.



Data Link Layer — Roadmap

Part 1

- **Framing**
- **Flow Control**
- **Guaranteed Delivery**
- **Sliding Window Protocols**

Part 2

- Error detection
- Error correction