



Computer Networks Term B10

Data Link Layer Outline

- Parallelism between Transport and Data Link Layer
- Tanenbaum's Treatment/Model of Data Link Layer
- Protocol 1: Utopia
- . Protocol 2: Stop-and-Wait
- Protocol 3: Positive Acknowledgment with Retransmission [PAR]
 - Old 'flawed version
 - Newer version



DL Layer Outline (cont)

- Pipelining and Sliding Windows
- . Protocol 4: One Bit Sliding Window
- Protocol 5: Go Back N
- Protocol 6: Selective Repeat
- Further Details and Decisions



Reliable Protocols at Two Layers



Tanenbaum's DL Layer Treatment

- Concerned with communication between two adjacent nodes in the subnet (node to node).
- Assumptions:
 - The bits are delivered in the order sent.
 - A rigid interface between the HOST and the node → the communications policy and the Host protocol (with OS effects) can evolve separately.
 - He uses a simplified model.



Tanenbaum's 'Simplified Model



Tanenbaum's Data Link Layer Model Assume the sending Host has *infinite* supply of messages. A node constructs a frame from a single packet message. The CRC is automatically appended in the hardware. The protocols are developed in increasing complexity to help students understand the data link layer issues.



Basic Elements of ARQ





Tanenbaum's Protocol Definitions

#define MAX_PKT 1024	/* determines packet size in bytes */
typedef enum {false, true} boolean; typedef unsigned int seq_nr; typedef struct {unsigned char data[MAX_PKT];} typedef enum {data, ack, nak} frame_kind;	/* boolean type */ /* sequence or ack numbers */ packet;/* packet definition */ /* frame_kind definition */
typedef struct { frame_kind kind; seq_nr seq; seq_nr ack; packet info; } frame:	/* frames are transported in this layer */ /* what kind of a frame is it? */ /* sequence number */ /* acknowledgement number */ /* the network layer packet */

Continued \rightarrow

Figure 3-9. Some definitions needed in the protocols to follow. These are located in the file protocol.h.



Protocol

Definitions (continued)

Figure 3-9. Some definitions needed in the protocols to follow. These are located in the file protocol.h. /* Wait for an event to happen; return its type in event. */
void wait_for_event(event_type *event);

/* Fetch a packet from the network layer for transmission on the channel. */ void from_network_layer(packet *p);

/* Deliver information from an inbound frame to the network layer. */
void to_network_layer(packet *p);

/* Go get an inbound frame from the physical layer and copy it to r. */ void from_physical_layer(frame *r);

/* Pass the frame to the physical layer for transmission. */ void to_physical_layer(frame *s);

/* Start the clock running and enable the timeout event. */
void start_timer(seq_nr k);

/* Stop the clock and disable the timeout event. */
void stop_timer(seq_nr k);

/* Start an auxiliary timer and enable the ack_timeout event. */ void start_ack_timer(void);

/* Stop the auxiliary timer and disable the ack_timeout event. */
void stop_ack_timer(void);

/* Allow the network layer to cause a network_layer_ready event. */ void enable_network_layer(void);

/* Forbid the network layer from causing a network_layer_ready event. */ void disable_network_layer(void);

/* Macro inc is expanded in-line: Increment k circularly. */ #define inc(k) if (k < MAX_SEQ) k = k + 1; else k = 0



Packet and Frame Definitions



physical layer



Computer Networks Data Link Layer

/* Protocol 1 (utopia) provides for data transmission in one direction only, from sender to receiver. The communication channel is assumed to be error free, and the receiver is assumed to be able to process all the input infinitely quickly. Consequently, the sender just sits in a loop pumping data out onto the line as fast as it can. */

typedef enum {frame arrival} event type; #include "protocol.h"

void sender1(void)

frame s; packet buffer;

void receiver1(void)

event_type event;

while (true) {

frame r;

```
while (true) {
    from_network_layer(&buffer);
    s.info = buffer;
    to_physical_layer(&s);
}
```

/* buffer for an outbound frame */
/* buffer for an outbound packet */

- ; /* go get something to send */
 - /* copy it into s for transmission */
 - /* send it on its way */
 - * Tomorrow, and tomorrow, and tomorrow, Creeps in this petty pace from day to day To the last syllable of recorded time - Macbeth, V, v */

/* filled in by wait, but not used here */

- /* only possibility is frame_arrival */
- /* go get the inbound frame */
- /* pass the data to the network layer */

Figure 3-10

Unrestricted Simplex Protocol



wait for event(&event);

from_physical_layer(&r); to network layer(&r.info); /* Protocol 2 (stop-and-wait) also provides for a one-directional flow of data from sender to receiver. The communication channel is once again assumed to be error free, as in protocol 1. However, this time, the receiver has only a finite buffer capacity and a finite processing speed, so the protocol must explicitly prevent the sender from flooding the receiver with data faster than it can be handled. */

typedef enum {frame_arrival} event_type; #include "protocol.h"

void sender2(void)

frame s; packet buffer; event_type event;

while (true) {
 from_network_layer(&buffer);
 s.info = buffer;
 to_physical_layer(&s);
 wait_for_event(&event);

/* buffer for an outbound frame */
/* buffer for an outbound packet */
/* frame_arrival is the only possibility */

/* go get something to send */ /* copy it into s for transmission */ /* bye bye little frame */ /* do not proceed until given the go ahead */

void receiver2(void)

}

frame r, s; event_type event; while (true) { wait_for_event(&event); from_physical_layer(&r); to_network_layer(&r.info); to_physical_layer(&s);

/* buffers for frames */ /* frame_arrival is the only possibility */

/* only possibility is frame_arrival */ /* go get the inbound frame */ /* pass the data to the network layer */ /* send a dummy frame to awaken sender */



Figure 3-11

Stop-and-Wait Protocol



Protocol 3: Positive Acknowledgement with Retransmissions [PAR]

Introduce Noisy Channels

- This produces:
 - 1. Damaged and lost frames
 - 2. Damaged and lost ACKs

PAR Protocol

- Tools and issues:
 - Timers
 - Sequence numbers
 - Duplicate frames



Stop-and-Wait [with errors]



In parts (a) and (b) transmitting station A acts the same way, but part (b) receiving station B accepts frame 1 twice.



Protocol 3 Positive ACK with Retransmission (PAR) [Old Tanenbaum Version]

```
#define MAX SEQ 1
typedef enum {frame_arrival, cksum_err, timeout} event_type;
include "protocol.h"
void sender_par (void)
 seq_nr next_frame_to_send;
 frame s:
 packet buffer;
 event_type event;
 next frame to send = 0;
 from_network_layer (&buffer);
 while (true)
 { s.info = buffer;
   s.seq = next_frame_to_send;
   to_physical_layer (&s);
   start_timer (s.seq);
   wait_for_event (&event);
   if (event == frame_arrival) {
      from_network_layer (&buffer);
      inc (next_frame_to_send); }
```



Protocol 3 Positive ACK with Retransmission (PAR) [Old Tanenbaum Version]

```
void receiver_par (void)
seq_nr next_frame_to_send;
frame r, s;
event_type event;
frame_expected = 0;
while (true)
{ wait_for_event (&event);
  if (event == frame arrival)
  { from_physical_layer (&r);
    if (r.seq == frame_expected) {
       to_network_layer(&r.info);
       inc (frame_expected);
                                   /* Note - no sequence number on ACK */
        to_physical_layer (&s);
```



PAR [OLD] problem

Ambiguities occur when ACKs are not numbered.



Leon-Garcia & Widjaja: *Communication Networks*



WPI	Computer Networks Data	Link Layer	18
Figure 3-12	A Positive Acknowledgeme. protocol.	nt with Retransmission Continued →	
	} }	Code ac	Ideo
Cuannel	}		
	from_network_layer(&buffer); inc(next frame to send);	/* get the next one to send */ /* invert next frame to send */	
Noisy	if (s.ack == next_frame_to_send) { stop_timer(s.ack);	/* turn the timer off */	
	from_physical_layer(&s);	/* get the acknowledgement */	
for a	wait_for_event(&event);	/* frame_arrival, cksum_err, timeout */	
11010001	to_physical_layer(&s); start_timer(s.seg);	/* send it on its way */ /* if answer takes too long, time out */	
Protocol	s.info = buffer; s.seq = next_frame_to_send;	/* construct a frame for transmission */ /* insert sequence number in frame */	
Simplex	from_network_layer(&buffer); while (true) {	/* fetch first packet */	
FAR	{ seq_nr next_frame_to_send; frame s; packet buffer; event_type event; next_frame_to_cond = 0;	/* seq number of next outgoing frame */ /* scratch variable */ /* buffer for an outbound packet */	,
	void sender3(void)		
	#define MAX_SEQ 1 typedef enum {frame_arrival, cksum_err, tim #include "protocol.h"	/* must be 1 for protocol 3 */ eout} event_type;	
	/* Protocol 3 (par) allows unidirectional data flow over an unreliable channel. */		

A Simplex Protocol for a Noisy Channel

```
void receiver3(void)
```

seq_nr frame_expected; frame r, s; event_type event;		
frame_expected = 0;		
while (true) {		
wait for event(&event);	/* possibilities: frame arrival, cksum err */	
if (event == frame arrival) {	/* a valid frame has arrived. */	
from physical laver(&r):	/* go get the newly arrived frame */	
if (r seq frame expected) {	/* this is what we have been waiting for */	
to notwork lover(&r info):	/* this is what we have been waiting for. /	
to_fietwork_layer(@f.inito),		
Inc(trame_expected);	/* next time expect the other sequence hr */	
}		
s.ack = 1 - frame_expected;	/* tell which frame is being acked */ Code	added
to_physical_layer(&s);	/* send acknowledgement */	
}		
}		

Figure 3-12.A Positive Acknowledgement with Retransmission protocol.



State Machine for Stop-and-Wait





Sliding Window Protocols [Tanen]

- Must be able to transmit data in <u>both</u> directions.
- Choices for utilization of the reverse channel:
 - mix DATA frames with ACK frames.
 - Piggyback the ACK
 - Receiver waits for DATA traffic in the opposite direction.
 - Use the ACK field in the frame header to send the sequence number of frame being ACKed.
 - \rightarrow better use of the channel capacity.



- ACKs introduce a new issue how long does receiver wait before sending ONLY an ACK frame?
 - → Now we need an ACKTimer !!
 - → The sender timeout period needs to be set longer.
- The protocol must deal with the premature timeout problem and be "robust" under pathological conditions.



- Each outbound frame must contain a sequence number. With n bits for the sequence number field, maxseq = 2ⁿ - 1 and the numbers range from 0 to maxseq.
- Sliding window :: the sender has a window of frames and maintains a list of consecutive sequence numbers for frames that it is permitted to send without waiting for ACKs.
- The receiver has a window of frames that has space for frames whose sequence numbers are in the range of frame sequence numbers it is permitted to accept.
- Note sending and receiving windows do NOT have to be the same size.
- The windows can be fixed size or dynamically growing and shrinking.



The Host is oblivious to sliding windows and the message order at the transport layer is maintained.

sender's DL window :: holds frames sent but
not yet ACKed.

- new packets from the Host cause the upper edge inside the sender's window to be incremented.
- acknowledged frames from the receiver cause the lower edge inside the sender's window to be incremented.



- All frames in the sender's window must be saved for possible retransmission and we need one timer per frame in the window.
- If the maximum sender window size is B, the sender needs at least B buffers.
- If the sender's window gets full (i.e., it reaches the maximum window size, the protocol must shut off the Host (the network layer) until buffers become available.



receiver's DL window

- Frames received with sequence numbers outside the receiver's window are not accepted.
- The receiver's window size is normally static.
- The set of acceptable sequence numbers is rotated as "acceptable" frames arrive.
- If a receiver's window size = 1, then the protocol only accepts frames in order.

This scheme is referred to as Go Back N.



- Selective Repeat :: receiver's window size > 1.
- . The receiver stores all correct frames within the acceptable window range.
- Either the sender times out and resends the missing frame, or
- Selective repeat receiver sends a NACK frame back the sender.



Choices in ACK Mechanisms

 The ACK sequence number indicates the last frame successfully received.
 OR -

2. ACK sequence number indicates the next_frame the receiver expects to receive.

Both schemes can be strictly individual ACKs or represent cumulative ACKs. Cumulative ACKs is the most common technique used.



/*	Protocol	4	(sliding	window)	is	bidirectional. */
----	----------	---	----------	---------	----	-------------------

#define MAX_SEQ 1 /* must be 1 for protocol 4 */ typedef enum {frame_arrival, cksum_err, timeout} event_type; #include "protocol.h" void protocol4 (void) { .

seq_n	r next_trame_to_send;	/* 0 or 1 only */			
seq_n	r frame_expected;	/* 0 or 1 onlý */			
frame	r, s;	/* scratch variables */			
packet	buffer;	<pre>/* current packet being sent */</pre>			
event_	type event;				
next_fr	rame_to_send = 0;	/* next frame on the outbound stream */			
frame_	expected = 0;	/* frame expected next */			
from_r	network_layer(&buffer);	/* fetch a packet from the network layer */			
s.info =	= buffer;	/* prepare to send the initial frame */			
s.seq =	= next_frame_to_send;	/* insert sequence number into frame */			
s.ack =	= 1 – frame_expected;	/* piggybacked ack */			
to_phy	sical_layer(&s);	/* transmit the frame */			
start_ti	mer(s.seq);	/* start the timer running */			
while (true) {				
wa	it for event(&event):	/* frame_arrival. cksum_err. or timeout */			
if (event == frame arrival) {	/* a frame has arrived undamaged. */			
	from_physical_layer(&r);	/* go get it */			
	<pre>if (r.seq == frame_expected) { to_network_layer(&r.info); inc(frame_expected); }</pre>	/* handle inbound frame stream. */ /* pass packet to network layer */ /* invert seq number expected next */			
	<pre>if (r.ack == next_frame_to_send) { stop_timer(r.ack); from_network_layer(&buffer); inc(next_frame_to_send); }</pre>	/* handle outbound frame stream. */ /* turn the timer off */ /* fetch new pkt from network layer */ /* invert sender's sequence number */			
} s.ir s.a s.a to_ sta }	nfo = buffer; eq = next_frame_to_send; ck = 1 - frame_expected; physical_layer(&s); rt_timer(s.seq);	/* construct outbound frame */ /* insert sequence number into it */ /* seq number of last received frame */ /* transmit a frame */ /* start the timer running */			

One-Bit Sliding Window Protocol



}

Figure 3-14. A 1-bit sliding window protocol.

+5

Go Back N





Go Back N with NAK error recovery





/* Protocol 5 (go back n) allows multiple outstanding frames. The sender may transmit up to MAX_SEQ frames without waiting for an ack. In addition, unlike in the previous protocols, the network layer is not assumed to have a new packet all the time. Instead, the network layer causes a network_layer_ready event when there is a packet to send.

```
#define MAX_SEQ 7 /* should be 2<sup>n</sup> - 1 */
typedef enum {frame_arrival, cksum_err, timeout, network_layer_ready} event.type;
#include "protocol.h"
```

```
static boolean between(seq_nr a, seq_nr b, seq_nr c)
{
/* Return true if a \leq b < c circularly; false otherwise. */
 if (((a \le b) \&\& (b < c)) || ((c < a) \&\& (a \le b)) || ((b < c) \&\& (c < a)))
     return(true);
   else
     return(false);
}
static void send_data(seq_nr frame_nr, seq_nr frame_expected, packet buffer[])
/* Construct and send a data frame. */
                                          /* scratch variable */
 frame s:
 s.info = buffer[frame_nr];
                                          /* insert packet into frame */
 s.seg = frame_nr;
                                          /* insert sequence number into frame */
 s.ack = (frame_expected + MAX_SEQ) % (MAX_SEQ + 1);/* piggyback ack */
                                          /* transmit the frame */
 to_physical_layer(&s);
 start_timer(frame_nr);
                                          /* start the timer running */
}
void protocol5(void)
                                          /* MAX_SEQ > 1; used for outbound stream */
 seq_nr next_frame_to_send;
                                          /* oldest frame as yet unacknowledged */
 seq_nr ack_expected;
                                          /* next frame expected on inbound stream */
 seq_nr frame_expected;
                                          /* scratch variable */
 frame r:
                                          /* buffers for the outbound stream */
 packet buffer[MAX_SEQ + 1];
                                          /* # output buffers currently in use */
 sea_nr nbuffered;
                                          /* used to index into the buffer array */
 seq_nr i;
 event_type event;
                                          /* allow network_layer_ready events */
 enable_network_layer();
                                          /* next ack expected inbound */
 ack\_expected = 0;
                                          /* next frame going out */
 next_frame_to_send = 0;
                                          /* number of frame expected inbound */
 frame_expected = 0;
 nbuffered = 0;
                                          /* initially no packets are buffered */
```

```
while (true) {
 wait_for_event(&event);
                                        /* four possibilities: see event_type above */
 switch(event) {
   case network_layer_ready:
                                        /* the network layer has a packet to send */
        /* Accept, save, and transmit a new frame. */
        from_network_layer(&buffer[next_frame_to_send]); /* fetch new packet */
        nbuffered = nbuffered + 1;
                                        /* expand the sender's window */
        send_data(next_frame_to_send, frame_expected, buffer);/* transmit the frame */
        inc(next_frame_to_send);
                                        /* advance sender's upper window edge */
        break;
   case frame_arrival:
                                       /* a data or control frame has arrived */
        from_physical_laver(&r);
                                       /* get incoming frame from physical layer */
        if (r.seq == frame_expected) {
             /* Frames are accepted only in order. */
            to_network_layer(&r.info); /* pass packet to network layer */
                                    /* advance lower edge of receiver's window */
             inc(frame_expected):
        }
        /* Ack n implies n – 1, n – 2, etc. Check for this. */
       while (between(ack_expected, r.ack, next_frame_to_send)) {
            /* Handle piggybacked ack. */
            nbuffered = nbuffered - 1; /* one frame fewer buffered */
            stop_timer(ack_expected); /* frame arrived intact; stop timer */
            inc(ack_expected);
                                       /* contract sender's window */
       break;
  case cksum_err: break;
                                      /* just ignore bad frames */
  case timeout:
                                      /* trouble; retransmit all outstanding frames */
       next_frame_to_send = ack_expected; /* start retransmitting here */
       for (i = 1; i <= nbuffered; i++) {
            send_data(next_frame_to_send, frame_expected, buffer);/* resend frame */
            inc(next_frame_to_send); /* prepare to send the next one */
       }
}
if (nbuffered < MAX_SEQ)
      enable_network_layer();
                                                   ....
else
      disable_network_layer();
```

}

Selective Repeat with NAK error recovery





Computer Networks Data Link Layer

/* Protocol 6 (selective repeat) accepts frames out of order but passes packets to the network layer in order. Associated with each outstanding frame is a timer. When the timer expires, only that frame is retransmitted, not all the outstanding frames, as in protocol 5. */ #define MAX_SEQ 7 /* should be 2ⁿ - 1 */ #define NR_BUFS ((MAX_SEQ + 1)/2) typedef enum {frame_arrival, cksum_err, timeout, network_layer_ready, ack_timeout} event_type; #include "protocol.h" boolean no_nak = true; /* no nak has been sent yet */ seq. nr oldest_frame = MAX.SEQ + 1; /* initial value is only for the simulator */ static boolean between(seq .nr a, seq_nr b, seq_nr c) /* Same as between in protocol5, but shorter and more obscure. */ return ((a <= b) && (b < c)) II ((c < a) && (a <= b)) II ((b < c) && (c < a)); } static void send_frame(frame_kind fk, seq_nr frame_nr, seq_nr frame_expected, packet buffer[]) /* Construct and send a data, ack, or nak frame. */ frame s; /* scratch variable */ s.kind = fk;/* kind == data, ack, or nak */ if (fk == data) s.info = buffer[frame_nr % NR_BUFS]; $s.seq = frame_nr$: /* only meaningful for data frames */ s.ack = (frame_expected + MAX_SEQ) % (MAX_SEQ + 1); if (fk == nak) no_nak = false; /* one nak per frame, please */ to_physical_layer(&s); /* transmit the frame */ if (fk == data) start_timer(frame_nr % NR_BUFS); stop_ack_timer(); /* no need for separate ack frame */ } void protocol6(void) { seq_nr ack_expected; /* lower edge of sender's window */ seq_nr next_frame_to_send; /* upper edge of sender's window + 1 */ seq_nr frame_expected; /* lower edge of receiver's window */ seq_nr too_far; /* upper edge of receiver's window + 1 */ int i: /* index into buffer pool */ frame r: /* scratch variable */ packet out_buf[NR_BUFS]; /* buffers for the outbound stream */ packet in_buf[NR_BUFS]; /* buffers for the inbound stream */ boolean arrived[NR_BUFS]; /* inbound bit map */ seq_nr nbuffered; /* how many output buffers currently used */ event_type event; enable_network_layer(); /* initialize */ ack_expected = 0: /* next ack expected on the inbound stream */ next_frame_to_send = 0; /* number of next outgoing frame */ $frame_expected = 0;$ too_far = NR_BUFS: nbuffered = 0: /* initially no packets are buffered */ for (i = 0; i < NR_BUFS; i++) arrived[i] = false;

```
while (true) {
                                              /* five possibilities: see event_type above */
 wait_for_event(&event);
 switch(event) {
                                              /* accept, save, and transmit a new frame */
   case network_layer_ready:
                                              /* expand the window */
        nbuffered = nbuffered + 1;
        from_network_layer(&out_buf[next_frame_to_send % NR_BUFS]); /* fetch new packet */
        send_frame(data, next_frame_to_send, frame_expected, out_buf);/* transmit the frame */
                                              /* advance upper window edge */
        inc(next_frame_to_send);
        break;
                                              /* a data or control frame has arrived */
   case frame_arrival:
                                              /* fetch incoming frame from physical layer */
        from_physical_layer(&r);
        if (r.kind == data) {
             /* An undamaged frame has arrived. */
             if ((r.seq != frame_expected) && no_nak)
               send_frame(nak, 0, frame_expected, out_buf); else start_ack_timer();
             if (between(frame_expected,r.seq,too_far) && (arrived[r.seq%NR_BUFS]==false)) {
                  /* Frames may be accepted in any order. */
                                                       /* mark buffer as full */
                  arrived[r.seq % NR_BUFS] = true;
                                                       /* insert data into buffer */
                  in_buf[r.seq % NR_BUFS] = r.info;
                  while (arrived[frame_expected % NR_BUFS]) {
                       /* Pass frames and advance window. */
                       to_network_layer(&in_buf[frame_expected % NR_BUFS]);
                       no_nak = true;
                       arrived[frame_expected % NR_BUFS] = false;
                                              /* advance lower edge of receiver's window */
                       inc(frame_expected);
                                              /* advance upper edge of receiver's window */
                       inc(too_far):
                                              /* to see if a separate ack is needed */
                       start_ack_timer();
                  }
             }
        if((r.kind==nak) && between(ack_expected,(r.ack+1)%(MAX_SEQ+1),next_frame_to_send))
             send_frame(data, (r.ack+1) % (MAX_SEQ + 1), frame_expected, out_buf);
        while (between(ack_expected, r.ack, next_frame_to_send)) {
                                              /* handle piggybacked ack */
             nbuffered = nbuffered - 1;
             stop_timer(ack_expected % NR_BUFS); /* frame arrived intact */
                                              /* advance lower edge of sender's window */
             inc(ack_expected);
         break;
    case cksum_err:
        if (no_nak) send_frame(nak, 0, frame_expected, out_buf); /* damaged frame */
        break;
    case timeout:
        send_frame(data, oldest_frame, frame_expected, out_buf); /* we timed out */
        break;
    case ack_timeout:
                                                       /* ack timer expired; send ack */
        send_frame(ack,0,frame_expected, out_buf);
  }
  if (nbuffered < NR_BUFS) enable_network_layer(); else disable_network_layer();
```

ł

}

Data Link Layer Summary

- Parallelism between Transport and Data Link Layer
- Tanenbaum's Treatment/Model of Data Link Layer
- Protocol 1: Utopia
- . Protocol 2: Stop-and-Wait
- Protocol 3: Positive Acknowledgment with Retransmission [PAR]
 - Old 'flawed version
 - Newer version



DL Layer Summary Outline (cont)

- Pipelining and Sliding Windows
- . Protocol 4: One Bit Sliding Window
- Protocol 5: Go Back N
- Protocol 6: Selective Repeat
- Further Details and Decisions

