


UDP. Comer Ch 12, p 197

User Datagram Protocol.

Ans.

UDP Packets are carried inside IP packets. (IPv4 or IPv6).



(ethernet carrying IPv4 carrying UDP)


Frame Type = 0800 = 0000 1000 0000 0000
 PROT = 17 = 0001 0001 .

Next: Proto Port Number.

What do ethernet Frame Type
 IPv4 Protocol Identifier
 UDP port number

have in common?

each points at next higher protocol.



Paradigm Shift:

The customer is a process.

This is process-to-process communication.

The UDP packet (H & F) does not change from end to end!

Remember: The IP header does change:

- (1) TTL --
- (2) Fragmentation if necessary.
- (3) ~~Recomputing~~ Check Options!
- (4) Recomputing Checksum.

UDP needs a method to identify the application process.

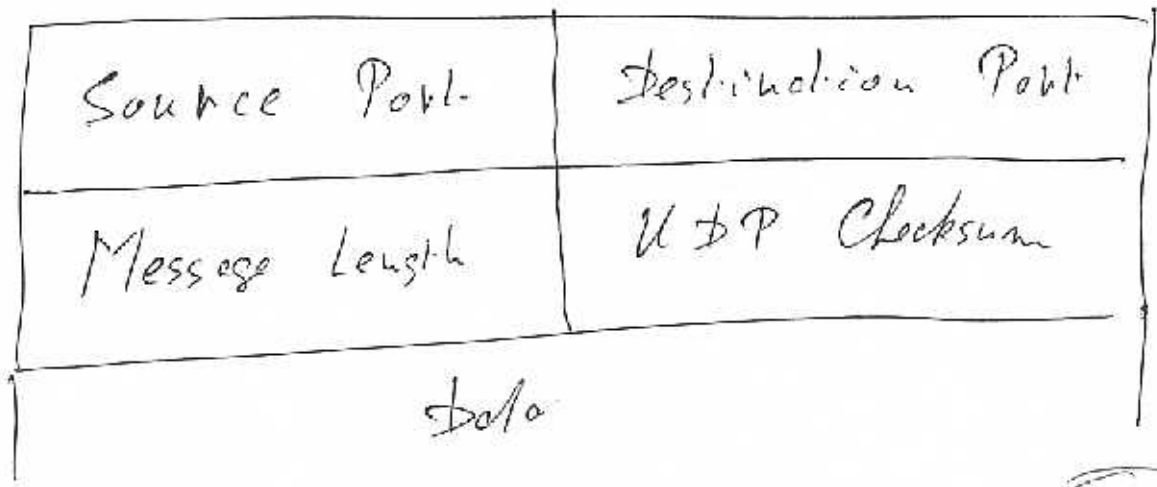
(In some sense: the socket. That is not part of the ~~stand~~ standard.)

How is the customer identified?

Port Number.

The source must know the port number of the destination.

Format of UDP packet.
Cormer p. 199



Ports: 16 bits each.
(0 - 65535).

NO
 UDP
 Options

Message Length:
H & D, in bytes.

Checksum: over Pseudo Header (later)
and data.

(unlike IP).
Use of UDP ~~header~~ checksum is optional.
zero if unused.

Checksum:

if used it is never zero!

Port numbers: identify processes.

(in most or all OSs: there is a
socket to help).

Port numbers:

0 - 1023	"Well known"	port numbers
	IANA	
1024 - 49,151	Registered	(?)(IANA?)
49,152 - 65535	Dynamic Ephemeral	

A process may "listen at a port".

Sometimes: "take any corner".

Sometimes: only if from specific
IP address / port number.


HTTP listens at port 80.
(mostly TCP, not UDP: later).

"Time of Day" listens at
UDP port 13.
(date).

Cornet p. 205 has list of
port numbers:
well-known.

What do ethernet Frame Type
IPv4 Protocol
UDP port number
have in common?

Client-Server Model.



- A. Client, Server are processes.
- B. Client, Server must already exist.
- C. Client process must know:
- (1) Its's local IP address.
 - (2) Server's IP address.
 - (3) Server's Port number.
- (thus for: UDP port number).



Client needs service from server.
(say data).

(1) Client asks its OS for a
UDP port number.
("next available").

(2) Client makes "data gram".

(3) Client sends to UDP:
Appl. data gram. (becomes UDP data).
+ Client IP address
+ Client port number
+ Server IP address
+ Server port number.

4. UDP makes "UDP datagram".




↓
later.

⑤ UDP gives IP:

- UDP datagram
- + Source IP address (yes!) (why?)
- + Dest IP address

6. IP makes IP datagram.



Protocol number in IP header is 17.

7. At this point fragmentation
may (already) occur!

Who determines size of ~~the~~
UDP datagram?

App! !

If the first network has small
MTU: fragmentation occurs!


8. IP datagram
(or fragments)
is/are sent.

Fragmentation:

How many IP ~~headers~~ headers?

How many UDP headers?

In case of LAN = ethernet



(LAN.H & LAN.Tv: LAN WRAPPER).

LAN Wrapper: completely new
every Hop

IP.H : moderate changes, every Hop

UDP.H : No changes!

Unchanged etc.

9. IP datagram
or fragments,
arrive at dest.

IP: re-assemble (if necessary).

Split off UDP datagram.

(10) IP hands to UDP:

UDP datagram.

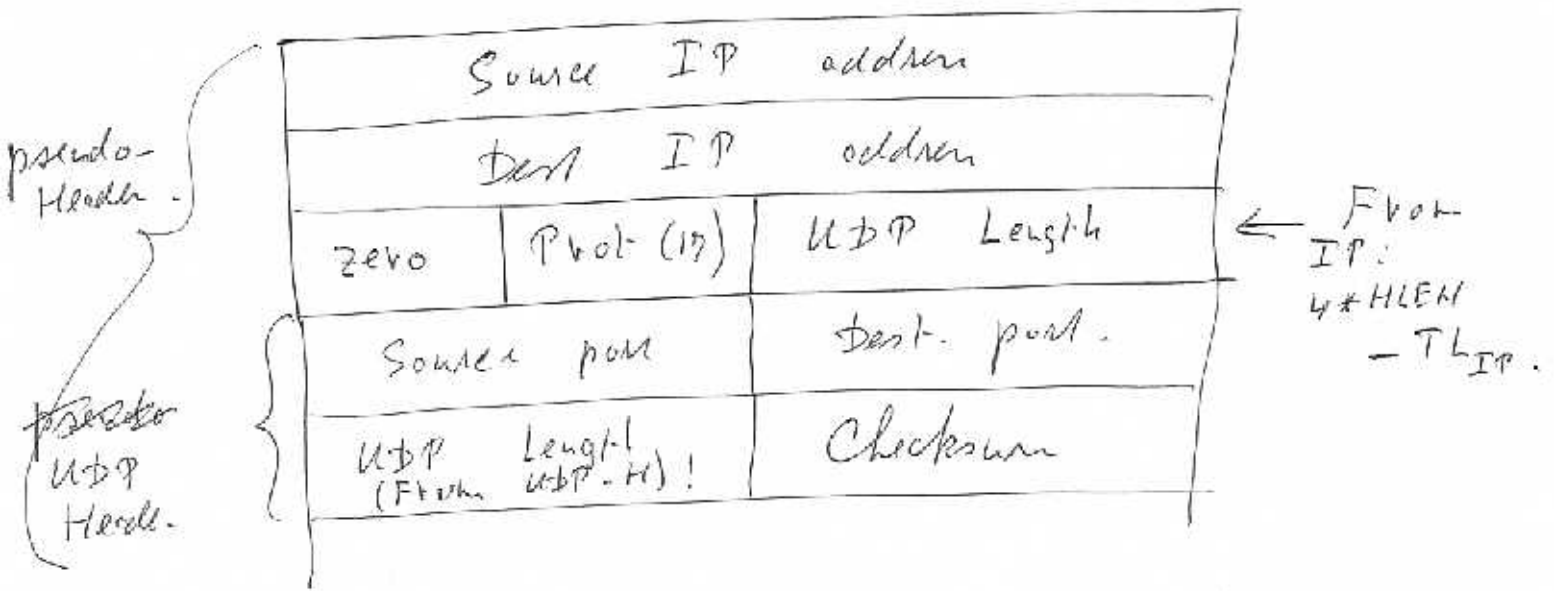
+ Client IP address

+ Server IP address (yes!)

+ UDP Total length (from IP Header!
 $HLEN * 4 - TL_{IP}$.)

(11) UDP ~~data~~ constructs
pseudo-header:

Pseudo Header.



(UDP Length: 2*! from different source)

The source did the same thing!
 (check: it led all the data to make the pseudo-header).

The source (optionally) computed checksum

The dest (if checksum != zero) co checks checksum.

If UDP data changed
 UDP header changed,
 Port changed (TCP → UDP on so)
 Length changed,
 IP address changed.

Check sum fails.
 Packet data given drops.


(12) If all OK:

UDP hands to Appl. (Server) the
~~data~~ Appl. data stream.

plus	Client	IP	address
	Client	Port	number
	Server	IP	address
	Server	Port	number.


If the server wants to send back:
 it has all that is needed.

Between UDP and "customer" (process)
 there is a queue



In Unix, Windows (also Java?)
 there is an entity called socket that
 "helps" in between TCP-UDP and the
 application processes.
Not part of standard.

UDP as multiplex or



UDP : Does not keep state
(per "flows").
(complete Alzheimer).

UDP : connectionless.
(Socket exists under control of
process).

- No Set-up
- No Tear-down
- No Connection
- No Error Correction
- No Re-transmission
(in case of loss, damage, ...)
- No Flow Control
(What is flow control?)
- No Overload Control.
(What is overload control?).

On the other hand!

139.
140

No Overhead.

UDP used for:

(1) RPC. Remote Procedure Call.

(2) Any time "Customer" (Process)

has its own
Flow Control
Error Control
Re-transmission
:

When is not worth the trouble?

Real Time Traffic.

VoIP Voice over IP.

Start on TCP.

Comer Ch. 13, p 269.

UDP

"No State" (per Flow).

Connectionless

No Set-up

No Tear-down

No Connection

No Error Correction

No Re Transmission
(in case of loss,
damage).

No Flow Control
(what is flow control?)

No Overload Control
(what is overload control?)

Data gram Oriented

TCP.

lots of "State".

Connection oriented

Set-up (SYN)

Tear-down (FIN)

Connection!

Retransmission

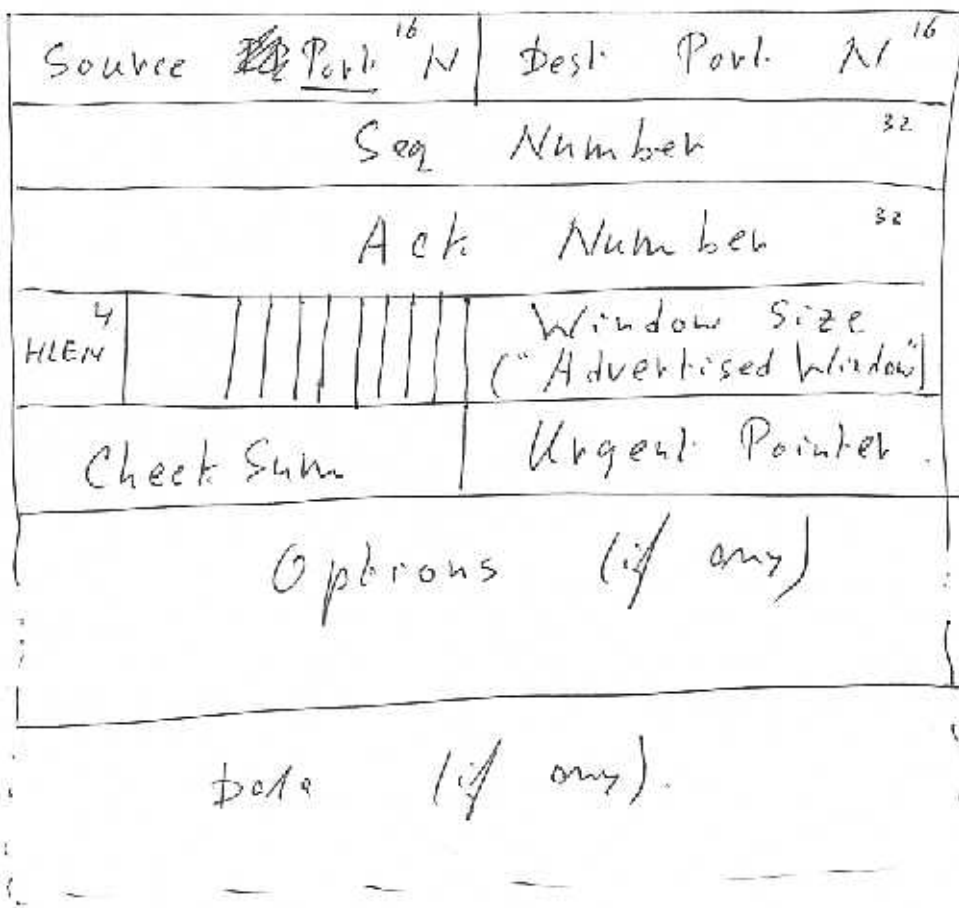
Flow Control

Overload Control

Byte- & Oriented


Byte-Stream
Oriented

TCP Header.



Control Flags
used to be 6.
now 8

IP data



TCP is byte oriented.

Bi-directional.

There is connection set-up,
tear-down.

The Bytes are numbered :

init + 1, init + 2, init + 3, ...

32 bit sequence number.

Repeats after $2^{32} = 4,294,967,296$.

Sequence number: the number of the first ^{TCP} data byte in the packet.

(or: the if there are no data bytes: the highest previous number).

Ack number: "I have received from you, contiguously, everything up to Ack - 1, but not Ack (next expected).

There may be received more, ~~the~~ higher numbers: gap.


Window WIN
 ("Advertised Window").

"You can send me bytes up to
 ACK + WIN.


(Flow Control!)

Receive ~~Window~~ ^{BUFFER}:

Free: WIN.




/// to appl:




It is legal to send in quick succession two packets, no data, some ACK, but increasing WIN.

Reliable



↓ receive packet
 "out of sequence"
 (previous one lost?)



Another ACK (this one compulsory)

Some	Ack number	} "duplicate acknowledgement"
Some	but I IV	
No	data	

Mean: } got something, past Next Expected,
 not what I wanted.

One of the ways to achieve
 reliability

Duplicate Acknowledgements:

Same Acknum

Same W I N

No data.

~~After~~ Count dup acks.

TCP Reno:

(source) After third duplicate acknowledgement,
an entire packet lost,

Re-transmit,

and do more (overload control)

When acknum increases:
reset counter to zero.


FS

Flags

There used to be 6 Control Flags.

In about 2001 more were added,
for ECN. Explicit Congestion
Notification.

RFCs 3168 etc.,
(Also RFC 3540)



CWR: Congestion Window Reduced
ECE: ECN Echo

URG: The urgent pointer is valid.

URGENT pointer: ^{Interrupt application} Relative to Seq.

ACK: Ack is valid.

("there has been previous data from you")
Always always = 1, unless ...

PSH: Push ("do not wait for full mess").

RST: Reset. (?)

SYN: This is a SYN packet

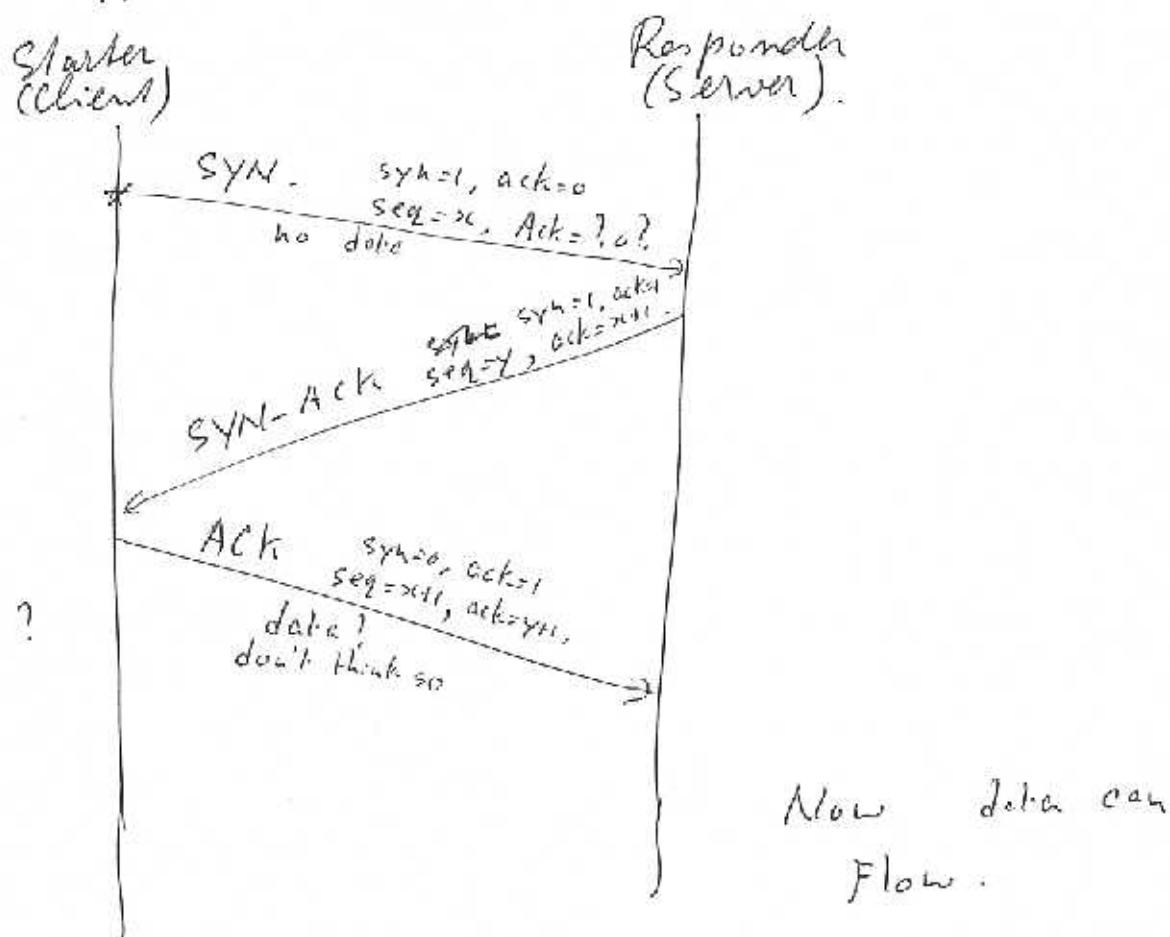
FIN: This is a FIN packet.

HLLEN: TCP header length,
in 32 bit words.

ECN: experimental standard.

Connection Set-up:

Three Way Handshake.



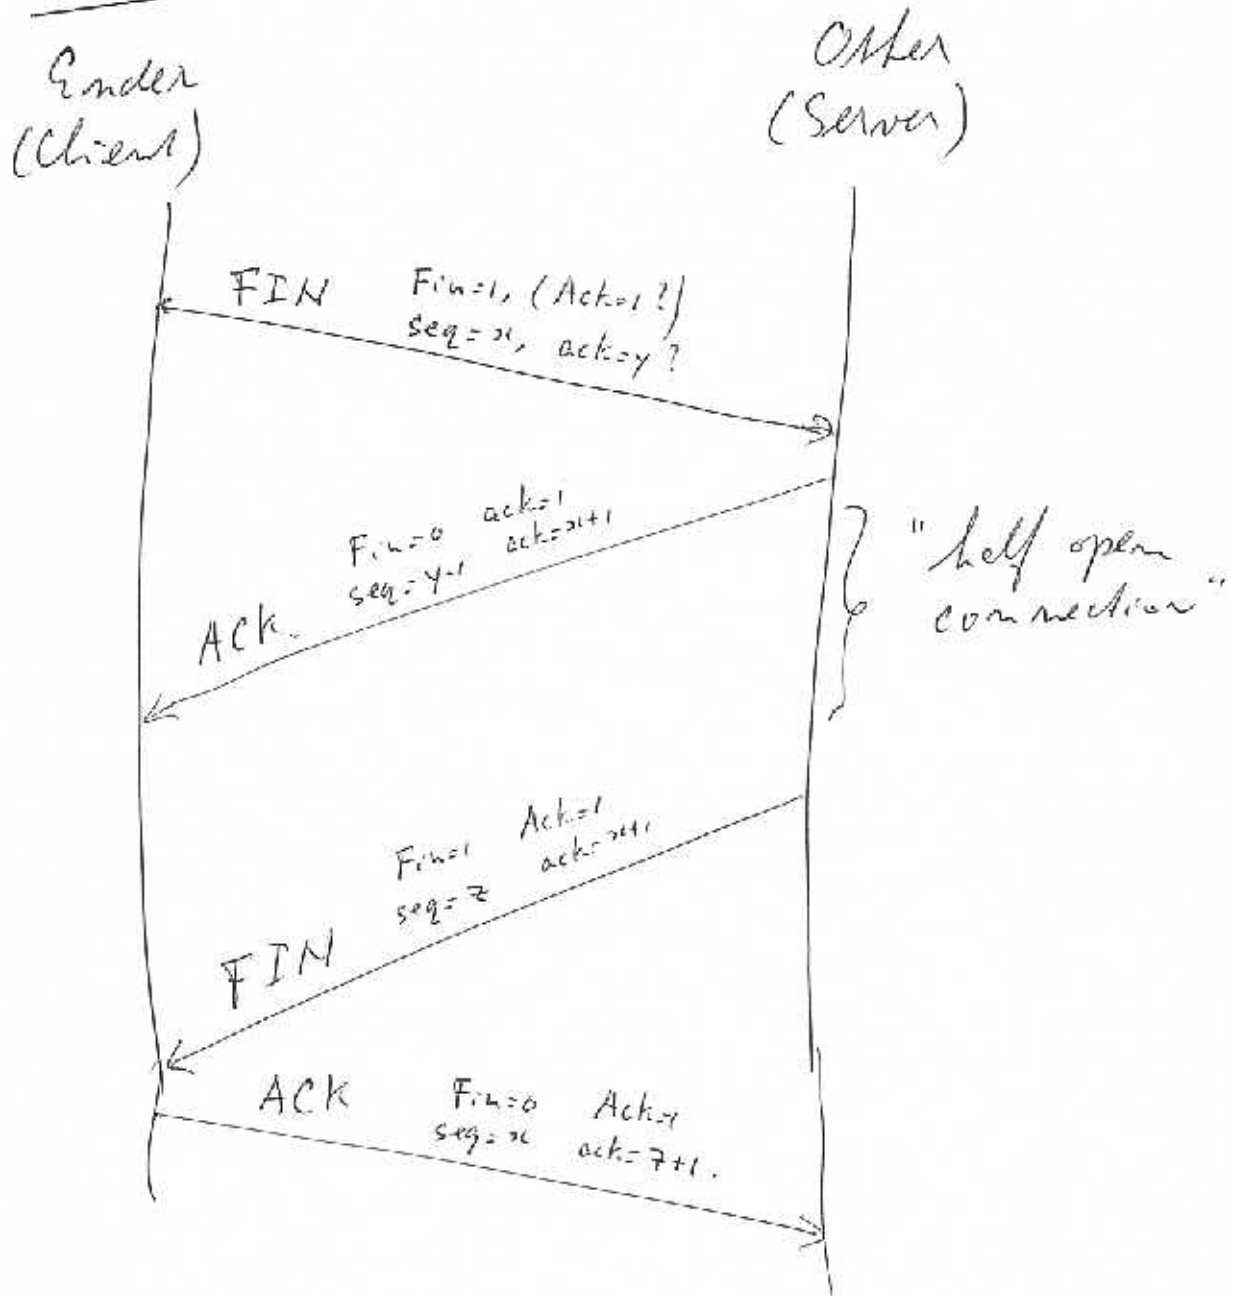
This set-up is Robust under:

Loss of packets.

Two parties starting at the "same same" time, etc.

Connection Tear-down.


Four way handshake -




However, in case of "no half open connection"
I have seen a different pattern!

TCP options -
Comer p. 223 etc.


(1) End of Options
E.O.O



(2) No-Op.



(3) MSS, Maximum Segment Size.



The next options are not in Comer.
You have to know them "temporarily",
Because they occur in the tcpdump data
you will look at.

MSS: "I am willing to accept from you
TCP packets with up to MSS data bytes".

Only in SYN, SYN-ACK packets.

Not necessarily symmetrical!
No negotiation!

5. Time Stamp.

(TCP Time Stamp).

kind
Code = 8

L=10

00001000	00001010
Time Stamp Value	
Time Stamp Echo Reply.	

Time in msec. since last
Universal Time midnight.

Could be

No. Op	No. Op	

Something Strange in tcpdump.

with NTP

⑥

SACK.

Selective Acknowledgement.

RFC 2018


RFC 2883.

Two components:

- (1) In SYN: "I am SACK capable".
- (2) Later: used to describe what data is missing.

First (2).

Suppose the receiver is in the situation (receive buffer)



Old way: Send dup ack: "N.E".

No details of what did, did not, arrive.

Now there is the TCP-SACK option

Code = 5 (kind = 5)	Length in <u>this</u> case
0000 0101	L = 18
1L	
1R	
2L	
2R	

N.E. (next expected) is in the acknowledgement field.

1L, ..., 2R : absolute numbers
(sequence numbers).

How many gaps can this describe?

Say n gaps: ^{2n entries.}
(HLEN ≤ 15 = (1111))
 $2 * n + 1 + 5 \leq 15$ $2 * n \leq 9$ $n \leq 4.5$


$n \leq 4$

Up to 4 gaps,
if no other options.

Problem: Source Behavior.

"When to Re-Resend?"

① "I understand SACK" or SYN.
(SYN only).



if both are SACK-capable,
it will be used.

SACK: often seen in tcpdump output
in lab.