

CS 925

Lecture 9

Congestion Control

Tuesday, February 20, 2024

Two Types of Congestion

▶ Receiver Congestion

- receiver is unable to keep up with incoming data
- solved by explicit feedback from receiver to sender

▶ Network Congestion

- nodes or links of the network are overloaded
- **explicit** congestion notification (few technologies)
- **implicit** congestion notification (Internet)

ECN in IP

▶ Last two bits in the IP TOS/DSCP field

From Wireshark:

Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)

0000 00.. = Differentiated Services Codepoint: Default (0)

.... ..00 = **Explicit Congestion Notification**: Not ECN-Capable Transport (0)

▶ ECN bits

- 00 – Non ECN-Capable Transport **Non-ECT**
- 01 or 10 – ECN-Capable Transport **ECT(0)** or **ECT(1)**
- 11 – Congestion Encountered **CE**

TCP session management

| Offsets Octet | | TCP Header | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|---------------|-----|---|----------|----|---|---|---|---|---|---|---|----|----|----|----|----|----|-----------------------------|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|---|---|---|-------------|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|
| Octet | Bit | 0 | | | | | | | | 1 | | | | | | | | 2 | | | | | | | | 3 | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Octet | Bit | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31 | | | | | | | | | | | | | | | | | | | |
| 0 | 0 | Source port | | | | | | | | | | | | | | | | Destination port | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 4 | 32 | Sequence number | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 8 | 64 | Acknowledgment number (if ACK set) | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 12 | 96 | Data offset | Reserved | NS | <table border="1"> <tr> <td>C</td><td>E</td><td>U</td><td>A</td><td>P</td><td>R</td><td>S</td><td>F</td> </tr> <tr> <td>W</td><td>C</td><td>R</td><td>C</td><td>S</td><td>S</td><td>Y</td><td>I</td> </tr> <tr> <td>R</td><td>E</td><td>G</td><td>K</td><td>H</td><td>T</td><td>N</td><td>N</td> </tr> </table> | | | | | | | | C | E | U | A | P | R | S | F | W | C | R | C | S | S | Y | I | R | E | G | K | H | T | N | N | Window Size | | | | | | | | | | | | | | | |
| C | E | U | A | P | R | S | F | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| W | C | R | C | S | S | Y | I | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| R | E | G | K | H | T | N | N | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 16 | 128 | Checksum | | | | | | | | | | | | | | | | Urgent pointer (if URG set) | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 20 | 160 | Options (if Data Offset > 5, padded at the end with "0" bytes if necessary) | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| ... | ... | ... | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

TCP Flags

Another image appropriated from Wikipedia...

ECN in TCP

- ▶ Two formerly reserved bits in the TCP Flags field

From Wireshark: Flags: 0x018 (PSH, ACK)

000. = Reserved: Not set

...0 = Nonce: Not set

.... 0... = Congestion Window Reduced (CWR): Not set

.... .0.. = ECN-Echo: Not set

.... ..0. = Urgent: Not set

.... ...1 = Acknowledgment: Set

.... 1... = Push: Set

....0.. = Reset: Not set

....0. = Syn: Not set

....0 = Fin: Not set

[TCP Flags:AP...]

- ▶ **ECN-Echo (ECE)** - receiver echoes back to sender received congestion indication (CE in IP)
- ▶ **Congestion Window Reduced (CWR)** - sender's acknowledgment of received ECE (and reduced Tx rate)

ECN

- ▶ ECN is not a solution in itself, it is apart of a larger mechanism
- ▶ It needs:
 - a mechanism in routers to detect congestion
 - a method to adjust sending rate in response to congestion indication