A final proposal related to TCP congestion control concerns UDP. Observe that although TCP reduces transmission when congestion occurs, UDP does not, which means that as TCP flows continue to back off, UDP flows consume more of the bandwidth. A solution known as *TCP Friendly Rate Control* (*TFRC*) was proposed. TFRC attempts to emulate TCP behavior by having a UDP receiver report datagram loss back to the sender and by having the sender use the reported loss to compute a rate at which UDP datagrams should be sent; TFRC has only been adopted for special cases.

11.21 Explicit Feedback Mechanisms (SACK and ECN)

Most versions of TCP use *implicit* techniques to detect loss and congestion. That is, TCP uses timeout and duplicate ACKs to detect loss, and changes in round trip times to detect congestion. Researchers have observed that slight improvements are possible if TCP includes mechanisms that provide such information *explicitly*. The next two sections describe two explicit techniques that have been proposed.

11.21.1 Selective Acknowledgement (SACK)

The alternative to TCP's cumulative acknowledgement mechanism is known as a *selective acknowledgement* mechanism. In essence, selective acknowledgements allow a receiver to specify exactly which data has been received and which is still missing. The chief advantage of selective acknowledgements arises in situations where occasional loss occurs: selective acknowledgements allow a sender to know exactly which segments to retransmit.

The *Selective ACKnowledgement* (*SACK*) mechanism proposed for TCP does not completely replace the cumulative acknowledgement mechanism, nor is it mandatory. Instead, TCP includes two options for SACK. The first option is used when the connection is established to allow a sender to specify that SACK is permitted. The second option is used by a receiver when sending an acknowledgement to include information about specific blocks of data that were received. The information for each block includes the first sequence number in a block (called the *left edge*) and the sequence number immediately beyond the block (called the *right edge*). Because the maximum size of a segment header is fixed, an acknowledgement can contain at most four SACK blocks. Interestingly, the SACK documents do not specify exactly how a sender responds to SACK; most implementations retransmit all missing blocks.

11.21.2 Explicit Congestion Notification

A second technique to avoid implicit measurement is intended to handle congestion in the network. Known as *Explicit Congestion Notification* (*ECN*), the mechanism requires routers throughout an internet to notify TCP as congestion occurs. The mechanism is conceptually straightforward: as a TCP segment passes through the internet, routers along the path use a pair of bits in the IP header to record congestion. Thus,