

In the case of reliable service, both approaches are implemented. In situations where errors are infrequent, end-to-end mechanisms are preferred. The TCP reliable stream service introduced later in this chapter provides an example of such an end-to-end mechanism. In situations where errors are likely, hop-by-hop error recovery becomes necessary. The HDLC data link control is an example of a hop-by-hop mechanism.

The end-to-end versus hop-by-hop option appears in other situations as well. For example, flow control and congestion control can be exercised on a hop-by-hop or an end-to-end basis. Security mechanisms provide another example in which this choice needs to be addressed. The approach selected typically determines which layer of the protocol stack provides the desired function. Thus, for example, mechanisms for providing congestion control and mechanisms for providing security are available at the data link layer, the network layer, and the transport layer.

5.2 ARQ PROTOCOLS

Automatic Repeat Request (ARQ) is a technique used to ensure that a data stream is delivered accurately to the user despite errors that occur during transmission. ARQ forms the basis for peer-to-peer protocols that provide for the reliable transfer of information. In this section we develop the three basic types of ARQ protocols, starting with the simplest and building up to the most complex. We also discuss the settings in which the three ARQ protocol types are applied.

This discussion assumes that the user generates a sequence of information blocks for transmission. The ARQ mechanism requires the block to contain a header with control information that is essential to proper operation, as shown in Figure 5.8. The transmitter will also append CRC check bits that cover the header and the information bits to enable the receiver to determine whether errors have occurred during transmission. We assume that the design of the CRC ensures that transmission errors can be detected with very high probability, as discussed in Chapter 3. Recall from Chapter 2 that we use the term **frame** to refer to the binary block that results from the combination of the header, user information, and CRC at the data link layer. We refer to the set of rules that govern the operation of the transmitter and receiver as the **ARQ protocol**.

The ARQ protocol can be applied in a number of scenarios. Traditionally, the most common application involves transmission over a single noisy communication channel. ARQ is introduced here to ensure a high level of reliability across a single transmission hop. As communication lines have become less noisy, the ARQ protocol has been implemented more often at the edges of the network to provide end-to-end reliability in the transmission of packets over multiple hops in a network, that is, over multiple communication channels and other network equipment. In this section we assume that the channel or

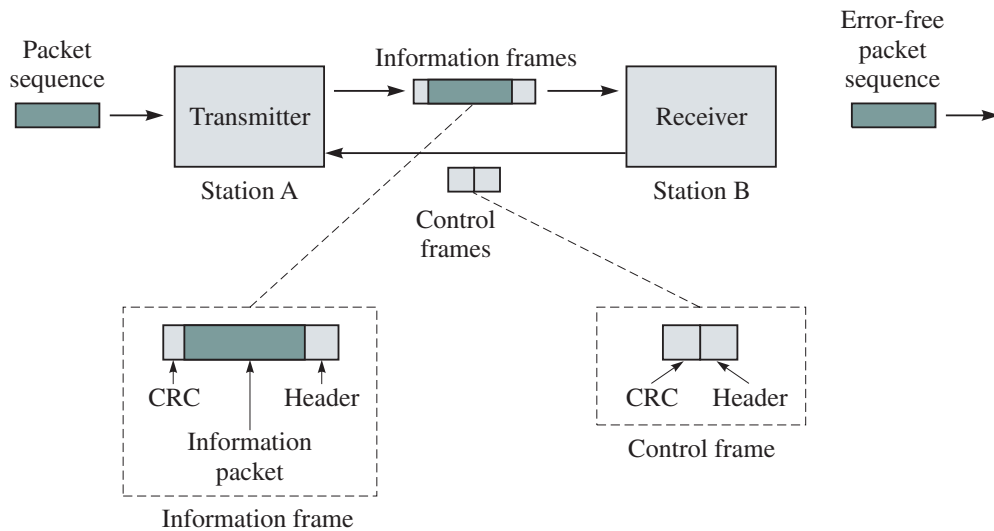


FIGURE 5.8 Basic elements of ARQ

sequence of channels is “wirelike” in the sense that the frames arrive at the receiver, if they arrive at all, in the same order in which they were sent. In the case of multiple hops over a network, this assumption would hold when a connection is set up and where all frames follow the same path, as in frame relay or ATM networks. In particular, we assume that frames, while in transit, cannot pass previously transmitted frames.³ In situations where these assumptions hold, the objective of the ARQ protocol is to ensure that packets are delivered error free to the destination, exactly once without duplicates, in the same order in which they were transmitted.

In addition to the error-detection code, the other basic elements of ARQ protocols consist of **information frames (I-frames)** that transfer the user packets, control frames, and time-out mechanisms, as shown in Figure 5.8. **Control frames** are short binary blocks that consist of a header that provides the control information followed by the CRC. The control frames include **ACKs**, which acknowledge the correct receipt of a given frame or group of frames; **NAKs**, which indicate that a frame has been received in error and that the receiver is taking certain action; and an **enquiry frame ENQ**, which commands the receiver to report its status. The time-out mechanisms are required to prompt certain actions to maintain the flow of frames. We can visualize the transmitter and receiver as working jointly on ensuring the correct and orderly delivery of the sequence of packets provided by the sender.


We begin with the simple case where information flows only in one direction, from the transmitter to the receiver. The reverse communication channel is used only for the transmission of control information. Later in the section we consider the case of bidirectional transfer of information.

³We consider the case where frames can arrive out of order later in this chapter.

5.2.1 Stop-and-Wait ARQ

The first protocol we consider is **Stop-and-Wait ARQ** where the transmitter and receiver work on the delivery of one frame at a time through an alternation of actions. In Figure 5.9a we show how ACKs and time-outs can be used to provide recovery from transmission errors, in this case a lost frame. At the initial point in the figure, stations A and B are working on the transmission of frame 0. Note that each time station A sends an I-frame, it starts an **I-frame timer** that will expire after some time-out period. The time-out period is selected so that it is greater than the time required to receive the corresponding ACK frame. Figure 5.9a shows the following sequence of events:

1. Station A transmits frame 0 and then waits for an ACK frame from the receiver.
2. Frame 0 is transmitted without error, so station B transmits an ACK frame.
3. The ACK from station B is also received without error, so station A knows the frame 0 has been received correctly.
4. Station A now proceeds to transmit frame 1 and then resets the timer.
5. Frame 1 undergoes errors in transmission. It is possible that station B receives frame 1 and detects the errors through the CRC check; it is also possible that frame 1 was so badly garbled that station B is unaware of the transmission.⁴ In either case station B does not take any action.
6. The time-out period expires, and frame 1 is retransmitted.



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

FIGURE 5.9 Possible ambiguities when frames are unnumbered


⁴In general, when errors are detected in a frame, the frame is ignored. The receiver cannot trust any of the data in the frame and, in particular, cannot take any actions based on the contents of the frame header.

The protocol continues in this manner until frame 1 is received and acknowledged. The protocol then proceeds to frame 2, and so on.

Transmission errors in the reverse channel lead to ambiguities in the Stop-and-Wait protocol that need to be corrected. Figure 5.9b shows the situation that begins as in Figure 5.9a, but where frame 1 is received correctly, and its acknowledgment undergoes errors. After receiving frame 1 station B delivers its contents to the destination. Station A does not receive the acknowledgment for frame 1, so the time-out period expires. Note that at this point station A cannot distinguish between the sequence of events in parts (a) and (b) of Figure 5.9. Station A proceeds to retransmit the frame. If the frame is received correctly by station B, as shown in the figure, then station B will accept frame 1 as a new frame and redeliver it to the user. Thus we see that the loss of an ACK can result in the delivery of a duplicate packet. The ambiguity can be eliminated by including a sequence number in the header of each I-frame. Station B would then recognize that the second transmission of frame 1 was a duplicate, discard the frame, and resend the ACK for frame 1.

A second type of ambiguity arises if the ACKs do not contain a sequence number. In Figure 5.10 frame 0 is transmitted, but the time-out expires prematurely. Frame 0 is received correctly, and the (unnumbered) ACK is returned. In the meantime station A has resent frame 0. Shortly thereafter, station A receives an ACK and assumes it is for the last frame. Station A then proceeds to send frame 1, which incurs transmission errors. In the meantime the second transmission of frame 0 has been received and acknowledged by station B. When station A receives the second ACK, the station assumes the ACK is for frame 1 and proceeds to transmit frame 2. The mechanism fails because frame 1 is not delivered. This example shows that *premature time-outs (or delayed ACKs) combined with loss of I-frames can result in gaps in the delivered packet sequence*. This ambiguity is resolved by providing a sequence number in the acknowledgment frames that enables the transmitter to determine which frames have been received.

The sequence numbers cannot be allowed to become arbitrarily large because only a finite number of bits are available in the frame headers. We now show that a one-bit sequence number suffices to remove the above ambiguities in the Stop-and-Wait protocol. Figure 5.11 shows the information or “state” that is main-



Transmitting station A misinterprets duplicate ACKs

FIGURE 5.10 Possible ambiguities when ACKs are unnumbered

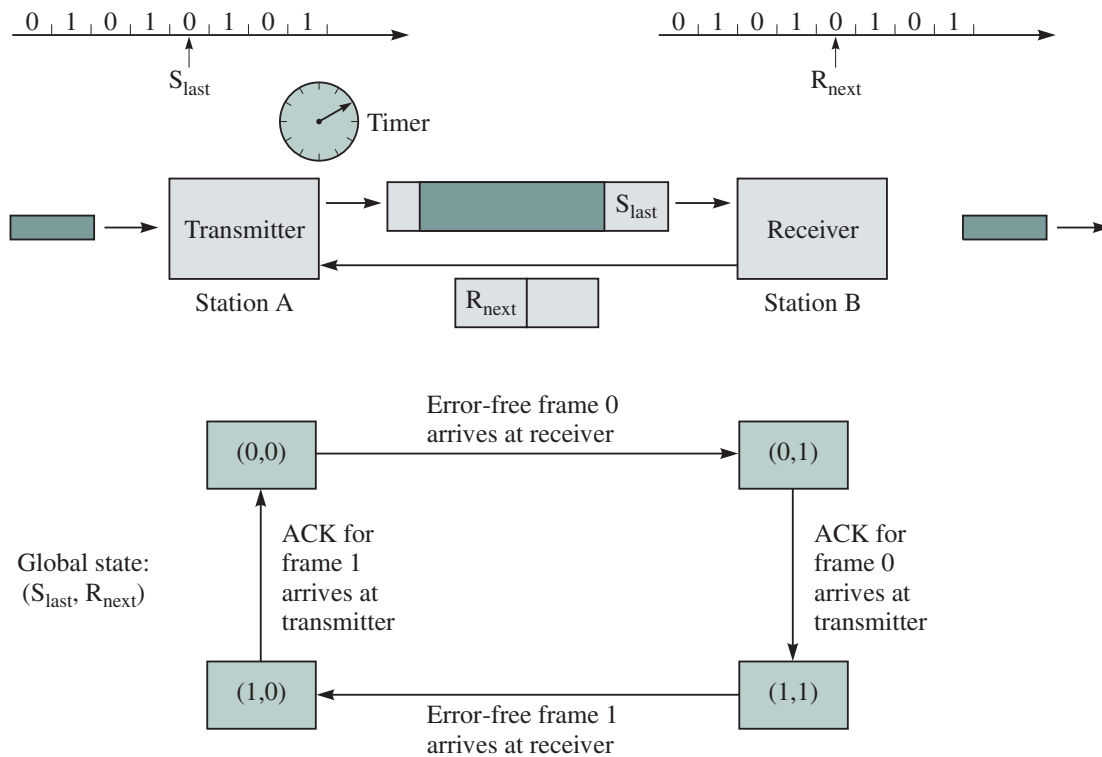


FIGURE 5.11 System state information in Stop-and-Wait ARQ

tained by the transmitter and receiver. The transmitter must keep track of the sequence number S_{last} of the frame being sent, its associated timer, and the frame itself in case retransmission is required. The receiver keeps track only of the sequence number R_{next} of the next frame it is expecting to receive.

Suppose that initially the transmitter and receiver are synchronized in the sense that station A is about to send a frame with $S_{last} = 0$ and station B is expecting $R_{next} = 0$. In Figure 5.11 the global state of the system is defined by the pair (S_{last}, R_{next}) , so initially the system is in state $(0,0)$.

The system state will not change until station B receives an error-free version of frame 0. That is, station A will continue resending frame 0 as dictated by the time-out mechanism. Eventually station B receives frame 0, station B changes R_{next} to 1 and sends an acknowledgment to station A with $R_{next} = 1$ implicitly acknowledging the receipt of frame 0. At this point the state of the system is $(0,1)$. Any subsequent received frames that have sequence number 0 are recognized as duplicates and discarded by station B, and an acknowledgment with $R_{next} = 1$ is resent. Eventually station A receives an acknowledgment with $R_{next} = 1$ and then begins transmitting the next frame, using sequence number $S_{last} = 1$. The system is now in state $(1,1)$. The transmitter and receiver are again synchronized, and they now proceed to work together on the transfer of frame 1. Therefore, a protocol that implements this mechanism that follows the well-defined sequence of states shown in Figure 5.11 can ensure the correct and orderly delivery of frames to the destination.

A number of modifications can be made to the above ARQ mechanism by the use of additional control frames. For example, in **checkpointing** the error recovery can be expedited through the use of a short control frame called the enquiry frame (ENQ). When a time-out expires, if the frame that needs to be retransmitted is very long, the transmitter can send an ENQ. When a station receives such a frame that station is compelled to retransmit its previous frame. As shown in Figure 5.12, station B then proceeds to retransmit its last ACK frame, which when received by station A resolves the ambiguity. Station A can then proceed with the transmission of the appropriate frame, as shown in the figure. By convention ACK and NAK frames contain R_{next} , indicating the next frame that is expected by the receiver and implicitly acknowledging delivery of all prior frames.

Stop-and-Wait ARQ becomes inefficient when the propagation delay is much greater than the time to transmit a frame. For example, suppose that we are transmitting frames that are 1000 bits long over a channel that has a speed of 1.5 megabits/second and suppose that the time that elapses from the beginning of the frame transmission to the receipt of its acknowledgment is 40 ms. The number of bits that can be transmitted over this channel in 40 ms is $40 \times 10^{-3} \times 1.5 \times 10^6 = 60,000$ bits. However, Stop-and-Wait ARQ can transmit only 1000 bits in this period time. This severe inefficiency is due to the requirement that the transmitter wait for the acknowledgment of a frame before proceeding with other transmissions. Later in the section we show that the situation becomes much worse in the presence of transmission errors that trigger retransmissions.

The **delay-bandwidth product** is the product of the bit rate and the delay that elapses before an action can take place. In the preceding example the delay-

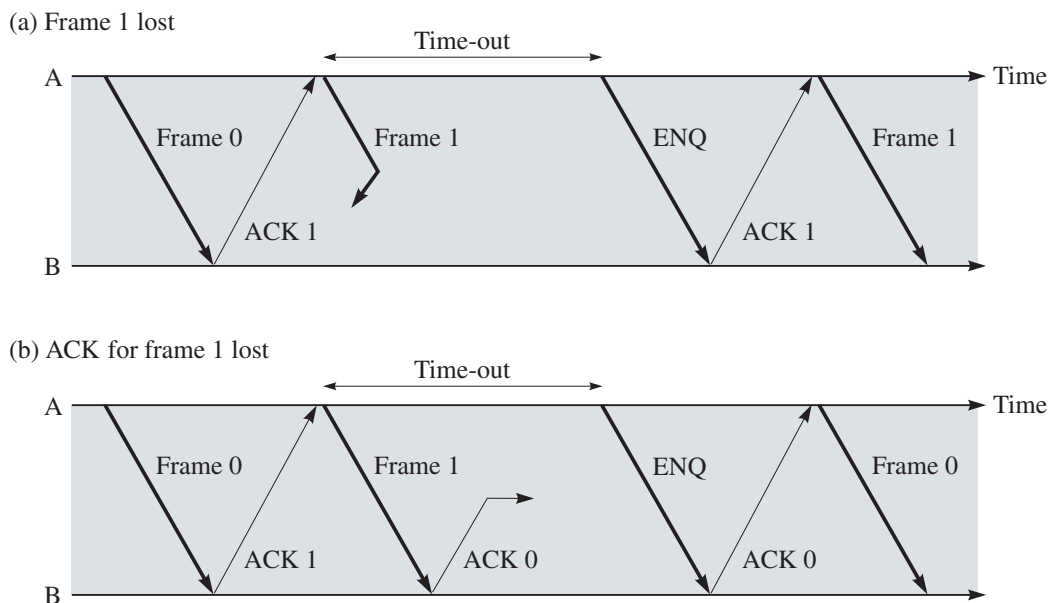


FIGURE 5.12 Stop-and-Wait ARQ enquiry frame

bandwidth product is 60,000 bits. In Stop-and-Wait ARQ the delay-bandwidth product can be viewed as a measure of lost opportunity in terms of transmitted bits. This factor arises as a fundamental limitation in many network problems.

Example—Bisync

Stop-and-wait ARQ was used in IBM's Binary Synchronous Communications (Bisync) protocol. Bisync is a character-oriented data link control that uses the ASCII character set, including several control characters. In Bisync, error detection was provided by two-dimensional block coding where each seven-bit character has a parity bit attached and where an overall block-check character is appended to a frame. Bisync has been replaced by data link controls based on HDLC, which is discussed later in this chapter.

Example—Xmodem

Xmodem, a popular file transfer protocol for modem, incorporates a form of Stop-and-Wait ARQ. Information is transmitted in fixed-length blocks consisting of a 3-byte header, 128 bytes of data, and a one-byte checksum. The header consists of a special start-of-header character, a one-byte sequence number, and a 2s complement of the sequence number. The check character is computed by taking the modulo 2 sum of the 128 data bytes. The receiver transmits an ACK or NAK character after receiving each block.

5.2.2 Go-Back-N ARQ

In this section we show that the inefficiency of Stop-and-Wait ARQ can be overcome by allowing the transmitter to continue sending enough frames so that the channel is kept busy while the transmitter waits for acknowledgments. We develop the Go-Back-N protocol that forms the basis for the HDLC data link protocol, which is discussed at the end of this chapter. Suppose for now that frames are numbered 0, 1, 2, 3, . . . The transmitter has a limit on the number of frames W_S that can be outstanding. W_S is chosen larger than the delay-bandwidth product to ensure that the channel can be kept busy.

The idea of the **Basic Go-Back-N ARQ** is as follows: Consider the transfer of a reference frame, say, frame 0. After frame 0 is sent, the transmitter sends $W_S - 1$ additional frames into the channel, optimistic that frame 0 will be received correctly and not require retransmission. If things turn out as expected, an ACK for frame 0 will arrive in due course while the transmitter is still busy sending frames into the channel, as shown in Figure 5.13. The system is now done with frame 0. Note, however, that the handling of frame 1 and subsequent frames is already well underway. A procedure where the processing of a new task is begun before the completion of the previous task is said to be **pipelined**. In




FIGURE 5.13 Basic Go-Back-N ARQ

effect Go-Back-N ARQ pipelines the processing of frames to keep the channel busy.

Go-Back-N ARQ gets its name from the action that is taken when an error occurs. As shown in Figure 5.13, after frame 3 undergoes transmission errors, the receiver ignores frame 3 and all subsequent frames. Eventually the transmitter reaches the maximum number of outstanding frames. It is then forced to “go back N” frames, where $N = W_S$, and begin retransmitting all packets from 3 onwards.

In the previous discussion of Stop-and-Wait, we used the notion of a global state (S_{lasts} , R_{next}) to demonstrate the correct operation of the protocol in the sense of delivering the packets error free and in order. Figure 5.14 shows the similarities between Go-Back-N ARQ and Stop-and-Wait ARQ in terms of error recovery. In Stop-and-Wait ARQ, the occurrence of a frame-transmission error results in the loss of transmission time equal to the duration of the time-out period. In Go-Back-N ARQ, the occurrence of a frame-transmission error results in the loss of transmission time corresponding to W_S frames. In Stop-and-Wait the receiver is looking for the frame with sequence number R_{next} ; in Go-Back-N the receiver is looking for a frame with a specific sequence number, say, R_{next} . If we identify the oldest outstanding (transmitted but unacknowledged) frame in Go-Back-N with the S_{last} frame in Stop-and-Wait, we see that the correct operation of the Go-Back-N protocol depends on ensuring that the oldest frame is eventually delivered successfully. The protocol will trigger a retransmission of S_{last} and the subsequent $W_S - 1$ frames each time the send window is exhausted. Therefore, as long as there is a nonzero probability of error-free frame transmission, the eventual error-free transmission of S_{last} is assured and the protocol will operate correctly. We next modify the basic protocol we have developed to this point.

The Go-Back-N ARQ as stated above depends on the transmitter exhausting its maximum number of outstanding frames to trigger the retransmission of a frame. Thus this protocol works correctly as long as the transmitter has an unlimited supply of packets that need to be transmitted. In situations where

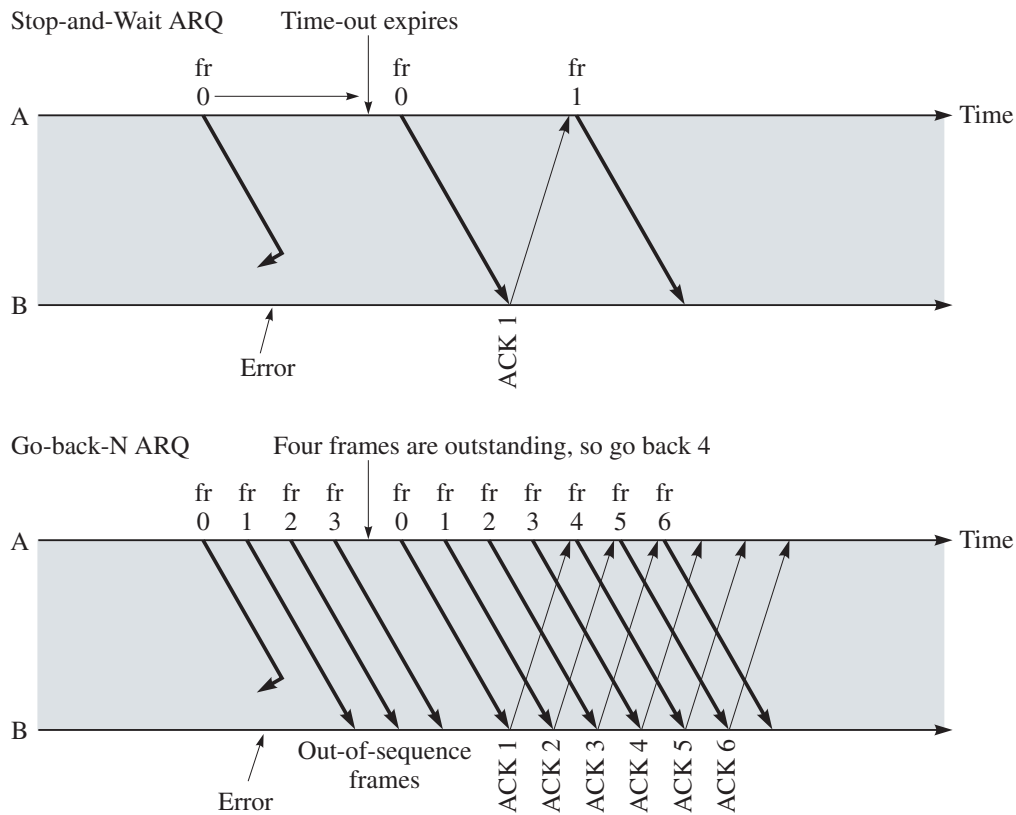


FIGURE 5.14 Relationship of Stop-and-Wait ARQ and Go-Back-N ARQ

packets arrive sporadically, there may not be $W_S - 1$ subsequent transmissions. In this case retransmissions are not triggered, since the window is not exhausted. This problem is easily resolved by modifying Go-Back-N ARQ such that a timer is associated with each transmitted frame.

Figure 5.15 shows how the resulting Go-Back-N ARQ protocol operates. The transmitter must now maintain a list of the frames it is processing, where S_{last} is the number of the last transmitted frame that remains unacknowledged and S_{recent} is the number of the most recently transmitted frame. The transmitter must also maintain a timer for each transmitted frame and must also buffer all frames that have been transmitted but have not yet been acknowledged. At any point in time the transmitter has a **send window** of available sequence numbers. The lower end of the window is given by S_{last} , and the upper limit of the transmitter window is $S_{\text{last}} + W_S - 1$. If S_{recent} reaches the upper limit of the window, the transmitter is not allowed to transmit further new frames until the send window slides forward with the receipt of a new acknowledgment.

The Go-Back-N protocol is an example of a **sliding-window protocol**. The receiver maintains a **receive window** of size 1 that consists of the next frame R_{next} it expects to receive. If an arriving frame passes the CRC check and has the correct sequence number, that is, R_{next} , then it is accepted and R_{next} is incremented. We say that the receive window slides forward. The receiver then sends an acknowledgment containing the incremented sequence number R_{next} , which

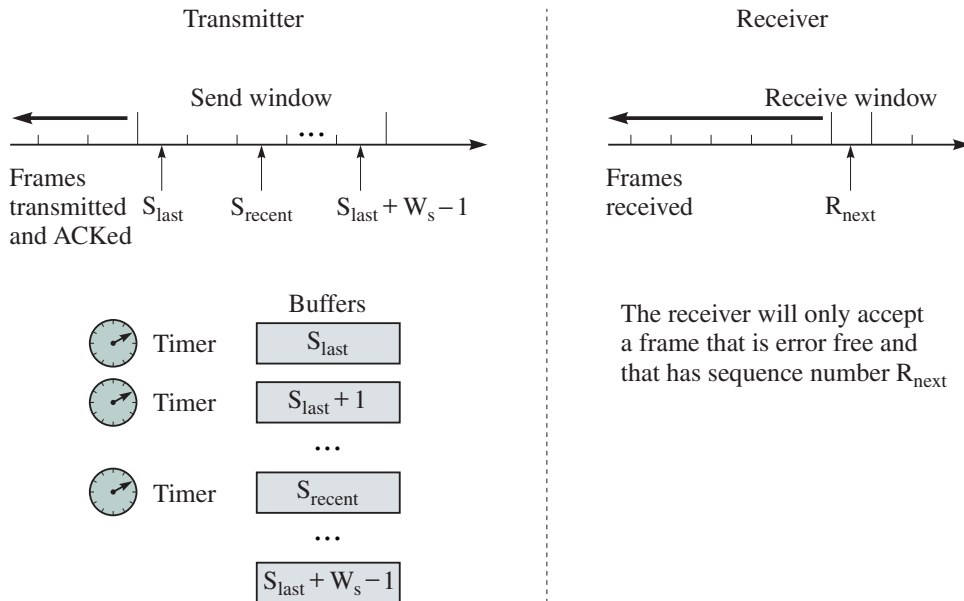


FIGURE 5.15 Go-Back-N ARQ

implicitly acknowledges receipt of all frames prior to R_{next} . Note how we are making use of the assumption that the channel is wirelike: When the transmitter receives an ACK with a given value R_{next} , it can assume that all prior frames have been received correctly, even if it has not received ACKs for those frames, either because they were lost or because the receiver chose not to send them. Upon receiving an ACK with a given value R_{next} , the transmitter updates its value of S_{last} to R_{next} and in so doing the send window slides forward. (Note that this action implies that $S_{last} \leq R_{next}$. Note as well that $R_{next} \leq S_{recent}$, since S_{recent} is the last in the transmission frames.)

The number of bits that can be allotted within a header per sequence number is limited to some number, say, m , which then allows us to represent at most 2^m possible sequence numbers, so the sequence numbers must be counted using modulo 2^m . Thus if $m = 3$, then the sequence of frames would carry the sequence numbers 0, 1, 2, 3, 4, 5, 6, 7, 0, 1, 2, 3, ... When an error-free frame arrives, the receiver must be able to unambiguously determine which frame has been received, taking into account that the sequence numbers wrap around when the count reaches 2^m . We will next show that the receiver can determine the correct frame if the window size is smaller than 2^m .

The example in Figure 5.16 shows the ambiguities that arise when the above inequality is not satisfied. The example uses $2^m = 2^2 = 4$ sequence numbers. The transmitter initially sends four frames in a row. The receiver sends four corresponding acknowledgments, which are all obliterated in the return channel. When the transmitter exhausts its available frame numbers, it goes back four and begins retransmitting from frame 0. When frame 0 arrives at the receiver, the receiver has $R_{next} = 0$, so it accepts the frame. However, the receiver does not know whether the ACK for the previous frame 0 was received. Consequently, the

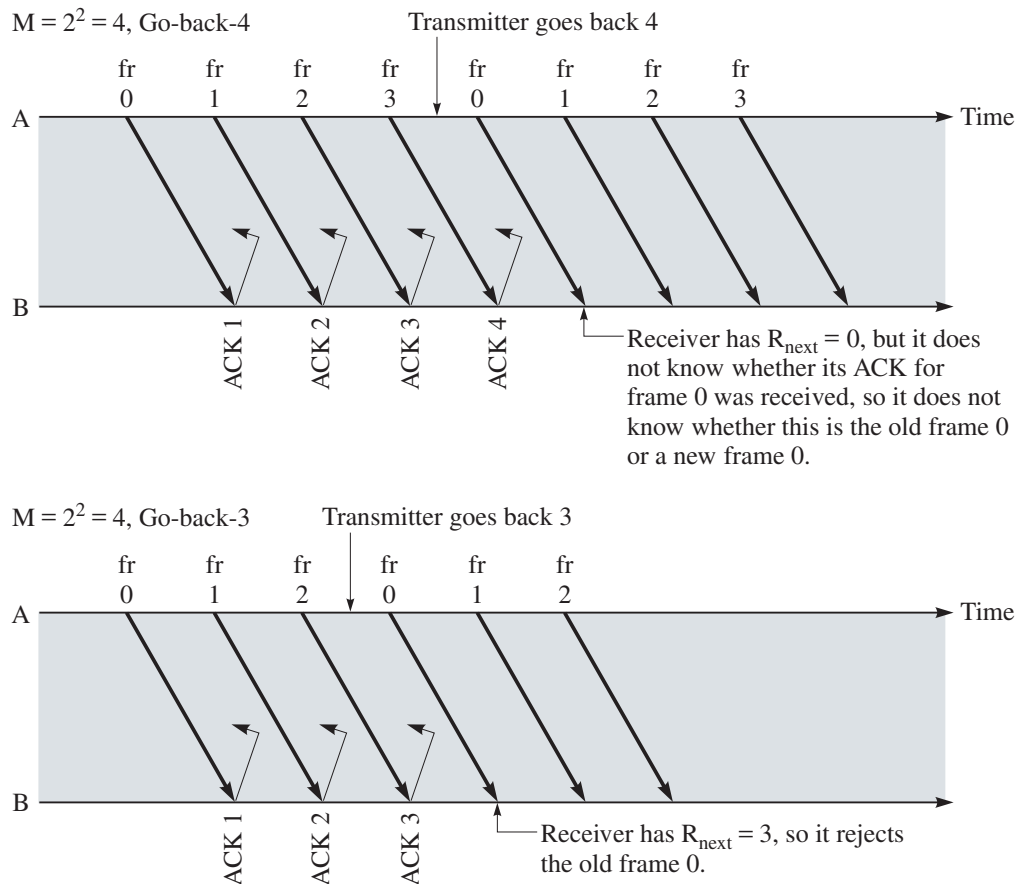


FIGURE 5.16 The window size should be less than 2^m

receiver cannot determine whether this is a new frame 0 or an old frame 0. The second example in Figure 5.16 uses $2^m = 4$ sequence numbering, but a window size of 3. In this case we again suppose that the transmitter sends three consecutive frames and that acknowledgments are lost in the return channel. When the retransmitted frame 0 arrives at the receiver, $R_{\text{next}} = 3$, so the frame is recognized to be an old one.

We can generalize Figure 5.16 as follows. In general, suppose the window size W_S is $2^m - 1$ or less and assume that the current send window is 0 up to $W_S - 1$. Suppose that frame 0 is received, but the acknowledgment for frame 0 is lost. The transmitter can only transmit new frames up to frame $W_S - 1$. Depending on which transmissions arrive without error, R_{next} will be in the range of 1 to W_S . Crucially, the receiver will not receive frame W_S until the acknowledgment for frame 0 has been received at the transmitter. Therefore, any receipt of frame 0 prior to frame W_S indicates a duplicate transmission of frame 0.

Figure 5.17 shows that the performance of Go-Back-N ARQ can be improved by having the receiver send a NAK message immediately after the first out-of-sequence frame is received. The NAK with sequence number R_{next} acknowledges all frames up to $R_{\text{next}} - 1$ and informs the transmitter that an error

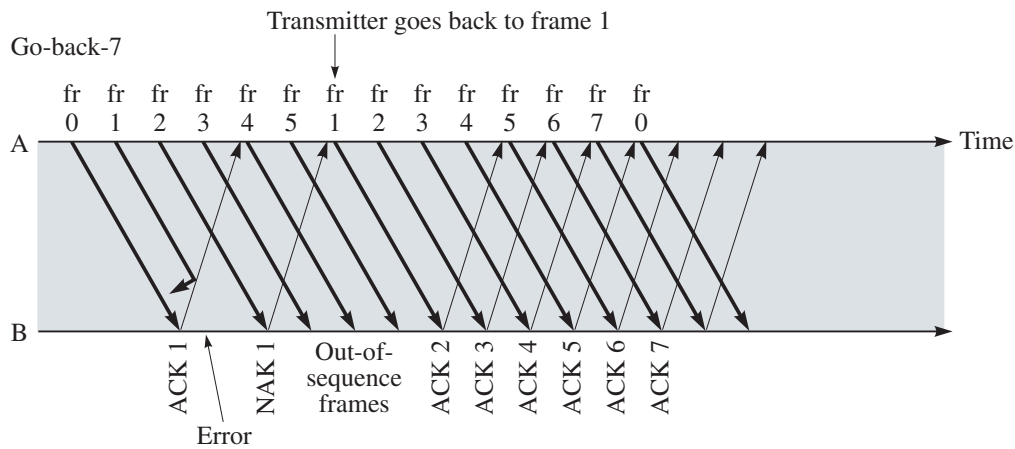


FIGURE 5.17 NAK error recovery

condition has been detected in frame R_{next} . The NAK informs the transmitter that the receiver is now discarding frames that followed R_{next} and instructs the transmitter to go back and retransmit R_{next} and all subsequent packets. By sending the NAK, the receiver sets a condition that cannot be cleared until frame R_{next} is received. For this reason, the receiver is allowed to send only one NAK for any given frame. If the NAK request fails, that is, the NAK is lost or the response transmission fails, then the receiver will depend on the timeout mechanism to trigger additional retransmissions. Note that in general the NAK procedure results in having the transmitter go back less than W_S frames.

Finally we consider the case where the information flow is bidirectional. The transmitter and receiver functions of the modified Go-Back-N protocol are now implemented by both stations A and B. In the direct implementation the flow in each direction consists of information frames and control frames. Many of the control frames can be eliminated by **piggybacking** the acknowledgments in the headers of the information frames as shown in Figure 5.18.

This use of piggybacking can result in significant improvements in the use of bandwidth. When a receiver accepts an error-free frame, the receiver can insert the acknowledgment into the next departing information frame. If no information frames are scheduled for transmission, the receiver can set an **ACK timer** that defines the maximum time it will wait for the availability of an information frame. When the time expires a control frame will be used to convey the acknowledgment.

In the bidirectional case the receiver handles out-of-sequence packets a little differently. A frame that arrives in error is ignored. Subsequent frames that are out of sequence but error free are discarded *after* the ACK (i.e., R_{next}) has been examined. Thus R_{next} is used to update the local S_{last} .

The I-frame time-out value should be selected so that it exceeds the time normally required to receive a frame acknowledgment. As shown in Figure 5.19, this time period includes the round-trip propagation delay $2T_{prop}$, two maximum-length frame transmission times on the reverse channel $2T_f^{max}$, and the frame

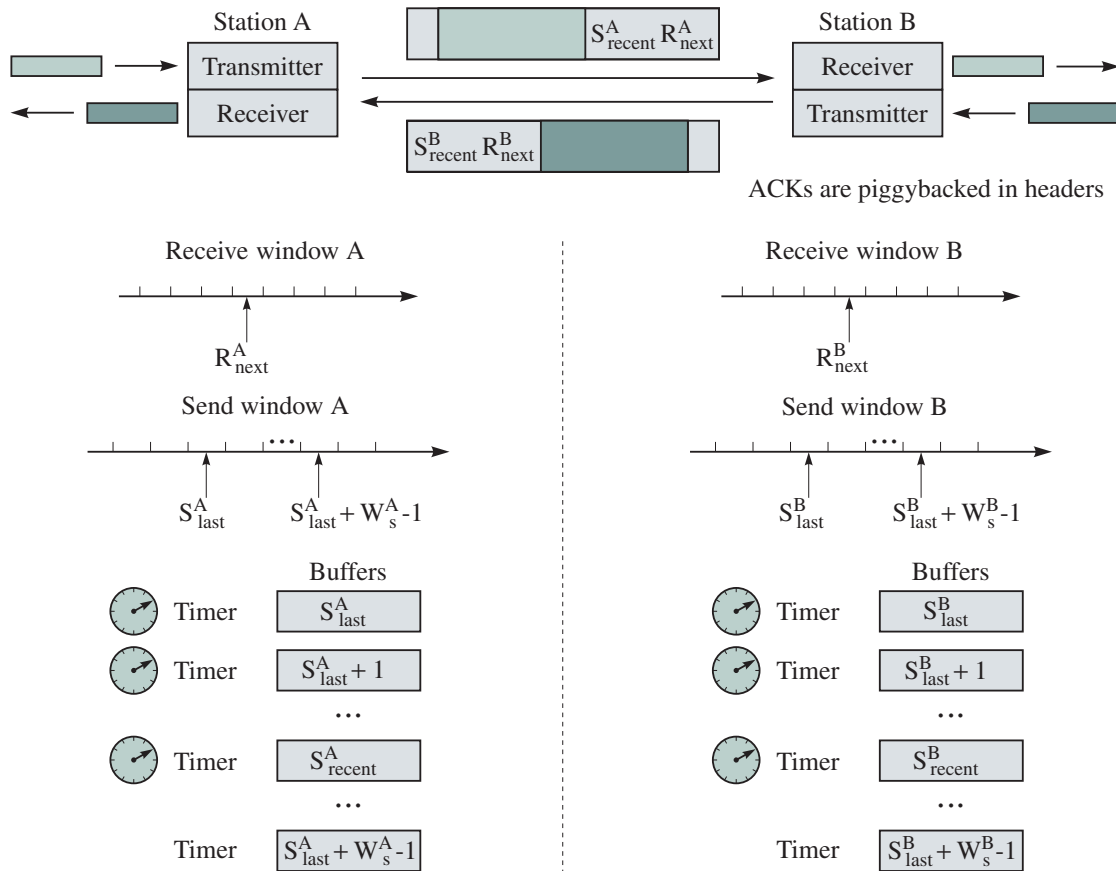


FIGURE 5.18 System parameters in bidirectional Go-Back-N ARQ

processing time T_{proc} .⁵ The transmission time of the frame in the forward direction and the ACK time-out are absorbed by the aforementioned frame transmission times.

$$T_{out} = 2T_{prop} + 2T_f^{max} + T_{proc}$$

As in the case of Stop-and-Wait ARQ, the performance of the modified Go-Back-N ARQ protocol is affected by errors in the information frames and errors in the acknowledgments. The transmission of long information frames at the receiver end can delay the transmission of acknowledgments, which can result in the expiry of time-outs at the transmitter end and in turn trigger unnecessary retransmissions. These long delays can be avoided by not using piggybacking and instead sending the short control frames ahead of the long information frames.

⁵Note that Figure 5.19 is more detailed than previous figures in that both the beginning and end of each frame transmission need to be shown explicitly.




FIGURE 5.19 Calculation of time-out values

Example—HDLC

The High-level Data Link Control (HDLC) is a data link control standard developed by the International Standards Organization. HDLC is a bit-oriented protocol that uses the CRC-16 and CRC-32 error-detection codes discussed in Chapter 3. HDLC can be operated so that it uses Go-Back-N ARQ, as discussed in this section, or Selective Repeat ARQ, which is discussed in section 5.2.3. HDLC is discussed in detail later in this chapter.

The most common use of HDLC is the Link Access Procedure—Balanced (LAPB) discussed later in this chapter. Several variations of the LAPB are found in different applications. For example, Link Access Procedure—Data (LAPD) is the data link standard for the data channel in ISDN. Another example is Mobile Data Link Protocol (MDLP), which is a variation of LAPD that was developed for Cellular Digital Packet Data systems that transmit digital information over AMPS and digital cellular telephone networks.

Example—V.42 Modem Standard

Error control is essential in telephone modem communications to provide protection against disturbances that corrupt the transmitted signals. The ITU-T V.42 standard was developed to provide error control for transmission over modem links. V.42 specifies the Link Access Procedure for Modems (LAPM) to provide the error control. LAPM is derived from HDLC and contains extensions for modem use. V.42 also supports the Microcom Network Protocol (MNP) error-correction protocols as an option. Prior to the development of V.42, these early protocols were de facto standards for error control in modems.

5.2.3 Selective Repeat ARQ

In channels that have high error rates, the Go-Back-N ARQ protocol is inefficient because of the need to retransmit the frame in error and all the subsequent

frames. A more efficient ARQ protocol can be obtained by adding two new features: first, the receive window is made larger than one frame so that the receiver can accept frames that are out of order but error free; second, the retransmission mechanism is modified so that only individual frames are retransmitted. We refer to this protocol as **Selective Repeat ARQ**. We continue to work under the constraint that the ARQ protocol must deliver an error-free and ordered sequence of packets to the destination.

Figure 5.20 shows that the send window at the transmitter is unchanged but that the receive window now consists of a range of frame numbers spanning from R_{next} to $R_{\text{next}} + W_R - 1$, where W_R is the maximum number of frames that the receiver is willing to accept at a given time. As before, the basic objective of the protocol is to advance the values of R_{next} and S_{last} through the delivery of the oldest outstanding frame. Thus ACK frames carry R_{next} , the oldest frame that has not yet been received. The receive window is advanced with the receipt of an error-free frame with sequence number R_{next} . Unlike the case of Go-Back-N ARQ, the receive window may be advanced by several frames. This step occurs when one or more frames that follow R_{next} have already been received correctly and are buffered in the receiver. R_{next} and the following consecutive packets are delivered to the destination at this point.

Now consider the retransmission mechanism in Selective Repeat ARQ. The handling of timers at the transmitter is done as follows. When the timer expires, only the corresponding frame is retransmitted. There is no longer a clear correspondence between the age of the timers and the sequence numbers. This situation results in a considerable increase in complexity. Whenever an out-of-sequence frame is observed at the receiver, a NAK frame is sent with sequence number R_{next} . When the transmitter receives such a NAK frame, it retransmits

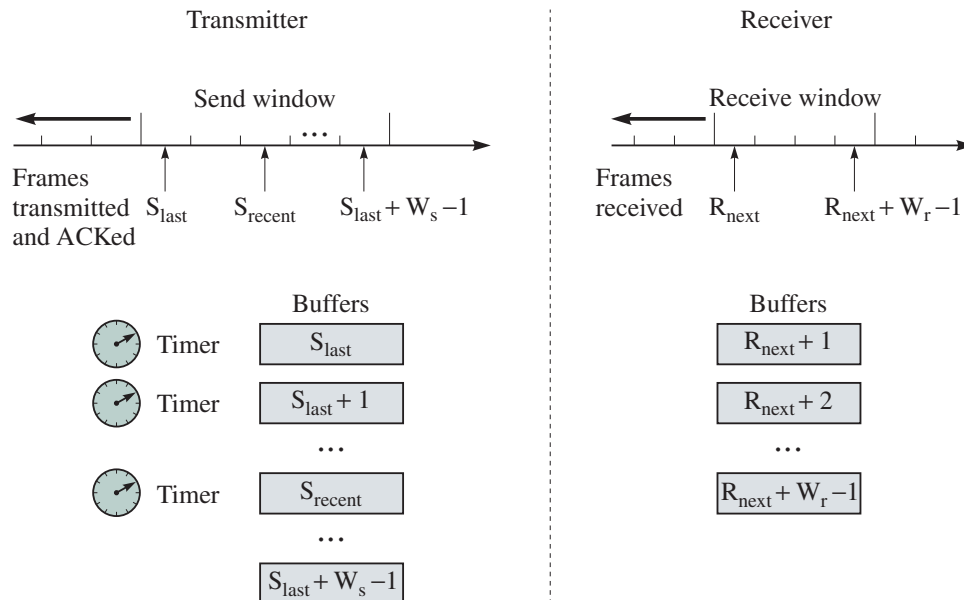


FIGURE 5.20 Selective Repeat ARQ

the specific frame, namely, R_{next} . Finally, we note that the piggybacked acknowledgment in the information frames continues to carry R_{next} .

For example, in Figure 5.21 when the frame with sequence number 2 finally arrives correctly at the receiver, frames 3, 4, 5, and 6 have already been received correctly. Consequently, the receipt of frame 2 results in a sliding of the window forward by five frames.

Consider now the question of the maximum send window size that is allowed for a given sequence numbering 2^m . In Figure 5.22 we show an example in which the sequence numbering is $2^2 = 4$ and in which the send windows and receive windows are of size 3. Initially station A transmits frames 0, 1, and 2.

All three frames arrive correctly at station B, and so the receive window is advanced to $\{3, 0, 1\}$. Unfortunately all three acknowledgments are lost, so when the timer for frame 0 expires, frame 0 is retransmitted. The inadequacy of the window size now becomes evident. Upon receiving frame 0, station B cannot determine whether it is the old frame 0 or the new frame 0. So clearly, send and receive windows of size $2^m - 1$ are too large.

It turns out that the maximum allowable window size is $W_S = W_R = 2^{m-1}$, that is, half the sequence number space. To see this we retrace the arguments used in the case of Go-Back-N ARQ. Suppose the window size W_S is 2^{m-1} or less and assume that the current send window is 0 to $W_S - 1$. Suppose also that the initial receive window is 0 to $W_S - 1$. Now suppose that frame 0 is received correctly but that the acknowledgment for frame 0 is lost. The transmitter can transmit new frames only up to frame $W_S - 1$. Depending on which transmissions arrive without error, R_{next} will be in the range between 1 and W_S while $R_{next} + W_R - 1$ will be in the range of 1 to $2W_S - 1$. The maximum value of R_{next} occurs when frames 0 through $W_S - 1$ are received correctly, so the value of R_{next} is W_S and the value of $R_{next} + W_R - 1$ increases to $2W_S - 1$. Crucially, the receiver will not receive frame $2W_S$ until the acknowledgment for frame 0 has been received at the transmitter. Any receipt of frame 0 prior to frame $2W_S$ indicates a duplicate transmission of frame 0. Therefore, the maximum size windows when $W_S = W_R$ is 2^{m-1} .

In the next section we develop performance models for the ARQ protocols. We show that Selective Repeat ARQ outperforms Go-Back-N and Stop-and-


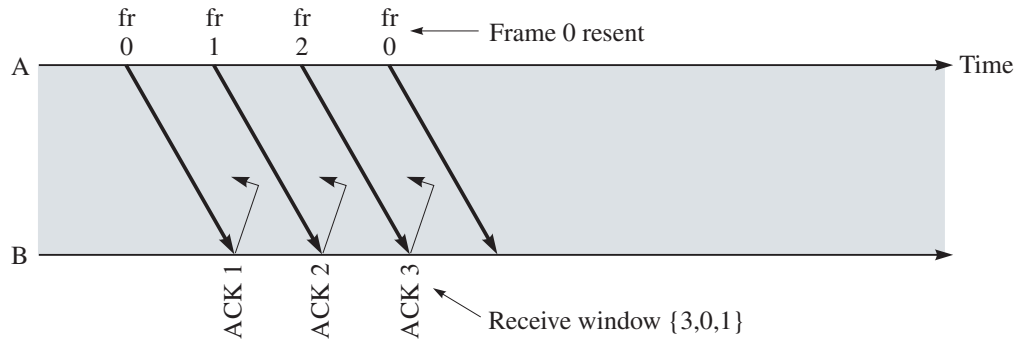


FIGURE 5.21 Error recovery in Selective Repeat ARQ

$M = 2^2 = 4$, selective repeat: send window = receive window = 3



Send window = receive window = 2

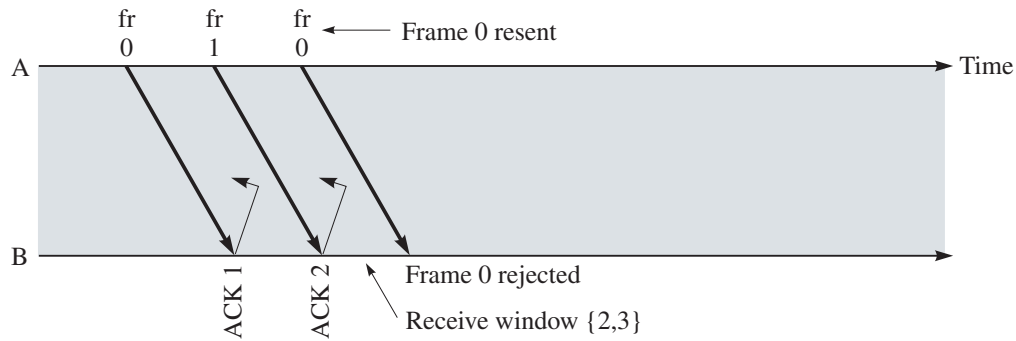


FIGURE 5.22 Maximum window size in Selective Repeat ARQ

Wait ARQ. Of course, this performance level is in exchange for significantly greater complexity.

Example—Transmission Control Protocol

The Transmission Control Protocol (TCP) uses a form of Selective Repeat ARQ to provide end-to-end error control across a network. TCP is normally used over internets that use IP to transfer packets in connectionless mode. TCP must therefore contend with the transfer of PDUs that may arrive out of order, may be lost, or may experience a wide range of transfer delays. We explain how TCP uses ARQ to provide the user of the TCP service with a reliable, connection-oriented stream service in section 5.3.3.

Example—Service Specific Connection Oriented Protocol

The Service Specific Connection Oriented Protocol (SSCOP) provides error control for signaling messages in ATM networks. SSCOP uses a variation of Selective Repeat ARQ that detects the loss of information by monitoring the sequence number of blocks of information that have been received and requesting selective retransmission. The ARQ protocol in SSCOP was originally designed for use in high-speed satellite links. These links have a large delay-bandwidth product, so

the protocol was also found useful in ATM networks, which also exhibit large delay-bandwidth product. SSCOP is discussed in Chapter 9.

ARQ: ROBUSTNESS AND ADAPTIVITY

Paul Green, in providing his list of remarkable milestones in computer communications in 1984, made the following remark about ARQ: “(ARQ) is one of those disarmingly simple ideas that seems so trivial (particularly in hindsight) that it really shouldn’t be on anyone’s list of exciting ideas.” ARQ was invented by H.C.A. van Duuren during World War II to provide reliable transmission of characters over radio. In his system each seven-bit character consisted of a combination of four marks (1s) and three spaces (0s). The reception of any other combination of bits led to a request for retransmission. This simple system led to the development of the many modern protocols that have ARQ as their basis.

Simplicity is not the only reason for the success of ARQ. Effectiveness and robustness are equally important attributes of ARQ. The retransmission mechanism gives ARQ the ability to adapt to variable and time-varying channel conditions. If a channel is clean, then ARQ operates in an efficient manner. If a channel becomes noisy, then ARQ adapts the transmission rate to the capability of the channel. This adaptivity makes ARQ relatively robust with respect to channel conditions and hence safe to deploy without detailed knowledge of the channel characteristics.

5.2.4 Transmission Efficiency of ARQ Protocols

In the previous sections we already introduced the performance differences between Stop-and-Wait, Go-Back-N, and Selective Repeat ARQ protocols. In this section we discuss performance results for these protocols and present a quantitative comparison of their transmission efficiency. We show that the delay-bandwidth product, the frame error rate, and the frame length are key parameters in determining system performance.

For simplicity we focus here on the case where the information is unidirectional and where the reverse channel is used only for acknowledgments. We also assume that all information frames have the same length and that the transmitter always has frames to transmit to the receiver. Figure 5.23 shows the components in the basic delay t_0 that transpires in Stop-and-Wait ARQ from the instant a frame is transmitted into the channel to the instant when the acknowledgment is confirmed. The first bit that is input into the channel appears at the output of the channel after a propagation time t_{prop} ; the end of the frame is received at station B after t_f additional seconds. Station B requires t_{proc} seconds to prepare an




FIGURE 5.23 Delay components in Stop-and-Wait ARQ

acknowledgment frame that will require t_{ack} seconds of transmission time. After an additional propagation delay, the acknowledgment frame is received at station A. Finally, t_{proc} additional seconds are required to carry out the CRC check. The basic time to send a frame and receive an ACK, in the absence of errors, is then given by

$$t_0 = 2t_{prop} + 2t_{proc} + t_f + t_{ack} = 2t_{prop} + 2t_{proc} + \frac{n_f}{R} + \frac{n_a}{R}$$

where n_f is the number of bits in the information frame, n_a is the number of bits in the acknowledgment frame, and R is the bit rate of the transmission channel.

The effective information transmission rate of the protocol in the absence of errors is then given by

$$R_{eff}^0 = \frac{\text{number of information bits delivered to destination}}{\text{total time required to deliver the information bits}} = \frac{n_f - n_o}{t_0}$$

where n_o is the number of overhead bits in a frame and is given by the total number of bits in the header and the number of CRC check bits. The transmission efficiency of Stop-and-Wait ARQ is given by the ratio R_{eff}^0 to R :

$$\eta_0 = \frac{\frac{n_f - n_o}{t_0}}{R} = \frac{1 - \frac{n_o}{n_f}}{1 + \frac{n_a}{n_f} + \frac{2(t_{prop} + t_{proc})R}{n_f}}$$

The preceding equation identifies clearly the sources of inefficiency in transmission. In the numerator the ratio n_o/n_f represents the loss in transmission efficiency due to the need to provide headers and CRC checks. In the denominator the term n_a/n_f is the loss in efficiency due to the time required for the acknowledgment message, where n_a is the number of bits in an ACK/NAK frame. Finally, the term $2(t_{prop} + t_{proc})R$ is the delay-bandwidth product, which was introduced earlier in the section.

Table 5.1 shows the efficiencies of Stop-and-Wait ARQ in a number of scenarios. We first assume a frame size of 1024 bytes (8192 bits) and $n_o = n_a = 8$ bytes. We consider three values of reaction time $t_{prop} + t_{proc}$: 5 ms, 50 ms, and 500 ms. Because the speed of light is 3×10^8 meters/second, these

		R (bps)				
		30 kbps	1.5 Mbps	45 Mbps	2.4 Gbps	
(a) $n_{frame} = 8192$ $n_{overhead} = 64, n_{ack} = 64$		0.005	0.95	0.35	1.77E-02	3.39E-04
$t_{prop} + t_{proc}$:	0.050	0.72	5.14E-02	1.80E-03	3.39E-05	3.39E-05
	0.500	0.21	5.39E-03	1.81E-04	3.39E-06	
(b) $n_{frame} = 524,288$		0.005	0.99	0.97	0.53	2.14E-02
$t_{prop} + t_{proc}$:	0.050	0.99	0.77	1.04E-01	2.18E-03	2.18E-03
	0.500	0.21	0.26	1.15E-02	2.18E-04	

TABLE 5.1 Efficiency of Stop-and-Wait ARQ in the absence of errors

delays correspond to distances of about 1500 km, 15,000 km, and 150,000 km links that could correspond roughly to moderate distance, intercontinental, and satellite links, respectively. We consider four values of bit rate: 30,000 bps for an ordinary telephone modem, 1.5 Mbps for a high-speed telephone line, 45 Mbps for high-speed access lines, and 2.4 Gbps for a high-speed backbone line. For the case of an ordinary telephone modem, we see that reasonable efficiencies are attained for reaction times up to 50 ms. This is why Stop-and-Wait ARQ is used extensively in data communication links. However, for the higher-speed lines the efficiencies quickly drop to the point where they are essentially zero. The efficiency can be improved by going to larger frame sizes. The second case considered in Table 5.1 is for $n_f = 64,000$ bytes = 524,288 bits. The efficiency is improved for the 1.5 Mbps line. Nevertheless, the weakness of the Stop-and-Wait protocol with respect to the delay-bandwidth product is again clearly evident in these examples.

We now consider the effect of transmission errors on the efficiency of Stop-and-Wait ARQ. If a frame incurs errors during transmission, the time-out mechanism will cause retransmission of the frame. Let P_f be the probability that a frame transmission has errors. We will assume that a frame transmission has errors independent of other frame transmissions.

Table 5.2 shows the transmission efficiencies of Stop-and-Wait, Go-Back-N, and Selective Repeat ARQ. Recall that Stop-and-Wait ARQ wastes $n_a/n_f + 2(t_{proc} + t_{prop})/n_f$ for every frame waiting for the acknowledgment. On the other hand, Go-Back-N ARQ wastes $W_S - 1$ only when an error occurs (with probability P_f). The superior efficiency of Go-Back-N can be seen by comparing the denominators in the expressions for Stop-and-Wait and Go-Back-N ARQ: For Stop-and-Wait we have $n_a/n_f + 2(t_{proc} + t_{prop})/n_f$, and for Go-Back-N we have $(W_S - 1)P_f$, which is much smaller in general. Note also that when $P_f = 0$, the efficiency of Go-Back-N ARQ is $1 - n_o/n_f$, which is the maximum possible given that the headers and CRC checks are essential.

The expressions in Table 5.2 for Go-Back-N and Selective Repeat ARQ are not affected explicitly by the delay-bandwidth product of the channel. Thus both schemes can achieve efficiency $1 - n_o/n_f$ in the absence of errors. Note, however, that this situation assumes that the window sizes have been selected to keep the

ARQ technique	Efficiency	Comments
<i>Stop-and-Wait</i>	$\eta = \frac{n_f - n_0}{R} = (1 - P_f) \frac{1 - \frac{n_0}{n_f}}{1 + \frac{n_a}{n_f} + \frac{2(t_{prop} + t_{proc})R}{n_f}} = (1 - P_f)\eta_0$	Delay-bandwidth product is main factor
<i>Go-Back-N</i>	$\eta = \frac{n_f - n_0}{R} = (1 - P_f) \frac{1 - \frac{n_0}{n_f}}{1 + (W_S - 1)P_f}$	Average wasted time: $(W_S - 1)P_f$
<i>Selective-Reject</i>	$\eta = (1 - P_f) \left(1 - \frac{n_0}{n_f}\right)$	Note impact of delay-bandwidth product through W_S

TABLE 5.2 Summary of performance results

channel busy while acknowledgments are returned. For Selective Repeat ARQ this assumption requires the buffering of up to W_S frames at the transmitter and at the receiver, and for Go-Back-N ARQ this assumption requires the buffering of up to W_S frames at the transmitter. The size W_S of the window required for each scenario considered in Table 5.2 is shown in Table 5.3. For the 30 kbps and the 1.5 Mbps channels, we obtain reasonable window sizes. For the 45 Mbps channel reasonable window sizes are obtained by going to longer frames, that is, 64 Kbytes frames instead of 1 Kbytes frames. However for the 2.4 Gbps channel, large window sizes are required even for 64 Kbytes frames.

The selection of optimum frame size must take the channel error rate into account. For a random bit error channel that introduces a bit error with probability p , we have

$$\begin{aligned} P_f &= 1 - P[\text{no bit errors in a frame}] \\ &= 1 - (1 - p)^{n_f} \\ &\approx 1 - (1 - n_f p) \approx n_f p \text{ for } p \ll 1 \end{aligned}$$

(a) $n_{frame} = 8192$

		R (bps)			
		30 kbps	1.5 Mbps	45 Mbps	2.4 Gbps
$t_{prop} + t_{proc}$:	0.005	2	4	57	2932
	0.050	2	20	551	29299
	0.500	6	185	5495	292,971

(b) $n_{frame} = 524,288$

$t_{prop} + t_{proc}$:	0.005	2	2	3	48
	0.050	2	2	11	460
	0.500	2	5	88	4580

TABLE 5.3 Required window sizes for different combinations of delay and bit rate

The preceding equation shows that the probability of frame error increases with frame size. For burst error channels the following approximation is sometimes used:

$$P_f \approx cn_f p \text{ where } 1/10 < c < 1/3.$$

As before, p is the bit error rate. We see that given the same p , the frame error rate is lower by a factor c for the burst error channel than for the random bit error channel. The reason is that the errors in the burst error channel tend to cluster in a few frames rather than be spread out randomly over many frames.

Figure 5.24 shows the relative performance of the three ARQ protocols in the presence of random bit errors that occur with probability p . The example considers the case of a 1024 byte frame, a 1.5 Mbps channel over a moderate-distance link, that is, 5 ms, where Go-Back-N and Selective Repeat ARQ use windows of size 4. It can be seen that Stop-and-Wait ARQ cannot achieve efficiencies higher than .35 regardless of the value of p . Go-Back-N ARQ achieves efficiencies comparable to Selective Repeat ARQ for p less than 10^{-6} and deteriorates to the performance of Stop-and-Wait ARQ when p increases to about 5×10^{-5} . Selective Repeat ARQ achieves high efficiencies over a greater range of p but also deteriorates as p becomes larger than 10^{-4} . Note that at this rate the probability of one or more errors in a frame is $1 - (1 - 10^{-4})^{8192} \approx 0.56$, so about half of the frames are in error.

Finally we return to the question of selecting the frame length. As the frame length is increased, the impact of the delay-bandwidth product is reduced. However, increasing the frame length also increases the probability of frame transmission error. These factors suggest that for any given bit error rate p , an optimum value of frame length n_f maximizes the transmission efficiency. In Figure 5.25 we show the throughput efficiency for Stop-and-Wait, Go-Back-N, and Selective Repeat ARQ as a function of n_f . We consider the case where $p = 10^{-4}$ and where the other parameters are the same as in the previous example. As n_f is varied, the window sizes are adjusted so that the channel can be kept busy. These window sizes are also shown in Figure 5.25. For Selective Repeat ARQ decreasing n_f reduces P_f and hence increases the efficiency. However if n_f

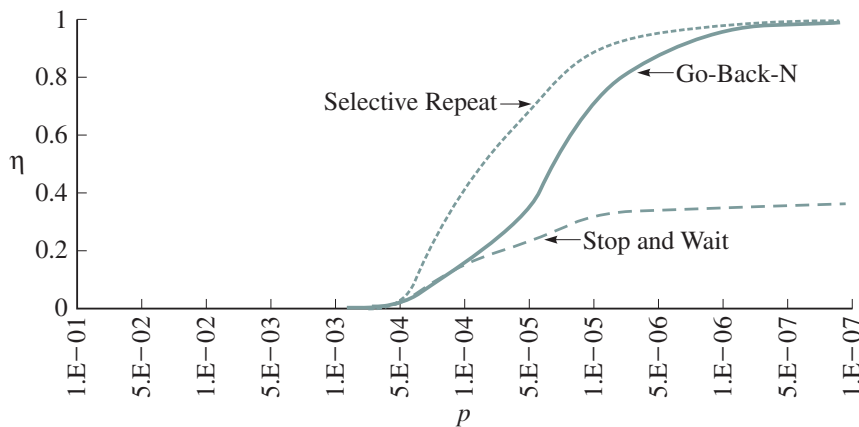


FIGURE 5.24 Transmission efficiency of ARQ protocols

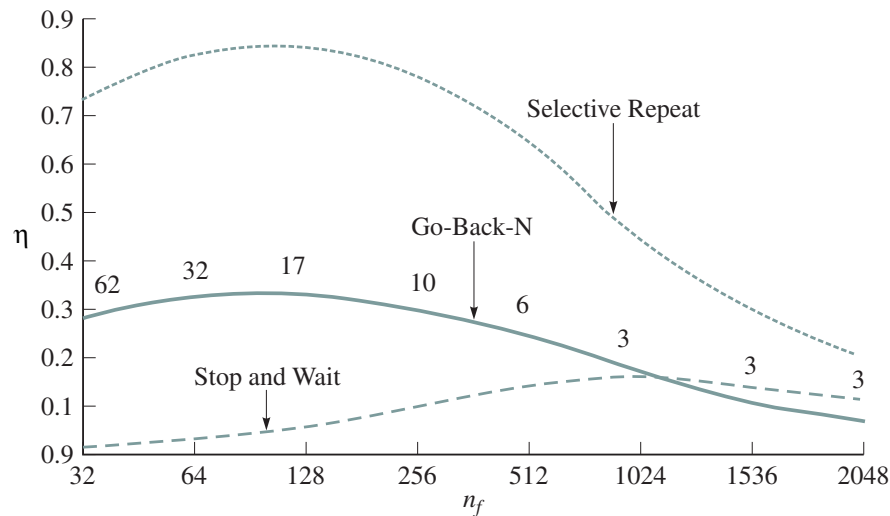


FIGURE 5.25 Optimum frame size

becomes too small, then the overhead term $(1 - n_0/n_f)$ starts decreasing the efficiency. Consequently, we see that the optimum n_f is in the range of 64 to 128 bytes. For Go-Back-N ARQ the optimum value of n_f balances the increases in efficiency due to decreasing P_f against the decrease in efficiency from the overhead term. Again the optimum value of n_f is in the range 64 to 128 bytes. Stop-and-Wait ARQ is more severely penalized for small values of n_f because of the effect of the delay-bandwidth product. Consequently, the optimum value of n_f is in the vicinity of 1024 bytes.

Example—Adapting Frame Size to Error Rate

Certain channels, such as cellular wireless telephone channels, can exhibit a wide variation in signal quality. These variations in quality are manifest in the rate at which transmitted frames incur errors. The results in Figure 5.25 suggest that significant improvements can be obtained by adjusting the frame length to the channel error rate. The MNP 10 error-control protocol was one of the first to implement this type of strategy of dynamically adjusting the frame size to increase the observed throughput in modem transmission over cellular telephone channels. More recent protocols incorporate the frame size adaptation feature as well as other physical layer adaptations to the channel conditions.

◆ DERIVATION OF EFFICIENCY OF ARQ PROTOCOLS

In this section we derive the results that were presented in Table 5.2. To calculate the transmission efficiency for Stop-and-Wait ARQ, we need to calculate the average total time required to deliver a correct frame. Let n_t be the number of transmissions required to deliver a frame successfully; then $n_t = i$ transmissions

are required if the first $i - 1$ transmissions are in error and the i th transmission is error free. Therefore,

$$P[n_t = i] = (1 - P_f)P_f^{i-1} \text{ for } i = 1, 2, 3 \dots$$

Each unsuccessful frame transmission will contribute a time-out period t_{out} to the total transmission time. The final successful transmission will contribute the basic delay t_0 . The average total time to transmit a frame is therefore

$$\begin{aligned} E[t_{total}] &= t_0 + \sum_{i=1}^{\infty} (i-1)t_{out}P[n_t = i] \\ &= t_0 + \sum_{i=1}^{\infty} (i-1)t_{out}(1 - P_f)P_f^{i-1} \\ &= t_0 + \frac{t_{out}P_f}{1 - P_f} \end{aligned}$$

The preceding equation takes into account the fact that in general the time-out period and the basic transmission time are different. To simplify the equation we will assume that $t_{out} = t_0$, then

$$E[t_{out}] = \frac{t_0}{1 - P_f}$$

Thus the effective transmission time for Stop-and-Wait ARQ is

$$R_{eff} = \frac{n_f - n_o}{E[t_{total}]} = (1 - P_f) \frac{n_f - n_o}{t_0} = (1 - P_f)R_{eff}^0$$

and the associated transmission efficiency is

$$\eta = \frac{E[t_{total}]}{R} = (1 - P_f) \frac{1 - \frac{n_o}{n_f}}{1 + \frac{n_a}{n_f} + \frac{2(t_{prop} + t_{proc})R}{n_f}} = (1 - P_f)\eta_0$$

In other words, the effect of errors in the transmission channel is to reduce the effective transmission rate and the efficiency by the factor $(1 - P_f)$.

Consider the basic Go-Back-N ARQ where the window size W_S has been selected so that the channel can be kept busy all the time. As before, let n_t be the number of transmissions required to deliver a frame successfully. If $n_t = i$, then as shown in Figure 5.14, $i - 1$ retransmissions of groups of W_S frames are followed by the single successful frame transmission. Therefore, the average total time required to deliver the frame is

$$\begin{aligned}
E[t_{total}] &= t_f \{1 + W_S \sum_{i=1}^{\infty} (i-1) P[n_t = i]\} \\
&= t_f \{1 + W_S \sum_{i=1}^{\infty} (i-1) (1 - P_f) P_f^{i-1}\} \\
&= t_f \left\{1 + W_S \frac{P_f}{1 - P_f}\right\} = t_f \left\{\frac{1 + (W_S - 1)P_f}{1 - P_f}\right\}
\end{aligned}$$

Thus the effective transmission time for basic Go-Back-N ARQ is

$$R_{eff} = \frac{n_f - n_o}{E[t_{total}]} = (1 - P_f) \frac{n_f - n_o}{t_f \{1 + (W_S - 1)P_f\}} = (1 - P_f) \frac{1 - \frac{n_o}{n_f}}{1 + (W_S - 1)P_f} R$$

and the associated transmission efficiency is

$$\eta = \frac{\frac{n_f - n_o}{E[t_{total}]}}{R} = (1 - P_f) \frac{1 - \frac{n_o}{n_f}}{1 + (W_S - 1)P_f}$$

Finally, for Selective Repeat ARQ each transmission error involves only the retransmission of the specific frame. Therefore, the average time required to transmit a frame is

$$\begin{aligned}
E[t_{total}] &= t_f \left\{1 + \sum_{i=1}^{\infty} (i-1) (1 - P_f) P_f^{i-1}\right\} \\
&= t_f \left\{1 + \frac{P_f}{1 - P_f}\right\} = t_f \frac{1}{1 - P_f}
\end{aligned}$$

Thus the effective transmission time for Selective Repeat ARQ is

$$R_{eff} = (1 - P_f) \left(1 - \frac{n_o}{n_f}\right) R$$

and the associated transmission efficiency is

$$\eta = (1 - P_f) \left(1 - \frac{n_o}{n_f}\right)$$

5.3 OTHER ADAPTATION FUNCTIONS

In section 5.2 we did a detailed study of ARQ sliding-window protocols. Our objective there was to provide reliable transfer of a sequence of messages over an unreliable communication channel. In this section we show how the various elements of the ARQ protocols can be used to provide other adaptation functions. In particular, we show how adaptation functions can provide flow control,