

Throughput Analysis of a Selective Repeat Retransmission Scheme for k -reliable Multicast

Bongtae Kim¹, Harry G. Perros², and Arne A. Nilsson¹

¹Department of Electrical and Computer Engineering &
Center for Communications and Signal Processing
North Carolina State University
Raleigh, NC 27695

²Department of Computer Science &
Center for Communications and Signal Processing
North Carolina State University
Raleigh, NC 27695

November 14, 1994

Abstract

This paper presents a selective repeat multicast retransmission scheme in which a message is considered to be successfully transmitted when a predefined number of acknowledgements (less than the number of all participating receivers) for the message is received. Analytical expressions for the maximum supportable throughput and the average number of receivers that receive a message correctly are derived. Numerical examples are provided to illustrate the effects of system parameters on the throughput.

1 Introduction

Retransmission based error control schemes are widely used in data communication networks in order to preserve data integrity. These schemes, called Automatic Repeat Request (ARQ), fall into one of three basic categories: stop-and-wait, go-back-N, and selective repeat. The stop-and-wait scheme allows only one message to be transmitted at a time and to remain unacknowledged. The second scheme, go-back-N, allows up to N messages to be in transit and unacknowledged. When a timeout occurs or a negative acknowledgement is received, the transmitter retransmits the timed-out or errored message and all subsequent messages already transmitted. In the selective repeat scheme, only timed-out or errored messages are retransmitted. The receiver keeps in its buffer the post errored messages until the errored messages is retransmitted and received correctly.

Many variants of the retransmission based error control schemes have been applied to a multicast environment, in which a single transmitter sends packets reliably to a set of receivers, see [1–3, 5, 7, 8]. It has been shown that the selective repeat schemes have the best

throughput performance, see [2, 6, 7]. These selective repeat schemes guarantee both error free and sequenced delivery of messages to all receivers participating in the multicast.

In this paper, we present and analyze the performance of a selective repeat multicast retransmission scheme, in which a message is considered to be correctly transmitted when a predefined number of acknowledgements are received. This predefined number of acknowledgements, hereafter referred to as the *ACK threshold*, is less than the number of receivers. Thus, if there are R receivers and the ACK threshold is k , $k < R$, the transmitter will assume that a message has been correctly transmitted at the moment when it receives the k -th acknowledgement.

This type of scheme is referred to as k -reliable, as opposed to a fully reliable scheme where all receivers must acknowledge positively before a message is considered to be correctly transmitted. When the ACK threshold is set to the number of receivers, i.e. $k = R$, the error control scheme becomes fully reliable. The k -reliable scheme is useful for applications which do not need acknowledgements from all receivers. Rather, they need a high throughput with an acceptable level of reliability. The k -reliable retransmission scheme was first proposed and analyzed by Kim, Nilsson, and Perros [4]. Four different stop-and-wait schemes were presented and the throughput performance of each schemes was analyzed.

This paper is organized as follows. In section 2, we briefly describe our proposed retransmission scheme. In section 3, we derive analytical expressions of the maximum supportable throughput of successfully transmitted messages and the average number of receivers that receive a message correctly. Numerical examples are given in section 4. Conclusions are given in section 5.

2 The Selective Repeat Scheme

In this section, we describe a selective repeat scheme for k -reliable multicasting. We consider one transmitter which transmits data to R receivers over a multicast network. All data are transmitted in the form of messages. Each message contains the following information: (1) a control field which includes a sequence number that uniquely identifies the message, and other information required by the error control scheme, (2) a cyclic redundancy check code which enables each receiver to detect transmission errors, and (3) an information field that contains the message to be transmitted.

The multicast network has many-to-one (R to 1) feedback paths on which acknowledgements from receivers are returned to the transmitter. An acknowledgement includes the receiver's identity and the sequence number of the acknowledged message. The acknowledgement is also protected against transmission errors.

The scheme uses only positive acknowledgements. That is, if no errors are detected in a received message, the receiver will send a positive acknowledgement. Erroneous or lost messages are not reported to the transmitter through negative acknowledgements. Rather, the transmitter detects erroneously transmitted message or lost messages using timeouts. An acknowledgement received in error is treated as if the corresponding message had been received in error by the receiver.

Transmitter Operation

The transmitter manages a timer for each transmitted message, a retransmission buffer and an *acknowledgement outstanding list* (AOL) [3]. When the transmitter starts to transmit a message, it also starts a timer for the message. If the message is transmitted for the first time, the transmitter initializes the AOL which contains the identity of all receivers from which acknowledgements are expected for that message. Once a message is transmitted, the message is stored in the retransmission buffer until it receives the k -th acknowledgement for the message. Upon receipt of an error free acknowledgement for a message from a receiver, that receiver is removed from the AOL. When k (or more) receivers are removed from the AOL before the timer for the message expires, the transmission of the message is completed and the message is removed from the retransmission buffer. If less than k acknowledgements for a message are received, the transmitter retransmits the message.

Receiver Operation

A receiver sends a positive acknowledgement to the transmitter whenever an error free message is received. Messages received with errors are discarded and no acknowledgement is sent back to the transmitter. All messages that are received correctly following an erroneous message are acknowledged, but they are not released to the upper layer. Instead, they are stored in the reordering buffer of the receiver. When the erroneous message is correctly retransmitted, the maximum set of messages with contiguous sequence numbers that were received correctly is released from the reordering buffer.

Because of the k -reliable scheme, it is possible that an erroneously received (or lost) message by a receiver may or may not be retransmitted. This depends upon whether this message has been correctly received and acknowledged by at least k other receivers. If the receiver assumes that the message will not be retransmitted, then it can release subsequent messages that it receives correctly to the upper layer. However, if the erroneous (or lost) message is retransmitted, it is possible that it will be received well after the receiver has released subsequent messages. Thus, the transmitted message will be released in an out-of-

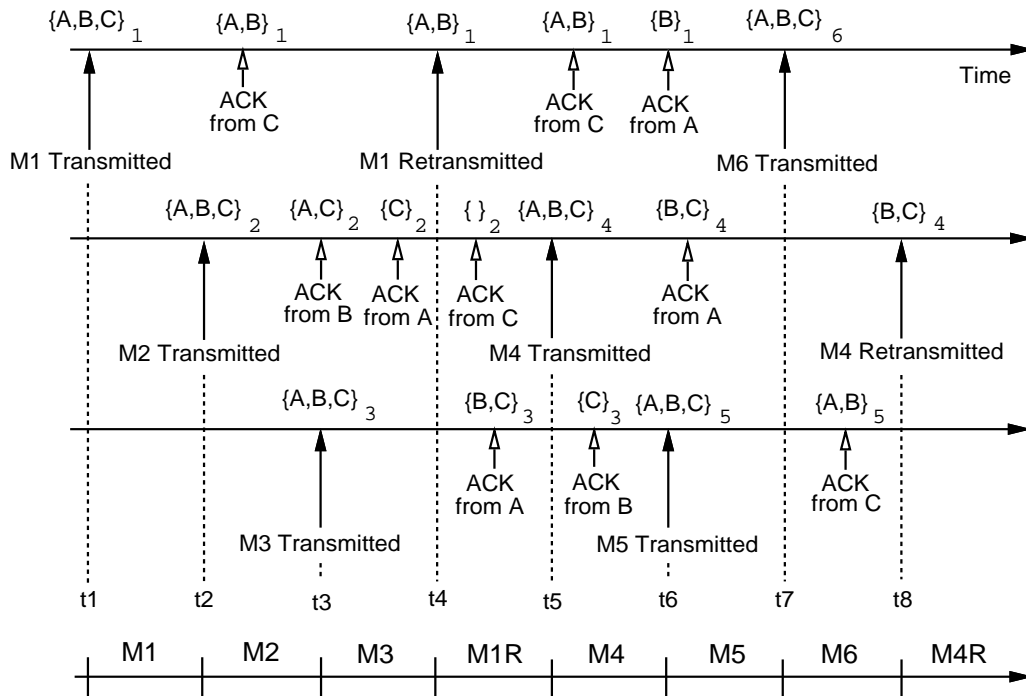


Figure 1: An example of the selective repeat retransmission scheme. 3 receivers : A,B,C, ACK threshold $k = 2$, and $N = 3$.

sequence order. To prevent this event, the transmitter includes status information in the control field of each transmitted message. Such information could be the sequence number of the oldest message held in its retransmission buffer. When a receiver receives a message correctly, the receiver can release all messages with lower sequence numbers than that of the oldest message.

Figure 1 shows an example of the selective repeat scheme with three receivers A,B,C, and the ACK threshold $k=2$. The number of messages N that can be transmitted without waiting acknowledgements is set to 3. In Figure 1, Message 1 (M1) is transmitted at time $t1$ followed by M2 and M3 at $t2$ and $t3$ respectively. Each horizontal line gives the time axis for each message. For each transmitted message, the AOL is initialized to $\{A,B,C\}$ and a timer for each message is started. At time $t4$, timer for M1 expires but only one acknowledgement for M1 has received so far. Thus, M1 is retransmitted. Note that the AOL is not reinitialized. At time $t5$, all three acknowledgements for M2 has received, and a new message M4 is transmitted, and so on.

3 Performance Analysis

3.1 Model Assumptions and Definitions

The round-trip delay for a particular receiver is defined as the time interval elapsing from the moment that the transmission of a message is started to the moment that an acknowledgement from this receiver is received. The round-trip delay, which includes delays such as message transmission time, round-trip propagation delays, and processing delays, is assumed to be independent among different receivers. In addition, we assume the followings:

- Messages are always waiting to be transmitted at the transmitter, and all messages are of the same length.
- The time axis is partitioned into slots, each equal to the transmission time of a message. Message transmission or retransmissions starts at the beginning of a slot.
- Each receiver receives a message correctly with probability p , $0 \leq p \leq 1$. The probability is identical and independent among different receivers. Acknowledgements are always received error free.
- The maximum round-trip delay is constant, equal to N slots. That is, it is equal to $N \times T_F$, where N is an integer and T_F is the transmission time of a message (which is equal to 1 slot). N can be seen as the number of messages that can be transmitted in one round-trip delay.
- The timeout period T_O is set equal to N slots.

Since the maximum round-trip delay is equal to N slots, k or more acknowledgements may arrive for the specific message M before the timer expires. Let us assume that they arrive within N' slots, $N' < N$. That is, after N' slots, message M has to be removed from the retransmission buffer. However, the transmitter has always messages to transmit, we assume that it does get round to handling message M until after the N -th slot.

3.2 Throughput Analysis

Let $\beta(R, k)$ be the average number of times a message is transmitted until at least k acknowledgements are received from R receivers. Then, the throughput $\eta(R, k)$ is

$$\eta(R, k) = \frac{1}{\beta(R, k)}. \tag{1}$$

For a particular receiver i ($1 \leq i \leq R$), the probability of receiving an acknowledgement on the x -th transmission is

$$\Pr\{X = x\} = (1-p)^{x-1}p, \quad x = 1, 2, 3, \dots \quad (2)$$

where X is a random variable denoting the number of times a message is transmitted until an acknowledgement from a particular receiver is received. From equation (2) we have

$$\begin{aligned} \Pr\{X \leq x\} &= \sum_{z=1}^x (1-p)^{z-1}p \\ &= 1 - (1-p)^x. \end{aligned} \quad (3)$$

Let Y be the random variable denoting the number of times a message is transmitted until at least k acknowledgements are received. The probability of receiving at least k acknowledgements within x transmissions is

$$\begin{aligned} \Pr\{Y \leq x\} &= \sum_{j=k}^R \binom{R}{j} \Pr\{X \leq x\}^j (1 - \Pr\{X \leq x\})^{R-j} \\ &= \sum_{j=k}^R \binom{R}{j} [1 - (1-p)^x]^j [(1-p)^x]^{R-j}. \end{aligned} \quad (4)$$

Notice that the maximum distribution of j independent identically distributed random variables is $\Pr\{X \leq x\}^j$. The expected number of times that a message is transmitted until k or more acknowledgements are received is

$$\begin{aligned} \beta(R, k) &= \sum_{x=0}^{\infty} (1 - \Pr\{Y \leq x\}) \\ &= \sum_{x=0}^{\infty} \left\{ 1 - \sum_{j=k}^R \binom{R}{j} [1 - (1-p)^x]^j [(1-p)^x]^{R-j} \right\} \\ &= \sum_{x=0}^{\infty} \sum_{j=0}^{k-1} \binom{R}{j} [1 - (1-p)^x]^j [(1-p)^x]^{R-j} \\ &= \sum_{j=0}^{k-1} \sum_{l=0}^j \binom{R}{j} \binom{j}{l} (-1)^l \sum_{x=0}^{\infty} [(1-p)^{R-j+l}]^x \\ &= \sum_{j=0}^{k-1} \sum_{l=0}^j \binom{R}{j} \binom{j}{l} (-1)^l \frac{1}{1 - (1-p)^{R-(j-l)}} \\ &= \sum_{l=0}^{k-1} \binom{R}{l} (-1)^l \sum_{j=l}^{k-1} \binom{R-l}{j-l} \frac{1}{1 - (1-p)^{R-(j-l)}} \\ &= \sum_{l=0}^{k-1} \binom{R}{l} (-1)^l \sum_{j=0}^{k-1-l} \binom{R-l}{j} \frac{1}{1 - (1-p)^{R-j}}. \end{aligned} \quad (5)$$

The throughput $\eta(R, k)$ is

$$\begin{aligned}\eta(R, k) &= \frac{1}{\beta(R, k)} \\ &= \frac{1}{\sum_{l=0}^{k-1} \binom{R}{l} (-1)^l \sum_{j=0}^{k-1-l} \binom{R-l}{j} \frac{1}{1-(1-p)^{R-j}}}.\end{aligned}\quad (6)$$

In equation (6), the alternating sign makes the calculation unstable when R is greater than 35. Equation (6) can be expressed in the following recursive form:

$$\begin{aligned}\eta(R, 1) &= 1 - q^R, \\ \eta(R, k) &= \frac{\eta(R, k-1)}{1 + \eta(R, k-1)\alpha(R, k-1)}\end{aligned}$$

where

$$\alpha(R, k) = \sum_{x=0}^{\infty} \binom{R}{k} (1 - q^x)^k (q^x)^{(R-k)}, \text{ and } q = 1 - p. \quad (7)$$

To compute $\alpha(R, k)$ in (7), let $B(x, j) = \binom{R}{j} (1 - q^x)^j (q^x)^{R-j}$. $B(x, j)$ can be put in the following recursive form:

$$B(x, j+1) = \frac{(R-j)(1-q^x)}{(j+1)q^x} B(x, j),$$

and

$$B(x, j-1) = \frac{j}{(R-j+1)(1-q^x)} B(x, j).$$

Now, we can compute $B(x, j)$ as follows:

1. Assume $B(x, j_{max})=1$ where $j_{max} = \lfloor Rq^x \rfloor$.
2. Compute $B(x, j)$ for $j = j_{max} - 1, \dots, 0$.
3. Compute $B(x, j)$ for $j = j_{max} + 1, \dots, R$.
4. Normalize such that $\sum_{j=0}^R B(x, j) = 1$.

$\alpha(R, k)$ is computed by summing $B(x, k)$ over x . The upper limit of x , x_{max} , is determined when $B(x, k)$ converges. That is, x_{max} is determined as a number that satisfies $|B(x+1, k) - B(x, k)| < \epsilon$, where ϵ is a positive small real number.

3.3 The Average Number of Happy Receivers

The k -reliable multicast retransmission scheme guarantees that at least k receivers receive a message correctly. $R - k$ or less receivers, however, may or may not receive a message correctly. A receiver that receives a message correctly is called a *happy* receiver. In this section, we obtain the average number of happy receivers for a message after a message has been correctly transmitted. That is, the transmitter has received at least k acknowledgements. As discussed above, it may take several retransmissions until at least k acknowledgements have been received.

We define a discrete time Markov chain $\{X_n\}$ with absorbing states. The Markov chain is associated with the transmission of a single message. X_n represents the cumulative number of acknowledgements received immediately after the n -th retransmission, and it takes the values of $0, 1, \dots, R$. The unit time of the Markov chain is equal to N slots, i.e. equal to the maximum round-trip delay or the timeout T_O . For $n = 0$, we have that $X_0 = 0$. That is, the message is about to be transmitted for the first time. Figure 2 shows the states of this Markov chain and the associated transition probability. The Markov chain gets absorbed when it enters a state with k or more acknowledgements. Thus, states 0 to $k - 1$ are transient, and states k to R are absorbing.

Let \mathbf{P} be the transition probability matrix of the process. We have

$$\mathbf{P} = \begin{pmatrix} P_{00} & P_{01} & P_{02} & \cdots & P_{0k-1} & P_{0k} & P_{0k+1} & \cdots & P_{0R} \\ 0 & P_{11} & P_{12} & \cdots & P_{1k-1} & P_{1k} & P_{1k+1} & \cdots & P_{1R} \\ 0 & 0 & P_{22} & \cdots & P_{2k-1} & P_{2k} & P_{2k+1} & \cdots & P_{2R} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & P_{k-1k-1} & P_{k-1k} & P_{k-1k+1} & \cdots & P_{k-1R} \\ 0 & 0 & 0 & \cdots & 0 & 1 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 0 & 0 & 1 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 1 \end{pmatrix}$$

where $P_{ij} = \binom{R-i}{j-i} p^{j-i} (1-p)^{R-j}$ for $0 \leq i \leq j$. The transition probability matrix \mathbf{P} has the form

$$\mathbf{P} = \begin{pmatrix} \mathbf{Q} & \mathbf{R} \\ \mathbf{O} & \mathbf{I} \end{pmatrix},$$

where \mathbf{O} is an $(R-k+1) \times k$ matrix whose entries are all zero, \mathbf{I} is an $(R-k+1) \times (R-k+1)$ identity matrix, \mathbf{Q} is a $k \times k$ upper triangular matrix whose nonzero entries $Q_{ij} = P_{ij}$ for $0 \leq i \leq j < k$, and \mathbf{R} is a $k \times (R-k+1)$ matrix whose entries $R_{ij} = P_{ij}$ for $0 \leq i < k$, $k \leq j \leq R$.




Figure 2: The Markov chain of the transmission process of a message.

Let U_{il} be the probability that the chain is absorbed at state l , given that it is at state i , $i = 0, 1, \dots, k-1$. Applying the law of total probability, we have

$$U_{il} = R_{il} + \sum_{j=0}^{k-1} Q_{ij}U_{jl}. \quad (8)$$

Equation (8) can be written in the following matrix form:

$$\mathbf{U} = \mathbf{R} + \mathbf{Q}\mathbf{U} \text{ or } \mathbf{U} = \mathbf{W}\mathbf{R} \text{ where } \mathbf{W} = (\mathbf{I} - \mathbf{Q})^{-1}.$$

Since the chain starts at state 0, we need to compute U_{0l} which denotes the probability that the transmitter receives exactly l ($l = k, k+1, \dots, R$) acknowledgements. The elements in the first row of the matrix \mathbf{U} represent U_{0l} . Thus, the average number of happy receivers $\gamma(R, k)$ that receive a message correctly is

$$\gamma(R, k) = \sum_{j=k}^R jU_{0j}. \quad (9)$$

The throughput can be also obtained by determining the mean time to absorption. For each state $i = 0, 1, 2, \dots, k-1$, let V_i be the expected time to absorption, given that the process started from state i . We have

$$V_i = 1 + \sum_{j=0}^{k-1} Q_{ij}V_j. \quad (10)$$

Let \mathbf{V} be a column vector with k elements and $\mathbf{1}$ be a column vector with k ones. Then, equation (10) can be written in the following matrix form:

$$\mathbf{V} = \mathbf{1} + \mathbf{Q}\mathbf{V} \text{ or } (\mathbf{I} - \mathbf{Q})\mathbf{V} = \mathbf{1}.$$

Since the matrix $\mathbf{I} - \mathbf{Q}$ is an upper triangular matrix, V_0 can be obtained by back substitution. The throughput $\eta(R, k)$ is equal to $\frac{1}{V_0}$.

4 Numerical Examples

We give plots of the throughput of the retransmission scheme in order to illustrate the effect of the system parameters. We assume that the message transmission time $T_F = 1$.

Figure 3 gives the throughput in terms of the probability p of receiving an acknowledgement for various ACK thresholds k when the number of receivers $R=10$. As the ACK threshold decreases for a fixed p , the throughput increases. If $p \geq 0.5$, the throughput of a multicast system with $k = 1$ is close to 1. This implies that the transmitter can transmit messages successfully at full utilization of the available bandwidth. The fully reliable multicast, $k = 10$, provides a lower bound on the throughput for all values of p .

Figure 4 shows the average number of happy receivers with different ACK thresholds. As p increases, the average number of happy receivers increases. As p decreases, the average number of happy receivers decreases.

In the fully reliable scheme, $k = R$, the throughput $\eta(R, k)$ of the transmitter is equal to the throughput $\zeta(R, k)$ of a receiver, i.e. the rate at which the receiver receives a message correctly. In the k -reliable scheme, $k < R$, the transmitter's throughput may not be equal to that of a receiver. Let $\gamma(R, k)$ be the average number of happy receivers. Then, we have

$$\zeta(R, k) = \frac{\eta(R, k)\gamma(R, k)}{R}.$$

We use the receiver's throughput to compare the performance of the multicast system for different ACK thresholds. For a multicast system with $k=5$ at $p=0.9$, the average number of happy receivers $\gamma(10, 5)$ is 9 and the throughput $\eta(10, 5)$ is close to 1. The receiver's throughput $\zeta(10, 5)$ for this system is 0.9. In the fully reliable multicast system with $k = 10$ at $p = 0.9$, $\zeta(10, 10) = \eta(10, 10) = 0.57$. The k -reliable scheme has a receiver's throughput 1.58 times higher than that of the fully reliable scheme (at the expense of reliability in the delivery of data). In Figure 5, the receiver's throughput is plotted in terms of p for various k . Notice that the maximum achievable throughput of any retransmission scheme, is p [9]. When $k = 1$, the receiver's throughput approaches this upper bound. $k = 10$ provides the lower bound on the throughput.

Figures 6 and 7 give the throughput in terms of the probability p of receiving a message correctly for various ACK thresholds k with the number of receivers $R = 30$ and $R = 50$ respectively.

5 Conclusions

We analyzed the performance of a selective repeat retransmission scheme for k -reliable multicasting in terms of maximum supportable throughput of successfully transmitted messages. An analytic expression for the average number of receivers that receive a message correctly is derived. An algorithm to calculate the throughput of a large multicast system is provided.

The fully reliable scheme provides a lower bound on the throughput. For a fixed probability of receiving a message correctly, as the ACK threshold decreases, the throughput increases. The k -reliable scheme can achieve high throughput at the expense of the reliability of delivered data. The correct balance between throughput and reliability must be evaluated for each application.

References

- [1] S. B. Calo and M. C. Easton, "A Broadcast Protocol for File Transfers to Multiple Sites," *IEEE Transactions on Communications*, vol. 29, pp. 670–678, November 1981.
- [2] S. R. Chandran and S. Lin, "Selective-repeat ARQ Schemes for Broadcast Links," *IEEE Transactions on Communications*, vol. 40, pp. 12–19, January 1992.
- [3] I. S. Gopal and J. M. Jaffe, "Point-to-multipoint Communication over Broadcast Links," *IEEE Transactions on Communications*, vol. 32, pp. 1034–1044, September 1984.
- [4] B. Kim, A. A. Nilsson, and H. G. Perros, "Throughput Analysis of Stop-and-wait Retransmission Schemes for k -reliable Multicast," in *Proceedings of International Conference on Local and Metropolitan Communication Systems*, (Kyoto, Japan), Chapman & Hall, December 1994.
- [5] K. Mase, T. Takenaka, H. Yamamoto, and M. Shinohara, "Go-back-N ARQ Schemes for Point-to-multipoint Satellite Communications," *IEEE Transactions on Communications*, vol. 31, pp. 583–589, April 1983.
- [6] N. Shacham and D. Towsley, "Resequencing Delay and Buffer Occupancy in a Selective-repeat ARQ with Multiple Receivers," in *Proceedings of IEEE INFOCOM '88*, pp. 515–524, 1988.
- [7] D. Towsley and S. Mithal, "A Selective-repeat ARQ Protocol for a Point-to-multipoint Channel," in *Proceedings of IEEE INFOCOM '87*, (San Francisco, CA), pp. 521–526, March 1987.
- [8] J. L. Wang and J. A. Silvester, "Optimal Adaptive Multireceiver ARQ Protocols," *IEEE Transactions on Communications*, vol. 41, pp. 1816–1829, 1993.
- [9] E. J. Weldon, "An Improved Selective-repeat ARQ Strategy," *IEEE Transactions on Communications*, vol. 30, pp. 480–486, March 1982.




Figure 3: Throughput vs probability p of receiving a message correctly : $R=10$.




Figure 4: Average number of happy receivers vs probability p of receiving a message correctly : $R=10$.




Figure 5: Receiver's throughput vs probability p of receiving a message correctly : $R=10$.




Figure 6: Throughput vs probability p of receiving a message correctly : $R=30$.




Figure 7: Throughput vs probability p of receiving a message correctly : $R=50$.