Analysis and Estimation of the Quality of Service of Group Communication Protocols

A. Coccoli[1], A. Bondavalli[2], F. Di Giandomenico[3]

1 Dip. di Ing. dell'Informazione, Univ. of Pisa, via Diotisalvi 2, 56126 Pisa, Italy - a.coccoli@guest.cnuce.cnr.it

2 Dip. di Sistemi e Informatica, Univ. of Firenze, Via Lombroso 6/17, Firenze, Italy - a.bondavalli@dsi.unifi.it

3 IEI-CNR, Area della Ricerca di Pisa, Via Alfieri 1, 56127 Pisa, Italy - digiandomenico@iei.pi.cnr.it

Abstract

QoS (defined as a proper set of quantitative characteristics) analysis is a necessary step for the early verification and validation of an appropriate design, and for taking design decisions about the most rewarding choice, in relation with user requirements. In this paper, we describe an analytical approach for the evaluation of the QoS offered by a family of group communication protocols in a wireless environment, and use experimental data for feeding our models.

Specific indicators have been defined and evaluated, which capture the main characteristics of the protocols and of the environment, focusing our attention on performance and dependability attributes. The main purpose of our analysis is to provide a fast, cost effective, and formally sound way to further analyze and understand the protocol behavior and its environment.

The physical characteristics of the system have been considered and the models introduced in this paper account for the correlation among successive packet transmissions due to fading and users mobility.

1. Introduction

The recent technological achievements such as the availability of the WEB and of mobile networks, constitute a fundamental basis for the development of distributed applications. An important problem that need to be solved for improving their success consists in devising designs providing the proper quality of service (QoS) as required by applications [11, 12]. QoS can be defined as a set of qualitative and quantitative characteristics of a distributed system necessary for obtaining the required functionality of an application. Therefore the term QoS encompasses many aspects such as reliability, availability, fault tolerance and also properties such as the atomicity or reliability of broadcast/multicast services.

It is clear that the usefulness and practical utilization of such (sub)system designs depend on the possibility to provide a QoS analysis of their offered features, in terms of proper defined dependability and performability related measures. When building a system, this is a necessary step for the early verification and validation of an appropriate design, and for taking design decisions about the most rewarding choice, in relation with user requirements.

This paper, addresses QoS analysis of a family of group communication protocols in wireless environment. Real-time reliable group communication in wireless local area networks has to deal, in particular, with the mobility of system components, and with the hostility of the environment that may cause a great loss of messages. The protocol defined in [7] tolerates erroneous and uncooperative behavior of the system components and provides the group of active components with a consistent view of its state. This protocol extends the IEEE 802.11 standard for wireless local area networks [1] and offers the interesting possibility to handle a trade-off between different aspects of the QoS offered, such as performances, delay time, reliability and formal properties of the broadcast. This flexibility in offering a different trade-off between performance, delay time, reliability and formal properties of the broadcast can be properly exploited if a fast, cost effective, and formally sound analysis of QoS can be performed.

In this paper, we analyze the protocol (family) and its environment, focusing our attention on typical performance indicators and on the coverage of the assumptions the correctness of the protocol is based on. We adopt an analytical approach based on Stochastic Activity Networks (SAN) [9, 10]. Experimental data previously collected in a representative context [3] have been used to provide parameters values for the model. The work described in this paper extends the preliminary analysis described in [3]. In fact, the models presented here are a much closer representation of the system and the environment. They account for physical characteristics such as the *fading channel* phenomenon, and for user mobility. Both of them affect a wireless communication and cause time correlation among successive messages, which is captured by our models.

The rest of the paper is organized as follows. Section 2 is devoted to the description of the considered communication protocols, together with the definition of relevant metrics representative of the QoS in the selected environment. In Section 3, our approach to modeling and the assumption made are described. . Section 4 contains the description of our models. The physical setting and the experiments used to derive proper values for the parameters are introduced in Section 5, while Section 6 contains the results of the models' evaluation. Finally, concluding remarks are outlined in Section 7.

2. A family of group communication protocols for wireless local area networks

2.1. Design of the protocols

Co-operation of autonomous mobile systems requires ability to communicate via wireless links. To achieve a real-time reliable group communication [4] in wireless local area networks requires coping with the mobility of the system components (definition of a co-operative group) and with the hostility of the environment (great loss of messages). The protocols presented in [7] provide reliable and efficient group communication services, based on extending the IEEE 802.11 standard for wireless local area networks.

In particular, they have been developed taking advantage of the centralized medium arbitration (Access Point, or AP), granted by the IEEE 802.11 during the Contention Free Period (CFP). The AP concedes exclusive access to the medium by transmitting a polling message to the stations in the group, according to a polling policy. The proposed protocols are based on the following fault assumptions:

1) Messages delivered during the CFP are delivered correctly within a fixed time-bound (t_m).
2) Messages may be lost (omission faults), possibly in an asymmetric way, i.e., some stations may receive a broadcast message and some may not. However, the number of consecutive message losses is bounded by the so-called *omission degree OD*.
3) Stations may suffer crash failures or leave the reach of the Access Point.
4) The Access Point is reliable; i.e., it is not subject to any kind of error.

A first protocol uses the AP as central co-ordinator and is structured into rounds in which the AP polls each station of the group exactly once. After being polled, a station returns a *broadcast request message* to the AP, which assigns a sequence number to the message and broadcasts it to the stations group. The broadcast request message is also used to acknowledge previous broadcasts. Its header contains some bits, each used to acknowledge one of the preceding broadcasts. Thus, one round after sending the i-th *broadcast message* m_i^x emitted by the generic station x, the AP is able to decide whether each group member has received m_i^x or not. In the latter case, the AP will retransmit m_i^x, otherwise $m_{(i+1)}^x$, if it exists. By the assumptions made above, a message is successfully received by all the stations after at most *OD*+1 rounds. If the AP does not receive any *broadcast request message* from station x within a certain period of time after polling x, the AP considers this message (or the polling message) to be lost. If AP does not receive any answer from station x after polling *OD* consecutive times, it considers x to have left the group and broadcasts a message indicating the change in the group membership.

This protocol implements a reliable group communication satisfying the properties of

i) *validity*, i.e., a message broadcast by a correct station is eventually delivered by every correct station;

ii) *agreement*, i.e., a message delivered by a station is eventually delivered by any other correct station;

iii) *integrity*, i.e., for any message m, every correct station delivers m at most once and only if m has been broadcast.

In order to improve the real time guarantees, a variant of the protocol has been proposed, which allows the user to specify the maximum number of message retransmissions, lower than *OD*. Such user-defined bound on message retransmissions (called *resiliency degree*, *res(c)*) can differ for different message classes. This variant makes use of *decision messages*. Whenever a message *m* is acknowledged by all stations (within *res(c)*+1 rounds), the AP broadcasts the decision message to deliver *m* to the applications, (retransmitted *OD*+1 consecutive times to guarantee reception by all the correct stations under assumption 2) above). If, however, m is not acknowledged by at least one station after *res(c)*+1 rounds, a decision not to deliver *m* is issued, again through the broadcast of a *decision message*. To make the implementation efficient, the access point piggybacks its decisions on broadcast messages, by properly extending their headers. The shorter delivery time for a message, obtained by reducing the maximum number of retransmissions for a broadcast message to *res(c)* times, is however paid in terms of violation of the validity property. In fact, a message issued by a correct station may be not received by all the other stations and therefore not delivered to any station. However, the agreement and integrity properties are retained, which is enough for significant application scenarios.

Because of the different characteristics shown by the two versions, they cannot be compared one against the other in an absolute way; the choice of which one is better suited to be employed in a system depends on the requirements of the specific application at hand.

2.2. Definition of appropriate QoS indicators

The protocols described above have been defined to provide reliable and efficient co-operation of autonomous mobile systems via wireless links. In order to prove such basic characteristics of reliability and efficiency, we estimate two groups of figures of interest: dependability related measures and performance related ones.

The dependability-related figures are directed to give an estimate of the coverage of the assumption on *OD*, the maximum number of consecutive message losses. They are:

i) $P_{R>OD}$, which indicates the probability that a *broadcast message* is not received by at least one of the receiving stations after *OD*+1 transmissions, in the time interval T_{CFP} (representing the duration of a CFP, i.e., the timing window during which the protocol operates). $P_{R>OD}$ applies to the first version of the protocol;

ii) ii) $P_{D>OD}$, which indicates the probability that a *decision message* (i.e., a message issued by the AP to commit or abort the delivery of a *broadcast message*) misses to be received by at least one station, again evaluated in the interval of time T_{CFP}. This measure is relative to the second version of the protocol.

Both $P_{R>OD}$ and $P_{D>OD}$ represent an estimation of the probability, for the protocols, to fail in an undetected way, a very undesirable event with possibly catastrophic consequences on the system and its users (we say, the protocol experiences a *catastrophic failure*).

The performance figures have to determine the technical limitations imposed by the communication system and the way the protocol behaves according to them. They are:

i) R_m, which indicates the average number of retransmissions for a single message;

ii) the throughput, as the number of delivered messages per second,

iii) for the second version of the protocol only, P_{UM} which indicates the probability that the AP does not receive acknowledgements on a message by all the stations in *res(c)* retransmissions, and therefore broadcasts to the active stations the decision not to deliver that message to the applications.

R_m and throughput are typical performance indicators, with R_m also useful to properly tune the protocol parameter *res(c)*. P_{UM} gives an indication of the extent of the violation of the validity property (point i) in section 2.1); again, desired values of P_{UM} can be obtained by proper tuning of *res(c)*.

3. Approach to Modeling

The behavior of the two versions of the communication protocol has been modeled by Stochastic Activity Networks (SAN) [9]. Instead of defining one single model for each version of the protocol, from which to derive all the QoS indicators identified, our choice has been to define two SANs tailored for the evaluation of specific measures. This allows limiting both the complexity and the size of the resulting models, with obvious benefits.

Before defining the assumptions made for the modeling and the models themselves, we highlight the approach we followed for capturing one of the main problems related to data block transmission in a wireless context: the impact of fading.

Fading is caused by interference between two or more versions of the transmitted signal which arrive at the receiver at slightly different times following several different paths. In urban areas, fading occurs because there is no single line-of-sight path between a mobile antenna and the base station (e.g. because of the difference of height between the mobile antenna and the surrounding structures). Even when a line-of-sight exists, the reflections from the ground and the surrounding structures cause the fading phenomenon. Another cause of the fading is the relative motion between the transmitting and the receiving antennas which originates a different shift in the received signal frequency called Doppler shift [8].

The model we considered for the representation of data transmission on fading channels is the one proposed by Zorzi et al. in [2, 13] This model considers the so called *flat fading* channel with relatively high data-rates (hundreds of kbit/s) and data blocks of hundreds of bits. In the literature this channel is modeled as a Gaussian random process and its correlation properties depend only on the normalized Doppler frequency $F_D \bullet T$ where:

F_D = Doppler Frequency = (speed of the antenna) / (signal wave-length) and

T = packet size / Data rate.

If $F_D \bullet T$ is < 0.1 then the process is very correlated. If $F_D \bullet T$ is > 0.2 then the correlation is practically negligible. The fading channel can be approximated by means of a first-order Markov model , depicted in Figure 1.

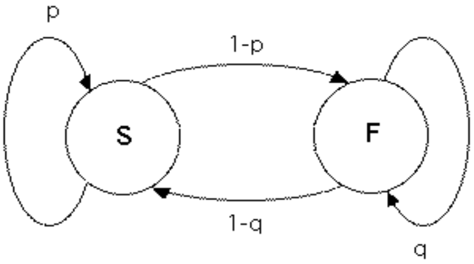

Figure 1: Markov Model of a Fading Channel

The transition matrix M(x) = M(1)^x that describes the Markov process is

$$M(x) = \begin{pmatrix} p(x) & 1 - p(x) \\ 1 - q(x) & q(x) \end{pmatrix}; M(1) = \begin{pmatrix} p & 1 - p \\ 1 - q & q \end{pmatrix}$$

where p(x) (q(x)) is the probability that the transmission i-th is successful (unsuccessful) given that the transmission (i-x)th was successful (unsuccessful). Note that 1/(1-q) is the average length of a burst of errors, while the steady state probability that an error occurs (that is that the process is in state F) is given by PE= (1-p)/(2-p-q)

Parameters p and q depend on the fading model and on the characteristics of the communication scheme. This model of correlation among successive transmissions has been included in the SAN models defined for the protocols and appropriate values for p and q have been derived using the experimental data available.

3.1. Assumptions

The assumptions under which the models have been defined are the following:

1) the time-bound for sending a message over the network is fixed and denoted with t_m. It represents a bound for both a) the time to exchange a message between two agents in the network (namely, the AP and any other mobile station), and b) the time to broadcast a message from the AP to all the other stations;

2) failures considered are only those affecting the messages, which may fail to be received by the mobile stations and/or by the AP (omission failure). Mobile stations are therefore reliable. However, a station may migrate from the group. The AP is assumed to be stable and reachable by all the stations belonging to the group;

3) each message exchanged among system components has the same marginal probability of failure PE. Message failures are correlated following the model described above;

4) the value of *res(c)* is the same for all the messages;

5) the models for the evaluation of the dependability-related figures assume that the group membership remains the same during the whole T_{CFP} interval, that is, no station misses *OD*+1 consecutive *poll-requests*, which is the condition for the AP to consider that station as migrated from the group;

6) failures of mobile stations are independent from each other.

4. The Models

Before introducing the models, we describe the way fading has been taken into account in the model themselves. As already said in Section 3, the channel alternates between two states (F, failure, and S, success) and the matrix M(1) gives the transition probabilities in one step.

Using the SAN formalism, this behavior can be modeled by a place (that we will call "SUCCESS") and by an activity whose case probabilities depend on the marking of this place. The marking of SUCCESS (0 or 1 in our models) will represent the state of the channel (F and S, respectively). The probabilities to be associated to the cases of the activity representing the reception of a message are derived from the probabilities associated to the transitions of the Markov Model in Figure 1. As an example the probability for the case accounting for a failure in the reception of a message, expressed in a C-like syntax, is:

if (marking of SUCCESS == 1) *return (1-P)* *else return q*

Once the outcome of the activity is determined, the subsequent action (executed by the corresponding output gate) consists in changing the marking of SUCCESS in a consistent way.

4.1. Model for evaluating $P_{D>OD}$, $P_{R>OD}$ and P_{UM}

We have defined a single model for the evaluation of $P_{D>OD}$, $P_{R>OD}$ and P_{UM}; from this model, the three indicators can be obtained by simply changing the value of some parameters.

Let's start considering the evaluation of $P_{D>OD}$. Since we are interested in *decision messages* (i.e., messages broadcast by the AP to commit or abort the delivery of a *broadcast message*), we have to consider the reception of consecutive messages by each station.

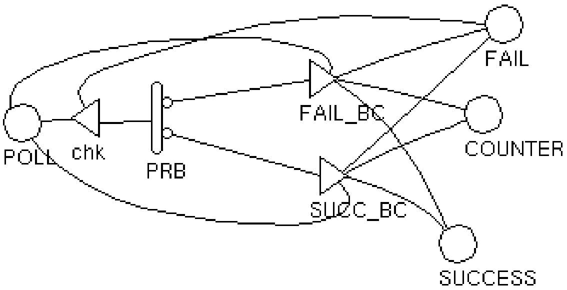

Figure 2: Sub-model station$_i$ used to evaluate $P_{D>OD}$, $P_{R>OD}$ and P_{UM}

In Figure 2, the sub-model representing the reception of a message by one generic station is shown. The overall model is obtained by replicating N-times this sub-model, where N is the number of stations belonging to the system; the place common to all the submodels is FAIL. A brief description of this model is given in the following.

The activity PRB represents the execution of the three actions: i) the AP polls a station; ii) the polled station sends to the AP a broadcast request message in reply to the poll; iii) the AP broadcasts the received message. These actions require the exchange of messages, which may be affected by the fading phenomenon. Because of the short information contained, the “poll” message is shorter than the “broadcast request message” and the “broadcast message” (these last two being approximately of the same length). Define M' as the transition matrix for the “poll” message and M the transition matrix for the other messages (as explained in Section 3). We obtain the probabilities of changing state (or of remaining in the same one) after the three actions, as the product of the matrices M'*M*M. These probabilities are associated to the cases of the activity PRB. If a failure occurs (case1) the output gate FAIL_BC will add one to the marking of COUNTER (this place traces the number of consecutive failures) and, unless the new marking exceeds OD, it also sets the marking of POLL to 1. If COUNTER exceeds OD, a token is put in FAIL and this event will stop any further action in the submodel. Moreover, since FAIL is in common with all the submodels, and since the input gate chk enables the activity PRB only when one token is in POLL and no tokens are in FAIL, all the submodels will stop their activity. When a success occurs, the marking of COUNTER is set to zero, and a new poll can be executed. From this model, $P_{D>OD}$ is obtained through a transient analysis at time T_{CFP}, by the use of a rate reward variable associated with the presence of a token in the place FAIL.

The evaluation of $P_{R>OD}$ shifts our attention to the reception of messages which are broadcast once per round. In fact, the event we are now interested in is the reception by the generic station_i of the same broadcast message, relative to station_j, which is broadcast by the AP once per round (in correspondence to the polling of station_j). The model in Figure 2 can be used to evaluate the occurrence of this event, by properly setting the rate of the activity PBR to a round duration (i.e., 3tm*N). Accordingly, keeping into account the fading phenomenon, the probabilities to change state or remain in the same after the three actions "poll-request-broadcast" have to be determined on the basis of a round interval. Therefore, the transition matrix which determines such probabilities is given by the product M'*M*M performed N times, that is, [M'*M*M]^N.

Concerning the indicator P_{UM}, it can be determined exactly as $P_{R>OD}$, but considering res(c) instead of OD (i.e., the threshold for COUNTER to put a token in FAIL is res(c))

4.3. Model for the evaluation of R_m.

The model for the evaluation of R_m has to take into account two different reasons of retransmission: a failure in the exchange of the couple poll-request message, and, once the couple has been successfully exchanged, a failure in the reception of the broadcast message. These are the reasons why, in this model, we have to distinguish between the station that wants to broadcast a message (sub-model *tx*) and the receiving ones (sub-model *rx*). By a *join* of the *tx* model and the *replication* of *rx* we obtain the model for the evaluation of R_m.

The model for the evaluation of R_m is derived from that used to evaluate $P_{R>OD}$. In fact, we are still interested in considering what happens to a specific broadcast message, so the related events occur once per round.

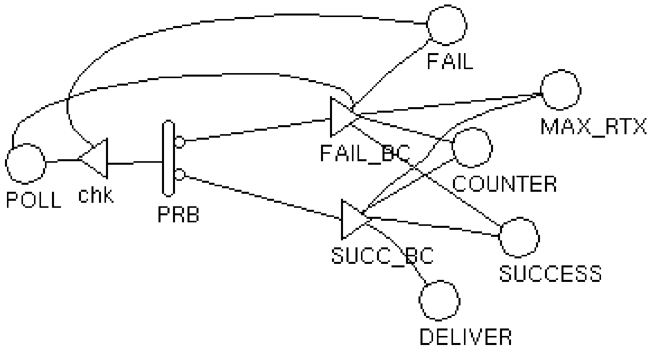

Figure 3: Sub-model station_i used to evaluate R_m

Figure 3 shows the submodel relative to the generic station_i used to evaluate R_m. Again, through the replicate operation, a (parametric) number of SANs as described in Figure3 are combined to form the complete SAN of our system, the places MAX_RTX and DELIVER being in common among all the sub-models.

To evaluate R_m, we need to know the maximum value reached by COUNTER for each broadcast message; for this reason, both the output gates FAIL_BC and SUCC_BC will update the common place MAX_RTX. Moreover, once the broadcast succeeds, the marking of DELIVER (again common to all the stations) is incremented (DELIVER counts the number of stations that successfully received the broadcast message). The number of average retransmissions will be given by the marking of MAX_RTX.

5. Derivation of input values from experimental settings.

This section describes the derivation of the reference values for the model parameters from an experimental setting and experimental measurements performed in [3]. An implementation of the second version of the protocol was set up on a system of Windows NT 4.0 Workstations and Laptops connected by an IEEE 802.11 Standard compliant wireless network. The settings were as follows:

Carrying frequency: 2.4 GHz

Packet size: 100-1000 bytes

Data Rate: 2 Mbit/sec

Some experiments have been carried out in an office environment under good physical conditions providing the following results:

Marginal probability of packet loss (PE): 1,60E-04

Time-bound for a message transmission: 7646 µsec (1000bytes), 2843 µsec (100bytes)

From these data it has been possible to derive values for our parameters, especially those related to the correlation of packet loss (p and q of Figure 1). A commonly adopted approximation in the presence of coding for data block transmission [6] considers the success determined by comparing the signal power to a threshold: if the received power is above a certain threshold the block is successfully decoded with probability 1, otherwise it is lost with probability 1. This threshold is sometimes called *fading margin F*. When a Rayleigh fading channel is considered, PE and q can be calculated as in [5]

$$P_E = 1 - e^{-1/F} \qquad 1 - q = \frac{Q(\theta, \rho\theta) - Q(\rho\theta, \theta)}{e^{1/F} - 1}$$

where

$$\theta = \sqrt{\frac{2/F}{1-\rho^2}}$$

$\rho = J_0(2\pi F_D * T)$ and

Q(.,.) is the Marcum Q function. J_0 is the modified Bessel function of 0-th order.

Recalling the equations for F_D and T reported in section 3, given the packet size, the speed of the mobile stations and the marginal error probability, one can compute F_D*T, p and q. A few values for packets of 1000 bytes are reported in Table 1.

Speed (m/s)	PE	F_D*T	p	q
0.05	1,60E-04	3,00E-03	0,9999125	0,4531810
0.5	1,60E-04	3,00E-02	0,9998414	0,0089858
5	1,60E-04	3,00E-01	0,9998400	0,0001622

Table 1. Derivation of F_D*T and of the correlation parameters p and q.

In the subsequent analyses we decided to consider also different values for PE and computed the corresponding values for p and q.

Table 2 summarizes the internal parameters of the protocols and of the models, together with the default values used in the subsequent numerical analysis (unless otherwise specified), as determined through experimental measurements.

Notation	Description	Value
1-q	probability of transitioning from the Failure state to the Success state	*f*(PE, F_D*T)
1-p	transition probability from the Success state to the Failure state	*f*(PE, F_D*T)
PE	probability of losing a message	[1.6E-4, 1E-2]
N	mobile stations in the group	4
t_m	Time-bound for a message transmission	7646 µsec
t_p	Time-bound for a poll message transmission	2380 µsec
OD	omission degree	[2, 10]
F_D*T	Normalized Doppler Frequency	[3E-3, 3E-1]
T_{CFP}	duration of a Contention Free Period	2400sec

Table 2. Notations and definitions

6. Numerical evaluations

A numerical evaluation of the SAN models presented in section 4 has been carried out, by using the tool UltraSAN [10].

Figures from 4 to 7 concern dependability-related indicators. Figure 4 and 5 show the values of the probability $P_{D>OD}$ as a function of PE and for a varying omission degree *OD* (the observation interval Tcfp is set to 2400 sec), each one for a different value of F_D*T.

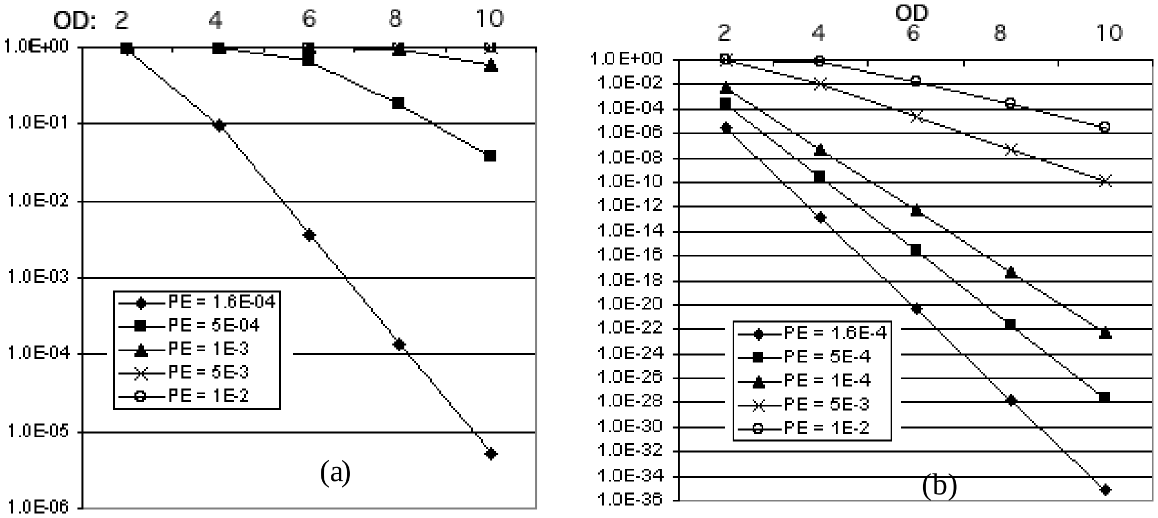

Figure 4: $P_{D>OD}$ for F_D*T = 3.00E-3 (a), $P_{D>OD}$ for F_D*T = 3.00E-2 (b)

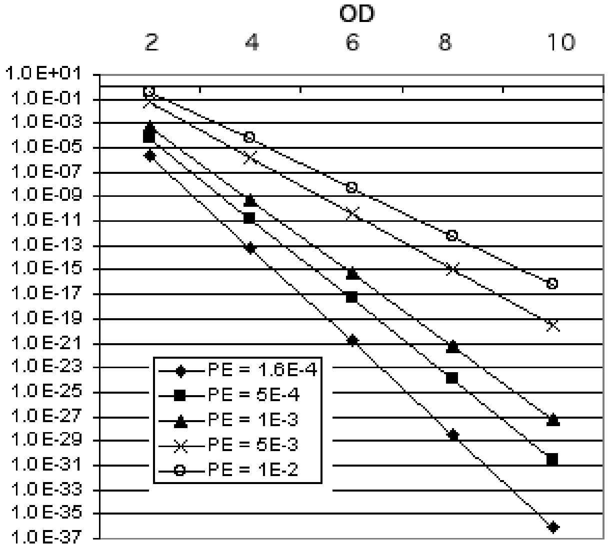

Figure 5: $P_{D>OD}$ for F_D*T = 3.00E-1

When PE increases, the values assumed by $P_{D>OD}$ increase, whereas for the same PE higher values of F_D*T determine lower values for $P_{D>OD}$. When F_D*T is equal to 3E-3, the fading phenomenon is quite strong and it affects the consecutive transmissions of messages, so to result in a higher failure probability for the protocol. The effects of the fading decreases as F_D*T increases, so that when it reaches the value of 3E-01, the correlation among the messages losses is negligible.

Figures 6 and 7 report the values obtained for $P_{R>OD}$ following the same approach. It can be noted that values for $P_{R>OD}$ are lower than those obtained for $P_{D>OD}$ for the same settings. This is due to the different time interval between retransmissions of two consecutive decision messages and broadcast messages, as already discussed in Section 4 so the fading does not affect the transmission of messages in the same way.

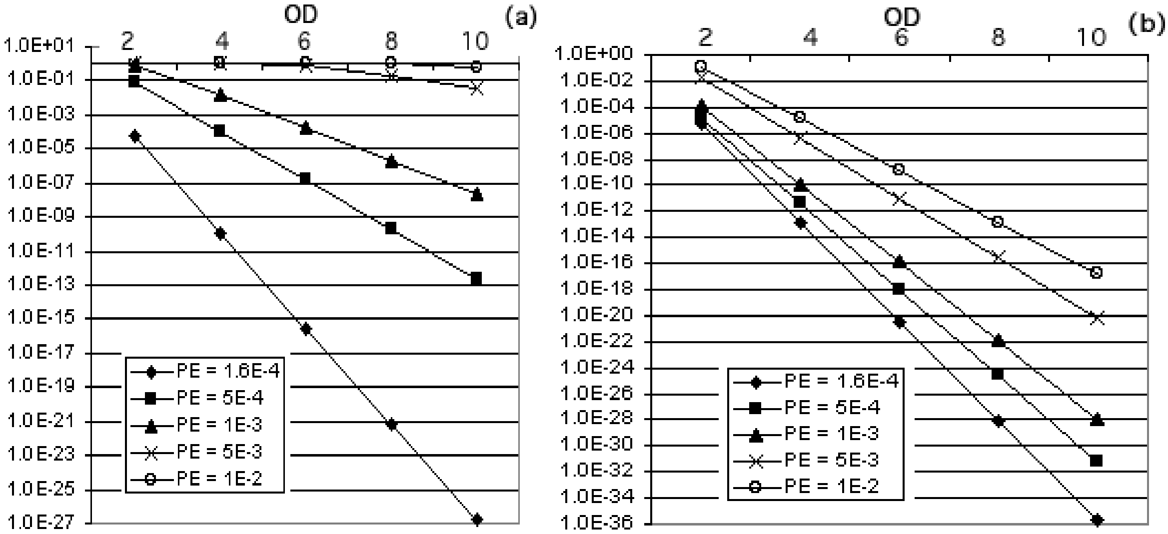

Figure 6a: $P_{R>OD}$ -- F_D*T = 3.00E-3 (a), $P_{R>OD}$ -- F_D*T = 3.00E-2 (b)

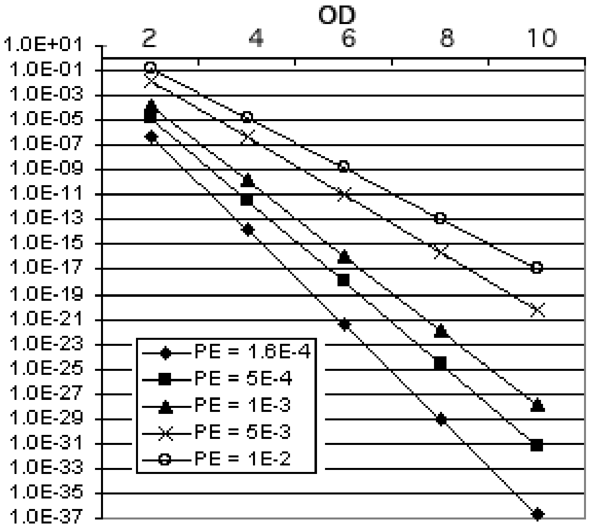

Figure 7: $P_{R>OD}$ -- F_D*T = 3.00E-1

Now we present the evaluation performed for P_{UM}. The values for P_{UM} have been obtained from the model for the evaluation of $P_{R>OD}$ by simply using the parameter res(c) instead of OD. In Figure 8 the values for P_{UM} with res(c) = 1..8, PE = (1.6E-4, 1E-2) and F_D*T = 3.00E-2 are shown. Obviously the same considerations already made for $P_{R>OD}$ apply.

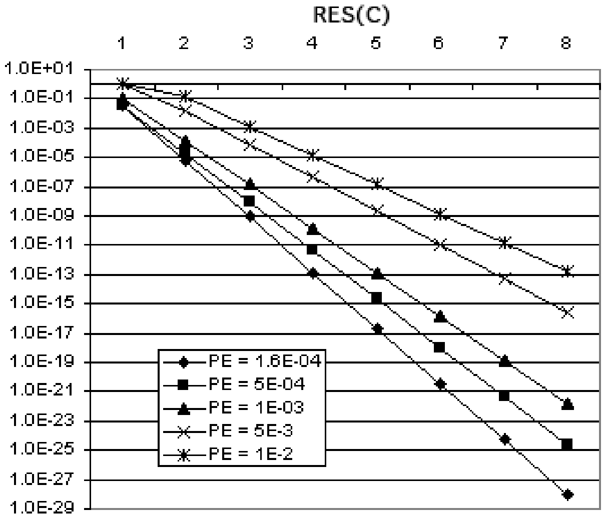

Figure 8: P_{UM} for varying values of res(c) (F_D*T = 3.00E-2)

Last we consider R_m and the throughput. The evaluation of the throughput is based on the average number of message retransmissions R_m and the average message delay t_m. Therefore, in our settings the expression for the throughput depends only on R_m since we assumed a constant message delay. The values for the throughput are given by: $\mathbf{1/((t_p + 2t_m)*R_m)}$.

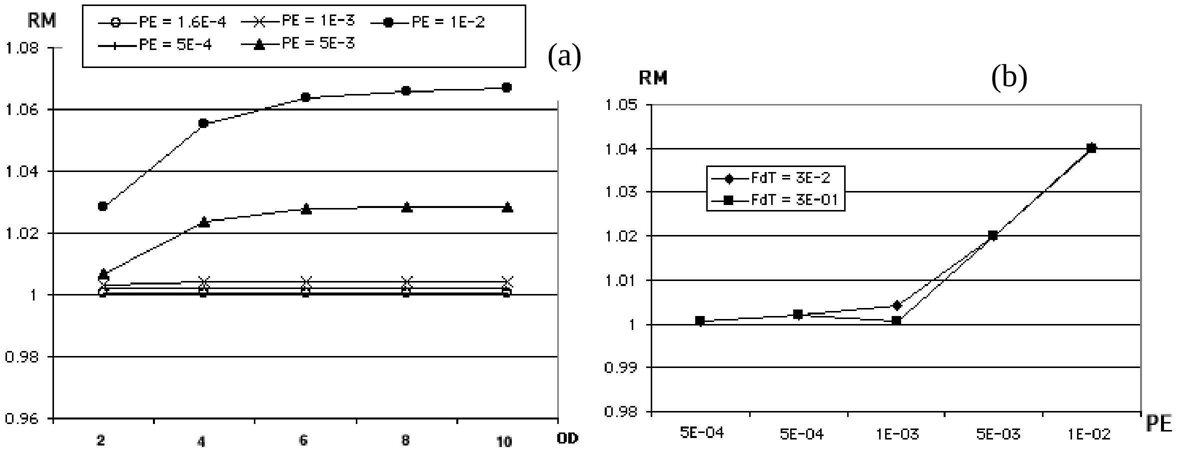

Figure 9a: R_m for different values of PE and OD (F_D*T = 3E-03) (a); R_m for different values of PE and F_D*T (b)

In our setting, the results obtained R_m vary depending on F_D*T. Figure 9a, (F_D*T = 3E-03) shows the variations on R_m at varying OD and PE. We observe while it depends on PE, some dependence on OD is observed only for high PE (that is in the left-top part of the Figure). When F_D*T = 3E-02 and F_D*T = 3E-01 as depicted in Figure 9b, OD does not impact on the obtained results (and their range of variations is very narrow [1.002, 1.04].

As a final observation we report together, in Figure 10, the values of $P_{D>OD}$ and $P_{R>OD}$ on a system with the same stations (2, 4, 6, 8 and 10) and parameter setting: F_D*T = 3E-2; OD = 5; PE = 5E-4

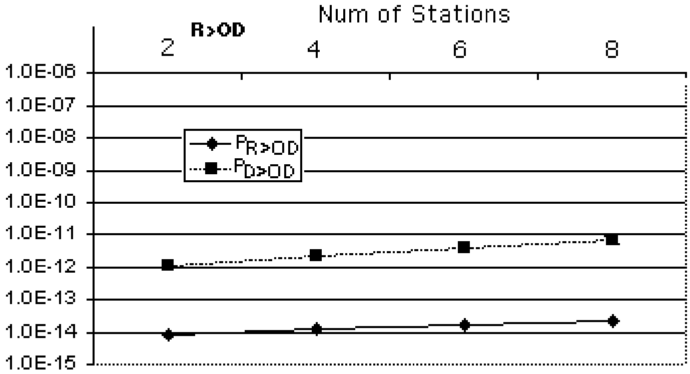

Figure 10: $P_{D>OD}$ and $P_{R>OD}$ for varying number of stations.

It is apparent that correlation makes the second protocol less resilient than the first. The correlation determines a higher failure probability since the second protocol uses consecutive messages while the first makes use of one message per cycle.

7. Concluding remarks

In this paper, we have performed an analysis of the QoS provided by family of group communication protocols in an experimental setting. The QoS metrics identified relate to both dependability and performance. Specifically, the dependability-related figures aim at giving an estimate of the coverage of the assumptions on which the protocols rely, while the performance figures can be used as indicators of the technical limitations imposed by the communication system and the way the protocol behaves according to them. We adopted an analytical approach and introduced models closely representing the system and the environment. They account for physical characteristics such as the *fading channel* phenomenon, and for user mobility. Both of them affect a wireless communication and cause time correlation among successive messages, which is captured by our models.

We used experimental data previously collected in a representative context to provide parameters values for our models. Then we performed several evaluations to highlight the behavior of the protocols depending on their settings and on the environment characteristics.

Several variations of the protocol may be devised (beside the two we have analyzed) providing a different trade-off between performance, delay time, reliability and formal properties of the broadcast. Our approach to the analysis of QoS being fast, cost effective, and formally sound allows to exploit this flexibility of the protocol suite.

References

[1] IEEE 802.11, "IEEE 802.11: Wireless LAN Medium Acces Control (MAC) and Physical Layer (PHY) specifications," 1997.

[2] A. Chockalingam, M. Zorzi and R.R. Rao, "Performance of TCP on Wireless Fading Links with Memory," in Proc. IEEE ICC'98, 1998, pp.

[3] A. Coccoli, S. Schemmer, F. Di Giandomenico, M. Mock and A. Bondavalli, "Analysis of Group Communication Protocols to Assess Quality of Service Properties," in Proc. HASE00 - 5th IEEE High Assurance System Engineering Symposium, Albuquerque, NM, USA, 2000, pp. 247-256.

[4] V. Hadzilacos and S. Toueg, "Fault-tolerant Broadcasts and Related Problems," in "Distributed Systems", S. J. Mullender Ed., Reading, Addison-Wesley, 1993, pp. 97-145.

[5] K.S.Miller, "Multidimensional Gaussian Distributions," New York, 1964.

[6] L.F.Chang, "Throughput Estimation Of Arq Protocols For A Rayleigh Fading Channel Using Fade- And Interfade-Durations Statistics," IEEE Trans. Veh. Tech., Vol. VT-40, pp. 23-229, 1991.

[7] M. Mock, E. Nett and S. Schemmer, "Efficient Reliable Real-Time Group Communication for Wireless Local Area Networks," in Proc. 3rd European Dependable Computimg Conference, Prague, Czech Republic, 1999, pp. 380-397.

[8] T.S. Rappaport, "Wireless Communications - Principles and Practice," 1996.

[9] W. H. Sanders and J. F. Meyer, "A Unified Approach for Specifying Measures of Performance, Dependability and Performability," in "Dependable Computing for Critical Applications, Vol. 4: of Dependable Computing and Fault-Tolerant Systems", Ed., Springer-Verlag, 1991, pp. 215-237.

[10] W. H. Sanders, W. D. Obal, M. A. Qureshi and F. K. Widjanarko, "The UltraSAN Modeling Environment," Performance Evaluation Journal, special issue on Performance Modeling Tools, Vol. 24, pp. 89-115, 1995.

[11] B. Teitelbaum, J. Sikora and T. Hanss, "Quality of Service for Internet2," in Proc. First Internet2 Joint Applications/Engineering Workshop: Enabling Advanced Applications Through QoS, Santa Clara, CA, 1998, pp. 5-16.

[12] A. Vogel, B. Kerherve and G. Von Bochmann, "Distributed Multimedia and QOS: A Survey," IEEE Multimedia, Vol. 2, pp. 10-19,

[13] M. Zorzi, R.R. Rao and L.B. Milstein, "On The Accuracy Of A First-Order Markov Model For Data Block Transmission On Fading Channels," in Proc. IEEE ICUPC'95, 1995, pp. 211-215.