Increasing the Throughput of the HomePNA MAC Protocol

Aurelio Amodei Jr., Luís Henrique M. K. Costa and Otto Carlos M. B. Duarte
Grupo de Teleinformática e Automação - PEE/COPPE - DEL/POLI
Universidade Federal do Rio de Janeiro
P.O. Box 68504 - CEP 21945-970 - Rio de Janeiro, Brazil

Abstract

This paper proposes a new mechanism to increase the HomePNA 2.0 MAC protocol throughput. First, we review the HomePNA MAC protocol and its collision resolution mechanism. Then, we use simulations to evaluate the throughput of HomePNA using a module that we implemented in the ns-2 simulator. We propose a priority aggregation mechanism that uses the 8 HomePNA priority levels in a more efficient way. The simulation results show that the proposed mechanism is able to increase the throughput up to 44%. Moreover, our mechanism does not require modifying the HomePNA specification, and can be easily implemented as a new sublayer above the MAC sublayer.

1. Introduction

Home networks aim to interconnect computers and network devices inside the houses, and to share Internet access and resources, such as printers and scanners. In 1998, a group of companies created the Home Phone-line Network Alliance (HomePNA) [10] to standardize home phone line networking technologies. Using the phone lines already deployed in the house does not require any new cabling infrastructure, reducing costs. This technology is usually referred to as “no new wires” technology, and it also includes the use of power lines and TV cables.

After delivering a first generation 1 Mbps standard in 1998, in the end of 1999 the group announced its second generation, HomePNA 2.0, which we focus on in this paper. The HomePNA specification [9] defines the physical layer and the MAC (Medium Access Control) sublayer, and provides transmission rates up to 32 Mbps. Based upon this specification, the ITU-T created the specifications G.989.1 [11], G.989.2 [12] and G.989.3 [13].

The HomePNA MAC protocol is an extension of CSMA/CD (Carrier Sense Multiple Access with Collision Detection) with some Quality of Service (QoS) support, such as eight priority levels for different classes of traffic. This traffic prioritization is absolute, thus a higher priority flow will always have preferential access to the medium over a lower priority one [4]. HomePNA also includes a new collision resolution mechanism, which is based on the ternary-tree algorithm [8].

In this paper, we present a new mechanism that increases the throughput of the HomePNA MAC protocol. This mechanism can be implemented without modifying the HomePNA standard, as an independent sublayer above the MAC sublayer. HomePNA defines eight priority levels. The idea is to aggregate some of the higher priorities in the highest one. Then, the outgoing high priority frames are assigned one of these priorities randomly. Since we spread the high priority frames over more than one priority level, many collisions can be avoided, specially as the network load increases. We implemented the HomePNA standard and the proposed mechanism in the ns-2 simulator (Network Simulator version 2) [6].

This paper is organized as follows. The related work is presented in Section 2. In Section 3, we present the HomePNA MAC protocol together with a throughput analysis. Section 4 presents the proposed mechanism. The simulation results are presented in Section 5, and in Section 6 we present some implementation details. Section 7 concludes the paper.

2. Related Work

Most work on HomePNA is related to channel modeling and to the physical layer. Bisaglia et al. [3] present a channel modeling and analyze the HomePNA physical layer. In [1] and [2], Bisaglia and Castle analyze HomePNA receiver architectures and equalization techniques.

Chung et al. [5] present a mathematical analysis of the HomePNA saturation throughput. The saturation throughput is the maximum throughput one can get in the network for a given number of nodes, when all nodes always have a frame to transmit. In this paper, we make this analysis by simulation, and use it as the basis for comparison of the proposed mechanism.

Kangude et al. [14] analyze the HomePNA MAC protocol and its collision resolution mechanism. The analysis shows that the number of collision resolution (CR) slots, defined as 3 by the specification, is not optimal in some cases. They suggest that a mechanism could be used to change the number of CR slots in cases where it is sub-optimal, to increase the efficiency of the MAC protocol. However, changing this value is difficult, as it is an established standard, and interoperability should also be considered.

3. The HomePNA MAC Protocol

The HomePNA MAC protocol is based on the Ethernet Carrier Sense Multiple Access with Collision Detection (CSMA/CD) for medium access control. It includes a priority mechanism with eight levels of priority for QoS support. Different classes of traffic can be labeled with priorities from 0 to 7, where 7 is the highest one. Based on the frame priority, its transmission occurs in a corresponding time interval preceded by the inter-frame gap (IFG) that separates the frames. The IFG lasts 29 μs.

As shown in Figure 1, the time intervals are organized in decreasing order of priority. Thus, higher priority flows start transmitting earlier, without contending with lower priority traffic. This results in an absolute traffic prioritization, similar to some priority schemes used in IEEE 802.11 ad hoc networks [15]. The duration of each priority slot, PRI_SLOT, is 21 μs. Stations must transmit their frames at the beginning of a slot whose number is equal to or less than the frame priority. Any transmission after slot 0 is assumed to happen at slot 0.

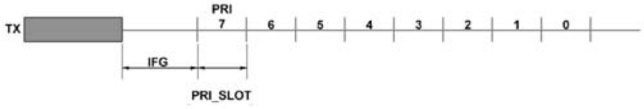

Figure 1. HomePNA priority slots.

Stations with frames to transmit must sense the carrier and defer the transmission if any carrier is detected before the time slot corresponding to the frame priority. In this case, the time slot counting is restarted.

All stations monitor the medium to detect collisions between the frames of other stations. A collision is detected through the transmission duration. The minimum duration of a valid frame is 92.5 μs and the maximum duration is 3122 μs. Any station that detects its frame collision shall cease transmission no later than 70 μs after the beginning of this frame. Therefore, any frame fragment too short or too long is considered as a collision.

If there is a collision, all stations start a distributed collision resolution algorithm called Distributed Fair Priority Queuing (DFPQ) [7]. After the algorithm execution, all stations involved in the collision are ordered in Backoff Levels (BL), which indicate the order these stations will transmit. The desired outcome is for only one station to be at BL 0, enabling it to acquire the channel. After a successful transmission, all other stations decrement their BL, and the new station(s) at BL 0 try to transmit. All stations, even the idle ones, must monitor the activity on the medium to keep track of the Maximum Backoff Level (MBL). By monitoring the MBL, stations that did not collide are not allowed to contend for access to the medium until all stations that collided have transmitted one frame successfully. The only exception is when a station has a frame with priority higher then the priority slot where the collision occurred. All stations must have 8 BL and 8 MBL counters, one for each priority.

As shown in Figure 2, after a collision occurs, there are 3 special collision resolution signaling slots (S0 to S2) before the priority slots. The BL and MBL counters are determined through these signaling slots. These slots have a duration of 32 μs, and are used only after a collision.

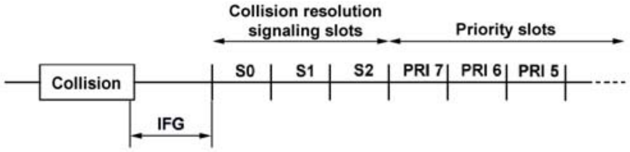

Figure 2. Collision resolution signaling slots.

After a collision, active stations - the ones involved on the collision resolution - randomly choose one of the three signaling slots to transmit a backoff signal. More than one station can transmit a backoff signal in the same signaling slot. If an active station sees a backoff signal in a slot preceding the one the station chose, it increments its BL counter. The MBL counter is incremented for each backoff signal seen and decremented when a successful transmission occurs. The MBL counter is non-zero whenever a collision resolution cycle is in progress. Passive stations keep their BL counter equal to the MBL counter, so that they transmit only when the collision resolution cycle ends.

The HomePNA standard can adaptively use rates from 4 to 32 Mbps, according to the channel conditions. This is the rate of transmission used only for the payload. The header and the trailer are always transmitted at 4 Mbps, with a more robust modulation and symbol rate, so that all stations can receive them correctly. The MAC frame is shown in Figure 3.

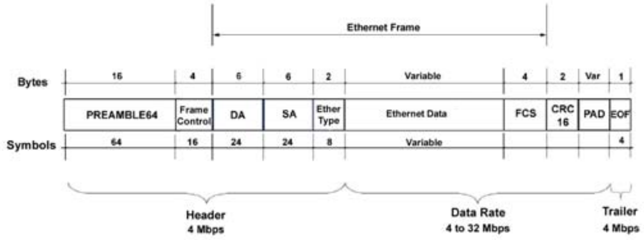

Figure 3. HomePNA MAC frame.

The frame is based on the IEEE 802.3 Ethernet frame. The Ethernet frame is preceded by a preamble of 16 bytes and a frame control field of 4 bytes, and followed by a Cyclic Redundancy Check field (CRC16) of 2 bytes, a variable-length padding and an end-of-frame (EOF) sequence of 1 byte. The padding is used when the transmission time of the complete frame is less than 92.5μs to guarantee the minimum valid frame duration, in order that this frame is not misinterpreted as a collision fragment.

3.1. Throughput Analysis

In [5], Chung et al. present a mathematical analysis of the HomePNA saturation throughput, i.e., where all transmitting nodes always have a frame to transmit, immediately after each successful transmission. We implemented a HomePNA module in the ns-2 simulator, and made the same analysis [5] by means of a simulation study. In the following, we compare the obtained results. We considered a simple physical layer, similar to the one used by the ns-2 simulator for Ethernet.

All the simulations had a duration of 100 seconds, and used 4 frame sizes, 160, 500, 1000, and 1500 bytes. The scenario is composed of one receiver node for all the transmitter nodes. For all results we calculated confidence intervals of 95% relative to the average. These intervals are represented as vertical error bars in the graphics.

Figure 4 shows the results for the 32 Mbps data transmission rate, varying the number of nodes from 1 to 30. We observed a small difference between these curves and the theoretical curves of [5]. In [5], the throughput increases when we go from 2 to 3 nodes, which does not happen in our analysis. This happens because in that analysis, the total transmission time of a collision resolution cycle includes one IFG preceding the beginning of the transmission, and one IFG after each collision or successful transmitted frame. Nevertheless, if there is one IFG after the last transmitted frame, the preceding IFG should not be included in this collision resolution cycle, because this IFG was already included after the last transmitted frame of the previous cycle. Except for that, the behavior of both curves is similar.

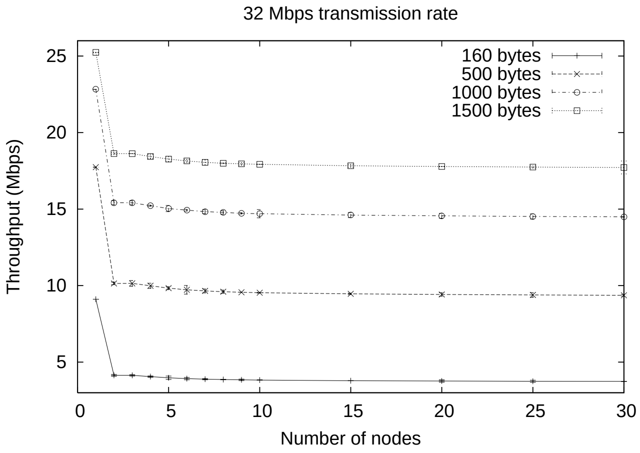

Figure 4. Throughput analysis at 32 Mbps.

Figure 5 shows the same results for a transmission rate of 4 Mbps. The MAC protocol efficiency at 4 Mbps is greater then the efficiency at 32 Mbps. For lower rates, the payload transmission time is greater, so the overhead time for header, trailer, IFG and priority slots represent a smaller fraction of the total transmission time, increasing the efficiency.

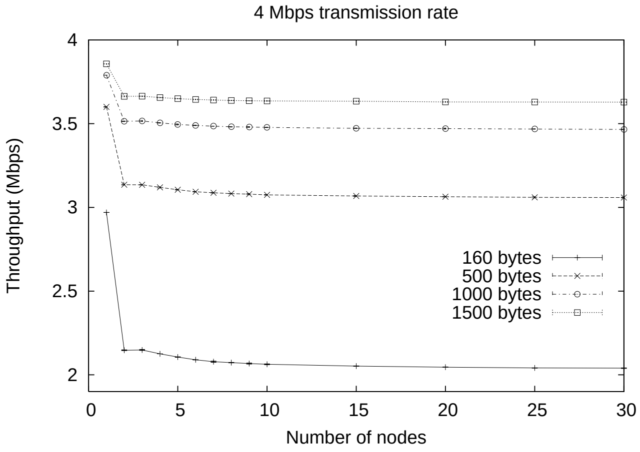

Figure 5. Throughput analysis at 4 Mbps.

We also analyze the frame size influence in the throughput. Figure 6 shows the results for a scenario with 1 transmitting node, where the frame size varies from 50 to 1500 bytes. The data transmission rate is 32 Mbps, and we use the same condition where the transmitting node always has a frame to transmit. For shorter frames, the payload transmission time is shorter, reducing the efficiency.

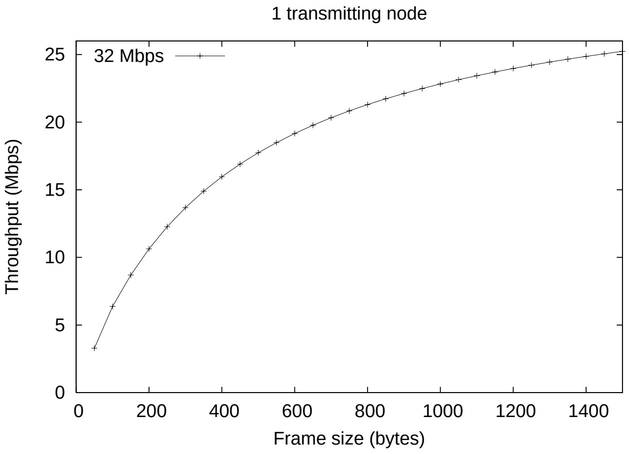

Figure 6. Frame size influence at 32 Mbps.

4. The Proposed Mechanism

We consider that in home networks, we only need few different classes of traffic, usually one for the highest priority traffic, such as video and voice, and another for the best effort traffic. At most, three different classes could be needed in some cases. Nevertheless, HomePNA defines eight priority levels, and in most cases some of these priorities are not used.

An important property of different priority frames is that they do not have to contend for the medium after a frame transmission, where the highest priority frame has preferential access. After its transmission, the other frames would be transmitted in order of priority. In general, to transmit several flows of video or voice, these flows are labeled with the highest priority, which is 7 in this case. What happens is that as the load increases, the collision probability in priority slot 7 will increase, and some time will be wasted in the resolution of the collisions.

The mechanism we propose aggregates some of these priority levels in the highest priority. The outgoing high priority frames are then mapped in one of these priorities randomly. Thus, this mechanism works like a random backoff inside the priority, where the small delay added to the frames can avoid many collisions, depending on the number of priority levels that were aggregated in the high priority.

In the same way the collision would be randomly resolved, the idea is that we can make this random choice before the transmission and try to avoid that collision. One of the benefits is that some collision resolution time will not be wasted. Hence, a network throughput increase can be achieved in the situations where it is most needed, that is, for high load and more than one station transmitting. On the other hand, the trade-off is that all frames transmitted will have to wait a few more priority slots to be transmitted, except if they get assigned priority 7. Each slot lasts 32 μs, and this extra delay could be up to 224 μs for a priority 0 frame, though we could use fewer slots for the aggregation, and the mean extra time could be reduced. The priority slots time added will reduce the efficiency in cases of few collisions, where this time will be wasted in trade of nothing. However, if the offered load is low, there is no need to increase the efficiency of the network. The only deficiency stands for the case of only one transmitting node. As there are no collisions, there will be a small degradation of the network efficiency.

As we have 8 priority slots, one of the parameters to be determined is the number of aggregated slots (AS) in the high priority. Larger values for AS would reduce the collision probability, but may also reduce the efficiency as we would have to wait more priority slots to transmit a frame. On the other hand, lower AS values would have the opposite effect. In this paper, we will consider AS varying from 2 to 7, and the simulation results will be shown in Section 5, evaluating the performance of each configuration.

5. Simulation Results

We implemented the proposed mechanism and a HomePNA simulation module in the ns-2 simulator [6] to analyze the efficiency of this mechanism, and what would be the optimal value for AS. We consider a scenario similar to the one used in the throughput analysis, where we vary the number of stations with high priority frames to transmit. Each station always has a frame to transmit. For all the results, except for the normalized curves, we calculated confidence intervals of 95% relative to the average. These intervals are represented as vertical error bars in the graphs.

Figure 7 shows the results for up to 30 nodes, 32 Mbps rate and 1500 bytes frames. The original results are plotted against the proposed mechanism results varying AS from 2 to 7. The results show that for more than 1 transmitting node a significant throughput gain is achieved for any value of AS. Moreover, we find that the optimal AS value is 4, since this configuration achieved the best throughput for any number of nodes above 1.

We can also notice the small degradation of the throughput for 1 transmitting node. With $AS = 4$, the throughput is reduced from 25.2 to 23.7 Mbps, or a 6% loss. However, this loss is relatively low, and not very significant as we usually have at least 2 transmitting nodes, so the 1 transmitting node situation is not very common.

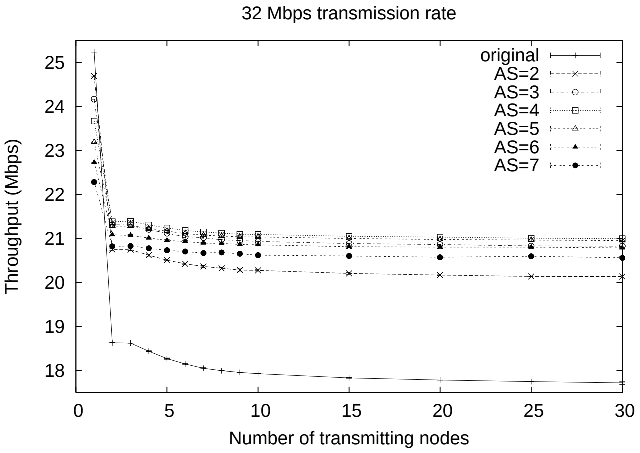

Figure 7. Throughput of the proposed mechanism against the original one for a 32 Mbps transmission rate.

In Figure 8 we show the same curves, but now normalized, dividing all the throughput curves by the original configuration throughput. This figure only shows the values from 2 to 30 nodes because that is the region of interest.

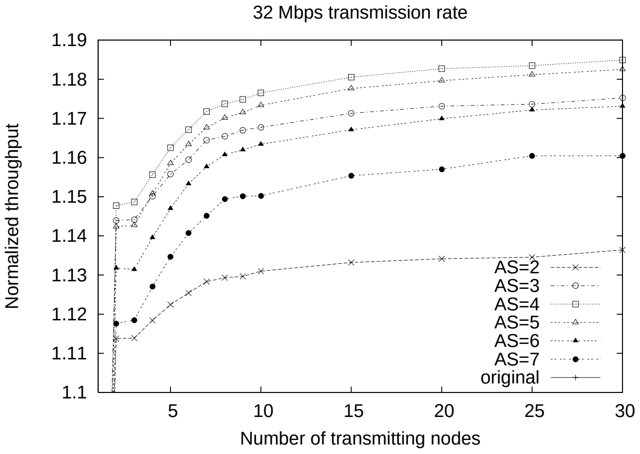

Figure 8. Normalized throughput of the proposed mechanism for a 32 Mbps transmission rate.

Note that $AS = 4$ is the best choice, laying above all the other curves. The throughput gain ranged from 14.8% for 2 nodes up to 18.6% for 30 nodes, achieving a throughput of about 21 Mbps even for a large number of nodes. We can also notice that for more than 30 nodes, the gain seems to be increasing, but 30 nodes was chosen as a practical limit for a home network. It is difficult to imagine a house with more than 30 devices connected to a network, or with more than 30 telephone outlets.

We can also notice that the throughput increases when the number of AS varies from to 2 to 4, and decreases as it varies from 4 to 7. This demonstrated that there is a tradeoff between the time added with the extra AS, and the reduction of the collision probability. Up to $AS = 4$, the collision probability reduction compensates the added time, achieving a maximum at $AS = 4$. Beyond this limit, increasing AS does not cause a significant collision probability reduction, and then the wasted time reduces the throughput.

The variation of the frame size also influences the gains achieved. For lower frame sizes, the gains increase. As the payload transmission time is shorter, the priority slots and the collision resolution time represent a larger fraction of the total transmission time of the frame. Figure 9 shows the results for the $AS = 4$ configuration as the frame size varies. We varied the frame size from 160 to 1500 and fixed the number of transmitting nodes to 15, because from 15 transmitting nodes on the original throughput tends to remain constant (Figure 4).

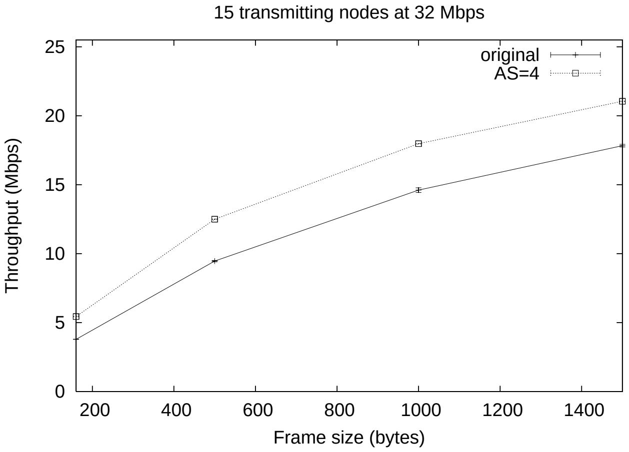

Figure 9. Throughput obtained varying the frame size.

Figure 9 shows that for smaller frames the absolute throughput may be lower, but the relative throughput gain increases. Figure 10 presents the same curves normalized by the original throughput. Note that the gains achieved increase as we use smaller frames. With 160 bytes frames we were able to achieve a throughput gain of about 44% with 15 transmitting nodes. With 30 transmitting nodes, this gain is 44.7%, and it was the maximum gain we could get for up to 30 nodes.

Another parameter that may affect the obtained gain is

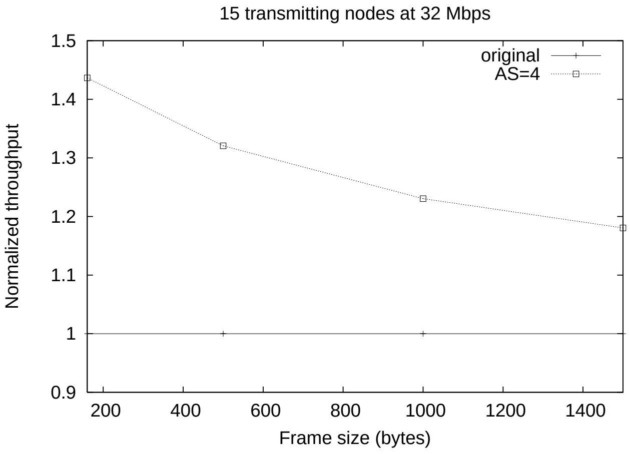

Figure 10. Normalized throughput varying the frame size.

the data transmission rate used. As this rate lowers, lower gains will be achieved, because the transmission time of the payload will increase, and the collision resolution and priority slots time will represent a smaller fraction of the total transmission time.

Figure 11 shows the throughput for the lower HomePNA data rate, 4 Mbps. Note that at this rate, the gains are lower, but the loss for only 1 transmitting node is also lower. It is important to notice that the MAC protocol efficiency at 4 Mbps is already high. For a large number of nodes it is about 90.7%, and with the proposed mechanism it is increased to 93.6%.

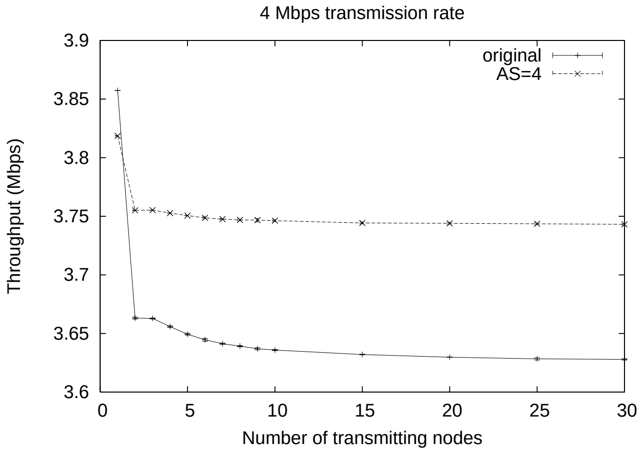

Figure 11. Throughput of the proposed mechanism for a 4 Mbps transmission rate.

The normalized throughput is shown in Figure 12. The gains obtained for this rate are only 2.5% for 2 transmitting nodes and 3.2% for 30 transmitting nodes. Nevertheless, the throughput loss for only 1 transmitting node is also reduced, with only a 1% loss. Thus, in the cases where the gains are lower, the loss for only 1 transmitting node is also lower.

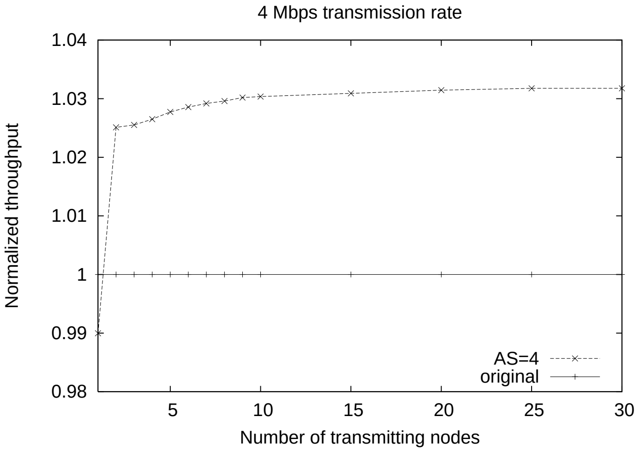

Figure 12. Normalized throughput of the proposed mechanism for a 4 Mbps transmission rate.

6. Implementation Details

A convenience of using this mechanism is the increase of the network throughput without changes to the HomePNA specification. It could be implemented in a new sublayer above the MAC sublayer. Therefore, before the packet goes to the MAC sublayer, its priority has also been randomly mapped to a correspondent MAC priority. The new sublayer can be considered as a priority mapping interface (PMI) for the MAC sublayer. Figure 13 shows a diagram of the new sublayer.

It is out of the scope of the HomePNA standard [9] to determine the priority mapping method, and it does not specify any related requirement. This mechanism could be implemented, for example, in the interface driver or in the firmware. Therefore, the idea is that the implementation would not require any hardware change or extra hardware, and the use of this mechanism could be even turned on or off.

As 4 of the MAC priority levels are mapped in the PMI highest priority, there should be a priority mapping of the remaining PMI priorities, in order to keep 8 priorities available at the PMI sublayer. Table 1 shows a suggested map-

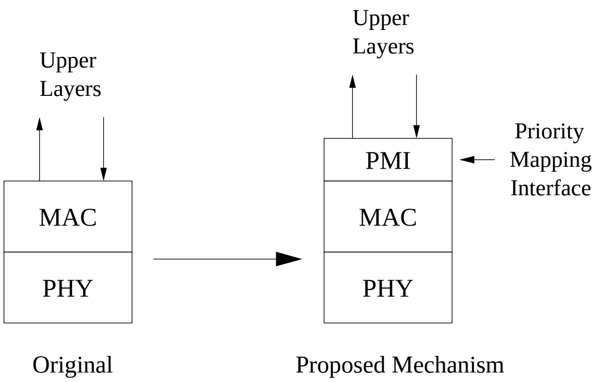

Figure 13. Implementation of the proposed mechanism.

ping for the PMI priorities. PMI priorities 6 and 5, 4 and 3, 2 and 1, and 0 are respectively mapped onto MAC priorities 3, 2, 1, and 0. Any other configuration would be possible, the only requirement is that the 4 highest MAC priorities are mapped into the PMI priority 7.

PMI priority (to the upper layers)	MAC priority (real priority)
7	7 6 5 4
6 5	3
4 3	2
2 1	1
0	0

Table 1. Priority mapping.

We also consider the mixing of interfaces that use and that do not use this mechanism in the same network. It would not be necessary that all the stations in the network implemented this mechanism, and there would still be some throughput gain, as collisions would be avoided in the same way. But if one of the stations that did not implement this mechanism wanted to use the priority 7 for an application that used all the bandwidth available, such as an FTP file transfer, there would be a problem. Any station that implemented the proposed mechanism would not be able to access the medium. At the first frame that was randomly assigned a priority below 7, that station would not access the medium anymore until the end of that application transmission, that would always use priority 7.

7. Conclusions

In this paper, we proposed a new mechanism to increase the HomePNA MAC throughput. First, we used simulations to evaluate the throughput of HomePNA using a module that we implemented in the ns-2 simulator. In the chosen scenario, all the stations always have a frame to transmit, so that we analyze the saturation throughput. We observe that the throughput slightly decreases as the number of nodes increases, tending to remain constant at approximately 15 nodes. For a 32 Mbps data rate and frames of 1500 bytes, the maximum throughput achieved is about 25.2 Mbps. When the number of nodes increases, the throughput tends to a constant value of about 18 Mbps.

The proposed mechanism uses some of the 8 priority levels available to improve the HomePNA throughput. We also implemented the proposed mechanism in ns-2, and the simulation results show that for more than one transmitting node in the network, throughput gains of up to 44.7% can be achieved. The throughput gain varies with the number of transmitting nodes and with the frame size. For higher data rates and smaller frames, the throughput gains increases, and it decreases with lower data rates and larger frames.

For a scenario of 32 Mbps transmission rate and 1500 bytes frames, the gain ranged from 14.8 to 18.6%, according to the number of transmitting nodes. Using frames of 160 bytes in a scenario with 15 nodes, we obtained a 44% gain. Furthermore, for the lower HomePNA transmission rate, 4 Mbps, the maximum gain achieved was 3.2% for 30 transmitting nodes, but on the other hand the loss for 1 transmitting node was only 1%. Hence, it is important to note that, except for only 1 transmitting node, in all the scenarios the proposed mechanism performed better than the original HomePNA MAC. Even for the 1 transmitting node scenario, the loss has not exceeded 6%.

Finally, we show that the proposed mechanism is implementable, and it does not require modifying the HomePNA specification. Furthermore, stations implementing this mechanism can operate with other stations that do not implement it on the same network. The mechanism can be implemented as a new sublayer above MAC sublayer, serving as a priority mapping interface (PMI) sublayer. It could be implemented in the software driver or firmware, therefore it could be easily deployed to HomePNA devices.

Acknowledgment

This work has been supported by CNPq, CAPES/PRODOC, FAPERJ, FUJB, COFECUB and RNP/FINEP/FUNTTEL.

References

[1] P. Bisaglia and R. Castle. A comparison of equalizer structures for frequency diverse QAM. Technical Report HPL-2001-186, Hewlett-Packard Laboratories, 2001.

[2] P. Bisaglia and R. Castle. Receiver architectures for HomePNA 2.0. Technical Report HPL-2001-247, Hewlett-Packard Laboratories, 2001.

[3] P. Bisaglia, R. Castle, and S. H. Baynham. Channel modeling and system performance for HomePNA 2.0. *IEEE Journal on Selected Areas in Communications*, 20(5):913–922, June 2002.

[4] W. Y. Chen. *Home Networking Basis: Transmission Environments and Wired/Wireless Protocols*. Prentice-Hall, 1^{st} edition, 2004. ISBN 0130165115.

[5] M. Y. Chung, H. C. Kim, and T.-J. Lee. HomePNA 2.0 - saturation throughput analysis. *IEEE Communications Letters*, 7(11):558–560, Nov. 2003.

[6] K. Fall and K. Varadhan. *The ns Manual*. UC Berkeley, LBL, USC/ISI, and Xerox PARC, Apr. 2002. http://www.isi.edu/nsnam/ns/ns-documentation.html.

[7] E. H. Frank and J. Holloway. Connecting the home with a phone line network chip set. *IEEE Micro*, 20(2):27–38, Apr. 2000.

[8] N. Golmie, Y. Saintillan, and D. H. Su. A review of contention resolution algorithms for ieee 802.14 networks. *IEEE Communication Surveys*, first quarter 1999. http://www.comsoc.org/pubs/surveys.

[9] Home Phoneline Network Alliance. Interface specification for HomePNA 2.02.7 10M8 technology. HomePNA 2.0 Standard, 1999.

[10] Home Phoneline Network Alliance. http://www.homepna.org, 2003.

[11] ITU-T. G.989.1: Phoneline networking transceivers - foundation. Feb. 2001.

[12] ITU-T. G.989.2: Phoneline networking transceivers - payload format and link layer requirements. Nov. 2001.

[13] ITU-T. G.989.3: Phoneline networking transceivers - isolation function. Mar. 2003.

[14] S. Kangude, J. Copeland, and M. Sherman. An analysis of the Home PNA collision resolution mechanism. In *28th IEEE Conference on Local Computer Networks - LCN'2003*, Bonn/Konigswinter, Germany, Oct. 2003.

[15] M. G. Rubinstein and J. F. de Rezende. Quality of service in ad hoc 802.11 networks. *Journal of the Brazilian Computer Society*, 9(1):41–52, July 2003.