

Effect of Non-Reciprocity on Infrared Wireless Local-Area Networks

Francis M. Chow and Joseph M. Kahn

Department of Electrical Engineering and Computer Sciences
University of California, Berkeley, CA 94720

Abstract

An infrared (IR) wireless network standard called Advanced Infrared (AIr) was proposed to the Infrared Data Association in April 1997. AIr provides a dynamic trade-off between bit rate and transmission range by employing four-slot pulse-position modulation (PPM) with variable-rate repetition coding (RC). At the media-access control (MAC) layer, carrier-sensing multiple access/collision avoidance (CSMA/CA) is utilized. To prevent collisions in case of hidden terminals, request-to-send and clear-to-send (RTS/CTS) packets are exchanged to reserve the channel. To insure that RTS/CTS exchange works reliably, it is necessary to maintain reciprocity, which means that the signal-to-noise ratio is symmetric between each pair of transceivers. In practice, however, reciprocity can be violated, due to poor transceiver design, manufacturing tolerances, or ambient light noise. We use a layered simulation approach to investigate the impact of non-reciprocity on the performance of IR wireless LANs. Studying the example of AIr in detail, we show that non-reciprocity can substantially reduce throughput and increase unfairness.

1. Introduction

For short-range, high-speed wireless links, infrared (IR) radiation [1],[2] offers several advantages over the traditional radio medium. IR offers a wide, unregulated bandwidth. IR radiation does not penetrate walls, permitting the same spectrum to be used in all rooms, leading to a potentially huge aggregate network capacity. An IR link employing intensity modulation with direct detection (IM/DD) does not suffer from multipath fading. IR does have several drawbacks, however. Communication from one room to another requires installation of IR access points interconnected via a wired backbone. In many environments, intense ambient light induces noise in IR receivers. In an IM/DD link, the receiver electrical signal-to-noise ratio (SNR) is proportional to the square of the received power. Hence, IR links must often transmit at relatively high power levels, and operate over a relatively limited range.

Over the past five years, the Infrared Data Association (IrDA) has established standards for short range, half-duplex line-of-sight (LOS), point-to-point links operating at bit rates up to 4 Mbps [3]. Hewlett-Packard and IBM are currently collaborating on a new proposed IrDA standard called Advanced Infrared (AIr). To date, both the AIr Physical Layer (AIr-IrPHY) [4] and the Media-Access Control (IrMAC) Layer [5] have been documented.

AIr-IrPHY [4] employs four-slot pulse-position modulation (PPM) with variable-rate repetition coding (4-PPM/VR)




Fig. 1. The use of RTS/CTS exchange to alleviate the hidden terminal problem. Without the RTS/CTS exchange, User 3 is unaware of User 1's transmission, and a collision will occur at User 4 if User 3 starts to transmit. Dashed lines denote schematically the range of angles within which the terminals transmit and receive.

to achieve a dynamic trade-off between transmission range and bit rate. The basic bit rate is 4 Mbps, and the repetition rate ranges from 1 to 16, corresponding to bit rates between 4 Mbps and 250 kbps. In AIr-IrPHY, the receiver estimates the SNR and informs the transmitter of the bit rate that can be supported.

IrMAC [5] provides for multi-access, peer-to-peer communication by using carrier-sensing multiple access with collision avoidance (CSMA/CA). In a multi-access wireless network, since each transceiver cannot generally receive from, nor even detect the presence of every other transceiver, carrier sensing alone cannot adequately regulate media access to prevent collisions at the receivers. This is the well-known "hidden terminal problem". In order to alleviate this problem, IrMAC utilizes a channel-reservation scheme based on the exchange of request-to-send and clear-to-send (RTS/CTS) packets [6], as illustrated by Fig. 1.

The effectiveness of RTS/CTS exchange in alleviating the hidden terminal problem depends on certain physical-layer properties being satisfied, including (a) the capability to reliably transmit packet headers over an extended range, and (b) the achievement of channel reciprocity, which means that the SNR between all transceivers is pairwise symmetric. In AIr, condition (a) is satisfied by using 16-fold repetition encoding on all packet headers. Satisfying condition (b) may prove more difficult, however. Deviations from reciprocity can be introduced by poor transceiver design or excessively large manufacturing tolerances. More fundamentally, different ambient light levels at various receiver locations lead to different shot-noise levels, upsetting reciprocity.




Fig. 2. Overview of the simulation flow.

Experimental measurements of the effect of non-reciprocity on packet reception are presented in [4]. To our knowledge, however, work to date has not provided a detailed analytical framework in which to study the effects of non-reciprocity, nor has it addressed the impact of non-reciprocity on the throughput of an entire realistic LAN. This paper studies non-reciprocity via analysis and simulation of the entire AIR network. The simulation methodology used in this paper is shown in Fig. 2. Through the physical layer simulator, the SNR and SIR of the transceivers as well as the required repetition rate (RR) are estimated according to the transceiver design and physical configuration of the LAN. The SNR, SIR and RR are then used to calculate the probability of successful packet transmission and carrier sensing, employing the theoretical expressions for error probability of repetition-coded PPM in the presence of co-channel interference. A discrete-event simulator is used to simulate generation and transmission of packets in the LAN, permitting us to quantify network throughput and fairness.

The remainder of this paper is organized as follows. In Section 2, we describe infrared channels and derive the conditions for channel reciprocity. In Section 3, we describe our modeling and analysis of the physical layer. In Section 4, we explain the interface between the physical layer simulation and the MAC layer simulation, and describe the MAC layer simulation itself. In Section 5, the simulation experiments and results are presented. Conclusions are given in Section 6.

In the interest of brevity, we have omitted some details of the analysis and many of our simulation results; these can be found in [7].

2. Infrared Channels and Channel Reciprocity

A. Intensity Modulation/Direct Detection Channel Model

The model of an IM/DD channel [2] is represented by

$$Y(t) = RX(t) \otimes h(t) + N(t) \quad (1)$$

where \otimes denotes convolution and R represents the receiver responsivity. The transmitted waveform $X(t)$ is the instantaneous optical power of the transmitter, while the received waveform $Y(t)$ is the instantaneous current in the photodetector of the receiver. $N(t)$ is modeled as Gaussian noise independent of $X(t)$. The average transmitted power, P_t , is

the time average of $X(t)$, and the average received optical power is given by $P = H(0)P_t$, where $H(0) = \int_{-\infty}^{\infty} h(t)dt$ is the channel d.c. gain. Assuming $N(t)$ is dominated by a white Gaussian component having double-sided power spectral density (PSD) N_o , we define the receiver electrical signal-to-noise ratio (SNR) as

$$SNR = \frac{R^2 P^2}{N_o B} \quad (2)$$

where B is the receiver noise bandwidth. When the link uses PPM with chip (time slot) duration T_c , the SNR becomes:

$$SNR = \frac{R^2 P^2 T_c}{N_o}. \quad (3)$$

B. Channel Reciprocity

To facilitate the discussion of the requirement of channel reciprocity, we define a coordinate system as shown in Fig. 3. The Cartesian coordinates (x_i, y_i, z_i) are local to transceiver i , which has its symmetry axis along z_i , while (θ_i, ϕ_i) are the spherical polar coordinates of an arbitrary direction with respect to the z_i axis. In these polar coordinates, the symmetry axis of the transceiver is simply denoted by $(0, 0)$. IR channel reciprocity is defined as the property that any pair of transceivers, irrespective of their separation and angular orientation, exhibit a pairwise-symmetric SNR, i.e.,

$$SNR_{ij} = SNR_{ji} \quad \forall i \neq j \quad (4)$$

where SNR_{ij} is the SNR achieved by receiver j when transmitter i is active. For a LOS link, the received power at receiver j when transmitter i is transmitting is given by

$$P_{ij} = \frac{I_i(\theta_i, \phi_i) A_{\text{eff}, j}(\theta_j, \phi_j)}{d_{ij}^2} \quad (5)$$

where I_i is the time-average radiant intensity of the transmitter, $A_{\text{eff}, j}$ is the effective area of the receiver, and d_{ij} is the separation. Using (2), (4) and (5), and assuming that all receivers have identical noise bandwidth B , the requirement for channel reciprocity for any user i is found to be

$$\frac{R_i A_{\text{eff}, i}(\theta_i, \phi_i)}{\sqrt{N_{o, i}}} \cdot \frac{1}{I_i(\theta_i, \phi_i)} = K \quad (6)$$

where K is a constant, independent of i , θ_i and ϕ_i , R_i is the responsivity of the photodetector and $N_{o, i}$ is the two-sided noise PSD. Note that in (6), the first factor represents properties of the receiver, while the second factor represents properties of the transmitter.

The channel reciprocity requirement can be conveniently broken down into two independent requirements [4], optical parity and system parity, from which simple trans-

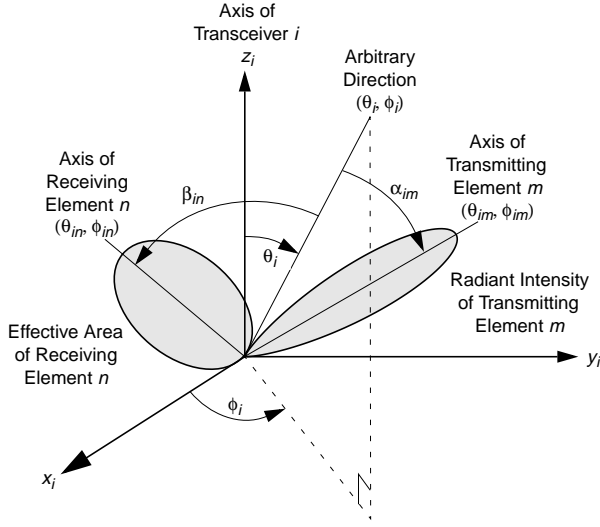


Fig. 3. Coordinate system used to describe transceivers.

ceiver design guidelines can be derived. Optical parity refers to the transceivers having identical emission and reception directivities:

$$\frac{I_i(\theta_i, \phi_i)}{I_i(0, 0)} = \frac{A_{\text{eff}, i}(\theta_i, \phi_i)}{A_{\text{eff}, i}(0, 0)} \quad \forall \theta_i, \phi_i. \quad (7)$$

System parity requires that for any user i , the magnitude of the radiant intensity emitted in the direction of the symmetry axis and the required magnitude of the average signal irradiance incident in the direction of the symmetry axis for achieving a given BER are in a constant relationship, i.e.,

$$\frac{R_i A_{\text{eff}, i}(0, 0)}{\sqrt{N_{o, i}}} \cdot \frac{1}{I_i(0, 0)} = K. \quad (8)$$

This implies that more powerful transmitters must be accompanied by more sensitive receivers. Note that since ambient shot noise varies with the receiver location, it may be difficult to maintain the condition (8) under all conditions. In this paper, we design the transceivers to satisfy (8) considering thermal noise only.

Once a global value of K is chosen, the transmitter power and receiver sensitivity are specified as follows. Consider a pair of transceivers of a particular design, which are pointed directly at each other. Both receivers must achieve the minimum required SNR (SNR_{min}) at the maximum range (d_{max}). Satisfying (8) requires

$$\begin{cases} I_i(0, 0) = d_{\text{max}}(R_b SNR_{\text{min}})^{1/4} K^{-1/2} \\ R_i A_{\text{eff}, i}(0, 0) / \sqrt{N_{o, i}} = d_{\text{max}}(R_b SNR_{\text{min}})^{1/4} K^{1/2} \end{cases} \quad (9)$$

which are sufficient to fully specify the transmitted power and the receiver sensitivity.

3. Physical Layer Modeling and Analysis

A. Transceiver Designs

In this study, we consider some transceivers that use multiple transmitting and receiving elements. Referring to the coordinate system of Fig. 3, the axes of transmitting element m and receiving element n are oriented at angles (θ_{im}, ϕ_{im}) and (θ_{in}, ϕ_{in}) , respectively, while α_{im} and β_{im} denote the angles between those respective axes and the arbitrary direction (θ_i, ϕ_i) .

We model the emission of the transmitters by assuming an axially symmetric radiation pattern and a generalized Lambertian radiant intensity. For a LOS link, the radiant intensity of a transmitting element tilted at angle (θ_{im}, ϕ_{im}) is related to the average transmitted power P_t by [8]

$$I_{im}(\alpha_{im}) = \frac{P_t}{M} \left[\frac{l+1}{2\pi} \right] \cos^l \alpha_{im} \quad (10)$$

where M is the number of transmitting elements, and l is the Lambertian order of the transmitter, related to the half angle at half radiant intensity, $\alpha_{1/2, i}$, by $l = -\ln 2 / \ln(\cos \alpha_{1/2, i})$. The total radiant intensity is $I_i(\theta_i, \phi_i) = \sum_{m=1}^M I_{im}(\alpha_{im})$.

We assume that receiver i employs an optical concentrator with a cutoff angle $\beta_{c, i}$ at which the effective area goes to zero. An idealized non-imaging optical concentrator [9] having an internal refractive index N_r achieves a gain of $g(\beta_{in}) = N_r^2 / \sin^2 \beta_{c, i}$ for $0 \leq \beta_{in} \leq \beta_{c, i}$. The effective area of a receiving element is therefore given by

$$A_{\text{eff}, in}(\beta_{in}) = \begin{cases} AT_s(\beta_{in})g(\beta_{in})\cos\beta_{in} & 0 \leq \beta_{in} \leq \beta_{c, i} \\ 0 & \beta_{in} > \beta_{c, i} \end{cases} \quad (11)$$

where A is the area of one detector element, $T_s(\beta_{in})$ is the signal transmission of the filter, and $g(\beta_{in})$ is the optical concentrator gain. The overall effective area of a receiver is $A_{\text{eff}, i}(\theta_i, \phi_i) = \sum_{n=1}^{N_i} A_{\text{eff}, in}(\beta_{in})$.

Using these emitter and receiver models, we have considered five classes of transceivers, as shown in Table 1. Classes A, B, C, D and E achieve optical and system parity. Classes D and E achieve system parity, but maintain only approximate optical parity in the x_i - z_i plane, since they use three narrow-angle emitters in that plane, but use a single wide-angle receiver. Class O violates optical parity altogether, while Class S satisfies optical parity but violates system parity. Based on the system requirement of 10^{-8} BER at 4 Mbps in a 4-m LOS link [4], using the error-probability expressions given in Section 3-E, we found the required SNR to be 9.33 dB. Using (9), this determined the emitter powers and receiver areas given in Table 1. The noise PSD N_o was calculated using expressions given in Section 3-B.

Table 1: Classes of transceivers considered. All transmitting elements lie in the x_i-z_i plane of Fig. 3. Each transceiver class employs $N = 1$ receiving element oriented along the z_i axis.

Transceiver Class	Transmitter				Receiver	
	Number of Elements M	Average Power per Element P/M (mW)	Half-Angle $\alpha_{1/2, i}$ ($^\circ$)	Tilt Angles θ_{im} , $m = 1, \dots, M$ ($^\circ$)	Detector Area per Element A (mm^2)	Concentrator Cutoff Angle $\beta_{c, i}$ ($^\circ$)
A	1	172	60	0	2.86	90
B	1	59.1	30	0	0.72	30
C	1	23.3	18	0	0.27	18
D	3	38.1	30	-40, 0, 40	2.86	90
E	4	19.2	18	-45, -15, 15, 45	2.86	90
O	1	59.1	30	0	2.86	90
S	1	344	60	0	1.43	90

B. Modeling Noise and Signal Propagation

Our physical-layer simulations assume a room of horizontal dimensions $6 \text{ m} \times 10 \text{ m}$, with a ceiling height of 3 m , as in [11]. Eight tungsten floodlights are mounted on the ceiling, and an entire $3 \text{ m} \times 6 \text{ m}$ wall consists of a window admitting bright skylight. The receiver d.c. current induced by ambient light is given by:

$$I_{dc} = \sum_{6 \text{ surfaces } F_k} RT_s \Delta\lambda \int_{F_k} S_k(x, y) \cdot \frac{\cos \alpha_k}{\pi d^2} \cdot A_{\text{eff}}(\beta_k) dx dy + \sum_{\text{lamp } w} \Delta\lambda p_{\text{lamp}} \cdot \frac{(I_{\text{lamp}} + 1)}{2\pi} \cdot \frac{\cos \alpha_w}{d^2} \cdot A_{\text{eff}}(\beta_w) \quad (12)$$

where $S_k(x, y)$ is the spectral radiant emittance at (x, y) of room surface k , α_k is the angle between the axis of the transmitting element k and the receiver-transmitter line, β_k is the angle between the axis of the receiving element k and the transmitter-receiver line, R is the receiver responsivity, T_s is the receiver filter transmission coefficient, and $\Delta\lambda$ is the bandwidth of the filter. We use the values of $S_k(x, y)$ for each surface and p_{lamp} given in [11]. We assume $R = 0.5 \text{ A/W}$ and $T_s = 1$, $\Delta\lambda = 283 \text{ nm}$, corresponding to a silicon photodiode equipped with a Schott RG-780 optical longpass filter. The total two-sided receiver noise PSD is given by $N_o = N_{\text{shot}} + N_{\text{thermal}}$. The ambient shot noise PSD is

$$N_{\text{shot}} = qI_{dc}. \quad (13)$$

Computation of the thermal noise PSD N_{thermal} is described in [7].

For a given choice of transceiver designs and transceiver locations within the room (these user configurations are

specified in Section 5-A), we compute the received signal powers P_{ij} , taking account of the LOS contribution given by (5), and of non-LOS contributions arising from up to three bounces off of the walls, ceiling, floor and window. The non-LOS components are simulated using the method described in [10].

C. Demodulation

A unit-energy matched filter is employed at the receiver to perform chip-by-chip demodulation of the 4-PPM signal. The incoming desired signal and the interfering signal are summed together with additive white Gaussian noise with a two-sided power spectral density N_o . After passing through the matched filter, the received signal is sampled at each chip period. The received sample is given by

$$r_k = r_{d,k} + r_{i,k} + n_k \quad (14)$$

where $r_{d,k}$ is the desired signal component, $r_{i,k}$ is the interfering signal component and n_k is the sampled value of the noise. In our modeling, we consider co-channel interference only from the single strongest interferer active at a given time. Accordingly, the desired and interfering signal components in (14) are given by

$$r_{d,k} \Big|_{s_{d,k}=1} = 4RP_d \sqrt{T_c}, \quad r_{d,k} \Big|_{s_{d,k}=0} = 0 \quad (15)$$

$$r_{i,k} = 4RP_i \sqrt{T_c} \left(s_{i,k-1} \cdot \frac{\tau}{T_c} + s_{i,k} \cdot \frac{T_c - \tau}{T_c} \right), \quad (16)$$

where $s_{d,k}$ and $s_{i,k}$ are the information sequences of the desired and interfering signal, $s_{d,k}, s_{i,k} \in \{0, 1\}$, P_d and P_i are the average power of the received desired and interfering signals, respectively, and τ is the delay of the interfering sig-

nal with respect to the start of frame of the desired signal, and is uniformly distributed from zero to the chip period, T_c . Since a unit-energy matched filter is used, the noise component has a zero-mean Gaussian distribution with variance N_o .

Exact calculation of the symbol-error probability requires the consideration of the correlated nature of the interference. To simplify the calculation, we neglect this correlated nature and reduce the combined effects of noise and interference to a discrete memoryless channel characterized by chip-error probabilities p_{01} and p_{10} . Neglecting correlation but taking into account the marginal probabilities of the joint occurrences of $s_{i,k-1}$ and $s_{i,k}$, when the desired signal consists of a “one” chip, the received sample is given by:

$$r_k|_{s_{d,k}=1} = \begin{cases} 4R\sqrt{T_c}P_d + n_k & \text{w.p.} \\ 4R\sqrt{T_c}\left(P_d + P_i\frac{T_c - \tau}{T_c}\right) + n_k & \text{w.p.} \\ 4R\sqrt{T_c}\left(P_d + P_i\frac{\tau}{T_c}\right) + n_k & \text{w.p.} \\ 4R\sqrt{T_c}(P_d + P_i) + n_k & \text{w.p.} \end{cases} \quad (17)$$

$$\frac{32RR_i + 1}{64RR_i}$$

$$\frac{16RR_i - 1}{64RR_i}$$

$$\frac{16RR_i - 1}{64RR_i}$$

$$\frac{1}{64RR_i}$$

When the desired signal consists of a “zero” chip, the received signal $r_k|_{s_{d,k}=0}$ is given by an expression like (17), but with the desired signal component $4R\sqrt{T_c}P_d$ subtracted from each term.

D. Threshold Selection and Chip-Error Probabilities

To simplify the implementation, chip-by-chip hard-decision decoding (as opposed to soft-decision decoding) is performed on the received samples. A threshold η is chosen such that if a received sample r_k is greater than the threshold, it is decoded as a “one” chip; otherwise, it is decoded as a “zero” chip. Determination of the optimal threshold is analytically intractable, and we have considered three different heuristic choices of threshold. Among these three choices, we have found through simulation [7] that the best threshold lies at the midpoint of the eye opening:

$$\eta_1 = \frac{1}{2}[\max(r_{d,k}|_{s_{d,k}=1}) + \min(r_{d,k}|_{s_{d,k}=0})] = 2(P_d + P_i)R\sqrt{T_c}. \quad (18)$$

We will assume the use of threshold η_1 in the remainder of this paper. We have calculated the chip-error probabilities p_{01} and p_{10} , which are given by equations (30) and (31) in [7]. These expressions for p_{01} and p_{10} depend only on four parameters: RR_i , SNR , SIR , and τ/T_c . The signal-to-interference ratio is defined here as

$$SIR = \frac{P_d^2}{P_i^2}. \quad (19)$$

In the absence of interference, these expressions reduce to

$$p = p_{01} = p_{10} = Q(\sqrt{4SNR}). \quad (20)$$

E. Symbol-, Bit-, and Packet-Error Probabilities

As stated above, by ignoring the correlation in the interfering chip sequence, we have reduced the combined effects of noise and co-channel interference to a discrete memoryless channel characterized by p_{01} and p_{10} . Under this approximation, we derive the symbol-error probability in terms of p_{01} and p_{10} . To decode a repetition-coded block consisting of $4RR_d$ chips, the decoder performs a majority vote among all the valid received symbols (symbols with exactly one “one” chip) [4]. Thus, the decoder can correct all chip-error patterns up to weight $RR_d - 1$. The exact symbol-error probability for 4-PPM/VR is given by

$$P_M = \sum_{i=0}^{RR_d} \sum_{j=0}^{RR_d} \gamma_{ji} p_{01}^{(RR_d-i+j)} p_{10}^i (1-p_{01})^{(2RR_d+i-j)} (1-p_{10})^{(RR_d-i)} + \sum_{j=2RR_d+1}^{3RR_d} \sum_{i=j-2RR_d}^{RR_d} [\gamma_{ji} p_{01}^{(RR_d-i+j)} p_{10}^i (1-p_{01})^{(2RR_d+i-j)} \times (1-p_{10})^{(RR_d-i)}], \quad (21)$$

where γ_{ji} is the number of non-correctable error patterns for weight- $(RR_d + j)$ errors with i errors on the “one” chips. Unfortunately, the complexity of the combinatorics involved makes it difficult to determine all the γ_{ji} 's, especially for large j and large RR_d . We have derived three approximate expressions for the symbol-error probability, which we refer to here as the “lower bound”, the “union upper bound” and the “tighter upper bound”. Expressions for these three bounds are given in [7]. As will be shown in the next subsection, the tighter upper bound gives the best estimate for P_M , and we have used that bound in our LAN simulations.

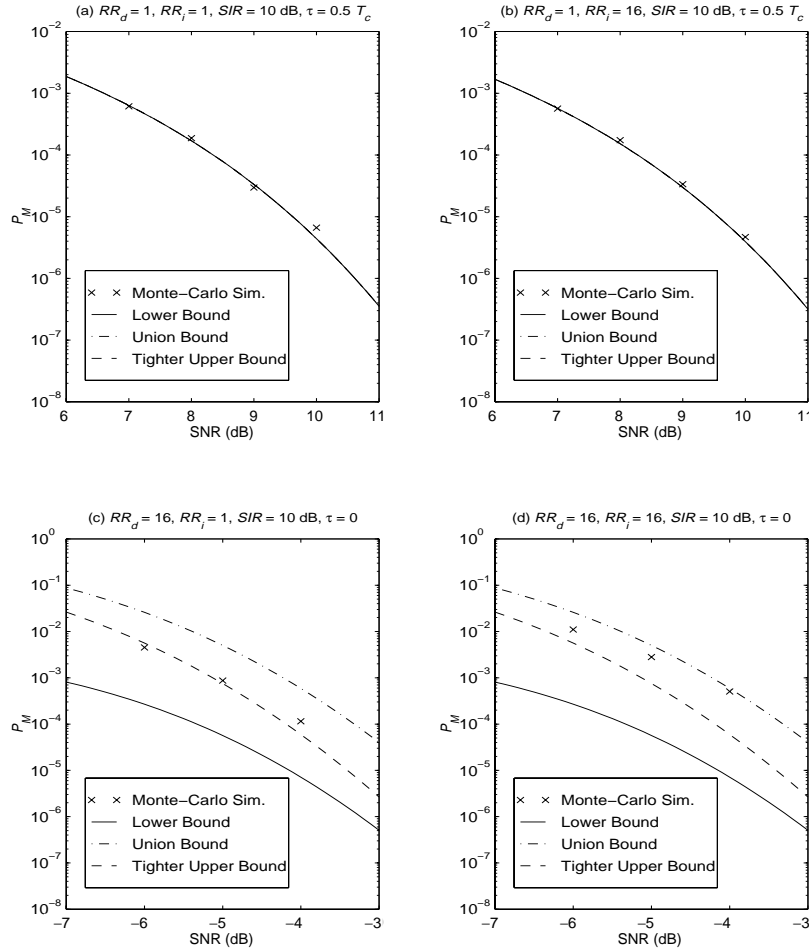


Fig. 4. Comparison between bounds on the symbol-error probability P_M and Monte-Carlo simulations.

Since 4-PPM symbols are orthogonal to each other, the bit-error probability is given by

$$P_b = \frac{M}{2(M-1)} P_M = \frac{2}{3} P_M \quad (22)$$

where $M = 4$. For a packet of size s bytes, the packet-error probability P_f is given by

$$P_f = 1 - (1 - P_M)^{4s}. \quad (23)$$

F. Verification of Approximate Interference Modeling

We have used Monte-Carlo simulations to test the validity of our approximate analytical expressions for the symbol-error probability P_M . Some illustrative examples are shown in Fig. 4. Parts (a) and (b) consider $\tau = T_c/2$, while parts (c) and (d) consider $\tau = 0$. In general, the analytical bounds agree well with simulation at low RR_d and high SIR . The discrepancies between the two increase with RR_i , and become more apparent at higher RR_d . In some cases, the Monte-Carlo-simulated error probability exceeds the tighter upper

bound. These discrepancies arise mainly because our analysis neglects correlation in the interfering signal. However, the tighter upper bound provides a good estimate of P_M when interference has the worst impact, i.e., when $\tau = 0$. For $SIR \geq 10$ dB, the discrepancy is less than 1 dB in terms of SNR (0.5 dB in terms of optical power) for all RR_i . For smaller SIR , the discrepancy increases, but such situations rarely occur when the MAC protocol works reliably. We have used the tighter upper bound in the LAN simulations described below.

G. Carrier Sensing

In IrMAC, carrier sensing is done through detection of the preamble field of a MAC layer frame. This preamble field consists of 128 repeated transmissions of the 4-PPM symbol “1000”. In our modeling of carrier sensing, for simplicity, we assume that all transmitters in the network are chip-synchronized. In order to perform carrier sensing, a receiver performs chip-by-chip hard decisions on the matched filter output. At each chip interval, a decoder con-

siders a block of the last $4N_{cs}$ chip decisions (corresponding to N_{cs} 4-PPM symbols), and determines the Hamming distance d_H between this block and the correct preamble codeword. If the Hamming distance d_H is less than a threshold η_{cs} , then the receiver has sensed a carrier. Analysis of the probabilities of successful carrier sensing and of false alarm is presented in [7]. In our LAN simulations, we chose $N_{cs} = 128$ and $\eta_{cs} = 126$. At SNR = -8 dB, these yield a 97% probability of successful carrier sensing, and a 3% probability of false alarm. This latter probability is acceptable, because at low SNRs, LAN throughput is impacted much more by failure of carrier sensing than by false alarm. Note that the AIr specifications state that transceivers should be able to perform reliable carrier sensing down to SNRs as low as -9 dB [4].

4. MAC Layer Modeling and Simulation

A. MAC Layer Simulation Assumptions

In order to simplify simulation of the MAC layer, we make several assumptions. (a) All users are assumed to be static. (b) When there is co-channel interference during any part of a packet, we regard that level of interference as being present during the entire packet. (c) We assume that the higher network layers, which we have not implemented in our simulations, pass data to the MAC layer in bursts, and that the burst length does not exceed the maximum burst time permitted by the IrMAC specifications [5]. (d) Each user has an infinite buffer for his input queue. (e) The modem turn-around time is assumed to be 0.1 ms, and the propagation delay is fixed at 30 ns, independent of the distance between transmitting and receiving nodes.

B. Interface Between Physical and MAC Layers

Since all users are assumed static, the communication channels are assumed to be fixed. Hence, for each LAN configuration, we perform a physical-layer simulation to determine SNR_{ij} (the SNR in receiver j when transmitter i is active), SIR_{ijk} (the SIR in receiver j when transmitters i and k are active, and k is the interfering transmitter), and RR_{ij} (the repetition rate used to transmit from i to j). The indices i, j, k run over all the nodes in the network. Using the analytical expressions given in Section 3, we then compute all of the relevant error probabilities, including those describing missed carrier sensing, false carrier sensing, reception of an erroneous robust packet header, and reception of an erroneous packet payload. During event-driven simulation of the MAC layer, the success or failure of each event is determined by performing a Bernoulli trial with the appropriate analytically determined probability.




Fig. 5. Configuration with two pairs of users. The four transceivers are aligned in a horizontal plane within a $6\text{ m} \times 10\text{ m}$, with a ceiling height of 3 m.

C. Event-driven MAC Simulator

To simulate the IrMAC protocol, we have implemented a custom event-driven simulator in C. In our simulator, simulation time is updated only when an event occurs. Each event consists of an event time and an event type. When an event occurs, the current simulation time is updated and according to the event type and the state of the users, new events are generated which are then put in the event queue with the appropriate event occurrence time. The event queue begins with nothing but data arrival events from each user. The simulation can be set to terminate based on criteria such as a specified number of data arrivals and simulation time. The event-driven-based simulation approach provides a good way to implement the finite state machine of the CSMA/CA protocol [5]. When an event occurs, appropriate actions are taken according to the state transition table and the state variables are updated.

D. MAC Layer Simulation Parameters

The traffic model assumes each user has data arriving in bursts independently. The interarrival time of data bursts is exponentially distributed, i.e., the data burst arrival time assumes a Poisson distribution. A data burst is made up of packets of two different kinds. 90% of the packets are modeled as “short” ones with a size of 256 bytes. 10% of the packets are modeled as “long” ones with a size of 1024 bytes. The number of packets in each burst randomly varies from one to any number, provided that the total burst time is less than the maximum allowed burst time. In the simulator, both the mean data burst interarrival time and maximum allowed burst time are adjustable parameters. In the AIr MAC specification, a transmitter can specify whether it requires an acknowledgment packet (ACK) from the receiver. For our simulations, all data packets require an ACK.

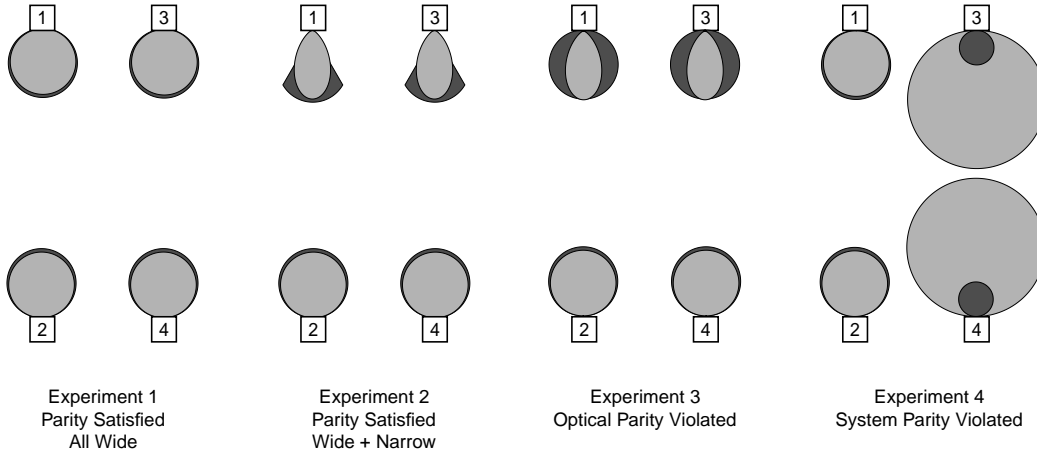


Fig. 6. Angular characteristics of the transceivers in Experiments 1 to 4. Light shading represents the transmitter radiant intensities, while dark shading represents the receiver effective light-collection areas.

We measure the network performance in terms of throughput, which is defined as the number of data bits successfully transmitted (and received) per unit time. Both the individual throughput and the aggregate throughput are monitored. We use a value of 250 ms for the maximum allowed burst time for all of our simulations. For each set of the physical layer simulation results, three sets of MAC layer simulations are performed with the mean data burst interarrival time chosen to be 100 ms, 500 ms and 5000 ms respectively. These values correspond to high, medium and low load traffic of the network. Each MAC layer simulation is terminated upon the arrival of the 1000th data burst of any user. This corresponds to real life run times of about 100 seconds to 5000 seconds depending on the choice of the mean data burst interarrival time. Due to the randomness introduced by the Poisson traffic, each set of MAC layer simulations is repeated ten times and the average throughput is taken.

5. LAN Simulations

A. User Configurations

We have considered five different user configurations, as described in [7]. In the interest of brevity, here we describe only the “two pairs of users” configuration, which is shown in Fig. 5. The two pairs of users are separated by a horizontal distance of D meters. Users 1 and 2 direct traffic only at one another, while users 3 and 4 direct traffic only at one another. All of the transceivers are oriented in the horizontal plane, i.e., the x_T - y_i plane is horizontal for each user.

We have performed four sets of experiments with this configuration, as described in Fig. 6. In Experiment 1, perfect optical parity and system parity is maintained, and wide-angle transceivers (Class A, Table 1) are employed by each user. In Experiment 2, approximate optical parity and perfect system parity is maintained, and two of the users employ

narrow-angle (Class B, Table 1) transceivers. In Experiment 3, two of the users violate optical parity by using narrow-angle transmitters and wide-angle receivers (Class O, Table 1). In Experiment 4, one pair of users violates system parity by using transmitters with twice the correct power and detectors with half the correct area (Class S, Table 1). Each experiment is performed under two different ambient light conditions no ambient light and intense ambient light. During each experiment, we have varied D from 3 m to 8 m.

B. Simulation Results

In this section, we present simulation results for the “two pairs of users” configuration. Results for the other four network configurations can be found in [7].

When the network is lightly loaded, we observed that the throughput in all experiments is approximately the same, except that there is a 10% drop of throughput when optical parity is violated in an intense ambient light environment. Under light loading, violation of channel reciprocity does not cause any significant adverse effects, because simultaneous transmissions occur rarely.

Fig. 7 shows simulation results obtained in a highly loaded network, i.e., with a mean data burst interarrival rate of 10 s^{-1} . In the absence of ambient light (Fig. 7 (a)), for D less than 4 m, the throughput is the same in all Experiments. As D increases, Experiment 2 gives the highest throughput, because the use of narrow-angle transceivers by User 1 and User 3 reduces the mutual interference between the two pairs of users, and the network essentially divides into two non-interfering segments. (A much smaller increase in throughput would be expected if User 1 and User 4 employed the narrow-beam transceivers instead). Compared to Experiment 2, Experiment 3 always gives a lower throughput, because of optical parity violation. Specifically, in Experiment 3, although User 1 and User 3 use narrow-angle

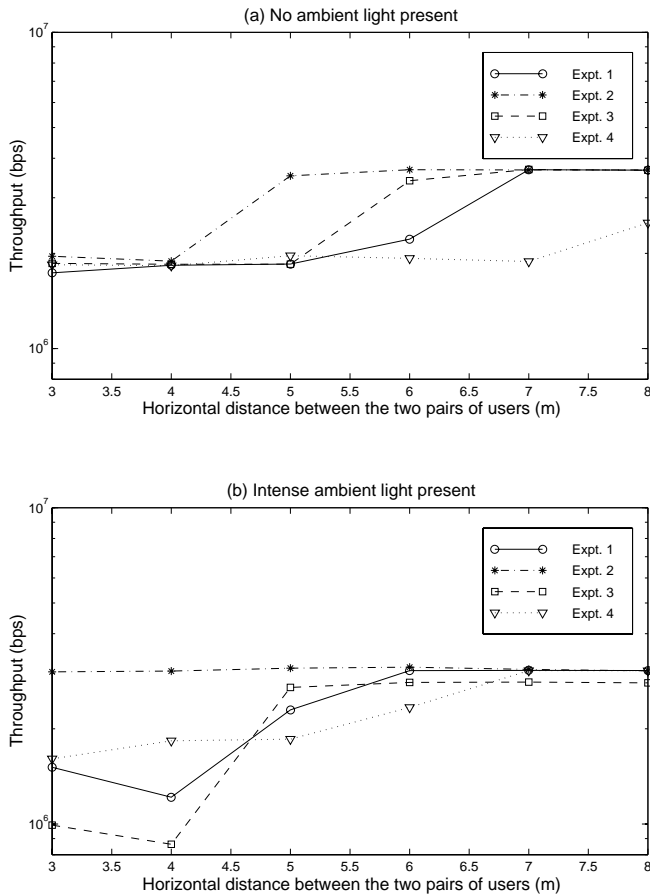


Fig. 7. Comparison of network throughputs for various experiments with two pairs of users, as shown in Fig. 5. The network is highly loaded, with a mean data burst interarrival rate of ten per second.

transmitters, their wide-angle receivers cause them to unnecessarily sense channel reservations by User 4 and User 2, respectively, thus inhibiting the ability of User 1 and User 3 to transmit. We might expect that Experiment 3 would yield lower throughput than Experiment 1, but this was not observed. Evidently, the channel non-reciprocity is not severe in the absence of ambient light, and the CSMA/CA protocol can still function. In Experiment 4, little rise in throughput is observed as D is increased, because of the system parity violation. The strong transmitters of User 3 and User 4 inhibit User 1 and User 2, even when D is large.

When there is intense ambient light (Fig. 7 (b)), we observe that in Experiment 2, the throughput is high for all D between 3 m and 8 m. This is because increased receiver noise desensitizes receiver carrier sensing, and the network is able to segment into two virtually non-interfering segments even for small D . At small D , Experiment 3 exhibits a 30% drop in throughput compared to Experiment 1, because

of optical parity violation. For example, because User 3 has a narrow-angle transmitter and a wide-angle receiver, User 2 cannot detect the reservation packets sent by User 3 but User 3 hears the transmissions from User 2 to User 1. As a result, collisions occur at User 3, and the throughput is reduced. Experiment 4 exhibits a throughput as much as 50% higher than Experiment 1 at small D . This is because the strong transmitters of User 3 and User 4 help overcome noise, allowing this pair to transmit to each other. However, this system parity violation leads to significant network unfairness, and the throughput between the stronger pair (Users 3 and 4) can be as much as 650 times the throughput of the weaker pair (Users 1 and 2).

6. Conclusions

We have presented a detailed analysis and modeling of the physical layer of the AIr standard. Physical layer and MAC layer simulations were performed to investigate the effect of non-reciprocity on network throughput. Our simulations show that the impact of non-reciprocity depends on the physical location of users and on the ambient light level. When multiple users are close to each other and there is intense ambient light, optical parity mismatch reduces the throughput whereas system parity violation introduces unfairness.

7. Acknowledgments

This work was supported by National Science Foundation ECS-9710065, the Sharp Corporation, and LG Electronics.

8. References

1. J. R. Barry, *Wireless Infrared Communications*, Kluwer Academic Publishers, Boston, 1994.
2. J. M. Kahn and J. R. Barry, "Wireless Infrared Communications," *Proceedings of the IEEE*, vol. 85, no. 2, pp. 265-298, February 1997.
3. I. Millar, M. Beale, B.J. Donoghue, K.W. Lindstrom and S. Williams "The IrDA Standards for High-Speed Infrared Communications", *Hewlett-Packard J.*, vol. 49, pp. 10-26, February, 1998.
4. Hewlett-Packard Company and IBM Corporation, "Request for Comments on Advanced Infrared (AIr) IrPHY Draft Physical Layer Specification," Version 0.4, January 1998.
5. Hewlett-Packard Company and IBM Corporation, "Request for Comments on Advanced Infrared (AIr) IrMAC Draft Protocol Specification," Version 0.2, July 1997.
6. P. Karn, "MACA - A New Channel Access Method for Packet Radio," *ARRL/CRRL Amateur Radio 9th Computer Networking Conference*, September 1990.
7. F. Chow, "Effect of Non-Reciprocity on Infrared Wireless Local-Area Networks", M.S. Report, University of California, Berkeley, December, 1998. Available on the web at: <http://www.eecs.berkeley.edu/~jmk/pubs/f.chow.ms.12.99.pdf>.

8. F. R. Gfeller and U. H. Bapst, "Wireless In-House Data Communication via Diffuse Infrared Radiation," *Proceedings of the IEEE*, vol. 67, pp. 1474-1486, November, 1979.
9. W.T. Welford and R. Winston, *High Collection Nonimaging Optics*, Academic Press, San Diego, 1989.
10. J. Barry, J. Kahn, W. Krause, E. Lee and D. Messerschmitt, "Simulation of Multipath Impulse Response for Indoor Wireless Optical Channels," *IEEE Journal on Selected Areas in Communications*, vol. 11, no. 3, pp. 367-379, April 1993.
11. J.B. Carruthers and J.M. Kahn, "Angle Diversity for Nondirected Wireless Infrared Communication", *Proc. of IEEE Intl. Conf. on Commun. (ICC 98)*, June 7-11, 1998, Atlanta, GA.