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

Francis M. Chow and Joseph M. Kahn

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

Phone: (510) 643-8848
Facsimile: (510) 642-2739
Internet: {fchow, jmk@eecs.Berkeley.EDU}

December 7, 1998

Abstract:

An infrared (IR) wireless network standard called Advanced Infrared (AIr) was proposed in April 1997. The new standard provides a dynamic tradeoff between bit rate and transmission range by employing four-slot pulse-position modulation with variable-rate repetition coding. 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. A physical-layer simulator models signal and noise reception, yielding values of the signal-to-noise ratio and signal-to-interference ratio in each receiver. Theoretical expressions for the error probability of RC-PPM, including the effect of co-channel interference, are used to obtain the probability of successful packet transmission and carrier sensing. A custom discrete-event simulator is used to simulate generation and transmission of packets in the LAN. Studying the example of AIr in detail, we show that non-reciprocity can substantially reduce throughput and increase unfairness.

CONTENTS

I. Introduction

The popularity of portable information appliances, such as notebook computers and personal digital assistants, has stimulated research and development aimed at low-cost, high-bandwidth wireless local-area networks (LANs). Infrared (IR) radiation [1],[2] offers several advantages over the traditional radio medium. Since the IR spectrum is not regulated worldwide, IR offers a virtually unlimited 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. IR emitters and detectors are available at low cost. An IR link employing intensity modulation with direct detection (IM/DD) does not suffer from multipath fading, simplifying receiver design. 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. Nonetheless, the commercial prospects for high-bandwidth infrared systems appear to be good.

Over the past five years, the Infrared Data Association (IrDA) has established standards [3]-[5] for short range, half-duplex line-of-sight (LOS) links operating at bit rates up to 4 Mbps. To date, the design of IrDA-standard transceivers has emphasized low cost and low power consumption. Current IrDA standards [3]-[5] describe directional transmitters and short-range, half-duplex, point-to-point links, and make no provision for multiple-access use of the IR medium. In view of this, Hewlett-Packard Company and IBM Corporation are currently collaborating on a new proposed IR standard called Advanced Infrared (AIr). As of this writing, both the AIr Physical Layer (AIr-IrPHY) and the Media-Access Control (IrMAC) Layer Specifications have been documented [6],[7], and are awaiting comment by IrDA members.

AIr-IrPHY [6] employs four-slot pulse-position modulation with variable-rate repetition coding (4-PPM/VR) to achieve a dynamic tradeoff 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 [7] 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 [8], 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.

Experimental measurements of the effect of non-reciprocity on packet reception are presented in [6]. 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 report studies non-reciprocity via analysis and simulation of the entire AIr network. We provide a detailed analysis of the physical layer and incorporate the physical layer model into the MAC layer network simulations to produce realistic results. The simulation flow employed in this report 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 an IR LAN. The SNR, SIR and RR are then used to calculate the frame-error rate, employing the theoretical expressions for error probability derived from the analysis of the physical layer. This methodology provides realistic input to the discrete-event MAC layer simulations. The network throughput is finally compared with different transceiver configurations to find out the effect of non-reciprocity on network performance. Five network configurations are carefully chosen to illustrate the impact of non-reciprocity under different LAN environments.

The remainder of this report is organized as follows. In Section II, we describe infrared channels and derive the conditions for channel reciprocity. In Section III, we describe our modeling and analysis of the physical layer. In Section IV, we explain the interface between the physical layer simula-

tion and the MAC layer simulation, and describe the MAC layer simulation itself. In Section V, the physical layer simulation experiments and the simulation results are presented. The effect of non-reciprocity is discussed, and transceiver design recommendations are given. Conclusions are given in Section VI.

II. Infrared Channels and Channel Reciprocity

A. Link Configuration

Infrared links are commonly classified according to two criteria, namely, degree of directionality of the transmitter and receiver and whether the link relies upon the existence of a line-of-sight (LOS) path between them. Fig. 3 shows this classification scheme. Directed links employ narrow field-of-view (FOV) transceivers that must be aimed in order to establish a communication link, while non-directed links employ wide FOV transceivers that alleviate the need for such positioning. LOS links rely upon a direct path between the transmitter and receiver for communication, whereas non-LOS links usually rely upon reflection of the light from the ceiling or some other diffusely reflecting surface. In general, directed links and LOS links minimize path loss and maximize power efficiency, and they can achieve higher transmission rates. However, they are less robust and less convenient to use. While suffering from lower transmission rates, non-directed and non-LOS links increase robustness and ease of use, allowing high user mobility and the links to operate even when there are barriers between the transmitter and the receiver.

B. Intensity Modulation/Direct Detection Channel Model

For a low cost infrared wireless system, the most viable modulation is intensity modulation (IM), in which the desired waveform is modulated onto the instantaneous power of the carrier; the most practical down-conversion technique is direct detection (DD), in which a photodetector produces a current proportional to the received instantaneous power [2]. Fig. 4 shows the modeling of infrared channels with IM/DD. The model of an IM/DD channel [2] is represented by

$$Y(t) = RX(t) \otimes h(t) + N(t) \tag{1}$$

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 is given by $P_t = \lim_{T \to \infty} \frac{1}{2T} \int_{-T}^{T} X(t) dt$, 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 N_o, we define the receiver electrical signal-to-noise ratio (SNR) as

$$SNR = \frac{R^2 P^2}{N_o B} \tag{2}$$

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

$$SNR = \frac{R^2 P^2 T_c}{N_o}. \tag{3}$$

C. Channel Reciprocity

To facilitate the discussion of the requirement of channel reciprocity, we first define a coordinate system as shown in Fig. 6. Using this notation, any arbitrary direction with respect to the i^{th} transceiver can be specified by the spherical polar coordinates (θ_i, ϕ_i) and 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 distance and angular orientation, mutually exhibit a similar link quality in both transmission directions. To achieve this, pairwise SNRs have to be symmetric, i.e.,

$$SNR_{ij} = SNR_{ji} \qquad \forall i \neq j \tag{4}$$

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

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

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), the requirement for channel reciprocity for any user i is found to be

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

where R_i is the responsivity of the photodetector and $N_{o, i}$ is the two-sided noise power spectral density. Note that the first factor is the property of the receiver while the second factor is the property of

the transmitter. As suggested by AIr-IrPHY [6], the channel reciprocity requirement can be conveniently broken down into two independent requirements, optical parity and system parity, by which simple transceiver design guidelines can be derived.

Optical parity refers to the transceivers having an approximate angular match of their emission and reception characteristics over all emission or reception angles. It requires that

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

For LOS transmissions, this requirement is reduced to apply only to angular directions within the horizontal plane corresponding to the normal usage of the device [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.,

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

This implies that more powerful transmitters must be accompanied by more sensitive receivers. Note that since the shot noise varies with respect to the location in an indoor environment, it is difficult to design transceivers that agree with (8) in general. When ambient light is present, $N_{o, i}$ is approximated by only considering the thermal noise. Now consider one particular transceiver design when all transceivers in the system are the same and the transmitter and the receiver both point directly at each other. Both receivers must achieve the minimum required SNR (SNR_{min}) at the maximum range (d_{max}). The system parity requirement is then expressed as

$$I_i(0, 0) = d_{\text{max}}(R_b SNR_{\text{min}})^{1/4} K^{-1/2} \tag{9}$$

$$\frac{R_i A_{\text{eff}, i}(0, 0)}{\sqrt{N_{o, i}}} = d_{\text{max}}(R_b SNR_{\text{min}})^{1/4} K^{1/2} \tag{10}$$

These are sufficient to fully specify the transmitted power and the receiver sensitivity. The constant K^1 is chosen globally for all the transceiver types. Equation (9) and (10) show that the larger K is, the less powerful the transmitter and the more sensitive the receiver will be.

[1.] Our constant K is related to the definition of system parity in [6] by $K = 16(\sqrt{SNR_{\text{min}} R_b})/(I_o E_o)$. I_o and E_o are terminologies presented in [6] and will not be defined in this report to avoid confusion.

III. Physical Layer Modeling and Analysis

In regard to the simulation flow shown in Fig. 2, this section introduces the physical layer simulations and gives a detailed account of the derivation of the theoretical expressions of the error probability. The MAC layer simulations are described in Section IV.

A. Modulation and Encoding

The AIr physical layer employs four-slot Pulse Position Modulation (4-PPM) as the modulation scheme. PPM is an orthogonal modulation technique that offers a decrease in required average power as compared to On-Off Keying (OOK). 4-PPM utilizes four symbols, each consisting of four time slots. We would refer each time slot as chips. The chip period, T_c, is 1/4 of the symbol period, T_m. A constant power $4P_t$ is transmitted during one of these chips (which represents logical 'one') and zero power is transmitted during the remaining three chips (which represent logical 'zero'). Since there are four different symbols, each symbol encodes two bits of information, and the relationship between bit rate, R_b, and T_c is given by $T_c = 1/(2R_b)$. Waveforms of 4-PPM are shown in Fig. 5. In the rest of the report, we will represent 4-PPM symbols as logical ones and zeros. This representation and the corresponding data bit pair mapping in shown in Table 1.

To enhance the flexibility of 4-PPM, variable repetition encoding (VR) is provided as a means of trade-off between data rate and transmission range. Using variable repetition encoding, each symbol is repeated RR times to make a block of RR symbols, where RR is the repetition rate. The resulting redundancy improves the probability of correctly decoding a symbol at the receiver. We will derive its effect in Section III-F. The advantage of VR is that it allows a virtual bandwidth reduction without having to physically change the receiver bandwidth. As a result, a simple and optimal receiver matched to the base data rate of 4 Mbps can be implemented. Table 2 illustrates the trade-off between data rate and transmission range.

B. Transceiver Designs

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 [9]

$$I_{im}(\alpha_{im}) = \frac{P_t}{M}\left[\frac{l+1}{2\pi}\right]\cos^l\alpha_{im} \tag{11}$$

where α_{im} is defined in Fig. 6, 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 $I_i(\theta_i, \phi_i) = \sum\limits_{m=1}^{M} I_{im}(\alpha_{im})$.

We assume the receivers employ an optical concentrator with a cutoff angle $\beta_{c, i}$ at which the effective area goes to zero. An idealized non-imaging optical concentrator [10] having an internal refractive index N_r achieves a gain of $g(\beta_{in}) = N_r^2 / \sin^2 \beta_{c, i}$ for $0 \le \beta_{in} \le \beta_{c, i}$, where β_{in} is defined in Fig. 6. The effective area of a receiving element is therefore given by

$$A_{\text{eff},in}(\beta_{in}) = \begin{cases} A T_s(\beta_{in}) g(\beta_{in}) \cos \beta_{in}, & 0 \le \beta_{in} \le \beta_{c, i} \\ 0, & \beta_{in} > \beta_{c, i} \end{cases} \tag{12}$$

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. Similarly, the overall effective area is $A_{\text{eff},i}(\theta_i, \phi_i) = \sum\limits_{n=1}^{N} A_{\text{eff},in}(\beta_{in})$.

Using these transceiver models, five classes of implementation examples (Classes A, B, C, D and E) that achieve optical and system parity are given in Table 3. In particular, the classes with multiple-element transmitters (Classes D, E) only maintain optical parity in the plane[1] that includes all the axes of the transmitting elements. The transmitting configurations for Class D and Class E are shown in Fig. 7. To help illustrate the effect of non-reciprocity, we also introduce an example of optical parity violation (Class O) and an example of system parity violation (Class S).

Based on the system requirement of 10^{-8} BER at 4 Mbps in a 4-m LOS link, as stated in the specification [6], we found the required SNR to be 9.33 dB (from which the constant K can be determined) from the theoretical symbol error probability derived in Section III-F. The values of the parameters in Table 3 are then calculated using (9), (10), (11), (12), where the noise power spectral density N_o is calculated using expressions given in Section III-C.

C. Modeling Noise and Channel Gains

The two-sided total noise power spectral density has a shot noise component and a thermal noise component, and is given by $N_o = N_{\text{shot}} + N_{\text{thermal}}$.

---

[1] This is referred as "horizontal" plane in [7] by assuming the IR devices are normally used horizontally.

For our physical layer simulations, we model the LAN environment as a room with dimensions $10\,\text{m} \times 6\,\text{m} \times 3\,\text{m}$ as in [12]. The ceiling height is 3 m. There are eight tungsten floodlights mounted on the ceiling, as shown in Fig. 8. The d.c. current is given by

$$
\begin{aligned}
I_{dc} = \sum_{6\ \text{surfaces}\ F_k} RT_s\Delta\lambda \int_{F_k} S_k(x, y) \cdot \frac{\cos\alpha_k}{\pi d^2} \cdot A_{\text{eff}}(\beta_k)dxdy \\
+ \sum_{\text{lamp}\ w} \Delta\lambda p_{\text{lamp}} \cdot \frac{(l_{\text{lamp}} + 1)}{2\pi} \cdot \frac{\cos^l\alpha_w}{d^2} \cdot A_{\text{eff}}(\beta_w)
\end{aligned}
\tag{13}
$$

where $S_k(x,y)$ is the spectral radiant emittance emitted at (x,y) of 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. The six surfaces correspond to the walls, ceiling, floor and window of the room. The values of $S_k(x,y)$ for each surface and p_{lamp} are specified in Fig. 8. For a silicon photodiode with a RG-780 optical longpass filter commonly used in IrDA receivers, we use $R = 0.5\,\text{A/W}$, $T_s = 1$ and $\Delta\lambda = 282.73\,\text{nm}$ according to the photodiode responsivity and filter transmission characteristics [2]. The two-sided power spectral density of the shot noise is given by

$$
N_{\text{shot}} = qI_{dc}
\tag{14}
$$

In general, thermal noise consists of a white component and a component proportional to the square of the frequency. This f^2 noise component is proportional to the square of the capacitance of the receiver. The receiver detector areas of our implementation examples in Table 3 are fairly small, which in turn gives small receiver capacitances. Consequently, the f^2 noise component becomes negligible compared to the white component. To simplify the model, it is justified to consider thermal noise as white. Assuming an op-amp is used to amplify the photocurrent at the receiver, the two sided thermal noise power spectral density is given by

$$
N_{\text{thermal}} = i_n^2 + \frac{e_n^2}{R^2_F} + \frac{2kT}{R_F}
\tag{15}
$$

where i_n^2 is the op-amp current noise, e_n^2 is the op-amp voltage noise, k is the Boltzmann constant, T is the absolute temperature, and R_F is the feedback resistance. The amplifier circuitry is illustrated in Fig. 9. To keep the thermal noise as small as possible, the feedback resistance should be as large as

possible. The value of the feedback resistance is limited by both the bandwidth of the amplifier and the voltage supply to the amplifier. First, we notice that the amplifier is also a lowpass filter, with 3 dB cutoff frequency given by

$$f_o = \frac{A_v}{2\pi(C_{in} + C_{det})R_F} \tag{16}$$

where A_v is the open-loop voltage gain of the op-amp, C_{in} is the input capacitance of the op-amp, and C_{det} is the capacitance of the photodetector. As stated in the next sub-section, the receiver employs a chip-by-chip demodulation scheme, which requires $f_o = 2R_b$. The maximum allowable feedback resistance is then given by

$$R_{Fmax} = \frac{A_v}{4\pi(C_{in} + C_{det})R_b} \tag{17}$$

Second, we want to make sure that the op-amp is not saturated by the photocurrent produced by the background light, i.e. $V_s \geq I_{dc}R_F$, where V_s is the supply voltage to the op-amp. Therefore,

$$R_{Fmax} = \min\{\frac{A_v}{4\pi(C_{in} + C_{det})R_b} , \frac{V_s}{I_{dcmax}}\} \tag{18}$$

The maximum d.c. current, I_{dcmax}, can be estimated by (13). Using $V_s = 3\text{V}$, $A_v = 1400$, $i_n^2 = 2.56 \times 10^{-24}\text{A}^2/\text{Hz}$, $e_n^2 = 1.18 \times 10^{-18}\text{V}^2/\text{Hz}$ and a capacitance of $1.9 \times 10^{-6}\text{F}/\text{m}^2$ per unit area of the detector, we estimate a value of $1.51 \times 10^{-24}\text{A}^2/\text{Hz}$ for the two-sided thermal noise power spectral density.

For a given choice of transceiver designs and transceiver locations within the room (these user configurations are specified in Section V-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 computed using 'IrSim', a program that implements the multipath channel simulation technique described in [11]. 'IrSim' uses the diffuse reflectivities of the various room surfaces, which are described in Fig. 8.

D. Demodulation

A unit-energy matched filter shown in Fig. 10 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

$$r_k = r_{d,k} + r_{i,k} + n_k \tag{19}$$

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. Since AIr uses a collision avoidance protocol in the MAC layer [7], the chances of having a nearby interferer is very small, and most interfering signals would be coming from users who are far away from an established communication link. We argue that we can provide a reasonably accurate model by considering only the strongest interferer, i.e., we will only consider one interfering signal in our analysis. Using this argument, the signal components in (19) are given by

$$r_{d,k}\big|_{s_{d,k}=1} = 4RP_d\sqrt{T_c}, \; r_{d,k}\big|_{s_{d,k}=0} = 0 \tag{20}$$

$$r_{i,k} = 4RP_i\sqrt{T_c}\left(s_{i,k-1}\cdot\frac{\tau}{T_c} + s_{i,k}\cdot\frac{T_c-\tau}{T_c}\right) \tag{21}$$

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 signal respectively. τ is the delay of the interfering signal with respect to the start of frame of the desired signal, and is uniformly distributed from zero to the chip period, T_c. The relationship between the desired signal and the interfering signal is illustrated by Fig. 11. Since a unit-energy matched filter is used, the noise component has a zero-mean Gaussian distribution with variance N_o [13].

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 effect of noise and interference to a discrete memoryless channel characterized by chip-error probabilities p_{01} and p_{10}. Calculation of p_{01} and p_{10} requires taking into account the marginal probabilities of the joint occurrences of $s_{i,k-1}$ and $s_{i,k}$, which is dependent on RR_i, the repetition rate of the interfering signal. Let $P(s_{i,k-1}, s_{i,k}, RR_i)$ be this marginal probability as a function of RR_i. As RR_i increases, there exists more correlation between neighboring chips and the approximation of a discrete memoryless channel becomes less accurate. Nevertheless, we will later justify the validity of this approximation. For a repetition encoded symbol with $4RR_i$ chip positions, let

$P_j(s_{i,k-1}, s_{i,k}, RR_i)$ be the probability of having $s_{i,k-1}$ at position j followed by $s_{i,k}$ at position $j+1$ given position j is $s_{i,k-1}$. We observe that for $1 \le j \le 4RR_i - 1$, $P_j(0,0,RR_i) = 2/3$, $P_j(0,1,RR_i) = 2/3$, $P_j(1,0,RR_i) = 1$ and $P_j(1,1,RR_i) = 0$. For $j = 4RR_i$, $P_j(0,0,RR_i) = 3/4$, $P_j(0,0,RR_i) = 1/4$, $P_j(0,0,RR_i) = 3/4$ and $P_j(0,0,RR_i) = 1/4$. Then the marginal probabilities $P(s_{i,k-1}, s_{i,k}, RR_i)$ are given by

$$P(0,0,RR_i) = \frac{3}{4} \cdot \frac{1}{4RR_i} \cdot \sum_{j=1}^{4RR_i} P_j(0,0,RR_i) = \frac{32RR_i + 1}{64RR_i}$$

$$P(0,1,RR_i) = \frac{3}{4} \cdot \frac{1}{4RR_i} \cdot \sum_{j=1}^{4RR_i} P_j(0,1,RR_i) = \frac{16RR_i - 1}{64RR_i}$$

$$P(1,0,RR_i) = \frac{1}{4} \cdot \frac{1}{4RR_i} \cdot \sum_{j=1}^{4RR_i} P_j(0,1,RR_i) = \frac{16RR_i - 1}{64RR_i}$$

$$P(1,1,RR_i) = \frac{1}{4} \cdot \frac{1}{4RR_i} \cdot \sum_{j=1}^{4RR_i} P_j(1,1,RR_i) = \frac{1}{64RR_i} \tag{22}$$

Using these marginal probabilities, the received sample is calculated to be

$$r_k \big|_{s_{d,k}=1} = \begin{cases} 4R\sqrt{T_c}P_d + n_k & \frac{32RR_i + 1}{64RR_i} \\ 4R\sqrt{T_c}\left(P_d + P_i \cdot \frac{T_c - \tau}{T_c}\right) + n_k & \frac{16RR_i - 1}{64RR_i} \\ 4R\sqrt{T_c}\left(P_d + P_i \cdot \frac{\tau}{T_c}\right) + n_k & \frac{16RR_i - 1}{64RR_i} \\ 4R\sqrt{T_c}(P_d + P_i) + n_k & \frac{1}{64RR_i} \end{cases} \quad \text{w.p.} \tag{23}$$

$$r_k\big|_{s_{d,k}=0} = \begin{cases} n_k & \dfrac{32RR_i+1}{64RR_i} \\[3ex] 4R\sqrt{T_c}P_i \cdot \dfrac{T_c-\tau}{T_c} + n_k & \dfrac{16RR_i-1}{64RR} \\[3ex] 4R\sqrt{T_c}P_i \cdot \dfrac{\tau}{T_c} + n_k & \dfrac{16RR_i-1}{64RR_i} \\[3ex] 4R\sqrt{T_c}P_i + n_k & \dfrac{1}{64RR_i} \end{cases} \quad \text{w.p.} \tag{24}$$

E. 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 is greater than the threshold, it is decoded as a 'one' chip; otherwise, it is decoded as a 'zero' chip. Since the probability distribution of the magnitude of the received sample is a sum of four Gaussian distributions, the optimal threshold is analytically intractable. We have to rely on some heuristic choices and compare their performances. We choose η_1 to be the center of the maximum eye-opening and η_2 to be the mid-point between the expected sampled value when $s_{d,k} = 1$ and that when $s_{d,k} = 0$. Next, by considering the fact that any sampled values greater than the minimum sampled value for $s_{d,k} = 1$ are going to help the decision of estimating the received chip to be 'one', η_3 is chosen to be the mid-point between the minimum sampled value when $s_{d,k} = 1$ and the expected sampled value when $s_{d,k} = 0$. These thresholds are given by

$$\eta_1 = \frac{1}{2}\left[\max\left(r_{d,k}\big|_{s_{d,k}=1}\right) + \min\left(r_{d,k}\big|_{s_{d,k}=0}\right)\right] = 2(P_d+P_i)R\sqrt{T_c} \tag{25}$$

$$\eta_2 = \frac{1}{2}\left\{E\left[r_{d,k}\big|_{s_{d,k}=1}\right] + E\left[r_{d,k}\big|_{s_{d,k}=0}\right]\right\} = (2P_d+P_i)R\sqrt{T_c} \tag{26}$$

$$\eta_3 = \frac{1}{2}\left\{\min\left(r_{d,k}\big|_{s_{d,k}=1}\right) + E\left[r_{d,k}\big|_{s_{d,k}=0}\right]\right\} = (2P_d+0.5P_i)R\sqrt{T_c} \tag{27}$$

For a general threshold η, recalling n_k is Gaussian, the chip-error probabilities are calculated by integrating the tails of the Gaussian distributions of the received sample. They are given by

$$p_{01} = \frac{32RR_i + 1}{64RR_i} Q\left(\frac{\eta}{\sqrt{N_o}}\right) + \frac{16RR_i - 1}{64RR_i} Q\left(\frac{\eta - 4R\sqrt{T_c}P_i \cdot (T_c - \tau)/T_c}{\sqrt{N_o}}\right) +$$

$$\frac{16RR_i - 1}{64RR_i} Q\left(\frac{\eta - 4R\sqrt{T_c}P_i \cdot \tau/T_c}{\sqrt{N_o}}\right) + \frac{1}{64RR_i} Q\left(\frac{\eta - 4R\sqrt{T_c}P_i}{\sqrt{N_o}}\right)$$

(28)

$$p_{10} = \frac{32RR_i + 1}{64RR_i} Q\left(\frac{4R\sqrt{T_c}P_d - \eta}{\sqrt{N_o}}\right) + \frac{16RR_i - 1}{64RR_i} Q\left(\frac{4R\sqrt{T_c}(P_d + P_i \cdot (T_c - \tau)/T_c) - \eta}{\sqrt{N_o}}\right) +$$

$$\frac{16RR_i - 1}{64RR_i} Q\left(\frac{4R\sqrt{T_c}(P_d + P_i \cdot \tau/T_c) - \eta}{\sqrt{N_o}}\right) + \frac{1}{64RR_i} Q\left(\frac{4R\sqrt{T_c}(P_d + P_i) - \eta}{\sqrt{N_o}}\right)$$

(29)

where $Q(x) = (1/\sqrt{2\pi}) \int_x^\infty e^{-y^2/2} dy$. In Section III-G, we will show that η_1 is the best among the three. Using η_1 as the threshold, the chip-error probabilities are given by

$$p_{01} = \frac{32RR_i + 1}{64RR_i} Q\left(\sqrt{4SNR\left(1 + \frac{2}{\sqrt{SIR}} + SIR\right)}\right) +$$

$$\frac{16RR_i - 1}{64RR_i} Q\left(\sqrt{4SNR\left(1 - \frac{2(1 - 2(\tau/T_c))}{\sqrt{SIR}} + SIR\left(1 - 2\frac{\tau}{T_c}\right)^2\right)}\right) +$$

$$\frac{16RR_i - 1}{64RR_i} Q\left(\sqrt{4SNR\left(1 + \frac{2(1 - 2(\tau/T_c))}{\sqrt{SIR}} + SIR\left(1 - 2\frac{\tau}{T_c}\right)^2\right)}\right) +$$

$$\frac{1}{64RR_i} Q\left(\sqrt{4SNR\left(1 - \frac{2}{\sqrt{SIR}} + SIR\right)}\right)$$

(30)

$$p_{10} = \frac{32RR_i + 1}{64RR_i} Q\left(\sqrt{4SNR\left(1 - \frac{2}{\sqrt{SIR}} + SIR\right)}\right) +$$

$$\frac{16RR_i - 1}{64RR_i} Q\left(\sqrt{4SNR\left(1 + \frac{2(1 - 2(\tau/T_c))}{\sqrt{SIR}} + SIR\left(1 - 2\frac{\tau}{T_c}\right)^2\right)}\right) +$$

$$\frac{16RR_i - 1}{64RR_i} Q\left(\sqrt{4SNR\left(1 - \frac{2(1 - 2(\tau/T_c))}{\sqrt{SIR}} + SIR\left(1 - 2\frac{\tau}{T_c}\right)^2\right)}\right) +$$

$$\frac{1}{64RR_i} Q\left(\sqrt{4SNR\left(1 + \frac{2}{\sqrt{SIR}} + SIR\right)}\right) \tag{31}$$

where SIR, the signal-to-interference ratio, is defined as $SIR = P_d^2/P_i^2$. The chip-error probabilities depend only on RR_i, SNR, SIR, τ and T_c. If τ is zero, i.e., when all signals are synchronized, the chip-error probabilities depend only on RR_i, SNR and SIR. In the absence of interference, these expressions reduce to

$$p = p_{01} = p_{10} = Q(\sqrt{4SNR}) \tag{32}$$

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

As described above, by ignoring the correlation of the interfering chip sequence, we have reduced the system to a discrete memoryless channel characterized by p_{01} and p_{10}. Using this approximation of the real system, we can derive the symbol-error probabilities in terms of p_{01} and p_{10}. The decoder performs a majority voting on all the valid received symbols (symbols with only one 'high' chip) to estimate the transmitted bit sequence out of the repetition encoded symbols, and can thus correct errors up to an error weight of $RR_d - 1$. We will first derive the more simple case of a binary symmetric channel where $p_{01} = p_{10}$. Then we will generalize the results to a binary asymmetric channel.

1. Symbol-error Probabilities of a Binary Symmetric Channel

The theoretical exact symbol error probability of a binary symmetric channel for a hard decision decoding system is given by

$$P_M = 1 - \sum_{k=0}^{4RR_d} h_k p^k (1 - p)^{4RR_d - k} \tag{33}$$

where h_k is the number of correctable error patterns for weight-k errors and p is the chip-error proba-bility given by (32). If we let g_k be the number of incorrectable error patterns for weight-k errors, and observe that for 4-PPM/VR, $h_k = 0$ for $k > 4RR_d - 1$ and $g_k = 0$ for $k < RR_d$, we can apply the property of binomial theorem and re-write this expression as

$$P_M = \sum_{k = RR_d}^{4RR_d} g_k p^k (1-p)^{4RR_d - k} \tag{34}$$

Consider for $RR_d = 16$, there are 64 chips and thus 2^{64} different error patterns. Because of the complexity of the combinatorics involved, it is extremely difficult to find out all the coefficients g_k, especially for large k and large RR_d. Fortunately, for small p, the summation is dominated by the first few terms. The expressions for g_k for $k = RR_d$ and $k = RR_d + 1$ are given as follows:

$$g_{RR_d} = \sum_{C_0 = 0}^{RR_d/2} \frac{RR_d!}{(RR_d - 2C_0)!(C_0!)^2} \cdot 4^{(RR_d - 2C_0)} \cdot [3 \cdot I_{\{C_0 \neq 0\}} + I_{\{C_0 = 0\}}] \tag{35}$$

$$g_{RR_d + 1} = \sum_{C_0 = 0}^{RR_d/2 - 1} \frac{RR_d!}{(C_0 + 1)!(RR_d - 2C_0 - 1)!C_0!} \cdot 4^{(RR_d - 2C_0 - 1)} \cdot \lambda_{C_0} \tag{36}$$

where I is the indicator function and λ_{C_0} is the number of incorrectable error patterns that are gener-ated solely by permutating the symbols with weight-2 errors when the number of correct symbols in the repetition coded block is C_0. It is found that $\lambda_0 = 6$, $\lambda_1 = 27$, and $\lambda_{C_0} = 18 + 15C_0$ for $C_0 \geq 2$.

A lower bound can be obtained by taking the first two terms of (34) as given above:

$$P_M \geq g_{RR_d} p^{RR_d} (1-p)^{3RR_d} + g_{RR_d + 1} p^{RR_d + 1} (1-p)^{3RR_d - 1} \tag{37}$$

A conservative Union (upper) Bound is obtained by assuming all error patterns with error weights greater than $RR_d - 1$ are incorrectable. This over-counting gives the expression

$$P_M \leq \sum_{k = RR_d}^{4RR_d} \binom{4RR_d}{k} p^k (1-p)^{4RR_d - k} \tag{38}$$

By replacing the first two terms of (38) using the coefficients given by (35) and (36), we arrive at a much tighter upper bound as compared to the Union Bound, expressed as

$$P_M \leq g_{RR_d} p^{RR_d}(1-p)^{3RR_d} + g_{RR_d+1} p^{RR_d+1}(1-p)^{3RR_d-1}$$
$$+ \sum_{k=RR_d+2}^{4RR_d} \binom{4RR_d}{k} p^k (1-p)^{4RR_d-k} \tag{39}$$

2. Symbol-error Probability of a Binary Asymmetric Channel

The above results can be extended to include the more general case when the communication channel is asymmetric. The chip-error probabilities p_{01} and p_{10} are given by (30) and (31). The theoretical exact symbol-error probability for a binary asymmetric channel is given by

$$P_M = \sum_{j=0}^{2RR_d} \sum_{i=0}^{RR_d} \gamma_{ji} p_{01}^{(RR_d-i+j)} p_{10}^{i}(1-p_{01})^{(2RR_d+i-j)}(1-p_{10})^{(RR_d-i)}$$
$$+ \sum_{j=2RR_d+1}^{3RR_d} \sum_{i=j-2RR_d}^{RR_d} \gamma_{ji} p_{01}^{(RR_d-i+j)} p_{10}^{i}(1-p_{01})^{(2RR_d+i-j)}(1-p_{10})^{(RR_d-i)} \tag{40}$$

where γ_{ji} is the number of incorrectable error patterns for weight-(RR_d+j) errors with i errors on the 'one' chips. Again, the complexity of the combinatorics involved makes it difficult to determine all the γ_{ji}'s, especially for large j and large RR_d; but the summation is dominated by the first few terms for small p_{01} and p_{10}. The expression for γ_{ji} for $j = 0, 1$ are given as follows:

$$\gamma_{0i} = \sum_{C_0=0}^{RR_d/2} a_{i, C_0} \tag{41}$$

and for $i < C_0$ or $i > RR_d - C_0$, $a_{i, C_0} = 0$ whereas for $C_0 \leq i \leq RR_d - C_0$,

$$a_{i, C_0} = \frac{RR_d!}{(RR_d-2C_0)!(C_0!)^2}\left[3^{RR_d-C_0-i+1}\left(\frac{1}{3} \cdot I_{\{C_0=0\}} + I_{\{C_0 \neq 0\}}\right)\frac{(RR_d-2C_0)!}{(RR_d-C_0-i)!(i-C_0)!}\right] \tag{42}$$

$$\gamma_{1i} = \sum_{C_0=0}^{RR_d/2-1} b_{i, C_0} \tag{43}$$

and for $i < C_0$ or $i > RR_d - C_0$, $b_{i, C_0} = 0$; for $i = C_0$,

$$b_{i,C_0} = \frac{RR_d!}{(RR_d - 2C_0 - 1)!(C_0 + 1)!C_0!}\left[3^{RR_d - 2C_0 + 1} \cdot \left(\frac{1}{3} \cdot I_{\{C_0 = 0\}} + I_{\{C_0 \neq 0\}}\right) \cdot (C_0 + 1)\right]; \quad (44)$$

for $C_0 < i < RR_d - C_0$,

$$b_{i,C_0} = \frac{RR_d!}{(RR_d - 2C_0 - 1)!(C_0 + 1)!C_0!}\left[3^{RR_d - C_0 - i} \cdot \mu_{C_0} \cdot \frac{(RR_d - 2C_0 - 1)!}{(RR_d - C_0 - i)!(i - C_0 - 1)!}\right.$$

$$\left. + 3^{RR_d - C_0 - i + 1} \cdot \left(\frac{1}{3} \cdot I_{\{C_0 = 0\}} + I_{\{C_0 \neq 0\}}\right) \cdot \frac{(RR_d - 2C_0 - 1)!}{(RR_d - C_0 - i - 1)!(i - C_0)!} \cdot (C_0 + 1)\right]; \quad (45)$$

and finally for $i = RR_d$,

$$b_{i,C_0} = \frac{RR_d!}{(RR_d - 2C_0 - 1)!(C_0 + 1)!C_0!}[\mu_{C_0}]. \quad (46)$$

In (45) and (46), μ_{C_0} is the number of incorrectable error patterns generated by the weight-2 error symbols with one of the two errors fixed on a 'one' chip, when the number of correct symbols within the VR-encoded block is C_0. $\mu_0 = 3$, $\mu_1 = 9$ and $\mu_{C_0} = 9 + 6C_0$ for $C_0 \geq 2$.

A lower bound can be obtained by taking only the terms of $j = 0, 1$ in (40), which are just described above. A union (upper) bound is obtained by taking (40) and replacing all the coefficients γ_{ji} by the terms given in (47). Similar to the binary symmetric case, this expression over-counts the number of incorrectable error patterns by assuming all error patterns for any error weights greater than $RR_d - 1$ are incorrectable.

$$\gamma_{ji} = \binom{3RR_d}{RR_d - i + j} \cdot \binom{RR_d}{i} \quad (47)$$

A tighter upper bound can be obtained by taking (40) and using the exact values of γ_{ji} for $j = 0, 1$ given in (41)-(46) and approximate values for $j > 1$ given in (47). This does not over-count the incorrectable error patterns of the dominant (first two) terms. As will be shown in the next sub-section, the tighter upper bound gives the best estimate for P_M. Following that subsection, for the remainder of the report, we use the tighter upper bound for P_M.

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

$$P_b = \frac{M}{2(M - 1)}P_M = \frac{2}{3}P_M \quad (48)$$

where M is the dimension of the signal space [13]. For 4-PPM, since there are four chips for each signal, M = 4. For a packet with size s bytes, the packet error probability P_f is given by

$$P_f = 1 - (1 - P_M)^{4s} \tag{49}$$

G. Threshold Selection and Verification of Approximate Interference Modeling

We want to optimize the performance of the receiver in terms of symbol error probability by choosing the right threshold. We compare the three choices (25)-(27) against the numerically determined optimal threshold at different SNRs and choices of τ. It is found that at high SNRs and when τ is close to zero, threshold η_1 always performs much better than η_2 and η_3. η_2 performs better than η_1 and η_3 when τ is close to 0.5 T_c. We observed that η_2 is far from optimal when the interfering signal gives a greater impact on the desired signal, i.e., when τ is close to zero, and averaging the performance in different cases, we choose η_1 as the threshold at the receiver. An illustrative comparison for the case of $RR_d = 4$, $RR_i = 16$, $SIR = 3$ dB with τ equals 0.1, 0.2, 0.3 and 0.4 T_c is shown in Fig. 12.

The verification of our approximation for P_M is performed through Monte-Carlo Simulations, in which a 'real-life' communication channel is modeled and simulated. White Gaussian noise is added to random chip sequences at the receiver and the chips are decoded to symbols. The decoder then counts the percentage of correct symbols and estimate the symbol-error probability. Simulations with different choices of parameters (RR_d, RR_i, SNR, SIR and τ) are performed and the results are compared against the lower bound, the union bound and the tighter upper bound. A rough guideline for determining the validity of the theoretical expressions is that the discrepancy between theory and simulations can be compensated by less than 1 dB in SNR. Due to the limitation of computing power, we can only verify the theoretical derivations up to a symbol error probability of about 10^{-5}.

In general, the simulation results agree with the theoretical results at low RR_d and high SIRs. Two illustrative examples are given in Fig. 13 (a) and (b). The discrepancy increases with RR_i, and becomes more apparent at higher RR_d. This is mainly due to the fact that our approximation failed to capture the effect of correlation between neighboring samples of the interfering signal. However, as shown in Fig. 13 (c) and (d), the tighter upper bound still gives a good estimate of the symbol error probability when the interfering signal gives the worst impact (when $\tau = 0$) to the desired signal. The estimation error increases as RR_i increases, but can still be compensated by less than 1 dB in SNR in the case of $RR_i = 16$. As SIR goes down, the results do not agree very well with the theoretical expressions. Fortunately, these situations rarely occur when the MAC protocol works reliably.

H. Carrier Sensing

Carrier sensing of the MAC layer is done through the preamble field of a MAC layer frame, of which the format is shown in Fig. 14. The preamble field consists of 128 repeated transmission of the 4-PPM symbol '1000'. For simplicity, we will assume all the clocks in the network are synchronized at the chip level. At the receiver, the same matched filter is used to perform chip-by-chip demodulation of the incoming preamble field. At each sampling time, the decoder determines the Hamming distance d_H between the correct preamble codeword (a certain number of symbols, N_{cs}, would be used as the correct codeword for this carrier sensing detection) and the incoming frame. The Hamming distance is then compared against a certain threshold η_{cs}. If the Hamming distance between the codewords is less than the threshold, carrier sensing is detected. We need to find a combination of N_{cs} and η_{cs} to provide reliable carrier sensing down to low SNRs, typically around -9 dB [6]. (The SNR requirement in the specification is 3 dB. By our definition (2), this is equivalent to -9 dB.)

First, we consider the case when no noise is present. The Hamming distance is zero for the correct preamble codeword with the correct delay, N_{cs} for the all zero (nothing transmitted) codeword, $1.5\,N_{cs}$ for a random data codeword, and $2\,N_{cs}$ for the correct preamble codeword with the wrong delay. When noise is present, the probability distribution of Hamming distance for the different cases are centered at those values. It appears that the all zero codeword would be the limiting case since its distribution is closest to the distribution when the correct preamble with the correct delay is present. Therefore, we can accurately approximate the situation by considering only the detection of the correct preamble codeword and the all zero codeword. Since the exact analysis of the probability of carrier sensing is not as critical as the analysis of the probability of symbol error, we do not include the effect of interference here and assume a binary symmetric channel with a chip-error probability p given by (32).

Let $\overrightarrow{C_1}$ be the correct preamble codeword and $\vec{0}$ be the all zero codeword. The probability distribution functions of the Hamming distance when $\overrightarrow{C_1}$ and $\vec{0}$ (i.e. nothing) is transmitted are given by

$$P(d_H = d \,|\, \overrightarrow{C_1}) \;=\; \binom{4N_{cs}}{d} p^d (1-p)^{4N_{cs}-d} \tag{50}$$

$$P(d_H = d \,|\, \vec{0}) \;=\; \sum_{k=0}^{d} \binom{N_{cs}}{d-k} \cdot \binom{3N_{cs}}{k} p^{N_{cs}-d+2k} (1-p)^{3N_{cs}+d-2k} \tag{51}$$

The probability of successful carrier sensing detection and probability of false alarm are given by

$$P(\text{success}) = \sum_{d=0}^{\eta_{cs}} P(d_H = d | \vec{C}_1) \tag{52}$$

$$P(\text{false alarm}) = \sum_{d=0}^{\eta_{cs}} P(d_H = d | \vec{0}) \tag{53}$$

We choose $N_{cs} = 128$ and $\eta_{cs} = 126$ to give a probability of successful carrier sensing detection of 96.8% at SNR = -8 dB. The maximum probability of false alarm is about 3%. This is acceptable because the main contribution to worse performance is not false carrier sensing detection but failing to detect the carrier at low SNRs.

IV. MAC Layer Modeling and Simulation

A. MAC Layer Simulation Assumptions

We make certain assumptions to the MAC layer simulations to simplify the simulation process. They are detailed and justified as follows. First, all users are static. This is justified when we compare the average time between users movement and the data rate. A considerable amount of data is transmitted each time before any users move and a static setting would suffice in accounting for the effect of non-reciprocity. Second, we assume that when there is interference, any part of a packet being interfered is regarded as the whole packet being interfered. This is a fair assumption because incorrectable error in any part of a packet is likely to corrupt the whole packet. Third, we assume the higher layer (not implemented) of the network passes data to the MAC layer in bursts, and the burst length does not exceed the maximum allowed burst time. This simplifies the functions of the simulator so that no segmentation of data burst is needed. Fourth, each user has an infinite buffer for its input queue. Although this might be less realistic, we are more concerned about network performance in terms of throughput, and queue length is one complication we choose to forgo. Finally, we assume a value of 0.1 ms for the modem turn around time (TAT) and a propagation delay of 30 ns (equivalent to a transmission distance of ten meters) for all the links.

B. The Physical/MAC Layer Simulations Interface

If all users are static, all communication channels are static as well. The SNRs and SIRs, which are the determining factors of error probabilities, are not changed throughout the MAC simulation. That means we could simplify the process of calculating the error probability of each packet during

simulation by 'pre-calculating' them all prior to simulation. This also saves us a lot of computational power. Using the theoretical expressions we derived in Section III-F with the SNRs and the SIRs obtained from the physical layer simulation, we can determine the repetition rate needed for every communication link to maintain a BER of 10^{-8} or below with and without interference. (Of course, the repetition rate cannot exceed 16, so that some links will fail to achieve 10^{-8} BER.) For the case when the main body of a data packet is interfered by the main body of another data packet, the repetition rate for both pairs of transceivers could change dynamically for some time depending on which pair starts the transmission first. However, the repetition rate is bound to settle down to a stable value after at most five packets exchange on both pairs (when *RR* goes from one to two, four, eight, and finally sixteen). To reduce the complexity of the repetition rate and error probability calculation, we will use the final settled value of repetition rate and apply it to MAC layer simulations whenever an interference occurs. This could introduce inappropriate repetition rate for a maximum number of three packets. Compared to tens or hundreds of packets in a data burst, this minor discrepancy is acceptable. The worst case situation is illustrated in Fig. 15.

Then, using the repetition rate under different conditions and the SNRs and the SIRs, we can calculate all the error probabilities, including the probability of missing carrier sensing detection, the probability of receiving an erroneous robust header, and the probability of receiving an erroneous main body. Each category is calculated under three different circumstances: no interference, interference by a robust header, and interference by a main body. By 'pre-calculating' all of these probabilities, we are able to 'assign' an error probability to each part of a packet in the MAC layer simulations. Upon a packet reception event, a Bernoulli trial of the error probability of that part of the packet is performed and whether the transmission is successful is determined.

C. Event-driven MAC Simulator

Realistic network simulations are required to investigate the effect of non-reciprocity of the physical layer on the performance of a multiple access network. Unfortunately, most of the packet-based network simulators available in the commercial arena and research community fail to address a realistic physical layer with appropriate modelings of the transceivers, interference, carrier sensing and impulse responses of the communication channels. To fully address the above issues of the physical layer, we wrote a customized MAC layer simulator which provides appropriate parameters for the modeling of the physical layer.

The simulator is event-driven, and the simulation time is only updated 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 begin with nothing but data arrival events from each user. The simulation can be set to finish at criteria such as a specified number of data arrivals and simulation time. Fig. 16 illustrates how it works.

D. The Finite State Machine of CSMA/CA

The finite state machine of the CSMA/CA protocol as specified in the AIr MAC specification [7] is fully implemented in the simulator. It is shown in Fig. 17 and is briefly described as follows.

The finite state machine consists of seven states, namely, *Idle*, *Contend*, *Wait For CTS*, *Wait For Data*, *Transmit*, *Receive* and *Idle Wait*. A user is in *Idle* state when it is waiting to receive or transmit data. When it has data to transmit, it is in *Contend* state to try to contend with other users for its chosen slot. After a user has transmitted a RTS and is waiting for the CTS to come back, it is in *Wait For CTS* state. When a user has successfully reserved the channel and is in the process of transmitting data, it is in *Transmit* state. A user also waits for the acknowledgment packet (ACK) to come back in *Transmit* state. On the receiver side, when a user has received a RTS and has replied a CTS, it is in *Wait For Data* state. After receiving the first data packet or detecting a valid reservation, a user is in *Receive* state to actively receive data or wait for the burst of data to complete. Finally, When a user has completed transmitting or receiving a burst of data and is waiting for either an End-of-Burst (EOB) or an End-of-Burst-Confirm (EOBC), it is in *Idle Wait* state. Every user runs the same finite state machine independently.

The event-driven-based simulation approach provides a nice way to implement the finite state machine. Only when an event occurs, appropriate actions are taken according to the state transition table and the state variables are updated.

E. Interference Modeling

We have devised an algorithm for our simulator for modeling interference occurrence. The algorithm can be implemented neatly in two functions, one to be executed when a packet is being transmitted, and one when the robust header or the main body of a packet is being received.

The interference information is contained in a four-dimensional array ([number of users]^4) called '**int_enabled**'. The four dimensions contain information of the interfering signal source, interfering signal destination, desired signal source, and desired signal destination respectively. When a packet is being transmitted, the simulator checks to see if any other users have started some other transmissions. For example, when User 1 is trying to transmit a packet to User 2, if User 3 has already started a transmission to User 4, the simulator will set **int_enabled[3][4][1][2]** and **int_enabled[1][2][3][4]** to 'True', because interference from User 3 on User 1 would also mean interference from User 1 on User 3. The interfering packet type is also recorded so that the appropriate error probability is known upon packet reception.

When a robust header or a main body is being received, say, at User 2 from User 1, the simulator checks to see if any *(x,y)* pair of **int_enabled[*x*][*y*][1][2]** is set to 'True'. Using this information, the repetition rate of the next data packet is updated if necessary, and the appropriate error probability category is picked to perform the Bernoulli trial to see if the packet is received correctly. In real-life, repetition rate update is done through CTS or ACK from the receiver, i.e., the receiver of the data informs the transmitter if any update of repetition rate is needed due to interference. For our simulation, we simplify this process by updating the information at the transmitter side upon the reception of a packet at the receiver side. Although this is not possible in real-life, it would not make a difference in our simulation results.

F. 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 as long as 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.

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 1000^{th} 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.

V. Network Simulation Results

A. User Configurations

We first attempt to illustrate the effect of non-reciprocity by a LAN with two pairs of users as shown in Fig. 18 (a). The users are situated in a 10 m $\times$ 6 m $\times$ 3 m room, as shown in Fig. 8. Two pairs of users are separated by a horizontal distance of D meters, and each pair attempts to establish an independent link. The transceivers are oriented in the horizontal plane, i.e., the x_i-y_i plane is horizontal for all users. We perform four sets of experiments with this configuration and vary D from three meters to eight meters in each set of the experiments. In Experiment 1, perfect optical parity and system parity is maintained, and wide angle transceivers (Class A, as in Table 3) are employed for each user. In Experiment 2, approximate optical parity and perfect system parity is maintained, and two of the users are replaced with narrow angle (Class B) transceivers. Optical parity is violated for two of the users in Experiment 3 by employing narrow angle transmitters and wide angle receivers (Class O), whereas system parity is violated for one pair of the users in Experiment 4 by doubling the transmitter power and halving the detector area (Class S). Each of the experiments are performed under two different background light levels: no ambient light and intense ambient light. The transceiver angular characteristics of different experiments are shown in Fig. 18 (b). The specification of the transceiver usages of the users are given in Table 4.

In addition to the configuration with two pairs of users, we also investigate the effect of non-reciprocity on some realistic network settings as shown in Fig. 19. We examine a "four users, short-range" configuration, a "four users, long-range" configuration, a "four users plus base-station" configuration and a "two groups of three users" configuration with different transceiver usages under intense ambient light in the same room as in Fig. 8. The base-station in the "four users plus base-sta-

tion" configuration is situated at the center of the ceiling with the transceiver pointing vertically downwards. It employs Class A transceiver in all experiments except in Experiment 2, which is described in more details below. The users in that particular configuration only communicate with the base-station and therefore aim their transceivers to the base-station at an angle. (The users are situated at a height of 1 m.) In all other configurations, the transceivers are oriented in the horizontal plane at a height of 1 m. For each configuration, four sets of generalized experiments are performed as described above: In Experiment 1, all users employ wide-angle transceivers with optical and system parity. Experiment 1N represents the only environment with no ambient light and is specifically designed for the "two groups of three users" configuration to illustrate the effect of ambient noise on LANs with several communicating groups. In Experiment 2, some of the users are replaced with narrow-angle transceivers but all users still employ transceivers that maintain optical and system parity (Class A, C, D, E transceivers are used). However, Experiment 2 in the "four users plus base station" configuration introduces a kind of optical parity mismatch which normally would not occur for users oriented horizontally. The base-station employs a multiple-element (Class D) transceiver and since optical parity is only maintained for the plane that includes all the axes of the transmitting elements [7], all users in this configuration have an optical parity mismatch with the base-station. In Experiment 3, two of the users violate optical parity (Classes A, C, D, E, O are used) and in Experiment 4, two of the users violate system parity (Classes A, D, S are used). The specifications of the parity violation of the users are given in Table 5. The specifications of the transceiver usages are given in Table 6 - Table 9.

In all configurations presented above except the "two groups of three users" configuration, traffic is generated only between users who would remain connected even when the transceivers are made narrow in some experiments. This ensures that packets are not simple thrown away during those experiments. Having the same traffic pattern applied across all experiments in the same configuration also ensures that the results are comparable. Random traffic is designated for the "two groups of three users" configuration to illustrate the different kind of results one would get if the users do not point their transceivers to the ones they try to establish communication with.

B. Simulation Results

1. Two Pairs of Users

When the load of the network is low, we observed that the throughput of all experiments are approximately the same, except that there is a 10% drop of throughput when optical parity is violated

in an intense ambient light environment. Violating channel reciprocity does not cause any significant adverse effects to the network throughput because simultaneous transmissions rarely occur in a lightly loaded network. These results are shown in Fig. 20.

On the other hand, a highly loaded network exhibits different behaviors under different experiments as shown in Fig. 21. When there is no ambient light, Experiment 2 gives higher throughput than Experiment 1 as D increases because the narrow-angle transmitters establish two independent links which do not interfere with each other. Compared to Experiment 2, Experiment 3 always gives a lower throughput due to the optical parity mismatch. In Experiment 3, although the transmitters of User 1 and User 3 are narrow-angle, their wide-angle receivers unnecessarily reserves the channel for the other communication link. We would also expect a drop in throughput in Experiment 3 compared to Experiment 1, but this was not observed since the degree of non-reciprocity is not terribly bad when there is no ambient light, and the CSMA/CA protocol can still barely work in this case. Experiment 4 does not exhibit a rise in throughput even though the two pairs of transceivers are well separated. This is because the system parity is violated by one pair of much stronger transmitters which are still able to introduce carrier sensing detection to the other pair when they are far away from them. When there is intense background light, we observed that the noise essentially segments the two pairs of transceivers in Experiment 2 and they communicate independently without interfering each other at all separation distances from three meters to eight meters. At short separation distances, there is a 30% drop in throughput in Experiment 3 compared to Experiment 1 due to optical parity mismatch. For instance, User 2 cannot detect the reservation packets sent by User 3 but its transmission to User 1 is overheard at User 3 because User 3 has a narrow-angle transmitter and a wide-angle receiver. Therefore, interference occurs at User 3 and the throughput is reduced. The stronger transmitters in Experiment 4 help to combat noise and give a higher throughput than Experiment 1 at short separation distances. Although the aggregate throughput could be as much as 50% higher, the stronger transmitters dominate the communication channel and the throughput of the stronger pair could be as high as 650 times the throughput of the weaker pair. This unfairness introduced by system parity mismatch is undesirable in most data networks.

2. Other User Configurations

Results from the other realistic LAN configurations show that when the users have a LOS link to most other users within their communication range, the impact of optical parity mismatch is relatively small. For example, in the "four users, short-range" configuration as shown in Fig. 19 (a), each user

has a LOS path to every other user. As a consequence, even when there is a optical parity violation, the users are still able to receive and decode the RTS/CTS packets correctly. This is illustrated in Fig. 22 (a), in which the throughput of Experiment 1 and Experiment 3 are similar.

Intuitively, unless the narrow-angle transceivers in Experiment 2 segment the network into several individual links, the throughput of Experiment 1 and Experiment 2 should be similar as well. However, this is not observed in the "four users, long-range" configuration. As shown in Fig. 22 (b), Experiment 2 has a factor of twenty reduction of throughput compared to Experiment 1 for long-range transmissions. This decrease is attributed to the narrow-angle transceivers employed in Experiment 2. For long-range communications, the SNR is so low that the portion of light reflected off the walls becomes significant. Wide-angle transceivers benefit from additional reflections and give a much higher throughput.

The increase in aggregate throughput in Experiment 4 is evident in all configurations simulated as described above. However, the unfairness is not as bad as in the configuration with two pairs of users, because there is not a particular pair of users with stronger transmitters dominating the weaker pair(s). On average, the stronger transmitters have an individual throughput four to five times higher than the weaker ones. The "four users plus base-station" configuration consists of users who tilt and point to the base-station on the ceiling. As mentioned above, the optical parity requirement only specifies the matching of angular characteristics on the plane that includes all axes of the transmitting elements (the "horizontal" plane). For users (and base-station) employing a multiple-element transmitter communicating to each other at a tilted angle, optical parity is indeed violated. This effect is easily seen in Experiment 2 of this configuration when the base-station violates optical parity with all the users. The throughput is reduced up to a factor of thirteen, as shown in Fig. 22 (c). So far we have assumed the users only send packets to those they have a LOS link with. If the users send random traffic to every other user, throughput reduces greatly due to packet loss. Interestingly, the ambient noise increases the aggregate throughput in this case by making the already poor quality links impossible to establish communications. This effect is most easily seen in the "two groups of three users" configuration in which the noise segments the room into two distinct groups. As illustrated in Fig. 22 (d), Experiment 1N is performed with no ambient noise and its throughput is lower than other experiments performed under intense ambient light.

C. Recommendations

We summarized our results in the form of transceiver design recommendations as follows.

- Transceivers should maintain both optical parity and system parity to optimize network throughput and fairness.

- If the traffic patterns of the users are very directed, narrow-angle transceivers generally give higher aggregate throughput for short range transmissions (up to about six meters).

- If transmitted power is not a great concern (for example, when the users do not have a high mobility such that they can use the wall outlets for power supply), wide-angle transceivers should be employed instead of narrow-angle ones for long range transmissions (eight to ten meters), because the reflections of the light from the walls, ceiling and floor increase the SNRs of the channels and this portion becomes significant for long range transmissions.

- The specification on optical parity only attempts to match the transmission and reception angular patterns in the "horizontal" plane. As a consequence, for network designs in which the users (and possibly a base-station) are not all located in the same horizontal plane, usage of multiple-element transmitters should be avoided.

- Given a Lambertian transmitter and a multiple-element transmitter, both maintaining optical parity and system parity, the former one always consumes a higher averaged transmitted power. In certain situations, it would introduce unfairness to the multiple-element users due to additional reflections off the ceiling and the floor, although this effect is subtle. However, when power consumption is a concern and usage is confined to the horizontal plane, the multiple-element configuration would be a better choice.

VI. Conclusions and Future Work

A. Conclusions

We have presented a detailed analysis and modeling of the physical layer of the upcoming AIr standard. Physical layer and MAC layer simulations were performed to investigate the effect of non-reciprocity on network throughput. We provide appropriate modeling of the physical layer *within* the MAC layer simulations to give realistic network simulations. The results show that the impact is sensitive to LAN configuration and background 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. It is recommended that transceivers should be designed to achieve channel reciprocity and different classes of transceivers are chosen based on the network usage and LAN configuration.

B. Future Work

The current MAC layer protocol employed in AIr is not optimal in solving the hidden terminal problem and could introduce additional impediments to network performance and fairness other than those from non-reciprocity. There are some studies [14],[15],[16] concerning the improvement of the CSMA/CA protocol. To further isolate the effect of non-reciprocity, these protocols could be implemented and tested with the same framework in our study.

There could also be improvement on the coding scheme, for instance, employing odd numbers for the repetition rate would make more effective decisions on majority voting. The impact of non-reciprocity on better coding schemes might differ.

Finally, the analysis and simulation of this study only involve the physical layer and the MAC layer of the network. It would be an interesting extension to include the functionalities of the higher layers of the network, for example, automatic retransmission request on the transport layer.

VII. Acknowledgments

VIII. 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] Infrared Data Association, "Serial Infrared Physical Layer Link Specification," Version 1.1, October 1995.

[4] Infrared Data Association, "Serial Infrared Link Access Protocol (IrLAP)," Version 1.1, June 1996.

[5] Infrared Data Association, "Link Management Protocol (IrLMP)," Version 1.1, January 1996.

[6] Hewlett-Packard Company and IBM Corporation, "Request for Comments on Advanced Infrared (AIr) IrPHY Draft Physical Layer Specification," Version 0.4, January 1998.

[7] Hewlett-Packard Company and IBM Corporation, "Request for Comments on Advanced Infrared (AIr) IrMAC Draft Protocol Specification," Version 0.2, July 1997.

[8] P. Karn, "MACA - A New Channel Access Method for Packet Radio," *ARRL/CRRL Amateur Radio 9th Computer Networking Conference*, September 1990.

[9] 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.

[10] W.T. Welford and R. Winston, *High Collection Nonimaging Optics*, Academic Press, San Diego, 1989.

[11] 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.

[12] 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.

[13] J. G. Proakis, *Digital Communications*, Third Edition, Mc-Graw-Hill, New York, 1995.

[14] V. Bharghavan, A. Demers, S. Shenker, L. Zhang, "MACAW: A Media Access Protocol for Wireless LANs," *ACM SIGCOMM '94 Conference on Communications Architectures, Protocols and Applications*, August 31 - September 2, 1994.

[15] C. L. Fullmer, J. J. Garcia-Luna-Aceves, "Floor Acquisition Multiple Access (FAMA) for Packet-Radio Networks," *ACM SIGCOMM '95*, August 28 - September 1, 1995.

[16] C. L. Fullmer, J. J. Garcia-Luna-Aceves, "Complete Single-Channel Solutions to Hidden Terminal Problems in Wireless LANs," *Proceedings of ICC'97 - International Conference on Communications*, vol. 2, pp. 575-579, June 1997.

IX. Tables

Table 1: Representation of 4-PPM symbols

Signal	Symbol	Data Bit Pair
$s_0(t)$	1000	00
$s_1(t)$	0100	01
$s_2(t)$	0010	10
$s_3(t)$	0001	11

Table 2: Dynamic trade-off between data rate and transmission range using Variable Repetition coding.

Repetition Rate (*RR*)	1	2	4	8	16
Bit Rate (Mb/s)	4	2	1	0.5	0.25
Range of Typical LOS Link (m)	4 ────────────────────▶ 8				

Table 3: Classes of transceivers designed to satisfy optical parity and system parity. Here, $\alpha_{1/2,\,i}$ denotes the half-angle at half-intensity of a single transmitting element. All transmitting elements lie in the x_i-z_i plane ($\phi_{im} = 0$) of Fig. 6. Each transceiver class employs $N = 1$ receiving element oriented along the z_i axis ($\theta_{i1} = 0$) of Fig. 6. In all classes, the receiver optical concentrator has refractive index $N_r = 1.44$.

Transceiver Class	Transmitter				Receiver	
	Number of Elements M	Average Power per Element P_t/M (mW)	Half-Angle $\alpha_{1/2,\,i}$ (°)	Tilt Angles θ_{im}, $m = 1, ..., M$ (°)	Detector Area per Element A (mm^2)	Concentrator Cutoff Angle $\beta_{c,\,i}$ (°)
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

Table 4: Transceiver classes of the users in the "two pairs of users" configuration. The classes are defined in Table 3.

	User 1	User 2	User 3	User 4
Expt. 1	A	A	A	A
Expt. 2	B	A	B	A
Expt. 3	O	A	O	A
Expt. 4	A	A	S	S

Table 5: Transceivers that violate optical or system parity in Experiments 3 and 4 with various LAN configurations shown in Fig. 19.

LAN Configuration	Violate Optical Parity	Violate System Parity
(a) Four users, short range	User 1, User 2	User 3, User 4
(b) Four users, long range	User 2, User 3	User 3, User 4
(c) Four users plus base station	User 2, User 3	User 3, Base Station
(d) Two groups of three users	User 1, User 4	User 3, User 4

Table 6: Transceiver classes of the users in the "four users, short range" configuration. The classes are defined in Table 3.

	User 1	User 2	User 3	User 4
Expt. 1	A	A	A	A
Expt. 2	A	D	C	E
Expt. 3	A	A	O	O
Expt. 4	S	S	A	A

Table 7: Transceiver classes of the users in the "four users, long range" configuration. The classes are defined in Table 3.

	User 1	User 2	User 3	User 4
Expt. 1	A	A	A	A
Expt. 2	C	E	D	A
Expt. 3	O	A	O	A
Expt. 4	S	S	A	A

Table 8: Transceiver classes of the users in the "four users plus base station" configuration. The classes are defined in Table 3.

	User 1	User 2	User 3	User 4	Base
Expt. 1	D	A	A	D	A
Expt. 2	E	D	A	C	D
Expt. 3	O	A	O	D	A
Expt. 4	S	A	A	D	S

Table 9: Transceiver classes of the users in the "two groups of three users" configuration. The classes are defined in Table 3.

	User 1	User 2	User 3	User 4	User 5	User 6
Expt. 1	D	D	D	A	A	A
Expt. 1N	D	D	D	A	A	A
Expt. 2	C	D	A	E	A	C
Expt. 3	C	D	O	E	A	O
Expt. 4	D	D	S	S	A	A

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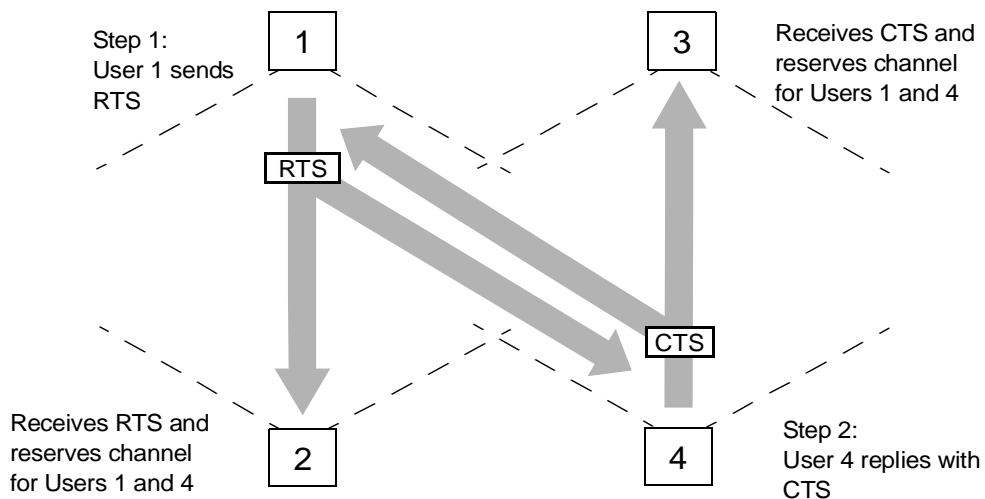
Fig. 1. The use of RTS/CTS exchange to alleviate the hidden terminal problem. Without the RTS/CTS exchange, User 3 is unaware of User 1's transmission, and collision will occur at User 4 if User 3 starts to transmit. Dashed lines denote schematically the range of angles within which the terminals transmit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

Room Setting
Transceiver Configuration
Transmitted Power
Ambient Light Levels

SNR
SIR
Repetition Rate

Frame-Error Rate P_f

Throughput
Fairness

| Physical Layer (Propagation and Noise) Simulator | Theoretical Expressions for Error Probability | Media-Access Control Simulator (Discrete-Event) |
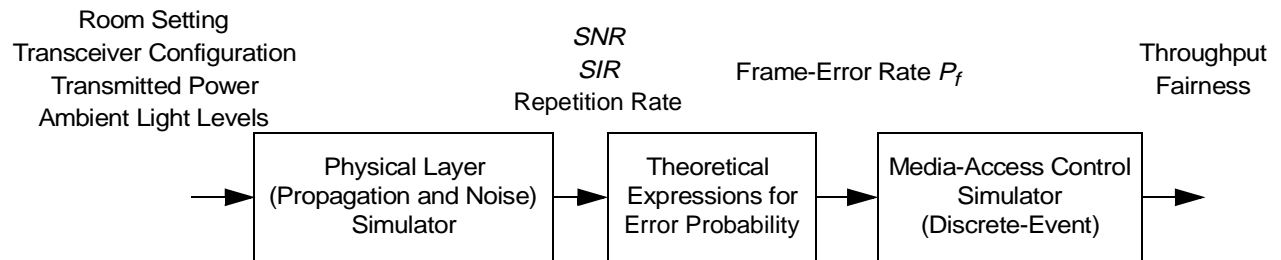
Fig. 2. Overview of the simulation flow.
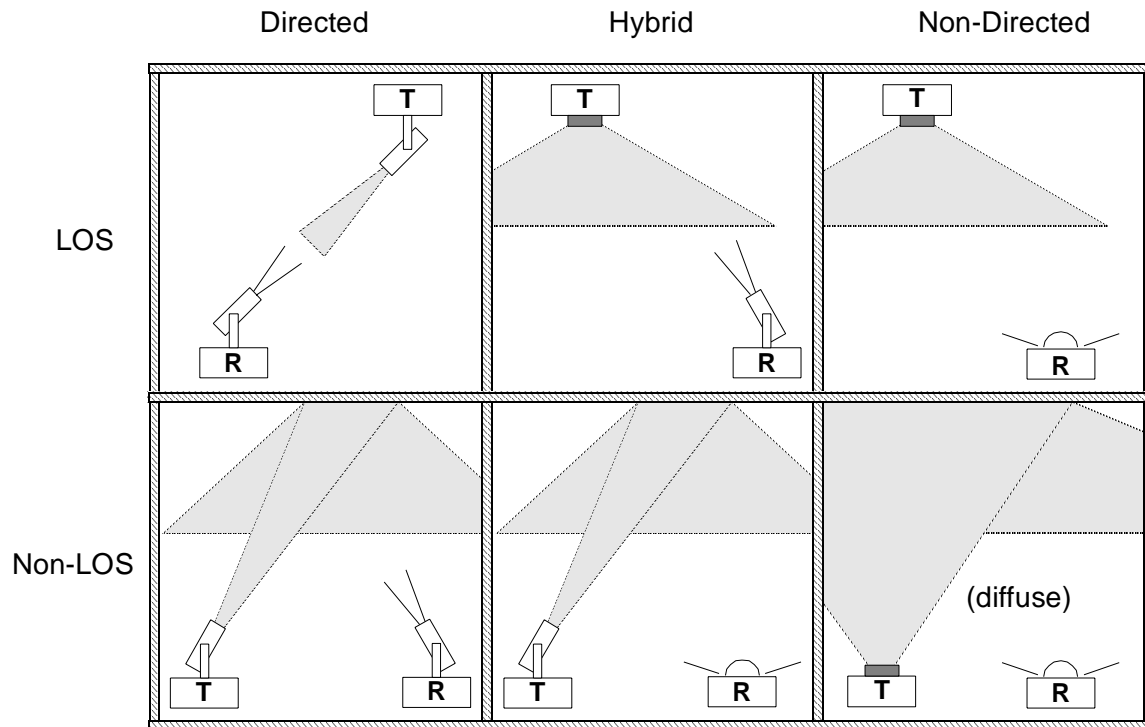
Fig. 3. Classification of infrared link configurations

Optical Power
$X(t)$ $\longrightarrow$ $\boxed{R\ h(t) \leftrightarrow R\ H(f)}$ $\longrightarrow$ $\oplus$ $\longrightarrow$ Photocurrent
$Y(t)$

Signal-Independent
Noise $N(t)$
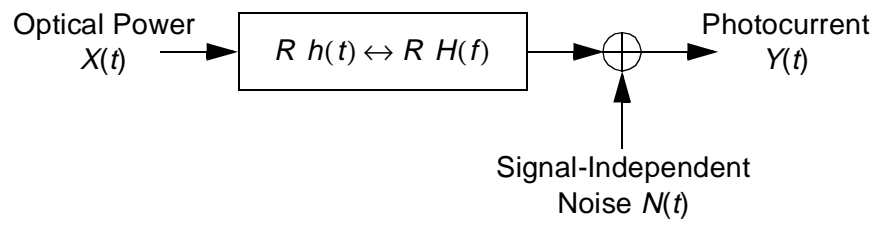
Fig. 4. Modeling an infrared link as a baseband linear, time-invariant system having impulse response $h(t)$, with signal-independent additive noise $N(t)$. The photodetector has responsivity R.
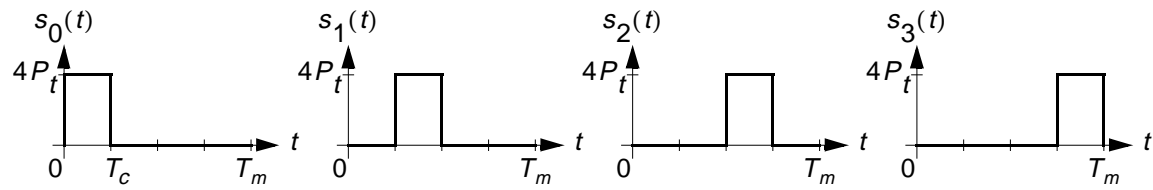
Fig. 5. 4-PPM signal set. Each symbol encodes two bits of information. The symbol period is T_m, while the chip period is T_c.
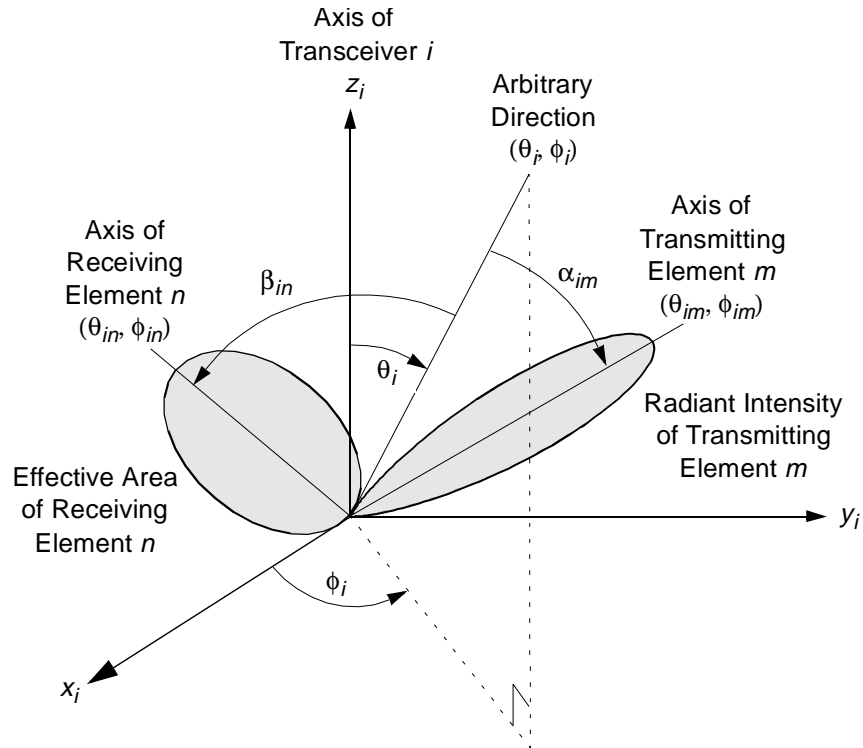
Fig. 6. Coordinate system used to describe transceivers. The Cartesian coordinates (x_i, y_i, z_i) are local to transceiver i, which has its symmetry axis along z_i, while (θ_i, ϕ_i) are the spherical polar coordinates of an arbitrary direction with respect to transceiver i. The axes of transmitting 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). In all transceiver designs considered in this report, $\phi_{im} = \phi_{in} = 0$, i.e., transmitting and receiving elements lie in the x_i-z_i plane.

(a)

(b)
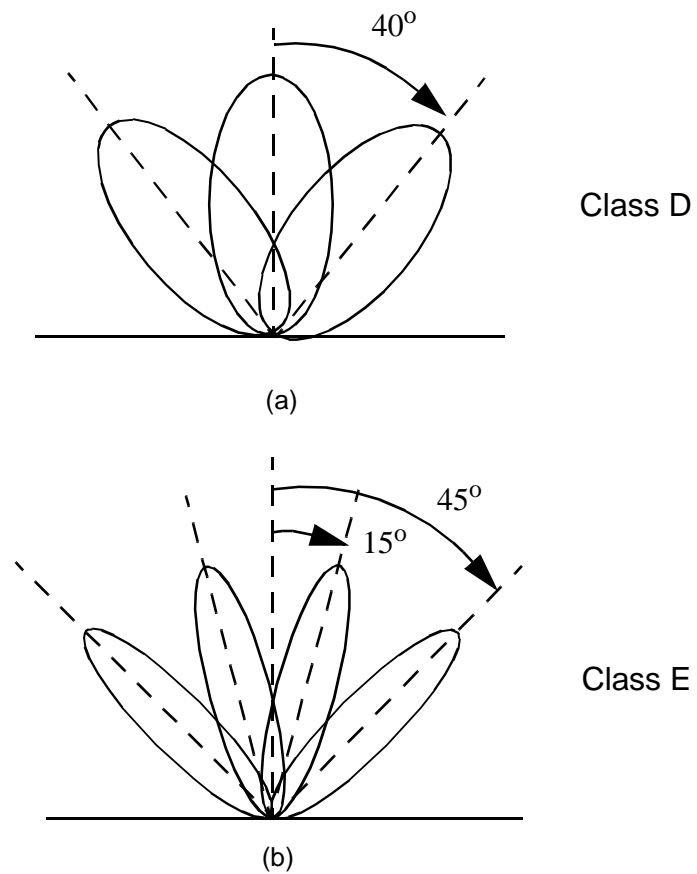
Fig. 7. Implementation examples involving multiple transmitting elements for achieving optical parity. (a) Three 30^o half-angle LEDs are used to match the Lambertian (90^o FOV) receiver; (b) Four 18^o half-angle LEDs are used to match the Lambertian receiver.
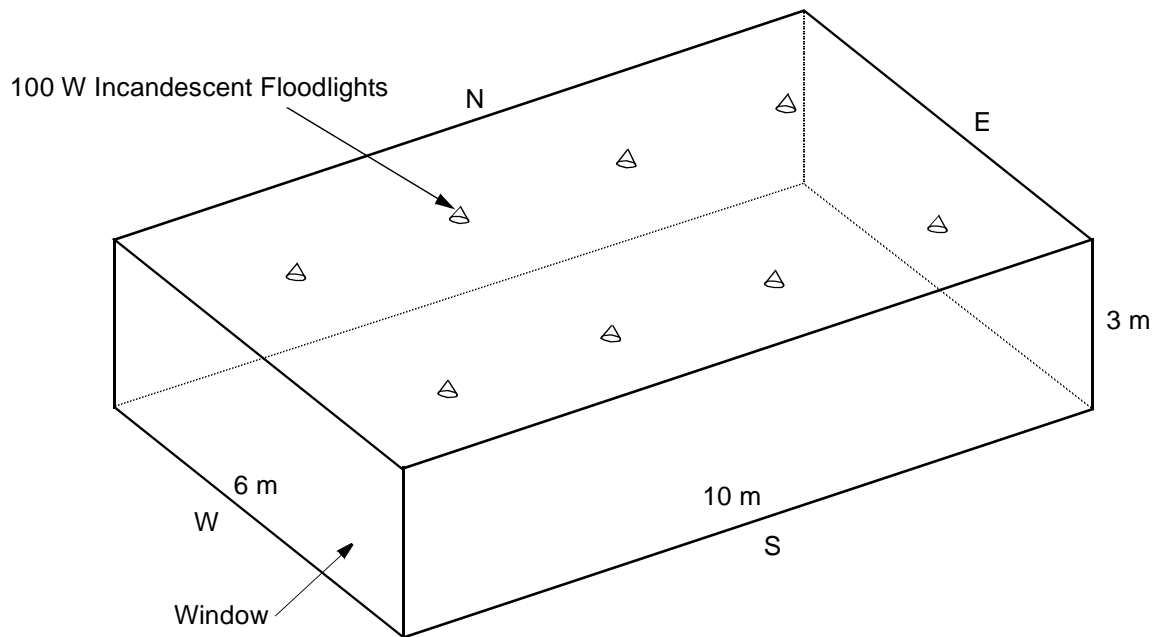
Fig. 8. Room in which various network configurations are simulated. The west wall is a single large window. The walls, ceiling, floor and window are modeled as Lambertian reflectors of reflectivity 0.6, 0.7, 0.2 and 0.04 respectively. These surfaces also act as ambient light sources. They are modeled as planar Lambertian transmitters, and measurements [12] suggest a spectral radiant emittance (S as in (13)) of 0.2 W/m^2/nm for the window, 0.01 W/m^2/nm for the east wall, linearly from 0.03 W/m^2/nm at the west edge to 0.01 W/m^2/nm at the east edge for the ceiling and linearly from 0.02 W/m^2/nm at the west edge to 0.01 W/m^2/nm at the east edge for the north wall, south wall and the floor. The ceiling floodlights are modeled to have generalized Lambertian radiation pattern of order $l_{lamp} = 2$ with optical spectral densities of $p_{lamp} = 0.037$ W/nm.
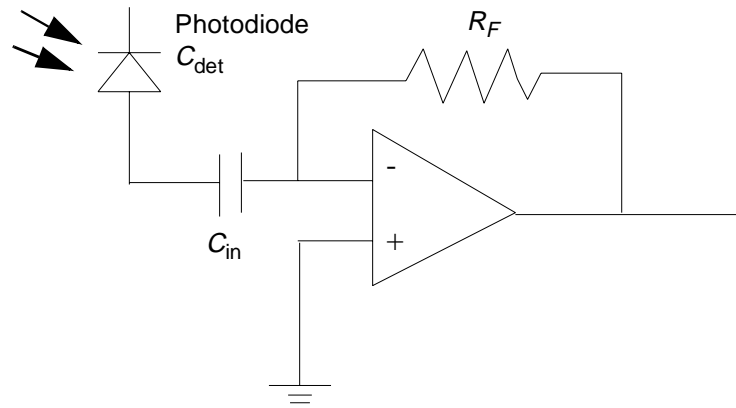
Fig. 9. The preamplifier circuitry of the receivers. We assume an op-amp is used for photocurrent amplification. R_F is the feedback resistance and the total capacitance is given by ($C_{in} + C_{det}$).
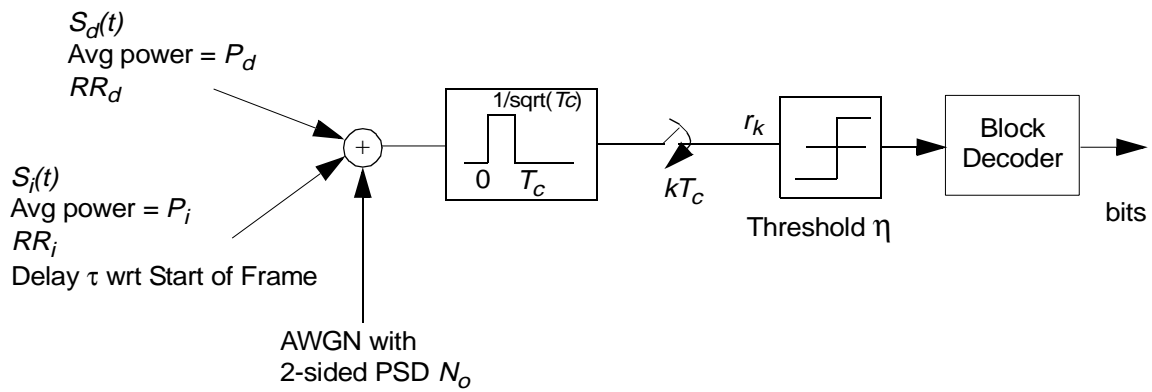
Fig. 10. Receiver Configuration: $S_d(t)$ is the desired signal and $S_i(t)$ is the interfering signal. Additive white Gaussian noise is added to the signals. A unit-energy matched filter is used to perform chip-by-chip demodulation. Hard decision decoding is employed for implementation simplicity. The block decoder uses majority voting on valid symbols to decode the repetition codes.
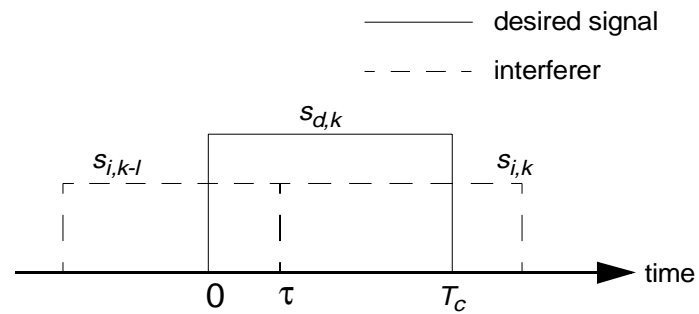
Fig. 11. Relationship between the desired signal and the interfering signal. The start of the interfering signal can fall anywhere within a chip period, i.e. τ is uniformly distributed between zero and T_c.
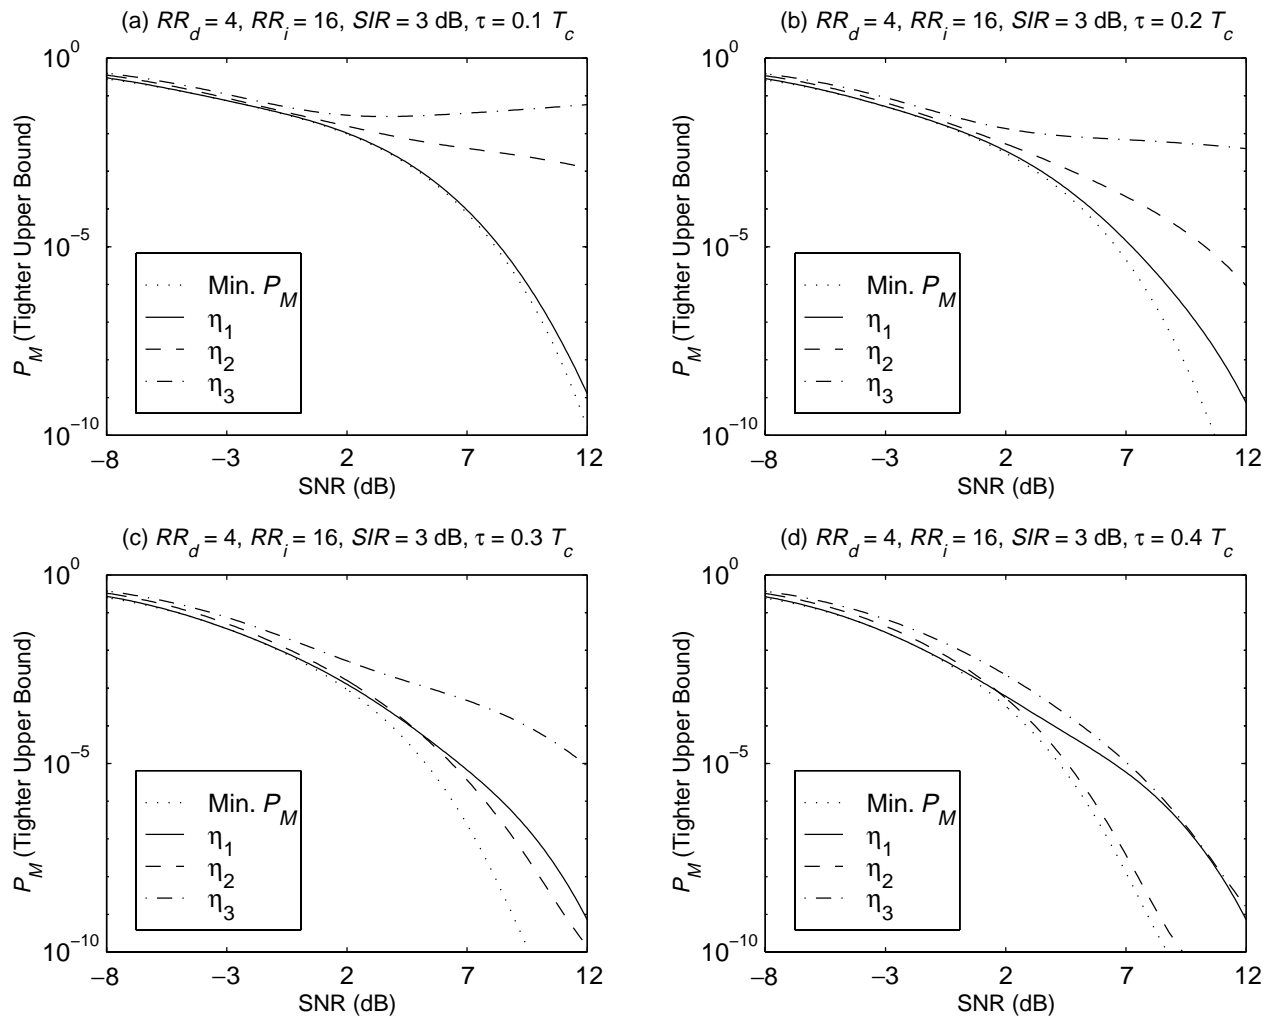
Fig. 12. Symbol-error probability obtained for different threshold choices at different values of the time delay τ. "Min. P_M" denotes the threshold, computed numerically, that minimizes P_M. Symbol-error probabilities are calculated using the tighter upper bound given by (40) and (41)-(47).
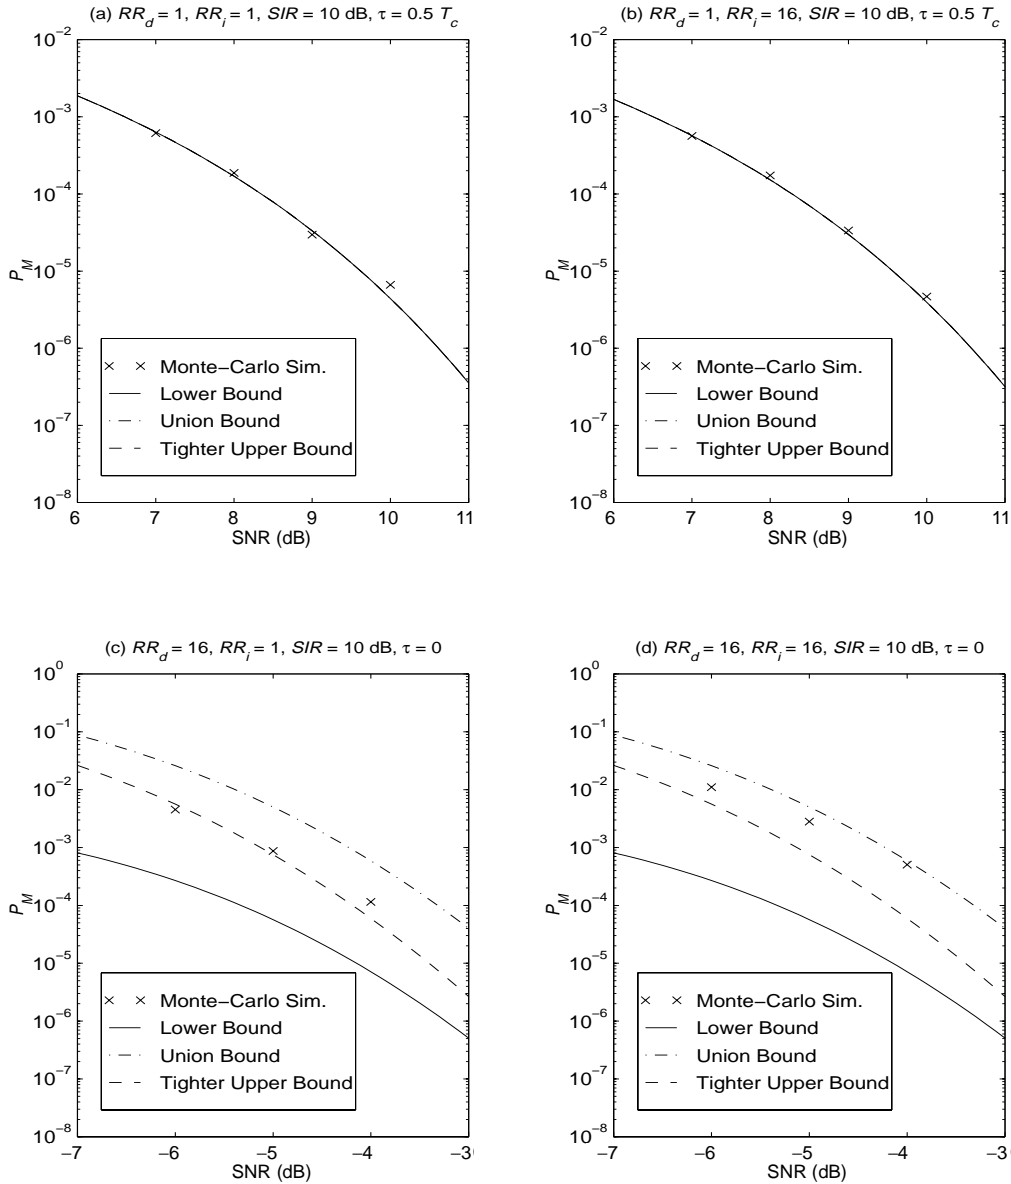
Fig. 13. Comparison of bounds (40)-(47) on symbol-error probability with Monte-Carlo simulations. The threshold η_1 is used for chip decisions in the simulations. In some cases, the simulated error probabilities are higher than the tighter upper bound, because that bound is derived based on an approximation that fails to capture the correlated nature of the interfering signal.

Preamble	Sync	Robust Header	Main Body
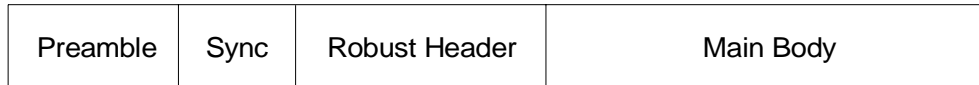
Fig. 14. IrMAC frame format. The preamble is used for initial carrier sensing, symbol clock synchronization and chip clock phase acquisition. The synchronization (sync) field identifies the first symbol slot in the robust header. The robust header contains all necessary information pertaining to the IrPHY and IrMAC layers. Data is put in the main body.
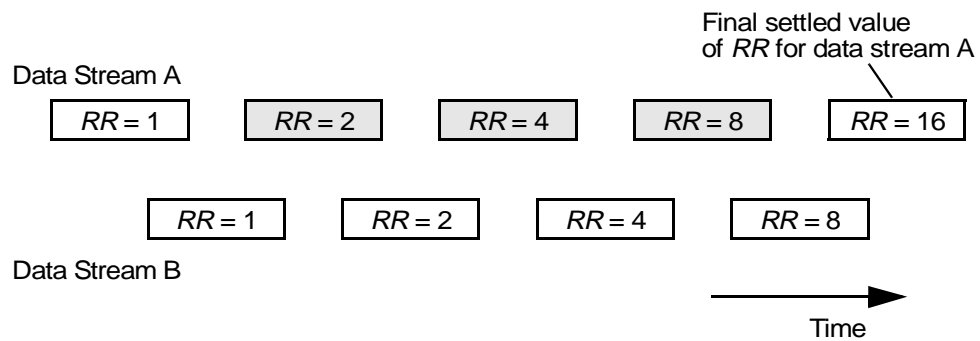
Fig. 15. Dynamic update of *RR* when a data packet stream is interfered by another data packet stream. Since the 'pre-calculation' of *RR* gives the final settled value in case of an interference, there could be a maximum of three packets (shaded above) using the incorrect *RR* in the simulation.
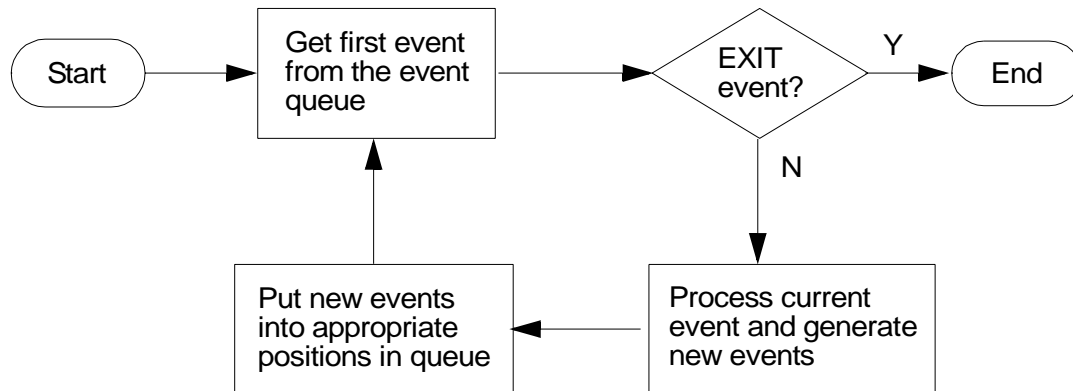
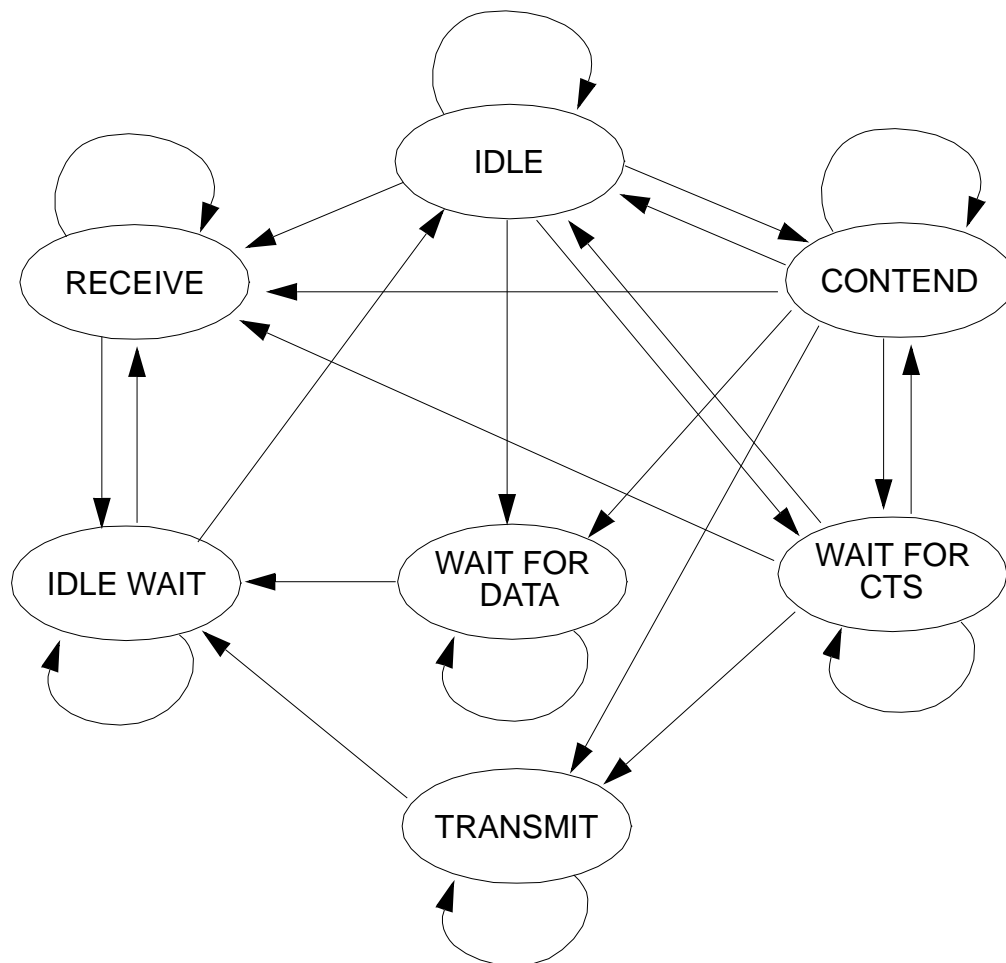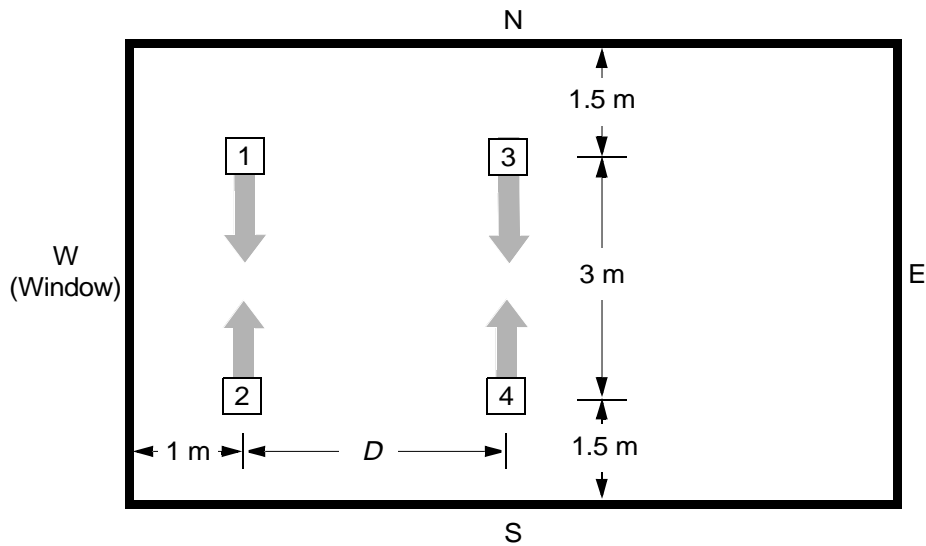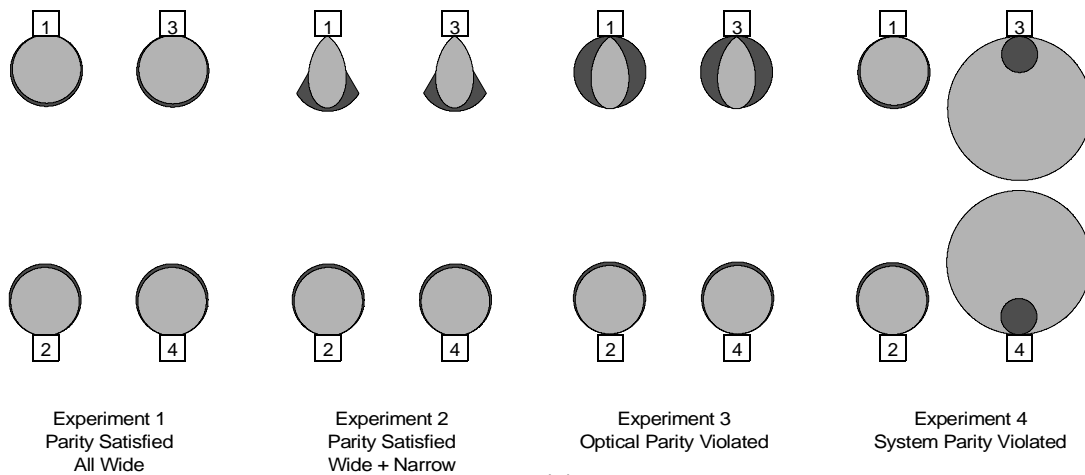**Fig. 16.** Flow chart of an event-driven simulator.

Fig. 17. Finite state machine of the CSMA/CA protocol employed in the IrMAC specification.

N

1.5 m

1

3

W
(Window)

3 m

E

2

4

1 m

D

1.5 m

S

(a)

1   3   1   3   1   3   1   3

2   4   2   4   2   4   2   4

Experiment 1
Parity Satisfied
All Wide

Experiment 2
Parity Satisfied
Wide + Narrow

Experiment 3
Optical Parity Violated

Experiment 4
System Parity Violated

(b)

Fig. 18. (a) Configuration with two pairs of users. The four transceivers are aligned in a horizontal plane within the room illustrated in Fig. 8. (b) Angular characteristics of the transceivers in Experiments 1 to 4. Light shading respresents the transmitter radiant intensities, while dark shading represents the receiver effective light-collection areas.
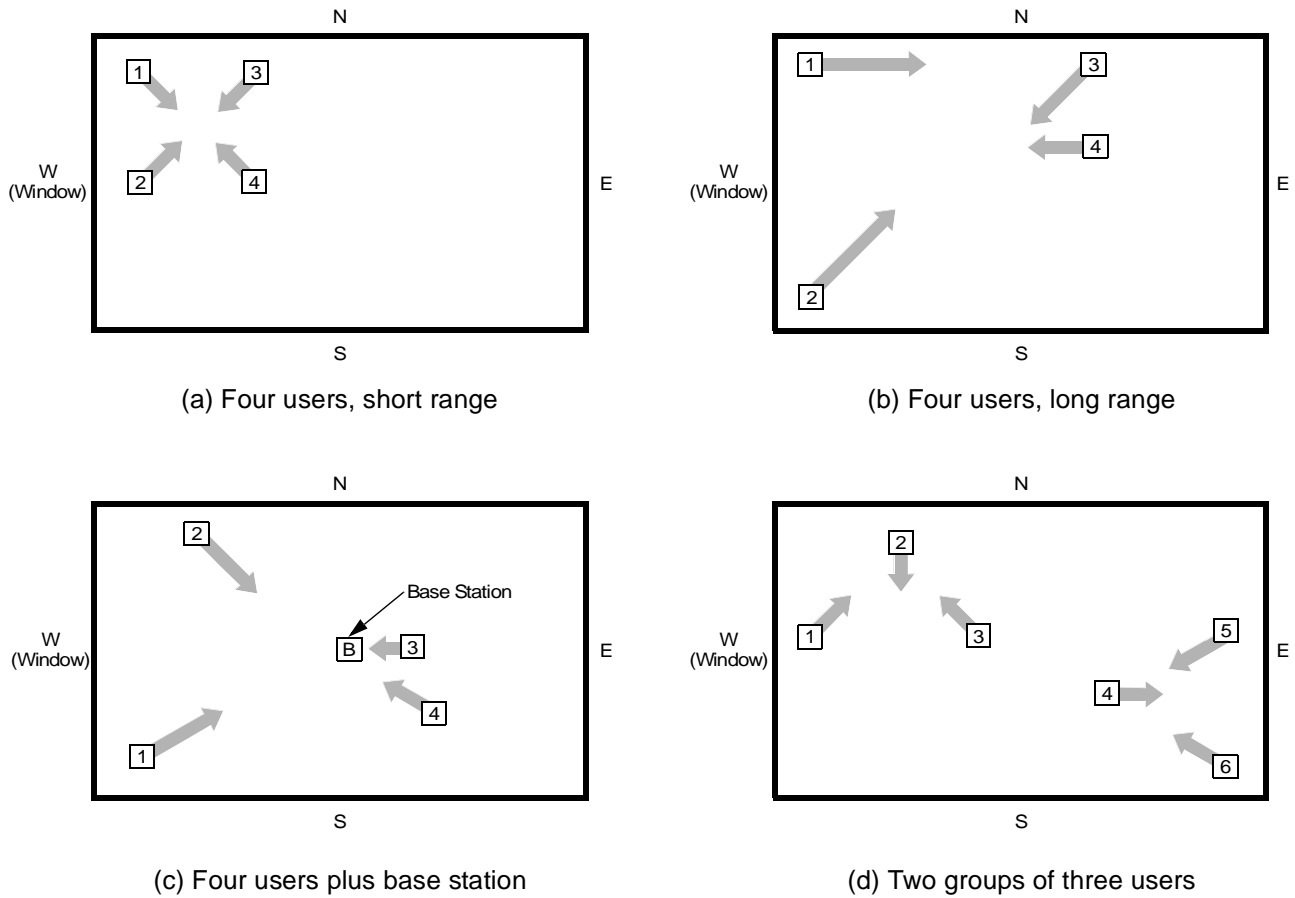
(a) Four users, short range

(b) Four users, long range

(c) Four users plus base station

(d) Two groups of three users

Fig. 19. Various configurations of users, drawn to scale within the room illustrated in Fig. 8. (a) Four users, short range, (b) Four users, long range, (c) Four users plus base station, (d) Two groups of three users.
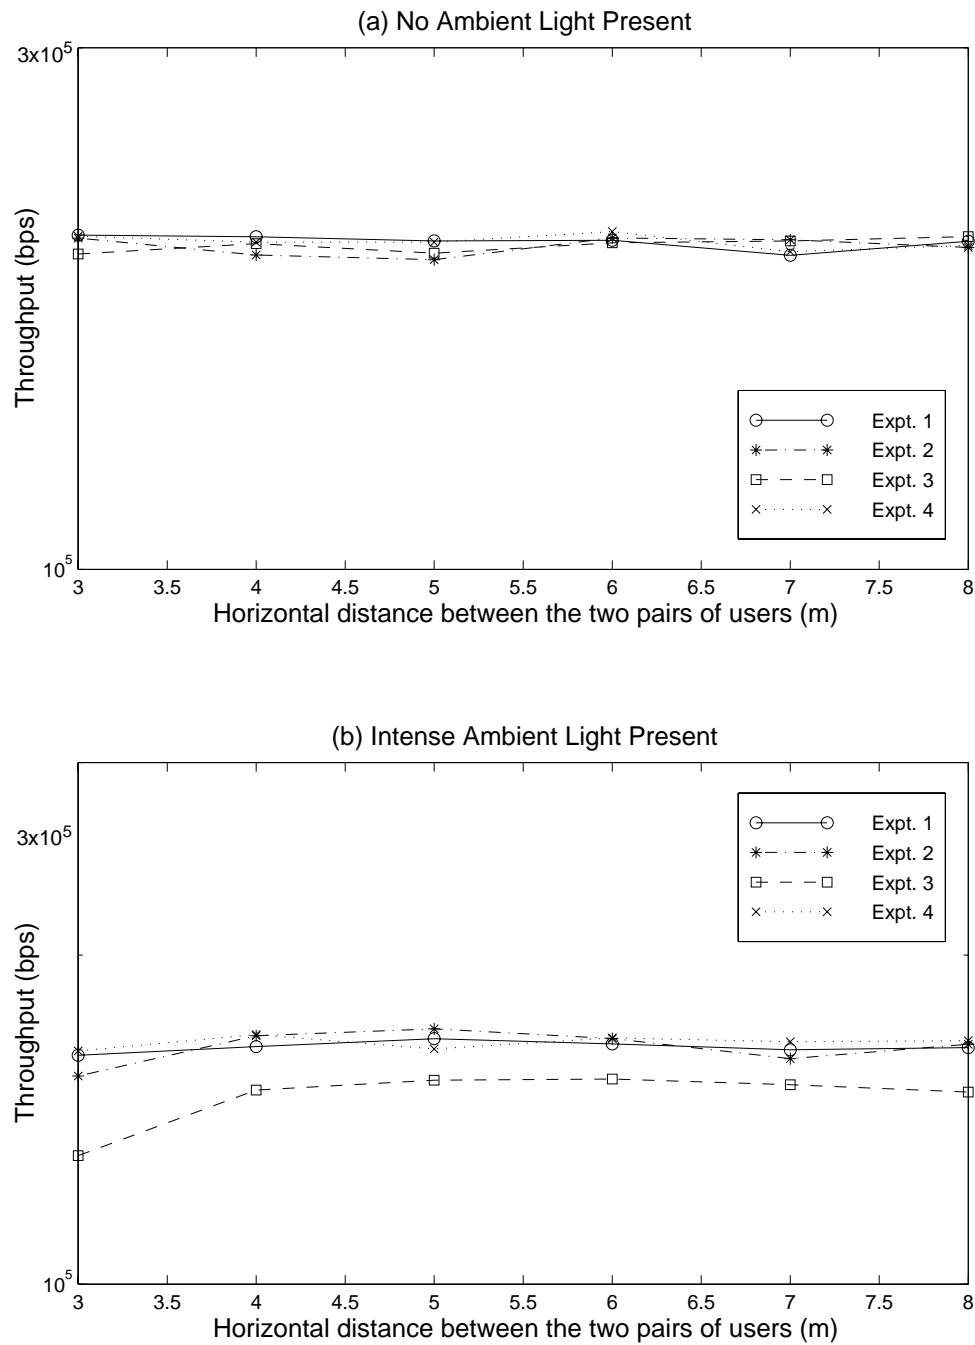
(a) No Ambient Light Present

(b) Intense Ambient Light Present

Fig. 20. Comparison of network throughputs for various experiments with two pairs of users, as shown in Fig. 18(a), in a lightly-loaded network.

(a) No ambient light present
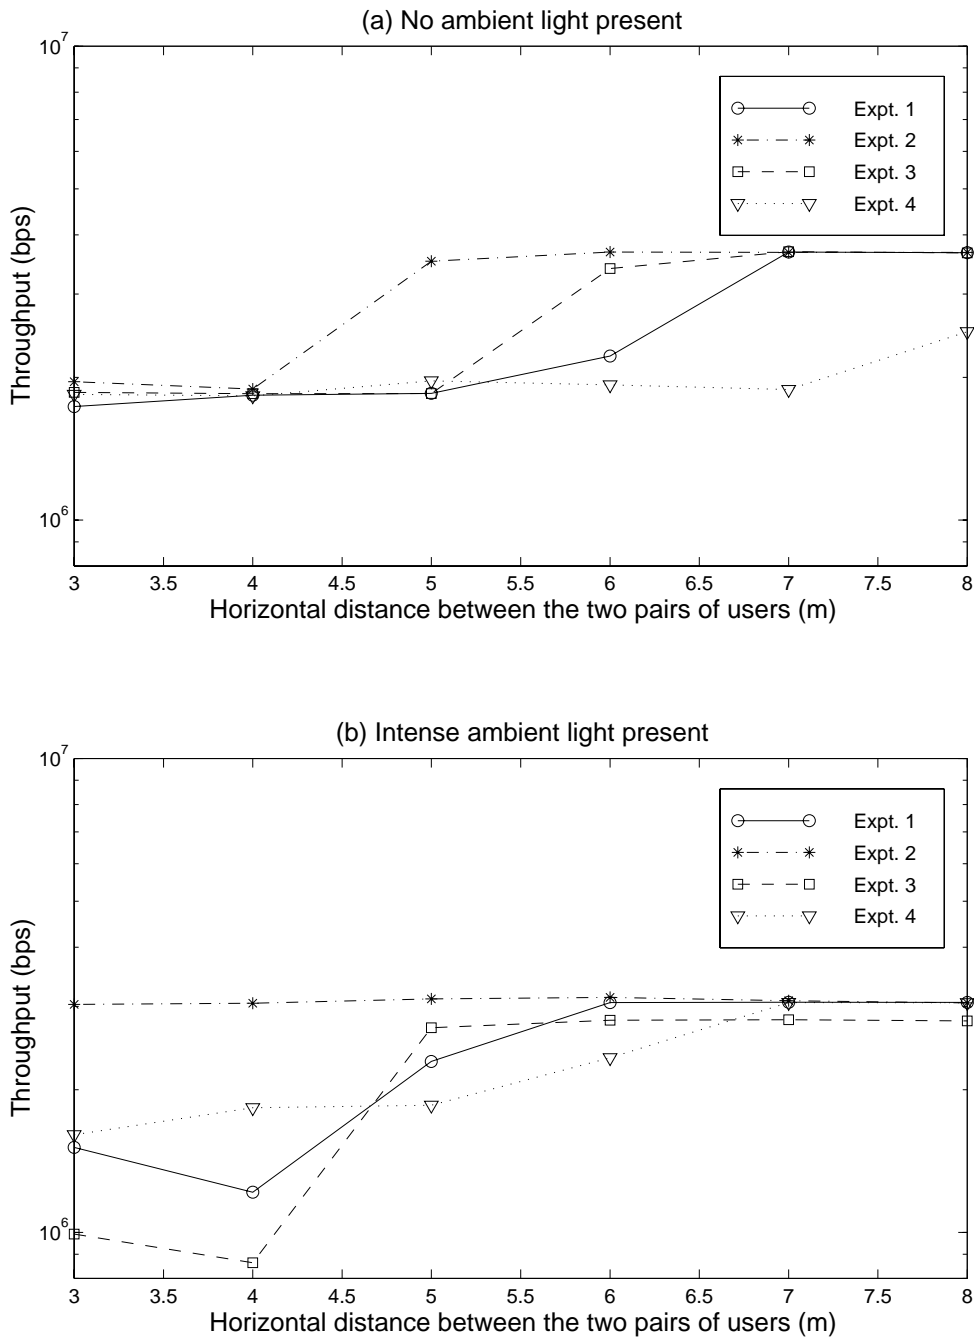

(b) Intense ambient light present

Fig. 21. Comparison of network throughputs for various experiments with two pairs of users, as shown in Fig. 18(a). The network is highly loaded, with a mean data burst interarrival rate of ten per second.
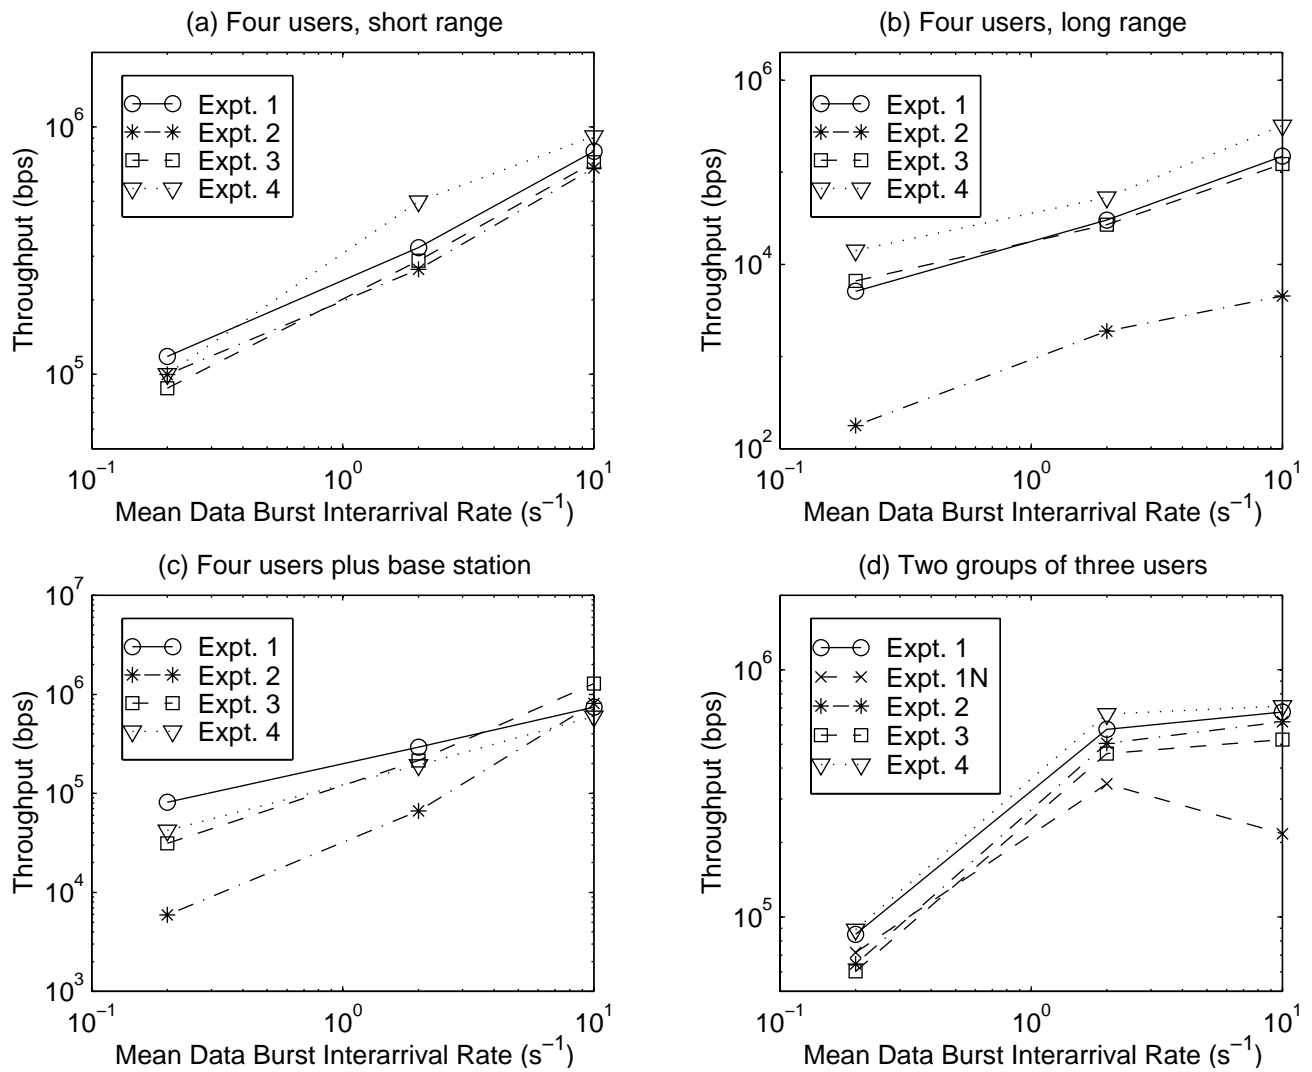
Fig. 22. Comparison of network throughputs for various experiments with the configurations shown in Fig. 19. (a) Four users, short range, (b) Four users, long range, (c) Four users plus base station, (d) Two groups of three users.