

On the Impact of Far-Away Interference on Evaluations of Wireless Multihop Networks

Douglas M. Blough
Georgia Institute of
Technology
Atlanta, US
doug.blough@gatech.edu

Claudia Canali
University of Modena and
Reggio Emilia
Modena, ITALY
claudia.canali@unimore.it

Giovanni Resta
IIT-CNR
Pisa, ITALY
g.resta@iit.cnr.it

Paolo Santi
IIT-CNR
Pisa, ITALY
paolo.santi@iit.cnr.it

ABSTRACT

It is common practice in wireless multihop network evaluations to ignore interfering signals below a certain threshold, which typically come from far-away interferers. However, the amount of error introduced by this approximation has been mostly disregarded in the literature. In this paper, we address this fundamental issue both theoretically and through extensive, packet-level simulations. We start by defining a bounded version of the well-known physical interference model, in which interference generated by transmitters located beyond a certain distance s from a receiver is ignored. We then formally prove that, under the assumption of constant transmitter density per unit area, the aggregate interference generated by far-away interferers converges to 0 if and only if s is an unbounded, increasing function of the number of nodes. This analysis also yields a lower bound on neglected interference that is approximately two orders of magnitude greater than the noise floor for typical parameter values and a surprisingly small number of nodes. These results suggest that the way interference is commonly dealt with in evaluation of wireless multihop networks could be highly inaccurate. We next validate this claim through extensive simulations done with a widely-used packet-level simulator (GTNetS), considering 802.11 MAC with both CBR and TCP traffic in networks of varying size and topology. The results of these simulations fully confirm the theoretical findings: errors in evaluating aggregate throughput when using the default interference model reached up to 210% with 100 nodes, and errors in individual flow throughputs were far greater.

1. INTRODUCTION

Multihop wireless networks are used in many different contexts, from mobile ad hoc networks to wireless sensor networks to wireless mesh networks. In such networks, spatial

reuse of a wireless communication channel, in which two links that are sufficiently far apart are used concurrently in the same channel, is possible. Maximizing spatial reuse is highly desirable in order to try to maximize achievable throughput in the network. Unfortunately, the amount of spatial reuse that can be exploited is strictly limited by interference that is generated when multiple simultaneous transmissions occur in the same channel.

Wireless interference has been traditionally modeled in terms of the communication graph, where nodes within 1 or 2 hops from a communicating link interfere with its transmission, and thus must be silent. Later, a somewhat more realistic model of interference, called the *protocol model* [6], has been considered. Here, a geometric notion of *interference range* is defined — the receiver of a link must not be within the interference range of another concurrent transmitter. These models suffer from the limitations that (i) interference is considered in a ‘binary’ sense (interference either totally eliminates the ability to communicate or is non-existent), and (ii) interference is considered only between *pairs* of nodes or links. In reality, wireless interference is much more complex. Whether a communication is successful depends on whether signal power exceeds the sum of the interference powers plus noise by a threshold that is a property of the physical layer radio design. This SINR (signal to interference plus noise ratio)-based model is known as the *physical interference model* [6]. The complexity here is that interference is neither binary nor pairwise; aggregated interference from all communicating nodes must be considered to decide whether a communication is successful. Theory aside, recent performance studies with 802.11-based mesh networks also demonstrate that multiple interferers must be considered to evaluate interference limited capacity of a link [4].

While physical interference models that account for all possible transmissions throughout the network are the most accurate, such models are very complex. This complexity manifests itself in several ways, e.g., NP-hardness of many important problems when considering physical interference [5], and use of approximate physical interference models in packet-level network simulators [16, 18, 21] to prevent simulation times from being slowed down by large factors. There are several approximations made in these simulators. All major simulators (ns2, GTNetS, and Glo-

mosim/QualNet) ignore individual interference contributions below a certain threshold (this can be implemented in different ways but can be thought of as an “interference range” or “radio range” beyond which interference is ignored). Inside the considered interference range, the simulators either accumulate all interference (Glomosim/QualNet) or consider only single interferers for a given transmission (ns2 and GTNetS). One of the main goals of this paper is to study the effects of these approximate interference models and to evaluate how much slowdown in simulation execution can be expected if more accurate models are used.

The first part of the paper concentrates on analytically evaluating the impact of the limited interference range assumption. We prove that, if the interference range is set to be a constant (independent of the number of nodes n in the network), the neglected interference remains large enough to cause significant errors in the accuracy of the model. This corresponds to the standard assumption in the protocol interference model and in all network simulators of which we are aware. For a constant transmitter density, we also prove that, if the interference range is an arbitrary unbounded increasing function of n , the neglected interference vanishes as $n \rightarrow \infty$, meaning that the approximate interference model approaches the accuracy of the true interference model asymptotically.

The second part of the paper presents simulation results from a packet-level simulator (GTNetS), which we modified to accumulate interference inside the interference range, while keeping the exact range as a tunable parameter. This allowed us both to validate our analytical results and to investigate trade-offs between simulation accuracy and simulation time by varying the interference range. The results of this evaluation demonstrate very large errors in evaluation of both packet delivery time and throughput for the commonly used interference models under several scenarios with 802.11 MAC and both CBR and TCP flows. Errors are considerable for CBR flows and even larger for TCP flows, reaching 200% in some cases. In terms of simulation time, using the most accurate model *does* have a significant negative impact. There is a roughly two orders of magnitude increase in simulation time for CBR flows and a roughly one order of magnitude increase for TCP flows.

Our results can be considered negative in that they indicate that (i) simulation-based evaluations of wireless multihop networks carried out to date are likely to be off by large amounts, and (ii) making these evaluations more accurate is likely to require substantially increased simulation times. However, at the end of the paper, we discuss several possibilities that can be explored, as future work, to improve simulation accuracy without a dramatic increase in simulation time.

2. RELATED WORK AND CONTRIBUTION

Several models for wireless interference have been proposed in the literature, ranging from pairwise models such as graph-based models [19] and the protocol model [6], to more complex and accurate models such as the physical [6] and generalized physical [11] interference models. A problem in using complex models such as physical and generalized physical is that interference is not confined within a bounded region around the transmitter, but it extends throughout the

network, possibly corrupting message reception at far-away receivers.

On one hand, the “unbounded” nature of complex interference models has hindered research on protocol design and network analysis. For instance, the computational complexity of scheduling under the physical interference model [5, 14, 15] and the derivation of scheduling algorithms with proven approximation bounds [1, 5] have been addressed only recently. At the same time, the unbounded range of complex interference models has discouraged their adoption in wireless network simulators. In fact, accounting for all possible interferers in the network when determining whether a single transmission is successful poses serious challenges in terms of scalability of the simulator, i.e., its ability to run simulations on medium to large scale networks in reasonable time. For this reason, commonly used wireless network simulators such as ns2 [16], GTNetS [18], and GloMoSim/Qualnet [21] use “bounded” interference models, in which low-power interfering signals (typically coming from far-away transmitters) are ignored. Intuitively, this is equivalent to confining interference generated by a transmitter on unintended receivers to a bounded region, which is in general smaller than the network deployment region.

Several studies have investigated the effects of different physical layer implementations on wireless network simulation accuracy [2, 8, 20]. In particular, in [20] Takai, et al. carried out a detailed analysis and comparison of the physical layer implementations of three commonly used wireless network simulators (ns2, GloMoSim, and OpNet). Among other things, they analyzed the relative impact of different physical layer sub-models (interference modeling, signal reception model, fading, path loss, etc.) on simulation results, and concluded that interference modeling is the sub-model with the strongest impact on simulation results.

Other research has focused on improving accuracy of ns2, the inaccuracies of which are well-known [12]. In particular, considerable efforts have been made in the vehicular networking community to extend the ns2 design to incorporate an accurate physical interference model, which is however tailored to vehicular scenarios [3]. This work considers an “unbounded” interference model, however it does not specifically evaluate the effects of choosing this model. Effects of interference modeling on wireless network simulation have been investigated also in [7, 9], which, however, are focused on ad hoc simulators developed by the authors. Furthermore, [9] considers only “bounded” interference models, while in [7] simulations are limited to the 802.15.4 protocol and do not model the MAC level, hence they do not consider important MAC-level issues that may have a significant impact on wireless network performance.

To the best of our knowledge, none of the previous work has investigated the issue of using “bounded” vs. “unbounded” interference models in widely used and existing simulators for wireless networks, and the effects on simulation accuracy and running time, which are the main focus of this paper. Furthermore, ours is the first paper that analytically shows that ignoring interference beyond a certain, fixed, interference range might result in ignoring a term as much as orders of magnitude larger than the noise, and that the

only way of obtaining accurate simulation results in general scenarios is to consider an interference range that grows in an unbounded fashion with the number of network nodes. We emphasize that we are concerned with interference modeling only, and not with other physical layer aspects such as packet reception model (SINR-based vs. BER-based), packet preamble capturing model, and so on. Our choice is motivated by the fact that previous work [3, 20] consistently found that the interference model is the PHY layer sub-model which has the greatest impact on simulation results.

3. A BOUNDED PHYSICAL INTERFERENCE MODEL

We now formally define a bounded physical interference model and compare it with the ‘true’ physical interference model. In the true physical interference model [6], successful reception of a packet sent by node u and destined to node v depends on the SINR at v . Denoting by $P_v(x)$ the received power at v of the signal transmitted by node x , a packet along link (u, v) is correctly received if and only if:

$$\frac{P_v(u)}{N + \sum_{w \in V' - \{u\}} P_v(w)} \geq \beta, \quad (1)$$

where N is the background noise, V' is the subset of nodes in V that are transmitting simultaneously, and β is a constant threshold (called the *SINR threshold* or *packet capture threshold*) that depends on the desired data rate, the modulation scheme, etc. In the true physical interference model, every concurrent transmission in the network must be considered when evaluating whether any single transmission is successful.

In the bounded physical interference model considered herein, henceforth referred to as the *BPI model*, we consider only concurrent transmissions within a given region enclosing the receiver of a particular transmission. We refer to the region considered by the BPI model for a particular receiver as its *interference region*. In the BPI model, concurrent transmissions outside of a receiver’s interference region are ignored. More formally, the BPI model is defined as follows:

DEFINITION 1. *A packet sent along link (u, v) is correctly received in the BPI model with interference region IR_v if and only if*

$$\frac{P_v(u)}{N + \sum_{w \in V' \cap \text{IR}_v - \{u\}} P_v(w)} \geq \beta,$$

We refer to the total interference from outside the interference region as *far-away interference*.

As is common in the literature, we consider circular interference regions, which can be characterized by a single parameter, referred to as the *interference range*. The BPI model is parameterized by the interference range, which we consider to be a function of the number of nodes in the network for the purposes of asymptotic analyses.

Note that both the BPI and the true physical interference models assume that interference from different transmitters on a given receiver is additive. Recent work has shown that, in some environments, interference is indeed additive [13].

4. ANALYSIS

4.1 Preliminaries

We are given a communication graph $G = (V, E)$, where V is the set of nodes in the multihop wireless network and E is the set of intended communication links, i.e., we assume that all traffic in the network will use links in the given set E even if other potentially viable links are present. A *feasible transmission set* T is a set of transmitter, receiver pairs such that: *i*) every node appears at most once as a transmitter or receiver in T ; *ii*) $\forall (u, v) \in T, (u, v) \in E$; and *iii*) $\forall (u, v) \in T$, the SINR inequality (for whichever physical interference model is under consideration) is satisfied at v when all transmitters in T are transmitting simultaneously and no other nodes are transmitting.

In the analysis, we assume the following concerning radio signal propagation: *a1*. Radio signal propagation obeys the log-distance path loss model, with path loss exponent $\alpha > 2$; and, *a2*. All nodes use the same transmit power P , which is such that the resulting transmission range r (in absence of interference) is at least 1. The last part of Assumption *a2* amounts to normalizing the unit of distance in the network to be equal to the transmission range.

The results presented in the following hold under the assumption of constant density of transmitters per unit of area, which reflects situations in which a typical CSMA MAC protocol (e.g., 802.11) is used to gain access to the channel, or in which transmissions are carefully scheduled with TDMA-like protocols. Our results can be extended also to the case of uniform random placement of transmitter nodes (these results are stated without a proof for lack of space).

4.2 Constant transmitter density

We start with proving the following upper bound on the interference generated by transmitters outside an interference range $s \geq 2r$ centered at the intended receiver of a communication.

THEOREM 1. *Assume a constant density of transmitter nodes, and let u be an arbitrary node in the network which is at the receiver end of a communication link; the interference generated by nodes located at distance $d > s$ from u , where $s \geq 2r$ is the interference range, is upper bounded by*

$$C(\alpha) = \frac{\pi \rho P}{s^{\alpha-2}} \cdot \frac{1}{4 - 5 \left(\frac{4}{5}\right)^{\frac{\alpha}{2}}},$$

where ρ is an upper bound to the transmitter density per unit area.

PROOF. Consider an infinite hierarchy $\mathcal{C} = \{\mathcal{C}_0, \mathcal{C}_1, \dots\}$ of circles centered at u , where the radius of \mathcal{C}_0 is s , and the radius of \mathcal{C}_k is sz^k , for some constant $z > 1$. It is easy to see that hierarchy \mathcal{C} covers all the plane, i.e., all the network nodes. Let us now bound the contribution to interference from transmitters located in $\Delta_{k+1} = \mathcal{C}_{k+1} - \mathcal{C}_k$. The area of Δ_{k+1} is $\pi s^2 z^{2k} (z^2 - 1)$, and the number of transmitters in Δ_{k+1} is at most $\rho \pi s^2 z^{2k} (z^2 - 1)$. The contribution to the interference level at node u due to each one of these transmitters (call it w) is upper bounded by $\frac{P}{s^\alpha z^{k\alpha}}$, since we are assuming a decay of power inversely proportional to $d(u, w)^\alpha$, and $d(u, w) > sz^k$. Hence, the total contribution

to interference due to transmitters in Δ_{k+1} is upper bounded by $\frac{P}{s^\alpha z^{k\alpha}} \cdot \rho \pi s^2 z^{2k} (z^2 - 1) = \frac{\pi \rho P}{s^{\alpha-2}} \cdot \frac{z^2 - 1}{z^{k(\alpha-2)}}$. The upper bound to the total interference level at u due to transmitters at distance larger than s from u can be obtained by evaluating

$$\sum_{k=0}^{+\infty} \frac{\pi \rho P}{s^{\alpha-2}} \cdot \frac{z^2 - 1}{z^{k(\alpha-2)}},$$

which converges to $\frac{\pi \rho P}{s^{\alpha-2}} \cdot \frac{z^\alpha (z^2 - 1)}{z^{\alpha - z^2}}$.

The bound above depends on the parameter z chosen to build the hierarchy of circles. If we set z in such a way that the area of Δ_1 (and, consequently, those of Δ_k , with $k > 1$) is at least equal to the area of a circle of radius $r = 1$ (this is to avoid artificially small areas, in which the constant transmitter density assumption might not be valid), by observing that $s \geq 2r$, we get $z = \frac{\sqrt{5}}{2}$. The upper bound on interference then becomes $C(\alpha) = \frac{\pi \rho P}{s^{\alpha-2}} \cdot \frac{1}{4-5(\frac{4}{5})^{\frac{\alpha}{2}}}$. \square

COROLLARY 1. *If the interference range s is chosen in such a way that $s = f(n)$, where $f(n)$ is an arbitrary unbounded increasing function of n , then the total interference at an arbitrary receiver node u due to nodes located at distance greater than s from u converges to 0 as $n \rightarrow \infty$.*

The next lemma shows that if the interference range s is chosen to be a constant, it is not asymptotically safe to ignore interference beyond distance s from the receiver, which, combined with Corollary 1, implies that $s = f(n)$ is a necessary and sufficient condition on the interference range for asymptotic accuracy of the BPI model.

THEOREM 2. *Assume s is set to an arbitrary constant $h > 1$. Then, there exists a node deployment and a transmission set T for that deployment such that: i) T is a feasible set under the BPI model with parameter s ; and ii) the total far-away interference at some receiver u converges to $C'(\alpha) > 0$ as $n \rightarrow \infty$, where*

$$C'(\alpha) = \frac{\pi P h^{2-\alpha}}{([\!|h|] + 2)^2} \cdot \frac{1}{2^{\frac{\alpha}{2}} - 2}.$$

PROOF. Assume nodes are placed on a square grid of side \sqrt{n} , and adjacent nodes in the grid are placed exactly at distance $r = 1$ (the transmission range). The interference range s is set to an arbitrary constant $h > 1$. Consider the following transmission set T : transmitters are located $[\!|h|] + 2$ hops away from each other in the grid along the horizontal and vertical direction (see Figure 1); for each transmitter node u in T , one of the neighbors of u in the grid is selected as the intended receiver. It is easy to see that transmission set T is feasible under the BPI model with interference range h : in fact, for every intended receiver node v in T , the closest interferer is at distance $> h$ from v ; hence, all interferers are ignored when computing the SINR at every receiver node in T , and the transmission set is feasible.

Let us now lower bound the interference at a certain receiver node v due to nodes outside the interference range under the




Figure 1: Transmission set used in the proof of Theorem 2. The figure refers to the case $s = 2$.

true physical model. By using a hierarchy similar to the one used in the proof of Theorem 1, by observing that the density of transmitter nodes per unit of area in transmission set T is $\frac{1}{([\!|h|] + 2)^2}$, and that each of these transmitters generates an interference level at least $\frac{P}{(h z^{k+1})^\alpha}$ when considering area Δ_{k+1} , we can lower bound the total interference at the node v in the center of the grid as follows:

$$\frac{\pi P h^{2-\alpha}}{([\!|h|] + 2)^2} \cdot \sum_{k=0}^{\bar{k}} \frac{z^2 - 1}{z^{k(\alpha-2)} z^\alpha},$$

where $\bar{k} = \log_z \left(\frac{\sqrt{n}}{2h} \right)$. For $n \rightarrow \infty$, the above lower bound converges to $\frac{\pi P h^{2-\alpha}}{([\!|h|] + 2)^2} \cdot \frac{z^2 - 1}{z^{\alpha - z^2}}$. The lemma follows by observing that, by setting $z = \sqrt{2}$, we guarantee that the area of Δ_1 is at least as large as the area of a circle of radius h , and the lower bound on transmitter density is satisfied. \square

It is easy to see that, for realistic parameter values, $C'(\alpha)$ is very high. For instance, by setting $P = 100mW$, $\alpha = 3$, $h = 2r$ (corresponding to the standard assumption that the interference range is twice the transmission range), and $r = 250m$, we get $C'(\alpha) = -55.2148dBm$, which is at least two orders of magnitude higher than the standard noise value. The above value is the limit of the total interference, which could actually be approached only for very large networks. But even if we consider relatively small networks and truncate the summation in the proof of Theorem 2 accordingly, we obtain a relatively high value of the total interference (e.g., $C'(\alpha) = -55.2166dBm$ for $n = 49$).

4.3 Random uniform scenario

In the random uniform scenario, it is assumed that a number $n = (8 + \varepsilon) C \ln C$ of (potentially transmitter) nodes is deployed uniformly at random in a square area R of side $l = \sqrt{C}$, where ε is an arbitrary positive constant. All the results in this section are stated without a proof due to lack of space.

THEOREM 3. *Assume the random uniform scenario, and let u be an arbitrary node in the network which is at the receiver end of a communication link; the interference generated by nodes located at distance $d > s$ from u , where $s \geq 2r$*

is the interference range, is upper bounded by

$$C_{rand}(\alpha) = \frac{6f(C)P}{s^{\alpha-2}} \cdot \frac{2^{\frac{\alpha}{2}}}{2^{\frac{\alpha}{2}} - 2},$$

w.h.p., where $f(C) \approx e \ln C$ is a properly defined function.

COROLLARY 2. *If the interference range s is chosen in such a way that $s = f(n)$, where $f(n)$ is an arbitrary function of n such that $\frac{\log n}{f(n)} \rightarrow 0$ as $n \rightarrow \infty$, then the total interference at an arbitrary node u due to nodes located at distance greater than s from u converges to 0 as $n \rightarrow \infty$, w.h.p.*

Observe that, compared to the constant transmitter density scenario, we have a slightly stronger requirement on the interference range to have convergence to 0 of the total interference generated from nodes outside the interference range. This additional constraint is due to the higher transmitter node density in the random uniform scenario ($\Theta(\log C)$ transmitter nodes per unit of area, instead of $O(1)$ as in the constant density scenario).

THEOREM 4. *Assume the random uniform scenario, and assume s is set to an arbitrary constant $h > 1$. Then, there exists a transmission set T for such that: i) T is a feasible set under the BPI model with parameter s ; and ii) the total interference at some receiver u due to nodes located at distance greater than s from u converges to $C'_{rand}(\alpha) > 0$ as $n \rightarrow \infty$ w.h.p., where*

$$C'_{rand}(\alpha) = \frac{4Ph^{2-\alpha}}{([h] + 1)^2(2^\alpha - 4)}.$$

Similarly to the constant density scenario, by using realistic parameters ($P = 100mW$, $\alpha = 3$, and $h = 2r$), we have $C'_{rand}(\alpha) \approx -61dBm$, which is at least two orders of magnitude larger than the typical noise value.

5. SIMULATIONS

5.1 Simulation setup

To evaluate different interference models, we use the packet-level simulator GTNetS [18]. In the following, we describe the basic GTNetS interference model and the modifications that we carried out to implement the true physical interference model.

5.1.1 Basic GTNetS interference model

The GTNetS simulator uses an approximate, bounded interference model where interference is considered only within a limited interference range (IR). Another significant approximation of the interference model used in GTNetS is that, similarly to ns2, it takes into account only one interferer at a time (without accumulation) among the nodes that are inside the IR, instead of accumulating all the interferences of concurrent transmissions. Note that this pairwise interference calculation is equivalent to considering only the *maximum* interferer among the nodes transmitting inside the IR.

Let us add some details about the basic GTNetS implementation of the interference model. A wireless link is associated to each node; the link can be in one of the following states:

- IDLE: the node is not involved in any transmission or reception;
- TX: the node is transmitting data;
- RX: the node is receiving valid data;
- CX: the node is receiving data that are not valid due to a collision occurred during the reception.

The physical carrier sense mechanism of the 802.11 MAC protocol implemented in GTNetS is based on the notion of wireless link state: the channel is considered idle if and only if the wireless link is in the IDLE state, otherwise is busy.

Whenever a node x has to transmit a packet, it checks the state of its wireless link; if the link state is IDLE, the node starts the transmission phase. All nodes outside the IR of x ignore the interference due to the transmission. On the other hand, for each node u within the IR, the state of the wireless link might change depending on the distance d between u and x .

If u is outside the carrier sensing range of node x , i.e. the range up to which a busy channel can be detected, no action is taken. On the other hand, for each node u inside the carrier sensing range, the simulator computes $P_u(x)$, that is the received power at u of the signal transmitted by x , and the instant $t_u(x)$ when the reception of the signal transmitted by x will end ($t_u(x)$ is calculated based on distance d , radio signal propagation time, and packet length). A reception event is scheduled at time $t_u(x)$ for the node u . If u is outside the *transmission* range of node x , an error flag will be set to indicate that the packet can not be correctly received.

The subsequent actions are determined by the current link state of node u . If the link is IDLE, the link state is changed to RX, and at the time $t_u(x)$ the link will be set to IDLE again. If the link is in RX or CX state, the node u is currently receiving a packet from another transmitting node y . In this case, the ratio of the received powers at u ($P_u(y)/P_u(x)$) is compared with the *SINR threshold* β to determine whether the transmission from x may invalidate the ongoing transmission from y . If $P_u(y)/P_u(x) < \beta$, the link state of node u is changed to CX to indicate that a collision occurred and the packet reception is not valid. The link state of node u will be set to IDLE again at the end of the transmission that lasts longer between the ongoing transmission from y and the transmission from x .

When a node handles a reception event, it considers the received packet as valid if the packet error flag is false and the wireless link state is not CX.

5.1.2 Modifications to GTNetS

We modified the GTNetS code to implement the true physical interference model. In contrast with the basic GTNetS design, we do not rely on the notion of wireless link state to detect channel conditions and correctness of the data transmission.

For each node u of the network, we implement a data structure $Interf(u)$ that keeps track of the cumulative interference generated by *all* simultaneously transmitting nodes

within the IR of u . When a node x starts a transmission, for each node u inside the IR of x three values are computed: $P_u(x)$, that is the received power at u of the signal transmitted by x ; $t1_u(x)$ and $t2_u(x)$, that are the times indicating the beginning and the end of the reception at u of the signal transmitted by x , respectively. $P_u(x)$ is recorded in the data structure $Interf(u)$ along with the times $t1_u(x)$ and $t2_u(x)$. Then, a reception event is scheduled for the node u at time $t2_u(x)$; in the reception event the initial time of the reception $t1_u(x)$ is also recorded. Finally, the error flag of the transmitted packet is set to true if u is outside the transmission range of x .

When a node u handles a reception event, it determines the validity of the received packet based on both the value of the error flag and the information about the cumulative interference recorded in $Interf(u)$. The interferences caused by simultaneous transmissions that occur during the reception of the packet are added up using $t1$ and $t2$ times to calculate the cumulative interference $cumInterf(u)$ at the receiver u , that is considered for SINR computation. A packet transmitted by node x is correctly received at node u if during the entire packet reception we have: $P_u(x)/(N + cumInterf(u)) \geq \beta$, where N is the background noise and β is the SINR threshold.

The information about the cumulative interference is used also for the physical carrier sensing phase at the transmitter. When a node x has data to transmit, it compares the carrier sensing threshold with the cumulative interference $cumInterf(x)$ due to all the ongoing transmissions from nodes inside the IR of x . A value of $cumInterf(x)$ lower than the threshold means that the channel is idle, and x starts the transmission phase.

It is worth noting that in our extended version of GTNetS each node u inside the IR of a node x will be affected by the interference caused by the transmission of x even if the distance d between u and x exceeds the carrier sensing range. On the other hand, in the basic version of GTNetS all the nodes outside the carrier sensing range of a transmitting node will ignore the interference, even if they are inside the IR.

Summarizing, the main features of the extended GTNetS version, which will be made available for download, are: *i*) accumulated interference calculation within the interference range, and *ii*) tunable interference range, which allows us to mimic the true physical interference model.

5.1.3 Simulation methodology

We carried out extensive simulations to evaluate the impact of different interference models on simulations referring to the same network scenario. We test the basic GTNetS interference model, denoted by *IR500-MI*, which considers only the maximum interferer within the interference range, that is set to the default value of $500m$ (twice the 802.11 nominal transmission range of $250m$). We compared the basic GTNetS version with our implementation of the physical interference model, which takes into account the cumulative interference from all transmitters within the IR. We consider a bounded model with IR set to $500m$, denoted by *IR500-CI* (which resembles GloMoSim interference model),

and the true physical interference model, denoted by *PhyI*, with unlimited IR (i.e., an interference range larger than the diameter of the network deployment region). In general, we use notation IRXX-MI to denote the interference model with IR equal to XXm , where only the maximum interferer is accounted for; notation IRXX-CI to denote the interference model with IR equal to XXm and cumulative interference within the interference range, and PhyI for the true physical interference model.

For our experiments, we consider two network topologies:

- *grid*: n nodes are located in a square, regularly-spaced grid; the distance between two adjacent nodes is $200m$;
- *random uniform*: n nodes are distributed uniformly at random in a square area of side $s = \sqrt{n} \times 100m$.

We evaluate the interference models in presence of two types of data traffic: CBR and TCP. In both cases, we use \sqrt{n} flows to generate traffic, with pairs of sources and destinations randomly chosen among the n nodes in the network. We ensure that a node can not be a source or a destination for more than one flow.

We assume nodes are equipped with 802.11b radios, and radio signal propagation obeys the log-distance path loss model. For routing packets between far-away nodes we use the DSR routing algorithm [10]. The parameters used in our simulations are shown in Table 1. The simulated time interval is 600 seconds: in a preliminary set of simulations, we have verified that this value of simulated time is high enough to ensure that steady-state conditions are reached in the considered network scenarios. For each simulation run, we began data collection only after 30 seconds to avoid transient effects due to initially empty routing tables. Results refer to data averaged over 15 runs. Simulations were run on Dual Intel Xeon 2Ghz machines with 2Gb of RAM.

Table 1: Simulation parameters

| | |
|----------------------------------|--------|
| Link data rate | 11Mbps |
| Transmission power | 100mW |
| Transmission range | 250m |
| Carrier sensing range | 500m |
| Carrier sensing threshold | -48dBm |
| Path loss exponent | 2.5 |
| Background noise | -90dBm |
| SINR threshold | 10dB |

We evaluate the interference models in terms of:

- *Throughput*: the total amount of data successfully delivered to destinations in the simulated time interval;
- *Packet delivery time*: the elapsed (simulated) time between the packet generation at the source and the packet arrival at the destination (only for CBR traffic);
- *Simulation running time*: the elapsed real time between the start and the termination of the simulation run.

We define *error* to be the percentage difference between one of the performance measures (either throughput or delivery time) under an approximate interference model and the same measure under the PhyI model.

5.2 Results for CBR traffic

In this set of simulations, we have randomly created \sqrt{n} CBR flows in a network with n nodes, and evaluated performance in terms of throughput and packet delivery time. The results for IR500-MI, IR500-CI, and PhyI for the grid topology with 100 nodes and varying rates of the CBR flows are reported in Figure 2. The figure reports also the average running time of a single simulation run (note the log-scale used in the y axis).

Results reported in Figure 2 clearly show that, while not accumulating interference within the shorter interference range has only a marginal effect on throughput and delivery time estimation, ignoring interference generated by far-away transmitters has a dramatic effect on performance estimation. More specifically, while the network appears to operate below saturation under the IR500-MI/CI models even for the highest data rate of $100Kbps$, it has already reached saturation at $80Kbps$ under the PhyI model. At a data rate of $100Kbps$, the IR500-MI model estimates a throughput above $99Kbps$, as compared to a throughput slightly above $82Kbps$ with the PhyI model, for an error of about 20%. The situation is even worse for the average delivery time, which is consistently underestimated by both the IR500-MI and the IR500-CI models. The relative error is as high as 60% at $80Kbps$. Note that the higher error in estimating delivery time as compared to throughput is due to the fact that more accurate interference modeling tends to increase not only the number of dropped packets due to low SINR values (which affects both throughput and delivery time), but also the average channel access time. In fact, a higher interference level in the network (which results from using a larger interference range) tends to result in a higher frequency of busy channel detection at the transmitter nodes, resulting in an increase of the average channel access time.

Unfortunately, the price to pay for accurately modeling wireless interference is a dramatic increase in the simulation running time, which increases by more than two orders of magnitude when using the PhyI model instead of the IR500-MI model. This increase in running time is due to the fact that, as the interference range increases, a larger number of events (one for every receiver in the transmitter's interference range) must be scheduled for a single transmission.

The results for the random topology scenario, which are not reported due to lack of space, confirm the trends observed in the grid topology, with an error of about 22% on the average throughput and of 65% on the delivery time when using the IR500-MI model instead of the PhyI model, and two order of magnitude increase in running time with the PhyI model.

Results reported in Figure 3 refer to a set of experiments with the grid topology where the CBR rate is set to $100Kbps$ and the number of network nodes varies from 36 to 144. As in the previous case, the IR500-MI model tends to overestimate throughput with respect to the more accurate PhyI model. In particular, IR500-MI estimates throughput degradation (to around $90Kbps$) only for the largest network size,

while with PhyI a similar degradation in network throughput can already be observed with 81 nodes. The error of IR500-MI is as high as 32% for $n = 144$. As in the previous case, relatively higher inaccuracy of the IR500-MI/CI models is observed in delivery time estimation, which is consistently underestimated (error up to 63% for $n = 144$). It is also worth mentioning that the three interference models predict comparable throughput and delivery time estimations only for values of $n \leq 64$. Hence, neglecting interference from far-away interferers might lead to very inaccurate results even for medium-size networks of 80 or more nodes. Observe that the fact that the error in throughput and delivery time estimation tends to increase with n for a fixed interference range validates the theoretical analysis reported in Section 4. Finally, concerning simulation running time, also in this case we observe an up to two orders of magnitude increase when using the PhyI model instead of the IR500-MI model.

In the case of random topology scenario, the IR500-MI model achieves an error in the performance estimations that is similar to the case of the grid topology, even if the observed throughput is generally lower for all the interference models. In particular, the IR500-MI estimates throughput degradation starting from a network size of $n = 81$, while a size of $n = 64$ is sufficient to have throughput degradation under the PhyI model. Also the simulation running time is slightly lower with respect to the grid scenario, due to the lower number of packets exchanged in the network. However, the difference of almost two orders of magnitude between the simulation running time under the PhyI and the IR500-MI/CI models is confirmed.

In Figure 4, we report the results of a set of experiments for the grid topology with 100 nodes and CBR rate of $100Kbps$, where the interference range varies from $500m$ to $3000m$. Note that, given that the diameter of the deployment region with these parameters is about $2550m$, interference model IR3000-CI is equivalent to the PhyI model. The plots report two curves, which represent the cases of maximal (MI) and cumulative (CI) interference. First, we observe that increasing the interference range in combination with MI interference modeling has almost no effect on throughput estimation and produces only a modest increase in delivery time estimation. This is due to the fact that the strongest interferer of a receiver is likely to be within the shorter $500m$ interference range, in which case models IR500-MI and IRXXXX-MI are equivalent. It is also worth observing that a good estimation accuracy is obtained only for relatively large values of the interference range ($\geq 2000m$), resulting in approximately the same running time as with the PhyI model. The trends observed in Figure 4 are valid also for the simulations under the random topology, that we do not report.

5.3 Results for TCP traffic

In the second set of experiments, we consider \sqrt{n} TCP flows randomly created in a the grid topology with n nodes. Here, we focused on evaluating throughput and average simulation running time.

Figure 5 reports the simulation results for networks of varying size. As with CBR traffic, the IR500-MI model consistently tends to overestimate average throughput. Errors are even higher than for CBR (up to 210% for $n = 100$).




Figure 2: Throughput (left), delivery time (center), and simulation running time (right) for CBR traffic with $n = 100$ and varying packet generation rate in the grid topology.




Figure 3: Throughput (left), delivery time (center), and simulation running time (right) for CBR traffic with $100Kbps$ packet generation rate and with varying network size in the grid topology.

Note that the average throughput of a TCP flow has a non-monotonic behavior with increasing network size. This behavior, which is consistently observed with all the interference models, can be explained as follows. With $n = 36$, we have relatively shorter TCP flows, and there was at least one one-hop flow in every simulated instance. Since one-hop TCP flows are not subject to intra-flow interference, they typically experience a much higher throughput than multihop flows, which explains the relatively high average throughput observed for $n = 36$. For larger values of n , no one-hop TCP flows occurred in the simulated instances, which explains the drop in throughput when passing from $n = 36$ to $n = 49$. After that value of n , the average throughput has an increasing trend, due to the fact that the number of TCP flows grows sub-linearly with n (we have \sqrt{n} flows with n nodes), implying a lower network congestion level observed by each TCP flow (on the average) as n increases.

Compared to the CBR case, the increase in simulation running time using the Phyl model instead of the IR500-MI model is less dramatic: for the largest network size, the relative increase is about *one* order of magnitude. This relatively less dramatic increase is also due to the congestion control TCP mechanism, which tends to reduce the number of transmitted packets (and, consequently, running time) as more congested network conditions are encountered.¹ Combining the fact that relatively less TCP packets are sent with the processing time of each packet (which is much longer in the Phyl model compared to the IR500-MI model), we have

¹This is in sharp contrast with the case of CBR traffic, where the number of generated packets *does not* depend on the network congestion level.

an overall increase of about one order of magnitude in simulation running time.

The statistics about the individual TCP flows are even more interesting than the aggregate data of \sqrt{n} flows reported in Figure 5. Figure 6 reports the average *absolute* error (in *KB*) between the amount of data sent by each TCP flow as estimated by the approximate interference models and as estimated by the Phyl model. The average absolute error (a.a.e.) is computed as follows:

$$\frac{\sum_{i=1, \dots, \sqrt{n}} |Thr_{x-Y}(i) - Thr(i)|}{\sqrt{n}},$$

where $Thr_{x-Y}(i)$ is the total amount of data sent by the i -th TCP flow as estimated by model RRx-Y, and $Thr(i)$ is the same statistic for the Phyl model.

Figure 6 reports the a.a.e. as defined above for three interference models: IR500-MI, IR500-CI and IR1500-CI. As expected, the a.a.e. tends to increase for growing values of n , although a sharp increase can be observed only for n larger than 80. As for the relative behavior of the considered interference models, we observe that IR500-CI reduces a.a.e. up to 24% w.r.t. IR500-MI, while IR1500-CI can reduce a.a.e. by as much as 65% w.r.t. IR500-MI. However, we recall that simulations with the IR1500-CI model have a running time comparable to that under the accurate Phyl interference model.

The choice of computing the absolute instead of percentage error is due to the fact that, for some simulated TCP flows, the percentage error can be extremely large. This occurs




Figure 4: Throughput (left), delivery time (center), and simulation running time (right) for CBR traffic with $n = 100$, $100Kbps$ packet generation rate, and varying size of the interference range in the grid topology.

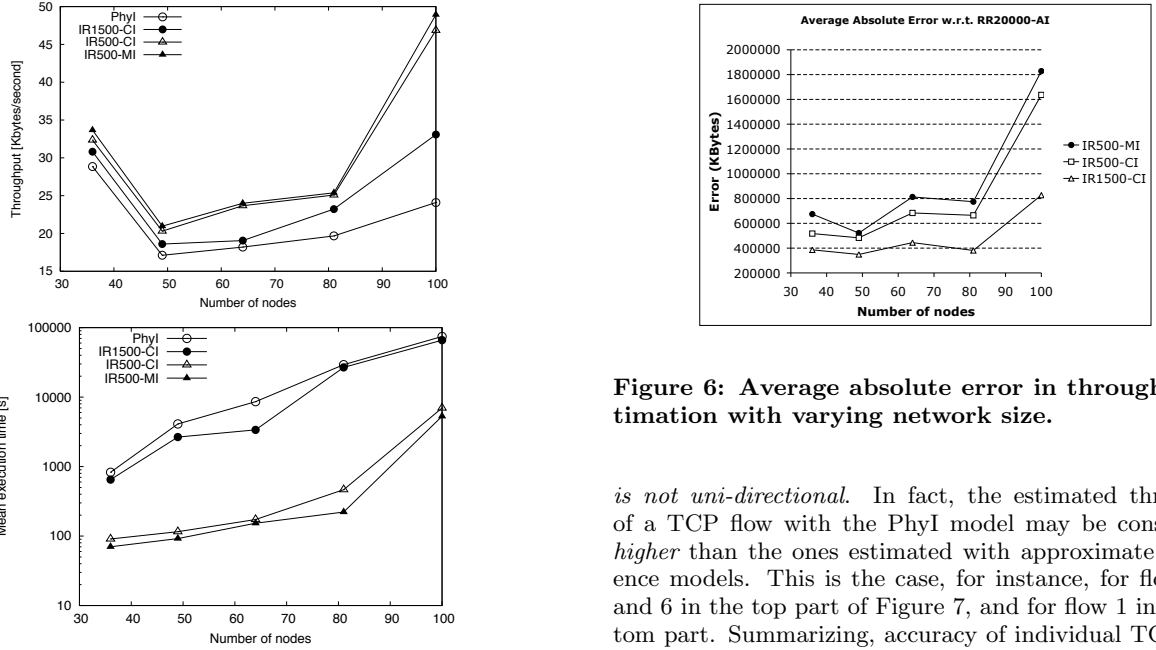


Figure 5: Average throughput (top), and simulation running time (bottom) for TCP traffic with varying network size.

Figure 6: Average absolute error in throughput estimation with varying network size.

when a certain TCP flow is estimated to have a certain, non-negligible throughput under model IRx-Y, while under the PhyI model it is starved. As shown by the disaggregated flow data in Figure 7, this situation occurs frequently. More specifically, Figure 7 reports the disaggregated statistics of the amount of data sent by individual TCP flows for two specific simulation instances, referring respectively to the simulation runs with minimal a.a.e. (top) and maximal a.a.e. (bottom) for the IR500-MI model. It is interesting to observe that, even in the relatively “more accurate” simulation run, there exists a flow (the 4th bar from the left in each interference model) that is consistently estimated as carrying a non-negligible amount of data by models IR500-MI/CI and IR1500-CI, while it was starved under the PhyI model. Thus, the percentage error for this specific flow is enormous for all the approximate interference models. Another interesting observation is that *error in estimating throughput of individual TCP flows with approximate interference models*

is not uni-directional. In fact, the estimated throughput of a TCP flow with the PhyI model may be considerably higher than the ones estimated with approximate interference models. This is the case, for instance, for flows 1, 5, and 6 in the top part of Figure 7, and for flow 1 in the bottom part. Summarizing, accuracy of individual TCP flow’s statistics with approximate interference models is very low due to: (i) incorrect prediction of starved flows, leading to virtually unlimited percentage error for such flows, and (ii) the bi-directional nature of the errors. Note that (i) and (ii) occur even in relatively “accurate” simulation runs, i.e., those corresponding to the lowest a.a.e. among the simulated scenarios.

6. DISCUSSION AND CONCLUSIONS

The main contribution of this paper is an in-depth investigation of the accuracy of approximate, bounded interference models in estimating relevant performance metrics of wireless multihop networks, and of the tradeoff between simulation accuracy and running time. Summarizing, our results suggest that ignoring far-away interference leads to highly inaccurate simulation results, but the price to pay for this accuracy is a one-two orders of magnitude increase in simulation running time.

Overall, the issue of how to address the fundamental tradeoff between simulation accuracy and running time is the main problem left open by this paper. We believe optimally addressing this tradeoff is a very difficult task, since it depends on the specific network setting at hand, the type of




Figure 7: Disaggregated TCP throughput for a scenario with minimal (top) and maximal (bottom) average absolute error.

transport-layer traffic, and so on. Still to be addressed is evaluating the impact on the simulation accuracy/running time tradeoff of other important parameters such as radio signal propagation models, non-additive interference, network topology, node mobility, and so on. A promising research direction is deriving analytical/stastical methods to accurately estimate aggregate interference generated by far-away interferers, so that relatively accurate estimates of the SINR values at single receivers can be obtained without actually computing each far-away interference signal level. An interesting first step in this direction is [17], where the authors introduce a Markov-chain based interference model which is fed by testbed measurement. While the idea of measurement driven interference model is interesting, the approach of [17] considers only 1-hop flows, and is based on Markov chains with a number of states exponential in the number of nodes (and is thus not scalable). Hence, further work is needed to derive scalable approaches that account also for multi-hop flows.

7. REFERENCES

- [1] G. Brar, D. Blough, and P. Santi, "Computationally Efficient Scheduling with the Physical Interference Model for Throughput Improvement in Wireless Mesh Networks," *Proc. ACM Mobicom*, pp. 2–13, 2006.
- [2] D. Cavin, Y. Sasson, and A. Schiper, "On the accuracy of MANET simulators", *Proc. ACM POMC*, pp. 38–43, 2002.
- [3] Q. Chen et al., "Overhaul of IEEE 802.11 Modeling and Simulation in NS-2", *Proc. ACM MSWiM*, pp. 159–168, 2007.
- [4] S. Das, D. Koutsonikolas, Y. Hu, D. Peroulis, "Characterizing Multi-Way Interference in Wireless Mesh Networks", *Proc. ACM Wintech*, 2005.
- [5] O. Goussevskaya, Y.V. Oswald, R. Wattenhofer, "Complexity in Geometric SINR", *Proc. ACM MobiHoc*, pp. 100–109, 2007.
- [6] P. Gupta and P.R. Kumar, "The Capacity of Wireless Networks," *IEEE Transactions on Information Theory*, Vol. 46, No. 2, pp. 388–404, 2000.
- [7] E. Hamida, G. Chelius, and J.-M. Gorce, "Scalable versus Accurate Physical Layer Modeling in Wireless Network Simulations", *Proc. PADS*, pp. 127–134, 2008.
- [8] J. Heidemann, N. Bulusu, J. Elson, C. Intanagonwiwat, K. Lan, Y. Xu, W. Ye, D. Estrin, and R. Govindan, "Effects of Detail in Wireless Network Simulation", *Proc. SCS Multiconference on Distributed Simulation*, pp. 3–11, 2001
- [9] A. Iyer, C. Rosenberg, A. Karnik, "What is the Right Model for Wireless Channel Interference?", *Proc. ACM QShine*, 2006.
- [10] D.B. Johnson, D.A. Maltz, "Dynamic Source Routing in Ad Hoc Wireless Networks", *Mobile Computing*, n. 353, pp. 153–181, 1996.
- [11] A. Keshavarz-Haddad, R. Riedi, "On the Broadcast Capacity of Multihop Wireless Networks: Interplay of Power, Density and Interference", *Proc. IEEE SECON*, pp. 314–323, 2007.
- [12] D. Kotz et al., "Experimental Evaluation of Wireless Simulation Assumptions", *Proc. ACM MSWiM*, pp. 78–82, 2004.
- [13] R. Maheshwari, S. Jain, S. Das, "A Measurement Study of Interference Modeling and Scheduling in Low-Power Wireless Networks," *Proc. ACM SenSys 2008*, pp. 141–154, 2008.
- [14] T. Moscibroda and R. Wattenhofer, "The Complexity of Connectivity in Wireless Networks", *Proc. IEEE Infocom*, pp. 1–13, 2006.
- [15] T. Moscibroda, R. Wattenhofer, A. Zollinger, "Topology Control Meets SINR: The Scheduling Complexity of Arbitrary Topologies", *Proc. ACM MobiHoc*, pp. 310–321, 2006.
- [16] <http://www.isi.edu/nsnam/ns2/>
- [17] L. Qiu, Y. Zhang, F. Wang, M. Han, and R. Mahajan, "A general model of wireless interference", *Proc. ACM MobiCom*, pp. 171–182, 2007.
- [18] G. Riley, "The Georgia Tech Network Simulator," *ACM SIGCOMM MoMeTools Workshop*, 2003.
- [19] G. Sharma, R.R. Mazumdar, and N.B.Shroff, "On the Complexity of Scheduling in Wireless Networks", *Proc. ACM Mobicom*, 2006.
- [20] M. Takai, J. Martin, R. Bagrodia, "Effects of Wireless Physical Layer Modeling in Mobile Ad Hoc Networks", *Proc. ACM MobiHoc*, pp. 87–94, 2001.
- [21] X. Zeng, R. Bagrodia, and M. Gerla, "GloMoSim: a Library for Parallel Simulation of Large-scale Wireless Networks", *Proc. Workshop on Parallel and Distributed Simulations*, (PADS), 1998.