This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

Coupled 802.11 Flows in Urban Channels: Model and Experimental Evaluation

Joseph Camp
Department of Electrical Engineering
Southern Methodist University, Dallas, TX
camp@lyle.smu.edu

Ehsan Aryafar and Edward Knightly
Department of Electrical and Computer Engineering
Rice University, Houston, TX
{ehsan, knightly}@ece.rice.edu

***Abstract*—Contending flows in multi-hop 802.11 wireless networks compete with two fundamental asymmetries: *(i)* channel asymmetry, in which one flow has a stronger signal, potentially yielding physical layer capture, and *(ii)* topological asymmetry, in which one flow has increased channel state information, potentially yielding an advantage in winning access to the channel. Prior work has considered these asymmetries independently with a highly simplified view of the other. However, in this work, we perform thousands of measurements on coupled flows in urban environments and build a simple, yet accurate model that *jointly* considers information and channel asymmetries. We show that if these two asymmetries are not considered jointly, throughput predictions of even two coupled flows are vastly distorted from reality when traffic characteristics are only slightly altered (e.g., changes to modulation rate, packet size, or access mechanism). These performance modes are sensitive not only to small changes in system properties, but also small-scale link fluctuations that are common in an urban mesh network. We analyze all possible capture relationships for two-flow sub-topologies and show that capture of the *reverse* traffic can allow a previously starving flow to compete fairly. Finally, we show how to extend and apply the model in domains such as modulation rate adaptation and understanding the interaction of control and data traffic.**

I. Introduction

In urban environments, IEEE 802.11 nodes interact in many ways, e.g., within and among paths in a multi-hop network and among deployments from different domains. Competing transmitters rarely have equal link quality to a given receiver, i.e., *channel asymmetries* are prevalent, especially in urban channels. When packets overlap in time, even slight link quality differences have been shown to cause physical layer capture such that the packet sent over the higher quality link is received correctly but the packet sent over the weaker link is dropped [1]. Moreover, transmitters or receivers of competing flows often have unequal channel state information, a situation termed *information asymmetry*. In such cases, a topological asymmetry results in a hidden node having inferior channel state information, forcing the hidden node to contend at random times guided by binary exponential backoff rather than at "idle times" driven by carrier sense. However, while the effects of information asymmetry and channel asymmetries are understood in isolation ([2], [3], [4], [5] and [6], [7], [8], respectively), their interdependencies have thus far been ignored.

In this paper, we *jointly* consider information and channel asymmetries with both analytical models and extensive, urban measurements. As in [9], we employ the two-flow enumeration technique of [4] and consider all topological couplings of paired flows. However, in contrast to [4], [9], we inform the model of channel asymmetries via a signal-strength matrix between nodes and a measurement characterization of physical layer capture events. By doing so, we reveal complex interdependencies of different system parameters that have been ignored by prior work. As an example, we show that not only do traffic parameters such as modulation rate and packet size change the timing of the model (e.g., backoff, carrier sense, and other factors that are most affected by information asymmetries), but they can also change the ability of traffic to perform physical layer capture at certain channel-asymmetry states. As a result, a new dimension emerges for throughput sharing as traffic parameter choices alter both capture relationships and information asymmetries.

In particular, our contributions are two-fold. First, we develop an analytical model that predicts the throughput of coupled flows with the input of signal strengths between nodes. Using an embedded Markov chain to characterize a broad set of link interaction states, we incorporate key system factors such as topology, modulation rate, packet size, channel conditions, and physical layer capture. This model is the first to jointly consider information and channel asymmetries when predicting the throughput of 802.11 flows. By doing so, we characterize our experimental finding that even with high-quality links, small-scale channel fluctuations common to urban environments can flip the throughput-sharing modes of coupled flows. One such example occurs when flows compete asymmetrically due to topological connectivity factors. Namely, as shown in [10] and modeled with idealized channels in [4], a flow can starve due to lack of knowledge at the sender about when to begin contention. The flow with full information "wins" the contention nearly all the time. However, we show that by incorporating channel asymmetries, a favorable channel state at the receiver for the information-poor transmitter can allow that flow to receive equal throughput compared to the information-rich transmitter. Conversely, we show that with flows previously assumed to have equal throughput distributions due to symmetric information, only slight channel asymmetries cause one flow to achieve nearly zero throughput.

Second, we design a set of urban experiments consisting of 1000's of measurements on a deployed urban mesh network. We first validate the analytical model and show that it is accurate in predicting the throughput of coupled flows for diverse channel conditions and topologies. Further, our measurements indicate that the throughput sharing of many coupled flows vary widely over time. By examining the channel conditions associated with the maximum and minimum differences in flow throughputs, we empirically identify the small-scale

978-1-4244-5837-0/10/$26.00 ©2010 IEEE

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

channel fluctuations that cause such changes in throughput-sharing modes. Throughout, we show that *reverse traffic* (acknowledgment and clear-to-send packets traveling in the reverse direction of data) has a critical impact that has not been studied. In contrast to the data direction, this reverse channel is *not* carrier sensed before transmitting. Thus, even within carrier-sense range, capture relationships have a critical impact yielding new link interdependencies, interactions with forward traffic, and vulnerable sub-topologies, all characterized by the model. Finally, we discuss extensions to the model and how to apply it to the two domains of modulation rate adaptation and predicting the effect of control traffic on data flows.

The paper is organized as follows. First, we present our analytical model in Section II. In Section III, we perform an extensive set of experiments on an urban mesh network to both validate and apply our model. We then compare to related work in Section IV. Lastly, we conclude in Section V.

II. Model

In this section, we develop a bi-dimensional discrete time Markov chain embedded over continuous time to study the throughput sharing behavior of two coupled sources. We explicitly account for different capture relationships that exist among competing nodes. The system state is the joint backoff stage of the coupled sources. The transition probability is determined by capture relationships and other system parameters. This allows different capture relationships to drive the steady-state distribution of the system. Using our model, we accurately predict the throughput as well as investigate the impact of capture relationships and other parameters on the system performance.

A. Background

We study the performance of coupled flows in multi-hop wireless networks. Fig. 1 depicts a snapshot in time of such flows with symmetric and asymmetric cross-flow connectivity, terms we will now define.

Coupled and Uncoupled Flows and Hidden Terminals. In most cases, a flow such as Aa that is interfering with flow Bb has backoff behavior that is *coupled* to that of flow Bb. In other cases, such as with broadcast traffic, the backoff processes of competing flows are *uncoupled*. When two transmitters such as A and B use 802.11, if *inter-sender* interference exists in which packets can be sensed or decoded between transmitters, one transmitter defers while the other transmits. The resulting behavior can be predicted using existing models such as [11] and extensions. However, where the two transmitters have no inter-sender interference (hidden terminals), no prior model jointly considers channel and information asymmetries.

Cross-Flow Connectivity. Flows with coupled backoff behavior can have cross-flow connectivity where the sender of one flow receives packets from the receiver of the competing flow. There can be *symmetric or asymmetric cross-flow connectivity* (refer to Fig. 1) if the senders of both flows can decode packets from the receiver of the competing flow or if only one sender is able to do so, respectively. The symmetry or asymmetry of this relationship has been shown to cause balanced or imbalanced throughput sharing with idealized channels due to the MAC layer behavior [4], [10].

Complexity of Capture Relationships. Node mobility or environmental movement can cause fluctuations in channel quality. Further, there are spatial differences between the senders of competing flows from a given receiver. The resulting difference in link qualities can cause physical layer capture for competing transmitters which are out-of-range and thus, can have simultaneous transmissions [1], [7].

Physical layer capture can occur for traffic in the forward direction (e.g., data or RTS packets from A and B overlapping at a in Fig.1) or for traffic in the reverse direction (e.g., CTS or ACK packets from a and b overlapping at A in Fig. 1). For a given flow, there can be *forward traffic capture* over the forward or reverse competing traffic and *reverse traffic capture* over the forward or reverse competing traffic. There are a total of four possible capture scenarios for a given flow with respect to competing flows and three possible results: winning capture, losing capture, or collision with loss.

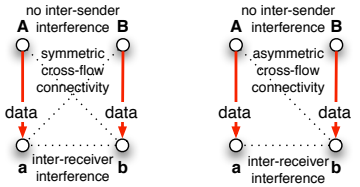

Fig. 1. A snapshot of coupled flows with symmetric and asymmetric cross-flow connectivity.

Timing Impact of Capture. Prior work has shown that not only the relative signal strength, but also the arrival time of the competing packets plays a critical role in physical layer capture properties [7]. Consider two overlapping packets arriving at a receiving node using the Prism 2.5 chipset (used in the TFA Network [12]). The stronger packet can only be captured if it arrives first or trails the weaker packet by less than the synchronization bits (6 mini-slots in 802.11b) of the weaker packet's preamble. In either case, the weaker packet is dropped [1]. Our analysis, models this family of capture behaviors. Recent measurement studies have shown that the later arriving packet can still be captured through the Message-In-Message function in the 802.11a standard and implemented in the Atheros chipset [7]. The same methodology can be extended to analyze such behavior.

B. Analytical Model

Joint Channel State Evolution. In order to correctly analyze the behavior of coupled sources, we consider the joint backoff evolution of the two flows. Fig. 2 shows an abstract representation of the joint channel state evolution where the arrows correspond to time instants in which both senders can potentially start transmitting the first packet of a new data exchange. We identify three main states: *(i)* an idle state, *(ii)* a single-access state, and *(iii)* an overlapping-packets state. In the single-access state either one flow transmits or both flows transmit but the first packet of the earlier flow does not overlap with other flow's later packet (e.g., in the right side of Fig. 1, A's RTS finishes and is receiving a CTS while B starts transmitting an RTS). In the overlapping-packets state, the first packets of both flows overlap in time.

The durations in which the channel remains in the above states are denoted by σ, T_S, and T_θ, respectively. While σ is a constant equal to one mini-slot in 802.11, the duration of the

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

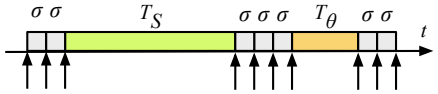

Fig. 2. Discrete time renewal process for the joint channel state evolution.

To State	Transition Probability
i , j	$(1-\gamma_i)(1-\gamma_j)$
0 , j	$\gamma_i(1-\gamma_j)^{f_A}(1-l_A)$
0 , $j+1$	$P_\theta S_{AB} C_{ABa,D}(1-l_A)$
$i+1$, $j+1$	$P_\theta S_{AB}(1-C_{ABa,D}(1-l_A))$

TABLE I
BASIC ACCESS WITH SYMMETRIC CONNECTIVITY.

other intervals (T_S and T_θ) depend on the modulation rates of transmitting nodes, the access mechanism, and the duration of overlapping packets. For example with RTS/CTS enabled, the duration of a successful single-access state with only one node transmitting is equal to:

$$T_{S_n} = \frac{\text{RTS} + \text{CTS} + \text{ACK}}{R_{\text{basic}}} + 3 \cdot \text{SIFS} + \text{DIFS} + \frac{l_n}{R_n} \quad (1)$$

In the above equation, the packet size and modulation rate of node n are denoted by l_n and R_n, respectively. In the overlapping-packets state, one or both packets are captured or both packets are dropped. Thus, the state duration is highly dependent upon capture relationships as well as other system parameters. These values are computed for each case once their corresponding probabilities are calculated.

System State. We represent the system state as the pair $(i,\ j)$, where i and j represent the current backoff stages of transmitters A and B, respectively. Note that $0 \leq (i,j) \leq m$, where $m+1$ is equal to the maximum retransmission limit. The key approximation in our model is that, at each switching time, the next state does not depend on the current state. This allows us to model the evolution of our bi-dimensional state process with a discrete time Markov chain embedded over continuous time at the time instants in which *both senders* can potentially start transmitting the first packet of a new data exchange. These packets could be RTS packets with the four-way handshake or data packets with basic access.

We further assume that a station's backoff counter is geometrically distributed over the contention window. This allows us to exploit the memoryless property of the geometric distribution without accounting for the remaining number of backoff slots. The geometrically-distributed backoff counter at stage i is given by $\gamma_i = \frac{2}{W_i - 1}$ where W_i is the window size of backoff stage i. Consequently, a station in stage i attempts a new transmission with probability γ_i.

Transition Probability Calculation. Nodes A and B have transmission probabilities of γ_i and γ_j, corresponding to backoff stages i and j, respectively. The transition probabilities stem from the generic state (i,j) and are summarized in Tables I through IV for both types of access mechanisms and cross-flow connectivities (refer to Fig. 1).[1]

Basic Access with Symmetric Cross-Flow Connectivity. We now calculate the transition probabilities with basic access (i.e., with RTS/CTS disabled) and symmetric cross-flow connectivity. We have summarized the transition probabilities for this group in Table I. The first row is an idle state in which neither of the nodes is transmitting a new packet. The second row refers to a single-access state leading to a successful transmission by A. The probability of this event is equal to the probability that: *(i)* sender A transmits and sender B does not attempt to transmit during A's transmission, $\gamma_i(1-\gamma_j)^{f_A}$, where f_A denotes A's data packet size in mini-slots, and *(ii)* A's data packet is received successfully by its receiver, $1-l_A$, where l_A is loss probability dependent on the link quality. The corresponding unsuccessful single-access state happens with the probability $\gamma_i(1-\gamma_j)^{f_A} l_A$ due to A's packet loss. Similar single-access state probabilities can be calculated for B due to symmetry. On the other hand, if the state is neither an idle state nor a single-access state, then it would be an overlapping-packets state. We denote the probability that the system enters such a state by P_θ, which is equal to:

$$P_\theta = 1-(1-\gamma_i)(1-\gamma_j)-(\gamma_i)(1-\gamma_j)^{f_A}-(\gamma_j)(1-\gamma_i)^{f_B} \quad (2)$$

Since the later packet can arrive anywhere during the transmission of the first packet, the probability that it arrives during the synchronization bits of the first packet is negligible. Thus, with basic access, we assume that the later arriving data packet is always dropped. The next two rows of Table I calculate the transmission probabilities when A's packets arrives earlier than B. With B arriving earlier than A, the transition probability can be calculated similarly due to symmetry.

Let $C_{XYx,D}$ denote the probability that X's data packet is captured over Y's transmission at receiver, x. The third row of our table denotes a successful transmission by A when it overlaps with B. This probability is equal to the probability that: *(i)* packets overlap, P_θ, *(ii)* A's packet arrives earlier or exactly at the same time as B conditioned that packets overlap, S_{AB}, *(iii)* A is captured, $C_{ABa,D}$, and *(iv)* A's packet is not lost due to poor channel conditions, $1-l_A$. The probability that B's packet arrives later than A conditioned that they overlap, S_{AB}, is equal to:

$$\frac{\gamma_i\gamma_j + \gamma_i\gamma_j \times \sum_{k=1}^{f_A-1}(1-\gamma_j)^k}{\gamma_i\gamma_j + \gamma_i\gamma_j \times \sum_{k=1}^{f_A-1}(1-\gamma_j)^k + \gamma_i\gamma_j \times \sum_{k=1}^{f_B-1}(1-\gamma_i)^k} \quad (3)$$

The first term in the numerator of Eq. 3 is the probability that both packets are transmitted at the same time. The second term calculates the probability that B's packet arrives later. The denominator of Eq. 3 is the overlapping-packets state probability which is equal to P_θ from Eq. 2.

A's packet would be lost if poor channel conditions exist or the packet is not captured. This results in a backoff stage of $(i+1, j+1)$ and is calculated in the fourth row of Table I. Finally, in order to calculate state duration for overlapping states, we assume that the later packet arrives in the middle of the first packet (which occurs on average). We select the overall length as the state duration.

Four-way Handshake with Symmetric Cross-Flow Connectivity. We now calculate the transition probabilities when the RTS/CTS mechanism is enabled. The idle and single-access states can be calculated the same as for basic access by replacing f with f', where f' is the transmission duration of an RTS in mini-slots. As a result, P_θ in Eq. 2 would correspond to the probability of overlapping RTS packets.

Fig. 3 depicts combinations in which B's RTS arrives earlier than A's RTS and at least one flow has a successful RTS/CTS exchange. In Cases 1 and 2, A has a winning RTS/CTS

[1]We assume that the packet loss for control messages is zero while data packets can be lost due to channel conditions and noise at the receiver. Considering control message losses, doubles the number of states which we have omitted for ease of presentation.

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

To State	Probability
$0\,,\,j+1$	$P_\theta S_{BA} O_{BA} C_{ABa}(1-C_{BAb})(1-l_A)$
$0\,,\,j+k$	$P_\theta S_{BA} O_{BA} C_{ABa} C_{BAb} C_{Aba} C_{abA}(1-C_{baB}) P_{k-1}(rts) C_{ABa,D}^{U(k-1)}(1-l_A)$
$i+k\,,\,0$	$P_\theta S_{BA} C_{BAb}(1-O_{BA} C_{ABa} C_{Aba} C_{baB}) P_{k-1}(rts) C_{BAb,D}^{U(k-1)}(1-l_B)$
$i+k\,,\,0$	$P_\theta S_{BA} C_{BAb} O_{BA} C_{ABa} C_{Aba} C_{baB}(1-C_{abA}) P_{k-1}(rts) C_{BAb,D}^{U(k-1)}(1-l_B)$
$0\,,\,0$	$P_\theta S_{BA} C_{BAb} O_{BA} C_{ABa} C_{Aba} C_{baB} C_{abA} C_{BAb,D} C_{ABa,D}(1-l_B)(1-l_A)$

TABLE II
USING THE FOUR-WAY HANDSHAKE WITH SYMMETRIC CONNECTIVITY.

exchange. Cases 3 and 4 refer to a winning RTS/CTS exchange by B. In the last case, both senders receive successful CTS packets, and hence, both transmit data packets.

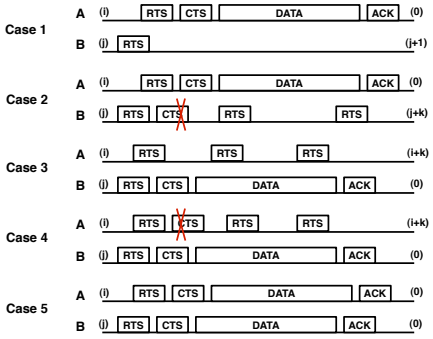

Fig. 3. Five cases for symmetric cross-flow connectivity based on timing and capture behaviors.

For each case in Fig. 3, the corresponding row in Table II calculates the transition probability. The transition probability of Case 1 is equal to the probability that: *(i)* A and B's RTS packets overlap, P_θ, *(ii)* B's RTS arrives earlier conditioned that they overlap, S_{BA}, calculated by Eq. 3 with f' instead of f, *(iii)* A's RTS arrives during the synchronization bits of B's RTS, conditioned that they overlap and B arrives earlier, O_{BA}, *(iv)* A's RTS is captured at its receiver a over B's RTS, C_{ABa}, *(v)* B's RTS is not captured at its receiver, 1-C_{BAb}, and *(vi)* finally, the data packet transmission is successful, 1-l_A, triggering an ACK transmission. Note that, in this case, the CTS packet transmitted by a is received by B, and hence, B defers for the rest of A's transmission. Thus, the only remaining probability to calculate is O_{BA}, which equals:

$$\frac{\gamma_i \gamma_j \sum_{k=0}^{k=s-1}(1-\gamma_i)^k}{\gamma_i \gamma_j \sum_{k=0}^{k=f_B-1}(1-\gamma_i)^k} \tag{4}$$

The numerator of Eq. 4 is the probability that packets overlap and the later packet arrives during the synchronization bits of the first packet where s is the duration of synchronization bits in mini-slots. The denominator calculates the probability that the later packet is transmitted anywhere during the first packet's transmission.

In Case 2 of Fig. 3, A's CTS is transmitted by b, but it is not received by B. Thus, sender B can retransmit RTS packets after a timeout. These RTS packets will not be captured at their receiver since the other flow is transmitting a data packet and RTS packets arrive in the middle of its transmission. If $k-1$ further RTS packets are transmitted by j, the final backoff stage of j at the end of the transmission of the other flow will increase by k.

The second row of Table II calculates the transition probability for Case 2. For a successful transmission by A, the data packet should be captured over the RTS retransmission, $C_{ABa,D}^{U(k-1)}$, where $U(k-1)$ is the unity function and is equal to one if any retransmissions happen and zero, otherwise. $P_n(rts)$ calculates the probability of n additional RTS packet transmissions and is equal to:

$$\begin{aligned} P(n \quad \text{retransmissions}) = {} & \textstyle\sum_{m_1=0}^{L-(n-1)r-1}(1-\gamma_j)^{m_1}(\gamma_j) \times [\\ & \textstyle\sum_{m_2=0}^{L-(n-1)r-m_1-1}(1-\gamma_{j+1})^{m_2}(\gamma_{j+1}) \times [\ldots \times [\\ & \textstyle\sum_{m_n=0}^{L-(n-1)r-1-\sum_{i=0}^{n-1} m_i}(1-\gamma_{j+n-1})^{m_n}(\gamma_{j+n-1}) \times \\ & (1-\gamma_{j+n})^{\frac{((L-nr-1-\sum_{i=1}^{n} m_i)+(|L-nr-1-\sum_{i=1}^{n} m_i|))}{2}}]]] \end{aligned} \tag{5}$$

In the above expression, L is equal to the number of available transmission opportunities in mini-slots. For example, in Case 2 of Fig. 3, this duration is equal to the duration of the data and ACK exchange of flow Aa minus the timeout duration. RTS retransmissions only occur after a fixed timeout, and r is equal to the RTS plus timeout duration in mini-slots.

Eq. 5 calculates the retransmission probability by adding all combinations in which n packets of size r can fit in L mini-slots. The above expression divides the whole duration into n parts, each of size $m_i + r$ mini-slots, where j transmits after the first m_i slots. Since the retransmitting RTS packets fail, the backoff stage of the transmitter increases and will be reset to the minimum backoff once the maximum is reached.

The remaining rows in Table II calculate the transition probabilities for the corresponding cases. These transition probabilities account for successful transmission by the flow with a successful RTS/CTS exchange. Hence, there must be a corresponding state that accounts for unsuccessful data transmissions which can be calculated from their corresponding successful transmissions. If flow A's RTS is leading, the overlapping-packets state probabilities can be easily calculated from Table II due to topological symmetry. Finally, a state should account for overlapping RTS transmissions where neither receives a successful CTS packet. The resulting backoff stage would be $(i+1, j+1)$, and its probability is one minus the summation of all other states.

To State	Probability
$0\,,\,j+1$	$P_\theta S_{AB}(1-l_A)$
$0\,,\,0$	$P_\theta S_{BA} C_{BAb,D} C_{Aba,D}(1-l_A)(1-l_B)$
$0\,,\,j+1$	$P_\theta S_{BA}(1-C_{BAb,D}+C_{BAb,D} l_B)(1-l_A)$
$i+1\,,\,0$	$P_\theta S_{BA} C_{BAb,D}(1-C_{Aba,D}+C_{Aba,D} l_A)(1-l_B)$

TABLE III
BASIC ACCESS WITH ASYMMETRIC CONNECTIVITY.

Basic Access with Asymmetric Cross-Flow Connectivity. We now move to the topology with asymmetric cross-flow connectivity and basic access. In this topology, a will not receive B's transmissions, whereas b receives transmissions by A and a (refer to Fig. 1). The overlapping-packets state probabilities are summarized in Table III, and the non-overlapping states remain the same as the symmetric scenario.

When packets overlap, a will only receive the packet sent by A, whereas b receives data packets transmitted by both A and B. As a result, if A's packet arrives earlier and is not lost

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

To State	Probability
$0\,,\, j+k$	$\gamma_i(1-\gamma_j)^{f'_A}P_k(rts)(1-l_A)$
$0\,,\, j+k$	$P_\theta S_{BA}(1-C_{BAb})(1-l_A)$
$0\,,\, j+k$	$P_\theta S_{AB}(1-O_{AB}C_{BAb}C_{Bab})P_{k-1}(rts)(1-l_A)$
$i+k\,,\, 0$	$P_\theta S_{BA}C_{BAb}(1-C_{Aba})O'_{BA}C_{bAa}P_{k-1}(rts)C_{BAb,D}^{U(k-1)}(1-l_A)$
$0\,,\, 0$	$P_\theta S_{BA}C_{BAb}(1-C_{Aba})(1-O'_{BA}C_{bAa})(1-P_0(rts))C_{BAb,D}C_{Bab,D}(1-l_B)C_{Aba}(1-l_A)$
$0\,,\, 0$	$P_\theta S_{BA}C_{BAb}C_{Aba}(1-C_{abA})(1-P_0(rts))C_{BAb,D}C_{Bab,D}(1-l_B)C_{Aba}(1-l_A)$
$i+1\,,\, 0$	$P_\theta S_{AB}O_{AB}C_{BAb}C_{Bab}C_{baA}(1-l_B)$
$0\,,\, 0$	$P_\theta S_{AB}O_{AB}C_{BAb}C_{Bab}(1-C_{baA})(1-C_{baA})(1-P_0(rts))C_{BAb,D}C_{Bab,D}(1-l_B)C_{Aba}(1-l_A)$

TABLE IV
USING THE FOUR-WAY HANDSHAKE WITH ASYMMETRIC CONNECTIVITY.

due to channel conditions, it will be successfully received at its receiver (first row of Table III).

If B's packet arrives earlier, three different scenarios can happen: simultaneous successful transmissions (second row of Table III), successful transmission only by A (third row), or successful transmission only by B (fourth row).

Any other overlapping-packets states will result in a backoff stage increase by both flows, where the probability is equal to one minus the summation of all other probabilities. Finally, we assume that with overlapping packets the later packet arrives in the middle of the other flow's transmission and take the overall length as the state duration for each case.

Four-way Handshake with Asymmetric Cross-Flow Connectivity. The main transition probabilities of this group are summarized in Table IV, and a sample of timeline graphs are plotted in Fig. 4. The first three rows of the table correspond to Case 1, in which A has a successful data transmission while the other flow retransmits RTS packets. This can happen in a single-access state or overlapping-packets state with either node transmitting earlier.

The fourth row of the table and Case 2 of Fig. 4, correspond to the probability that B has a successful data packet transmission and arrives earlier. Furthermore, the CTS packet transmitted by b is captured over A's RTS at a. As a result, future RTS transmissions by A will not be replied to by a since it is able to set its NAV timer correctly. Hence, the backoff stage of A will increase if it retransmits RTS packets. Note that the CTS packet should arrive during the synchronization bits of A's RTS. Denoted by O'_{BA}, this probability can be derived from Eq. 2. On the other hand, if the CTS packet transmitted by B is not captured over A's RTS at a, different states happen depending on additional attempts of A and on successful or failed transmissions of each flow. One such example is plotted in Case 3 and calculated in the fifth row, where we have presented the probability of a additional attempt which results in a simultaneous successful transmission. Other states include no additional attempt by A, loss by either flow, or loss by both flows. The sixth row of the table calculates the same probability of Case 3 except the CTS packet is unsuccessful.

In the last case of Fig. 4, node A transmits earlier and captures the CTS packet transmitted by b. As a result, no further attempt is made. This probability is calculated in the seventh row of Table IV. If the CTS packet is not captured, different outcomes can happen similar to Case 3. The last row of the table calculates the probability if A makes another attempt and is successful. With certain capture relationships, two flows can have simultaneously winning RTS/CTS transmissions and hence, data packet transmissions. These probabilities can be calculated by plotting the timeline graphs which we have omitted due to space limitations.

We emphasize that all the probabilities presented in Table IV assume a successful transmission by the flow winning RTS/CTS, and unsuccessful transmissions can be derived from them. Finally, any other overlapping-packets states will result in backoff stage increases by both flows, where the probability is equal to one minus the summation of all other probabilities.

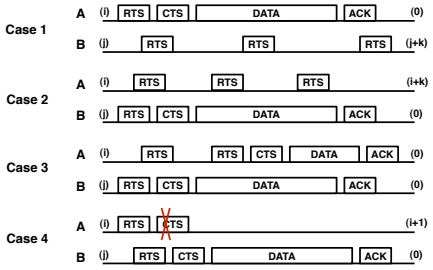

Fig. 4. Four cases for asymmetric cross-flow connectivity based on timing and capture behaviors.

Throughput Calculation. By numerically solving the Markov chain for each access mechanism and topology, which is ergodic for any choice of parameters, we obtain the stationary distribution $\pi = \pi_i, \forall i$. Long-term performance metrics such as throughput can be obtained directly from the solution of the Markov chain. From renewal-reward theory, the throughput of either flow is given by:

$$T = \frac{\sum_n \pi_n P_{S_n}}{\Delta} \tag{6}$$

Here, P_{S_n} is the probability of successful transmission of either flow at state n, and Δ is the average duration of a step. Δ is computed from the average duration of all possible events in all states, weighted by their respective probabilities.

Handling Non-Saturated Flows. So far in our analysis, we have assumed that when the backoff counter of a flow reaches zero, the transmitter always sends a data packet, i.e. the senders are fully backlogged. We now extend our analysis to the case that the packet arrival rate of each flow i is λ_i. We define a new probability ρ_i, which is the probability that the sender has a data packet to send when it is attempting to transmit a packet. To do so, we replace γ_i in our prior equations with $\gamma_i \times \rho_i$. With saturated throughput, ρ_i is equal to one. With unsaturated throughput, the achieved throughput of a flow i is less than or equal to λ_i. With coupled sources, a closed form expression for ρ_i that yields throughput equal to λ_i does not exist. Hence, we approximate ρ_i^{new} as $\alpha \times \rho_i^{old} + (1-\alpha) \times \frac{\lambda_i}{T_i^{old}}$ and adopt a global iterative procedure to update it. During each iteration, we utilize the throughput analysis to update the variables of every node as a function of its neighbor's ρ values (as computed in the previous iteration). The procedure ends when the throughput achieved by each flow is less than or equal to its demand.

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

III. Urban Experimental Evaluation of Information and Channel Asymmetries

In this section, we perform thousands of measurements of coupled flows in an urban mesh network to both validate our model and experimentally analyze the complex factors that contribute to different throughput-sharing modes. With our model, we explore the full set of interdependencies that lead to this behavior and show that reverse capture plays a critical role. Further, we experimentally and analytically show that this inversion can be based on small-scale channel fluctuations common to urban environments.

A. Experimental Set-up and Measured Model Inputs

Throughout the TFA experiments, we activate two fully-backlogged UDP flows (Aa and Bb) with 1500B packets. We repeat the experiment in 120-second intervals for all combinations of 802.11b rates and for both access mechanisms. Each sender to receiver link is strong, enabling the highest modulation rate to achieve a high delivery ratio (i.e., 11 Mbps performs well). Before the experiment, we measure the data packet loss probability per modulation rate for each flow in isolation for our model. During the throughput experiments, we perform per-second SNR measurements and use the average relative SNR per link pair. Our capture measurements from [13] are then used to find the corresponding capture probability. Tables V and VI describe the average relative SNR for each possible competing link pair for the topologies. In some cases, one of the two competing links lacks connectivity which results in a capture probability of 1 for the other link. We denote this as +∞ dB or -∞ dB in Tables V and VI.

B. Channel Asymmetries with Symmetric Connectivity

As a baseline for our model validation, we first consider the throughput of coupled flows with symmetric cross-flow connectivity which has been shown to fairly share bandwidth in idealized channel conditions with equal modulation rates [4], [10]. While this topology has *symmetric* cross-flow connectivity, there is vast heterogeneity in channel conditions between the flows, resulting in diverse capture characteristics based on the packet size and modulation rate [13]. As an example, Table V shows the SNR matrix for the left topology of Fig. 1. Sender A sends to receiver a, and B sends to b. The SNR difference between A and B ($A - B$) is -3.2 dB at b, and $A - B$ is +0.6 dB at a. Hence, with overlapping control packets transmitted by A and B, the probability for B's control packets to win capture at b is 0.98 from [13]. The same packets from B are likely to collide with A's packets at a since the channel differences are small, and the resulting ability of A to capture is negligible. We now consider how accurate our model is at predicting the throughput sharing of these coupled flows with such topological complexities.

Fig. 5(a) and 5(b) depict the throughput achieved by each flow with RTS/CTS enabled. We observe that our model provides an excellent match with measurement results for all combinations of modulation rates. Observe that the throughput for flow Aa is near zero for all rate combinations. Our model reveals the reason for the severely imbalanced throughput. In order to have a successful transmission by flow Aa, its RTS transmission should not overlap with RTS transmissions of B.

	Relative SNR (dB)			
RX	$A-B$	$A-b$	$a-B$	$a-b$
A	-	-	+∞	+1.0
a	+0.6	-1.4	-	-
B	-	-∞	-	+3.9
b	-3.2	-	+1.0	-

TABLE V
Symmetric cross-flow connectivity sub-topology where a positive value favors Aa and negative, Bb.

Moreover, if B's RTS arrives earlier, it will be captured while if A's RTS overlaps with B, it will be dropped. As a result, A's backoff stage will continuously increase whereas B's backoff stage remains close to zero. We have performed similar analysis with RTS/CTS disabled to show that our model is accurate and found that flow Aa achieves nearly zero throughput in this case as well [13]. In summary, for both access mechanisms with symmetric cross-flow connectivity, forward traffic capture causes a bi-modal shift in the throughput-sharing mode where the channel asymmetry overwhelms the symmetry in information at each transmitter. Later, we discuss the additional effects of capture in the reverse direction in this topology.

C. Channel Asymmetries with Asymmetric Connectivity

We now consider coupled flows that compete with asymmetric cross-flow connectivity. Under perfect channels, this case yields one flow starving due to lack of information [4], [10]. Similar to the table for the symmetric sub-topology, Table VI describes the competing links in the sub-topology. Recall that with asymmetric cross-flow connectivity $a-b$ at B is $-\infty$ dB. We repeat the same experiment with this grouping of nodes and channel configuration.

	Relative SNR (dB)			
RX	$A-B$	$A-b$	$a-B$	$a-b$
A	-	-	+∞	+3.8
a	-7.1	-2.8	-	-
B	-	-∞	-	-∞
b	-1.6	-	-2.6	-

TABLE VI
Asymmetric cross-flow connectivity sub-topology (positive value favors Aa and negative, Bb).

Fig. 5(c) and 5(d) report the model's prediction and corresponding throughput measurement for a pair of coupled flows in TFA with RTS/CTS and asymmetric cross-flow connectivity. Surprisingly, the flow without information (Bb) is able to achieve the same throughput as the flow with information (Aa). Thus, the coupled flows have an inverted throughput-sharing mode from [4], [10]. This can be explained by the much larger ability to capture at b for B and A's inability to capture at its own receiver versus b. As opposed to basic access, there are many more dependencies that are required to allow equal sharing that we explore in the next section. In short, the joint presence of forward *and* reverse traffic is required to invert the imbalanced sharing of the topology, making it balanced. With RTS/CTS *dis*abled, we have found that capture in the forward traffic direction alone is sufficient to balance the throughput with asymmetric cross-flow connectivity [13]. *Finding: With asymmetric cross-flow connectivity, forward capture inverts the throughput-sharing mode for basic access. However, forward and reverse capture is required to invert the throughput-sharing mode for the four-way handshake.*

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

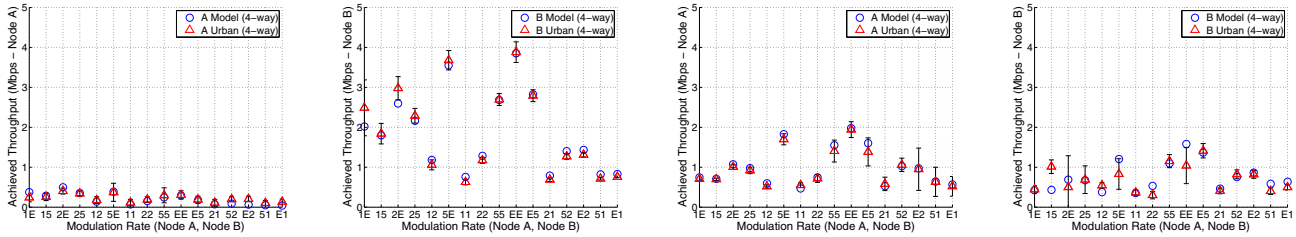

(a) Node A: 4-way Symmetric (b) Node B: 4-way Symmetric (c) Node A: 4-way Asymmetric (d) Node B: 4-way Asymmetric

Fig. 5. Inverted throughput-sharing modes from model and validated by urban experiments on the topology and channel conditions represented by Table V and Table VI with symmetric and asymmetric cross-flow connectivity, respectively. Here, an E represents 11 Mbps.

D. Reverse Capture Shifts the Symmetric Profile

We now explore the full range of the aforementioned interdependencies to invert the throughput-sharing modes. To do so, we use our model and Jain's Fairness Index, defined as $(\sum x_i)^2/(n \cdot \sum x_i^2)$ where x_i is the achieved throughput of flow i and n is the total number of flows [14]. The fairness index of 1 corresponds to an equal throughput sharing whereas a fairness index of 0.5 corresponds to one flow starving and the other obtaining all the throughput.

Even when coupled flows have symmetry in information, the throughput-sharing modes can be inverted for both types of access mechanisms. We now explore the impact that capture relationships have on the throughput sharing of the symmetric cross-flow connectivity topology. While we expect that the forward traffic capture would dominate the behavior of basic access, the role of reverse traffic on sharing is previously unstudied, especially with RTS/CTS enabled. Here, we present the results from our model where two coupled UDP, fully-backlogged flows compete with a modulation rate of 5.5 Mbps.

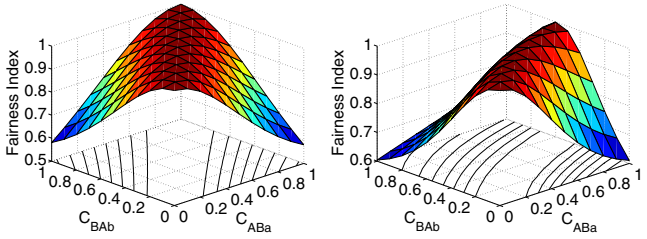

(a) 2-way Symmetric (b) 4-way Symmetric

Fig. 6. Symmetric cross-flow connectivity with (a) 2-way access and (b) 4-way access with both reverse captures of Bb to 1.

Fig. 6(a) depicts the fairness achieved by the two transmitters (A and B) based on their ability to capture at their own receivers (C_{BAb} and C_{ABa}) using basic access. The fairness property depends on the symmetry of the forward traffic capture of the two flows. Prior work has predicted three points of this figure: [2] predicted the left and right corners of Fig. 6(a) (starvation mode) and [4], [10] predicted the result with no capture (fair-sharing mode).

With the four-way handshake, the fairness index remains nearly identical to Fig. 6(a). In fact, if the graphs for both access mechanisms were overlaid, the differences are almost indistinguishable. However, the capture requirements for each access mechanism is very different (i.e., for a given channel condition, RTS packets have much lower capture thresholds than data packets). *Finding: For a given channel condition, use of two- vs. four-way handshake can yield a bi-modal shift because of the increasing ability to perform physical layer capture by the lower modulation rate and smaller size of the RTS packet as compared to the data packet.*

Fig. 6(b) depicts our model's throughput prediction for the symmetric case with the four-way handshake where the reverse traffic is fully able to capture (C_{bAa} and C_{baA} equal 1). We observe that a shift in the sharing occurs, favoring flow Bb (which is able to capture in the reverse direction). To achieve a balanced throughput in this case, flow Aa must have a greater forward traffic capture than Bb. *Finding: Reverse traffic capture shifts the throughput-sharing mode with symmetric cross-flow connectivity and the four-way handshake.*

E. Forward and Reverse Capture with Information Asymmetry

With asymmetric cross-flow connectivity, fairness occurred when the information-poor flow Bb (i.e., the flow which lacks information) is able to capture in the forward traffic direction (C_{BAb}) for basic access. However, we have not yet considered the effect on the sharing when the information-rich flow Aa also has forward traffic capture (C_{Aba}).

Fig. 7(a) depicts a three-dimensional diagram of the fairness of two transmitters, A and B, according to their respective ability to capture at receivers a and b. In the left part of the figure, the information-poor node (B) is able to completely capture at b and A is unable to capture at a. This is the scenario that leads to perfect sharing for basic access. As A's forward traffic capture (C_{Aba}) increases, the fairness index decreases rapidly and independent of B's forward traffic capture value (C_{BAb}). *Finding: With asymmetric cross-flow connectivity and basic access, inversion of the throughput-sharing mode primarily depends on the information-rich flow losing forward traffic capture and secondarily depends on the information-poor flow winning forward traffic capture.*

Now, we consider the asymmetric case with the four-way handshake. Here, all four directions of capture in the forward and reverse directions must be considered since the RTS/CTS exchange preempts any data transmission. For the information-poor flow (Bb), the most important relationships to equalize throughput sharing is the forward over forward traffic capture (C_{BAb}) and the reverse over forward traffic capture (C_{bAa}). We first present the results from the model with these two capture relationships. We later show other relationships that contribute to increased throughput of Bb.

Fig. 7(b) depicts the fairness index for asymmetric cross-flow connectivity with the four-way handshake based upon the ability of the information-poor flow Bb to capture in the forward direction versus competing forward traffic (C_{BAb}) and in the reverse direction versus competing forward traffic (C_{bAa}). On the left and right corners of the figure, near starvation of flow Bb occurs with the complete capture of the reverse or forward direction, respectively. However, in the middle of the figure, both relationships winning capture

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

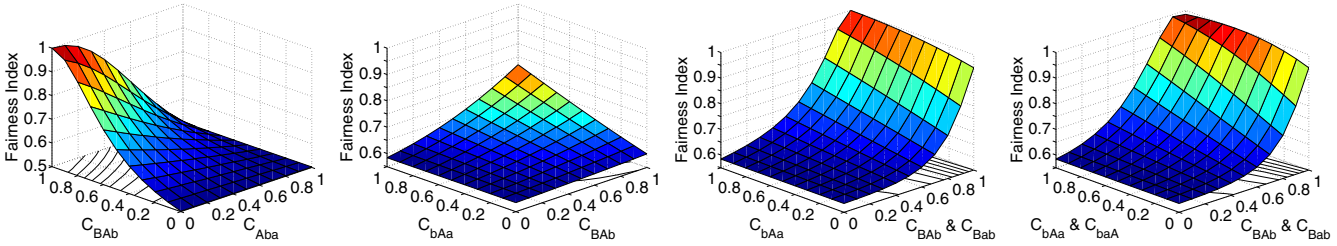

(a) 2-way Asymmetric (b) 4-way Asymmetric (c) 4-way Asymmetric (d) 4-way Asymmetric

Fig. 7. Asymmetric cross-flow connectivity with 2-way handshake (a) and 4-way handshake (b-d) where the information-poor flow has increased chance to win the contention due to favorable capture relationships.

contribute to a much more equivalent throughput sharing. In Fig. 7(c), we add the ability of Bb to capture in the forward direction versus the reverse traffic (C_{Bab}). In the middle of the figure, we observe that a completely fair distribution of throughput (and complete inversion) can now be experienced. Finally, in Fig. 7(d), we add the ability of Bb to capture in the reverse direction versus reverse traffic (C_{baA}). Where these four capture relationships are 1 (middle of the figure), we observe that the flow Bb actually achieves greater throughput than Aa. *Finding: With asymmetric cross-flow connectivity and the four-way handshake, the information-poor flow requires a confluence of link capture relationships to cause the throughput-sharing mode to invert. Yet, when the throughput-sharing mode does invert, the information-poor transmitter can obtain even* higher *throughput than the information-rich transmitter, a behavior that does not occur with basic access.*

F. Small-scale Channel Fluctuations Driving Modes

From our thousands of measurements over the course of a month on multiple topologies, we found many topologies have highly varying throughput sharing. The vast differences are despite the use of off-peak times for our experiments and limited activity of other nodes in the mesh network. In a particular grouping of four nodes with asymmetric cross-flow connectivity (described in Table VI), we found that the throughput sharing over a month's time period went from a starvation mode to a fair-sharing mode.

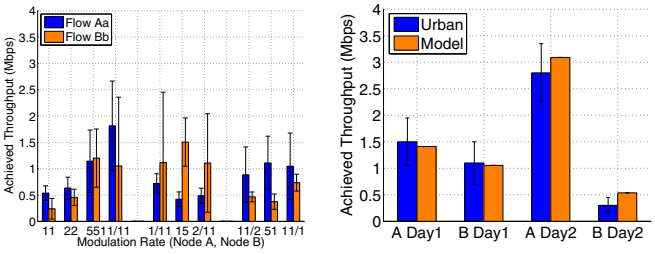

(a) High Variation (over a month). (b) Throughput (on two different days).

Fig. 8. Asymmetric topology (from Table VI) over a month of measurements and measurements from the two days where the difference in throughput achieved is minimum (day 1) and maximum (day 2).

Fig. 8(a) depicts each flow's throughput over the course of a month where the average is represented by a bar, and the standard deviation is represented by error bars above and below the average. We observe that across many different modulation rate combinations, both flows have highly varying achieved throughput. Fig. 8(b) shows single-day measurements where both the minimum and maximum difference between throughput sharing were achieved. With the minimum difference in sharing (Day 1), nearly fair throughput is achieved, while the case with the maximum difference (Day 2) is severely imbalanced. By examining the differences in average SNR values between the two experiments, we observed a 1 dB relative difference in the pair of competing links in the forward traffic direction with negligible differences for all other links. Namely, when flow Bb is able to win forward traffic capture, it can achieve approximately equal throughput with flow Aa (as described with both access mechanisms in the validation experiment above). However, when B is unable to do so (e.g., when $A - B$ at b in Table VI goes from -1.6 dB to -0.6 dB), there is a large throughput difference between flows Bb and Aa. Recall that reverse traffic capture is present in this topology allowing the forward traffic capture relationship to make a difference. Therefore, small-scale channel fluctuations allow switching between the fair-sharing and starving modes. While the change in channel state is between different days of the month (larger time-scale), it exposes a corresponding throughput sharing behavior that are caused by the small-scale fluctuations in the TFA network presented in [13]. *Finding: Even common, small-scale channel fluctuations can cause a change in mode to invert the throughput-sharing mode.*[2]

G. Extensions and Applications

Thus far, we have analyzed the interactions of packet size, modulation rate, capture, topology, and channel conditions on the throughput sharing of two coupled flows. In [13], we have extended our work by using an approximate decoupling technique to investigate these interactions on a single flow's throughput competing with many uncoupled sources. Furthermore, we apply our analysis and model to two domains.

First, we consider the well-studied problem of rate adaptation to reveal how the joint consideration of channel and information asymmetries change protocol decisions. Prior work has considered the problem of choosing the modulation rate that achieves the highest throughput based on the channel condition from the sender to receiver [15]. However, no prior work has considered the interdependence of physical layer capture and modulation rate selection. We find through our analysis and experiments the following. *With interfering uncoupled sources, the throughput-maximizing modulation rate may be lower than the throughput-maximizing modulation rate allowed for the flow in isolation and higher than the throughput-maximizing modulation rate for a channel with constant noise.*

Second, we predict and explain the disproportionate effect of low-rate control traffic on data flows within a mesh network observed in [12]. Control traffic such as routing announcements with low average rate (10's of Kbps) has a disproportionate impact on data traffic throughput, potentially reducing data throughput by 100's of Kbps. Our analysis

[2]We also experimentally found that 1 dB caused a bi-modal shift with symmetric cross-flow connectivity, but this is not presented here due to space limitations.

and experiments reveal the following. *The largest throughput reduction factors are due to the joint factors of control traffic originating from a hidden terminal and the data traffic's inability to win capture over the control traffic. Furthermore, we show that since the control traffic has an uncoupled backoff with broadcast traffic, its ability to win capture does not affect the throughput of the data flow as with the interfering source with coupled backoff behavior.* Details of our work on both application domains can be found in [13].

IV. Related Work

Analytical Models of 802.11 and CSMA. There is a rich body of work on modeling CSMA, dating back to the seminal work by Kleinrock and Tobagi [16]. Other models include a perfect capture assumption which was based upon the timing of the packet as opposed to the channel condition [2], [17]. With the introduction of 802.11, Bianchi presented a simplified model that used the assumptions of single-rate, single-clique, fully-backlogged traffic with fixed packet size [11]. Reference [18] considered physical layer features such as hidden terminals and capture without topological asymmetries. More general topologies and scenarios were later explored with an idealized channel and interference model [4], [5]. Recent measurement-based models use $O(n)$ measurements of n nodes to predict throughput for use in applications such as online network management [19], [20], [21].

Measurement Studies of Multi-Hop 802.11 Networks. A number of works have identified the channel conditions and timing under which physical layer capture occurs for pairs of nodes [1], [7], [22]. Others have proposed modifying physical layer properties to address the lack of fairness that results [23], [24]. Measurements have been performed on indoor multi-hop wireless topologies to characterize interactions of flows [9]. Finally, measurement studies have been performed in mesh networks to explore the link behavior [25], flow performance [26], rate adaptation [27], and overhead effects [12].

In contrast, our work is the first to jointly study information and channel asymmetries via modeling and experimentation to reveal that small-scale channel fluctuations common to an urban mesh deployment can yield bi-modal performance shifts.

V. Conclusion

In summary, we perform extensive measurements on coupled flows within an urban mesh network and analytically model the complex factors that exist with information and channel asymmetries. Our experimental analysis and model reveal that small-scale channel fluctuations common to urban environments are able to completely invert the throughput-sharing mode. Using our model, we explore the interdependencies of these complex factors and find that reverse capture plays a critical role in defining these performance modes. Finally, we show how to extend and apply our model and experimentation to two different problem domains: modulation rate selection and the interaction of control and data traffic.

VI. Acknowledgements

This research was supported by NSF grants CNS-0751173, CNS-08102501, CNS-0619797, the Cisco Collaborative Research Initiative, and the Texas Instruments Leadership University Program.

References

[1] A. Kochut, A. Vasan, A. U. Shankar, and A. Agrawala, "Sniffing out the correct physical layer capture model in 802.11b," in *Proc. of IEEE ICNP*, Oct. 2004.
[2] F. Tobagi, "Modeling and performance analysis of multihop packet radio networks," *Proc. of the IEEE*, vol. 75, no. 1, pp. 135–155, Jan. 1987.
[3] T. Nandagopal, T. Kim, X. Gao, and V. Bharghavan, "Achieving MAC layer fairness in wireless packet networks," in *Proc. of ACM MobiCom*, Boston, MA, Aug. 2000.
[4] M. Garetto, J. Shi, and E. Knightly, "Modeling media access in embedded two-flow topologies of multi-hop wireless networks," in *Proc. of ACM MobiCom*, Aug. 2005.
[5] M. Garetto, T. Salonidis, and E. Knightly, "Modeling per-flow throughput and capturing starvation in csma multi-hop wireless networks," *IEEE/ACM Transactions on Networking*, vol. 16, no. 10, pp. 864–877, Aug. 2008.
[6] C. Ware, J. Judge, J. Chicharo, and E. Dutkiewicz, "Unfairness and capture behavior in 802.11 adhoc networks," in *Proc. of IEEE International Conference on Communications*, New Orleans, LA, 2000.
[7] J. Lee, W. Kim, S.-J. Lee, D. Jo, J. Ryu, T. Kwon, and Y. Choi, "An experimental study on the capture effect in 802.11a networks," in *Proc. of ACM WiNTECH*, Sept. 2007.
[8] G. Judd and P. Steenkiste, "Understanding link-level 802.11 behavior: Replacing convention with measurements," in *Proc. of WICON*, Oct. 2007.
[9] J. Lee, S.-J. Lee, W. Kim, D. Jo, T. Kwon, and Y. Choi, "Understanding interference and carrier sensing in wireless mesh networks," *IEEE Communications Magazine*, vol. 47, no. 7, July 2009.
[10] V. Bharghavan, S. Demers, S. Shenker, and L. Zhang, "MACAW: A media access protocol for wireless LANs," in *Proc. of ACM SIGCOMM '94*, London, UK, Aug. 1994.
[11] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE JSAC*, vol. 18, no. 3, pp. 535–547, Mar. 2000.
[12] J. Camp, V. Mancuso, O. Gurewitz, and E. Knightly, "A measurement study of multiplicative overhead effects in wireless networks," in *Proc. of IEEE INFOCOM*, 2008.
[13] J. Camp, E. Aryafar, and E. Knightly, "Understanding contending CSMA flows: Analytical modeling and urban measurements," Tech. Rep., http://lyle.smu.edu/∼camp/papers/twoflowTR.pdf.
[14] R. Jain, W. Hawe, and D. Chiu, "A quantitative measure of fairness and discrimination for resource allocation in shared computer systems," Ohio State University, Columbus, Tech. Rep. DEC-TR-301, 1984.
[15] G. Judd, X. Wang, and P. Steenkiste, "Efficient channel-aware rate adaptation in dynamic environments," in *Proc. of ACM MobiSys*, 2008.
[16] L. Kleinrock and F. A. Tobagi, "Packet switching in radio channels: Part I - carrier sense multiple-access modes and their throughput-delay characteristics," *IEEE Transactions on Communications*, vol. 23, no. 12, pp. 1400–1416, 1975.
[17] R. R. Boorstyn, A. Kershenbaum, B. Maglaris, and V. Sahin, "Throughput analysis in multihop csma packet radio networks," *IEEE Transactions on Communications*, vol. 35, no. 3, pp. 267–274, mar 1987.
[18] M. Carvalho and J. J. Garcia-Luna-Aceves, "A scalable model for channel access protocols in multihop ad hoc networks," in *Proc. of ACM MobiCom*, Sept. 2004.
[19] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, "Measurement-based models of delivery and interference in static wireless networks," in *Proc. of ACM SIGCOMM*, Pisa, Italy, Sept. 2006.
[20] L. Qiu, Y. Zhang, F. Wang, M.-K. Han, and R. Mahajan, "A general model of wireless interference," in *Proc. of ACM MobiCom*, Sept. 2007.
[21] Y. Li, L. Qiu, Y. Zhang, R. Mahajan, and E. Rozner, "Predictable performance optimization for wireless networks," in *Proc. of ACM SIGCOMM*, Aug. 2008.
[22] G. Judd and P. Steenkiste, "Understanding link-level 802.11 behavior: Replacing convention with measurement," in *Proc. of IEEE WICON*, Oct. 2007.
[23] S. Ganu, K. Ramachandran, M. Gruteser, I. Seskar, and J. Deng, "Methods for restoring MAC layer fairness in IEEE 802.11 networks with physical layer capture," in *Proc. of ACM REALMAN*, May 2006.
[24] T.-Y. Lin and J. Hou, "Interplay of spatial reuse and SINR-determined data rates in CSMA/CA-based, multi-hop, multi-rate wireless networks," in *Proc. of IEEE INFOCOM*, May 2007.
[25] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11 mesh network," in *Proc. of ACM SIGCOMM*, Portland, OR, 2004.
[26] J. Bicket, S. Biswas, D. Aguayo, and R. Morris, "Architecture and evaluation of the MIT Roofnet mesh network," in *Proc. of ACM MobiCom*, Cologne, Germany, August 2005.
[27] J. C. Bicket, "Bit-rate selection in wireless networks," M.S. Thesis, MIT, February 2005.