

Design and Experimental Evaluation of Context-Aware Link-Level Adaptation

Jialin He, Hui Liu, Pengfei Cui, Jonathan Landon, Onur Altintas,
Rama Vuyyuru, Dinesh Rajan, and Joseph Camp

Abstract— Context awareness has received increasing attention with the proliferation of various types of sensors on mobile devices. However, while wireless performance is known to be highly correlated with environmental settings, mobile devices have yet to fully exploit the awareness of context to improve wireless performance. In this paper, we leverage available context information to improve link-level adaptation via decision-tree classifiers and extensively evaluate its performance over emulated channels as well as with in-field trials. We first propose a classification method based on decision trees to select the optimal transmission parameters such as modulation, coding rate and packet size. We then quantify the throughput improvement using the proposed scheme and show that in some scenarios the throughput increases by over 100% compared to traditional SNR-based rate adaptation protocols. Second, we analyze the amount of training to assess the classification scheme. Third, we validate classification-based method by implementation on two different test platforms for extensive experimentation. We reveal the importance of the various contextual attributes used and identify channel type as a key parameter that affects classification performance. Finally, we study and quantify the use of context information across multiple different frequency bands and demonstrate the significant throughput gains that can be obtained.

I. INTRODUCTION

Sensing information is becoming increasingly available and exploited on mobile devices for various applications such as light sensors for brightness control, accelerometers for auto-rotation, and proximity sensors to prevent key presses. These innovations follow extensive work on context awareness in the field of computing [1], [2]. Yet, even with location or environmental information being readily available and wireless performance being shown to heavily depend on surrounding conditions [3], wireless protocols have just begun to use context information to improve performance. In particular, recent work has enabled context awareness to choose which wireless technology is most energy-efficient [4], determine the application and energy requirements of the wireless connection [5], calculate the relative distances and speed of two nodes of a wireless link [6], and respond to mobility patterns for rate adaptation [7]. However, each work has not explicitly considered wireless channel characteristics and the resulting performance due to the surrounding environmental setting.

In this paper, we present a context-aware link-level adaptation framework that uses a decision-tree classifier to improve

wireless performance. We evaluate our methods on a diverse set of wireless channels in the field and on a channel emulator and show that in certain scenarios our approach can nearly double the throughput compared to traditional SNR-based adaptation schemes. The specific contextual information that we consider in this paper are the measured SNR, node velocity, and *channel type*. A simplistic characterization of channel type could be a pedestrian or vehicular channel. The basic problem of interest is as follows: *given the channel type, SNR, and velocity, find the transmission mode that achieves the highest throughput*. To solve this problem, we first classify the wireless performance using a decision tree trained across a wide range of different contexts or environmental scenarios. Based on this training data, the classifier can quickly and efficiently arrive at optimal decisions for the link-level parameters to use if the node encounters a context *similar* to the ones encountered in the training. The algorithm also has a tunable threshold parameter that determines when the present context is different enough from others in the past and consequently initiates training in this scenario. The results of this training are subsequently appended to the decision tree to assist with future context-aware link adaptation. Lastly, motivated by increasingly frequency-agile hardware, we consider the viability of using training from one frequency band to another and find that a yet unused band's performance can be improved with training across frequency bands.

The main contributions of this paper are as follows:

- We propose a simple and efficient decision-tree-based classification framework that provides a way to select the optimal transmission parameters in each context. The use of the classification is shown to result in significant throughput improvement over traditional SNR-based rate adaptation protocols.
- We quantify the amount of training required to construct the decision tree and study the robustness of the decision tree to accommodate partial context information (such as context information with missing attributes). Results indicate that knowledge of the *channel type* has a significant impact on the throughput gains that can be obtained.
- We demonstrate via experimentation that the decision tree generated for one band can aid link adaptation on another band. In other words, the rate adaptation scheme can quickly adapt in a new frequency band without re-training, thereby reducing the total training overhead for emerging multiband cognitive systems to become context aware. Finally, we compare the performance of the *implemented* context-aware adaptation with purely SNR-based adaptation schemes.

J. He, H. Liu, P. Cui, J. Landon, D. Rajan, and J. Camp are with Southern Methodist University, Dallas, TX. O. Altintas is with Toyota InfoTechnology Center, Co. Ltd., Tokyo, Japan. R. Vuyyuru is with Toyota InfoTechnology Center, U.S.A. Inc., Mountain View, CA.

The SMU authors were supported on this work by NSF grant CNS-0958436 and Toyota InfoTechnology Center, U.S.A. Inc.

We experimentally analyze our framework with custom and off-the-shelf hardware platforms to show that the context-aware link-level adaptation is platform-agnostic. The custom hardware platform we use is the Wireless Open-Access Research Platform (WARP) [8], while the off-the-shelf platform is a Gateworks 2358 board with four 802.11-based Ubiquiti radios (XR7, XR9, XR2, and XR5) with separate frequency bands. Each of the platform's radios have a physical layer based on the IEEE 802.11a/g standard, but the latter platform offers built-in GPS, increased transmission power, physical layer coding, and fulfills the timing requirements of 802.11a/g whereas the former offers advanced programmability and observability across the network stack.

The paper is organized as follows. We first formulate the problem and present the classification-based context-aware scheme in Section II. We extensively evaluate our methods in a diverse set of channels and across multiple frequency bands in Section III. Lastly, we consider related work in Section IV and conclude in Section V.

II. WIRELESS CLASSIFICATION BASED ON CONTEXT

In this section, we describe the formulation of the main problem of interest, introduce decision tree classifiers, and describe how they can be applied to the current scenario.

A. Problem Formulation

The objective of this work is to demonstrate the improvements in performance by leveraging context information to make link-level rate adaptation decisions. As noted before, the specific context information that we consider are the measured SNR, node velocity and channel type. In the proposed architecture to support context awareness, channel type can be acquired from a combination of the current location (e.g., via GPS) and a land-use map generated from past measurements and stored in a database [9]. In this paper, we do not assume the existence of a database with channel type information stored at each location. Instead, we focus on inferring the channel type information from in-situ and on-line training.

Let set S denote the n -tuple of link-level parameters that can be selected. In the numerical results, we consider two different sets S for the two different platforms. For the WARP platform, $S_{WARP} = \{m, p\}$, where $m \in \{BPSK, QPSK, 16-QAM\}$ represents the modulation that is selected and $p \in \{100, 1000\}$ represents the possible packet sizes in bytes. For the Ubiquiti platform, $S_{Ubiquiti} = \{q, p\}$, where $q = \{c, m\}$ represents the 802.11a/g compliant coding and modulation pairs that are supported, and p is the packet sizes being considered. Let set C represent the possible context information that is available at a particular node at a given time.

In this paper, the optimization metric of interest is the measured throughput G . The optimization problem is stated as follows: Given a particular context $c \in C$, select the optimal $s^* \in S$ that maximizes the throughput, G_{th} , where R_{th} is the physical layer data rate. Formally, the problem is posed as follows:

$$\begin{aligned} \max_{s \in S} \quad & G_{th} = (1 - PER) * R_{th} * \frac{l_{payload}}{l_{packet}} \quad (1) \\ \text{given} \quad & SNR, \text{velocity}, \text{channel type} \end{aligned}$$

where PER is the measured packet error rate, and $l_{payload}$ and l_{packet} are the sizes of the payload and packet, respectively. The optimization problem is solved using a look-up table generated as the output of a classifier, as explained in Section II-B. We also consider two different variations of (1). (i) For comparison and to act as a baseline, we consider a solution to (1) in which the system has knowledge of SNR only. (ii) We generate the values that maximize throughput with extensive evaluation over the entire combination of transmission strategies. The corresponding maximal throughput, G^* , serves as an upper bound to the performance that can be achieved by context-aware adaptation. Note that this upper bound is computed off-line via exhaustive search and cannot be implemented in real-time on practical systems.




Fig. 1. Performance per mode (modulation, packet size) in WARP on a 4-tap channel with a velocity of 60 kmph.

As a visual illustration of the problem statement, Fig. 1 depicts the throughput of our custom hardware platform, WARP, across three modulation orders (BPSK, QPSK, 16-QAM) and two packet sizes (100B, 1000B) composing six different transmission modes (while the off-the-shelf platform has 18 total modes). The test is performed on a channel emulator with an experimental set-up as discussed in Section III. The channel model used has a power-delay profile consisting of four taps with a relatively low delay spread and little multipath, meaning that the direct path from the sender to receiver predominantly characterizes the channel condition. The node is moving at a velocity of 60 kmph. In the figure, we can observe that different modes achieve the highest throughput for a given SNR region and would be said to be the *ideal* rate choices in such regions. Imagine a three-dimensional space that has such curves across velocities and SNRs. Considering that multiple channel types exist, searching over the entire space for the highest performing regions becomes intractable. Thus, we use a decision-tree classifier in the next section to find the regions of SNR and velocity across a broad range of channel types.

B. Decision-Tree Classification

Classification algorithms attempt to extract the relationship among similar data points such as the wireless performance of various transmission modes in different environmental contexts. Decision trees [10], [11] are one of the most popular classifiers due to their ability to deduce a set of simpler




Fig. 2. Decision Tree for WARP Measurements.

decisions from a complex classifying problem, providing a decision which is easily interpreted. In particular, in our framework, the decision-tree classification scheme is used to establish the connection between the sensing information and the optimal link setting. If the combination of different sensing data such as environmental type, node velocities, and channel qualities can be classified into the appropriate category, the wireless transmitter can eliminate poor choices for a particular setting and quickly converge to optimality. Also, once abundant sensing information is available, the decision tree can select only the context data it needs, to conserve resources.

A decision tree is a widely-used information-entropy-based classification algorithm. In this scheme, the connection between classes and features is represented in the form of a tree structure. Classification starts at the root of the tree and moves to leaf nodes until a class is encountered. Due to its low complexity and reasonable classification performance, this scheme is easy to implement and adapts to real situations well without a high computational load. In the presence of large amounts of contextual information, the decision tree classifies the cases in such a way that it maximizes the information metric, which can be either an information gain or a gain ratio. Thus, the classification scheme can streamline the large contextual-information data set.

Before integrating the decision-tree classification model into context-aware wireless systems, in-situ training is needed. In one setting, we record the set of the attenuation, velocity, channel type, and all the modes with achieved throughput. We identify the best mode having the highest throughput for each point and use it with other contextual information for training. We choose the C4.5 algorithm [12], a widely-used algorithm to generate decision trees, to derive the classification scheme used in our system. After an initial training set, subsequent data points are used to test the performance of the decision tree to attempt to find the transmission mode that performs best according to the sensing information tuple. The learned tree for the WARP platform is shown in Fig. 2 (see Section III

for details about the training data). In the tree, we observe that some channels have different context information that is more valuable denoted by its relative distance to the root. For example, velocity is more important in making decisions in Channel D than the other channel types. The decision tree structure forms the primary decision engine for context-aware link adaptation.

III. EXPERIMENTAL EVALUATION

We now evaluate our context-aware link-level adaptation by performing a wide range of tests using both WARP and Ubiquiti platforms. Using both platforms, we show that our algorithm can: (i) accurately select the best transmission mode and (ii) outperform existing SNR-based mechanisms.

A. Experimental Set-up for Controlled, Repeatable Channels

In our experiments, we seek to compare a set of transmission modes against one another to evaluate which setting produces the highest level of performance for a specific wireless channel. Thus, we use an Azimuth ACE-MX channel emulator for repeatability and control through each transmission configuration and corresponding measurement of performance on a given wireless channel.

To understand the performance of each transmission mode, we use both the custom (WARP) and off-the-shelf (Ubiquiti) hardware platforms. For each, the transmitter and receiver are connected with the Azimuth ACE-MX via cabling from the radio to the emulator input. A PC runs Director-II software which controls the emulator (over Ethernet) and allows configuration of the channel characteristics such as model type, path loss, Doppler, and input or output attenuation. For the purposes of the channel emulator experiments, the two platforms have slightly different characteristics. The Ubiquiti radios match the physical layer coding and modulations of IEEE 802.11a/g whereas WARP lacks physical layer coding, currently allowing only BPSK, QPSK, and 16-QAM. With two packet sizes, there are 18 total transmission modes that we study with the Ubiquiti radios and six total modes with WARP. In both cases, the platforms send packets over the emulated channels that have predefined attenuations and velocities, as shown in Table III.

Parameter	Values
Channel Models	Static, Ch. A, Ch. B, Ch. C, Ch. D
Attenuations (dB)	0, 6, 12, 18, 24, 30, 36, 42
Velocities (km/hr)	0, 30, 60, 90, 120
802.11a/g Rates (Mbps)	6, 9, 11, 12, 18, 24, 36, 48, 54
Packet Size (Bytes)	100, 1000

TABLE III

CHANNEL EMULATOR PARAMETERS

For training the decision tree, the five velocities and eight attenuations shown in Table III are used, composing a set of 40 measurement points per channel type. The channel models A-D are a set of broadly used channel types defined by the International Telecommunication Union (ITU). Each measurement is approximately a minute in duration. Using this training, we first evaluate the performance improvement of using our decision-tree classification on randomized channel scenarios. Thus, we randomize velocity and attenuation to form a test set on each channel. We then compare the throughput achieved by the rate decision given by the classification against both

Missing Attributes	No Missing Attributes			Channel Type			Velocity			SNR		
	A	I	G	A	I	G	A	I	G	A	I	G
WARP	76.3%	40.2%	4.2%	53.8%	7.4%	36.7%	72.2%	32.6%	9.3%	62.8%	29.7%	11.3%
Ubiquiti XR2	51.9%	60.8%	66.0%	35.6%	14.3%	75.8%	50.0%	58.0%	66.5%	48.1%	52.9%	67.6%

TABLE I

DECISION TREE PERFORMANCE: ACCURACY (A) IS PERCENTAGE OF TIME CLASSIFICATION CHOOSES IDEAL RATE, IMPROVEMENT (I) IS ITS THROUGHPUT GAIN OVER A SNR-BASED METHOD, AND GAP (G) IS ITS THROUGHPUT PERCENTAGE FROM THE MAXIMUM ACHIEVABLE.

Band	Training: 700 MHz			Training: 900 MHz			Training: 2.4 GHz			Training: 5.8 GHz		
	A	I	G	A	I	G	A	I	G	A	I	G
Testing: 700 MHz	57.7%	35.2%	13.9%	38.5%	99.6%	19.0%	30.8%	33.9%	18.4%	3.9%	6.9%	42.8%
Testing: 900 MHz	46.2%	-6.1%	37.1%	73.1%	17.1%	16.6%	47.1%	-2.4%	28.2%	22.1%	-20.2%	44.9%
Testing: 2.4 GHz	40.4%	50.8%	66.6%	45.2%	52.7%	67.2%	51.9%	60.8%	66.0%	26.9%	46.6%	70.2%
Testing: 5.8 GHz	16.3%	6.2%	24.6%	20.9%	55.1%	19.6%	16.3%	67.1%	22.8%	60.5%	99.8%	6.75%

TABLE II

TRAINING ACROSS FREQUENCY BANDS: ACCURACY (A) IS PERCENTAGE OF TIME CLASSIFICATION CHOOSES IDEAL RATE, IMPROVEMENT (I) IS ITS THROUGHPUT GAIN OVER A SNR-BASED METHOD, AND GAP (G) IS ITS THROUGHPUT PERCENTAGE FROM THE MAXIMUM ACHIEVABLE.

the decision of an SNR-based scheme and the maximum achievable throughput found via exhaustive search. For the SNR-based scheme, the thresholds of ranges are defined by the highest performing rate on a particular SNR in a static topology. In Table I in Section III-B, we more thoroughly describe the results, considering the performance for varying levels of attributes available to the decision tree.

B. Analysis of Tree Structure

By applying the decision-tree-based rate adaptation, the system performance has been greatly improved on both experimental platforms as shown in Table I. Here, accuracy (A) is the percentage of all decisions that the decision tree chooses the rate which has the highest throughput (i.e., the ideal rate). The improvement (I) is the throughput gain by using context-aware classification to make the decision versus using the SNR-based decision for a node in a static channel. The gap (G) is the percentage of throughput from the maximum achievable which is computed via exhaustive search. Table I shows that the throughput with the proposed rate adaption is much closer to the maximum achievable throughput compared with the SNR-based scheme.

To establish the importance of the context attributes, we calculate the classification accuracy with each missing attribute. In other words, we solve the optimization problem in (1) given only knowledge of (i) SNR and velocity, or (ii) SNR and channel type, or (iii) velocity and channel type. This experiment is equivalent to the case in a real system where the system fails to fetch one attribute, or the system receives incorrect context information which cannot be used. The experiment has been done on both the WARP and Ubiquiti platforms, illustrating that the importance of the attributes are platform agnostic. Table I also lists the classification performance with one missing attribute.

As shown in Table I for both platforms, the classification accuracy and the throughput drops if one attribute is missing, and the gap between the throughput of the selected transmission mode and the maximum achievable throughput increases. Note that a missing channel type causes more severe performance loss than the other two attributes. For example, on the WARP platform, a 32% gap occurs if the channel type information is missing, while missing SNR or velocity only leads to 7% or 5% throughput reduction, respectively.

These results demonstrate that knowledge of the channel type is very important to the performance of the context-aware rate adaptation scheme. This is confirmed by the decision tree in Fig. 2 where the relative importance of the attributes defines the tree structure. We have investigated this issue further and found a scheme to accurately determine the channel type, which will be presented in future work.

C. Classification Performance Across Frequency Bands

To further demonstrate the utility of context information in next-generation multiband cognitive networks, we present the following scenario and related question. Consider a setting in which a particular node is operating at a particular frequency band and has developed a decision tree for determining the best operating mode for each context. Now, if the node were to shift transmissions to a different band for which it has no *a priori* training information, can it make coarse estimates of the performance of the various modes in the new frequency band based on the decision tree of the prior context? Our results indicate that, indeed, significant gains are possible in several cases. Table II shows the results when we train the decision tree for a particular band and test the performance on another frequency band (and the same band for comparison). Here, we use the four Ubiquiti radios on 700 MHz, 900 MHz, 2.4 GHz, and 5.8 GHz, corresponding with XR7, XR9, XR2, and XR5, respectively. The experiments are performed on randomized channels as described above.

From Table II, we find classification across different bands has higher throughput in most scenarios as compared to the SNR-based method. For instance, training the tree at 700 MHz and using this to make link-level rate adaptation decisions at 2.4 GHz results in a 50% improvement over SNR-based schemes. The channels of different frequencies share part of their properties. Thus, classification across different bands can gain relative improvement over the SNR-based method. In Table II, when we use the test set of one frequency and the training set of the same frequency, the gap from the optimal performance is smaller than the gap when training is performed at other frequencies. The improvement over SNR-based methods does not show a similar trend because when we use the training set from a different frequency, the SNR-based selection is from the training frequency. Otherwise, the SNR-based selection could have a poorer performance than

the selection of the same frequency as that of the test set, which make the improvements larger than that from training with the set of the same frequency.

IV. RELATED WORK

Machine learning has been applied to communication systems in different propagation environments. The applications of machine learning in cognitive radios have been developed to optimize performance, such as capacity maximization and dynamic spectrum access [13], [14]. In [15], a genetic algorithm is used to optimize the radio parameters for the physical layer. In order to explore the potential of different learning algorithms on cognitive radios, the authors use a loop to determine when to apply which specific learning algorithm in [16]. In machine learning, classification normally refers to the supervised learning algorithm. Based on the information from training data, classification algorithms express the relationship between input features and output categories determined by their properties. Then, incoming data can be assigned to one of the given categories according to the relationship. Classification methods have been widely used in context awareness. In [17], a hidden-Markov classifier uses environmental noise to determine the context (e.g., office or city street). In [18], a probabilistic classification framework based on decision trees and Bayesian Compositional Hierarchies is used for object classification. However, we believe that this is the first application of decision-tree classifiers in link-level adaptation.

In terms of link-level adaptation, adjusting transmission settings to increase the performance of a link has been studied with respect to transmission rate (MCS) adaptation [19], power control [20], and frequency switching [21]. When solved jointly, however, the parameter's ordering of importance is unclear. Our approach classifies the performance based upon in-situ measurements of different parameters in a set of environments. Thus, we allow the possibility of multiple parameter combinations to have the highest performance and have a comprehensive view of how each unique combination will perform against another. Lastly, there are works that use context information to aid rate adaptation [5], [6], [7]. In each work, context is defined in slightly different manners to perform the transmission parameter adaptation: (i) application type, spectrum availability, and energy budget [5], (ii) distance and vehicular speed [6], and (iii) user mobility patterns [7]. In contrast, we use classification that is based on a context consisting of channel type, SNR, and velocity.

V. CONCLUSION

In this work, we present a context-aware framework for link-level adaptation that leverages decision-tree classification. We evaluate the framework across custom and off-the-shelf hardware platforms and achieve nearly double the throughput of SNR-based adaptation. We also explore the role that missing context attributes play in achieving such gains and find that the channel type to be a key factor. Lastly, we evaluate the effectiveness of training across multiple frequency bands and find that the training from one band can lead to informed decisions on another band.

There are three immediate implications from our work. First, while we place the context-aware link-level decision in a protocol similar to SNR-based rate adaptation, the decision could be broadly used in various types of rate adaptation including loss-based mechanisms. Second, since our mechanism has been implemented using MadWifi drivers and Atheros chipsets, which is a common set-up with 802.11 hardware, our scheme can be used by practitioners as well as researchers. Finally, while we study the importance of different attributes and find that the channel type is critical in the rate selection, a separate work would discuss how to infer channels and add channels via training.

REFERENCES

- [1] C. Bolchini, C. Curino, E. Quintarelli, F. Schreiber, and L. Tanca, "A data-oriented survey of context models," *ACM SIGMOD Record*, vol. 36, no. 4, pp. 19–26, 2007.
- [2] B. Schilit, N. Adams, and R. Want, "Context-aware computing applications," in *IEEE Workshop on Mobile Computing Systems and Applications*, Dec. 1994.
- [3] G. Stuber, *Principles of Mobile Communication*, 4th ed. Boston: Kluwer, 2000.
- [4] A. Rahmati and L. Zhong, "Context-based network estimation for energy-efficient ubiquitous wireless connectivity," *IEEE Transactions on Mobile Computing*, vol. 10, pp. 54–66, 2011.
- [5] V. Pejovic and E. M. Belding, "A context-aware approach to wireless transmission adaptation," in *IEEE SECON*, June 2011.
- [6] P. Shankar, T. Nadeem, J. Rosca, and L. Ifode, "CARS: Context-aware rate selection for vehicular networks," in *IEEE ICNP*, Oct. 2008.
- [7] L. Ravindranath, C. Newport, H. Balakrishnan, and S. Madden, "Improving wireless network performance using sensor hints," in *USENIX NSDI*, Apr. 2010.
- [8] P. Murphy, A. Sabharwal, and B. Aazhang, "Design of WARP: Wireless open-access research platform," in *European Signal Processing Conference*, Florence, Italy, June 2006.
- [9] J. Camp, O. Altintas, R. Vuyyuru, and D. Rajan, "Context-aware collection, decision, and distribution (C2D2) engine for multi-dimensional adaptation in vehicular networks," in *ACM VANET Workshop (to appear)*, Sept. 2011.
- [10] J. R. Quinlan, "Induction of decision trees," *Machine Learning*, vol. 1, pp. 81–106, March 1986.
- [11] S. R. Safavian, "A survey of decision tree classifier methodology," *IEEE Trans. on Systems, Man and Cybernetics*, vol. 21, pp. 660–674, 1991.
- [12] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. Witten, "The weka data mining software: an update," *ACM SIGKDD Explorations Newsletter*, vol. 11, no. 1, pp. 10–18, 2009.
- [13] S. Haykin, "Cognitive radio: brain-empowered wireless communications," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 2, pp. 201–220, 2005.
- [14] C. Clancy, J. Hecker, E. Stuntebeck, and T. O'Shea, "Applications of machine learning to cognitive radio networks," *Wireless Communications, IEEE*, vol. 14, no. 4, pp. 47–52, 2007.
- [15] T. Rondeau, B. Le, C. Rieser, and C. Bostian, "Cognitive radios with genetic algorithms: Intelligent control of software defined radios," in *SDR Forum Technical Conference*, Nov. 2004.
- [16] Y. Zhao, J. Gaeddert, K. Bae, and J. Reed, "Radio environment map enabled situation-aware cognitive radio learning algorithms," in *SDR Forum Technical Conference*, Nov. 2006.
- [17] L. Ma, D. Smith, and B. Milner, "Context awareness using environmental noise classification," in *ISCA EUROSPEECH*, Sept. 2003.
- [18] A. Kreuzmann, K. Terzić, and B. Neumann, "Context-aware classification for incremental scene interpretation," in *ACM UCVP*, Nov. 2009.
- [19] J. Camp and E. Knightly, "Modulation rate adaptation in urban and vehicular environments: Cross-layer implementation and experimental evaluation," *IEEE/ACM Transactions on Networking*, vol. 18, no. 6, pp. 1949–1962, Dec. 2010.
- [20] M. Chiang, P. Hande, T. Lan, and C. Tan, "Power control in wireless cellular networks," *Found. Trends Netw.*, vol. 2, pp. 381–533, April 2008.
- [21] H. Rahul, F. Edalat, D. Katabi, and S. G. Sodini, "Frequency-aware rate adaptation and MAC protocols," in *ACM MobiCom*, Sept. 2009.