

RESEARCH

Open Access



Architecture and experimental evaluation of context-aware adaptation in vehicular networks

Hui Liu^{*} , Jialin He, John Wensowitch, Dinesh Rajan and Joseph Camp

Abstract

In vehicular networks, the propagation environment changes rapidly for mobile nodes. To achieve high throughput, wireless devices need to be highly adaptive to these environmental changes by altering their transmission parameters across different layers of the network stack. Sensors in mobile and vehicular nodes can be used to form an understanding of the surrounding context. Such contextual awareness is particularly important in vehicular networks as the frequent context switching and increased channel fluctuations can cause existing adaptation protocols to fail to converge to the optimal transmission parameters. In this paper, we leverage information about the environmental context to enable improved rate adaptation performance in vehicular networks. In particular, we propose a classification-based link-level adaptation framework, which can effectively learn the relationship between context information (such as velocity, SNR, and channel type) and the throughput of various transmission modes. We then quantify the throughput improvement using the proposed scheme and show that our proposed framework can significantly enhance the performance of rate adaptation. With experiments on emulated and in-field channels, we observe that the throughput increases by up to 245% over protocols which use SNR alone to make rate decisions. Based on an analysis of attribute importance, we identify channel type as a key parameter that affects classification performance. Since channel type often cannot be directly obtained, we propose a multi-dimensional channel inference method for use when knowledge about the channel type is not available. We demonstrate that the proposed channel inference achieves an accuracy of up to 94% in previously encountered channels and can quickly signal that a channel has not yet been encountered. The robustness of the proposed methods are demonstrated using experimental data from two different hardware platforms and three different carrier frequency bands. Lastly, we evaluate the most predominant Linux-based rate selection algorithm (Minstrel), study the relative rate selection accuracy of our approach, and analyze the key role that the retry mechanism in Minstrel plays on its performance.

Keywords: Context awareness, Rate adaptation, Vehicular networks

1 Introduction

Mobility has a significant impact on channel characteristics at the physical layer and media access layer [1, 2]. Moreover, vehicular obstructions such as trucks, buildings, and trees can significantly attenuate the signal compared to line-of-sight conditions [3–5]. As today's vehicular and mobile devices are equipped with an increasing number of sensors, vehicular nodes in communication with one another can begin to understand the

environment (e.g., urban, suburban, downtown, residential) and scenario (e.g., relative velocity and distance) in which they are operating. *Context awareness* offers new opportunities for performance improvement in vehicular networks. Prior work has extensively considered using context to improve performance, but primarily in the area of computing [6, 7]. In the area of wireless and vehicular networking, using context is still in the early stages of development.

When nodes travel at vehicular speeds from scenario to scenario, variations of the wireless channel exist such as signal attenuation, channel fading, and interference [8]. An awareness of context could give hints as to the

^{*}Correspondence: huil@smu.edu
Department of Electrical Engineering, Southern Methodist University, Dallas, TX, USA

likely channel state and facilitate quick convergence to the optimal transmission parameters at many layers of the network stack. Hence, the availability of context awareness on today's vehicular nodes provides unprecedented opportunities for adaptation improvement when a node transitions across environments. On the other hand, context awareness poses additional challenges not present in more traditional adaptation protocols: (1) how to incorporate various data streams from diverse sensors or other system resources in vehicular environments in a context-aware system and (2) how to leverage context information specifically for better adaptation choices. The intrinsic problem we seek to solve is to exploit the correlation between the context information and the knowledge of channel behavior to improve rate adaptation decisions. Classification is used extensively in the machine learning community and has increasingly attracted attention in wireless networks. Classification has been proven to be efficient and effective in pattern extraction in wireless traffic classification, cognitive radio applications, and network security [9–11]. The intelligent learning ability of classifiers provides a new perspective for exploiting context information and serves as a valuable tool for context-aware rate adaptation.

Wireless performance of different transmission rates is known to be highly dependent upon the channel conditions, having a fundamental trade-off between robustness and throughput. The high mobility of vehicular nodes exacerbates the fluctuation of the wireless channel from Doppler, multipath, and shadowing effects due to passing cars and structures. Thus, to remain connected and/or improve spectral efficiency in a vehicular scenario, frequent adjustments to the data transmission rate (e.g., modulation rate and coding scheme) is imperative. Various mechanisms have been proposed for adapting transmission parameters of a wireless link such as use of packet statistics at the transmitter [12, 13] and direct channel quality measurement combined with feedback [14]. The former method has been shown to suffer from a confusion of channel and congestion loss types (which motivated further protocol design [15, 16]) whereas the latter can suffer from improper and/or excessive training [17].

In this paper, we propose a classification-based context-aware rate adaptation framework for vehicular and mobile networks and a mechanism to recognize previously encountered channels or trigger new training in an unencountered setting. The specific context information that we consider in this paper are the measured vehicle velocity, signal-to-noise-ratio (SNR), and channel type. Channel modeling work has well established that it is not only the direct path from sender to receiver that has impact but also the paths that reflect from fixed and moving obstacles in a given environment. Hence, channel type abstracts the notion of being in a certain neighborhood, rural area,

downtown region, or vehicular setting and leads to the corresponding channel characteristics of multipath, fading, and Doppler effects. A formal definition of the channel type is discussed in Section 3. We design and conduct experiments on both emulated and in-field channels. The experimental results demonstrate that the proposed rate adaptation framework increases the throughput by up to 245% over traditional SNR-based rate adaptation protocols and the default scheme in the Linux device driver, *Minstrel*. Also, we analyze the effect of context attributes and find the channel type to indeed be an important factor. However, the channel type is not directly known in real scenarios, especially when vehicular nodes frequently transition across different environments. As a result, we propose a multi-dimensional channel inference algorithm to predict the channel type. We also quantify the performance of this channel inference algorithm.

The main contributions of this paper are the following:

1. We identify the intrinsic problem of context-aware rate adaptation and present a context-aware link-level adaptation framework in vehicular environments. In particular, a classification-based scheme is used to establish the connection between the context information and the optimal link settings automatically. To the best of our knowledge, we are the first to propose a context-aware rate adaptation framework which can be generalized to incorporate different kinds of context information without modifying the core algorithm.
2. We study the importance of the channel type as context information. Since channel type is not easily obtainable in real scenarios, we propose a multi-dimensional channel inference algorithm that uses context information and link-layer performance to predict the channel type.
3. We implement our design with custom and off-the-shelf hardware platforms. Experiments both on emulated and in-field channels show that our proposed framework can significantly improve the performance of rate adaptation. For example, in certain scenarios our approach can achieve more than three times the throughput compared to traditional SNR-based adaptation schemes. The proposed channel inference algorithm can achieve a high accuracy of up to 94%. We also compare our algorithm with a preinstalled rate selection scheme in the current Linux wireless framework. In-field results indicate a significant throughput gain up to 245% in some scenarios.
4. In addition to analyzing the most predominant Linux-based rate selection protocol, *Minstrel*, we disentangle the relationship between the optimality

of the rate choice for a given scenario and Multi-Rate Retry (MRR). MRR is a protocol that the rate selection methods use when the initial transmission fails. In doing so, we show that while Minstrel benefits greatly from the use of MRR, our protocol has negligible differences when MRR is employed due to the high level of rate selection accuracy already employed by our scheme.

The rest of the paper is organized as follows. In Section 2, we describe the classification-based rate adaptation framework. Using this framework, we identify the importance of channel type as context information and propose the multi-dimensional channel inference algorithm in Section 3. Then, in Section 4, we evaluate the context-aware rate adaptation and channel inference on emulated channels to directly compare with prior schemes in repeatable conditions. In Section 5, we perform experiments in the field to test the proposed framework. Finally, we present related work in Section 6 and conclude in Section 7.

2 Classification-based context-aware rate adaptation

In this section, we first formulate the fundamental problem of rate adaptation studied in this paper. The problem considers choosing the transmission rate between two nodes in wireless transmission range of one another regardless of the end-to-end source and destination of a particular flow (i.e., we consider the link performance of each sender-receiver pair which could be applied to each link in the entire network). Then, we present our solution based on a classification-based framework.

2.1 Problem formulation

First, let us consider the case where a sender-receiver pair attempts to choose the transmission rate according to SNR information alone. To isolate this variable, we compare the throughput performance across six different transmission modes for a programmable hardware platform (WARP), including the combinations of three modulation orders (BPSK, QPSK, 16-QAM) and two packet sizes (100 B, 1000 B). The test is performed on a channel emulator with an experimental set-up as discussed in Section 4.1. The channel model used has a power-delay profile consisting of four taps with a relatively small 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. Figure 1 is a visual illustration of rate selection problem base on SNR information alone. In the figure, we can observe that one optimal transmission mode achieves the highest throughput for a given SNR region and would be said to be the *ideal* mode

choice in such region: 16QAM at the highest SNR and BPSK at the lowest SNR. Now consider that a wireless node that has knowledge of not just SNR but also potentially the vehicular velocity and the type of channel of the environment. Two different channel conditions are said to belong to the same channel model if the throughput of the various transmission modes are *sufficiently identical* (formally defined in Section 3.1) in both of those channel conditions. From this contextual tuple of information, imagine a three-dimensional space that has such performance curves across velocities and SNRs per channel type. In other words, the relationship between the throughput performance and SNR is subject to change when the channel model or velocity changes. Considering that multiple channel types and velocity values exist, the additional context information can help improve the throughput performance, but searching over the entire multi-dimensional space for the highest performing regions becomes intractable. In this adaptive system, the transmitter needs to select the desired modulation order, coding rate, and packet size based on the available knowledge. The objective of this adaptation is to maximize the throughput, G_{th} . Formally, the problem is posed as follows:

$$\max_{m \in M, r \in C, s \in S} G_{th} \text{ given } SNR, v, c \quad (1)$$

where M , C , and S are, respectively, the set of available modulation orders, coding rates, and packet sizes constituting different transmission modes in the system. Further, v , c , and SNR represent, respectively, the vehicular velocity, channel type, and SNR measured at the received side. The vehicular velocity refers to the relative velocity between the transmitter and receiver nodes.

We make the following remarks regarding the problem formulated in (1):

1. We use throughput as the indicator of link performance, although other metrics such as BER, PER, delay, or jitter could also be used.
2. The channel type c represents a main characterization of the propagation environment to which the current channel belongs. We assume that any channel can be classified into a countable number of channel types to assist in the identification of the optimal rate selection¹. Further, we assume that the transmitter is able to infer the current channel type in advance. In our case, the transmitter is able to infer the channel type with a multi-dimensional algorithm for channel inference introduced in Section 3.
3. There are various ways in which throughput could be calculated. For example, a traffic generator could report the application-level throughput or a device

¹Bounded only by system memory and complexity bounds.




Fig. 1 Performance of different modes in WARP. The throughput performance for different modulation schemes and packet sizes on a 4-tap channel with a velocity of 60 kmph

driver could count the total packets per second to calculate the MAC-level throughput. Often, the information available to the implementation governs the most appropriate notion of throughput. The exact throughput calculation in our experiment is introduced in Section 4.1.

4. To stay in focus, we consider only SNR, v , and c as the context information in this paper. Other environment-related parameters or the direct use of sensor data can be considered in future work.

2.2 Solution methodologies

We have formulated the link-level adaptation problem using an optimization framework. Traditional analytical solutions typically require concavity of the objective function. However, in our problem, it is difficult to find a closed form expression for the throughput. Further, the variables are continuously valued. Therefore, we use a classification-based algorithm to identify the optimal solution.

2.2.1 Decision-tree classification

Classification algorithms attempt to extract the relationship between the objective class labels and features given for classification such as various transmission modes and different environmental contexts, respectively. The decision tree [18, 19] is one of the most popular, entropy-based classifiers due to its ability to deduce a set of simple rules that are easily interpreted for a complex classification problem. In our framework, a decision-tree classifier is used to establish the connection between the context

information and the optimal link setting. If a combination of various sensing information such as the vehicular velocity, and channel quality can be classified into the appropriate category, the wireless transmitter can eliminate poor choices for a particular setting and quickly converge to the optimal choice.

In a decision-tree classifier, the relationship between classes and features is represented in the form of a tree structure. Classification starts at the root of the tree and moves to the leaf nodes until a class is encountered. Due to its low complexity, this scheme is easy to implement and adapts to real situations well without a high computational load. To derive the classification scheme used in our system, we use the C4.5 algorithm [20], a widely-used algorithm to generate decision trees. In the presence of large amounts of context data, the C4.5 algorithm chooses an attribute with the highest information gain to split the training data each time. This splitting of the training data can be based on either the information entropy gain or gain ratio, until each resulting subset contains training data of a single class. A subsequent pruning procedure then simplifies the tree to avoid overfitting. Thus, the classification scheme can streamline a large data set that contains both the wireless performance information and contextual information.

2.2.2 Application of the decision-tree classifier to link-level adaptation

In order to incorporate the decision-tree classification model into context-aware wireless systems, we use real

RF chains over actual RF channels (emulated and in-field). We use a combination of custom and off-the-shelf hardware platforms in our experimentation. The custom hardware platform we use is the Wireless Open-Access Research Platform (WARP) [21], while the off-the-shelf platform is a Gateworks 2358 board with three 802.11-based Ubiquiti radios (XR2, XR5, and XR9) that operate on different frequency bands. Each of the platform's radios has a physical layer based on the IEEE 802.11 standard. The Ubiquiti platform offers a built-in GPS, increased transmission power, physical layer coding, and fulfills the timing requirements of 802.11, while the WARP offers advanced programmability and observability across the network stack.

In one setting, we compare experimentally (both in the lab using a channel emulator for controlled, repeatable experimentation and in the field) the throughput of the various transmission modes for a given SNR, velocity, and channel type. We collect the set of SNR $\mathcal{S} = \{s_1, s_2, \dots, s_{N_S}\}$, velocity $\mathcal{V} = \{v_1, v_2, \dots, v_{N_V}\}$, channel type $\mathcal{C} = \{c_1, c_2, \dots, c_{N_C}\}$, and the achieved throughput values for all available transmission modes $\mathcal{M} = \{m_1, m_2, \dots, m_{N_{\text{mode}}}\}$. In the training phase, we measure the achievable throughput G_{c_i, v_i, s_i, m_i} using transmission mode m_i in the channel type c_i , with the velocity v_i and signal strength s_i . We identify the optimal transmission mode m_{best} with the highest throughput and use it along

with the context information for training. A training set contains $N_C \cdot N_V \cdot N_S$ data points, which include the optimal transmission mode m_{best} for a given channel type c_i , velocity v_i , and SNR s_i . Then, we use the C4.5 algorithm to derive the decision-tree classifier based on the training set as discussed in Section 2.2.1. An example of the decision-tree structure that we obtained with the C4.5 algorithm based on training data for the WARP platform is shown in Fig. 2 (see Section 4 for details about the training process). The derived decision-tree classifier is implemented as a look-up table in the context-aware system, so the context-aware system could make packet-by-packet rate decisions.

As soon as the system collects the current velocity, SNR, and channel type information, the classification-based adaptation scheme attempts to find the optimal transmission mode m_{best} with the highest throughput. The detailed process of the context-aware rate adaptation algorithm is shown in Algorithm 1. In particular, we consider the contextual information as the input (e.g., velocity and instantaneous SNR) and attempt to output the ideal transmission rate, m_{best} . To do so, we loop through the modulation and channel types, collecting each of the throughput values that correspond to that contextual tuple. Then, we apply the channel inference approach as described in Section 3.2. If a channel is recognized, we use the existing look-up structure to find and use the optimal transmission




Fig. 2 Decision tree for WARP Measurements. A decision-tree classifier is trained for rate adaptation based on WARP measurements

Algorithm 1 Context-Aware Rate Adaptation**Input:** Context Information $\{v, s\}$ **Output:** m_{best}

```

1: Initialize  $m = m_0, c = c_0$ 
2: Collect  $G_{c_0, v, s, m_0}$ 
3: Apply channel inference
4: if Convergent to an existing channel  $c_i$  then
5:   Set  $c = c_i$ 
6:   Predict  $m_{best}$  with decision tree classifier
7: else
8:   Add a new channel type  $c_{new}$ 
9:   Collect training data  $\mathcal{C}, \mathcal{V}, \mathcal{S}$  and  $G$  for  $c_{new}$  of
     transmission modes  $\mathcal{M}$ 
10:  Adapt the decision tree classifier
11: end if

```

rate. If the channel is not recognized, a new channel type is established and the decision structure begins training for that channel to learn the optimal transmission rate. The training phase could be completed beforehand to avoid introducing computational load to the inter-vehicle communication.

3 Multi-dimensional channel inference

Channel modeling is an area with a rich history, and several models have been proposed that offer a trade-off between complexity and accuracy. The International Telecommunication Union (ITU) has standardized channel models that are relevant for pedestrian and vehicular scenarios [22]. For example, the channel model Pedestrian A in the ITU Channel Model suite reflects a possible channel representation in a case where the relative velocity between the transmitter and the receiver is low, and the channel has a relatively small delay spread [22].

However, the channel type information is difficult to directly obtain in practice. In this section, we first motivate the importance of channel type. Then, we introduce a multi-dimensional channel inference algorithm that classifies an unknown channel into one of the channel models in the training set or detects if the channel type has yet to be encountered and should be added as a new channel type.

3.1 Importance of channel type

First, we assume a multi-dimensional pattern that relates the throughput performance of a given set of rate choices over a number of contextual parameters. The dimensions are composed of the contextual attributes that the link performance depends. Then, the modulation and coding scheme has a certain throughput per contextual tuple. This multi-dimensional link performance-based shape that forms can be compared against other shapes from

other contextual tuples. Two different measurement sets from two different channel conditions could be said to belong to the same channel type if their effective performance (as measured by the throughput) of the available transmission modes, exhibit similar behavior restricted by a pre-defined threshold on the similarity metric for all values of the various contextual attributes. Standardization bodies (e.g., ITU) have produced representative power-delay profiles that are representative of different environments. However, two different power-delay profiles could form a similar multi-dimensional shape and be similar for our purposes. Hence, we define a similarity index that allows the ability to distinguish the multi-dimensional link performance-based shape that forms from environment to environment. We elaborate the concept in Section 3.2. The experimental validations are in Section 4.6 and Section 5.1.

The decision tree in Fig. 2 indicates the relative importance of the attributes, where the relative importance of an attribute (signified by branches) is denoted by its relative distance from the root of the tree. For example, velocity is more important in making decisions in Channel D than in the other channel types. The channel type turns out to be more informative than other attributes in the tree-building process and is selected as the first level of the tree which is the closest to the root. We thoroughly quantify the importance of different context information via experimentation in Section 4.2.

3.2 Multi-dimensional classification for channel inference

The performance of wireless systems has long been known to depend on the channel conditions. In particular and as it specifically relates to this work, each transmission mode's achievable throughput depends on several different factors including the power-delay profile of a given environment. These factors go beyond the traditional indicators of signal quality such as the signal strength and the noise level. Specifically, we pointed out in a previous work that coherence time affects the SNR-rate thresholds [17]. Hence, since velocity directly affects Doppler and coherence time, velocity is also performance related since it affects the maximum Doppler bandwidth. From the perspective of rate selection, the important information from channel inference or channel modeling is the target rate that can give the highest throughput in the current channel. In other words, where the underlying channel parameters are sufficiently similar to allow two channels to have the same multi-dimensional pattern of the available transmission modes, these two channels can be treated as the same. This leads to the proposed channel inference technique based on link-level performance.

Figure 3a presents an example of the best transmission modes for various velocity and SNR values and different channel types. In this example, there are 8 available




Fig. 3 Throughput performance comparison of transmission modes. **a** The transmission mode with highest throughput in the context space. Dots of the same color represent that the context share the common optimal transmission mode. **b** The variance of the throughput performance with the SNR and velocity for 4 known channels and an unknown channel

transmission modes ($N_{\text{mode}} = 8$). The granularity of the SNR is 5 dBm and that of the velocity is 15 kmph. For the representation of the throughput variance, we visualize the optimal transmission mode with the highest throughput for the context. In Fig. 3a, different colors represent different optimal transmission modes for the context, and dots of the same color represent the common optimal transmission mode. The three plots depict the performance data set from 3 different channels. As can be seen in Fig. 3a, Channel 1 and Channel 2 are similar in terms of common optimal transmission modes, while Channel 3 is dramatically different. We now propose a more formal approach to channel type inference based on the multi-dimensional representation of the throughput of all transmission modes on a given channel, which more efficiently utilizes in-field measurements.

The sets of channel types, velocities, and signal strength values in the training set are represented by $\mathcal{C} = \{c_1, c_2, \dots, c_{N_C}\}$, $\mathcal{V} = \{v_1, v_2, \dots, v_{N_V}\}$, and $\mathcal{S} = \{s_1, s_2, \dots, s_{N_S}\}$, respectively. Similarly, the transmission modes available in a system can be represented by $\mathcal{M} = \{m_1, m_2, \dots, m_{N_{\text{mode}}}\}$. The achievable throughput $G_{c,v,s,m}$ is measured using transmission mode m in the channel model c , with the velocity v and signal strength s . For each channel model c , the $N_V \cdot N_S \cdot N_{\text{mode}}$ throughput values constitute the performance data d_c . A complete training set $\mathcal{R} = \{d_1, d_2, \dots, d_C\}$ contains $N_C \cdot N_V \cdot N_S \cdot N_{\text{mode}}$ throughput values, which describes the system performance in these N_C channels.

With the training set $\mathcal{R}$, we now use N throughput values (where $N \geq 2$) measured in the field on an unknown channel type and its related context information to classify

the unknown channel into one of the known channels in $\mathcal{C}$. Each record i ($i = 1, \dots, N$), $\hat{G}_i$ denotes the measured throughput using mode $\hat{m}_i$, and $\hat{v}_i, \hat{s}_i$ represent the corresponding velocity and signal strength, respectively. Then, we search the training set and find the throughput of all the channel models with the same context information and transmission mode. In other words, we look for $G_{c,\hat{v}_i,\hat{s}_i,\hat{m}_i}$ of all the N_C channel models. It is possible that $\hat{V}_i$ is not in $\mathcal{V}$ (the same holds for $\hat{s}_i$). In this case, bi-linear interpolation is used to calculate the throughput $\tilde{G}_{c,\hat{v}_i,\hat{s}_i,\hat{m}_i}$ based on the throughput of the nearest velocity and signal strength values.

As an example of a multi-dimensional space, if the rate performance of the contextual parameters were plotted in a 3-D space where the x , y , and z axis are signal strength, velocity, and throughput, respectively, the training set for a particular channel (d_c) is a family of meshes, where each mesh represents the throughput of a certain transmission mode with respect to signal strength and velocity. With interpolation, the meshes can be filled and turned into surfaces. By labeling all the $\tilde{G}_{c,\hat{v}_i,\hat{s}_i,\hat{m}_i}$, we can plot the data as shown in Fig. 3b corresponding to N_C known channel types, indicating the throughput variation trend on the order of i from 1 to N . Similarly, we can also plot the throughput for the unknown testing channel. Let TSV_c denote the training set vector as:

$$(\hat{s}_{i+1} - \hat{s}_i) I_x + (\hat{v}_{i+1} - \hat{v}_i) I_y + (\tilde{G}_{c,\hat{v}_{i+1},\hat{s}_{i+1},\hat{m}_{i+1}} - \tilde{G}_{c,\hat{v}_i,\hat{s}_i,\hat{m}_i}) I_z \quad (2)$$

where I_x, I_y, I_z are unit vectors in the x, y , and z directions. Likewise for the unknown channel type c , the measured

vector, $\overline{MV}_c$, can be represented by:

$$(\hat{s}_{i+1} - \hat{s}_i) I_x + (\hat{v}_{i+1} - \hat{v}_i) I_y + \left(\hat{G}_{c, \hat{v}_{i+1}, \hat{s}_{i+1}, \hat{m}_{i+1}} - \hat{G}_{c, \hat{v}_i, \hat{s}_i, \hat{m}_i} \right) I_z \quad (3)$$

Let $\gamma_{c,i}$ denote the angle between $\overline{TSV}_c$ and $\overline{MV}_c$ for channel type c . The similarity index γ_c is defined as

$$\gamma_c = \frac{\sum_{i=1}^{N-1} \gamma_{c,i}}{N-1}. \quad (4)$$

The channel inference algorithm computes the channel model $c_{\text{fit}} = c_{j_0}$ that most likely has the same performance as the unknown channel. We can formally represent this as $j_0 = \arg_j \min \gamma_{c,j}$. In this way, we find one existing channel type that has most similar performance compared to the unknown channel in terms of throughput variation in different context. Subsequently, the decision structure for channel model c_{j_0} can be used for the unknown channel. If the similarity index γ_c is sufficiently large, then the behavior is not guaranteed to converge to an existing channel type and should form a new channel type which has not been previously encountered. One might seek to form an exact threshold for γ_c (which could be done in future work). However, we take a more practical approach in leveraging γ_c to perceive if over a number of in-field measurements, the channel type is converging to the behavior of a previously observed channel or not. The formulation of similarity index γ_c reflects both reasonable complexity and affordable convergence. Adding more terms or constraints, such as the absolute throughput value into the similarity index is feasible depending on required precision of the application scenarios. To evaluate whether we are correct in inferring the channel type to a previously observed channel, we use an accuracy metric as discussed in Section 4.6. In Section 5.2, we show experimental examples of both converging and diverging from previously encountered channels, which forms a new channel type.

Figure 3b is a representative 3-D plot for $N_C = 4$ and $N = 5$. In the figure, the surfaces are not shown because they would block the inner points. We can see that the throughput of the channel model A follows the trend of that of the unknown channel, which results in the minimum γ_c compared with other channel models. The multi-dimensional channel inference approach exploits the channel characteristic information lying in the throughput measurements. When N grows, more features are provided by $\{\hat{G}_i\}$, and the resulting inference is more accurate. Thus, there is a natural trade-off between computation complexity and estimation accuracy, which is quantified in Fig. 8a and discussed in Section 4.6. The context-aware system could collect adequate data samples in a limited time and make a timely channel-inference decision.

4 Experimental results and discussion on emulated channels

We now evaluate our context-aware link-level adaptation by performing a wide range of in-lab tests using both WARP and Ubiquiti platforms. Using both platforms, we show that our algorithm (i) accurately selects the best transmission mode, (ii) precisely infers channel characteristics, and (iii) outperforms existing rate adaptation mechanisms. We design a testbed implementation over both in-field and emulated channels. We focus on emulated channels so that we could replay them to test the different algorithms in a fair manner.

4.1 Emulator-based experimental settings

As an initial analysis of the proposed adaptation which selects data rates according to various pieces of context information, we use a channel emulator for controlled, repeatable channels to quantify the gains we could expect over existing schemes. Hence, in our experiments, we connect the transmitter and the receiver via an Azimuth ACE-MX [23]. The Azimuth ACE-MX has a set of predefined and widely-used ITU pedestrian and vehicular channel models. Table 2 describes the power-delay profile for ITU channel models implemented on the emulator. We now briefly describe each of the channels to relate them to actual environments. Channel A represents a channel that lacks a strong multipath component (e.g., an environment that lacks tall buildings or structures) and corresponds to ITU Pedestrian A. Channel B represents an environment where there is some multipath and corresponds to ITU Pedestrian B. Channel C has multiple strong multipath reflections (e.g., a downtown environment with a line-of-sight component) and represents ITU Vehicular A. Channel D actually has a multi-path component that is stronger than the line-of-sight component and represents ITU Vehicular B. The custom channel was created from randomly generated power and delay values and will be used for experimentation purposes later in the paper.

Figure 4 shows the set-up of the experiment on emulated channels. A PC controls the emulator (over Ethernet) and allows configuration of the channel characteristics such as model type, path loss, Doppler, and input or output attenuation. The Ubiquiti radios on the Gateworks boards match the physical layer coding rates and modulation orders of IEEE 802.11a/g whereas WARP lacks physical layer coding, currently allowing only uncoded 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 emulated channels that have predefined attenuation levels and velocities, as shown in Table 1.




Fig. 4 Set-up of the emulator-based experiment. This figure shows the set-up of the emulated environment for experiments, including the equipments and the connections between equipments

On the WARP platform, the traffic consists of packets with pre-determined content. Therefore, the BER and PER information is available at the receiver side, and the throughput, G_{th} , can be calculated as:

$$G_{th} = (1 - PER) * R_{th} * \frac{l_{payload}}{l_{packet}} \quad (5)$$

where R_{th} is the physical data layer rate, and $l_{payload}$ and l_{packet} are the lengths of the payload and the entire packet, respectively. On the Ubiquiti platform, we use *iperf* to generate traffic. Thus, the original payload of a packet in error is unknown at the receiver, precluding BER calculation. However, the time duration of a packet transmission can be calculated based on the specification in the IEEE 802.11 standard and data rate. The probability of successful transmissions can be obtained via investigating the response from the receiver (i.e., the reception of an ACK packet). Therefore, the throughput, G_{th} , can be calculated as:

$$G_{th} = \frac{P * l_{payload}}{t_{SIFS} + t_{pre} + l_{packet}/R_{th}} \quad (6)$$

where P is the probability of a successful transmission, t_{SIFS} and t_{pre} are, respectively, the duration of the short interframe space (SIFS) and preamble. The latter two terms are specified in the 802.11 standard [24].

4.2 Performance of context-aware rate adaptation

For training the decision tree, the five velocities and eight attenuations shown in Table 1 are used, composing a set of 40 measurement points per channel type in Table 2. Each measurement is approximately 1 minute in duration to get an accurate throughput value. The well-trained decision

tree functions as a look-up table in a real-time implementation for the practical application without much computational load. Using this training, we first evaluate the performance improvement of using our classification-based scheme on randomized channel scenarios. In particular, we randomize the velocity and attenuation attributes so as to create different values from our training set to form a testing set on each channel type. We then compare the throughput achieved by the rate decision given by the classification against both the decision of an SNR-based scheme and the decision which achieves maximum throughput found via exhaustive search. For the SNR-based scheme, the thresholds are defined by the highest performing rate on a particular SNR in a static topology. Results in Table 3 include the performance of the proposed rate adaptation framework. With the trained decision structure, the proposed framework outperforms the SNR-based method by 40.2% on WARP and 60.8% on the Ubiquiti platform, respectively.

The decision-tree structure forms the primary decision engine for context-aware link adaptation. The decision

Table 1 Transmission modes and channel emulator parameters

Parameter	Values
Channel models	Ch. A, Ch. B, Ch. C, Ch. D
Attenuations (dB)	0, 6, 12, 18, 24, 30, 36, 42
Velocities (kmph)	0, 30, 60, 90, 120
Ubiquiti 802.11a/g rates (Mbps)	6, 9, 11, 12, 18, 24, 36, 48, 54
Packet size (B)	100, 1000
WARP modulation schemes	BPSK, QPSK, 16QAM

Table 2 Power delay profiles of the channel models used on the emulator

Model	Channel A		Channel B		Channel C		Channel D		Custom Channel	
Tap	Delay (μ s)	Power (dB)	Delay (μ s)	Power (dB)	Delay (μ s)	Power (dB)	Delay (μ s)	Power (dB)	Delay (μ s)	Power (dB)
1	0	0	0	0	0	0	0	− 2.5	0	0
2	.11	− 9.7	.2	− .9	.31	− 1.0	.30	0	.15	− 4.0
3	.19	− 19.2	.8	− 4.9	.71	− 9.0	8.9	− 12.8	.38	− 8.7
4	.41	− 22.8	1.2	− 8.0	1.09	− 10.0	12.9	− 0.0	.85	− 12.9
5	-	-	2.3	− 7.8	1.73	− 15.0	17.1	− 25.2	1.40	− 15.3
6	-	-	3.7	− 23.9	2.51	− 20.0	20	− 16.0	-	-

tree in Fig. 2 implies that the channel type is more important than other attributes in making rate decisions. To validate this observation, we solve the optimization problem of rate adaptation given only knowledge of (i) SNR and velocity, (ii) SNR and channel type, or (iii) velocity and channel type. This experiment is equivalent to the case in a real system where one attribute cannot be obtained or the system receives erroneous sensor information for that attribute, which cannot be used. The experiment has been done on both the WARP and Ubiquiti platforms. Table 3 lists the classification performance with one missing attribute along with the performance when all attributes are present. The experiments are performed on a channel emulator for repeatability and control. Hence, the gap is the percentage of throughput difference from the maximum achievable, which is computed via exhaustive search after channel conditions are repeated for all modes.

As shown in Table 3 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 a more severe performance loss than the other two attributes. For example, on the WARP platform, a 32% reduction in throughput 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 the greatest contributing factor of the three attributes in making rate decisions.

To verify the performance algorithm in a real-time scenario, we implement the proposed classification-based rate adaptation protocol under the Linux wireless framework *mac80211* [25]. In the original framework, *Minstrel* is the rate selection algorithm that is enabled by default. It is the most powerful algorithm compared with other available options in the Linux framework [26]. *Minstrel* will attempt to select the rate with highest historical throughput at the point when the packet is transmitted. In addition, for hardware that supports multi-rate retry (MRR), *Minstrel* will choose up to two other candidate rates in case the first choice fails. We load our trained decision-structure-based rate control module onto the Ubiquiti platform and compare its performance in the four emulated channels with *Minstrel* with and without MRR.

We simulate a typical process in a vehicular scenario by altering the parameters of the channel emulator. The parameters that fed into the emulator are listed in Table 4. We use a script to control the emulator in each round during the experiment to ensure that every time the system is experiencing the same channel variance and experiment duration. Table 5 shows the throughput comparison between *Minstrel* and the proposed classification-based algorithm. Channel A has fewer number of taps and lower tap delay compared with the other three, which means that it represents a simpler channel with almost no obstacles between the transmitter and the receiver. In this case, the rate selection logic needs not be as adaptive in order to achieve high throughput. Hence, the two algorithms perform similarly. This result can be attributed to the lack of a strong multipath component that experiences far less

Table 3 Classification-based context-aware performance: (A)accuracy is percentage of time classification choosing ideal rate, (I)mprovement is its throughput gain over a SNR-only method, and (G)ap is its throughput percentage from the maximum achievable

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

Table 4 Transition of attributes in the emulated channels

Index	1	2	3	4	5	6
Power attenuation (db)	10	18	20	30	11	0
Velocity (kmph)	0	10	15	30	18	0
Duration (s)	20	10	10	10	10	10

Table 5 Performance comparison in the emulated channels

Algorithms		Throughput (Mbps)			
		Channel A	Channel B	Channel C	Channel D
With MR	Minstrel	20.83	4.64	7.65	5.56
	Context-aware	20.39	5.66	8.83	6.15
Without MR	Minstrel	20.18	2.42	2.90	3.02
	Context-aware	20.14	6.05	7.59	6.03

channel fluctuations in a mobile context than the remaining three channels. With the use of Channels B through D, the performance of our rate adaptation using our decision structure significantly improves. For now, we focus on the accuracy of rate choices made by each scheme. Later, we will explore the effect of MRR on the performance of rate selection in Section 4.4

To more easily observe the aforementioned behavior, we have plotted the rate selection process. Figure 5 a and b show the variation of signal strength and throughput reported by hardware across time. Figure 5 c and d plots the rate selection behavior and corresponding throughput performance when the channel condition fluctuates. The Channel A graph reveals that both algorithms are performing well and constantly choosing the best rate when the channel condition is good. The throughputs shown in Table 5 indicate that there is nearly identical performance across both algorithms with and without MRR (approximately 20 Mbps). However, the Channel C graph indicates that the decision structure is able to accurately choose the best rate as the first rate choice whereas *Minstrel* cannot quickly adapt to the channel variation and has to try a large number of rates. The advantage of our algorithm is twofold: it achieves a higher throughput, and it does not rely on MRR. Figure 5 e and f describes the rate selection progress of the two algorithms in Channel C. In Table 5, we observe that in both Channels B and D, similar effects hold as Channel C. Namely, *Minstrel* has a dramatic performance degradation without MRR, whereas the use of MRR has almost no influence on our context-aware rate selection.

4.3 Effect of training set size

In this subsection, we evaluate the effect of training set size on the proposed classification-based context-aware rate adaptation. We use the same training data and testing data as we used in Section 4.2. For each channel type, we have 40 data points in total and evaluate the effect of the training set size by varying the number of data points in the training set. In Fig. 6a and b, we show the effect of the training set size on the accuracy of rate prediction and the gap from the maximum achievable throughput by gradually increasing the number of data points in the training set. We can observe that:

- As the size of the training set increases, the accuracy of the rate prediction first increases and then decreases. Before reaching the best performance, enlarging the training set size can provide more information for learning the embedded pattern. Increasing the training set also includes wrong information from sensor errors, which can make the classifier confused in selecting the right transmission mode for a given context. Also, the overfitting problem exists with an increasing training set size [27].
- The same size of the training set has different performance on different channel types. The main characteristics (e.g., the multipath effect and resulting delay spreads) are different in the four channel models. Thus, to achieve the same prediction performance, the learning process needs differing amounts of training data to extract the embedded patterns of channel complexity. Also, the training data collection process might introduce incorrect information, which is implemented independently for the four channel types. Hence, the four training sets provide differing amounts of useful information for training the decision tree.

4.4 Effect of MRR

Multi-rate retry (MRR) relies on functionality that is available in only some wireless chipsets. The performance of MRR-based rate selection algorithms decays when MRR is not enabled [28]. Also MRR is not available on all the wireless devices and hardware, for example, BCM43XX, Rt2x00, RTL8180 (Realtek)/RTL8185, RTL8187, and RT2800 (RaLink). There are works investigating the performance of *Minstrel* [29]. However, to the best of our knowledge, no work has isolated how MRR affects *Minstrel*. In this section, we evaluate the impact of MRR on *Minstrel* and show that the proposed classification-based context-aware rate adaptation does not rely on MRR.

In Fig 7, without MRR, *Minstrel* suffers severe performance degradation because *Minstrel* is a trial-based algorithm and relies on the multiple rate choices to ensure successful transmission. *Minstrel* is sensitive to changes in the channel condition and keeps trying a broad range





Fig. 5 Signal strength (SS), throughput (TH), and rate selection comparison on emulated channels. Rate selection behavior varies with the signal strength and the multi-rate setting. **c, d** The top is Minstrel, and the bottom is context-aware. **e, f** The top is with MR, and the bottom is without MR

of rates. In most cases, Minstrel's first rate choice is not correct. The classification-based context-aware scheme selects the best rate with the consideration of context information, which leads to a more precise result: the first choice is the correct choice in most cases and thus, not having MRR has little impact on performance.


As discussed in Section 4.2, Table 5 shows the throughput comparison between Minstrel and the context-aware framework. By comparing the throughput performance with and without MRR, we can see that the performance of Minstrel relies heavily on MRR, since its rate choices are based on historical data. When experiencing a channel that fluctuates dramatically, its first rate choice for the

next packet is already obsolete. On hardware that does not support MRR, its throughput will rapidly decay in such scenarios. The context-aware rate adaptation makes the rate selection based on the real-time environmental condition. Therefore, our context-based scheme can quickly adapt to the environmental variations.

Table 6 shows the performance comparison between the Minstrel and the context-aware framework with and without MRR in in-field experiments. Minstrel relies more on MRR and its performance decreases significantly without MRR. The context-aware framework has a higher percentage of selecting the most appropriate rate upon the first rate selection without the reliance on MRR.



(a) Accuracy of rate prediction.



(b) Gap from the maximum achievable throughput.


Fig. 6 Effect of training set size on rate prediction performance. **a** The impact of the training set size on prediction accuracy. **b** the impact of the training set size on throughput performance

4.5 Multiple frequency bands adaptation


To further demonstrate the utility of context information in next-generation multiband cognitive networks, we discuss the following scenario. Consider a particular node that 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 transmission modes in the new frequency band based on the decision tree of the prior band? Our results indicate that, indeed, significant

gains are possible in several cases. Table 7 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 Ubiquiti radios on 900 MHz, 2.4 GHz, and 5.8 GHz, corresponding with XR9, XR2, and XR5, respectively. The experiments are performed with randomized velocities and attenuations of the channel types described above.

From Table 7, we find classification across different bands has higher throughput in most scenarios as compared to the throughput when making rate decisions based on SNR alone. For instance, training the tree at



(a) MRR effect on Minstrel



(b) MRR effect on Context-aware framework

Fig. 7 Effect of MRR. The comparison of the effect of MRR on the throughput performance of Minstrel and the proposed context-aware framework

Table 6 In-field performance comparison

Place	MRR	Throughput (Mbps)		
		Minstrel	Context-aware	Improvement
Bishop	With	5.53	5.77	4.34%
	Without	4.74	5.25	10.75%
Park	With	5.00	5.54	10.80%
	Without	3.35	4.68	39.70%

900 MHz and using this training to make rate adaptation decisions at 2.4 GHz results in a 52% improvement over SNR-based schemes. The channels of different frequency bands share similar behavior with respect to transmission modes, throughput, and context information. Thus, classification across different bands can achieve large gains versus using the SNR-based method across frequency bands. In Table 7, when we use the testing set of one frequency band and the training set of the same frequency band, the gap from the optimal performance is smaller than the gap when training is performed at other frequency bands. The improvement over SNR-based methods does not show a similar trend because when we use the training set from a different frequency band, the SNR-based selection is from the training frequency band. Otherwise, the SNR-based selection could have a poorer performance than the selection of the same frequency band as that of the testing set, which would make the improvements larger than that from training with the set of the same frequency band.

4.6 Performance of channel inference

We implement two different types of tests to validate our channel inference algorithm. In the first test, we select N sets of (c_i, v_i, s_i, m_i) where $c_i = c_0, c_0 \in \mathcal{C}$ for all $i = 1, \dots, N$, $m_i \in \mathcal{M}$ but $v_i \notin \mathcal{V}$, $s_i \notin \mathcal{S}$. We measure the throughput $G_i = G_{c_i, v_i, s_i, m_i}$, and using $\{G_i\}$ and the corresponding $\{v_i\}$, $\{s_i\}$, $\{m_i\}$ as the input to the algorithm, we check if the inferred value equals c_0 .

In the second test, we generate a new channel $c_{\text{new}} \notin \mathcal{C}$ with a randomized power-delay profile and infer which channel type it resembles in terms of link-level performance. To do so, we select N sets of $(c_{\text{new}}, s_i, v_i, m_i)$, measure the resulting throughput $\{\eta_i\}$, and check if the

algorithm can choose the model $c_{\text{fit}} \in \mathcal{C}$ that is most similar to c_{new} . As stated in Section 2.1, we seek to select the transmission mode that maximizes throughput. Therefore, the definition of c_{fit} is the channel model that, once fed into the decision tree along with the context information, outputs the transmission mode $m_{\text{best}} \in \mathcal{M}$ that achieves the maximum throughput in c_{new} . Next, we will introduce the proposed method to determine c_{fit} and then discuss the evaluation results.

To determine c_{fit} , we first find r_{best} by measuring the throughput of all the $m_k \in \mathcal{M}$ in the unknown channel with the same context information. We can more formally state this as:

$$m_{\text{best}} = \arg \max_{M_k} G_{c_{\text{new}}, v_i, s_i, m_k}, k = 1, \dots, N_{\text{mode}}. \quad (7)$$

Then, we consider N_C scenarios such that $c_{\text{fit}} = c_j, j = 1, \dots, N_C$. For each scenario, we input (c_j, v_i, s_i) to the decision tree and check if the output transmission mode m_j matches with m_{best} . We can repeat the procedure by choosing different v and s and find the c_j which has the most matching instances. In fact, the method stated above can also be recognized as one example of a channel inference algorithm. We will call it the *exhaustive search* algorithm and compare it with the performance of the multi-dimensional channel inference algorithm.

We use the same training set and the decision tree in Section 4.5 where four channel models are involved (i.e., $C = 4$ in both evaluations). In the first type of evaluation, for the exhaustive search algorithm, we choose $N = 16$ records to investigate the effectiveness of the algorithm. The inference accuracy is 40%, on average. Since our channel inference algorithm can work with N records (where $N \geq 2$), and it has better performance with more records, we have done an iterative evaluation while increasing the record number. Table 8 shows the result of the evaluation. The superiority of the proposed multi-dimensional channel inference algorithm is clearly evident from the accuracy of 74.8% with only two testing points and over 94% accuracy with just seven testing points. As a practical note, a lower number of testing points translates to less training data required when entering an unknown region in the field to perform channel inference.

In the second evaluation, among 24 results that are given by the exhaustive search algorithm, 66.7% select channel

Table 7 Training across frequency bands: (A)ccuracy is percentage of time classification chooses ideal rate, (I)mprovement is its throughput gain over a SNR-based method, and (G)ap is its throughput percentage from the maximum achievable

Band	Training, 900 MHz			Training, 2.4 GHz			Training, 5.8 GHz		
	A	I	G	A	I	G	A	I	G
Testing, 900 MHz	73.1%	17.1%	16.6%	47.1%	— 2.4%	28.2%	22.1%	— 20.2%	44.9%
Testing, 2.4 GHz	45.2%	52.7%	67.2%	51.9%	60.8%	66.0%	26.9%	46.6%	70.2%
Testing: 5.8 GHz	20.9%	55.1%	19.6%	16.3%	67.1%	22.8%	60.5%	99.8%	6.75%

Table 8 Geometry-based channel inference in known channels

N	2	3	4	5	6	7
Accuracy (%)	74.8	81.9	87.5	90.0	91.7	94.3

model C in $\mathcal{C}$ as c_{fit} . For our channel inference algorithm, we have done the iterative evaluation as shown in Fig. 8a. The confidence of selecting channel model C is the highest out of all channels with just two records and has a monotonically increasing confidence level with additional records. On the other hand, the other channel models have low confidence for seven records.

Figure 8b shows that more training data will lead to more accurate channel inference. A large training set will provide more detail about the throughput relationships across different SNR and velocity combinations. Thus, channel inference can reduce the effect of errors in angular calculation introduced by interpolation.

5 In-field experimental results and discussion

In this section, we evaluate the rate adaptation framework by performing experiments on in-field wireless channels.


5.1 Channel inference evaluation with in-field trials

We perform an in-field evaluation of the channel inference algorithm and the classification-based context-aware scheme using data collected from experiments at two different fixed locations. The context-aware algorithm is trained as in Section 4.2. At each location, we vary the distance between the transmitter and receiver to alter the signal attenuation levels at the receiver, which will be reflected in the receiver's SNR. The antennas are either on the top of cars (2 m in height) or on tops of buildings or

parking garages as pictured in Fig. 9, where the four antennas (each approximately 5 dBi) are spaced approximately a third of a meter apart. For each attenuation, we measure the SNR and achievable throughput in each transmission mode. With the complete set of throughput data, we are able to monitor and evaluate the performance of channel inference algorithm. We pick several sets of SNR values and their corresponding throughput and apply our channel inference algorithm. Since we have the throughput of all the transmission modes and can find the best performing mode via exhaustive search, we can compare it with the output of the rate adaptation framework.

Our first location is a residential apartment complex. Two WARP nodes operating at 2.4 GHz are used to measure the achievable throughput across the established link. One node sends fully backlogged packets to the other node, which logs SNR and throughput. Figure 10a shows the channel inference performance. We can see the confidence of the classification as Channel A quickly converges to 1, while others remain at a much lower level. With this channel inference result as the input of channel type to our decision tree trained as in Section 4.2, the throughput improvement over the SNR-based mechanism is 24.15%.

The second location is in downtown Dallas. This time we use a 5.8-GHz radio (XR5) on the Gateworks board to see if our scheme can operate in a different frequency band. The result of channel inference is shown in Fig. 10b and indicates that Channel type A is closer to the channel in downtown than other models. Although the confidence is not as high as in the residential area, the inference of Channel A as the channel type into the decision tree achieves an even higher throughput improvement of 245%. The reason for this large improvement is that the



(a) Change of selection confidence with available samples. (b) Effect of training set size on channel inference accuracy .

Fig. 8 Effect of data volume on channel inference performance. **a** The impact of number of available samples on channel-inference confidence. **b** The effect of training set size on channel inference accuracy



Fig. 9 Gateworks board with four Ubiquiti radios and antennas at 900 MHz, 2.4 GHz, and 5.8 GHz on top of a parking garage on campus. This figure shows the in-field experiment set-up

SNR-based algorithm cannot perform as well as in the residential area, showing that context-aware scheme tends to do better as the complexity of the environment increases.

Figure 10c shows the channel inference performance at the same location but at a 900-MHz carrier frequency (XR9). The confidence indices of all the channel models are similar and below a threshold of 40%. This result indicates that the current channel should be considered as a separate, newly encountered channel type needing training for the decision tree and channel inference algorithm.

5.2 In-field mobility: context-aware link-level adaptation

We now use the same verification method as in Section 4.1 in the field to compare our algorithm with *Minstrel*. Our results indicate that the proposed algorithm can outperform *Minstrel* in the outdoor environment similar to the performance over emulated channels.

We use the Gateworks 2358 with our algorithm implemented in the *mac80211* module. The experiment is performed on a 2.412-GHz channel and *iperf* is used for traffic generation. We place an omni-directional, 2.4-GHz antenna on top of two cars. During the experiment, one vehicle's location is fixed and the other vehicle traverses a fixed route. We attempt to maintain a constant speed across different experiment with respect to the traffic laws (approximately 20 mph).

The context-aware rate adaptation learns about the environment during the transmission the first time the node enters an area. According to our protocol a pre-knowledge (training set) of the channel should be available. Here, we assume that we have the training set around the regions where we perform the experiment. Building up such historical training knowledge is possible if the transmitter saves the context information of a region and uses the trained decision structure the next time that it enters the same region.

As shown in Fig. 11, one experiment location is on Bishop Boulevard, the main street of our university campus. The maximum distance of the two nodes is 180 m (based on GPS measurements at the two nodes). There are




Fig. 10 Confidence of inference at different locations. The confidence of selecting a known channel for in-field channels is calculated with the channel inference algorithm




Fig. 11 Experiment locations. A screen shot of the Google map shows an experiment location

some trees on the grassy median inside the route, but the two nodes predominantly have line-of-sight links. There is light vehicular traffic but many pedestrians. The other location is a public park surrounded by a parking lot. The maximum distance of the two nodes is 190 m. Due to many other parked vehicles between the two nodes, there is a greater degree of multipath and path loss as compared to on Bishop Boulevard. The vehicular traffic is also heavier since the park is off campus.

Table 9 shows the performance comparison between the two algorithms. Similar to what we have seen on the emulator (Section 4.1), the advantage of the context-aware framework over *Minstrel* becomes apparent when the channel conditions worsen, mainly in the form of more severe channel fluctuations.

From Fig. 12, it appears that our algorithm can fill the “rate selection gap” compared with *Minstrel*. In fact, as shown in Fig. 12 a and c, our algorithm can quickly adapt to channel fluctuations. For example, at the tenth second, a lower rate choice is made which achieves higher throughput. Such downward transitions of the channel

are where our context-aware approach achieves the most gains over a historical algorithm such as *Minstrel*.

In addition to the results shown, we have captured data in several other scenarios and they all demonstrate similar improvements in performance.

6 Related work

6.1 Context-aware wireless systems

In the computing area, the notion of using context information to improve performance is well-established [30, 31]. However, applying context awareness to wireless systems is still a developing area of study [6, 32]. As systems become more complex, such as multi-user MIMO, performance of the protocol will heavily rely on knowledge of the environment of propagation [33]. In particular, works have considered the file size and whether to use cellular or IEEE 802.11 to reduce energy [34], application requirements, spectral activity and bandwidth available [35] or user directionality and speed and knowledge of indoor/outdoor [8]. Furthermore, relative distances and speeds have been considered in [36, 37] for the purposes of link adaptation. Other work has shown that by utilizing user location and movement, channel state information is predictable enough to improve power control [38]. Vehicular networks have historically used several channel metrics such as fading, delay spread, and Doppler spread to characterize the channel [39]. In contrast, we consider how to infer wireless channels based on link-level performance for use with classification algorithms based on SNR, velocity, and channel type to choose the most appropriate rate for the particular setting.

6.2 Classification algorithms

Decision-tree classification [18] and its variants [19] are used in many diverse areas such as sensor networks [40], traffic classification [10], radar signal classification, and speech recognition. In recent years, machine learning techniques have improved the classification of modulation schemes in complex environments leading to increased spectral efficiency [41]. In contrast to prior work using decision-tree classification, we use it as a means to classify wireless performance of different environmental contexts and transmission modes to lead to better decisions for rate adaptation protocols.

6.3 Rate adaptation algorithms

Adjusting modulation and/or coding rate to increase the performance of a link has used various mechanisms for adapting the link from packet statistics at the transmitter [12, 13] to direct channel quality measurement combined with feedback [14]. The former method has been shown to suffer from a confusion of channel and congestion loss types, whereas the latter can suffer from improper and/or excessive training [17]. Other work has studied rate

Table 9 In-field performance comparison

Place	MR	Throughput (Mbps)		
		Minstrel	Context-aware	Improvement
Bishop	Enabled	5.53	5.77	4.34%
	Disabled	4.74	5.25	10.75%
Park	Enabled	5.00	5.54	10.80%
	Disabled	3.35	4.68	39.70%




Fig. 12 In-field signal strength, throughput, and rate selection comparison. The rate selection performance is analyzed in in-field experiments

adaptation for MIMO systems using both window-based methods [42, 43] and SNR directly [44]. Other work has studied the accuracy of throughput estimation using rate adaptation with Minstrel and multi-rate retry under various channel conditions in 802.11 WLAN [45]. In contrast, we leverage a multi-dimensional channel inference algorithm that uses link-layer performance and classification to perform rate selection.

6.4 Geometrical channel characterization

The patterns in geometrical representation has been used to characterize propagation models of wireless channels. A geometrical structure based channel model proposed in [46] uses three parameters of the signal to characterize a channel: the power of the multipath component, the time of arrival of the component, and the angle of arrival of the components. An interference classification approach using the angular difference between the current measurement and the stored reference power values of the interference to identify interference is introduced in [12]. Also, the authors choose the transmission channel based on the identified interference. In contrast to prior

work, we use a multi-dimensional model to infer the channel type to lead to better decisions for link adaptation protocols.

7 Conclusion

In this work, we presented a context-aware framework for link-level adaptation that leveraged decision-tree classification and geometry-based channel inference. With custom and off-the-shelf hardware platforms, we evaluated the framework across emulated and in-field channels and showed that the proposed adaptation can achieve more than three times the throughput of SNR-based adaptation in certain conditions. We evaluated the effectiveness of training across multiple frequency bands and found 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 how to infer channels and add channels via training, a separate work would examine the tradeoff in a perfunctory versus exhaustive set of known channels.

Abbreviations

MRR: Multi-rate retry

Acknowledgements

Not applicable.

Authors' contributions

HL carried out the training and testing of the context-aware rate adaptation and drafted the manuscript. JH carried out the training and testing of the channel inference method. JW participated in the study of related work. DR participated in the design of the study and helped to draft the manuscript. JC conceived of the study, participated in its design and coordination, and helped to draft the manuscript. All authors read and approved the final manuscript.

Authors' information

Not applicable.

Funding

Not applicable.

Availability of data and materials

The datasets supporting the conclusions of this article are available in the DART-CARs repository [ark:/13960/t3vt8gh06](https://doi.org/10.13960/t3vt8gh06).

Competing interests

The authors declare that they have no competing interests.

Received: 30 May 2018 Accepted: 13 February 2020

Published online: 24 February 2020

References

- H. Hartenstein, K.P. Laberteaux, A tutorial survey on vehicular ad hoc networks. *IEEE Commun. Mag.* **46**(6), 164–171 (2008)
- W. Alasmary, W. Zhuang, in *Proceedings of IEEE Vehicular Technology Conference*. The mobility impact in IEEE 802.11p infrastructureless vehicular networks (IEEE, Ottawa, 2010)
- G. Karagiannis, O. Altintas, E. Ekici, G. Heijenk, B. Jarupan, K. Lin, T. Weil, Vehicular networking: A survey and tutorial on requirements, architectures, challenges, standards and solutions. *IEEE Commun. Surv. Tutor.* **13**(4), 584–616 (2011)
- H. Kreml, I. Seskar, P. Spasojevic, in *Proceedings of IEEE Vehicular Networking Conference*. Experimental modeling of the effect of adjacent lane traffic on the vehicular channel (IEEE, Tokyo, 2009)
- R. Meireles, M. Boban, P. Steenkiste, O. Tonguz, J. Barros, in *Proceedings of IEEE Vehicular Networking Conference*. Experimental study on the impact of vehicular obstructions in VANETs (IEEE, Jersey City, 2010)
- M. Baldauf, S. Dustdar, F. Rosenberg, A survey on context-aware systems. *Int. J. Ad Hoc Ubiquit. Comput.* **2**(4), 263–277 (2007)
- J. Hong, E. Suh, S.J. Kim, Context-aware systems: A literature review and classification. *Expert Syst. Appl.* **36**(4), 8509–8522 (2009)
- T.Y.-H. Chen, A. Sivaraman, S. Das, L. Ravindranath, H. Balakrishnan, in *MIT Computer Science and Artificial Intelligence Laboratory*. Designing a context-sensitive context detection service for mobile devices MIT-CSAIL-TR-2015-029, (Cambridge, 2015)
- C. Clancy, J. Hecker, E. Stuntebeck, T. O'Shea, Applications of machine learning to cognitive radio networks. *IEEE Wirel. Commun.* **14**(4), 47–52 (2007)
- T.T.T. Nguyen, G. Armitage, A survey of techniques for internet traffic classification using machine learning. *IEEE Commun. Surv. Tutor.* **10**(4), 56–76 (2008)
- S. Roy, C. Ellis, S. Shiva, D. Dasgupta, V. Shandilya, Q. Wu, in *Proceedings of Hawaii International Conference on System Sciences*. A survey of game theory as applied to network security (IEEE, Honolulu, 2010)
- K.R. Chowdhury, I.F. Akyildiz, in *Proceedings of International Conference on Communications*. Interferer classification, channel selection and transmission adaptation for wireless sensor networks (IEEE, Dresden, 2009)
- S.H.Y. Wong, H. Yang, S. Lu, V. Bharghavan, in *Proceedings of International Conference on Mobile Computing and Networking*. Robust rate adaptation for 802.11 wireless networks (ACM, Los Angeles, 2006)
- W. Gong, H. Liu, J. Liu, X. Fan, K. Liu, Q. Ma, X. Ji, Channel-aware rate adaptation for backscatter networks. *IEEE/ACM Trans. Netw.* **26**(2), 751–764 (2018)
- S. Jangsher, S. A. Khayam, Q. Chaudhari, in *Proceedings of INFOCOM*. Application-aware mimo video rate adaptation (IEEE, Orlando, 2012)
- J. Kim, S. Kim, S. Choi, D. Qiao, in *Proceedings of IEEE INFOCOM*, vol. 6. CARA: Collision-aware rate adaptation for IEEE 802.11 WLANs (IEEE, Barcelona, 2006)
- J. Camp, E. Knightly, Modulation rate adaptation in urban and vehicular environments: Cross-layer implementation and experimental evaluation. *IEEE/ACM Trans. Netw.* **18**(6), 1949–1962 (2010)
- J.R. Quinlan, Induction of decision trees. *Mach. Learn.* **1**, 81–106 (1986)
- F.H. Botes, L. Leenen, R. DeLaHarpe, Ant tree miner amyntas: Automatic, cost-based feature selection for intrusion detection. *J. Inf. Warf.* **16**(4), 73–92 (2017)
- M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, I.H. Witten, The weka data mining software: an update. *ACM SIGKDD Explor. Newsl.* **11**(1), 10–18 (2009)
- P. Murphy, A. Sabharwal, B. Aazhang, in *Proceedings of European Signal Processing Conference*. Design of WARP: Wireless open-access research platform (IEEE, Florence, 2006)
- T.B. Sorensen, P.E. Mogensen, F. Frederiksen, in *Proceedings of Vehicular Technology Conference*. Extension of the ITU channel models for wideband (OFDM) systems (IEEE, Dallas, 2005)
- Aximuth ACE - MIMO Channel Emulator. <http://www.aximuthsystems.com>. Accessed 10 May 2018
- IEEE Computer Society LAN MAN Standards Committee and others, Wireless LAN Medium Access Control and Physical Layer Specification. ANSI/IEEE Standard 802.11 (2003)
- Mac80211 document. <http://wireless.kernel.org/en/developers/Documentation/mac80211>. Accessed 10 May 2018
- W. Yin, P. Hu, J. Indulska, K. Bialkowski, in *Proceedings of the ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*. Performance of MAC80211 rate control mechanisms (ACM, Miami, 2011)
- C. M. Bishop, *Pattern Recognition and Machine Learning*, vol. 4. (Springer, New York, 2006)
- N. Koçi, M.K. Marina, in *Proceedings of the Conference Local Computer Networks*. Understanding the role of multi-rate retry mechanism for effective rate control in 802.11 wireless LANs, (Zurich, 2009)
- D. Xia, J. Hart, Q. Fu, in *Proceedings of the International Conference On*. Evaluation of the minstrel rate adaptation algorithm in IEEE 802.11 g WLANs (IEEE, Budapest, 2013)
- Q.N. Nguyen, M. Arifuzzaman, K. Yu, A context-aware green information-centric networking model for future wireless communications. *IEEE Access*. **6**, 22804–22816 (2018)
- E. Ngai, M.B. Srivastava, J. Liu, in *Proceedings of INFOCOM*. Context-aware sensor data dissemination for mobile users in remote areas, (Orlando, 2012)
- T.E. Bogale, X. Wang, L.B. Le, Machine intelligence techniques for next-generation context-aware wireless networks. *CoRR abs/1801.04223* (2018)
- E.G. Larsson, O. Edfors, F. Tufvesson, T.L. Marzetta, Massive mimo for next generation wireless systems. *IEEE Commun. Mag.* **52**(2), 186–195 (2014)
- A. Rahmati, L. Zhong, Context-based network estimation for energy-efficient ubiquitous wireless connectivity. *IEEE Trans. Mob. Comput.* **10**, 54–66 (2011)
- V. Pejovic, E.M. Belding. Whiterate: A context-aware approach to wireless rate adaptation, vol. 13, (2014), pp. 921–934
- P. Shankar, T. Nadeem, J. Rosca, L. Iftode, in *Proceedings of the IEEE International Conference on Network Protocols*. CARs: Context-aware rate selection for vehicular networks (IEEE, Orlando, 2008)
- G. Judd, X. Wang, P. Steenkiste, in *Proceedings of the ACM International Conference on Mobile Systems, Applications, and Services*. Efficient

- channel-aware rate adaptation in dynamic environments (ACM, Breckenridge, 2008)
38. L. Huang, C.W. Sung, C.S. Chen, in *Proceedings of the IEEE International Conference on Communication Systems*. Context-aware wireless broadcast for next generation cellular networks (IEEE, Macau, 2014)
 39. C.F. Mecklenbrauker, A.F. Molisch, J. Karedal, F. Tufvesson, A. Paier, L. Bernado, T. Zemen, O. Klemp, N. Czink, Vehicular channel characterization and its implications for wireless system design and performance. *Proc. IEEE*. **99**(7), 1189–1212 (2011)
 40. J. Yim, Introducing a decision tree-based indoor positioning technique. *Expert Syst. Appl.* **34**(2), 1296–1302 (2008)
 41. S.J. Kim, D. Yoon, in *International Conference on Information and Communication Technology Convergence*. Automatic modulation classification in practical wireless channels (IEEE, Jeju, 2016)
 42. I. Pefkianakis, Y. Hu, S.-B. Lee, C. Peng, S. Sakellari, S. Lu, Window-based rate adaptation in 802.11n wireless networks. *Mob. Netw. Appl.* **18**(1), 156–169 (2013)
 43. S.-B. Lee, S. Lu, Towards MIMO-Aware 802.11n rate adaptation. *IEEE/ACM Trans. Netw.* **21**(3), 692–705 (2013)
 44. W.-L. Shen, Y.-C. Tung, K.-C. Lee, K. Lin, S. Gollakota, D. Katabi, M.-S. Chen, Rate adaptation for 802.11 multiuser MIMO networks. *IEEE Trans. Mob. Comput.* **13**(1), 35–47 (2012)
 45. I. Kim, Y.T. Kim, Realistic modeling of IEEE 802.11 WLAN considering rate adaptation and multi-rate retry. *IEEE Trans. Consum. Electron.* **57**(4), 1496–1504 (2011)
 46. P. Petrus, J. H. Reed, T. S. Rappaport, Geometrical-based statistical macrocell channel model for mobile environments. *IEEE Trans. Commun.* **50**(3), 495–502 (2002)

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)