Scout: An Asymmetric Vehicular Network Design over TV Whitespaces

Tan Zhang Sayandeep Sen Suman Banerjee

Department of Computer Sciences
University of Wisconsin-Madison, USA
{tzhang, sdsen, suman}@cs.wisc.edu

ABSTRACT

We explore the use of TV whitespaces based communication systems for providing robust connectivity to vehicles. We envision a setup where whitespaces base stations mounted by roadside that communicate with whitespaces gateway nodes placed on vehicles. A key challenge in this setup is the *asymmetry in transmission power limits* – the fixed base station is allowed to communicate at up to 4W, while the mobile gateways on vehicles are limited to 100mW. This paper presents a specific network design called *Scout* to deal with this asymmetry in which whitespaces transceivers are used in the downlink direction while a more traditional cellular path is used in the uplink one. As the key component of this system, we describe a novel channel probing mechanism that sends a forward radio to look ahead and identify the best channel parameters to be used when the rear radio eventually reaches the forward post. We report various challenges and experience in this design and our ongoing plans to use it for providing Internet access to public city buses. Our initial results indicate a 4× coverage improvement and 1.4× throughput gain achieved by *Scout*.

1. INTRODUCTION

TV whitespaces (512-698MHz), recently released by FCC for unlicensed usage in the U.S. [6], provides wireless communication systems with abundant spectrum resource (180MHz) and excellent propagation characteristics (1.9km). In this work we intend to leverage this whitespaces spectrum for providing long-range and high-speed network connectivity to moving vehicles.

The application we target at is to provide Internet access to the commuters of a city metro with about 250 buses. Figure 1 shows our target deployment, which consists of base stations mounted along roadside, connected to a proxy server. The base stations backhaul Internet traffic generated by the whitespaces gateway nodes (clients) mounted on top of the buses. Each gateway node is connected to a WiFi hotspot inside the bus for providing Internet access to commuters.

Ideally, we would like to use a whitespaces-only solution. However, due to various scalability and cost reasons, we instead adopt an asymmetric network design called *Scout* which uses whitespaces for downlink communications and the cellular path for uplink communications. In particular, we present two key ideas in *Scout*: (i) the use of a cellular link to significantly extend the coverage of each whitespaces base station, (ii) the use of an extra scouting radio to compensate for the effects of feedback delay over the cellular link in taking protocol decisions.

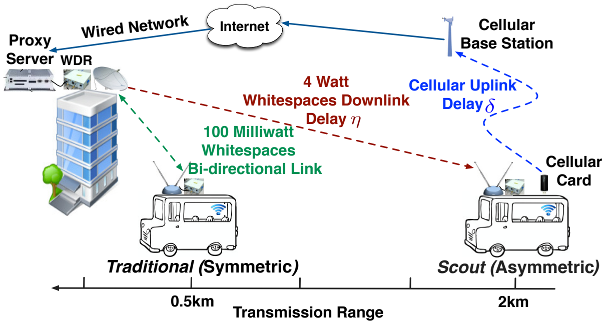

Figure 1: *Traditional* (symmetric) Whitespaces Network v.s. *Scout* (asymmetric) Whitespaces Network. The symmetric network (left vehicle) uses whitespaces for both uplink and downlink communications. In contrast, *Scout* (right vehicle) uses whitespaces for downlink communications, while using the cellular service for uplink communications. The delay of the cellular feedback path $\delta \gg \eta$, the delay of whitespaces downlink.

1.1 Motivation and design of Scout

One major design goal of this vehicular network is to leverage the good propagation characteristics of TV whitespaces frequencies to reduce the deployment and management cost by having a small number of base stations to cover a large area. Unfortunately, the *asymmetric transmission power limit* in TV whitespaces has significantly limited the reachability of the base station. According to the FCC's recent ruling [6], the mobile client is allowed to transmit at *much lower* power–$100mW$ [1] compared to the static base station allowed to–$4W$. The $40\times$ difference in the power limit is to protect against the potential interference resulting from the mobile clients roaming into some unpredictable locations. Thus, a symmetric, whitespaces-only network solution as proposed in [8] (Figure 1 left vehicle) would limit the transmission range of the base stations to that of the "weak" mobile clients when bi-directional communication applications are commonly supported. Measurements in our whitespaces testbed have shown 4× reduction in the transmission range of the base station ($1.9km$ and $0.5km$ for 4W and 100mW

ACM HotMobile'13, February 26–27, 2013, Jekyll Island, Georgia, USA.

---

[1] This maximum power includes the gain of antennas and regardless of number of channels used for communications.

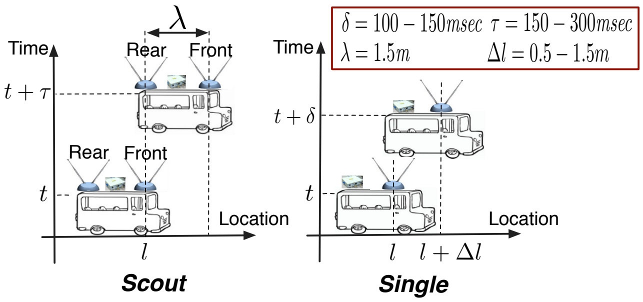

Figure 2: Diagram of *Scout* and *Single*. λ is the antenna separation, τ is the antenna alignment period and δ is the cellular delay. $\tau \geq \delta$.

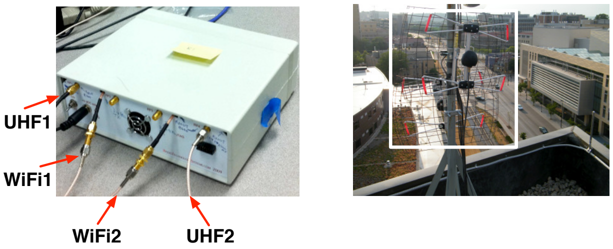

(a) Wideband digital radio.

(b) The base station.

Figure 3: Experiment setup.

respectively) in a symmetric network design. To address this problem, we design *Scout* which has the following characteristics.

(i) An asymmetric network design for improving the coverage of base stations: Figure 1 (right vehicle) shows *Scout*, an asymmetric network architecture to deal with the problem of power asymmetry. *Scout* uses different technologies for uplink and downlink communications. The downlink connection (from the base station to the bus) is over whitespaces frequencies, while the uplink (from the bus to the base station) is over cellular. Thus, *Scout* leverages the ubiquitous cellular connectivity to circumvent the "weak" whitespaces uplink, in turn enabling each whitespaces base station to extend the bi-directional communication range to a maximum. Furthermore, since the downlink load is reported to be dominant in many network applications (10× in WiRover [7]), *Scout* is efficient in utilizing whitespaces spectrum to transmit the majority of data. In our current implementation, we use a WiFi-based whitespaces hardware in [16] for downlink and 3G for uplink communications.

Unfortunately, the operation of the asymmetric network is plagued by the *high feedback latency* in the cellular uplink, which is 10–500× of the packet transmission duration in whitespaces communications. The problem is exacerbated by vehicle mobility which causes rapid change of communication environment. Since most of wireless communication systems rely on channel feedback for making protocol decisions at different layers of the stack, e.g., PHY data rate, FEC at the MAC layer as well as congestion control, error recovery at the transport layer, the poor feedback accuracy can significantly degrade backhaul performance. To tackle the problem of slow uplink, we present our second technique of using an extra radio to measure channel condition for a location in advance. This information is then used for optimizing various transmission decisions for another radio by the time it reaches that specific location.

(ii) A scouting radio for channel estimation: The core intuition of our work comes from the observation that the location of a vehicle has a profound effect on the channel characteristics experienced by a radio mounted on the vehicle [11, 12, 14]. For instance, a vehicle traveling behind a building is observed to have much worse reception than driving in line-of-sight to the base station. To leverage this channel property, we place a scouting radio at the head of the vehicle (front radio) as shown in Figure 2. This front radio passively monitors the downlink traffic at the current location l. It sends back the observed channel information, e.g., packet loss rates, to the base station over the cellular link. After time τ, the receiving radio placed at the rear of the vehicle (rear radio) reaches the same location l. The base station can utilize the channel observation made at l, τ time ago to adjust its transmission for the rear radio.

While the proposed scheme is incapable of tracking fast fading, we will show in Section 2 that channel estimates made by the scouting radio is still beneficial in estimating different link characteristics. This in turn leads to better protocol decisions and ultimately improves the overall performance of the vehicular Internet access. Furthermore, the idea of adding one scouting radio is complementary to various diversity combining techniques used in existing multi-antenna systems, which we will discuss in Section 6.

In this paper, we highlight the efficacy of the scouting radio based mechanism by showing that it can enhance the performance of rate adaptation. We emphasize that *while the problem of rate adaptation has been well studied in a traditional, symmetric communication system, it is unique in our asymmetric network with a feedback link hundred times slower than the forward path.*

Comparison with cellular technologies: Finally, we note that cellular technologies (3G, LTE, WiMax, etc.) also successfully address the problem of providing long-range network coverage to mobile clients with weak transmission power. Unfortunately, translating the gains achieved in *licensed bands* by the cellular technologies to operate over *opportunistic* TV whitespaces appears nontrivial. One of the major challenges is that traditional cellular technologies require centralized coordination by the base station to tightly control the transmission power of each mobile station [5, 18] for managing interference. This allows low power mobile clients to communicate successfully with cellular base stations. However, in the whitespace domain the base station would be incapable of controlling other unlicensed devices potentially using an entirely different communication protocol. The uncoordinated interference is likely to overwhelm the weak uplink signal. *Scout* is robust to interference by communicating the uplink traffic over the licensed spectrum.

In contrast to a symmetric cellular solution over the licensed spectrum, *Scout* offloads the majority of traffic over TV whitespaces. The abundant free spectrum leads to better performance and lower spectrum cost if a pre-paid cellular billing option is available.

1.2 Contributions

Our contribution in this work is two-fold. First, we present an asymmetric network architecture–*Scout* which efficiently uses TV whitespaces for providing wide-area network connectivity to vehicles. Second, as a key component in *Scout*, we propose a novel channel probing framework to address the problem of feedback delay, enabling our system to extract most benefits out of TV whitespaces. Based on our trace driven simulation, *Scout* can extend the coverage of the base station by 4× and achieve a median throughput improvement of 1.4×.

2. FEASIBILITY OF SCOUT

In this section, we validate the feasibility of *Scout*. We begin by explaining our experiment setup.

2.1 Experiment setup

Radio platform: Figure 3(a) shows our radio platform, called the Wide Band Digital Radio (WDR). It performs a frequency translation function similar to the KNOWS platform [1]. With two independent signal processing paths, the WDR can simultaneously translate two signals between the UHF band and the ISM band. This enables us to use a single WDR radio to process signals received from both antennas at the client. The converted signals are fed to WiFi cards for the baseband processing. Due to some hardware limitations, the current version of WDR can support 802.11 b/g data rates up to 18Mbps.

Testbed: Our testbed currently includes a base station and a mobile client. The base station consists of a host machine, a WDR radio, a high-gain power amplifier and a directional TV antenna, as shown in Figure 3(b). The total transmission power at the base station is 3.8W. For the mobile client, we use a vehicle carrying one WDR radio and two omni-directional antennas. The downlink communications are configured at a center frequency of 662MHz with a bandwidth of 20MHz according to a spectrum occupancy database [19].

Metrics: We measure the loss rate calculated for every 10 contiguous packets at the same fixed data rate. We use it as an indicator of channel quality for a given location and at a given time. We then calculate the magnitude of difference in loss rates under different time separations to classify whether channel condition has changed with varying location or time (or both). We denote this time separation as a *lag* in the following discussion.

2.2 Validation of intuition

The feasibility of *Scout* can be ascertained by comparing the following two approaches, which are shown in Figure 2. In *Scout*, suppose the front radio measures the loss property at location l, time t. How accurate is this measurement in predicting the channel condition for the rear radio when it reaches the same location l at time $t + \tau$? We contrast this with the other alternative of a single radio (*Single*). In the single radio scenario, the only radio will estimate the loss property at location l at time t, and use this estimate to predict its performance at location $l + \Delta l$ at time $t + \delta$.

To understand the stability of channel loss properties as only a function of time, we present the variation of loss rates for the same locations with different lags. We measure this by placing a single radio mounted atop a car at 12 equally spaced locations on a 200m road stretch. Figure 4 (*Single static*) shows that loss variation remains small with a lag below 300*msec* for all the measured locations.

We next determine the stability of loss measurements done by the same single radio as a function of both time and location. The speed of the vehicle is between 5m/s (18 km/hr) and 10m/s (35 km/hr), which is typical for urban area due to the 40km/hr speed limit. As can be seen from Figure 4 (*Single* 10m/s and 5m/s), the difference in loss rates increases drastically with increasing lags. The degree of variation is expected as the single radio system is measuring the loss rates at different locations and different time. When using the stale channel observation to predict the loss rate, *Single* would make an estimation error of over 30% under the delay of 100–150*msec* in a 3G uplink.

We finally benchmark the mismatch in the loss rates under *Scout* setup with two radios (front and back) *at the same location* with different lags. The result is again shown in Figure 4 (*Scout variable speed*). We note that the difference in loss rates between two radios at the same location with a lag of 300*msec* remains a fifth of a single radio traveling at 10m/s speed. This demonstrates that *Scout* can indeed improve the channel estimation at a given location.

Based on this maximum lag, we choose the antenna spacing λ to be 1.5m and the driving speed 5–10m/s. This ensures that the highest τ $300msec$ (1.5m$/5$*m/s*) to be within this maximum threshold, and the lowest τ $150msec$ (1.5m$/10$*m/s*) to be greater than the typical cellular delay δ (100*msec*).

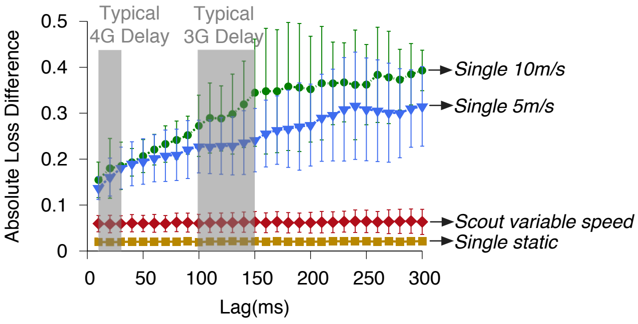

Figure 4: Average of absolute loss difference at various lags for 12Mbps packets. The lag is the elapsed time between two measurements. *Single* denotes one radio and *Scout* is the two radio setup.

Finally, we note that adopting 4G LTE as the feedback link would introduce much lower latency (10–25ms). However, Figure 4 shows that it still leads to 15% - 20% estimation error which *Scout* can avoid. This justifies the effectiveness of *Scout* along the advancement of the cellular technology.

3. DESIGN OF SCOUT

In *Scout* we use a single radio at the base station for downlink transmission. At the client, we use the rear radio for packet reception, while leveraging the front radio to passively monitor the downlink traffic for channel estimation. The packet acknowledgments generated by both radios are sent over the cellular link to the base station.

To select a data rate for the rear radio at a given location l, we use the feedback previously generated by the front radio at l, combined with the feedback from the rear radio currently available at the base station. We use a GPS device at the client to relate a feedback from the front radio to the monitored location. The GPS device reports its reading once every second over the cellular uplink to the base station. We present the pseudo-code of *Scout* in Algorithm 3 and give a detailed explanation next.

Combining feedback from dual radios: In our deployment we observe mismatched reception performance between two radios when their antennas are mounted at different locations on top of the vehicle, e.g., different orientation and tilting. This difference is usually small enough for the two radios to experience a common channel trend. However, we find that selecting data rates solely based on the feedback from the front radio leads to sub-optimum performance.

In *Scout* we utilize the front radio to detect channel variation, while using the rear radio to identify a set of rate candidates to select from. Specifically, we use one of the existing algorithms, e.g., SampleRate, to select a preliminary data rate based on the stale feedback from the rear radio. This is shown in Line 1 of Algorithm 3. This rate is expected to work reasonably well under stable channel condition. Once we detect the change of channel condition at the front radio, we use the detected channel trend and the set of candidate rates successfully received by the rear radio to derive a more appropriate rate decision.

Relating feedback to the observed location: The accuracy of associating the feedback to its monitored location significantly affects

Algorithm 3.1: SCOUT$(t, \lambda, v, w, \mathcal{F}, \mathcal{R})$

```
INPUT $t$: Current time, $\lambda$: Antenna spacing.
$v$: Vehicle speed, $w$: Window size.
$\mathcal{F}$: Set of ACKs from the front antenna.
$\mathcal{R}$: Set of successful rates from the rear antenna.
OUTPUT $rate$: Selected rate.
$\tau = \lambda / v$
$rate \leftarrow set\_origin\_rate(\mathcal{R})$                                   (1)
$W_{cur} \leftarrow \{f_j : f_j \in \mathcal{F}, |j - (t - \tau)| <= w/2\}$                  (2)
$W_{prev} \leftarrow \{f_j : f_j \in \mathcal{F}, |j - (t - \tau - w)| <= w/2\}$             (3)
$var\_detected \leftarrow detect\_variation(W_{cur}, W_{prev})$
if $var\_detected < 0$
  then $rate \leftarrow set\_lower\_successful\_rate(rate_{prev}, \mathcal{R})$      (4)
  else if $var\_detected > 0$
  then $rate \leftarrow set\_higher\_successful\_rate(rate_{prev}, \mathcal{R})$     (5)
return $(rate)$
```

the performance of *Scout*. For example, Figure 4 shows that a location error of 3m (corresponding to 300ms lag at 10m/s speed) can lead to 40% off in loss estimation. Unfortunately, our low priced GPS modules have a positioning inaccuracy on the order of 10m.

To circumvent this problem, we use the speed reading v reported by the GPS instead of the geo-location reading since it has a much higher accuracy–0.1m/s. We calculate the radio alignment period τ, which is the time elapsed for the rear radio to reach the same location l since the front radio was previously at l (Figure 2). We calculate this period with the formula $\tau = \lambda/v$ where λ is the fixed antenna separation. Note that v remains constant in our calculation considering the short GPS updating interval (1*second*). We then retrieve a window of ACK packets W_{cur} generated by the front radio τ time ago, which is the desired feedback observed at location l (Line 2 in Alg. 3). Since τ is small (< 300*msec*), the positioning inaccuracy of *Scout* is below $3cm$ (0.1*m/s* $\times$ 300*msec*).

Detecting channel variation: To this purpose, we analyze the error performance between the feedback W_{cur} generated by the front radio at location l, and another window of ACKs W_{prev} generated prior to W_{cur} (Line 3 in Alg. 3). We compare the packet loss of each common data rate between these two windows of packets. If the loss rate of a given data rate in W_{cur} increases (decreases) by at least α (β) fold over W_{prev}, we conclude that the error performance of that data rate has changed. We use a voting mechanism to combine results for all the data rates to determine the trend of channel variation. We empirically set α and β at a large value–0.5 to prevent *Scout* from reacting to random channel fluctuation. We configure the time duration of W_{cur} and W_{prev} to be 25ms to collect sufficient, yet relevant feedback.

Adjusting data rates in response to channel variation: Based on the detected trend of channel variation, we select the next higher or lower data rate to the previous rate decision, but only from a set of candidate rates recently succeeding at the rear radio (Line 4-5 of Alg. 3). By doing so, we can choose a rate not only suitable for the current channel condition but also consistent with the reception performance of the rear radio.

4. IMPLEMENTATION

We implement *Scout* at a 3.5 layer on top of the cellular and whitespaces links as shown in Figure 6. To provide a single link abstraction, we leverage a virtual network device in Linux called *TUN*, and have the base station and the client exchange application data through it. We create a user-space program passing packets between *TUN* and one of the underlying network interfaces, i.e., cellular and whitespaces. For downlink communications, the program running at the base station reads the application data from *TUN*, sending them through the whitespaces interface (ath0). At the client, the program delivers downlink packets received by the rear whitespaces radio (ath1), and sends them via *TUN* to the application. This client program also generates ACKs for the packets received by both whitespaces radios (ath0 and ath1), sending them with the GPS readings to the base station over the cellular link.

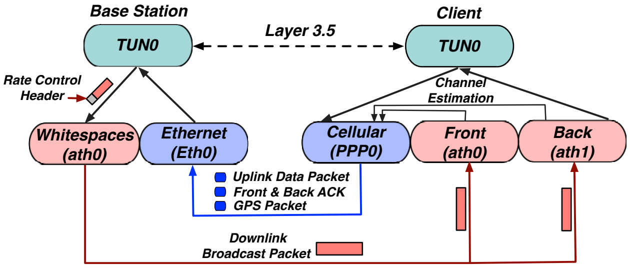

Figure 6: Implementation of *Scout* with *TUN* devices.

We configure the whitespaces radio at the base station to operate in the WiFi broadcast mode. This prevents any client radio from generating MAC layer ACK over the whitespaces, while having both client radios to receive the downlink packets for *Scout*. We enable the base station to control the broadcast data rate by appending a control header in each downlink packet, and modify the Ath5k driver to accept this control information.

We implement the user-space program in 4500 lines of C++ and add about 20 lines of C code in the Ath5k driver.

5. EVALUATION

We evaluate the performance of *Scout* after integrating it with two popular rate adaptation algorithms, SampleRate [4] and RRAA [21]. We denote the enhanced algorithms as Scout-Sample and Scout-RRAA. We have found that *Scout* achieves median throughput improvement of 38% and 39% over SampleRate and RRAA.

Methodology: We use trace-driven emulation as a preliminary evaluation of *Scout*, with real trace captured in TV whitespaces. For trace collection, we use a similar approach as in AccuRate [15] by instructing the base station to transmit short back-to-back packets (200 byte), using 8 802.11 b/g data rates up to 18Mbps alternatively. We then use two radios at the client to capture packet traces from 10 drives along a 1.5km bus route. The distance between the base station and the client is about 200–750m, and the vehicle speed is about 18 – 35 km/hr, which is typical for our city metro. We choose the antenna spacing to be $1.5m$ as described in Section 2. We term each set of 8 contiguous packets at all different rates as a *packet train*.

For our emulation, we make each algorithm select one data rate in each packet train. If the chosen rate belongs to one of the successful rates in the current train, we conclude that this rate succeeds in the duration of this train and vice versa. We then calculate the throughput of different algorithms based on these results. To emulate the feedback delay (typically 100ms in our testbed), we provide each algorithm with feedback generated by both radios 100ms ago at the client for rate selection.

We empirically adjust the sampling interval of SampleRate from 10 packets to once every train (8 packets). We further set the size of estimation window for both algorithms to be 10 trains, which is found to perform best in our emulation.

Throughput: We calculate the throughput of different algorithms in each 50m road segment to evaluate performance improvement

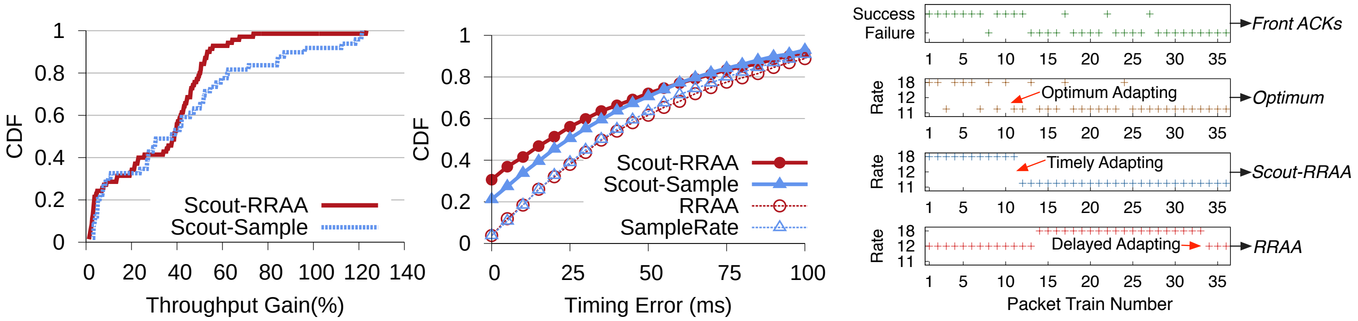

(a) Throughput gain.

(b) Detection accuracy.

(c) Snapshot of rate decisions.

Figure 5: Experiment results. (a) CDF of throughput improvement of Scout-RRAA and Scout-Sample over RRAA and SampleRate. (b) CDF of timing errors in detecting channel variation. (c) A snapshot for rate decisions made by different algorithms.

of *Scout* under different road conditions [2]. The average throughput from 10 drives is 5.5Mbps and 5.2Mbps for Scout-Sample and Scout-RRAA respectively. Figure 5[a] shows the CDF of throughput improvement achieved by Scout-Sample and Scout-RRAA over SampleRate and RRAA. We observe a median throughput gain of 38% and 39%, and an upper quartile gain of 57% and 48%. We expect more improvement when higher data rates are available.

Timeliness in adapting to channel variation: To this purpose, we benchmark the time offset of different algorithms to an optimum algorithm in adapting to channel variation. The optimum algorithm selects the highest successful data rate in each train. Specifically, we record the time of data rate changes made by different algorithms, and measure the absolute time difference for each rate change between algorithms under study to the optimum algorithm. We denote each time difference as a timing error. Figure 5[b] shows the CDF of timing errors of different algorithms. We observe 31% and 21% of the case when Scout-RRAA and Scout-Sample have detected channel variation at the exact time. In contrast, RRAA and SampleRate have only detected 4% of channel variation in a timely fashion. Thus, Scout-RRAA and Scout-Sample have achieved a 8× and 5× improvement in detection accuracy. We further report that the two *Scout* algorithms incur 5% more false positives than their original counterpart for detecting channel variation.

Snapshot of different algorithms in operation: We finally present a snapshot of rate decisions made by different algorithms in Figure 5[c]. First, we can observe a similar trend of channel condition observed by the two radios at the same location, when referring to ACKs previously generated by the front radio (Front ACKs) and the data rates selected by the optimum algorithm for the rear radio (Optimum). Second, Scout-RRAA accurately detects the channel variation, and adapts to a lower rate at the same time as the optimum algorithm does. In contrast, RRAA selects a lower rate at a much later time of 23 trains (about 115ms) due to the feedback delay. Finally, Scout-RRAA directly selects 11Mbps rate based on those candidate rates successfully received at the rear radio, rather than choosing the ineffective sequential rate–12Mbps as RRAA does.

6. DISCUSSION AND FUTURE WORK

Improving channel estimation: In *Scout* the front radio at each mobile client monitors the downlink traffic broadcasted to *all the clients* for channel estimation. We expect this downlink traffic to be highly available when the number of clients served by a base station is high. As a performance enhancement, we intend to develop a mechanism for the base station to periodically send probe traffic upon detecting the downlink to be idle. Furthermore, we plan to investigate more sophisticated indicators, e.g., channel response for detecting channel variation.

Fully leveraging the dual radios: The *Scout* design currently uses one extra radio (front) solely for channel estimation. In our future work we intend to investigate different schemes in combining signals received by both radios to improve reception diversity, e.g., via maximum ratio combining [20] at the PHY layer or packet level combining [9] at the MAC layer.

Scouting at different layers: We intend to generalize the application of *Scout* feedback to enhance the performance of all layers of the network stack. For instance, at the MAC layer we can potentially apply forward error correction codes (FEC) and proactively duplicate downlink packets based on this feedback information. Similarly, at the transport layer we will leverage *Scout* to detect connection "blackout", and predicatively offload the traffic over the cellular link to prevent the slow start of TCP.

7. RELATED WORK

We divide the related work into five categories, i.e., whitespaces networking, vehicular networking, multi-antenna systems, channel estimation and rate adaptation.

Whitespaces networking: A few research efforts [1, 10] have explored designing networks over TV whitespaces. WhiteFi [1], as the first whitespaces network, uses a symmetric network design with a WiFi-like protocol. SenseLess [10] is a network design that purely relies on a spectrum occupancy database to determine available channels. Built on these two pieces of work, *Scout* addresses the challenge of extending the coverage of a vehicular network, constrained by the unbalanced transmit power limit in a given free channel. In contrast to WhiteFi, *Scout* uses an asymmetric architecture consisting of one whitespaces link and one cellular link.

Vehicular networking: A large body of prior work [2, 3, 7, 11, 13] has been done for providing Internet connectivity into vehicles. The existing approaches can be categorized into cellular based solutions (MAR [13], WiRover [7]), WiFi based solutions (ViFi [3], MobiSteer [11]) or a combination of both (Wiffler [2]). In contrast to all these approaches, *Scout* explores a whitespaces based backhauling solution backed up by a cellular uplink. It achieves long-range and high-speed network connectivity by harvesting both the long propagation range (1.9km) and the abundant spectrum resource (180MHz) in TV whitespaces.

Multi-antenna systems: The design of multi-radio, multi-antenna

---

[2] The reason for choosing the road segment to be 50m is to ensure that at least 5s trace is available for calculating each throughput sample.

systems [9, 17, 20] have been well studied in the past work. These MIMO systems harness the path diversity in wireless channel for robust reception [9, 20], scaling throughput [20], or simultaneous communications with multiple users [17]. *Scout* is different from all these MIMO techniques by using an extra radio system for channel estimation. More significantly, *Scout* is complementary to all these techniques.

Location-based channel estimation: *Scout* leverages a location dependent channel property, which has been reported in the prior work [11, 12, 14]. Bartendr [14] and BreadCrumbs [12] use the location of the mobile client to predict network connectivity in cellular networks and WiFi respectively. MobiSteer [11] uses the location of a vehicle to select the best AP and the directional beam to serve the vehicle. All the above approaches require a training database which is updated on a coarse timescale, e.g., order of days [14]. In contrast, *Scout* can obtain "fresh" channel information collected shortly ago (150–300ms), achieving higher accuracy without any training overhead.

Rate adaptation: We integrate *Scout* with two popular rate adaptation algorithms, i.e., SampleRate [4] and RRAA [21], to demonstrate its efficiency. SampleRate periodically picks a random rate to probe the channel, and selects the rate with the highest throughput. RRAA tracks the packet loss in a short time duration, and uses the predetermined loss threshold to adapt rate. To be effective, both require timely feedback, which is hampered by the cellular delay, yet is benefited by the scouting radio.

8. CONCLUSIONS

In this work we explore an asymmetric network design called *Scout* for providing wide-area vehicular network connectivity over TV whitespaces. The proposed architecture circumvents the bottleneck of the whitespaces uplink with the cellular technology, and significantly extends the transmission range of the base stations. To deal with the feedback delay in the cellular uplink, we have designed a novel channel estimation framework that uses one scouting radio to measure the channel condition at a location beforehand. It provides the base station with this more accurate channel information to select transmission parameters for the receiving radio. Based on our trace-driven simulation, *Scout* leads to 4× improvement in network coverage and 1.4× improvement in downlink throughput.

Acknowledgments

We would like to thank Peter W. Wolniansky for providing us with the radio prototypes operating in TV whitespaces. We are also grateful to our shepherd Matt Welsh whose feedback helped bring the paper to its final form. Tan Zhang, Sayandeep Sen and Suman Banerjee are supported in part by the US National Science Foundation through awards CNS–1040648, CNS–0916955, CNS–0855201, CNS–0747177, CNS–1064944, and CNS–1059306.

9. REFERENCES

[1] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White space networking with wi-fi like connectivity. In *SIGCOMM*, 2009.

[2] A. Balasubramanian, R. Mahajan, and A. Venkataramani. Augmenting mobile 3g using wifi. In *MobiSys*, 2010.

[3] A. Balasubramanian, R. Mahajan, A. Venkataramani, B. N. Levine, and J. Zahorjan. Interactive wifi connectivity for moving vehicles. In *SIGCOMM*, 2008.

[4] J. C. Bicket. Bit-rate selection in wireless networks. In *Master's thesis, Massachusetts Institute of Technology*, 2005.

[5] FCC. Ieee standard for local and metropolitan area networks, 2009.

[6] FCC. Unlicensed operation in the tv broadcast bands, second memorandum opinion and order, 2010.

[7] J. Hare, L. Hartung, and S. Banerjee. Beyond deployments and testbeds: experiences with public usage on vehicular wifi hotspots. In *MobiSys*, 2012.

[8] Microsoft Research. Networking over white spaces. http://research.microsoft.com/en-us/projects/KNOWS/deployment.aspx.

[9] A. Miu, H. Balakrishnan, and C. E. Koksal. Improving loss resilience with multi-radio diversity in wireless networks. In *MobiCom*, 2005.

[10] R. Murty, R. Chandra, T. Moscibroda, and P. V. Bahl. Senseless: A database-driven white spaces network. *IEEE Transactions on Mobile Computing*, 2012.

[11] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. Das. Mobisteer: using steerable beam directional antenna for vehicular network access. In *MobiSys*, 2007.

[12] A. J. Nicholson and B. D. Noble. Breadcrumbs: forecasting mobile connectivity. In *MobiCom*, 2008.

[13] P. Rodriguez, I. Pratt, J. Chesterfield, R. Chakravorty, and S. Banjeree. Mar: A commuter router infrastructure for the mobile internet. In *Mobisys*, 2004.

[14] A. Schulman, V. Navda, R. Ramjee, N. Spring, P. Deshpande, C. Grunewald, K. Jain, and V. N. Padmanabhan. Bartendr: a practical approach to energy-aware cellular data scheduling. In *MobiCom*, 2010.

[15] S. Sen, N. Santhapuri, R. R. Choudhury, and S. Nelakuditi. Accurate: constellation based rate estimation in wireless networks. In *NSDI*, 2010.

[16] S. Sen, T. Zhang, M. Buddhikot, S. Banerjee, D. Samardzija, and S. Walker. A dual technology femto cell architecture for robust communication using whitespaces. In *DySpan*, 2012.

[17] C. Shepard, H. Yu, N. Anand, E. Li, T. Marzetta, R. Yang, and L. Zhong. Argos: practical many-antenna base stations. In *Mobicom*, 2012.

[18] C. Smith and D. Collins. *3G Wireless Networks*. McGraw-Hill, 2001.

[19] Spectrum Bridge Inc. Spectrum bridge database. http://www.spectrumbridge.com/Home.aspx.

[20] P. Thornycroft. Designed for speed: Network infrastructure in an 802.11n world. In *Aruba Networks white paper*, 2007.

[21] S. H. Y. Wong, H. Yang, S. Lu, and V. Bharghavan. Robust rate adaptation for 802.11 wireless networks. In *MobiCom*, 2006.