Wireless channel dynamics for relay selection under ultra-reliable low-latency communication

Rigge, Paul; Swamy, Vasuki Narasimha; Nelson, Christian; Tufvesson, Fredrik; Sahai, Anant; Nikolic, Borivoje

Published in:
2020 IEEE 31st Annual International Symposium on Personal, Indoor and Mobile Radio Communications, PIMRC 2020

DOI:
10.1109/PIMRC48278.2020.9217210

2020

Wireless Channel Dynamics for Relay Selection under Ultra-Reliable Low-Latency Communication

Paul Rigge*, Vasuki Narasimha Swamy†, Christian Nelson‡, Fredrik Tufvesson‡, Anant Sahai*, Borivoje Nikolić*
*University of California, Berkeley, CA, USA, †Intel, Santa Clara, CA, USA, ‡Lund University, Lund, Sweden

***Abstract*—Ultra-reliable, low-latency communication (URLLC) is being developed to support critical control applications over wireless networks. Exploiting spatial diversity through relays is a promising technique for achieving the stringent requirements of URLLC, but coordinating relays reliably and with low overhead is a challenge. Adaptive relay selection techniques have been proposed as a way to simplify implementation while still achieving the requirements of URLLC. Identifying good relays with low overhead and high confidence is critical for such adaptive relay selection techniques.**

Channel dynamics must be taken into account by adaptive relay selection algorithms because channel quality may degrade in the time it takes to estimate the relay's channel and schedule a transmission. Spatial channel dynamics are well studied in many settings such as RADAR and the fast-fading wireless channels, but less so in the URLLC context where rare events neglected in other models may be important. In this work, we perform measurements to validate channel models in the slow fading regime of interest. We compare measurements to Jakes's model and discuss the appropriateness of Jakes's model for URLLC relay selection. This is further applied to demonstrate that easily implementable relay selection techniques perform well in practical settings.

Polynomial interpolation and neural-net-based algorithms were evaluated as channel prediction algorithms. These techniques perform orders of magnitude better than relay selection on average (nominal) SNR.

I. Introduction

Ultra-reliabile low-latency communication (URLLC) is of increasing importance, and has been introduced in the 5G standardization efforts [1]. Low-latency communication is an important requirement for many cyber-physical and distributed systems with machine-to-machine (M2M) communication [2]. Ultra-reliability is important in safety-critical settings like robotics or vehicle platooning [2]. Furthermore, URLLC can provide a path for existing control systems designed with wired networks (such as Sercos) to be converted to wireless. Wireless networks are desirable in industrial settings because they promise to reduce weight and cost in routing wires. 5G URLLC proposals [1] have a target latency of 1-10 ms and reliability of 10^{-5}. This reliability may be sufficient to enable many applications, but is much lower than the reliability offered by wired networks and is insufficient for many safety-critical applications.

We consider network sizes of < 100 nodes with short packets on the order of 20 bytes, which in practice corresponds to a position and a velocity in robotic systems. The information flow consists of sensor measurements being sent to a central controller, followed by the central controller distributing actuation messages to each node. Typically, an entire cycle needs to be completed with a latency on the order of $2\,\mathrm{ms}$ and with a probability that every packet in the cycle is delivered successfully of 10^{-9}.

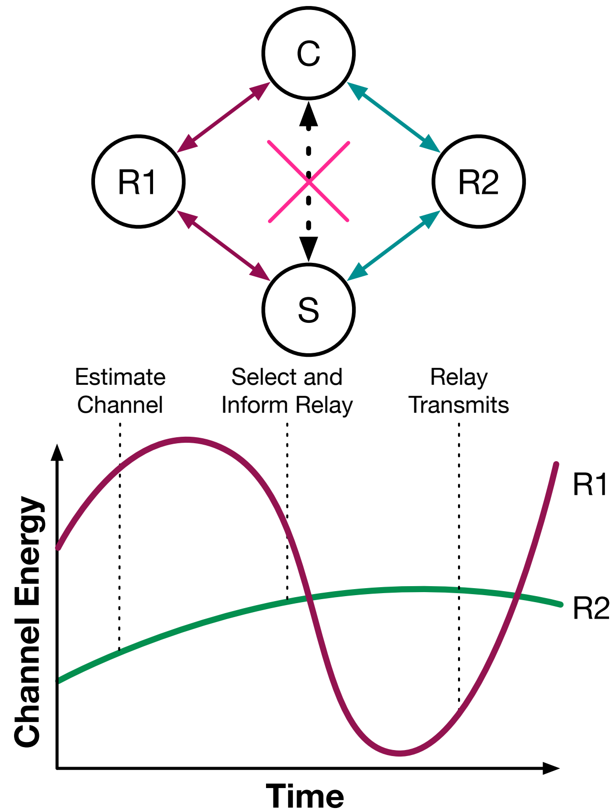

Fig. 1: The top figure depicts a sensor S and controller C attempting to communicate via a channel in fade. They cannot directly communicate because of the fade, so they need other users in the network to act as relays. The sensor and controller estimate the channels to potential relays $R1$ and $R2$, but the process of estimating the channel, selecting a relay, informing the selected relay to transmit, and scheduling the transmission takes time. In that time, the channels may change and a channel that was good may move into a fade. Prediction algorithms that account for channel dynamics are needed to achieve high reliability.

To achieve high reliability and low latency, point-to-point information flows need to have high diversity, otherwise they become a latency/reliability bottleneck. In most scenarios of interest, point-to-point channels do not natively have the

978-1-7281-4490-0/20/$31.00 ©2020 IEEE

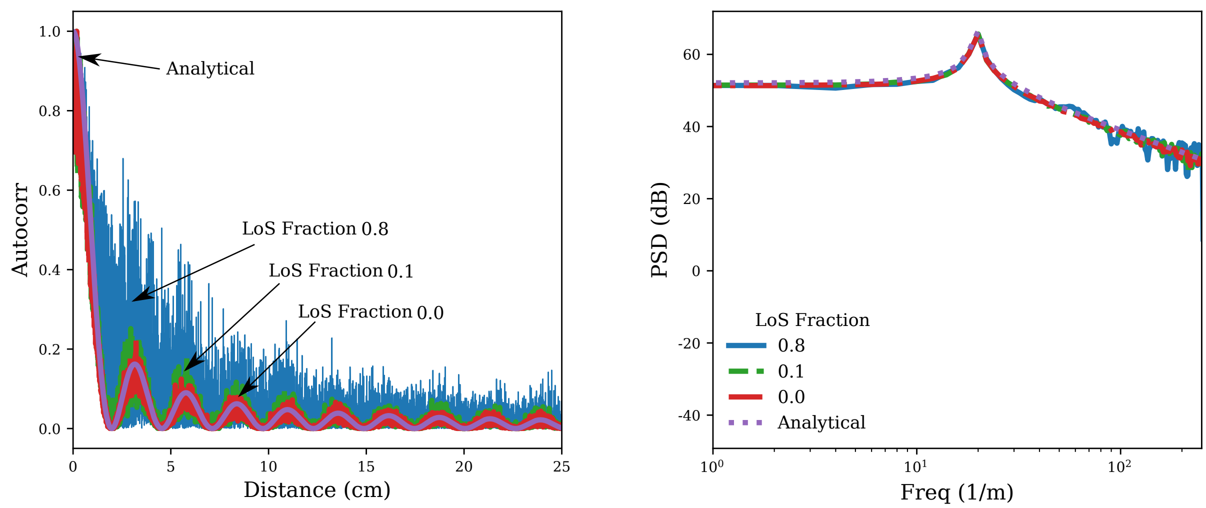

Fig. 2: Simulated autocorrelation and power spectral density (PSD) for channels with $f_c = 5.8\,\text{GHz}$. "LoS Fraction" is the fraction of the signal energy in the line-of-sight path.

amount of diversity necessary for URLLC. Fortunately, diversity can be harnessed through various means. Frequency hopping is simple and implementation friendly, but does not scale well [3]. Spatial diversity is more scalable, but is less straightforward to implement. Previous work in [3] presents a relay-based scheme that harnesses spatial diversity. It relies on simultaneous transmission by all relays using a space-time code, which imposes challenging requirements for implementing synchronization and channel coding.

Later work in [4], [5] presents a modification of [3] that reduces implementation complexity by removing all simultaneous transmissions, instead selecting a small number of high-quality relays. If a small number of relays can be used with confidence that they will have good channel realizations, this will dramatically reduce requirements for synchronization and channel coding. Time-division multiplexing of relays is attractive because of its simplicity, and if high reliability can be achieved with a small number of relays, the overhead is not too high. Of course, this approach relies critically on reliably choosing relays with good channels, which in turn depends on the characteristics of the wireless channel as depicted in Fig. 1, which shows how the channel quality of a relay can change in time between channel estimation and when a relay actually transmits.

Relay selection algorithms in the literature typically use some long-term average behavior of a link as optimization criteria [6]. However, these algorithms are less suitable for ultra-reliable communication because of multipath fading. Links with the same average SNR may have very different instantaneous SNRs because of multi-path fading; or, even worse, a link with lower average SNR may have higher instantaneous SNR because the normally better link may be fading. Studies in [4], [5] investigate channel dynamics in this regime and propose using algorithms to predict whether a relay's channel is going into or out of a fade.

The performance of relay selection algorithms depends critically on channel dynamics. Therefore, it is fair to question if a model captures all the behavior salient to selecting a good relay. In an ultra-reliable setting, rare events neglected in a model may result in failures a model would not predict. It is not clear if standard models like Jakes' model and Rayleigh fading are useful in evaluating relay selection algorithms for URLLC.

This work examines the validity of the modelling assumptions behind adaptive relay selection for URLLC. The main contributions of this paper are:

1) Channel measurements in scenarios relevant to URLLC from over-the-air transmissions.
2) Analysis of channel measurements and conformity to models.
3) Demonstration that channel prediction algorithms perform well on channel measurements.

II. Channel Models

The Rayleigh-fading model is commonly employed for the indoor environments considered here. Rayleigh-fading channels are typically thought of as of sum-of-sinusoids, which is made explicit in models that capture dynamics like Jakes's model [7]. In our previous work [4], we have studied the problem of relay selection in URLLC under these models with both analytical methods and simulations, with both showing good agreement. In particular, [4] concluded that the channels studied could be predicted well for the timescales of URLLC relaying.

Of course, Jakes's and Rayleigh-fading models are relatively simple and may not fully capture the relevant phenomena in selecting good relays. The study in [5] begins to address this by pointing out some ways in which the Rician channel is more favorable than Rayleigh in this setting. The authors in [8] show that channel dynamics are similar for Rayleigh, Rician, and Nakagami channels. They present a model for channel variation where the key parameter is maximum Doppler shift, not a model-specific channel parameter, implying that more nuanced channel models give rise to similar channel dynamics. In most of their cases this model showed good fit with measured data. The finite-SNR diversity-multiplexing tradeoff is explored in [9]. On the extreme side of the tradeoff that maximizes diversity, the diversity gain doesn't depend on the Rician K-factors at all. From this we may conclude that, at least in the case of the Rician and Nakagami channels, adding more detail to the channel model does not affect our ability to predict which relays will be good.

The standard Jakes's model is known to have spatial covariance given by

$$\mathrm{cov}(t) = J_0\left(\frac{2\pi}{\lambda_c} vt\right), \tag{1}$$

where λ_c is the carrier wavelength, v is the maximum Doppler shift, and $J_0(\cdot)$ is the Bessel function of the first kind [4]. We extend the standard Jakes model to include a line-of-sight (LoS) component. We simulate this in a way similar in spirit to [10], which adds a specular component to a sum-of-sinusoids model. The simulation consists of a two-dimensional space with n stationary scatterers distributed uniformly at random. A single, stationary transmitter is placed uniformly at random, and transmits a tone at frequency f_c (wavelength λ_c). A single receiver is placed uniformly at random and moves at constant velocity $\vec{v}$ with uniformly random direction. The channel at time t is given by

$$h(t) = \frac{K}{(K+1)\sqrt{n}} \sum_{i=1}^{n} \exp\left(j \frac{2\pi\left(d_i^{(Rx)}(t) + d_i^{(Tx)}(t)\right)}{\lambda_c}\right) + \frac{1}{K+1} \exp\left(2\pi d^{Rx \leftrightarrow Tx}(t)\right), \tag{2}$$

where n is the number of scatterers, $d_i^{(Rx)}$ is the distance between scatterer i and the receiver, $d_i^{(Tx)}$ is the distance between scatterer i and the transmitter, $d^{(Rx \leftrightarrow Tx}$ is the distance between transmitter and receiver, and K is a parameter determining the fraction of energy in the LoS path.

The first term is the same as when there is no LoS path, and has the same expectation as before: a Bessel function of the first kind. The second term is a tone whose frequency depends on the relative motion between the transmitter and receiver and the carrier wavelength.

Fig. 2 shows simulation results for various channel parameterizations. The addition of a LoS component does not add significant energy beyond the maximum Doppler shift. Thus, adding a LoS component does not change the conclusions in [4] about performing relay selection by predicting channel quality.

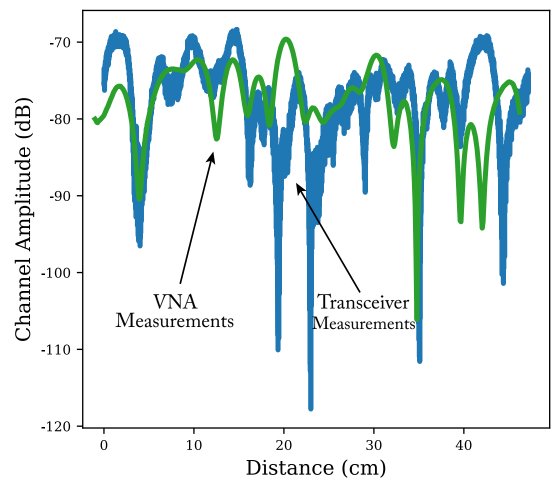

Fig. 3: Representative examples of time-varying channel amplitude from the transceiver-based setup and VNA-based setup.

III. Channel Measurements

We perform measurements in a way conceptually similar to the simulation in Sec. II. The experimental setup consists of a receive antenna and a transmit antenna in an indoor environment. The receive antenna travels along a path with a receiver taking channel estimates. The transmit antenna is fixed for a series of measurements, although it is moved in-between measurements. Line-of-sight between the transmitter and receiver is suppressed with a sheet of aluminum foil at the transmitter.

We use two different measurement setups. The first uses a commercial transceiver with RF frontend, ADC, and DAC to estimate channel coefficients. The second setup uses a vector network analyzer (VNA) to achieve the same task. Both setups are described and analyzed in the following sections.

A. Transceiver-Based Setup

The measurement setup was located in an office building. The hardware platform was implemented using a ADI FMCOMMS-3 frontend with Xilinx ZC706 FPGAs for the baseband. The transmitter and receiver utilized the same board in order to share a local oscillator, removing the need to correct for a frequency offset. Performing no frequency offset correction between transmitter and receiver avoids separation of the LO contribution from the Doppler shift. The ADI reference designs [11] were modified to add custom hardware [12], including stream $\leftrightarrow$ memory DMAs and logic to control timing of a capture. A Gold sequence with period 4095 is truncated to length 1029 and transmitted with period 1029[1]. The sequence was transmitted repeatedly, and raw samples were captured at the receiver, saved to a file, and postprocessed to find a series of channel estimates.

[1]Truncating the sequence is a consequence of our FPGA implementation, nothing that confers a particular benefit.

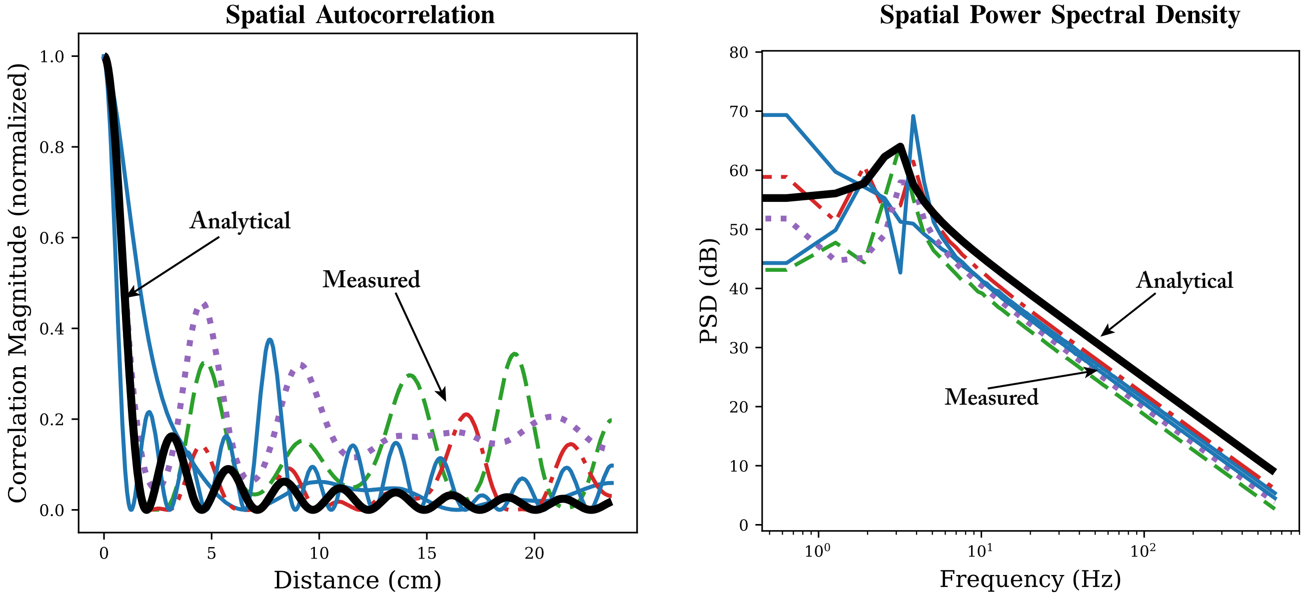

Fig. 4: Normalized correlation and power spectral density of channel measurements at $5.8\,\mathrm{GHz}$. Note that the spatial autocorrelation plot is in distance units whereas the power spectral density plot is in temporal frequency units. The scaling of the PSD plot emphasizes that the peak is at the expected frequency corresponding to the maximum Doppler shift. Both figures are normalized so that lag 0 of the correlation is 1. Bold black lines are the expected analytical result. All other lines are different measured results with the transmit antenna moved in-between measurements.

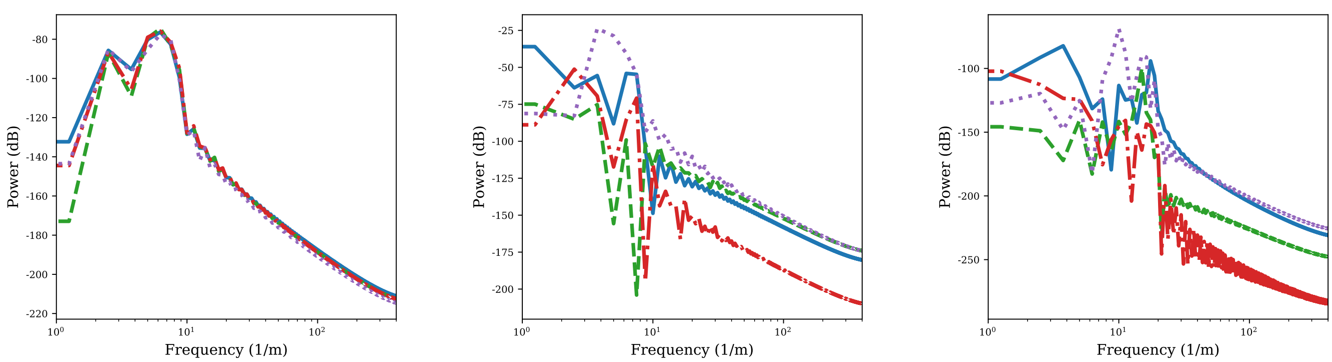

Fig. 5: The figures above show measured power spectral densities from the VNA-based setup. Each line corresponds to a different measurement. The left figure has $f_c = 915\,\mathrm{MHz}$, the center has $f_c = 2.45\,\mathrm{GHz}$, and the right has $f_c = 5.8\,\mathrm{GHz}$.

The position of the receiver is controlled by an XY-table with Parker Automation 404XE linear actuators, Ares servos, and 6K6 motion controller. This apparatus positions the receive antenna with accuracy $<0.1\,\mathrm{mm}$. A computer coordinates the collection of channel estimates and movement of the receive antenna, and a web application running on the FPGA provides an API for initiating and downloading channel estimates [13]. A serial interface controls the motion controller of the XY table.

The antenna is moved on a linear path approximately $50\,\mathrm{cm}$ long at a velocity of $150\,\mathrm{mm/s}$. The carrier frequency was $5.8\,\mathrm{GHz}$, approximately the highest frequency for common WiFi deployments. Each measurement is 3 seconds long with samples taken at $10\,\mathrm{MS/s}$. The AGC is fixed to a common, constant value for every measuremnt. The transmitted signal is 1029-periodic and the channel is estimated for each period of the sequence, so the effective rate at which the channel is estimated is approximately $9.7\,\mathrm{kHz}$. Channel estimation recovers a complex value for the dominant channel coefficient.

B. VNA-Based Setup

The second measurement setup consists of a Rohde & Schwarz ZNB8 2-port vector network analyzer (VNA), SkyCross 2-2931-A wide-band antennas, and a positioner that moved $1\,\mathrm{mm}$ between channel measurements. The channels are captured in a rich scattering environment with no line-of-sight component. To achieve a high SNR, the IF bandwidth was set to $1\,\mathrm{kHz}$ which resulted in a SNR in the range of $30\,\mathrm{dB}$ to $40\,\mathrm{dB}$. The VNA captured the channel at frequencies

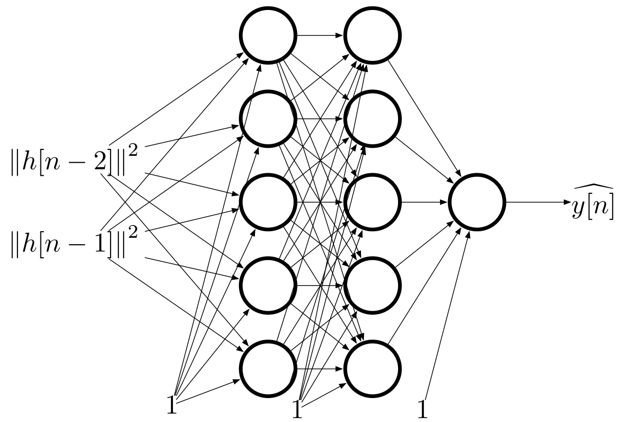

Fig. 6: Architecture of neural network used to predict channel quality. The two inputs are the power of the past two channel measurements. All layers are fully connected with bias offsets. The sigmoid function is used as the activation function for all nodes. There is an input layer, two hidden layers, and an output layer. $\widehat{y[n]}$ gives a score that can be interpreted as the estimated probability that $h[n]$ will be above the decoding threshold.

from $2.3\,\mathrm{GHz}$ to $6.0\,\mathrm{GHz}$ with a linear frequency spacing of $500\,\mathrm{kHz}$. Several campaigns were conducted to prove that the measurements are repeatable.

C. Measurement Results

Fig. 3 is the power of the estimated channel for a typical measurement from the transceiver-based setup. It is evident that the experimental setup is able to observe meaningful channel dynamics. Channel autocorrelations (specifically, circular autocorrelation) and power spectral densities are shown for the transceiver-based measurements in Fig. 4.

The spatial autocorrelation plot does not show perfect agreement, but they have similar drop-offs from the main lobe with periodic ripples after the initial drop. The power spectral density plots show good agreement. The PSD is relatively flat at low frequencies until it peaks at the maximum Doppler shift. All measurements showed peaking at approximately the same frequency. The peaks at $3\,\mathrm{Hz}$ are consistent with the velocity of the antenna, which is given by $v\frac{f_c}{c}$ $=2.9\,\mathrm{Hz}$ for the parameters in this experiment. After the peak, all measurements show the expected roll-off of $20\,\mathrm{dB/decade}$.

Captured power spectral densities are shown for the VNA-based measurements in Fig. 5. As in Fig. 4, the measurements have the expected shape with a peak at the maximum Doppler shift. Unlike Fig. 4, the figures are presented with spatial (rather than temporal) frequency. In spatial frequency, the peak location is given by $\frac{f_c}{c}$, so we expect the peak to be $8\,\mathrm{Hz}$ for the $2.45\,\mathrm{GHz}$ measurement and approximately $19\,\mathrm{Hz}$ for the $5.8\,\mathrm{GHz}$ measurement. The measurements are consistent with this. Overall, these measurements show good consistency with Jakes's model and the simulations in Fig. 2.

IV. Relay Selection

A. Problem Setup

We evaluate some simple relay selection algorithms. A time horizon of $\Delta = 6.67\,\mathrm{ms}$ is chosen, which corresponds to $1\,\mathrm{mm}$ of movement at the chosen velocity. Each prediction algorithm uses m past channel estimates spaced Δ apart, i.e. $\{h(t-\Delta), h(t-2\Delta) \ldots h(t-m\Delta)\}$, to predict if $h(t)$ is a good channel. Channel powers are used in all cases for prediction. Channel measurements are normalized and a decoding threshold is chosen such that each channel measurement has outage probability of 27%, which corresponds to a nominal SNR of $4\,\mathrm{dB}$ and decoding threshold of $0\,\mathrm{dB}$.

The measured data is preprocessed as described in Sec. IV. The input to the network is a m-tuple of channel qualities $(|h(t-i\Delta)|)_{0<i\leq m}$The expected output is a $0/1$ variable that indicates if the channel to be predicted was above or below the decoding threshold. For each position the transmitter is placed at, which we consider to be a distinct emulated relay, there are 941594 input/output pairs.

B. Prediction Algorithms

Polynomial (Lagrangian) interpolation and a simple neural net were evaluated as channel prediction algorithms. The fully-connected neural net uses $m = 2$ past channel estimates. The architecture used in [14] did not yield good results on measured data, so another hidden layer was added and the hidden layers were widened from 2 to 5. The neural net is fully connected and has two hidden layers with five nodes each, as depicted in Fig. 6. The sigmoid function is used as the activation function at each node. The output can be considered an estimate of the probability that the channel will be satisfactory, given the past channel qualities.

The data is collected with the same hardware and procedure and then divided into training, validation, and testing sets, which consist of 9, 8, and 30 relays respectively. Training is performed via standard stochastic gradient descent on the training set. Hyperparameters are optimized with a tree-structured Parzen estimator on the validation set. Each channel prediction algorithm is evaluated in two ways: single-link estimation and paired-link estimation. In the single-link case, the prediction algorithm looks at a pool of n links corresponding to n potential relays. The estimator gives a score to each channel based on past measurements and the link with the highest score is selected. The fraction of time the selected link is good is an estimate of the prediction algorithm's ability to select good links on an individual basis.

In a relay selection scenario, there are two channels that need to be good for the relay to succeed: the controller-to-relay channel and the sensor-to-relay channel. The paired-link estimation evaluation simulates this scenario. Channel measurements are assigned to each other such that two channels form a pair. The channels are individually scored by the chosen prediction algorithm and a combined score is computed as the minimum of the two scores. The pair of channels with the best combined score is evaluated, and if it is below the

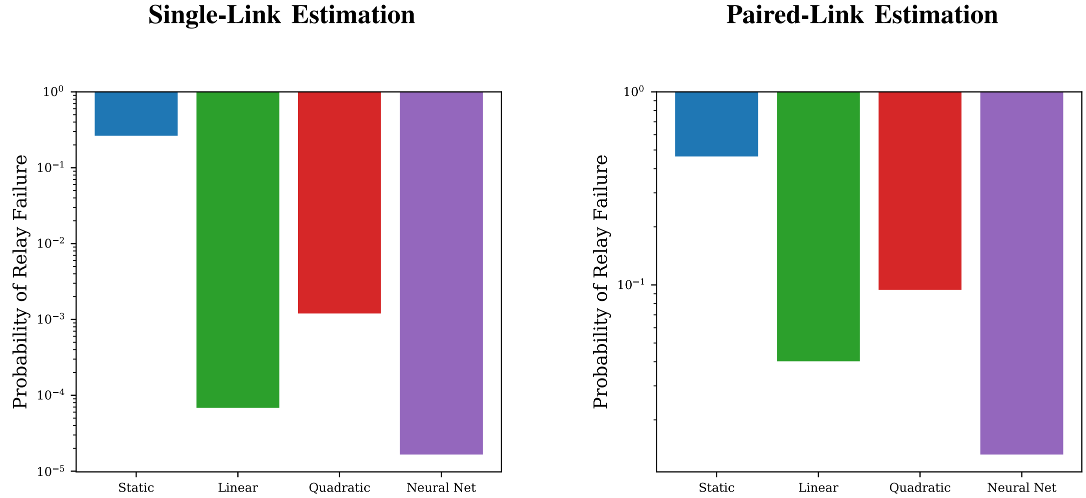

Fig. 7: Results of various channel prediction algorithms. The single-link estimation results have $n = 30$ relays and the paired-link estimation results have $n = 15$ relays for a total of 30 distinct channels.

decoding threshold we call it a failure. The fraction of time that a prediction algorithm fails is therefore an estimate for the probability that the relay selection algorithm would fail to choose a good relay.

C. Results

Both single- and paired-link estimation results are shown in Fig. 7. Static relay selection performs poorly because the outage probability is relatively high. Linear interpolation works better than quadratic interpolation, which may be explained by Lagrange interpolation's sensitivity to noise. The neural net performs best, but linear interpolation performs well too.

The paired-link estimation in Fig. 7 has worse failure probabilities than the single-link estimation. There are only 15 relays to choose between as opposed to 30 relays for single-link). Furthermore, there are two different ways to fail. However, the same trend is present for the paired-link estimation and the linear interpolation and neural net estimators are dramatically better than static relay selection.

V. Conclusion

Adaptive relay selection is a valuable technique for achieving the requirements of URLLC. Selecting good relays depends on prediction algorithms that take channel dynamics into account. This work presents measurements that show good agreement with the analytical and simulation-based models for channel dynamics. Furthermore, channel selection algorithms were applied to those measurements to show that simple algorithms can reliably select a good relay. This demonstrates that dynamic relay selection techniques can be used to achieve high reliability in real-world environments.

VI. Acknowledgements

The authors acknowledge the NSF GRFP, CNS-1321155, and ECCS-1343398 for support, as well as ADEPT and the Berkeley Wireless Research Center. This work was partially supported by the Wallenberg AI and Autonomous Systems and Software Program (WASP). Thanks to Rimas Avižienis for sharing his work with the ADI board and FPGA platform. Thanks to Simon Scott for helpful documentation and starter code for the XY table, as well as Fred Burghardt for assembly and initial setup of the XY table.

References

[1] 3GPP TR 38.913, "Study on scenarios and requirements for next generation access technologies," 2018.

[2] G.P. Fettweis, "The Tactile Internet: Applications and Challenges," *IEEE Vehicular Technology Magazine*, vol. 9, no. 1, pp. 64–70, mar 2014.

[3] V.N. Swamy *et al.*, "Cooperative communication for high-reliability low-latency wireless control," *IEEE International Conference on Communications*, vol. 2015-Septe, pp. 4380–4386, 2015.

[4] ——, "Predicting Wireless Channels for Ultra-Reliable Low-Latency Communications," in *IEEE International Symposium on Information Theory - Proceedings*, 2018.

[5] ——, "Wireless Channel Dynamics and Robustness for Ultra-Reliable Low-Latency Communications," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 4, pp. 705–720, 2018.

[6] Y. Jing and H. Jafarkhani, "Single and multiple relay selection schemes and their achievable diversity orders," *IEEE Transactions on Wireless Communications*, 2009.

[7] W.C. Jakes, *Microwave mobile communications*. Wiley-IEEE Press, 1994.

[8] A. Abdi *et al.*, "Comparison of the level crossing rate and average fade duration of Rayleigh, Rice and Nakagami fading models with mobile channel data," *Vehicular Technology Conference Fall 2000. IEEE VTS Fall VTC2000. 52nd Vehicular Technology Conference (Cat. No.00CH37152)*, vol. 4, pp. 1850–1857, 2000.

[9] R. Narasimhan, "Finite-SNR diversity-multiplexing tradeoff for correlated Rayleigh and Rician MIMO channels," *IEEE Transactions on Information Theory*, 2006.

[10] C. Xiao, Y.R. Zheng, and N.C. Beaulieu, "Novel sum-of-sinusoids simulation models for rayleigh and rician fading channels," *IEEE Transactions on Wireless Communications*, 2006.

[11] Analog Devices Inc., "HDL Reference Designs," 2014. [Online]. Available: https://github.com/analogdevicesinc/hdl

[12] P. Rigge and R. Avižienis, "adi-bbp," 2017. [Online]. Available: https://github.com/grebe/adi-bbp

[13] P. Rigge and C. Nelson, "XYTable," 2018. [Online]. Available: https://github.com/grebe/xytable

[14] V. Narasimha Swamy, "Real-time Communication Systems For Automation Over Wireless: Enabling Future Interactive Tech," Ph.D. dissertation, University of California Berkeley, 2018.