Using robot mobility to exploit multipath fading — **Source link** ⬚

Magnus Lindhé, Karl Henrik Johansson

Related papers:

- On the Spatial Predictability of Communication Channels

- Co-Optimization of Communication and Motion Planning of a Robotic Operation under Resource Constraints and in Fading Environments

- Robotic Router Formation in Realistic Communication Environments

- Communication-Aware Motion Planning in Mobile Networks

- Path Planning for Networked Robotic Surveillance

Using Robot Mobility to Exploit Multipath Fading

Magnus Lindhé and Karl Henrik Johansson[*]
ACCESS Linnaeus Centre
School of Electrical Engineering
Royal Institute of Technology
SE-100 44 Stockholm, Sweden
{lindhe|kallej}@ee.kth.se

November 10, 2008

Abstract

Communication-aware motion control allows mobile networked robots to increase the average communication throughput. We exploit that in a multipath fading channel, robots can measure the signal-to-noise ratio (SNR) and adapt their motion to spend slightly more time at positions where the channel is good. Two new such cross-layer strategies are analyzed and evaluated: periodic stopping, where the stop duration is a function of the SNR, and controlled stopping, where the robot stops when the communication buffer is filling up. It is shown that the expected average channel capacity can be twice as high as when no cross-layer information is utilized. Experimental evaluation of the strategies confirms the theoretical results.

1 Introduction

Imagine a fleet of mobile surveillance robots, guarding an oil refinery by night. The robots patrol along randomized routes, chosen to cover the whole facility, and send camera imagery back to a manned control room. Powerful computers scan the images for intruders, fire or damages and can alert the guards if anything is out of the ordinary. If an alarm is triggered, the robots can go to the contaminated area and search it in a coordinated fashion, so that an intruder cannot get back into the cleared area undetected. This allows the guards to focus on responding to more complex situations, while these dull or potentially dangerous tasks are handled by the robots.

A system like the one described above employs advanced communications and robotics. To enable coordinated control, and collection of sensor data, the robots need communication links that provide high quality of service (QoS). Important metrics are throughput, network delay and outage probability. With cameras becoming smaller and cheaper by the day, nodes in the network can be expected to deliver high bandwidth information with low delay tolerance. Achieving this in a mobile robot network poses several challenges that must be handled, including the following: First, the nodes are resource-constrained, both in terms of energy and computation power. This requires methods of scheduling sleep for sensors and transceivers, as well as cross-layer design of medium access (MAC) protocols that can adaptively trade QoS for energy [1]. Second, since the nodes are moving and are often spread out geographically, ad-hoc routing mechanisms are required that can adapt to changing network topologies [2]. Third, to improve the reliability of delivery without too many retransmissions that cost power and cause varying delays, new transport layer protocols must be developed [3].

Coordination strategies for multi-robot systems are often formulated as decentralized control laws. Information is in these systems only exchanged between neighboring robots, so local controllers have to be designed to converge to some global behavior despite the limited communication [4]. Recently, there has been a growing interest in the robotics community to study such distributed control problems under the QoS constraints described above. An example closely related to our motivating scenario is to make the robots map the environment while cooperatively searching it in minimal time, without getting too far apart [5]. A coordination problem with another type of communication constraint is considered in this paper, but first we briefly review some common models of communication in mobile robotics.

1.1   Communication Models for Robotics

Selecting a communication model for robot coordination is a tradeoff between accurately representing reality and keeping the model complexity down so that the design problem is feasible to solve. Four communication models are schematically illustrated in Figure 1, where robot 1 is exploring an office floor while staying in contact with robot 2, acting as a relay at the entrance. Figure 1a shows a binary disc graph: robots within a radius R are assumed to have perfect communication, otherwise they are disconnected. Such a model lends itself well to graph-based algorithms, popular for flocking or formation control. Using algebraic graph theory, it is possible to give conditions on the network topology that guarantee convergence to a common position or formation [6]. Figure 1b illustrates a free-space model, where communication quality decays as a function of distance. This model allows to formulate continuous tradeoffs between sensing and communication [7] or to control relay nodes that try to maintain connectivity between an exploring robot and a base station [8]. It also allows to infer the posi-
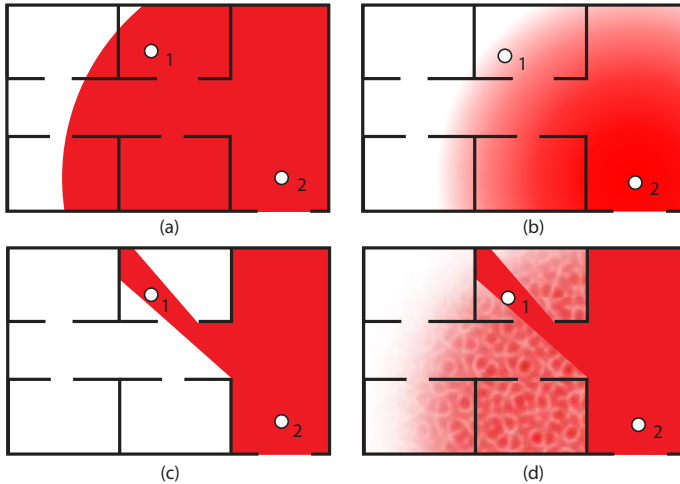
Figure 1: Robot 1 explores an office floor while staying connected with robot 2, acting as a relay at the entrance. Four communication models are illustrated: (a) binary disc graph, (b) free-space propagation, (c) shadowing and (d) multipath fading.

tion of sensors, using measurements of the signal strength [9]. Others consider shadowing, formulated as a constraint that only robots with a clear line of sight can communicate, as shown in Figure 1c. This geometric approach represents radio propagation in urban environments, especially at higher frequencies [10, 11]. A drawback is its complex dependence on the environment. In some sense, multipath fading, shown in Figure 1d, is a complement to the line of sight model, since it is most pronounced when there are obstacles between the transmitter and receiver. It accurately represents the fast changes in the signal strength due to fading, but is normally only formulated as a stochastic model since the fading is hard to predict deterministically. There are motion planning algorithms that use this model [12], adapting the step size of the algorithm so as to decrease the risk of getting stuck in deep fades. As shown in our earlier work [14], it is also possible to explicitly make use of when the fading gives a good channel. We have previously considered stopping the robot at points where the signal quality is high, depending on the amount of data waiting to be transmitted. We now propose a simpler strategy that only uses the signal quality, and compute the expected link capacity for each, compared to the case of not stopping at all. Further, the channel model is experimentally validated and the strategies are tested, using channel properties recorded from measurements.

1.2 Outline

The main objective of this paper is to describe how multipath fading can be useful for communication between robots. We will show this

Figure 2: Our example scenario: a group of robots patrolling an office floor. Each robot is streaming video to a base station, either directly or by relaying through another robot. The robots need to adapt their motion to maintain high radio throughput.

by using a scenario, inspired by the motivating applications described above. This paper is organized as follows: In the coming section, we present the scenario and the basic idea of how to exploit the fading. Then we state our models for the robot kinematics and the wireless channel. With this in place, we analyze three possible motion strategies in a multipath fading environment: no stopping, periodic stopping and controlled stopping. For each strategy, we compute the resulting average link capacity. We then illustrate these strategies by experiments, using actual recorded channel properties. Finally, we end with some conclusions.

2 Exploiting Multipath Fading in Robot Coordination

To illustrate how multipath fading can be exploited, we will use the scenario in Figure 2: A group of robots is used for surveillance of an office floor during night. The robots patrol the offices along given paths and stream camera images to a base station, which in turn feeds the data to an operator. Robots far from the base station use multihop relaying to be able to cover all rooms. The exact motion timing is not important, as long as each floor is cleared on time. It is, however, crucial that the robots maintain low latency and high throughput for the link to the base station. Otherwise, the video images will be noisy and intruders will have time to escape before being detected.

This freedom in time allows to modify the motion of each robot to improve communications, at the expense of not perfectly tracking the pre-planned trajectory. In an environment with multipath fading, this could mean spending slightly more time at positions that offer low channel attenuation and quickly passing points where the channel is worse. Since the fading varies over distances of a wavelength, finding such positions only requires small deviations along the trajectory. We will present and analyze methods for doing this tradeoff between com-

munication and tracking, under different assumptions on what feedback the robot gets from its radio. To simplify the presentation, we will consider the case of a single robot communicating with a base station. The strategies presented here for point-to-point communication provide the basic functionality needed by higher-layer protocols for maintaining connectivity within a whole group of robots.

3 Robot and Channel Models

In this section, we state our model of the robot and reduce it to one-dimensional motion along the reference trajectory. We also introduce a channel model of static Rayleigh fading. The model is validated through measurements in our lab. Finally, we define the link capacity as the byte reception rate, which will be used to compare different motion strategies.

3.1 Robot Model

The position of the robot is $q \in \mathbb{R}^2$. We assume that it has a pre-planned time-stamped reference trajectory $q_{\text{ref}}(t)$, moving at a velocity $v_{\text{ref}}(t)$, and a controller for following it. This allows us to reduce the problem to considering the one-dimensional motion of the robot along the reference trajectory. Let Δ be the position of the robot along the trajectory, relative to the reference, so that $\Delta > 0$ means that the robot is going ahead of the reference. Also let φ be the relative velocity. Stopping the robot can be done by applying breaks or short-circuiting the motors, which does not consume battery power. We model the motion control as a hybrid system, where the robot can be in one of two modes, $\sigma = \texttt{stop}$ or $\sigma = \texttt{drive}$. The dynamics of the one-dimensional motion are then

$$\sigma = \texttt{stop:} \begin{cases} \dot{\Delta} = \varphi \\ \dot{\varphi} = -k_v(\varphi + v_{\text{ref}}) \end{cases} \qquad \sigma = \texttt{drive:} \begin{cases} \dot{\Delta} = \varphi \\ \dot{\varphi} = a, \end{cases}$$

where the controls are $u = (a, \sigma)$ and $k_v \gg 1$ is chosen to model the robot stopping quickly.

3.2 Channel Model

We assume a Rayleigh fading environment where the SNR, γ, is exponentially distributed with average γ_0. We consider trajectories where the distance to the receiver does not change significantly and there is constant shadowing, so γ_0 does not change over time. In a deployed system, the approach presented here could therefore be complemented by other components that avoid shadowing from obstacles and adapt the large-scale motion of the robot to limit the path loss.

Since multipath fading is caused by multiple reflections of the signal against objects, one can expect that if nothing in the environment moves, the resulting fading should not change over time, but only as
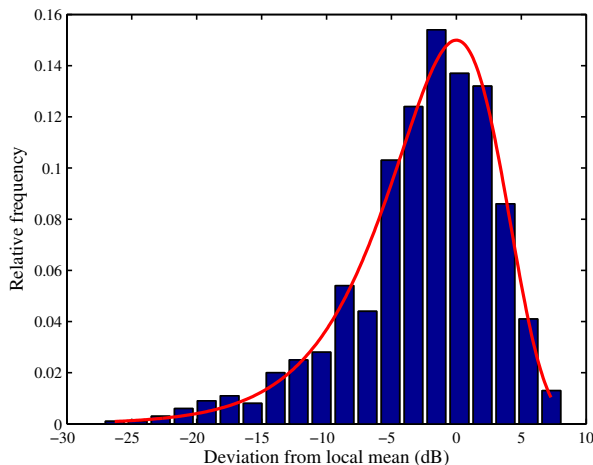
Figure 3: Normalized histogram of 1000 samples, taken 1 cm apart in our lab. The probability distribution function of Rayleigh fading (in dB) is included as a reference.

a function of the position of the transmitter and receiver. Successive minima occur about every half wavelength [13]. This is a reasonable assumption in applications such as nighttime surveillance, rescue missions in collapsed buildings or military exploration of possibly hostile environments. As described in the experiment section, we have validated the model by measurements which gave the histogram in Figure 3. The figure also shows the ideal Rayleigh distribution function, which fits the measurements well. By measuring the change in SNR over time when the transmitter and receiver were not moving, it was confirmed that the fading does not change in a static environment.

To allow comparison between different control strategies, we define the normalized link capacity $c(\gamma)$ as the probability of correct reception of one byte when the SNR is γ. Other choices could be to study the bit error probability or packet reception rate for packets of several bytes, but we believe that the byte capacity is an illustrative measure of the link performance. It also does not need assumptions on protocol issues such as packet size, error-correcting codes or retransmission schemes, which can be used to improve the performance on a packet level.

4 Communication-Aware Motion Control

In this section, we describe an architecture for the interconnection of the robot platform and radio hardware. We then analyze three levels of feedback from the radio link and the sensor buffer to the robot, yielding different control strategies: First, no stopping. Second, periodic stopping where the SNR can be sampled when standing still. Third, controlled stopping, using continuous measurement of the SNR
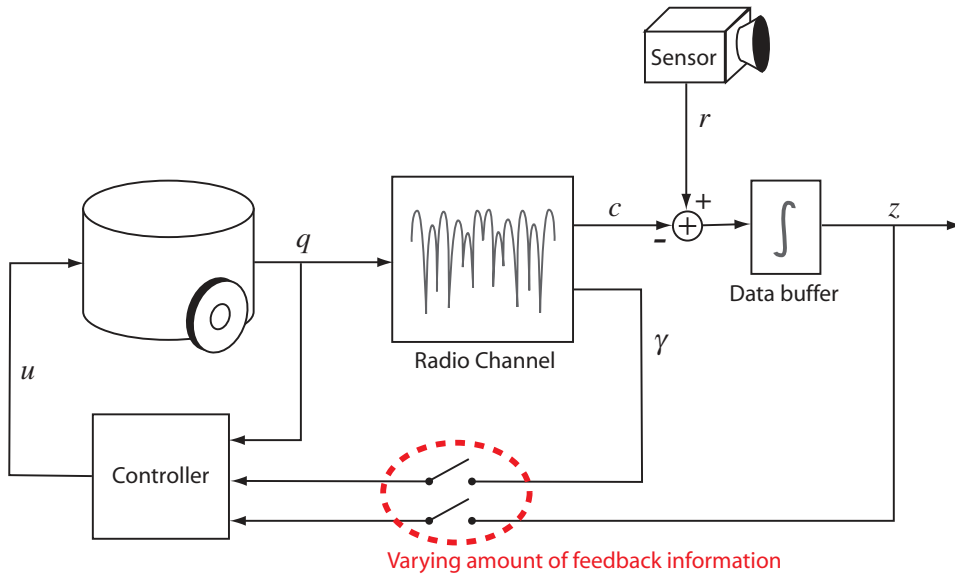
6

Figure 4: Architecture of the system, with a robot platform, a radio and a controller. The radio buffers data from a sensor and sends it through a wireless channel, whose capacity depends on the position of the robot. We present strategies for the controller with and without feedback on the SNR γ and the state z of the data buffer.

and the amount of data waiting to be transmitted. For each strategy, we analyze the expected link capacity.

4.1 System Architecture

The overall system is comprised of the robot platform and the radio. As illustrated in Figure 4, the position q of the platform determines the SNR, γ, and channel capacity, c. A buffer in the radio stores data arriving at a rate r from a sensor. The size of the buffer is $z \geq 0$ and its dynamics are

$$\dot{z} = r - c.$$

We will first consider the case when the controller has no feedback from the radio or buffer, and thus follows the reference trajectory without stopping. Then we will study the periodic stopping strategy, where the loop is closed between the radio channel and the motion of the platform. It is assumed that γ can be sampled only when standing still, as is the case in many slowly sampling radio transceivers. Finally, we will consider controlled stopping, where the controller has access to continuous measurements of both γ and z, so it can stop when needed and find local maxima of γ.

7

4.2 No Stopping

With no feedback from the radio, we assume that the robot drives along the desired path without adjusting its motion to the radio channel. The control law is

$$u = u(q).$$

The nominal link capacity as a function of the average SNR can be described as $c_{\mathrm{drive}} = E\{c(\gamma)|\sigma = \mathtt{drive}\}$. We will use this as a baseline, to compare against the more advanced motion strategies below.

4.3 Periodic Stopping

If the radio hardware needs some time to sample the capacity, the robot has to stand still to avoid the channel changing. A possible approach is then to schedule periodic stops and use the measured SNR to determine the length of the stop. The control law is now

$$u = u(q, \gamma).$$

We suggest the following strategy: The robot drives at velocity $2v_{\mathrm{ref}}$ for a constant time, τ_{drive}. Then it stops, measures the SNR and determines the length τ_{stop} of the stop. After waiting τ_{stop}, it starts driving again. To get the desired average velocity, we require that the expected stop time is equal to the drive time, $i.e.$, $E\{\tau_{\mathrm{stop}}(\gamma)\} = \tau_{\mathrm{drive}}$.

We have investigated two candidates for the function $\tau_{\mathrm{stop}}(\gamma)$: a linear policy and a threshold policy. The linear policy can be expressed as

$$\tau_{\mathrm{stop}}(\gamma) = \max\{0, \lambda(\gamma - \gamma_0)\},$$

which achieves the desired average velocity if $\lambda = \tau_{\mathrm{drive}}\gamma_0^{-1}e$, where e is the base of the natural logarithm.

The threshold policy is to have a constant stopping time and stop if γ is higher than some threshold value γ_{th}:

$$\tau_{\mathrm{stop}}(\gamma) = \left\{ \begin{array}{ll} \alpha\tau_{\mathrm{drive}} & \text{if } \gamma > \gamma_{th} \\ 0 & \text{else,} \end{array} \right.$$

where $\alpha > 1$. The choice $\gamma_{th} = \gamma_0 \ln \alpha$ ensures that we get the desired average velocity. Increasing the parameter α means making fewer but longer stops, which increases the resulting link capacity at the expense of larger deviations from the reference trajectory. The extreme policy $\alpha \to \infty$ corresponds to making a single very long stop at the point where the signal strength is the highest. However, high values of α make the strategy very sensitive to errors in the channel model, so we have found $\alpha = 4$ to work well in our experiments. The expected resulting link capacity for each policy will be computed in the comparison below and typical trajectories will be illustrated in the experiments.

4.4 Controlled Stopping

If the controller measures the SNR and the amount of buffered sensor data continuously, it can choose to stop the robot at local maxima

of the link capacity whenever it needs to communicate with a higher capacity than c_{drive}. But stopping also carries a cost in that the robot falls behind the reference position. This suggests an adaptive strategy that makes the robot stop to communicate when the buffer starts filling up and then makes it drive to catch up again with the reference. The control law can be written as

$$u = u(q, \gamma, z).$$

Since local maxima of the capacity are less than a wavelength apart, we assume that it takes negligible time to find a point where the capacity is greater than or equal to some value c_{stop}. For the problem to be feasible, we also assume $c_{\text{drive}} < r < c_{\text{stop}}$.

The system can be approximated as a linear system, switching between two modes: Either it drives and communicates with link capacity c_{drive}, or it stands still and communicates with a higher capacity c_{stop}. To compute a controller for both the mode switching and the continuous control when driving, we formulate an optimal control problem:

$$\min_{\sigma, a} \quad \int_0^\infty \left(x^T Q x + R a^2 \right) e^{-\beta t} dt$$
$$\text{s.t.} \quad \dot{x} = A_\sigma x + B_\sigma a,$$

where the state vector x contains the buffer size z and relative position Δ and the controls are the acceleration a and mode σ. The weights $Q \succeq 0$ and $R > 0$ are used to penalize buffer size, deviation from the reference position and acceleration. If R is large compared to Q, the resulting controller will give a smooth trajectory with little acceleration, while if R is small, the result will be longer stops and higher acceleration to catch up with the reference. The factor $e^{-\beta t}$, where $\beta > 0$, is needed for the numerical procedure to converge. We solve the problem using relaxed dynamic programming, which results in a feedback controller, $(a^*(x), \sigma^*(x))$, given in the form of a look-up table that can be precomputed and stored on a resource-constrained robot [14].

The controller drives the system towards periodically switching between `drive` and `stop`. The controller adapts the duty cycle of the switching to balance the outflow and inflow of the buffer. While doing this, it also makes a tradeoff between deviation from the reference position and control effort, which affects the switching frequency. The resulting link capacity will be computed below, and an example trajectory will be shown in the experiment section.

4.5 Comparison

The presented strategies above, no stopping, periodic stopping and controlled stopping, are examples of the tradeoff between communication performance and reference tracking. Reference tracking imposes timing constraints on the motion, so the robot can only stop and communicate long enough for it to be able to catch up with the reference

afterwards. Power is also a concern, since stopping often or for long times requires more energy for catching up. To assist an application developer in making the proper tradeoff, the expected link capacities for each strategy are illustrated in Figure 5, as a function of the average SNR, γ_0. We have assumed non-coherent frequency shift keying, with a bit error rate of $\frac{1}{2}e^{-\frac{\gamma}{2}}$, so

$$c(\gamma) = \left(1 - \frac{1}{2}e^{-\frac{\gamma}{2}}\right)^8.$$

Periodic stopping with a threshold policy is illustrated for $\alpha = 4$, *i.e.*, when the fixed stop time is four times as long as the drive time. As expected, the more complex controller for stopping on demand can achieve the highest capacity, at the expense of reference tracking and locomotion power. The graph also illustrates how all strategies using feedback from the radio give capacity improvements in the transition region where the signal is getting weaker, but make no difference if the signal is very strong or very weak. For example, the periodic stopping strategy with a linear stop time policy gives improvements of over 100% compared to constant motion in the interval -6 dB$< \gamma_0 <4$ dB.

The graphs for periodic stopping can be derived by first computing the expected amount of data transmitted per stop, which is $E\{c(\gamma)\tau_{\text{stop}}(\gamma)\}$. This gives an expected average link capacity of

$$E\{c(\gamma)\} = \frac{E\{c(\gamma)\tau_{\text{stop}}(\gamma)\} + c_{\text{drive}}\tau_{\text{drive}}}{2\tau_{\text{drive}}}.$$

For the strategy of controlled stopping, the controller can switch between two link capacities: c_{drive} and a higher c_{stop}. It adaptively sacrifices reference tracking and locomotion power to adapt the average link capacity to the buffer inflow, so if the control signal is unlimited, the asymptotic maximum average capacity is c_{stop}. This capacity is achieved by the controller finding local maxima of the SNR and stopping there. Based on experience from measurements, we assume that these local maxima correspond to the 90th percentile of the capacity, which is the level plotted in Figure 5.

5   Experimental Results

To evaluate the strategies presented above under realistic channel conditions, we have used a robot to measure the actual signal strength fluctuations as a function of position in our lab. Each strategy has then been simulated, using the measured sequence to compute the channel capacity along the trajectory. As described in the modeling section, we have also used the collected data to validate the model of static Rayleigh fading.

5.1   Experimental Setup

Our measurement robot has unicycle kinematics and a laptop onboard for control. To the laptop was connected a TMote Sky sensor node,
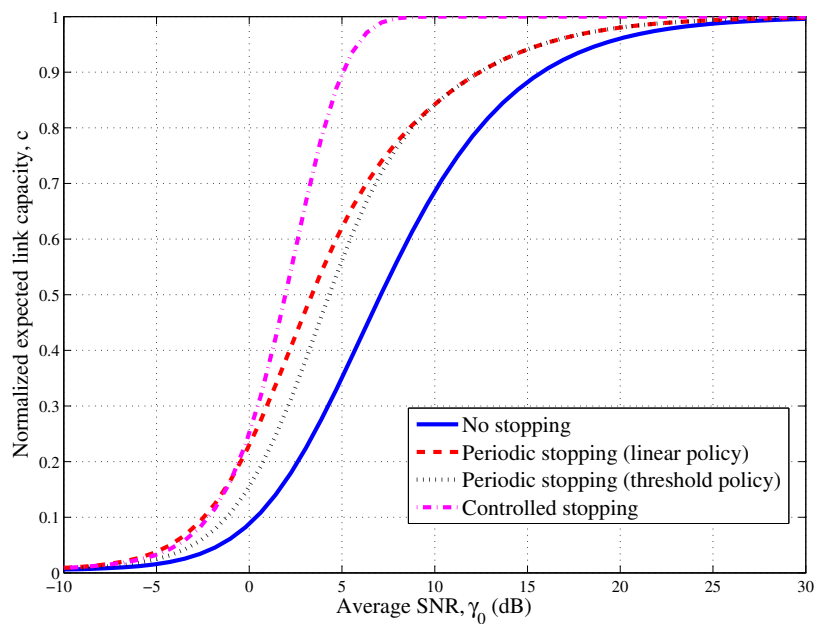
Figure 5: Expected normalized link capacity for the strategy of periodic stopping, with either linear or threshold stop time policies, and for controlled stopping. The strategy of no stopping is included as a reference. The curve for controlled stopping is an asymptotic upper bound since the controller adapts the capacity to what is needed to keep the buffer size bounded.

equipped with a IEEE 802.15.4 compliant CC2420 2.4 GHz transceiver. The CC2420 has a detector for received signal strength (RSS) with a stated accuracy of 6 dB, but our experience is that the relative accuracy of the detector is in the same range as the resolution, which is 1 dB. Another TMote Sky was used as test transmitter, sending 64 50-byte packets per second at a data rate of 250 kbit/s. The transmitter was placed on a height of about 2 m in one end of the lab and the robot was placed in the other end, with the receiver at a height of about 0.3 cm. There was no direct signal path, since the lab is full of computers and other metal equipment that effectively scatters the signal. The lab was unoccupied during the measurements, to mimic the conditions in an office at night.

To record the signal strength, the robot was driven 1 cm at a time, then stopped and recorded the number of received packets for 1 s. At each position, the average received signal strength (RSS) of the packets was recorded. Due to lack of space, the robot drove along a straight line, stopped after 2.5 m and was turned manually in place to follow a new line. We recorded 1000 samples with an average RSS of -67 dBm. A histogram of the RSS, with the average subtracted, was plotted in Figure 3.

Finally, we also tested the assumption on the fading being static over time, by moving the transmitter to different locations and measuring the RSS over 3 min with the robot standing still. Then the RSS was registered separately for each packet, with no averaging. The standard deviation was 1.1 dB or less for all measurement series. No outliers were more than 2.7 dB from the average.

5.2   Results

Figure 6 shows the simulation results for periodic stopping with a linear stop time policy (left) and a threshold policy (middle), as well as for controlled stopping (right). For the threshold policy, we used $\alpha = 4$. It is assumed that the reference position is moving along the x-axis at 0.1 m/s and for each control strategy, we have plotted the position q_x of the robot, as well as the reference position (dashed). The buffer size z is also illustrated, for an inflow of $r = 0.6$.

Periodic stopping with a linear stop time policy gave a 71% improvement in average link capacity compared to no stopping. The corresponding improvement for periodic stopping with a threshold policy was 69%. The strategy of controlled stopping gave a link capacity equal to the inflow, since the buffer was almost empty at the end of the simulation. This means a 67% improvement over no stopping. Figure 6 also illustrates how controlled stopping actively keeps the buffer size low and that it results in better reference tracking than the periodic stopping strategies. It can be noted that, for a FIFO buffer, the size of the buffer is a measure of the data latency. So controlled stopping also achieves the lowest data latency.

Controlled stopping assumes that a high signal strength can be found immediately when stopping. To simulate this, whenever the
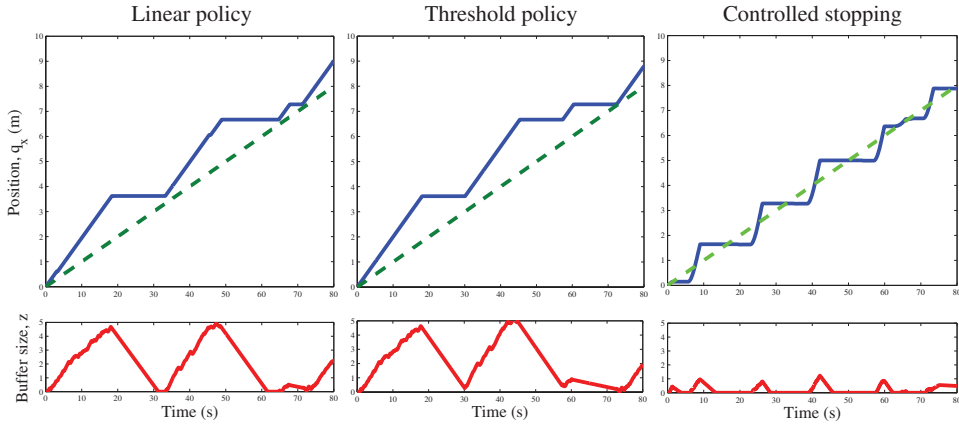
Figure 6: Resulting trajectories when simulating the proposed motion strategies, using channel properties recorded by measurements. The dashed lines represent the reference position, moving at constant velocity, and the lower curves illustrate the buffer size. Periodic stopping with a linear stop time policy (left) or a threshold policy (middle) give larger buffer sizes and worse reference tracking than the controlled stopping strategy (right).

controller decided to switch to the `stop` mode, the robot sampled the channel every 0.1 s and was forced to stay in the `drive` mode until it found a position where $c \geq r$. At some parts of the trajectory the SNR was lower, which caused the robot to overshoot the reference trajectory some when driving in search of a good enough position. Guided by Figure 5, we used the parameters $c_{\mathrm{drive}} = 0.36$ and $c_{\mathrm{stop}} = 0.9$.

To reduce the influence of possible interference, the measurements were performed using the highest possible transmission power, 0 dBm. But, as commented above, stopping strategies give the best result when the link is on the limit of losing contact. To better illustrate this, we have assumed a high noise level so that the average SNR for the simulation becomes 5 dB.

6 Conclusions

We have analyzed and evaluated methods to improve the capacity of wireless robot communication, in environments that exhibit multipath fading. The main idea is to exploit freedom in moving the robot, to make it stop and communicate at positions where the channel is better. Two main strategies were considered: Periodic stopping and controlled stopping. These strategies make different assumptions on the information available to the controller from the radio and also yield controllers of different complexity.

Theoretical analysis, assuming Rayleigh fading, shows that both strategies can give significant improvements of the channel capacity

compared to no stopping, using no feedback from the radio. The more complex strategy, controlled stopping, can achieve the highest improvement by adaptively sacrificing reference tracking. It is also important to note that these methods contribute the most in the transition region where the channel capacity starts to decay, but make no difference if the signal is very strong or very weak. Simulations, using actual channel properties, show that the strategies work also under more realistic conditions. The periodic stopping strategy appears to be more robust to errors in the channel model than the controlled stopping strategy.

This method could be applied also in high bandwidth systems such as video links, where the fading may be frequency selective. The receiver could then be equipped with an equalizer and the SNR after equalization could be fed back to the motion controller. The resulting SNR may not be Rayleigh distributed, but the statistical analysis could be adapted to the new distribution.

We end by noting that the methods presented here could be combined with other approaches to mitigate multipath fading, such as antenna or frequency diversity. Antenna diversity is achieved by placing multiple antennas far enough apart for them to experience uncorrelated fading. That closely parallels the presented methods, where instead a single antenna is moved between sampling instances. Depending on the available mounting space for antennas on the robot and the tracking error that can be tolerated, antenna diversity and communication-aware motion therefore complement each other to improve communication performance.

References

[1] A. Fallahi and E. Hossain, "QoS provisioning in wireless video sensor networks: a dynamic power management framework," *IEEE Wireless Communications*, vol. 14, no. 6, 2007.

[2] K. Akkaya and M. Younis, "A survey on routing protocols for wireless sensor networks," *Ad Hoc Networks*, vol. 3, no. 3, 2005.

[3] I. A. Akyildiz, T. Melodia, and K. R. Chowdury, "Wireless multimedia sensor networks: a survey," *IEEE Wireless Communications*, vol. 14, no. 6, 2007.

[4] R. M. Murray, "Recent research in cooperative control of multivehicle systems," *Journal of Dynamic Systems, Measurement, and Control*, vol. 129, no. 5, 2007.

[5] W. Burgard, M. Moors, C. Stachniss, and F. Schneider, "Coordinated multi-robot exploration," *IEEE Transactions on Robotics*, vol. 21, no. 3, 2005.

[6] M. M. Zavlanos and G. J. Pappas, "Distributed connectivity control of mobile networks," in *Proceedings of the IEEE Conference on Decision and Control*, New Orleans, USA, 2007.

[7] T. Chung, J. Burdick, and R. Murray, "A decentralized motion coordination strategy for dynamic target tracking," in *Proceedings*

of the IEEE International Conference on Robotics and Automation, Orlando, USA, 2006.

[8] E. Stump, A. Jadbabaie, and V. Kumar, "Connectivity management in mobile robot teams," in *Proceedings of the IEEE International Conference on Robotics and Automation*, Pasadena, USA, 2008.

[9] E. Elnahrawy, X. Li, and R. Martin, "The limits of localization using signal strength: a comparative study," in *Proceedings of the IEEE Conference on Sensor and Ad Hoc Communications and Networks*, Santa Clara, USA, 2004.

[10] A. Ganguli, J. Cortés, and F. Bullo, "Distributed deployment of asynchronous guards in art galleries," in *American Control Conference*, Minneapolis, USA, 2006.

[11] J. M. Esposito and T. W. Dunbar, "Maintaining wireless connectivity constraints for swarms in the presence of obstacles," in *Proceedings of the IEEE International Conference on Robotics and Automation*, Orlando, USA, 2006.

[12] Y. Mostofi, "Communication-aware motion planning in fading environments," in *Proceedings of the IEEE International Conference on Robotics and Automation*, Pasadena, USA, 2008.

[13] W. C. Jakes, Ed., *Microwave Mobile Communications*. IEEE Press, 1974.

[14] M. Lindhé and K. H. Johansson, "Communication-aware trajectory tracking," in *Proceedings of the 2008 IEEE International Conference on Robotics and Automation*, Pasadena, USA, 2008.