Rate adaptation using acknowlegement feedback: throughput upper bounds

Rate Adaptation using Acknowledgement Feedback: Throughput Upper Bounds

Chin Keong Ho
Institute for Infocomm Research, A*STAR
1 Fusionopolis Way, Singapore
Email: hock@i2r.a-star.edu.sg

Job Oostveen
TNO Information and Comm. Technology
Groningen, The Netherlands
Email: job.oostveen@tno.nl

Jean-Paul Linnartz
Philips Research Laboratories
Eindhoven, The Netherlands
Email: j.p.linnartz@philips.com

Abstract—We consider packet-by-packet rate adaptation to maximize the throughput over a finite-state Markov channel. To limit the amount of feedback data, we use past packet acknowledgements (ACKs) and past rates as channel state information. It is known that the maximum achievable throughput is computationally prohibitive to determine. Thus, in this paper we derive two upper bounds on the maximum achievable throughput, which are tighter than previously known ones. We compare the upper bounds with a known myopic rate-adaptation policy. Numerical studies over a wide range of SNR suggest that the myopic rate-adaptation policy is close to the upper bounds and may be adequate in slowly time-varying channels.

I. INTRODUCTION

Automatic repeat request (ARQ) [1], [2] is commonly used to enhance the reliability or the throughput of packet-switched systems. When the channel experiences an instantaneous deep fade or strong interference, a packet cannot be recovered. An explicit negative acknowledgement (NACK), or a missing positive ACK (PACK), is then used to request a retransmission by the receiver to the transmitter. To efficiently use the channel, the rate at which each packet is encoded, i.e., the modulation constellation and code rate used, should ideally match the instantaneous channel condition. This rate adaptation poses a challenging tracking problem for time-varying channels.

To perform rate adaptation, channel state information (CSI) is needed. Although more informative CSI leads to better channel tracking and hence higher throughput, in practice the availability of CSI is limited by the communication system employed. Rate adaptation can be implemented for any ARQ system if only the history of ACKs is used as CSI, e.g., [3]–[5], without assuming channel reciprocity and availability of additional CSI. Additionally, rates used for previous transmissions can be considered as CSI [6], [7], since the rates are already known at the transmitter and need only to be stored in memory. In [6], the CSI is limited to past rates and ACKs in the same *frame*, where typically a frame consists of several packets. In [7], *all* past rates and ACKs are used as CSI, which improves the tracking of the channel quality; for brevity we refer to this as *ACK-rate CSI*.

In this paper, to match variations in channel conditions, we employ rate adaptation and seek to maximize the throughput averaged over an infinite time horizon. To limit the feedback, ACK-rate CSI is employed, similar to [7]. This is, however, a PSPACE-complete problem [7], [8], which is considered

at least as hard as an NP-complete problem. This means that optimal rate adaptation schemes cannot be implemented practically. Hence, rate adaptation policies based on heuristics are typically devised, see e.g., [7], [9]. So far, the maximum achievable throughput has not been quantified numerically. As such, an upper bound on the maximum achievable throughput is obtained in [7]. A tighter upper bound may be desirable to provide a more accurate indication of how close a rate adaptation scheme performs with respect to the optimal one.

Our contribution pertains in that we establish two new upper bounds that are tighter than currently known ones. To obtain these upper bounds, we let the transmitter receive a CSI that is more informative than ACK-rate CSI. Specifically, we periodically update the transmitter with a delayed version of the exact channel coefficient, in addition to the ACK-rate CSI. We compare the upper bounds with a myopic rate-adaptation policy [9], that seek to maximize only the current throughput, without regarding how the future throughput is affected. Numerical studies show that this policy performs within one dB of signal-to-noise ratio (SNR) to the derived upper bounds over a wide range of throughputs at a channel power correlation of 0.99. Moreover, these upper bounds tighten the closest known bound [7] by about half a dB.

For our analysis, we use a first-order finite-state Markov channel (FSMC) [10] to model the variations of the channel over time, which allows us to obtain tractable results and to build insights. For simplicity, we assume that a packet is recovered if and only if the SNR exceeds a threshold related to the rate. This model is reasonably accurate for a large class of coding and modulation schemes. Finally, we assume that the buffer for storing information bits at the transmitter has infinite size and always contains sufficient bits. This is appropriate if many information bits are already pre-stored at the transmitter, such as in streaming applications. In contrast, in [7] the buffer is finite in size and arriving bits may subsequently be dropped.

This paper is organized as follows. Section II describes the system model. Section III proposes two upper bounds for the problem of maximizing the throughput. Numerical results are obtained in Section IV. Section V concludes the paper.

II. SYSTEM MODEL

The system model is depicted in Fig. 1. At time $k = 1, 2, \cdots$, the CSI available for rate adaptation is denoted as

Fig. 1. System model for the rate adaptation process.

$\mathbf{c}_k$, which will be specified subsequently. The time index coincides with the packet index for simplicity. Based on the CSI, the rate adaptation block selects the rate R_k from a finite set $\mathcal{S}_R$ to transmit packet k. The rate determines the coding rate and modulation scheme used. We express the rate in bits per symbol, i.e., in bits per channel use. After the rate is determined, the buffer collects $R_k N_s$ bits from a source. The bits are then encoded and decoded independently for each packet, even in retransmissions. This ARQ scheme is commonly known as a Type I hybrid ARQ scheme [11]. The buffered bits are encoded as a codeword $\mathbf{x}_k \in \mathbb{C}^{N_s}$ consisting of N_s symbols, each of unit power, then sent as packet k.

We consider a flat-fading channel with channel amplitude $H_k \geq 0$ that varies (slowly) between packets but is time invariant during each packet duration. The received codeword is given by

$$\mathbf{y}_k = H_k e^{j\theta_k} \mathbf{x}_k + \mathbf{n}_k, \quad k = 1, 2, \cdots, \tag{1}$$

where $\mathbf{n}_k$ is a i.i.d. circularly symmetric zero-mean unit-variance complex additive white Gaussian noise (AWGN) vector of length N_s. We consider coherent detection by a receiver that knows and corrects the channel phase variations, so for simplicity we let $\theta_k = 0$. The rate is assumed to be known at the receiver for decoding, e.g., via the packet header. To determine whether a packet error occurs at the receiver, error detection is carried out at the packet level, usually by a cyclic redundancy check (CRC). The transmitter is informed whether the previous packet $k-1$ has been received correctly via an ACK bit A_{k-1}, which is either a positive ACK (PACK) $A_{k-1} = 1$ or a negative ACK (NACK) $A_{k-1} = 0$. We assume that all ACKs are received error free.

A. Channel Statistics

For analytical tractability, we use the FSMC to model temporal variations of the channel amplitude H_k [10]. That is, H_k is in a discrete set $\mathcal{H} = \{H^1, \cdots, H^{N_{\mathcal{H}}}\}$, where $N_{\mathcal{H}}$ is the total number of channel states. Further, H_k has the Markov property that $p(H_k|H_1, \cdots, H_{k-1}) = p(H_k|H_{k-1})$, where $p(\cdot)$ is the probability mass function (pmf).

The bivariate (continuous) Rayleigh distribution $p(\tilde{H}_k, \tilde{H}_{k-1})$ is fully determined by the power correlation coefficient [12] $\bar{\rho} = \text{cov}(\tilde{H}_k^2, \tilde{H}_{k-1}^2)/\sqrt{\text{var}(\tilde{H}_k^2)\text{var}(\tilde{H}_{k-1}^2)}$. The closer $\bar{\rho}$ is to one, the slower the channel variation. For our simulations, we model the FSMC such that the channel is approximately Rayleigh distributed. We assume a stationary

fading channel *discretized* from the Rayleigh distribution with average SNR $\bar{\gamma} = \mathbb{E}[|\tilde{H}_k|^2]$, according to [10]. Specifically, the Rayleigh distribution is first truncated to lie between $\tau_0 = 0$ and a maximum value of $\tau_{N_{\mathcal{H}}} = \alpha$. Then, the discrete set of channel states $\mathcal{H}$ is chosen such that each channel state occurs with equal probability. To accurately model the Rayleigh distribution in our simulations, both parameters $N_{\mathcal{H}}$ and α are taken to be large. Our subsequent analysis, however, applies generally for any FSMC.

In our analysis, we assume that the local-mean parameters $\bar{\gamma}$ and $\bar{\rho}$ are known to the transmitter. These parameters describe the long-term statistics of the channel and can be accurately estimated given sufficient time.

B. CSI

We initialize the ACK and rate as $A_0 = \emptyset, R_0 = \emptyset$, respectively, where $\emptyset$ is the null value. The initial channel amplitude H_0 is randomly generated using the discretized Rayleigh distribution. We collect all ACKs until time k as vector $\mathbf{a}_k \triangleq [A_0, A_1, \cdots, A_k]$, and similarly all rates and channel amplitudes until time k as $\mathbf{r}_k, \mathbf{h}_k$, respectively.

The receiver always has exact knowledge of the channel state. We study the maximum achievable throughput when the *ACK-rate CSI* $\mathbf{c}_k$ is available at the transmitter at time k:

$$\mathbf{c}_k = \{A_{k-1}, R_{k-1}, \mathbf{c}_{k-1}\}, \tag{2a}$$

or equivalently
$$\mathbf{c}_k = \{\mathbf{a}_{k-1}, \mathbf{r}_{k-1}\}. \tag{2b}$$

This CSI, also depicted in Fig. 1, is the primary focus of our study. The CSI consists of a recent update of the rate and the ACK and past CSIs, all available in a causal manner.

C. Throughput

If packet k is received correctly, the contribution of the packet to the throughput equals the data rate R_k. This occurs when $A_k = 1$. If $A_k = 0$, the packet is lost in an outage and it is discarded (a common practice in delay-sensitive applications) or retransmitted. In both cases the instantaneous throughput is zero. Given H_k, the expected throughput for packet k encoded at rate R_k is thus

$$t(R_k, H_k) = R_k p(A_k = 1|R_k, H_k). \tag{3}$$

Since the channel is not known exactly but instead a CSI $\mathbf{c}_k$ is given, the expected throughput becomes

$$T(R_k; \mathbf{c}_k) = \sum_{H_k} p(H_k|\mathbf{c}_k)\, t(R_k, H_k) = \mathbb{E}_{H_k|\mathbf{c}_k}[t]. \tag{4}$$

Let $C(H) = \log_2\left(1 + |H|^2\right)$ denote the channel capacity. For our simulations, we assume that the packet is recovered correctly, i.e., $p(A_k = 1|R_k, H_k) = 1$, if and only if $R_k < C(H)$.

2

(a) ACK-rate CSI; see Fig. 1.

(b) Full CSI.

(c) Delayed CSI.

(d) No CSI, $K = 3$.

(e) Periodic CSI with period $P = 2$.

Fig. 2. Causal diagrams illustrates the dependence of channel H_k, ACK A_k and rate R_k as time progresses initially during rate adaptation. Different CSIs are available at the transmitter (a)-(e). In all cases, the channel is Markovian and the ACK depends on the rate and channel, but the rate depends on the specific type of CSI available.

III. MAXIMIZING INFINITE-HORIZON THROUGHPUT

A. Problem Formulation

A rate adaptation policy assigns the rate R_k to use given the CSI $\mathbf{c}_k$ for all k. Given policy π, the throughput averaged over an infinite-time horizon is given by

$$\mathcal{T}(\pi) = \lim_{K \to \infty} \frac{1}{K} \mathbb{E} \left[\sum_{k=1}^{K} T(R_k; \mathbf{c}_k) \right]. \quad (5)$$

The expectation is carried out over all random variables, namely channel amplitudes $\{H_k\}$, rates $\{R_k\}$ and CSI $\{\mathbf{c}_k\}$.

We consider the problem of maximizing $\mathcal{T}(\pi)$ by varying π. We denote the optimal rate adaptation policy as π^*, and the maximum achievable throughput as $T^*_{\text{ACK-rate}}$, i.e.,

$$T^*_{\text{ACK-rate}} = \max_{\pi} \mathcal{T}(\pi) = \mathcal{T}(\pi^*). \quad (6)$$

In general, the superscript * denotes optimality while the subscript denotes the type of CSI used. The computation of (6) is, however, a PSPACE-complete problem [7], [8]. In this paper, we seek to instead find upper bounds for $T^*_{\text{ACK-rate}}$.

B. Other CSI

To upper bound $T^*_{\text{ACK-rate}}$, we consider the use of more informative CSI, namely full CSI, delayed CSI and periodic CSI. To illustrate the potential gain of having CSI, we also consider the case when no CSI is available.

1) Full CSI: $\mathbf{c}_k = H_k$. The instantaneous channel amplitude H_k is provided as the CSI.

2) Delayed CSI: $\mathbf{c}_k = H_{k-1}$. Due to causality, full CSI cannot be provided in practice. Here, a delayed version of the channel amplitude (where the delay is one packet long) is provided as the CSI.

3) Periodic CSI: The transmitter is updated periodically at time k, with period P, a delayed channel amplitude H_{k-1}. Additionally, some ACK-rate CSI is available. Specifically,

$$\mathbf{c}_k = \big\{ H_{\tau(k)}, [A_{\tau(k)+1}, \cdots, A_{k-1}], [R_{\tau(k)+1}, \cdots, R_{k-1}] \big\}. \quad (7)$$

Here, $\tau(k) = P\lfloor(k-1)/P\rfloor$, where $\lfloor x \rfloor$ denotes the largest integer less than x. We interpret $\tau(k)$ as the most recent time index prior to k at which the channel amplitude is known.

For example, for $P = 1$, we get $\mathbf{c}_k = \{H_{k-1}\}$, which reduces to the delayed CSI; for $P = 2$, $\mathbf{c}_k = \{H_{k-1}\}$ for $k = 1, 3, \cdots$, while $\mathbf{c}_k = \{H_{k-2}, A_{k-1}, R_{k-1}\}$ for $k = 2, 4, \cdots$.

4) No CSI: $\mathbf{c}_k = \emptyset$. Finally, CSI is not available in this case (besides knowing the long-term parameters $\bar{\gamma}$ and $\bar{\rho}$).

C. Causal Diagrams to Illustrate Rate Adaptation

We treat each rate as a random variable and the sequence of rates over time as a stochastic process. The causal relationship of the channel amplitudes $\{H_k\}$, the CSI $\{\mathbf{c}_k\}$ and the rates $\{R_k\}$ can be represented with a directed graph [13], which can be established rigorously as a *causal diagram* [14]. Let $\mathcal{A}_k, \mathcal{B}_k$ be sets of random variables at time k or earlier, and let x_k be a random variable at time k. If $p(x_k|\mathcal{A}_k, \mathcal{B}_k) = p(x_k|\mathcal{A}_k)$, we draw an arrow from each of the random variables in $\mathcal{A}_k$ to x_k. Intuitively we may say that x_k is *caused* by the random variables in $\mathcal{A}_k$, but not in $\mathcal{B}_k$.

Besides providing a graphical overview of how the random variables interact, a causal diagram allows any conditional independence to be easily established [13], [14]. The causal diagrams for different CSIs are illustrated in Fig. 2 for K packets, based on the following considerations. For all k, we have $H_k \to H_{k+1}$ since the channel amplitude is Markovian. Moreover, the probability of a PACK or NACK depends only on the present channel and rate, thus we have $\{H_k, R_k\} \to A_k$. Finally, we let each rate depend only on its corresponding CSI, thus we have $\mathbf{c}_k \to R_k$. For any CSI, the joint pdf of $\mathbf{h}_K, \mathbf{a}_K, \mathbf{r}_K$ can then be factored as

$$p(\mathbf{h}_K, \mathbf{a}_K, \mathbf{r}_K) = \prod_{k=1}^{K} p(H_k|H_{k-1}) p(R_k|\mathbf{c}_k) p(A_k|R_k, H_k). \quad (8)$$

D. Main Analytical Results and Discussions

The maximum achievable throughput, given by (6) for ACK-rate CSI, is also defined similarly for other types of CSIs by an appropriate substitution of $\mathbf{c}_k$ according to Section III-B. We denote the maximum achievable throughput for full CSI, delayed CSI, periodic CSI and no CSI as $T^*_{\text{full}}, T^*_{\text{delayed}}, T^*_{\text{periodic}}$ and T^*_{no}, respectively.

We say that a rate-adaptation policy is a *myopic policy*, if the rate for each packet is adapted to maximize only the current throughput, without concerns about the effect on future

achievable throughput. Mathematically, given CSI $\mathbf{c}_k$, the rate obtained by the myopic policy can be expressed as

$$R_k^* = \arg\max T(R_k; \mathbf{c}_k). \tag{9}$$

Thus, a myopic policy achieves a throughput of $\mathbb{E}[T(R_k^*; \mathbf{c}_k)]$.

Our first main result is stated in Theorem 1.

Theorem 1: The maximum throughput for full CSI, delayed CSI or no CSI is achieved by a stationary myopic policy, which can be expressed, respectively, as

$$T_{\text{full}}^* = \mathbb{E}_{H_k}[\max_{R_k} T(R_k; H_k)] = \mathbb{E}_H[C(H)], \tag{10}$$

$$T_{\text{delayed}}^* = \mathbb{E}_{H_{k-1}}[\max_{R_k} T(R_k; H_{k-1})], \tag{11}$$

$$T_{\text{no}}^* = \max_{R_k} T(R_k; \emptyset) = \max_R \mathbb{E}_H[t(R, H)]. \tag{12}$$

The maximum achievable throughput for periodic CSI with period P can be expressed as

$$T_{\text{periodic}}^* = \mathbb{E}_{\mathbf{c}_1}[J_1(\mathbf{c}_1)]/P. \tag{13}$$

Here, J_1 can be expressed recursively using $J_k(\mathbf{c}_k)$ with decreasing $k = P, P-1, \cdots, 1$, where

$$J_K(\mathbf{c}_K) = \max_{R_k} T(R_k; \mathbf{c}_k), \tag{14}$$

while for $k = P - 1, \cdots, 1$,

$$J_k(\mathbf{c}_k) = \max_{R_k}\{T(R_k; \mathbf{c}_k) + \mathbb{E}[J_{k+1}(\mathbf{c}_{k+1})]\}. \tag{15}$$

Proof: We employ Bellman's equations [15] for our proof. Due to space constraint, we have delegated this and subsequent proofs to [16]. ∎

Discussions: The results (10)–(12) has been obtained previously, e.g., [7], but has been included here for completeness. The new result (13) follows from using dynamic programming.

Our next result establishes upper bounds for $T_{\text{ACK-rate}}^*$.

Theorem 2: The maximum achievable throughput for full CSI, delayed CSI, periodic CSI, ACK-rate CSI and no CSI are ordered decreasingly, i.e.,

$$T_{\text{full}}^* \geq T_{\text{delayed}}^* \geq T_{\text{periodic}}^* \geq T_{\text{ACK-rate}}^* \geq T_{\text{no}}^*. \tag{16}$$

Proof: Our proof relies on Theorem 1 and on the Markovian relationship of the channel amplitudes, rates and CSIs over time as captured in (8). Details are given in [16]. ∎

Discussions: With more "informative" CSI, we expect that a larger throughput can be achieved. Theorem 2 strengthens this intuition, since we expect that full CSI is more informative than delayed CSI, delayed CSI is more informative than periodic CSI, and so on. Consequently, T_{periodic}^* serves as a *new* upper bound for $T_{\text{ACK-rate}}^*$. We note that T_{delayed}^* has been used as an upper bound for $T_{\text{ACK-rate}}^*$ in [7], but it cannot be tighter than T_{periodic}^*.

Our last result establishes another new tight upper bound.

Theorem 3: Let T_{ub} be defined by[1]

$$T_{\text{ub}} = \mathbb{E}_{H_{k-2}}\left[\max_{R_{k-1}} \mathbb{E}_{R_{k-1}, A_{k-1}|H_{k-2}}\left[\max_{R_k} T(R_k; \bar{\mathbf{c}}_k)\right]\right], \tag{17}$$

---

[1]The subscript "ub" represents "upper bound" in short.

where $\bar{\mathbf{c}}_k = \{R_{k-1}, A_{k-1}, H_{k-2}\}$ is the CSI available at time k. Then T_{ub} is an upper bound for the maximum achievable throughput with ACK-rate CSI $T_{\text{ACK-rate}}^*$. Moreover, it is a tighter upper bound than T_{delayed}^*, i.e.,

$$T_{\text{delayed}}^* \geq T_{\text{ub}} \geq T_{\text{ACK-rate}}^*. \tag{18}$$

Proof: Details can be found in [16]. ∎

Discussions: Theorem 3 introduces another upper bound T_{ub} for $T_{\text{ACK-rate}}^*$ that is tighter than T_{delayed}^*. The superscript $*$ is omitted in this notation T_{ub}, because the (genie-aided) policy in (17) that achieves T_{ub} cannot be implemented in practice. From numerical simulations in Section IV, T_{ub} can be even tighter than T_{periodic}^*.

We can interpret (17) as a maximization of the current throughput $T(R_k; \bar{\mathbf{c}}_k)$. To maximize the throughput, the past rate R_{k-1} has been optimized given CSI H_{k-2} (i.e., an channel amplitude delayed by one unit of time), while the current rate R_k is optimized given CSI $\bar{\mathbf{c}}_k$. This CSI consists of the past rate, past ACK and a channel amplitude delayed by *two* units of time.

Intuitively, two aspects make T_{ub} achieve a higher throughput than $T_{\text{ACK-rate}}^*$ where ACK-rate CSI is available. Firstly, the CSI $\bar{\mathbf{c}}_k$ available for adapting R_k is more informative than in the case of ACK-Rate CSI, with H_{k-2} being an additional CSI. Secondly, both past rate R_{k-1} and current rate R_k are used to maximize the current throughput (for packet k), without regarding how past throughput (for packet $k-1$) and future throughput (for packet $k+1$ onwards) are affected. However, since past and present rates are *always* used to optimize for the current packet, this genie-aided scheme cannot be implemented in practice.

We can generalize T_{ub}. For a delay-related parameter $P \geq 1$, we make the common CSI H_{k-P} available to all P packets, namely packet $k - P + 1$ to packet k, in addition to their ACK-rate CSI. We then concurrently adapt the rates for these packets to maximize the throughput of packet k. Clearly, (17) corresponds to the case of $P = 2$. This generalized T_{ub} still satisfies Theorem 3 for $P \geq 2$, which can be proven similarly as for $P = 2$. However, the complexity of computing T_{ub} increases exponentially with P, hence for our numerical results we restrict to $P = 2$.

IV. NUMERICAL STUDY

For our numerical results, we discretized the channel amplitude H by using $N_{\mathcal{H}} = 100$ channel states. We set the maximum channel amplitude as $\alpha = 5$ for an average SNR of $\bar{\gamma} = 0$ dB; experiments showed that the throughput for full CSI is affected negligibly by these choices of α and $\bar{\gamma}$. At other SNRs, we scaled the Rayleigh distribution to $\bar{\gamma} = 0$ dB, then discretized the distribution similarly. The power correlation coefficient is fixed at $\bar{\rho} = 0.99$. To reduce the effects of rate quantization and to observe the full dynamic behavior of rate adaptation, we used finely quantized rates from the set $\mathcal{S}_R = \{0, 0.05, 0.1, \cdots, 12\}$ for all SNRs.

We numerically obtain the maximum achievable throughput for various types of CSI given by (10)–(13) in Theorem 1, as

Fig. 3. The PRA at different SNR compared to the benchmarks and upper bound, $\bar{\rho} = 0.99$.

well as the upper bound (17) in Theorem 3. In addition, for the case when ACK-rate CSI is available, we consider the throughput achieved by a myopic policy, i.e., the rate is adapted according to (9). The myopic policy is sub-optimal with ACK-rate CSI, although it is optimal with full CSI, delayed CSI and no CSI. The throughput for ACK-rate CSI is obtained by Monte Carlo simulations by averaging the throughput over the first 1000 packets and over 1000 independent runs. To implement the myopic policy, we have to track the *a posteriori* channel pmf $p(H_k|\mathbf{c}_k)$. The complexity of tracking can be very high, so we employ the particle filter to approximate $p(H_k|\mathbf{c}_k)$ over time, according to [9]. This approach is accurate when the number of particles is large, while at a low complexity.

From Fig. 3, we see that the maximum achievable throughput T^*_{delayed} for delayed CSI incurs an SNR loss of around 2 dB compared to T^*_{full} for full CSI. This fundamental loss results from the temporal variation of the channel and the causality constraint imposed in practice, and is irrecoverable. Moreover, T^*_{delayed} serves as an upper bound for $T^*_{\text{ACK-rate}}$ (which cannot be directly computed). However, we see that both of our proposed bounds are tighter than T^*_{delayed}. Specifically, the upper bound T^*_{periodic} that uses periodic CSI is tighter than T^*_{delayed} by about 0.2 dB, while the upper bound T_{ub} is tighter than T^*_{delayed} by about 0.5 dB. In both of these new tight bounds we have used an equivalent period of $P = 2$. If P is increased further, the upper bound can be tightened further, but the computational complexity quickly becomes prohibitive. Finally, we observe that the myopic policy that uses ACK-rate CSI achieves a throughput that is within one dB of the tightest upper bound (namely T_{ub}) over a wide range of SNRs.

For benchmarking, we also plot the maximum achievable throughput T^*_{no} for no CSI. From Fig. 3, we observe that the myopic policy that exploits ACK-rate CSI performs significantly better than this benchmark. For example, the myopic policy is about 3.5 dB better at a throughput of 2 bit/symbol.

From these numerical results, we may conclude that even

with a lean ACK-rate CSI, substantial improvement can be achieved when the channel is slowly time varying. Moreover, a myopic policy already achieves a throughput that is close to what can be maximally achieved. If further throughput improvement is desired, alternatively more informative CSI can be provided, e.g., in the form of periodic CSI.

V. CONCLUSION AND OUTLOOK

We have considered packet-by-packet rate adaptation to improve the average throughput, based on past ACKs and past rates as channel state information. We have proposed two new upper bounds which are tighter than currently known ones. We have shown that the myopic policy, which maximizes only the current throughput, is already fairly close to the maximum achievable throughput, yet at a reasonable complexity. This suggests that the myopic policy may be sufficient for some practical purposes. Future research will focus on improving the myopic policy for a more general fading channel with multi-user interference [16]. Other interesting directions include quantifying and optimizing the delay performance.

REFERENCES

[1] J. M. Wozencraft and M. Horstein, "Coding for two-way channels," Res. Lab. Electron., MIT, Cambridge, MA, Tech. Rep. 383, Jan. 1961.
[2] P. Sindhu, "Retransmission error control with memory," *IEEE Trans. Commun.*, vol. 25, no. 5, pp. 473–479, May 1977.
[3] P. Chevillat, J. Jelitto, A. N. Barreto, and H. L. Truong, "A dynamic link adaptation algorithm for IEEE 802.11a wireless LANs," in *Proc. IEEE ICC'03*, Anchorage, May 2003, pp. 1141–1145.
[4] D. Qiao and S. Choi, "Fast-responsive link adaptation for IEEE 802.11 WLANs," in *Proc. IEEE ICC'05*, Seoul, Korea, May 2005, pp. 3583–3588.
[5] A. Kamerman and L. Monteban, "WaveLAN-II: A high-performance wireless LAN for the unlicensed band," *Bell Labs Technical Journal*, pp. 118–133, Summer 1997.
[6] M. Rice and S. B. Wicker, "Adaptive error control for slowly varying channels," *IEEE Trans. Commun.*, vol. 42, no. 234, pp. 917–926, Feb/Mar/Apr 1994.
[7] A. K. Karmokar, D. V. Djonin, and V. K. Bhargava, "POMDP-based coding rate adaptation for type-I hybrid ARQ systems over fading channels with memory," *IEEE Trans. Wireless Commun.*, vol. 5, no. 12, pp. 3512–3523, Dec. 2006.
[8] C. H. Papadimitriou and J. N. Tsitsiklis, "The complexity of Markov decision processes," *Mathematics of Operations Research*, vol. 12, no. 3, pp. 441–450, 1987.
[9] C. K. Ho and J. Oostveen, "Rate adaptation in time varying channels using acknowledgement feedbacks," in *Proc. 63rd IEEE Vehicular Technology Conf.*, vol. 4, Melbourne, Australia, May 2006, pp. 1683–1687.
[10] C. C. Tan and N. C. Beaulieu, "On first-order Markov modeling for the Rayleigh fading channels," *IEEE Trans. Commun.*, vol. 48, no. 12, pp. 2032–2040, Dec. 2000.
[11] J. D. J. Costello, J. Hagenauer, H. Imai, and S. B. Wicker, "Applications of error-control coding," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2531–2560, Oct. 1998.
[12] C. C. Tan and N. C. Beaulieu, "Infinite series representations of the bivariate Rayleigh and Nakagami-m distributions," *IEEE Trans. Commun.*, vol. 45, no. 10, pp. 1159–1161, Oct. 1997.
[13] T. P. Minka, "From hidden Markov models to linear dynamical systems," Tech. Rep., revised 7/18/99.
[14] J. Pearl, *Causality: Models, Reasoning, and Inference*. Cambridge, U.K.: Cambridge University Press, 2000.
[15] D. P. Bertsekas, *Dynamic Programming and Optimal Control Vol. 1*. Belmont, MA: Athena Scientific, 1995.
[16] C. K. Ho, J. Oostveen, and J.-P. Linnartz, "Rate adaptation using acknowledgement feedback in finite-state Markov channels with collisions," *IEEE Trans. Wireless Commun.*, submitted 2008.