

Modeling TCP in a Multi-Rate Multi-User CDMA System

Majid Ghaderi^{1*}, Ashwin Sridharan², Hui Zang²,
Don Towsley^{1*}, and Rene Cruz³

¹ University of Massachusetts Amherst, {mghaderi,towsley}@cs.umass.edu

² Sprint Advanced Technology Labs, {ashwin.sridharan,hui.zang}@sprint.com

³ University of California San Diego, cruz@ece.ucsd.edu

Abstract. Modern CDMA wireless channels support multiple transmission rates, which can be dynamically assigned to users based on traffic demand. However, in practice, assignment of high rate channels comes with the penalty of increased power as well as smaller orthogonal codes, which constrains their assignment to only a *subset* of active users. This motivates the need to carefully control high rate channel assignments so as to minimize power and achieve fairness among users. In this work, we propose a simple class of channel allocation policies to achieve this goal for TCP sessions. We develop an analytical model that explicitly captures both TCP dynamics *and* the impact of multiple users contending for a shared resource to evaluate the performance of the allocation policy. The model is shown to be accurate by comparing against *ns-2* simulations and its utility demonstrated in computing the minimum number of required high rate channels to minimize contention.

1 Introduction

Modern CDMA networks incorporate powerful mobile processors and digital communication techniques that allow wireless channel schedulers to rapidly allocate (or de-allocate) high data rate channels in response to channel conditions or user traffic demands. Although an eminently desirable feature, in practice, high data rate channels are scarce resources in a CDMA wireless system that cannot be simply assigned to all users. This is due to the increased power requirements as well as reduced codeword set making such an allocation infeasible. Consequently, the high rate channels must be carefully utilized so as to fairly

* This research is continuing through participation in the International Technology Alliance sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the U.S. Government, the UK Ministry of Defense, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

service users as well as minimize power. While the existing mechanisms in commercial CDMA networks based on buffer occupancy are reasonable for inelastic applications, it is not clear as to how to select high data rate channels or share them among multiple TCP sessions. In particular, TCP, which is the dominant transport protocol in the Internet, incorporates adaptive rate mechanisms and is quite sensitive to the available channel rates and channel errors. Hence, it can be expected to demonstrate complex behavior to dynamic rate allocations by the wireless scheduler. While rate allocation in CDMA networks for elastic applications in general has been previously studied [1,2], they do not incorporate TCP dynamics.

Our work takes the first step in addressing this issue by modeling a class of simple allocation policies in such a system while explicitly incorporating TCP dynamics of each session. Recently we demonstrated in [3] that for a *single* TCP session, significant throughput gains are possible if the scheduler selects and allocates channel rates as a function of the TCP transmission rate. In this paper, we extend this framework to include the presence of multiple users along with practical constraints on assignment of high rate wireless channels. Furthermore, we propose a class of allocation policies for sharing the high rate data channels among concurrent TCP sessions, whose performance is explored both analytically and through simulations. The utility of the model is demonstrated by evaluating the impact of various design features like the number of high-rate channels and their data rates on TCP throughput, which is useful for planning purposes.

The rest of this paper is organized as follows. Section 2 describes the high rate contention problem in CDMA wireless systems. Section 3, presents our system model. The TCP-aware allocation policy is presented in Section 4 and analyzed in Section 5. Section 6 validates the accuracy of the model against *ns-2* simulations. Our conclusions as well as future work are discussed in Section 7.

2 CDMA Wireless System

In this section, we describe the wireless framework considered in the paper as well as motivate the problem. For ease of exposition, we use the CDMA2000 1xRTT [4] system as a reference example of a practical system although our analysis applies equally to other wireless systems that dynamically adapt their wireless channel rate in response to user traffic. The focus of this work is on mobile sessions in a single cell that are concurrently involved in TCP bulk transfers on the downlink, *i.e.*, from the access network (or Base Station Controller (BSC)) to the mobile device. The throughput experienced by such sessions on the downlink is strongly influenced by the wireless channel scheduler residing at the BSC because it controls the channel rate allocated to each session. For example, CDMA2000 standards [4] specify up to six different channel rates that can be dynamically assigned to a user by the BSC. The lowest channel rate, called the *fundamental channel* can be simultaneously assigned to *all* users and hence defines the capacity in terms of the number of users. The higher channel

rates are called *supplemental channels* and dynamically assigned to only a *subset* of users at any given instant in order to boost data rates.

Given that the scheduler can choose from up to six rates, the obvious question arises as to why the scheduler does not always assign the highest channel rate to all users? Indeed, this issue forms the crux of the resource allocation problem investigated in this paper. As was briefly alluded to in Section 1, allocation of a high rate channel incurs penalties in two ways. First, in order to maintain a reasonably low frame error rate (FER), a high rate channel requires the wireless antenna to transmit at higher power which can introduce excess interference in neighboring cells and shrink the available power budget at the cell itself for transmission to other users. Doing so for *all users* would further exacerbate the problem. Second, and perhaps more importantly, in practice, higher rate channels in a CDMA system are usually achieved by reducing the orthogonal code length (called Walsh codes). This in turn implies that not all active users can *simultaneously* transmit at a higher rate, simply because fewer orthogonal code-words are available at that rate, making it inherently infeasible. Furthermore, as a direct corollary of the above argument, the *higher* the supplemental channel rate, the shorter the code length, and hence smaller the subset of concurrently assignable users.

For both these reasons, it is desirable that the wireless channel scheduler judiciously assigns high rate channels such that a) it is utilized only when required so as to minimize power and b) *cycled* among all users desiring the high rate channel so as to achieve fair resource allocation. Since TCP incorporates an adaptive feedback mechanism that is sensitive to channel errors and available channel rate, its behavior to dynamic rate adaptation by the wireless scheduler can be quite different from that of inelastic applications. This immediately raises the issue of how to allocate supplemental channels to TCP sessions in a CDMA system and its impact on achieved throughput. We address this issue in this paper.

3 System Model

We assume a single cell with N continuously backlogged TCP sessions transferring data over the downlink. Hence, there are N fundamental channels that can be simultaneously used by all users.

1. The cell is assumed to have K high rate supplemental channels, such that at any instant in time up to $K \leq N$ users can transmit at the higher rate.
2. We assume that the fundamental channel has rate C_0 and the supplemental channel a rate of C_1 . Clearly $C_0 \leq C_1$. At each point in time, the scheduler decides which type of channel is to be assigned to each TCP session. While the fundamental channel is always available, the allocation of high rate channels is arbitrated by the wireless scheduler because $K \leq N$.
3. The packet error probability is implicitly assumed to be a function of the assigned channel and denoted by $p_0(p_1)$ when the assigned channel rate

is $C_0(C_1)$. Note that typically $p_0 \leq p_1$, reflecting the impact of smaller codewords to achieve higher channel rates.

4. We assume the presence of power control to primarily combat fast fading and interference effects. This is true in current systems where fast closed-loop power control tracks a specified target SINR (or equivalently target FER).
5. We assume no (or a very small) buffer at the base station. Hence, TCP experiences congestion if its sending rate exceeds the maximum allocated channel rate.

We note that modern cellular systems can support up to four or more different supplemental channel rates. However, obtaining succinct analytical expressions for TCP throughput even for three rates is quite difficult. Moreover, our simulations show that there is only a small improvement in system performance by considering more than two rates [5]. Finally, we emphasize that no specific assumptions have been made regarding how the two channel rates are achieved nor how they result in the specific channel error probabilities.

4 TCP-Aware Rate Allocation

We now present our TCP-aware channel allocation scheme. It essentially comprises of a simple preemptive allocation discipline with parameter α that utilizes the TCP transmission rate to make preemption decisions. Specifically, we assume that the wireless scheduler can monitor the TCP rate of each session and, based upon it, decides what channel to allocate. The mechanism works in the following fashion:

1. Upon starting, a TCP session i is assigned a low-rate channel with rate C_0 .
2. Since the TCP rate increases additively (in the absence of loss), when the transmission rate of the session reaches C_0 , the capacity of the fundamental channel, the scheduler executes the following decision mechanism:
 - (a) If a supplemental channel C_1 is free *i.e.*, less than K supplemental channels are occupied, then the session is always assigned the supplemental channel.
 - (b) Otherwise, it is allocated a supplemental channel with probability α . Since at most K users can utilize the high rate channel, this means that with probability α a session already utilizing the high-rate channel is *preempted* to accommodate the requesting session. The session to be preempted is chosen randomly with probability $1/K$. We note that the preempted session faces a congestion loss with probability *one* since its current transmission rate must have been greater than C_0 and hence immediately halves its sending rate.
 - (c) If the supplemental channel request is *denied*, which happens with probability $1-\alpha$, then the session experiences a congestion loss with probability *one* and immediately halves its sending rate, in this case to $C_0/2$.

3. Finally, if a session occupying the supplemental channel drops its sending rate to less than that of the supplemental channel rate, either due to congestion or channel errors, its supplemental channel is de-allocated and it is assigned a fundamental channel⁴. This makes sense, since it frees up resources that would not be utilized by the TCP session.

The preemption probability α is a system parameter that needs to be optimized in order to maximize TCP throughput. The model we present in the next section explicitly captures the impact of α on TCP throughput and hence, can be utilized to find the optimal value of α .

5 TCP Throughput Model

In the previous section, we proposed a simple scheme for TCP-aware channel allocation. We now present an analytical model to compute the throughput in such a system. Rather than list all the details of the analysis, due to lack of space, we highlight the salient features of the system as well as our approach to model it. The reader is referred to the technical report [6] for further information.

From the perspective of a TCP session, such a multi-rate system results in the session being in either of two distinct states corresponding to whether it is allocated a fundamental or supplemental channel. Each type of channel has different channel capacities, round trip times and packet error probabilities (C_0, R_0, p_0) and (C_1, R_1, p_1) respectively. Given the operation of the proposed scheduler, the state of a session, is decided by two factors: a) its own sending rate or *instantaneous throughput*, and b) *the supplemental channel occupancy*, both of which the scheduler utilizes in channel allocation. The latter is critical since it reflects the impact of the multi-user resource sharing aspect of the system. Specifically, whether there are K or less supplemental channel allocations influences both the probability that the session *receives* a supplemental channel, as well as the rate at which it itself is *preempted* from a supplemental channel.

Now, the *supplemental channel occupancy* in turn depends on the rate at which sessions in the fundamental channel state request the supplemental channel, which is actually a function of the time a session spends in each state (fundamental or supplemental). This essentially closes the loop in the sense that the state of the TCP session is a function of the supplemental channel occupancy, which itself depends on the state of each session. A natural approach to analyze such a system with coupled variables is the fixed-point method, which in this case, works as follows. We first derive the instantaneous throughput distribution equations for a TCP session *given* the supplemental channel occupancy statistics. Note that there is a one-to-one relation between the instantaneous TCP throughput and session state. Next, we utilize the newly computed throughput equations to update the supplemental channel occupancy statistics. These iterations are repeated till a desired convergence criteria is met. Both components are described in the next two sub-sections.

⁴ This implies that $C_1 \leq 2C_0$ which is true in practical systems. In practice, most wireless rates follow the relation $C_1 = 2C_0$.

5.1 TCP Session Throughput

This sub-section computes the throughput of a single TCP session *given* the supplemental channel occupancy statistics. In order to capture the evolution of TCP dynamics in the dynamic two-rate environment, we utilize the model developed in [3] as a starting point and incorporate *preemption* from a supplemental channel as well as denial of the supplemental channel. The reader is referred to the technical report [6] for further information. We first state some assumptions regarding TCP dynamics as well as our notation:

1. Let $i = 0(1)$ denote that the session is allocated a fundamental (supplemental) channel.
2. All TCP sessions are assumed to be homogeneous in the sense that they have the same propagation delay a . Hence, their round trip time in state i is given by $R_i = a + L/C_i$ where L is the packet size. We note that the assumption of homogeneity is made in this work only for purposes of illustration of the model. As discussed in Section 5.2 our model can readily tackle heterogeneity as well.
3. In congestion avoidance mode, the TCP sending rate of a session, when in state i , grows at a linear rate of L/R_i^2 bits/sec² in the absence of loss.
4. We assume that the channel losses can be modeled by an inhomogeneous Poisson process with rate $p_i X(t)$, $i = 0, 1$ at time t , where p_i is the packet loss probability.
5. Let $f_i(x, t)$ denote the density function of the instantaneous rate $X(t)$ in state i . Then $f(x, t) = f_0(x, t) + f_1(x, t)$, where $f(x, t)$ denotes the density function of $X(t)$.
6. Let q denote the probability that the session is *allowed* to transition from $i = 0$ to $i = 1$, *i.e.*, its request for a supplemental channel is honored. Let ν denote the rate at which a session is *preempted* from a supplemental channel, *i.e.*, transitions from $i = 1$ to $i = 0$. We note that these quantities are computed based on the supplemental channel occupancy.

We further define the terms $\phi_{ij} = \frac{p_i R_j^2}{L^2}$ and $\delta_i = \frac{L^2}{R_i^2}$ for $i = 0, 1$ and $g = \frac{R_0}{R_1}$. Using techniques from fluid analysis we can write differential equations for the evolution of the steady-state distribution $f(x)$ of the instantaneous throughput $X(t)$. Then, by applying Mellin transforms [7, 8], which is defined as

$$\hat{f}_i(u) = \int_0^\infty f_i(x) x^{u-1} dx, \quad (1)$$

we can solve the differential equations (see [6]) to obtain the steady state probability ($\hat{f}_i(1)$) and throughput ($\hat{f}_i(2)$) in each state i . We may then solve for the mean TCP throughput given by:

$$\bar{X} = \hat{f}_0(2) + \hat{f}_1(2) = \frac{\hat{\Delta}(2) + \sum_{k \geq 0} (\phi_{00})^k \Pi_k(2) \psi(2 + 2k)}{\hat{\Delta}(1) + \sum_{k \geq 0} (\phi_{00})^k \Pi_k(1) \psi(1 + 2k)} \quad (2)$$

where,

$$\begin{aligned} \psi(u) = & \left(\frac{g^2}{q} - g^u - \frac{g^2}{2^u} \left(\frac{1-q}{q} \right) \right) (C_0)^u \Delta(gC_0) - \left(\frac{g^2}{(2g)^u} - 1 \right) (C_1)^u \Delta(C_1) \\ & - u \widehat{\Delta}(u) - \nu \left(\frac{1}{\delta_0 (2g)^u} - \frac{1}{\delta_1} \right) \widehat{\Delta}(u+1) - \left(\frac{\phi_{10}}{(2g)^u} - \phi_{11} \right) \widehat{\Delta}(u+2), \end{aligned}$$

and,

$$\begin{aligned} \Delta(x) &= e^{-\left(\frac{\phi_{11}}{2}\right)x^2 - \left(\frac{\nu}{\delta_1}\right)x}, \\ \widehat{\Delta}(u) &= \int_{gC_0}^{C_1} \Delta(u)x^{u-1} du. \end{aligned}$$

5.2 Supplemental Channel Occupancy

In the previous sub-section, we computed the TCP throughput of a single session given supplemental channel assignment probability q and preemption rate ν . Our goal in this sub-section is to show how these values may be computed.




Fig. 1. Number of users in high-rate region.

We model the supplemental channel occupancy process as a Markov chain where the state, i , is equal to the number of concurrently active supplemental channels. Clearly, $0 \leq i \leq K$. Now, let λ and μ denote, respectively, the rates at which a *single* user requests and releases a supplemental channel conditioned upon there being one available. We assume that both these processes are exponentially distributed, which is a reasonable assumption for large N and K . The transition rates for the Markov chain are then given as shown in Fig. 1. Denoting the stationary distribution of the Markov chain by π , it is easily seen that

$$\pi_i = \binom{N}{i} \left(\frac{\lambda}{\mu} \right)^i \pi_0, \quad (3)$$

where $\sum_{i=0}^K \pi_i = 1$.

It now remains to compute λ and μ . Since λ is the rate at which a session requests a high-rate channel, it is essentially equal to the rate at which the TCP rate crosses the low channel rate boundary $X(t) = C_0$ conditioned on a

supplemental channel being available. This is simply equal to $\lambda\Delta t$, $\Delta t \rightarrow 0$ and can be computed as follows:

$$\begin{aligned}
\lambda\Delta t &= \mathbb{P}\{(C_0 - \delta_0\Delta t \leq X(t) \leq C_0) \wedge (\text{no loss in } [t, t + \Delta t]) \mid X(t) \leq C_0\} \\
&= \frac{\mathbb{P}\{(C_0 - \delta_0\Delta t \leq X(t) \leq C_0) \wedge (\text{no loss in } [t, t + \Delta t])\}}{\mathbb{P}\{X(t) \leq C_0\}} \\
&= \frac{(1 - p_0 C_0 \Delta t) f_0(C_0 - \delta_0 \Delta t, t) (\delta_0 \Delta t)}{\widehat{f}_0(1)}
\end{aligned} \tag{4}$$

In the limit, as $\Delta t \rightarrow 0$ and $t \rightarrow \infty$, we obtain

$$\lambda = \frac{\delta_1 f_1((gC_0)^+)}{q \widehat{f}_0(1)}. \tag{5}$$

Please refer to [6] for further details.

To compute μ , we note that from the Mellin transform analysis in previous subsection, we readily have expressions for $\widehat{f}_0(1)$ and $\widehat{f}_1(1)$, the steady-state probabilities of being in low rate and high rate regions respectively. It is then straightforward to see that

$$\mu = \lambda \frac{\widehat{f}_0(1)}{\widehat{f}_1(1)}.$$

The relevant supplemental channel occupancy statistics can now be computed in a straightforward fashion as follows:

1. The probability q that a session in state $i = 0$ requesting the supplemental channel, will be assigned one is given by:

$$q = (1 - \pi_K^0) + \alpha \pi_K^0 = 1 - (1 - \alpha) \pi_K^0, \tag{6}$$

where, α is the scheduler parameter and π_K^0 is the conditional probability that all K supplemental channels are occupied given that a user is about to request a supplemental channel. In [6], we have shown that

$$\pi_i^0 = \frac{(N - i) \pi_i}{\sum_{j=0}^K (N - j) \pi_j}. \tag{7}$$

2. The rate at which a user in state $i = 1$ is preempted from the supplemental channel is given by:

$$\nu = \frac{\alpha}{K} \lambda (N - K) \pi_K^1, \tag{8}$$

where, π_K^1 is the conditional probability that all K supplemental channels are occupied given that a user has a supplemental channel and is given by [6]

$$\pi_i^1 = \frac{i \pi_i}{\sum_{j=0}^K j \pi_j}. \tag{9}$$

Finally, using a fixed-point iteration technique we can compute mean TCP throughput in multi-user systems as follows. We start with $(q = 1, \nu = 0)$ and compute $(\bar{X}, f_1(gC_0), \hat{f}_0(1), \hat{f}_1(1))$. Using these values, we then compute (λ, π_K) from (5) and (3). By substituting in (6) and (8), we compute new values for (q, ν) . The iteration continues until computed TCP throughput \bar{X} converges with some desired precision.

In passing, we emphasize that the model can easily accommodate heterogeneous TCP sessions. Specifically, sessions can be classified into (say) L classes of users, based on similarity of propagation delay, error characteristics, *etc.* for each of which we can compute the TCP throughput following the approach in Section 5.1. The high-rate channel occupancy process is then readily seen to be an L -dimensional birth-death Markov chain with transition rates λ_l, μ_l . It is well known that such a chain has a product type form for the stationary distribution which can be easily solved to obtain the occupancy probabilities.

6 Numerical Results

This section is devoted to evaluating the accuracy of our model by comparing against *ns-2* simulations as well as demonstrating its utility in design issues. We consider a CDMA downlink scenario with a fundamental channel rate of $C_0 = 76.8$ Kbps and supplemental channel rate of $C_1 = 102.4$ Kbps resulting in packet error probabilities of $p_0 = 7.7 \times 10^{-4}$ and $p_1 = 7.9 \times 10^{-4}$ respectively⁵. This scenario is reflective of the fact that higher rate channels typically incur higher frame errors.

We begin by determining the accuracy of the model. Fig. 2 plots the per-user TCP throughput as a function of α for $N = 10$ users with $K = 2, 5$ and 8 supplemental channels for both the model and simulation. We have also plotted 95% confidence intervals where each simulation point is averaged over 20 runs with different seeds. We note that the model closely tracks the simulation results for all values of K . More importantly, both the model and simulations indicate the same throughput *behavior* as a function of α . Specifically, for $K = 2$, there is an $\alpha^* \in (0, 1)$ that maximizes the TCP throughput, while larger values of K results in the optimum $\alpha^* = 1$. However, as K increases, the throughput becomes less sensitive to α , *e.g.*, for $K = 8$ in Fig. 2.

6.1 Impact of Number of Supplemental Channels

Next, we study the impact of the number of supplemental channels K on user throughput. As shown previously, the model gives reasonable estimates and hence we use analytical results from the model only.

⁵ The rate values were chosen to reflect those of a CDMA2000 1xRTT system. The packet error probabilities were computed based on BER of 10^{-5} and 1/5 error coding. The reader is referred to [3] for more details.

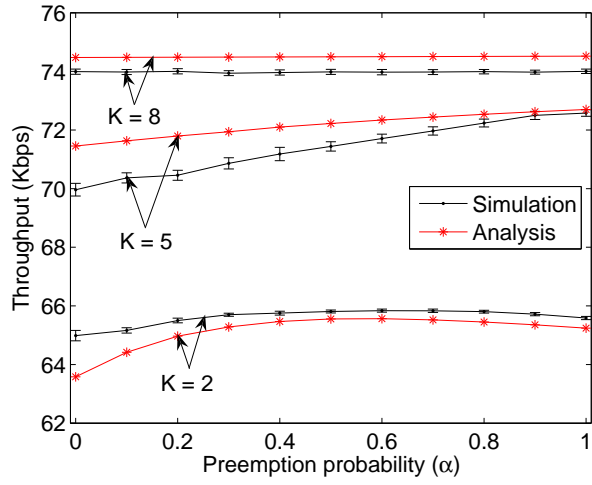


Fig. 2. Throughput as function of α for $N = 10$.

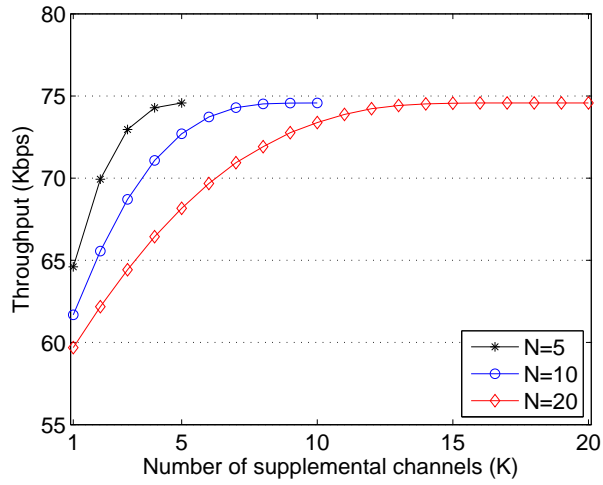


Fig. 3. Max throughput as function of K .

Fig. 3 plots the TCP throughput as a function of K for $N = 5, 10$ and 20 users. For each K , the maximum throughput achieved at the optimal α^* was chosen. The curves in Fig. 3 indicate that the throughput *gain* by adding additional supplemental channels follows a law of diminishing returns, *i.e.*, beyond a certain K , addition of extra supplemental channels does not yield significant gains. For example, for $N = 20$, allocation of only $K = 10$ supplemental channels achieves nearly the same throughput as that with a dedicated high rate channel per user, *i.e.*, $K = 20$. This stands to reason, since TCP increases its rate only additively. Thus not all users require high rate channels simultaneously and the resulting statistical multiplexing gain can be exploited to utilize less than N high rate channels while achieving almost the same effect as dedicating a supplemental channel to each user.

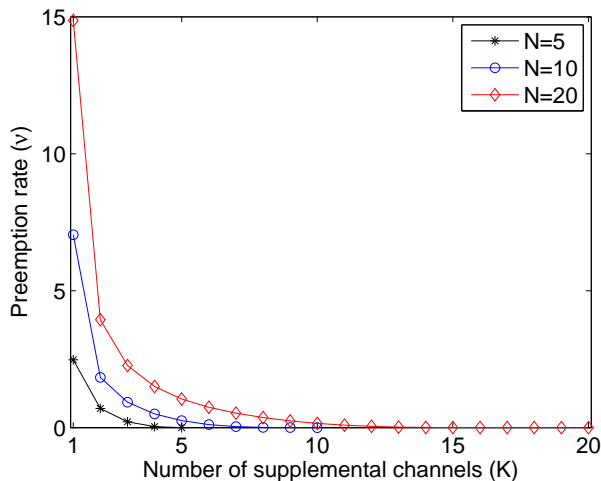


Fig. 4. Max preemption rate (ν) as function of K .

Importantly, the preemption rate ν computed by the model from (8), which is plotted in Fig. 4 as a function of K at $\alpha = 1$, captures this phenomenon. Specifically, beyond a certain K (as a function of N) the contention rate goes to zero which indicates that additional supplemental channels have very little impact. This is an important feature since it allows for computation of the minimum number of supplemental channels to allocate in a cell as a function of number of active TCP sessions, preventing resource wastage.

7 Conclusion

In this work, we addressed the issue of allocation of CDMA wireless channels to concurrently active TCP sessions. Due to inherent system constraints, high data

rate channels cannot be assigned simultaneously to all sessions. This creates an interesting resource contention problem. We proposed a simple channel allocation scheduler for this problem in the context of TCP sessions and constructed an analytical model to study its performance. Through comparison with simulations the model was shown to provide reasonable estimates of TCP throughput and characterize interesting design parameters like the minimum required number of high data rate channels for a set of TCP sessions.

There are many different future directions possible from our current work. Some of avenues that we are exploring include incorporation of more sophisticated schedulers as well as utilizing the model to address resource planning issues like selection of both the *number* and *rates* of supplemental channels and impact of heterogeneous users.

References

1. Fodor, G., Telek, M., Badia, L.: On the Tradeoff Between Blocking and Dropping Probabilities in CDMA Networks Supporting Elastic Services. In: Proc. IFIP/NETWORKING, Coimbra, Portugal (May 2006)
2. Altman, E.: Capacity of a Multi-Service Cellular Network with Transmission Rate Control : A Queueing Analysis. In: Proc. ACM MOBICOM, Atlanta, GA (September 2002)
3. Ghaderi, M., Sridharan, A., Zang, H., Towsley, D., Cruz, R.: TCP-Aware Resource Allocation in CDMA Networks. In: Proc. ACM MOBICOM, Los Angeles (September 2006)
4. Association, T.I.: TIA EIA IS-2000. Available at www.tiaonline.org/standards/sfg/imt2k/cdma2000/ (March 2000)
5. Ghaderi, M., Sridharan, A., Zang, H., Towsley, D., Cruz, R.: Modeling and Optimization TCP Aware Cross-Layer of Resource Allocation in CDMA Networks. Technical Report *RR06-ATL-030566*, Sprint Advanced Technology Labs (March 2006)
6. Ghaderi, M., Sridharan, A., Zang, H., Towsley, D., Cruz, R.: Performance Modelling of TCP in a Multi-Rate Multi-User CDMA System. Technical Report *RR06-ATL120729*, Sprint ATL (December 2006) Available at <http://research.sprint.com/publications/uploads/RR06-ATL-120729.pdf>.
7. Baccelli, F., Cruz, R., Nucci, A.: CDMA channel parameters maximizing TCP throughput. In: Proc. Workshop on Information Theory and its Applications, La Jolla, CA, USA (February 2006)
8. Baccelli, F., McDonald, D.: Mellin transforms for TCP throughput with applications to cross layer optimization. In: CISS, Princeton (2006)