

On the Influence of User Behaviour and Admission Control on System Performance in HS-DSCH

Mats Folke and Ulf Bodin

Division of Systems and Interaction

Department of Computer Science and Electrical Engineering

Luleå University of Technology

Email: {mats.folke, ulf.bodin}@ltu.se

Abstract—In this paper we investigate the need for admission control for the high-speed downlink shared channel (HS-DSCH) through the evaluation of two admission control mechanisms. One mechanism uses the number of active users in a cell as a metric and the other one uses the mean downlink throughput of a user. We also introduce a model for user behaviour in which the goodput of a completed file download decides if further downloads are made. In order to measure user-experienced quality we use a utility function for transforming per-flow goodput into a user satisfaction index.

System performance, measured by total user satisfaction and total goodput, is evaluated for a range of session arrival rates and admission control limits. This evaluation is done using the ns-2 simulator, together with extensions of our own.

If the objective is to maximise goodput, our results show that no admission control is needed. Maximising user satisfaction benefits from an admission control. We also note that the impact of user behaviour is not insignificant.

I. INTRODUCTION

The High-Speed Downlink Shared Channel (HS-DSCH) in Wideband CDMA (WCDMA) release 6 has theoretical peak bitrates for data services of 14 Mbps [1], [2]. Moreover, delays considerably shorter than for other shared data channel technologies in previous releases of WCDMA are possible. HS-DSCH is primarily shared in the time domain, where users are assigned time slots according to a scheduling algorithm that runs independently at each Node B. The short Transfer Time Interval (TTI) of 2 ms enables fast link adaptation, fast scheduling and fast Hybrid Automatic Repeat reQuest (HARQ). The channel is designed for bursty Internet traffic, such as web traffic.

Four traffic classes are defined for HSDPA. *Conversational* is for streaming audio (i.e. VoIP), which requires low delays and strict requirements for minimum bandwidth. *Streaming* is for streaming video, which also requires low delays but higher and slightly more varying bandwidth demands than the conversational class. The *interactive* class is for interactive traffic (i.e. web surfing). The demand for bandwidth is elastic, and there are no tight bounds in delay. The final class is *background*. This is for traffic with low demands on delay, such as e-mails.

In this paper we investigate the need for admission control in the interactive traffic class to improve and optimise total goodput and user satisfaction. These metrics represent the total value provided by the system, but do not capture the

quality perceived by individual users. The variance in quality experienced by users can however be kept sufficiently low through a properly weighted proportional fair scheduler. We therefore exclude the aspect of per-user quality in this study.

TCP reduces the send rate upon congestion and users not experiencing high enough bit-rates are likely to become dissatisfied and quit ongoing sessions prematurely. Users finishing sessions prematurely decrease the system load in a similar manner as a pre-emption mechanism, which would drop users with low experienced quality. We evaluate whether or not these existing load control features provided by TCP and the users themselves are sufficient to optimise the system utilisation¹. For this evaluation we examine the performance of two admission control mechanisms in combination with a simple model of user behaviour. The metrics used to judge this performance are average total goodput and user satisfaction.

In addition to optimise system utilisation admission control can be used to ensure the stability of the system. An admission control mechanism ensuring stability typically limits the maximum number of users in a cell. We use this approach for the first mechanism tested. The second admission control mechanism keeps a running average for the mean throughput in a cell and denies new users when this metric drops below a given limit.

For the evaluation we use an application model based on user experienced quality. Depending on the goodput of a completed file transfer a user may choose to download additional files. In order to estimate user quality we employ a utility function. Utility functions have been widely used to model user behaviour in both wired and wireless networks (e.g., in analysing reservations vs. best-effort [3] and for studies of fairness in wireless networks [4], [5]). For our evaluation we study the two admission control mechanisms for various loads from a system perspective as well as from a user perspective.

In [6] Hosein presents an algorithm that provides guaranteed levels of throughput. Multiple such levels are maintained through a combined scheduler and admission control mechanism. While the guaranteed levels of throughput may meet the needs of the conversational and streaming traffic classes, they are not sufficient for the interactive class. This is because users

¹We assume all interactive traffic to use TCP.

of this class typically become increasingly satisfied with the increasing goodput instead of immediately being satisfied only when an expected goodput is obtained. Hence, maintaining high system throughput and user satisfaction are important for the interactive class.

This paper consists of four sections. Section I contains the introduction. It is followed by Section II which describes the method and metrics we have used. Section III contains our results and the paper is concluded in Section IV.

II. METHOD

In this section our simulation environment is presented. This includes the radio model, the user mobility model and the load model. After the simulation environment is described we present the metrics used for the evaluation.

A. Simulation Environment

1) *Radio model*: We have implemented a model of HS-DSCH in the Network Simulator version 2.28 (ns-2) [7]. The radio model includes lognormal shadow fading with a standard deviation of 8 dB and exponential path loss with a propagation constant of 3.5. Self interference is assumed to be 10 percent and the interference from simultaneous transmissions within a user's own cell is approximated to 40 percent. The interference from transmissions in other cells than a user's own cell is dampened through distance. Wrap-around for interference is supported. All cells have omnidirectional antennas and a radius of 500 m. Code multiplexing for up to three users in the same time slot for a given cell is supported. The available coding and modulation combinations are accounted for in Table I. Block errors are uniformly distributed. When SIR is less than -3.5 dB approximately every second block is received in error, for better SIR conditions the block error rate is 10 percent.

The choice of scheduling algorithm is important. We have implemented both a Round-Robin (RR) scheduler and a SIR scheduler. In this investigation we only use the RR scheduler. It distributes the transmission slots fairly among users. No fast HARQ is implemented; instead, damaged radio-blocks are immediately retransmitted. Multi-path fading is not implemented. Both of these mechanisms would give variations in RTT if modelled. We expect however these variations to be too small to have an impact on the system-level properties.

TABLE I

LINK ADAPTATION PARAMETERS. SIZE REFERS TO THE RADIO-BLOCK SIZES.

Coding (rate)	Modulation (type)	SIR (dB)	Bitrate (Mbps)	Size (bytes)
0.25	QPSK	-3.5	1.44	360
0.50	QPSK	0.0	2.88	720
0.38	16QAM	3.5	4.32	1080
0.63	16QAM	7.5	7.20	1800

2) *User mobility*: When starting a simulation the UEs are randomly distributed according to a uniform distribution for the x-axis and the y-axis on a cell plan consisting of seven cells each having a radius of 500 m. Traffic sources are at equal distance (50 ms) from the Node Bs and the UEs are associated with the closest Node B. The links connecting the Node Bs with the traffic sources are overprovisioned, hence packets will only be dropped due to congestion at the IP-buffers. During each session the UE moves with a speed drawn from a low-speed mobility model [8]. All directions are equally likely to be taken when beginning a new session. Wrap-around for the moving users is supported. This means that a UE moving off the cell plan re-appears at the other end.

3) *Application and load model*: A session generator uses a Poisson arrival process to initiate new sessions. When a session begins, the UE starts downloading a file. Upon completion the goodput of the transfer is calculated. If the goodput is above a certain threshold, the user is satisfied and will begin a new download. However, if the goodput is less than the predetermined threshold, the user is dissatisfied and will not download any more files. At most four file transfers will be performed in a session.

The file sizes are drawn from a Pareto distribution with a mean of 30458 bytes and the shape parameter set to 1.7584 [9]. At the endpoints we use TCP Sack [10], as implemented using the TCPSack1 agent in ns-2. This includes support for limited transmit [11] and a variant of SACK loss recovery, as specified in [12]. Upon completion of a transfer, the endpoints are reset, thus no teardown is performed, but connection establishment takes place for every flow. When a session is started a UE is said to be active and when the session is ended, the UE is inactive.

The load of the system is varied by setting the arrival rate of new users initiating sessions. This value ranges from 10 to 40 new sessions per second. Our system also employs two algorithms for admission control. The first is based on the maximum number of simultaneously active users per cell, which ranges from 10 to 40. The other algorithm is based on mean user throughput. The throughput of a cell is calculated using a moving average filter. When a new user arrives at the cell the mean user throughput is calculated by dividing the total throughput with the number of active users. If this mean user throughput is greater than a predefined threshold, the new user is admitted. The threshold ranges from 100 to 400 kbit/s. We also use a moving average filter to keep track of the utilisation² of a cell. Regardless of admission control algorithm, a user is admitted to the system if the utilisation is less than 98%. A data transfer not admitted gets a goodput of 0 kbit/s and thus a satisfaction index equal to 0.

Apart from varying the load, we also experiment with the user behaviour. We model this as a threshold mechanism in goodput. We have selected two values, 0 kbit/s and 100 kbit/s. 0 kbit/s means that the user will perform four file transfers in

²We define the utilisation over a period of time as the ratio of used timeslots over the total number of timeslots.

each session, no matter what the goodput is. The other value, 100 kbit/s, means that the user will abort its session if the goodput of a file transfer is below that value. A satisfied user will however not perform more than four file transfers before ending its session.

Each scenario is run five times with different initial values for the positions and velocities of the users as well as the starting times of the session and the sizes of the file transfers. The simulation runs for 150 simulated seconds. For the analysis, the first 20 and final 20 seconds of the simulation is discarded.

B. Evaluation metrics

We use two metrics for this evaluation. Goodput is defined as the ratio of file size over transmission time. An unnecessary retransmission may prolong the transmission time and may thus affect the goodput negatively. We study goodput from a system perspective (total goodput).

We use a utility function to transform the goodput of a file transfer to an index describing the level of user satisfaction. The function can be seen in Equation 1. This Satisfaction Index is calculated for each completed file transfer and is summed up for all transfers in a session. Thus, the satisfaction index for a session will be between 0 and 4. The Satisfaction Index of a file transfer will be 0 if the goodput of the transfer is below 100 kbit/s. If the goodput is above 400 kbit/s the Satisfaction Index will be set to 1. Since a denied user has no goodput at all, the corresponding file transfer will receive a Satisfaction Index of 0.

$$f(x) = \begin{cases} 0 & x < 100000 \\ \frac{x-100000}{300000} & 100000 \leq x < 400000 \\ 1 & x \geq 400000 \end{cases} \quad (1)$$

As a support for these two metrics we also present the mean number of active users at any given time during a simulation run. This metric can be used to see if the admission control, which seeks to limit the number of users, performs the way it should.

Studying admission control algorithms over different loads, employing user models, in order to maximise the potential revenue of 3GPP system is desirable. We analyse whether the optimal load differs between optimising the system for maximal total goodput and for maximal total user satisfaction respectively. This analysis is performed with and without the impact from user behaviour (i.e., users finishing their sessions prematurely or not).

III. RESULTS

In this section we present and discuss our results. All results are averaged over five simulations. Though variance is not shown, we have calculated it and judged it sufficiently small and will therefore not discuss it further. We begin with presenting the results without modelling user behaviour.

A. Without user behaviour modelled

In Figure 1 we see the resulting metrics depicted for various loads and limits in admission control. The total goodput does not increase with increasing arrival rates above 20 arriving sessions per second, though some increase can be seen as the maximal number of users is increased. The reason for this is obvious when looking at Figure 1(c) which shows that the number of active users does not increase with increased arrival rates for loads above that level. In fact, the limit in maximum number of users is reached when the arrival rate is increased above 20 new sessions per second. This has an effect in total satisfaction (as seen in Figure 1(b)) which drops to a minimum when the arrival rate and the maximum number of users is increased. Since the satisfaction index is based on per-user goodput which must decrease if the total goodput is kept constant and the number of users is increased this behaviour is expected. We believe that this is the reason that maximising total goodput and maximising total satisfaction requires completely different settings to the admission control. Thus, in order to maximise total goodput no restriction should be set to the number of users. The influence of admission control is obvious when looking at the number of active users and total satisfaction, but less so when referring to total goodput.

When we use the minimum mean throughput admission control mechanism and still not model user behaviour we obtain the results showed in Figure 2. A comparison between Figure 2(a) and Figure 1(a) shows that for both setups, the system is saturated in total goodput for loads above 20 arriving sessions per second. The total goodput seems slightly less when the admission control is based on minimal mean throughput though (Figure 2(a)). This fact is further supported as Figure 2(c) reveals that the number of active users does not increase with loads above 20 arriving sessions per second. This also shows that the admission control mechanisms works in that way that it limits the number of active users in the system. Setting the minimum mean throughput to a low value maximises total goodput.

Comparing Figures 2(c) and 1(c) we see that the latter accepts many more users. This suggests that the range for the minimum mean throughput mechanism could have included values below 100 kbit/s in order to accept more users.

B. With user behaviour modelled

When user behaviour is modelled the goodput is higher than when user behaviour is not modelled. This can be seen when comparing Figures 1(a) and 3(a). This effect is expected since users experiencing poor radio conditions will end their sessions prematurely, giving resources to users with better radio conditions. The user behaviour is actually working as a pre-emption mechanism, which raises the question whether implementing such a mechanism in the Node B is needed, given that our model of user behaviour is accurate enough.

In Figure 3(b) there is a "ridge" when the maximum number of users is set to 15. When the maximum number of users is set to 10 the users are too few to generate enough satisfaction

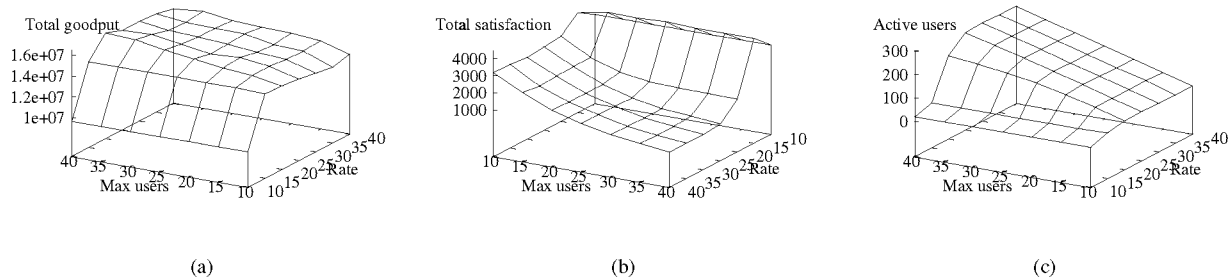


Fig. 1. The images show the total goodput (in bits/s), the total user satisfaction and the mean number of active users during a simulation. These are all averaged over all of the simulations when the admission control is based on the maximum number of users and no user behaviour is modelled. Figure 1(b) is rotated to increase the readability.

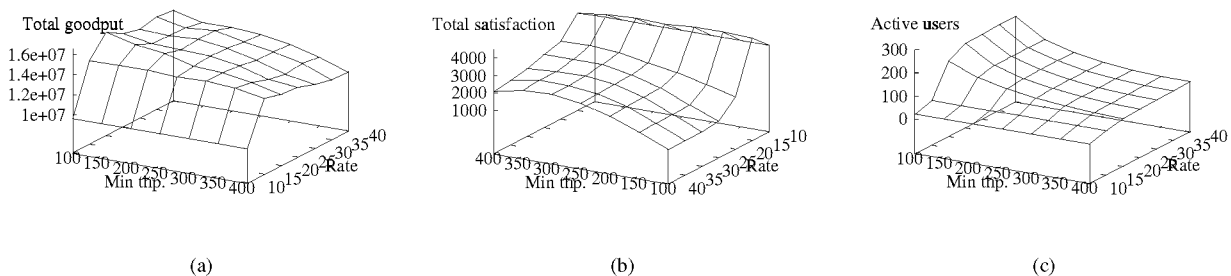


Fig. 2. The images show the total goodput (in bits/s), the total user satisfaction and the mean number of active users during a simulation. These are all averaged over all of the simulations when the admission control is based on the minimum mean throughput with no user behaviour modelled. Figure 2(b) is rotated to increase the readability.

and when the limit is set above 15 the increased competition immediately influences the total satisfaction negatively. This ridge also appears in Figures 2(b) and 4(b) around 300 kbit/s but not as prominent. We think that if we had set the maximum number of users below 10, the ridge might have appeared in Figure 1(b) as well. If the objective is to maximise user satisfaction, a minimum mean throughput of 300 kbit/s seems like a good setting.

Figure 3(c) tells us that the admission control only affects the number of active users for low limits (less than a maximum of 25 users) and high loads (arrival rates above 25–30 sessions per second). The reason that this figure does not look like Figure 1(c) is because of the user behaviour which dampens the number of active users by shortening the sessions of the users with poor goodput.

Minimum mean throughput has been used as the admission control algorithm to produce the results shown in Figure 4. In the first figure, Figure 4(a), we can see that the total goodput is as high as in Figure 3(a). Clearly, the choice of admission control mechanism does not affect total goodput much. Instead, the choice of including a model for user behaviour or not has a far greater impact on total goodput. We also notice that the system reaches a saturation point for high loads and the minimum mean throughput set high. This can also be seen in Figure 4(c). For a minimum mean throughput of 250 kbit/s the number of active users does not increase beyond arrival rates of 30 new sessions per second.

Looking at both Figures 2(a) and 4(a) we notice that the

admission control has a larger impact on total goodput during higher loads than during lighter loads. This effect is the result of our limit in utilisation. For lighter loads the utilisation never reaches 98% for longer periods of time, which is the limit for the admission control to be employed.

The ridge in total satisfaction appears for this setup as well. In Figure 4(b) we can see that the total satisfaction is kept relatively constant when the minimum mean throughput is set to 300 kbit/s. The reason behind this is the same as before. For all three figures we note that admission control helps keep the total user satisfaction high, but its effect on total goodput is questionable.

IV. CONCLUSIONS

This paper investigates the need for admission control in the interactive traffic class for HS-DSCH to improve and optimise total goodput and user satisfaction. We test through simulations the assumption that existing load control features provided by TCP and the users themselves are sufficient to optimise the system utilisation. Goodput is measured over completed file transfers and user satisfaction is computed using a simple utility function.

The simulations show that admission control may not improve the total system goodput. Instead, in our simulations, more users in the system result in higher total goodput. It should however be noted that optimising the total goodput alone is not feasible since the average goodput will then drop below what is acceptable for the users of the system.

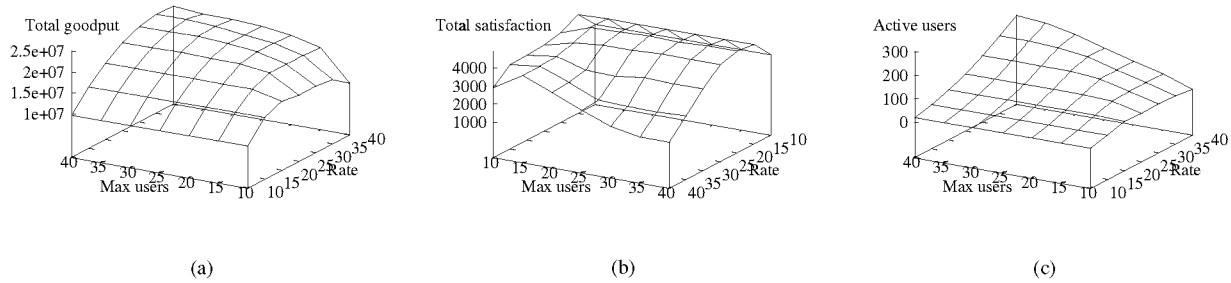


Fig. 3. The images show the total goodput (in bits/s), the total user satisfaction and the mean number of active users during a simulation. These are all averaged over all of the simulations when the admission control is based on the maximum number of users with user behaviour modelled. Figure 3(b) is rotated to increase the readability.

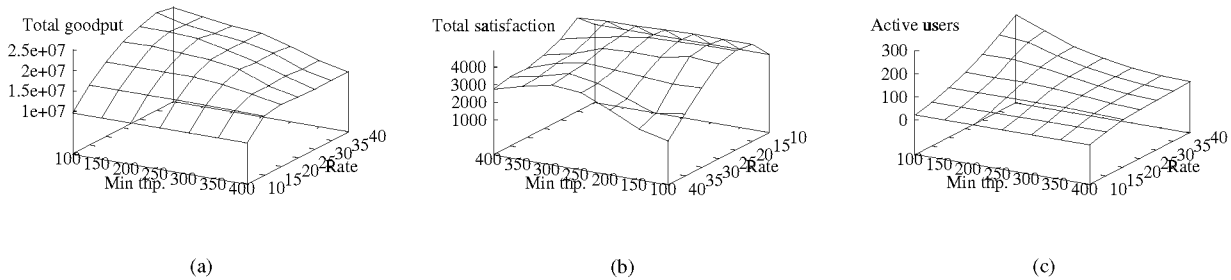


Fig. 4. The images show the total goodput (in bits/s), the total user satisfaction and the mean number of active users during a simulation. These are all averaged over all of the simulations when the admission control is based on the minimum mean throughput with user behaviour modelled. Figure 4(b) is rotated to increase the readability.

In contrast to total goodput, the total user satisfaction in an HS-DSCH system can be optimised using admission control as illustrated by the simulations. Hence, our assumption that existing load control features provided by TCP and the users themselves are sufficient to optimise the system utilisation is wrong.

Users finish their sessions prematurely due to low transmission quality means that the number of users experiencing bad radio conditions is reduced. Consequently, the system goodput is higher when this user behaviour is modelled compared to when users always finish their sessions. This means for the evaluated admission control mechanisms that admission limits should be set to accept more users when it can be assumed that users experiencing low goodput finish before sessions are fully completed. Our model for user behaviour is probably not the correct one, but we think it is reasonable and thus it serves to prove that knowledge about user behaviour is essential.

A proportional fair scheduler weighted to account for the SIR of each individual user would offer higher variance in throughput over time and thus render different results for the two admission control mechanisms tested herein. We are currently implementing such a scheduler and intend to include new results analysing this issue in the future.

REFERENCES

- [1] Technical Specification Group Radio Access Network (TSG-RAN), "Medium Access Control (MAC) protocol specification (Release 6)," 3rd Generation Partnership Project (3GPP), Tech. Rep. TS 25.321, Dec. 2004.
- [2] T. E. Kolding, K. I. Pedersen, J. Wigard, F. Frederiksen, and P. E. Mogensen, "High Speed Downlink Packet Access: WCDMA Evolution," *IEEE Vehicular Technology Society News*, vol. 50, no. 1, pp. 4–10, Feb. 2003.
- [3] L. Breslau and S. Shenker, "Best-effort versus reservations: A simple comparative analysis," in *SIGCOMM*, 1998, pp. 3–16. [Online]. Available: citeseer.ist.psu.edu/breslau98besteffort.html
- [4] X. Gao, T. Nandagopal, and V. Bharghavan, "Achieving application level fairness through utility-based wireless fair scheduling," in *Global Telecommunications Conference*. San Antonio: IEEE, Nov. 2001, pp. 3257–3261.
- [5] R. R.-F. Liao and A. T. Campbell, "A utility-based approach for quantitative adaptation in wireless packet networks," *Wireless Networks*, vol. 7, no. 5, pp. 541–557, 2001. [Online]. Available: citeseer.ist.psu.edu/liao01utilitybased.html
- [6] P. A. Hosein, "QoS Control for WCDMA High Speed Packet Data," in *International Workshop on Mobile and Wireless Communications Network*, San Diego, CA, USA, Sept. 2002, pp. 169–173.
- [7] S. McCame and S. Floyd, "The Network Simulator – ns-2," <http://www.isi.edu/nsnam/ns>.
- [8] Motorola, "Evaluation Methods for High Speed Downlink Packet Access (HSDPA)," 3GPP, Tech. Rep., Jul. 2000.
- [9] A. Reyes-Lecuona, E. Gonzalez-Parada, E. Casilari, J. C. Casasola, and A. Daz-Estrella, "A page-oriented WWW traffic model for wireless simulations," in *16th ITC*, Edinburgh, Jun. 1999, pp. 1271–1280.
- [10] K. Fall and S. Floyd, "Simulation-based Comparisons of Tahoe, Reno and SACK TCP," *Computer Communications Review*, vol. 26, no. 1, pp. 5–21, Jul. 1996.
- [11] M. Allman, H. Balakrishnan, and S. Floyd, "Enhancing TCP's Loss Recovery Using Limited Transmit," IETF, RFC Standards track 3042, Jan. 2001.
- [12] E. Blanton, M. Allman, K. Fall, and L. Wang, "A Conservative Selective Acknowledgment (SACK)-based Loss Recovery Algorithm for TCP," IETF, RFC Standards Track 3517, Apr. 2003.