Optimal Radio Access Technology Selection on Heterogeneous Networks

Diego Pacheco-Paramo, Vicent Pla, Vicente Casares-Giner and Jorge Martinez-Bauset

Dept. Comunicaciones, Universidad Politecnica de Valencia
Camino de Vera s/n, 46022, Valencia, Spain
Email: diepacpa@posgrado.upv.es,{vpla,vcasares,jmartinez}@dcom.upv.es

Abstract

The joint management of radio resources in heterogeneous networks is considered to improve their capacity. We propose joint schemes for admission control and access technology selection with vertical handoffs. Optimal policies are found for wireless networks that support two access techniques and cover the same geographical area. In addition, the system under study also supports heterogeneous traffic of two types: streaming and elastic. We explore the optimization of different functions expressed in terms of blocking probabilities and throughput. An exhaustive numerical analysis allows us to characterize the optimal admission policies according to the arrival type and system state. Based on this characterization, heuristic policies are designed and their performance is compared to the one obtained by previously proposed schemes. This analysis is also done when constraints, expressed in terms of blocking probability bounds, are added. An extension of the previous system that includes vertical handoffs, in order to evaluate their impact on the system performance, is also studied. For the four types of vertical handoffs considered, we determine and characterize the optimal policies according to the arrival type, system state and vertical handoff action. Since it is not computationally feasible to calculate the optimal policies online, new heuristic policies with vertical handoffs are design and evaluated. It is found that the heuristic policies scale with the system size without requiring any adjustment, their performance is very close to the one obtained by the optimal policies and they are simple to implement, and, therefore, can be used online. In addition, we find that the heuristic policies are insensitive to the service time of the voice sessions and the elastic flow sizes beyond the mean. Finally, in order to take into into account the cost of performing vertical handoffs, a new optimization problem is formulated that relates the costs of voice and data blocking with the costs of vertical handoffs.

Keywords: joint call admission control, vertical handoff, resource allocation, Markov decision process

1. Introduction

The new generation of wireless systems allow the subscribers to exploit the availability of multiple radio access technologies (RATs) in the same geographical area, in order to achieve a permanent connection with the best possible performance [1]. The coexistence of various RATs has promoted the interest in the joint management of their radio resources. Common radio resource management (CRRM) has been proposed as a model to control several RATs in a coordinated way, exploiting the fact that this coordination allows an improvement in the use of the limited resources [2]. Among the CRRM functions, there is a special interest in the joint call admission control (JCAC), which in this new scenario not only defines if a session should be accepted or not, but also in what technology [3].

Some JCAC proposals in heterogeneous networks, such as those found in [4] and [5], were designed to achieve a certain degree of load balancing among the different access networks, which translates into a better resource utilization. More specifically, in [5] several JCAC policies have been studied for a scenario similar to ours and their performance measured in terms of throughput and blocking probability. However, while in [5] only simple heuristic policies are considered, we follow an optimization approach.

In [6] a Markov model is used for an heterogeneous network (WLAN and CDMA), and linear programming is used to solve the optimization problem. This study explicitly acknowledges the inherent complexity of the problem which could make it computationally in-tractable for large systems. This fact leads the authors to propose reinforcement learning as a possible solution method. Our study differs from the one in [6] in that our main objective is the characterization of the optimal policies and the design of heuristic policies. In [7] they formulate an optimization problem as a Markov decision process and solve it using genetic algorithms. However, the objective of their study is different from ours, and therefore the optimization functions depend on different parameters.

Finally, fuzzy-based solutions are used in [8] and [9]. These solutions take different user and system parameters, weight them, and produce an online decision. On the other hand, our interest is to find optimal solutions, characterize them and synthesize heuristic solutions that can minimize the blocking probabilities, or bound them, and maximize the throughput.

Our work is in part motivated by the study in [10]. Although we share similar evaluation scenarios, the study in [10] focuses on the analysis of a set of heuristic policies, while we follow an optimization approach.

Another important difference is that we explore the use of vertical handoffs to improve further the system performance, which was not considered in the studies mentioned above.

Six are the main contributions of our study. First, we apply an optimization approach based on the formalism of Markov Decision Processes (MDPs) to search for the optimal policy, instead of common heuristic approaches. The large cardinality of the state space of the MDPs that model our system makes it unfeasible to find numerical solutions online. Therefore, our interest is the characterization of the optimal policies and the design of heuristic ones. Policy iteration is used to solve the MDPs, a method that does not depend on the initial conditions and always finds the optimal solution in a finite, and usually small, number of steps.

Second, we design heuristic policies whose performance scales very well with the system size. They are also very simple to implement, and therefore can be used online, and their performance is very close to the optimal policies performance.

Third, we compare the performance of the heuristic policies obtained with our approach with the performance of heuristic policies proposed in the literature. This is also done when constraints, expressed in terms of blocking probability bounds, are added. The results show that, in general, our policies clearly outperform the heuristic policies proposed previously.

Fourth, we study an extension of the previous system that includes vertical handoffs, in order to evaluate their impact on the system performance. Most vertical handoff schemes, like those in [11, 12], privilege user preferences. However, we consider that the operator's point of view, that focuses on improving the resource utilization, is also important. For the four types of vertical handoffs considered, we determine and characterize the optimal policies according to the arrival type, system state and vertical handoff action. Since it is not computationally feasible to calculate the optimal policies online, new heuristic policies with vertical handoffs are designed and evaluated.

Fifth, we validate the analytical study by simulation. We evaluate the impact that the distributions of the service time of the voice sessions and of the elastic flow length (in bits) have on the performance of the heuristic policies. An interesting and very important finding is that their performance is insensitive to distributions of these random variables beyond the mean.

Sixth, since vertical handoffs increase the system complexity [3], it is necessary to evaluate it. In order to take into into account the costs of performing vertical handoffs, a new optimization problem is formulated that relates the costs of voice and data blocking with the costs of vertical handoffs. The solution of this optimization problem allows to determine the range of costs for which it is advisable to use vertical handoffs and of what types. Note that in previous studies like [13], a monetary cost is defined, which allows the subscribers to decide which technology is more suitable. However, we believe that our approach is more general.

The rest of the paper is structured as follows. In Section 2 we describe the Markov model according to the characteristics of the system. The solution method is described in Section 3. In Section 4 the results of the optimization are shown, and analyzed for a static scenario. The optimal solutions for various traffic values are studied in Section 5. In Section 6, we add QoS restrictions to the optimization problem and analyze the results. Section 7 introduces vertical handoff in the action set, new optimal policies are found and their characterization is obtained. New heuristic solutions are described and evaluated in different scenarios and simulations are performed in order to validate analytic results. To understand the cost impact of vertical handoffs, in Section 8 we formulate a new optimization problem that relates the costs of voice and data blocking with the costs of vertical handoffs. Finally, in Section 9 we present the conclusions of the study.

2. System description and Markov model

We study a system with two wireless access networks that use two multiple access techniques: TDMA and

WCDMA. Their characteristics allow us to model technologies such as GSM and UMTS. However, the proposed model and analysis can be extended to other multiple access techniques like OFDMA, which has been defined to support, for example, WiMAX and LTE. Additionally, both access networks provide voice and data services in the same area, as it was proposed in [10]. When a session arrives, a decision has to be made on whether it should be accepted or not, and also in what technology should be served. Initially, we consider that a decision is made once a session starts and will be held until it terminates, i.e. there are no vertical handoffs.

In order to obtain an analytically tractable model, we assume that voice (data) session arrivals follow a Poisson process with rate $\lambda_v(\lambda_d)$. We also assume that the service time for voice sessions is exponentially distributed with mean $1/\mu_v$. On the other hand, as data sessions generate elastic traffic, their sojourn time will depend on the available resources. The size of the flows generated by the data sessions is exponentially distributed with mean σ (in bits). If BR_d is the data bit rate experienced for a given user, then the service time will be exponentially distributed with mean $1/\mu_d = \sigma/BR_d$. Clearly, BR_d might depend on the system state as discussed later.

2.1. State space

The system is modeled as a 4 dimensional continuous time Markov chain (CTMC), with states represented by the vector $\boldsymbol{s} = (s_1,s_2,s_3,s_4)$ where s_1 represents the number of ongoing voice sessions on TDMA, s_2 the data sessions on TDMA, s_3 the voice sessions on WCDMA and s_4 the data sessions on WCDMA.

We define C as the fixed number of channels in TDMA. A voice session will always use a whole channel, so there can only be C simultaneous voice sessions on this technology. On the other hand, data sessions can share a channel when TDMA is at full capacity, in such a way that n_c data sessions can be served per channel. This means that a maximum of $C \cdot n_c$ simultaneous data sessions can be active in TDMA. According to this, the first condition that a state must fulfill to be feasible is given by

$$s_1 \cdot n_c + s_2 \leq n_c \cdot C. \tag{1}$$

The capacity on WCDMA is defined by

$$s_3\left(\frac{W/BR_{w,v}}{(E_b/N_0)_v} + 1\right)^{-1} + s_4\left(\frac{W/BR_{w,d}}{(E_b/N_0)_d} + 1\right)^{-1} \leq \eta_{ul}, \tag{2}$$

where W is the chip rate, $BR_{w,x}$ is the bit rate used for transmitting service x in WCDMA, $(E_b/N_0)_x$ is the bit energy to noise density required for service x, and η_{ul} is the uplink cell load factor. This is the same expression used in [14].

Considering that each technology has independent resources, the feasible combination of data and voice users can be determined for each technology separately. We define S as the set of feasible states, that is all the state vectors $\boldsymbol{s}$ that fulfill the conditions defined in (1) and (2).

OFDMA-based systems commonly operate as FDMA-TDMA systems, where the resources to be shared are organized as a two-dimensional matrix of sub-channels and time-slots that repeats every frame (e.g., every 5 ms in WiMAX) [15, 16]. In addition, the same resources (sub-channels and time-slots) can be shared by different sessions in consecutive frames. Although the support of preambles imposes some difficulties to share the resource efficiently, effective algorithms have been proposed to solve this problem [16]. Then, the admission control problem in OFDMA-based systems can be formulated as in (1), where each accepted session of each service category consumes a certain constant amount of resources in the two- dimensional resource matrix [17, 18].

2.2. System metrics

Different performance parameters can be determined, once the steady state probabilities of the CTMC have been obtained. In particular, we are interested in the voice blocking probability, the data blocking probability, and the total throughput. The blocking probability refers to the fraction of sessions initiation requests that are blocked. To calculate the throughput we have to consider that the bit rate is independent for each service and technology and that data sessions in TDMA can share a channel, which is reflected in the min $(C - s_1, s_2)$ factor of the following equation:

$$\begin{aligned} Th = \sum_{s \in S} P(\boldsymbol{s})(&s_1 BR_{t,v} + s_3 BR_{w,v} \\ &+ \min(C - s_1, s_2) BR_{t,d} + s_4 BR_{w,d}), \end{aligned} \tag{3}$$

where $P(s)$ is the steady state probability of being in state s and $BR_{x,y}$ is the bit rate used for transmitting service y (voice or data) in technology x (TDMA or WCDMA).

3. Optimization problem

Voice and data arrivals can occur at any state of the CTMC defined in the last section. Then, at each state,

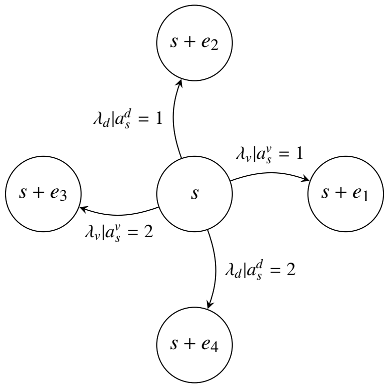

Figure 1: Transitions for voice and data arrivals.

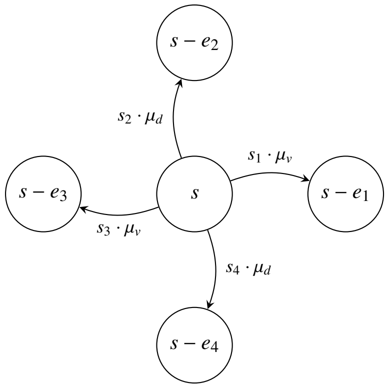

Figure 2: Transitions for voice and data departures.

a decision must be made on whether an arriving session should be admitted and the appropriate technology. In a Markov decision process (MDP), a policy π defines which actions a=(a_s^v,a_s^d) should be taken at each state s when voice (a_s^v) or data (a_s^d) arrivals occur. It should be clear that decision epochs occur only at arrivals. The main objective is to find among the possible policies the one that optimizes a chosen function.

In this system, the set of actions A defines the possible values for a_s^v and a_s^d, and it is defined in Table 1. Figure 1 shows the transitions from state $\boldsymbol{s}$ for arrivals. These transitions are clearly conditioned by a_s^v and a_s^d, where e_i is a 4 dimensional vector of zeros with a 1 on the i–th position. The transitions from state $\boldsymbol{s}$ for departures are shown in Figure 2.

We define now three different admission schemes and set the values taken by a_s^x accordingly, where x could be voice or data. These policies were suggested and analyzed in [10]. Let AC_{TDMA} and AC_{WCDMA} be the available free capacity in TDMA and WCDMA, respectively.

Table 1: Set of possible actions A

value	*action*
0	Block session
1	Send session to TDMA
2	Send session to WCDMA

Scheme i : sessions are sent to TDMA (if possible)

$$a_s^x = \begin{cases} 1 & \text{if } AC_{TDMA} > 0 \\ 2 & \text{if } AC_{TDMA} = 0 \text{ and } AC_{WCDMA} > 0 \\ 0 & \text{if } AC_{TDMA} = AC_{WCDMA} = 0 \end{cases}$$

Scheme ii: sessions are sent to WCDMA (if possible)

$$a_s^x = \begin{cases} 2 & \text{if } AC_{WCDMA} > 0 \\ 1 & \text{if } AC_{WCDMA} = 0 \text{ and } AC_{TDMA} > 0 \\ 0 & \text{if } AC_{TDMA} = AC_{WCDMA} = 0 \end{cases}$$

Scheme iii: sessions are sent to the technology with lower occupation

$$a_s^x = \begin{cases} 1 & \text{if } AC_{TDMA} > AC_{WCDMA} \\ 2 & \text{if } AC_{WCDMA} > AC_{TDMA} \\ \text{random} & \text{if } AC_{WCDMA} = AC_{TDMA} > 0 \\ 0 & \text{if } AC_{WCDMA} = AC_{TDMA} = 0 \end{cases}$$

Using these schemes we define three policies:

Policy	Scheme for voice	Scheme for data
1	*i*	*ii*
2	*ii*	*i*
3	*iii*	*iii*

It should be noted that Policies #1 and #2 exploit the hypothesis that each technology is more appropriate for a specific service, which is not the case of Policy #3 that focuses on improving radio resource utilization.

3.1. Cost function

The state space of the MDP is defined by (1) and (2), and the possible set of actions and the transition rates associated to them are defined in Table 1 and Figures 1 and 2. Since our interest relies on data and voice blocking probabilities, as well as the total throughput, we have define two different objective functions. The first one, is the weighted sum of the voice and data blocking probabilities (BP_v and BP_d),

$$F_{BP} = BP_v \cdot \alpha + BP_d \cdot (1 - \alpha). \quad (4)$$

The parameter α, $0 \leq \alpha \leq 1$, is the one responsible for giving more or less weight to each blocking probability. When α is closer to 0, minimizing the objective function will have a bigger impact on data blocking probability than on voice blocking probability, and the opposite will happen when α is closer to 1.Thus, α relates the way in which the blocking probabilities will be minimized. The cost function associated to the objective function for each feasible state $\boldsymbol{s}$ is

$$\text{cost}(\boldsymbol{s}) = 1 - (\alpha \cdot F_v(a_s) + (1 - \alpha) \cdot F_d(a_s)), \quad (5)$$

where $F_x(a_s)$=1 if a_s is 1 or 2, and 0 otherwise, being x the service.

The second objective function is the aggregated throughput, so in that case we try to maximize the value defined by (3). The reward for each state $\boldsymbol{s}$ is

$$\begin{aligned} \text{cost}(\boldsymbol{s}) = \; & s_1 BR_{t,v} + s_3 BR_{w,v} \\ & + \min(C - s_1, s_2) BR_{t,d} + s_4 BR_{w,d}. \end{aligned} \quad (6)$$

Therefore, the reward associated to each state when maximizing the throughput does not depend on the actions taken on that state. It must be noted that we use the term *cost* for notational purposes, even when a reward is expected from throughput.

3.2. Solution method

Policy iteration [19] to find the optimal policy π_{opt}. This method can search among the finite group of possible policies for the MDP and find the optimal in a finite number of steps. The relative values V allow to relate the cost obtained in the actual state with costs expected from future actions, and are found using the next equation:

$$\mathbf{c}^{\pi} - c_{\pi} \cdot e + V_{\pi} R_{\pi}^T = 0. \quad (7)$$

The $\mathbf{c}^{\pi}$ in the previous expression is the vector of costs associated to being in each state, R_π is the transition matrix, and c_π is the value of the objective function for policy π.

Once V and c_π are found using (7), it is possible to find the action on each state that will minimize the objective function using the next expression:

$$\min_a \{c_s^a - c_{\pi} + \sum_{s \neq u} r_{s,u}^a (v_{\pi}^u - v_{\pi}^s)\}, \quad (8)$$

Table 2: Policy iteration algorithm.

Step 1.	Choose an arbitrary policy π_i
Step 2.	Calculate relative values V_{π_i} and the mean cost (revenue) c_{π_i} for the initial policy π_i using expression 7
Step 3.	Find action a for each state $\boldsymbol{s}$ that minimize (maximize) the expression in 8. The resulting policy is called π_f.
Step 4.	If the resulting policy π_f differs from the initial policy π_i, go back to step 2, using π_f as the initial policy. If not, π_f is the optimal policy.

where v_π^u and v_π^s are the relative values for states u and s respectively when the policy π is used, c_s^a is the cost associated to state s when action a is taken, $r_{s,u}^a$ is the transition rate from state s to state u, when the action for state s is a. The set of actions will define a new policy and the process is repeated until the optimal policy is found. In Table 2 the policy iteration algorithm is shown. It is important to recall that the optimal policy will be found no matter what initial policy is used. However, the number of iterations could change.

4. Optimal Policy Analysis: Static Scenario

The parameters of the system are defined in Table 3. As it can be seen, bit rates are independent of the technology used, and are higher for data than for voice. Also, the (E_b/N_0) required for both services in WCDMA are the same. The maximum voice and data capacities are of 4 and 8 users in TDMA and of 13 and 4 users in WCDMA. The system parameter values have been chosen in order to maintain a similar capacity on both technologies and to keep the optimization problem computationally tractable. In Table 11 we define a much larger system, which is analyzed in Section 7. The objective of this section is to explore the optimal policy when the two optimization functions are used: blocking probability and throughput.

4.1. Blocking function optimization

Unless otherwise stated, throughout the document we use a value of $\alpha = 0.5$. That is, the objective function will be defined by $F_{BP}(\pi)$=0.5 $\cdot$ BP_v+ 0.5 $\cdot$ BP_d. Fig. 3 shows the value of the objective function for every iteration until the optimal policy is achieved and results

Table 3: Initial Scenario for Policy Iteration.

WCDMA	TDMA
W=3.84 Mcps	$C = 4$
$(E_b/N_0)_v$=14 dB	n_c= 2
$(E_b/N_0)_d$=14 dB	$BR_{t,v}$=12.2 Kbps
$BR_{w,v}$=12.2 Kbps	$BR_{t,d}$=44.8 Kbps
$BR_{w,d}$=44.8 Kbps	
η_{ul}= 1	

Clients
λ_v= 0.025
λ_d= 0.134
μ_v= 0.0083
σ= 1 Mb

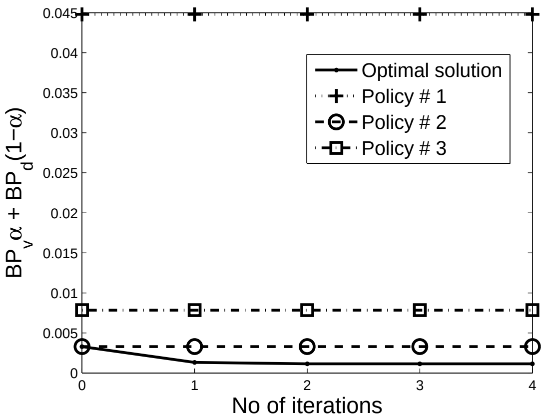

Figure 3: Blocking probability function optimization.

are compared with those obtained with the fixed policies described in Section 3. We can summarize the main characteristics of the optimal policy as follows:

Service	*Action*
Voice	scheme *ii*
Data	• if there is no channel sharing on TDMA: scheme *i* • if there is channel sharing on TDMA: ≅ scheme *iii*

The initial policy, Policy #2, has the closest performance to the optimal policy because they both send voice sessions to WCDMA (scheme *ii*). The decisions for data sessions depend on channel sharing since this reduces the transmission rate, which raises the residence

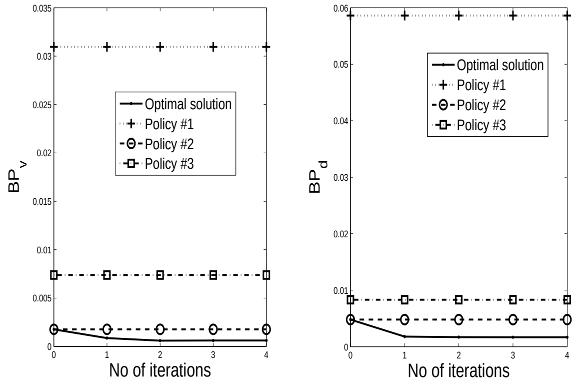

(a) *VoiceBlockingProbability*. (b) *DataBlockingProbability*.

Figure 4: Blocking Probabilities.

time and thus the blocking probability. Therefore, channel sharing produces a policy very similar to scheme *iii* where resources are used according to occupation.

Only 22.3% of a_s^d from Policy #2 differ from the optimal, and there is only a 6.1% probability for these actions to occur. However, results differ considerably. The cost for the optimal policy is 34.78% of the one obtained with Policy #2, and of course similar results are found on BP_v and BP_d, as it is shown in Figs. 4(a) and 4(b). This improvement also added 462 *bps* to the throughput obtained with Policy #2.

4.2. Throughput optimization

The characteristics of the optimal policy for throughput optimization, can be summarized as:

Service	*Action*
Voice	scheme *ii* with blocking
Data	• if there is no channel sharing on TDMA: scheme *i* • if there is channel sharing on TDMA: ≅ scheme *iii*

Although voice sessions are sent to WCDMA (scheme *ii*), there may be blocking when there is available capacity on TDMA. This happens in order to save capacity for data sessions, which contribute more to the total throughput. In the optimal policy, nearly 25 % of the states where no more capacity is available on WCDMA, even if there is capacity in TDMA, voice sessions are blocked. Also, since channel sharing on TDMA directly

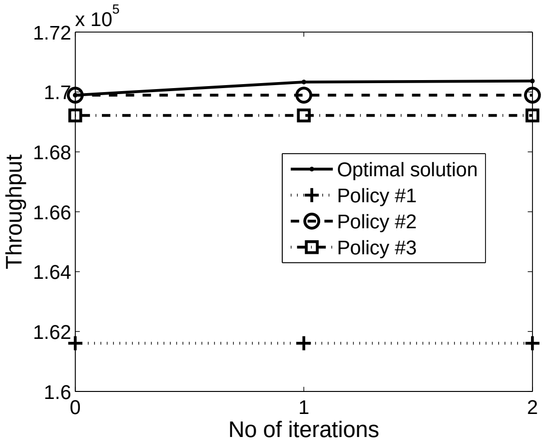

Figure 5: Throughput optimization.

reduces the total throughput, a policy similar to scheme *iii* is used for data sessions.

Figure 5 shows the throughput until the optimal policy is reached. After two iterations, the throughput raises to 170.366 Kbps, a gain of 474 bps over Policy #2. The policy obtained with the first iteration is very similar to the optimal, since the second iteration only improves the total throughput in 35.8 *bps*, about 8% of the 438.9 *bps* gained before. Therefore, it is possible to use this sub-optimal solution in order to solve larger systems, in a similar approach as it was used in [20]. An interesting fact of throughput optimization, is that it also improves BP_v and BP_d. Since these parameters have very low values in the initial policy the improvement is small, going from 0.17 % to 0.079 % for BP_v and from 0.48 % to 0.15 % for BP_d, but it shows their correlation with the total throughput. Actually, the BP_d is lower than the one obtained optimizing the blocking function (0.167 %), indicating the importance of accepting data sessions for the improvement of the total throughput.

5. Optimal Policy Analysis: Parameter Variation

In this section we evaluate the performance of the optimal policies and compare it with the one obtained by the heuristic policies defined in Section 3, which were propose previously in the literature. We also characterize the optimal policies obtained for each load point, and based on their dynamic behavior, we define a heuristic policy.

5.1. Voice users arrival rate variation

Although they are not identical, the main characteristics of optimal policies for both optimization functions as λ_v grows can be summarized as:

Service	*Lower λ_v*	*Higher λ_v*
Voice	scheme *ii*	scheme *ii* with blocking
Data	If there is no channel sharing on TDMA: scheme *i*	If there is no channel sharing on TDMA: scheme *i*
	if there is channel sharing on TDMA: ≅ scheme *iii*	if there is channel sharing on TDMA: ≅ scheme *iii*, preferably TDMA

Figure 6 shows the results obtained with the optimal policy for the blocking optimization function, and the values obtained with the fixed policies. When λ_v reaches 0.047, the offered load for voice is 5.64 *Erl*, almost two times the offered load for data (3 *Erl*). Given the great impact on data sessions, the optimal policy blocks voice sessions, and this is more intense for higher values of λ_v. For λ_v=0.047 only 3 of 80 states where voice sessions were sent to TDMA according to scheme *ii*, decide to block in the optimal policy. At this point, BP_v=0.2 % and BP_d=1 %. When λ_v grows to 0.065, 26 of those 80 states block voice sessions, that is eight times more states than before. As expected, BP_d only grows to 2.5 %, but the impact on BP_v is bigger, since it grows almost ten times, to 2.056 %.

Dependence on channel sharing for a_s^d is maintained for all λ_v, but once sharing is mandatory, some changes occur. As λ_v grows, the number of states that decide to send data sessions to TDMA under this circumstances, grows from 180 of 400 when λ_v=0.005 to 225 of 400 when λ_v=0.065. This is done in order to reduce BP_v, since WCDMA capacity is saved for voice sessions. It is also worth noting that when TDMA is almost at full capacity (>95%), data sessions are blocked. This effect can be seen for the full range of λ_v values that appear on Fig. 6.

Figure 7 shows the throughput when it is used as the optimization function for the optimal policy and the three fixed policies as λ_v grows. As λ_v grows, it becomes mandatory to save resources for data sessions since they have a higher throughput. For this reason, voice sessions may be blocked even when there is room on TDMA. For λ_v=0.005, 14 of 80 states where occupancy on WCDMA is full but there is still room on TDMA, block voice sessions. When λ_v grows to 0.065, 36 of the same 80 states block voice sessions. This affects BP_v, and when λ_v>0.041, the optimal policy has a higher BP_v than Policy #2.

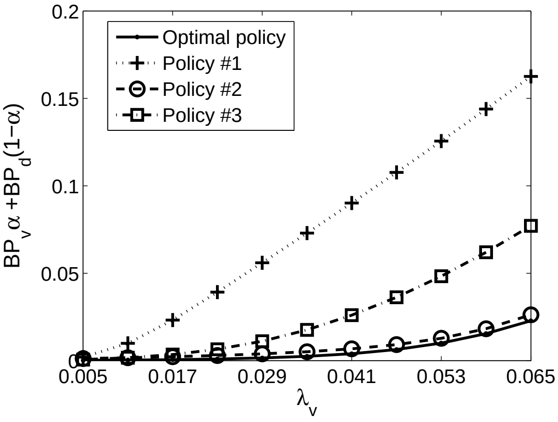

Figure 6: Blocking function for various λ_v.

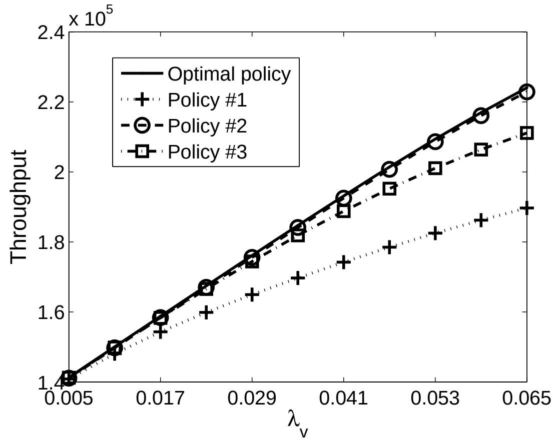

Figure 7: Throughput for various λ_v.

Dependence on channel sharing for a_s^d is maintained for all λ_v, but when sharing is mandatory, optimal policies vary. As λ_v grows, the occupation of WCDMA does too, so more data sessions must be sent to TDMA in order to increase the total throughput. When λ_v=0.005, 144 of the 400 states that use TDMA channel sharing send data sessions to TDMA. For λ_v=0.065, 212 of the same 400 states send sessions to TDMA.

Therefore, for both optimization criteria, voice sessions are sent to WCDMA, with more blocking states as λ_v grows, and data sessions are sent to TDMA while no sharing is needed. When sharing is mandatory, data sessions may be sent to WCDMA or TDMA with more states using TDMA as λ_v grows.

5.2. Data users arrival rate variation.

Although not identical, the main characteristics of the optimal policies for both optimization function as λ_d grows can be summarized as:

Service	*Lower* λ_d	*Higher* λ_d
Voice	scheme *ii*	scheme *ii*
Data	If there is no channel sharing on TDMA: scheme *i*	If there is no channel sharing on TDMA: scheme *i*
	if there is channel sharing on TDMA: $\cong$ scheme *iii*	if there is channel sharing on TDMA: $\cong$ scheme *iii*, preferably WCDMA

Figure 8 compares the blocking function values for the optimal policy with different values of λ_d with those obtained with the three fixed policies. In this scenario, voice sessions follow scheme *ii* for every λ_d, showing that the load of data sessions has no impact on a_s^v. On the other hand, the big influence of TDMA occupancy over a_s^d, suggest the opposite. In fact, when λ_d=0.2, data sessions arrive with a higher rate than voice sessions, so in order to maintain a low objective function value, it is necessary to send more data sessions to WCDMA so TDMA's capacity is not exhausted too fast. This is the opposite case of the one we saw in the last section where as λ_v grew, more data sessions were sent to TDMA. The second effect is that as λ_d grows, it is necessary to block sessions in order to minimize the optimizing function. For the values of Fig. 8, the system has 1000 states. Only 70 of the total do not allow data sessions, thus these are data blocking states. For the optimal policy, while $\lambda_d \leq 0.095$, data sessions are only blocked on those 70 states. Once this value is surpassed, the optimal policy may block data sessions even when there is capacity left. When λ_d=0.11, 4 of the 930 states with capacity left decide to block data sessions, and this number grows to 30 of those 930 states when λ_d=0.2, more than seven times the initial number of blocking states.

Figure 9 shows the throughput when it is used as the optimization function for various λ_d. Since data sessions have a higher throughput, the optimal policy may penalize some voice sessions to maximize the objective function. When λ_d=0.2, BP_v is higher in the optimal policy than in Policy #2, reaching 2%. In fact, if we let λ_d=0.275, BP_v for the optimal policy would be higher than any of the fixed policies. In order to maximize the

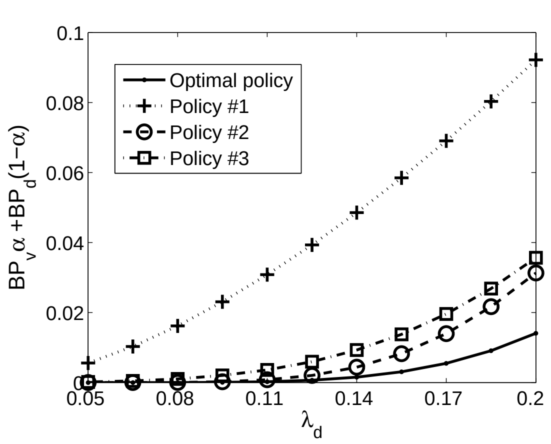

Figure 8: Blocking function for various λ_d.

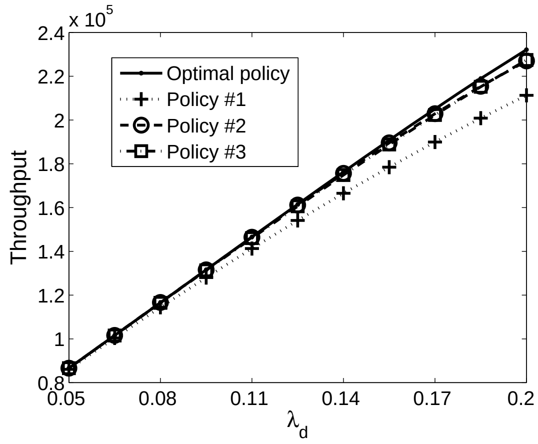

Figure 9: Throughput for various λ_d.

throughput, the opposite happens for BP_d, which is always the lowest in the optimal policy. As usual, voice sessions are treated according to scheme *ii*, but as λ_d grows higher than 0.05, instead of sending sessions to TDMA when WCDMA is full, they may be blocked. The number of states doing this grows slowly, but when λ_d=0.3 no voice sessions are ever sent to TDMA, saving the space for data sessions.

TDMA occupancy influence on a_s^d is increased by λ_d. When sharing becomes mandatory, as λ_d grows there seems to be a preference for sending sessions to WCDMA using the capacity saved by voice sessions. As a consequence of the importance of data sessions for throughput, for the values of λ_d in Fig.9, the optimal policy never blocks data sessions unless there is no other chance. That is, as λ_d grows voice sessions are always sent to WCDMA, but some blocking may appear when we optimize the throughput. On the other hand λ_d does not affect decisions for voice sessions when the blocking function is optimized. Data sessions are sent to TDMA while no sharing is needed, but when sharing is mandatory, sessions may be sent to TDMA or WCDMA, but as λ_d grows, a clear tendency towards WCDMA is seen for both optimization criteria.

5.3. Voice users service rate variation

Figure 10(a) shows the blocking function value while maintaining a constant offered traffic of 3 *Erl* for both services (as the static scenario), and changing μ_v from 0.0833 to 1.833. It is easy to realize how changes on μ_v have little or no impact in the blocking function for different policies. This behavior is a result of what happens with the blocking probabilities and the throughput, where the results are almost constant for the full range of μ_v. The optimal policies can be summarized as:

Service	*Lower* μ_v	*Higher* μ_v
Voice	scheme *ii*	scheme *ii* with blocking
Data	If there is no channel sharing on TDMA: scheme *i*	If there is no channel sharing on TDMA: scheme *i*
	if there is channel sharing on TDMA: ≅ scheme *iii*	if there is channel sharing on TDMA: ≅ scheme *iii*.

As μ_v grows, some voice sessions are blocked. However, it only differs on 0.2% of the states when μ_v changes from 0.258 to 1.833. Also, μ_v has little influence over a_s^d. As μ_v grows, a small percentage of states (<5%) where sessions used to be sent to WCDMA are sent to TDMA. In Fig.10(b) it is shown the throughput as the optimizing function for a constant load of 3 *Erl* in both services, while μ_v changes. When throughput is optimized, BP_v can be higher in the optimal policy than in other policies. The main characteristics of the optimal policies can be summarized in the same way as in the blocking function optimization shown before in this section.

5.4. Data service rate variation

Fig.11(a) shows the blocking function value for a constant load of 3 *Erl* on both services while μ_d changes. To change this parameter, we kept $BR_{x,d}$ in

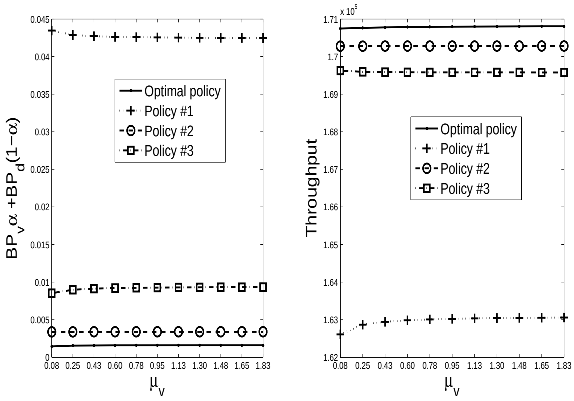

(a) *Blocking Optimization.* (b) *Throughput Optimization.*

Figure 10: Optimization for various μ_v.

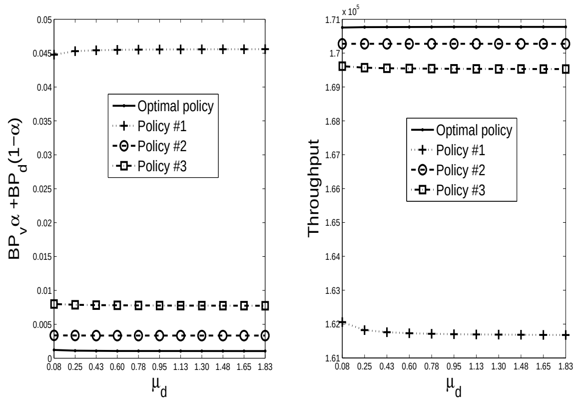

(a) *Blocking Optimization.* (b) *Throughput Optimization.*

Figure 11: Optimization for various μ_d.

44.8 Kbps, and changed the mean data length σ. The main characteristics of the optimal policies can be summarized as:

Service	*Lower μ_d*	*Higher μ_d*
Voice	scheme *ii*	scheme *ii* with blocking
Data	If there is no channel sharing on TDMA: scheme *i*	If there is no channel sharing on TDMA: scheme *i*
	if there is channel sharing on TDMA: ≅ scheme *iii*	if there is channel sharing on TDMA: ≅ scheme *iii.*

In Fig.11(b) appears the throughput (used as the optimization function) for the optimal and the fixed policies when μ_d changes and the load for each service is kept constant (3 *Erl*). The optimal policies are very similar to those obtained with the blocking function. For the lowest values of μ_d in Fig.11(b) most of the states where the capacity of WCDMA is full, decide to block sessions even when there is space left on TDMA. As μ_d grows, these states decide to send voice sessions to TDMA. This blocking is done in order to protect data sessions, whose contribution to throughput is higher.

5.5. Heuristic policy

According to the analysis realized in previous sections, it is possible to define a heuristic policy based on those characteristics that were common for all the scenarios.

Therefore, the heuristic policy is summarized as:

Service	*action*
Voice	scheme *ii*
Data	• if no channel sharing on TDMA: scheme *i* • if channel sharing on TDMA: scheme *iii*

The heuristic policy sends voice sessions to WCDMA which is a very simple solution that is consistent with optimal solutions of Section 5.2 and for the lower values of λ_v in Section 5.1. On the other hand, the optimal solution for data sessions is more complex and composed by two stages as previous analysis showed. On the first stage, data sessions are sent to TDMA (scheme 1) until the arrival of a new session would force channel sharing. On the second stage decisions depend on voice and data loads, as can be seen from Sections 5.1 and 5.2. When λ_v grows, data sessions are sent to TDMA, making the optimal policy for data sessions as scheme 1. On the other hand, when λ_d is the one that grows, more data sessions are sent to WCDMA. Therefore the heuristic policy uses scheme 1 for the first stage, and scheme 3 for the second stage, that is, compares the occupancy on each technology and sends sessions to the one with the lowest value.

Figure 12 compares the blocking function of the heuristic policy and Policy #2 for the same values of Fig.5. It is evident that for the lowest values of λ_v, the

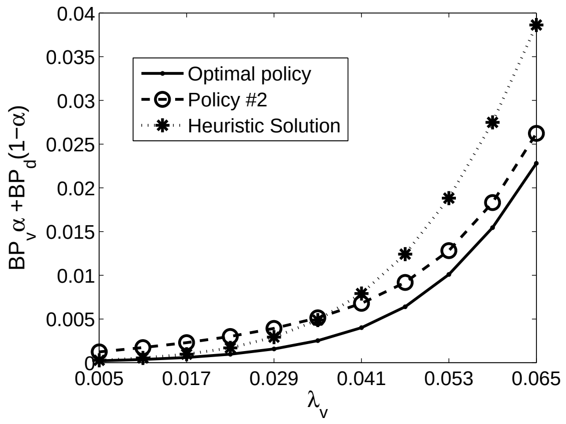

Figure 12: Blocking function for various λ_v.

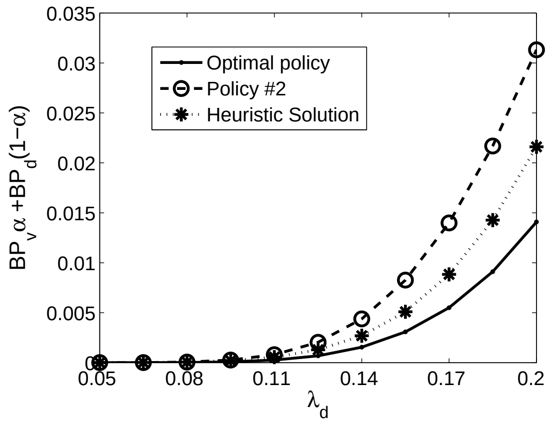

Figure 14: Blocking function for various λ_d.

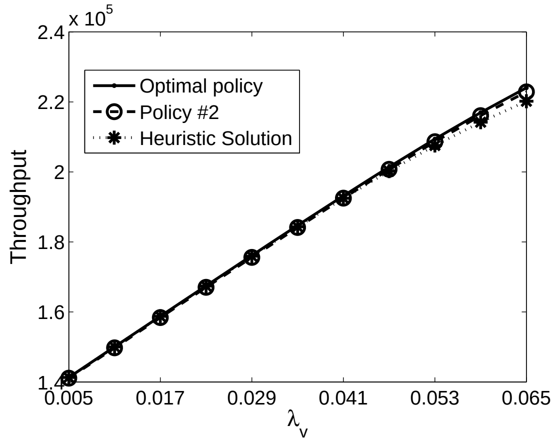

Figure 13: Throughput for various λ_v.

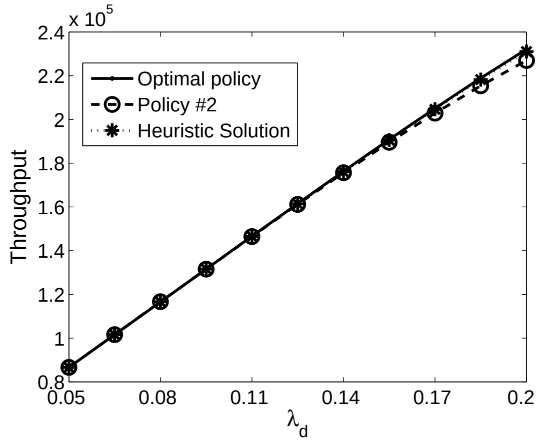

Figure 15: Throughput for various λ_d.

heuristic policy is better than Policy #2, but this changes when $\lambda_v > 0.035$. This occurs because Policy #2 always sends sessions to TDMA, and this behavior is similar to what is done by the optimal solution as λ_v grows (Section 5.1). The same could be said for Fig.13, where the throughput is the optimization function. In this case the improvement of the heuristic solution over Policy #2 is very small, but it is maintained while $\lambda_v < 0.041$. When this value is reached, Policy #2 is better than the heuristic solution because of its similarity to the optimal solution.

Figure 14 compares the heuristic solution with Policy #2 and the optimal solution, when the blocking function is optimized and λ_d changes. In this case, for the lowest values of λ_d Policy #2 is better than the heuristic, but once λ_d>0.08, the heuristic is better. This can be explained by the fact that the optimal solution sends more data sessions to WCDMA as λ_d grows, and while Policy #2 always sends data sessions to TDMA, the heuristic policy may send some sessions to WCDMA according to occupation on each technology.This behavior has the same effect over the throughput as can be seen on Fig.15, where it is compared the throughput for the heuristic solution and Policy #2. For the lowest values of λ_d, Policy #2 is better than the heuristic, but when $\lambda_d > 0.65$ the heuristic solution has a higher throughput. Therefore the heuristic policy represents some advantages over the best fixed policy when data load is high, but its behavior is not as good as that of Policy #2 when voice load is high.

6. Throughput optimization with QoS constraints

In this section we compare the policies in terms of the maximum arrival they can support subject to QoS

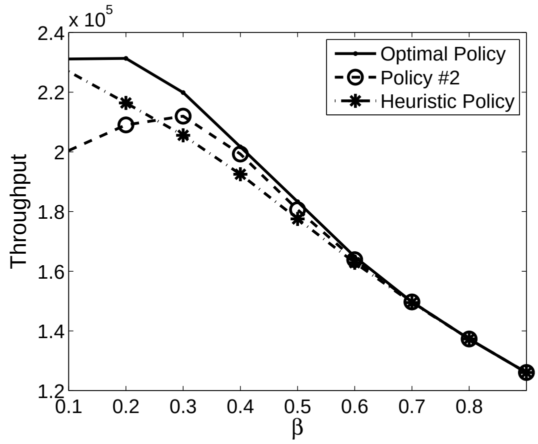

Figure 16: Throughput as a function of β.

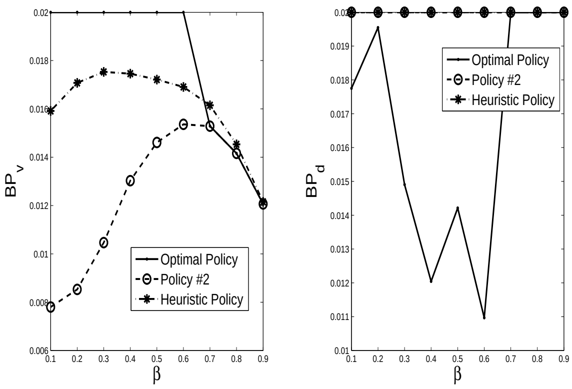

(a) *VoiceBlocking Probability.* (b) *DataBlocking Probability.*

Figure 17: Blocking probabilities for various β.

constraints. Until now, we considered throughput optimization independent of other factors, but to provide QoS, it is necessary to limit both PB_v and PB_d. To do this, we defined the parameter β, which is the fraction of λ_v over the total arrival rate λ_T, that is, λ_v=$\lambda_T \cdot \beta$ and λ_d=$\lambda_T \cdot (1-\beta)$. In Fig.16 it is shown the throughput as a function of β for three policies: The optimal policy when we optimize the throughput, Policy #2, and the heuristic policy defined in the last section. It has to be considered that even when they share the same value of β, the arrival rates differ for each policy according to the restrictions imposed by the blocking probabilities (max 2 %). For example, when β=0.2, the total arrival rate λ_T that the optimal policy can handle is 0.2158, while it is of 0.2018 for the heuristic policy and 0.1945 for Policy #2. Therefore, the optimal policy is able to receive a higher arrival rate for voice and data sessions while

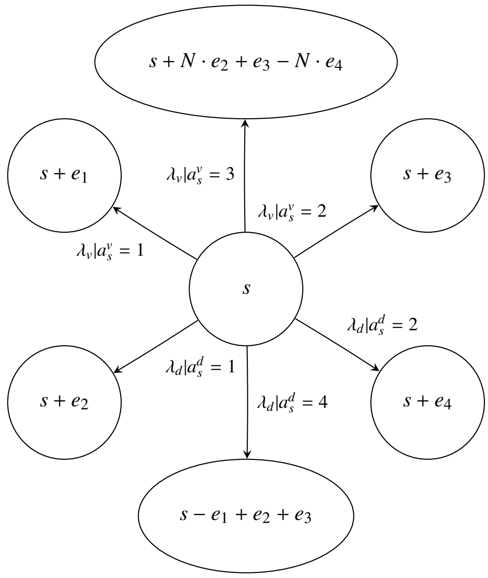

Figure 18: Transitions for voice and data arrivals using VH.

mantaining PB_v and PB_d below 2 %. Figures 17(a) and 17(b) show the voice and data blocking probabilities for different values of β. It can be seen that the optimal policy is the only one able to change its active constraint between PB_v and PB_d, taking advantage of the contribution of data sessions on the throughput. This cannot be done by the other two policies, which penalize PB_d=0.2 for every β. As for the heuristic policy, it shows a higher maximum throughput for low values of β than Policy #2, and this behaviour is reversed as β grows when voice sessions are more important in order to maximize throughput, a behaviour consistent with the results of the previous section.

7. Optimization with vertical handoffs

So far, we have analyzed a system where no vertical handoffs are used (now referred to as NVH), i.e. a session is served in the same technology where it was accepted until it terminates. However, vertical handoff is considered an important feature in CRRM that improves the use of resources. In Table 4 we introduce four types of vertical handoffs based on the analysis done in previous sections, where N is the necessary number of data sessions such that one voice session can be accepted in WCDMA.

All types of vertical handoffs are based on the results obtained in Sections 4 and 5, but the first two types

Type	*Triggering Event*	*Conditions*	*Number / Session class*	*From ⇒ To*	*NVH*	*VH-A*	*VH-B*
I	voice **arrival**	Full occupation on WCDMA	*N* **data** sessions	WCDMA⇒TDMA	–	✓	✓
II	data **arrival**	Channel sharing on TDMA	One **voice** session	TDMA⇒WCDMA	–	✓	✓
III	voice or data **departure** from WCDMA	Channel sharing on TDMA	One **voice** session	TDMA⇒WCDMA	–	–	✓
IV	voice or data **departure** from TDMA	VH does not produce channel sharing on TDMA	One **data** session	WCDMA⇒TDMA	–	–	✓

Table 4: Vertical Handoff Types

are triggered by the arrival of a session, while the other two occur when a session terminates. This distinction is important, since it affects the construction of the new Markov decision processes. The addition of types I and II of vertical handoff to NVH defines VH-A. The Markov model for this system has the same state space S of NVH, since restrictions (1) and (2) are not affected by vertical handoffs. The same applies for the Markov model of system VH-B, which uses all types of vertical handoffs (I-IV). However, the set of actions A and the transitions associated to them do change. For VH-A, a decision is made each time a session arrives as can be seen in Table 5 and in Figure 18, where actions 0, 1 and 2 do not include a vertical handoff, while actions 3 and 4 correspond to vertical handoff types I and II respectively. For action 3, the value N is the number of data sessions that need to be moved so that a new voice session can access WCDMA according to (2). The same applies for system VH-B, which includes vertical handoff types III and IV. The reason why handoff types III and IV do not affect the set A is that decisions are made only when a session terminates. In our model, handoffs triggered by service completion are done every time their specific conditions are fulfilled, so there are no decisions involved.

With this in mind, and using the same cost functions of (4) and (5), two new MDPs are constructed. These new MDPs will show the improvement when vertical handoffs are included, where the MDP based on VH-A uses handoffs when sessions arrive and the MDP based on VH-B uses handoffs for both arrivals and departures.

7.1. Optimal policy using vertical handoffs

In the last section were introduced two new systems that included vertical handoffs: VH-A and VH-B. The difference between them is that VH-B includes vertical handoffs for users departures. I n this section we define two new MDPs that are based on VH-A (MDP VH-A) and VH-B (MDP VH-B), and compare the performance of their optimal policies with the performance of the policies obtained with the MDP for the NVH system (MDP NVH). It is our interest to study the impact that vertical handoffs have on the optimal policies and their performance. Using the scenario specified in Table 3 for MDP VH-A and MDP VH-B, it is possible characterize the optimal policies as different parameters vary. Table 6 summarizes the main characteristics of the policies that optimize the blocking function defined in (4), when λ_v varies from 0.005 to 0.095. In general, the behavior for voice and data sessions are very similar in both MDPs. Nevertheless, some minor differences change the structure of the optimal solutions, and that has an important impact on the results obtained. Optimal solutions for MDP VH-A show a lower percentage of unused states than those solutions obtained using MDP VH-B. This is because Vertical Handoff types III and IV organize sessions in such a way that a larger number of states from set S are left unused. When λ_v=0.005 the percentage of unused states is three times higher for the solution obtained with MDP VH-B than for the one obtained with MDP VH-A, and while this difference shortens when λ_v= 0.095, the percentage of unused states grows. This growth shows that an organized use of resources is necessary in order to obtain the optimal solution in

Table 5: set of actions A for MDPs with vertical handoff

value	*action*
0	Block session
1	Send session to TDMA
2	Send session to WCDMA
3	VH for N data sessions from WCDMA to TDMA and the voice session is sent to WCDMA.
4	VH for 1 voice session from TDMA to WCDMA and the data session is sent to TDMA.

both MDPs. Also, it may occur some voice blocking when there are resources available, but this behavior is expected only for the highest values of λ_v, when its influence is higher on the optimization function.

Similar results are obtained for optimization of the blocking function defined in (4) when λ_d varies from 0.05 to 0.225. Voice and data sessions use Vertical Handoff types I and II respectively, but the percentage of unused states follows a different pattern. As λ_d grows, the percentage of unused states for the optimal solutions obtained with MDP VH-A remains constant, and it decreases for the solutions of MDP VH-B. This decrease shows that Vertical Handoff types III and IV allow more flexibility on how resources are managed, because a higher data arrival rate demands that more data sessions share channels in order to reduce the blocking probability, something that cannot be done as successfully by only using Vertical Handoff types I and II. Also, for the highest values of λ_d, some states decide to block data sessions even when there is capacity left, in order to save space for voice sessions. However, the percentage of states doing this is below 3% for both MDPs. A summary of these results is shown in Table 7.

Using the scenario specified in Table 3 and the throughput optimization function defined in (5) while λ_v varies from 0.005 to 0.095, the optimal policies for MDP VH-A and MDP VH-B are summarized in Table 8. Again, voice and data sessions use vertical handoff types I and II respectively in both MDPs. For the lowest values of λ_v, both MDPs decide to block voice sessions (<1%) in order to enhance throughput, since data sessions contribute more to the total throughput. However, this situation changes as λ_v grows when voice sessions are accepted as long as it is possible. On the other hand, the percentage of unused states differs for each MDP. For MDP VH-A, the percentage of unused states grows with λ_v from 16% to 41.7%. On the other hand, for MDP VH-B, the percentage of unused states diminishes from 61.2% to 49.5% as λ_v grows.

The optimal policies for MDP VH-A and MDP VH-B as λ_d varies from 0.05 to 0.225 are summarized in Table 9.Voice and data sessions use vertical handoff types I and II respectively in both MDPs. For the highest values of λ_d MDPs decide to block some voice sessions, even though in a very small percentage (<1%). This is done in order to accept more data sessions since they contribute more to the total throughput. For both MDPs the percentage of unused states grows with λ_d, and it is always higher for MDP VH-B.

7.2. Result analysis for vertical handoff MDPs

In this section we propose two new heuristic policies, which exploit the characterization of the optimal policies obtained from MDP VH-A and MDP VH-B that was described in the previous section. Heuristic policy VH-A makes use of vertical handoff types I and II, as it was seen for the optimal policies found using MDP VH-A. On the other hand, heuristic policy VH-B also includes vertical handoff types III and IV, as it was done by the optimal policies obtained using MDP VH-B. Both heuristic policies are defined in Table 10.

Our objective is to compare the performance obtained by the two new heuristic policies with the one obtained by the optimal policies of the MDPs. The evaluation scenario is defined by Table 11. The comparative study also includes the three heuristic policies previously proposed in the literature, that were defined earlier in Section 3. In this new scenario, the maximum capacity for voice and data sessions in WCDMA are 71 and 28 respectively, so the capacity has grown over 6 times over the system described in Table 3. However, the computational cost does not grow linearly. Using a desktop personal computer, it took around 1 minute to find the optimal policy of a single load point for the small system, while it took around 6 hours in this system. Therefore, the calculation time of the optimal policies makes them unfeasible for online use. It should be noted that this is not an issue for the heuristic schemes that can be used as a JCAC algorithm.

Figure 19 shows the blocking function for the three MDPs, the two heuristics and the three fixed policies, when λ_v varies from 0.09996 to 0.4998 and λ_d=0.448.

		MDP VH-A	MDP VH-B
Voice sessions	Lowest λ_v	• Sent to WCDMA or VH type I is used. • 10.4% of states are not used.	• Sent to WCDMA or VH type I is used. • 32.1% of states are not used.
	Highest λ_v	• Sent to WCDMA or VH type I is used. • 41.7% of states are not used. • Blocked in 1.00% of usable states.	• Sent to WCDMA or VH type I is used. • 50.4% of states are not used. • Blocked in 1.2% of usable states.
Data sessions	Lowest λ_v	• Sent to TDMA while no sharing is needed. If so, VH type II is used or sent to WCDMA. • 10.4% of states are not used.	• Sent to TDMA while no sharing is needed. If so, VH type II is used or sent to WCDMA. • 32.1% of states are not used.
	Highest λ_v	• Sent to TDMA while no sharing is needed. If so, VH type II is used or sent to WCDMA. • 41.7% of states are not used.	• Sent to TDMA while no sharing is needed. If so, VH type II is used or sent to WCDMA. • 50.4% of states are not used.

Table 6: Main characteristics of the optimal solutions for the blocking function for various λ_v.

		MDP VH-A	MDP VH-B
Voice sessions	Lowest λ_d	• Sent to WCDMA or VH type I is used. • 41.6% of states are not used.	• Sent to WCDMA or VH type I is used. • 49.4% of states are not used.
	Highest λ_d	• Sent to WCDMA or VH type I is used. • 41.6% of states are not used.	• Sent to WCDMA or VH type I is used. • 39.6% of states are not used.
Data sessions	Lowest λ_d	• Sent to TDMA while no sharing is needed. If so, VH type II is used or sent to WCDMA. • 41.6% of states are not used.	• Sent to TDMA while no sharing is needed. If so, VH type II is used or sent to WCDMA. • 49.4% of states are not used.
	Highest λ_d	• Sent to TDMA while no sharing is needed. If so, VH type II is used or sent to WCDMA. • 41.6% of states are not used. • Blocked in 2.56% of usable states.	• Sent to TDMA while no sharing is needed. If so, VH type II is used or sent to WCDMA. • 39.6% of states are not used. • Blocked in 2.98% of usable states.

Table 7: Main characteristics of the optimal solutions for the blocking function for various λ_d.

The blocking function value for the heuristic policies when λ_v=0.4998 is less than half the value achieved by Policy #2. This improvement is larger when compared to the other fixed policies. Also, it should be noted that the heuristic policies may even improve over the optimal policy MDP NVH as λ_v grows. The improvement

		MDP VH-A	MDP VH-B
Voice sessions	Lowest λ_v	• Sent to WCDMA or VH type I is used. • 16% of states are not used. • Blocked in 0.71% of usable states.	• Sent to WCDMA or VH type I is used. • 61.2% of states are not used. • Blocked in 0.25% of usable states.
	Highest λ_v	• Sent to WCDMA or VH type I is used. • 41.7% of states are not used.	• Sent to WCDMA or VH type I is used. • 49.5% of states are not used.
Data sessions	Lowest λ_v	• Sent to TDMA while no sharing is needed. If so, VH type II is used or sent to WCDMA. • 16% of states are not used.	• Sent to TDMA while no sharing is needed. If so, VH type II is used or sent to WCDMA. • 61.2% of states are not used.
	Highest λ_v	• Sent to TDMA while no sharing is needed. If so, VH type II is used or sent to WCDMA. • 41.7% of states are not used.	• Sent to TDMA while no sharing is needed. If so, VH type II is used or sent to WCDMA. • 49.5% of states are not used.

Table 8: Main characteristics of the optimal solutions for throughput for various λ_v.

		MDP VH-A	MDP VH-B
Voice sessions	Lowest λ_d	• Sent to WCDMA or VH type I is used. • 41.6% of states are not used.	• Sent to WCDMA or VH type I is used. • 45.4% of states are not used.
	Highest λ_d	• Sent to WCDMA or VH type I is used. • 43% of states are not used. • Blocked in 0.7% of usable states.	• Sent to WCDMA or VH type I is used. • 47.2% of states are not used. • Blocked in 0.56% of usable states.
Data sessions	Lowest λ_d	• Sent to TDMA while no sharing is needed. If so, VH type II is used or sent to WCDMA. • 41.6% of states are not used.	• Sent to TDMA while no sharing is needed. If so, VH type II is used or sent to WCDMA. • 45.4% of states are not used.
	Highest λ_d	• Sent to TDMA while no sharing is needed. If so, VH type II is used or sent to WCDMA. • 43% of states are not used.	• Sent to TDMA while no sharing is needed. If so, VH type II is used or sent to WCDMA. • 47.2% of states are not used.

Table 9: Main characteristics of the optimal solutions for throughput for various λ_d.

of the blocking function has an effect in the aggregated throughput, which raises to 1.156 Mbps for heuristic VH-A while it only reaches 1.116 Mbps for Policy #2 when λ_v=0.4998. Therefore under these conditions, using an heuristic policy represents an improvement of 40 Kbps, close to one data channel, or 3 voice chan-

Heuristic policy	*Event*	*Action*
VH-A	Voice arrival	• Send to WCDMA. • If it is not possible, use VH type I. • If it is not possible, send to TDMA.
	Data arrival	• Send to TDMA if no channel sharing is needed. • If there is channel sharing, use VH type II. • If it is not possible, send to WCDMA. • If it is not possible, send to TDMA.
VH-B	Voice arrival	• Send to WCDMA. • If it is not possible, use VH type I. • If it is not possible, send to TDMA.
	Data arrival	• Send to TDMA if no channel sharing is needed. • If there is channel sharing, use VH type II. • If it is not possible, send to WCDMA. • If it is not possible, send to TDMA.
	Voice departure	• Use VH type III if departs from WCDMA • Use VH type IV if departs from TDMA
	Data departure	• Use VH type III if departs from WCDMA • Use VH type IV if departs from TDMA

Table 10: Definition for heuristic policies with vertical handoff.

Table 11: New scenario of study.

WCDMA	TDMA
W=3.84 Mcps	C = 8
$(E_b/N_0)_v$=6.5 dB	n_c= 3
$(E_b/N_0)_d$=5 dB	$BR_{t,v}$=12.2 Kbps
$BR_{w,v}$=12.2 Kbps	$BR_{t,d}$=44.8 Kbps
$BR_{w,d}$=44.8 Kbps	
η_{ul}= 1	
Clients	
μ_v= 0.0083	
σ= 1 Mb	

nels, with blocking probabilities of 2.2 % and 1.8 % for data and voice, while Policy #2 raises them to 6.5% and 4.6%, respectively.

Figure 20 shows the blocking function for the MDPs, the heuristics and the fixed policies when λ_d varies from 0.3584 to 1.792 and λ_v is 0.0833. The blocking function values are lower for the heuristic functions when compared to the fixed policies, but they are higher than those obtained by MDP NVH when λ_d>1.0752. However, when λ_d=1.792, the total throughput for heuristic VH-A is of 1.4896 Mbps and of 1.466 Mbps for MDP NVH. That is, the heuristic policies cannot diminish the

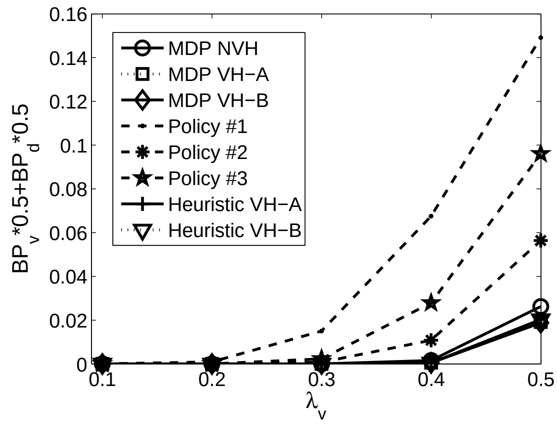

Figure 19: Blocking function for various λ_v.

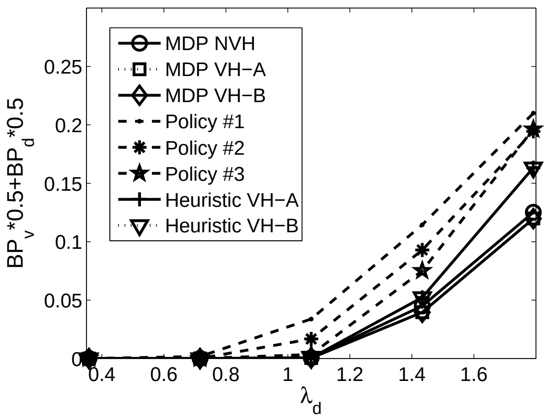

Figure 20: Blocking function for various λ_d.

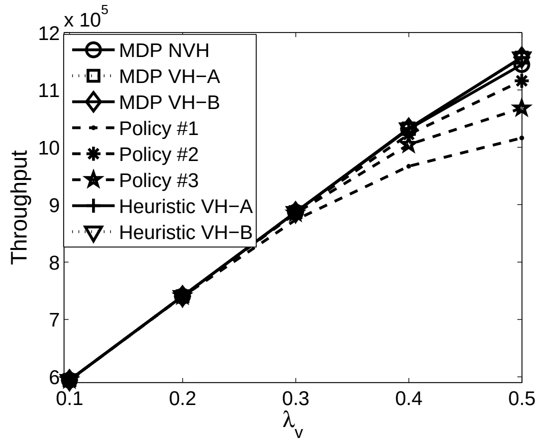

Figure 21: Throughput for various λ_v.

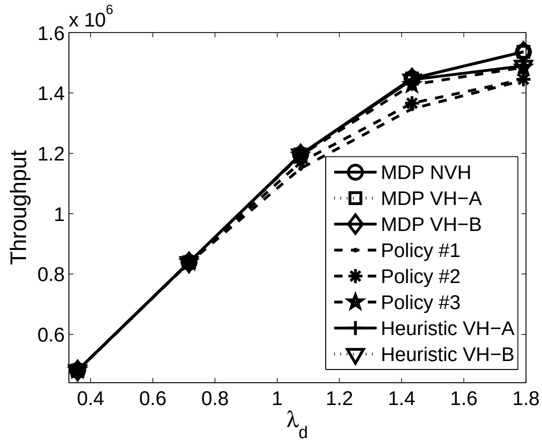

Figure 22: Throughput for various λ_d.

value of the blocking function for the higher values of λ_d as much as MDP NVH does, but they still raise the throughput, even though this is not the objective function. As in the case when λ_v varies, the differences for both heuristics are not very significant.

The optimal throughput for the MDPs, as well as the throughput obtained by the heuristic policies and the fixed policies as λ_v grows from 0.09996 to 0.4998, keeping λ_d=0.448, is shown in Fig. 21. There is an improvement of the heuristic policies throughput not only over the fixed policies, but also over the optimal policies of MDP NVH. The improvement in throughput is small when λ_v=0.4998, 12 Kbps over MDP NVH, and more when compared to the fixed policies. However, it is significant when we consider that the voice and data blocking probabilities of heuristic VH-A are of 2.2% and 1.8% respectively, while these probabilities are of 1.5% and 3.8% for MDP NVH, and raise to a range of 6.5–19.5% and 4.6–10.5% for the fixed policies. That is, the heuristic policies are able to improve throughput while maintaining low blocking probabilities (around 2%), and the other policies are not.

Figure 22 shows the throughput for the MDPs, the heuristics and the fixed policies when λ_d varies from 0.3584 to 1.792 and λ_v is 0.0833. As observed, there is not too much room for the improvement of the throughput when λ_d varies, since most policies obtain similar values. The only real difference is seen when λ_d=1.792, where the MDPs raise around 6 Kbps over the other policies. Since this is achieved by blocking all the voice sessions, this policies cannot be considered useful. It is interesting to see that that the heuristics still achieve a higher throughput than all fixed policies, while keeping lower voice and data blocking probabilities. However, when λ_d=1.4336, the voice and data blocking probabilities of the heuristic policies are of 5% and 2.9% for data and voice. Thus this load point could be considered a practical load limit.

In summary, we may conclude that the difference between the performance achieved by the heuristic and optimal policies is negligible, and this negligible difference is consistent for both optimization criteria and a wide range of system parameters values.

7.3. Model Validation

In order to validate the performance of the heuristic policies found, which were determined analytically, we evaluate them also by discrete event simulation. Three different distributions for the service time of voice sessions T_v and data size (σ) were used, namely, exponential, hyper-exponential and Erlang. We set the coeffi-

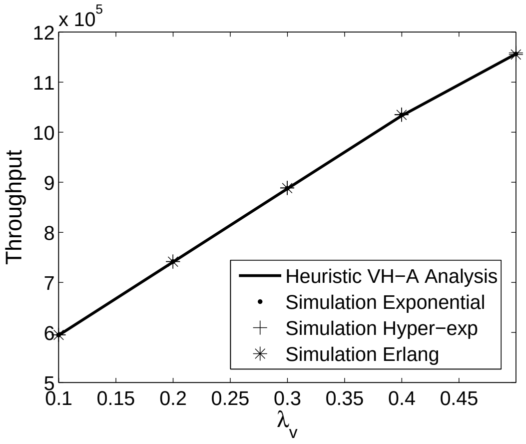

Figure 23: Simulation vs Analysis results of throughput as λ_v varies.

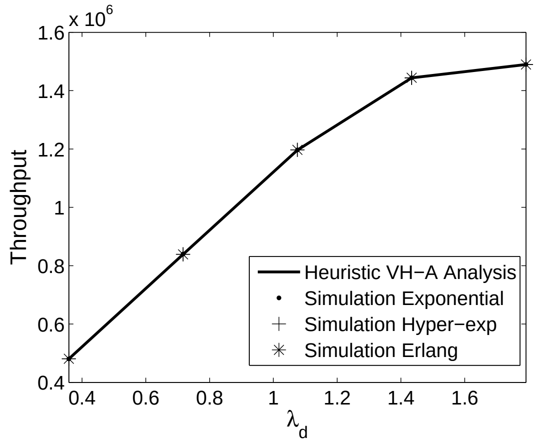

Figure 24: Simulation vs Analysis results of throughput as λ_d varies.

cients of variation of the last two to 2 and 0.5, respectively. The exponential distribution is known to have a coefficient of variation equal to 1. Clearly, the mean of all distributions coincides for each load point. Our purpose here is not only to establish the correctness of our mathematical analysis, but also to asses the impact of the exponentially distributed data size and voice service time assumptions.

The simulation results for different distributions of the data size using the heuristic policy VH-A are shown in Fig. 23 and Fig. 24. As observed, there is an excellent agreement between the analytical and simulation results for all distributions. The same applies for the simulation results of the heuristic policy VH-B, shown in Fig. 25 and Fig. 26. An interesting and very important finding is that the performance of the heuristic policies is insensitive to distribution of the data size beyond its mean. Although not shown, an identical conclusion is obtained with respect to the insensitivity of the system performance with respect to the distribution of the voice service time.

8. Cost of the Vertical Handoff

In the previous section we introduced four types of vertical handoffs and used them to find new optimal policies. We also designed two new heuristic policies and showed that they outperform the optimal policies obtained for a system without vertical handoffs. The main issue in this section is to explore the impact that session blocking and vertical handoff costs have on the optimal policies.

In order to do this, we define the objective function

$$F_{VH} = \theta \cdot \zeta_{VB} + (1 - \theta) \cdot \zeta_{DB} + C_{VH} \cdot \zeta_{VH}, \quad (9)$$

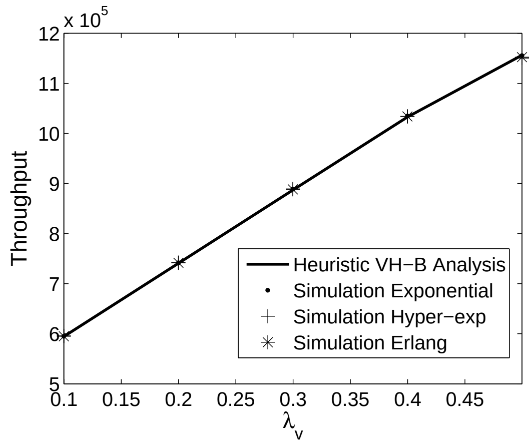

Figure 25: Simulation vs Analysis results of throughput as λ_v varies.

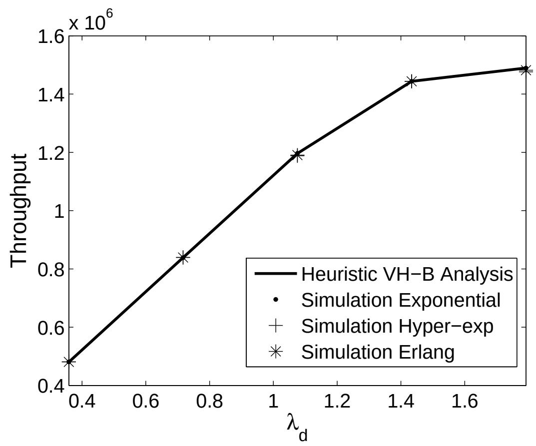

Figure 26: Simulation vs Analysis results of throughput as λ_d varies.

where θ, $0 \leq \theta \leq 1$, is the factor that defines the cost of blocking a single voice or data session, and ζ_{VB} and ζ_{DB} are the mean voice and data blocking rates. In the same way, C_{VH} is the cost of performing a single vertical handoff, and ζ_{VH} is the mean rate of vertical handoffs performed. Hence, by assigning values to θ and C_{VH}, a new optimal policy that minimizes F_{VH} can be found by solving the appropriate MDP.

8.1. Markov decision process for VH cost

In this section we define three new MDPs based on (9). The first one does not use vertical handoffs (MDP BR), the second one uses vertical handoff types I and II (MDP C1), and the last one uses all four types of vertical handoffs (MDP C2). These MDPs are different from those defined in Section 7.1, as their objective functions are different. The state space for all of them S is defined by (1) and (2). The set of actions a is defined in Table 1 for the MDP BR and in Table 5 for the MDP C1 and MDP C2. The cost function associated to the objective function for each feasible state $\boldsymbol{s}$ for the MDP BR is

$$\text{cost}(\boldsymbol{s}) = \lambda_v \cdot G(a_s^v) \cdot \theta + \lambda_d \cdot G(a_s^d) \cdot (1-\theta), \qquad (10)$$

for the MDP C1 is

$$\begin{aligned}\text{cost}(\boldsymbol{s}) &= \lambda_v \cdot G(a_s^v) \cdot \theta + \lambda_d \cdot G(a_s^d) \cdot (1-\theta)\\ &\quad + C_{VH} \cdot (\lambda_v \cdot R(a_s^v) + \lambda_d \cdot R(a_s^d)), \qquad (11)\end{aligned}$$

and for the MDP C2 is

$$\begin{aligned}\text{cost}(\boldsymbol{s}) &= \lambda_v \cdot G(a_s^v) \cdot \theta + \lambda_d \cdot G(a_s^d) \cdot (1-\theta)\\ &\quad + C_{VH} \cdot \left(\lambda_v \cdot R(a_s^v) + \lambda_d \cdot R(a_s^d)\right)\\ &\quad + C_{VH} \cdot \left(T(\boldsymbol{s}_{TDMA}^v) + T(\boldsymbol{s}_{TDMA}^d)\right.\\ &\quad \left. + T(\boldsymbol{s}_{WCDMA}^v) + T(\boldsymbol{s}_{WCDMA}^d)\right), \qquad (12)\end{aligned}$$

where the coefficients are explained in Table 12.

8.2. Result analysis

As before, policy iteration is used to solve the MDPs. The reference scenario is defined in Table 3. We set θ= 0.5, i.e. the cost of blocking voice and data sessions is the same.

Figure 27 shows the optimal cost for each of the different MDPs studied as C_{VH} varies. MDP C2 has the lowest optimal cost when C_{VH}=0, and is followed closely by MDP C1. However, as C_{VH} grows, the cost of MDP C2 grows rapidly, surpassing MDP C1 and MDP BR. This is explained by the large amount of vertical handoffs of types III and IV performed by MDP C2 policies. In fact, for values of C_{VH} as low as 0.01, the optimal policy obtained by MDP C2 is more costly than the one obtained by MDP BR. As the cost of blocking voice and data sessions is 0.5, we could say that it would make sense to use the policies obtained by MDP C2 only when the cost of a vertical handoff is 50 times lower than the cost of blocking a voice or data sessions.

On the other hand, the costs of MDP C1 policies vary slowly and are bounded by those of MDP BR. This occurs because if the objective function growth is caused by a vertical handoff, then the MDP C1 policy can always choose not to use it, which makes the MDP C1 policies tend to the ones obtained by MDP BR. The interesting point here is to find for which C_{VH} value both MDPs reach the same optimal policy, and therefore the same costs. This value is C_{VH}=0.9. Then, as before, we could say that it would make sense to use the policies obtained by MDP C1 only when the cost of a vertical handoff is not bigger that 1.8 times the cost of blocking a voice or data sessions.

It is interesting to notice that when C_{VH}=0, the optimal policies of MDP C1 and MDP C2 are very similar to those of MDP VH-A and MDP VH-B, even though the objective functions are different. Therefore, performance parameters, such as throughput and blocking probabilities (BP_v and BP_d), are similar as well. As C_{VH} grows, the performance degrades, and this happens at a faster rate for the MDP C2 policies than for thew MDP C1 ones. Hence, while the ratio of voice/data blocking cost to C_{VH} is high, it is expected that the heuristic policies VH-A and VH-B fairly represent the optimal policies. This last remark is true even for higher ratio values in the case of MDP C1 and the heuristic VH-A, for the reasons explained earlier.

9. Conclusions

We have studied optimal policies for the selection of collocated wireless networks with heterogeneous access techniques (WCDMA and TDMA) and offered services (voice and data). Two different optimization criteria were used, one based on the blocking probabilities of each service, and the other on the total throughput. We formulated the optimization problem using the formalism of Markov decision processes and used policy iteration to solve it.

Optimal policies have been found for various traffic scenarios using both optimization functions, and their performance was compared to the ones obtained by

SYMBOL	DEFINITION	VALUE
$G(a_s^x)$	Indicates if $\boldsymbol{s}$ is a blocking state for service x	• If action a_s^x is blocking, $G(a_s^x)$=1. • Otherwise $G(a_s^x)$=0.
$R(a_s^x)$	Indicates the number of sessions that suffer vertical handoff when a session of service x arrives while the system is on state $\boldsymbol{s}$.	• If a_s^x=3, and all the conditions for vertical handoff type I are fulfilled, $R(a_s^v) = N$. • If a_s^x=4, and all the conditions for vertical handoff type II are fulfilled, $R(a_s^d) = 1$. • Otherwise, $R(a_s^x) = 0$.
$T(\boldsymbol{s}_{TDMA}^v)$	Indicates the rate of sessions that suffer vertical handoff when a voice session is served on TDMA while the system is on state $\boldsymbol{s}$.	• If the conditions for vertical handoff type IV are fulfilled once a voice session is served on TDMA, $T(\boldsymbol{s}_{TDMA}^v) = s_1 \cdot \mu_v$. • Otherwise $T(\boldsymbol{s}_{TDMA}^v) = 0$.
$T(\boldsymbol{s}_{TDMA}^d)$	Indicates the rate of sessions that suffer vertical handoff when a data session is served on TDMA while the system is on state $\boldsymbol{s}$.	• If the conditions for vertical handoff type IV are fulfilled once a data session is served on TDMA, $T(\boldsymbol{s}_{TDMA}^d) = \min(C\text{-}s_1, s_2) \cdot BR_{t,d}/\sigma$. • Otherwise $T(\boldsymbol{s}_{TDMA}^d) = 0$.
$T(\boldsymbol{s}_{WCDMA}^v)$	Indicates the rate of sessions that suffer vertical handoff when a voice session is served on WCDMA while the system is on state $\boldsymbol{s}$.	• If the conditions for vertical handoff type III are fulfilled once a voice session is served on WCDMA, $T(\boldsymbol{s}_{WCDMA}^v) = s_3 \cdot \mu_v$. • Otherwise $T(\boldsymbol{s}_{WCDMA}^v) = 0$.
$T(\boldsymbol{s}_{WCDMA}^d)$	Indicates the rate of sessions that suffer vertical handoff when a data session is served on WCDMA while the system is on state $\boldsymbol{s}$.	• If the conditions for vertical handoff type III are fulfilled once a data session is served on WCDMA, $T(\boldsymbol{s}_{WCDMA}^d) = s_4 \cdot BR_{w,d}/\sigma$. • Otherwise $T(\boldsymbol{s}_{WCDMA}^d) = 0$.

Table 12: Coefficients for vertical handoff cost optimization functions.

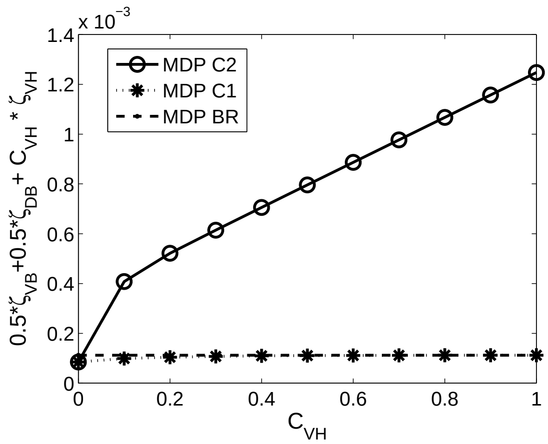

Figure 27: Optimization values for various C_{VH}.

three fixed policies. We were able to characterize the optimal policies after an exhaustive analysis of the their behavior in different scenarios and with different traffic profiles. Based on this characterization, heuristic policies were proposed and their performance was analyzed. We showed that they outperform the three fixed policies, particularly when the data arrival rate is larger than voice arrival rate. We also studied the optimization of the throughput considering QoS restrictions based on bounds on the blocking probabilities. The results showed that the advantages of the heuristic policies over the others are not affected by adding these constraints.

In order to improve further the resource occupancy, we introduced different types of vertical handoffs based on the knowledge gained about the behavior of the system. We determined and characterize new optimal policies according to the arrival type, system state and vertical handoff action. Since it is not computationally feasible to calculate the optimal policies online, new heuristic policies with vertical handoffs were designed and evaluated. The evaluation of the new heuristic policies was done in a system larger than the one used to characterize the optimal policies from which the heuristic ones derive. However, we found that their performance scale very well with the cardinality of the state space, which is very close to the performance achieved by optimal policies.

We also analyzed the performance of these heuristic policies by simulation. We found that the simulation results practically overlap the analytical ones, which allowed us to validate the last ones. In addition to exponential distributions, we also evaluated the performance of the heuristic policies by simulation with the hyper-exponential and Erlang distributions. An interesting and very important finding is that their performance is insensitive to distributions of the service time of the voice sessions and the elastic flow length (in bits) beyond the mean.

Finally, the cost of vertical handoff was studied in relation to those of voice and data blocking, and some interesting remarks were done in order to understand the impact that vertical handoff has.

Acknowledgements

This work was supported by the Spanish Government and the European Commission through projects TIN2010-21378-C02-02 and TIN2008-06739-C04-02.

REFERENCES

[1] Leijia Wu, Kumbesan Sandrasegaran,"A survey on Common Radio Resource Management", The 2nd conference on Wireless Broadband and U-Wideband Communications, AusWireless, August, 2007.

[2] A. Tolli, P. Hakalin, H. Holma, "Performance Evaluation of Common Radio Resource Management", IEEE International Conference on Communications (ICC), pp. 3429–3433, Vol 5, April, 2002.

[3] Oklabisi E. Falowo, H. Anthony Chan,"Joint Call Admission Control Algorithms: Requirements, Approaches, and Design Considerations", Computer Communications, pp. 1200–1217, Vol 31, 2008.

[4] A. Pillekeit, F. Derakhshan, E. Jugl, A. Mitschele-Thiel, "Force-based Load Balancing in Co-located UMTS/GSM Networks", IEEE 60th VTC 2004-Fall, pp. 4402–4406, 2004.

[5] K. Suleiman, A. Chan and M. Dlodlo, "Load Balancing in the call admission control of heterogeneous wireless networks", International Wireless Communications and Mobile Computing Conference, pp. 245–250, July, 2006.

[6] Fei Yu, Vikram Krishnamurthy, "Optimal Joint Session Admission Control in Integrated WLAN and CDMA Cellular Networks with Vertical Handoff", IEEE Transactions on Mobile Computing, pp. 126–139, Vol 6, 2007.

[7] D. Karabudak, C. Hung, B. Bing, "A call admission control scheme using genetic algorithms", ACM Symposium on Applied Computing, pp. 1151–1158, 2004.

[8] R. Agusti, O. Sallent, J. Perez-Romero, L. Giupponi, "A fuzzy-neural based approach for joint radio resource management in a beyond 3G framework", 1st International Conference on Quality of Service in heterogeneous in Wire/Wireless networks, pp. 216–224, 2004.

[9] O. Falowo, H. Chan, "Fuzzy Logic Based Call Admission Control for Next Generation Wireless Networks", 3rd International Symposium on Wireless Communication Systems, pp. 574–578, 2006.

[10] Xavier Gelabert and Jordi Perez-Romero and Oriol Sallent and Ramon Agusti, "A Markovian Approach to Radio Access Technology Selection in Heterogeneous Multiaccess/Multiservice Wireless Networks", IEEE Transactions on Mobile Computing, number 10, Vol 7, 2008.

[11] Enrique Stevens-Navarro E., Vincent Wong, "Comparison between Vertical Handoff Decision Algorithms for Heterogeneous Wireless Networks", IEEE Vehicular Technology Conference, VTC 2006-Spring, pp. 947–951, 2006.

[12] Pramod Goyal, S.K. Saxena , "A Dynamic Decision Model for Vertical Handoffs across Heterogeneous Wireless Networks". Proceedings of World Academy of Science, Engineering and Technology, Vol.31, July, 2008.

[13] Ahmed Hasswa, Nidal Nasser, Hossam Hassanein, "Generic Vertical Handoff Decision Function for Heterogeneous Wireless Networks," 2nd IFIP International Conference on Wireless and Optical Communications Networks, pp.239–243, 2005.

[14] Klein S. Gilhousen, Irwin M. Jacobs, Roberto Padovani, Andrew J. Viterbi, Lindsay A. Weaver, Charles E. Wheatley, "On the Capacity of a Cellular CDMA System", IEEE Transactions on Vehicular Technology, number 2, Vol 40, 1991.

[15] I. Koffman, V. Roman, "Broadband wireless access solutions based on OFDM access in IEEE 802.16," IEEE Communications Magazine, vol.40, no.4, pp.96–103, Apr. 2002.

[16] Y. Ben-Shimol, I. Kitroser and Y. Dinitz, "Two-Dimensional Mapping for Wireless OFDMA Systems," IEEE Transactions on Broadcasting, vol. 52, no. 3, pp. 388–396, September 2006.

[17] V. Pla, J. Martinez-Bauset, V. Casares-Giner, "Comments on call blocking probability and bandwidth utilization of OFDM subcarrier allocation in next-generation wireless networks," IEEE Communications Letters, vol.12, no.5, pp.349, May 2008.

[18] Ch. Paik, Y.S. Suh, "Generalized Queueing Model for Call Blocking Probability and Resource Utilization in OFDM Wireless Networks," IEEE Communications Letters, vol.15, no.7, pp.767–769, July 2011.

[19] Bertsekas,D.P., "Dynamic Programming and Optimal Control", Athena Scientific, 2001.

[20] Van Leeuwaarden, J. and Aalto, S. and Virtamo, J., "Load balancing in cellular networks using first policy iteration", 5th Finnish Teletraffic Seminar, Citeseer, 2002.