ANALYTICAL ASSESSMENT OF THE PERFORMANCE IMPACT OF SPECTRUM SHARING IN CELLULAR NETWORKS CONSIDERING DIFFERENT TIME SCALES

Remco Litjens and Haibin Zhang

TNO
P.O. Box 5050, 2600 GB Delft, The Netherlands
phone: +31(0)888 667184, fax: +31(0)888 667349, email: remco.litjens@tno.nl
web: www.tno.nl

ABSTRACT

We present an analytical approach to assess the performance gains from spectrum sharing in a wireless communications network with two co-sited cells, with specific focus on the performance impact of the timescale at which spectrum is shared. Primary focus is on the analysis of spectrum sharing mechanisms that operate on a timescale of minutes or higher, based on average load information, while a comparison is made with an idealistic reference scheme of full and continuous spectrum sharing. The numerical results indicate the strong superiority of spectrum sharing at a timescale of seconds, with lower (but still appreciated) performance gains observed for spectrum sharing mechanisms operating at a minute (or higher) timescale.

1. INTRODUCTION

Future wireless network deployment and operation need to be energy- and cost-efficient, while ensuring competition among operators. To achieve these goals, an important way is spectrum and/or infrastructure sharing among operators [1], which among others enables quick launch of new technology networks, eases the site constraints, mitigate the problem of spectrum scarcity, and reduce CAPEX (CAPital EXpenditure) and OPEX (OPerational EXpenditure).

Spectrum sharing may be classified in the following ways: (1) spectrum sharing for *unlicensed bands* [2] or *licensed band* [3][4][5], according to the type of spectrum shared. (2) *Equal sharing* [5] or *unequal sharing* [4][6], depending on the role of participants sharing the spectrum. Here "equal" indicates that each participant have either the same priority in using the shared spectrum, or priority according to e.g. the portion of spectrum it brings in for sharing. On the other hand, "unequal" means that there are primary and secondary users of the shared spectrum, where the secondary users only have access to the part of shared spectrum not used by the primary users (e.g. in cognitive radio-based spectrum sharing [6]). (3) *Interference cancellation-based* spectrum sharing, where users of multiple operators use the same frequency simultaneously and thus interfere with each other (e.g. in [7]), or *interference avoidance-based* spectrum sharing, where users of different operators are scheduled to not simultaneously use the same frequency (or other resources such as codes in [4]) at small timescale (e.g. milliseconds, seconds or minutes). Note that, in the latter case spectrum is still shared among operators at relatively large timescales (e.g. hours, days). In this paper, we focus on interference avoidance-based sharing of licensed spectrum bands among mobile operators, where the users of each operator have access to the shared spectrum with equal priority.

Kumar S. *et al.* [3] studied the challenges at physical layer and radio resource management faced by spectrum sharing among operators where services of different operators do not interfere with each other. They also proposed centralized and distributed approaches at high level, taking into account traffic profiles, QoS requirements, timescale and sequence of operation, and information availability, etc. Salami G. et al. [4] proposed dynamic spectrum allocation algorithms for spectrum sharing (in the form of code sharing) between two Universal Mobile Telecommunication System (UMTS) operators, with the assumption that each operator has priority in access to its own spectrum shared (not *all* the shared spectrum). The FP7 SAPHYRE project [5] works on equal-priority spectrum and infrastructure sharing among operators, where operators voluntarily bring (some or all of) their spectrum and/or infrastructure for sharing in a cooperative method. The objectives of SAPHYRE include among others the development of holistic solutions for efficient spectrum/infrastructure sharing, covering cross-layer (physical, medium access control and network layers) technical aspects, and business and regulatory respects. The work presented in this paper is done within the SAPHYRE project.

The principal contribution of this paper is to present a purely analytical assessment of the impact of different spectrum sharing schemes and the applied sharing timescale, on the experienced performance in a system with two cells/operators, characterised by time-varying traffic loads.

The outline of the paper is as follows. In Section 2 we outline the system model considered in the assessment of different spectrum sharing schemes. Subsequently, four distinct spectrum sharing schemes are defined in Section 3. Section 4 then outlines the stochastic analysis applied to evaluate the achieved performance of the different schemes. The scenarios and results associated with some numerical experiments are presented in Section 5 and Section 6, respectively. Section 7 ends this paper with some concluding remarks.

2. SYSTEM MODEL

We consider two co-sited omni-directional cells possibly corresponding to two distinct operators, as depicted in Figure 1. Although the applied analysis allows a more generic setting, each cell is assumed to operate the same radio access technology, e.g. LTE[1].

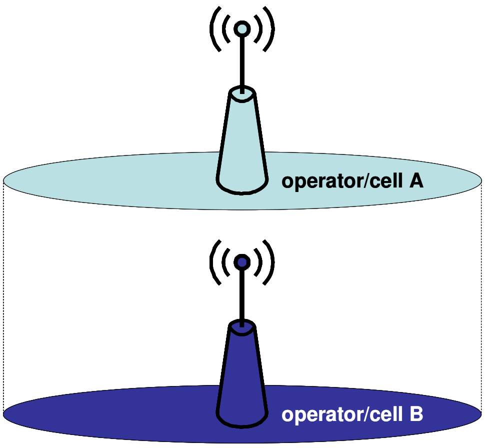

Figure 1: Considered system model.

In order to model the distance-dependent radio link quality and hence attainable data rates, each cell is segmented into K concentric, equal-area zones. See Figure 2 (K = 10). To ensure equal-area zones, the radius of zone k is chosen such that $(r_k / r_K)^2 = k / K$, with r_K the cell radius. Users located in zone k experience rate r_k (in Mbps) per assigned MHz and assuming exclusive use of the shared channel.

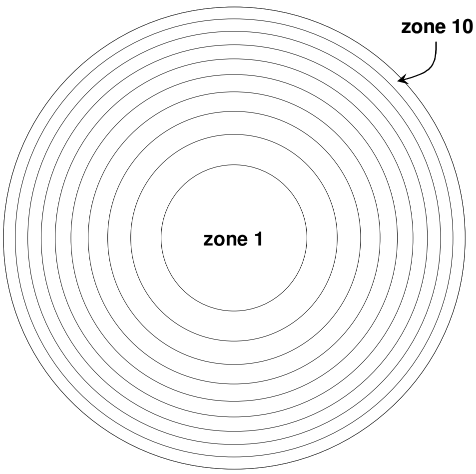

Figure 2: Cell segmentation in concentric zones (K = 10).

Users (flows) are generated according to a time-varying, spatially uniform Poisson arrival process with arrival rate $\lambda_j(t)$ in cell j at time t. A flow is modelled as a file download with a generally distributed size with mean denoted Ω (in Mb). At any time, the flows present in a cell fairly share the available resources. Assuming proportionality of data rates to assigned spectrum, a user in zone k experiences an effective data rate of $S\, r_k / n$, with S the assigned spectrum (in MHz) and n the total number of active flows in the cell.

[1] 'Long-Term Evolution', the unofficial name for the E-UTRAN (Evolved UMTS Terrestrial Radio Access Network) technology.

The total amount of available spectrum is denoted S (in MHz) which is to be (orthogonally) split into cell-specific spectrum assignments denoted S_A and S_B. In the non-sharing reference case, we assume a fair split, i.e. $S_A = S_B = ½S$. A number of spectrum sharing schemes will be studied that aim at periodically redetermining the cell-specific spectrum assignments to match the time-varying load asymmetry and thereby enhancing overall performance. The considered spectrum sharing schemes are outlined in Section 3.

3. SPECTRUM SHARING SCHEMES

The spectrum sharing scheme periodically assesses the load asymmetry between cells A and B, and optimises the split of S over S_A and S_B accordingly. Considering the paper's focus on the impact of the sharing timescale, the period length ϑ is a key parameter of the spectrum sharing scheme. In the quantitative assessment presented in Section 5, ϑ will be varied from 1 minute to 1440 minutes. Considering the assumption of a daily traffic profile, the latter effectively corresponds with a fixed and constant spectrum split.

A distinction is made between three spectrum sharing schemes, which are orthogonal in nature in the sense that at all times any spectrum slice is assigned either to cell A or to cell B. The considered spectrum sharing schemes are characterised as follows:

- **SSS-I** In each period S is split in relative proportion to the offered loads (in Mbps) per cell.
- **SSS-II** In each period S is split such that the achieved average throughput is optimised.
- **SSS-III** In each period S is split such that the achieved cell edge throughput is optimised.

Besides the reference case of a fixed spectrum split, which is effectively considered by choosing ϑ equal to 1440 for any of the schemes, another reference case is considered by considering full sharing at the finest timescale. In a practical implementation, this degree of spectrum sharing requires an algorithm which monitors, on a timescale of (milli)seconds, the actual number of active flows in each cell as well as their actual link qualities (including the effects of slow and fast fading), and shift spectrum between the cells based on this information. In this paper, the potential performance that can be achieved by this idealistic scheme (denoted **SSS-IV**) is estimated by analysing a single cell serving the aggregate of the two traffic profiles and considering a constant availability of S MHz.

Regardless of the applied sharing scheme, the achieved performance is expressed in terms of the average or cell edge user throughput, which is determined both on a minute-by-minute basis and averaged of a whole day.

The quantitative assessment of Section 5 is formulated for a scenario with two LTE cells, assuming a total spectrum availability of 20 MHz, which considering guard bands translates to an effective bandwidth of S = 18 MHz. We assume that spectrum can be shared in units of 180 kHz, i.e. the spectral width of a physical resource block, the unit of resource assignment in LTE.

4. STOCHASTIC ANALYSIS

In the assessment approach, the experienced performance is evaluated on a minute-by-minute basis, considering the time-varying cell-specific flow arrival rates and spectrum assignments, as determined by the spectrum sharing scheme.

Consider a given such minute t, characterised by flow arrival rates $\lambda_A(t)$, $\lambda_B(t)$ and spectrum assignments $S_A(t)$, $S_B(t)$. In the following, we omit the time variable for ease of notation. Given the assumed equal area of the zones and the spatial uniformity of the flow arrival process, the zone-specific arrival rate λ_{jk} is equal to λ_j / K, for $j \in \{A,B\}$ and $k = 1,\ldots,K$.

The flow-level performance of each individual cell can be analysed by means of a multi-class M/G/1 processor sharing model, which belongs to class of product-form 'networks' and is analytically tractable (see e.g. [8]). In particular, the joint probability distribution of the number n_{jk} in each zone (class) k of cell $j \in \{A,B\}$ is given by

$$\Pr\{n_{j1} = \nu_1, \cdots, n_{jK} = \nu_K\} = (1-\rho_j)\frac{(\nu_1 + \cdots + \nu_K)!}{\nu_1! \cdots \nu_K!}\prod_{k=1}^{K} \rho_{jk}^{\nu_k},$$

where $\rho_{jk} \equiv \lambda_{jk}\Omega / (S_j r_k)$ denotes the effective traffic load offered to cell j in zone k, $k = 1,\ldots,K$, and where $\rho_j \equiv \sum_j \rho_{jk}$ is the aggregate traffic load in cell j. Using Little's formula the expected download time of a flow in cell j's zone k can then be derived to be equal to

$$E\{T_{jk}\} = \frac{1}{\lambda_{jk}} E\{n_{jk}\} = \frac{\Omega}{S_j r_k (1-\rho_j)},$$

for $k = 1,\ldots,K$. These expressions are known to be insensitive to the specific form of the flow size distribution, depending on its mean only.

The expected throughput experienced by a flow in zone k of cell j can then be estimated by

$$E\{R_{jk}\} = \frac{\Omega}{E\{T_{jk}\}},$$

and hence the average throughput and cell edge throughput in cell j are given by

$$E\{R_j\} = \sum_{k=1}^{K} E\{R_{jk}\} \text{ and } E\{R_{jK}\}$$

respectively, which are used in our assessment as the key metrics indicating the attained performance in cell $j \in \{A,B\}$ in the considered minute. Appropriate weighting of these metrics with the aggregate or edge-specific flow arrival rates yields the overall performance aggregated over both cells.

The above analysis is also readily applied to the idealistic spectrum sharing case of $\vartheta\downarrow 0$ where the system is effectively considered as a single cell facing an aggregate flow arrival rate of $\lambda_A(t) + \lambda_B(t)$ and a fixed spectrum assignment of $S_A(t) + S_B(t) = S$. The derived performance serves as a reference for the performance achieved by other schemes.

It is noted that the proposed analytical approach is unsuitable for the assessment of spectrum sharing schemes that operate at other timescales other than $\vartheta \geq 1$ or $\vartheta\downarrow 0$. This is due to the fact that the approach is based on performance in stochastic equilibrium, which given the traffic characteristics is likely to be an unrealistic assumption for $\vartheta \in (0,1)$. For such interesting scenarios, dynamic system-level simulations seem to be the most appropriate approach, which we intend to pursue in our continued research.

5. SCENARIOS

For the presented quantitative study, the different system and traffic parameters need to set to specific values. See Table 1 and Table 2 for the general parameters and the zone-specific normalised data rates, respectively.

Table 1: Parameter settings.

Parameter	Value	Description
K	10	Number of zones per cell
Ω	1 Mb	Average flow size
S	18 MHz	Aggregate spectrum availability

Table 2: Zone-specific data rates (Mbps/MHz).

r_1	r_2	r_3	r_4	r_5	r_6	r_7	r_8	r_9	r_{10}
5.0	3.7	2.5	1.6	1.0	0.6	0.4	0.25	0.18	0.1

The cell-specific arrival rates are assumed to be given by the exemplary bi-modal profile

$$\lambda_j(t) = 1.5 + 1.25\sin\left(\pi\left(\frac{t + \tau_j + 540}{360}\right)\right),$$

with $\tau_A = 0$ and $\tau_B \in \{0,5,15,90,360\}$, specifying five distinct shifts of the daily traffic profiles of cells A and B. For the example case of $\tau_B = 360$, Figure 3 (top) depicts the considered cell-specific and aggregate traffic profiles, expressed in terms of the offered load, i.e. $\lambda_A(t)\Omega$, $\lambda_B(t)\Omega$ and $(\lambda_A(t) + \lambda_B(t))\,\Omega$ (in Mbps).

6. NUMERICAL RESULTS

In this section we present and discuss the quantitative results that have been derived applying the analytical approach to the scenarios discussed above.

First, as an example, consider the scenario with $\tau_B = 360$, spectrum sharing scheme SSS-I and sharing period $\vartheta \in \{15, 240, 1440\}$. See Figure 3. The top chart shows the scenario's cell-specific traffic profiles, as well as the sum of both traffic profiles, which is used to analyse the impact of idealised (continuous) spectrum sharing. The middle chart shows the optimised setting of S_A for each minute of the day. The reference case with $\vartheta = 1440$ assumes a fixed spectrum split, which considering the equivalence of the cell-specific traffic profiles is optimal at $S_A = S_B = 9$ MHz. For the cases of $\vartheta =$

240 and ϑ = 15, the individual periods with period-specific S_A levels are clearly recognisable. Observe that the maximum (minimum) level of S_A indeed corresponds with the maximum (minimum) level of $\lambda_A(t)$. At the other extreme, the idealised case of 'continuous' spectrum sharing is depicted by a constant assignment of 18 MHz to the single cell serving the aggregated traffic of cells A and B.

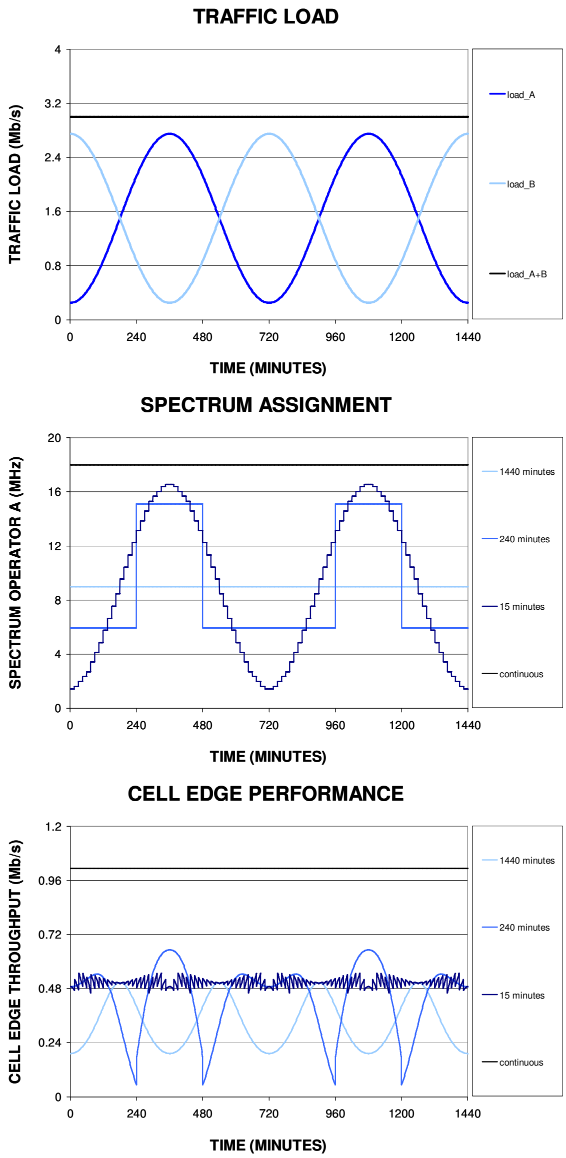

Figure 3: Traffic profiles, optimised spectrum assignments (SSS-I) and induced cell edge performance.

The bottom chart shows for each minute the expected user throughput as experienced at the cell edge. For the fixed spectrum assignment of 9 MHz per cell (ϑ = 1440), the cell edge performance appears to be best at those moments when the traffic loads in cells A and B are identical, i.e. $\lambda_A(t) = \lambda_B(t)$, and worst when their absolute difference $|\lambda_A(t) - \lambda_B(t)|$ is largest. In such a case, the cell edge performance is good in the cell with light traffic, poor in the other cell with heavy traffic, and hence the appropriately weighted (based on the cell-specific cell edge traffic loads) average of the cell edge throughputs is also rather poor. For the cases with time-varying spectrum assignments, the achieved performance in each minute is jointly determined by the current traffic load and the current spectrum split. For instance, a slight increase in the traffic load of cell A could lead to a slight reduction or improvement of the cell A performance, depending on whether in that period the spectrum assignment to cell A was decreased or increased. Similarly for cell B, of course, and the overall performance is then a weighted average of these cell-specific performance levels. This joint influence of time-varying traffic loads and spectrum assignments makes that the performance curve may seem a bit erratic, particularly for the case of 15 minute-based sharing actions (ϑ = 15), although the case of ϑ = 240 shows a similar pattern, but stretched in both the time and performance dimensions.

Observe that the attained performance for the case of ϑ = 15 is fairly constant over time, which indicates that the spectrum sharing scheme does a pretty decent job of matching supply (spectral resources) and demand (traffic load). Also, note that the performance is significantly better than for the default case with fixed spectrum sharing. Recall further that SSS-I is a load-based sharing scheme, which is not directly targeted to optimising cell edge performance, which explains why the case with ϑ = 15 does not necessarily outperform the case with ϑ = 240 in each individual minute.

The chart further clearly reveals that the idealised spectrum sharing scheme significantly outperforms any of the spectrum sharing schemes that periodically shifts spectrum between cells based on average traffic load estimates rather than instantaneous knowledge of actual numbers of active flows. This insight advocates that it may be worth the additional complexity required in the form of near-continuously exchanging actual load information and shifting spectrum accordingly. As mentioned before, the applied analytical approach is unsuitable to assess such cases, for which a dynamic system-level simulation approach is recommended.

Figure 4 shows the average (left chart column) and cell edge (right chart column) performance, appropriately averaged over the considered day, for different spectrum sharing schemes. In each chart, the performance achieved under idealised or fixed spectrum sharing is shown as well. These results show again the superiority of spectrum sharing based on actual traffic information at the finest timescale. The gains from spectrum sharing based on average load information ($\vartheta \geq 15$) are significant only if τ_B is relatively large, i.e. if the traffic profiles of the sharing cells are significantly different. We note that this does not need to mean that the different cells need to have their peak hours far apart, as was considered in the numerical experiments, as long as the typically less smooth (more bursty) traffic curves are sufficiently different in the various spectrum sharing periods. Naturally, when spectrum is shared based on instantaneous rather than average loads, i.e. if spectrum is shared based on the fluctuations of the actual number of on-going sessions, additional gains are expected to be achievable, even for low τ_B. As mentioned before, the applied analytical model is unsuitable to assess such spectrum sharing schemes, for which a simulation approach is recommended.

The results further show that the achieved cell edge throughputs are indeed better if the spectrum sharing scheme is directly targeted at optimising cell edge performance, rather then 'simply' trying to match spectrum split to the relative cell loads. The same holds for the average performance.

7. CONCLUDING REMARKS

In this paper we have presented an analytical approach to assess the performance gains from spectrum sharing in a system with two co-sited cells. In the numerical experiments, particular focus has been placed on the performance impact of the timescale at which spectrum is shared.

The results indicate the superiority of spectrum sharing on a timescale at which flows arrive and depart (typically seconds), where sharing actions are based on actual numbers of active flows and their respective link qualities. Implementationally less complex spectrum sharing mechanisms that operate on a timescale of minutes or higher, based on average load information, is also shown to bring significant performance gains, as long as the cell-specific traffic curves are sufficiently different at the minute (or higher) timescale.

In light of the obtained insights, our continued research will primarily concentrate on the development and (simulation-based) assessment of spectrum sharing algorithms at sub-minute timescales. Furthermore, in cooperation with SAPHYRE, we intend to concentrate on the system-level evaluation of non-orthogonal spectrum sharing schemes.

ACKNOWLEDGMENT

This work has been performed in the framework of the European research project SAPHYRE [5], which is partly funded by the European Union under its FP7 ICT Objective 1.1 - The Network of the Future.

REFERENCES

[1] ITU: ICT Regulation Toolkit, available online: http://www.ictregulationtoolkit.org/en/Section.3451.html.

[2] R. Etkin, A. Parekh and D. Tse, 'Spectrum sharing for unlicensed bands', *IEEE JSAC*, vol. 25, no. 3, 2007.

[3] S. Kumar et al., 'Spectrum sharing for next generation wireless communication networks', *Proceedings of CogART '08*, Aalborg, Denmark, 2008.

[4] G. Salami et al., 'Nonpool based spectrum sharing for two UMTS operators in the UMTS Extension Band', *Proceedings of PIMRC '08*, Cannes, France, 2008.

[5] FP7 SAPHYRE (*'Sharing Physical Resources - Mechanisms and Implementations for Wireless Networks'*) project, see www.saphyre.eu, 2011

[6] H. Salameh, M. Krunz and O. Younis, 'Cooperative adaptive spectrum sharing in cognitive radio networks', *IEEE Trans. on Networking*, vol. 18, no. 4, 2010.

[7] J. Lee, J. G. Andrews and D. Hong, 'The effect of interference cancelation on spectrum-sharing transmission capacity', *Proceedings of ICC '11*, Kyoto, Japan, 2011.

[8] J.W. Cohen, 'The multiple phase service network with generalized processor sharing', *Acta Informatica*, vol. 12, 1979.

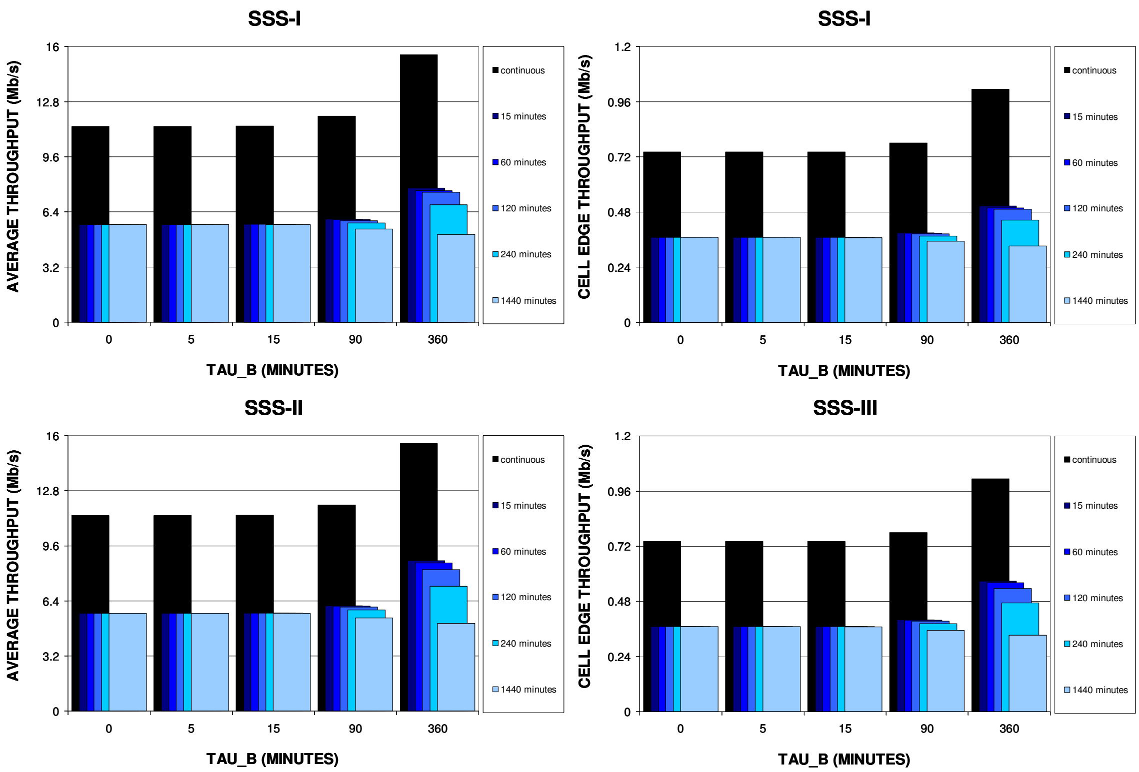

Figure 4: Day-averaged (average and cell edge) user throughput performance versus τ_B and the spectrum sharing timescale for different spectrum sharing schemes.