What and How Much to Gain by Spectral Agility?

Chun-Ting Chou, Hyoil Kim, Kang G. Shin and Sai Shankar N

CSE-TR-494-04

August 2004

Computer Science and Engineering Division
Room 3402 EECS Building

What and How Much to Gain by Spectral Agility?

Chun-Ting Chou, Hyoil Kim and Kang G. Shin
Real-Time Computing Laboratory
The University of Michigan
Ann Arbor, MI 48109-2122, U.S.A.
{choujt,hyoilkim,kgshin}@umich.edu

Sai Shankar N
Philips Research USA
345 Scarborough Road,
Briarcliff Manor, NY 10510, USA
sai.shankar@philips.com

Abstract

Static spectrum allocation has resulted in low spectrum efficiency in licensed bands and poor performance of radio devices in crowded unlicensed bands. To remedy these problems, we exploit the concept of "spectral agility" such that radio devices can *dynamically* utilize idle spectral bands. We establish a mathematical model for the performance gain made by spectral agility, and use the model to evaluate important performance metrics such as spectrum efficiency, throughput of a spectral-agile network, and packet blocking/waiting time of spectral-agile devices.

We propose three basic mechanisms to realize spectral-agile networks: spectrum opportunity discovery, spectrum opportunity management, and spectrum use coordination. These mechanisms are implemented in the *ns-2*, and the control overhead incurred by using spectral agility is evaluated. Our simulation results have shown that the throughput of a spectral-agile network is improved by up to 90%, and the improvement is very close to the performance bound predicted by our analytical (mathematical) model. These results demonstrate and confirm the spectral agility's capability of improving spectral utilization in an efficient, distributed, and automatic manner.

I. Introduction

Conventional wireless devices are only allowed to operate in designated spectrum bands primarily because of regulatory restrictions. Within each designated band, radio devices adopt specific communication protocols and use fixed modulation schemes and medium access control. Even though such designs simplify protocol and hardware developments, there exists one major potential problem with them — inefficient utilization of precious spectral resources. The main cause of inefficiency is that radio devices in crowded spectral bands are prohibited from (due to regulations) and incapable of (due to hardware limitation) using other idle or sparsely-used spectrum bands. Such spectrum inefficiency is becoming a serious problem as more and more communication protocols and commercial wireless devices are being developed to operate in crowded unlicensed spectrum bands.

Recently, the concept of spectral agility has been drawing considerable attention for its potential to alleviate the inefficiency problem. For example, the US Federal Communications Commission (FCC) has issued a Notice of Public Rulemaking and Order regarding so-called *cognitive radio* technologies [1]. The Defense Advanced Research Projects Agency (DARPA) has also started the neXt Generation (XG) Communications Program to develop new technologies which allow multiple users to share the spectrum through adaptive mechanisms [2]. The US Army has also been researching the so-called "Adaptive Spectrum Exploitation" (ASE) for real-time spectrum management in the battlefield [3][4]. Although the focus of these programs are somewhat different, the basic principles are the same: if radio devices can explore the wireless spectrum and locate sparsely-used spectral bands, they can exploit them opportunistically to improve not only the devices' performance but also the overall spectrum utilization. In the long run, such "spectral agility" can also facilitate secondary markets in spectrum use (e.g., a licensee may allow secondary spectrum uses by a third party) and automated frequency coordination among different radio systems [1].

Of course, such spectral agility cannot be realized without developing new hardware/software and changing the current spectrum allocation policies. Fortunately, the advances in software defined radio (SDR) [5][6] have enabled the development of flexible and powerful radio interfaces for supporting spectral agility. Also, the FCC's ongoing review of the current spectrum regulations is expediting the adoption of more flexible spectrum allocation policies for spectral agility. However, there remain many open questions that we need to answer before realizing spectral agility. The first and the foremost question is to what extent the improvement can be, in terms of spectrum utilization and individual devices' performance. Without a clear understanding of the achievable improvement, one cannot justify the use of spectral agility since controlling a spectral-agile network may incur a considerable amount of overhead. This leads to several implementation questions, including how individual devices discover and identify sparsely-used spectrum bands, how to characterize or prioritize these spectrum bands, and how and when to utilize them. Obviously, different implementations incur different amounts of control overhead. Thus, the final question is how the control overhead may degrade the improvement achieved with spectral agility. Without answering these questions, it is meaningless and difficult to develop spectral-agile communication protocols and networks in an effective way.

In this paper, we address some of these questions. First, we establish an analytical model and provide an upper-bound performance analysis for radio networks with spectral agility. The analysis sets the benchmark of an ideal spectral-agile network's performance, and thus, enables the evaluation of different implementations. Then, we propose a set of spectral agility-related functionalities, including *spectrum opportunity discovery, spectrum opportunity management, and spectrum use coordination*, which constitute the basic building blocks of a spectral-agile network [7]. Based on these functionalities, a variety of spectral-agile networks can be developed. We implement these functionalities on the ns-2 implementation of the IEEE 802.11 wireless LANs which are currently operating in very crowded unlicensed bands (e.g., 2.4GHz bands for the 802.11b,g standards and 5GHz for the 802.11a standard). Finally, we conduct the *ns-2*-based simulation to demonstrate the benefits of using spectral agility in both existing and emerging wireless networks.

The rest of this paper is organized as follows. In Section II, we describe the system model and assumptions while the analytical model and some numerical results are presented in Section III. Section IV introduces the spectral-agility-related functionalities, and the simulation results are analyzed and discussed in Section V. Finally, conclusions are drawn in Section VI.

II. System Model

We consider two types of networks, namely *primary* and *secondary* networks. A primary network has exclusive access to designated spectral bands while a secondary network only accesses a spectral band when the corresponding primary network does not use that band. For example, a primary network can be any licensed-band network, and a secondary network is an unlicensed-band network such as an IEEE 802.11 wireless LAN. To realize such an opportunistic use of primary networks' idle spectral resources, we assume that a secondary network has spectral agility, which is enabled by the SDR. It is then a secondary network's responsibility to locate available resources, in both spectral and temporal domains, as shown in Figure 1.

Even though it is desirable to have the entire spectrum accessible to a secondary network, hardware limitations (such as antenna design) usually determine the accessible range. Therefore, the "wireless spectrum" in this paper is referred to as the portion of the wireless spectrum which can be accessed by a secondary network. The spectrum is divided into "channels," each of which is the smallest unit of a spectral band. We assume that each secondary network only uses a single channel for basic communication, but it can also use multiple channels for better performance. For example, the SDR makes it possible to adopt a modulation scheme requiring more bandwidth when several adjacent channels are available simultaneously. Moreover, it is also possible to use discrete channels as sub-carriers of a multi-carrier modulation scheme (such as the Orthogonal Frequency Division Multiplexing (OFDM)) or use these channels in adaptive frequency hopping [3] for transmission in a multi-path fading environment.

We assume that the temporal usage of each channel (by the primary network/user of that channel) can be characterized by a random process. When a primary network does not always use its designated channel, it leaves some "holes," or idle time slots, in the channel's usage schedule which may be exploited by secondary networks. As shown in Figure 1, the blank slots represent such holes each of which is referred to as a *spectral opportunity* in the rest of the paper. For example, there exists a spectral opportunity in channel 4 after $t = t_1$. Moreover, the entire spectrum is regarded as providing a spectral opportunity during $[t_2, t_3]$. Depending on the primary network's spectrum usage pattern, the duration of a spectral opportunity can be up to several hours or even days in spectral bands reserved for emergency, or can be only few milliseconds in heavily-used spectral bands. It is relatively easy for a secondary network to use long-lasting opportunities. However, for short-lasting or "ephemeral" opportunities, a secondary network may not be able to detect their existence and then utilize them before they "expire." Therefore, we only focus on the case when spectral opportunities last in the order of seconds.

In order to exploit spectral opportunities, a secondary network has to first scan the spectrum, either periodically or randomly, to discover the opportunities. It should be noted that our problem differs significantly from the problems of using dynamic frequency selection mechanisms in the existing systems, such as Dynamic Channel Selection (DCS) [8] in cellular networks, Dynamic Frequency Selection (DFS) [9] in the IEEE 802.11h standard or Auto Frequency Allocation (AFA) [10] in the HiperLAN. These schemes address the problem of choosing a good channel (either a frequency in the Frequency Division Multiple Access (FDMA) system, or time slots in the Time Division Multiple Access (TDMA) system) so that transmission in that channel may experience less interference or cause less interference to other transmissions in the same channel. In our problem, a spectral-agile network seeks both spectral and temporal opportunities in the wireless spectrum, and utilizes these opportunities opportunistically. Among the thus-found opportunities, a spectral-agile secondary network decides on which opportunities to use and when to utilize them. If and when activities of a primary network are detected, the secondary network must vacate the channel in order not to interfere with the primary network. Obviously, all wireless nodes (i.e., radio devices) in a secondary network must always take the same spectral opportunity to maintain their inter-connectivity. Therefore, the wireless nodes in a spectral-agile network must disseminate the information of spectral opportunities and the decision of switching to different opportunities. These procedures are detailed in Section IV.

III. Analytical Model for Performance Improvements

We establish a mathematical model to analyze the potential improvements by using spectral agility, in terms of a secondary network's spectral utilization and packet blocking/waiting time.

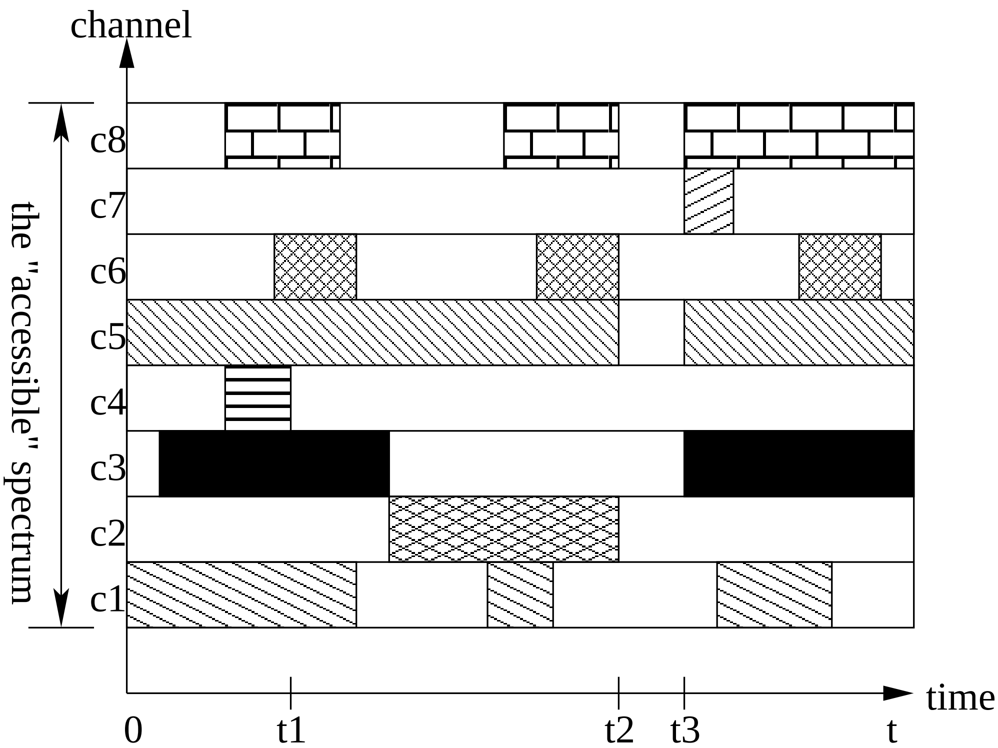

Fig. 1. Spectrum opportunities for spectral-agile devices

The spectral utilization of a secondary network is measured by the total amount of time a secondary network can access a channel for transmission. One can convert the channel access time to the network's actual throughput once the underlying medium access control (MAC) and modulation mechanisms are specified. Therefore, we use channel access time so as not to be confined to any specific MAC and modulation schemes. The packet blocking time is defined as the time interval during which a secondary network has no spectral opportunity to utilize (thus, it has to suspend all transmissions).

Suppose there are N primary networks each with one designated channel, and there are M secondary networks seeking spectral opportunities. The usage pattern of the primary network in each channel is assumed to be an *i.i.d.* ON/OFF random process with independent ON- and OFF-periods. An ON-period represents that a channel is busy while an OFF-period is regarded as a potential spectral opportunity for secondary networks. To simplify our analysis, we assume that the distributions of both ON- and OFF-periods in each channel are exponentially-distributed with means equal to T_{on} and T_{off}, respectively. We will explore different distributions using simple simulations at the end of this section.

In order to provide a performance upper-bound, we assume that each secondary network has an infinite amount of traffic to transmit. Moreover, each secondary network can scan a channel, switch to a channel, and vacate a channel instantly (when claimed by the primary network) without incurring any control overhead or delay. The control overhead and delays are implementation-dependent, and their impact on the improvement will be investigated in Section IV. In order to provide a comparative feel for the performance improvement of using spectral agility, we introduce and use a "naive" secondary network which listens to a fixed channel (i.e., without spectral agility), and transmits only when that channel is not used by the primary network. The spectral utilization of such a naive secondary network can easily be computed as $\frac{T_{off}}{T_{on}+T_{off}}$, and the average blocking time is T_{on}.

A. A Special Case: $M = 1$

We first consider a special case when there is only one secondary network. As shown in Figure 2, the only time interval during which a secondary network has no channel for traffic transmission is when all channels are occupied by the primary networks. Such blocking intervals, denoted as t_{block}, always begin when a channel switches from an OFF-period to an ON-period and ends when one channel switches from an ON-period to an OFF-period. Therefore, t_{block} is computed as

$$t_{block} = \min_{i=1,2,\cdots,N} (T^{(i)}_{remain}), \tag{1}$$

where $T^{(i)}_{remain}$ is the remaining ON-period in channel i. Assuming that the ON-periods are independent and exponentially distributed, one can compute the distribution of t_{block} as

$$P(t_{block} = t) = \frac{N \cdot e^{-\frac{T_{on}}{N}t}}{T_{on}}. \tag{2}$$

Eq. (2) shows that with spectral agility, a secondary network can reduce the average blocking time to $\frac{T_{on}}{N}$, as compared to T_{on} in the naive secondary network without agility. The spectral utilization of such a spectral-agile secondary network is obtained by

$$U = 1 - \frac{N(p^{N-1} \cdot \frac{T_{on}}{N})}{T_{on} + T_{off}}, \tag{3}$$

where $p = \frac{T_{on}}{T_{on}+T_{off}}$ is the probability that a channel is occupied by the primary network. Eq. (3) is derived based on the fact that a blocking interval starts only if a channel switches from an OFF-period to an ON-period while all other channels have already been in the ON-periods. Eq. (3) can be simplified further to

$$U = 1 - (\frac{T_{on}}{T_{on} + T_{off}})^N, \tag{4}$$

showing that the spectral utilization of a secondary network is a simple function of the primary network's channel utilization. Finally, the improvement of the spectral utilization achieved by a spectral-agile secondary network is computed as

$$I = \frac{U}{1 - \frac{T_{on}}{T_{on}+T_{off}}} - 1, \tag{5}$$

as compared to the naive secondary network.

B. The General Case: $M > 1$

Eq. (4) shows that the spectral utilization of a spectral-agile secondary network is simply a function of the primary network's channel utilization, $\tau = \frac{T_{on}}{T_{on}+T_{off}}$. We can generalize this simple equation for the case when different channels have different utilizations, say, channel i with utilization $\tau_i = \frac{T^{(i)}_{on}}{T^{(i)}_{on}+T^{(i)}_{off}}$. Based on Eq. (4), the fraction of time during which there are k channels available simultaneously is computed as

$$r_k = \sum_{c=1}^{\frac{N!}{k!(N-k)!}} \left[\prod_{i \in S^k_c} (1 - \tau_i) \prod_{j \in \{1,2,\cdots,N\} - S^k_c} \tau_j \right], \tag{6}$$

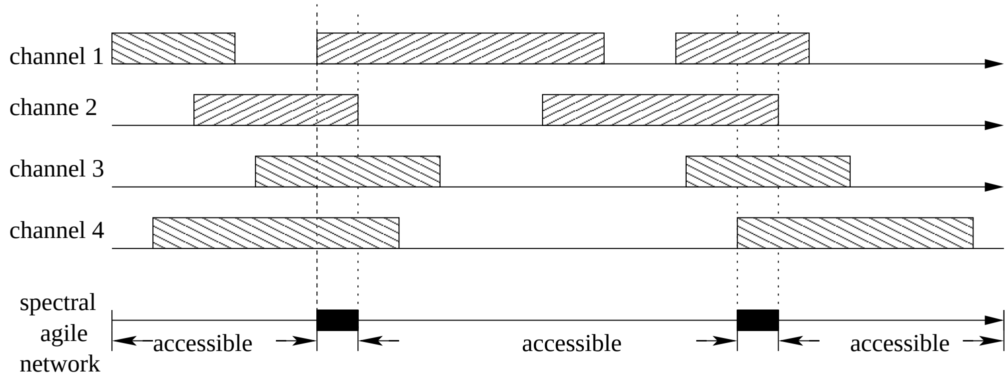

Fig. 2. A special case: N=4

where S_c^k is a set of k channels, chosen from N channels, which are available for spectral-agile secondary networks. For example, we can set $S_1^k = \{1, 2, \cdots, k\}$, $S_2^k = \{2, 3, \cdots, k+1\}$, and so on.

To further generalize our analysis, we assume that there are $M > 1$ spectral-agile secondary networks trying to exploit available spectral opportunities. Obviously, each secondary network obtains exactly one channel if there are no less than M channels available. Otherwise, these M secondary networks have to share less than M available channels. The spectral utilization of each spectral-agile network is then computed by

$$U_{agile} = \sum_{k=0}^{N} \frac{\min(M, k) r_k}{M}. \tag{7}$$

As we mentioned in Section II, the SDR enables a radio device to dynamically use a variety of MAC and modulation schemes, depending on the underlying wireless environment. Therefore, a spectral-agile network can use multiple channels simultaneously, thus achieving more channel access time for better performance. We will describe how to analyze this type of spectral-agile networks at the end of this section.

Since there are $M > 1$ secondary networks, each aforementioned naive (i.e., non-agile) secondary network can use two approaches to selection of a channel: (1) each network randomly selects its own channel independently of others, and (2) all secondary networks cooperate in a way that no more than one secondary network uses the same channel, if possible. The advantage of the first approach is the simplicity while the advantage of the second approach is that each secondary network obtains more channel access time.

1) Random channel selection: Given that a secondary network chooses channel i, the probability that the other k secondary networks also choose the same channel is

$$p_k = \frac{(M-1)!}{k!(M-1-k)!}(\frac{1}{N})^k(\frac{N-1}{N})^{M-1-k}. \tag{8}$$

Therefore, the average channel access time that a spectral-agile network can acquire, given that it has chosen channel i, is

$$T_i = \sum_{k=0}^{M-1} p_k \frac{T_{off}^{(i)}}{(k+1)(T_{on}^{(i)} + T_{off}^{(i)})}. \tag{9}$$

The fraction of time in which each (no-agility) secondary network has a channel for its traffic transmission is then computed as

$$U_{random} = \frac{1}{N}\sum_{i=1}^{N} T_i. \tag{10}$$

2) Coordinated channel selection: If each secondary network coordinates its selection of a channel with the others in order to avoid the case of more than one network trying to use the same channel, the fraction of channel access time is computed as

$$U_{coordinated} = \frac{\sum_{c=1}^{\frac{N!}{M!(N-M)!}} \frac{1}{M}\sum_{i\in S_c^M} \frac{T_{off}^{(i)}}{T_{on}^{(i)}+T_{off}^{(i)}}}{\frac{N!}{M!(N-M)!}}. \tag{11}$$

Here, we simply average all the possibilities of choosing M channels from N channels for naive secondary networks. We set $\frac{N!}{M!(N-M)!} = 1$ in case of $M > N$.

We can now compare the spectral utilization of a secondary network using (1) spectral agility, (2) no agility with random channel selection (Approach I), and (3) no agility with coordinated channel selection (Approach II) based on Eqs. (7), (10), and (11). We investigate two scenarios with $N = 12$ and $N = 3$. The main reason for choosing these numbers is that there are 12 (non-overlapping) channels in the 5-GHz band for the IEEE 802.11a wireless LAN and 3 (non-overlapping) channels in the 2.4-GHz band for the IEEE 802.11b wireless LAN.[1] Therefore, even though spectral agility cannot be applied immediately to the licensed bands due to the current regulations, the 802.11 wireless LAN may use spectral agility to improve performance in the crowded, unlicensed bands.

Figure 3 shows the case of $N = 12$ and $M = 9$ with different average channel loads generated by the primary networks. For each given channel load, we choose the loads of these 12 channels to be homogeneous or heterogeneous. In case of homogeneous loads, each channel is assigned a load equal to the average channel load, while, in case of heterogeneous loads, different channels are assigned different loads with their variance maximized (i.e., the utilization of each channel differs significantly from each other). The improvement shown in Figure 3 is defined as

$$\text{improvement (\%)} = \left(\frac{U_{agile}}{U_{random/coordinated}} - 1\right)\cdot 100\%, \tag{12}$$

where U_{agile}, U_{random}, and $U_{coordinated}$ are given in Eqs. (7), (10), and (11), respectively. The results demonstrate that use of spectral agility always achieves a higher spectral utilization for a secondary network than the case of no agility with random channel selection or coordinated channel selection. Of course, the improvement by using spectral agility is much less (still more than 25% in most cases) than the case of no agility with coordinated channel selection (Figure 3-(b)). Note, however, that coordinated channel selection needs off-line channel information. If the channel loads range widely, it is possible that a secondary network may choose busy channels (unless it scans all channels for a long period of time). In contrast, using spectral agility allows a secondary network to dynamically choose the channel with the least activities. Such advantages are also illustrated in Figure 3, where we achieve an extra 8-10% improvement for heterogeneous loads when the channel load is around 0.2 $\sim$ 0.3.

An interesting observation is that the improvement ratio (i.e., Eq. (12)) saturates when the average channel load of the primary network is greater than 0.5. This can be explained by Figure 4, which shows the fraction of time for which a secondary network can access a channel. The fraction of time a spectral-agile network can access a channel linearly decreases with the increase in

[1] According to the US regulation, there will be more released channels in the 5-GHz band.

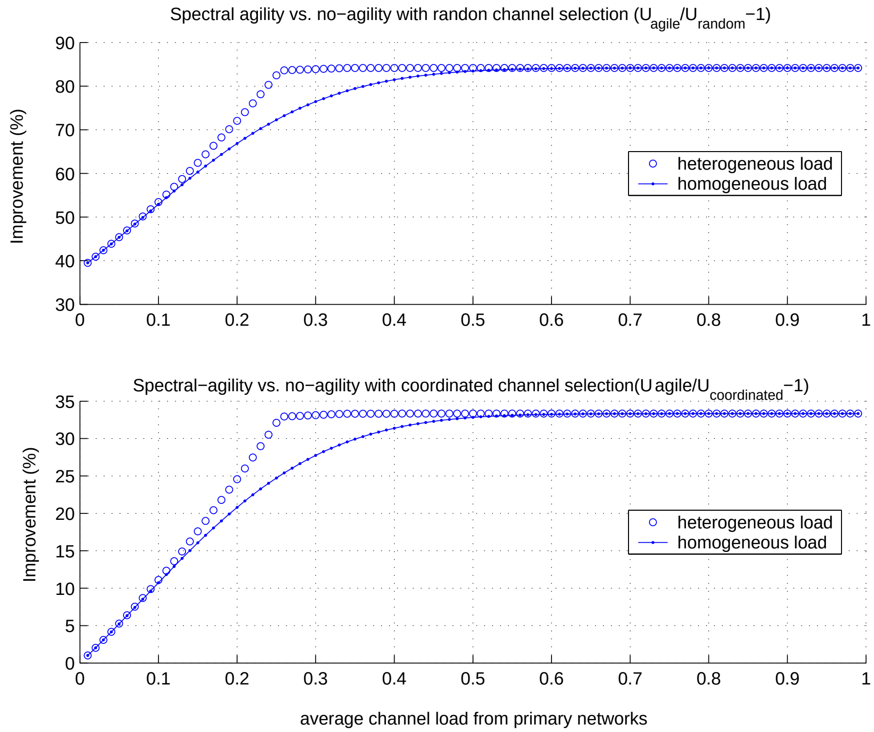

Fig. 3. Improvement of spectral utilization for spectral-agile networks: $N = 12$ and $M = 9$. *Although the figure shows the maximal improvement percentage (82%) occurs when the channel load approaches 1, it does not suggest that spectral agility generates the greatest amount of spectral opportunities. Instead, it shows that, for example, with load of 0.99, the average channel access time for a spectral-agile node increases from 0.01=1-0.99 (i.e., no-agility) to 0.0182 sec out of an one-second period as also shown in Figure 4

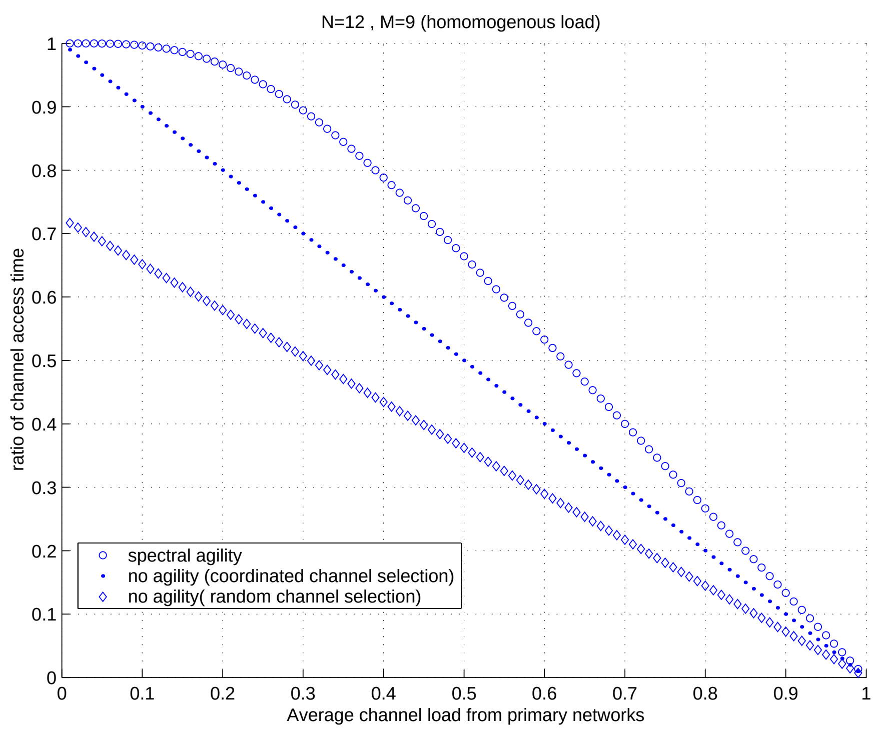

Fig. 4. Fraction of channel access time: $N = 12$ and $M = 9$. *This figure, together with Figure 3, suggest that a spectral-agile secondary network benefits most from spectral agility when the channel utilization of a primary network is lightly-(0.2) or moderately-loaded ($0.7 \sim 0.8$).

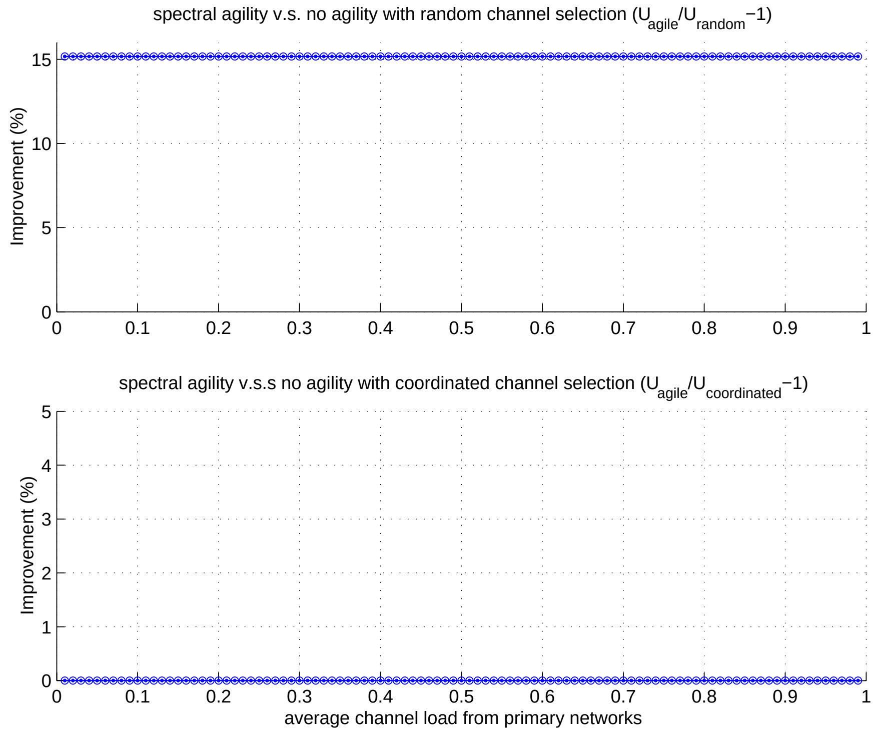

Fig. 5. Improvement of spectral utilization for spectral-agile networks: $N = 3$ and $M = 5$. *The figures shows that when the number of available channels is less than the number of secondary networks, using spectral agility generates the same performance as that of using static coordinated channel selection. However, spectral agility still outperforms static random channel selection.

the average channel load from primary networks beyond 0.3 in all three cases (i.e., with spectral agility, no agility with coordinated channel selection, and no agility with random channel selection). Because of such linearity, the improvement ratio of using spectral agility, as compared to no-agility cases, remains unchanged when the channel load is greater than 0.3 in Figure 3. Figure 4 also suggests that when the average channel load of the primary network is very large, it does not make much sense to use spectral agility as indicated by Figure 3 (even though it shows an 80% improvement with the load of 0.9). This is because when the channel is extremely busy, the amount of access time that each spectral-agile network can obtain is very small (less than 10% of the total time with the channel load of 0.9). Therefore, the control overhead (incurred by using spectral agility) may exhaust most of the channel access time a secondary network acquires, hence, easily offsetting the improvement gained with spectral agility.

Next, we consider the case of $M > N$ and choose $N = 3$ and $M = 5$ as an example. Figure 5-(b) shows that using spectral agility and using no agility with coordinated channel selection achieve exactly the same performance (i.e., no improvement). The results make sense because when $M > N$, there are simply not enough channels for all secondary networks (so they have to share idle channels with each other). In fact, one can simplify both Eqs. (7) and (11) as

$$U_{agile} = U_{coordinated} = \frac{1}{M} \sum_{i=1}^{N} \frac{T_{off}^{(i)}}{T_{on}^{(i)} + T_{off}^{(i)}}, \tag{13}$$

when $M > N$ and verify the result in Figure 5-(b). There are some marginal improvements by using spectral agility as compared to using no agility with random channel selection as shown in in Figure 5-(a). This is simply because some idle channels may be left unused in the case of random channel selection.

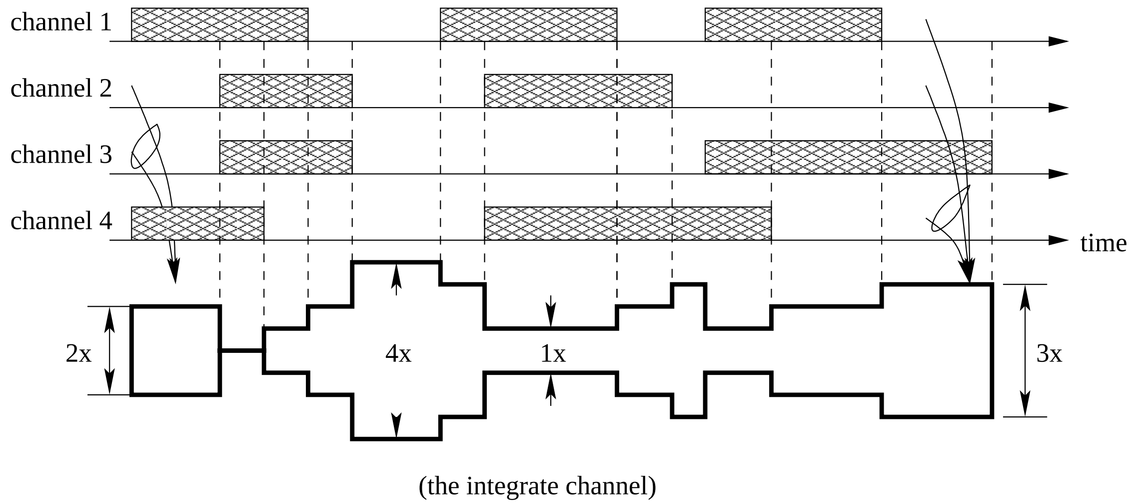

Fig. 6. Use of multiple channels: $N = 4$

Figures 3 and 5 show that radio devices can only benefit from spectral agility when there are enough resources for opportunistic uses (i.e., $M < N$). Fortunately, field studies have shown that there are many under-utilized spectral resources in some wireless spectral band [11][12]. Moreover, there are two additional advantages of using spectral agility that we have not yet discussed when $M > N$. First, Eq. (2) shows that when the spectral agility is used, the average blocking time is reduced by a factor of N in the special case or reduced from $\frac{\sum T_{on}^{(i)}}{N}$ to $\frac{1}{\sum \frac{1}{T_{on}^{(i)}}}$ in the general case. Thus, even though the spectral utilization is not improved by using spectral agility when $M > N$, the packet delays are reduced significantly by using spectral agility. Another advantage is the spectral-agile network's capability of using multiple channels. In the above analysis, we assumed that a spectral-agile network (or more precisely, the wireless nodes in the network) always uses a single channel, even when more than one channel are available. We can expect that if a spectral-agile network can use all available channels, the performance must be improved. Figure 6 illustrates this scenario in which each spectral-agile network aggregates all available channels into a single, higher-capacity spectral opportunity. Then, all spectral-agile networks use this aggregated opportunity, instead of using separate channels for transmission as discussed earlier. This will provide a multiplexing gain just as we can obtain by multiplexing several traffic flows on a high-capacity transmission line in conventional wired networks.

Before analyzing the multiplexing gain of using multiple channels, we would like to investigate the effects of different ON/OFF distributions on the improvement of spectral utilization by using spectral agility. The main purpose of this study is to verify the applicability of our model, which is established based on the assumption of exponentially-distributed ON-/OFF periods. Here, we use Matlab to simulate the random ON/OFF periods and calculate the total time intervals of overlapping ON-periods (i.e., the blocking intervals for a spectral-agile network) for the case of $N = 3$ and $M = 1$. We use exponential (as in our earlier derivation), uniform, and Rayleigh distributions. Figure 7 shows a very good match between our analytical results and the simple simulation results, demonstrating the applicability of our analytical model. The reason why the improvement ratios (again as defined in Eq. (12)) are much higher (up to 200%) is that there is only one spectral-agile network seeking spectral opportunities, and thus, it need not share spectral opportunities with other spectral-agile networks. However, as we discussed earlier, such a large improvement ratio, in fact, represents only a very small increase of channel access time for a secondary network if the average channel load of the primary network is extremely high. Therefore, one should not expect improve-

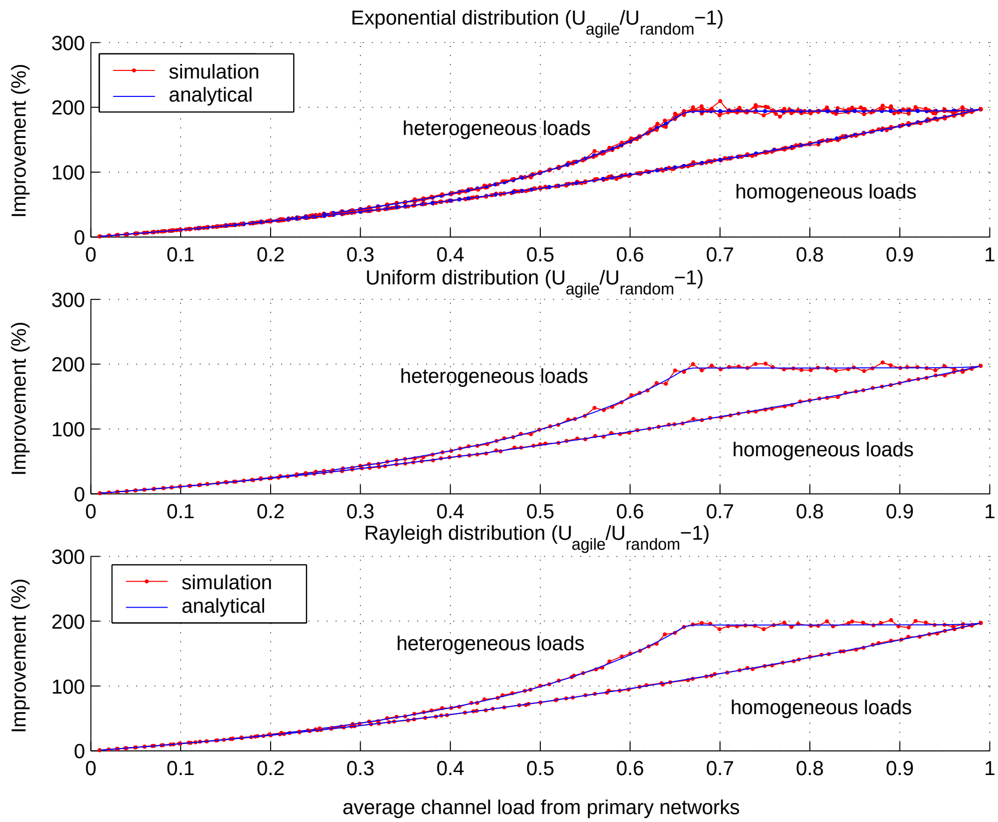

Fig. 7. Improvement of spectral utilization for spectral-agile networks: different ON/OFF distributions *Although the figure shows the maximal improvement percentage (200%) occurs when the channel load approaches 1, it does not suggest that spectral agility generates the greatest amount of spectral opportunities. Instead, it shows that, for example, with load of 0.99, the average channel access time for a spectral-agile node increases from 0.01=1-0.99 (i.e., no-agility) to 0.03 sec out of an one-second period, similar to what shows in Figure 3.

ment in reality, given the control overhead incurred by spectral agility, when the average channel load of the primary network is very high.

C. *Multiplexing Gain of Using Multiple Channels*

If all spectral-agile networks use the aggregated spectral opportunities, packets from all spectral-agile networks share the same aggregate "channel" with a varying transmission capacity as shown in Figure 6. The transmission capacity depends on how many primary networks are using the channels, and the distribution of the transmission capacity is determined by Eq. (6). If the arrival process in each spectral-agile network is assumed to follow a Poisson process with rate λ, the aggregated arrival process is also Poisson with rate $M\lambda$. Therefore, the system can be modeled as an M/G/1 queueing system. However, it is possible that all channels are occupied by primary networks with probability r_0 in Eq. (6) and for an average duration of $\frac{1}{\sum_i^N \frac{1}{T_{on}^{(i)}}}$, so that the transmission capacity is 0 from the spectral-agile networks' perspectives. This blocking process is modeled as another arrival process with rate $\frac{1}{r_0}$, and the "packet" with an average service time of $\frac{1}{\sum_i^N \frac{1}{T_{on}^{(i)}}}$. The resulting M/G/1 queue with preemptive priority is illustrated in Figure 8-(b). The average packet waiting time of a spectral-agile network is then computed by using the results in [13] as

$$T_{SA} = \frac{\frac{1}{\mu_p}(1-\rho_p-\rho_{SA}) + R_{SA}}{(1-\rho_p)(1-\rho_p-\rho_{SA})}, \tag{14}$$

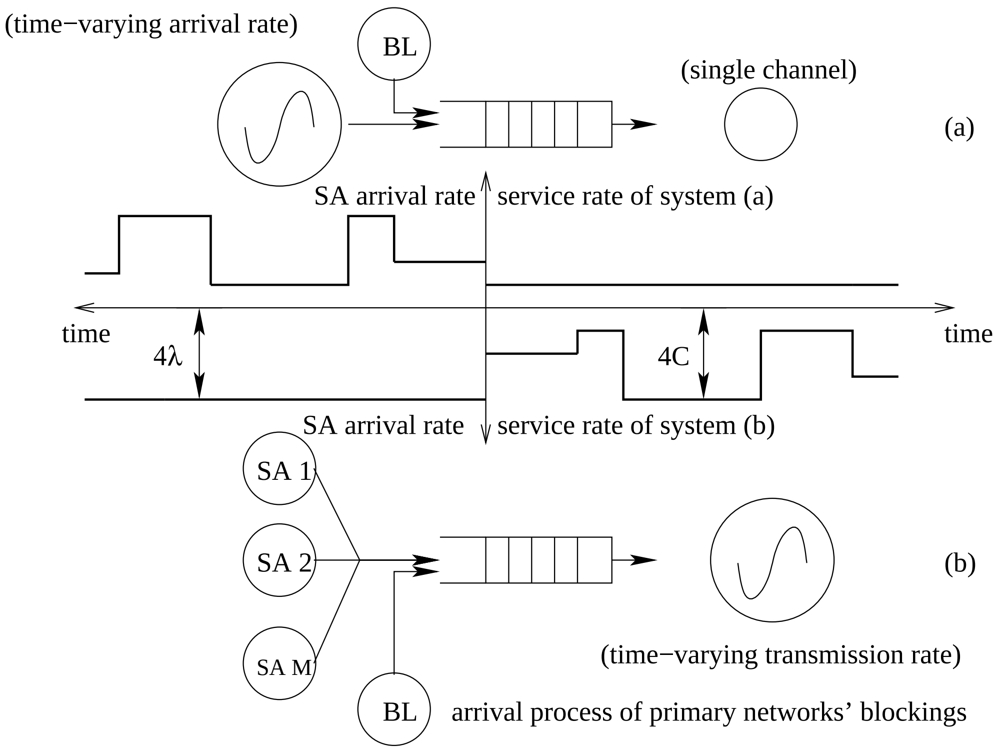

Fig. 8. Queueing models for statistical multiplexing gain: $N = 4$

where $\mu_p = \sum_i^N \frac{1}{T_{on}^{(i)}}$, $\rho_p = \frac{1}{r_0 \mu_p}$, ρ_{SA} represents the server utilization of the spectral-agile network, and R_{SA} represents the average residual service time seen by the packets of spectral-agile networks.

If we assume the average packet size is L and the transmission capacity of a single channel is C, ρ_{SA} is computed as

$$\rho_{SA} = \frac{M\lambda}{\mu_{SA}}, \tag{15}$$

where $\frac{1}{\mu_{SA}} = \sum_{i=1}^N \frac{L}{i \cdot C} r_i$ is the average service time of a packet from spectral-agile networks. Finally, the residual time R_{SA} is computed as

$$R_{SA} = \frac{1}{2} \left[M\lambda \sum_{i=1}^{N} (\frac{L}{i \cdot C})^2 r_i + \frac{2}{r_0 \mu_p^2} \right], \tag{16}$$

as derived in [13].

We can use the M/G/1 queueing model with preemptive priority for the case when each spectral-agile network uses at most one channel. In this case, the "service rate" is constant (from the perspective of packets of a secondary network), and is equal to the transmission capacity of a single channel unless all the channels are occupied by the primary networks. However, packet arrivals in a channel changes with the number of active primary networks. That is, the packet arrivals in a channel are dependent on the state of the primary network's occupation of the spectrum. The less the number of idle channels, the greater the arrival rate in each idle channel. We can model this arrival process as a Markov-Modulated Poisson Process (MMPP) using Eq. (6), but for the sake of simplicity we approximate the arrival process as a Poisson process, which gives us an M/D/1 queue with preemptive priority. In order to model it as a Poisson process, we need to calculate the average arrival rate.

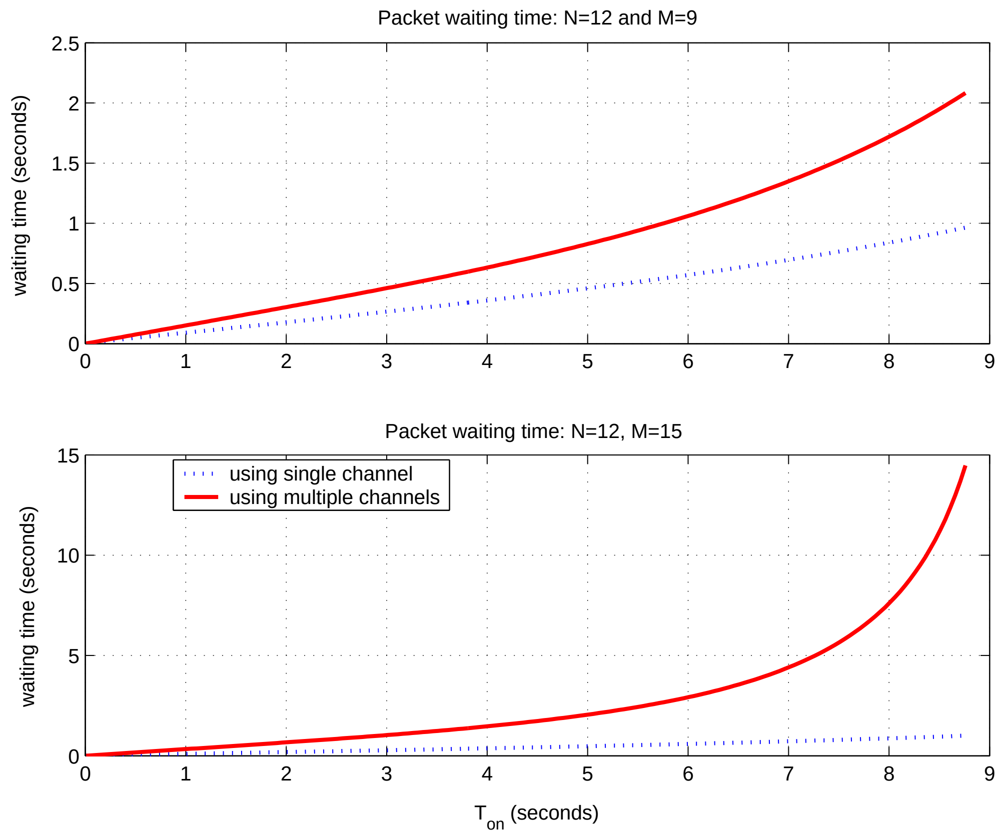

Fig. 9. Average waiting time of packets from a spectral-agile network: $T_{off} = 1$ and $\frac{L}{C} = 0.1$

1) Case I: $M < N$: If there are at least M channels available, then the arrival rate at the $M/G/1$ queue is just λ. If the number of available channels is $M - 1$, then one of the spectral-agile networks joins the channel which has already been "occupied" by another spectral-agile network. That is, multiple spectral-agile networks share one channel. The average arrival rate is then $\frac{M}{M-1}\lambda r_k$. Proceeding similarly, we have the average arrival rate computed as

$$\lambda_{new} = \sum_{i=M}^{N} r_i \lambda + \sum_{i=1}^{M-1} r_{M-i} \left(\frac{M}{M-i} \right) \lambda. \tag{17}$$

2) Case II: $M > N$: In this case, we have more spectral-agile networks than the total number of channels. If none of the channels are occupied by the primary network, then the best-case arrangement occurs when each channel has $\lceil \frac{M}{N} \rceil$ spectral-agile networks. Proceeding similarly to the previous subsection, we have the arrival rate λ computed as

$$\lambda_{new} = \sum_{i=0}^{N-1} r_{N-i} \left(\frac{M}{N-i} \right). \tag{18}$$

Finally, we can use Eqs. (14) and (16) with the new average arrival rate and constant packet service time $\frac{L}{C}$.

Figure 9 plots the average packet waiting time of a spectral-agile network when it uses a single channel and multiple channels. We fix the value of T_{off} at 1 second while varying the value of T_{on}, so as to vary each channel's average load imposed by the primary network. Obviously, the average

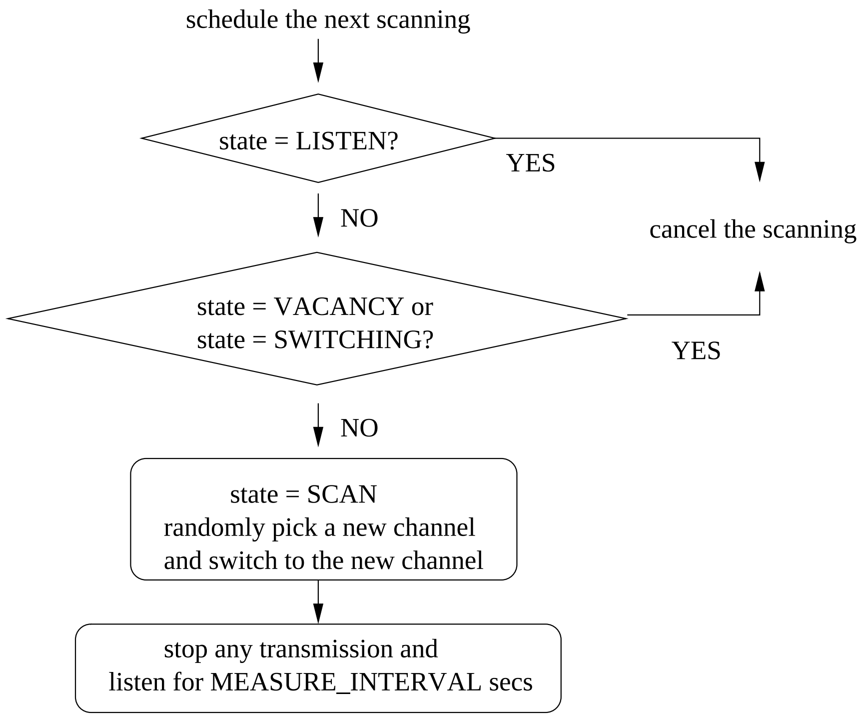

Fig. 10. Spectral opportunity discovery: before scanning

packet waiting time in the case of $M < N$ is less than that in the case of $M > N$ as there are less spectral-agile networks seeking spectral opportunities in the case of $M < N$. However, the packet waiting time of using multiple channels is always less than that of using a single channel in both cases. The improvement is even more significant in the case of $M > N$ as expected. These numerical results demonstrate the potential advantages of using multiple channels in a spectral-agile network, especially when $M > N$.

IV. Spectral-agility Functionalities

As mentioned in the Introduction, a spectral-agile secondary network needs three basic functionalities for spectral agility: spectral opportunity discovery, spectral opportunity management, and coordination of spectral opportunity uses. There are two basic principles to follow in realizing these functionalities. First, the concept of *Listen-Before-Talk* is applied whenever a spectral-agile network wants to exploit a spectral opportunity. Thus, primary networks will not be affected. For example, a wireless node in a spectral-agile network may stay in what we call the *LISTEN* state for *LISTEN_INTERVAL* seconds after switching to a new channel, or after switching back from the scanned channel to the original communicating channel (i.e., after scanning other channels). Second, whenever a spectral-agile network decides to switch to a different channel, all wireless nodes in that spectral-agile network should be notified of this switching and then switch to the same channel at the same time.

A. Spectral Opportunity Discovery

In order to utilize spectral opportunities, a spectral-agile network must be aware of the presence of spectral opportunities. Therefore, a wireless node of the spectral-agile network should scan the spectrum regularly. Figure 10 illustrates the basic scanning procedures for discovering spectral opportunities. At the beginning of each scheduled scan, a wireless node first schedules the next

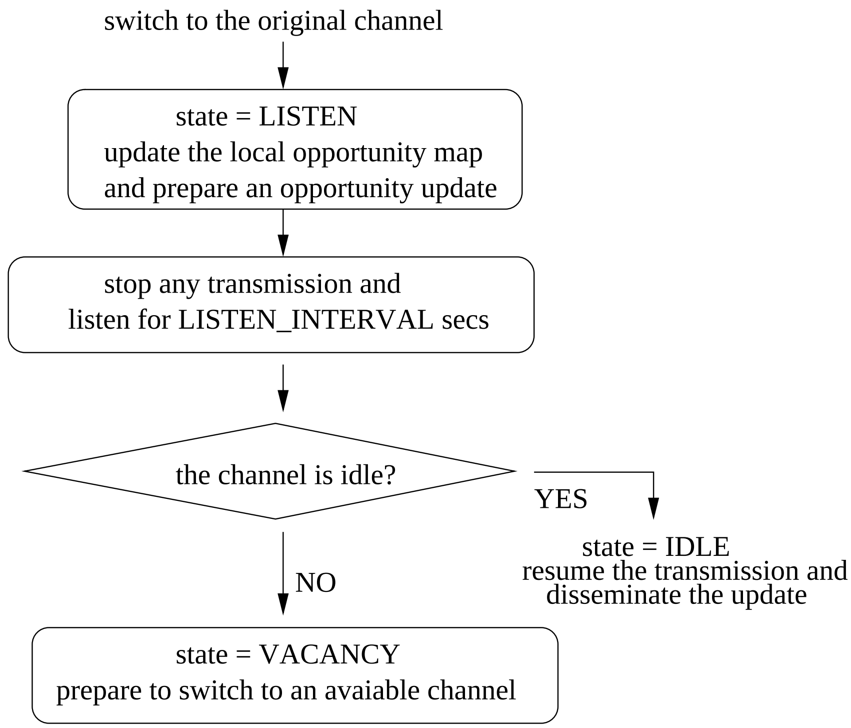

Fig. 11. Spectral opportunity discovery: after scanning

opportunity scan. If the node has been in the *LISTEN* state, this scan is canceled because the node must keep silent and listen on the current channel. Moreover, if the wireless node detects any activity of primary networks (i.e., in the *VACANCY* state) or is switching to a different channel (i.e., in the *SWITCHING* state), the scan is also canceled because if the node leaves the current channel for scanning, it may lose connectivity with the other nodes as they may also detect the presence of the primary networks and are about to switch to another channel (details are in the next subsections). If none of these two situations occurs, the wireless node randomly chooses a channel to scan and enters the *SCAN* state. At this time, the *Listen-Before-Talk* principle is applied.

During a channel scan, the wireless node records "activities" on that scanned channel. These activities are characterized by several parameters, including the fraction of time that the channel is deemed busy during the scan interval, the average received power and if possible, the activity type (either primary or secondary). These parameters will later be used by the spectral opportunity management to identify potential spectral opportunities. Upon completion of scanning, the wireless node switches back to the previous channel and enters the *LISTEN* state before resuming the normal transmission. In the meantime, that wireless node updates its database of spectral opportunities — called the *spectral opportunity map* (SOM)— based on the collected parameters and prepares to disseminate the latest opportunity information to the other wireless nodes in the same spectral-agile network. While staying in the *LISTEN* state, if a wireless node detects any activity of primary networks on the current channel, the wireless node prepares to vacate the current channel, which is managed by the spectrum use coordination as we will discuss later. This procedure is illustrated in Figure 11.

B. Spectral Opportunity Management

Each wireless node in a spectral-agile network maintains a SOM, which stores the status of each channel in the wireless spectrum. There are two methods for updating the SOM: by scanning

2	8	8	8	8
Index	Duration	P_utilization	S_utilization	P_power

(a) Spectral opportunity update

Index	Idle	T_Duration	avg_P_util	avg_S_util	P_power
0	1	10	0.85	0.23	−20 db
1	0	24	0.17	0	−10 db
2	0	N/A	N/A	N/A	N/A
N	1	N/A	N/A	N/A	N/A

(b) Spectral opportunity map

Fig. 12. Spectral opportunity management (SOM)

a channel and by receiving spectral opportunity updates from the other wireless nodes. As we mentioned in the previous subsection, a wireless node disseminates the opportunity update after resuming transmission in the original channel. The information contained in an opportunity update is listed in Figure 12-(a), where the "Index" field represents the channel index, the "Duration" field represents the scanning duration, the "P_/S_utilization" field represents the percentage of the scanning duration when activities from primary/secondary networks are detected, and the last field represents the average detected power of primary networks' transmissions.

Figure 12-(b) shows an example of a wireless node's SOM. The "Idle" field indicates if a channel is available or not. For example, a value of 1 means that the channel is idle and considered as a spectral opportunity. This field is set to 0 when the latest spectral opportunity update contains a non-zero P_utilization. The "T_Duration field" represents the accumulative amount of time that a spectral-agile network has scanned for that channel. T_Duration is used to compute the average channel utilization of primary and secondary networks (i.e., the "avg_P_util" and "avg_S_util" fields in the SOM). The value of avg_P_util is updated by

$$avg\_P\_util = \frac{\textit{T\_Duration} \cdot \textit{avg\_P\_util} + \textit{Duration} \cdot \textit{P\_utilization}}{\textit{T\_Duration} + \textit{Duration}}, \tag{19}$$

and so are the values of avg_S_util and avg_P_power. The average channel utilization and average power are useful when multiple spectral opportunities are available, and thus, help a spectral-agile network choose a "good" opportunity. One should note that the time duration of each opportunity is not included in the SOM simply because it is difficult to predict or estimate such information, given that the primary networks may reclaim the channels at any time. As we will explain in the next section, spectral-agile nodes have to vacate the channel immediately upon detection of any activity from the primary network. Therefore, a spectral-agile node needs to know whether or not a channel is available, instead of how long it may last.[2]

Note that different wireless nodes in the same spectral-agile network may have different SOMs because a wireless node can miss some opportunity updates from the other nodes. However, it is

[2]Of course, any additional information, such as the duration of channel availability, if available, may help a node make a better decision on spectral opportunity use.

not essential for all nodes to maintain a network-wide, unique SOM as long as all wireless nodes in the same spectral-agile network coordinate their channels switching. This will be detailed in the next subsection.

C. Spectral Use Coordination

To enable automatic and cooperative use of available opportunities among radio devices of a spectral-agile network or among different spectral-agile networks, we need a resource-use coordination mechanism to resolve any potential conflict/contention in utilizing these opportunities. Based on the participants involved in resource-use coordination, we propose two control mechanisms as described below.

1) Intra-network Resource-use Synchronization: The most challenging task in realizing a spectral-agile network is to maintain inter-node connectivity in a spectral-agile network. For example, if some wireless nodes decide to switch to channel 1 while the others decide to switch to channel 2, then these nodes will lose their connectivity to each other. Figure 13 depicts the operations of channel switching when a spectral-agile network detects a primary network activity. Upon detection of a primary network activity, the wireless node enters the *VACANCY* state, and searches its SOM for any spectral opportunity. If there is not any available spectral opportunity, then the wireless node remains in the *VACANCY* state and cancels the next upcoming scanning. The reason for cancelling the upcoming scanning is that the other wireless nodes (in the same spectral-agile network) could have found a spectral opportunity and are about to disseminate a switch notification. If this wireless node now leaves the current channel for scanning, it may miss the notification and lose the connection with the others. In case a wireless node locates some spectral opportunities, the node prepares a switch notification and waits for *VACANCY_INTERVAL* seconds before sending such a notification. This ensures the other wireless nodes which have been in the *SCAN* state to have enough time to finish the scanning, switch back to the original channel, and still receive this notification given that

$$\mathit{VACANCY\_INTERVAL} > \mathit{MEASURE\_INTERVAL} + \mathit{LISTEN\_INTERVAL}. \tag{20}$$

Once the node successfully sends the switch notification, it enters the *SWITCHING* state and prepares to switch to the new channel.

To avoid disseminating a switch notification at the same time as the others, each wireless node waits for extra *OFFSET* seconds before sending a notification. Obviously, each node must have a unique value of *OFFSET*. Finally, to further avoid receiving different switch notifications from different nodes, a node with a pending transmission of a switch notification cancels its own notification after receiving a switch notification from the others. Together with the transmission offset, only one unique switch notification will be disseminated and received by all wireless nodes in a spectral-agile network. This operation is depicted in Figure 14.

Note that it is always possible that a wireless node may miss a switch notification due to transmission errors. Therefore, there is no absolute guarantee for synchronized switches even if one applies other sophisticated retransmission and handshaking mechanisms. One may try to establish a network-wide, unified SOM so that, whenever a spectral-agile network needs to vacate a channel, all nodes in that network choose the same spectral opportunity without requiring the need to notify each other. By doing so, the difficulty shifts from securely disseminating a switch notification to securely disseminating *all* spectral opportunity updates. Since updating the SOM is more frequent than sending a switch notification, our current implementation should be more reliable. In any case, all wireless nodes may either switch back to the previous communicating channel or a pre-defined channel for re-synchronization, when perceiving the existence of a missing node (from the same spectral-agile network) after switching to a new channel.

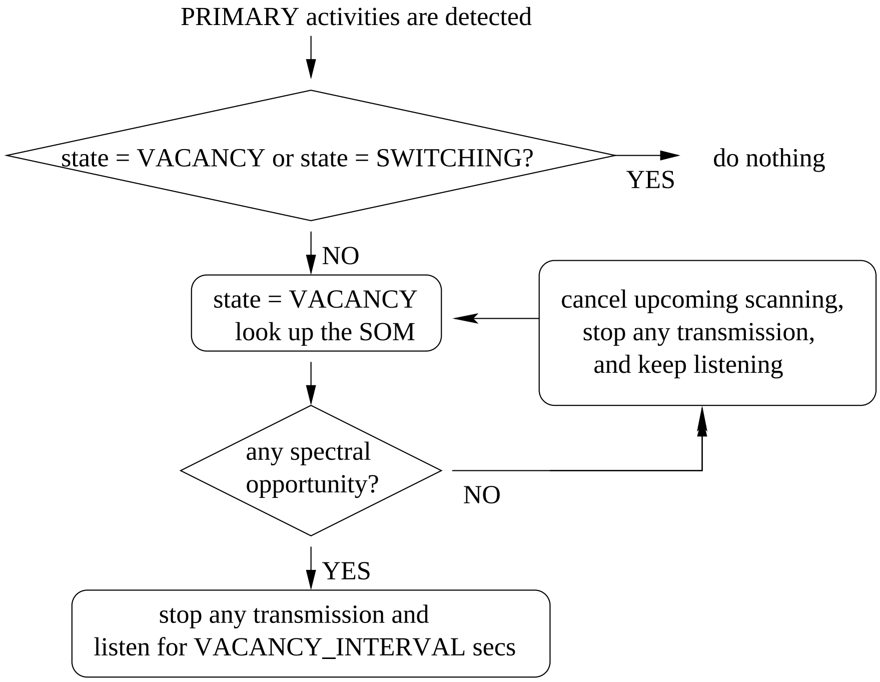

Fig. 13. Spectral opportunity use: preparation for vacating a channel

2) Inter-network Resource-use Cooperation: To make two different spectral-agile networks cooperatively utilize a resource opportunity, we need (i) a multiple access control so that the networks may fairly share the spectrum, and (ii) a “load-balancing” mechanism so that each spectral-agile network may utilize a different opportunity, if multiple opportunities exist. The first goal can be easily achieved by using the IEEE 802.11 standard-like carrier-sense-multiple-access/collision avoidance (CSMA/CA) with exponential random backoffs. To achieve the second goal, we propose a load-balancing algorithm to coordinate the use of multiple opportunities among different spectral-agile networks. When a spectral-agile network detects the presence of any other spectral-agile network, wireless nodes in that spectral-agile network should immediately check the availability of other opportunities in its own SOM. If any opportunity other than the currently-utilized opportunity is located, the wireless node follows the intra-network resource-use synchronization procedure to switch to that opportunity. To prevent all involved spectral-agile networks from renouncing the currently-utilized opportunity, a delay (unique for each spectral-agile network) is introduced so that only one of the spectral-agile network actually changes its use of the opportunity. If a spectral-agile network does renounce the current opportunity (i.e., a channel), those spectral-agile networks that have not vacated yet will cancel their intra-network synchronization procedure, after perceiving the absence of that leaving spectral-agile network. These “staying” spectral-agile networks may update the channel status in their SOMs and repeat the above procedure, if they are able to locate other opportunities. This way, we can achieve a balanced (and maximum) resource utilization in a distributed manner.

V. Evaluation

The three basic components of a spectral-agile network in Section IV are implemented in *ns-2* so that we can evaluate the performance (as compared to analytical upper bounds) and the effects of overhead associated with spectral agility. We use the IEEE 802.11 standard as the MAC-layer protocol for spectral-agile secondary networks. The wireless nodes in a primary network also use

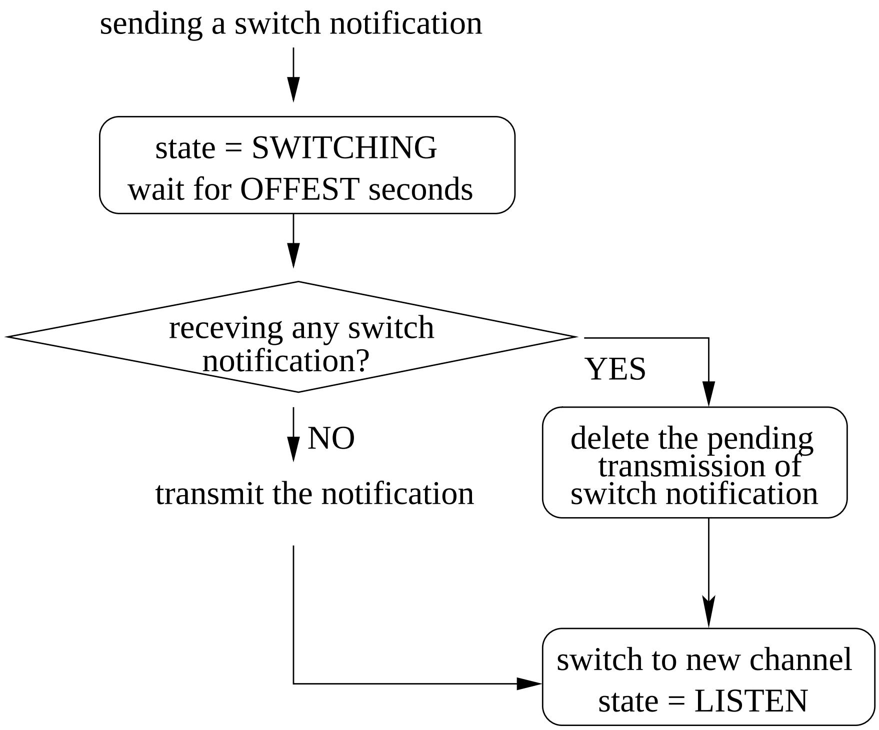

Fig. 14. Spectral opportunity use: dissemination of a switching notification

the IEEE 802.11 MAC standard but they have exclusive access of the designated channel. If an IEEE 802.11 node in a spectral-agile network detects any activity of an IEEE 802.11 node in a primary network, the node in the spectral-agile network suspends any transmission as explained before.

We assume that there is only one primary network in each channel, and there are two wireless nodes in each primary network. One of these two nodes has an *ns-2* ON/OFF traffic generator and transmits packets to the other node in the same primary network. The average channel utilization of a primary network is then determined by the mean values of ON- and OFF-periods. We also assume that there are 3 wireless nodes in a spectral-agile network. To fully utilize spectral opportunities, we use the *ns-2* constant-bit-rate (CBR) traffic generator so that nodes in the spectral-agile network always have packets to transmit as we assumed in Section III. Finally, we assume that the packet size from all traffic generators is 500 bytes and all wireless nodes use 1-Mbps for data transmission. Figure 15 shows the simulation setup for the case of three channels (i.e., channels 1, 6 and 11 in the IEEE802.11 standard) with a single spectral-agile network.

As explained in Section IV, several parameters are needed to control a spectral-agile network, namely *MEASURE_PERIOD*, *MEASURE_INTERVAL*, *VACANCY_INTERVAL*, and *LISTEN_INTERVAL*. The value of *MEASURE_PERIOD* determines the frequency of seeking a spectral opportunity map (SOM). Obviously, the smaller a node's *MEASURE_PERIOD*, the more accurate the SOM becomes. However, a small value of *MEASURE_PERIOD* incurs more control overhead (e.g., frequent dissemination of opportunity updates to other nodes), and interrupts normal transmission more frequently. The value of *MEASURE_INTERVAL* determines time granularity of the spectral opportunities that a spectral-agile network can detect. If the duration of a spectral opportunity is less than *MEASURE_INTERVAL*, a spectral-agile network cannot detect the existence of such a spectral opportunity because the scanned channel becomes "busy" before the scanning is completed. However, choosing too small a *MEASURE_INTERVAL* value is not a good idea either, sim-

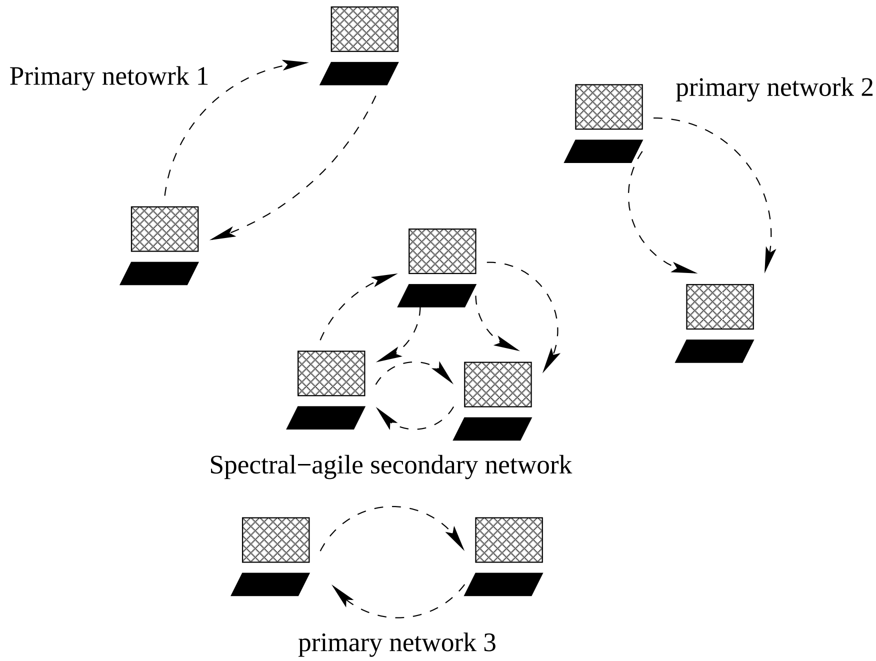

Fig. 15. Simulation setup for single spectral-agile network: $N = 3$ and $M = 1$

ply because not enough "activities" will be collected. The same criteria can be applied to choose the value of *LISTEN_INTERVAL* since choosing too small or too large a value results in either interfering primary networks (resuming transmission too fast) or wasting a spectral opportunity (waiting too long). Finally, we choose the value of *VACANCY_INTERVAL* according to Eq. (20).

Based on the transmission rate and packet size chosen above, we let *MEASURE_INTERVAL* = 20 ms, *LISTEN_INTERVAL* = 10 ms, and *VACANCY_INTERVAL* = 40 ms in all of the simulation runs.[3] However, we change the value of *MEASURE_PERIOD* in order to investigate its impact on both performance improvement and control overhead. In the following simulation, we use $N = 3$ as we want to simulate the case of using spectral agility in the current IEEE 802.11b wireless LAN in the 2.4-GHz band. Of course, these mechanisms can be applied to other types of networks and other spectral bands, once the regulatory restriction is removed.

A. Throughput Improvement for a Single Spectral-agile Network

We choose *MEASURE_PERIOD* = 0.5 second, $T_{on} = 10 * channel\ load$ seconds, and $T_{off} = 10 * (1 - channel\ load)$ seconds in this simulation. Figure 16 shows the improvements of a spectral-agile network's throughput as compared to a network without spectral agility. Here, we use throughput as the performance metric since the MAC protocol (i.e., the IEEE 802.11b standard) is specified. We consider both homogeneous and heterogeneous loads, and the simulation results are compared with the analytical results (in solid lines). The improvement obtained from the simulation is shown to be very close to the analytical upper bound in some cases, especially when the average channel load ranges between 0.3 and 0.6. Within this region, the improvement ranges between 40 and 80% for homogeneous loads, and ranges between 50 and 90% for heterogeneous loads. Considering the control overhead incurred by spectral agility, the results verify the effectiveness of our implementation.

One interesting observation is that the improvement is much less than the analytical results as the channel load increases, and using spectral agility is even worse (-22%) than without using spectral

[3]It should be noted that we only focus on the case when the average duration of a spectral opportunity is in the order of seconds.

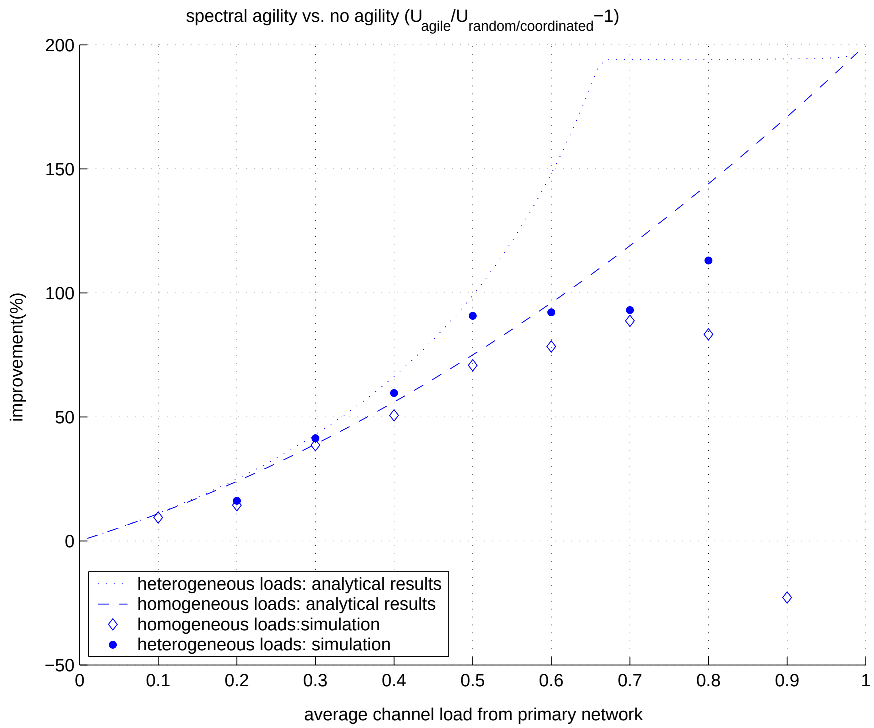

Fig. 16. A single spectral-agile network: spectral agility vs. no agility with random/coordinated channel selection. *The substantial discrepancy between the analytical and simulation results when the channel load approaches 1 results from that our analytical model does not consider any scanning/control overhead. However, these overheads easily consume the minuscule channel access time (as shown in Figure 4) gained by spectral agility when the load is close to 1.

agility when the channel is extremely busy (because of primary networks' activities). The main reason for this is that when the channel is heavily-loaded, a spectral-agile network has few spectral opportunities. The scanning, listening, and switching simply interrupt the network's normal transmission without finding many opportunities. Under this circumstance, staying with a fixed channel is better. That is, one should not use spectral agility in extremely busy spectral bands in the first place.

Figure 16 also confirms that when the loads of the channels are diverse, a spectral-agile network achieves better performance as shown in Section III. One can make an extra 10 to 15% improvement since a spectral-agile network dynamically searches for the least-utilized channels and makes use of them more efficiently.

B. Throughput Improvement of Multiple Spectral-agile Networks

The previous simulation shows that the throughput of a single spectral-agile network increased by up to 90%. We now use $N = 3$ and $M = 2$ to investigate how spectral-agile networks interact with each other when seeking and utilizing spectral opportunities as shown in Figure 17. For an illustrative purpose, we only simulate the case of homogeneous channel loads and set *MEASURE_PERIOD*=0.5 second. In order to make these two spectral-agile networks share the spectral opportunities, instead of letting them compete for these opportunities, we assign different priorities to each spectral-agile network. The priority is used by a spectral-agile network to determine the value of delay in the inter-group resource-use cooperation algorithm. If a lower-priority spectral-agile network detects the existence of a higher-priority spectral-agile network, the lower-priority network vacates the current channel first ***if and only if*** the SOM indicates that there exist other available spectral opportunities. This way, the lower-priority network is not discriminated in terms

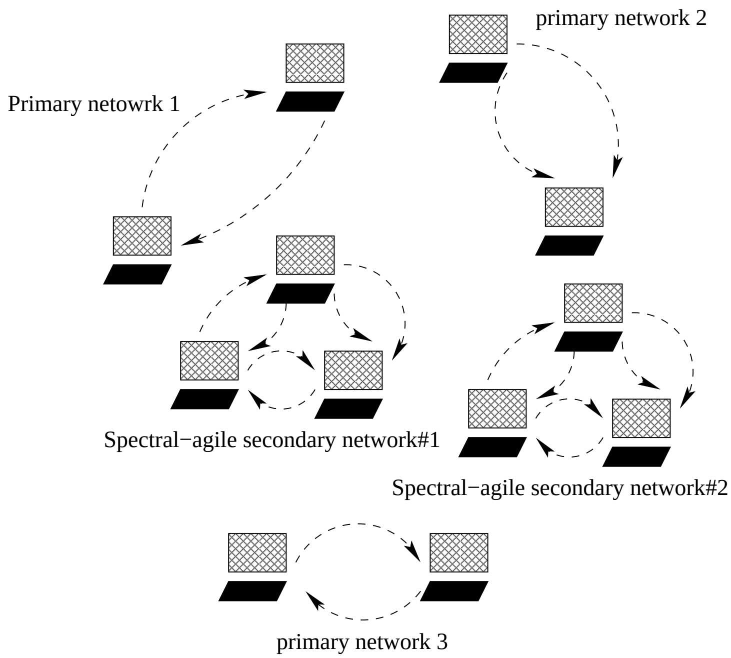

Fig. 17. Simulation setup for multiple spectral-agile networks: $N = 3$ and $M = 2$

of using spectral opportunities. Our simulation results show that these two spectral-agile networks always achieve almost the same throughput.

Figure 18 shows the improvement of spectral-agile networks' average throughput, as compared to the case of using no agility with coordinated channel selection. In general, the improvements are very close to the analytical results (within a 13% margin). One reason why the simulation gives more improvements than the analytical bound (uder moderate channel loads) is that a non-agile secondary network also suspends the transmission for *VACANCY_INTERVAL* seconds before detecting that channel again, if the network has detected any activity of the primary network in the assigned channel. For a spectral-agile network, it is less likely to encounter a busy channel because of spectral agility, especially when the channels are moderately-loaded. That is, the overhead of detecting the (channel) idleness in a non-agile network is higher than a spectral-agile network when the channel is moderately-loaded, and so is the amount of time wasted on waiting. One can also observe that using spectral agility results in poorer performance (-9%) than without using agility, when the channels are heavily-loaded. Again, it does not make any sense to use spectral agility in those heavily-loaded channels as virtually no opportunity exists in those channels. Thus, the overhead easily offsets any improvement made by spectral agility as in the case of a single spectral-agile network.

The simulation results also demonstrate a very important advantage of using spectral agility: *by using spectral agility, we can achieve a higher throughput (more than 30% in many cases, as compared to using no agility with coordinated channel selection, let alone an even higher improvement as compared to using random channel selection) without any off-line planning on spectral resource allocation*. That is, using spectral agility easily achieves the *automated frequency use coordination* as we mentioned in Section I and results in a much higher spectral utilization.

C. *Improvements vs.* MEASURE_PERIOD

We now investigate the effects of *MEASURE_PERIOD* on the improvement of a spectral-agile network's throughput. We choose three different loads for the primary network, 0.2, 0.5 and 0.8,

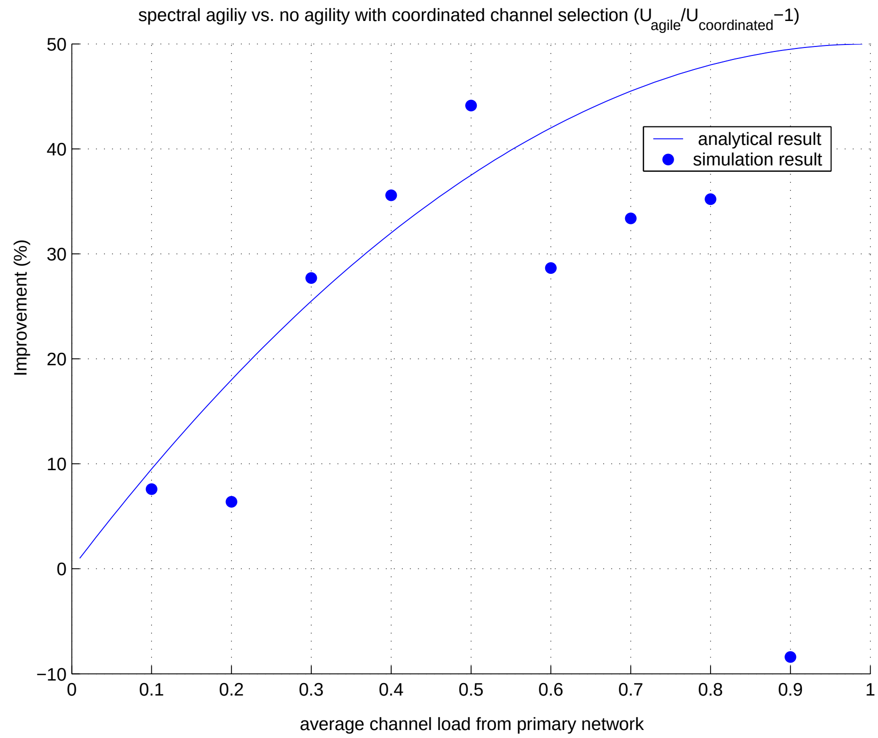

Fig. 18. Multiple spectral-agile networks: spectral agility vs. no agility with coordinated channel selection. *The substantial discrepancy between the analytical and simulation results when the channel load approaches 1 results from that our analytical model does not consider any scanning/control overhead. However, these overheads easily consume the minuscule channel access time (as shown in Figure 4) gained by spectral agility when the load is close to 1.

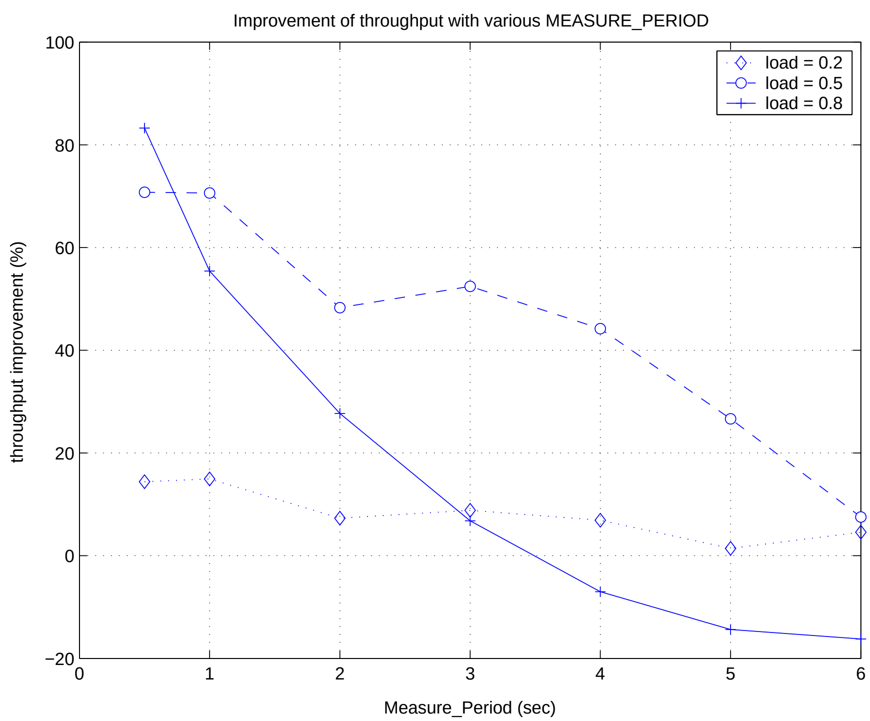

Fig. 19. Impacts of *MEASURE_PERIOD* on the spectral-agile network's throughput improvement

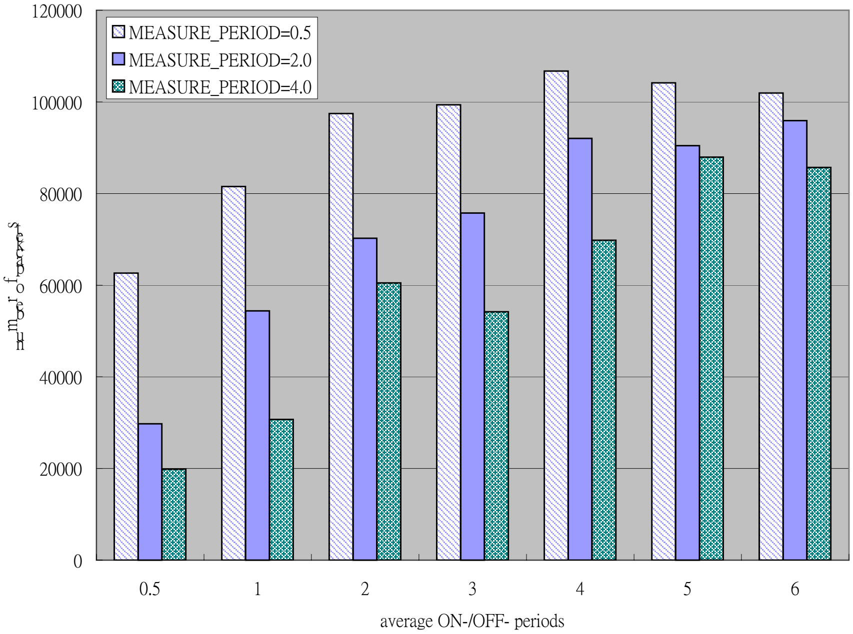

Fig. 20. Effects of *MEASURE_PERIOD* vs. Effects of average *ON-/OFF-period* on the spectral-agile network throughput

still use $T_{on} = 10 * channel\ load$ seconds and $T_{off} = 10 * (1 - channel\ load)$ seconds, and change the value of *MEASURE_PERIOD*. Figure 19 shows that for a fixed channel load, the improvement decreases with the increase of *MEASURE_PERIOD*. This is because the less frequently a spectral-agile network scans the spectrum, with a lower probability an available channel can be found. Therefore, it is very important for a spectral-agile network to choose an appropriate *MEASURE_PERIOD* value since choosing too large a value of *MEASURE_PERIOD* may result in poor performance, especially when the channel is heavily-loaded with the traffic of primary networks. It is when the channel is very busy that a spectral-agile network needs spectral opportunities most. Thus, using a large value of *MEASURE_PERIOD* degrades the improvements most when the channel load is high. This explains the decrease of throughput improvement when the load is 0.8.

In fact, one can conclude that the improvement of a spectral-agile network is primarily determined by the value of *MEASURE_PERIOD*. A spectral-agile network should choose a *MEASURE_PERIOD* based on the channel loads, and more importantly, the duration of ON-/OFF-period in each channel. If the channels switch between ON- and OFF-periods very often, a smaller *MEASURE_PERIOD* is required. That is, the degree of agility that a spectral-agile network needs, depends on the dynamics of the scanned spectrum. Therefore, using an adaptive *MEASURE_PERIOD* should achieve better performance.

D. Improvements vs. Duration of a Spectral Opportunity

As discussed above, the throughput improvement of a spectral-agile network is determined by *MEASURE_PERIOD* and the average duration of ON-/OFF-periods of primary networks. To be on the safe side, one may choose a very small *MEASURE_PERIOD* in order to exploit the spectral agility. A potential problem with this is that too frequent scanning interrupts too often normal transmission of the spectral-agile network and also incurs high overhead. We investigate such a trade-off as follows. We choose 3 different values of *MEASURE_PERIOD*. For each *MEASURE_PERIOD* value, we change the T_{on} and T_{off} values but keep the channel load ($= \frac{T_{on}}{T_{on}+T_{off}}$=0.5) unchanged.

The total number of packets transmitted (by the spectral-agile network) within a 1000-second interval is plotted in Figure 20.

For any given value of *MEASURE_PERIOD*, the number of transmitted packets generally increases with the average duration of ON-/OFF-periods (i.e., T_{on} and T_{off}). Of course, a spectral-agile network need not scan the channels too frequently when T_{on}/T_{off} is relatively large (compared to *MEASURE_PERIOD*) since the switching also occurs less frequently. This explains the slight decrease for the case of *MEASURE_PERIOD*=0.5 after the average ON-/OFF-periods are larger than 4.0 seconds. However, as compared to using a larger *MEASURE_PERIOD*, using a smaller *MEASURE_PERIOD* always achieves much better performance even though the overhead increases linearly with the scanning frequency. This is because the overhead incurred by scanning is relatively small in our implementation (only *MEASURE_INTERVAL*+*LISTEN_INTERVAL* =0.03 second for every *MEASURE_PERIOD*=0.5 second).

VI. Conclusions and Future Work

In this paper, we have investigated the issue of using spectral agility to improve both spectral utilization efficiency and secondary networks performance. We established a simple mathematical model to analyze spectral-agile networks, and provided a performance benchmark by which different implementations of spectral-agile networks can be evaluated. The results (based on this model) have shown that the channel utilization of a spectral-agile network is improved by 35 to 200% when compared to the cases of no agility, depending on the primary network's channel utilization and the number of spectral-agile networks.

In order to realize a spectral-agile network, we proposed three basic functionalities, namely spectral opportunity discovery, spectral opportunity management, and spectral use coordination. These functionalities have been added to the IEEE 802.11 wireless LAN in the *ns-2*. The simulation results show that (1) the improvement of a secondary network's throughput can be up to 90% by using spectral agility, (2) the improvement is close to the performance bound predicted by our analytical model, and (3) the improvement is achieved in a distributed and automated way with little overhead, and outperforms the improvement of non-agile networks using static, coordinated channel selection.

We are currently examining spectral-agile networks which use multiple spectral opportunities simultaneously, and studying its improvement of packet waiting time.

References

[1] "Facilitating Opportunities for Flexible, Efficient, and Reliable Spectrum Use Employing Cognitive Radio Technologies," *the FCC Notice of Proposed Rulemaking and Order — ET Docket No. 03-108*.
[2] http://www.darpa.mil/ato/programs/XG/rfc_vision.pdf — *Vision RFC*.
[3] C. Bergstrom, S. Chuprun and D. Torrieri, "Adaptive spectrum exploitation using emerging software defined radios", *IEEE Radio and Wireless Conference 1999*, pp.113-116.
[4] P. K. Lee, "Joint frequency hopping and adaptive spectrum exploitation", *IEEE MILCOM 2001*, Vol.1, pp.566-570.
[5] J. Mitola, "The software radio architecture", *IEEE Communications*, Vol.33, No.5, 1995, pp.26-38.
[6] "White paper on regulatory aspects of software defined radio", *SDR forum document number SDRF-00-R-0050-v0.0*.
[7] http://www.darpa.mil/ato/programs/XG/rfc_af.pdf — *Architecture RFC*.
[8] J.B. Punt, D. Sparreboom, F. Brouwer and R. Prasad, "Mathematical analysis of dynamic channel selection in indoor mobile wireless communication systemse", *IEEE Transactions onVehicular Technology*, Vol.47, No.4, Nov. 1998, pp.1302-1313.
[9] "IEEE 802.11h standard (Amendment to IEEE 802.11 Standard, 1999 Edition)", 2003.
[10] J. Khun-Jush, G. Malmgren, P. Schramm and J. Torsner, "Overview and performance of HIPERLAN type 2-a standard for broadband wireless communications", *IEEE VTC 2000-Spring, Tokyo*, Vol.1, pp.112 - 117.
[11] "Additional spectrum for unlicensed devices below 900 MHz and in the 3 GHz band", *FCC ET Docket No. 02-380*.
[12] "Future trends in defence antenna technology", http://www.bcba15324.pwp.blueyonder.co.uk/consulting/Bibliography/Paper-16.pdf.
[13] D. Bertsekas and R. Gallager, "Data Networks", *Prentice-Hall*, New York, 1992.